Lecture 1: January 28

Overview. The purpose of this course is to give an introduction to the theory of algebraic $\mathscr{D}$-modules. I plan to cover roughly the following topics:

- modules over the Weyl algebra A_n
- $\mathscr{D}$-modules on smooth algebraic varieties
- functors on $\mathscr{D}$-modules (and how they relate to PDE)
- holonomic $\mathscr{D}$-modules, regularity (with a focus on what it means)
- b-functions, localization along a hypersurface
- $\mathscr{D}$-modules of normal crossing type (as a class of examples)
- Riemann-Hilbert correspondence (with proofs in the normal crossing case)
- some applications, either to representation theory or to algebraic geometry

The website for the course,

`http://www.math.stonybrook.edu/~cschnell/mat615`,

contains a list of useful references.

Introduction. Very briefly, $\mathscr{D}$-modules were invented in Japan (by Mikio Sato, Masaki Kashiwara, and others) and France (by Alexander Grothendieck, Zogman Mebkhout, and others). It has its origins in the field of "algebraic analysis", which means the study of partial differential equations with algebraic tools. The theory of algebraic $\mathscr{D}$-modules was further developed by Joseph Bernstein.

Systems of linear equations. $\mathscr{D}$-modules arise naturally from systems of linear partial differential equations. To get a better understanding of how this works, let us first look at the example of a system of linear equations

$$\sum_{j=1}^{q} a_{i,j} x_j = 0, \quad i = 1, \dots, p, \tag{1.1}$$

with coefficients $a_{i,j}$ in a field K (such as $\mathbb{R}$ or $\mathbb{C}$). In linear algebra, one usually transforms such a system in various ways, for example by making a substitution in the unknowns $x_1, \dots, x_q$, or by taking linear combinations of the equations. One can associate to the system in (1.1) a single K-vector space that is invariant under such transformations. Consider the linear mapping

$$\varphi \colon K^p \to K^q, \quad \varphi(y_1, \dots, y_p) = \left(\sum_{i=1}^{p} y_i a_{i,1}, \dots, \sum_{i=1}^{p} y_i a_{i,q} \right),$$

built from the coefficient matrix of the system in (1.1), and define the K-vector space $V = \ker \varphi = K^q / \varphi(K^p)$. It sits in the exact sequence

$$K^p \xrightarrow{\varphi} K^q \xrightarrow{\pi} V \longrightarrow 0,$$

and the solution space to (1.1) can be recovered from V as

$$\operatorname{Hom}_K(V, K) = \left\{ f \colon K^q \to K \;\middle|\; f \circ \varphi = 0 \right\}.$$

Indeed, a linear mapping from V to K is the same thing as a linear mapping $f \colon K^q \to K$ whose composition with φ is equal to zero.

$$\begin{tikzcd} K^p \arrow[r, "\varphi"] \arrow[dr, "0"'] & K^q \arrow[r, "\pi"] \arrow[d, "f"] & V \arrow[dl] \\ & K & \end{tikzcd}$$

Now f is uniquely determined by the q scalars $x_j = f(e_j) \in K$, where e_j denotes the j-th coordinate vector in K^q. Since $f \circ \varphi = 0$, we get

$$\sum_{i,j} y_i a_{i,j} x_j = 0$$

for every $(y_1, \dots, y_p) \in K^p$. This means exactly that $(x_1, \dots, x_q) \in K^q$ is a solution to the system of linear equations in (1.1).

The same construction can be applied to systems of linear equations with coefficients in other rings. For example, let $R = K[x_1, \dots, x_n]$ be the polynomial ring in n variables, and consider the system of linear equations

$$\sum_{j=1}^{q} f_{i,j} u_j = 0, \quad i = 1, \dots, p, \tag{1.2}$$

with polynomial coefficients $f_{i,j} \in R$. As before, we can associate to the system an R-module $M = \operatorname{coker} \varphi$, defined as the cokernel of the morphism of R-modules

$$\varphi \colon R^p \to R^q, \quad \varphi(v_1, \dots, v_p) = \left(\sum_{i=1}^{p} v_i f_{i,1}, \dots, \sum_{i=1}^{p} v_i f_{i,q} \right),$$

and the space of solutions $(u_1, \dots, u_q) \in R^q$ to the system in (1.2) can be recovered from M as $\operatorname{Hom}_R(M, R)$. This formulation has the advantage that we can describe solutions over other R-algebras S, such as the ring of formal power series $K[[x_1, \dots, x_n]]$, in the same way, by taking $\operatorname{Hom}_R(M, S)$.

Note. The polynomial ring R is noetherian, meaning that every ideal of R is finitely generated. This implies that every submodule of a finitely generated R-module is again finitely generated. In particular, every finitely generated R-module is isomorphic to the cokernel of $\varphi \colon R^p \to R^q$ for some $p, q \in \mathbb{N}$. Studying systems of linear equations with coefficients in R is therefore the same thing as studying finitely generated R-modules.

Systems of linear partial differential equations. We now apply the same construction to systems of linear partial differential equations with coefficients in the polynomial ring. The role of the polynomial ring $R = K[x_1, \dots, x_n]$ is played by the *Weyl algebra* $A_n = A_n(K)$. The elements of A_n are linear partial differential operators of the form

$$P = \sum_{i_1, \dots, i_n} f_{i_1, \dots, i_n}(x_1, \dots, x_n) \frac{\partial^{i_1}}{\partial x_1^{i_1}} \cdots \frac{\partial^{i_n}}{\partial x_n^{i_n}},$$

where $f_{i_1, \dots, i_n} \in R$, and the sum is finite. To simplify the notation, we put $\partial_j = \partial / \partial x_j$, and write the above sum in multi-index notation as

$$P = \sum_{\alpha, \beta} c_{\alpha,\beta} x^\alpha \partial^\beta,$$

where $x^\alpha = x_1^{\alpha_1} \cdots x_n^{\alpha_n}$, and $\partial^\beta = \partial_1^{\beta_1} \cdots \partial_n^{\beta_n}$. We can multiply two differential operators in the obvious way, using the relations

$$[x_i, x_j] = 0, \quad [\partial_i, \partial_j] = 0, \quad [\partial_i, x_j] = \delta_{i,j}, \tag{1.3}$$

where $\delta_{i,j} = 1$ if $i = j$, and $\delta_{i,j} = 0$ otherwise. The relation $[\partial_i, \partial_j] = 0$ expresses the equality of mixed partial derivatives; the relation $[\partial_i, x_j] = \delta_{i,j}$ is a consequence of the product rule:

$$\frac{\partial}{\partial x_i}(x_j f) = \frac{\partial x_j}{\partial x_i} f + x_j \frac{\partial f}{\partial x_j} = \delta_{i,j} f + x_j \frac{\partial}{\partial x_i} f$$

Multiplication of differential operators turns A_n into a non-commutative ring. Differential operators naturally act on polynomials, by the usual (algebraic) rules for computing derivatives of polynomials; if we denote the action of a differential operator P on a polynomial f by the symbol Pf, we obtain a linear mapping

$$A_n \times R \to R, \quad (P, f) \mapsto Pf.$$

This makes the polynomial ring R into a left module over the Weyl algebra A_n.

The action on polynomials leads to the following more intrinsic description of the Weyl algebra: A_n is the smallest subring of the ring of K-linear endomorphisms

$$\operatorname{Hom}_K\bigl(K[x_1, \dots, x_n], K[x_1, \dots, x_n]\bigr)$$

that contains $K[x_1, \dots, x_n]$ and the partial derivative operators $\partial_1, \dots, \partial_n$. Algebraically, one can also describe the Weyl algebra by generators and relations: A_n is the non-commutative K-algebra generated by the $2n$ symbols $x_1, \dots, x_n, \partial_1, \dots, \partial_n$, subject to the relations in (1.3).

Now suppose that we have a system of linear partial differential equations

$$\sum_{j=1}^{q} P_{i,j} u_j = 0, \quad i = 1, \dots, p, \tag{1.4}$$

with differential operators $P_{i,j} \in A_n$. As before, we consider the morphism of left A_n-modules

$$\varphi \colon A_n^p \to A_n^q, \quad \varphi(Q_1, \dots, Q_p) = \left(\sum_{i=1}^{p} Q_i P_{i,1}, \dots, \sum_{i=1}^{p} Q_i P_{i,q} \right),$$

and associate to the system in (1.4) the left A_n-module

$$M = \operatorname{coker} \varphi = A_n^q / \varphi(A_n^p).$$

Note that it becomes necessary to distinguish between left and right A_n-modules, because A_n is non-commutative. We can again recover the solutions to the system in (1.4) directly from M, as follows. Let S be any commutative K-algebra with an action by differential operators, meaning that S is a left A_n-module. Examples are the polynomial ring $R = K[x_1, \dots, x_n]$, the ring of formal power series $K[[x_1, \dots, x_n]]$, etc. For $K = \mathbb{R}$ or $K = \mathbb{C}$, we might also be interested in the ring of convergent power series, the ring of C^∞-functions, etc. In any of these examples, the solutions in S are given by the formula

$$\operatorname{Hom}_{A_n}(M, S) = \left\{ f \colon A_n^q \to S \mid f \circ \varphi = 0 \right\}$$

Indeed, a morphism of left A_n-modules from M to S is the same thing as a morphism of left A_n-modules $f \colon A_n^q \to A_n$ whose composition with φ is equal to zero.

$$\begin{tikzcd}
A_n^p \arrow[r, "\varphi"] \arrow[dr, "0"'] & A_n^q \arrow[r, "\pi"] \arrow[d, "f"] & M \arrow[dl] \\
& S &
\end{tikzcd}$$

Once again, f is uniquely determined by the q functions $u_j = f(e_j) \in S$, where e_j denotes the j-th coordinate vector in A_n^q. Since $f \circ \varphi = 0$, we get

$$\sum_{i,j} Q_i P_{i,j} u_j = 0$$

for every $(Q_1, \dots, Q_p) \in A_n^p$, and so $(u_1, \dots, u_q) \in S^q$ solves the system of linear partial differential equations in (1.4).

Note. The Weyl algebra A_n is again left noetherian, meaning that every left ideal of A_n is finitely generated. (We will prove this next time.) This implies that every submodule of a finitely generated left A_n-module is again finitely generated. Studying systems of linear partial differential equations with coefficients in R is therefore the same thing as studying finitely generated left A_n-modules.

One advantage of this point of view is that we can describe the solutions to the system in a uniform way, by applying the *solution functor* $\operatorname{Hom}_{A_n}(M, -)$. We shall see later on that the solution functor is not exact (in the sense of homological algebra), and that it is natural to consider its derived functors. We shall also see that for so-called "regular holonomic" systems, one can recover the system up to isomorphism from its solutions (in the derived sense); this is the content of the famous Riemann-Hilbert correspondence.

Example 1.5. The exponential function $u = e^x$ solves the ordinary differential equation $u' = u$, which we can write in the form $(\partial - 1)u = 0$. The corresponding left A_1-module is $A_1/A_1(\partial - 1)$. The function $v = e^{1/x}$ solves the ordinary differential equation $-x^2 v' = v$, whose associated A_1-module is $A_1/A_1(x^2\partial + 1)$. Later on, when we discuss regularity, we shall see how the essential singularity of v at the point $x = 0$ affects the properties of the A_1-module $A_1/A_1(x^2\partial + 1)$.

Another advantage is that we can transform the system in (1.4) without changing the isomorphism class of the A_n-module M.

Example 1.6. Consider the second-order equation $a(x)u'' + b(x)u' + c(x)u = 0$, where $a, b, c \in K[x]$. A standard trick is to transform this into a system of two first-order equations $u_1' - u_2 = 0$ and $au_2' + bu_2 + cu_1 = 0$, by setting $u_1 = u$ and $u_2 = u'$. The first-order system leads to the left A_1-module

$$M_1 = \operatorname{coker}\left(A_1^2 \xrightarrow{\left(\begin{smallmatrix} \partial & -1 \\ c & a\partial + b \end{smallmatrix}\right)} A_1^2 \right)$$

and the second-order system to the left A_1-module

$$M_2 = A_1/A_1(a\partial^2 + b\partial + c)$$

Can you find an isomorphism between M_1 and M_2 as left A_1-modules?

Left and right A_n-modules. I already mentioned that it is necessary to distinguish between left A_n-modules and right A_n-modules, due to the non-commutativity of the Weyl algebra. Left A_n-modules naturally arise from functions, whereas right A_n-modules arise naturally from distributions. Let us look at the example of distributions in more detail. The $\mathbb{R}$-algebra $C_0^\infty(\mathbb{R}^n)$ of all compactly supported C^∞-functions on $\mathbb{R}^n$ is naturally a left $A_n(\mathbb{R})$-module; as before, we denote the action of a differential operator $P \in A_n$ on a test function $\varphi \in C_0^\infty(\mathbb{R}^n)$ by the symbol $P\varphi$. With the topology of uniform convergence of all derivatives on compact subsets, $C_0^\infty(\mathbb{R}^n)$ becomes a topological $\mathbb{R}$-vector space, and we denote by $\operatorname{Db}(\mathbb{R}^n)$ its topological dual. In other words, a distribution $D \in \operatorname{Db}(\mathbb{R}^n)$ is a continuous linear functional from $C_0^\infty(\mathbb{R}^n)$ to $\mathbb{R}$. We write the natural pairing between distributions and test functions as

$$\operatorname{Db}(\mathbb{R}^n) \times C_0^\infty(\mathbb{R}^n) \to \mathbb{R}, \quad (D, \varphi) \mapsto \langle D, \varphi \rangle.$$

In analysis, it is also common to use the more suggestive notation

$$\langle D, \varphi \rangle = \int_{\mathbb{R}^n} D\varphi \, d\mu,$$

where $d\mu$ is Lebesgue measure. Using formal integration by parts, $\mathrm{Db}(\mathbb{R}^n)$ naturally becomes a right A_n-module, by defining

$$\langle DP, \varphi \rangle = \langle D, P\varphi \rangle$$

for $D \in \mathrm{Db}(\mathbb{R}^n)$, $P \in A_n$, and $\varphi \in C_0^\infty(\mathbb{R}^n)$. For example, $D\partial_j$ is the distribution obtained by applying D to the test function $\partial\varphi/\partial x_j$. If we take any distribution, and act on it by differential operators, we obtain a right A_n-module inside $\mathrm{Db}(\mathbb{R}^n)$.

Example 1.7. Consider the delta function $\delta_0 \in \mathrm{Db}(\mathbb{R}^n)$, defined by $\langle \delta_0, \varphi \rangle = \varphi(0)$. Clearly, $\delta_0 x_1 = \cdots = \delta_0 x_n = 0$, and in fact, one can show that the right A_n-module generated by δ_0 is isomorphic to

$$A_n/(x_1, \ldots, x_n)A_n.$$

As an $\mathbb{R}$-vector space, this is just $\mathbb{R}[\partial_1, \ldots, \partial_n]$, but the A_n-action is nontrivial.

Exercises.

Exercise 1.1. Construct an isomorphism between the two left A_1-modules M_1 and M_2 in Example 1.6.

Exercise 1.2. Show that if $P \in A_n(\mathbb{R})$ satisfies $(P\varphi)(0) = 0$ for every test function $\varphi \in C_0^\infty(\mathbb{R}^n)$, then $P \in (x_1, \ldots, x_n)A_n$.

Lecture 2: January 30

Recall that the Weyl algebra $A_n = A_n(K)$ is generated by $x_1, \dots, x_n, \partial_1, \dots, \partial_n$, subject to the relations

$$[x_i, x_j] = 0, \quad [\partial_i, \partial_j] = 0, \quad [\partial_i, x_j] = \delta_{i,j},$$

Today, we begin studying A_n-modules in detail. One interesting difference between modules over A_n and modules over the polynomial ring $R = K[x_1, \dots, x_n]$ is the absence of nilpotents.

Example 2.1. As a $K[x]$-module, $K[x]/(x^2)$ is not isomorphic to two copies of K, because the action by x is nilpotent but not trivial. On the other hand, it is a fun exercise to show that the left A_1-module A_1/A_1x^2 is actually isomorphic to two copies of A_1/A_1x.

Left and right A_n-modules. The crucial difference between the Weyl algebra and the polynomial ring is that $A_n(K)$ is non-commutative. This means that we need to distinguish between left and right A_n-modules. In fact, there are no interesting two-sided A_n-modules.

Proposition 2.2. *$A_n(K)$ is a simple algebra, meaning that the only two-sided ideals of $A_n(K)$ are the zero ideal and $A_n(K)$.*

Proof. This follows from the commutator relations in A_n. We can write any $P \in A_n$ in multi-index notation as

$$P = \sum_{\alpha,\beta} c_{\alpha,\beta} x^\alpha \partial^\beta.$$

One can easily show by induction that

$$[\partial_j, x^\alpha \partial^\beta] = \alpha_j x^{\alpha - e_j} \partial^\beta \quad \text{and} \quad [x_j, x^\alpha \partial^\beta] = -\beta_j x^\alpha \partial^{\beta - e_j},$$

where $e_j \in \mathbb{N}^n$ is the j-th coordinate vector. Now suppose that $I \subseteq A_n$ is a nonzero two-sided ideal. Choose any nonzero $P \in I$, and write it as $P = \sum c_{\alpha,\beta} x^\alpha \partial^\beta$. Let

$$m = \max\bigl\{ \alpha_1 \bigm| c_{\alpha,\beta} \neq 0 \bigr\}$$

be the largest power of x_1 that appears in P. If $m \geq 1$, then by the formulas from above, the commutator

$$[\partial_1, P] = \partial_1 P - P \partial_1$$

is nonzero, and the maximal power of x_1 that appears is now $m-1$. Because I is a two-sided ideal, we still have $[\partial_1, P] \in I$. After repeating this operation m times, we obtain a nonzero element $P_1 \in I$ in which x_1 does not appear. Continuing in this way, we can successively eliminate $x_1, \dots, x_n$ by taking commutators with $\partial_1, \dots, \partial_n$, and then eliminate $\partial_1, \dots, \partial_n$ by taking commutators with $x_1, \dots, x_n$, until we arrive at a non-zero constant contained in I. But then $I = A_n(K)$. □

For reasons of notation, we usually work with left A_n-modules. This is no loss of generality, because one can convert left modules into right modules and vice versa. Before I explain this, let me first show you how to describe left (or right) A_n-modules in very simple terms.

Example 2.3. A left A_n-module is the same thing as a $K[x_1, \dots, x_n]$-module M, together with a family of commuting K-linear endomorphisms $d_1, \dots, d_n \in \operatorname{End}_K(M)$, subject to the condition that

$$d_i(x_j m) - x_j d_i(m) = \delta_{i,j} m$$

for every $m \in M$ and every $i, j = 1, \dots, n$. From this data, we can reconstruct the left A_n-module structure using the generators and relations for the Weyl algebra. Indeed, if we define $\partial_j m = d_j(m)$ for $m \in M$, then the condition on $d_1, \dots, d_n$ says exactly that $[\partial_i, \partial_j]$ and $[\partial_i, x_j] - \delta_{i,j}$ act trivially on M, and so we obtain a left A_n-module.

Example 2.4. A right A_n-module is a $K[x_1, \dots, x_n]$-module M, together with a family of commuting K-linear endomorphisms $d_1, \dots, d_n \in \mathrm{End}_K(M)$, such that

$$d_i(x_j m) - x_j d_i(m) = -\delta_{i,j} m$$

for every $m \in M$ and every $i, j = 1, \dots, n$. From this data, we can reconstruct the right A_n-module structure by setting $m\partial_j = d_j(m)$ for $m \in M$. As before, the condition on $d_1, \dots, d_n$ says that $[\partial_i, \partial_j]$ and $[\partial_i, x_j] - \delta_{i,j}$ act trivially on M, and so we obtain a right A_n-module.

Since the only difference in the two descriptions is the minus sign, we can easily convert left A_n-modules into right A_n-modules (and back) by changing the sign.

Example 2.5. Suppose that M is a left A_n-module. Define $d_1, \dots, d_n \in \mathrm{End}_K(M)$ by setting $d_i(m) = -\partial_i m$ for $m \in M$. The sign change means that

$$d_i(x_j m) - x_j d_i(m) = -\partial_i(x_j m) + x_j \partial_i m = -[\partial_i, x_j] = -\delta_{i,j} m,$$

and so this defines a right A_n-module structure on M. Concretely, a differential operator $P = \sum c_{\alpha,\beta} x^\alpha \partial^\beta$ now acts on an element $m \in M$ as

$$mP = \sigma(P)m,$$

where $\sigma(P) = \sum (-1)^{|\beta|} c_{\alpha,\beta} x^\alpha \partial^\beta$ and $|\beta| = \beta_1 + \dots + \beta_n$. The resulting involution $\sigma\colon A_n \to A_n$ also swaps the left and right module structure on A_n itself.

Filtrations on algebras. Recall that the order of a partial differential operator $P = \sum c_{\alpha,\beta} x^\alpha \partial^\beta \in A_n(K)$ is the maximal number of partial derivatives that appear in P; in symbols,

$$\mathrm{ord}(P) = \max\left\{ \beta_1 + \dots + \beta_n \;\middle|\; c_{\alpha,\beta} \neq 0 \right\}$$

Because of the relation $[\partial_i, x_i] = \delta_{i,j}$, the commutator between a differential operator of order d and a differential operator of order e is a differential operator of order at most $d + e - 1$. In this sense, the Weyl algebra is only mildly non-commutative. In fact, A_n is an example of a filtered algebra, in the following sense.

Definition 2.6. Let R be a K-algebra, not necessarily commutative. A filtration $F_\bullet = F_\bullet R$ on R is a sequence of linear subspaces

$$\{0\} = F_{-1} \subseteq F_0 \subseteq F_1 \subseteq \dots \subseteq R,$$

such that $F_j \cdot F_k \subseteq F_{j+k}$ and $R = \bigcup F_k$.

In particular, $F_0 R$ is a subalgebra of R, and each $F_k R$ is a left (and right) module over $F_0 R$. In many cases of interest, the $F_k R$ are finitely generated as $F_0 R$-modules.

Example 2.7. The *order filtration* on A_n is defined by

$$F_j^{\mathrm{ord}} A_n = \left\{ P = \sum c_{\alpha,\beta} x^\alpha \partial^\beta \;\middle|\; \mathrm{ord}(P) = |\beta| \leq j \right\}$$

In this case, $F_0^{\mathrm{ord}} A_n = K[x_1, \dots, x_n]$, and each $F_j^{\mathrm{ord}} A_n$ is a finitely generated $K[x_1, \dots, x_n]$-module. Note that we have $F_j^{\mathrm{ord}} \cdot F_k^{\mathrm{ord}} = F_{j+k}^{\mathrm{ord}}$ for every $j, k \geq 0$.

Example 2.8. The *Bernstein filtration* on A_n is defined by

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

In this case, $F_0^B A_n = K$, and each $F_j^B A_n$ is a K-vector space of finite dimension. Note that we have $F_j^B \cdot F_k^B = F_{j+k}^B$ for every $j, k \geq 0$.

The advantage of the Bernstein filtration is that each F_j^B is finite dimensional. The advantage of the order filtration is that it generalizes to the case of $\mathscr{D}$-modules on arbitrary smooth algebraic varieties (whereas the Bernstein filtration only makes sense on affine space).

Definition 2.9. Given a filtration $F_\bullet R$ on a K-algebra R, the *associated graded algebra* is defined as

$$\operatorname{gr}^F R = \bigoplus_{j=0}^{\infty} F_j/F_{j-1}.$$

It inherits a multiplication from R in the natural way: for $r \in F_j$ and $s \in F_k$, the product $(r + F_{j-1}) \cdot (s + F_{k-1}) = rs + F_{j+k-1}$ is well-defined.

For both the order filtration and the Bernstein filtration, the associated graded algebra of A_n is simply the polynomial ring in $2n$ variables. In particular, the associated graded algebra is commutative.

Proposition 2.10. *Let $A_n = A_n(K)$.*

(a) If $F_\bullet A_n$ is the Bernstein filtration, then

$$\operatorname{gr}^F A_n \cong K[x_1, \dots, x_n, \partial_1, \dots, \partial_n],$$

with the usual grading by the total degree in $x_1, \dots, x_n, \partial_1, \dots, \partial_n$.

(b) If $F_\bullet A_n$ is the order filtration, then

$$\operatorname{gr}^F A_n \cong K[x_1, \dots, x_n, \partial_1, \dots, \partial_n],$$

with the grading by the total degree in $\partial_1, \dots, \partial_n$.

Proof. We prove this only for the Bernstein filtration, the other case being similar. From the definition of the Bernstein filtration as

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

it is obvious that $x_1, \dots, x_n, \partial_1, \dots, \partial_n \in F_1$. For clarity, we use $\bar{x}_1, \dots, \bar{x}_n, \bar{\partial}_1, \dots, \bar{\partial}_n$ to denote their images in F_1/F_0. It is also obvious that F_j/F_{j-1} is generated by all monomials of degree j in $\bar{x}_1, \dots, \bar{x}_n, \bar{\partial}_1, \dots, \bar{\partial}_n$. It remains to analyze the relations. Obviously, $\bar{x}_1, \dots, \bar{x}_n$ commute, and $\bar{\partial}_1, \dots, \bar{\partial}_n$ commute. Since

$$\partial_i x_j - x_j \partial_i = [\partial_i, x_j] = \delta_{i,j} \in F_0,$$

we have $\bar{\partial}_i \bar{x}_j - \bar{x}_j \bar{\partial}_i = 0$ as elements of F_2/F_1. Therefore, all $2n$ elements commute with each other; as there are no further relations, we obtain the desired isomorphism with the polynomial ring. □

Filtrations on A_n-modules. For the time being, we only consider left A_n-modules. Let $F_\bullet A_n$ be either the Bernstein filtration or the order filtration.

Definition 2.11. Let M be a left A_n-module. A *compatible filtration* $F_\bullet M$ on M is a sequence of linear subspaces

$$\{0\} F_{-1} M \subseteq F_0 M \subseteq F_1 M \subseteq \dots \subseteq M,$$

with $F_j A_n \cdot F_k M \subseteq F_{j+k} M$ and $M = \bigcup F_k M$, such that each $F_k M$ is finitely generated as an $F_0 A_n$-module.

Given a compatible filtration on M, one forms the associated graded module

$$\operatorname{gr}^F M = \bigoplus_{k=0}^{\infty} F_k M / F_{k-1} M,$$

which again inherits the structure of a graded module over $\operatorname{gr}^F A_n$ by defining $(r + F_{j-1} A_n) \cdot (m + F_{k-1} M) = rm + F_{j+k-1} M$. Since $\operatorname{gr}^F A_n$ is a polynomial ring in $2n$-variables, this puts us back in the world of commutative algebra. At least for finitely generated modules, one can use this device to transfer properties of modules over the polynomial ring to modules over the Weyl algebra.

Definition 2.12. A compatible filtration $F_\bullet M$ is called *good* if $\mathrm{gr}^F M$ is finitely generated over $\mathrm{gr}^F A_n$.

The following proposition gives a useful necessary and sufficient criterion for a filtration to be good.

Proposition 2.13. *Let M be a left A_n-module. A compatible filtration $F_\bullet M$ is good if, and only if, there exists $j_0 \geq 0$ such that $F_i A_n \cdot F_j M = F_{i+j} M$ for every $i \geq 0$ and every $j \geq j_0$.*

Proof. To simplify the notation, we put

$$\mathrm{gr}_j^F A_n = F_j A_n / F_{j-1} A_n \quad \text{and} \quad \mathrm{gr}_k^F M = F_k M / F_{k-1} M.$$

Let us first prove that the condition is sufficient. Taking $j = j_0$ and $i = j - j_0$, we see that $F_j M = F_{j-j_0} A_n \cdot F_{j_0} M$ for every $j \geq j_0$. This implies almost immediately that $\mathrm{gr}^F M$ is generated, over $\mathrm{gr}^F A_n$, by the direct sum of all $\mathrm{gr}_j^F M$ with $j \leq j_0$. Now each $F_j M$ is finitely generated over $F_0 A_n$, which means that $\mathrm{gr}_j^F M$ is finitely generated over $\mathrm{gr}_0^F A_n$. In total, we therefore get a finite number of elements that generate $\mathrm{gr}^F M$ as a $\mathrm{gr}^F A_n$-module.

The more interesting part is to show that the condition is sufficient. Here it is enough to prove the existence of an integer $j_0 \geq 0$ such that $F_j M = F_{j-j_0} A_n \cdot F_{j_0} M$ for every $j \geq j_0$; the general case follows from this by induction on $j \geq j_0$. Since everything is graded, the fact that $\mathrm{gr}^F M$ is finitely generated over $\mathrm{gr}^F A_n$ implies that it can be generated by finitely many homogeneous elements; let j_0 be the maximum of their degrees. For every $j \geq j_0$, we then have

$$\mathrm{gr}_j^F M = \sum_{i=0}^{j_0} \mathrm{gr}_{j-i}^F A_n \cdot \mathrm{gr}_i^F M,$$

which translates into the relation

$$F_j M = F_{j-1} M + \sum_{i=0}^{j_0} F_{j-i} A_n \cdot F_i M = F_{j-1} M + F_{j-j_0} A_n \cdot F_{j_0} M,$$

using the fact that $F_{j-i} A_n = F_{j-j_0} A_n \cdot F_{j_0-i} A_n$. At this point, we can prove the desired equality $F_j M = F_{j-j_0} A_n \cdot F_{j_0} M$ by induction on $j \geq j_0$. □

We can now show that the existence of a good filtration characterizes finitely generated A_n-modules.

Corollary 2.14. *Let M be a left A_n-module. Then M admits a good filtration if, and only if, it is finitely generated over A_n.*

Proof. Suppose that M is generated, over A_n, by finitely many elements $m_1, \dots, m_k$. Then we can define a compatible filtration $F_\bullet M$ by setting

$$F_j M = F_j A_n \cdot m_1 + \dots + F_j A_n \cdot m_k.$$

Note that each $F_j M$ is finitely generated over $F_0 A_n$, due to the fact that $F_j A_n$ is finitely generated over $F_0 A_n$. With this definition, we have $F_j M = F_j A_n \cdot F_0 M$ for every $j \geq 0$, and therefore the filtration is good by Proposition 2.13.

Conversely, suppose that M admits a good filtration $F_\bullet M$. By Proposition 2.13, there is an integer $j_0 \geq 0$ such that $F_j M = F_{j-j_0} A_n \cdot F_{j_0} M$ for every $j \geq j_0$. Since $M = \bigcup F_j M$, and since $F_{j_0} M$ is finitely generated over $F_0 A_n$, it follows pretty directly that M is finitely generated over A_n. □

The following result is useful for comparing different good filtrations.

Corollary 2.15. *Let M be a left A_n-module with a good filtration $F_\bullet M$. Then for every compatible filtration $G_\bullet M$, there exists some $j_1 \geq 0$ such that $F_j M \subseteq G_{j+j_1} M$ for all $j \geq 0$.*

Proof. As before, choose $j_0 \geq 0$ with the property that $F_j M = F_{j-j_0} A_n \cdot F_{j_0} M$ for every $j \geq j_0$. Since $F_{j_0} M$ is finitely generated over the commutative noetherian ring $F_0 A_n$, and since $G_\bullet M$ is an exhaustive filtration of M by finitely generated $F_0 A_n$-modules, there is some $j_1 \geq 0$ such that $F_{j_0} M \subseteq G_{j_1} M$. But then

$$F_j M \subseteq F_{j+j_0} M = F_j A_n \cdot F_{j_0} M \subseteq F_j A_n \cdot G_{j_1} M \subseteq G_{j+j_1} M,$$

as claimed. □

Let us conclude the discussion of good filtrations by proving that the Weyl algebra is left noetherian. Notice how, during the proof, passing to the associated graded algebra/module allows us to transfer the noetherian property from the commutative ring $\mathrm{gr}^F A_n$ to the non-commutative ring A_n.

Proposition 2.16. *Let M be a finitely generated left A_n-module. Then every submodule of M is again finitely generated. In particular, A_n itself is left noetherian.*

Proof. Let $N \subseteq M$ be a left A_n-submodule. Since M is finitely generated, it admits a good filtration $F_\bullet M$. If we define

$$F_j N = N \cap F_j M,$$

then it is easy to see that $F_i A_n \cdot F_j N \subseteq F_{i+j} N$. Moreover, each $F_j N$ is finitely generated over $F_0 A_n$: this follows from the fact that $F_j M$ is finitely generated over $F_0 A_n$ because $F_0 A_n$ is commutative and noetherian. Therefore $F_\bullet N$ is a good filtration. By construction, we have

$$\mathrm{gr}_j^F N \subseteq \mathrm{gr}_j^F M,$$

which says that $\mathrm{gr}^F N$ is a submodule of $\mathrm{gr}^F M$. Since the original filtration was good, $\mathrm{gr}^F M$ is a finitely generated module over the commutative noetherian ring $\mathrm{gr}^F A_n$, and so $\mathrm{gr}^F N$ is also finitely generated over $\mathrm{gr}^F A_n$. This proves that N is finitely generated over A_n. □

Exercises.

Exercise 2.1. Consider the left A_1-module $M = A_1/A_1 x$. As a K-vector space, M is isomorphic to $K[\partial]$. Write down a formula for the resulting A_1-action on $K[\partial]$.

Exercise 2.2. Show that the left A_1-module $A_1/A_1 x^2$ is isomorphic to the direct sum of two copies of $A_1/A_1 x$.

Exercise 2.3. $M = K[x, x^{-1}]$ is a left A_1-module, with the usual differentiation rule $\partial \cdot x^k = k x^{k-1}$. Show that M is generated, as an A_1-module, by x^{-1}. What does the associated graded module for the good filtration $F_j M = F_j A_1 \cdot x^{-1}$ look like?

Lecture 3: February 11

Dimension and multiplicity. We are going to introduce two important invariants of modules over the Weyl algebra, namely dimension and multiplicity. They are defined using good filtrations. For this, we need to work with the Bernstein filtration on A_n, so in today's lecture, $F_\bullet A_n = F_\bullet^B A_n$ will always mean the Bernstein filtration. Recall that each $F_j^B A_n$ has finite dimension over K.

Let M be a finitely generated A_n-module, where $A_n = A_n(K)$ and K is a field. Choose a good filtration $F_\bullet M$ on M, compatible with the Bernstein filtration $F_\bullet A_n$. We saw last time that the existence of such a filtration is equivalent to M being finitely generated. Since $F_0 A_n = K$, each subspace $F_j M$ in the good filtration is a K-vector space of finite dimension. Consider its dimension

$$\dim_K F_j M = \sum_{i=0}^{j} \dim_K F_i M / F_{i-1} M$$

as a function of $j \geq 0$. Here are some examples:

(1) For $M = A_n$ with the Bernstein filtration, we have
$$F_j A_n = \left\{ \sum c_{\alpha,\beta} x^\alpha \partial^\beta \;\middle|\; |\alpha| + |\beta| \leq j \right\}$$
and therefore
$$\dim F_j A_n = \binom{2n+j}{2n} = \frac{1}{(2n)!} j^{2n} + \cdots$$
is a polynomial of degree $2n$ in the variable j, at least for $j \geq 0$.

(2) For $M = K[x_1, \dots, x_n]$, with the usual filtration by degree, we have
$$\dim F_j M = \binom{n+j}{n} = \frac{1}{n!} j^n + \cdots$$
is a polynomial of degree n in the variable j.

(3) Consider $M = A_n / A_n(x_1, \dots, x_n)$, with the filtration induced by the Bernstein filtration on A_n. As a K-vector space, M is isomorphic to $K[\partial_1, \dots, \partial_n]$, and the filtration is just the filtration by degree. So again,
$$\dim F_j M = \binom{n+j}{n} = \frac{1}{n!} j^n + \cdots$$

(4) Consider the A_1-module $M = K[x, x^{-1}]$, with the filtration $F_j M = F_j A_n \cdot x^{-1}$. Clearly, $F_0 M$ is spanned by x^{-1}, and it is easy to see that $F_j M$ is spanned by $x^{j-1}, x^{j-2}, \dots, x^{-j-1}$ for every $j \geq 0$. So
$$\dim F_j M = 2j + 1$$
for $j \geq 0$, which is again a polynomial of degree 1.

In fact, at least for sufficiently large values of j, the function $\dim_K F_j M$ always grows like a polynomial.

Proposition 3.1. *There is a polynomial $\chi(M, F_\bullet M, t) \in \mathbb{Q}[t]$, called the* Hilbert polynomial *of $(M, F_\bullet M)$, with the property that*
$$\dim_K F_j M = \chi(M, F_\bullet M, j)$$
for all sufficiently large values of j.

Proof. The point is that $\operatorname{gr}^F A_n$ is a polynomial ring in $2n$ variables, and so we can use the theory of Hilbert functions for finitely generated modules over the polynomial ring. (This is explained very well in Eisenbud's book *Commutative Algebra.*) Let me sketch the proof. Set $S = \operatorname{gr}^F A_n$, and recall that this is isomorphic to the polynomial ring in $2n$ variables, with the usual grading by degree. The fact that $F_\bullet M$ is a good filtration means that $\operatorname{gr}^F M$ is a finitely generated graded S-module. By Hilbert's syzygy theorem, every finitely generated graded S-module admits a finite resolution by graded free S-modules; the length of such a resolution is at most the number of variables in the polynomial ring, so $2n$ in our case. Choose such a resolution

$$0 \to E_{2n} \to E_{2n-1} \to \cdots \to E_1 \to E_0 \to \operatorname{gr}^F M \to 0.$$

Denoting by $S(q)$ the graded S-module with $S(q)_i = S_{q+i}$, we have

$$E_p = \bigoplus_{q \in \mathbb{N}} S(-q)^{\oplus b_{p,q}}$$

for certain natural numbers $b_{p,q} \in \mathbb{N}$, all but finitely many of which are of course zero. By counting monomials, we have

$$\dim S_i = \binom{i+2n-1}{2n-1}$$

for $i \geq 0$, and so if we take dimensions in the resolution from above, we get

$$\dim F_i M / F_{i-1} M = \sum_{p=0}^{2n} (-1)^p \sum_q b_{p,q} \dim S_{i-q} = \sum_{p=0}^{2n} (-1)^p \sum_q b_{p,q} \binom{i-q+2n-1}{2n-1}.$$

At least for $i \gg 0$, this is a polynomial of degree at most $2n-1$ in the variable i, whose coefficients are rational numbers. It follows that

$$\dim F_j M = \sum_{i=0}^{j} \dim F_i M / F_{i-1} M$$

is a polynomial of degree at most $2n$ in the variable j, at least for $j \gg 0$. $\square$

If $M \neq 0$, then the Hilbert polynomial is not the zero polynomial; let $d \geq 0$ be its degree. The proof shows that $d \leq 2n$. Since $\dim F_j M$ is of course always a non-negative integer, it is not hard to see that the leading coefficient of the polynomial $\chi(M, F_\bullet M, t)$ must be of the form

$$\frac{m}{d!}$$

for some integer $m \geq 1$. (See the exercises.) Both d and m are actually invariants of the module M itself.

Lemma 3.2. *The two numbers d and m only depend on M, but they do not depend on the choice of good filtration on M.*

Proof. Let $\chi_F(t) = \chi(M, F_\bullet M, t)$ be the Hilbert polynomial for the good filtration $F_\bullet M$. Suppose that $G_\bullet M$ is another good filtration, with Hilbert polynomial $\chi_G(t) = \chi(M, G_\bullet M, t)$. By Corollary 2.15, there is an integer $k \geq 0$ such that

$$F_{j-k} M \subseteq G_j M \subseteq F_{j+k} M$$

for every $j \geq 0$. This gives

$$\dim F_{j-k} M \leq \dim G_j M \leq \dim F_{j+k} M,$$

and therefore we obtain the inequality

$$\chi_F(t-k) \leq \chi_G(t) \leq \chi_F(t+k)$$

for the Hilbert polynomials. Since $\chi_F(t \pm k)$ has the same leading term as $\chi_F(t)$, it follows that $\chi_G(t)$ is also a polynomial of degree d with leading coefficient $m/d!$. $\square$

The number $d = d(M)$ is called the *dimension* of the A_n-module M, and the number $m = m(M)$ is called the *multiplicity*. As long as $M \neq 0$, we have $d(M) \geq 0$ and $m(M) \geq 1$. If $M = 0$, we use the convention that $m(M) = 0$. We will see later what the geometric significance of these two numbers is. Going back to the four examples from above, we see that A_n has dimension $2n$ and multiplicity 1; both $K[x_1, \dots, x_n]$ and $A_n/A_n(x_1, \dots, x_n)$ have dimension n and multiplicity 1; and the A_1-module $K[x, x^{-1}]$ has dimension 1 and multiplicity 2.

Let us investigate the behavior of dimension and multiplicity for submodules and quotient modules. Recall that a short exact sequence of A_n-modules

$$0 \to M' \to M \to M'' \to 0$$

means that M' is a submodule of M, and that M'' is isomorphic to the quotient module M/M'. Given a filtration $F_\bullet M$, we can induce filtrations on M' and M'' by setting

$$F_j M' = M' \cap F_j M \quad \text{and} \quad F_j M'' = \operatorname{im}(F_j M \to M'').$$

With this definition, the associated graded modules form a short exact sequence

$$0 \to \operatorname{gr}^F M' \to \operatorname{gr}^F M \to \operatorname{gr}^F M'' \to 0,$$

now in the category of $\operatorname{gr}^F A_n$-modules.

Proposition 3.3. *Let M be a finitely generated A_n-module, and $F_\bullet M$ a good filtration. Suppose that*

$$0 \to M' \to M \to M'' \to 0$$

is a short exact sequence of A_n-modules. Then the induced filtration $F_\bullet M'$ and $F_\bullet M''$ are both good, and

$$0 \to \operatorname{gr}^F M' \to \operatorname{gr}^F M \to \operatorname{gr}^F M'' \to 0$$

is a short exact sequence of finitely generated graded $\operatorname{gr}^F A_n$-modules. Moreover:

(a) *One has $\chi(M, F_\bullet M, t) = \chi(M', F_\bullet M', t) + \chi(M'', F_\bullet M'', t)$.*
(b) *One has $d(M) = \max\{d(M'), d(M'')\}$.*
(c) *If $d(M') = d(M'')$, then $m(M) = m(M') + m(M'')$.*

Proof. The short exact sequence follows from the definition of the filtrations on M' and M''. Since $F_\bullet M$ is a good filtration, $\operatorname{gr}^F M$ is finitely generated over the polynomial ring $\operatorname{gr}^F A_n$. The polynomial ring is commutative and noetherian, and so both the submodule $\operatorname{gr}^F M'$ and the quotient module $\operatorname{gr}^F M''$ are again finitely generated, which means that $F_\bullet M'$ and $F_\bullet M''$ are also good filtrations. Taking dimensions in the short exact sequence, we get the relation

$$\chi(M, F_\bullet M, t) = \chi(M', F_\bullet M', t) + \chi(M'', F_\bullet M'', t)$$

among the three Hilbert polynomials. The other two assertions are obvious consequences. □

Example 3.4. The calculation in the proposition explains for example why the multiplicity of the A_1-module $K[x, x^{-1}]$ should be 2. Indeed, we have a short exact sequence

$$0 \to K[x] \to K[x, x^{-1}] \to K[x, x^{-1}]/K[x] \to 0.$$

The class of x^{-1} generates the quotient module, but since $x \cdot x^{-1} = 1$, it is also annihilated by x, and so the quotient module is actually isomorphic to $A_1/A_1(x)$. Both the submodule and the quotient module have multiplicity 1, and therefore $K[x, x^{-1}]$ must have multiplicity 2.

Bernstein's inequality. In our discussion of Hilbert functions, we have only used properties of the polynomial ring $\operatorname{gr}^F A_n$. Now comes the first place where A_n-modules are genuinely different from modules over the polynomial ring. The following important result is due to Joseph Bernstein.

Theorem 3.5 (Bernstein's inequality)**.** *Let $M \neq 0$ be a finitely generated A_n-module. Then $d(M) \geq n$.*

Choose a filtration $F_\bullet M$, compatible with the Bernstein filtration on A_n; after a shift in the indexing, we can assume that $F_0 M \neq 0$.

Lemma 3.6. *The multiplication map*

$$F_j^B A_n \to \operatorname{Hom}_K(F_j M, F_{2j} M), \quad P \mapsto (m \mapsto Pm),$$

is injective for every $j \geq 0$.

Proof. We argue by induction on $j \geq 0$. For $j = 0$, the statement is clearly true: $F_0^B A_n = K$, and since $F_0 M \neq 0$, the multiplication map $K \to \operatorname{Hom}_K(F_0 M, F_0 M)$ is obviously injective. Now suppose that the result is known for $j - 1 \geq 0$. Assume for the sake of contradiction that there is a nonzero differential operator $P \in F_j^B A_n$ that lies in the kernel of the multiplication map, so that $Pm = 0$ for every $m \in F_j M$. Clearly, P cannot be constant (because $F_j M$ is nonzero), and so P has to contain x_i or ∂_i for some $i = 1, \dots, n$. If x_i appears in P, then by a calculation we did in Lecture 1, the commutator $[P, \partial_i] \in F_{j-1}^B A_n$ is still nonzero. But then

$$[P, \partial_i]m = P(\partial_i m) - \partial_i(Pm) = 0$$

for every $m \in F_{j-1} M$; indeed, both m and $\partial_i m$ belong to $F_j M$, and P annihilates $F_j M$ by assumption. This contradicts the inductive hypothesis. If ∂_i appears in P, then we use the same argument with $[P, x_i]$ instead. $\square$

Now suppose that $F_\bullet M$ is a good filtration, and let $\chi(t) = \chi(M, F_\bullet M, t)$ be the Hilbert polynomial. The lemma gives

$$\dim F_j^B A_n \leq \dim \operatorname{Hom}_K(F_j M, F_{2j} M) = \dim F_j M \cdot \dim F_{2j} M,$$

and therefore

$$\binom{j + 2n}{2n} \leq \chi(j) \cdot \chi(2j)$$

for all sufficiently large values of j. Since $\chi(t)$ is a polynomial of degree $d(M)$, we conclude that $2n \leq 2d(M)$, or $n \leq d(M)$. This proves Bernstein's inequality.

Holonomic modules. Bernstein's inequality suggests the following definition.

Definition 3.7. A finitely generated A_n-module M is called *holonomic* if either $M \neq 0$ and $d(M) = n$, or if $M = 0$.

Holonomic modules are those for which the dimension takes the minimal value allowed by Bernstein's inequality. We also consider the zero module to be holonomic for convenience. In the special case of holonomic modules, Proposition 3.3 has many nice consequences. The following result would be cumbersome to state if we did not consider the zero module to be holonomic.

Corollary 3.8. *Suppose that*

$$0 \to M' \to M \to M'' \to 0$$

is a short exact sequence of A_n-modules. Then M is holonomic if and only if M' and M'' are holonomic. In particular, submodules and quotient modules of holonomic modules are again holonomic.

Proof. This follows from the fact that $d(M) = \max\{d(M'), d(M'')\}$ and Bernstein's inequality. $\square$

Now suppose that M is a nonzero holonomic module, with a certain multiplicity $m(M) \geq 1$. If we have any chain of submodules

$$M_1 \subseteq M_2 \subseteq M_3 \subseteq \cdots \subseteq M_\ell \subseteq M,$$

then each M_j is again holonomic, hence of dimension n. By Proposition 3.3, the multiplicities add, and so

$$m(M) = m(M_1) + m(M/M_1) = m(M_1) + m(M_2/M_1) + \cdots + m(M_\ell/M_{\ell-1}).$$

If the chain is strictly increasing, then each term in the sum is ≥ 1, and so we get $\ell \leq m(M)$. In other words, the length of any strictly increasing (or decreasing) chain of submodules is bounded by $m(M)$.

Corollary 3.9. *Let M be a holonomic A_n-module.*

(a) M is both noetherian and artininian, meaning that every increasing or decreasing chain of submodules stabilizes.

(b) M has finite length, meaning that it admits a finite filtration whose subquotients are simple A_n-modules.

Proof. The first assertion follows from the calculation we just did. For the second assertion, see the exercises. $\square$

We have already seen a few simple examples of holonomic modules; for instance, $K[x_1, \dots, x_n]$ is a holonomic A_n-module, and $K[x, x^{-1}]$ is a holonomic A_1-module. Here is a more interesting class of holonomic A_n-modules.

Proposition 3.10. *Let $p \in K[x_1, \dots, x_n]$ be a nonzero polynomial. Then*

$$M = K[x_1, \dots, x_n, p^{-1}],$$

with the structure of left A_n-module given by formal differentiation, is a holonomic A_n-module.

Unlike the example of $K[x, x^{-1}]$, it is not even obvious that M is finitely generated. Fortunately, we can use the following numerical criterion for holonomicity.

Lemma 3.11. *Let M be a A_n-module, and $F_\bullet M$ a filtration compatible with the Bernstein filtration on A_n. If*

$$\dim_K F_j M \leq \frac{c}{n!} j^n + c_1 (j+1)^{n-1}$$

for some constants $c, c_1 \geq 1$, then M is holonomic and $m(M) \leq c$. In particular, M is finitely generated.

Proof. The idea is to study finitely generated submodules of M. These are easy to construct: simply take any finite number of elements of M and look at the submodule they generate. Let $N \subseteq M$ be any nonzero finitely generated submodule, and $F_\bullet N$ a good filtration of N. The filtration $N \cap F_\bullet M$ is compatible with the Bernstein filtration, but of course not necessarily good. Still, according to Corollary 2.15, there is an integer $k \geq 0$ such that

$$F_j N \subseteq N \cap F_{j+k} M \subseteq F_{j+k} M$$

for every $j \geq 0$. Taking dimensions, we get

$$\dim F_j N \leq \dim F_{j+k} M \leq \frac{c}{n!} (j+k)^n + c_1 (j+k+1)^{n-1},$$

and therefore $d(N) \leq n$. Since $d(N) \geq n$ by Bernstein's inequality, we see that $d(N) = n$, and so N is holonomic. It also follows that $m(N) \leq c$, by looking at the leading terms on both sides. Therefore any finitely generated submodule of M is holonomic and has multiplicity at most c.

This implies now that M itself must be finitely generated, hence holonomic. To see this, choose any nonzero element $m_1 \in M$, and let N_1 be the submodule generated by m_1. If $N_1 = M$, then we are done; otherwise, choose an element $m_2 \in M \setminus N_1$, and let N_2 be the submodule generated by m_1 and m_2. If $N_2 = M$, then we are done; otherwise, choose an element $m_3 \in M \setminus N_2$, and let N_3 be the submodule generated by m_1, m_2, m_3. Continuing in this way, we produce an chain of submodules $N_1 \subset N_2 \subset N_3 \subset \cdots$. Because each N_j is holonomic with $m(N_j) \leq c$, this chain has to stabilize after at most c steps, and so M is in fact generated by at most c elements. In particular, M is holonomic and $m(M) \leq c$. $\square$

Note that the filtration $F_\bullet M$ is not necessarily good. The lemma is quite remarkable: it allows us to prove that M is finitely generated simply by computing the dimensions of $F_j M$.

Now we apply this to study the A_n-module $M = K[x_1, \dots, x_n, p^{-1}]$. The action by A_n is by formal differentiation:

$$\partial_j(fp^{-\ell}) = -\ell f \frac{\partial p}{\partial x_j} p^{-(\ell+1)} + \frac{\partial f}{\partial x_j} p^{-\ell} = \left(-\ell f \frac{\partial p}{\partial x_j} + p \frac{\partial f}{\partial x_j} \right) p^{-(\ell+1)}.$$

Let $m = \deg p$, and consider the filtration

$$F_j M = \left\{ fp^{-\ell} \;\middle|\; \deg f \leq (m+1)\ell \right\}.$$

Each $F_j M$ is a finite-dimensional K-vector space. If $fp^{-\ell} \in F_j M$, then $\deg f \leq (m+1)\ell$, and so $x_j f p^{-\ell}$ and $\partial_j(fp^{-\ell})$ again belong to $F_{j+1}M$ (by the above formula). In other words, the filtration is compatible with the Bernstein filtration on A_n. Lastly, we have $M = \bigcup F_j M$; indeed, given any element $fp^{-\ell} \in M$, we have

$$fp^{-\ell} = (fp^k) p^{-(\ell+k)},$$

and since $\deg(fp^k) = \deg f + km \leq (m+1)(\ell+k)$ for sufficiently large k, the element eventually belongs to $F_{\ell+k}M$. Taking dimensions, we have

$$\dim F_j M = \binom{(m+1)j + n}{n},$$

which is a polynomial of degree n in j with leading coefficient $(m+1)^n/n!$. So the lemma shows that M is holonomic with $m(M) \leq (m+1)^n$.

Exercises.

Exercise 3.1. Suppose that $\chi(t) \in \mathbb{Q}[t]$ has the property that $\chi(j) \in \mathbb{Z}$ for all sufficiently large values of $j \in \mathbb{Z}$. Show that $\chi(t)$ can be written as a linear combination, with integer coefficients, of the polynomials

$$\chi_n(t) = \frac{t(t-1)\cdots(t-n+1)}{n!}$$

for $n \geq 0$. Conclude that the leading coefficient of $\chi(t)$ has the form $m/d!$ for some $m \in \mathbb{Z}$, where d is the degree of $\chi(t)$.

Exercise 3.2. Show that $A_1/A_1 P$ is holonomic for every nonzero $P \in A_1$.

Exercise 3.3. Recall that a (left) A_n-module M is said to be *simple* if it does not have any A_n-submodules besides $\{0\}$ and M. Show that every simple A_n-module is *cyclic*, meaning that it be generated by a single element.

Exercise 3.4. The goal of this exercise is to prove that every holonomic A_n-module is cyclic. This phenomenon is very different from the case of modules over the polynomial ring.

(a) Let M be a nonzero holonomic A_n-module. Show that M has finite length, meaning that it admits a filtration by A_n-submodules whose subquotients are simple modules. Let $\ell \geq 1$ be the length of such a filtration.
(b) Show that the result is true if $\ell = 1$.
(c) If $\ell \geq 2$, let $N \subseteq M$ be a simple submodule, generated by some $m_0 \in N$. By induction, M/N is cyclic, so let $m \in M$ be any element that maps to a generator of M/N. Show that the left ideal $I = \left\{ P \in A_n \;\middle|\; Pm = 0 \right\}$ is nonzero.
(d) Show that there is some $Q \in A_n$ such that IQ is not contained in the left ideal $\left\{ P \in A_n \;\middle|\; Pm_0 = 0 \right\}$. (Hint: A_n is a simple algebra.)
(e) Now choose $P \in I$ such that $PQm_0 \neq 0$. Show that the element $m + Qm_0$ generates M as a left A_n-module.

Lecture 4: February 13

Last time, somebody asked what happens to chains of submodules when the dimension is greater than n. Here is an example to show that there can be infinite descending chains. (Since A_n is noetherian, there are no infinite ascending chains in finitely generated A_n-modules.)

Example 4.1. Consider the chain of submodules

$$A_1 \supset A_1 x \supset A_1 x^2 \supset \cdots$$

All modules in this chain are isomorphic to A_1, and all subquotients are isomorphic to $A_1/A_1 x$. What happens is that, in the short exact sequence

$$0 \to A_1 \xrightarrow{x} A_1 \to A_1/A_1 x \to 0,$$

the first two modules have dimension 2 and multiplicity 1, whereas the third module has dimension 1 and multiplicity 1.

Distributions and polynomials. Today, we are going to look at an application of holonomic A_n-modules to the study of certain integrals. This was in fact one of the reasons why the theory was developed in the first place. For the time being, we take $K = \mathbb{R}$. Let $p \in \mathbb{R}[x_1, \dots, x_n]$ be a nonzero polynomial with the property that $p(x_1, \dots, x_n) \geq 0$ for every $(x_1, \dots, x_n) \in \mathbb{R}^n$. (We can always achieve this by replacing p by its square.)

Let $S(\mathbb{R}^n)$ be the Schwartz space of all rapidly decreasing functions. A complex-valued function $\varphi \in C^\infty(\mathbb{R}^n)$ is rapidly decreasing if the quantity

$$p_{\alpha,\beta}(f) = \sup_{x \in \mathbb{R}^n} \lvert x^\alpha \partial^\beta \varphi(x) \rvert$$

is finite for every pair of multi-indices $\alpha, \beta \in \mathbb{N}^n$. Then $S(\mathbb{R}^n)$ is a topological vector space, with the topology defined by the family of semi-norms $p_{\alpha,\beta}$. A *tempered distribution* T is a continuous linear functional $T \colon S(\mathbb{R}^n) \to \mathbb{C}$.

Now fix a rapidly decreasing function $\varphi \in S(\mathbb{R}^n)$, and consider the integral

$$T_s(\varphi) = \int_{\mathbb{R}^n} p(x)^s \varphi(x) \, d\mu(x),$$

as a function of the complex parameter $s \in \mathbb{C}$. For $\operatorname{Re} s > 0$, the integral makes sense and has a finite value, due to the fact that φ is rapidly decreasing (and p only takes nonnegative real values). Differentiation under the integral sign shows that $T_s(\varphi)$ is actually a holomorphic function of s for $\operatorname{Re} s > 0$.

Example 4.2. The Gamma function

$$\Gamma(s) = \int_0^\infty x^{s-1} e^{-x} \, dx$$

is a typical example of such an integral. The integral only makes sense for $\operatorname{Re} s > 0$, but in fact, $\Gamma(s)$ can be analytically continued to a meromorphic function on $\mathbb{C}$ with simple poles along $\{0, -1, -2, \dots\}$. This is done step by step, using integration by parts. One has

$$\frac{d}{dx}\bigl(x^s e^{-x}\bigr) = s x^{s-1} e^{-x} - x^s e^{-x},$$

and therefore

$$s\Gamma(s) = x^s e^{-x} \Big|_0^\infty + \int_0^\infty x^s e^{-x} \, dx = \Gamma(s+1)$$

for $\operatorname{Re} s > 0$; now the identity $\Gamma(s) = \Gamma(s+1)/s$ provides an extension of the Gamma function to $\operatorname{Re} s > -1$, with a simple pole at $s = 0$.

Now the question is whether $T_s(\varphi)$ can always be extended to a meromorphic function on the entire complex plane. Bernstein discovered that the answer is yes. The reason is that one always has a functional equation of the form

$$D(s)p(x)^{s+1} = b(s)p(x)^s, \tag{4.3}$$

where $b(s) \in \mathbb{R}[s]$ is a monic polynomial, and $D(s) \in A_n\big(\mathbb{R}[s]\big)$ is a differential operator with coefficients in the ring $\mathbb{R}[s]$. This sort of relation gives the desired meromorphic extension, again step by step. Indeed, after substituting into the integral and integrating by parts, we get

$$b(s)T_s(\varphi) = \int_{\mathbb{R}^n} D(s)p(x)^{s+1}\varphi(x)\,d\mu = \int_{\mathbb{R}^n} p(x)^{s+1}\sigma\big(D(s)\big)\varphi(x)\,d\mu,$$

where $\sigma\big(D(s)\big)$ is the differential operator obtained from $D(s)$ by the left-to-right transformation in Lecture 2. (The reason is that each time we integrate by parts to move ∂_j from the first to the second factor, we get an additional minus sign.) The new integral is again holomorphic for $\operatorname{Re} s > -1$, and after dividing by $b(s)$, we obtain a meromorphic extension of $T_s(\varphi)$ to the half plane $\operatorname{Re} s > -1$, possibly with poles along the zero set of $b(s)$. Continuing in this manner, we can extend $T_s(\varphi)$ to a meromorphic function on the entire complex plane, with poles contained in the set

$$\big\{\, s \in \mathbb{C} \;\big|\; b(s+k) = 0 \text{ for some } k \geq 0 \,\big\}.$$

For this reason, we obviously want to choose the polynomial $b(s)$ in (4.3) to be of minimal degree.

Example 4.4. In the case of the Gamma function, we have $p(x) = x$, and the desired relation is simply that $\partial x^{s+1} = (s+1)x^s$.

Bernstein polynomials. Let us now investigate the existence of the relation in (4.3). This works over any field K, and so we relax the assumptions and allow $p \in K[x_1, \dots, x_n]$ to be any nonzero polynomial. Set $m = \deg p$. Since we are going to work algebraically, we let s be an independent variable, and consider the field of rational functions $K(s)$, and the Weyl algebra $A_n\big(K(s)\big)$ with coefficients in $K(s)$. We now endow the $K(s)$-vector space

$$M = K(s)[x_1, \dots, x_n, p^{-1}]$$

with the structure of a left $A_n\big(K(s)\big)$-module, as follows. Multiplication by polynomials with coefficients in $K(s)$ is defined as usual; and

$$\partial_j\big(fp^{-\ell}\big) = \frac{\partial f}{\partial x_j}p^{-\ell} + (s-\ell)f\frac{\partial p}{\partial x_j}p^{-(\ell+1)}.$$

One can check, based on the discussion in Lecture 2, that this defines a left action by the Weyl algebra with coefficients in $K(s)$. The formulas are easier to remember if we introduce a formal symbol p^s, with the property that

$$\partial_j p^s = sp^{-1}\frac{\partial p}{\partial x_j} \cdot p^s,$$

and write elements of Mp^s in the form $fp^{s-\ell}$. Then the formula from above is simply the (formally correct) differentiation rule

$$\partial_j\big(fp^{s-\ell}\big) = \frac{\partial f}{\partial x_j}p^{s-\ell} + (s-\ell)f\frac{\partial p}{\partial x_j}p^{s-(\ell+1)}. \tag{4.5}$$

The same calculation as in Lecture 3 shows that the filtration

$$F_jM = \big\{\, fp^{-\ell} \;\big|\; \deg f \leq (m+1)\ell \,\big\}$$

is compatible with the Bernstein filtration on $A_n\big(K(s)\big)$, and

$$\dim_{K(s)} F_j M = \binom{(m+1)\ell + n}{n}.$$

According to Lemma 3.11, M is therefore a holonomic module, of multiplicity at most $(m+1)^n$.

Now consider, for $k \geq 0$, the submodule $M_k \subseteq M$ generated by p^k; concretely,

$$M_k = A_n\big(K(s)\big) \cdot p^k \subseteq M.$$

Clearly $M_0 \supseteq M_1 \supseteq M_2 \supseteq$, and because M is holonomic, each M_k is holonomic, and the chain has to stabilize after at most $m(M)$ many steps. So there exists some $k \geq 0$ such that $M_{k+1} = M_k$. This means concretely that there is a differential operator $Q(s) \in A_n\big(K(s)\big)$ with the property that $Q(s)p^{k+1} = p^k$. Note that $Q(s)$ has coefficients in the field of rational functions $K(s)$, so there may be denominators. Let $d(s) \in K[s]$ be a nonzero polynomial such that $R(s) = d(s)Q(s)$ has coefficients in $K[s]$. Then we get $R(s)p^{k+1} = d(s)p^k$, which we can write symbolically as

$$R(s)p^{s+k+1} = d(s)p^{s+k}.$$

After replacing s by $s-k$ everywhere (which is compatible with the differentiation rule in (4.5), and therefore okay), we obtain the identity

$$R(s-k)p^{s+1} = d(s-k)p^s,$$

which has the same shape as (4.3). Now let $b(s) \in K[s]$ be the monic polynomial of minimal degree that satisfies a relation of the form

$$D(s)p^{s+1} = b(s)p^s$$

for some differential operator $D(s) \in A_n\big(K[s]\big)$.

Definition 4.6. The polynomial $b(s) \in K[s]$ is called the *Bernstein polynomial* of $p \in K[x_1, \dots, x_n]$, and $D(s) \in A_n\big(K[s]\big)$ is called a *Bernstein operator* for p.

In fact, the set of all polynomials for which such a relation holds is closed under addition and multiplication by elements of $K[s]$, and therefore an ideal in $K[s]$. The Bernstein polynomial is then simply the unique monic generator of this ideal, keeping in mind that $K[s]$ is a principal ideal domain.

Note. The relation $D(s)p = b(s)$ in the module M implies (by induction on the exponent of p in the denominator) that $M_0 = M$, in the notation from above. Here is another way of looking at the Bernstein polynomial: Multiplication by s defines an endomorphism of the quotient module

$$M_0/M_1 = M/A_n\big(K(s)\big)p,$$

and $b(s)$ is the minimal polynomial for this endomorphism.

Let us finish by computing a few examples of Bernstein polynomials.

Example 4.7. In one variable, let $p = x$. Here $\partial x^{s+1} = (s+1)x^s$, and so we have $b(s) = s+1$ and $D(s) = \partial$.

Example 4.8. Still in one variable, take $p = x^2$. Now $\partial p^{s+1} = (s+1)2xp^s$, and

$$\partial^2 p^{s+1} = (s+1)\big(2p^s + 4x^2 sp^{s-1}\big) = (s+1)(2p^s + 4sp^s) = (s+1)(4s+2)p^s,$$

and therefore $b(s) = (s+1)(s+\frac{1}{2})$.

Example 4.9. The previous example generalizes to $p = x^m$; after applying ∂^m, one finds that $b(s) = (s+1)(s+\frac{m-1}{m})\cdots(s+\frac{1}{m})$.

Example 4.10. In n variables, we can take $p = x_1^{m_1} \cdots x_n^{m_n}$, and after applying the differential operator $\partial_1^{m_1} \cdots \partial_n^{m_n}$, we get

$$b(s) = \prod_{j=1}^{n} \prod_{k=1}^{m_j} \left(s - \frac{k}{m_j} \right)$$

Example 4.11. Another case that can be computed by hand is $p = x_1^2 + \cdots + x_n^2$. Here we again have

$$\partial_j^2 p^{s+1} = (s+1)\big(2p^s + 4x_j^2 s p^{s-1}\big)$$

by the calculation in the second example, and therefore

$$(\partial_1^2 + \cdots + \partial_n^2) p^{s+1} = (s+1)(2n + 4s) p^s.$$

So the Bernstein polynomial in this case is $b(s) = (s+1)(s + \frac{n}{2})$.

These examples suggest that $s = -1$ is always a root of the Bernstein polynomial. It can be proved (using resolution of singularities) that all roots of the Berstein polynomial are negative rational numbers. In general, the Bernstein polynomial can be found using computer algebra systems (such as *Macaulay 2*); except when p is homogeneous, the shape of the Bernstein operator $D(s)$ is not easy to guess in advance, however. Here is a more complicated example for algebraic geometers.

Example 4.12. Consider the polynomial $p = x_1^2 + x_2^3$; this has a so-called cusp singularity at the origin. One can show that

$$\left(\frac{1}{27} \partial_2^3 + \frac{x_2}{6} \partial_1^2 \partial_2 + \frac{x_1}{8} \partial_1^3 \right) p^{s+1} = (s + \tfrac{5}{6})(s+1)(s + \tfrac{7}{6}) p^s,$$

and so the Bernstein polynomial is $b(s) = (s + \frac{5}{6})(s+1)(s + \frac{7}{6})$.

The Bernstein polynomial is of interest in the study of hypersurface singularities. Indeed, the zero set of the polynomial p defines a hypersurface in affine space, to use the terminology from algebraic geometry, and many invariants of its singularities are related to the roots of the Bernsteint polynomial. For example, the largest root of the Bernstein polynomial is the so-called "log canonical threshold" of p.

Lecture 5: February 18

Basic facts about algebraic geometry. The goal of today's class is to give a geometric interpretation for the dimension $d(M)$ from last time. Suppose for the time being that K is an algebraically closed field (such as $\mathbb{C}$). We can then think of the polynomial ring $K[x_1, \dots, x_n]$ as being the ring of algebraic functions on the affine space K^n. If $A_n = A_n(K)$ is the Weyl algebra, and $F_\bullet A_n$ is either the Bernstein filtration or the degree filtration, then $\operatorname{gr}^F A_n \cong K[x_1, \dots, x_n, \xi_1, \dots, \xi_n]$, where ξ_j is the class of ∂_j. We can think of this polynomial ring in $2n$ variables as the ring of algebraic functions on $K^{2n} = K^n \times K^n$, viewed as the cotangent bundle of K^n. The additional variables $\xi_1, \dots, \xi_n$, are linear functions on the fibers of the cotangent bundle. We will see below that $d(M)$ can be interpreted as the "dimension" of a certain subset of K^{2n}, called the *characteristic variety* of M.

Since algebraic geometry language will be useful for this, we start with a brief review of the basic correspondence between closed algebraic subsets of K^n and ideals in the polynomial ring $K[x_1, \dots, x_n]$. To any ideal $I \subseteq K[x_1, \dots, x_n]$, we can associate a closed subset

$$Z(I) = \left\{ (a_1, \dots, a_n) \in K^n \;\middle|\; f(a_1, \dots, a_n) = 0 \text{ for every } f \in I \right\}$$

Since the polynomial ring is noetherian, every ideal is finitely generated, and so every closed subset of this type can in fact be defined by finitely many polynomial equations. Conversely, to a closed subset $Z \subseteq K^n$ defined by polynomial equations, we can associate the ideal

$$I_Z = \left\{ f \in K[x_1, \dots, x_n] \;\middle|\; f(a_1, \dots, a_n) = 0 \text{ for every } (a_1, \dots, a_n) \in Z \right\}$$

of all polynomials that vanish on Z. If $f^m \in I_Z$ for some $m \geq 1$, then of course also $f \in I_Z$ (because K is a field), and so I_Z is always a radical ideal. Here the radical of an ideal I is defined as

$$\sqrt{I} = \left\{ f \in K[x_1, \dots, x_n] \;\middle|\; f^m \in I \text{ for some } m \geq 1 \right\},$$

and an ideal is called a *radical ideal* if $I = \sqrt{I}$. One can show that

$$Z(I_Z) = Z \quad \text{and} \quad I_{Z(I)} = \sqrt{I}.$$

The second assertion is usually called the *Nullstellensatz*. One can summarize this by saying that $I \mapsto Z(I)$ and $Z \mapsto I_Z$ sets up a one-to-one correspondence

$$\bigl(\text{closed algebraic subsets of } K^n\bigr) \longleftrightarrow \bigl(\text{radical ideals in } K[x_1, \dots, x_n]\bigr)$$

This correspondence reverses the order, meaning that $I_1 \subseteq I_2$ iff $Z(I_2) \subseteq Z(I_1)$. The quotient ring $K[x_1, \dots, x_n]/I_Z$ can be viewed as the ring of algebraic functions on the algebraic variety Z, where a polynomial determines a function on Z by restriction (and I_Z is the ideal of functions whose restriction to Z is zero).

Since K is algebraically closed, every maximal ideal in $K[x_1, \dots, x_n]$ is of the form $(x_1 - a_1, \dots, x_n - a_n)$ for some $(a_1, \dots, a_n) \in K^n$, and so under the above correspondence, maximal ideals in the polynomial ring correspond to points of K^n. More generally, prime ideals correspond to irreducible algebraic subsets, where irreducible means that the set cannot be written as a union of two strictly smaller algebraic sets. One can define the *dimension* of a closed algebraic subset $Z \subseteq K^n$ in two equivalent ways: geometrically, as the length of the longest strictly decreasing chain of irreducible closed algebraic subsets

$$Z \supseteq Z_0 \supset Z_1 \supset \cdots \supset Z_d$$

contained in Z; algebraically, as the length of the longest strictly increasing chain of prime ideals

$$I_Z \subseteq P_0 \subset P_1 \subset \cdots \subset P_d$$

containing I_Z. This notion of dimension is known as the *Krull dimension*, and is denoted by $\dim Z$. The geometric picture of the chain is that Z_0 has dimension d, Z_1 has dimension $d-1$, and so on, down to Z_d, which has dimension 0 (and hence is a point). Since ideals in $K[x_1, \dots, x_n]$ containing I_Z are in one-to-one correspondence with ideals in the quotient ring $K[x_1, \dots, x_n]/I_Z$, one also has

$$\dim Z = \dim\bigl(K[x_1, \dots, x_n]/I_Z\bigr),$$

where the dimension $\dim R$ of a commutative ring R is by definition the length of the longest strictly increasing chain of prime ideals in R. The polynomial ring $K[x_1, \dots, x_n]$ has dimension n, of course.

We shall also need the notion of the support of a module. Let M be a finitely generated module over $K[x_1, \dots, x_n]$. Then

$$\operatorname{Supp} M \subseteq K^{2n}$$

is the set of all points $(a_1, \dots, a_n) \in K^n$ such that the localization of M at the maximal ideal $(x_1 - a_1, \dots, x_n - a_n)$ is nontrivial. The geometric picture is that M corresponds to a (coherent) sheaf on K^n, and the support of M is the set of points where the stalk of this sheaf is nontrivial. (In other words, the complement of $\operatorname{Supp} M$ is the largest open set on which the sheaf is trivial.) The support of M is a closed algebraic subset, defined by the annihilator ideal

$$\operatorname{Ann} M = \operatorname{Ann}_{K[x_1,\dots,x_n]} M = \bigl\{\, f \in K[x_1, \dots, x_n] \bigm| fm = 0 \text{ for every } m \in M \,\bigr\}.$$

We have $\dim \operatorname{Supp} M = \dim K[x_1, \dots, x_n]/\operatorname{Ann} M$.

Characteristic varieties. Now we return to modules over the Weyl algebra. Let M be a finitely generated left A_n-module. If we choose a good filtration $F_\bullet M$, compatible with the Bernstein filtration on A_n, then the associated graded module $\operatorname{gr}^F M$ is finitely generated over $\operatorname{gr}^F A_n$, the polynomial ring in $2n$ variables. One of the basic facts about Hilbert polynomials is that the degree $d(M)$ of the Hilbert polynomial of $\operatorname{gr}^F M$ is equal to the dimension of $\operatorname{Supp} \operatorname{gr}^F M$; in symbols,

$$d^B(M) = \dim \operatorname{Supp}(\operatorname{gr}^F M) = \dim \operatorname{gr}^F A_n / \operatorname{Ann}(\operatorname{gr}^F M).$$

I have added the superscript B to indicate that this notion of dimension is related to the Bernstein filtration on A_n. We would now like to have an analogous definition for the degree filtration on the Weyl algebra, since that is the case that generalizes to arbitrary $\mathscr{D}$-modules.

From now on, we use the notation $F_\bullet A_n$ for the filtration by the degree of differential operators. Let M be a finitely generated left A_n-module, and choose a good filtration $F_\bullet M$ compatible with the degree filtration on A_n. We define

$$I(M, F_\bullet) = \operatorname{Ann}_{\operatorname{gr}^F A_n}(\operatorname{gr}^F M)$$

as the annihilator of $\operatorname{gr}^F M$, and use the notation

$$J(M) = \sqrt{I(M, F_\bullet M)}$$

for the radical ideal. We will see in a moment that $J(M)$ only depends on M, but not on the particular good filtration chosen, justifying the notation. As we said earlier, the closed subset of K^{2n} corresponding to the radical ideal $J(M)$ is the support of the module $\operatorname{gr}^F M$.

Definition 5.1. The *characteristic variety* $\operatorname{Ch}(M)$ is the closed algebraic subset of K^{2n} corresponding to the radical ideal $J(M)$. Let

$$d^{\deg}(M) = \dim \operatorname{Ch}(M) = \dim\bigl(\operatorname{gr}^F A_n / J(M)\bigr)$$

be the dimension of the characteristic variety.

Examples show that the ideal $I(M, F_\bullet M)$ depends on the filtration. Nevertheless, the radical ideal $J(M)$ and the characteristic variety $\mathrm{Ch}(M)$ only depend on M.

Proposition 5.2. *The ideal $J(M)$ only depends on M, but not on the choice of good filtration $F_\bullet M$. The same is therefore true for $\mathrm{Ch}(M)$.*

Proof. We first need to describe the annihilator of $\mathrm{gr}^F M$ more concretely. For a differential operator $P \in F_k A_n$ of order exactly k, we denote by $[P]$ its image in $\mathrm{gr}_k^F A_n$; this is usually called the (principal) symbol of P. Likewise, if $m \in F_j M$, we write $[m] \in \mathrm{gr}_j^F M$ for its image in the associated graded module. The module structure on $\mathrm{gr}^F M$ is then defined by setting

$$[P] \cdot [m] = [Pm] \in \mathrm{gr}_{k+j}^F M$$

for $[P] \in \mathrm{gr}_k^F A_n$ and $[m] \in \mathrm{gr}_j^F M$. Thus $[P] \cdot [m] = 0$ means that $Pm \in F_{k+j-1} M$ (but it does not mean that $Pm = 0$). Since $\mathrm{gr}^F M$ is a graded module, the annihilator ideal $\mathrm{Ann}(\mathrm{gr}^F M)$ is a homogeneous ideal; by what we just said, it is generated by all those homogeneous elements $[P] \in \mathrm{gr}_k^F A_n$ with the property that

$$P \cdot F_j M \subseteq F_{k+j-1} M$$

for every $j \geq 0$. The radical ideal $\sqrt{I(M, F_\bullet M)}$ is therefore generated by those homogeneous elements $[P] \in \mathrm{gr}_k^F A_n$ such that, for some $m \geq 1$, one has

$$P^m \cdot F_j M \subseteq F_{mk+j-1} M \tag{5.3}$$

for every $j \geq 0$.

Now let $G_\bullet M$ be another good filtration. By Corollary 2.15, the two good filtrations are comparable, and so there is some $j_0 \geq 0$ such that

$$F_j M \subseteq G_{j+j_0} M \quad \text{and} \quad G_j M \subseteq F_{j+j_0} M$$

for every $j \geq 0$. Suppose that $[P] \in \mathrm{gr}_k^F A_n$ belongs to the radical of $I(M, F_\bullet M)$, hence that we have (5.3) for some $m \geq 1$. Let $\ell \geq 1$ be any integer. We have

$$P^{\ell m} \cdot G_j M \subseteq P^{\ell m} \cdot F_{j+j_0} M \subseteq F_{\ell mk+j+j_0-\ell} M \subseteq G_{\ell mk+j+2j_0-\ell} M.$$

If we take $\ell = 2j_0 + 1$ and $m' = \ell m$, then we have

$$P^{m'} \cdot G_j M \subseteq G_{m'k+j-1} M$$

for every $j \geq 0$, and so P belongs to the radical of $I(M, G_\bullet M)$. Since the situation is symmetric, we conclude that $\sqrt{I(M, G_\bullet)} = \sqrt{I(M, F_\bullet M)}$, and hence that $J(M)$ is independent of the choice of good filtration. □

Example 5.4. One can tell from the characteristic variety whether or not a finitely generated A_n-module M is actually finitely generated over the polynomial ring $K[x_1, \dots, x_n]$. Suppose that M is finitely generated over $K[x_1, \dots, x_n]$. Then setting $F_{-1} M = \{0\}$ and $F_j M = M$ for $j \geq 0$ defines a good filtration, and since $\mathrm{gr}_j^F M = 0$ for $j \neq 0$, every element in $\mathrm{gr}^F A_n$ of strictly positive degree annihilates $\mathrm{gr}^F M$. This means that $\mathrm{Ch}(M)$ is defined by the ideal $(\xi_1, \dots, \xi_n)$ in the polynomial ring $\mathrm{gr}^F A_n = K[x_1, \dots, x_n, \xi_1, \dots, \xi_n]$; in other words, $\mathrm{Ch}(M)$ is the "zero section".

Conversely, if $\mathrm{Ch}(M)$ is the zero section, then M is actually finitely generated over $K[x_1, \dots, x_n]$. Here is the reason. Choose a good filtration $F_\bullet M$, so that $\mathrm{gr}^F M$ is finitely generated over $\mathrm{gr}^F A_n = K[x_1, \dots, x_n, \xi_1, \dots, \xi_n]$. By assumption, some power of each ξ_j belongs to the annihilator, which means that $\xi_1^{e_1} \cdots \xi_n^{e_n}$ acts trivially on $\mathrm{gr}^F M$ as long as $e_1 + \cdots + e_n$ is sufficiently large. Thus the finitely many generators of $\mathrm{gr}^F M$ over $\mathrm{gr}^F A_n$, together with their finitely many images under the elements $\xi_1^{e_1} \cdots \xi_n^{e_n}$ for $e \in \mathbb{N}^n$, generate $\mathrm{gr}^F M$ over $K[x_1, \dots, x_n]$. But this implies that M itself is finitely generated over $K[x_1, \dots, x_n]$.

Equality of dimensions. In the next few lectures, we are going to prove that the two notions of dimension (with respect to the Bernstein filtration and with respect to the degree filtration) agree: for any finitely generated A_n-module, one has

$$d^B(M) = d^{\deg}(M).$$

This will tell us in particular that the Bernstein inequality $d(M) \geq n$ also holds with respect to the degree filtration. The geometric interpretation is that the characteristic variety $\operatorname{Ch}(M)$ always has dimension at least n. The strategy for proving this is to relate two kinds of dimension to a third invariant of M, which is of a more homological nature and can be defined without reference to good filtrations. The invariant is defined in terms of the Ext-modules $\operatorname{Ext}_R^j(M,R)$, namely

$$j(M) = \min\bigl\{\, j \geq 0 \bigm| \operatorname{Ext}_R^j(M,R) \neq 0 \,\bigr\}.$$

The precise result that we are going to prove is that

$$d^B(M) = 2n - j(M) = d^{\deg}(M).$$

Let me end with a brief reminder about Ext-modules. Recall that if R is any ring, and if M and N are two left R-modules, we can form the group

$$\operatorname{Hom}_R(M,N)$$

of all left R-linear morphisms from M to R. This defines a contravariant functor $\operatorname{Hom}_R(-,N)$ from left R-modules to groups, and $\operatorname{Ext}_R^j(M,N)$ is by definition the j-th derived functor. Concretely, one computes $\operatorname{Ext}_R^j(M,N)$ by choosing a resolution of M by free left R-modules,

$$\cdots \to L_2 \to L_1 \to L_0 \to M \to 0,$$

and then applying the functor $\operatorname{Hom}_R(-,N)$ to this resolution. Thus $\operatorname{Ext}_R^j(M,N)$ is the j-th cohomology group of the complex

$$0 \to \operatorname{Hom}_R(L_0,N) \to \operatorname{Hom}_R(L_1,N) \to \operatorname{Hom}_R(L_2,N) \to \cdots$$

In particular, $\operatorname{Ext}_R^0(M,N) = \operatorname{Hom}_R(M,N)$. Note that unless R is commutative, $\operatorname{Hom}_R(M,N)$ typically no longer has the structure of a left or right R-module. But in the special case where $N = R$, we can use the right R-module structure on the ring R to endow $\operatorname{Hom}_R(M,R)$ with the structure of a right R-module. Concretely, for $f \in \operatorname{Hom}_R(M,R)$, and for $r \in R$, we define $f \cdot r \in \operatorname{Hom}_R(M,R)$ by the formula

$$(f \cdot r)(x) = f(x)r.$$

Since the multiplication in R is associative, $f \cdot r$ is again left R-linear. Using a resolution as above, it follows that each $\operatorname{Ext}_R^j(M,R)$ is naturally a right R-module. (Similar comments apply if we work with right R-modules.)

Exercises.

Exercise 5.1. Let $M = A_1/A_1(x)$ be the left A_1-module related to the δ-function. Show that the image of $1 \in A_1$ and the image of $\partial \in A_1$ both generate M, but that the two resulting good filtrations $F_\bullet M$ and $G_\bullet M$ give rise to different annihilator ideals: $I(M, F_\bullet M) \neq I(M, G_\bullet M)$.

Exercise 5.2. Let $I \subseteq A_n$ be a left ideal, and let $F_j I = I \cap F_j A_n$ be the induced filtration. Describe the ideal $\operatorname{Ann}(\operatorname{gr}^F I)$ inside $\operatorname{gr}^F A_n$ in concrete terms.

Lecture 6: February 20

General setup. We start working on the proof of the theorem from last time, comparing the two notions of dimension $d^B(M)$ (with respect to the Bernstein filtration) and $d^{\deg}(M)$ (with respect to the degree filtration). In order to make the result more useful, and to simplify the notation, we are going to work in the following more general setting.

Let R be a ring with 1. We assume that R is filtered; as before, this means that R comes with an exhaustive increasing filtration $F_\bullet R$, with

$$\{0\} = F_{-1}R \subseteq F_0 R \subseteq F_1 R \subseteq \cdots,$$

such that $1 \in F_0 R$ and $F_i R \cdot F_j R \subseteq F_{i+j} R$ for all $i, j \geq 0$. This makes $F_0 R$ a subring of R. We define $S = \operatorname{gr}^F R$ to be the associated graded ring, with $S_j = F_j R/F_{j-1} R$, and with the product defined by $(r + F_i R) \cdot (r' + F_j R) = (rr' + F_{i+j} R)$; note that $F_0 R = S_0$ is also a subring of S. Generalizing from what happens in the case $R = A_n$, we make the following two assumptions about S:

(A) S is a commutative noetherian ring.
(B) S is regular of dimension $\dim S = 2n$.

As in Lecture 2, the assumption (A) implies that R is left noetherian; moreover, the subring $F_0 R = S_0$ is also commutative and noetherian. The condition in (B) means concretely that for every maximal ideal $\mathfrak{m} \subseteq S$, the localization $S_\mathfrak{m}$ is a regular local ring of dimension $2n$, in the sense that

$$\dim_{S/\mathfrak{m}} \mathfrak{m}/\mathfrak{m}^2 = \dim S_\mathfrak{m} = 2n.$$

This implies that every finitely generated $S_\mathfrak{m}$-module has a free resolution of length at most $2n$; in fact, by a theorem of Serre, the two things are equivalent to each other. The geometric meaning of the condition in (B) is of course that the scheme $\operatorname{Spec} S$ is nonsingular of dimension $2n$.

Example 6.1. Take $R = A_n$, either with the Bernstein filtration $F^B_\bullet A_n$ or the degree filtration $F^{\deg}_\bullet A_n$. In both cases, S is the polynomial ring in $2n$ variables.

Now let M be a finitely generated left R-module. As in Lecture 3, we have the notion of a *compatible filtration* $F_\bullet M$. Recall that this means that $F_\bullet M$ is an exhaustive increasing filtration of M, such that $F_i R \cdot F_j M \subseteq F_{i+j} M$ for every $i, j \geq 0$, and such that each $F_j M$ is finitely generated over the commutative ring $F_0 R$. As before, the filtration is called *good* if the associated graded module $\operatorname{gr}^F M$ is finitely generated over $S = \operatorname{gr}^F R$. Every finitely generated R-module has a good filtration. As in the case of A_n, one shows that the ideal

$$J(M) = \sqrt{\operatorname{Ann}_S(\operatorname{gr}^F M)}$$

is independent of the choice of good filtration $F_\bullet M$. It is easy to see that a prime ideal $P \subseteq S$ contains $J(M)$ if and only if the localized module $M_P = S_P \otimes_S M$ is nonzero. The geometric interpretation is that the finitely generated S-module $\operatorname{gr}^F M$ defines a coherent sheaf on the scheme $\operatorname{Spec} S$, and the closed subscheme defined by the ideal $J(M)$ is the support of this sheaf.

Definition 6.2. Let M be a finitely generated left R-module. We set

$$d(M) = \dim S/J(M) = \dim \operatorname{Supp}(\operatorname{gr}^F M)$$
$$j(M) = \min\{\, j \geq 0 \mid \operatorname{Ext}^j_R(M, R) \neq 0 \,\}$$

The theorem I stated last time holds in this generality.

Theorem 6.3. *Let $(R, F_\bullet R)$ be a filtered ring satisfying the two conditions in (A) and (B). Then one has*

$$d(M) + j(M) = \dim S$$

for every finitely generated left R-module M.

This immediately implies the result I stated last time. Take $R = A_n$, and suppose that M is a finitely generated left A_n-module. The definition of the invariant $j(M)$ does not mention any filtrations, and so it is the same no matter what filtration on R we consider. If we take $F_\bullet R = F_\bullet^B A_n$, we get

$$d^B(M) + j(M) = 2n,$$

and if we take $F_\bullet R = F_\bullet^{\deg} A_n$, we get

$$d^{\deg}(M) + j(M) = 2n.$$

The two equations together give us the desired equality $d^B(M) = d^{\deg}(M)$.

The commutative case. The proof of Theorem 6.3 is going to take some time. Let us first consider what happens in the commutative case. In the general setting from above, R is of course allowed to be commutative; but to avoid any confusion, let me stick to the notation S for the commutative noetherian ring.

Proposition 6.4. *Let S be a commutative noetherian ring, regular of dimension $2n$. For any finitely generated S-module M, set $J(M) = \sqrt{\operatorname{Ann}_S M}$ and define*

$$d(M) = \dim S/J(M) \quad \text{and} \quad j(M) = \min\bigl\{\, j \geq 0 \bigm| \operatorname{Ext}_S^j(M, S) \neq 0 \,\bigr\}$$

Then the following is true:

(a) *If $\operatorname{Ext}_S^j(M, S) \neq 0$, then $2n - d(M) \leq j \leq 2n$.*
(b) *One has $d\bigl(\operatorname{Ext}_S^j(M, S)\bigr) \leq 2n - j$ for every $j \geq 0$.*
(c) *One has $d\bigl(\operatorname{Ext}_S^{j(M)}(M, S)\bigr) = d(M)$.*
(d) *The identity $d(M) + j(M) = 2n$ holds.*

Proof. Let me try to give at least an idea of the proof (without dotting all the i's). The first step is to reduce to the case where S is a regular local ring. We can test whether or not $\operatorname{Ext}_S^j(M, S)$ is zero by localizing at all maximal ideals of M. Let $\mathfrak{m} \subseteq S$ be any maximal ideal containing $J(M)$; in terms of the scheme $\operatorname{Spec} S$, we are choosing a closed point on the support of M. Then one has

$$S_\mathfrak{m} \otimes_S \operatorname{Ext}_S^j(M, S) \cong \operatorname{Ext}_{S_\mathfrak{m}}^j\bigl(S_\mathfrak{m} \otimes_S M, S_\mathfrak{m}\bigr).$$

After replacing S by its localization, and M by $S_\mathfrak{m} \otimes_S M$, we can therefore assume that S is a regular local ring of dimension $2n$. Geometrically, this means that we are working locally near a point of $\operatorname{Supp} M$.

We prove (a) and (b) by induction on $d = \dim S/J(M) \geq 0$. When $d = 0$, the fact that S is local implies that $J(M) = \mathfrak{m}$. Since M is finitely generated, one has $\mathfrak{m}^\ell M = 0$ for some $\ell \geq 0$. By considering the chain of submodules $M \supseteq \mathfrak{m}M \supseteq \mathfrak{m}^2 M \supseteq \cdots \supseteq \mathfrak{m}^\ell M = \{0\}$ and the long exact sequence for Ext-modules, we reduce to the case where $\mathfrak{m}M = 0$. Now M is finitely generated over the field $S/\mathfrak{m}$, and so we further reduce to the case where $M = S/\mathfrak{m}$ is the residue field of the local ring. Since S is regular, the Koszul complex (for any system of $2n$ generators for the maximal ideal) resolves $S/\mathfrak{m}$; from this resolution, one obtains

$$\operatorname{Ext}_S^j(S/\mathfrak{m}, S) = \begin{cases} S/\mathfrak{m} & \text{if } j = 2n, \\ 0 & \text{if } j \neq 2n. \end{cases}$$

This establishes (a) and (b) in the case $d = 0$. For the inductive step, it suffices (with a little bit of extra work) to consider the case where there is an element $f \in \mathfrak{m}$ that is not a zero-divisor on M. We then have a short exact sequence

$$0 \to M \xrightarrow{f} M \to M/fM \to 0,$$

and $d(M/fM) = d - 1$. The geometric picture is that $\operatorname{Supp} M$ is a closed subset of dimension d, and that the hypersurface defined by f intersects it in a subset of dimension $d - 1$; the S-module M/fM is of course representing the restriction of M to the hypersurface. Define

$$E^j = \operatorname{Ext}^j_S(M, S) \quad \text{and} \quad F^j = \operatorname{Ext}^j_S(M/fM, S).$$

By induction, we have $F^j = 0$ unless $2n - d - 1 \leq j \leq 2n$, and $d(F^j) \leq 2n - j$. The long exact cohomology sequence for Ext-modules gives

$$\cdots \to F^j \to E^j \xrightarrow{f} E^j \to F^{j+1} \to \cdots.$$

If $j \notin \{2n - d, \dots, 2n\}$, then we have $F^j = F^{j+1} = 0$, and so multiplication by f is an isomorphism from E^j to itself. Since E^j is a finitely generated S-module, and $f \in \mathfrak{m}$, this implies $E^j = 0$ by Nakayama's lemma. This proves (a). Also from the exact sequence, E^j/fE^j is isomorphic to a submodule of F^{j+1}, and therefore

$$2n - (j + 1) \geq d(F^{j+1}) \geq d(E^j/fE^j) \geq d(E^j) - 1,$$

which proves (b).

Now we turn to (c). From (a), we get $j(M) \geq 2n - d(M)$. Combined with (b), this gives

$$d(E^j) \leq 2n - j \leq 2n - j(M) \leq d(M),$$

with strict inequality for $j > j(M)$. Assume for the sake of contradiction that $d(E^{j(M)}) < d(M)$. Then $d(E^j) < d(M)$ for every $j \geq 0$. Setting

$$E = \bigoplus_{j=2n-d(M)}^{2n} E^j,$$

this gives $d(E) < d(M)$, and therefore the ideal $J(E)$ must be strictly bigger than $J(M)$. After localizing at an element $f \in J(E) \setminus J(M)$, we achieve that $M \neq 0$ but $\operatorname{Ext}^j_S(M, S) = 0$ for every $j \geq 0$. Now one can show (as an exercise) that this contradicts the fact that M is finitely generated.

It remains to deduce (d). We have already seen that $j(M) \leq 2n - d(M)$. The reverse inequality follows from (c) and (b), because

$$d(M) = d(E^{j(M)}) \leq 2n - j(M).$$

This completes the proof. $\square$

Filtered resolutions. Now we return to the case where M is a finitely generated left R-module. Choose a good filtration $F_\bullet M$. Proposition 6.4, applied to the finitely generated S-module $\operatorname{gr}^F M$, gives

$$d(\operatorname{gr}^F M) + j(\operatorname{gr}^F M) = 2n.$$

Obviously, we have $J(M) = \sqrt{\operatorname{Ann}_S(\operatorname{gr}^F M)} = J(\operatorname{gr}^F M)$, and therefore

$$d(M) = \dim S/J(M) = d(\operatorname{gr}^F M).$$

The identity $d(M) + j(M) = 2n$ in Theorem 6.3 is therefore equivalent to

$$j(M) = j(\operatorname{gr}^F M).$$

In order to prove the theorem, we therefore need to understand the relationship between $\operatorname{Ext}^j_R(M, R)$ and $\operatorname{Ext}^j_{\operatorname{gr}^F M}(\operatorname{gr}^F M, \operatorname{gr}^F R)$. We will see next time that this involves a spectral sequence. To set it up, we need a resolution of M by free R-modules that takes into account the good filtration $F_\bullet M$.

Proposition 6.5. *Let $(M, F_\bullet M)$ be a finitely generated R-module with a good filtration. Then there exists a free resolution*

$$\cdots \to L_2 \to L_1 \to L_0 \to M \to 0$$

where each $(L_j, F_\bullet L_j)$ is a free R-module with a good filtration, and the differentials in the resolution respect the filtrations. Moreover,

(a) each $\operatorname{gr}^F L_j$ is free over S, of the same rank as L_j, and
(b) the complex of S-modules

$$\cdots \to \operatorname{gr}^F L_2 \to \operatorname{gr}^F L_1 \to \operatorname{gr}^F L_0 \to \operatorname{gr}^F M \to 0$$

is exact.

Proof. For any $e \in \mathbb{Z}$, define $R(e) = R$, but with the good filtration $F_j R(e) = F_{j+e} R$. We are going to construct a resolution in which each L_j is a direct sum of copies of $R(e)$ for various values of e.

We start by building L_0. Since $\operatorname{gr}^F M$ is a finitely generated graded S-module, we can choose homogeneous generators $[m_1], \dots, [m_r]$, of degrees $e_1, \dots, e_r$, meaning that $m_i \in F_{e_i} M$. Then

$$\operatorname{gr}_j^F M = \sum_{i=1}^r S_{j-e_i}[m_i],$$

and an easy argument shows that therefore

$$F_j M = \sum_{i=1}^r F_{j-e_i} R \cdot m_i$$

for every $j \geq 0$. This means exactly that we have a surjective morphism of left R-modules

$$L_0 = \bigoplus_{i=1}^r R(-e_i) \to M$$

compatible with the good filtrations on both terms, such that $\operatorname{gr}^F L_0 \to \operatorname{gr}^F M$ is also surjective. Let M' be the kernel of $L_0 \to M$, with the induced filtration. Then the sequence

$$0 \to \operatorname{gr}^F M' \to \operatorname{gr}^F L_0 \to \operatorname{gr}^F M \to 0$$

is short exact, and since S is noetherian, it follows that $\operatorname{gr}^F M'$ is finitely generated; in other words, M' is finitely generated, and $F_\bullet M'$ is a good filtration. Now apply the same argument to $(M', F_\bullet M')$ to construct L_1, and continue step by step to create the desired free resolution for M. □

Let $\cdots \to L_2 \to L_1 \to L_0$ be a filtered free resolution of M with the properties in the proposition. If we set $L_j^* = \operatorname{Hom}_R(L_j, R)$, then the complex of right R-modules

$$0 \to L_0^* \to L_1^* \to L_2^* \to \cdots$$

can be used to compute $\operatorname{Ext}_R^j(M, R)$. In fact, each term in this complex again has a natural compatible filtration (in the sense of right R-modules).

Definition 6.6. Let L be a finitely generated left R-module with a good filtration $F_\bullet L$. On the right R-module $L^* = \operatorname{Hom}_R(L, R)$, we define

$$F_j L^* = \left\{ \phi \in L^* \;\middle|\; \phi(F_i L) \subseteq F_{i+j} R \text{ for every } i \geq 0 \right\}$$

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

Lemma 6.7. *Suppose that L is a finitely generated left R-module with a good filtration $F_\bullet L$. Then L^* is a finitely generated right R-module, and the filtration $F_\bullet L^*$ is again good.*

Proof. Since L is finitely generated, L^* is clearly again finitely generated. It is easy to see that $F_j L^* \cdot F_k R \subseteq F_{j+k} L^*$. Indeed, if $\phi \in F_j L^*$ and $r \in F_k R$, then we have

$$(\phi \cdot r)(x) = \phi(x) \cdot r$$

and this belongs to $F_{i+j} R \cdot F_k R \subseteq F_{i+(j+k)} R$. We also need to prove that the filtration on L^* is exhaustive. Let $\phi \in \operatorname{Hom}_R(L, R)$ be arbitrary. Since the filtration on L is good, there exists some $j_0 \geq 0$ such that $F_{j+j_0} L = F_j R \cdot F_{j_0} L$ for every $j \geq 0$. Since ϕ is left R-linear, we get

$$\phi(F_{j+j_0} L) \subseteq F_j R \cdot \phi(F_{j_0} L).$$

Now $F_{j_0} L$ is finitely generated over $F_0 R$, and therefore $\phi(F_{j_0} L) \subseteq F_{j_1} R$ for some $j_1 \geq 0$. We now obtain

$$\phi(F_{j+j_0} L) \subseteq F_j R \cdot F_{j_1} R \subseteq F_{j+j_1} R,$$

which is enough to conclude that $\phi \in F_{j_1} L^*$. The proof that the filtration $F_\bullet L^*$ is good is left as an exercise. □

Exercises.

Exercise 6.1. Let S be a local ring, M a finitely generated S-module. Suppose that $\operatorname{Ext}^j_S(M, S) = 0$ for every $j \geq 0$. Prove that $M = 0$.

Exercise 6.2. Let $L = R(\ell)$, as a left R-module. Show that L^* is isomorphic to $R(-\ell)$ as a right R-module (with the filtration defined in class).

Exercise 6.3. Let L be a finitely generated left R-module with a good filtration $F_\bullet L$. Show that the natural morphism

$$\operatorname{gr}^F L^* \to \operatorname{Hom}_S\big(\operatorname{gr}^F L, S\big)$$

is injective, and use this to prove that $\operatorname{gr}^F L^*$ is finitely generated over S.

Lecture 7: February 25

Review from last time. Let me briefly recall where we are at. The general setting is that R is a (non-commutative) ring with 1, endowed with a filtration $F_\bullet R$, such that the associated graded ring $S = \operatorname{gr}^F R$ is commutative and nonsingular of dimension $\dim S = 2n$. The prototypical example is of course $R = A_n(K)$, with S being the polynomial ring in $2n$ variables. Given a finitely generated left R-module M, together with a good filtration $F_\bullet M$, we are trying to compare

$$\operatorname{Ext}_R^j(M, R) \quad \text{and} \quad \operatorname{Ext}_S^j(\operatorname{gr}^F M, S).$$

More precisely, we want to show that the two integers

$$j(M) = \min\left\{ j \geq 0 \;\middle|\; \operatorname{Ext}_R^j(M, R) \neq 0 \right\}$$
$$j(\operatorname{gr}^F M) = \min\left\{ j \geq 0 \;\middle|\; \operatorname{Ext}_S^j(\operatorname{gr}^F M, S) \neq 0 \right\}$$

are always equal to each other. To this end, we had constructed a resolution

$$\cdots \to L_2 \to L_1 \to L_0 \to M \to 0 \tag{7.1}$$

of M by free left R-modules, such that (1) each L_j has a good filtration; (2) the morphisms in the resolution respect the filtrations; (3) the induced complex

$$\cdots \to \operatorname{gr}^F L_2 \to \operatorname{gr}^F L_1 \to \operatorname{gr}^F L_0 \to \operatorname{gr}^F M \to 0 \tag{7.2}$$

is still exact, and therefore gives a resolution of $\operatorname{gr}^F M$ by free S-modules. In fact, each L_j was a direct sum of copies of $R(e)$, for different values of $e \in \mathbb{Z}$, where $R(e) = R$ as a left R-module, but with the good filtration $F_i R(e) = F_{e+i} R$.

Now each $L_j^* = \operatorname{Hom}_R(L_j, R)$ is a right R-module, and the j-th cohomology of the complex of right R-modules

$$0 \to L_0^* \to L_1^* \to L_2^* \to \cdots$$

is equal to $\operatorname{Ext}_R^j(M, R)$. We further showed that each L_j^* again has a good filtration (as a right R-module) – in fact, each L_j^* is again a direct sum of copies of $R(e)$, viewed as a right R-module, by one of the exercises from Lecture 6. One has

$$\operatorname{gr}^F L_j^* \cong \operatorname{Hom}_S\bigl(\operatorname{gr}^F L, S\bigr),$$

and because of the exactness of (7.2), it follows that the j-th cohomology of the complex of graded S-modules

$$0 \to \operatorname{gr}^F L_0^* \to \operatorname{gr}^F L_1^* \to \operatorname{gr}^F L_2^* \to \cdots$$

is equal to $\operatorname{Ext}_S^j(\operatorname{gr}^F M, S)$. So our problem comes down to comparing the cohomology of a filtered complex to the cohomology of the associated graded complex. This can be done using the formalism of spectral sequences.

The spectral sequence of a filtered complex. Generally speaking, a *spectral sequence* is a sequence of complexes

$$\bigl(E_\ell^\bullet, d_\ell\bigr),$$

indexed by $\ell \in \mathbb{N}$. Here each $E_\ell^\bullet$ is a complex of vector spaces, modules, or whatever, and the differentials $d_\ell \colon E_\ell^\bullet \to E_\ell^{\bullet+1}$ are morphisms in the appropriate category. The complex $E_\ell^\bullet$ is often called the "ℓ-th page" of the spectral sequence. What makes a sequence of complexes into a spectral sequence is that each $E_{\ell+1}^\bullet$ is obtained from the previous complex $E_\ell^\bullet$ by taking cohomology:

$$E_{\ell+1}^n \cong H^n\bigl(E_\ell^\bullet\bigr) = \frac{\ker\bigl(d_\ell \colon E_\ell^n \to E_\ell^{n+1}\bigr)}{\operatorname{im}\bigl(d_\ell \colon E_\ell^{n-1} \to E_\ell^n\bigr)}$$

Of course, taking cohomology kills the differentials, and so the new differential $d_{\ell+1}$ has to come from somewhere else.

Typically, there is some quantity that one would like to compute, and the initial page of the spectral sequence is a known (or easily obtained) "approximation" to this quantity. As ℓ gets larger, the approximation gets better and better, and things eventually "converge" to the quantity one is trying to compute. This is of course just a rough description; I am going to make it more precise later on.

In my opinion, the best example for understanding spectral sequences is the spectral sequence of a filtered complex. Suppose then that we have a complex $(K^\bullet, d)$, consisting of vector spaces, modules, or whatever:

$$\cdots \to K^{n-1} \xrightarrow{d} K^n \xrightarrow{d} K^{n+1} \to \cdots$$

We are interested in computing the cohomology

$$H^n(K^\bullet) = \frac{\ker\bigl(d\colon K^n \to K^{n+1}\bigr)}{\operatorname{im}\bigl(d\colon K^{n-1} \to K^n\bigr)}$$

of this complex. Suppose also that the complex is filtered, meaning that each K^n has an increasing filtration $F_\bullet K^n$, possibly infinite in both directions,

$$\cdots \subseteq F_j K^n \subseteq F_{j+1} K^n \subseteq \cdots$$

that is compatible with the differentials in the complex, meaning that $d(F_j K^n) \subseteq F_j K^{n+1}$. We also assume that

$$\bigcup_{j \in \mathbb{Z}} F_j K^n = K^n \quad \text{and} \quad F_j K_n = 0 \text{ for } j \ll 0. \tag{7.3}$$

The compatibility with the differential means that each $F_j K^\bullet$ is a subcomplex of $K^\bullet$, and so we obtain a filtration on the cohomology of $K^\bullet$ by setting

$$F_j H^n(K^\bullet) = \operatorname{im}\Bigl(H^n(F_j K^\bullet) \to H^n(K^\bullet)\Bigr).$$

In fact, it is not hard to see that

$$F_j H^n(K^\bullet) = \frac{F_j K^n \cap \ker d + d(K^{n-1})}{d(K^{n-1})} \cong \frac{F_j K^n \cap \ker d}{F_j K^n \cap d(K^{n-1})}.$$

and hence that that the associated graded object is given by

$$\operatorname{gr}_j^F H^n(K^\bullet) \cong \frac{F_j K^n \cap \ker d}{F_{j-1} K^n \cap \ker d + F_j K^n \cap d(K^{n-1})}.$$

The spectral sequence is going to let us compute not $H^n(K^\bullet)$ itself, but the graded pieces for the above filtration. The first approximation to this – and the starting point for the spectral sequence – is the associated graded complex $\operatorname{gr}^F K^\bullet$, with the induced differential, and terms

$$\cdots \to \operatorname{gr}^F K^{n-1} \xrightarrow{d} \operatorname{gr}^F K^n \xrightarrow{d} \operatorname{gr}^F K^{n+1} \to \cdots$$

Again, it is not hard to show that

$$H^n\bigl(\operatorname{gr}_j^F K^\bullet\bigr) = \frac{\ker\bigl(d\colon \operatorname{gr}_j^F K^n \to \operatorname{gr}_j^F K^{n+1}\bigr)}{\operatorname{im}\bigl(d\colon \operatorname{gr}_j^F K^{n-1} \to \operatorname{gr}_j^F K^n\bigr)} \cong \frac{F_j K^n \cap d^{-1}(F_{j-1} K^{n+1})}{F_{j-1} K^n + d(F_j K^n)}.$$

Note that this is usually not the same as $\operatorname{gr}_j^F H^n(K^\bullet)$.

Example 7.4. Here is a typical example of a filtered complex. Let $(A, \mathfrak{m})$ be a local ring, and suppose that $K^\bullet$ is a complex of free A-modules of finite rank. We can filter each K^n by powers of the maximal ideal,

$$K^n \supseteq \mathfrak{m} K^n \supseteq \mathfrak{m}^2 K^n \supseteq \cdots,$$

which amounts to setting $F_0K^n = K^n$ and $F_{-j}K^n = \mathfrak{m}^j K^n$ for $j \geq 0$. Here the second condition in (7.3) does not hold, but it turns out that one can weaken this to the condition that

$$\bigcap_{j\in\mathbb{Z}} (F_jK^n + L) = L$$

for every submodule $L \subseteq K_n$, which does hold in this example (by Krull's theorem). In particular, the intersection of all F_jK^n equals zero, which makes sense if we think of elements of $\mathfrak{m}^j$ as functions that vanish to order j; going further down in the filtration on K^n therefore means getting closer to zero.

Example 7.5. The long exact sequence in cohomology is a toy example of a spectral sequence. Suppose that we just have one subcomplex $K_0^\bullet \subseteq K^\bullet$. Together with the quotient complex, this makes a short exact sequence

$$0 \to K_0^\bullet \to K^\bullet \to K_1^\bullet \to 0$$

and so we get a long exact sequence in cohomology:

$$\cdots \to H^{n-1}(K_1^\bullet) \to H^n(K_0^\bullet) \to H^n(K^\bullet) \to H^n(K_1^\bullet) \to H^{n+1}(K_0^\bullet) \to \cdots$$

This tells us how the cohomology of $K^\bullet$ is related to the cohomology of the subcomplex and the quotient complex: there are additional maps $H^n(K_1^\bullet) \to H^{n+1}(K_0^\bullet)$, and the two graded pieces of $H^n(K^\bullet)$ are the cokernel respectively kernel of these maps. If the filtration is longer, then the picture is still similar, but it takes more steps to get from the cohomology of the associated graded complex to the associated graded of the cohomology of $K^\bullet$.

As explained above, we may think of elements of F_jK^n as being "close to zero" when $j \ll 0$. The idea behind the spectral sequence is to "approximate" the condition $x \in F_jK^n$ and $dx = 0$ by the weaker condition $dx \in F_{j-\ell}K^n$, and then increasing the value of $\ell \geq 0$. In other words, we are approximating $F_jK^n \cap \ker d$ by the decreasing sequence of submodules $F_jK^n \cap d^{-1}(F_{j-\ell}K^{n+1})$ for $\ell \geq 0$; this makes sense because of the condition in (7.3). With this in mind, we can now give the precise definition of the spectral sequence of a filtered complex.

For each $n, j \in \mathbb{Z}$ and each $\ell \in \mathbb{N}$, we define

$$Z_{\ell,j}^n = F_jK^n \cap d^{-1}(F_{j-\ell}K^{n+1}).$$

In other words, an element $x \in F_jK^n$ belongs to $Z_{\ell,j}^n$ iff $dx \in F_{j-\ell}K^{n+1}$. By construction, the differential $d\colon K^n \to K^{n+1}$ induces a morphism

$$d_\ell\colon Z_{\ell,j}^n \to Z_{\ell,j-\ell}^{n+1}, \quad x \mapsto dx.$$

Similarly, for each $n, j \in \mathbb{Z}$ and each $\ell \in \mathbb{N}$, we define

$$\begin{aligned} B_{\ell,j}^n &= Z_{\ell,j}^n \cap \Big(F_{j-1}K^n + d\big(F_{j+\ell-1}K^{n-1}\big)\Big) \\ &= F_{j-1}K^n \cap d^{-1}(F_{j-\ell}K^{n+1}) + F_jK^n \cap d(F_{j+\ell-1}K^{n-1}) \\ &= Z_{\ell-1,j-1}^n + d(Z_{\ell-1,j+\ell-1}^{n-1}). \end{aligned}$$

We can then form the quotient

$$E_{\ell,j}^n = Z_{\ell,j}^n / B_{\ell,j}^n,$$

and observe that d_ℓ maps $B_{\ell,j}^n$ into $B_{\ell,j-\ell}^{n+1}$, and therefore induces a morphism

$$d_\ell\colon E_{\ell,j}^n \to E_{\ell,j-\ell}^{n+1}$$

with the property that $d_\ell \circ d_\ell = 0$.

To obtain a complex $(E_\ell^\bullet, d_\ell)$, we consider the graded modules

$$E_\ell^n = \bigoplus_{j \in \mathbb{Z}} E_{\ell,j}^n.$$

By construction, the differential $d_\ell \colon E_\ell^n \to E_\ell^{n+1}$ reduces the degree by ℓ.

Example 7.6. For $\ell = 0$, we have

$$Z_{0,j}^n = F_j K^n \quad \text{and} \quad B_{0,j}^n = F_{j-1} K^n,$$

since $d(F_j K^n) \subseteq F_j K^{n+1}$ by assumption. Consequently,

$$E_{0,j}^n = \frac{F_j K_n}{F_{j-1} K_n} = \operatorname{gr}_j^F K^n,$$

with differential d_0 induced by d. Given (7.3), it also makes sense to set

$$Z_{\infty,j}^n = F_j K^n \cap \ker d \quad \text{and} \quad B_{\infty,j}^n = F_{j-1} K^n \cap \ker d + F_j K^n \cap d\bigl(K^{n-1}\bigr),$$

which extends the above notation (formally) to $\ell = \infty$. Then

$$E_{\infty,j}^n = \frac{F_j K^n}{F_{j-1} K^n \cap \ker d + F_j K^n \cap d(K^{n-1})} \cong \operatorname{gr}_j^F H^n(K^\bullet),$$

according to our earlier calculation.

Now let us show that the complexes $(E_\ell^\bullet, d_\ell)$ really form a spectral sequence.

Proposition 7.7. *For each $n, j \in \mathbb{Z}$ and each $\ell \in \mathbb{N}$, one has*

$$E_{\ell+1,j}^n \cong H^n\bigl(E_{\ell,j}^\bullet, d_\ell\bigr).$$

Proof. Set $H_{\ell,j}^n = H^n(E_{\ell,j}^\bullet)$, and recall that this is the cohomology of the complex

$$Z_{\ell,j+\ell}^{n-1}/B_{\ell,j+\ell}^{n-1} \xrightarrow{d_\ell} Z_{\ell,j}^n/B_{\ell,j}^n \xrightarrow{d_\ell} Z_{\ell,j-\ell}^{n+1}/B_{\ell,j-\ell}^{n+1}.$$

We start by defining a function

$$\phi \colon E_{\ell+1,j}^n \to H_{\ell,j}^n.$$

Suppose that $x \in Z_{\ell+1,j}^n$. Then also $x \in Z_{\ell,j}^n$ and

$$d_\ell x = dx \in d\bigl(Z_{\ell+1,j}^n\bigr) \subseteq B_{\ell,j-\ell}^{n+1},$$

and so x defines a class $\phi(x) \in H_{\ell,j}^n$. This class does not depend on the choice of representative, because

$$B_{\ell+1,j}^n = Z_{\ell+1,j}^n \cap \Bigl(B_{\ell,j}^n + d\bigl(Z_{\ell,j+\ell}^{n-1}\bigr)\Bigr)$$

by the lemma below. Indeed, we see that $x \in B_{\ell+1,j}^n$ if and only if its image in $H_{\ell,j}^n$ is zero, and so ϕ is well-defined and injective.

It remains to argue that ϕ is also surjective. Any class in $H_{\ell,j}^n$ can be represented by an element $x \in Z_{\ell,j}^n$ with $d_\ell x \in B_{\ell,j-\ell}^{n+1}$. After unwinding the definitions, this is saying that $x \in F_j K^n$ and $dx \in F_{j-\ell} K^{n+1}$ and

$$dx = dx' + y$$

for some $x' \in F_{j-1} K^n$ with $dx' \in F_{j-\ell} K^{n+1}$ and some $y \in F_{j-\ell-1} K^{n+1}$. Thus

$$x - x' \in F_j K^n \cap d^{-1}\bigl(F_{j-\ell-1} K^{n+1}\bigr) = Z_{\ell+1,j}^n,$$

and after replacing x by $x - x'$, we can assume from the beginning that $x \in Z_{\ell+1,j}^n$. But this means exactly that the given class is in the image of ϕ. □

Lemma 7.8. *One has*

$$B_{\ell+1,j}^n = Z_{\ell+1,j}^n \cap \Bigl(B_{\ell,j}^n + d\bigl(Z_{\ell,j+\ell}^{n-1}\bigr)\Bigr)$$

for every $j, n \in \mathbb{Z}$ and every $\ell \in \mathbb{N}$.

Proof. Unwinding the definitions shows that

$$B^n_{\ell,j} + d\bigl(Z^{n-1}_{\ell,j+\ell}\bigr) = F_{j-1}K^n \cap d^{-1}\bigl(F_{j-\ell}K^{n+1}\bigr) + F_jK^n \cap d\bigl(F_{j+\ell}K^{n-1}\bigr)$$

and so the intersection with $Z^n_{\ell+1,j} = F_jK^n \cap d^{-1}(F_{j-\ell-1}K^{n+1})$ equals

$$F_{j-1}K^n \cap d^{-1}\bigl(F_{j-\ell-1}K^{n+1}\bigr) + F_jK^n \cap d\bigl(F_{j+\ell}K^{n-1}\bigr) = B^n_{\ell+1,j}. \qquad \square$$

In what sense does the spectral sequence of a filtered complex "converge"? Note that the $Z^n_{\ell,j}$ form a decreasing chain of submodules of F_jK^n with

$$Z^n_{\infty,j} = \bigcap_{\ell\in\mathbb{N}} Z^n_{\ell,j}.$$

Proposition 7.7 shows that $E^n_{\ell+1,j}$ is a subquotient of $E^n_{\ell,j}$, but there is in general no natural morphism from one to the other, which means that one cannot take a (direct or inverse) limit in the algebraic sense. Fortunately, what happens almost always in practice is that, for each fixed $j, n \in \mathbb{Z}$, the modules $E^n_{\ell,j}$ stabilize for sufficiently large ℓ. In fact, one has the following necessary and sufficient condition for stabilization, in terms of the filtration on the complex.

Proposition 7.9. *Fix some $n \in \mathbb{Z}$. The differential $d_\ell \colon E^n_\ell \to E^{n+1}_\ell$ vanishes for every $\ell \geq \ell_0$ if, and only if, the filtration satisfies*

$$F_jK^{n+1} \cap d\bigl(K^n\bigr) = F_jK^{n+1} \cap d\bigl(F_{j+\ell_0-1}K^n\bigr)$$

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

Proof. The differential $d_\ell \colon E^n_\ell \to E^{n+1}_\ell$ vanishes for every $\ell \geq \ell_0$ exactly when $d(Z^n_{\ell,j}) \subseteq B^{n+1}_{\ell,j-\ell}$ for every $\ell \geq \ell_0$ and every $j \in \mathbb{Z}$. After replacing j by $j+\ell$, this translates into the condition that

$$\begin{aligned} &F_jK^{n+1} \cap d\bigl(F_{j+\ell}K^n\bigr) \\ &\qquad \subseteq F_{j-1}K^{n+1} \cap d^{-1}\bigl(F_{j-\ell}K^{n+2}\bigr) + F_jK^{n+1} \cap d\bigl(F_{j+\ell-1}K^n\bigr), \end{aligned}$$

or after intersecting with $d(F_{j+\ell}K^n)$,

$$F_jK^{n+1} \cap d\bigl(F_{j+\ell}K^n\bigr) = F_{j-1}K^{n+1} \cap d\bigl(F_{j+\ell}K^n\bigr) + F_jK^{n+1} \cap d\bigl(F_{j+\ell-1}K^n\bigr).$$

Recursively applying this identity (for $\ell \geq \ell_0$), and using the fact that the filtration on K^n is exhaustive, we can rewrite this in the equivalent form

$$F_jK^{n+1} \cap d\bigl(K^n\bigr) = F_{j-1}K^{n+1} \cap d\bigl(K^n\bigr) + F_jK^{n+1} \cap d\bigl(F_{j+\ell_0-1}K^n\bigr).$$

According to (7.3), there is some $j_0 \in \mathbb{Z}$ with $F_{j_0}K^{n+1} = 0$. We now get the desired conclusion by recursively applying the identity above (for $j \geq j_0$). $\square$

Corollary 7.10. *If there is some $\ell_0 \in \mathbb{N}$ with the property that*

$$\begin{aligned} F_jK^{n+1} \cap d\bigl(K^n\bigr) &= F_jK^{n+1} \cap d\bigl(F_{j+\ell_0-1}K^n\bigr) \\ F_jK^n \cap d\bigl(K^{n-1}\bigr) &= F_jK^n \cap d\bigl(F_{j+\ell_0-1}K^{n-1}\bigr) \end{aligned}$$

for every $j \in \mathbb{Z}$, then one has $E^n_{\ell_0} = E^n_\infty$.

For example, one has $E^n_1 = E^n_\infty$ exactly when the differential d is strictly compatible with the filtration, in the sense that $F_jK^n \cap d(K^{n-1}) = d(F_jK^{n-1})$ (and the same condition with $n+1$ in place of n).

Note. I have been using the "natural" indexing for the spectral sequence, where n is the position in the complex $K^\bullet$, and j the degree with respect to the filtration on K^n. For historical reasons, people usually index their spectral sequences differently, and our $E^n_{\ell,j}$ is usually denoted by $E^{-j,n+j}_\ell$. (This looks more natural in the special case of a double complex.)

Application to our problem. Now we return to the case of a finitely generated left R-module M, endowed with a good filtration $F_\bullet M$. If we apply the spectral sequence formalism to the complex of right R-modules

$$0 \to L_0^* \to L_1^* \to L_2^* \to \cdots,$$

with the good filtration $F_\bullet L_j^*$ constructed earlier, we obtain a spectral sequence with $E_0^j = \operatorname{gr}^F L_j^*$ and with differential d_0 induced by the differential in the original complex. It follows that

$$E_1^j = \operatorname{Ext}_S^j\big(\operatorname{gr}^F M, S\big),$$

because the complex in (7.2) is a free resolution of $\operatorname{gr}^F M$. On the other hand, the complex in (7.1) is a free resolution of M, and so we get

$$E_\infty^j = \operatorname{gr}^F \operatorname{Ext}_R^j(M, R).$$

Recall that we are trying to prove the identity $j(M) = j(\operatorname{gr}^F M)$. The first thing we should do is check that the spectral sequence converges, in the sense that each E_ℓ^j stabilizes for $\ell \gg 0$. This is a consequence of the following lemma about good filtrations.

Lemma 7.11. *Let $(K^\bullet, d)$ be a complex of left (or right) R-modules, and suppose that each K^n has a good filtration $F_\bullet K^n$ such that $d(F_j K^n) \subseteq F_j K^{n+1}$ for every $j, n \in \mathbb{Z}$. Then for every $n \in \mathbb{Z}$, there is some $j_0 \in \mathbb{N}$ such that*

$$F_j K^{n+1} \cap d\big(K^n\big) = F_j K^{n+1} \cap d\big(F_{j+j_0} K^n\big).$$

Proof. On the submodule $d(K^n) \subseteq K^{n+1}$, we have two good filtrations, one induced by the good filtration on K^{n+1}, the other by the good filtration on K^n. Let us denote these by

$$F_j d(K^n) = F_j K^{n+1} \cap d(K^n) \quad \text{and} \quad G_j d(K^n) = d(F_j K^n).$$

The first filtration is good because $\operatorname{gr}^F d(K^n)$ is a submodule of the finitely generated S-module $\operatorname{gr}^F K^{n+1}$; the second filtration is good because $\operatorname{gr}^G d(K^n)$ is a quotient module of the finitely generated S-module $\operatorname{gr}^F K^n$. In both cases, we are using the fact that S is noetherian. By Corollary 2.15, there is an integer $j_0 \geq 0$ such that

$$F_j d(K^n) \subseteq G_{j+j_0} d(K^n)$$

for every $j \in \mathbb{Z}$. We get the result by intersecting both sides with $F_j K^{n+1}$. $\square$

Together with the convergence criterion in Corollary 7.10, this shows that $E_\ell^n = E_\infty^n$ for $\ell \gg 0$, and so our spectral sequence does indeed converge. Now recall that

$$E_1^j = \operatorname{Ext}_S^j(\operatorname{gr}^F M, S).$$

We can use the results about E_1^j from Proposition 6.4, plus the spectral sequence, to prove the following theorem.

Theorem 7.12. *Let M be a finitely generated R-module with a good filtration $F_\bullet M$.*

(a) One has $j(\operatorname{gr}^F M) = j(M)$, and thus $\operatorname{Ext}_R^j(M, R) = 0$ for $j < j(\operatorname{gr}^F M)$.
(b) One has $d\big(\operatorname{Ext}_R^j(M, R)\big) \leq 2n - j$ for every $j \geq 0$.
(c) One has $d\big(\operatorname{Ext}_R^{j(M)}(M, R)\big) = 2n - j(M)$.

Proof. To simplify the notation, let me set $j_0 = j(\operatorname{gr}^F M)$, which means that $E_1^j = 0$ for all $j < j_0$. According to Proposition 6.4, we have

$$d(E_1^j) \leq 2n - j$$

for every $j \geq 0$, with equality for $j = j_0$. Here $d(M) = \dim S/J(M)$ is the dimension of the support.

Since $E^j_{\ell+1}$ is a subquotient of E^j_ℓ, it follows that $E^j_\ell = 0$ for $j < j_0$ and $\ell \geq 1$. But $E^j_\infty = E^j_\ell$ for $\ell \gg 0$, and so $E^j_\infty = 0$ for $j < j_0$. Remembering that

$$E^j_\infty = \operatorname{gr}^F \operatorname{Ext}^j_R(M, R),$$

we deduce that $\operatorname{Ext}^j_R(M, R) = 0$ for $j < j_0$, and hence that $j(M) \geq j_0$. This gives us one half of (a), namely

$$j(M) \geq j(\operatorname{gr}^F M).$$

By the same reasoning, $d(E^j_1) \leq 2n - j$ implies that $d(E^j_\infty) \leq 2n - j$, and therefore

$$d\bigl(\operatorname{Ext}^j_R(M, R)\bigr) \leq 2n - j$$

for every $j \geq 0$, which is (b). Lastly, we have $d(E^{j_0}_1) = 2n - j_0$, but $E^{j_0-1}_1 = 0$ and $d(E^{j_0+1}_1) \leq 2n - j_0 - 1$. Therefore

$$E^{j_0}_2 \cong \ker\bigl(d_1 \colon E^{j_0}_1 \to E^{j_0+1}_1\bigr),$$

and since $d(E^{j_0+1}_1) \leq 2n - j_0 - 1$, we see that $d(E^{j_0}_2) = 2n - j_0$. Continuing in this way, we get $d(E^{j_0}_\ell) = 2n - j_0$ for every $\ell \geq 1$, and therefore

$$d\bigl(\operatorname{Ext}^{j_0}_R(M, R)\bigr) = 2n - j_0.$$

In particular, $\operatorname{Ext}^{j_0}_R(M, R) \neq 0$, and so $j_0 \geq j(M)$. This gives us the other inequality

$$j(\operatorname{gr}^F M) \geq j(M),$$

and so (a) and (c) are proved. □

Exercises.

Exercise 7.1. Generalize the proof of Proposition 7.9 to the case where the filtration on each module K^n in the complex satisfies

$$\bigcap_{j \in \mathbb{Z}} (F_j K^n + L) = L$$

for every submodule $L \subseteq K_n$.

Lecture 8: February 27

Holonomic modules and duality. Recall that R is a filtered ring, whose associated graded ring $S = \operatorname{gr}^F R$ is commutative, noetherian, and nonsingular of dimension $\dim S = 2n$. Last time, we proved the following theorem about finitely generated (left or right) R-modules.

Theorem. *Let M be a finitely generated R-module with a good filtration $F_\bullet M$.*

(a) One has $j(\operatorname{gr}^F M) = j(M)$, and thus $\operatorname{Ext}_R^j(M,R) = 0$ for $j < j(\operatorname{gr}^F M)$.
(b) One has $d\big(\operatorname{Ext}_R^j(M,R)\big) \leq 2n - j$ for every $j \geq 0$.
(c) One has $d\big(\operatorname{Ext}_R^{j(M)}(M,R)\big) = 2n - j(M)$.

As I explained before, the fact that $j(\operatorname{gr}^F M) = j(M)$, together with the identity $d(\operatorname{gr}^F M) + j(M) = 2n$, implies that

$$d(M) + j(M) = 2n$$

for every finitely generated R-module.

Example 8.1. In the case of the Weyl algebra A_n, this says that the two notions of dimension (with respect to the Bernstein filtration and the degree filtration) are the same. Since we know from Bernstein's inequality that $d^B(M) \geq n$ for every nonzero finitely generated A_n-module M, it follows that also $d^{\deg}(M) \geq n$.

Let us now assume that Bernstein's inequality holds: Every finitely generated left or right R-module M satisfies $d(M) \geq n$, provided that $M \neq 0$. We saw earlier that this holds when $R = A_n$. An equivalent formulation is that every finitely generated left or right R-module satisfies $j(M) \leq n$, meaning that $\operatorname{Ext}_R^j(M,R) \neq 0$ for some $j \leq n$, again provided that $M \neq 0$. Bernstein's inequality, together with the above theorem, has some remarkable consequences.

Corollary 8.2. *If M is a finitely generated R-module, then $\operatorname{Ext}_R^j(M,R) = 0$ for $j > n$.*

Proof. Let M be a finitely generated left (or right) R-module. Then each $E^j = \operatorname{Ext}_R^j(M,R)$ is a finitely generated right (or left) R-module, and the theorem gives $d(E^j) \leq 2n - j$. But Bernstein's inequality says that $d(E^j) \geq n$ whenever $E^j \neq 0$, and so the conclusion is that $E^j = 0$ for $j > n$. ∎

Note that this is completely false for finitely generated S-modules, where Ext^j can be nonzero in the range $0 \leq j \leq 2n$.

The most interesting R-modules are clearly those for which the dimension $d(M)$ is minimal (or where the quantity $j(M) = 2n - d(M)$ is maximal). By analogy with the case $R = A_n$, we call such modules holonomic.

Definition 8.3. A finitely generated left (or right) R-module M is called *holonomic* if either $M = 0$, or $M \neq 0$ and $d(M) = n$.

An equivalent definition is that M is holonomic if either $M = 0$, or $M \neq 0$ and $j(M) = n$. Since $\operatorname{Ext}_R^j(M,R) = 0$ for $j > n$, we obtain the following alternative characterization of holonomic R-modules.

Corollary 8.4. *A finitely generated R-module M is holonomic if and only if $\operatorname{Ext}_R^j(M,R) = 0$ for every $j \neq n$.*

Given any holonomic left (or right) R-module M, we therefore get another right (or left) R-module

$$M^* = \operatorname{Ext}_R^j(M,R).$$

This is called the *holonomic dual*. Let us investigate the properties of M^*.

Lemma 8.5. *If M is holonomic, then M^* is also holonomic.*

Proof. Since $j(M) = n$, the theorem from last time shows that

$$d(M^*) = d\big(\operatorname{Ext}_R^{j(M)}(M, R)\big) = 2n - j(M) = n.$$

This says that M^* is again holonomic. □

The association $M \mapsto M^*$ is contravariant functor from the category of holonomic left (or right) R-modules to the category of holonomic right (or left) R-modules. Indeed, given a morphism of left R-modules $f\colon M \to N$ between two holonomic R-modules M and N, the functoriality of Ext shows that we have a morphism of right R-modules

$$f^*\colon \operatorname{Ext}_R^n(N, R) \to \operatorname{Ext}_R^n(M, R)$$

in the opposite direction, and it is not hard to see that $(f \circ g)^* = g^* \circ f^*$. As a contravariant functor, the holonomic dual is also exact: if

$$0 \to M_1 \to M_2 \to M_3 \to 0$$

is a short exact sequence of holonomic left (or right) R-modules, then the long exact sequence for $\operatorname{Ext}_R^j(-, R)$ becomes a short exact sequence

$$0 \to \operatorname{Ext}_R^n(M_3, R) \to \operatorname{Ext}_R^n(M_2, R) \to \operatorname{Ext}_R^n(M_1, R) \to 0,$$

due to the vanishing of $\operatorname{Ext}_R^j(M_i, R)$ for $j \neq n$. In other words,

$$0 \to M_3^* \to M_2^* \to M_1^* \to 0$$

is again a short exact sequence.

Proposition 8.6. *We have $M \cong M^{**}$ for every holonomic left (or right) R-module M, and hence the holonomic dual gives an equivalence of categories*

$$(\textit{holonomic left } R\textit{-modules}) \cong (\textit{holonomic right } R\textit{-modules})^{\mathrm{op}}.$$

Proof. Let M be a holonomic left R-module. Choose a free resolution

$$\cdots \to L_2 \to L_1 \to L_0 \to M \to 0$$

by free left R-modules of finite rank. The complex of right R-modules

$$0 \to L_0^* \to L_1^* \to L_2^* \to \cdots$$

is then exact except in degree n, where the cohomology is $M^* = \operatorname{Ext}_R^n(M, R)$. Choose another free resolution

$$\cdots \to K_2 \to K_1 \to K_0 \to M^* \to 0$$

by free right R-modules of finite rank. By a general lemma in homological algebra, there is a morphism of complexes of right R-modules

$$\begin{tikzcd}
\cdots \arrow[r] & K_1 \arrow[r, "d"] \arrow[d, "f_1"] & K_0 \arrow[r] \arrow[d, "f_0"] & 0 \arrow[r] \arrow[d] & \cdots \\
\cdots \arrow[r] & L_{n-1}^* \arrow[r, "d"] & L_n^* \arrow[r, "d"] & L_{n+1}^* \arrow[r] & \cdots
\end{tikzcd}$$

that induces an isomorphism on cohomology. (Such morphisms are called quasi-isomorphisms.) Let me briefly recall the construction. Since M^* is the cohomology in degree n of the complex, we have $M^* = \ker d / \operatorname{im} d$, and so the submodule $\ker d \subseteq L_n^*$ maps onto M^*. Because K_0 is a free R-module, we can find a lifting

$$\begin{tikzcd}
K_0 \arrow[d, dashed] \arrow[dr, two heads] & \\
\ker d \arrow[r, two heads] & M^*
\end{tikzcd}$$

indicated by the dashed arrow, and we denote by $f_0 \colon K_0 \to L_n^*$ the composition. By construction, $d \circ f_0 = 0$, and so the first square in the diagram below commutes:

$$\begin{tikzcd} \cdots \arrow[r] & K_1 \arrow[r, "d"] & K_0 \arrow[r] \arrow[d, "f_0"] & 0 \arrow[r] \arrow[d] & \cdots \\ \cdots \arrow[r] & L_{n-1}^* \arrow[r, "d"] & L_n^* \arrow[r, "d"] & L_{n+1}^* \arrow[r] & \cdots \end{tikzcd}$$

Since the composition $K_1 \to K_0 \to M^*$ is zero, the morphism $f_0 \circ d$ maps K_1 into the submodule $\operatorname{im} d \subseteq \ker d \subseteq L_n^*$. This submodule is the image of L_{n-1}^*, and because K_1 is a free R-module, and so we can again find a lifting

$$\begin{tikzcd} K_1 \arrow[d, dashed, "f_1"] \arrow[dr, two heads] & \\ L_{n-1}^* \arrow[r, two heads, "d"] & \operatorname{im} d \end{tikzcd}$$

which now makes the second square in the diagram commute:

$$\begin{tikzcd} \cdots \arrow[r] & K_1 \arrow[r, "d"] \arrow[d, "f_1"] & K_0 \arrow[r] \arrow[d, "f_0"] & 0 \arrow[r] \arrow[d] & \cdots \\ \cdots \arrow[r] & L_{n-1}^* \arrow[r, "d"] & L_n^* \arrow[r, "d"] & L_{n+1}^* \arrow[r] & \cdots \end{tikzcd}$$

Continuing in this manner produces the desired morphism of complexes. If we now apply the functor $\operatorname{Hom}_R(-, R)$ a second time, we obtain a morphism of complexes of left R-modules

$$\begin{tikzcd} \cdots \arrow[r] & L_{n+1} \arrow[r] \arrow[d] & L_n \arrow[r] \arrow[d] & L_{n-1} \arrow[r] \arrow[d] & \cdots \\ \cdots \arrow[r] & 0 \arrow[r] & K_0^* \arrow[r] & K_1^* \arrow[r] & \cdots \end{tikzcd}$$

One can show that this morphism still induces an isomorphism on cohomology. Now the complex in the first row is a resolution of M, and therefore only has cohomology at L_0. Likewise, because M^* is holonomic, the complex in the second row only has cohomology at K_n^*, where the cohomology is M^{**}. In this way, we obtain a morphism of left R-modules $M \to M^{**}$, which is an isomorphism by the comment above. □

We can use this result to compare the characteristic varieties of M and M^*.

Corollary 8.7. *If M is holonomic, then* $\operatorname{Ch}(M) = \operatorname{Ch}(M^*)$.

Proof. Choose a good filtration $F_\bullet M$ and recall that $\operatorname{Ch}(M)$ is the closed subset of $\operatorname{Spec} S$ defined by the radical of $\operatorname{Ann}_S(\operatorname{gr}^F M)$, or equivalently, the support of the finitely generated S-module $\operatorname{gr}^F M$. The filtered free resolution from last time induces a good filtration on $M^* = \operatorname{Ext}_R^n(M, R)$; in fact, using the spectral sequence from last time, $E_\infty^n = \operatorname{gr}^F \operatorname{Ext}_R^n(M, R) = \operatorname{gr}^F M^*$. Since the spectral sequence converges, E_∞^n is a subquotient of $E_1^n = \operatorname{Ext}_S^n(\operatorname{gr}^F M, S)$, and therefore

$$\operatorname{Ch}(M^*) = \operatorname{Supp} E_\infty^n \subseteq \operatorname{Supp} E_1^n \subseteq \operatorname{Supp}(\operatorname{gr}^F M) = \operatorname{Ch}(M).$$

But then we also have $\operatorname{Ch}(M) = \operatorname{Ch}(M^{**}) \subseteq \operatorname{Ch}(M^*)$, and so the two characteristic varieties are in fact equal. □

The existence of the holonomic dual gives another explanation for the fact that the category of holonomic A_n-modules is both artinian and noetherian. In fact, recall that we showed earlier, using the notion of multiplicity, that every ascending or descending chain of submodules of a holonomic A_n-module M has finite length (bounded by the multiplicity of M). Since the holonomic dual takes ascending chains of submodules of M to descending chains of submodules of M^*, both chain conditions are equivalent in this case. This is again unlike the commutative case.

Exercises.

Exercise 8.1. Let R be a ring with 1. Let $A_\bullet$ and $B_\bullet$ be two complexes of free R-modules of finite rank. Suppose that we have a morphism of complexes

$$\begin{tikzcd}
\cdots \arrow[r] & A_{n-1} \arrow[r] \arrow[d, "f_{n-1}"] & A_n \arrow[r] \arrow[d, "f_n"] & A_{n+1} \arrow[r] \arrow[d, "f_{n+1}"] & \cdots \\
\cdots \arrow[r] & B_{n-1} \arrow[r] & B_n \arrow[r] & B_{n+1} \arrow[r] & \cdots
\end{tikzcd}$$

that induces isomorphisms on cohomology. Show that the same thing is true after applying the functor $(-)^* = \mathrm{Hom}_R(-, R)$: the induced morphism of complexes

$$\begin{tikzcd}
\cdots \arrow[r] & B^*_{n+1} \arrow[r] \arrow[d, "f^*_{n+1}"] & B^*_n \arrow[r] \arrow[d, "f^*_n"] & B^*_{n+1} \arrow[r] \arrow[d, "f^*_{n-1}"] & \cdots \\
\cdots \arrow[r] & A^*_{n+1} \arrow[r] & A^*_n \arrow[r] & A^*_{n+1} \arrow[r] & \cdots
\end{tikzcd}$$

is again a quasi-isomorphism. (Hint: Use the mapping cone. Show that the mapping cone of f is an exact complex of free R-modules, and therefore homotopic to zero. Show that this property is preserved by the functor $\mathrm{Hom}_R(-, R)$, and conclude that the morphism between the dual complexes is also a quasi-isomorphism.)

Lecture 9: March 4

Local coordinates on algebraic varieties. Let X be an algebraic variety over a field k, with structure sheaf $\mathscr{O}_X$. More precisely, X is a scheme of finite type over k, meaning that for every affine open subset $U \subseteq X$, the ring of functions $\Gamma(U, \mathscr{O}_X)$ is a finitely generated k-algebra, or in other words, a quotient of a polynomial ring. We say that X is *nonsingular* of dimension n if, at each closed point $x \in X$, the stalk

$$\mathscr{O}_{X,x} = \lim_{U \ni x} \Gamma(U, \mathscr{O}_X)$$

is a regular local ring of dimension n; in other words, if $\mathfrak{m}_x \subseteq \mathscr{O}_{X,x}$ denotes the maximal ideal, then

$$\dim_{\mathscr{O}_{X,x}/\mathfrak{m}_x} \mathfrak{m}_x/\mathfrak{m}_x^2 = n = \dim \mathscr{O}_{X,x}.$$

When the field k is perfect (which is always the case in characteristic zero), an equivalent condition is that the sheaf of Kähler differentials $\Omega^1_{X/k}$ is locally free of rank n.

Since we are going to need this in a moment, let me briefly review derivations and Kähler differentials. Let A be a finitely generated k-algebra. A *derivation* from A into an A-module M is a k-linear mapping $D\colon A \to M$ such that $\delta(fg) = f\delta(g) + g\delta(f)$ for every $f, g \in A$. We denote by $\operatorname{Der}_k(A, M)$ the set of all such derivations; this is an A-module in the obvious way. In the special case $M = A$, we use the notation $\operatorname{Der}_k(A)$ for the derivations from A to itself. In view of the formula $\delta(fg) = f\delta(g) + g\delta(f)$, such a derivation is the algebraic analogue of a vector field, acting on the set of functions in A. We have $\operatorname{Der}_k(A) \subseteq \operatorname{End}_k(A)$, and one can check that if $\delta_1, \delta_2 \in \operatorname{Der}_k(A)$, then their commutator

$$[\delta_1, \delta_2] = \delta_1 \circ \delta_2 - \delta_2 \circ \delta_1 \in \operatorname{End}_k(A)$$

is again a derivation. It is the analogue of the Lie bracket on complex manifolds.

The module of *Kähler differentials* $\Omega^1_{A/k}$ represents the functor $M \mapsto \operatorname{Der}_k(A, M)$, in the sense that one has a functorial isomorphism

$$\operatorname{Der}_k(A, M) \cong \operatorname{Hom}_A\bigl(\Omega^1_{A/k}, M\bigr).$$

In other words, $\Omega^1_{A/k}$ is an A-module, together with a derivation $d\colon A \to \Omega^1_{A/k}$, such that every derivation $\delta \in \operatorname{Der}_k(A, M)$ factors uniquely as $\delta = \tilde{\delta} \circ d$ for a unique A-linear map $\tilde{\delta}\colon \Omega^1_{A/k} \to M$. Concretely, $\Omega^1_{A/k}$ can be constructed by taking the free A-module on the set of generators df, for $f \in A$, and imposing the relations $d(fg) = fdg + gdf$ and $d(f+g) = df + dg$ for every $f, g \in A$, and $df = 0$ for every $f \in k$. By construction, one has

$$\operatorname{Der}_k(A) \cong \operatorname{Hom}_A(\Omega^1_{A/k}, A),$$

which makes the module of Kähler differentials dual to the module of derivations.

Globally, $\Omega^1_{X/k}$ is a coherent sheaf of $\mathscr{O}_X$-modules, such that for every affine open subset $U \subseteq X$, one has $\Gamma(U, \Omega^1_{X/k}) = \Omega^1_{A/k}$, where $A = \Gamma(U, \mathscr{O}_X)$. There is again a universal derivation $d\colon \mathscr{O}_X \to \Omega^1_{X/k}$. Think of $\Omega^1_{X/k}$ as an algebraic analogue of the sheaf of holomorphic one-forms on a complex manifold. The *tangent sheaf*

$$\mathscr{T}_X = \mathcal{H}om_{\mathscr{O}_X}(\Omega^1_{X/k}, \mathscr{O}_X)$$

is defined as the dual of the sheaf of Kähler differentials; on affines, one has $\Gamma(U, \mathscr{T}_X) = \operatorname{Der}_k(A)$, using the notation from above. This is an algebraic analogue of the sheaf of holomorphic tangent vector fields on a complex manifold.

Now suppose that X is nonsingular of dimension n, or equivalently, that $\Omega^1_{X/k}$ is locally free of rank n. At every closed point $x \in X$, one can choose *local coordinates* in the following way: there is an affine open neighborhood U of x, together with n regular functions $x_1, \ldots, x_n \in \Gamma(U, \mathscr{O}_X)$, such that

$$\Omega^1_{X/k}\big|_U \cong \bigoplus_{i=1}^n \mathscr{O}_X\big|_U \cdot dx_i.$$

Dually, we have derivations $\partial_1, \ldots, \partial_n \in \mathrm{Der}_k\big(\Gamma(U, \mathscr{O}_X)\big)$, such that

$$\mathscr{T}_X\big|_U \cong \bigoplus_{i=1}^n \mathscr{O}_X\big|_U \cdot \partial_i.$$

This says that $df = \partial_1(f) \cdot dx_1 + \cdots + \partial_n(f) \cdot dx_n$ for every $f \in \Gamma(U, \mathscr{O}_X)$, and so the derivation ∂_i plays the role of the partial derivative operator $\partial/\partial x_i$. One can choose the functions $x_1, \ldots, x_n \in \Gamma(U, \mathscr{O}_X)$ in such a way that they generate the maximal ideal $\mathfrak{m}_x \subseteq \mathscr{O}_{X,x}$. Keep in mind that the morphism $U \to \mathbb{A}^n_k$ defined by the local coordinates is étale, but not usually an embedding (because open sets in the Zariski topology are too big).

The sheaf of differential operators. Let X be a nonsingular algebraic variety. Our goal is to define the sheaf of differential operators $\mathscr{D}_X$, which is a global analogue of the Weyl algebra $A_n(k)$. This will be a quasi-coherent sheaf of $\mathscr{O}_X$-modules $\mathscr{D}_X$, together with an increasing filtration $F_\bullet \mathscr{D}_X$ by coherent $\mathscr{O}_X$-modules, where $F_j \mathscr{D}_X$ consists of differential operators of order $\leq j$.

We start by considering the affine case. So let $U \subseteq X$ be an affine open subset, and set $A = \Gamma(U, \mathscr{O}_X)$, which is a finitely generated k-algebra. We are going to define an A-module $D(A) \subseteq \mathrm{End}_k(A)$, whose elements are the algebraic differential operators of finite order on A. It will satisfy

$$D(A) = \bigcup_{j=0}^\infty F_j D(A),$$

where $F_j D(A)$ is the submodule of operators of order $\leq j$. The idea is that operators of order 0 should be multiplication by elements in A, and that if $P \in F_i D(A)$ and $Q \in F_j D(A)$, then their commutator $[P, Q] = P \circ Q - Q \circ P \in \mathrm{End}_k(A)$ should belong to $F_{i+j-1} D(A)$. This is consistent with what happens for the Weyl algebra.

For an element $f \in A$, we also use the symbol $f \in \mathrm{End}_k(A)$ to denote the operator of multiplication by f. Observe that $P \in \mathrm{End}_k(A)$ is multiplication by the element $P(1) \in A$ if and only if P is A-linear if and only if $[P, f] = 0$ for every $f \in A$. We can therefore define

$$F_0 D(A) = \big\{\, P \in \mathrm{End}_k(A) \,\big|\, [P, f] = 0 \text{ for every } f \in A \,\big\} \cong A.$$

We then define $F_j D(A)$ recursively by saying that

$$F_j D(A) = \big\{\, P \in \mathrm{End}_k(A) \,\big|\, [P, f] \in F_{j-1} D(A) \text{ for every } f \in A \,\big\}.$$

This construction of differential operators is due to Grothendieck.

Example 9.1. Let us work out the relationship between $F_1 D(A)$ and $\mathrm{Der}_k(A)$. Every derivation $\delta \in \mathrm{Der}_k(A)$ is also a differential operator of order 1, because

$$[\delta, f](g) = \delta(fg) - f\delta(g) = \delta(f) \cdot g$$

for every $f, g \in A$, which shows that $[\delta, f] = \delta(f) \in F_0 D(A)$. Conversely, suppose that we have some $P \in F_1 D(A)$. By definition, for every $f \in A$, there exists some $p_f \in A$ such that $[P, f] = p_f$. Concretely, this means that

$$P(fg) - fP(g) = p_f \cdot g$$

for every $f, g \in A$. Taking $g = 1$, we get $p_f = P(f) - fP(1)$, and so

$$P(fg) - fP(g) - gP(f) + fgP(1) = 0.$$

It is then easy to check that $P - P(1)$ is a derivation. The conclusion is that

$$F_1 D(A) \cong A \oplus \operatorname{Der}_k(A)$$

with $P \in F_1 D(A)$ corresponding to the pair $\big(P(1), P - P(1)\big)$.

It is easy to see that each $F_j D(A)$ is a finitely generated A-module, and that composition in $\operatorname{End}_k(A)$ has the following effect: if $P \in F_i D(A)$ and $Q \in F_j D(A)$, then $P \circ Q \in F_{i+j} D(A)$ and $[P, Q] \in F_{i+j-1} D(A)$. With some more work, one can prove the following result.

Proposition 9.2. *Let A be a finitely generated k-algebra. If A is nonsingular of dimension n, then the following is true:*

(a) *As an A-algebra, $D(A) \subseteq \operatorname{End}_k(A)$ is generated by $\operatorname{Der}_k(A)$, subject to the relations $[\delta, f] = \delta(f)$ for every $\delta \in \operatorname{Der}_k(A)$ and every $f \in A$.*
(b) *One has $F_j D(A)/F_{j-1} D(A) \cong \operatorname{Sym}^j \operatorname{Der}_k(A)$ for $j \geq 0$.*
(c) *One has an isomorphism of graded A-algebras*

$$\operatorname{gr}^F D(A) = \bigoplus_{j=0}^{\infty} F_j D(A)/F_{j-1} D(A) \cong \operatorname{Sym} \operatorname{Der}_k(A)$$

between the associated graded algebra of $D(A)$ and the symmetric algebra on $\operatorname{Der}_k(A)$.

Here, for any A-module M, the j-th symmetric power $\operatorname{Sym}^j M$ is the A-module obtained by quotienting $M \otimes_A \cdots \otimes_A M$ by the submodule generated by elements of the form $m_1 \otimes \cdots m_j - m_{\sigma(1)} \otimes \cdots m_{\sigma(j)}$, for all permutations $\sigma \in S_j$. The *symmetric algebra* on M is the graded A-algebra

$$\operatorname{Sym} M = \bigoplus_{j=0}^{\infty} \operatorname{Sym}^j M.$$

It has the following universal property: if B is any A-algebra, then every morphism of A-modules $M \to B$ extends uniquely to a morphism of A-algebras $\operatorname{Sym} M \to B$. For example, one has $\operatorname{Sym} A^{\oplus r} \cong A[x_1, \dots, x_r]$.

Let us give a concrete description of differential operators in local coordinates. Let $U \subseteq X$ be an affine open, with local coordinates $x_1, \dots, x_n$, and set $A = \Gamma(U, \mathscr{O}_X)$. The A-module $\operatorname{Der}_k(A)$ is free of rank n, generated by the derivations $\partial_1, \dots, \partial_n$, and so $D(A)$ is freely generated over A by products of these. In other words, every $P \in F_j D(A)$ can be written uniquely in the form

$$P = \sum_{|\alpha| \leq j} f_\alpha \partial^\alpha,$$

where $\partial^\alpha = \partial_1^{\alpha_1} \cdots \partial_n^{\alpha_n}$ and where $f_\alpha \in A$. The only difference with the case of the Weyl algebra is that the coefficients now belong to the ring A, instead of to the polynomial ring.

Example 9.3. In the case $A = k[x_1, \dots, x_n]$, we have $D(A) = A_n(k)$, and the filtration $F_\bullet D(A)$ agrees with the order filtration.

Now we would like to say that $\mathscr{D}_X$ is the unique sheaf of $\mathscr{O}_X$-modules with the property that $\Gamma(U, \mathscr{D}_X) = D\big(\Gamma(U, \mathscr{O}_X)\big)$ for every affine open $U \subseteq X$. For this to work, one needs the following compatibility result.

Proposition 9.4. *Let A be a finitely generated k-algebra that is nonsingular of dimension n. For nonzero $f \in A$, set $A_f = A[f^{-1}]$. Then one has isomorphisms*

$$D(A_f) \cong A_f \otimes_A D(A) \quad \textit{and} \quad F_j D(A_f) \cong A_f \otimes_A F_j D(A).$$

The content of this is that every differential operator on A_f extends, after multiplication by a sufficiently large power of f, to a differential operator on A. (The analogous result for Kähler differentials is that $\Omega^1_{A_f/k} \cong A_f \otimes_A \Omega^1_{A/k}$; you can find this in Hartshorne, who quotes Matsumura for the proof.)

Note. Unless X is affine, $\Gamma(X, \mathscr{D}_X)$ does not embed into the k-linear endomorphisms of $\Gamma(X, \mathscr{O}_X)$. For example, we shall see below that there are many algebraic differential operators on $\mathbb{P}^n_k$, but since $\mathbb{P}^n_k$ is proper, every regular function on $\mathbb{P}^n_k$ is constant. This is why differential operators are defined locally.

The proposition implies that $\mathscr{D}_X$ is a quasi-coherent sheaf of $\mathscr{O}_X$-modules, and that each $F_j\mathscr{D}_X$ is coherent. Indeed, recall that a sheaf of $\mathscr{O}_X$-modules $\mathscr{F}$ is called *quasi-coherent* if, for every affine open subset $U \subseteq X$, the restriction of $\mathscr{F}$ to U is the sheaf of $\mathscr{O}_X$-modules associated with the $\Gamma(U, \mathscr{O}_X)$-module $\Gamma(U, \mathscr{F})$. On an affine scheme $\operatorname{Spec} A$, a necessary and sufficient condition for $\mathscr{F}$ to be quasi-coherent is that

$$\Gamma\big(D(f), \mathscr{F}\big) \cong A_f \otimes_A \Gamma(\operatorname{Spec} A, \mathscr{F})$$

for every $f \in A$, where $D(f) \subseteq \operatorname{Spec} A$ denotes the principal affine open defined by f. When X is noetherian, which is the case for schemes of finite type over a field, $\mathscr{F}$ is *coherent* if each $\Gamma(U, \mathscr{F})$ is finitely generated over $\Gamma(U, \mathscr{O}_X)$. So the proposition says exactly that $\mathscr{D}_X$ is quasi-coherent and that each $F_j\mathscr{D}_X$ is coherent.

The isomorphisms in Proposition 9.2 globalize as follows. One has $F_0\mathscr{D}_X = \mathscr{O}_X$, and for every $j \geq 0$, one has

$$\operatorname{gr}^F_j \mathscr{D}_X = F_j\mathscr{D}_X / F_{j-1}\mathscr{D}_X \cong \operatorname{Sym}^j \mathscr{T}_X,$$

where $\mathscr{T}_X$ is the tangent sheaf. One also has an isomorphism of graded $\mathscr{O}_X$-algebras

$$\operatorname{gr}^F \mathscr{D}_X \cong \operatorname{Sym} \mathscr{T}_X,$$

and so the associated graded algebra of $\mathscr{D}_X$ is again commutative, as in the case of the Weyl algebra. Since X is nonsingular, $\mathscr{T}_X$ is locally free of rank n, and the symmetric algebra on $\mathscr{T}_X$ can be interpreted as the sheaf of algebraic functions on the cotangent bundle. Let us denote by $p \colon T^*X \to X$ the cotangent bundle of X, with its natural projection to X. This is again a nonsingular algebraic variety, now of dimension $2n$, locally isomorphic to the product of X and affine space $\mathbb{A}^n_k$. By the correspondence between vector bundles and locally free sheaves (from Hartshorne's book), one has an isomorphism

$$T^*X \cong \mathbb{V}(\mathscr{T}_X) = \mathbf{Spec}_X \operatorname{Sym} \mathscr{T}_X,$$

and therefore $p_*\mathscr{O}_{T^*X} \cong \operatorname{Sym} \mathscr{T}_X$ as $\mathscr{O}_X$-algebras. This is why people sometimes refer to $\mathscr{D}_X$ as a "noncommutative deformation" of the cotangent bundle.

Example 9.5. Let us consider the example $X = \mathbb{P}^n_k$. The k-vector space $\Gamma(X, \mathscr{D}_X)$ of global differential operators on projective space is infinite-dimensional. There are several ways to see this. One way is by diagram chasing. We have $F_0\mathscr{D}_X = \mathscr{O}_X$, and therefore $\Gamma(X, F_0\mathscr{D}_X) = k$. For each $j \geq 1$, we have a short exact sequence

$$0 \to F_{j-1}\mathscr{D}_X \to F_j\mathscr{D}_X \to \operatorname{Sym}^j \mathscr{T}_X \to 0.$$

One can show by induction that $H^1(X, F_j\mathscr{D}_X) = 0$ for $j \geq 0$, and so

$$H^0(X, F_j\mathscr{D}_X)/H^0(X, F_{j-1}\mathscr{D}_X) \cong H^0(X, \operatorname{Sym}^j \mathscr{T}_X).$$

These vector spaces can then be computed using the Euler sequence

$$0 \to \mathscr{O}_X \to \mathscr{O}_X(1)^{\oplus(n+1)} \to \mathscr{T}_X \to 0.$$

For example, $\dim H^0(X, \mathscr{T}_X) = (n+1)^2 - 1$, and so $\dim H^0(X, F_1\mathscr{D}_X) = (n+1)^2$.

Another way is to use the standard open covering $X = U_0 \cup U_1 \cup \cdots \cup U_n$. Since each U_i is isomorphic to $\mathbb{A}_k^n$, one has $\Gamma(U_i, \mathscr{D}_X) \cong A_n(k)$, and so an element of $\Gamma(X, \mathscr{D}_X)$ can be described by $(n+1)$ elements of the Weyl algebra that are related to each other by the coordinate transformations among the U_i. (See the exercises.)

The third way is to use the presentation of X as a quotient of $\mathbb{A}_k^{n+1}$ minus the origin, by identifying points of $\mathbb{P}_k^n$ with lines in $\mathbb{A}_k^{n+1}$. Recall how this works in the case of the Euler sequence. Once $n \geq 1$, a vector field on $\mathbb{A}_k^{n+1}$ minus the origin is the same thing as a vector field on $\mathbb{A}_k^{n+1}$, hence of the form

$$f_0 \partial_0 + f_1 \partial_1 + \cdots + f_n \partial_n,$$

for polynomials $f_0, \dots, f_n \in k[x_0, \dots, x_n]$. Such a vector field descends to X if and only if it is homogeneous of degree 0, where $\deg x_j = 1$ and $\deg \partial_j = -1$. At the same time, the Euler vector field

$$x_0 \partial_0 + x_1 \partial_1 + \cdots + x_n \partial_n$$

is tangent to the lines through the origin, and therefore descends to the zero vector field. This shows that $\Gamma(X, \mathscr{T}_X)$ is generated by the $(n+1)^2$ vector fields $x_i \partial_j$, subject to the single relation $x_0 \partial_0 + \cdots + x_n \partial_n = 0$. In the same way, one can show that $\Gamma(X, \mathscr{D}_X)$ is isomorphic to the space of differential operators on $\mathbb{A}_k^{n+1}$ that are homogeneous of degree 0, modulo the ideal generated by the Euler vector field. Concretely, an element $P \in \Gamma(X, F_j \mathscr{D}_X)$ can be written in the form

$$P = \sum_{|\alpha| = |\beta| \leq j} c_\alpha x_0^{\alpha_0} \cdots x_n^{\alpha_n} \partial_0^{\beta_0} \cdots \partial_n^{\beta_n}$$

and this expression is unique modulo multiples of $x_0 \partial_0 + \cdots + x_n \partial_n$. The restriction of P to the standard affine open U_0 is obtained by setting $x_0 = 1$ and using the relation $\partial_0 = -(x_1 \partial_1 + \cdots + x_n \partial_n)$.

Algebraic $\mathscr{D}_X$-modules. Let me end with the following definition. An *algebraic $\mathscr{D}$-module* on a nonsingular algebraic variety X is a quasi-coherent sheaf of $\mathscr{O}_X$-modules $\mathcal{M}$, together with a (left or right) action by the sheaf of differential operators $\mathscr{D}_X$. In other words, for every affine open subset $U \subseteq X$, with $A = \Gamma(U, \mathscr{O}_X)$, we get an A-module M, together with a (left or right) action by the module of differential operators $D(A)$.

Exercises.

Exercise 9.1. Show that one has $\operatorname{Der}_k(A_f) \cong A_f \otimes_A \operatorname{Der}_k(A)$ for every $f \in A$.

Exercise 9.2. For $X = \mathbb{P}_k^n$, compute $\dim_k \Gamma(X, F_j \mathscr{D}_X)$ as a function of $j \geq 0$.

Exercise 9.3. Consider the example $X = \mathbb{P}_k^1$. If we use the symbol x_0 for the coordinate on $U_0 = \mathbb{A}_k^1$, and x_1 for the coordinate on $U_1 = \mathbb{A}_k^1$, then $\Gamma(U_0, \mathscr{D}_X)$ is the Weyl algebra on x_0 and ∂_0, and $\Gamma(U_1, \mathscr{D}_X)$ is the Weyl algebra on x_1 and ∂_1. Using the coordinate change $x_1 = x_0^{-1}$, decide when two differential operators

$$P = \sum_{i,j} a_{i,j} x_0^i \partial_0^j \quad \text{and} \quad Q = \sum_{i,j} b_{i,j} x_1^i \partial_1^j$$

have the same restriction to $U_0 \cap U_1$. Use this to describe the space $\Gamma(X, \mathscr{D}_X)$ of global differential operators on $\mathbb{P}_k^1$.

Lecture 10: March 6

Algebraic $\mathscr{D}$-modules. Let me first recall the definition of an algebraic $\mathscr{D}$-module from last time. As before, X is an algebraic variety over a field k, nonsingular of constant dimension n. We denote by $\mathscr{D}_X$ the sheaf of algebraic differential operators on X, and by $F_j\mathscr{D}_X$ the subsheaf of operators of order $\leq j$. Then each $F_j\mathscr{D}_X$ is a coherent sheaf of $\mathscr{O}_X$-modules, and $\mathscr{D}_X$ itself is quasi-coherent.

Definition 10.1. An *algebraic $\mathscr{D}$-module* is a quasi-coherent sheaf of $\mathscr{O}_X$-modules $\mathcal{M}$, together with a left (or right) action by $\mathscr{D}_X$.

Since $\mathscr{D}_X$ is noncommutative, we again have to distinguish between left and right modules. In the case of a left $\mathscr{D}$-module $\mathcal{M}$, the set of sections $M = \Gamma(U, \mathcal{M})$ over any affine open subset $U \subseteq X$ is thus a left module over the algebra of differential operators $D(A)$, where $A = \Gamma(U, \mathscr{O}_X)$. The quasi-coherence condition means that the restriction of $\mathcal{M}$ to the open set U is uniquely determined by this $D(A)$-module. Recall from Lecture 9 that the algebra $D(A)$ is generated, as an A-subalgebra of $\operatorname{End}_k(A)$, by the derivations $\operatorname{Der}_k(A)$, subject to the relation $[\delta, f] = \delta(f)$ for all $\delta \in \operatorname{Der}_k(A)$ and all $f \in A$. The left $D(A)$-action on M is therefore the same thing as a k-linear mapping

$$\operatorname{Der}_k(A) \otimes_k M \to M, \quad \delta \otimes m \mapsto \delta m,$$

such that $(f\delta)m = f(\delta m)$, $\delta(fm) = f\delta(m) + \delta(f)m$ and $\delta(\eta m) - \eta(\delta m) = [\delta, \eta]m$ for all $\delta, \eta \in \operatorname{Der}_k(A)$, all $f \in A$, and all $m \in M$. Globally, to turn a quasi-coherent sheaf of $\mathscr{O}_X$-modules $\mathcal{M}$ into a left $\mathscr{D}_X$-module, we need a k-linear morphism

$$\mathscr{T}_X \otimes_k \mathcal{M} \to \mathcal{M}$$

that satisfies those three conditions locally. (You can work out for yourself what happens for right $\mathscr{D}$-modules.)

Example 10.2. Since the algebra of differential operators on the affine space $\mathbb{A}^n_k$ is the Weyl algebra $A_n(k)$, an algebraic $\mathscr{D}$-module on $\mathbb{A}^n_k$ is (up to the equivalence between quasi-coherent sheaves and modules) the same thing as a left (or right) module over $A_n(k)$.

Here are some examples of left and right $\mathscr{D}$-modules.

Example 10.3. The structure sheaf $\mathscr{O}_X$ is a left $\mathscr{D}_X$-module. Indeed, for every affine open subset $U \subseteq X$, the algebra of differential operators $D(A)$ acts on $A = \Gamma(U, \mathscr{O}_X)$ by construction.

Example 10.4. Every algebraic vector bundle with integrable connection is a left $\mathscr{D}_X$-module. Let $\mathscr{E}$ be the corresponding locally free sheaf of $\mathscr{O}_X$-modules; in Hartshorne's notation, the vector bundle is then $\mathbb{V}(\mathscr{E}^*)$. A *connection* is a k-linear morphism $\nabla \colon \mathscr{E} \to \Omega^1_{X/k} \otimes_{\mathscr{O}_X} \mathscr{E}$ that satisfies the Leibniz rule. In other words, for every affine open subset $U \subseteq X$ and every pair of sections $s \in \Gamma(U, \mathscr{E})$ and $f \in \Gamma(U, \mathscr{O}_X)$, the connection should satisfy

$$\nabla(fs) = f\nabla(s) + df \otimes s.$$

We can also regard the connection as a k-linear morphism $\nabla \colon \mathscr{T}_X \otimes_k \mathscr{E} \to \mathscr{E}$, but we use the differential geometry notation $\nabla_\theta(s)$ instead of $\nabla(\theta \otimes s)$ for $\theta \in \Gamma(U, \mathscr{T}_X)$ and $s \in \Gamma(U, \mathscr{E})$. In this notation, we have

$$\nabla_{f\theta}(s) = f\nabla_\theta(s), \tag{10.5}$$

and the Leibniz rule becomes

$$\nabla_\theta(fs) = f\nabla_\theta(s) + \theta(f)s. \tag{10.6}$$

The connection is called *integrable* if

$$\nabla_\theta \circ \nabla_\eta - \nabla_\eta \circ \nabla_\theta = \nabla_{[\theta,\eta]} \tag{10.7}$$

for every pair of vector fields $\theta, \eta \in \Gamma(U, \mathscr{T}_X)$. This is equivalent to the vanishing of the curvature operator in $\Omega^2_{X/k} \otimes_{\mathscr{O}_X} \mathscr{E}nd_{\mathscr{O}_X}(\mathscr{E})$. The conditions in (10.5), (10.6) and (10.7) are exactly saying that the action of $\mathscr{T}_X$ on $\mathscr{E}$ extends to a left action by the sheaf of differential operators $\mathscr{D}_X$, and so $\mathscr{E}$ becomes a left $\mathscr{D}$-module.

In general, the left action of $\mathscr{D}_X$ on a left $\mathscr{D}$-module $\mathcal{M}$ may be considered (formally) as a connection operator $\nabla\colon \mathcal{M} \to \Omega^1_{X/k} \otimes_{\mathscr{O}_X} \mathcal{M}$ that satisfies the Leibniz rule and is integrable, in the sense that it locally satisfies the conditions expressed in (10.5), (10.6) and (10.7).

Example 10.8. Unlike in the case of affine space, we cannot turn left $\mathscr{D}$-modules into right $\mathscr{D}$-modules by changing signs, since we might not be able to do this consistently on all affine open subsets. Instead, the primary example of a right $\mathscr{D}$-module is the canonical bundle $\omega_X = \bigwedge^n \Omega^1_{X/k}$, whose sections are the algebraic n-forms. If $U \subseteq X$ is an affine open subset with local coordinates $x_1, \dots, x_n$, then ω_X is locally free of rank one, spanned by $dx_1 \wedge \cdots \wedge dx_n$. The tangent sheaf $\mathscr{T}_X$ acts on ω_X by Lie differentiation. Given $\omega \in \Gamma(U, \omega_X)$ and $\theta, \theta_1, \dots, \theta_n \in \Gamma(U, \mathscr{T}_X)$, the formula for the Lie derivative is

$$(\operatorname{Lie}_\theta \omega)(\theta_1, \dots, \theta_n) = \theta \cdot \omega(\theta_1, \dots, \theta_n) - \sum_{i=1}^n \omega\bigl(\theta_1, \dots, [\theta, \theta_i], \dots, \theta_n\bigr).$$

One can check quite easily that the following relations hold:

$$\begin{aligned} \operatorname{Lie}_\theta(f\omega) &= f \operatorname{Lie}_\theta \omega + \theta(f)\omega = \operatorname{Lie}_{f\theta} \omega \\ \operatorname{Lie}_{[\theta,\eta]} \omega &= \operatorname{Lie}_\theta \operatorname{Lie}_\eta \omega - \operatorname{Lie}_\eta \operatorname{Lie}_\theta \omega \end{aligned}$$

This almost looks like ω_X should be a left $\mathscr{D}_X$-module, but note that (10.5) is not satisfied since $\operatorname{Lie}_{f\theta} \omega \neq f \operatorname{Lie}_\theta \omega$. But if we instead define

$$\omega_X \otimes_k \mathscr{T}_X \to \omega_X, \quad \omega \otimes \theta \mapsto \omega \cdot \theta = -\operatorname{Lie}_\theta(\omega)$$

and also write the $\mathscr{O}_X$-action on ω_X on the right, we obtain

$$\begin{aligned} &\omega \cdot \theta(f) = (-\operatorname{Lie}_\theta \omega) f + \operatorname{Lie}_\theta(\omega f) = (\omega \cdot \theta) f - (\omega f) \cdot \theta \\ &\omega \cdot [\theta, \eta] = -\operatorname{Lie}_{[\theta,\eta]} \omega = \operatorname{Lie}_\theta \operatorname{Lie}_\eta \omega - \operatorname{Lie}_\eta \operatorname{Lie}_\theta \omega = (\omega \cdot \theta) \cdot \eta - (\omega \cdot \eta) \cdot \theta. \end{aligned}$$

These are exactly the relations defining $\mathscr{D}_X$, and so we obtain on ω_X the structure of a right $\mathscr{D}_X$-module. In local coordinates, we have

$$(f\, dx_1 \wedge \cdots \wedge dx_n) \cdot P = (P^\sigma f)\, dx_1 \wedge \cdots \wedge dx_n,$$

where $P^\sigma = \sum (-\partial)^\alpha f_\alpha$ is the formal adjoint of $P = \sum f_\alpha \partial^\alpha$. In local coordinates, the left $\mathscr{D}$-module structure on $\mathscr{O}_X$ and the right $\mathscr{D}$-module structure on ω_X are therefore related to each other exactly as in the case of the Weyl algebra.

Good filtrations and characteristic variety. As in the case of the Weyl algebra, we study $\mathscr{D}$-modules using filtrations. Let $\mathcal{M}$ be a left $\mathscr{D}_X$-module. We consider increasing filtrations $F_\bullet \mathcal{M}$ by coherent $\mathscr{O}_X$-submodules $F_j \mathcal{M}$ such that

$$F_i \mathscr{D}_X \cdot F_j \mathcal{M} \subseteq F_{i+j} \mathcal{M}$$

for all $i, j \in \mathbb{Z}$. We also assume that the filtration is exhaustive, meaning that

$$\bigcup_{j \in \mathbb{Z}} F_j \mathcal{M} = \mathcal{M}.$$

Note that each $F_j\mathcal{M}$ is assumed to be coherent over $\mathscr{O}_X$. We say that such a filtration is *good* if the associated graded module

$$\operatorname{gr}^F \mathcal{M} = \bigoplus_{j \in \mathbb{Z}} F_j\mathcal{M}/F_{j-1}\mathcal{M}$$

is locally finitely generated over $\operatorname{gr}^F \mathscr{D}_X$. This implies that $F_j\mathcal{M} = 0$ for $j \ll 0$.

Now suppose that $U \subseteq X$ is an affine open subset, and set $A = \Gamma(U, \mathscr{O}_X)$ and $M = \Gamma(U, \mathcal{M})$. By the same argument as in the case of the Weyl algebra, one shows that M is finitely generated over $D(A)$ if and only if admits a good filtration $F_\bullet M$ by finitely generated A-modules; again, this means that $F_i D(A) \cdot F_j M \cdot F_{i+j} M$ and $\operatorname{gr}^F M$ is finitely generated over $\operatorname{gr}^F D(A)$.

Definition 10.9. We say that a left (or right) $\mathscr{D}_X$-module is *coherent* if it is locally finitely generated over $\mathscr{D}_X$.

Note that this is not the same thing as being $\mathscr{O}_X$-coherent; in fact, most coherent $\mathscr{D}_X$-modules are not coherent over $\mathscr{O}_X$. Every coherent $\mathscr{D}_X$-module has a good filtration locally, meaning on each affine open subset; in fact, we will see next time that coherent $\mathscr{D}_X$-modules always admit a global good filtration $F_\bullet\mathcal{M}$.

Given a good filtration $F_\bullet\mathcal{M}$ (globally or locally), the associated graded $\operatorname{gr}^F\mathcal{M}$ is coherent over the sheaf of $\mathscr{O}_X$-algebras

$$\operatorname{gr}^F \mathscr{D}_X \cong \operatorname{Sym} \mathscr{T}_X \cong p_* \mathscr{O}_{T^*X},$$

where $p \colon T^*X \to X$ again means the cotangent bundle. By the correspondence between coherent sheaves on T^*X and finitely generated modules over $p_*\mathscr{O}_{T^*X}$, we thus obtain a coherent sheaf of $\mathscr{O}_{T^*X}$-modules on the cotangent bundle that we denote by the symbol $\widetilde{\operatorname{gr}^F\mathcal{M}}$.

Definition 10.10. The *characteristic variety* $\operatorname{Ch}(\mathcal{M})$ is the closed algebraic subset of T^*X given by the support of $\widetilde{\operatorname{gr}^F\mathcal{M}}$, with the reduced scheme structure.

As in the case of the Weyl algebra, any two good filtrations on $\mathcal{M}$ are comparable; for the same reason as before, this implies that the subsheaf

$$\sqrt{\operatorname{Ann}_{\operatorname{gr}^F \mathscr{D}_X} \operatorname{gr}^F \mathcal{M}} \subseteq \operatorname{gr}^F \mathscr{D}_X$$

is independent of the choice of good filtration. If we denote by $\mathcal{J}_{\mathcal{M}} \subseteq \mathscr{O}_{T^*X}$ the corresponding coherent sheaf of ideals on the cotangent bundle, then $\operatorname{Ch}(\mathcal{M})$ is the closed subscheme defined by $\mathcal{J}_{\mathcal{M}}$. We are going to show later on that Bernstein's inequality carries over to arbitrary coherent $\mathscr{D}$-modules: as long as $\mathcal{M} \neq 0$, every irreducible component of $\operatorname{Ch}(\mathcal{M})$ has dimension at least n.

Example 10.11. If $\mathscr{E}$ is the left $\mathscr{D}_X$-module determined by a vector bundle with integrable connection, then $\operatorname{Ch}(\mathscr{E})$ is the zero section. The reason is that $\mathscr{E}$ is coherent over $\mathscr{O}_X$, which means that setting $F_j\mathscr{E} = 0$ for $j < 0$ and $F_j\mathscr{E} = \mathscr{E}$ for $j \geq 0$ gives a good filtration. Here

$$\operatorname{Ann}_{\operatorname{gr}^F \mathscr{D}_X} \operatorname{gr}^F \mathscr{E} = \bigoplus_{j \geq 1} \operatorname{gr}_j^F \mathscr{D}_X,$$

and so $\mathcal{J}_{\mathscr{E}}$ is the ideal of the zero section. Of course, this works more generally for any $\mathscr{D}$-module that is coherent over $\mathscr{O}_X$.

The example has a useful converse.

Proposition 10.12. *Let $\mathcal{M}$ be a coherent $\mathscr{D}_X$-module. If $\mathcal{M}$ is coherent over $\mathscr{O}_X$, then $\mathcal{M}$ is actually a locally free $\mathscr{O}_X$-module of finite rank (and therefore comes from a vector bundle with integrable connection).*

Proof. Since $\mathcal{M}$ is a quasi-coherent $\mathscr{O}_X$-module, it suffices to check that the localization $\mathscr{O}_{X,x} \otimes_{\mathscr{O}_X} \mathcal{M}$ at every closed point $x \in X$ is a free $\mathscr{O}_{X,x}$-module of finite rank. This reduces the problem to the following special case: A is a regular local ring of dimension n, containing a field k, with maximal ideal $\mathfrak{m}$ and residue field $A/\mathfrak{m} \cong k$, and M is a left $D(A)$-module that is finitely generated over A. Here $D(A)$ is again the algebra of k-linear differential operators on A. We need to prove that M is a free A-module of finite rank.

First, some preparations. Since A is regular of dimension n, the maximal ideal $\mathfrak{m}$ is generated by n elements $x_1, \dots, x_n$ whose images in $\mathfrak{m}/\mathfrak{m}^2$ are linearly independent over k. Let $\partial_1, \dots, \partial_n \in \operatorname{Der}_k(A)$ be the corresponding derivations, which freely generate $\operatorname{Der}_k(A)$ as an A-module. For every nonzero $f \in A$, we define the order of vanishing as

$$\operatorname{ord}(f) = \max\bigl\{ \ell \geq 0 \bigm| f \in \mathfrak{m}^\ell \bigr\};$$

this makes sense because the intersection of all powers of the maximal ideal is trivial. If $f = 0$, we formally set $\operatorname{ord}(f) = +\infty$. The key point is that we can reduce the order of vanishing of f by applying a suitable derivation. Indeed, suppose that $\operatorname{ord}(f) = \ell$. The ideal $\mathfrak{m}^\ell$ is generated by all monomials of degree ℓ in $x_1, \dots, x_n$, and so we can write

$$f = \sum_{|\alpha| = \ell} f_\alpha x^\alpha,$$

with at least one $f_\alpha \in A$ being a unit (because otherwise $f \in \mathfrak{m}^{\ell+1}$). Choose a multi-index α such that f_α is a unit, and then choose $i = 1, \dots, n$ such that $\alpha_i \geq 1$. Since $\partial_i(x_j) = \delta_{i,j}$, we get

$$\partial_i(f) = \sum_{|\alpha| = \ell} \Bigl(\partial_i(f_\alpha) x^\alpha + f_\alpha \alpha_i x^{\alpha - e_i} \Bigr),$$

and this expression clearly belongs to $\mathfrak{m}^{\ell-1}$ but not to $\mathfrak{m}^\ell$. Hence $\operatorname{ord}(\partial_i(f)) = \ell - 1$.

As I said, we need to prove that M is a free A-module of finite rank. To do this, pick a minimal set of generators $m_1, \dots, m_r \in M$, whose images in $M/\mathfrak{m}M$ are linearly independent over k. This gives us a surjective morphism of A-modules

$$A^{\oplus r} \to M, \quad (f_1, \dots, f_r) \mapsto f_1 m_1 + \dots + f_r m_r,$$

and we are going to show that it is also injective, hence an isomorphism. Suppose that there was a nontrivial relation $f_1 m_1 + \dots + f_r m_r = 0$. Then $f_1, \dots, f_r \in \mathfrak{m}$, because $m_1, \dots, m_r$ are linearly independent modulo $\mathfrak{m}M$. In other words, we have

$$\ell = \min\bigl\{ \operatorname{ord}(f_1), \dots, \operatorname{ord}(f_r) \bigr\} \geq 1.$$

Now the idea is to use the $D(A)$-module structure to create another relation for which the value of ℓ is strictly smaller. By repeating this, we eventually arrive at a relation with $\ell = 0$, contradicting the fact that $m_1, \dots, m_r$ are linearly independent modulo $\mathfrak{m}M$. Here we go. If we apply ∂_i to our relation, we obtain

$$0 = \partial_i \cdot \sum_{j=1}^r f_j m_j = \sum_{j=1}^r [\partial_i, f_j] m_j + \sum_{j=1}^r f_j (\partial_i m_j) = \sum_{j=1}^r \partial_i(f_j) m_j + \sum_{j=1}^r f_j (\partial_i m_j).$$

We can write each $\partial_i m_j$ in terms of the generators $m_1, \dots, m_r$ as

$$\partial_i m_j = \sum_{k=1}^r a_{i,j,k} m_k,$$

and after reindexing, we get the new relation

$$\sum_{j=1}^r \Bigl(\partial_i(f_j) + \sum_{k=1}^r a_{i,k,j} f_k \Bigr) m_j = 0.$$

If we now choose j such that $\operatorname{ord}(f_j) = \ell$, and then choose i such that $\operatorname{ord}(\partial_i(f_j)) = \ell - 1$, then the j-th coefficient in the new relation belongs to $\mathfrak{m}^{\ell-1}$ but not to $\mathfrak{m}^{\ell}$, as desired. $\square$

We showed in Lecture 5 that $\mathcal{M}$ is coherent over $\mathscr{O}_X$ if and only if its characteristic variety is contained in the zero section of the cotangent bundle. This means that if $\mathcal{M}$ is a coherent $\mathscr{D}_X$-module with $\operatorname{Ch}(\mathcal{M})$ contained in the zero section, then $\mathcal{M}$ is a locally free $\mathscr{O}_X$-module of finite rank, and the $\mathscr{D}_X$-module structure is the same as the datum of an integrable connection on $\mathcal{M}$.

Lecture 11: March 11

Global good filtrations. Let us return for the moment to the topic of good filtrations. I said last time that, by the same argument as in the case of A_n, every coherent $\mathscr{D}_X$-module locally admits a good filtration. But in fact, good filtrations also exist globally, because of the finiteness inherent in the definitions.

Lemma 11.1. *Let $\mathcal{M}$ be an algebraic $\mathscr{D}_X$-module. If $\mathcal{M}$ is coherent, then there exists a good filtration $F_\bullet \mathcal{M}$ by coherent $\mathscr{O}_X$-modules.*

Proof. It will be enough to construct an $\mathscr{O}_X$-submodule $\mathscr{F} \subseteq \mathcal{M}$ that is coherent over $\mathscr{O}_X$ and that generates $\mathcal{M}$ as a $\mathscr{D}_X$-module. Once we have that, we can define a filtration by setting

$$F_j \mathcal{M} = F_j \mathscr{D}_X \cdot \mathscr{F} \subseteq \mathcal{M},$$

and for the same reason as in the case of the Weyl algebra, each $F_j \mathcal{M}$ is coherent over $\mathscr{O}_X$, and the filtration $F_\bullet \mathcal{M}$ is good.

Since X is of finite type over k, it is quasi-compact, and so we can cover X by finitely many affine open subsets $U_1, \dots, U_m$. Then $\Gamma(U_i, \mathcal{M})$ is finitely generated over $\Gamma(U_i, \mathscr{D}_X)$, and after choosing a finite set of generators and taking the $\Gamma(U_i, \mathscr{O}_X)$-submodule of $\Gamma(U_i, \mathcal{M})$ generated by this set, we certainly obtain a coherent $\mathscr{O}_{U_i}$-module $\mathscr{F}_{U_i} \subseteq \mathcal{M}\big|_{U_i}$ that has the desired properties on U_i.

To turn these locally defined subsheaves into global objects, we use the following fact from Hartshorne's book: Suppose that $\mathscr{G}$ is a quasi-coherent sheaf on an algebraic variety X. If we have a nonempty open subset $U \subseteq X$, and a coherent subsheaf $\mathscr{F}_U \subseteq \mathscr{G}\big|_U$, then there is a coherent subsheaf $\mathscr{F} \subseteq \mathscr{G}$ such that $\mathscr{F}\big|_U = \mathscr{F}_U$. When applied to our situation, this says that there are coherent $\mathscr{O}_X$-modules $\mathscr{F}_1, \dots, \mathscr{F}_n$ such that $\mathscr{F}_i\big|_{U_i} = \mathscr{F}_{U_i}$. Then the image of

$$\mathscr{F}_1 \oplus \dots + \oplus \mathscr{F}_m \to \mathcal{M}$$

is an $\mathscr{O}_X$-submodule of $\mathcal{M}$ that is coherent over $\mathscr{O}_X$ (because it is the image of a coherent $\mathscr{O}_X$-module) and generates $\mathcal{M}$ as a $\mathscr{D}_X$-module. □

This result is peculiar to the algebraic setting, and does not hold at all for analytic $\mathscr{D}$-modules.

Characteristic varieties are involutive. Recall the definition of the characteristic variety from last time. If $\mathcal{M}$ is a coherent $\mathscr{D}_X$-module, we can choose a global good filtration $F_\bullet \mathcal{M}$, which makes the associated graded module $\operatorname{gr}^F \mathcal{M}$ coherent over $\operatorname{gr}^F \mathscr{D}_X \cong \operatorname{Sym} \mathscr{T}_X$. If $\widetilde{\operatorname{gr}^F \mathcal{M}}$ denotes the corresponding coherent sheaf on the cotangent bundle T^*X, then

$$\operatorname{Ch}(\mathcal{M}) = \operatorname{Supp} \widetilde{\operatorname{gr}^F \mathcal{M}}.$$

Equivalently, the characteristic variety is the reduced closed subscheme of the cotangent bundle corresponding to the homegeneous ideal

$$\sqrt{\operatorname{Ann}_{\operatorname{gr}^F \mathscr{D}_X} \operatorname{gr}^F \mathcal{M}} \subseteq \operatorname{gr}^F \mathscr{D}_X.$$

The most important result about the characteristic variety is the following theorem.

Theorem 11.2. $\operatorname{Ch}(\mathcal{M})$ *is* involutive *with respect to the natural symplectic structure on T^*X. In particular, every irreducible component of* $\operatorname{Ch}(\mathcal{M})$ *has dimension* $\geq n$.

Note that this gives a lot more information about the characteristic variety than Bernstein's inequality. This result was first proved by analytic methods, but Gabber later discovered an algebraic proof. Bernstein's inequality can of course be proved by more elementary means. We are not going to prove Theorem 11.2; instead, I will review some basic facts about symplectic geometry, so that we can understand at least the statement, and where the difficulties lie.

Symplectic vector spaces. Let us start with a brief discussion of symplectic vector spaces. Let V be a finite-dimensional vector space over a field k. Usually, k will be field of real or complex numbers, but the definition works over any field of characteristic $\neq 2$. A *symplectic form* is a bilinear form

$$\sigma \colon V \otimes_k V \to k$$

that is anti-symmetric and non-degenerate. In other words, one has $\sigma(v,w) = -\sigma(w,v)$ for every $v, w \in V$, and if we denote by $V^* = \operatorname{Hom}_k(V,k)$ the dual vector space, then the induced linear mapping

$$V \to V^*, \quad w \mapsto \sigma(-,w),$$

is an isomorphism (called the "Hamiltonian isomorphism"). For every linear functional $\theta \in V^*$, one therefore has a unique element $H_\theta \in V$ such that $\theta(v) = \sigma(v, H_\theta)$ for all $v \in V$.

The dimension of a symplectic vector space is always an even number. One way to see this is as follows. Pick a nonzero vector $w \in V$, and consider the linear subspace $L = k \cdot w \subseteq V$. Since $\sigma(w,w) = 0$, one has L contained in the subspace

$$L^\perp = \big\{\, v \in V \;\big|\; \sigma(v,w) = 0 \,\big\}.$$

The fact that σ is nondegenerate implies that $L^\perp = \dim V - 1$. One easily checks that the quotient space $L^\perp/L$, with the bilinear form induced by σ, is again a symplectic vector space. Since $\dim V = 2 + \dim L^\perp/L$, the claim now follows by induction.

Example 11.3. If V is any finite-dimensional k-vector space, then $V \oplus V^*$ is a symplectic vector space, with symplectic form given by

$$\big((v_1,\theta_1),(v_2,\theta_2)\big) \mapsto \theta_1(v_2) - \theta_2(v_1).$$

In fact, every symplectic vector space is isomorphic to this model (after a suitable choice of basis).

Given a subspace $W \subseteq V$, one defines

$$W^\perp = \big\{\, v \in V \;\big|\; \sigma(v,w) = 0 \text{ for every } w \in W \,\big\}.$$

Under the Hamiltonian isomorphism $V \cong V^*$, the subspace $W^\perp$ corresponds exactly to the kernel of the restriction homomorphism $V^* \to W^*$, and therefore

$$\dim W + \dim W^\perp = \dim V.$$

Definition 11.4. Let $W \subseteq V$ be a linear subspace.

(1) W is called *involutive* if $W^\perp \subseteq W$; then $\dim W \geq \frac{1}{2}\dim V$.
(2) W is called *Lagrangian* if $W^\perp = W$; then $\dim W = \frac{1}{2}\dim V$.
(3) W is called *isotropic* if $W^\perp \supseteq W$; then $\dim W \leq \frac{1}{2}\dim V$.

Note that an involutive (or isotropic) subspace is Lagrangian iff $\dim W = \frac{1}{2}\dim V$.

Example 11.5. Consider the symplectic vector space $V \oplus V^*$. If $W \subseteq V$ is any linear subspace, then $W \oplus \ker(V^* \to W^*)$ is always a Lagrangian subspace of $V \oplus V^*$. It is clearly isotropic: if v_1, v_2 are vectors in W, and θ_1, θ_2 are linear functionals whose restriction to W is trivial, then $\theta_1(v_2) - \theta_2(v_1) = 0$. Since

$$\dim W + \dim \ker(V^* \to W^*) = \dim V$$

is exactly half the dimension of $V \oplus V^*$, it follows that the subspace is Lagragian.

Symplectic algebraic varieties. A nonsingular algebraic variety X is called *symplectic* if the tangent space $T_x X$ at every closed point $x \in X$ is a symplectic vector space, and the symplectic forms vary in an algebraic way from point to point. More precisely, there should exist a global algebraic two-form $\sigma \in \Gamma(X, \Omega^2_{X/k})$ whose restriction $\sigma_x \colon T_x X \otimes_k T_x X \to k$ gives a symplectic form on $T_x X$ for every closed point $x \in X$. Of course, this implies that $\dim X$ is even.

Example 11.6. The example we care about is the cotangent bundle T^*X of a nonsingular algebraic variety X of dimension n. Note that $\dim T^*X = 2n$. If we choose local coordinates $x_1, \dots, x_n$ on X, then the differentials $dx_1, \dots, dx_n$ give a local trivialization for $\Omega^1_{X/k}$, and so we obtain local coordinates $x_1, \dots, x_n, \xi_1, \dots, \xi_n$ on the cotangent bundle. In these coordinates,

$$\sigma_X = \sum_{i=1}^{n} d\xi_i \wedge dx_i$$

is a symplectic form. Indeed, at any closed point $(x, \xi) \in T^*X$, we have

$$T_{(x,\xi)}\bigl(T^*X\bigr) = T_x X \oplus (T_x X)^*,$$

because the fiber of $p \colon T^*X \to X$ over the point x is the cotangent space $(T_x X)^*$, and because a vector space is isomorphic to its own tangent space. Under this isomorphism, the two-form σ_X corresponds exactly to the standard symplectic form in Example 11.3. In more functorial language, one can describe σ_X as follows. As with any vector bundle, the pullback $p^*\Omega^1_{X/k}$ has a tautological global section, whose image under $p^*\Omega^1_{X/k} \to \Omega^1_{T^*X/k}$ gives a one-form

$$\alpha_X \in \Gamma\bigl(T^*X, \Omega^1_{T^*X/k}\bigr).$$

In local coordinates as above, one has $\alpha_X = \sum_i \xi_i \, dx_i$. Then

$$\sigma_X = d\alpha_X \in \Gamma\bigl(T^*X, \Omega^2_{T^*X/k}\bigr)$$

is the symplectic form from above.

Let X be a nonsingular algebraic variety with a symplectic form σ. Then σ_x induces an isomorphism between the tangent space $T_x X$ and the cotangent space $(T_x X)^*$ at every closed point $x \in X$, and this allows us to convert one-forms into vector fields and vice versa. In particular, every function $f \in \Gamma(U, \mathscr{O}_X)$ determines a vector field $H_f \in \Gamma(U, \mathscr{T}_X)$, with the property that $df = \sigma(-, H_f)$ as one-forms on U. The *Poisson bracket* of two functions $f, g \in \Gamma(U, \mathscr{O}_X)$ is defined by

$$\{f, g\} = H_f(g) = dg(H_f) = \sigma(H_f, H_g) \in \Gamma(U, \mathscr{O}_X).$$

If $d\sigma = 0$, then one has $[H_f, H_g] = H_{\{f,g\}}$.

Example 11.7. In local coordinates $x_1, \dots, x_n, \xi_1, \dots, \xi_n$ on the cotangent bundle, the Hamiltonian vector field of a function f is given by

$$H_f = \sum_{i=1}^{n} \left(\frac{\partial f}{\partial \xi_i} \frac{\partial}{\partial x_i} - \frac{\partial f}{\partial x_i} \frac{\partial}{\partial \xi_i} \right),$$

and consequently, the Poisson bracket can be calculated as

$$\{f, g\} = \sum_{i=1}^{n} \left(\frac{\partial f}{\partial \xi_i} \frac{\partial g}{\partial x_i} - \frac{\partial f}{\partial x_i} \frac{\partial g}{\partial \xi_i} \right).$$

We can extend the notion of involutive (or Lagrangian or isotropic) to subvarieties of X by looking at their tangent spaces at nonsingular points. Thus a reduced algebraic subvariety $Y \subseteq X$ is called *involutive* (or *Lagrangian* or *isotropic*) if at every nonsingular closed point $x \in Y$, the tangent space $T_x Y \subseteq T_x X$ is involutive (or Lagrangian or isotropic).

Example 11.8. In the case of the cotangent bundle T^*X, the conormal bundle of a nonsingular subvariety $Z \subseteq X$ is a nonsingular Lagrangian subvariety. At a closed point $x \in Z$, the fiber of the conormal bundle consists of all those cotangent vectors in $(T_x X)^*$ that vanish on the subspace $T_x Z$. As a subspace of

$$T_{(x,\xi)}\big(T^*X\big) = T_x X \oplus (T_x X)^*,$$

the tangent space to the conormal bundle is therefore

$$T_x Z \oplus \ker\big((T_x X)^* \to (T_x Z)^*\big),$$

and this is a Lagrangian subspace by Example 11.5 from above. If we choose local coordinates $x_1, \dots, x_n$ on X such that Z is defined by $x_{k+1} = \dots = x_n = 0$, then the conormal bundle is defined by $\xi_1 = \dots = \xi_k = x_{k+1} = \dots = x_n = 0$ in the corresponding coordinates on the cotangent bundle.

The following lemma gives a way to check whether a reduced subvariety $Y \subseteq X$ is involutive by using the ideal sheaf $\mathcal{I}_Y \subseteq \mathscr{O}_X$.

Lemma 11.9. *Let X be a nonsingular algebraic variety with a symplectic form, and $Y \subseteq X$ a reduced algebraic subvariety. The following conditions are equivalent:*

(a) The subvariety Y is involutive.
(b) The ideal sheaf $\mathcal{I}_Y$ is closed under the Poisson bracket, $\{\mathcal{I}_Y, \mathcal{I}_Y\} \subseteq \mathcal{I}_Y$.

Proof. Without loss of generality, we may assume that X is affine, and that Y is the closed subvariety defined by an ideal $I \subseteq \Gamma(X, \mathscr{O}_X)$. Note that Y is assumed to be reduced. We start with a general observation. Let $x \in Y$ be a nonsingular point, and let σ_x be the symplectic form on $T_x X$. Then

$$(T_x Y)^\perp = \big\{\, v \in T_x X \;\big|\; \sigma_x(v, w) = 0 \text{ for every } w \in T_x Y \,\big\}$$

is spanned by the values at x of the Hamiltonian vector fields H_f, as f ranges over the elements of the ideal I. Indeed, since $x \in Y$ is a nonsingular point, a tangent vector $v \in T_x X$ belongs to the subspace $T_x Y$ exactly when $df(v) = 0$ for every $f \in I$. Under the Hamiltonian isomorphism, this condition becomes

$$\sigma_x(v, H_f) = df(v) = 0,$$

whence the claim.

Now let us show that $\{I, I\} \subseteq I$ implies that Y is involutive. If $x \in Y$ is a nonsingular point, we need to argue that $(T_x Y)^\perp \subseteq T_x Y$. In light of the observation from above, this amounts to saying that, for every $f, g \in I$, the function $dg(H_f)$ vanishes at the point x. But this is the case, because $dg(H_f) = H_f(g) = \{f, g\} \in I$.

For the converse, suppose that Y is involutive, so that $(T_x Y)^\perp \subseteq T_x Y$ at every nonsingular point $x \in Y$. Then we again have $\{f, g\} = dg(H_f) = 0$ at every nonsingular point of Y, and hence on all of Y because $\{f, g\}$ is a regular function and the set of nonsingular points is Zariski-open and dense in Y. Because Y is reduced, it follows that $\{f, g\} \in I$. □

Involutivity of the characteristic variety. We return to the characteristic varieties of coherent $\mathscr{D}_X$-modules. If $p\colon T^*X \to X$ is the cotangent bundle, then

$$p_* \mathscr{O}_{T^*X} \cong \operatorname{gr}^F \mathscr{D}_X,$$

and one can use this isomorphism to describe the Poisson bracket in terms of differential operators. For each $j \geq 0$, we denote by

$$\sigma_j\colon F_j \mathscr{D}_X \to \operatorname{gr}_j^F \mathscr{D}_X$$

the "principal symbol" operator. If P is a local section of $F_i\mathscr{D}_X$, and Q a local section of $F_j\mathscr{D}_X$, then their commutator $[P,Q]$ is a local section of $F_{i+j-1}\mathscr{D}_X$. One can show, using the description of the Poisson bracket in local coordinates, that

$$\{\sigma_i(P), \sigma_j(Q)\} = \sigma_{i+j-1}\big([P,Q]\big).$$

Now suppose that $\mathcal{M}$ is a coherent left $\mathscr{D}_X$-module, and that $F_\bullet\mathcal{M}$ is a good filtration. It is easy to see, using the alternative description of the Poisson bracket, that the ideal

$$\operatorname{Ann}_{\operatorname{gr}^F\mathscr{D}_X}\operatorname{gr}^F\mathcal{M} \subseteq \operatorname{gr}^F\mathscr{D}_X$$

is closed under the Poisson bracket. This is a local question, and so we may restrict everything to an affine open subset $U \subseteq X$. If we set $A = \Gamma(U, \mathscr{O}_X)$ and $R = D(A)$, we then have a finitely generated left R-module M, together with a good filtration $F_\bullet M$, such that $\operatorname{gr}^F M$ is finitely generated over $S = \operatorname{gr}^F R$. The claim is that the homogeneous ideal

$$I = I(M, F_\bullet M) = \operatorname{Ann}_S \operatorname{gr}^F M$$

is closed under the Poisson bracket on S. Suppose that we have two elements $P \in F_iR$ and $Q \in F_jR$ such that $\sigma_i(P)$ and $\sigma_j(Q)$ belong to the ideal I. Recall from Lecture 5 that this is equivalent to having

$$P \cdot F_kM \subseteq F_{i+k-1}M \quad \text{and} \quad Q \cdot F_kM \subseteq F_{j+k-1}M$$

for every $k \in \mathbb{Z}$. But then

$$[P,Q] \cdot F_kM \subseteq P \cdot F_{j+k-1}M + Q \cdot F_{i+k-1}M \subseteq F_{(i+j-1)+k-1}M,$$

and therefore $\sigma_{i+j-1}([P,Q]) \in I$. This shows that $\{I, I\} \subseteq I$.

Why does this argument not prove Theorem 11.2? The issue is that the ideal of the characteristic variety is not I itself, but $\sqrt{I}$, because the characteristic variety is by definition reduced. For non-reduced ideals, being closed under the Poisson bracket does not correspond to the geometric notion of being involutive, because all points of a nonreduced subscheme can be singular. And the fact that an ideal is closed under the Poisson bracket does not imply the same property for its radical. This is what makes Theorem 11.2 nontrivial.

Exercises.

Exercise 11.1. Let X be a nonsingular affine variety with a symplectic form. Prove the following three identities involving the Poisson bracket: for all $f, g, h \in \Gamma(X, \mathscr{O}_X)$,

$$\begin{aligned} \{f,g\} + \{g,f\} &= 0 \\ \{\{f,g\},h\} + \{\{g,h\},f\} + \{\{h,f\},g\} &= 0 \\ \{f, gh\} &= \{f,g\}h + g\{f,h\}. \end{aligned}$$

The first two identities are saying that $\Gamma(X, \mathscr{O}_X)$ is a Lie algebra under the operation $(f,g) \mapsto \{f,g\}$. The third identity is saying that $\{f, -\}$ is a derivation of $\Gamma(X, \mathscr{O}_X)$.

Exercise 11.2. Show that if $d\sigma = 0$, then one has $[H_f, H_g] = H_{\{f,g\}}$.

Exercise 11.3. Let X be a nonsingular affine variety with local coordinates $x_1, \dots, x_n$. Use the description of the Poisson bracket on T^*X to prove that

$$\{\sigma_i(P), \sigma_j(Q)\} = \sigma_{i+j-1}\big([P,Q]\big),$$

for every $P \in F_iD(A)$ and every $Q \in F_jD(A)$, where $A = \Gamma(X, \mathscr{O}_X)$.

Lecture 12: March 13

Gabber's theorem. Last time, we talked about the result that the characteristic variety $\mathrm{Ch}(\mathcal{M})$ of a coherent $\mathscr{D}_X$-module $\mathcal{M}$ is involutive (with respect to the natural symplectic structure on the cotangent bundle). We saw that the ideal

$$\mathrm{Ann}_{\mathrm{gr}^F \mathscr{D}_X} \mathrm{gr}^F \mathcal{M} \subseteq \mathrm{gr}^F \mathscr{D}_X$$

is closed under the Poisson bracket, and that Theorem 11.2 is equivalent to the radical being closed under the Poisson bracket. This is a problem in algebra, albeit a very difficult one, and there is a purely algebraic proof, due to Gabber.

In fact, Gabber works in the following more general setup. Suppose that R is a $\mathbb{Q}$-algebra, with an increasing algebra filtration $F_\bullet R$, such that the associated graded ring $S = \mathrm{gr}^F R$ is commutative and noetherian. This means that if $u \in F_i R$ and $v \in F_j R$, then their commutator $[u, v] = uv - vu \in F_{i+j-1} R$. If we again use the notation $\sigma_i \colon F_i R \to S_i$ for the "symbol" homomorphism, we can therefore define the *Poisson bracket* of two homogeneous elements of S by the formula

$$\{\sigma_i(u), \sigma_j(v)\} = \sigma_{i+j-1}\big([u, v]\big).$$

After extending this bilinearly, we obtain a Poisson bracket $\{-,-\} \colon S \otimes_{\mathbb{Q}} S \to S$, and one can check that it satisfies the same identities as the Poisson bracket on a symplectic manifold. But note that this is more general than the case $R = D(A)$, because Gabber is not assuming that S is nonsingular.

Theorem 12.1 (Gabber)**.** *Using the notation from above, suppose that M is a finitely generated R-module with a good filtration $F_\bullet M$, and consider the ideal*

$$J = \sqrt{\mathrm{Ann}_{\mathrm{gr}^F R} \mathrm{gr}^F M} \subseteq \mathrm{gr}^F R.$$

If $P \subseteq \mathrm{gr}^F R$ is minimal among prime ideals containing J, then $\{P, P\} \subseteq P$. In particular, one has $\{J, J\} \subseteq J$.

The minimal primes containing the ideal J correspond, geometrically, to the irreducible components of $\mathrm{Supp}\, \mathrm{gr}^F M$ inside the scheme $\mathrm{Spec}\, S$. So Gabber's theorem is saying that every irreducible component of the support is "involutive", in the sense that its ideal is closed under the Poisson bracket. In the case of $\mathscr{D}$-modules, this is saying that every irreducible component of the characteristic variety of a coherent $\mathscr{D}$-module is involutive.

Holonomic $\mathscr{D}$-modules. One consequence of Theorem 11.2 is that Bernstein's inequality holds for algebraic $\mathscr{D}$-modules: If X is a nonsingular algebraic variety of dimension n, and $\mathcal{M}$ a coherent $\mathscr{D}_X$-module, then either $\mathcal{M} = 0$, or every irreducible component of $\mathrm{Ch}(\mathcal{M})$ has dimension $\geq n$. As in the case of the Weyl algebra, the most important $\mathscr{D}$-modules are those for which the dimension of the characteristic variety is as small as possible.

Definition 12.2. A coherent $\mathscr{D}_X$-module $\mathcal{M}$ is called *holonomic* if $\mathcal{M} \neq 0$ and $\dim \mathrm{Ch}(\mathcal{M}) = n$, or if $\mathcal{M} = 0$.

If $\mathcal{M}$ is nonzero and holonomic, then each irreducible component of its characteristic variety has dimension n, and is therefore (by Theorem 11.2) a Lagrangian subvariety of T^*X. Since the ideal defining $\mathrm{Ch}(\mathcal{M})$ is homogeneous, these Lagrangians are moreover *conical*, that is, closed under the natural $\mathbb{G}_m$-action on T^*X by rescaling in the fiber direction. Here are some typical examples of conical Lagrangian subvarieties.

Example 12.3. If $Y \subseteq X$ is a nonsingular subvariety, then the conormal bundle $N^*_{Y|X}$ is a nonsingular Lagrangian subvariety of T^*X. Since it is a vector bundle of rank $\dim X - \dim Y$ over Y, it is clearly conical. More generally, suppose that $Y \subseteq X$ is an arbitrary reduced and irreducible subvariety. The set of nonsingular points Y_{reg} is Zariski-open and dense in Y, and so the conormal bundle $N^*_{Y_{\mathrm{reg}}|X}$ is locally closed, conical, and Lagrangian. Its Zariski closure

$$T^*_Y X = \overline{N^*_{Y_{\mathrm{reg}}|X}}$$

is therefore a conical Lagrangian subvariety of T^*X. It is called the *conormal variety* of Y in X.

In fact, every conical Lagrangian subvariety of T^*X is a conormal variety.

Proposition 12.4. *Let $W \subseteq T^*X$ be an irreducible subvariety that is conical and Lagrangian. Then $Y = p(W)$ is an irreducible subvariety of X, and $W = T^*_Y X$.*

Proof. The statement is local, and so we may assume that $X = \operatorname{Spec} A$ is affine and that $T^*X = X \times \mathbb{A}^n_k$. Since $W \subseteq X \times \mathbb{A}^n_k$ is conical, it is defined by an ideal in $A[\xi_1, \dots, \xi_n]$ that is homogeneous in the variables $\xi_1, \dots, \xi_n$. This ideal also defines a closed subvariety $\tilde{W} \subseteq X \times \mathbb{P}^{n-1}_k$, and since the projection $p_1 \colon X \times \mathbb{P}^{n-1}_k \to X$ is proper, it follows that $Y = p(W) = p_1(\tilde{W})$ is an irreducible subvariety of X. It remains to show that $W = T^*_Y X$. Since both subvarieties are irreducible of dimension n, it will be enough to show that the general point of W is contained in the conormal bundle to Y_{reg}.

Let $(x, \xi) \in W$ be a general nonsingular point. By generic smoothness, we have $x \in Y_{\mathrm{reg}}$ and the map on tangent spaces $T_{(x,\xi)} W \to T_x Y$ is surjective. Choose local coordinates $x_1, \dots, x_n$ in a neighborhood of the point x, such that Y is defined by the equations $x_{k+1} = \dots = x_n = 0$. If we again denote by $x_1, \dots, x_n, \xi_1, \dots, \xi_n$ the resulting coordinates on T^*X,then the conormal bundle to Y_{reg} is defined by the equations $\xi_1 = \dots = \xi_k = x_{k+1} = \dots = x_n = 0$. Since W is a Lagrangian subvariety, the subspace

$$T_{(x,\xi)} W \subseteq T_{(x,\xi)} \bigl(T^*X\bigr) = T_x X \oplus (T_x X)^*$$

is n-dimensional and Lagrangian. Its image under the projection to $T_x X$ is the subspace $T_x Y$. If we denote vectors in $T_x X \oplus (T_x X)^*$ by $(a_1, \dots, a_n, b_1, \dots, b_n)$, then this image is the set of vectors with $a_{k+1} = \dots = a_n = 0$. For dimension reasons, $T_{(x,\xi)} W$ must contain an $(n-k)$-dimensional space of vectors of the form $(0, \dots, 0, b_1, \dots, b_n)$, and from the Lagrangian condition, we get $b_1 = \dots = b_k = 0$.

Now we use the fact that W is conical. Since $(x, \xi) \in W$, the entire line $(x, k \cdot \xi)$ is contained in W, and so the tangent vector to the line, which is $(0, \dots, 0, \xi_1, \dots, \xi_n)$, must belong to $T_{(x,\xi)} W$. But as we saw, this implies that $\xi_1 = \dots = \xi_k = 0$, and so (x, ξ) lies on the conormal bundle to Y_{reg}. Since (x, ξ) was a general point of W, we deduce that $W \subseteq T^*_Y X$, which suffices to conclude the proof. □

This proposition has interesting implications for holonomic $\mathscr{D}$-modules. Suppose that $\mathcal{M}$ is a nonzero holonomic $\mathscr{D}_X$-module. Its characteristic variety is a finite union of conical Lagrangian subvarieties, and so there are finitely many irreducible subvarieties $Y_1, \dots, Y_m \subseteq X$, without loss of generality distinct, such that

$$\operatorname{Ch}(\mathcal{M}) = \bigcup_{i=1}^m T^*_{Y_i} X.$$

Now there are two possibilities. If say $Y_1 = X$, then $U = X \setminus (Y_2 \cup \dots \cup Y_m)$ is a dense Zariski-open subset, and the restriction of $\mathcal{M}$ to U has its characteristic variety equal to the zero section. By Proposition 10.12, it follows that $\mathcal{M}\big|_U$ is locally free of finite rank, and therefore a vector bundle with integrable connection. The connection acquires some kind of singularities at the remaining subvarieties $Y_2, \dots, Y_n$. The other possibility is that $Y_1, \dots, Y_n \neq X$. In that case, the restriction of $\mathcal{M}$ to $X \setminus (Y_1 \cup \cdots \cup Y_n)$ is trivial, which says that $\mathcal{M}$ is supported on the union $Y_1 \cup \cdots \cup Y_n$. Either way, $\mathcal{M}$ is generically a vector bundle with integrable connection.

Holonomic $\mathscr{D}$-modules and duality. Our earlier results about duality for holonomic modules still hold in this context; indeed, the assumptions we made in Lecture 6 apply to the case $R = D(A)$. In general, if $\mathcal{M}$ is a coherent left (or right) $\mathscr{D}_X$-module, then each

$$\mathcal{E}xt^j_{\mathscr{D}_X}(\mathcal{M}, \mathscr{D}_X)$$

is again a coherent right (or left) $\mathscr{D}_X$-module. On an affine open subset $U \subseteq X$ with $A = \Gamma(U, \mathscr{O}_X)$, the corresponding $D(A)$-module is of course $\operatorname{Ext}^j_{D(A)}\big(M, D(A)\big)$, where $M = \Gamma(U, \mathcal{M})$. One then has

$$\mathcal{E}xt^j_{\mathscr{D}_X}(\mathcal{M}, \mathscr{D}_X) = 0 \quad \text{for } j \geq n+1,$$

as well as the useful identity

$$\min\Big\{ j \geq 0 \;\Big|\; \mathcal{E}xt^j_{\mathscr{D}_X}(\mathcal{M}, \mathscr{D}_X) \neq 0 \Big\} + \dim \operatorname{Ch}(\mathcal{M}) = 2n.$$

If $\mathcal{M}$ is a nonzero holonomic $\mathscr{D}_X$-module, then $\mathcal{E}xt^j_{\mathscr{D}_X}(\mathcal{M}, \mathscr{D}_X) = 0$ for every $j \neq n$, and one can again define the *holonomic dual* by

$$\mathcal{M}^* = \mathcal{E}xt^n_{\mathscr{D}_X}(\mathcal{M}, \mathscr{D}_X).$$

As before, one has $(\mathcal{M}^*)^* \cong \mathcal{M}$, and $\operatorname{Ch}(\mathcal{M}^*) = \operatorname{Ch}(\mathcal{M})$. The holonomic dual is again an exact contravariant functor from the category of left (or right) holonomic $\mathscr{D}_X$-modules to the category of right (or left) holonomic $\mathscr{D}_X$-modules.

Direct images under closed embeddings. In the next few lectures, we are going to look at various operations on algebraic $\mathscr{D}$-modules, such as pushing forward or pulling back along a morphism of algebraic varieties. This will also give us many new examples of $\mathscr{D}$-modules. We will be especially interested in the effect of these functors on holonomic $\mathscr{D}$-modules. Things are somewhat similar to the case of coherent sheaves, formally, but there are also some interesting differences. Let us start with the simplest case, namely pushing forward along a closed embedding.

Example 12.5. Consider the closed embedding $i\colon \mathbb{A}^{n-1}_k \to \mathbb{A}^n_k$ defined by the equation $x_n = 0$. If $\mathcal{M}$ is a $\mathscr{D}$-module on $\mathbb{A}^{n-1}_k$, then its pushforward $i_*\mathcal{M}$ is not a $\mathscr{D}$-module on $\mathbb{A}^n_k$. The problem is that $x_1, \dots, x_n$ and $\partial_1, \dots, \partial_{n-1}$ act in a natural way on $i_*\mathcal{M}$, but we don't know what to do with ∂_n. In terms of rings and modules, the closed embedding corresponds to the quotient morphism $k[x_1, \dots, x_{n-1}, x_n] \to k[x_1, \dots, x_{n-1}]$, and the $\mathscr{D}$-module to a module M over the Weyl algebra $A_{n-1}(k)$. We can consider M as a module over $k[x_1, \dots, x_n]$, with x_n acting trivially, but we cannot let ∂_n act trivially this would violate the commutator relation $[\partial_n, x_n] = 1$.

Suppose that $i\colon X \to Y$ is a closed embedding between two nonsingular algebraic varieties, and $\mathcal{M}$ an algebraic $\mathscr{D}_X$-module. For the same reason as above, $i_*\mathcal{M}$ is not in general a $\mathscr{D}_Y$-module. To motivate the correct definition, let us first look at the example of distributions.

Example 12.6. Consider the closed embedding

$$i\colon \mathbb{R}^k \to \mathbb{R}^n, \quad i(x_1, \dots, x_k) = (x_1, \dots, x_k, 0, \dots, 0).$$

Suppose that we have a distribution D on $\mathbb{R}^k$; recall that D is a continuous linear functional on the space of compactly supported smooth functions $C^\infty_0(\mathbb{R}^k)$, and that $\langle D, \varphi \rangle$ denotes the real number obtained by evaluating D on a test function φ. The pushforward distribution i_*D is defined in the obvious way:

$$\langle i_*D, \psi \rangle = \langle D, \psi\big|_{\mathbb{R}^k} \rangle,$$

for any $\psi \in C_0^\infty(\mathbb{R}^n)$. The point is of course that we know how to pull back functions. Now suppose that D satisfies a system of partial differential equations. Can we figure out the partial differential equations satisfied by $i_* D$?

Recall that the Weyl algebra $A_k(\mathbb{R})$ acts on the space of distributions by formal integration by parts: if $\varphi \in C_0^\infty(\mathbb{R}^k)$ and $P \in A_k(\mathbb{R})$, then

$$\langle D \cdot P, \varphi \rangle = \langle D, P\varphi \rangle.$$

Therefore D determines a right ideal

$$I(D) = \left\{ P \in A_k(\mathbb{R}) \;\middle|\; D \cdot P = 0 \right\} \subseteq A_k(\mathbb{R}),$$

and also a right $A_k(n)$-module $A_k(n)/I(D)$. In these terms, we are trying to find the right ideal $I(i_* D)$ from $I(D)$. This is actually fairly easy.

First, the functions $x_{k+1}, \dots, x_n$ vanish on $\mathbb{R}^k$, and so every differential operator of the form $Q = x_{k+1}Q_{k+1} + \dots + x_n Q_n \in A_n(\mathbb{R})$ annihilates $i_* D$, because

$$\left\langle i_* D \cdot Q, \psi \right\rangle = \sum_{j=k+1}^{n} \left\langle i_* D \cdot x_j Q_j, \psi \right\rangle = \sum_{j=k+1}^{n} \left\langle D, x_j Q_j \psi\big|_{\mathbb{R}^k} \right\rangle = 0.$$

We can write any $Q \in A_n(\mathbb{R})$ in the form

$$Q = x_{k+1}Q_{k+1} + \dots + x_n Q_n + \sum_{\alpha \in \mathbb{N}^{n-k}} P_\alpha \partial_{k+1}^{\alpha_{k+1}} \cdots \partial_n^{\alpha_n}$$

where $P_\alpha \in A_k(\mathbb{R})$ only involves $x_1, \dots, x_k, \partial_1, \dots, \partial_k$. Suppose that $Q \in I(i_* D)$. If we act on a test function of the form $\varphi\eta$, with $\varphi \in C_0^\infty(\mathbb{R}^k)$ and $\eta \in C_0^\infty(\mathbb{R}^{n-k})$, we obtain

$$\left\langle i_* D \cdot Q, \varphi\eta \right\rangle = \sum_{\alpha \in \mathbb{N}^{n-k}} \frac{\partial^{\alpha_{k+1} + \dots + \alpha_n} \eta}{\partial x_{k+1}^{\alpha_{k+1}} \cdots \partial x_n^{\alpha_n}}(0) \cdot \left\langle D, P_\alpha \varphi \right\rangle.$$

By choosing η appropriately, we can pick out the individual terms, and so

$$0 = \langle D, P_\alpha \varphi \rangle = \langle D \cdot P_\alpha, \varphi \rangle$$

for every $\alpha \in \mathbb{N}^{n-k}$ and every $\varphi \in C_0^\infty(\mathbb{R}^k)$. In other words, each P_α belongs to $I(D)$. It is easy to see that the converse is also true, and so we conclude that

$$I(i_* D) = (x_{k+1}, \dots, x_n) A_n(\mathbb{R}) + I(D) A_n(\mathbb{R}).$$

Here is another way to put this. Remembering that right (and left) ideals in the Weyl algebra are finitely generated, we have $I(D) = (P_1, \dots, P_r) A_k(\mathbb{R})$, and so the right $A_k(\mathbb{R})$-module determined by the distribution D is

$$A_k(\mathbb{R}) / (P_1, \dots, P_r) A_k(\mathbb{R}).$$

Then the right $A_n(\mathbb{R})$-module determined by the distribution $i_* D$ is

$$A_n(\mathbb{R}) / (P_1, \dots, P_r, x_{k+1}, \dots, x_n) A_n(\mathbb{R}).$$

This is much larger than the other module, but has a natural action by $A_n(\mathbb{R})$.

The example suggest that pushing forward works naturally for right $\mathscr{D}$-modules. The reason is that distributions give rise to right $\mathscr{D}$-modules, whereas functions give rise to left $\mathscr{D}$-modules, and one can push forward distributions, but not functions. It also suggests how to define the pushforward, at least in the special case of modules over the Weyl algebra.

The transfer module. Let me now show you the actual definition. Suppose that $i\colon X \to Y$ is a closed embedding between two nonsingular algebraic varieties; since X and Y are both nonsingular, X is locally a complete intersection in Y. We will see next time that

$$\mathscr{D}_{X\to Y} = \mathscr{O}_X \otimes_{i^{-1}\mathscr{O}_Y} i^{-1}\mathscr{D}_Y$$

is a $(\mathscr{D}_X, i^{-1}\mathscr{D}_Y)$-bimodule, which is to say that it has both a left action by $\mathscr{D}_X$ and a right action by $i^{-1}\mathscr{D}_Y$, and the two actions commute. The right action by $i^{-1}\mathscr{D}_Y$ is the obvious one; the left action by $\mathscr{D}_X$ is less obvious and involves both factors in the tensor product. Given a right $\mathscr{D}_X$-module $\mathcal{M}$, one then defines its pushforward as

$$i_+\mathcal{M} = i_*\bigl(\mathcal{M} \otimes_{\mathscr{D}_X} \mathscr{D}_{X\to Y}\bigr);$$

this becomes a right $\mathscr{D}_Y$-module through the natural morphism $\mathscr{D}_Y \to i_* i^{-1}\mathscr{D}_Y$. We will see next time that, in local coordinates, this definition agrees with what happens for distributions.

Exercises.

Exercise 12.1. Let $\mathcal{M}$ be a left $\mathscr{D}_X$-module and $\mathcal{N}$ a right $\mathscr{D}_X$-module. Show that the tensor product $\mathcal{N} \otimes_{\mathscr{O}_X} \mathcal{M}$ is naturally a right $\mathscr{D}_X$-module.

Exercise 12.2. Recall that the canonical line bundle ω_X is a right $\mathscr{D}_X$-module. Show that the tensor product $\mathscr{D}_X^\omega = \omega_X \otimes_{\mathscr{O}_X} \mathscr{D}_X$ is a right $\mathscr{D}_X$-module in *two* different ways. Show that the two right $\mathscr{D}_X$-module structures commute with each other, and that there is an automorphism of $\mathscr{D}_X^\omega$ that interchanges them.

Exercise 12.3. The previous exercise gives a way to convert left $\mathscr{D}$-modules into right $\mathscr{D}$-modules and back. Show that if $\mathcal{M}$ is a left $\mathscr{D}_X$-module, then

$$\mathscr{D}_X^\omega \otimes_{\mathscr{D}_X} \mathcal{M}$$

is a right $\mathscr{D}_X$-module; here one right $\mathscr{D}_X$-module structure on $\mathscr{D}_X^\omega$ is used to define the tensor product, and the other one is used to turn the tensor product into a right $\mathscr{D}_X$-module. Conversely, show that if $\mathcal{N}$ is a right $\mathscr{D}_X$-module, then

$$\mathcal{H}om_{\mathscr{D}_X}\bigl(\mathscr{D}_X^\omega, \mathcal{N}\bigr)$$

is a left $\mathscr{D}_X$-module; here one right $\mathscr{D}_X$-module structure on $\mathscr{D}_X^\omega$ is used to define $\mathcal{H}om_{\mathscr{D}_X}$, and the other one is used to turn $\mathcal{H}om_{\mathscr{D}_X}$ into a left $\mathscr{D}_X$-module. Finally, show that the obvious morphism

$$\mathcal{M} \to \mathcal{H}om_{\mathscr{D}_X}\bigl(\mathscr{D}_X^\omega, \mathscr{D}_X^\omega \otimes_{\mathscr{D}_X} \mathcal{M}\bigr)$$

is an isomorphism of left $\mathscr{D}_X$-modules.

Lecture 13: March 25

The transfer module. Last time, we looked at the example of distributions to understand what the pushforward of an algebraic $\mathscr{D}$-module under a closed embedding should be. In the case of $i \colon \mathbb{R}^k \hookrightarrow \mathbb{R}^n$, defined by $i(x_1, \dots, x_k) = (x_1, \dots, x_k, 0, \dots, 0)$, we concluded that the pushforward of a right $A_k(\mathbb{R})$-module of the form

$$A_k(\mathbb{R})/(P_1, \dots, P_m)A_k(\mathbb{R})$$

should be the right $A_n(\mathbb{R})$-module

$$A_n(\mathbb{R})/(P_1, \dots, P_m, x_{k+1}, \dots, x_n)A_n(\mathbb{R}).$$

Let me know explain how to define the pushforward under a closed embedding in general. Let $i \colon Y \hookrightarrow X$ be a closed embedding, with X nonsingular of dimension n and Y nonsingular of dimension r. The definition uses the *transfer module*

$$\mathscr{D}_{Y \to X} = \mathscr{O}_Y \otimes_{i^{-1}\mathscr{O}_X} i^{-1}\mathscr{D}_X,$$

which is a $(\mathscr{D}_Y, i^{-1}\mathscr{D}_X)$-bimodule. In other words, $\mathscr{D}_{Y \to X}$ is both a left $\mathscr{D}_Y$-module and a right $i^{-1}\mathscr{D}_X$-module, and the two structures commute with each other. The right $i^{-1}\mathscr{D}_X$-module structure is the obvious one, induced by right multiplication on the second factor of the tensor product. The left $\mathscr{D}_Y$-module structure is less obvious, and involves both factors. Remember that since X and Y are both nonsingular, we have a short exact sequence

$$0 \to \mathscr{T}_Y \xrightarrow{\delta_i} i^*\mathscr{T}_X = \mathscr{O}_Y \otimes_{i^{-1}\mathscr{O}_X} i^{-1}\mathscr{T}_X \to N_{Y|X} \to 0,$$

where $N_{Y|X}$ is the normal bundle of Y in X, a locally free $\mathscr{O}_Y$-module of rank $\dim X - \dim Y$. Now $\mathscr{T}_Y$ acts on $\mathscr{D}_{Y \to X}$ as follows:

$$\theta \cdot (f \otimes P) = \theta(f) \otimes P + f \cdot \delta_i(\theta) \cdot (1 \otimes P),$$

where $\theta \in \mathscr{T}_Y$, $f \in \mathscr{O}_Y$, and $P \in i^{-1}\mathscr{D}_X$ are local sections. I will leave it as an exercise to show that this extends to a left $\mathscr{D}_Y$-module structure.

Example 13.1. Let us write out everything in local coordinates. Choose local coordinates $x_1, \dots, x_n$ on X, in such a way that Y is defined by the equations $x_{r+1} = \dots = x_n = 0$. We write $\partial_1, \dots, \partial_n$ for the corresponding vector fields on X; then $y_1 = x_1, \dots, y_r = x_r$ are local coordinates on Y, with vector fields $\partial_{y_1}, \dots, \partial_{y_r}$. The morphism $\delta_i \colon \mathscr{T}_Y \to i^*\mathscr{T}_X$ sends ∂_{y_j} to $1 \otimes \partial_j$, and so we get

$$\partial_{y_j} \cdot (f \otimes P) = \partial_{y_j} f \otimes P + f \otimes \partial_j P,$$

where $\partial_j P$ is the product in $\mathscr{D}_X$.

Lemma 13.2. *The transfer module $\mathscr{D}_{Y \to X}$ contains a copy of $\mathscr{D}_Y$ and is a locally free left $\mathscr{D}_Y$-module of infinite rank.*

Proof. Since $\mathscr{D}_{Y \to X} = \mathscr{O}_Y \otimes_{i^{-1}\mathscr{O}_X} i^{-1}\mathscr{D}_X$, the transfer module has a global section given by $1 \otimes 1$. This embeds a copy of $\mathscr{D}_Y$ into $\mathscr{D}_{Y \to X}$, by letting $\mathscr{D}_Y$ act on $1 \otimes 1$. In local coordinates as above, we have

$$\partial_{y_j} \cdot (1 \otimes 1) = 1 \otimes \partial_j.$$

More generally, for any differential operator $Q = \sum_\alpha f_\alpha \partial_y^\alpha$ on Y, we get

$$Q \cdot (1 \otimes 1) = \sum_\alpha f_\alpha \otimes \partial^\alpha = \sum_\alpha f_\alpha \otimes \partial_1^{\alpha_1} \cdots \partial_r^{\alpha_r}.$$

This shows that the resulting morphism $\mathscr{D}_Y \to \mathscr{D}_{Y \to X}$ is injective.

Since we are working locally, every differential operator P on X can be written uniquely in the form $P = \sum_\beta g_\beta \partial^\beta$, where $\beta \in \mathbb{N}^n$. By restriction, each $g_\beta \in \Gamma(X, \mathscr{O}_X)$ defines an element $\bar{g}_\beta \in \Gamma(Y, \mathscr{O}_Y)$, and we have

$$f \otimes P = \sum_\beta f \otimes g_\beta \partial^\beta = \sum_{\beta_{r+1},\dots,\beta_n} \left(\sum_{\beta_1,\dots,\beta_r} f \bar{g}_\beta \otimes \partial_1^{\beta_1} \cdots \partial_r^{\beta_r} \right) \cdot \partial_{r+1}^{\beta_{r+1}} \cdots \partial_n^{\beta_n}.$$

This shows that the morphism $\mathscr{D}_Y \otimes_k k[\partial_{r+1}, \dots, \partial_n] \to \mathscr{D}_{Y\to X}$, given by multiplication, is an isomorphism. More formally, consider the subalgebra

$$\mathscr{D}_X^Y = \bigoplus_{\alpha \in \mathbb{N}^r} \mathscr{O}_X \cdot \partial_1^{\alpha_1} \cdots \partial_r^{\alpha_r} \subseteq \mathscr{D}_X.$$

Then we have $\mathscr{D}_X \cong \mathscr{D}_X^Y \otimes_k k[\partial_{r+1}, \dots, \partial_n]$, and therefore

$$\mathscr{D}_{Y\to X} \cong (\mathscr{O}_Y \otimes_{i^{-1}\mathscr{O}_X} i^{-1}\mathscr{D}_X^Y) \otimes_k k[\partial_{r+1}, \dots, \partial_n],$$

and the discussion above shows that $\mathscr{O}_Y \otimes_{i^{-1}\mathscr{O}_X} i^{-1}\mathscr{D}_X^Y$ identifies with the copy of $\mathscr{D}_Y$ inside $\mathscr{D}_{Y\to X}$. ☐

Definition 13.3. The *pushforward* of a right $\mathscr{D}_Y$-module is defined as

$$i_+\mathcal{M} = i_*\big(\mathcal{M} \otimes_{\mathscr{D}_Y} \mathscr{D}_{Y\to X}\big);$$

it becomes a right $\mathscr{D}_X$-module through the morphism $\mathscr{D}_X \to i_* i^{-1}\mathscr{D}_X$.

Note that the pushforward is an exact functor, in the sense that if

$$0 \to \mathcal{M}' \to \mathcal{M} \to \mathcal{M}'' \to 0$$

is a short exact sequence of right $\mathscr{D}_Y$-modules, then

$$0 \to i_+\mathcal{M}' \to i_+\mathcal{M} \to i_+\mathcal{M}'' \to 0$$

is a short exact sequence of right $\mathscr{D}_X$-modules. The reason is that the tensor product over $\mathscr{D}_Y$ is exact (because $\mathscr{D}_{Y\to X}$ is locally free as a left $\mathscr{D}_Y$-module) and that i_* is exact (because $i \colon Y \hookrightarrow X$ is a closed embedding).

The inclusion $\mathscr{D}_Y \hookrightarrow \mathscr{D}_{Y\to X}$ induces an inclusion of $i_*\mathcal{M}$ into the pushforward $i_+\mathcal{M}$. In local coordinates as in the lemma, we get

$$i_+\mathcal{M} \cong i_*\mathcal{M} \otimes_k k[\partial_{r+1}, \dots, \partial_n],$$

and so the problem that $i_*\mathcal{M}$ is not a $\mathscr{D}_X$-module is solved by simply creating a new copy of $i_*\mathcal{M}$ for every monomial in $\partial_{r+1}, \dots, \partial_n$. Note the the submodule $i_*\mathcal{M}$ is annihilated by the equations $x_{r+1}, \dots, x_n$ of Y, but because of the relation $[\partial_j, x_j] = 1$, this is no longer true for $i_+\mathcal{M}$. In general, every section of $i_*\mathcal{M}$ is annihilated by the ideal sheaf $\mathcal{I}_Y \subseteq \mathscr{O}_X$, and every section of $i_+\mathcal{M}$ is annihilated by some power of $\mathcal{I}_Y$.

Example 13.4. Let's compute the pushforward of $\mathscr{D}_Y$. We have

$$i_+\mathscr{D}_Y = i_*\big(\mathscr{D}_Y \otimes_{\mathscr{D}_Y} \mathscr{D}_{Y\to X}\big) = i_*\mathscr{D}_{Y\to X} = i_*\big(\mathscr{O}_Y \otimes_{i^{-1}\mathscr{O}_X} i^{-1}\mathscr{D}_X\big).$$

The natural morphism $\mathscr{D}_X \to i_+\mathscr{D}_Y$, given by sending $P \in \mathscr{D}_X$ to $1 \otimes P$, is clearly surjective, and its kernel is exactly the right ideal $\mathcal{I}_Y\mathscr{D}_X$. Thus $i_+\mathscr{D}_Y \cong \mathscr{D}_X/\mathcal{I}_Y\mathscr{D}_X$.

Example 13.5. Let us compare the definition with the calculation from last time. Consider the closed embedding $i \colon \mathbb{A}_k^r \hookrightarrow \mathbb{A}_k^n$, corresponding to the quotient morphism $k[x_1, \dots, x_n] \to k[x_1, \dots, x_r]$. Let's compute the pushforward of the right A_r-module $M = A_r/(P_1, \dots, P_m)A_r$. By the previous example, the pushforward of A_r itself is given by $A_n/(x_{r+1}, \dots, x_n)A_n$. Using the presentation

$$A_r^{\oplus m} \xrightarrow{(P_1,\dots,P_m)} A_r$$

for M and the exactness of i_+, we see that the pushforward of M is the cokernel of the induced morphism

$$\big(A_n/(x_{r+1},\dots,x_n)A_n\big)^{\oplus m} \longrightarrow A_n/(x_{r+1},\dots,x_n)A_n.$$

One then checks that for the endomorphism of A_r given by left multiplication by a differential operator $P \in A_r$, the induced endomorphism of $A_n/(x_{r+1},\dots,x_n)A_n$ is still left multiplication by P. Thus that the pushforward of M is isomorphic to

$$A_n/(P_1,\dots,P_m,x_{r+1},\dots,x_n)A_n,$$

in agreement with the calculation we did for distributions last time.

Coherence and characteristic variety. Now let us study the effect of the pushforward functor on coherence and on the characteristic variety.

Lemma 13.6. *If $\mathcal{M}$ is a coherent right $\mathscr{D}_Y$-module, then $i_+\mathcal{M}$ is a coherent right $\mathscr{D}_X$-module.*

Proof. Since $\mathcal{M}$ is coherent over $\mathscr{D}_Y$, we can find a coherent $\mathscr{O}_Y$-module $\mathscr{F} \subseteq \mathcal{M}$ such that $\mathscr{F} \cdot \mathscr{D}_Y = \mathcal{M}$. Using the embedding of $i_*\mathcal{M}$ into $i_+\mathcal{M}$, the coherent $\mathscr{O}_X$-module $i_*\mathscr{F}$ embeds into $i_+\mathcal{M}$, and one checks in local coordinates that it generates $i_+\mathcal{M}$ as a right $\mathscr{D}_X$-module. Therefore $i_+\mathcal{M}$ is coherent. □

To understand the effect of pushing forward on the characteristic variety, we need to investigate in more detail what happens to a good filtration. Suppose that $\mathcal{M}$ is a coherent right $\mathscr{D}_Y$-module, and choose a good filtration $F_\bullet\mathcal{M}$, so that each $F_j\mathcal{M}$ is a coherent $\mathscr{O}_Y$-module. Using the embedding of $i_*\mathcal{M}$ into the pushforward $i_+\mathcal{M} = i_*(\mathcal{M} \otimes_{\mathscr{D}_Y} \mathscr{D}_{Y\to X})$, each $F_j\mathcal{M}$ therefore defines a subsheaf $i_*(F_j\mathcal{M}) \subseteq i_+\mathcal{M}$. To get a filtration that is compatible with the $\mathscr{D}_X$-module structure, we now define

$$F_j(i_+\mathcal{M}) = i_*(F_j\mathcal{M}) + i_*(F_{j-1}\mathcal{M}) \cdot F_1\mathscr{D}_X + i_*(F_{j-2}\mathcal{M}) \cdot F_2\mathscr{D}_X + \cdots \tag{13.7}$$

Since $F_j\mathcal{M} = 0$ for $j \ll 0$, there are only finitely many terms, and so each $F_j(i_+\mathcal{M})$ is a coherent $\mathscr{O}_X$-module. To check that this gives a good filtration, we work in local coordinates. So let $U \subseteq X$ be an affine open subset, with local coordinates $x_1,\dots,x_n \in A = \Gamma(U,\mathscr{O}_X)$, such that Y is defined by the ideal $I = (x_{r+1},\dots,x_n)$. Set $B = A/I$, and let $M = \Gamma(U \cap Y, \mathscr{D}_Y)$; this is a right $D(B)$-module, of course, but we may also consider it as an A-module on which I acts trivially. From our earlier discussion,

$$\Gamma(U, i_+\mathcal{M}) \cong M \otimes_k k[\partial_{r+1},\dots,\partial_n] \underset{\text{def}}{=} \tilde{M},$$

and the above filtration is given by

$$F_j\tilde{M} = F_jM \otimes 1 + (F_{j-1}M \otimes 1) \cdot F_1D(A) + (F_{j-2}M \otimes 1) \cdot F_2D(A) + \cdots.$$

We can write this in more compact notation as

$$F_j\tilde{M} = \sum_\alpha F_{j-|\alpha|}M \otimes \partial_{r+1}^{\alpha_{r+1}} \cdots \partial_n^{\alpha_n}.$$

The associated graded module is therefore given by

$$\operatorname{gr}^F \tilde{M} = \operatorname{gr}^F M \otimes_k k[\partial_{r+1},\dots,\partial_n], \tag{13.8}$$

with the grading in which every ∂_j has degree 1. Concretely,

$$\operatorname{gr}^F_j \tilde{M} = \bigoplus_\alpha \operatorname{gr}^F_{j-|\alpha|} M \otimes \partial_{r+1}^{\alpha_{r+1}} \cdots \partial_n^{\alpha_n}.$$

Now $\operatorname{gr}^F \tilde{M}$ is a graded module over $\operatorname{gr}^F D(A) \cong A[\partial_1, \dots, \partial_n]$. Let us describe the module structure in more detail. Recall that $\operatorname{gr}^F M$ is a finitely generated graded module over $\operatorname{gr}^F D(B) \cong B[\partial_1, \dots, \partial_r]$. From (13.8), we get

$$\operatorname{gr}^F \tilde{M} \cong \operatorname{gr}^F M \otimes_{B[\partial_1,\dots,\partial_r]} B[\partial_1, \dots, \partial_n],$$

and since $I \subseteq A$ acts trivially on $\operatorname{gr}^F \tilde{M}$ by construction, this is actually an isomorphism of $A[\partial_1, \dots, \partial_n]$-modules. Since $\operatorname{gr}^F M$ is finitely generated over $B[\partial_1, \dots, \partial_r]$, this shows that $\operatorname{gr}^F \tilde{M}$ is finitely generated over $\operatorname{gr}^F D(A)$, and so the filtration in (13.7) is indeed good.

The calculation we have just done has the following geometric interpretation. The closed embedding $i\colon Y \hookrightarrow X$ gives rise to two morphisms between the cotangent bundles of X and Y:

$$\begin{tikzcd} Y \times_X T^*X \arrow[r, "di"] \arrow[d, "p_2"] & T^*Y \\ T^*X \end{tikzcd}$$

Here the morphism $di\colon Y \times_X T^*X \to T^*Y$ corresponds to the pullback morphism $i^*\Omega^1_{X/k} \to \Omega^1_{Y/k}$ between Kähler differentials, and is therefore a morphism of vector bundles, with kernel the conormal bundle of Y in X. In particular, it is a smooth morphism of relative dimension $\dim X - \dim Y$. If we denote by $\widetilde{\operatorname{gr}^F \mathcal{M}}$ the coherent $\mathscr{O}_{T^*Y}$-module corresponding to $\operatorname{gr}^F \mathcal{M}$, then the above isomorphism takes the form

$$\widetilde{\operatorname{gr}^F(i_+\mathcal{M})} \cong (p_2)_* di^* \widetilde{\operatorname{gr}^F \mathcal{M}}. \tag{13.9}$$

The reason is that, in local coordinates, the morphisms of k-algebras corresponding to the morphisms between cotangent bundles are

$$\begin{tikzcd} B[\partial_1, \dots, \partial_n] & B[\partial_1, \dots, \partial_r] \arrow[l] \\ A[\partial_1, \dots, \partial_n] \arrow[u] \end{tikzcd}$$

and so pulling back via di corresponds to tensoring the $B[\partial_1, \dots, \partial_r$-module $\operatorname{gr}^F M$ by $B[\partial_1, \dots, \partial_n]$, and pushing forward via p_2 corresponds to consider the result as a module over $A[\partial_1, \dots, \partial_n]$. The calculation from above shows that the result is isomorphic to $\operatorname{gr}^F \tilde{M}$. Let us summarize the conclusion.

Proposition 13.10. *Let $i\colon Y \to X$ be a closed embedding, and $\mathcal{M}$ a coherent right $\mathscr{D}_Y$-module. Then the pushforward $i_+\mathcal{M}$ satisfies*

$$\operatorname{Ch}(i_+\mathcal{M}) = p_2\bigl(di^{-1} \operatorname{Ch}(\mathcal{M})\bigr),$$

and so $\dim \operatorname{Ch}(i_+\mathcal{M}) = \dim \operatorname{Ch}(\mathcal{M}) + \dim X - \dim Y$.

Proof. Since the characteristic variety of $\mathcal{M}$ is the support of $\widetilde{\operatorname{gr}^F \mathcal{M}}$, the formula for the characteristic variety is an immediate consequence of (13.9). Because di is a smooth morphism of relative dimension $\dim Y - \dim X$, whereas p_2 is a closed embedding, the asserted formula for the dimension of the characteristic variety follows from this. □

The formula for the characteristic variety of the pushforward has several useful consequences. Firstly, it implies that $\mathcal{M}$ is holonomic if and only if $i_+\mathcal{M}$ is holonomic. The reason is of course that $\dim \operatorname{Ch}(i_+\mathcal{M}) - \dim X = \dim \operatorname{Ch}(\mathcal{M}) - \dim Y$. Secondly, it gives another proof for Bernstein's inequality $\dim \operatorname{Ch}(\mathcal{M}) \geq \dim X$, independently of symplectic geometry. Recall that, back in Lecture 3, we proved Bernstein's inequality for finitely generated modules over the Weyl algebra, by looking at Hilbert functions. We can now deduce from this that Bernstein's inequality holds for all algebraic $\mathscr{D}$-modules. Suppose then that $\mathcal{M}$ is a finitely generated right $\mathscr{D}_X$-module, where X is a nonsingular algebraic variety. Since the question is local, we may assume that X is affine. We can then choose a closed embedding $i \colon X \hookrightarrow \mathbb{A}_k^m$ into affine space. By Proposition 13.10, we have

$$\dim \mathrm{Ch}(\mathcal{M}) - \dim X = \dim \mathrm{Ch}(i_+\mathcal{M}) - m \geq 0,$$

where the inequality is a consequence of Bernstein's inequality for the Weyl algebra. Thus $\dim \mathrm{Ch}(\mathcal{M}) \geq \dim X$.

Kashiwara's equivalence. Let $i \colon Y \hookrightarrow X$ be a closed embedding. We had already noted that

$$i_+ \colon (\text{coherent right } \mathscr{D}_Y\text{-modules}) \to (\text{coherent right } \mathscr{D}_X\text{-modules})$$

is an exact functor. One of the first results that Kashiwara proved in his thesis is a description of the image of this functor. Clearly, every right $\mathscr{D}_X$-module of the form $i_+\mathcal{M}$ is supported on Y, in the following sense.

Definition 13.11. The *support* of a coherent right $\mathscr{D}_X$-module $\mathcal{N}$ is defined as

$$\operatorname{Supp} \mathcal{N} = p\big(\mathrm{Ch}(\mathcal{N})\big),$$

where $p \colon T^*X \to X$ is the projection.

Since $\mathrm{Ch}(\mathcal{N})$ is conical, its image in X is always a closed algebraic subset. It follows that $\operatorname{Supp} \mathcal{N}$ is the complement of the largest Zariski-open subset $U \subseteq X$ such that $\mathcal{N}\big|_U$ is trivial. Since every section of $i_+\mathcal{M}$ is annihilated by a sufficiently large power of $\mathcal{I}_Y$, it is clear that $\operatorname{Supp}(i_+\mathcal{M}) \subseteq Y$. (This allows follows from Proposition 13.10, of course.)

Theorem 13.12 (Kashiwara's equivalence)**.** *The functor i_+ is an equivalence of categories between the category of (coherent) right $\mathscr{D}_Y$-modules and the category of (coherent) right $\mathscr{D}_X$-modules with support contained in Y.*

We will give the proof next time.

Exercises.

Exercise 13.1. Suppose that $X = \operatorname{Spec} A$ is affine, and that Y is the closed subscheme defined by an ideal $I \subseteq A$, so that $Y = \operatorname{Spec} B$ for $B = A/I$. Show that the morphism $\mathrm{Der}_k(B) \to B \otimes_A \mathrm{Der}_k(A)$ puts a left $D(B)$-module structure on $B \otimes_A D(A)$, and that it commutes with the natural right $D(A)$-module structure.

Exercise 13.2. Let $X = \operatorname{Spec} A$, with local coordinates $x_1, \dots, x_n \in A$, and let $I = (x_{r+1}, \dots, x_n)$. Show that if M is a finitely generated right $D(B)$-module, where $B = A/I$, then $M \otimes_k k[\partial_{r+1}, \dots, x_n]$ is finitely generated as a right $D(A)$-module.

Exercise 13.3. Let M be a graded $B[\partial_1, \dots, \partial_r]$-module. Show that

$$\begin{aligned} &\operatorname{Ann}_{A[\partial_1,\dots,\partial_n]}\big(M \otimes_k k[\partial_{r+1}, \dots, \partial_n]\big) \\ &\qquad\qquad = (x_{r+1}, \dots, x_n) + A[\partial_1, \dots, \partial_n] \cdot \operatorname{Ann}_{B[\partial_1,\dots,\partial_r]} M, \end{aligned}$$

as ideals in $A[\partial_1, \dots, \partial_n]$.

Lecture 14: March 27

Kashiwara's equivalence. Let us start by giving the proof of Kashiwara's equivalence from last time. Here is the statement again.

Theorem (Kashiwara's equivalence)**.** *Let $i\colon Y \hookrightarrow X$ be a closed embedding between nonsingular algebraic varieties. The functor i_+ gives an equivalence between the category of coherent right $\mathscr{D}_Y$-modules and the category of coherent right $\mathscr{D}_X$-modules with support cotained in Y.*

Proof. Recall that if $\mathcal{M}$ is a coherent right $\mathscr{D}_Y$-module, we defined

$$i_+\mathcal{M} = i_*\big(\mathcal{M} \otimes_{\mathscr{D}_Y} \mathscr{D}_{Y\to X}\big),$$

where the transfer module $\mathscr{D}_{Y\to X} = \mathscr{O}_Y \otimes_{i^{-1}\mathscr{O}_X} i^{-1}\mathscr{D}_X$ is a $(\mathscr{D}_Y, i^{-1}\mathscr{D}_X)$-bimodule. The first step is to construct an inverse for the functor i_+. We have seen that $i_+\mathcal{M}$ always contains a copy of the $\mathscr{O}_X$-module $i_*\mathcal{M}$, and from the local description, it is clear that $i_*\mathcal{M}$ is exactly the subsheaf of $i_+\mathcal{M}$ that is annihilated by the ideal sheaf $\mathcal{I}_Y \subseteq \mathscr{O}_X$. Thus the inverse functor should take a coherent right $\mathscr{D}_X$-module $\mathcal{N}$ to the subsheaf of sections that are annihilated by $\mathcal{I}_Y$. An efficient way to do this is as follows. Given a coherent right $\mathscr{D}_X$-module $\mathcal{N}$, we define

$$i^\sharp\mathcal{N} = \mathcal{H}om_{i^{-1}\mathscr{D}_X}\big(\mathscr{D}_{Y\to X}, i^{-1}\mathcal{N}\big).$$

Here we use the right $i^{-1}\mathscr{D}_X$-module structure on the transfer module for $\mathcal{H}om_{i^{-1}\mathscr{D}_X}$. The left $\mathscr{D}_Y$-module on $\mathscr{D}_{Y\to X}$ then induces a right $\mathscr{D}_Y$-module structure on $i^\sharp\mathcal{N}$. We can rewrite the above definition as

$$i^\sharp\mathcal{N} = \mathcal{H}om_{i^{-1}\mathscr{D}_X}\big(\mathscr{O}_Y \otimes_{i^{-1}\mathscr{O}_X} i^{-1}\mathscr{D}_X, i^{-1}\mathcal{N}\big) \cong \mathcal{H}om_{i^{-1}\mathscr{O}_X}(\mathscr{O}_Y, i^{-1}\mathcal{N}),$$

using the adjunction between $\mathcal{H}om$ and the tensor product. From the short exact sequence $0 \to i^{-1}\mathcal{I}_Y \to i^{-1}\mathscr{O}_X \to \mathscr{O}_Y \to 0$, we obtain an exact sequence

$$0 \to i^\sharp\mathcal{N} \to i^{-1}\mathcal{N} \to \mathcal{H}om_{i^{-1}\mathscr{O}_X}(i^{-1}\mathcal{I}_Y, i^{-1}\mathcal{N})$$

and so $i^\sharp\mathcal{N}$ is exactly the subsheaf of $i^{-1}\mathcal{N}$ annihilated by $i^{-1}\mathcal{I}_Y$. I will leave it as an exercise to check that this isomorphism is compatible with the natural $\mathscr{D}_Y$-module structure on both sides.

Now the claim is that the natural morphism $i^\sharp i_+\mathcal{M} \to \mathcal{M}$ is an isomorphism for every coherent right $\mathscr{D}_Y$-module $\mathcal{M}$, and that the natural morphism $\mathcal{N} \to i_+ i^\sharp\mathcal{N}$ is an isomorphism for every coherent right $\mathscr{D}_X$-module $\mathcal{N}$ such that $\operatorname{Supp}\mathcal{N} \subseteq Y$. This can be checked locally, and so we may assume without loss of generality that $X = \operatorname{Spec} A$ is affine, with coordinates $x_1, \dots, x_n \in A$, and that the closed embedding is defined by the ideal $I = (x_{r+1}, \dots, x_n) \subseteq A$. If we set $B = A/I$, we then have $Y = \operatorname{Spec} B$. In this setting, the pushforward of a right $D(B)$-module M is isomorphic to $M \otimes_k k[\partial_{r+1}, \dots, \partial_n]$, and it is easy to see from this description that the submodule annihilated by the ideal I is exactly $M \otimes 1 \cong M$. This proves the first isomorphism.

The proof of the second isomorphism is more interesting. Suppose that N is a right $D(A)$-module with $\operatorname{Supp} N$ contained in the closed subscheme $V(I)$. This means that every $s \in N$ is annihilated by a sufficiently large power of I. Our goal is to prove that $N \cong N_0 \otimes_k k[\partial_{r+1}, \dots, \partial_n]$, where $N_0 = \big\{\, s \in N \,\big|\, sI = 0 \,\big\}$. For this, we consider the effect of the operators

$$T_j = x_j \partial_j$$

on the module N. The point is that

$$T_j \cdot \partial_{r+1}^{e_{r+1}} \cdots \partial_n^{e_n} = \partial_{r+1}^{e_{r+1}} \cdots \partial_n^{e_n} \cdot (T_j - e_j),$$

and since T_j acts trivially on the submodule N_0, we have

$$s \otimes \partial_{r+1}^{e_{r+1}} \cdots \partial_n^{e_n} \cdot (T_j - e_j) = 0$$

for every $s \in N_0$. This means that we can read off the exponents of each monomial from the eigenvalues of the operators $T_{r+1}, \dots, T_n$.

Now let us make this precise. The operators $T_{r+1}, \dots, T_n$ commute, and a short calculation shows that

$$T_j(T_j - 1) \cdots (T_j - e) = x_j^{e+1} \partial_j^{e+1}$$

for every $e \geq 0$. For any $s \in N$, we have $s x_j^{e+1} = 0$ for $e \gg 0$, and therefore

$$s T_j(T_j - 1) \cdots (T_j - e) = s x_j^{e+1} \partial_j^{e+1} = 0.$$

This means that s can be written as a sum of eigenvectors of T_j with eigenvalues in $\mathbb{N}$. Since $T_{r+1}, \dots, T_n$ commute, we therefore obtain a decomposition

$$N = \bigoplus_{e_{r+1}, \dots, e_n \in \mathbb{N}} N_{e_{r+1}, \dots, e_n}$$

into simultaneous eigenspaces, where T_j acts on $N_{e_{r+1}, \dots, e_n}$ as multiplication by e_j. Now the claim is that $N_{0,\dots,0} = N_0$, and that this decomposition gives us an isomorphism $N \cong N_0 \otimes_k k[\partial_{r+1}, \dots, \partial_n]$ between N and the pushforward of N_0.

To simplify the notation, let me assume that $r = n - 1$, meaning that $I = (x_n)$ is principal. Then the eigenspace decomposition becomes

$$N = \bigoplus_{e \in \mathbb{N}} N_e,$$

where the operator $T_n = x_n \partial_n$ acts on N_e as multiplication by e. Since T_n commutes with $x_1, \dots, x_{n-1}, \partial_1, \dots, \partial_{n-1}$, each N_e is a $D(B)$-module. Suppose that we have $s \in N_e$. Then we get $s\partial_n \in N_{e+1}$, because

$$s\partial_n T_n = s(\partial_n x_n)\partial_n = s(x_n \partial_n + 1)\partial_n = s\partial_n(e + 1);$$

likewise, we get $s x_n \in N_{e-1}$, because

$$s x_n T_n = s x_n (x_n \partial_n) = s x_n (\partial_n x_n - 1) = s x_n e - s x_n = s x_n (e - 1).$$

Since N_e is trivial for $e \leq -1$, we conclude that $N_0 = \{ s \in N \mid s x_n = 0 \}$; moreover, we see that for $e \geq 0$, the morphism

$$N_0 \to N_e, \quad s \mapsto s\partial_n^e,$$

is an isomorphism of $D(B)$-modules. It is now easy to check that

$$N_0 \otimes_k k[\partial_n] \to N, \quad \sum_{e \in \mathbb{N}} s_e \otimes \partial_n^e \mapsto \sum_{e \in \mathbb{N}} s_e \partial_n^e,$$

is an isomorphism of $D(A)$-modules. This proves the second isomorphism. □

Example 14.1. Kashiwara's equivalence implies that $\mathscr{D}$-modules, unlike $\mathscr{O}$-modules, never have nontrivial nilpotents. For example, the A_1-module $A_1/x^3 A_1$ is isomorphic to three copies of A_1/xA_1.

Kashiwara's equivalence suggests the following definition of the category of algebraic $\mathscr{D}$-modules on a singular algebraic variety. Suppose that X is a nonsingular algebraic variety, and $Y \subseteq X$ any closed subvariety. Then an algebraic $\mathscr{D}_Y$-module is defined to be an algebraic $\mathscr{D}_X$-module whose support is contained in Y. One can use Kashiwara's equivalence to show that the resulting category is, up to equivalence, independent of the choice of nonsingular ambient variety X.

Pulling back. Suppose that $f\colon X \to Y$ is a morphism between two nonsingular algebraic varieties. It is not hard to construct a pullback functor from algebraic $\mathscr{D}_Y$-modules to algebraic $\mathscr{D}_X$-modules. Recall that we have a natural morphism

$$\delta_f\colon \mathscr{T}_X \to f^*\mathscr{T}_Y = \mathscr{O}_X \otimes_{f^{-1}\mathscr{O}_Y} f^{-1}\mathscr{T}_Y,$$

dual to the pullback morphism $f^*\Omega^1_{Y/k} \to \Omega^1_{X/k}$ on Kähler differentials. Now if $\mathcal{M}$ is any left $\mathscr{D}_Y$-module, then this morphism gives

$$f^*\mathcal{M} = \mathscr{O}_X \otimes_{f^{-1}\mathscr{O}_Y} f^{-1}\mathcal{M}$$

the structure of a left $\mathscr{D}_X$-module. The formula is the same as in the case of the transfer module: one has

$$\theta \cdot (g \otimes u) = \theta(g) \otimes u + g \cdot \delta_f(\theta) \cdot (1 \otimes u),$$

where $\theta \in \mathscr{T}_X$, $g \in \mathscr{O}_X$, and $u \in f^{-1}\mathcal{M}$ are local sections. We can say this more compactly by noting that

$$f^*\mathcal{M} \cong \left(\mathscr{O}_X \otimes_{f^{-1}\mathscr{O}_Y} f^{-1}\mathscr{D}_Y\right) \otimes_{f^{-1}\mathscr{D}_Y} f^{-1}\mathcal{M} = \mathscr{D}_{X\to Y} \otimes_{f^{-1}\mathscr{D}_Y} f^{-1}\mathcal{M}.$$

The transfer module $\mathscr{D}_{X\to Y}$ is a $(\mathscr{D}_X, f^{-1}\mathscr{D}_Y)$-bimodule, and $f^*\mathcal{M}$ becomes a left $\mathscr{D}_X$-module through the left $\mathscr{D}_X$-module structure on $\mathscr{D}_{X\to Y}$. Since the pullback of a quasi-coherent $\mathscr{O}_Y$-module is a quasi-coherent $\mathscr{O}_X$-module, it is clear that $f^*\mathcal{M}$ is again an algebraic $\mathscr{D}_X$-module.

Now the functor f^{-1} is exact, but tensor product is only right-exact, and so makes sense to consider also the right derived functors.

Definition 14.2. We define the *inverse image* of a left $\mathscr{D}_Y$-module $\mathcal{M}$ by the formula $f^*\mathcal{M} = \mathscr{D}_{X\to Y} \otimes_{f^{-1}\mathscr{D}_Y} f^{-1}\mathcal{M}$. For $j \geq 0$, we define $L^{-j}f^*\mathcal{M}$ as the j-th right derived functor of f^*.

As usual, $L^{-j}f^*\mathcal{M}$ is computed by choosing a resolution of $\mathcal{M}$ by $\mathscr{D}_Y$-modules that are locally free (or flat) over $\mathscr{O}_Y$; alternatively, we can choose a resolution of $\mathscr{D}_{X\to Y}$.

Example 14.3. Suppose that $\mathscr{E}$ is a locally free $\mathscr{O}_Y$-module with an integrable connection $\nabla\colon \mathscr{E} \to \Omega^1_{Y/k} \otimes_{\mathscr{O}_Y} \mathscr{E}$, viewed as a left $\mathscr{D}_Y$-module. The inverse image is then simply the usual pullback $f^*\mathscr{E}$, together with the integrable connection

$$f^*\nabla\colon f^*\mathscr{E} \to f^*\Omega^1_{Y/k} \otimes_{\mathscr{O}_X} f^*\mathscr{E} \to \Omega^1_{X/k} \otimes_{\mathscr{O}_X} f^*\mathscr{E},$$

viewed as a left $\mathscr{D}_X$-module.

Example 14.4. Consider the left A_1-module $M = A_1/A_1x$ and its pullback to the origin in $\mathbb{A}^1_k$. The corresponding morphism of k-algebras is $k[x] \to k$; using the free resolution

$$k[x] \xrightarrow{x} k[x]$$

for k, the derived functors of the pullback are computed by the complex

$$A_1/A_1x \xrightarrow{x} A_1/A_1x,$$

where the map is $P \mapsto xP$. The kernel is isomorphic to k, generated by the image of $1 \in A_1$; the cokernel is trivial, because $1 = -x\partial$ modulo A_1x. Thus $L^0i^*M = 0$ and $L^{-1}i^*M = k$.

In Lecture 12, I said that the definition of the pushforward functor (in the case of a closed embedding) was motivated by the pushforward of distributions. So why do I not talk about pulling back functions before introducing the pullback functor? The reason is that pulling back $\mathscr{D}$-modules does not correspond to pulling back functions; as we will see next week, the actual meaning is much more interesting. For now, let me just point out one difference between the two functors: pulling back does not necessarily preserve coherence.

Example 14.5. Consider the embedding $\operatorname{Spec} k \hookrightarrow \mathbb{A}^1_k$ of the origin, corresponding to the morphism of k-algebras $k[x] \to k$. The pullback of $\mathscr{D}_{\mathbb{A}^1_k}$ is the k-module $k \otimes_{k[x]} A_1(k) = A_1(k)/xA_1(k)$. This is infinite-dimensional, because the elements $1, \partial, \partial^2, \dots$ are all linearly independent, and in particular, it is not coherent over k.

In general, the pullback of a $\mathscr{D}_X$-module of the form $\mathscr{D}_X/\mathscr{D}_X(P_1, \dots, P_m)$ is not coherent, and so we cannot interpret it as pulling back functions and looking at the differential equations they satisfy.

The following lemma is obvious from the definition.

Lemma 14.6. *If $f\colon X \to Y$ and $g\colon Y \to Z$ are morphisms between nonsingular algebraic varieties, then one has a natural isomorphism of functors $(g \circ f)^* = f^* g^*$.*

We can factor any morphism $f\colon X \to Y$ through its graph as

$$X \xrightarrow{i_f} X \times Y \xrightarrow{p_2} Y$$

as a closed embedding i_f followed by a smooth morphism p_2 (actually, a projection in a product). Because of the lemma, this means that it suffices to understand the pullback functor in the case of closed embeddings and smooth morphisms.

Non-characteristic inverse image. I am now going to describe a condition under which f^* preserves coherence. This will also help us understand what the pullback functor is doing in terms of differential equations. To do this, we revisit a very pretty classical result about differential equations, called the *Cauchy-Kovalevskaya theorem.* Let's begin with the case of ordinary differential equations.

Theorem 14.7 (Cauchy-Kovalevskaya)**.** *Consider the initial value problem*

$$\frac{du}{dt} = F(u), \quad u(0) = 0,$$

for a real function u. If $F\colon (-\varepsilon, \varepsilon) \to \mathbb{R}$ is real-analytic near 0, then the solution u is also real-analytic near 0.

Proof. Although it is not directly connected with $\mathscr{D}$-modules, let me show you the proof, because it is very beautiful. The proof is basically Cauchy's original proof. How do we show that u is real-analytic? We have to prove that the Taylor series

$$\sum_{n=0}^{\infty} u^{(n)}(0) \frac{t^n}{n!}$$

converges in a neighborhood of 0, and for that, we need to compute the values of all the derivatives $u^{(n)}(0)$. The differential equation gives

$$\begin{aligned}
u' &= F(u) \\
u'' &= F'(u)u' = F'(u)F(u) \\
u''' &= F''(u)u'F(u) + (F'(u))^2 u' = F''(u)(F(u))^2 + (F'(u))^2 F(u).
\end{aligned}$$

and so on. In principle, we can compute $u^{(n)}(0)$ for every $n \geq 0$, but the formulas get very complicated, and so trying to prove the convergence of the series looks pretty hopeless. Still, what we get is that

$$u^{(n)} = P_n\big(F(u), F'(u), \dots, F^{(n-1)}(u)\big),$$

where P_n is a polynomial with nonnegative integer coefficients. These polynomials are universal, in the sense that they do not depend on the given function F. For example, $P_2(x, y) = yx$ and $P_3(x, y, z) = zx^2 + y^2 x$. Because P_n has nonnegative coefficients, this gives us an upper bound

$$|u^{(n)}(0)| \leq P_n\big(|F(0)|, |F'(0)|, \dots, |F^{(n-1)}(0)|\big)$$

on the derivatives of u, using the initial condition $u(0) = 0$. Now Cauchy makes the following brilliant observation. Suppose that we have another function G with the property that $|F^{(n)}(0)| \leq G^{(n)}(0)$ for every $n \geq 0$. Then

$$|u^{(n)}(0)| \leq P_n\big(G(0), G'(0), \dots, G^{(n-1)}(0)\big) = v^{(n)}(0),$$

where v is the solution to the initial value problem

$$\frac{dv}{dt} = G(v), \quad v(0) = 0.$$

The reason is again that P_n has nonnegative coefficients, and that the same polynomial P_n works for both F and G. Such a function G is called a "majorant", and the proof is known as the *method of majorants*. Suppose that we manage to find G in such a way that the function v is real-analytic. Then the Taylor series

$$\sum_{n=0}^{\infty} v^{(n)}(0) \frac{t^n}{n!}$$

has a positive radius of convergence, and since $|u^{(n)}(0)| \leq v^{(n)}(0)$ for every $n \geq 0$, the same is true for the series

$$\sum_{n=0}^{\infty} |u^{(n)}(0)| \frac{t^n}{n!}.$$

This is sufficient to conclude that u is real-analytic in a neighborhood of 0.

It remains to construct a suitable majorant G. By assumption, F is real-analytic near 0, and so its Taylor series

$$\sum_{n=0}^{\infty} F^{(n)}(0) \frac{t^n}{n!}$$

has a positive radius of convergence. By comparing this series with a geometric series, we find that there are constants $C > 0$ and $r > 0$ such that $|F^{(n)}(0)| \leq Cn!/r^n$ for every $n \geq 0$. We can then take

$$G(t) = C \sum_{n=0}^{\infty} \left(\frac{t}{r}\right)^n = \frac{Cr}{r-t},$$

because $G^{(n)}(0) = Cn!/r^n \geq |F^{(n)}(0)|$ by construction. The solution of the corresponding initial value problem

$$\frac{dv}{dt} = \frac{Cr}{r-v}, \quad v(0) = 0,$$

is easily found using separation of variables; the result is that $v = r - r\sqrt{1 - 2Ct/r}$. This is evidently real-analytic for $|t| < r/2C$, and so we are done. $\square$

Exercises.

Exercise 14.1. Let $X = \operatorname{Spec} A$ and $Y = \operatorname{Spec} B$, where $B = A/I$ for an ideal $I \subseteq A$ and both A and B are nonsingular. Let N be a right $D(A)$-module.

(a) Show that $N_0 = \{\, s \in N \mid sI = 0 \,\}$ is a B-module, and that the map

$$N_0 \otimes_B T_B \to N_0, \quad s \otimes \theta \mapsto s \cdot \delta(\theta),$$

makes N_0 into a right $D(B)$-module, where $\delta \colon \operatorname{Der}_k(B) \to B \otimes_A \operatorname{Der}_k(A)$ is the induced morphism between derivations.

(b) Check that the isomorphism of B-modules

$$\operatorname{Hom}_{D(A)}\big(B \otimes_A D(A), N\big) \cong \operatorname{Hom}_A(B, N) \cong N_0$$

is actually an isomorphism of right $D(B)$-modules.

Exercise 14.2. If $T = x\partial$, prove the identities

$$T\partial^e = \partial^e(T-e) \quad \text{and} \quad T(T-1)\cdots(T-e) = x^{e+1}\partial^{e+1}$$

for every $e \geq 0$.

Lecture 15: April 1

The Cauchy-Kovalevskaya theorem. Last time, we showed that the solution to the initial value problem

$$\frac{du}{dt} = F(u), \quad u(0) = 0,$$

is real-analytic near $t = 0$, provided that this is true for the function F. I also showed you Cauchy's proof, using the "method of majorants". Today, we are going to generalize this result to partial differential equations. We work on $\mathbb{R}^n$, with coordinates $x_1, \dots, x_n$, and consider a partial differential equation of the form

$$Pu = \sum_{|\alpha| \leq k} f_\alpha \partial^\alpha u = 0,$$

where each f_α is a real-analytic function in a neighborhood of the origin, say. (And $\partial_j = \partial/\partial x_j$, as usual.) In other words, P is a linear differential operator of order k with real-analytic coefficients. We will specify the initial conditions on the hyperplane $x_n = 0$, which is a copy of $\mathbb{R}^{n-1}$. They are

$$u\big|_{\mathbb{R}^{n-1}} = g_0, \quad \partial_n u\big|_{\mathbb{R}^{n-1}} = g_1, \quad \dots, \quad \partial_n^{k-1} u\big|_{\mathbb{R}^{n-1}} = g_{k-1},$$

where $g_0, g_1, \dots, g_{k-1}$ are real-analytic in a neighborhood of the origin in $\mathbb{R}^{n-1}$. From this data, we can of course compute all partial derivatives of u of order at most $k-1$ on $\mathbb{R}^{n-1}$; indeed, if $\alpha \in \mathbb{N}^n$ is a multi-index, then

$$\partial^\alpha u\big|_{\mathbb{R}^{n-1}} = \partial_1^{\alpha_1} \cdots \partial_{n-1}^{\alpha_{n-1}} g_{\alpha_n}, \tag{15.1}$$

provided that $\alpha_n \leq k-1$.

The goal is to show that the solution u is real-analytic near the origin. For that to be true, the Taylor series of u at the origin needs to be determined by the equation $Pu = 0$ plus the initial conditions, and so we had better be able to compute *all* partial derivatives of u at the origin. Since we can always differentiate along $\mathbb{R}^{n-1}$, the real question is how to find

$$\partial_n^j u\big|_{\mathbb{R}^{n-1}}$$

for $j \geq k$. Clearly, this information has to come from $Pu = 0$. Since P has order k, we can rewrite $Pu = 0$ as

$$f_{(0,\dots,0,k)} \cdot \partial_n^k u = -\sum_{\alpha_n \leq k-1} f_\alpha \partial^\alpha u,$$

and in view of (15.1), we can solve this for $\partial_n^k u\big|_{\mathbb{R}^{n-1}}$ if and only if the restriction of the coefficient function $f_{(0,\dots,0,k)}$ to $\mathbb{R}^{n-1}$ is everywhere nonzero. (If we only care about what happens at the origin, then the condition is that $f_{(0,\dots,0,k)}$ should be nonzero at the origin.) If that is the case, we can of course divide through by $f_{(0,\dots,0,k)}$ and arrange that ∂_n^k appears with coefficient 1.

Definition 15.2. We say that P is *non-characteristic* with respect to the hypersurface $x_n = 0$ if the coefficient function $f_{(0,\dots,0,k)}$ is everywhere nonzero on $\mathbb{R}^{n-1}$.

Assuming that P is non-characteristic (and $f_{(0,\dots,0,k)} = 1$), we can rewrite the equation $Pu = 0$ in the form

$$\partial_n^k u = Qu,$$

where Q is a differential operator of order k in which ∂_n^k does not appear. We can now use this equation recursively, together with (15.1), to compute $\partial^\alpha u\big|_{\mathbb{R}^{n-1}}$ for every $\alpha \in \mathbb{N}^n$. In particular, assuming that P is non-characteristic, the equation $Pu = 0$ together with the initial conditions on $\mathbb{R}^{n-1}$ give enough information to compute the Taylor series for u at the origin. We can now state the PDE version of the Cauchy-Kovalevskaya theorem.

Theorem 15.3 (Cauchy-Kovalevskaya)**.** *Let P be a linear partial differential operator of order k whose coefficients are real-analytic near the origin in $\mathbb{R}^n$. If P is non-characteristic with respect to $x_n = 0$, then the boundary-value problem*

$$Pu = 0, \quad u\big|_{\mathbb{R}^{n-1}} = g_0, \quad \partial_n u\big|_{\mathbb{R}^{n-1}} = g_1, \quad \dots, \quad \partial_n^{k-1} u\big|_{\mathbb{R}^{n-1}} = g_{k-1},$$

has a unique real-analytic solution u near the origin in $\mathbb{R}^n$, for every choice of functions $g_0, g_1, \dots, g_{k-1}$ real-analytic near the origin in $\mathbb{R}^{n-1}$.

Example 15.4. Here is an example to show that the solution can fail to be real-analytic if P is "characteristic". This example is due to Kovalevskaya herself. Consider the heat equation $\partial_t u = \partial_x^2 u$ in $\mathbb{R}^2$, with coordinates (x, t). Since the equation is first-order in t, we only need a single initial condition $u(x, 0) = g(x)$. Note that the operator $P = \partial_t - \partial_x^2$ is characteristic with respect to $t = 0$, because it has order 2, but no term involving ∂_t^2. Here is a heuristic reason why we cannot expect u to be real-analytic in general. From the equation, we get

$$\partial_t^n u = \partial_x^{2n} u,$$

and at $(x, t) = (0, 0)$, this evaluates to $g^{(2n)}(0)$. If the Taylor series of g at the origin has a finite radius of convergence, then

$$|g^{(2n)}(0)| \geq C \frac{(2n)!}{r^{2n}}$$

for some $C, r > 0$. But this means that the function $h(t) = u(0, t)$ cannot be real-analytic in t: indeed, from the above, we deduce that

$$|h^{(n)}(0)| \geq C \frac{(2n)!}{r^{2n}},$$

and since $(2n)!$ grows so much faster than $n!$, the Taylor series of $h(t)$ has radius of convergence equal to zero. For an actual example, take $g(x) = 1/(x^2 + 1)$.

Now let me give an outline of the proof of Theorem 15.3. As explained above, we can rewrite the equation $Pu = 0$ in the form

$$\partial_n^k u = Qu,$$

where Q is a differential operator of order k with real-analytic coefficients, such that Q has order at most $k-1$ in ∂_n. Moreover, we can subtract a suitable real-analytic function from u to arrange that $g_0 = g_1 = \ldots = g_{k-1} = 0$. We now rewrite the problem as a system of first-order PDE for $N = \binom{n+k-1}{n} + 1$ unknown functions $u_1, \dots, u_N$. These functions are the $N - 1$ partial derivatives $\partial^\alpha u$ for $|\alpha| \leq k - 1$, and the auxiliary $u_N = x_n$. In vector notation, the system takes the form

$$\frac{\partial \underline{u}}{\partial x_n} = \sum_{j=1}^{n-1} B_j(x_1, \dots, x_{n-1}) \frac{\partial \underline{u}}{\partial x_j} + B_0(x_1, \dots, x_{n-1}) \underline{u}, \tag{15.5}$$

where $\underline{u} = (u_1, \dots, u_N)$, and where the coefficient matrices $B_0, \dots, B_{n-1}$ are derived from Q, hence real-analytic near the origin. Note that we threw in the function $u_N = x_n$ in order to make the coefficients be independent of x_n; of course, the corresponding equation is simply $\partial u_N / \partial x_n = 1$. The initial condition is that $\underline{u}$ is the zero vector for $x_n = 0$.

Now one can again use the method of majorants to prove that $\underline{u}$ is real-analytic near the origin in $\mathbb{R}^n$. From (15.5), all partial derivatives of $\underline{u}$ at the origin are given by (very complicated) universal polynomials with nonnegative integer coefficients in the partial derivatives of $B_0, \dots, B_{n-1}$ at the origin. Using the fact that the coefficient matrices are real-analytic near the origin, one can again write down simple majorants for each of them, and then explicitely solve the resulting system of first-order PDE to show that its solution $\underline{v}$, and hence also $\underline{u}$, is real-analytic near the origin.

Non-characteristic $\mathscr{D}$-modules. Here is a geometric interpretation for the condition that P is non-characteristic with respect to $x_n = 0$. If $P = \sum_\alpha f_\alpha \partial^\alpha$ has order k as above, then its principal symbol

$$\sigma_k(P) = \sum_{|\alpha|=k} f_\alpha(x_1, \dots, x_n) \cdot \xi_1^{\alpha_1} \cdots \xi_n^{\alpha_n}$$

is a homogeneous polynomial of degree k in the variables $\xi_1, \dots, \xi_n$. We said that P is non-characteristic iff $f_{(0,\dots,0,k)}(x_1, \dots, x_{n-1}, 0) \neq 0$ for every $x_1, \dots, x_{n-1}$. Another way of saying this is that if we set $x_n = 0$ and assign arbitrary values to the variables $x_1, \dots, x_{n-1}, \xi_1, \dots, \xi_{n-1}$, then $\sigma_k(P)$, considered as a polynomial in the remaining variable ξ_n, always has degree exactly k. The geometric meaning of this condition is as follows. We have the usual maps between the cotangent bundles $T^*\mathbb{R}^n = \mathbb{R}^n \times \mathbb{R}^n$ and $T^*\mathbb{R}^{n-1} = \mathbb{R}^{n-1} \times \mathbb{R}^{n-1}$:

$$\begin{tikzcd} \mathbb{R}^{n-1} \times_{\mathbb{R}^n} T^*\mathbb{R}^n \arrow[r, "di"] \arrow[d, "p_2"] & T^*\mathbb{R}^{n-1} \\ T^*\mathbb{R}^n & \end{tikzcd}$$

Using $x_1, \dots, x_n, \xi_1, \dots, \xi_n$ as coordinates on $T^*\mathbb{R}^n$, the maps are just

$$p_2(x_1, \dots, x_{n-1}, \xi_1, \dots, \xi_n) = (x_1, \dots, x_{n-1}, 0, \xi_1, \dots, \xi_n)$$
$$di(x_1, \dots, x_{n-1}, \xi_1, \dots, \xi_n) = (x_1, \dots, x_{n-1}, \xi_1, \dots, \xi_{n-1}).$$

Consider the subset $\mathrm{Ch}(P) \subseteq T^*\mathbb{R}^n$ defined by the equation $\sigma_k(P) = 0$. Setting $x_n = 0$ and prescribing values for $x_1, \dots, x_{n-1}, \xi_1, \dots, \xi_{n-1}$ amounts to looking at the fibers of $p_2^{-1}\,\mathrm{Ch}(P)$ over $T^*\mathbb{R}^{n-1}$, and so P is non-characteristic exactly when the projection from $p_2^{-1}\,\mathrm{Ch}(P)$ to $T^*\mathbb{R}^{n-1}$ is a finite morphism of degree k. If we observe that $\mathrm{Ch}(P)$ is the characteristic variety of the $\mathscr{D}$-module $A_n(\mathbb{R})/A_n(\mathbb{R})P$, this finiteness condition makes sense for arbitrary coherent $\mathscr{D}$-modules.

Let me now give the general definition. Suppose that $f\colon X \to Y$ is a morphism between two nonsingular algebraic varieties. Here is the diagram of the induced morphisms between cotangent bundles:

$$\begin{tikzcd} X \times_Y T^*Y \arrow[r, "df"] \arrow[d, "p_2"] & T^*X \\ T^*Y & \end{tikzcd}$$

Definition 15.6. Let $\mathcal{M}$ be a coherent left $\mathscr{D}_Y$-module. We say that $\mathcal{M}$ is *non-characteristic* with respect to $f\colon X \to Y$ if the morphism

$$df\colon p_2^{-1}\,\mathrm{Ch}(\mathcal{M}) \to T^*X$$

is finite over its image.

Example 15.7. Consider the closed embedding $i\colon \mathbb{A}_k^{n-1} \hookrightarrow \mathbb{A}_k^n$, defined by $x_n = 0$. Our earlier discusion shows that if $P \in A_n$ is nonzero, then the left A_n-module A_n/A_nP is non-characteristic with respect to i if and only if the differential operator P is non-characteristic with respect to $x_n = 0$ in the classical sense.

Example 15.8. If $f\colon X \to Y$ is a smooth morphism, then every coherent $\mathscr{D}_Y$-module is non-characteristic with respect to f. Indeed, smoothness means that we have a short exact sequence

$$0 \to f^*\Omega^1_{Y/k} \to \Omega^1_{X/k} \to \Omega^1_{X/Y} \to 0,$$

with $\Omega^1_{X/Y}$ locally free of rank $\dim X - \dim Y$. But this says that

$$df\colon X \times_Y T^*Y \to T^*X$$

is a closed embedding (of codimension $\dim X - \dim Y$), and so $p_2^{-1}\operatorname{Ch}(\mathcal{M})$ is trivially finite over its image in T^*X.

In the following example, we compute the pullback of an A_n-module of the form A_n/A_nP to the hypersurface $x_n = 0$, in the case where P is non-characteristic.

Example 15.9. Consider the left A_n-module $M = A_n/A_nP$, where $P \in A_n$ is a nonzero differential operator of order $r \geq 0$. Suppose that M is non-characteristic with respect to the closed embedding $i\colon \mathbb{A}^{n-1} \hookrightarrow \mathbb{A}^n$ defined by the equation $x_n = 0$. We claim that, in this case, the pullback i^*M is not only coherent, but actually a free A_n-module of rank r. The definition of the pullback gives

$$i^*M = k[x_1,\dots,x_{n-1}] \otimes_{k[x_1,\dots,x_n]} M \cong A_n/(x_nA_n + A_nP), \tag{15.10}$$

where the right-hand side is a left A_{n-1}-module in the obvious way. We have a morphism of left A_{n-1}-modules

$$\varphi\colon A_{n-1}^{\oplus r} \to A_n/(x_nA_n + A_nP)$$
$$(Q_0, Q_1, \dots, Q_{r-1}) \mapsto Q_0 + Q_1\partial_n + \dots + Q_{r-1}\partial_n^{r-1}.$$

We will show that φ is an isomorphism. Let us first argue that ∂_n^r is in the image. We can write our differential operator $P \in A_n$ uniquely in the form

$$P = f\partial_n^r - P_{r-1}\partial_n^{r-1} - \dots - P_1\partial_n - P_0,$$

where $f \in k[x_1,\dots,x_n]$ and where $P_0,\dots,P_{r-1} \in A_n$ do not involve ∂_n. The fact that P is non-characteristic means that f is nowhere vanishing on $\mathbb{A}^{n-1}$; after rescaling, we can assume that $f = 1 - x_ng$. Writing $P_j = Q_n + x_nR_j$, with $Q_j \in A_{n-1}$, we get

$$\partial_n^r = \sum_{j=0}^{r-1} Q_j\partial_n^j + x_n\Big(g\partial_n^r + \sum_{j=0}^{r-1} R_j\partial_n^j\Big) + P, \tag{15.11}$$

and so ∂_n^r belongs to the image of φ. Using the relation in (15.11) repeatedly, we see that this is true for all powers of ∂_n, and so φ is surjective.

It remains to prove that φ is injective. This is equivalent to saying that if

$$Q_0 + Q_1\partial_n + \dots + Q_{r-1}\partial_n^{r-1} = x_nS + TP$$

for some $Q_0,\dots,Q_{r-1} \in A_{n-1}$ and $S, T \in A_n$, then actually $Q_0 = \dots = Q_{r-1} = 0$. We can write $T = x_nT_0 + T_1$, in such a way that x_n does not appear in T_1; since $x_nS + TP = x_n(S + T_0) + T_1P$, we can therefore assume without loss of generality that T does not involve x_n. Now suppose, for the sake of contradiction, that $T \neq 0$. On the right-hand side of the equation, ∂_n^r appears with a nonzero coefficient: indeed, P contains $(1 - x_ng)\partial_n^r$, and since T does not involve x_n, it is not possible to cancel this term against anything from x_nS. But this clearly contradicts the fact that ∂_n^r does not appear on the left-hand side of the equation. The conclusion is that $T = 0$; and then also $Q_0 = \dots = Q_{r-1} = 0$, because the right-hand side is divisible by x_n, whereas the left-hand side does not involve x_n.

The preceding example, together with the Cauchy-Kovalevskaya theorem, sheds some light on what the pullback of $\mathscr{D}$-modules has to do with differential equations.

Example 15.12. Continuing with the previous example, let us take $k = \mathbb{R}$. Set $\mathcal{M} = \mathscr{D}_{\mathbb{R}^n}/\mathscr{D}_{\mathbb{R}^n}P$. Let us denote by $\mathscr{R}_{\mathbb{R}^n}$ the sheaf of real-analytic functions on $\mathbb{R}^n$; it is a left $\mathscr{D}_{\mathbb{R}^n}$-module in the obvious way. Recall from Lecture 1 that real-analytic solutions to the equation $Pu = 0$ on an open subset $U \subseteq \mathbb{R}^n$ correspond naturally to morphisms of left $\mathscr{D}_{\mathbb{R}^n}$-modules $\mathcal{M} \to \mathscr{R}_{\mathbb{R}^n}$ over U; here the morphism takes the generator $1 \in \Gamma(U, \mathscr{D}_{\mathbb{R}^n})$ to the corresponding function $u \in \Gamma(U, \mathscr{R}_{\mathbb{R}^n})$.

In this notation, the Cauchy-Kovalevskaya theorem says that if $V \subseteq \mathbb{R}^{n-1}$ is an open subset, and $g_0, g_1, \dots, g_{r-1} \in \Gamma(V, \mathscr{R}_{\mathbb{R}^{n-1}})$ are arbitrary real-analytic functions on V, there is an open subset $U \subseteq \mathbb{R}^n$ with $U \cap \mathbb{R}^{n-1} = V$, and a real-analytic function $u \in \Gamma(U, \mathscr{R}_{\mathbb{R}^n})$, such that $Pu = 0$ and

$$\partial_n^j u\big|_{\mathbb{R}^{n-1}} = g_j \quad \text{for } j = 0, 1, \dots, r-1.$$

By what we have just said, u may be viewed as a section of the sheaf

$$i^{-1}\mathcal{H}om_{\mathscr{D}_{\mathbb{R}^n}}\big(\mathcal{M}, \mathscr{R}_{\mathbb{R}^n}\big)$$

on the open subset V. Now we have a natural morphism of sheaves

$$i^{-1}\mathcal{H}om_{\mathscr{D}_{\mathbb{R}^n}}\big(\mathcal{M}, \mathscr{R}_{\mathbb{R}^n}\big) \to \mathcal{H}om_{\mathscr{D}_{\mathbb{R}^{n-1}}}\big(i^*\mathcal{M}, i^*\mathscr{R}_{\mathbb{R}^n}\big) \to \mathcal{H}om_{\mathscr{D}_{\mathbb{R}^{n-1}}}\big(i^*\mathcal{M}, \mathscr{R}_{\mathbb{R}^{n-1}}\big);$$

it works by applying the pullback functor i^* to a morphism of left $\mathscr{D}_{\mathbb{R}^n}$-modules $\mathcal{M} \to \mathscr{R}_{\mathbb{R}^n}$, and then composing with the restriction morphism $i^*\mathscr{R}_{\mathbb{R}^n} \to \mathscr{R}_{\mathbb{R}^{n-1}}$. The preceding example shows that $i^*\mathcal{M}$ is a free $\mathscr{D}_{\mathbb{R}^{n-1}}$-module of rank r, generated by the images of $1, \partial_n, \dots, \partial_n^{r-1}$. Thus

$$\mathcal{H}om_{\mathscr{D}_{\mathbb{R}^{n-1}}}\big(i^*\mathcal{M}, \mathscr{R}_{\mathbb{R}^{n-1}}\big) \cong \mathscr{R}_{\mathbb{R}^{n-1}}^{\oplus r},$$

and one checks that the resulting morphism

$$i^{-1}\mathcal{H}om_{\mathscr{D}_{\mathbb{R}^n}}\big(\mathcal{M}, \mathscr{R}_{\mathbb{R}^n}\big) \to \mathscr{R}_{\mathbb{R}^{n-1}}^{\oplus r}$$

takes u to its boundary values

$$u\big|_{\mathbb{R}^{n-1}}, \partial_n u\big|_{\mathbb{R}^{n-1}}, \cdots, \partial_n^{r-1} u\big|_{\mathbb{R}^{n-1}}.$$

This means that we can interpret the Cauchy-Kovalevskaya theorem, in more fancy language, as the statement that the morphism

$$i^{-1}\mathcal{H}om_{\mathscr{D}_{\mathbb{R}^n}}\big(\mathcal{M}, \mathscr{R}_{\mathbb{R}^n}\big) \to \mathcal{H}om_{\mathscr{D}_{\mathbb{R}^{n-1}}}\big(i^*\mathcal{M}, \mathscr{R}_{\mathbb{R}^{n-1}}\big)$$

is an isomorphism of sheaves on $\mathbb{R}^{n-1}$. This tells us that the $\mathscr{D}$-module pullback $i^*\mathcal{M}$ has to do with the *boundary conditions* for the partial differential equation $Pu = 0$; the fact that $i^*\mathcal{M}$ is free of rank r means that we can specify r independent real-analytic functions as boundary conditions.

Non-characteristic pullback. Our next goal is to show that if $f\colon X \to Y$ is a morphism between nonsingular algebraic varieties, and if $\mathcal{M}$ is a coherent left $\mathscr{D}_Y$-module that is non-characteristic with respect to f, then the pullback $f^*\mathcal{M}$ is coherent over $\mathscr{D}_X$. To simplify the analysis, we are going to factor f through its graph. Let us see how this factorization interacts with being non-characteristic.

Suppose for a moment that we have an arbitrary factorization

$$\begin{tikzcd} X \arrow[r, "g"] \arrow[rr, bend right, "f"'] & Z \arrow[r, "h"] & Y \end{tikzcd}$$

with Z nonsingular. We can then draw the following big diagram of induced morphisms between cotangent bundles:

$$\begin{tikzcd} X \times_Y T^*Y \arrow[r, "q"] \arrow[d, "g \times \mathrm{id}"] \arrow[rr, bend left, "df"] & X \times_Z T^*Z \arrow[r, "dg"] \arrow[d, "p_2"] & T^*X \\ Z \times_Y T^*Y \arrow[r, "dh"] \arrow[d, "p_2"] & T^*Z & \\ T^*Y & & \end{tikzcd}$$

If $h\colon Z \to Y$ is a smooth morphism, then dh is a closed embedding, and so its base change along $g\colon X \to Z$, which is denoted by q in the diagram above, is also a closed embedding. Since $df = dg \circ q$, we see that the subset $p_2^{-1}\operatorname{Ch}(\mathcal{M})$ of $X \times_Y T^*Y$ is finite over T^*Z if and only if its image under q is finite over T^*Z. This observation can be used to reduce the study of non-characteristic pullback to two special cases: smooth morphisms and closed embeddings.

Exercises.

Exercise 15.1. On $\mathbb{R}^n$, we use coordinates $x_1, \dots, x_n$. Let $\mathcal{M} = \mathscr{D}_{\mathbb{R}^n}/\mathscr{D}_{\mathbb{R}^n}P$, where P is a differential operator of order r that is non-characteristic with respect to $x_n = 0$. Show that the morphism

$$i^{-1}\mathcal{H}om_{\mathscr{D}_{\mathbb{R}^n}}\big(\mathcal{M}, \mathscr{R}_{\mathbb{R}^n}\big) \to \mathscr{R}_{\mathbb{R}^{n-1}}^{\oplus r}$$

in Example 15.12 takes a real-analytic solution to the equation $Pu = 0$ to the r-vector of its normal derivatives

$$u\big|_{\mathbb{R}^{n-1}}, \partial_n u\big|_{\mathbb{R}^{n-1}}, \cdots, \partial_n^{r-1}u\big|_{\mathbb{R}^{n-1}}.$$

Lecture 16: April 8

Non-characteristic pullback and coherence. Recall that if $f\colon X \to Y$ is a morphism between nonsingular algebraic varieties, we have the following morphisms between cotangent bundles:

$$
\begin{tikzcd}
X \times_Y T^*Y \arrow[r, "df"] \arrow[d, "p_2"] & T^*X \\
T^*Y
\end{tikzcd}
\tag{16.1}
$$

We said last time that a coherent left $\mathscr{D}_Y$-module $\mathcal{M}$ is called *non-characteristic* with respect to f if $p_2^{-1}\operatorname{Ch}(\mathcal{M})$ is finite over its image in T^*X (under the morphism df). Here are three typical examples.

Example 16.2. If f is a smooth morphism, then df is a closed embedding, and so every coherent left $\mathscr{D}_Y$-module is noncharacteristic with respect to f.

Example 16.3. If $\mathcal{M}$ is a vector bundle with integrable connection, then $\operatorname{Ch}(\mathcal{M})$ is the zero section in T^*Y. Since the zero section in $X \times_Y T^*Y$ and in T^*X are both isomorphic to X, the restriction of df to $p_2^{-1}\operatorname{Ch}(\mathcal{M})$ is an isomorphism, and so $\mathcal{M}$ is non-characteristic with respect to any morphism f. So being non-characteristic is really a condition on the other components of the characteristic variety.

Example 16.4. The left $\mathscr{D}_Y$-module $\mathscr{D}_Y$ is *never* non-characteristic with respect to a closed embedding $f\colon X \hookrightarrow Y$ (as long as $\dim X < \dim Y$). Indeed, $\operatorname{Ch}(\mathcal{M}) = T^*Y$ in this case, and since df has positive-dimensional fibers, $p_2^{-1}\operatorname{Ch}(\mathcal{M})$ is not finite over its image.

Our goal for today is to show that pulling back preserves coherence in the non-characteristic setting.

Theorem 16.5. *Let $f\colon X \to Y$ be a morphism between nonsingular algebraic varieties, and $\mathcal{M}$ a coherent left $\mathscr{D}_Y$-module. If $\mathcal{M}$ is non-characteristic with respect to f, then the following is true.*

(a) The pullback $f^*\mathcal{M}$ is a coherent left $\mathscr{D}_X$-module.
(b) One has $L^{-j}f^*\mathcal{M} = 0$ for $j \geq 1$.
(c) One has $\operatorname{Ch}(f^*\mathcal{M}) = df\big(p_2^{-1}\operatorname{Ch}(\mathcal{M})\big)$.

Note that since $df\colon p_2^{-1}\operatorname{Ch}(\mathcal{M}) \to T^*X$ is a finite morphism, the image is again a closed algebraic subset of T^*X. Thus the statement in (c) makes sense.

For the proof, the idea is to factor $f\colon X \to Y$ as a closed embedding followed by a smooth morphism, and to analyze the two cases separately.

Smooth morphisms. Suppose that $f\colon X \to Y$ is a smooth morphism. In the diagram in (16.1), the morphism p_2 is then also smooth, and the morphism df is a closed embedding. Now let $\mathcal{M}$ be a coherent left $\mathscr{D}_Y$-module. We have

$$f^*\mathcal{M} = \mathscr{D}_{X\to Y} \otimes_{f^{-1}\mathscr{D}_Y} f^{-1}\mathcal{M} \cong \mathscr{O}_X \otimes_{f^{-1}\mathscr{O}_Y} f^{-1}\mathcal{M},$$

and since smooth morphisms are flat, the tensor product with $\mathscr{O}_X$ is exact. In particular, the higher derived functors of the tensor product are zero, and so $L^{-j}f^*\mathcal{M} = 0$ for $j \geq 1$. This proves (b). Next, we show that $f^*\mathcal{M}$ is coherent over $\mathscr{D}_X$. By assumption, $\mathcal{M}$ is coherent over $\mathscr{D}_Y$, and so $f^{-1}\mathcal{M}$ is coherent over $f^{-1}\mathscr{D}_Y$. Since the left $\mathscr{D}_X$-module structure on $f^*\mathcal{M}$ comes from $\mathscr{D}_{X\to Y}$, it is therefore enough to show that the morphism

$$\mathscr{D}_X \to \mathscr{D}_{X\to Y} = \mathscr{O}_X \otimes_{f^{-1}\mathscr{O}_Y} f^{-1}\mathscr{D}_Y, \quad P \mapsto P\cdot(1\otimes 1)$$

is surjective. This can be done locally. We can therefore assume that X and Y are affine, and we can choose local coordinates $x_1, \dots, x_{n+r} \in \Gamma(X, \mathscr{O}_X)$ and $y_1, \dots, y_n \in \Gamma(Y, \mathscr{O}_Y)$, in such a way that the morphism on tangent sheaves

$$\mathscr{T}_X \to f^* \mathscr{T}_Y = \mathscr{O}_X \otimes_{f^{-1}\mathscr{O}_Y} f^{-1}\mathscr{T}_Y$$

maps ∂_{x_j} to $1 \otimes \partial_{y_j}$ for $1 \leq j \leq n$, and to zero otherwise. (This means that $\partial_{x_{n+1}}, \dots, \partial_{x_{n+r}}$ generate the relative tangent sheaf $\mathscr{T}_{X/Y}$.) Now every element of $\Gamma(X, \mathscr{D}_{X \to Y})$ can be written in the form

$$\sum_{\alpha \in \mathbb{N}^n} g_\alpha \otimes \partial_{y_1}^{\alpha_1} \cdots \partial_{y_n}^{\alpha_n},$$

with $g_\alpha \in \Gamma(X, \mathscr{O}_X)$, and because of how we defined the $\mathscr{D}_X$-module structure on the transfer module, this expression equals

$$\sum_{\alpha \in \mathbb{N}^n} g_\alpha \partial_{x_1}^{\alpha_1} \cdots \partial_{x_n}^{\alpha_n} \cdot (1 \otimes 1).$$

Thus $\mathscr{D}_X \to \mathscr{D}_{X \to Y}$ is indeed surjective, with kernel generated by the relative tangent sheaf $\mathscr{T}_{X/Y}$.

It remains to prove that $\operatorname{Ch}(f^*\mathcal{M}) = df\big(p_2^{-1}\operatorname{Ch}(\mathcal{M})\big)$. Choose a good filtration $F_\bullet\mathcal{M}$, and observe that because f is flat, we have $f^*F_j\mathcal{M} \subseteq f^*\mathcal{M}$. If we set $\mathcal{N} = f^*\mathcal{M}$, we thus get a filtration with terms $F_j\mathcal{N} = f^*F_j\mathcal{M}$. It is clear that each $F_j\mathcal{N}$ is a coherent $\mathscr{O}_X$-module; moreover, flatness of f gives

$$\operatorname{gr}_j^F \mathcal{N} = F_j\mathcal{N}/F_{j-1}\mathcal{N} \cong f^* \operatorname{gr}_j^F \mathcal{M}.$$

Once we check that $F_\bullet\mathcal{N}$ is a good filtration, we can use it to compute $\operatorname{Ch}(\mathcal{N})$. Working locally, we can assume that X and Y are affine, and that we have local coordinates $x_1, \dots, x_{n+r} \in \Gamma(X, \mathscr{O}_X)$ and $y_1, \dots, y_n \in \Gamma(Y, \mathscr{O}_Y)$ as above. To abbreviate, set $A = \Gamma(X, \mathscr{O}_X)$ and $B = \Gamma(Y, \mathscr{O}_Y)$; then A is a smooth B-algebra. We shall use the same symbol ∂_j to denote both ∂_{x_j} and ∂_{y_j}; then the morphism on tangent sheaves takes ∂_j to $1 \otimes \partial_j$ for $1 \leq j \leq n$, and to zero otherwise.

Let us set $M = \Gamma(Y, \mathcal{M})$ and $N = \Gamma(X, \mathcal{N})$. By construction,

$$N = A \otimes_B M \quad \text{and} \quad F_jN = A \otimes_B F_jM \quad \text{and} \quad \operatorname{gr}_j^F N = A \otimes_B \operatorname{gr}_j^F M.$$

As the filtration on M is good, the associated graded $\operatorname{gr}^F M$ is finitely generated over $\operatorname{gr}^F D(B) = B[\partial_1, \dots, \partial_n]$. The left $D(A)$-module structure on N is given by

$$\partial_j(a \otimes m) = \begin{cases} \partial_j a \otimes m + a \otimes \partial_j m & \text{if } 1 \leq j \leq n, \\ \partial_j a \otimes m & \text{if } n+1 \leq j \leq n+r. \end{cases}$$

This formula shows that the filtration $F_\bullet N$ is compatible with the action by $D(A)$. It also shows that $\partial_{n+1}, \dots, \partial_{n+r}$ act trivially on

$$\operatorname{gr}^F N = A \otimes_B \operatorname{gr}^F M,$$

and that $\partial_1, \dots, \partial_n$ only act on the second factor. Said differentialy, we have an isomorphism of graded $A[\partial_1, \dots, \partial_{n+r}]$-modules

$$\operatorname{gr}^F N \cong A[\partial_1, \dots, \partial_n] \otimes_{B[\partial_1, \dots, \partial_n]} \operatorname{gr}^F M, \tag{16.6}$$

with $A[\partial_1, \dots, \partial_{n+r}]$ acting on the first factor in the obvious way. This says that $\operatorname{gr}^F N$ is finitely generated over $A[\partial_1, \dots, \partial_{n+r}]$, and so $F_\bullet N$ is a good filtration.

It is now easy to compute the characteristic variety $\operatorname{Ch}(\mathcal{N})$. If we rewrite the diagram in (16.1) in terms of rings, we get

$$\begin{tikzcd} \operatorname{Spec} A[\partial_1, \dots, \partial_n] \arrow[r, "df"] \arrow[d, "p_2"] & \operatorname{Spec} A[\partial_1, \dots, \partial_{n+r}] \\ \operatorname{Spec} B[\partial_1, \dots, \partial_n] \end{tikzcd}$$

with p_2 induced by the morphism of rings $B \to A$, and df induced by the quotient morphism $A[\partial_1, \dots, \partial_{n+r}] \to A[\partial_1, \dots, \partial_n]$. Thus (16.6) says that the coherent sheaf on $T^*X = \operatorname{Spec} A[\partial_1, \dots, \partial_{n+r}]$ corresponding to $\operatorname{gr}^F N$ is obtained by first pulling back $\operatorname{gr}^F M$ along p_2, and then pushing forward along df. Globally,

$$\widetilde{\operatorname{gr}^F \mathcal{N}} \cong df_* p_2^* \widetilde{\operatorname{gr}^F \mathcal{M}},$$

and since p_2 is surjective and df a closed embedding, we get

$$\operatorname{Ch}(\mathcal{N}) = df\big(p_2^{-1} \operatorname{Ch}(\mathcal{M})\big),$$

proving (c) for all smooth morphisms.

Factorizing through the graph. Using the graph embedding, we can write any morphism $f\colon X \to Y$ as the composition of a closed embedding $i\colon X \hookrightarrow Z$ and a smooth morphism $g\colon Z \to Y$. (Here $Z = X \times Y$, of course, but let me write Z to simplify the notation.) We already know that $\mathcal{N} = g^*\mathcal{M}$ is coherent over $\mathscr{D}_Z$, and that $\operatorname{Ch}(\mathcal{N}) = dg\big(p_2^{-1} \operatorname{Ch}(\mathcal{M})\big)$. Using the big diagram

$$\begin{tikzcd} X \times_Y T^*Y \arrow[r, hook] \arrow[d, "i \times \mathrm{id}"] \arrow[rr, bend left, "df"] & X \times_Z T^*Z \arrow[r, "di"] \arrow[d, "p_2"] & T^*Z \\ Z \times_Y T^*Y \arrow[r, hook, "dg"] \arrow[d, "p_2"] & T^*Z \\ T^*Y \end{tikzcd}$$

from last time, we see that $p_2^{-1} \operatorname{Ch}(\mathcal{N})$ is finite over its image in T^*X (under the morphism di); this says that $\mathcal{N}$ is non-characteristic with respect to the closed embedding $i\colon X \hookrightarrow Z$. As $f^*\mathcal{M} \cong i^*\mathcal{N}$, this reduces the proof of Theorem 16.5 to the case of a closed embedding.

Closed embeddings. Suppose now that $f\colon X \to Y$ is a closed embedding. We are only going to treat the case where $\dim X = \dim Y - 1$; to go from there to the general case, one uses the fact that f can be locally factored as a composition of $\dim Y - \dim X$ closed embeddings of codimension one (because closed embeddings between nonsingular algebraic varieties are locally complete intersections).

The problem is local, and so we can assume that Y is affine, with $B = \Gamma(Y, \mathscr{O}_Y)$. Choose local coordinates $y_0, y_1, \dots, y_n \in B$, in such a way that X is defined by the equation $y_0 = 0$; then $A = \Gamma(X, \mathscr{O}_X) \cong B/By_0$, and the images $x_1, \dots, x_n \in A$ of $y_1, \dots, y_n \in B$ are local coordinates on X. The morphism on tangent sheaves

$$\mathscr{T}_X \to f^*\mathscr{T}_Y = \mathscr{O}_X \otimes_{f^{-1}\mathscr{O}_Y} f^{-1}\mathscr{T}_Y$$

takes ∂_{x_j} to $1 \otimes \partial_{y_j}$ for $1 \leq j \leq n$. (The remaining vector field ∂_{y_0} is not in the image; it generates the normal bundle.) We again write ∂_j for both ∂_{x_j} and ∂_{y_j}, so that the morphism on tangent sheaves takes ∂_j to $1 \otimes \partial_j$. With this notation, the diagram in (16.1) becomes

$$\begin{tikzcd} \operatorname{Spec} A[\partial_0, \dots, \partial_n] \arrow[r, "df"] \arrow[d, "p_2"] & \operatorname{Spec} A[\partial_1, \dots, \partial_n] \\ \operatorname{Spec} B[\partial_0, \dots, \partial_n]. \end{tikzcd} \tag{16.7}$$

This time, p_2 is a closed embedding and df is smooth of relative dimension one.

We are going to use the following basic fact from algebraic geometry.

Lemma 16.8. *Let B be a finitely generated A-algebra.*

(1) If B is integral over A, then every finitely generated B-module M is also finitely generated as an A-module.

(2) If M is a finitely generated B-module such that $\operatorname{Supp} M$ is finite over $\operatorname{Spec} A$, then M is also finitely generated as an A-module.

Proof. The first assertion follows from the fact that B itself is finitely generated as an A-module. To prove the second assertion, we may replace B by the quotient ring $B/\operatorname{Ann}_B(M)$ and assume without loss of generality that $\operatorname{Ann}_B(M) = 0$. The support of M is then the reduced closed subscheme defined by the nilradical of B, and so the hypothesis says that $B/\operatorname{Nil} B$ is integral over A. This means that for every $b \in B$, there is a monic polynomial $h(t) \in A[t]$ such that $h(b) \in \operatorname{Nil} B$. But then $h(b)^m = 0$ for some $m \geq 1$, and so b is integral over A. We now conclude from the first assertion that M is finitely generated as an A-module. □

Now let $\mathcal{M}$ be a coherent left $\mathscr{D}_Y$-module that is non-characteristic with respect to f. Set $M = \Gamma(Y, \mathcal{M})$, which is a finitely generated module over the ring of differential operators $D(B) = \Gamma(Y, \mathscr{D}_Y)$. The following lemma expresses the non-characteristic property of M in terms of differential operators.

Lemma 16.9. *For every $u \in M$, there exists a nontrivial differential operator $P \in D(B)$ that is non-characteristic with respect to $y_0 = 0$ and satisfies $Pu = 0$.*

Proof. The submodule $D(B)u \subseteq M$ is isomorphic to $D(B)/I$, where

$$I = \left\{ P \in D(B) \;\middle|\; Pu = 0 \right\}$$

is a left ideal in $D(B)$. The characteristic variety of $D(B)/I$ is contained in that of M, and so $D(B)/I$ is again non-characteristic with respect to f. As a subset of $T^*Y = \operatorname{Spec} B[\partial_0, \dots, \partial_n]$, the characteristic variety of $D(B)/I$ is cut out by the principal symbols $\sigma(P) \in B[\partial_0, \dots, \partial_n]$ of all the differential operators $P \in I$. Its preimage under p_2 is therefore cut out by their images in $A[\partial_0, \dots, \partial_n]$. Because this subset is finite over $\operatorname{Spec} A[\partial_1, \dots, \partial_n]$, we can argue as in the preceding lemma to show that there is a monic polynomial $h(t)$ of some degree $d \geq 1$, with coefficients in the ring $A[\partial_1, \dots, \partial_n]$, such that $h(\partial_0) \in A[\partial_0, \dots, \partial_n]$ belongs to the ideal generated by $\sigma(P)$ for $P \in I$. Keeping all terms in $h(\partial_0)$ that are homogeneous of degree d, we conclude that there exists a differential operator $P \in I$ of order d, such that the image of $\sigma(P)$ in $A[\partial_0, \dots, \partial_n]$ contains the term ∂_0^d. But this says exactly that P is non-characteristic with respect to $y_0 = 0$. □

Note. Since M is finitely generated over $D(B)$, the lemma implies that there exist finitely many differential operators $P_1, \dots, P_r \in D(B)$, all non-characteristic with respect to $y_0 = 0$, and a surjective morphism

$$\bigoplus_{i=1}^{r} D(B)/D(B)P_i \to M.$$

By applying the same observation to the kernel, one can in fact show that M admits a resolution by non-characteristic $D(B)$-modules of the form $D(B)/D(B)P$.

Now let us continue with the proof of Theorem 16.5. The derived functors $L^{-j} f^* \mathcal{M}$ are computed, in our local coordinates, by the complex of $D(A)$-modules

$$M \xrightarrow{y_0} M.$$

To show that $L^{-j} f^* \mathcal{M} = 0$ for every $j \geq 1$, we only have to argue that multiplication by y_0 is injective. Suppose that we have some $u \in M$ with $y_0 u = 0$. By the lemma, we can find a differential operator $P \in D(B)$, say of degree $d \geq 0$, such that $Pu = 0$ and such that P is non-characteristic with respect to $y_0 = 0$. Concretely, this means that the coefficient of ∂_0^d is constant modulo y_0. As $y_0 u = 0$, we can therefore assume without loss of generality that ∂_0^d appear with coefficient 1 in P. Let us choose P in such a way that d is minimal. The commutator $[y_0, P]$ contains the term $-d\partial_0^{d-1}$, and since

$$[y_0, P]u = y_0 P u - P y_0 u = 0,$$

we conclude by minimality that $d = 0$, and hence that $u = 0$. This proves (b).

To prove the other two assertions, we choose a good filtration $F_\bullet M$, with $\mathrm{gr}^F M$ finitely generated over $\mathrm{gr}^F D(B) = B[\partial_0, \dots, \partial_n]$. Set $\mathcal{N} = f^* \mathcal{M}$ and $N = \Gamma(X, \mathcal{N})$, so that

$$N = A \otimes_B M.$$

This time, tensoring with A is no longer an exact functor, but we can still define a filtration on N by setting

$$F_j N = \mathrm{im}\big(A \otimes_B F_j M \to A \otimes_B M\big).$$

With this definition, each $\mathrm{gr}_j^F N$ is a quotient of $B \otimes_A \mathrm{gr}_j^F M$, and by exactly the same calculation as before, the $A[\partial_1, \dots, \partial_n]$-module $\mathrm{gr}^F N$ is a quotient of $A \otimes_B \mathrm{gr}^F M$, considered as an $A[\partial_1, \dots, \partial_n]$-module through the morphism in (16.7).

Now I claim that $A \otimes_B \mathrm{gr}^F M$ is finitely generated over $A[\partial_1, \dots, \partial_n]$. Indeed, $\mathrm{gr}^F M$ is finitely generated over $B[\partial_0, \dots, \partial_n]$ (because $F_\bullet M$ is good), and so $A \otimes_B \mathrm{gr}^F M$ is finitely generated over $A[\partial_0, \dots, \partial_n]$. By the non-characteristic property, the support inside $\operatorname{Spec} A[\partial_0, \dots, \partial_n]$ is finite over $\operatorname{Spec} A[\partial_1, \dots, \partial_n]$, and so the claim follows from Lemma 16.9. Therefore $\mathrm{gr}^F N$, which is a quotient, is also finitely generated over $A[\partial_1, \dots, \partial_n]$, proving that $\mathcal{N} = f^* \mathcal{M}$ is coherent over $\mathscr{D}_X$. This argument also shows that

$$\mathrm{Ch}(\mathcal{N}) \subseteq df\big(p_2^{-1} \, \mathrm{Ch}(\mathcal{M})\big),$$

because the support of $A \otimes_B \mathrm{gr}^F M$ contains the support of the quotient module $\mathrm{gr}^F N$. Some extra work is required to show that the two sides are actually equal. (In brief, one has to construct a good filtration $F_\bullet M$ such that $\mathrm{gr}_j^F N = A \otimes_B \mathrm{gr}_j^F M$.)

Exercises.

Exercise 16.1. Suppose that $X \subseteq \mathbb{A}^n$ is a nonsingular subvariety. Determine the set of hyperplanes $H \subseteq \mathbb{A}^n$ such that $p_2^{-1}(T_X^* \mathbb{A}^n)$ is finite over its image in $T^* H$.

$$\begin{tikzcd} H \times_{\mathbb{A}^n} T^*\mathbb{A}^n \arrow[r] \arrow[d, "p_2"] & T^*H \\ T^*\mathbb{A}^n & \end{tikzcd}$$

Lecture 17: April 10

Direct images in general. We are now going to define the direct image functor for (right) $\mathscr{D}$-modules for an arbitrary morphism $f\colon X \to Y$ between nonsingular algebraic varieties. Let $\mathcal{M}$ be a right $\mathscr{D}_X$-module. By analogy with the case of closed embeddings, the direct image should be

$$f_*\big(\mathcal{M} \otimes_{\mathscr{D}_X} \mathscr{D}_{X\to Y}\big).$$

Recall that the transfer module $\mathscr{D}_{X\to Y} = \mathscr{O}_X \otimes_{f^{-1}\mathscr{O}_Y} f^{-1}\mathscr{D}_Y$ is a $(\mathscr{D}_X, f^{-1}\mathscr{D}_Y)$-bimodule, and so the direct image is again a right $\mathscr{D}_Y$-module. The problem with this definition is that the resulting functor is neither right nor left exact, and therefore not suitable from a homological algebra standpoint. (The reason is that we are mixing the right exact functor $\otimes$ with the left exact functor f_*.) This problem can be fixed by working in the derived category; in fact, Sato, who founded algebraic analysis, independently invented the theory of derived categories for his needs.

Derived categories. Let me very briefly review some basic facts. Let X be a topological space, and $\mathscr{R}_X$ a sheaf of (maybe noncommutative) rings on X. We denote by $\mathrm{Mod}(\mathscr{R}_X)$ the category of (sheaves of) left $\mathscr{R}_X$-modules; this is an abelian category. Note that right $\mathscr{R}_X$-modules are the same thing as left modules over the opposite ring $\mathscr{R}_X^{op}$. We use the notation

$$D^b(\mathscr{R}_X)$$

for the *derived category* of cohomologically bounded complexes of left $\mathscr{R}_X$-modules. The objects of this category are complexes of left $\mathscr{R}_X$-modules, with the property that only finitely many of the cohomology sheaves are nonzero. The set of morphisms between two objects takes more time to describe, and this is where the action is happening. Recall that when we compute a derived functor, we have to replace a sheaf (or complex of sheaves) by a suitable resolution: injective resolutions in the case of pushforward, flat resolutions in the case of tensor product, etc. The reason for introducing the derived category is that one wants to have a place where a sheaf (or complex of sheaves) is isomorphic to any of its resolutions.

Example 17.1. Suppose that we choose an injective resolution

$$0 \to \mathscr{F} \to \mathcal{I}^0 \to \mathcal{I}^1 \to \cdots$$

for a sheaf of $\mathscr{O}_X$-modules, say. Homological algebra shows that any two such resolutions are the same up to homotopy, meaning that if $\mathcal{J}^\bullet$ is another injective resolution of $\mathscr{F}$, then there is a morphism of complexes $\mathcal{I}^\bullet \to \mathcal{J}^\bullet$, unique up to homotopy; and its composition with the morphism going the other way is homotopic to the identity morphism. But $\mathscr{F}$ is not isomorphic to the complex $\mathcal{I}^\bullet$; all one has is a *quasi-isomorphism*, meaning a morphism of complexes that induces isomorphisms on cohomology sheaves. So if we want $\mathscr{F}$ to be isomorphic to $\mathcal{I}^\bullet$, then we need to work up to homotopy and somehow create an inverse for the morphism $\mathscr{F} \to \mathcal{I}^\bullet$.

Back to $D^b(\mathscr{R}_X)$. The set of morphisms between two objects is obtained by a two-step procedure: starting from all morphisms of complexes, one first identifies morphisms that are homotopy equivalent, and then one formally adjoins inverses for all quasi-isomorphisms. As I said, this construction makes sure that a sheaf (or complex of sheaves) is isomorphic to any of its resolutions by a unique isomorphism.

Concerning the existence of resolutions, one has the following basic fact:

(1) Every $\mathscr{R}_X$-module can be embedded into an injective $\mathscr{R}_X$-module.
(2) Every $\mathscr{R}_X$-module is a quotient of a flat $\mathscr{R}_X$-module.

One can then use the Cartan-Eilenberg construction to show that every cohomologically bounded complex of $\mathscr{R}_X$-modules has both injective and flat resolutions.

The direct image functor. We can now define the direct image functor for an arbitrary morphism $f\colon X \to Y$ between nonsingular algebraic varieties. The construction is done in two stages. First, we have a functor

$$D^b(\mathscr{D}_X^{op}) \to D^b(f^{-1}\mathscr{D}_Y^{op}), \quad \mathcal{M}^\bullet \mapsto \mathcal{M}^\bullet \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_X} \mathscr{D}_{X\to Y},$$

obtained by taking the derived tensor product with the transfer module $\mathscr{D}_{X\to Y}$. Concretely, this means that we choose a flat resolution for the complex of right $\mathscr{D}_X$-modules $\mathcal{M}^\bullet$, and then tensor this resolution with $\mathscr{D}_{X\to Y}$. For the time being, we do not make any quasi-coherence assumptions. Second, we have a functor

$$D^b(f^{-1}\mathscr{D}_Y^{op}) \to D^b(\mathscr{D}_Y^{op}), \quad \mathcal{N}^\bullet \mapsto \mathbf{R}f_*\mathcal{N}^\bullet,$$

obtained by applying the derived pushforward functor for sheaves. Concretely, this means that we choose an injective resolution for the complex of right $f^{-1}\mathscr{D}_Y$-modules $\mathcal{N}^\bullet$, and then apply the usual pushforward functor f_* to each sheaf in the complex. Each sheaf in the resulting complex is naturally a right $\mathscr{D}_Y$-module through the morphism $\mathscr{D}_Y \to f_*f^{-1}\mathscr{D}_Y$.

One has to show that both functors are well-defined and "exact", meaning that they preserve distinguished triangles (which are the derived category version of short exact sequences of complexes). We define the pushforward functor as the composition of the two functors above.

Definition 17.2. Let $f\colon X \to Y$ be a morphism between nonsingular algebraic varieties. The *pushforward* is the exact functor

$$f_+\colon D^b(\mathscr{D}_X^{op}) \to D^b(\mathscr{D}_Y^{op}), \quad f_+\mathcal{M}^\bullet = \mathbf{R}f_*\bigl(\mathcal{M}^\bullet \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_X} \mathscr{D}_{X\to Y}\bigr)$$

between derived categories.

Note that the general definition involves first choosing a flat resolution for the complex $\mathcal{M}^\bullet$, and then a second injective resolution for $\mathcal{M}^\bullet \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_X} \mathscr{D}_{X\to Y}$. Of course, this is only for theoretical purposes; in practice, we factor f into a closed embedding followed by a projection, and there are simple formulas for computing the pushforward in both cases.

Example 17.3. Another word about resolutions. In the case of $\mathscr{D}_X$-modules, one can use results about $\mathscr{O}_X$-modules to get resolutions very easily. For example, suppose that we want to represent a quasi-coherent right $\mathscr{D}_X$-module $\mathcal{M}$ as a quotient of a flat $\mathscr{D}_X$-module. Pick a quasi-coherent $\mathscr{O}_X$-module $\mathscr{F} \subseteq \mathcal{M}$ that generates $\mathcal{M}$ as a $\mathscr{D}_X$-module. If $\mathcal{M}$ is a coherent $\mathscr{D}_X$-module, we can choose $\mathscr{F}$ to be a coherent $\mathscr{O}_X$-module; in general, $\mathscr{F} = \mathcal{M}$ will always do the job. Now pick a flat $\mathscr{O}_X$-module $\mathscr{E}$ that maps onto $\mathscr{F}$. Then the composition

$$\mathscr{E} \otimes_{\mathscr{O}_X} \mathscr{D}_X \to \mathscr{F} \otimes_{\mathscr{O}_X} \mathscr{D}_X \to \mathcal{M}$$

is surjective, and $\mathscr{E} \otimes_{\mathscr{O}_X} \mathscr{D}_X$ is flat as a right $\mathscr{D}_X$-module.

Here are some concrete examples of the pushforward functor.

Example 17.4. Suppose that $i\colon X \hookrightarrow Y$ is a closed embedding. In this case, the transfer module $\mathscr{D}_{X\to Y}$ is locally free (as a left $\mathscr{D}_X$-module), and tensoring with $\mathscr{D}_{X\to Y}$ is therefore exact. The pushforward functor i_* is also exact, and so we have

$$i_+\mathcal{M}^\bullet = i_*\bigl(\mathcal{M}^\bullet \otimes_{\mathscr{D}_X} \mathscr{D}_{X\to Y}\bigr).$$

This agrees with our earlier definition in the case of a single $\mathscr{D}_X$-module; in the case of a complex, we simply apply the naive pushforward functor for a closed embedding term by term.

Example 17.5. Suppose that $j\colon U \hookrightarrow Y$ is an open embedding. Then

$$\mathscr{D}_{U\to Y} = \mathscr{O}_U \otimes_{j^{-1}\mathscr{O}_Y} j^{-1}\mathscr{D}_Y \cong \mathscr{D}_U,$$

by the basic properties of $\mathscr{D}_Y$ from Lecture 9. This shows that the pushforward functor agrees with $\mathbf{R}j_*$ in this case. Generally speaking, j_* is exact when the complement $Y \setminus U$ is a divisor; otherwise, there might be higher derived functors. The localization $k[x_1, \dots, x_n, p^{-1}]$ that we analyzed in Lecture 3 is a concrete example, namely the pushforward of $k[x_1, \dots, x_n]$ along the open embedding $\mathbb{A}^n \setminus Z(p) \hookrightarrow \mathbb{A}^n$.

Example 17.6. Let's consider the case where $f\colon X \to \operatorname{Spec} k$ is the morphism to a point. In this case, the pushforward $f_+\mathcal{M}$ should be viewed as something like the cohomology of X with coefficients in a right $\mathscr{D}_X$-module $\mathcal{M}$. The transfer module

$$\mathscr{D}_{X\to \operatorname{Spec} k} = \mathscr{O}_X \otimes_{f^{-1}\mathscr{O}_{\operatorname{Spec} k}} f^{-1}\mathscr{D}_{\operatorname{Spec} k} \cong \mathscr{O}_X$$

is just $\mathscr{O}_X$ in this case; it has the structure of a left $\mathscr{D}_X$-module (and a right k-module). To compute the pushforward

$$f_+\mathcal{M} = \mathbf{R}f_*\bigl(\mathcal{M} \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_X} \mathscr{O}_X\bigr),$$

we can use a resolution of $\mathscr{O}_X$ by left $\mathscr{D}_X$-modules. Such a resolution is furnished by the *Spencer complex*

$$\operatorname{Sp}(\mathscr{D}_X) = \Bigl[\mathscr{D}_X \otimes_{\mathscr{O}_X} \bigwedge^n \mathscr{T}_X \to \dots \to \mathscr{D}_X \otimes_{\mathscr{O}_X} \bigwedge^2 \mathscr{T}_X \to \mathscr{D}_X \otimes_{\mathscr{O}_X} \mathscr{T}_X \to \mathscr{D}_X\Bigr],$$

which lives in degrees $-n, \dots, -1, 0$. The Spencer complex maps to $\mathscr{O}_X$ via the $\mathscr{D}_X$-linear map $\mathscr{D}_X \to \mathscr{O}_X$ that takes $P \in \mathscr{D}_X$ to $P(1) \in \mathscr{O}_X$. This is surjective, and the kernel is generated by $\mathscr{T}_X$. The general formula for the differentials

$$d\colon \mathscr{D}_X \otimes_{\mathscr{O}_X} \bigwedge^{k+1} \mathscr{T}_X \to \mathscr{D}_X \otimes_{\mathscr{O}_X} \bigwedge^{k} \mathscr{T}_X$$

in the Spencer complex is as follows:

$$\begin{aligned} d\bigl(P \otimes \theta_0 \wedge \theta_1 \wedge \dots \wedge \theta_k\bigr) &= \sum_{i=0}^k (-1)^i (P\theta_i) \otimes \theta_0 \wedge \dots \wedge \widehat{\theta_i} \wedge \dots \wedge \theta_k \\ &+ \sum_{0 \le i < j \le k} (-1)^{i+j} P \otimes [\theta_i, \theta_j] \wedge \theta_0 \wedge \dots \wedge \widehat{\theta_i} \wedge \dots \wedge \widehat{\theta_j} \wedge \dots \wedge \theta_k \end{aligned}$$

In local coordinates $x_1, \dots, x_n$, the tangent sheaf is a free $\mathscr{O}_X$-module with basis $\partial_1, \dots, \partial_n$, and the above formula simplifies to

$$d\bigl(P \otimes \partial_{i_0} \wedge \partial_{i_1} \wedge \dots \wedge \partial_{i_k}\bigr) = \sum_{j=0}^k (-1)^j (P\partial_{i_j}) \otimes \partial_{i_0} \wedge \dots \wedge \widehat{\partial_{i_j}} \wedge \dots \wedge \partial_{i_k}.$$

Except for the fact that $\mathscr{D}_X$ is noncommutative, this is the same formula as for the differentials in a Koszul complex. Let us check that the Spencer complex resolves $\mathscr{O}_X$. From the formula for the differentials, it is clear that we can filter $\operatorname{Sp}(\mathscr{D}_X)$ by the family of subcomplexes

$$F_p \operatorname{Sp}(\mathscr{D}_X) = \Bigl[F_{p-n}\mathscr{D}_X \otimes_{\mathscr{O}_X} \bigwedge^n \mathscr{T}_X \to \dots \to F_{p-1}\mathscr{D}_X \otimes_{\mathscr{O}_X} \mathscr{T}_X \to F_p\mathscr{D}_X\Bigr].$$

The description of the differential in local coordinates shows that the associated graded complex

$$\operatorname{gr}_\bullet^F \operatorname{Sp}(\mathscr{D}_X) = \Bigl[\operatorname{gr}_{\bullet-n}^F \mathscr{D}_X \otimes_{\mathscr{O}_X} \bigwedge^n \mathscr{T}_X \to \dots \to \operatorname{gr}_{\bullet-1}^F \mathscr{D}_X \otimes_{\mathscr{O}_X} \mathscr{T}_X \to \operatorname{gr}_\bullet^F \mathscr{D}_X\Bigr]$$

identifies with the Koszul complex for the regular sequence $\partial_1, \dots, \partial_n \in \mathrm{gr}_1^F \mathscr{D}_X$, and is therefore a resolution of $\mathscr{O}_X$ as a graded $\mathrm{gr}^F \mathscr{D}_X$-module. This proves that the Spencer complex resolves $\mathscr{O}_X$ as a left $\mathscr{D}_X$-module.

Since each term of the Spencer complex is a locally free $\mathscr{D}_X$-module, we get

$$f_+ \mathcal{M} \cong \mathbf{R} f_* \big(\mathcal{M} \otimes_{\mathscr{D}_X} \mathrm{Sp}(\mathscr{D}_X) \big) = \mathbf{R} f_* \, \mathrm{Sp}(\mathcal{M}),$$

where the Spencer complex of $\mathcal{M}$ is defined analogously by

$$\mathrm{Sp}(\mathcal{M}) = \Big[\mathcal{M} \otimes_{\mathscr{O}_X} \bigwedge^n \mathscr{T}_X \to \dots \to \mathcal{M}_X \otimes_{\mathscr{O}_X} \bigwedge^2 \mathscr{T}_X \to \mathcal{M}_X \otimes_{\mathscr{O}_X} \mathscr{T}_X \to \mathcal{M} \Big],$$

with the same formula for the differentials. The pushforward of a right $\mathscr{D}_X$-module is therefore equal to the hypercohomology of its Spencer complex $\mathrm{Sp}(\mathcal{M})$.

Example 17.7. In the case of ω_X, you can check that the Spencer complex $\mathrm{Sp}(\omega_X)$ is isomorphic to the algebraic de Rham complex

$$\mathrm{DR}(\mathscr{O}_X) = \Big[\mathscr{O}_X \to \Omega^1_{X/k} \to \dots \to \Omega^n_{X/k} \Big].$$

The j-th hypercohomology group of the de Rham complex is denoted by $H^j_{dR}(X/k)$ and is called the *j-th algebraic de Rham cohomology* of X. When X is defined over the complex numbers, Grothendieck's comparison theorem tells us that $H^j_{dR}(X/\mathbb{C}) \cong H^j(X, \mathbb{C})$ is isomorphic to the singular cohomology of X, considered as a complex manifold.

Let us check that the pushforward functor is compatible with composition of morphisms.

Proposition 17.8. *Let $f \colon X \to Y$ and $g \colon Y \to Z$ be morphisms between nonsingular algebraic varieties. Then one has $g_+ \circ f_+ \cong (g \circ f)_+$, as functors from $D^b(\mathscr{D}_X^{op})$ to $D^b(\mathscr{D}_Z^{op})$.*

Proof. Let $\mathcal{M}^\bullet \in D^b(\mathscr{D}_X^{op})$ be any complex of right $\mathscr{D}_X$-modules. By definition,

$$\begin{aligned} g_+\big(f_+ \mathcal{M}^\bullet\big) &= \mathbf{R} g_* \Big(\mathbf{R} f_* \big(\mathcal{M}^\bullet \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_X} \mathscr{D}_{X \to Y} \big) \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_Y} \mathscr{D}_{Y \to Z} \Big) \\ (g \circ f)_+ \mathcal{M}^\bullet &= \mathbf{R}(g \circ f)_* \big(\mathcal{M}^\bullet \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_X} \mathscr{D}_{X \to Z} \big). \end{aligned}$$

We clearly need a relation among the three transfer modules to compare these two expressions. Here is the relevant computation:

$$\begin{aligned} \mathscr{D}_{X \to Z} &= \mathscr{O}_X \otimes_{(g \circ f)^{-1} \mathscr{O}_Z} (g \circ f)^{-1} \mathscr{D}_Z \\ &\cong \mathscr{O}_X \otimes_{f^{-1} \mathscr{O}_Y} \big(f^{-1} \mathscr{O}_Y \otimes_{f^{-1} g^{-1} \mathscr{O}_Z} f^{-1} g^{-1} \mathscr{D}_Z \big) \\ &\cong \mathscr{O}_X \otimes_{f^{-1} \mathscr{O}_Y} f^{-1} \big(\mathscr{O}_Y \otimes_{g^{-1} \mathscr{O}_Z} g^{-1} \mathscr{D}_Z \big) \\ &= \mathscr{O}_X \otimes_{f^{-1} \mathscr{O}_Y} f^{-1} \mathscr{D}_{Y \to Z} \\ &\cong \big(\mathscr{O}_X \otimes_{f^{-1} \mathscr{O}_Y} f^{-1} \mathscr{D}_Y \big) \otimes_{f^{-1} \mathscr{D}_Y} f^{-1} \mathscr{D}_{Y \to Z} \\ &= \mathscr{D}_{X \to Y} \otimes_{f^{-1} \mathscr{D}_Y} f^{-1} \mathscr{D}_{Y \to Z} \end{aligned}$$

In fact, since $\mathscr{D}_Z$ is locally free as an $\mathscr{O}_Z$-module, the higher derived functors of all the tensor products in the above calculation are trivial, and we even have

$$\mathscr{D}_{X \to Z} \cong \mathscr{D}_{X \to Y} \overset{\mathbf{L}}{\otimes}_{f^{-1} \mathscr{D}_Y} f^{-1} \mathscr{D}_{Y \to Z}. \tag{17.9}$$

Because $\mathbf{R}(g \circ f)_* \cong \mathbf{R} g_* \circ \mathbf{R} f_*$, it will therefore be enough to show that

$$\mathbf{R} f_* \big(\mathcal{M}^\bullet \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_X} \mathscr{D}_{X \to Y} \big) \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_Y} \mathscr{D}_{Y \to Z} \to \mathbf{R} f_* \big(\mathcal{M}^\bullet \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_X} \mathscr{D}_{X \to Y} \overset{\mathbf{L}}{\otimes}_{f^{-1} \mathscr{D}_Y} f^{-1} \mathscr{D}_{Y \to Z} \big)$$

is an isomorphism (in the derived category of right $g^{-1}\mathscr{D}_Z$-modules). Setting

$$A = \mathcal{M}^\bullet \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_X} \mathscr{D}_{X\to Y} \in D^b(f^{-1}\mathscr{D}_Y^{op}) \quad \text{and} \quad B = \mathscr{D}_{Y\to Z} \in D^b(\mathscr{D}_Y),$$

this is a consequence of the "projection formula" in the following lemma. ☐

Lemma 17.10. *If $A \in D^b(f^{-1}\mathscr{D}_Y^{op})$ and $B \in D^b(\mathscr{D}_Y)$, then*

$$\mathbf{R}f_* A \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_Y} B \to \mathbf{R}f_*\big(A \overset{\mathbf{L}}{\otimes}_{f^{-1}\mathscr{D}_Y} f^{-1}B\big)$$

is an isomorphism.

Proof. This is a local question, and so we can assume that Y is affine. We can then resolve B by a complex of free $\mathscr{D}_Y$-modules, and thereby reduce the problem to the case where B is a free $\mathscr{D}_Y$-module. But the result is obvious in that case because all the functors preserve direct sums. ☐

Exercises.

Exercise 17.1. The de Rham complex of a left $\mathscr{D}_X$-module $\mathcal{M}$ is defined as

$$\mathrm{DR}(\mathcal{M}) = \Big[\mathcal{M} \to \Omega^1_{X/k} \otimes_{\mathscr{O}_X} \mathcal{M} \to \cdots \to \Omega^n_{X/k} \otimes_{\mathscr{O}_X} \mathcal{M}\Big],$$

with differentials given in local coordinates $x_1, \dots, x_n$ by the formula

$$d(\alpha \otimes m) = d\alpha \otimes m + (-1)^{\deg \alpha} \sum_{j=1}^{n} dz_j \wedge \alpha \otimes (\partial_j m).$$

Here $n = \dim X$. Recall from Lecture 12 that $\mathscr{D}_X^\omega \otimes_{\mathscr{D}_X} \mathcal{M} \cong \omega_X \otimes_{\mathscr{O}_X} \mathcal{M}$ has the structure of a right $\mathscr{D}_X$-module. Show that the Spencer complex of $\mathscr{D}_X^\omega \otimes_{\mathscr{D}_X} \mathcal{M}$ is isomorphic to the de Rham complex of $\mathcal{M}$.

Exercise 17.2. Continuing from the previous exercise, show that

$$\mathbb{H}^{-n}\,\mathrm{DR}(\mathcal{M}) = \big\{\, s \in \Gamma(X, \mathcal{M}) \;\big|\; \partial_1 s = \cdots = \partial_n s = 0 \,\big\}$$

is the space of global sections of $\mathcal{M}$ that are annihilated by all vector fields.

Lecture 18: April 15

Direct images and coherence. Last time, we defined the direct image functor (for right $\mathscr{D}$-modules) as the composition

$$\begin{tikzcd} D^b(\mathscr{D}_X^{op}) \arrow[r, "\overset{\mathbf{L}}{\otimes}\mathscr{D}_{X\to Y}"] \arrow[rr, bend right, "f_+"'] & D^b(f^{-1}\mathscr{D}_Y^{op}) \arrow[r, "\mathbf{R}f_*"] & D^b(\mathscr{D}_Y^{op}) \end{tikzcd}$$

where $f\colon X \to Y$ is any morphism between nonsingular algebraic varieties. We also showed that $g_+ \circ f_+ \cong (g \circ f)_+$.

Today, our first task is to prove that direct images preserve quasi-coherence and, in the case when f is proper, coherence. The definition of the derived category $D^b(\mathscr{D}_X^{op})$ did *not* include any quasi-coherence assumptions. We are going to denote by $D^b_{qc}(\mathscr{D}_X^{op})$ the full subcategory of $D^b(\mathscr{D}_X^{op})$, consisting of those complexes of right $\mathscr{D}_X$-modules whose cohomology sheaves are quasi-coherent as $\mathscr{O}_X$-modules. Recall that we included the condition of quasi-coherence into our definition of algebraic $\mathscr{D}$-modules in Lecture 10. Similarly, we denote by $D^b_{coh}(\mathscr{D}_X^{op})$ the full subcategory of $D^b(\mathscr{D}_X^{op})$, consisting of those complexes of right $\mathscr{D}_X$-modules whose cohomology sheaves are coherent $\mathscr{D}_X$-modules (and therefore quasi-coherent $\mathscr{O}_X$-modules). This category is of course contained in $D^b_{qc}(\mathscr{D}_X^{op})$.

Theorem 18.1. *Let $f\colon X \to Y$ be a morphism between nonsingular algebraic varieties. Then the functor f_+ takes $D^b_{qc}(\mathscr{D}_X^{op})$ into $D^b_{qc}(\mathscr{D}_Y^{op})$. When f is proper, it also takes $D^b_{coh}(\mathscr{D}_X^{op})$ into $D^b_{coh}(\mathscr{D}_Y^{op})$.*

We are going to deduce this from the analogous result for $\mathscr{O}_X$-modules. Recall that if $\mathscr{F}$ is a quasi-coherent $\mathscr{O}_X$-module, then the higher direct image sheaves $R^j f_*\mathscr{F}$ are again quasi-coherent $\mathscr{O}_Y$-modules. Moreover, if $\mathscr{F}$ is coherent and f is a proper morphism, then each $R^j f_*\mathscr{F}$ is a coherent $\mathscr{O}_Y$-module. The first result is fairly elementary; the second one, due to Grauert in the analytic setting and to Grothendieck in the algebraic setting, takes more work to prove.

To go from $\mathscr{O}_X$-modules to $\mathscr{D}_X$-modules, we work with "induced $\mathscr{D}$-modules". The construction is straightforward. Given any $\mathscr{O}_X$-module $\mathscr{F}$, the tensor product

$$\mathscr{F} \otimes_{\mathscr{O}_X} \mathscr{D}_X$$

is a right $\mathscr{D}_X$-module in the obvious way. Right $\mathscr{D}_X$-modules of this form are called *induced $\mathscr{D}$-modules.* If $\mathscr{F}$ is quasi-coherent, then $\mathscr{F} \otimes_{\mathscr{O}_X} \mathscr{D}_X$ is quasi-coherent as an $\mathscr{O}_X$-module; if $\mathscr{F}$ is coherent, then $\mathscr{F} \otimes_{\mathscr{O}_X} \mathscr{D}_X$ is a coherent $\mathscr{D}_X$-module.

Lemma 18.2. *Every (quasi)coherent $\mathscr{D}_X$-module admits a resolution by (quasi)-coherent induced $\mathscr{D}_X$-modules. The same thing is true for complexes.*

Proof. The point is that every (quasi)coherent $\mathscr{D}_X$-module is the quotient of a (quasi)coherent induced $\mathscr{D}_X$-module. Indeed, if $\mathcal{M}$ is a right $\mathscr{D}_X$-module that is quasi-coherent over $\mathscr{O}_X$, then we can use the obvious surjection

$$\mathcal{M} \otimes_{\mathscr{O}_X} \mathscr{D}_X \to \mathcal{M}.$$

If $\mathcal{M}$ is a coherent right $\mathscr{D}_X$-module, we showed in Lecture 11 that there exists a coherent $\mathscr{O}_X$-module $\mathscr{F} \subseteq \mathcal{M}$ with the property that $\mathscr{F} \cdot \mathscr{D}_X = \mathcal{M}$. This says that

$$\mathscr{F} \otimes_{\mathscr{O}_X} \mathscr{D}_X \to \mathcal{M}$$

is surjective. The kernel of the morphism is again either quasi-coherent or coherent, and so we can iterate the construction to produce the desired resolution

$$\cdots \to \mathscr{F}_1 \otimes_{\mathscr{O}_X} \mathscr{D}_X \to \mathscr{F}_0 \otimes_{\mathscr{O}_X} \mathscr{D}_X \to \mathcal{M} \to 0.$$

Keep in mind that the morphisms $\mathscr{F}_k \otimes_{\mathscr{O}_X} \mathscr{D}_X \to \mathscr{F}_{k-1} \otimes_{\mathscr{O}_X} \mathscr{D}_X$ are typically *not* induced by morphisms of $\mathscr{O}_X$-modules $\mathscr{F}_k \to \mathscr{F}_{k-1}$.

To deduce the result for complexes, one can then apply the usual Cartan-Eilenberg construction. □

Direct images of induced $\mathscr{D}$-modules are very easy to compute. Indeed,

$$\begin{aligned}(\mathscr{F} \otimes_{\mathscr{O}_X} \mathscr{D}_X) \overset{\mathbf{L}}{\otimes}_{\mathscr{D}_X} \mathscr{D}_{X\to Y} &\cong \mathscr{F} \overset{\mathbf{L}}{\otimes}_{\mathscr{O}_X} \mathscr{D}_{X\to Y} = \mathscr{F} \overset{\mathbf{L}}{\otimes}_{\mathscr{O}_X} \bigl(\mathscr{O}_X \otimes_{f^{-1}\mathscr{O}_Y} f^{-1}\mathscr{D}_Y\bigr) \\ &\cong \mathscr{F} \overset{\mathbf{L}}{\otimes}_{f^{-1}\mathscr{O}_Y} f^{-1}\mathscr{D}_Y = \mathscr{F} \otimes_{f^{-1}\mathscr{O}_Y} f^{-1}\mathscr{D}_Y,\end{aligned}$$

due to the fact that $\mathscr{D}_Y$ is locally free, hence flat, over $\mathscr{O}_Y$. Now the usual projection formula (for $\mathscr{O}_Y$-modules) gives

$$f_+(\mathscr{F} \otimes_{\mathscr{O}_X} \mathscr{D}_X) \cong \mathbf{R}f_*\bigl(\mathscr{F} \otimes_{f^{-1}\mathscr{O}_Y} f^{-1}\mathscr{D}_Y\bigr) \cong \mathbf{R}f_*\mathscr{F} \otimes_{\mathscr{O}_Y} \mathscr{D}_Y.$$

All cohomology modules of this complex are therefore again induced $\mathscr{D}_Y$-modules of the form $R^j f_*\mathscr{F} \otimes_{\mathscr{O}_Y} \mathscr{D}_Y$. They are quasi-coherent as $\mathscr{O}_Y$-modules if $\mathscr{F}$ is quasi-coherent; and coherent as $\mathscr{D}_Y$-modules if $\mathscr{F}$ is coherent and f is proper. This proves the theorem for all induced $\mathscr{D}$-modules.

Proof of Theorem 18.1. Let us first prove the assertion about quasi-coherence. By the lemma, every object in $D^b_{qc}(\mathscr{D}_X^{op})$ is isomorphic to a complex of of quasi-coherent induced $\mathscr{D}_X$-modules, of the form

$$\cdots \to \mathscr{F}^p \otimes_{\mathscr{O}_X} \mathscr{D}_X \to \mathscr{F}^{p+1} \otimes_{\mathscr{O}_X} \mathscr{D}_X \to \cdots$$

let me stress again that the differentials in this complex are $\mathscr{D}_X$-linear, but not induced by $\mathscr{O}_X$-linear morphisms from $\mathscr{F}^p$ to $\mathscr{F}^{p+1}$. If we apply the direct image functor f_+ to this complex, and use our calculation for induced $\mathscr{D}$-modules from above, we obtain a spectral sequence with

$$E_1^{p,q} = (R^q f_*\mathscr{F}^p) \otimes_{\mathscr{O}_Y} \mathscr{D}_Y$$

that converges to the cohomology sheaves of $f_+\bigl(\mathscr{F}^\bullet \otimes_{\mathscr{O}_X} \mathscr{D}_X\bigr)$. Each $E_1^{p,q}$ is quasi-coherent as an $\mathscr{O}_Y$-module, and so the cohomology sheaves of the direct image are also quasi-coherent as $\mathscr{O}_Y$-modules.

The proof for coherence is similar. By the lemma, every object in $D^b_{coh}(\mathscr{D}_X^{op})$ is isomorphic to a complex of coherent induced $\mathscr{D}_X$-modules; this means that we can choose all the $\mathscr{F}^p$ as coherent $\mathscr{O}_X$-modules. If $f\colon X \to Y$ is proper, then each $R^q f_*\mathscr{F}^p$ is a coherent $\mathscr{O}_Y$-module. But then each $E_1^{p,q}$ is a coherent $\mathscr{D}_Y$-module, and the spectral sequence implies that the cohomology sheaves of the direct image are also coherent $\mathscr{D}_Y$-modules. □

Example 18.3. Suppose that X is proper over $\operatorname{Spec} k$. Then Theorem 18.1 says in particular that the hypercohomology groups of $\operatorname{Sp}(\mathcal{M})$ are finite-dimensional k-vector spaces for every coherent right $\mathscr{D}_X$-module $\mathcal{M}$. In particular, the algebraic de Rham cohomology groups $H^j_{dR}(X/k)$ are finite-dimensional whenever X is proper over $\operatorname{Spec} k$. (We will see later that this is actually true without properness!)

Example 18.4. Our calculation for induced $\mathscr{D}$-modules shows that the direct image of a coherent $\mathscr{D}_X$-module by a non-proper morphism is usually not coherent. For example, if $f\colon X \to \operatorname{Spec} k$ is not proper, the j-th cohomology module of $f_+\mathscr{D}_X$ is isomorphic to $H^j(X, \mathscr{O}_X)$, which is typically not finite-dimensional over k.

Preservation of holonomicity. The direct and inverse image functors

$$f_+ \colon D^b_{qc}(\mathscr{D}_X^{op}) \to D^b_{qc}(\mathscr{D}_Y^{op}) \quad \text{and} \quad \mathbf{L}f^* \colon D^b_{qc}(\mathscr{D}_Y) \to D^b_{qc}(\mathscr{D}_X)$$

only preserve coherence with some extra assumptions. For $\mathbf{L}f^*$, we need the non-characteristic property; for f_+, we need properness. A small miracle of the theory is that both functors nevertheless preserve the most interesting class of $\mathscr{D}$-modules, namely the holonomic ones. We have already seen one special case of this phenomenon back in Lecture 3, namely that the localization $k[x_1, \dots, x_n, p^{-1}]$ along a nonzero polynomial $P \in k[x_1, \dots, x_n]$ is holonomic over the Weyl algebra $A_n(k)$.

By analogy with quasi-coherent and coherent $\mathscr{D}$-modules, we use the notation $D^b_h(\mathscr{D}_X)$ for the full subcategory of $D^b_{coh}(\mathscr{D}_X)$, whose objects are those complexes of $\mathscr{D}_X$-modules whose cohomology sheaves are holonomic. This category contains all bounded complexes of holonomic $\mathscr{D}_X$-modules, of course, but also injective or flat resolutions of such complexes; we need to work in this larger category in order to define f_+ or $\mathbf{L}f^*$. Fortunately, Beilinson has shown that the inclusion functor

$$D^b\big(\mathrm{Mod}_h(\mathscr{D}_X)\big) \to D^b_h(\mathscr{D}_X)$$

is an equivalence of categories. This means concretely that every complex of $\mathscr{D}_X$-modules with holonomic cohomology sheaves is isomorphic, in $D^b_h(\mathscr{D}_X)$, to a bounded complex of holonomic $\mathscr{D}_X$-modules.

Theorem 18.5. *Let $f \colon X \to Y$ be a morphism of nonsingular algebraic varieties.*
(a) The functor f_+ takes $D^b_h(\mathscr{D}_X^{op})$ into $D^b_h(\mathscr{D}_Y^{op})$.
(b) The functor $\mathbf{L}f^*$ takes $D^b_h(\mathscr{D}_Y)$ into $D^b_h(\mathscr{D}_X)$.

Let me remind you about the case of closed embeddings.

Lemma 18.6. *Let $i \colon X \hookrightarrow Y$ be a closed embedding, and $\mathcal{M}^\bullet \in D^b_{coh}(\mathscr{D}_X^{op})$. Then one has $\mathcal{M}^\bullet \in D^b_h(\mathscr{D}_X^{op})$ if and only if $i_+\mathcal{M}^\bullet \in D^b_h(\mathscr{D}_Y^{op})$.*

Proof. The naive direct image functor $i_+\mathcal{M} = i_*\big(\mathcal{M} \otimes_{\mathscr{D}_X} \mathscr{D}_{X\to Y}\big)$ is exact, and so

$$\mathcal{H}^k(i_+\mathcal{M}^\bullet) \cong i_+(\mathcal{H}^k\mathcal{M}^\bullet).$$

This reduces the problem to the case of a single coherent right $\mathscr{D}_X$-module $\mathcal{M}$. We showed back in Lecture 13 that $i_+\mathcal{M}$ is a coherent right $\mathscr{D}_Y$-module, and that

$$\dim \mathrm{Ch}(i_+\mathcal{M}) = \dim \mathrm{Ch}(\mathcal{M}) + \dim Y - \dim X.$$

It follows that $\mathcal{M}$ is holonomic if and only if $i_+\mathcal{M}$ is holonomic. ☐

The proof of Theorem 18.5 is done in two stages. First, there are a certain number of (formal) steps that reduce the general problem to the case of modules over the Weyl algebra. Second, one uses the Bernstein filtration to do the required work for modules over the Weyl algebra. Let me go over the reduction steps rather quickly, without paying too much attention to the details.

The crucial observation is that (a) follows from the special case of a coordinate projection $\mathbb{A}^{n+1}_k \to \mathbb{A}^n_k$. Let me explain how this works. First, we observe that it is enough to consider a single holonomic $\mathscr{D}_X$-module $\mathcal{M}$. The reason is that, as with any complex, one has a convergent spectral sequence

$$E_2^{p,q} = \mathcal{H}^p f_+(\mathcal{H}^q\mathcal{M}^\bullet) \Longrightarrow \mathcal{H}^{p+q} f_+\mathcal{M}^\bullet,$$

and as long as each $\mathcal{H}^p f_+(\mathcal{H}^q\mathcal{M}^\bullet)$ is holonomic, it follows that all cohomology sheaves of $f_+\mathcal{M}^\bullet$ are holonomic. Second, we can factor any morphism as

$$\begin{tikzcd} X \arrow[r, "i_f"] \arrow[rr, bend right, "f"'] & X \times Y \arrow[r, "p_2"] & Y \end{tikzcd}$$

into a closed embedding followed by a projection. Since we already know that $(i_f)_+\mathcal{M}$ is again holonomic, we only need to consider the case where $X = Z \times Y$ and $f\colon Z \times Y \to Y$ is the second projection.

Third, we can further reduce the problem to the case where $X = Z \times Y$ and Y are both affine. Since the statement is local on Y, we can obviously assume that Y is affine. Choose an affine open covering $Z = Z_1 \cup \cdots \cup Z_n$, such that each $Z \setminus Z_j$ is a nonsingular divisor in Z. Set $U_j = Z_j \times Y$, and for each subset $\alpha \subseteq \{1, \dots, n\}$, denote the resulting open embedding by

$$j_\alpha\colon U_\alpha = \bigcup_{j \in \alpha} U_j \hookrightarrow X.$$

For any sheaf of $\mathscr{O}_X$-modules, and in particular for our holonomic right $\mathscr{D}_X$-module $\mathcal{M}$, we have the Cech resolution

$$0 \to \mathcal{C}^0(M) \to \mathcal{C}^1(M) \to \cdots,$$

whose terms are given by

$$\mathcal{C}^k(M) = \bigoplus_{|\alpha|=k} (j_\alpha)_*\big(\mathcal{M}\big|_{U_\alpha}\big).$$

Since j_α is an affine morphism, we have

$$(j_\alpha)_*\big(\mathcal{M}\big|_{U_\alpha}\big) = \mathbf{R}(j_\alpha)_*\big(\mathcal{M}\big|_{U_\alpha}\big) \cong (j_\alpha)_+\big(\mathcal{M}\big|_{U_\alpha}\big),$$

and so the Cech complex is actually a resolution of $\mathcal{M}$ by right $\mathscr{D}_X$-modules. It is therefore enough to show that each

$$f_+(j_\alpha)_+\big(\mathcal{M}\big|_{U_\alpha}\big) \cong (f \circ j_\alpha)_+\big(\mathcal{M}\big|_{U_\alpha}\big)$$

is a complex of $\mathscr{D}_Y$-modules with holonomic cohomology sheaves. Since the restriction of $\mathcal{M}$ to the affine open subset U_α is holonomic, this reduces the problem to the case of a morphism between nonsingular affine varieties.

Fourth, the result for coordinate projections on affine space implies the result for all morphisms $f\colon X \to Y$ between nonsingular affine varieties. To see this, let us choose closed embeddings $i_X\colon X \hookrightarrow \mathbb{A}^m$ and $i_Y\colon Y \hookrightarrow \mathbb{A}^n$. We then have a commutative diagram

$$\begin{tikzcd}
X \arrow[r, "f"] \arrow[d, hook, "i_f"] & Y \arrow[dd, hook, "i_Y"] \\
X \times Y \arrow[d, hook, "i_X \times i_Y"] & \\
\mathbb{A}^m \times \mathbb{A}^n \arrow[r, "p_2"] & \mathbb{A}^n
\end{tikzcd}$$

where all vertical morphisms are closed embeddings. The lemma says that $f_+\mathcal{M}$ belongs to $D_h^b(\mathscr{D}_Y^{op})$ if and only if $(i_Y \circ f)_+\mathcal{M}$ belongs to $D_h^b(\mathscr{D}_{\mathbb{A}^n}^{op})$. Since we already know that the closed embeddings i_f and $i_X \times i_Y$ preserve holonomicity, we only have to consider what happens for $p_2\colon \mathbb{A}^m \times \mathbb{A}^n \to \mathbb{A}^m$. This can be factored as a composition of m coordinate projections, and so we have successfully reduced the proof of (a) to the special case of a coordinate projection $\mathbb{A}^{n+1} \to \mathbb{A}^n$.

The second observation is that the statement for the inverse image functor in (b) is a formal consequence of (a). As before, we only have to consider a single holonomic left $\mathscr{D}_Y$-module $\mathcal{M}$, and since we know that pulling back along a smooth morphism preserves holonomicity, the general problem reduces to the case of closed embeddings. Locally, we can factor any closed embedding as a composition of closed embeddings of codimension one, and so we only have to prove that if $\mathcal{M}$ is a holonomic left $\mathscr{D}_Y$-module, and $i\colon X \hookrightarrow Y$ a closed embedding of codimension one, then $\mathbf{L}i^*\mathcal{M} \in D_h^b(\mathscr{D}_X)$. Let $j\colon U \hookrightarrow Y$ be the open embedding of the complement $U = Y \setminus X$. Ignoring the difference between left and right $\mathscr{D}$-modules,

$$j_*(\mathcal{M}\big|_U) \cong j_+(\mathcal{M}\big|_U)$$

is again a $\mathscr{D}_Y$-module, due to the fact that j is affine. Provided that we know (a) for the open embedding $j\colon U \hookrightarrow Y$, it follows that $j_*(\mathcal{M}\big|_U)$ is affine. We will show next time that we have an exact sequence of $\mathscr{D}_Y$-modules

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

where I am again ignoring the difference between left and right $\mathscr{D}$-modules. It follows that each $i_+(L^{-j}i^*\mathcal{M})$ is a holonomic $\mathscr{D}_Y$-module, and by the case of closed embeddings, this implies that $L^{-j}i^*\mathcal{M}$ is a holonomic $\mathscr{D}_X$-module. This is what we wanted to show.

Excercises.

Exercise 18.1. Morihiko Saito observed that every right $\mathscr{D}_X$-module $\mathcal{M}$ has a *canonical* resolution by induced $\mathscr{D}_X$-modules. Recall that the Spencer complex $\mathrm{Sp}(\mathscr{D}_X)$ is a resolution of $\mathscr{O}_X$ by locally free left $\mathscr{D}_X$-modules.

(a) Show that each term of the complex

$$\mathrm{Sp}(\mathcal{M}) \otimes_{\mathscr{O}_X} \mathscr{D}_X$$

has the structure of a right $\mathscr{D}_X$-module. (Hint: See Lecture 12.)

(b) Construct an isomorphism of right $\mathscr{D}_X$-modules

$$\mathcal{M} \otimes_{\mathscr{O}_X} \Big(\mathscr{D}_X \otimes_{\mathscr{O}_X} \bigwedge^k \mathscr{T}_X \Big) \cong \Big(\mathcal{M} \otimes_{\mathscr{O}_X} \bigwedge^k \mathscr{T}_X \Big) \otimes_{\mathscr{O}_X} \mathscr{D}_X$$

to show that each term in above complex is an induced $\mathscr{D}_X$-module.

(c) Show that the above complex is a resolution of $\mathcal{M}$ by induced $\mathscr{D}_X$-modules.

Exercise 18.2. Let $\mathscr{F}$ and $\mathscr{G}$ be two $\mathscr{O}_X$-modules. We have a morphism

$$\mathrm{Hom}_{\mathscr{D}_X}\big(\mathscr{F} \otimes_{\mathscr{O}_X} \mathscr{D}_X, \mathscr{G} \otimes_{\mathscr{O}_X} \mathscr{D}_X\big) \to \mathrm{Hom}_{\mathscr{O}_X}\big(\mathscr{F}, \mathscr{G} \otimes_{\mathscr{O}_X} \mathscr{D}_X\big) \to \mathrm{Hom}_k(\mathscr{F}, \mathscr{G}),$$

obtained by composing with $\mathscr{G} \otimes_{\mathscr{O}_X} \mathscr{D}_X \to \mathscr{G}$, $u \otimes P \mapsto u \cdot P(1)$. Show that this morphism is injective. The image is called the space of *differential morphisms* from $\mathscr{F}$ to $\mathscr{G}$.

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

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

$$M_0 = \bigl\{\, m \in M \bigm| my_0 = 0 \,\bigr\}$$

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

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

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*

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

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

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

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 $f: X \to \operatorname{Spec} k$ 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 $i: Z \hookrightarrow X$. 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

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

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.

Lecture 20: April 22

Fuchsian differential equations. Our next topic is regularity. Let me try to motivate the definition by talking about another classical topic, namely Fuchsian differential equations. We work over the complex numbers, and take X to be a small open disk containing the origin in $\mathbb{C}$. Consider a differential equation of the form $Pu = 0$, where

$$P = a_0(x)\partial^m + a_1(x)\partial^{m-1} + \cdots + a_m(x)$$

is a differential operator of order m with holomorphic coefficients $a_j(x)$. If $a_0(0) \neq 0$, then the equation has m linearly independent holomorphic solutions, determined by the initial conditions $u(0), u'(0), \dots, u^{(m-1)}(0)$. Another way to say this is that the $\mathscr{D}_X$-module $\mathscr{D}_X/\mathscr{D}_X P$ is isomorphic to $\mathscr{O}_X^{\oplus m}$, where the isomorphism takes a vector $(u_0, \dots, u_{m-1})$ to the image of $u_0 + u_1\partial + \cdots + u_{m-1}\partial^{m-1}$. Here $\mathscr{D}_X$ is the sheaf of linear differential operators with holomorphic coefficients.

If $a_0(0) = 0$, then the story becomes more complicated.

Example 20.1. Suppose that $P = x\partial - \alpha$ for some $\alpha \in \mathbb{C}$. Here the solution $u = x^\alpha = e^{\alpha \log x}$ is really only defined on sectors, because of the term $\log x$.

Example 20.2. Suppose that $P = x^2\partial - 1$. Here the solution $u = e^{-1/x}$ is single-valued, but has an essential singularity at the origin. This is bad.

We need some terminology to talk about the solutions to the equation $Pu = 0$. Let us denote by R the ring of holomorphic functions on X, and by K its fraction field; elements of K are meromorphic functions. Further, we use $\tilde{R}$ to denote the ring of multi-valued holomorphic functions on $X \setminus \{0\}$; by this we mean holomorphic functions on the universal covering space. Using the exponential function

$$\mathbb{C} \to \mathbb{C}^*, \quad t \mapsto e^{2\pi i t},$$

the universal covering space of a disk of radius r minus the origin is the half-plane $\operatorname{Im} t > \frac{1}{2\pi}\log(1/r)$. This means that $\tilde{R}$ is the ring of holomorphic functions on a suitable half-plane. For example, $\log x = 2\pi i t$ and $x^\alpha = e^{2\pi i \alpha t}$ belong to $\tilde{R}$.

We want to avoid essential singularities; this can be done by controlling the rate of growth of solutions near the origin. We say that a multi-valued holomorphic function $f \in \tilde{R}$ has *moderate growth* near the origin if on any sector

$$S = \{\, x \in \mathbb{C} \mid 0 < |x| < \varepsilon \text{ and } \theta_0 \leq \arg x \leq \theta_1 \,\},$$

there is an integer $k \geq 0$ and a constant $C \geq 0$ such that

$$|f(x)| \leq \frac{C}{|x|^k}$$

for every $x \in S$. Let $\tilde{R}^{mod} \subseteq \tilde{R}$ be the subring of multi-valued functions with moderate growth near the origin. The functions x^α and $(\log x)^\ell$ belong to $\tilde{R}^{mod}$ for every $\alpha \in \mathbb{C}$ and $\ell \in \mathbb{N}$.

Example 20.3. Suppose that f is a single-valued holomorphic function on the punctured disk $X \setminus \{0\}$. Then f has moderate growth near the origin iff f is meromorphic; the reason is that $x^k f$ extends to a holomorphic function on X by Riemann's extension theorem. Thus moderate growth prevents essential singularities.

Let me now remind you of the classical theorem by Fuchs. After shrinking X, if necessary, we can assume that the origin is the only zero of $a_0(x)$; we can then divide through by $a_0(x)$ to get a differential operator with meromorphic coefficients.

Theorem 20.4 (Fuchs)**.** *Let $P = \partial^m + a_1(x)\partial^{m-1} + \cdots + a_m(x)$ be a differential operator of order m with $a_j(x) \in K$. The following two conditions are equivalent:*

(a) *All multi-valued solutions $u \in \tilde{R}$ of the differential equation $Pu = 0$ have moderate growth near the origin.*
(b) *For $j = 1, \ldots, n$, the function $a_j(x)$ has a pole of order at most j at the origin.*

If the conditions in the theorem are satisfied, the differential equation $Pu = 0$ is said to be *regular at the origin.* There is another way to formulate the algebraic condition in (b). Using identity $x^j \partial^j = (x\partial)(x\partial - 1) \cdots (x\partial - j + 1)$, we get

$$x^m P = (x\partial)^m + b_1(x)(x\partial)^{m-1} + \cdots + b_m(x),$$

and (b) becomes the condition that $b_1(x), \cdots, b_m(x)$ are holomorphic functions.

Systems of differential equations and regularity. We will prove Theorem 20.4 by turning the problem into a system of first-order differential equations. If we set $u_1 = u, u_2 = \partial u, \ldots, u_m = \partial^{m-1} u$, then $Pu = 0$ is of course equivalent to the system of m first-order equations

$$\begin{aligned} \partial u_1 &= u_2 \\ \partial u_2 &= u_3 \\ &\vdots \\ \partial u_{m-1} &= u_m \\ \partial u_m &= -(a_m u_1 + \cdots + a_1 u_m) \end{aligned}$$

More generally, let us consider a first-order system of the form

$$\partial u_i = \sum_{j=1}^{m} a_{i,j} u_j, \quad i = 1, \ldots, m,$$

with m unknown functions $u_1, \ldots, u_m$ and meromorphic coefficients $a_{i,j} \in K$. We can also write this in the form $\partial U = AU$, where U is the column vector with entries $u_1, \ldots, u_m$, and A is an $m \times m$-matrix whose entries are meromorphic functions.

Example 20.5. If condition (b) is satisfied, we can instead look at the m functions $v_1 = u, v_2 = x\partial u, \ldots, v_m = (x\partial)^{m-1} u$; the equation $Pu = 0$ is then also equivalent to the following system:

$$\begin{aligned} x\partial v_1 &= v_2 \\ x\partial v_2 &= v_3 \\ &\vdots \\ x\partial v_{m-1} &= v_m \\ x\partial v_m &= -(b_m v_1 + \cdots + b_1 v_m) \end{aligned}$$

In matrix notation, this becomes $x\partial V = BV$, where the entries of the $m \times m$-matrix B are now holomorphic functions.

Now let us describe the multi-valued solutions of such a system $\partial U = AU$. We can pull the system back to the universal covering space of $X \setminus \{0\}$, which amounts to setting $x = e^{2\pi i t}$. This gives us a system of first-order equations with holomorphic coefficients on a half-space; by Cauchy's theorem, it has m linearly independent holomorphic solutions $\tilde{u}^1, \ldots, \tilde{u}^m$; here each $\tilde{u}^j$ is a column vector with entries in $\tilde{R}$. Let us denote by $\tilde{S}(t)$ the $m \times m$-matrix whose columns are $\tilde{u}^1, \ldots, \tilde{u}^m$. Since the coefficients of the system are invariant under the substitution $t \mapsto t + 1$, the columns of $\tilde{S}(t + 1)$ form another basis for the vector space of solutions, and so

$$\tilde{S}(t + 1) = \tilde{S}(t) C$$

for a certain matrix $C \in \mathrm{GL}_n(\mathbb{C})$. This matrix is called the *monodromy matrix* of the system, because it describes how the multi-valued solutions to the system transform when going around the origin.

Choose a matrix Γ with the property that $C = e^{2\pi i \Gamma}$; such a matrix always exists, and is unique if we require that the eigenvalues of Γ have their real part in a fixed interval of unit lengt, such as $[0, 1)$. The matrix $\tilde{S}(t)e^{-2\pi i t \Gamma}$ is now invariant under the substitution $t \mapsto t + 1$, and so

$$\tilde{S}(t)e^{-2\pi i t \Gamma} = \Sigma(e^{2\pi i t}),$$

where $\Sigma(x)$ is an $m \times m$-matrix whose entries are holomorphic functions on $X \setminus \{0\}$. Replacing $2\pi i t$ by $\log x$, we see that the columns of the matrix

$$S(x) = \Sigma(x)e^{\log x \, \Gamma}$$

form a basis for the space of multi-valued solutions to the system $\partial U = AU$.

Changing the basis in the vector space of solutions amounts to conjugating C and Γ by the change-of-basis matrix. Since we are working over $\mathbb{C}$, we can therefore choose our basis in such a way that Γ is in Jordan canonical form. Thus Γ is block-diagonal, with blocks of the type

$$\begin{pmatrix} \alpha & 1 & & & \\ & \alpha & 1 & & \\ & & \ddots & \ddots & \\ & & & \alpha & 1 \\ & & & & \alpha \end{pmatrix}$$

which means that $e^{\log x \, \Gamma}$ is block-diagonal, with blocks of the type

$$x^\alpha \cdot \begin{pmatrix} 1 & L_1(x) & L_2(x) & \cdots & L_{m-1}(x) \\ & 1 & L_1(x) & \cdots & L_{m-2}(x) \\ & & \ddots & \ddots & \vdots \\ & & & 1 & L_1(x) \\ & & & & 1 \end{pmatrix}$$

where now $L_j(x) = \frac{1}{j!}(\log x)^j$. This gives a fairly concrete description of what multi-valued solutions look like.

Example 20.6. A corollary of the discussion so far is that any m-th order differential equation of the form $Pu = 0$ has a solution of the form $x^\alpha h(x)$, where $h(x)$ is holomorphic outside the origin, and $\alpha \in \mathbb{C}$ has the property that $e^{2\pi i \alpha}$ is an eigenvalue of the monodromy matrix C.

Now our goal is to prove a version of Theorem 20.4 for systems.

Definition 20.7. We say that two systems $\partial U = AU$ and $\partial V = BV$ are *equivalent* if there is a matrix $M(x) \in \mathrm{GL}_m(K)$ with meromorphic entries such that

$$B = \partial M \cdot M^{-1} + MAM^{-1}.$$

This means that U solves the first system iff $V = MU$ solves the second one.

Here is the analogue of Theorem 20.4 for systems.

Theorem 20.8. *Let A be an $m \times m$-matrix with entries in K. The following three conditions are equivalent:*

(a) All multi-valued solutions of $\partial U = AU$ have moderate growth near the origin, meaning that the individual components of U do.

(b) The system $\partial U = AU$ is equivalent to a system of the form $\partial V = x^{-1}\Gamma V$, where Γ is an $m \times m$-matrix with constant entries.

(c) The system $\partial U = AU$ is equivalent to a system of the form $\partial V = x^{-1}BV$, where B is an $m \times m$-matrix with holomorphic entries.

A system satisfying these equivalent conditions is called *regular* at the origin.

Proof. Let us show that (a) implies (b). We already know that a fundamental system of solutions is of the form $S(x) = \Sigma(x)e^{\log x\,\Gamma}$. By assumption, the entries of the matrix $S(x)$ have moderate growth near the origin. Since powers of $\log x$ have moderate growth, it follows that the entries of

$$\Sigma(x) = S(x)e^{-\log x\,\Gamma}$$

also have moderate growth near the origin. The entries of $\Sigma(x)$ are therefore meromorphic functions, and so $\Sigma(x) \in \mathrm{GL}_m(K)$. After replacing U by $V = \Sigma^{-1}(x)U$, we obtain the equivalent system

$$\partial V = \frac{1}{x}\Gamma V,$$

which is what we wanted to show.

It is clear that (b) implies (c), and so it remains to prove that (c) implies (a). Let V be any multi-valued solution of the system $x\partial V = BV$. Here V is a column vector with entries $v_1, \dots, v_m \in \tilde{R}$. To prove that $v_1, \dots, v_m \in \tilde{R}^{mod}$, we need to understand their asymptotic behavior on any sector

$$S = \left\{\, x \in \mathbb{C} \;\middle|\; 0 < |x| < \varepsilon \text{ and } \theta_0 \leq \arg(x) \leq \theta_1 \,\right\}.$$

Let us set $\|V\|^2 = |v_1|^2 + \cdots + |v_m|^2$ and $x = re^{i\theta}$. Since the entries of B are holomorphic, they are bounded on S. A short calculation using $\partial V = x^{-1}BV$ gives

$$\frac{\partial}{\partial r}\|V\| \leq \frac{1}{2\|V\|}\sum_{j=1}^{m} 2|v_j|\left|\frac{\partial v_j}{\partial r}\right| \leq \sqrt{\sum_{j=1}^{m}\left|\frac{\partial v_j}{\partial x}\right|^2} \leq \frac{C}{r}\|V\|,$$

where $C \geq 0$ is an upper bound for the matrix norm of B on the sector S. After integrating over r, this becomes

$$\|V(re^{i\theta})\| \leq \|V(r_0e^{i\theta})\| + \int_r^{r_0} \frac{C}{s}\|V(se^{i\theta})\|\,ds,$$

for any $0 < r \leq r_0 < \varepsilon$. Now we apply Grönwall's inequality to conclude that

$$\|V(re^{i\theta})\| \leq \|V(r_0e^{i\theta})\| \exp \int_r^{r_0} \frac{C}{s}\,ds = \|V(r_0e^{i\theta})\| \left(\frac{r_0}{r}\right)^C.$$

This means exactly that all entries of V have moderate growth at the origin. $\square$

Note. Grönwall's inequality says that an integral inequality of the form

$$f(t) \leq C + \int_{t_0}^{t} g(s)f(s)\,ds$$

for a real function $f(t)$ implies that

$$f(t) \leq C \exp \int_{t_0}^{t} g(s)\,ds.$$

We are now in a position to prove the theorem of Fuchs from the beginning.

Proof of Theorem 20.4. Consider a differential operator

$$P = \partial^m + a_1(x)\partial^{m-1} + \cdots + a_m(x),$$

with $a_j \in K$. Suppose that each a_j has a pole of order at most j at the origin. As we remarked before, we can rewrite $x^mP = (x\partial)^m + b_1(x)(x\partial)^{m-1} + \cdots + b_m(x)$, with $b_j \in R$ holomorphic. Setting $v_1 = u, v_2 = x\partial u, \dots, v_m = (x\partial)^{m-1}u$, it follows that the column vector $V = (v_1, \dots, v_m)$ solves a system of the form $\partial V = x^{-1}BV$. By Theorem 20.8, the multi-valued functions $v_1, \dots, v_n$ have moderate growth near the origin, and so in particular $u \in \tilde{R}^{mod}$.

Let us prove the converse. Suppose that all multi-valued solutions of $Pu = 0$ have moderate growth near the origin. If we write the corresponding system in the form $\partial U = AU$, then we have

$$t^m + a_1 t^{m-1} + \cdots + a_m = \det(t\,\mathrm{id} - A),$$

and so we can recover the coefficients of P from the characteristic polynomial of the matrix A. It is not hard to see that all solutions of $\partial U = AU$ also have moderate growth near the origin. By Theorem 20.8, our system is equivalent to a system of the form $\partial V = x^{-1}\Gamma V$, where Γ is an $m \times m$-matrix with constant entries. Consequently, there exists a matrix $M \in \mathrm{GL}_m(K)$ such that

$$A = \partial M \cdot M^{-1} + \frac{1}{x} M \Gamma M^{-1},$$

After clearing denominators, we get $M = x^\ell N$, with $N \in \mathrm{GL}_m(R)$. Then

$$A = \frac{1}{x}\Big(N\Gamma N^{-1} + \ell\,\mathrm{id}\Big) + \partial N \cdot N^{-1},$$

and if we now compute the characteristic polynomial, we find that the j-th coefficient a_j has a pole of order at most j at $x = 0$ (being equal to a sum of $j \times j$-minors of the matrix on the right-hand side). ☐

The theorem we have just proved has another interesting consequence.

Corollary 20.9. *Two regular systems are equivalent if and only if their monodromy matrices are conjugate.*

Proof. The proof of Theorem 20.8 shows that any regular system is equivalent to a system of the form

$$\partial U = \frac{1}{x}\Gamma U,$$

where Γ is an $m \times m$-matrix with constant entries, such that the monodromy matrix of the system is $e^{2\pi i \Gamma}$. If two such systems have conjugate monodromy matrices, then they are easily seen to be equivalent (via a constant matrix M.) To prove the converse, it is of course enough to consider systems of this special type. Suppose that two such systems with matrices Γ and Γ' are equivalent. This means that there exists a matrix $M \in \mathrm{GL}_m(K)$ such that

$$\frac{1}{x}\Gamma' = \partial M \cdot M^{-1} + \frac{1}{x} M \Gamma M^{-1}.$$

Write $M = x^\ell N$, with $N \in \mathrm{GL}_m(R)$. After clearing denominators, we get

$$\Gamma' = x\partial N \cdot N^{-1} + N(\Gamma + \ell\,\mathrm{id})N^{-1},$$

and since Γ and Γ' are constant, we can now set $x = 0$ to obtain

$$\Gamma' = N(\Gamma + \ell\,\mathrm{id})N^{-1}.$$

Since $e^{2\pi i \ell} = 1$, this implies that $e^{2\pi i \Gamma'} = N e^{2\pi i \Gamma} N^{-1}$. ☐

Exercises.

Exercise 20.1. Show directly that if two systems $\partial U = AU$ and $\partial V = BV$ are equivalent, then their monodromy matrices are conjugate to each other.

Lecture 21: April 24

Regularity for holonomic $\mathscr{D}$-modules. Last time, we considered differential equations of the form $Pu = 0$, where $P = a_0(x)\partial^m + a_1(x)\partial^{m-1} + \cdots + a_m(x)$ is a differential operator of order m with holomorphic coefficients, such that $a_0(0) = 0$. We showed that all multi-valued solutions have moderate growth near the origin iff

$$x^m P = (x\partial)^m + b_1(x)(x\partial)^{m-1} + \cdots + b_m(x), \tag{21.1}$$

with $b_1(x), \dots, b_m(x)$ holomorphic. In that case, one says that the equation $Pu = 0$ has a regular singularity at the origin. Let us now reformulate this algebraic condition in terms of the left $\mathscr{D}_X$-module $\mathcal{M} = \mathscr{D}_X/\mathscr{D}_X P$. For the time being, $\mathscr{D}_X$ again means the sheaf of linear differential operators with holomorphic coefficients.

We first observe that the characteristic variety of $\mathcal{M}$ is defined by the principal symbol $\sigma_m(P) = a_0(x)\xi^m$, where x and ξ are the natural coordinates on the cotangent bundle. Since $a_0(0) = 0$, it follows that $\mathrm{Ch}(\mathcal{M})$ is the subset defined by the equation $x\xi = 0$. This means that if $F_\bullet\mathcal{M}$ is any good filtration of $\mathcal{M}$, for example the one induced by the order filtration on $\mathscr{D}_X$, then some power of $x\xi$ annihilates $\mathrm{gr}^F\mathcal{M}$. Let me now show you how (21.1) can be used to construct a particular good filtration with better properties.

Suppose that we have (21.1) with $b_1(x), \dots, b_m(x)$ holomorphic. Then we can define a good filtration $F_\bullet\mathcal{M}$ by setting

$$F_k\mathcal{M} = \sum_{j=0}^{m-1} F_k\mathscr{D}_X \cdot (x\partial)^j + \mathscr{D}_X P.$$

It is not hard to see that this is indeed a good filtration; moreover,

$$x\partial \cdot F_k\mathcal{M} \subseteq F_k\mathcal{M}$$

for every $k \in \mathbb{N}$, by virtue of (21.1). This means that $\mathrm{gr}^F\mathcal{M}$ is annihilated by the first power of $x\xi$.

Kashiwara and Kawai introduced the notion of holonomic $\mathscr{D}$-modules with regular singularities as a generalization of this case. From now on, we let X be a nonsingular algebraic variety (over a field k of characteristic zero). For a coherent left $\mathscr{D}_X$-module $\mathcal{M}$, we denote by $\mathcal{I}_{\mathrm{Ch}(\mathcal{M})} \subseteq \mathscr{O}_{T^*X}$ the ideal sheaf of the characteristic variety. Recall that

$$\mathcal{I}_{\mathrm{Ch}(\mathcal{M})} = \sqrt{\mathrm{Ann}_{\mathrm{gr}^F\mathscr{D}_X}\,\mathrm{gr}^F\mathcal{M}},$$

where $F_\bullet\mathcal{M}$ is any good filtration. It follows that there is some (usually large) integer N such that $\mathcal{I}^N_{\mathrm{Ch}(\mathcal{M})} \cdot \mathrm{gr}^F\mathcal{M} = 0$. Roughly speaking, we say that $\mathcal{M}$ is regular if we can find a good filtration for which $N = 1$. For technical reasons, we have to be slightly more careful. Suppose first that X is proper over $\operatorname{Spec} k$.

Definition 21.2. Let X be a nonsingular algebraic variety that is proper over $\operatorname{Spec} k$. A holonomic left $\mathscr{D}_X$-module $\mathcal{M}$ is called *regular* (in the sense of Kashiwara and Kawai) if it admits a good filtration $F_\bullet\mathcal{M}$ such that $\mathcal{I}_{\mathrm{Ch}(\mathcal{M})} \cdot \mathrm{gr}^F\mathcal{M} = 0$.

If $P \in F_k\mathscr{D}_X$ is a differential operator of order k, then $\sigma_k(P)$ belongs to $\mathcal{I}_{\mathrm{Ch}(\mathcal{M})}$ if and only if $\sigma_k(P)$ vanishes along the characteristic variety of $\mathcal{M}$. The condition in the definition is therefore saying that whenever P is a differential operator of order k such that $\sigma_k(P)$ vanishes along $\mathrm{Ch}(\mathcal{M})$, then

$$P \cdot F_j\mathcal{M} \subseteq F_{j+k-1}\mathcal{M}$$

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

The original definition by Kashiwara and Kawai is only asking that a good filtration with $\mathcal{I}_{\mathrm{Ch}(\mathcal{M})} \cdot \operatorname{gr}^F \mathcal{M} = 0$ should exist locally on X; but they show that $\mathcal{M}$ then actually has a globally defined good filtration with this property.

One can prove (with a lot of work) that direct images by proper morphisms, and inverse images by arbitrary morphisms, preserve regularity. If we used the above definition to define regularity when X is not proper, we would run into the problem that direct images by open embeddings do not necessarily preserve regularity.

Example 21.3. Consider the holonomic A_1-module $M = A_1/A_1(\partial - 1)$. The filtration induced by the order filtration certainly has the property in the definition (and the differential equation $\partial u = u$ has a regular singularity at the origin). The problem occurs near the point at infinity. Indeed, if we consider the open embedding $\mathbb{A}^1 \hookrightarrow \mathbb{P}^1$, and look at M in the other affine chart with coordinate $y = x^{-1}$, we get $\partial_x - 1 = -y^2\partial_y - 1$. The A_1-module

$$A_1/A_1(y^2\partial_y + 1)$$

is *not* regular in the above sense; indeed, the differential equation $y^2\partial_y u + u = 0$ does not satisfy the condition in Theorem 20.4.

Example 21.4. A more well-behaved example is $M = A_1/A_1(x\partial - \alpha)$, for $\alpha \in k$. Since $x\partial x = -y\partial_y$, this becomes $A_1/A_1(y\partial_y + \alpha)$ in the chart at infinity, which again has a regular singularity.

Since we would like direct images by arbitrary morphisms to preserve regularity, we need to include open embeddings into the definition. Let X be a nonsingular algebraic variety. Since k has characteristic zero, Nagata's theorem implies that we can always embed X into a nonsingular algebraic variety $\bar{X}$ that is proper over $\operatorname{Spec} k$. We can always arrange that $\bar{X} \setminus X$ is a divisor; using embedded resolution of singularities, we can moreover achieve that this divisor only has normal crossing singularities. In either case, $j \colon X \hookrightarrow \bar{X}$ is an affine morphism, and so if $\mathcal{M}$ is a holonomic left $\mathscr{D}_X$-module, the direct image $j_+\mathcal{M} = j_*\mathcal{M}$ is again a holonomic left $\mathscr{D}_{\bar{X}}$-module.

Definition 21.5. Let X be a nonsingular algebraic variety. A holonomic left $\mathscr{D}_X$-module $\mathcal{M}$ is called *regular* (in the sense of Kashiwara and Kawai) if, for any affine open embedding $j \colon X \hookrightarrow \bar{X}$ into a nonsingular algebraic variety $\bar{X}$ that is proper over $\operatorname{Spec} k$, the direct image $j_+\mathcal{M}$ is regular on $\bar{X}$.

In fact, it suffices to check this for a single embedding $j \colon X \hookrightarrow \bar{X}$. Here is why. Given any two affine open embeddings $j \colon X \hookrightarrow \bar{X}$ and $j' \colon X \hookrightarrow \bar{X}'$, one can take the closure of the image of $(j, j') \colon X \hookrightarrow \bar{X} \times \bar{X}'$, and resolve the resulting singularities to obtain a third embedding $j'' \colon X \hookrightarrow \bar{X}''$ such that $j = f \circ j''$ and $j' = f' \circ j''$ for two proper morphisms $f \colon \bar{X}'' \to X$ and $f' \colon \bar{X}'' \to \bar{X}'$. Since direct images by proper morphisms preserve regularity, it follows that $j_+\mathcal{M}$ is regular on $\bar{X}$ if and only if $j'_+\mathcal{M}$ is regular on $\bar{X}'$.

Regularity and solutions. Over the complex numbers, one can also detect regularity by looking at solutions. The idea is that a left $\mathscr{D}_X$-module $\mathcal{M}$ is regular if and only if all formal power series solutions of $\mathcal{M}$ are convergent. Let us make this precise. We now assume that X is a complex manifold of dimension n, and we denote by $\mathscr{D}_X$ the sheaf of differential operators with holomorphic coefficients. If $\mathcal{M}$ is a holonomic left $\mathscr{D}_X$-module, we can define regularity as above by the (local) existence of a good filtration such that $\mathcal{I}_{\mathrm{Ch}(\mathcal{M})} \cdot \operatorname{gr}^F \mathcal{M} = 0$. Fix a point $x \in X$, and denote by $\mathscr{O}_{X,x}$ the local ring of holomorphic functions that are defined in some neighborhood of x, and by $\hat{\mathscr{O}}_{X,x}$ its completion with respect to the maximal ideal. Concretely, $\hat{\mathscr{O}}_{X,x}$ are formal power series in local coordinates $x_1, \dots, x_n$, and the subring $\mathscr{O}_{X,x}$ consists of those power series that actually converge in a neighorhood of the given point. The stalk $\mathcal{M}_x$ is a holonomic left $\mathscr{D}_{X,x}$-module. In particular, it is coherent, and so we can think of $\mathcal{M}_x$ as being obtained from a system of linear partial differential equations (by choosing a presentation of $\mathcal{M}_x$). As we discussed in Lecture 1, the space of holomorphic solutions to the system can be described as

$$\operatorname{Hom}_{\mathscr{D}_{X,x}}(\mathcal{M}_x, \mathscr{O}_{X,x}).$$

Roughly speaking, regularity of $\mathcal{M}$ means that the natural morphism

$$\operatorname{Hom}_{\mathscr{D}_{X,x}}(\mathcal{M}_x, \mathscr{O}_{X,x}) \hookrightarrow \operatorname{Hom}_{\mathscr{D}_{X,x}}(\mathcal{M}_x, \hat{\mathscr{O}}_{X,x})$$

is an isomorphism. In other words, every convergent power series solution actually converges. This is not quite true, but it becomes true if we replace the naive solution functor by its derived version

$$\mathbf{R}\operatorname{Hom}_{\mathscr{D}_{X,x}}(\mathcal{M}_x, \mathscr{O}_{X,x}).$$

Concretetly, this is computed by choosing a resolution of $\mathcal{M}_x$ by free $\mathscr{D}_{X,x}$-modules of finite rank, and then applying the functor $\operatorname{Hom}_{\mathscr{D}_{X,x}}(-, \mathscr{O}_{X,x})$.

Theorem 21.6 (Kashiwara-Kawai)**.** *Let X be a complex manifold, and $\mathcal{M}$ a holonomic left $\mathscr{D}_X$-module. Then $\mathcal{M}$ is regular, in the sense that it (locally) admits a good filtration $F_\bullet\mathcal{M}$ with $\mathcal{I}_{\operatorname{Ch}(\mathcal{M})} \cdot \operatorname{gr}^F \mathcal{M} = 0$, iff the morphism*

$$\mathbf{R}\operatorname{Hom}_{\mathscr{D}_{X,x}}(\mathcal{M}_x, \mathscr{O}_{X,x}) \to \mathbf{R}\operatorname{Hom}_{\mathscr{D}_{X,x}}(\mathcal{M}_x, \hat{\mathscr{O}}_{X,x})$$

is an isomorphism in the derived category, for every point $x \in X$.

We do not have the tools to prove this, so let me instead illustrate the result by a simple example.

Example 21.7. On $X = \mathbb{C}$, consider the left $\mathscr{D}$-module $\mathcal{M} = \mathscr{D}/\mathscr{D}(x^2\partial - 1)$, which is clearly not regular at the point $x = 0$. Let us see how the solution functor detects this. A free resolution of $\mathcal{M}$ is given by

$$\mathscr{D} \xrightarrow{x^2\partial - 1} \mathscr{D}$$

and so we need to compare the cohomology of the two complexes

$$\begin{array}{ccc} \mathscr{O} & \xrightarrow{x^2\partial-1} & \mathscr{O} \\ \big\downarrow & & \big\downarrow \\ \hat{\mathscr{O}} & \xrightarrow{x^2\partial-1} & \hat{\mathscr{O}} \end{array}$$

The horizontal differential takes a (convergent) power series $\sum_{n=0}^\infty a_n x^n$ to the (convergent) power series

$$(x^2\partial - 1)\sum_{n=0}^\infty a_n x^n = \sum_{n=0}^\infty \big((n-1)a_{n-1} - a_n\big)x^n$$

where $a_{-1} = 0$ (to simplify the notation). It is easy to see that the kernel of $x^2\partial - 1$ is trivial: from the relations $(n-1)a_{n-1} - a_n = 0$ for every $n \in \mathbb{N}$, one obtains $a_0 = a_1 = a_2 = \cdots = 0$.

The behavior of the cokernel is more interesting. On $\hat{\mathscr{O}}$, the operator $x^2\partial - 1$ is surjective. Indeed, if $\sum_{n=0}^\infty b_n x^n$ is any formal power series, then the equation

$$\sum_{n=0}^\infty b_n x^n = (x^2\partial - 1)\sum_{n=0}^\infty a_n x^n$$

means that $(n-1)a_{n-1} - a_n = b_n$, and this can be solved recursively. But on $\mathscr{O}$, the operator is no longer surjective. For instance, if we try to solve

$$x = (x^2\partial - 1)\sum_{n=0}^{\infty} a_n x^n,$$

we obtain $a_0 = 0$, $a_1 = -1$, and $a_n = (n-1)a_{n-1}$ for $n \geq 2$, from which it follows that $a_n = -(n-1)!$ for $n \geq 1$. The resulting series

$$-\sum_{n=1}^{\infty}(n-1)! \cdot x^n$$

clearly has radius of convergence equal to zero.

Lecture 22: April 29

Today, I would like to discuss a very useful class of examples, namely regular holonomic $\mathscr{D}$-modules of "normal crossing type". We will show that these objects have a simple combinatorial description in terms of vector spaces and certain linear maps between them. We will describe them both on affine space and on projective space. Before we can do that, we need to review a few basic results about $\mathscr{D}$-modules on projective space.

$\mathscr{D}$-affine varieties. We have already seen that algebraic $\mathscr{D}$-modules on affine space are the same thing as modules over the Weyl algebra $A_n(k)$. Somewhat surprisingly, a similar result holds on projective space. In fact, projective space turns out to be $\mathscr{D}$-affine, in the following sense.

Definition 22.1. A nonsingular algebraic variety X is called *$\mathscr{D}$-affine* if it satisfies the following two conditions:

(a) The global section functor
$$\Gamma(X, -)\colon \operatorname{Mod}_{qc}(\mathscr{D}_X) \to \operatorname{Mod}\bigl(\Gamma(X, \mathscr{D}_X)\bigr)$$
is exact.

(b) If $\Gamma(X, \mathcal{M}) = 0$ for some $\mathcal{M} \in \operatorname{Mod}_{qc}(\mathscr{D}_X)$, then $\mathcal{M} = 0$.

Here $\operatorname{Mod}_{qc}(\mathscr{D}_X)$ denotes the category of left $\mathscr{D}_X$-modules that are quasi-coherent as $\mathscr{O}_X$-modules; earlier on, we used the term "algebraic $\mathscr{D}$-modules".

Example 22.2. Any nonsingular affine variety is $\mathscr{D}$-affine; in fact, the global sections functor is exact on all quasi-coherent $\mathscr{O}_X$-modules in that case.

Suppose that $\mathcal{M}$ is a left $\mathscr{D}_X$-module. The space of global sections $\Gamma(X, \mathcal{M})$ is then naturally a left module over the ring of global differential operators $\Gamma(X, \mathscr{D}_X)$. On a $\mathscr{D}$-affine variety, this gives an equivalence of categories between algebraic $\mathscr{D}$-modules and modules over the ring $\Gamma(X, \mathscr{D}_X)$.

Theorem 22.3. *Let X be a nonsingular algebraic variety that is $\mathscr{D}$-affine.*

(1) Any $\mathcal{M} \in \operatorname{Mod}_{qc}(\mathscr{D}_X)$ is generated by its global sections.

(2) The global sections functor
$$\Gamma(X, -)\colon \operatorname{Mod}_{qc}(\mathscr{D}_X) \to \operatorname{Mod}\bigl(\Gamma(X, \mathscr{D}_X)\bigr)$$
is an equivalence of categories, with inverse $\mathscr{D}_X \otimes_{\Gamma(X, \mathscr{D}_X)} (-)$.

Proof. To simplify the notation, set $R = \Gamma(X, \mathscr{D}_X)$. For (1), we need to show that the natural morphism $\mathscr{D}_X \otimes_R \Gamma(X, \mathcal{M}) \to \mathcal{M}$ is surjective. Let $\mathcal{M}_0 \subseteq \mathcal{M}$ be the image. Since the global sections functor is exact by (a), we get a short exact sequence
$$0 \to \Gamma(X, \mathcal{M}_0) \to \Gamma(X, \mathcal{M}) \to \Gamma(X, \mathcal{M}/\mathcal{M}_0) \to 0.$$
The first two spaces are equal by construction, and so $\Gamma(X, \mathcal{M}/\mathcal{M}_0) = 0$, from which it follows by (b) that $\mathcal{M}_0 = \mathcal{M}$. This proves (1).

Now we turn to (2). The claim is that the inverse functor is given by sending a left $\Gamma(X, \mathscr{D}_X)$-module V to the left $\mathscr{D}_X$-module $\mathscr{D}_X \otimes_R V$. It suffices to show that the two natural morphisms
$$\alpha_{\mathcal{M}}\colon \mathscr{D}_X \otimes_R \Gamma(X, \mathcal{M}) \to \mathcal{M}$$
$$\beta_V\colon V \to \Gamma(X, \mathscr{D}_X \otimes_R V)$$
are isomorphisms for every $\mathcal{M} \in \operatorname{Mod}_{qc}(\mathscr{D}_X)$ and every $V \in \operatorname{Mod}(R)$. Let us first prove that β_V is an isomorphism. This is clearly the case when V is a direct sum of copies of R. When V is an arbitrary R-module, we choose a presentation
$$R^{\oplus I} \longrightarrow R^{\oplus J} \longrightarrow V \longrightarrow 0$$

where I and J are two (possibly infinite) sets. We then get the following diagram with exact rows:

$$\begin{tikzcd}
R^{\oplus I} \arrow[r] \arrow[d, "\cong"] & R^{\oplus J} \arrow[r] \arrow[d, "\cong"] & V \arrow[r] \arrow[d, "\beta_V"] & 0 \\
\Gamma(X, \mathscr{D}_X^{\oplus I}) \arrow[r] & \Gamma(X, \mathscr{D}_X^{\oplus J}) \arrow[r] & \Gamma(X, \mathscr{D}_X \otimes_R V) \arrow[r] & 0
\end{tikzcd}$$

The bottom row is exact because tensor product is right-exact, and because the global sections functor is exact by condition (a) in the definition. Now the 5-lemma implies that β_V is an isomorphism.

It remains to show that $\alpha_{\mathcal{M}}$ is an isomorphism. We already know that $\alpha_{\mathcal{M}}$ is surjective; setting $\mathcal{K} = \ker \alpha_{\mathcal{M}}$, we have a short exact sequence of $\mathscr{D}_X$-modules

$$0 \to \mathcal{K} \to \mathscr{D}_X \otimes_R \Gamma(X, \mathcal{M}) \to \mathcal{M} \to 0$$

and therefore, again by (a), a short exact sequence of R-modules

$$0 \to \Gamma(X, \mathcal{K}) \to \Gamma\big(X, \mathscr{D}_X \otimes_R \Gamma(X, \mathcal{M})\big) \xrightarrow{\beta} \Gamma(X, \mathcal{M}) \to 0.$$

Since we have already shown that $\beta = \beta_{\Gamma(X,\mathcal{M})}$ is an isomorphism, it follows that $\Gamma(X, \mathcal{K}) = 0$, and hence by (b) that $\mathcal{K} = 0$. This concludes the proof of (2). $\square$

As you would expect, coherent $\mathscr{D}_X$-modules correspond to finitely generated $\Gamma(X, \mathscr{D}_X)$-modules.

Corollary 22.4. *If X is $\mathscr{D}$-affine, then*

$$\Gamma(X, -) \colon \operatorname{Mod}_{coh}(\mathscr{D}_X) \to \operatorname{Mod}_{fg}\big(\Gamma(X, \mathscr{D}_X)\big)$$

is also an equivalence of categories.

Proof. We keep the notation $R = \Gamma(X, \mathscr{D}_X)$. If V is a finitely generated R-module, then $\mathscr{D}_X \otimes_R V$ is clearly a coherent $\mathscr{D}_X$-module. Thus we only have to show that $\Gamma(X, \mathcal{M})$ is a finitely generated R-module whenever $\mathcal{M} \in \operatorname{Mod}_{coh}(\mathscr{D}_X)$. Concretely, we have to find finitely many global sections that generate $\mathcal{M}$ as a $\mathscr{D}_X$-module.

Since $\mathcal{M}$ is coherent, the restriction of $\mathcal{M}$ to any affine open subset $U \subseteq X$ is generated as a $\mathscr{D}_U$-module by finitely many sections in $\Gamma(U, \mathcal{M})$. The isomorphism $\mathscr{D}_X \otimes_R \Gamma(X, \mathcal{M}) \cong \mathcal{M}$ in the theorem gives

$$\Gamma(U, \mathscr{D}_X) \otimes_R \Gamma(X, \mathcal{M}) \cong \Gamma(U, \mathcal{M}),$$

and so $\mathcal{M}\big|_U$ is generated as a $\mathscr{D}_U$-module by finitely many sections in $\Gamma(X, \mathcal{M})$. Now X is quasi-compact, hence covered by finitely many affine open subsets; it follows that finitely many global sections generate $\mathcal{M}$ as a $\mathscr{D}_X$-module. In other words, we have a surjective morphism

$$\mathscr{D}_X^{\oplus r} \to \mathcal{M} \to 0.$$

Because the global sections functor is exact by (a), we get a surjection

$$R^{\oplus r} = \Gamma(X, \mathscr{D}_X^{\oplus r}) \to \Gamma(X, \mathcal{M}) \to 0,$$

and so $\Gamma(X, \mathcal{M})$ is a finitely generated R-module. $\square$

We are now going to show that projective spaces are $\mathscr{D}$-affine.

Theorem 22.5. *The projective space $\mathbb{P}_k^n$ is $\mathscr{D}$-affine.*

Proof. Let me begin with a preliminary discussion about global sections on $\mathbb{P}^n$. On $\mathbb{A}^{n+1}$, we have coordinates $x_0, x_1, \dots, x_n$. Let $X \subseteq \mathbb{A}^{n+1}$ be the open complement of the origin. Then $\mathbb{P}^n$ is the quotient of X by the $\mathbb{G}_m$-action that rescales the coordinates. We denote the quotient morphism by $\pi \colon X \to \mathbb{P}_k^n$; the open embedding by $j\colon X \hookrightarrow \mathbb{A}^{n+1}$; and the closed embedding of the origin by $i\colon \operatorname{Spec} k \hookrightarrow \mathbb{A}^{n+1}$. Here are the three morphisms in diagram form:

$$\begin{tikzcd} X \arrow[r, "j"] \arrow[d, "\pi"] & \mathbb{A}^{n+1} & \operatorname{Spec} k \arrow[l, hook', "i"'] \\ \mathbb{P}^n & & \end{tikzcd}$$

The Euler vector field $\theta = x_0\partial_0 + x_1\partial_1 + \cdots + x_n\partial_n$ is tangent to the fibers of π. Now suppose that $\mathcal{M}$ is a left $\mathscr{D}_{\mathbb{P}^n}$-module. Then $\mathbb{G}_m$ acts on the space of global sections of $\pi^*\mathcal{M} = \mathscr{O}_X \otimes_{\pi^{-1}\mathscr{O}_{\mathbb{P}^n}} \pi^{-1}\mathcal{M}$, and this gives us a direct sum decomposition

$$\Gamma(X, \pi^*\mathcal{M}) = \bigoplus_{\ell\in\mathbb{Z}} \Gamma_\ell(X, \pi^*\mathcal{M});$$

here $\mathbb{G}_m$ acts on the subspace $\Gamma_\ell(X, \pi^*\mathcal{M})$ with the character $z \mapsto z^\ell$. It follows that θ operates on $\Gamma_\ell(X, \pi^*\mathcal{M})$ as multiplication by ℓ. We have

$$\Gamma(\mathbb{P}^n, \mathcal{M}) \cong \Gamma(X, \pi^*\mathcal{M})^{\mathbb{G}_m} = \Gamma_0(X, \pi^*\mathcal{M}); \tag{22.6}$$

indeed, pullbacks of global sections from $\mathbb{P}^n$ are clearly $\mathbb{G}_m$-invariant, and conversely, any $\mathbb{G}_m$-invariant section on X descends to a global section on $\mathbb{P}^n$. Also note that multiplication by x_j takes Γ_ℓ into $\Gamma_{\ell+1}$, and multiplication by ∂_j takes Γ_ℓ into $\Gamma_{\ell-1}$; the reason is that $[\theta, x_j] = x_j$ and $[\theta, \partial_j] = -\partial_j$.

Now let us start proving that $\mathbb{P}^n$ satisfies the two conditions in (a) and (b). We first show that the global sections functor is exact. Let

$$0 \to \mathcal{M}_1 \to \mathcal{M}_2 \to \mathcal{M}_3 \to 0$$

be a short exact sequence of quasi-coherent $\mathscr{D}_{\mathbb{P}^n}$-modules. Since π is smooth, the pullback functor π^* is exact, which means that

$$0 \to \pi^*\mathcal{M}_1 \to \pi^*\mathcal{M}_2 \to \pi^*\mathcal{M}_3 \to 0$$

is a short exact sequence of quasi-coherent $\mathscr{D}_X$-modules. Because $j\colon X \hookrightarrow \mathbb{A}^{n+1}$ is an open embedding, $j_+ \cong \mathbf{R}j_*$ (after the appropriate conversion between left and right $\mathscr{D}$-modules). Thus we get an exact sequence of quasi-coherent $\mathscr{D}_{\mathbb{A}^{n+1}}$-modules

$$0 \to j_*\pi^*\mathcal{M}_1 \to j_*\pi^*\mathcal{M}_2 \to j_*\pi^*\mathcal{M}_3 \to R^1 j_*\pi^*\mathcal{M}_1 \to \cdots$$

The global sections functor on the affine space $\mathbb{A}^{n+1}$ is exact, and so we finally obtain an exact sequence of A_{n+1}-modules

$$0 \to \Gamma(X, \pi^*\mathcal{M}_1) \to \Gamma(X, \pi^*\mathcal{M}_2) \to \Gamma(X, \pi^*\mathcal{M}_3) \to \Gamma\big(\mathbb{A}^{n+1}, R^1 j_*\pi^*\mathcal{M}_1\big) \to \cdots$$

Now $R^1 j_*\pi^*\mathcal{M}_1$ is a quasi-coherent $\mathscr{D}_{\mathbb{A}^{n+1}}$-module supported on the origin, and so by Kashiwara's equivalence (from Lecture 13), it must be the direct image of a quasi-coherent $\mathscr{D}_{\operatorname{Spec} k}$-module. Concretely, we have

$$\Gamma\big(\mathbb{A}^{n+1}, R^1 j_*\pi^*\mathcal{M}_1\big) \cong k[\partial_0, \partial_1, \dots, \partial_n] \otimes_k V,$$

where V is a k-vector space. The key point is now that θ acts on the right-hand side with strictly negative eigenvalues. Indeed, for any $\alpha \in \mathbb{N}^{n+1}$, we have

$$\theta \cdot \partial^\alpha \otimes v = \sum_{j=0}^n x_j\partial_j \cdot \partial^\alpha \otimes v = \sum_{j=0}^n -(\alpha_j + 1)\partial^\alpha \otimes v = -\big(|\alpha| + n + 1\big) \cdot \partial^\alpha \otimes v.$$

The conclusion is that

$$0 \to \Gamma_0(X, \pi^*\mathcal{M}_1) \to \Gamma_0(X, \pi^*\mathcal{M}_2) \to \Gamma_0(X, \pi^*\mathcal{M}_3) \to 0$$

is short exact; because of (22.6), this proves that $\Gamma(\mathbb{P}^n, -)$ is an exact functor.

All that is left is to show that $\Gamma(\mathbb{P}^n, \mathcal{M}) = 0$ implies $\mathcal{M} = 0$. Here we argue by contradiction and assume that $\mathcal{M} \neq 0$. Since $\pi \colon X \to \mathbb{P}^n$ has a section over each of the $n+1$ basic affine open subsets, we must have $\pi^*\mathcal{M} \neq 0$, and therefore

$$\Gamma(X, \pi^*\mathcal{M}) = \Gamma\big(\mathbb{A}^{n+1}, j_*\pi^*\mathcal{M}\big) \neq 0.$$

It follows that there is some $\ell \in \mathbb{Z}$ such that $\Gamma_\ell(X, \pi^*\mathcal{M}) \neq 0$. On the other hand, we have $\Gamma_0(X, \pi^*\mathcal{M}) = 0$ by (22.6). We will show that this leads to a contradiction. Suppose first that $\ell \geq 1$. Take any nonzero element $s \in \Gamma_\ell(X, \pi^*\mathcal{M})$. Then

$$\theta s = \sum_{j=0}^{n} x_j \partial_j s = \ell s \neq 0,$$

and so at least one $\partial_j s \in \Gamma_{\ell-1}(X, \pi^*\mathcal{M})$ must be nonzero. Repeating this argument, we eventually arrive at $\Gamma_0(X, \pi^*\mathcal{M}) \neq 0$, which is a contradiction. The remaining possibility is that $\ell \leq -1$. Since $s \in \Gamma(X, \pi^*\mathcal{M})$ and $\pi^*\mathcal{M}$ is quasi-coherent, we cannot have $x_j s = 0$ for every j. It follows that $\Gamma_{\ell+1}(X, \pi^*\mathcal{M}) \neq 0$, and as before, this leads to a contradiction after finitely many steps. □

This result says, in particular, that coherent $\mathscr{D}_{\mathbb{P}^n}$-modules are the same thing as finitely generated modules over the ring of differential operators $\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})$. Let us briefly discuss the structure of this ring. We have

$$\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n}) \cong \Gamma_0\big(X, \mathscr{D}_{X\to\mathbb{P}^n}\big),$$

where $\mathscr{D}_{X\to\mathbb{P}^n} = \pi^*\mathscr{D}_{\mathbb{P}^n}$ is the transfer module. Recall from Lecture 16 that, in the case of a smooth morphism, $\mathscr{D}_{X\to\mathbb{P}^n}$ is the quotient of $\mathscr{D}_X$ by the submodule generated by the relative tangent bundle. In our setting, $\mathscr{D}_{X\to\mathbb{P}^n} \cong \mathscr{D}_X/\mathscr{D}_X\theta$, and so we recover the fact, already stated in Lecture 9, that $\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})$ consists of all differential operators on $\mathbb{A}^{n+1}$ that are homogenous of degree 0, modulo multiples of the Euler vector field θ.

One can turn this into a very concrete presentation by generators and relations, as follows. For $i, j \in \{0, 1, \dots, n\}$, set $D_{i,j} = x_i\partial_j$. A short calculation gives

$$[D_{i,j}, D_{k,\ell}] = \begin{cases} D_{i,i} - D_{j,j} & \text{if } k = j \text{ and } \ell = i, \\ D_{i,\ell} & \text{if } k = j \text{ and } \ell \neq i, \\ -D_{k,j} & \text{if } k \neq j \text{ and } \ell = i, \\ 0 & \text{if } k \neq j \text{ and } \ell \neq i. \end{cases} \tag{22.7}$$

We also have $\theta = D_{0,0} + D_{1,1} + \dots + D_{n,n}$. Then $\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})$ is generated as a non-commutative k-algebra by the $D_{i,j}$, and all the relations are generated by the above commutator relations and the additional relation $D_{0,0}+D_{1,1}+\dots+D_{n,n} = 0$.

Regular holonomic $\mathscr{D}$-modules of normal crossing type. We now turn to the classification of regular holonomic $\mathscr{D}$-modules of normal crossing type. Let me first explain what I mean by "normal crossing type". On $\mathbb{A}^n$, we can intersect the various components of the normal crossing divisor $x_1 \cdots x_n = 0$ to obtain a total of 2^n nonsingular closed subvarieties. (Here we use the convention that the empty intersection equals $\mathbb{A}^n$.) Their conormal bundles give us 2^n conical Lagrangian subvarieties of the cotangent bundle $T^*\mathbb{A}^n$. In the usual coordinate system $x_1, \dots, x_n, \xi_1, \dots, \xi_n$ on the cotangent bundle, the union of all these Lagrangians is exactly the closed subset

$$Z(x_1\xi_1, \dots, x_n\xi_n);$$

indeed, on each component, we have either $x_j = 0$ or $\xi_j = 0$, for every $j = 1, \dots, n$. We say that a (necessarily holonomic) $\mathscr{D}_{\mathbb{A}^n}$-module $\mathcal{M}$ is *of normal crossing type* if its characteristic variety satisfies

$$\mathrm{Ch}(\mathcal{M}) \subseteq Z(x_1\xi_1, \dots, x_n\xi_n).$$

Example 22.8. On $\mathbb{A}^2$, the condition is that the characteristic variety has at most four irreducible components: the zero section, the conormal bundles to the two axes, and the cotangent space to the origin.

Here is a typical example, to get started.

Example 22.9. Consider the A_n-module $M = A_n/A_n(x_1\partial_1 - \alpha_1, \dots, x_n\partial_n - \alpha_n)$, where $\alpha_1, \dots, \alpha_n \in k$ are scalars. The characteristic variety is defined by the principal symbols of the n operators, hence is exactly the set $Z(x_1\xi_1, \dots, x_n\xi_n)$. In particular, M is holonomic; I will leave it as an exercise to check that M is regular in the sense of Kashiwara and Kawai.

The analogous definition on $\mathbb{P}^n$ has to include the hyperplane at infinity. In homogeneous coordinates $x_0, x_1, \dots, x_n$, we are therefore looking at the closed subset

$$Z(x_0\xi_0, x_1\xi_1, \dots, x_n\xi_n) \subseteq T^*\mathbb{P}^n;$$

note that even though the cotangent bundle is not trivial, the notation still makes sense because each $x_j\partial_j$ is a globally defined vector field on $\mathbb{P}^n$. We then say that a (necessarily holonomic) $\mathscr{D}_{\mathbb{P}^n}$-module $\mathcal{M}$ is *of normal crossing type* if

$$\mathrm{Ch}(\mathcal{M}) \subseteq Z(x_0\xi_0, x_1\xi_1, \dots, x_n\xi_n).$$

Our goal is to describe explicitly all regular holonomic $\mathscr{D}_{\mathbb{P}^n}$-modules of normal crossing type, at least when k is algebraically closed. It will help us that $\mathbb{P}^n$ is $\mathscr{D}$-affine. Our starting point is the following lemma.

Lemma 22.10. *Let $\mathcal{M}$ be a holonomic left $\mathscr{D}_{\mathbb{P}^n}$-module that is regular and of normal crossing type. Then there is a finite-dimensional k-vector space $V \subseteq \Gamma(\mathbb{P}^n, \mathcal{M})$ that generates $\Gamma(\mathbb{P}^n, \mathcal{M})$ as a $\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})$-module, and is preserved by $x_0\partial_0, \dots, x_n\partial_n$.*

Proof. Regularity means that there is a global good filtration $F_\bullet\mathcal{M}$ such that $\mathcal{I}_{\mathrm{Ch}(\mathcal{M})}$ annihilates $\mathrm{gr}^F\mathcal{M}$. Since $\mathrm{Ch}(\mathcal{M}) \subseteq Z(x_0\xi_0, x_1\xi_1, \dots, x_n\xi_n)$, this says concretely that we have

$$x_j\partial_j \cdot F_i\mathcal{M} \subseteq F_i\mathcal{M}$$

for every $j = 0, 1, \dots, n$ and $i \in \mathbb{Z}$. Since $F_i\mathcal{M}$ is a coherent $\mathscr{O}_{\mathbb{P}^n}$-module,

$$\Gamma(\mathbb{P}^n, F_i\mathcal{M}) \subseteq \Gamma(\mathbb{P}^n, \mathcal{M})$$

is a finite-dimensional k-vector space that is preserved by $x_0\partial_0, \dots, x_n\partial_n$. We showed during the proof of Corollary 22.4 that $\mathcal{M}$ is generated as a $\mathscr{D}_{\mathbb{P}^n}$-module by finitely many global sections. If we choose i large enough, these sections will be global sections of $F_i\mathcal{M}$, and so the subspace $V = \Gamma(\mathbb{P}^n, F_i\mathcal{M})$ actually generates $\Gamma(\mathbb{P}^n, \mathcal{M})$ as a module over $\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})$. □

Now $x_0\partial_0, \dots, x_n\partial_n$ are commuting endomorphisms of the finite-dimensional k-vector space V. Assuming that k is algebraically closed, we get a decomposition

$$V = \bigoplus_{\alpha \in k^{n+1}} V_\alpha$$

into generalized eigenspaces, where $V_\alpha \subseteq V$ consists of all vectors $v \in V$ such that $(x_j\partial_j - \alpha_j)^m v = 0$ for $j = 0, 1, \dots, n$ and $m \gg 0$. In other words, $x_j\partial_j - \alpha_j$ acts nilpotently on the subspace V_α. Of course, only finitely many of the V_α are actually nonzero; also note that we must have $\alpha_0 + \alpha_1 + \dots + \alpha_n = 0$, due to the fact that $\theta = x_0\partial_0 + \dots + x_n\partial_n$ acts trivially on V. If we define

$$A = \left\{ \alpha \in k^{n+1} \;\middle|\; \alpha_0 + \alpha_1 + \dots + \alpha_n = 0 \right\},$$

then the direct sum above is actually indexed by a finite subset of A. Since V generates $\Gamma(\mathbb{P}^n, \mathcal{M})$, we get a similar decomposition for the entire space of global sections.

Lemma 22.11. *Let $\mathcal{M}$ be a holonomic left $\mathscr{D}_{\mathbb{P}^n}$-module that is regular and of normal crossing type, and set $M = \Gamma(\mathbb{P}^n, \mathcal{M})$. We have a decomposition*

$$M = \bigoplus_{\alpha \in A} M_\alpha$$

into finite-dimensional k-vector spaces M_α, such that the operator $x_j \partial_j - \alpha_j$ acts nilpotently on M_α for $j = 0, 1, \dots, n$.

Proof. To be completely precise, we define, for every $\alpha \in A$, the subspace

$$M_\alpha = \left\{ s \in M \;\middle|\; (x_j \partial_j - \alpha_j)^m s = 0 \text{ for } j = 0, 1, \dots, n \text{ and } m \gg 0 \right\}.$$

Since different M_α are easily seen to be linearly independent, it suffices to prove that every $s \in M$ can be written as a sum of elements in finitely many M_α. This is true for elements of V by the discussion above; and for other elements, it follows from the fact that M is generated by V as a $\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})$-module. Indeed, $\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})$ is generated as a k-algebra by the operators $D_{i,j} = x_i \partial_j$, and since we already have the desired decomposition for elements of V, we only have to prove that

$$D_{i,j} \cdot M_\alpha \subseteq M_{\alpha + e_i - e_j},$$

where e_i is the i-th coordinate vector in k^{n+1}. But as $x_k \partial_k = D_{k,k}$, this follows quite easily from the commutator relations

$$[D_{i,j}, D_{k,k}] = \begin{cases} 0 & \text{if } k = i = j \text{ or } k \neq i, j, \\ D_{i,j} & \text{if } k = j \text{ and } k \neq i, \\ -D_{i,j} & \text{if } k = i \text{ and } k \neq j. \end{cases}$$

that we had proved earlier. $\square$

Exercises.

Exercise 22.1. Prove that $D_{i,j} \cdot M_\alpha \subseteq M_{\alpha + e_i - e_j}$.

Exercise 22.2. Verify the relations in (22.7), and prove that $\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})$ does have the claimed presentation by generators and relations.

Lecture 23: May 1

Regular holonomic $\mathscr{D}$-modules of normal crossing type. Let me briefly recall what we did last time. We first showed that $\mathbb{P}^n$ is $\mathscr{D}$-affine, which meant that the global sections functor

$$\Gamma(\mathbb{P}^n, -)\colon \operatorname{Mod}_{qc}(\mathscr{D}_{\mathbb{P}^n}) \to \operatorname{Mod}\bigl(\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})\bigr)$$

is an equivalence of categories. In other words, algebraic $\mathscr{D}$-modules on $\mathbb{P}^n$ are uniquely determined by their space of global sections, which is a module over the ring $\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})$. We also showed that the ring of differential operators on $\mathbb{P}^n$ is generated by the $(n+1)^2$ operators $D_{i,j} = x_i \partial_j$, subject to the commutator relations

$$[D_{i,j}, D_{k,\ell}] = \begin{cases} D_{i,i} - D_{j,j} & \text{if } k = j \text{ and } \ell = i, \\ D_{i,\ell} & \text{if } k = j \text{ and } \ell \neq i, \\ -D_{k,j} & \text{if } k \neq j \text{ and } \ell = i, \\ 0 & \text{if } k \neq j \text{ and } \ell \neq i, \end{cases} \tag{23.1}$$

and the extra relation $\theta = D_{0,0} + D_{1,1} + \dots + D_{n,n} = 0$. We then showed that if $\mathcal{M}$ is a regular holonomic $\mathscr{D}_{\mathbb{P}^n}$-module whose characteristic variety is contained in the set $Z(x_0\xi_0, x_1\xi_1, \dots, x_n\xi_n) \subseteq T^*\mathbb{P}^n$, then we get a decomposition

$$\Gamma(\mathbb{P}^n, \mathcal{M}) = \bigoplus_{\alpha \in A} M_\alpha,$$

where $A = \bigl\{\, \alpha \in k^{n+1} \bigm| \alpha_0 + \alpha_1 + \dots + \alpha_n = 0 \,\bigr\}$. Here each M_α is a finite-dimensional k-vector space, consisting of those global sections of $\mathcal{M}$ on which the $n+1$ operators $D_{j,j} - \alpha_j$ act nilpotently.

How about the converse? Suppose we are given a collection of finite-dimensional k-vector spaces M_α, indexed by $\alpha \in A$. What extra information is needed to turn the direct sum

$$M = \bigoplus_{\alpha \in A} M_\alpha$$

into (the space of global sections of) a regular holonomic $\mathscr{D}_{\mathbb{P}^n}$-module of normal crossing type? First, M should be a left module over the ring $\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})$, and so we need to have linear operators

$$D_{i,j}\colon M_\alpha \to M_{\alpha + e_i - e_j}$$

for every $\alpha \in A$ and every $i, j \in \{0, 1, \dots, n\}$. These operators should satisfy the commutator relations above, as well as the identity $D_{0,0} + D_{1,1} + \dots + D_{n,n} = 0$. We also want M to be finitely generated, which means that finitely many of the M_α should generate M as a $\Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})$-module. Finally, the operator $D_{j,j} - \alpha_j$ should act nilpotently on M_α for every $j \in \{0, 1, \dots, n\}$. It is then not hard to show that the corresponding $\mathscr{D}_{\mathbb{P}^n}$-module is regular holonomic of normal crossing type.

Other variants. There are some useful variants of the classification above. One is regular holonomic $\mathscr{D}$-modules of normal crossing type on affine space $\mathbb{A}^n$. Let $\mathcal{M}$ be a holonomic $\mathscr{D}_{\mathbb{A}^n}$-module with the property that

$$\operatorname{Ch}(\mathcal{M}) \subseteq Z(x_1\xi_1, \dots, x_n\xi_n) \subseteq T^*\mathbb{A}^n.$$

In that case, we say that $\mathcal{M}$ is *of normal crossing type*. Recall that $\mathcal{M}$ is regular, in the sense of Kashiwara and Kawai, if the direct image $j_+\mathcal{M}$ is regular on $\mathbb{P}^n$, where $j\colon \mathbb{A}^n \hookrightarrow \mathbb{P}^n$ is the open embedding. One can show that if $\mathcal{M}$ is regular holonomic of normal crossing type on $\mathbb{A}^n$, then $j_+\mathcal{M}$ is regular holonomic of normal crossing type on $\mathbb{P}^n$. Thus we obtain a decomposition

$$\Gamma(\mathbb{A}^n, \mathcal{M}) = \Gamma(\mathbb{P}^n, j_+\mathcal{M}) = \bigoplus_{\alpha \in k^n} M_\alpha,$$

which we are now indexing by $\alpha \in k^n$. (This is okay because $\alpha_0 = -(\alpha_1 + \dots + \alpha_n)$, so there is no loss of information.) Again, each M_α is a finite-dimensional k-vector space, consisting of all global sections of $\mathcal{M}$ on which the n commuting operators $x_j\partial_j - \alpha_j$ act nilpotently. This time, we have

$$x_j \colon M_\alpha \to M_{\alpha+e_j} \quad \text{and} \quad \partial_j \colon M_\alpha \to M_{\alpha-e_j}$$

for every $j = 1, \dots, n$; this follows from the commutator relation $[\partial_j, x_j] = 1$. Conversely, given a collection of finite-dimensional k-vector spaces M_α, indexed by $\alpha \in k^n$, and a collection of linear operators $x_j \colon M_\alpha \to M_{\alpha+e_j}$ and $\partial_j \colon M_\alpha \to M_{\alpha-e_j}$ subject to the relations $[\partial_i, x_j] = \delta_{i,j}$, the direct sum

$$M = \bigoplus_{\alpha \in k^n} M_\alpha$$

becomes a module over the Weyl algebra $\Gamma(\mathbb{A}^n, \mathscr{D}_{\mathbb{A}^n})$; if this module is finitely generated, and if each $x_j\partial_j - \alpha_j$ acts nilpotently on M_α, then the corresponding $\mathscr{D}_{\mathbb{A}^n}$-module is regular holonomic of normal crossing type.

There is also a local analytic version of the classification, for $k = \mathbb{C}$. Let $\mathscr{D}_{\mathbb{C}^n,0}$ denote the ring of linear differential operators with holomorphic coefficients that are defined in some neighborhood of the origin in $\mathbb{C}^n$. We say that a holonomic $\mathscr{D}_{\mathbb{C}^n,0}$-module $\mathcal{M}$ is of normal crossing type if its characteristic variety $\mathrm{Ch}(\mathcal{M})$ is contained in the set $Z(x_1\xi_1, \dots, x_n\xi_n)$. We say that $\mathcal{M}$ is regular if it satisfies the condition from Lecture 21, meaning if there exists a good filtration $F_\bullet\mathcal{M}$ such that each $F_k\mathcal{M}$ is a finitely generated $\mathscr{O}_{\mathbb{C}^n,0}$-module stable under the action by $x_1\partial_1, \dots, x_n\partial_n$. Define

$$M_\alpha = \left\{ s \in \mathcal{M} \;\middle|\; (x_j\partial_j - \alpha_j)^m s = 0 \text{ for } j = 0, 1, \dots, n \text{ and } m \gg 0 \right\}.$$

Each M_α is a finite-dimensional $\mathbb{C}$-vector space, and their direct sum

$$M = \bigoplus_{\alpha \in \mathbb{C}^n} M_\alpha$$

is a regular holonomic module over the Weyl algebra $A_n(\mathbb{C})$, of normal crossing type. Then one can show (with a lot of extra work) that

$$\mathcal{M} \cong \mathscr{D}_{\mathbb{C}^n,0} \otimes_{A_n(\mathbb{C})} M.$$

In other words, the $\mathscr{D}_{\mathbb{C}^n,0}$-module structure on $\mathcal{M}$ is completely determined by the much simpler algebraic $\mathscr{D}$-module M. Note that this result is only true in the local analytic setting. The following example explains why.

Example 23.2. Consider the $\mathscr{D}_{\mathbb{A}^1}$-module $\mathcal{M} = \mathscr{D}_{\mathbb{A}^1}/\mathscr{D}_{\mathbb{A}^1}(\partial - 1)$. It is easy to see that $\mathrm{Ch}(\mathcal{M})$ is the zero section, and that $\mathcal{M}$ is actually a line bundle with integrable connection. Except for regularity at infinity, $\mathcal{M}$ is therefore regular holonomic of normal crossing type. But it is not true, not even Zariski-locally, that $\Gamma(\mathbb{A}^1, \mathcal{M}) = A_1/A_1(\partial - 1)$ has a decomposition into generalized eigenspaces for $x\partial$; in fact, you can check for yourself that $x\partial$ does not have any nontrivial eigenvectors. What goes wrong is that we need a solution to $\partial u = u$ to get an isomorphism between $\mathcal{M}$ and $\mathscr{O}_{\mathbb{A}^1}$. But the solution is $u = e^x$, which is not an algebraic function, because it has an essential singularity at infinity. Another way to say this is that $\mathcal{M}$ is *not* regular at infinity.

Solutions. Let us discuss a few more properties of the classification on $\mathbb{A}^n$. For simplicity, I will assume from now on that $k = \mathbb{C}$. Consider a regular holonomic $\mathscr{D}$-module of normal crossing type, with decomposition

$$M = \bigoplus_{\alpha \in \mathbb{C}^n} M_\alpha.$$

Here each M_α is a finite-dimensional $\mathbb{C}$-vector space. By construction, $x_j\partial_j - \alpha_j$ acts nilpotently on M_α, and so $x_j\partial_j$ is an isomorphism as long as $\alpha_j \neq 0$. Consequently,

$$\partial_j \colon M_\alpha \to M_{\alpha - e_j} \quad \text{and} \quad x_j \colon M_{\alpha - e_j} \to M_\alpha$$

are injective respectively surjective for $\alpha_j \neq 0$. Likewise, $\partial_j x_j - \alpha_j - 1$ acts nilpotently on M_α, and so $\partial_j x_j$ is an isomorphism as long as $\alpha_j \neq -1$. Thus

$$\partial_j \colon M_{\alpha + e_j} \to M_\alpha \quad \text{and} \quad x_j \colon M_\alpha \to M_{\alpha + e_j}$$

are surjective respectively injective for $\alpha_j \neq -1$. We can summarize this by saying that $\partial_j \colon M_\alpha \to M_{\alpha - e_j}$ is an isomorphism for $\alpha_j \neq 0$, and that $x_j \colon M_\alpha \to M_{\alpha + e_j}$ is an isomorphism for $\alpha_j \neq -1$.

This implies of course that those vector spaces M_α with

$$-1 \leq \operatorname{Re} \alpha_j \leq 0 \quad \text{for every } j = 1, \dots, n$$

determine all the others. Since M is finitely generated over $A_n(\mathbb{C})$, the set

$$F = \left\{ \alpha \in \mathbb{C}^n \;\middle|\; M_\alpha \neq 0 \text{ and } -1 \leq \operatorname{Re} \alpha_j \leq 0 \text{ for all } j \right\}$$

must be finite. Thus M is generated as an $A_n(\mathbb{C})$-module by the direct sum of those M_α with $\alpha \in F$.

Recall that any holonomic A_n-module has finite length, meaning that it has a finite composition series whose subquotients are simple. Let us describe more explicitly what simple regular holonomic $\mathscr{D}$-modules of normal crossing type look like. Suppose that M is simple but nonzero. Choose some $\alpha \in F$, so that $M_\alpha \neq 0$ and $-1 \leq \operatorname{Re} \alpha_j \leq 0$ for all j. Since each $x_j\partial_j - \alpha_j$ acts nilpotently on M_α, we can find a common eigenvector $s \in M_\alpha$ such that $x_j\partial_j s = \alpha_j s$ for every $j = 1, \dots, n$. Since M is simple, we must have $A_n s = M$. Because s is an eigenvector, it is not hard to see that $A_n s$ intersects M_α exactly in the subspace $\mathbb{C}s$. Thus $M_\alpha = \mathbb{C}s$ is one-dimensional. Now there are two special cases:

(1) One case is that $\alpha_j = 0$. Then $x_j\partial_j s = 0$, and so the submodule $A_n(\mathbb{C})\partial_j s$ does not contain s. Since M is simple, this forces $\partial_j s = 0$.
(2) The other case is that $\alpha_j = -1$. Then $\partial_j x_j s = 0$, and for the same reason as before, this forces $x_j s = 0$.

We conclude that M is generated as an A_n-module by $s \in M_\alpha$, and that s is annihilated by $(x_j\partial_j - \alpha_j)$ for $\alpha_j \neq -1, 0$, by ∂_j for $\alpha_j = 0$, and by x_j for $\alpha_j = -1$. It is easy to see that there cannot be any other relations, and so we get

$$M \cong A_n/I,$$

where $I_\alpha \subseteq A_n$ is the left ideal generated by the n differential operators

$$\begin{cases} x_j\partial_j - \alpha_j & \text{for } \alpha_j \neq -1, 0, \\ \partial_j & \text{for } \alpha_j = 0, \\ x_j & \text{for } \alpha_j = -1. \end{cases}$$

We see that M is supported on the linear subspace

$$\operatorname{Supp} M_\alpha = \bigcap_{\alpha_j = -1} Z(x_j),$$

and so by Kashiwara's equivalence, it is the pushforward of a regular holonomic $\mathscr{D}$-module of normal crossing type on $\operatorname{Supp} M_\alpha$. Outside of the union of the hyperplanes $Z(x_j)$ with $\alpha_j \neq -1, 0$, the latter is a line bundle with integrable connection; this connection has a regular singularity at each of the hyperplanes in question, with monodromy $e^{2\pi i \alpha_j}$.

Now let see what we can say about the solutions of regular holonomic $\mathscr{D}$-modules of normal crossing type on $\mathbb{C}^n$. Since algebraic differential equations typically do not have algebraic solutions, we need to work in the analytic topology; we use the notation $\mathscr{O}_{\mathbb{C}^n}$ for the sheaf of holomorphic functions on $\mathbb{C}^n$, and the notation $\mathscr{D}_{\mathbb{C}^n}$ for the sheaf of differential operators with holomorphic coefficients. Let us write $\mathcal{M} = \mathscr{D}_{\mathbb{C}^n} \otimes_{A_n} M$ for the analytic $\mathscr{D}_{\mathbb{C}^n}$-module determined by the $A_n(\mathbb{C})$-module M. Recall that we have the (derived) solutions functor

$$\operatorname{Sol}(\mathcal{M}) = \mathbf{R}\mathcal{H}om_{\mathscr{D}_{\mathbb{C}^n}}\big(\mathcal{M}, \mathscr{O}_{\mathbb{C}^n}\big).$$

It can be computed for example by choosing a resolution of $\mathcal{M}$ by free $\mathscr{D}_{\mathbb{C}^n}$-modules, and then applying the usual solutions functor term by term. For simple modules of normal crossing type, this is easily done. Fix a multi-index $\alpha \in F$ as above. To keep the notation simple, let me set

$$P_j = \begin{cases} x_j \partial_j - \alpha_j & \text{if } \alpha_j \neq -1, 0, \\ \partial_j & \text{if } \alpha_j = 0, \\ x_j & \text{if } \alpha_j = -1. \end{cases}$$

Then our simple $\mathscr{D}_{\mathbb{C}^n}$-module has the form

$$\mathcal{M}_\alpha = \mathscr{D}_{\mathbb{C}^n} / \mathscr{D}_{\mathbb{C}^n}(P_1, \dots, P_n),$$

The Koszul complex for $P_1, \dots, P_n$ gives a resolution by free $\mathscr{D}_{\mathbb{C}^n}$-modules:

$$\mathscr{D}_{\mathbb{C}^n} \to \mathscr{D}_{\mathbb{C}^n}^{\oplus n} \to \cdots \to \mathscr{D}_{\mathbb{C}^n}^{\oplus \binom{n}{2}} \to \mathscr{D}_{\mathbb{C}^n}^{\oplus n} \to \mathscr{D}_{\mathbb{C}^n}$$

Consequently, $\operatorname{Sol}(\mathcal{M}_\alpha)$ is represented by the complex

$$\mathscr{O}_{\mathbb{C}^n} \to \mathscr{O}_{\mathbb{C}^n}^{\oplus n} \to \mathscr{O}_{\mathbb{C}^n}^{\oplus \binom{n}{2}} \to \cdots \to \mathscr{O}_{\mathbb{C}^n}^{\oplus n} \to \mathscr{O}_{\mathbb{C}^n}, \tag{23.3}$$

placed in degrees $0, 1, \dots, n$, and with a Koszul-type differential, induced by the n operators $f \mapsto T_j f$. We are interested in computing the cohomology sheaves of this complex.

Example 23.4. For $n = 1$, there are three cases. If $\alpha = 0$, the complex looks like

$$\mathscr{O}_{\mathbb{C}} \xrightarrow{\partial} \mathscr{O}_{\mathbb{C}}.$$

By the holomorphic Poincaré lemma (or by a direct computation with power series), this complex only has cohomology in degree 0, where we get the constant sheaf $\mathbb{C}$. If $\alpha = -1$, the complex looks like

$$\mathscr{O}_{\mathbb{C}} \xrightarrow{x} \mathscr{O}_{\mathbb{C}}.$$

It only has cohomology in degree 1, where we get a one-dimensional skyscraper sheaf at the origin. Lastly, if $\alpha \neq -1, 0$, the complex looks like

$$\mathscr{O}_{\mathbb{C}} \xrightarrow{x\partial - \alpha} \mathscr{O}_{\mathbb{C}}.$$

This only has cohomology in degree 0. Away from the origin, the multi-valued holomorphic function x^α solves the equation $(x\partial - \alpha)f = 0$, and so we get a locally constant sheaf on $\mathbb{C}^*$, with monodromy $e^{2\pi i \alpha}$. At the origin, the function x^α does not make sense, and in fact, the equation $(x\partial - \alpha)f = 0$ does not have a solution that is holomorphic in a neighborhood of the origin. So in this case, the 0-th cohomology sheaf of the complex is a so-called constructible sheaf: it is locally constant on $\mathbb{C}^*$, but with a different stalk at the origin. Note that in each case, exactly one cohomology sheaf is nonzero; and if the nonzero cohomology sheaf occurs in degree 0, it is supported on all of $\mathbb{C}$; if it occurs in degree 1, then it is supported at the origin.

By working with power series, one can show that the complex in (23.3) is (locally) quasi-isomorphic to a product; thus its cohomology is described by what happens for each of the n operators T_j individually. The conclusion is that (23.3) has exactly one nonzero cohomology sheaf, say in degree k (where k is the number of j such that $\alpha_j = -1$); moreover, that cohomology sheaf is supported on the linear subspace

$$\bigcap_{\alpha_j=-1} Z(x_j),$$

whose codimension is exactly k. It is also a constructible sheaf, meaning locally constant (of rank 0 or 1) on each stratum of the natural stratification on $\mathbb{C}^n$.

From this, we can deduce what happens for $\operatorname{Sol}(\mathcal{M})$ in general. Recall that $\mathcal{M}$ has a finite composition series whose subquotients $\mathcal{M}_1, \dots, \mathcal{M}_r$ are simple.

Example 23.5. Suppose that $\mathcal{M}$ has a composition series of length two:

$$0 \to \mathcal{M}_1 \to \mathcal{M} \to \mathcal{M}_2 \to 0$$

Since the solutions functor is contravariant, we obtain a long exact sequence

$$\mathcal{H}^{i-1} \operatorname{Sol}(\mathcal{M}_1) \to \mathcal{H}^i \operatorname{Sol}(\mathcal{M}_2) \to \mathcal{H}^i \operatorname{Sol}(\mathcal{M}) \to \mathcal{H}^i \operatorname{Sol}(\mathcal{M}_1) \to \mathcal{H}^{i+1} \operatorname{Sol}(\mathcal{M}_2)$$

Since $\operatorname{Sol}(\mathcal{M}_1)$ and $\operatorname{Sol}(\mathcal{M}_2)$ each have only a single nonzero cohomology sheaf, it follows that $\operatorname{Sol}(\mathcal{M})$ can have at most two nonzero cohomology sheaves, both constructible with respect to the natural stratification on $\mathbb{C}^n$. Moreover, $\dim \operatorname{Supp} \mathcal{H}^i \operatorname{Sol}(\mathcal{M}) \geq i$. The inequality can be strict, for example if $\mathcal{H}^i \operatorname{Sol}(\mathcal{M}_2) \neq 0$ and $\mathcal{H}^{i-1} \operatorname{Sol}(\mathcal{M}_1) \neq 0$; then $\mathcal{H}^i \operatorname{Sol}(\mathcal{M})$ is a quotient of the constructible sheaf $\mathcal{H}^i \operatorname{Sol}(\mathcal{M}_2)$, whose support is a linear subspace of codimension i. It follows that $\mathcal{H}^i \operatorname{Sol}(\mathcal{M})$ is still constructible, but its support may be smaller than than of $\mathcal{H}^i \operatorname{Sol}(\mathcal{M}_2)$.

In general, we have a spectral sequence

$$E_1^{p,q} = \mathcal{H}^{p+q} \operatorname{Sol}(\mathcal{M}_p) \Longrightarrow \mathcal{H}^{p+q} \operatorname{Sol}(\mathcal{M}).$$

Each $\operatorname{Sol}(\mathcal{M}_p)$ has exactly one nonzero cohomology sheaf, which is constructible for the natural stratification on $\mathbb{C}^n$; if $\mathcal{H}^j \operatorname{Sol}(\mathcal{M}_p) \neq 0$, then it is supported on a linear subspace of codimension j. Since kernels and cokernels of morphisms between constructible sheaves are again constructible, we see that all cohomology sheaves of $\operatorname{Sol}(\mathcal{M})$ are constructible; it also follows, as in the example, that

$$\operatorname{codim} \operatorname{Supp} \mathcal{H}^j \operatorname{Sol}(\mathcal{M}) \geq j.$$

Exercises.

Exercise 23.1. Suppose that we are given a family of k-vector spaces M_α, indexed by $\alpha \in A$, and a family of linear mappings $D_{i,j} \colon M_\alpha \to M_{\alpha+e_i-e_j}$.

(1) Show that if the relations in (23.1) hold, and $D_{0,0} + D_{1,1} + \dots + D_{n,n} = 0$, then the direct sum

$$M = \bigoplus_{\alpha \in A} M_\alpha$$

becomes a left module over $R = \Gamma(\mathbb{P}^n, \mathscr{D}_{\mathbb{P}^n})$.

(2) Suppose that M is finitely generated as an R-module, and that each operator $D_{j,j} - \alpha_j$ acts nilpotently on M_α. Show that the characteristic variety of $\mathcal{M} = \mathscr{D}_{\mathbb{P}^n} \otimes_R M$ is contained in the set $Z(x_0\xi_0, x_1\xi_1, \dots, x_n\xi_n)$.

(3) Show that $\mathcal{M}$ is a regular holonomic $\mathscr{D}_{\mathbb{P}^n}$-module of normal crossing type.

Exercise 23.2. Find the decomposition of $\Gamma(\mathbb{P}^n, \mathcal{M})$ in the following cases:

(1) $\mathcal{M} = \mathscr{O}_{\mathbb{P}^n}$

(2) $\mathcal{M} = j_+\mathscr{O}_U$, where $U = \mathbb{P}^n \setminus Z(x_0x_1 \cdots x_n)$

(3) $\mathcal{M} = i_+\mathscr{O}_{\mathbb{P}^{n-1}}$, where $\mathbb{P}^{n-1} = Z(x_0)$.

Exercise 23.3. Let $\mathcal{M}$ be a regular holonomic $\mathscr{D}_{\mathbb{P}^n}$-module of normal crossing type. Given the decomposition for $\Gamma(\mathbb{P}^n, \mathcal{M})$, determine the resulting decomposition for the holonomic dual of $\mathcal{M}$.

Lecture 24: May 6

The Riemann-Hilbert correspondence. Last time, we showed that the solution complex of a regular holonomic $\mathscr{D}$-module of normal crossing type has several special properties: its cohomology sheaves are locally constant on the strata of the divisor, and the dimensions of their supports satisfy a collection of inequalities. This is a special case of the Riemann-Hilbert correspondence, which relates regular holonomic $\mathscr{D}$-modules and constructible sheaves.

Let us begin with a few basic definitions. Let X be a nonsingular algebraic variety over the complex numbers. A *stratification* is a decomposition

$$X = \bigsqcup_{\alpha \in A} X_\alpha$$

into locally closed algebraic subsets, called *strata*, such that each X_α is nonsingular, and such that the Zariski-closure of each X_α is a union of finitely many other strata. The same definition makes sense on complex manifolds, taking each X_α to be a locally closed complex submanifold.

Example 24.1. The divisor $x_1 \cdots x_n = 0$ induces a natural stratification on $\mathbb{A}^n$ with 2^n strata, indexed by subsets $I \subseteq \{1, \dots, n\}$. The stratum corresponding to the subset I consists of those points where $x_i = 0$ for every $i \in I$, and $x_i \neq 0$ for every $i \notin I$.

Example 24.2. We can stratify $\mathbb{A}^2$ according to the singularities of the nodal curve $C = Z(y^2 - x^2 - x^3)$, into a 2-dimensional stratum $\mathbb{A}^2 \setminus C$, a 1-dimensional stratum $C \setminus \{(0,0)\}$, and a 0-dimensional stratum $\{(0,0)\}$.

We also need the notion of constructibility for sheaves. It is necessary to work in the classical (or analytic) topology, because the Zariski topology is too coarse to allow for interesting locally constant sheaves. Given a nonsingular algebraic variety X, we denote by X^{an} the associated complex manifold, with the topology induced by the usual Euclidean topology on $\mathbb{C}^n$. Let F be a sheaf of finite-dimensional $\mathbb{C}$-vector spaces on X^{an}. This means that for every open subset $U \subseteq X^{an}$, the space of sections $\Gamma(U, F)$ is a finite-dimensional $\mathbb{C}$-vector space. We say that F is *constructible* if there is a stratification

$$X = \bigsqcup_{\alpha \in A} X_\alpha$$

such that the restriction of F to each stratum X_α^{an} is a locally constant sheaf. (Constructible sheaves on arbitrary complex manifolds are defined in a similar way.) Every locally constant sheaf is constructible, of course, but the point is that the usual functors on sheaves preserve constructibility. (Going from locally constant sheaves to constructible sheaves is very similar to going from locally free sheaves to coherent sheaves, in that sense.)

Example 24.3. If $f \colon X \to Y$ is a proper morphism between nonsingular algebraic varieties, and if $f^{an} \colon X^{an} \to Y^{an}$ denotes the resulting proper holomorphic mapping between complex manifolds, then the sheaves $R^i f_*^{an} \mathbb{C}_{X^{an}}$ are constructible. The reason is that one can find a stratification for Y, in such a way that the restriction of f to each stratum of Y is a topological fiber bundle.

Example 24.4. If $j \colon U \hookrightarrow X$ is an open embedding, and $j^{an} \colon U^{an} \hookrightarrow X^{an}$ denotes the resulting embedding of complex manifolds, then the sheaves $R^i j_*^{an} \mathbb{C}_{U^{an}}$ are constructible. This is easy to show in the case where $X \setminus U$ is a normal crossing divisor; the general case follows by using resolution of singularities and the result in the previous example.

More generally, one can show that the usual direct and inverse image functors on sheaves preserve constructibility: if $f\colon X \to Y$ in any morphism between nonsingular algebraic varieties, and F any constructible sheaf on X^{an}, then $R^i f_*^{an} F$ is a constructible sheaf on Y^{an}. Likewise, if G is any constructible sheaf on Y^{an}, then $(f^{an})^{-1}G$ is a constructible sheaf on X^{an}. One can say the same thing in the language of derived categories. Denote by $D_c^b(\mathbb{C}_{X^{an}})$ the derived category of (cohomologically) constructible sheaves; its objects are complexes of sheaves of $\mathbb{C}$-vector spaces on X^{an} whose cohomology sheaves are constructible (and zero in all but finitely many degrees). Then if $f\colon X \to Y$ is any morphism between nonsingular algebraic varieties, the usual derived pushforward of sheaves gives an exact functor

$$\mathbf{R}f_*^{an}\colon D_c^b(\mathbb{C}_{X^{an}}) \to D_c^b(\mathbb{C}_{Y^{an}}),$$

and the usual inverse image of sheaves gives an exact functor

$$(f^{an})^{-1}\colon D_c^b(\mathbb{C}_{Y^{an}}) \to D_c^b(\mathbb{C}_{X^{an}}).$$

We can now state the first general result about solution complexes of regular holonomic $\mathscr{D}$-modules. Let $\mathscr{D}_X$ be the usual sheaf of differential operators on X, and denote by $\mathscr{D}_{X^{an}}$ the sheaf of differential operators with holomorphic coefficients on the complex manifold X^{an}. Given a coherent $\mathscr{D}_X$-module $\mathcal{M}$, we denote by $\mathcal{M}^{an}$ the associated analytic $\mathscr{D}_{X^{an}}$-module; this can be constructed using local presentations of $\mathcal{M}$, for example. The following result was proved by Kashiwara in his thesis; it is usually called "Kashiwara's constructibility theorem".

Theorem 24.5. *Let X be a nonsingular algebraic variety, and $\mathcal{M}$ a holonomic left $\mathscr{D}_X$-module. Then the solution complex*

$$\operatorname{Sol}(\mathcal{M}) = \mathbf{R}\mathcal{H}om_{\mathscr{D}_{X^{an}}}\big(\mathcal{M}^{an}, \mathscr{O}_{X^{an}}\big)$$

is constructible, hence an object of $D_c^b(\mathbb{C}_{X^{an}})$.

In fact, Kashiwara proves this result for holonomic $\mathscr{D}$-modules on complex manifolds. One consequence is that one has an exact (contravariant) solutions functor

$$\operatorname{Sol}\colon D_h^b(\mathscr{D}_X) \to D_c^b(\mathbb{C}_{X^{an}})^{op}$$

that associates to every complex of $\mathscr{D}_X$-modules with holonomic cohomology a constructible complex of solutions. We saw a very special case of this result last time, namely solutions of regular holonomic $\mathscr{D}$-modules of normal crossing type.

The solution functor is contravariant, but there is also a covariant version of Kashiwara's theorem. Recall that the Spencer complex

$$\operatorname{Sp}(\mathscr{D}_X) = \Big[\mathscr{D}_X \otimes_{\mathscr{O}_X} \bigwedge^n \mathscr{T}_X \to \cdots \to \mathscr{D}_X \otimes_{\mathscr{O}_X} \bigwedge^2 \mathscr{T}_X \to \mathscr{D}_X \otimes_{\mathscr{O}_X} \mathscr{T}_X \to \mathscr{D}_X\Big]$$

is a resolution of $\mathscr{O}_X$ by locally free left $\mathscr{D}_X$-modules; likewise, $\operatorname{Sp}(\mathscr{D}_{X^{an}})$ is a resolution of $\mathscr{O}_{X^{an}}$ by locally free left $\mathscr{D}_{X^{an}}$-modules. Thus

$$\begin{aligned}\operatorname{Sol}(\mathcal{M}) &= \mathbf{R}\mathcal{H}om_{\mathscr{D}_{X^{an}}}\big(\mathcal{M}^{an}, \mathscr{O}_{X^{an}}\big)\\ &\cong \mathbf{R}\mathcal{H}om_{\mathscr{D}_{X^{an}}}\big(\mathcal{M}^{an}, \operatorname{Sp}(\mathscr{D}_{X^{an}})\big)\\ &\cong \mathbf{R}\mathcal{H}om_{\mathscr{D}_{X^{an}}}\big(\mathcal{M}^{an}, \mathscr{D}_{X^{an}}\big) \otimes_{\mathscr{D}_{X^{an}}} \operatorname{Sp}(\mathscr{D}_{X^{an}}).\end{aligned}$$

Now suppose that $\mathcal{M}$ is holonomic. Then the complex $\mathbf{R}\mathcal{H}om_{\mathscr{D}_X}(\mathcal{M}, \mathscr{D}_X)$ only has cohomology in degree n, and

$$\mathcal{E}xt^n_{\mathscr{D}_X}(\mathcal{M}, \mathscr{D}_X) = \mathcal{M}^*$$

is the holonomic dual (which is a holonomic right $\mathscr{D}_X$-module). Consequently,

$$\mathbf{R}\mathcal{H}om_{\mathscr{D}_{X^{an}}}\big(\mathcal{M}^{an}, \mathscr{D}_{X^{an}}\big) \cong \mathcal{M}^{*,an}[-n],$$

and after plugging this into the relation from above, we get

$$\mathrm{Sol}(\mathcal{M}) \cong \mathcal{M}^{*,an}[-n] \otimes_{\mathscr{D}_{X^{an}}} \mathrm{Sp}(\mathscr{D}_{X^{an}}) \cong \mathrm{Sp}(\mathcal{M}^{*,an})[-n]. \tag{24.6}$$

Under the conversion between right and left $\mathscr{D}$-modules, the Spencer complex of a right $\mathscr{D}$-module goes to the de Rham complex of a left $\mathscr{D}$-module. This leads to the following equivalent formulation of Kashiwara's constructibility theorem: If $\mathcal{M}$ is a holonomic left $\mathscr{D}_X$-module on a nonsingular algebraic variety X, then the de Rham complex

$$\mathrm{DR}(\mathcal{M}^{an}) = \Big[\mathcal{M}^{an} \to \Omega^1_{X^{an}} \otimes \mathcal{M}^{an} \to \cdots \to \Omega^n_{X^{an}} \otimes \mathcal{M}^{an}\Big],$$

placed in degrees $-n, \dots, 0$, is constructible. More generally, the de Rham functor

$$\mathrm{DR}\colon D^b_h(\mathscr{D}_X) \to D^b_c(\mathbb{C}_{X^{an}})$$

is an exact covariant functor.

Kashiwara's theorem makes no assumptions about regularity, but the price to pay is that many different $\mathscr{D}$-modules can have the same solution complex.

Example 24.7. Here is the simplest example of this phenomenon. On $\mathbb{A}^1$, consider the family of $\mathscr{D}_{\mathbb{A}^1}$-modules $\mathcal{M}_\lambda = \mathscr{D}_{\mathbb{A}^1}/\mathscr{D}_{\mathbb{A}^1}(\partial - \lambda)$, indexed by $\lambda \in \mathbb{C} \setminus \{0\}$. We have already seen that these $\mathscr{D}$-modules have an irregular singularity at infinity. The solution complex of $\mathcal{M}_\lambda$ is

$$\mathscr{O}_\mathbb{C} \xrightarrow{\partial - \lambda} \mathscr{O}_\mathbb{C}.$$

The kernel of $\partial - \lambda$ is clearly spanned by the function $e^{\lambda x}$, while the cokernel is trivial; this means that the solution complex is always isomorphic to the constant sheaf $\mathbb{C}$, independently of λ. On the other hand, $\mathcal{M}_\lambda$ and $\mathcal{M}_\mu$ are not isomorphic as $\mathscr{D}_{\mathbb{A}^1}$-modules for $\lambda \neq \mu$.

If one imposes the condition of regularity, then this problem goes away, and the solutions functor (as well as the de Rham functor) becomes an equivalence of categories. This is the content of the famous Riemann-Hilbert correspondence.

Theorem 24.8. *Let X be a nonsingular algebraic variety. Then the functors*

$$\mathrm{Sol}\colon D^b_h(\mathscr{D}_X) \to D^b_c(\mathbb{C}_{X^{an}})^{op}$$
$$\mathrm{DR}\colon D^b_h(\mathscr{D}_X) \to D^b_c(\mathbb{C}_{X^{an}})$$

are equivalences of categories.

This result again holds more generally on complex manifolds. There are three proofs: an analytic proof by Kashiwara; a more algebraic proof by Mebkhout; and a completely algebraic proof by Bernstein (which only works on algebraic varieties). The Riemann-Hilbert correspondence also respects the various functors on both sides: for example,

$$\mathrm{DR} \circ f_+ \cong \mathbf{R}f_* \circ \mathrm{DR} \quad \text{and} \quad \mathrm{DR} \circ \mathbf{L}j^* \cong j^{-1} \circ \mathrm{DR}.$$

These isomorphisms do not hold without the assumption of regularity. The Riemann-Hilbert correspondence therefore establishes a direct link between algebraic objects (regular holonomic $\mathscr{D}$-modules) and topological objects (constructible sheaves).

Example 24.9. The holonomic dual also has a natural interpretation in terms of the Riemann-Hilbert correspondence. On $D^b_c(\mathbb{C}_{X^{an}})$, one has Verdier's duality functor

$$\mathbb{D}_{X^{an}}\colon D^b_c(\mathbb{C}_{X^{an}}) \to D^b_c(\mathbb{C}_{X^{an}})^{op}, \quad F \mapsto \mathbf{R}\mathcal{H}om_{\mathbb{C}_{X^{an}}}\big(F, \mathbb{C}_{X^{an}}[2n]\big),$$

where $n = \dim X$. One can show that, for any holonomic $\mathscr{D}_X$-module $\mathcal{M}$, one has an isomorphism

$$\mathbb{D}_{X^{an}}\big(\mathrm{DR}(\mathcal{M}^{an})\big) \cong \mathrm{Sp}(\mathcal{M}^{*,an})$$

which means that the Riemann-Hilbert correspondence turns holonomic duality into Verdier duality.

Perverse sheaves. The Riemann-Hilbert correspondence works on the level of the derived category. Where do regular holonomic $\mathscr{D}$-modules go under the equivalence of categories? We saw last time that the solution complex of a regular holonomic $\mathscr{D}$-module of normal crossing type satisfies a collection of inequalities: the j-th cohomology sheaf of $\mathrm{Sol}(\mathcal{M})$ is supported on a union of strata of codimension at least j. Kashiwara proved that this is true for arbitrary holonomic $\mathscr{D}$-modules: if $\mathcal{M}$ is a holonomic $\mathscr{D}$-module on a nonsingular algebraic variety (or, more generally, on a complex manifold), then

$$\operatorname{codim} \operatorname{Supp} R^j \operatorname{Sol}(\mathcal{M}) \geq j$$

for every $j \in \mathbb{Z}$. Using the identity in (24.6), an equivalent formulation is that

$$\dim \operatorname{Supp} \mathcal{H}^j \operatorname{DR}(\mathcal{M}^{an}) \leq -j$$

for every $j \in \mathbb{Z}$. One gets a similar collection of inequalities also for the Verdier dual $\mathbb{D}_{X^{an}} \operatorname{DR}(\mathcal{M}^{an})$, because of the identity in Example 24.9. This motivates the following definition.

Definition 24.10. A complex $F \in D_c^b(\mathbb{C}_{X^{an}})$ is called a *perverse sheaf* if

$$\dim \operatorname{Supp} \mathcal{H}^j F \leq -j \quad \text{and} \quad \dim \operatorname{Supp} \mathcal{H}^j \mathbb{D}_{X^{an}}(F) \leq -j$$

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

Example 24.11. If $\mathcal{M}$ is a holonomic $\mathscr{D}_X$-module, then $\operatorname{DR}(\mathcal{M}^{an})$ is a perverse sheaf. This is simply a rewording of Kashiwara's theorem. Note that regularity is not needed here.

The definition (and the somewhat strange name) of perverse sheaves is due to Beilinson, Bernstein, Deligne, and Gabber. They showed that the collection of perverse sheaves forms an abelian category contained in $D_c^b(\mathbb{C}_{X^{an}})$. The collection of inequalities in the definition had actually appeared in two completely independent places: once in Kashiwara's study of holonomic $\mathscr{D}$-modules, and then again in the intersection homology theory of Goresky and Macpherson. This circumstance is of course explained by the Riemann-Hilbert correspondence. In fact, once Theorem 24.8 is known, purely formal reasoning implies that the de Rham functor

$$\operatorname{DR} \colon D_h^b(\mathscr{D}_X) \to D_c^b(\mathbb{C}_{X^{an}})$$

takes the abelian category of regular holonomic $\mathscr{D}_X$-modules isomorphically to the abelian category of perverse sheaves. Unfortunately, I cannot offer you any good explanation of what perverse sheaves really are, other than saying that they are the image of the regular holonomic $\mathscr{D}$-modules under the Riemann-Hilbert correspondence. From this point of view, the crucial result is the equivalence between the two derived categories; the collection of inequalities is just what one gets when one goes from one side to the other.

Exercises.

Exercise 24.1. Show that if $\lambda \neq \mu$, then $\mathcal{M}_\lambda = \mathscr{D}_{\mathbb{A}^1}/\mathscr{D}_{\mathbb{A}^1}(\partial - \lambda)$ is not isomorphic to $\mathcal{M}_\mu$ as a $\mathscr{D}_{\mathbb{A}^1}$-module.

Lecture 25: May 8

Meromorphic connections. Before the full Riemann-Hilbert correspondence was proved, Deligne established an important special case. It has to do with the relationship between locally constant sheaves and vector bundles with integrable connection. Suppose that X is a nonsingular and proper algebraic variety over the complex numbers. If we are given a vector bundle of rank r with integrable connection, then the subsheaf of flat sections is a locally constant sheaf of rank m (with respect to the analytic topology). Conversely, given a locally constant sheaf of rank m, say E, we can form the holomorphic vector bundle $\mathscr{E} = \mathscr{O}_X \otimes_{\mathbb{C}} E$, which has the same (locally constant) transition functions as E. The formula

$$\nabla(f \otimes s) = df \otimes s$$

defines an integrable connection on $\mathscr{E}$, and the subsheaf of ∇-flat sections is of course isomorphic to E. Lastly, X is proper, and so the pair $(\mathscr{E}, \nabla)$ actually comes from an *algebraic* vector bundle with intgrable connection (by a version of Serre's GAGA theorem). The conclusion is that the (a priori topological) object E is actually algebraic.

Deligne's version of the Riemann-Hilbert correspondence generalizes this to not necessarily proper varieties. It goes through an intermediate class of objects, called meromorphic connections. Here is the definition. Let X be a complex manifold, and $D \subseteq X$ a divisor. For simplicity, we are only going to consider the case where D has simple normal crossing singularities: D is a union of nonsingular hypersurfaces meeting transversely. In suitable local coordinates $x_1, \dots, x_n$, the equation defining D is of the form $x_1 \cdots x_r = 0$. We let

$$\mathscr{O}_X(*D)$$

be the sheaf of meromorphic functions on X that are holomorphic on $X \setminus D$; it is naturally a subsheaf of $j_* \mathscr{O}_{X \setminus D}$, where $j \colon X \setminus D \hookrightarrow X$ is the inclusion of the complement. The notation $*D$ is supposed to remind you of the pole order along D. Locally, $\mathscr{O}_X(*D)$ is isomorphic to $\mathscr{O}_X[t]/(ht-1)$, where h is a local equation for D; it follows that $\mathscr{O}_X(*D)$ is still a coherent sheaf of $\mathscr{O}_X$-algebras.

Definition 25.1. A *meromorphic connection* is a coherent $\mathscr{O}_X(*D)$-module M, together with an integrable connection

$$\nabla \colon M \to \Omega_X^1 \otimes_{\mathscr{O}_X} M$$

that satisfies the Leibniz rule $\nabla(fs) = df \otimes s + f \nabla s$ and the integrability condition $[\nabla_\theta, \nabla_{\theta'}] = \nabla_{[\theta, \theta']}$.

Note. In the Leibniz rule, we are considering only $f \in \mathscr{O}_X$, but the same formula works for every $f \in \mathscr{O}_X(*D)$. To make this precise, define $\Omega_X^1(*D)$ as the sheaf of meromorphic one-forms on X that are holomorphic on $X \setminus D$, so that

$$\Omega_X^1(*D) = \Omega_X^1 \otimes_{\mathscr{O}_X} \mathscr{O}_X(*D).$$

We can then consider ∇ as a $\mathbb{C}$-linear morphism

$$\nabla \colon M \to \Omega_X^1(*D) \otimes_{\mathscr{O}_X(*D)} M,$$

and now the Leibniz rule makes sense for $f \in \mathscr{O}_X(*D)$.

A meromorphic connection is naturally a left $\mathscr{D}_X$-module, since the two identities imply that the left action by $\mathscr{T}_X$ extends to a left action by $\mathscr{D}_X$ (see the discussion in Lecture 10). On $X \setminus D$, the $\mathscr{D}$-module is coherent, and therefore a holomorphic vector bundle with integrable connection. In that sense, a meromorphic connection is an extension of a vector bundle with integrable connection on $X \setminus D$ to an object on X with singularities along D.

Definition 25.2. If (M, ∇) and (N, ∇) are two meromorphic connections, then a *morphism* $\varphi\colon (M,\nabla) \to (N,\nabla)$ is a morphism of $\mathscr{O}_X(*D)$-module $\varphi\colon M \to N$ that is compatible with the connections, in the sense that

$$\nabla\big(\varphi(s)\big) = (\mathrm{id} \otimes \varphi)(\nabla s).$$

We denote by $\mathrm{Conn}(X, D)$ the category of meromorphic connections on (X, D). It is an abelian category. There are two simple but useful observations about morphisms in $\mathrm{Conn}(X, D)$. The first says that morphisms are determined by what their restriction to $X \setminus D$.

Proposition 25.3. *Let $\varphi\colon (M,\nabla) \to (N,\nabla)$ be a morphism of meromorphic connections. If $\varphi\big|_{X\setminus D}$ is an isomorphism, then φ is an isomorphism.*

Proof. The kernel and cokernel of φ are meromorphic connections whose support is, by construction, contained inside D. It is therefore enough to prove that a meromorphic connection (M, ∇) such that $\operatorname{Supp} M \subseteq D$ must be trivial. Let s be any local section of M, and h a local equation for D. The subsheaf $\mathscr{O}_X \cdot s \subseteq M$ is coherent over $\mathscr{O}_X$, and its support is contained inside D, and so $h^m s = 0$ for $m \gg 0$. But then $s = h^{-m}(h^m s) = 0$, proving that $M = 0$. ☐

The second observation is useful for functoriality questions. Suppose that (M, ∇) and (N, ∇) are two meromorphic connections. Then

$$\mathcal{H}om_{\mathscr{O}_X(*D)}(M, N)$$

is again an $\mathscr{O}_X(*D)$-module in a natural way, and the formula

$$(\nabla\varphi)(s) = (\mathrm{id} \otimes \varphi)(\nabla s) - \nabla\big(\varphi(s)\big)$$

defines an integrable connection that makes $\mathcal{H}om_{\mathscr{O}_X(*D)}(M, N)$ into a meromorphic connection. You should check that morphisms of meromorphic connections $\varphi\colon (M,\nabla) \to (N,\nabla)$ are exactly the same thing as ∇-flat global sections of $\mathcal{H}om_{\mathscr{O}_X(*D)}(M, N)$.

Deligne's theorem on meromorphic connections. Deligne proved that locally constant sheaves on $X \setminus D$ correspond to meromorphic connections on (X, D) that are *regular* along D. Regularity was originally defined by restricting to curves, but in the case where D is a normal crossing divisor, we can use another definition that is closer to the Kashiwara-Kawai notion of regularity for $\mathscr{D}$-modules.

Definition 25.4. A meromorphic connection (M, ∇) is called *regular* if there is a locally free $\mathscr{O}_X$-module L with

$$M \cong \mathscr{O}_X(*D) \otimes_{\mathscr{O}_X} L,$$

such that in any local trivialization of L, the connection has at worst logarithmic poles along D.

More precisely, suppose that $e_1, \dots, e_m$ form a local trivialization for L. Then the condition is that

$$\nabla e_i = \sum_{j,k} a_{i,j}^k \frac{dx_k}{x_k} \otimes e_j,$$

for certain holomorphic functions $a_{i,j}^k$. Since L is then preserved by the differential operators $x_1\partial_1, \dots, x_n\partial_n$, this means that M, viewed as a left $\mathscr{D}_X$-module, is regular in the sense of Kashiwara and Kawai. The letter L comes from the fact that L is traditionally called a *lattice.*

Keeping the notation from above, we let $A^k \in \mathrm{Mat}_{m\times m}(\mathscr{O}_X)$ be the matrix with entries $a_{i,j}^k$. The restriction of A^k to the divisor D_k, defined by the equation $x_k = 0$, is a well-defined endomorphism of the locally free sheaf $L\big|_{D_k}$, called the *residue* of ∇ along D_k. We use the symbol

$$\operatorname{Res}^L_{D_k}(\nabla) = A^k\big|_{D_k}$$

to denote the residue. We may drop the superscript L when the lattice is clear from the context.

Lemma 25.5. *Let (M,∇) be a meromorphic connection with lattice L.*
(a) On $D_k \cap D_\ell$, the residues $\operatorname{Res}_{D_k}(\nabla)$ and $\operatorname{Res}_{D_\ell}(\nabla)$ commute.
(b) The eigenvalues of $\operatorname{Res}_{D_\ell}(\nabla)$ are locally constant along D_ℓ.

Proof. In the notation from above, we have

$$\nabla e_i = \sum_{j,k} a^k_{i,j} \frac{dx_k}{x_k} \otimes e_j,$$

and A^k is the $m \times m$-matrix with entries $a^k_{i,j}$. With respect to the trivialization $e_1, \dots, e_m$, we therefore have $\nabla_{\partial_k} = A^k/x_k$. The integrability condition for the connection is $[\nabla_{\partial_k}, \nabla_{\partial_\ell}] = 0$, which expands out to

$$\frac{\partial}{\partial x_k}\left(\frac{A^\ell}{x_\ell}\right) + \frac{A^\ell}{x_\ell}\frac{A^k}{x_k} = \frac{\partial}{\partial x_\ell}\left(\frac{A^k}{x_k}\right) + \frac{A^k}{x_k}\frac{A^\ell}{x_\ell}.$$

After rearranging the terms, this becomes

$$x_k\partial_k(A^\ell) + A^\ell A^k = x_\ell \partial_\ell(A^k) + A^k A^\ell,$$

and so the restriction of the two matrices A^k and A^ℓ to the set $x_k = x_\ell = 0$ commute with each other.

For the proof of the second assertion, denote by $\bar{L}$ the restriction of L to the divisor D_ℓ; similarly, $\bar{A}^k$ is the restriction of A^k, and so on. The formula

$$\nabla \bar{e}_i = \sum_{j,k\neq \ell} \bar{a}^k_{i,j} \frac{dx_k}{x_k} \otimes \bar{e}_j$$

defines an integrable connection with logarithmic poles on $\bar{L}$, and one checks that $\bar{A}^\ell$ is a horizontal section of $\mathcal{H}om_{\mathscr{O}_{D_\ell}}(\bar{L},\bar{L})$. It follows that the eigenvalues of $\bar{A}^\ell$ must be locally constant. $\square$

Deligne's main theorem is that every bundle with integrable connection on U can be uniquely extended to a regular meromorphic connection on (X,D); in fact, even the lattice is more or less unique, except for a small ambiguity in the eigenvalues of the residues.

Theorem 25.6. *Let X be a complex manifold, and $D \subseteq X$ a divisor with simple normal crossing singularities. Set $U = X \setminus D$, and fix a section $\tau \colon \mathbb{C}/\mathbb{Z} \to \mathbb{C}$ of the projection $\mathbb{C} \to \mathbb{C}/\mathbb{Z}$. Given $(M,\nabla) \in \operatorname{Conn}(U)$, there is a unique locally free sheaf L_τ on X with the following three properties:*
(a) One has $L_\tau\big|_U = M$.
(b) The connection $\nabla \colon M \to \Omega^1_U \otimes_{\mathscr{O}_U} M$ extends to

$$\nabla \colon M_\tau \to \Omega^1_X \otimes_{\mathscr{O}_X} M_\tau,$$

where $M_\tau = \mathscr{O}_X(*D) \otimes_{\mathscr{O}_X} L_\tau$.
(c) At each irreducible component of D, the residue of ∇ has eigenvalues in the set $\tau(\mathbb{C}/\mathbb{Z}) \subseteq \mathbb{C}$.

Moreover, with the above choice of L_τ, the restriction mapping

$$\Gamma(X, M_\tau)^\nabla \to \Gamma(U, M)^\nabla$$

from ∇-flat sections of M_τ to ∇-flat sections of M is an isomorphism.

Proof of Deligne's theorem. The proof of Deligne's theorem has two parts. The first part is to prove that L_τ is unique (up to isomorphism). The second part is to construct a suitable lattice L_τ locally on X; the local objects can then be glued together into a global lattice using uniqueness.

Let us start with the local existence, since that is easier. Since we are working locally, we can assume that $X = \Delta^n$, where $\Delta \subseteq \mathbb{C}$ is an open disk containing the origin. The divisor D will be given by the equation $x_1 \cdots x_r = 0$, and so $U = (\Delta^*)^r \times \Delta^{n-r}$. By the correspondence between vector bundles with integrable connection and locally constant sheaves, $(M, \nabla) \in \operatorname{Conn}(U)$ corresponds to a locally constant sheaf on U, hence to a representation $\pi_1(U) \to \mathrm{GL}_m(\mathbb{C})$, where m is the rank of M. Since the fundamental group of U is abelian, this is equivalent to giving r commuting matrices $C^1, \dots, C^r \in \mathrm{GL}_m(\mathbb{C})$. (These are the monodromy matrices of the locally constant sheaf.)

It is a simple exercise to show that there are matrices $\Gamma_1, \dots, \Gamma_r \in \mathrm{Mat}_{m\times m}(\mathbb{C})$, uniquely determined by the following three conditions:

(1) $e^{2\pi i\, \Gamma^j} = C^j$,
(2) the eigenvalues of Γ^j lie in the set $\tau(\mathbb{C}/\mathbb{Z})$,
(3) $\Gamma^1, \dots, \Gamma^r$ commute.

We can now define $L_\tau = \mathscr{O}_X^{\oplus m}$, and put a meromorphic connection on the free $\mathscr{O}_X(*D)$-module $M_\tau = \mathscr{O}_X(*D)^{\oplus m}$ by the formula

$$\nabla e_i = \sum_{j,k} \Gamma^k_{i,j} \frac{dx_k}{x_k} \otimes e_j.$$

From the construction, it is clear that this has the three properties in the statement of the theorem. What about flat sections? A ∇-flat section of M is the same thing as a monodromy invariant vector $v \in \mathbb{C}^m$, meaning one with $C^1 v = \cdots = C^r v = v$. This is equivalent to $\Gamma^1 v = \cdots = \Gamma^r v = 0$, and so v also represents a ∇-flat section of M_τ.

The more demanding part of the proof is the uniqueness of L_τ. You will see that the argument is very similar to what we did for the theorem of Fuchs (in Lecture 20). The problem is local, and so we continue to assume that $X = \Delta^n$, with coordinates $x_1, \dots, x_n$, and D defined by $x_1 \cdots x_r = 0$. Suppose that L and L' are two lattices that both have the three properties stated in the theorem. Denote by ∇ and ∇' the logarithmic connections on L and L'. With respect to a trivialization $e_1, \dots, e_m$ for L, we can write

$$\nabla e_i = \sum_{j,k} a^k_{i,j} \frac{dx_k}{x_k} \otimes e_j,$$

where $a^k_{i,j}$ are holomorphic functions on X; we set

$$\omega = \sum_k A^k \frac{dx_k}{x_k},$$

which is an $m \times m$-matrix of logarithmic one-forms. We use primes to denote the corresponding objects for (L', ∇').

By assumption, $(L, \nabla)\big|_U \cong (L', \nabla')\big|_U$. After a short calculation, the isomorphism between the two bundles with connection translates into the existence of an invertible matrix $S \in \mathrm{GL}_m(\mathscr{O}_U)$ such that

$$dS = S\omega - \omega' S.$$

The entries of S are holomorphic functions on $U = X \setminus D$, possibly with essential singularities along D. To prove the uniqueness statement, it is enough to show that $S \in \mathrm{GL}_m(\mathscr{O}_X)$, meaning that the entries of S should extend to holomorphic functions on X. By Hartog's theorem, holomorphic functions extend over subsets of codimension ≥ 2, and so we only need to to prove that the entries of S extend over the generic point of each irreducible component of D. To keep the notation simple, we will check this at points of

$$D_1 \setminus \bigcup_{k \neq 1} D_k,$$

meaning at points where $x_1 = 0$ but $x_2 \cdots x_r \neq 0$. Write

$$\omega = A^1 \frac{dx_1}{x_1} + \sum_{k \geq 2} A^k \frac{dx_k}{x_k}$$
$$\omega' = A'^1 \frac{dx_1}{x_1} + \sum_{k \geq 2} A'^k \frac{dx_k}{x_k}$$

The relation $dS = S\omega - \omega' S$ gives

$$x_1 \frac{\partial S}{\partial x_1} = SA^1 - A'^1 S, \tag{25.7}$$

and after taking the matrix norm of both sides, we obtain

$$|x_1| \cdot \left\| \frac{\partial S}{\partial x_1} \right\| \leq C \cdot \|S\|,$$

where $C > 0$ is a constant that depends on the size of the (holomorphic) entries of the two matrices A^1 and A'^1. As in Lecture 20, we can now apply Grönwall's inequality to deduce that the entries of S have moderate growth near x_1, hence are meromorphic functions on the set where $x_2 \cdots x_r \neq 0$.

It remains to show that the entries of S are actually holomorphic functions for $x_2 \cdots x_r \neq 0$. Consider the Laurent expansion

$$S = \sum_{j=p}^{\infty} S_j x_1^j,$$

where $S_p \neq 0$ is the leading term. After substituting this into (25.7), we get

$$\sum_{j=p}^{\infty} j S_j x_1^j = \sum_{j=p}^{\infty} \bigl(S_j A^1 - A'^1 S_j\bigr) x_1^j.$$

The coefficients at x_1^p equate to

$$pS_p = S_p \cdot A^1\big|_{x_1=0} - A'^1\big|_{x_1=0} \cdot S_p = S_p \cdot \operatorname{Res}_{D_1}^{L}(\nabla) - \operatorname{Res}_{D_1}^{L'}(\nabla') \cdot S_p.$$

Since both $\operatorname{Res}_{D_1}^{L}(\nabla)$ and $\operatorname{Res}_{D_1}^{L'}(\nabla')$ have their eigenvalues contained in the set $\tau(\mathbb{C}/\mathbb{Z})$, this relation forces $p = 0$. Indeed, suppose that v is a nontrivial eigenvector for $\operatorname{Res}_{D_1}^{L}(\nabla)$, with eigenvalue λ. Then

$$p(S_p v) = \lambda(S_p v) - \operatorname{Res}_{D_1}^{L'}(\nabla')(S_p v),$$

and so $S_p v$ is an eigenvector for $\operatorname{Res}_{D_1}^{L'}(\nabla')$, with eigenvalue $\lambda - p$. (Since S is invertible, we must have $S_p v \neq 0$). As the difference of the two eigenvalues is an integer, this can only happen for $p = 0$. The conclusion is that S extends holomorphically to all of X, proving the desired uniqueness.

Deligne's Riemann-Hilbert correspondence. We are now ready for Deligne's version of the Riemann-Hilbert correspondence. Let $\mathrm{Loc}(X \setminus D)$ denote the category of locally constant sheaves (of finite-dimensional $\mathbb{C}$-vector spaces) on $X \setminus D$.

Theorem 25.8. *Let X be a complex manifold, and $D \subseteq X$ a divisor with simple normal crossing singularities. Then the restriction functor*

$$\mathrm{Conn}(X, D)^{\mathrm{reg}} \to \mathrm{Loc}(X \setminus D)$$

is an equivalence of categories.

Here we associate to a meromorphic connection $(M, \nabla) \in \mathrm{Conn}(X, D)$ the locally constant sheaf of ∇-flat sections of $M\big|_U$, where $U = X \setminus D$. The proof is very easy at this point. First, every locally constant sheaf on $X \setminus D$ is the sheaf of ∇-flat sections of some $(M, \nabla) \in \mathrm{Conn}(U)$. By Theorem 25.6, there is an extension of (M, ∇) to a regular meromorphic connection on (X, D): for any choice of τ, the pair (M_τ, ∇) will do. This shows that the restriction functor is essentially surjective.

It remains to prove that it is also fully faithful. The functor of ∇-flat sections gives an equivalence of categories between $\mathrm{Conn}(U)$ and $\mathrm{Loc}(U)$, and so it suffices to prove that $\mathrm{Conn}(X, D)^{\mathrm{reg}} \to \mathrm{Conn}(U)$ is fully faithful. Let (M, ∇) and (N, ∇) be meromorphic connections, and set $H = \mathcal{H}om_{\mathscr{O}_X(*D)}(M, N)$; recall that (H, ∇) is again a meromorphic connection. As we saw earlier, we have an isomorphism

$$\mathrm{Hom}_{\mathrm{Conn}(X,D)}\Big((M, \nabla), (N, \nabla)\Big) \cong \Gamma(X, H)^{\nabla}$$

between the set of morphisms in the category $\mathrm{Conn}(X, D)$ and the set of ∇-flat sections of H. Similarly,

$$\mathrm{Hom}_{\mathrm{Conn}(U)}\Big((M, \nabla)\big|_U, (N, \nabla)\big|_U\Big) \cong \Gamma(U, H)^{\nabla},$$

and so the problem reduces to showing that

$$\Gamma(X, H)^{\nabla} \to \Gamma(U, H)^{\nabla}$$

is an isomorphism.

Lemma 25.9. *Let $(M, \nabla) \in \mathrm{Conn}(X, D)$ be a regular meromorphic connection. Then the restriction morphism*

$$\Gamma(X, M)^{\nabla} \to \Gamma(U, M)^{\nabla}$$

is an isomorphism, where $U = X \setminus D$.

Proof. Since (M, ∇) is regular, there is a lattice L with $M \cong \mathscr{O}_X(*D) \otimes_{\mathscr{O}_X} L$, such that ∇ has logarithmic poles. Pick any section $\tau \colon \mathbb{C}/\mathbb{Z} \to \mathbb{C}$, for example with $\mathrm{Re}\, \tau \in [0, 1)$. By Theorem 25.6, there exists L_τ with $(L, \nabla)\big|_U \cong (L_\tau, \nabla)\big|_U$. Arguing as in the proof of Theorem 25.6, we find that the isomorphism is locally given by a matrix with meromorphic entries, and hence that (M, ∇) is isomorphic to (M_τ, ∇) as a meromorphic connection. Now the assertion about flat sections follows from the last sentence of Theorem 25.6. □

Deligne's Riemann-Hilbert correspondence again leads to an interesting algebraicity result. Suppose that X is a nonsingular proper variety. Then every locally constant sheaf on $X \setminus D$ comes from a meromorphic connection on (X, D), and hence (by a version of Serre's GAGA theorem) from an algebraic object. Since we have resolution of singularities, we can write every nonsingular algebraic variety in the form $X \setminus D$. Thus every locally constant sheaf on a nonsingular algebraic variety comes from an algebraic vector bundle with integrable connection.

Exercises.

Exercise 25.1. Let (M, ∇) and (N, ∇) be meromorphic connections. Check that $\bigl(\mathcal{H}om_{\mathscr{O}_X(*D)}(M, N), \nabla\bigr)$ is a meromorphic connection, and that $\varphi\colon (M, \nabla) \to (N, \nabla)$ is a morphism of meromorphic connections if and only if, when viewed as a global section of $\mathcal{H}om_{\mathscr{O}_X(*D)}(M, N)$, it satisfies $\nabla\varphi = 0$.

Exercise 25.2. Let $C \in \mathrm{GL}_m(\mathbb{C})$. Show that there is a unique $\Gamma \in \mathrm{Mat}_{m\times m}(\mathbb{C})$ such that $e^{2\pi i\,\Gamma} = C$ and such that the eigenvalues of Γ lie in the set $\tau(\mathbb{C}/\mathbb{Z})$.

Lecture 26: May 10

One-forms on varieties of general type. In the final two lectures, I am going to show you an application of $\mathscr{D}$-module theory to a problem in algebraic geometry. It has to do with holomorphic one-forms and their zero loci. Recall that on a smooth projective curve of genus $g \geq 1$, every holomorphic one-form has exactly $2g-2$ zeros, counted with multiplicity. The situation for surfaces is less clear, but one can still show that every holomorphic one-form on a surface of general type must have a non-empty zero locus. (We'll see a proof of this fact in a second.) This lead Christopher Hacon and Sándor Kovács (and, independently, Tie Luo and Qi Zhang) to conjecture that the same result should hold on any variety of general type; they also proved their conjecture for threefolds. A few years ago, Mihnea Popa and I used $\mathscr{D}$-modules to prove the conjecture in general. The proof I am going to present is a simplified version of our original argument that Chuanhao Wei and I found sometime afterwards.

Theorem 26.1. *Let X be a smooth projective variety over the complex numbers. If X is of general type, then every holomorphic one-form on X has a non-empty zero locus.*

To be precise, for any $\omega \in H^0(X, \Omega_X^1)$, we define the zero locus to be

$$Z(\omega) = \left\{\, x \in X \;\middle|\; \omega(T_x X) = 0 \,\right\}.$$

Then the theorem is claiming that if X is of general type, in the sense that $\dim H^0(X, \omega_X^m)$ grows like a constant times $m^{\dim X}$, then necessarily $Z(\omega) \neq \emptyset$ for every $\omega \in H^0(X, \Omega_X^1)$. Another motivation for thinking that this might be true is that one-forms are dual to vector fields, and zero loci of vector fields are of course related to the topology of X. (For example, if X admits an everywhere nonzero vector field, then its Euler characteristic must be zero.)

Example 26.2. Let us consider the case of surfaces. Suppose that X is a smooth projective surface of general type. Suppose that there was a holomorphic one-form $\omega \in H^0(X, \Omega_X^1)$ with empty zero locus. We will use some of the many numerical identities for surfaces to produce a contradiction.

First, we observe that X must be minimal. Otherwise, X would be the blowup of a smooth projective surface Y at some point, and since $H^0(X, \Omega_X^1) \cong H^0(Y, \Omega_Y^1)$, the one-form ω would be the pullback of a one-form from Y. But then ω has to vanish at some point of the exceptional divisor, contradiction. Now the fact that X is a of general type means that $c_1(X)^2 \geq 1$; together with the Bogomolov-Miyaoka-Yau inequality, we get

$$3c_2(X) \geq c_1(X)^2 \geq 1.$$

But $c_2(X)$ is the topological Euler characteristic of X, and so $e(X) \neq 0$.

Now the contradiction comes from the fact that a surface with a nowhere vanishing holomorphic one-form must have $e(X) = 0$. To see this, consider the complex

$$0 \longrightarrow \mathscr{O}_X \xrightarrow{\omega} \Omega_X^1 \xrightarrow{\omega} \Omega_X^2 \longrightarrow 0$$

where the differential is wedge product with ω. This is a Koszul complex, and since $Z(\omega) = \emptyset$, the complex is exact, and so its hypercohomology is trivial. The hypercohomology spectral sequence

$$E_1^{p,q} = H^q(X, \Omega_X^p)$$

therefore converges to zero. This gives

$$e(X) = \sum_{p,q} (-1)^{p+q} \dim H^q(X, \Omega_X^p) = \sum_{p,q} (-1)^{p+q} \dim E_1^{p,q}$$
$$= \sum_{p,q} (-1)^{p+q} \dim E_\infty^{p,q} = 0,$$

since the alternating sum of the dimensions is preserved under taking cohomology.

Let us make a few general observations about Theorem 26.1. The condition that X is of general type can be restated as follows: for any ample line bundle L on X, there is some $m \geq 1$ such that $\omega_X^m \otimes L^{-1}$ has a section.

Example 26.3. In the special case $m = 1$, we can use the Nakano vanishing theorem to give a simple proof of Theorem 26.1. Suppose that $H^0(X, \omega_X \otimes L^{-1}) \neq 0$, and that there is a holomorphic one-form $\omega \in H^0(X, \Omega_X^1)$ with $Z(\omega) = \emptyset$. Let $n = \dim X$. As before, the complex

$$0 \longrightarrow \mathscr{O}_X \xrightarrow{\omega} \Omega_X^1 \xrightarrow{\omega} \cdots \xrightarrow{\omega} \Omega_X^n \longrightarrow 0$$

is exact, and so the hypercohomology spectral sequence

$$E_1^{p,q} = H^q(X, \Omega_X^p \otimes L^{-1})$$

converges to zero. Since L is ample, the Nakano vanishing theorem tells us that $E_1^{p,q} = 0$ for $p + q < n$. In particular, all the differentials going into the term in the position $(n, 0)$ vanish. But then

$$E_\infty^{n,0} = E_1^{n,0} = H^0(X, \omega_X \otimes L^{-1}) \neq 0,$$

which is a contradiction. Unfortunately, this simple argument totally breaks down once $m \geq 2$. But we will see that it is still basically a vanishing theorem that is responsible for Theorem 26.1.

Another observation is that holomorphic one-forms are closely related to abelian varieties. Indeed, we always have the Albanese mapping

$$\mathrm{alb}\colon X \to \mathrm{Alb}(X) = H^0(X, \Omega_X^1)^* / H_1(X, \mathbb{Z})$$

to an abelian variety of dimension $h^0(X, \Omega_X^1)$, and by construction,

$$H^0(X, \Omega_X^1) \cong H^0\big(\mathrm{Alb}(X), \Omega_{\mathrm{Alb}(X)}^1\big).$$

It thus makes sense to consider more generally an arbitrary morphism $f\colon X \to A$ to an abelian variety A, and to ask about the zero loci of the holomorphic one-forms $f^*\omega$, for $\omega \in H^0(A, \Omega_A^1)$. Of course, we should replace the assumption "X of general type" by the condition that $\omega_X^m \otimes f^*L^{-1}$ has sections for $m \gg 1$, where L is an ample line bundle on A. This suggests the following more general result.

Theorem 26.4. *Let $f\colon X \to A$ be a morphism from a smooth projective variety to an abelian variety. If $H^0(X, \omega_X^m \otimes f^*L^{-1}) \neq 0$ for some $m \geq 1$ and some ample line bundle L on A, then one has $Z(f^*\omega) \neq \emptyset$ for every $\omega \in H^0(A, \Omega_A^1)$.*

Set $W = H^0(A, \Omega_A^1)$, and consider the incidence variety

$$Z_f = \big\{ (x, \omega) \in X \times W \bigm| x \in Z(f^*\omega) \big\} \subseteq X \times W.$$

The theorem is claiming that the second projection $p_2\colon Z_f \to W$ is surjective. Since A is an abelian variety, we have $T^*A = A \times W$, and so the usual diagram of morphisms between cotangent bundles becomes:

$$\begin{array}{ccc} X \times W & \xrightarrow{df} & T^*X \\ \Big\downarrow{\scriptstyle f \times \mathrm{id}} & & \\ A \times W & & \end{array}$$

With this notation, we have $Z_f = df^{-1}(0)$. When we looked at direct images for $\mathscr{D}$-modules (in Lecture 13), we encountered the set

$$S_f = (f \times \mathrm{id})\bigl(df^{-1}(0)\bigr) = (f \times \mathrm{id})(Z_f).$$

It contains the characteristic varieties of the direct image $\mathscr{D}$-modules $\mathcal{H}^j f_+ \omega_X$. (In Lecture 13, we proved this for closed embeddings.) Concretely,

$$S_f = \bigl\{ (a, \omega) \in A \times W \bigm| f^{-1}(a) \cap Z(f^* \omega) \neq \emptyset \bigr\},$$

and so $Z(f^*\omega) \neq \emptyset$ for every $\omega \in W$ is equivalent to the surjectivity of $p_2 \colon S_f \to W$. This suggests the following strategy for proving Theorem 26.4: find a $\mathscr{D}_A$-module whose characteristic variety $\mathrm{Ch}(\mathcal{M})$ is contained in the set S_f, and then use results about $\mathscr{D}$-modules to show that $p_2 \colon \mathrm{Ch}(\mathcal{M}) \to W$ must be onto.

We could not actually get this idea to work, but we found a good replacement for it, based on work of Viehweg and Zuo. Here is a rought outline for the proof of Theorem 26.4. On the cotangent bundle $T^*A = A \times W$, we construct a morphism $\mathscr{F} \to \mathscr{G}$ between two coherent sheaves, with the following three properties:

(a) The support of $\mathscr{F}$ is contained in the set S_f.
(b) The induced morphism $H^0(A \times W, \mathscr{F}) \to H^0(A \times W, \mathscr{G})$ is nontrivial.
(c) The coherent sheaf $(p_2)_* \mathscr{G}$ on W is torsion-free.

Here $p_1 \colon A \times W \to A$ and $p_2 \colon A \times W \to W$ are the two projections. We will see next time that $\mathscr{G}$ is (almost) the coherent sheaf coming from a $\mathscr{D}_A$-module $\mathcal{M}$ with a good filtration $F_\bullet \mathcal{M}$.

Lemma 26.5. *Such a morphism $\mathscr{F} \to \mathscr{G}$ can only exist if $p_2(S_f) = W$.*

Proof. Consider the induced morphism

$$(p_2)_* \mathscr{F} \to (p_2)_* \mathscr{G}.$$

Both sheaves are coherent (by properness of p_2), and the support of $(p_2)_* \mathscr{F}$ is contained in the set $p_2(S_f)$. Now suppose that $p_2(S_f) \neq W$. Then $(p_2)_* \mathscr{F}$ is a torsion sheaf, and so the morphism to the torsion-free sheaf $(p_2)_* \mathscr{G}$ must be trivial. Taking global sections, we find that

$$H^0(A \times W, \mathscr{F}) = H^0\bigl(W, (p_2)_* \mathscr{F}\bigr) \to H^0\bigl(W, (p_2)_* \mathscr{G}\bigr) = H^0(A \times W, \mathscr{G})$$

is trivial; but this is a contradiction. □

Filtered $\mathscr{D}$-modules and the Rees construction. For the proof of Theorem 26.4, it is important to work with pairs $(\mathcal{M}, F_\bullet \mathcal{M})$, where $\mathcal{M}$ is a coherent $\mathscr{D}$-module, and $F_\bullet \mathcal{M}$ a good filtration. Here the filtration is not just a tool to study $\mathscr{D}$-modules, but an essential piece of data. One can define the direct image and duality functors for filtered $\mathscr{D}$-modules by analogy with the unfiltered case, as follows.

Let X be a nonsingular algebraic variety over a field k (of characteristic zero). We can combine $\mathscr{D}_X$ with its order filtration $F_\bullet \mathscr{D}_X$ into a single sheaf of algebras

$$\tilde{\mathscr{D}}_X = \bigoplus_{k=0}^{\infty} F_k \mathscr{D}_X,$$

called the *Rees algebra* of $\mathscr{D}_X$. This is a sheaf of non-commutative graded algebras, with multiplication defined in the obvious way. We denote by $z \in \tilde{\mathscr{D}}_{X,1}$ the image of $1 \in F_1 \mathscr{D}_X$; then $\tilde{\mathscr{D}}_X$ contains a copy of $\mathscr{O}_X[z]$. It is easy to see that

$$\tilde{\mathscr{D}}_X / \tilde{\mathscr{D}}_X (z - z_0) \cong \mathscr{D}_X$$

for every $z_0 \neq 0$, because in the quotient, each $F_k \mathscr{D}_X$ gets identified with its image in $F_{k+1} \mathscr{D}_X$. Likewise,

$$\tilde{\mathscr{D}}_X / \tilde{\mathscr{D}}_X z \cong \operatorname{gr}^F \mathscr{D}_X,$$

because in the quotient, the image of $F_k\mathscr{D}_X$ in $F_{k+1}\mathscr{D}_X$ goes to zero. We can therefore think of the Rees algebra $\tilde{\mathscr{D}}_X$ as a family of algebras over the affine line $\operatorname{Spec} k[z]$, in which $\mathscr{D}_X$ deforms into $\operatorname{gr}^F\mathscr{D}_X$.

Given a coherent left (or right) $\mathscr{D}_X$-module $\mathcal{M}$ and a good filtration $F_\bullet\mathcal{M}$, we can form the *Rees module*

$$\tilde{\mathcal{M}} = R_F\mathcal{M} = \bigoplus_{k\in\mathbb{Z}} F_k\mathcal{M}.$$

This is a graded left (or right) module over $\tilde{\mathscr{D}}_X$ in the obvious way; since the filtration is good, $\tilde{\mathcal{M}}$ is coherent over $\tilde{\mathscr{D}}_X$. As before, one checks that

$$\tilde{\mathcal{M}}/(z-z_0)\tilde{\mathcal{M}} \cong \mathcal{M}$$

for every $z_0 \neq 0$, whereas

$$\tilde{\mathcal{M}}/z\tilde{\mathcal{M}} \cong \operatorname{gr}^F\mathcal{M}.$$

An important point is that not every graded $\tilde{\mathscr{D}}_X$-module comes from a filtered $\mathscr{D}_X$-module.

Lemma 26.6. *A graded $\tilde{\mathscr{D}}_X$-module $\tilde{\mathcal{M}}$ is the Rees module of a filtered $\mathscr{D}_X$-module if and only if it has no z-torsion.*

Graded $\tilde{\mathscr{D}}_X$-modules without z-torsion are called *strict*. Since $\operatorname{Spec} k[z]$ is one-dimensional, this condition is equivalent to flatness over $k[z]$.

Proof. It is easy to see that a graded $\tilde{\mathscr{D}}_X$-module of the form $R_F\mathcal{M}$ does not have any z-torsion. Let us prove the converse. Suppose for the time being that $\tilde{\mathcal{M}}$ is any graded left $\tilde{\mathscr{D}}_X$-module. Define

$$\mathcal{M} = \tilde{\mathcal{M}}/(z-1)\tilde{\mathcal{M}},$$

which is a left module over $\tilde{\mathscr{D}}_X/\tilde{\mathscr{D}}_X(z-1) \cong \mathscr{D}_X$. The image of the k-th graded piece $\tilde{\mathcal{M}}_k$ defines a subsheaf $F_k\mathcal{M} \subseteq \mathcal{M}$, with the property that $F_j\mathscr{D}_X \cdot F_k\mathcal{M} \subseteq F_{j+k}\mathcal{M}$. It follows that the Rees module $R_F\mathcal{M}$ is a graded $\tilde{\mathscr{D}}_X$-module without z-torsion.

Now we have a morphism of graded $\tilde{\mathscr{D}}_X$-modules

$$\varphi\colon \tilde{\mathcal{M}} \to R_F\mathcal{M},$$

that takes $\tilde{\mathcal{M}}_k$ to $F_k\mathcal{M}$; by construction, this morphism is surjective. One checks that $\ker\varphi$ consists exactly of those sections of $\tilde{\mathcal{M}}$ that are killed by some power of z. In particular, φ is an isomorphism whenever $\tilde{\mathcal{M}}$ does not have any z-torsion. $\square$

Functors for Rees modules. One can define all the usual functors for $\mathscr{D}$-modules also for modules over the larger algebra $\tilde{\mathscr{D}}$. The two functor we need are the direct image functor and the duality functor. Given a morphism $f\colon X \to Y$, we define the *transfer module*

$$\tilde{\mathscr{D}}_{X\to Y} = \mathscr{O}_X \otimes_{f^{-1}\mathscr{O}_Y} f^{-1}\tilde{\mathscr{D}}_Y$$

by the same formula as for $\mathscr{D}$-modules. It is again a $(\tilde{\mathscr{D}}_X, f^{-1}\tilde{\mathscr{D}}_Y)$-bimodule. We can then define the direct image functor

$$f_+(-) = \mathbf{R}f_*\big(- \overset{\mathbf{L}}{\otimes}_{\tilde{\mathscr{D}}_X} \tilde{\mathscr{D}}_{X\to Y}\big)\colon D^b_{g,qc}(\tilde{\mathscr{D}}_X^{op}) \to D^b_{g,qc}(\tilde{\mathscr{D}}_Y^{op})$$

between the derived categories of quasi-coherent graded right $\tilde{\mathscr{D}}$-modules. As in the case of $\mathscr{D}$-modules, one can use induced $\tilde{\mathscr{D}}$-modules to show that the direct image by a proper morphism preserves coherence.

If we specialize to $z = 1$, for example, by taking the (derived) tensor product with $\tilde{\mathscr{D}}/\tilde{\mathscr{D}}(z-1)$, we recover the usual direct image functor for right $\mathscr{D}$-modules.

On the other hand, we can specialize to $z = 0$, by taking the (derived) tensor product with $\tilde{\mathscr{D}}/\tilde{\mathscr{D}}z$. This gives us a functor

$$\operatorname{gr}^F \colon D^b_{g,qc}(\tilde{\mathscr{D}}_X^{op}) \to D^b_{g,qc}(\operatorname{gr}^F \mathscr{D}_X),$$

with takes a Rees module of the form $R_F\mathcal{M}$ to the associated graded module $\operatorname{gr}^F \mathcal{M}$. By computing what happens to the transfer module, one checks that the following diagram is commutative:

$$\begin{array}{ccc} D^b_{g,qc}(\tilde{\mathscr{D}}_X^{op}) & \xrightarrow{f_+} & D^b_{g,qc}(\tilde{\mathscr{D}}_Y^{op}) \\ \big\downarrow{\scriptstyle \operatorname{gr}^F} & & \big\downarrow{\scriptstyle \operatorname{gr}^F} \\ D^b_{g,qc}(\operatorname{gr}^F \mathscr{D}_X) & \longrightarrow & D^b_{g,qc}(\operatorname{gr}^F \mathscr{D}_Y) \end{array}$$

Here the arrow on the bottom is the functor

$$\mathbf{R}f_*\bigl(- \overset{\mathbf{L}}{\otimes}_{\operatorname{gr}^F \mathscr{D}_X} f^*(\operatorname{gr}^F \mathscr{D}_Y)\bigr) \colon D^b_{g,qc}(\operatorname{gr}^F \mathscr{D}_X) \to D^b_{g,qc}(\operatorname{gr}^F \mathscr{D}_Y).$$

If we forget about the grading, then quasi-coherent sheaves of $\operatorname{gr}^F \mathscr{D}_X$-modules are the same thing as quasi-coherent sheaves of $\mathscr{O}_{T^*X}$-modules on the cotangent bundle. The geometric interpretation of the above functor is then

$$\mathbf{R}(p_2)_* \circ \mathbf{L}(df)^* \colon D^b_{qc}(\mathscr{O}_{T^*X}) \to D^b_{qc}(\mathscr{O}_{T^*Y}),$$

where the morphisms between cotangent bundles are as in the diagram below.

$$\begin{array}{ccc} X \times_Y T^*Y & \xrightarrow{df} & T^*X \\ \big\downarrow{\scriptstyle p_2} & & \\ T^*Y & & \end{array}$$

The direct image functor for Rees modules therefore interpolates between the usual direct image functor for $\mathscr{D}$-modules, and the natural functor on the level of cotangent bundles. One subtle point is that even if we start from a Rees module $R_F\mathcal{M}$, the direct image

$$f_+(R_F\mathcal{M}) \in D^b_{g,qc}(\tilde{\mathscr{D}}_Y^{op})$$

might have z-torsion (= not be strict). If that happens, it means that $f_+(R_F\mathcal{M})$ has more cohomology that the complex of right $\mathscr{D}_Y$-modules $f_+\mathcal{M}$. (The extra cohomology is z-torsion, of course.) Equivalently, it means that the complex of graded $\operatorname{gr}^F \mathscr{D}_Y$-modules

$$\mathbf{R}f_*\bigl(\operatorname{gr}^F \mathcal{M} \overset{\mathbf{L}}{\otimes}_{\operatorname{gr}^F \mathscr{D}_X} f^*(\operatorname{gr}^F \mathscr{D}_Y)\bigr)$$

has some additional cohomology that is not visible to the direct image $f_+\mathcal{M}$ of the underlying $\mathscr{D}$-module.

One can also define a duality functor for $\tilde{\mathscr{D}}$-modules. As with $\mathscr{D}$-modules, the tensor product $\omega_X \otimes_{\mathscr{O}_X} \tilde{\mathscr{D}}_X$ has two commuting structures of right $\tilde{\mathscr{D}}_X$-modules. If $\tilde{\mathcal{M}}$ is a right $\tilde{\mathscr{D}}_X$-module, then

$$\mathcal{H}om_{\tilde{\mathscr{D}}_X}(\tilde{\mathcal{M}}, \omega_X \otimes_{\mathscr{O}_X} \tilde{\mathscr{D}}_X)$$

still has the structure of a right $\tilde{\mathscr{D}}_X$-module. Passing to derived categories, we obtain the (contravariant) duality functor

$$\mathbb{D}_X = \mathbf{R}\mathcal{H}om_{\tilde{\mathscr{D}}_X}(-, \omega_X \otimes_{\mathscr{O}_X} \tilde{\mathscr{D}}_X)[n] \colon D^b_{g,qc}(\tilde{\mathscr{D}}_X^{op}) \to D^b_{g,qc}(\tilde{\mathscr{D}}_X^{op})^{op}.$$

Here $[n]$ means shifting to the left by $n = \dim X$ steps. If we specialize to $z = 1$, we recover the usual duality functor for $\mathscr{D}_X$-modules; if we specialize instead to $z = 0$, we obtain the functor

$$\mathbf{R}\mathcal{H}om_{\operatorname{gr}^F \mathscr{D}_X}(-, \omega_X \otimes_{\mathscr{O}_X} \operatorname{gr}^F \mathscr{D}_X)[n] \colon D^b_{g,qc}(\operatorname{gr}^F \mathscr{D}_X) \to D^b_{g,qc}(\operatorname{gr}^F \mathscr{D}_X)^{op}.$$

We can again express this in geometric terms: if $\mathscr{G}$ denotes the coherent sheaf on T^*X corresponding to $\operatorname{gr}^F \mathcal{M}$, then the above functor is

$$\mathbf{R}\mathcal{H}om_{\mathscr{O}_{T^*X}}(\mathscr{G}, p^*\omega_X)[n],$$

where $p\colon T^*X \to X$ is the projection. As before, $\mathbb{D}_X(R_F\mathcal{M})$ can acquire z-torsion. For instance, suppose that $\mathcal{M}$ is a holonomic right $\mathscr{D}_X$-module. Then

$$\mathbf{R}\mathcal{H}om_{\mathscr{D}_X}(\mathcal{M}, \omega_X \otimes_{\mathscr{O}_X} \mathscr{D}_X)[n]$$

only has cohomology in degree zero (where we get the holonomic dual $\mathcal{M}^*$). But the complex $\mathbb{D}_X(R_F\mathcal{M})$ might have cohomology in other degrees as well (which will then be z-torsion). In fact, one can show that $\mathbb{D}_X(R_F\mathcal{M})$ is again strict if and only if the complex

$$\mathbf{R}\mathcal{H}om_{\operatorname{gr}^F \mathscr{D}_X}(\operatorname{gr}^F \mathcal{M}, \omega_X \otimes_{\mathscr{O}_X} \operatorname{gr}^F \mathscr{D}_X)[n]$$

only has cohomology in degree zero; in commutative algebra terminology, this is equivalent to $\operatorname{gr}^F \mathcal{M}$ being a Cohen-Macaulay module over $\operatorname{gr}^F \mathscr{D}_X$.

Hodge modules. You can think of *Hodge modules* as being a special class of filtered $\mathscr{D}$-modules that behave well under the various functors. More precisely, a Hodge module on a nonsingular algebraic variety X is a (regular holonomic) right $\mathscr{D}_X$-module $\mathcal{M}$ together with a good filtration $F_\bullet\mathcal{M}$. There is some extra data, too, and several very restrictive conditions have to be satisfied, which make sure that the pair $(\mathcal{M}, F_\bullet\mathcal{M})$ comes from a polarizable variation of Hodge structure.

Example 26.7. The pair $(\omega_X, F_\bullet\omega_X)$, with the filtration defined by $F_{-n-1}\omega_X = 0$ and $F_{-n}\omega_X = \omega_X$, is an example of a Hodge module. That this is so is a deep theorem by Morihiko Saito, who created this theory.

For our purposes, the following three facts are important. (Again, all three are difficult theorems due to Saito.) First, if $(\mathcal{M}, F_\bullet\mathcal{M})$ is a Hodge module on X, and if $f\colon X \to Y$ is a *proper* morphism between nonsingular algebraic varieties, then all cohomology modules of the complex $f_+(R_F\mathcal{M})$ are strict, and the resulting filtered $\mathscr{D}_Y$-modules are again Hodge modules on Y. In particular, we can compute their associated graded modules:

$$\operatorname{gr}^F \mathcal{H}^j f_+\mathcal{M} \cong R^j f_*\big(\operatorname{gr}^F \mathcal{M} \overset{\mathbf{L}}{\otimes}_{\operatorname{gr}^F \mathscr{D}_X} f^*(\operatorname{gr}^F \mathscr{D}_Y)\big).$$

Second, the duality functor preserves Hodge modules: the complex $\mathbb{D}_X(R_F\mathcal{M})$ only has cohomology in degree zero, which is strict, and the resulting filtered $\mathscr{D}_X$-module $(\mathcal{M}', F_\bullet\mathcal{M}')$ is again a Hodge module on X. Once again, this means that we can compute the associated graded module:

$$\operatorname{gr}^F \mathcal{M}' \cong R^n\mathcal{H}om_{\operatorname{gr}^F \mathscr{D}_X}(\operatorname{gr}^F \mathcal{M}, \omega_X \otimes_{\mathscr{O}_X} \operatorname{gr}^F \mathscr{D}_X).$$

Third, Hodge modules on projective varieties satisfy a vanishing theorem similar to the Kodaira vanishing theorem. Given a Hodge module $(\mathcal{M}, F_\bullet\mathcal{M})$, we can form the Spencer complex

$$\operatorname{Sp}(\mathcal{M}) = \Big[\mathcal{M} \otimes \bigwedge^n \mathscr{T}_X \to \cdots \to \mathcal{M} \otimes \mathscr{T}_X \to \mathcal{M}\Big]$$

which lives in degrees $-n, \dots, 0$. (Since $\mathcal{M}$ is regular holonomic, $\operatorname{Sp}(\mathcal{M})$ is actually a perverse sheaf, by Kashiwara's theorem.) The Spencer complex is filtered by the family of subcomplexes

$$F_k \operatorname{Sp}(\mathcal{M}) = \Big[F_{k-n}\mathcal{M} \otimes \bigwedge^n \mathscr{T}_X \to \cdots \to F_{k-1}\mathcal{M} \otimes \mathscr{T}_X \to F_k\mathcal{M}\Big],$$

and the k-th subquotient

$$\mathrm{gr}_k^F \mathrm{Sp}(\mathcal{M}) = \Big[\mathrm{gr}_{k-n}^F \mathcal{M} \otimes \bigwedge^n \mathscr{T}_X \to \dots \to \mathrm{gr}_{k-1}^F \mathcal{M} \otimes \mathscr{T}_X \to \mathrm{gr}_k^F \mathcal{M}\Big]$$

is a complex of coherent $\mathscr{O}_X$-modules. For example, for the pair $(\omega_X, F_\bullet \omega_X)$, the Spencer complex is the holomorphic de Rham complex, and the $(-p)$-th subquotient is Ω_X^p, placed in degree $n-p$.

Theorem 26.8 (Saito's vanishing theorem)**.** *Let X be a nonsingular projective variety, and L an ample line bundle. If $(\mathcal{M}, F_\bullet \mathcal{M})$ is a Hodge module on X, then*

$$H^i\big(X, \mathrm{gr}_k^F \mathrm{Sp}(\mathcal{M}) \otimes L\big) = 0 \quad \textit{for every } i > 0,$$
$$H^i\big(X, \mathrm{gr}_k^F \mathrm{Sp}(\mathcal{M}) \otimes L^{-1}\big) = 0 \quad \textit{for every } i < 0.$$

Hodge modules on abelian varieties. Let us now return to abelian varieties. Suppose that A is an abelian variety and L an ample line bundle on A. Since the tangent bundle of A is trivial, one can prove a much stronger vanishing theorem. Let me explain how this works. Fix a Hodge module $(\mathcal{M}, F_\bullet \mathcal{M})$ on A, and for simplicity, suppose that $F_{-1}\mathcal{M} = 0$ and $F_0 \mathcal{M} \neq 0$. Then

$$\mathrm{gr}_0^F \mathrm{Sp}(\mathcal{M}) = \mathrm{gr}_0^F \mathcal{M},$$

and so Saito's vanishing theorem gives

$$H^i(A, \mathrm{gr}_0^F \mathcal{M} \otimes L) = 0 \quad \text{for all } i > 0. \tag{26.9}$$

The next subquotient of the Spencer complex is

$$\mathrm{gr}_1^F \mathrm{Sp}(\mathcal{M}) = \Big[\mathrm{gr}_0^F \mathcal{M} \otimes \mathscr{T}_A \to \mathrm{gr}_1^F \mathcal{M}\Big].$$

Since $\mathscr{T}_A \cong \mathscr{O}_A^{\oplus g}$, where $g = \dim A$, the term $\mathrm{gr}_0^F \mathcal{M} \otimes \mathscr{T}_A$ has no higher cohomology by (26.9). On the other hand, the vanishing theorem says that

$$H^i\big(A, \mathrm{gr}_1^F \mathrm{Sp}(\mathcal{M}) \otimes L\big) = 0 \quad \text{for all } i > 0.$$

If we put these two facts together, we find that

$$H^i(A, \mathrm{gr}_1^F \mathcal{M} \otimes L) = 0 \quad \text{for all } i > 0. \tag{26.10}$$

Continuing in this manner, we arrive at the conclusion that

$$H^i(A, \mathrm{gr}_k^F \mathcal{M} \otimes L) = 0 \quad \text{for all } i > 0, \tag{26.11}$$

and so all graded quotients $\mathrm{gr}_k^F \mathcal{M}$ satisfy the same Kodaira-type vanishing theorem.

Now recall that $T^*A = A \times W$, where $W = H^0(A, \Omega_A^1)$. The vanishing theorem can be used to produce torsion-free sheaves on W. Suppose that $(\mathcal{M}, F_\bullet \mathcal{M})$ is a Hodge module on A. Denote by $\mathscr{G}$ the coherent sheaf on the cotangent bundle corresponding to the associated graded module $\mathrm{gr}^F \mathcal{M}$. Also let $p_1 \colon A \times W \to A$ and $p_2 \colon A \times W \to W$ be the two projections.

Lemma 26.12. *If L is an ample line bundle on A, then $(p_2)_*\big(\mathscr{G} \otimes p_1^* L^{-1}\big)$ is a torsion-free coherent sheaf on W.*

Proof. Coherence is clear (because p_2 is proper). Let us first analyze what happens when we tensor by L instead of L^{-1}. The higher direct images sheaves

$$R^i(p_2)_*(\mathscr{G} \otimes p_1^* L)$$

are coherent, and since W is affine, we have

$$H^0\big(W, R^i(p_2)_*(\mathscr{G} \otimes p_1^* L)\big) = H^i\big(A \times W, \mathscr{G} \otimes p_1^* L\big) = H^i\big(A, (p_1)_*\mathscr{G} \otimes L\big).$$

This vanishes for every $i > 0$ because of (26.11) and the fact that $(p_1)_*\mathscr{G} = \mathrm{gr}^F \mathcal{M}$. The conclusion is that the complex

$$\mathbf{R}(p_2)_*(\mathscr{G} \otimes p_1^* L)$$

is actually a single coherent sheaf in degree zero.

Now let us turn to the sheaf $(p_2)_*(\mathscr{G} \otimes p_1^* L^{-1})$. If we apply Grothendieck duality for the proper morphism p_2, we get

$$\mathbf{R}\mathcal{H}om_{\mathscr{O}_W}\bigl(\mathbf{R}(p_2)_*(\mathscr{G} \otimes p_1^* L^{-1}), \mathscr{O}_W\bigr) \cong \mathbf{R}(p_2)_* \mathbf{R}\mathcal{H}om_{\mathscr{O}_{A\times W}}\bigl(\mathscr{G} \otimes p_1^* L^{-1}, p_1^* \omega_A[g]\bigr),$$

since the relative dualizing sheaf is $\omega_{A\times W/W} = p_1^*\omega_A$. We can rewrite the right-hand side in the more compact form

$$\mathbf{R}(p_2)_*(\mathscr{G}' \otimes p_1^* L),$$

where we have introduced the new complex

$$\mathscr{G}' = \mathbf{R}\mathcal{H}om_{\mathscr{O}_{A\times W}}(\mathscr{G}, p_1^*\omega_A)[g].$$

We can now use the results about the duality functor. They imply that $\mathscr{G}'$ is actually a coherent sheaf; more precisely, we have $\mathbb{D}_X(R_F\mathcal{M}) = R_F\mathcal{M}'$ for a Hodge module $(\mathcal{M}', F_\bullet\mathcal{M}')$, and $\mathscr{G}'$ is the coherent sheaf associated to $\operatorname{gr}^F\mathcal{M}'$. According to the discussion above,

$$\mathbf{R}\mathcal{H}om_{\mathscr{O}_W}\bigl(\mathbf{R}(p_2)_*(\mathscr{G} \otimes p_1^* L^{-1}), \mathscr{O}_W\bigr) \cong (p_2)_*\bigl(\mathscr{G}' \otimes p_1^* L\bigr)$$

is therefore a coherent sheaf in degree zero. After dualizing again, we get

$$(p_2)_*(\mathscr{G} \otimes p_1^* L^{-1}) \cong \mathcal{H}om_{\mathscr{O}_W}\bigl((p_2)_*(\mathscr{G}' \otimes p_1^* L), \mathscr{O}_W\bigr),$$

which is reflexive, hence torsion-free. □

Exercise.

Exercise 26.1. Let $\tilde{\mathcal{M}}$ be a coherent graded left $\tilde{\mathscr{D}}_X$-module. Define $\mathcal{M} = \tilde{\mathcal{M}}/(z-1)\tilde{\mathcal{M}}$, and let $F_k\mathcal{M}$ be the image of $\tilde{\mathcal{M}}_k$.

(a) Show that $F_\bullet\mathcal{M}$ is a good filtration.

(b) Show that the kernel of the morphism $\varphi\colon \tilde{\mathcal{M}} \to R_F\mathcal{M}$ consists exactly of those sections that are killed by some power of z.

Lecture 27: May 17

Today is the last class of the semester. We are going to finish the proof of Theorem 26.4. Let me state the result again.

Theorem. *Let $f\colon X \to A$ be a morphism from a smooth projective variety to an abelian variety. If $H^0(X, \omega_X^m \otimes f^*L^{-1}) \neq 0$ for some $m \geq 1$ and some ample line bundle L on A, then one has $Z(f^*\omega) \neq \emptyset$ for every $\omega \in H^0(A, \Omega_A^1)$.*

Last time, we introduced the set

$$S_f = \big\{ (a, \omega) \in A \times W \bigm| f^{-1}(a) \cap Z(f^*\omega) \neq \emptyset \big\} = (f \times \mathrm{id})\big(df)^{-1}(0)\big),$$

where the notation is as follows:

$$\begin{tikzcd} X \times W \arrow[r, "df"] \arrow[d, "f \times \mathrm{id}"] & T^*X \\ A \times W \end{tikzcd}$$

We also observed that the result about one-forms is equivalent to the surjectivity of $p_2\colon S_f \to W$. Finally, we talked briefly about filtered $\mathscr{D}$-modules and Hodge modules, and we showed that if $(\mathcal{M}, F_\bullet\mathcal{M})$ is a Hodge module on the abelian variety A, and if $\mathscr{G}$ is the coherent sheaf on $T^*A = A \times W$ corresponding to $\mathrm{gr}^F \mathcal{M}$, then for any ample line bundle L,

$$(p_2)_*(\mathscr{G} \otimes p_1^* L^{-1})$$

is a *torsion-free* coherent sheaf on W. This was a consequence of Saito's vanishing theorem, ultimately. Today, I will show you how to construct the required objects from the hypothesis that $\omega_X^m \otimes f^*L^{-1}$ has a section.

Base change. Whenever the m-th power of a line bundle has a section, one can construct a cyclic covering. We can put ourselves in this situation with the help of a very useful small trick. On the abelian variety A, we have the multiplication homomorphism

$$[m]\colon A \to A, \quad a \mapsto \underbrace{a + \cdots + a}_{m \text{ times}},$$

for any $m \in \mathbb{Z}$. It is finite and étale, of degree equal to $m^{2\dim A}$, which is the same as the number of m-torsion points in A. The effect of pulling back by $[m]$ is to make line bundles more divisible. In fact, if L is symmetric, in the sense that $[-1]^*L \cong L$, then one has $[m]^*L \cong L^{m^2}$; if L is anti-symmmetric, in the sense that $[-1]^*L \cong L^{-1}$, then one still has $[m]^*L \cong L^m$. Since we can write any line bundle as the product of a symmetric and an anti-symmetric one, it follows that

$$[2m]^*L \cong L'^m$$

for some other line bundle L'. Now consider the fiber product diagram

$$\begin{tikzcd} X' \arrow[r, "\psi"] \arrow[d, "f'"] & X \arrow[d, "f"] \\ A \arrow[r, "{[2m]}"] & A. \end{tikzcd}$$

Because ψ is finite and étale, we get $\psi^*\omega_X \cong \omega_{X'}$, and therefore

$$\psi^*\big(\omega_X^m \otimes f^*L^{-1}\big) \cong \big(\omega_{X'} \otimes f'^*L'^{-1}\big)^m.$$

Again because ψ is finite and étale, it does not affect the zero loci of holomorphic one-forms; more precisely, we have

$$\psi^{-1}Z(f^*\omega) = Z(f'^*\omega),$$

because $[2m]^*\omega = 2m\cdot\omega$. For the purpose of proving Theorem 26.4, we can therefore safely replace $f\colon X \to A$ by its base change $f'\colon X' \to A$; this allows us to assume that the m-th power of the line bundle $B = \omega_X \otimes f^*L^{-1}$ has a nontrivial global section.

Cyclic coverings. Suppose for a moment that we have a nonsingular algebraic variety X and a line bundle B, as well as a nontrivial global section $s \in H^0(X, B^m)$ for some $m \geq 2$. In that case, one can construct a finite morphism

$$\pi\colon Y \to X$$

with the property that π^*B has a global section s_0 such that $s_0^m = \pi s$. Since the group of m-th roots of unities naturally acts on Y, this is called the *cyclic covering* determined by the section s.

Example 27.1. When B is the trivial bundle, s is just a regular function on X; in that case, Y is the closed subscheme of $X \times \mathbb{A}^1$ defined by the equation $t^m = s$, where t is the coordinate on $\mathbb{A}^1$. Here t serves as the m-th root of s.

The construction in the general case is similar. Let $p\colon V = \mathbb{V}(B) \to X$ be the algebraic line bundle (whose sheaf of sections is the locally free sheaf B). The pullback π^*B has a tautological section $s_0 \in H^0(V, \pi^*B)$, and one defines $Y \subseteq V$ as the closed subscheme cut out by the section $s_0^m - \pi^*s$ of the line bundle π^*B^m. By construction, the morphism $\pi\colon Y \to X$ is finite of degree m, and π^*B has a global section s_0 such that $s_0^m = \pi^*s$. (This construction has a simple universal property, which I will leave to you to formulate and prove.)

Unless the divisor of s is nonsingular, the cyclic covering Y will be singular, but we can resolve its singularities. In this way, we obtain a proper morphism

$$\varphi\colon Z \to X,$$

generically finite of degree m, from a nonsingular algebraic variety Z, such that the line bundle φ^*B has a section $s_0 \in H^0(Z, \varphi^*B)$ with $s_0^m = \varphi^*s$.

Sheaves. If we apply the cyclic covering construction to $B = \omega_X \otimes f^*L^{-1}$, we obtain the following diagram:

$$\begin{tikzcd} Z \arrow[r, "\varphi"] \arrow[dr, bend right, "h"'] & X \arrow[d, "f"] \\ & A \end{tikzcd}$$

Here Z is a nonsingular projective variety of dimension $\dim Z = \dim X = n$, and φ is generically finite of degree m. We may view the resulting nontrivial section of $\varphi^*B = \varphi^*\omega_X \otimes h^*L^{-1}$ as a nontrivial morphism

$$h^*L \to \varphi^*\omega_X. \tag{27.2}$$

We can use the morphism from Z to A to construct a filtered $\mathscr{D}$-module on the abelian variety. The underlying $\mathscr{D}_A$-module is simply the direct image $\mathcal{M} = \mathcal{H}^0 h_+ \omega_Z$. Since $(\omega_Z, F_\bullet \omega_Z)$ is actually a Hodge module on Z, the graded $\tilde{\mathscr{D}}_A$-module $\tilde{\mathcal{M}} = \mathcal{H}^0 h_+(R_F \omega_Z)$ is strict, and so there is a good filtration $F_\bullet \mathcal{M}$ such that $\tilde{\mathcal{M}} = R_F \mathcal{M}$. Moreover, $(\mathcal{M}, F_\bullet \mathcal{M})$ is again a Hodge module on A. If we denote by $\mathscr{G}$ the associated coherent sheaf on $T^* = A \times W$, then we know from last time that

$$(p_2)_*\big(\mathscr{G} \otimes p_1^* L^{-1}\big)$$

is a torsion-free coherent sheaf on W.

Since we constructed $\mathscr{G}$ from the morphism $h\colon Z \to A$, which is more singular than the original morphism $f\colon X \to A$, the support of $\mathscr{G}$ has nothing to do with the set $S_f \subseteq T^*A$ that we are interested in; in fact, one has $\operatorname{Supp} \mathscr{G} \subseteq S_h$, which is much larger in general. But we can use the existence of (27.2) to construct another coherent sheaf $\mathscr{F}$ with $\operatorname{Supp}\mathscr{F} \subseteq S_f$. Consider again the "big" diagram

$$
\begin{tikzcd}
Z \times W \arrow[r] \arrow[d, "\varphi\times\mathrm{id}"] \arrow[rr, bend left, "dh"] \arrow[dd, bend right, "h\times\mathrm{id}"'] & Z \times_X T^*X \arrow[r, "d\varphi"] \arrow[d, "p_2"] & T^*Z \\
X \times W \arrow[r, "df"] \arrow[d, "f\times\mathrm{id}"] & T^*X & \\
A \times W. & &
\end{tikzcd}
$$

Last time, we said that for direct images of Hodge modules, one can compute the corresponding sheaves on the cotangent bundle very explicitly. The characteristic variety of ω_Z is the zero section in T^*Z, and the resulting coherent sheaf is $i_*\omega_Z$, where $i\colon Z \hookrightarrow T^*Z$ is the zero section. In the case of $\mathcal{M} = \mathcal{H}^0 h_+\omega_Z$, the formula from last time says that $\mathscr{G}$ is the 0-th cohomology sheaf of the complex

$$\mathbf{R}(h \times \mathrm{id})_*\mathbf{L}(dh)^*(i_*\omega_Z).$$

Let $p\colon T^*Z \to Z$ be the projection. Since the zero section is exactly the vanishing locus of the tautological section of $p^*\Omega_Z^1$, the Koszul complex

$$p^*\Omega_Z^{n+\bullet} = \Big[p^*\mathscr{O}_Z \to p^*\Omega_Z^1 \to \cdots \to p^*\Omega_Z^n\Big]$$

is a locally free resolution of the coherent sheaf $i_*\omega_Z$ on T^*Z. Consequently,

$$\mathbf{L}(dh)^*(i_*\omega_Z) = \Big[p_1^*\mathscr{O}_Z \to p_1^*\Omega_Z^1 \to \cdots \to p_1^*\Omega_Z^n\Big],$$

which means that $\mathscr{G}$ is the 0-th cohomology sheaf of the complex

$$\mathbf{R}(h \times \mathrm{id})_*\Big[p_1^*\mathscr{O}_Z \to p_1^*\Omega_Z^1 \to \cdots \to p_1^*\Omega_Z^n\Big].$$

Now consider the morphism $\varphi\colon Z \to X$. For each $p \geq 0$, we have a pullback morphism $\varphi^*\Omega_X^p \to \Omega_Z^p$; these fit together into a morphism of complexes

$$\Big[p_1^*\varphi^*\mathscr{O}_X \to p_1^*\varphi^*\Omega_X^1 \to \cdots \to p_1^*\varphi^*\Omega_X^n\Big] \to \Big[p_1^*\mathscr{O}_Z \to p_1^*\Omega_Z^1 \to \cdots \to p_1^*\Omega_Z^n\Big].$$

In derived category notation, this means that we have a morphism

$$\mathbf{L}(\varphi \times \mathrm{id})^*\mathbf{L}(df)^*(i_*\omega_X) \to \mathbf{L}(dh)^*(i_*\omega_Z).$$

Here $i\colon X \hookrightarrow T^*X$ is the zero section, and $p\colon T^*X \to X$ the projection. Since $i_*\mathscr{O}_X \otimes p^*\omega_X \cong i_*(\mathscr{O}_X \otimes i^*p^*\omega_X) \cong i_*\omega_X$ by the projection formula, we can rewrite this morphism in the more convenient form

$$p_1^*(\varphi^*\omega_X) \otimes \mathbf{L}(\varphi \times \mathrm{id})^*\mathbf{L}(df)^*(i_*\mathscr{O}_X) \to \mathbf{L}(dh)^*(i_*\omega_Z).$$

Now we compose this with (27.2) to obtain a morphism

$$p_1^*(h^*L) \otimes \mathbf{L}(\varphi \times \mathrm{id})^*\mathbf{L}(df)^*(i_*\mathscr{O}_X) \to \mathbf{L}(dh)^*(i_*\omega_Z).$$

Move the line bundle factor to the other side, and use the adjunction between the two functors $\mathbf{L}(\varphi \times \mathrm{id})^*$ and $\mathbf{R}(\varphi \times \mathrm{id})_*$. This gives an equivalent morphism

$$\mathbf{L}(df)^*(i_*\mathscr{O}_X) \to \mathbf{R}(\varphi \times \mathrm{id})_*\Big(p_1^*(h^*L^{-1}) \otimes \mathbf{L}(dh)^*(i_*\omega_Z)\Big).$$

Now push forward to $A \times W$ and use the projection formula to pull out the line bundle factor. This finally gives us the following morphism

$$\mathbf{R}(f \times \mathrm{id})_*\mathbf{L}(df)^*(i_*\mathscr{O}_X) \to \mathbf{R}(h \times \mathrm{id})_*\mathbf{L}(dh)^*(i_*\omega_Z) \otimes p_1^*L^{-1} \tag{27.3}$$

in the derived category $D^b_{coh}(\mathscr{O}_{A\times W})$. If we take cohomology in degree zero, we therefore obtain a morphism of coherent sheaves

$$\mathscr{F} \to \mathscr{G} \otimes p_1^* L^{-1}. \tag{27.4}$$

Here $\mathscr{F}$ is the 0-th cohomology sheaf of the complex $\mathbf{R}(f\times \mathrm{id})_* \mathbf{L}(df)^*(i_*\mathscr{O}_X)$, and as such, it is obviously supported inside the set

$$(f\times \mathrm{id})\big(df^{-1}(0)\big) = S_f.$$

Now all the pieces are in place to prove the theorem about one-forms.

Proof of Theorem 26.4. We are trying to show that $p_2\colon S_f \to W$ is surjective. Suppose, for the sake of argument, that $p_2(S_f) \neq W$. Then $(p_2)_*\mathscr{F}$ is a coherent sheaf on W whose support is contained inside a proper closed subset, hence a torsion sheaf. Because $(p_2)_*(\mathscr{G}\otimes p_1^*L^{-1})$ is torsion-free, the morphism

$$(p_2)_*\mathscr{F} \to (p_2)_*(\mathscr{G}\otimes p_1^*L^{-1})$$

must be trivial. Taking global sections, this means that the morphism

$$H^0(A\times W, \mathscr{F}) \to H^0(A\times W, \mathscr{G}\otimes p_1^*L^{-1})$$

is also trivial. Now both sides are actually graded modules, due to the fact that (27.3) is constructed from sheaves on the zero section (which are stable under the natural $\mathbb{C}^*$-action on the cotangent bundle). The first nontrivial graded piece (in degree $-n$) comes out to be

$$H^0(X, \mathscr{O}_X) \to H^0(Z, \omega_Z \otimes h^*L^{-1})$$

But now we have a contradiction, because the composition $h^*L \to \varphi^*\omega_X \to \omega_Z$ is not the zero morphism, due to the fact that (27.2) is nontrivial by assumption. This means that we have a nontrivial section of $\omega_Z \otimes h^*L^{-1}$, and so the above morphism cannot have been zero. The conclusion is that $p_2(S_f) = W$. $\square$