BASIC THEORY OF ALGEBRAIC $\mathcal{D}$-MODULES

NICHOLAS SWITALA

I taught a 12-week mini-course on algebraic $\mathcal{D}$-modules at UIC during the autumn of 2016. After each week, I posted lecture notes. What follows is simply a compilation of these weekly notes, which means there is more repetition in them than a polished, unified document would tolerate.

I thank Kevin Tucker, Wenliang Zhang, Chris Skalit, Eric Riedl, and especially Matthew Woolf for comments, questions, and corrections in the weekly seminar. Because of their vigilance, some false or confusing statements made in the seminar do not appear in these notes. The responsibility for any remaining errors rests, of course, with me.

Week 1: motivation from local cohomology; Weyl algebras

If R is a commutative Noetherian ring, $I = (f_1, \dots, f_r)$ is an ideal, and M is an R-module, the *local cohomology modules* $H^i_I(M)$ of M supported at I are the cohomology objects of the *Čech complex*

$$0 \to M \to \oplus_i M_{f_i} \to \oplus_{i,j} M_{f_i f_j} \to \cdots \to M_{f_1 \cdots f_r} \to 0,$$

and one can show that the modules so obtained are independent of the choice of generators $f_1, \dots, f_r$ for I. (This is the most concrete of myriad equivalent definitions of local cohomology.)

As an example, let $R = k[x]$ where k is a field. For the ideal I, take the principal ideal $(x) \subset R$, and for the module M, take R itself. In this case, the Čech complex is

$$0 \to R \xrightarrow{\delta} R_x \to 0,$$

and the map $\delta : R \to R_x$ ($r \mapsto \frac{r}{1}$) is injective, since R is a domain. Therefore $H^0_{(x)}(R) = \ker \delta = 0$, and $H^1_{(x)}(R) = \operatorname{coker} \delta = R_x/R$. As a k-space, R_x/R is identified with the direct sum $\oplus_{i>0} k \cdot x^{-i}$ of "inverse polynomials" in x; the R-module structure is defined by the usual distributive laws with the additional relations $x^j = 0$ for all $j \geq 0$.

Observe that, in the above example, $H^1_{(x)}(R)$ is *not* a finitely generated R-module. This is a theme of research in this area: local cohomology modules tend to be huge, and it is frequently useful to find additional structures on them with respect to which they are "smaller". One such additional structure is that of a $\mathcal{D}$-*module*: a module over a ring $\mathcal{D}$ of differential operators.

The basic examples of such rings $\mathcal{D}$ are *Weyl algebras*. Let k be a field of characteristic 0, and let $R = k[x_1, \dots, x_n]$. The ring $\mathcal{D} = \mathcal{D}(R, k)$ of k-*linear differential operators* on R (called the nth *Weyl algebra* and also denoted A_n) is the k-subalgebra of $\operatorname{End}_k(R)$ generated by $\{x_i, \partial_i\}_{i=1}^n$ where x_i denotes the endomorphism of R defined by multiplication with x_i and ∂_i denotes partial differentiation with respect to x_i. Clearly, the x_i commute with each other, as do the ∂_i; furthermore, x_i commutes with ∂_j if $i \neq j$. However, due to the product rule for differentiation, $\partial_i x_i = x_i \partial_i + 1$ (here 1 is the identity endomorphism of R).

In particular, $\mathcal{D}$ is a *non-commutative* ring, which contains a copy of R as a commutative subring of "multiplication" endomorphisms. $\mathcal{D}$ is a simple ring (has no nontrivial two-sided ideals), is left- and right-Noetherian, and has the further surprising property that every left (or right) ideal can be generated by at most two elements. Because $\mathcal{D}$ is non-commutative, by a "$\mathcal{D}$-module" we can mean either a left or right module over this ring, and must specify which. By convention, we always mean *left* module unless we explicitly say otherwise.

Recall our example of local cohomology in the case $n = 1$: the R-module $H^1_{(x)}(R)$ was not finitely generated. Using our concrete identification $H^1_{(x)}(R) = \oplus_{i>0} k \cdot x^{-i}$, we can view this object as a module over the first Weyl algebra, $\mathcal{D} = A_1 = k\langle x, \partial\rangle/(\partial x - x\partial - 1)$: we only need to say how $\partial = \frac{d}{dx}$ acts, and for this, we can just use the quotient rule. As a $\mathcal{D}$-module, $H^1_{(x)}(R)$ is not only finitely generated, it is generated by one element: $\frac{1}{x}$. Indeed, we can differentiate $\frac{1}{x}$ enough times to get any negative power of x; there will be a scalar numerator, but in characteristic 0, this doesn't matter.

This example illustrates two things that happen more generally. First, if $I \subset R$ is any ideal, then $H^i_I(R)$ (or even $H^i_I(M)$, for any $\mathcal{D}$-module M) can be viewed as a $\mathcal{D}$-module: the partial derivatives ∂_i act via the quotient rule, since the differentials in the Čech complex are just sums of localization maps. Moreover, the $\mathcal{D}$-modules $H^i_I(R)$ are finitely generated (even generated by one element), and of finite length in the category of $\mathcal{D}$-modules. These properties follow from the fact that $H^i_I(R)$ is a *holonomic* $\mathcal{D}$-module, as observed by Lyubeznik in 1993. Lyubeznik used the holonomy of $H^i_I(R)$ to deduce strong finiteness properties, for example that $H^i_I(R)$ has only finitely many associated prime ideals in R. (A holonomic $\mathcal{D}$-module is one that is as small as possible, among nonzero, finitely generated $\mathcal{D}$-modules, with respect to a certain measure of *dimension*. Holonomic $\mathcal{D}$-modules are sometimes called *maximally overdetermined* or *of Bernstein class*; the term *holonomic* comes, roughly, from the Greek for "everywhere law-abiding".)

Some treatments of the theory of Weyl algebras and holonomic modules over them are the book "A primer of algebraic $\mathcal{D}$-modules" by Coutinho, the first chapter of the notes "Lectures on the algebraic theory of $\mathcal{D}$-modules" by Milicic, and the first chapter of the book "Rings of differential operators" by Björk. Chapter 3 of Björk's book describes the case where $R = k[[x_1, \ldots, x_n]]$ is a formal power series ring, which is similar in many ways to the polynomial case, though harder. Also useful for learning about the power series case are the early papers and Ph.D. thesis of van den Essen. (Rather than listing them all, in an act of shameless self-promotion, I refer you to the bibliography of my own paper, "Van den Essen's theorem on the de Rham cohomology of a holonomic $\mathcal{D}$-module over a formal power series ring".)

The goal of this mini-course is to cover the sophisticated (sheafified) versions of some of the highlights of a course on Weyl algebras. It is possible to read many research papers on local cohomology having only worked with $\mathcal{D}$-modules in the concrete settings of polynomial or power series rings, and the learning curve from such concrete settings to $\mathcal{D}$-modules over smooth schemes can be steep.

This mini-course will be based on the first two or three chapters of the book "$\mathcal{D}$-modules, perverse sheaves, and representation theory" by Hotta, Takeuchi, and Tanisaki, hereafter "HTT". We will discuss the theory of $\mathcal{D}_X$-modules, where X is a smooth scheme of pure dimension over an algebraically closed field k of characteristic 0, with structure sheaf $\mathcal{O}_X$. (HTT assumes $k = \mathbb{C}$ throughout, but they only begin using this assumption in chapter 4.) Eventually, we will need to make an additional mild assumption on X, namely that every coherent $\mathcal{O}_X$-module is the quotient of a locally free $\mathcal{O}_X$-module. Since X is already assumed to be smooth, this assumption (the "resolution property") amounts to requiring that X have affine diagonal (the intersection of any two open affines in X is affine). HTT, and other sources on $\mathcal{D}$-modules, usually assume that X is quasi-projective in order to guarantee the resolution property. This discussion only becomes relevant when functors between derived categories are considered.

Besides covering the basic constructions and operations on $\mathcal{D}$-modules (most importantly, push-forward and pull-back), our goals will be to cover the proofs of *Kashiwara's theorem* (if $Z \hookrightarrow X$

is a closed immersion, the category of $\mathcal{D}_Z$-modules is equivalent to the category of $\mathcal{D}_X$-modules supported on Z) and *Bernstein's inequality*, which is the fundamental result on *dimensions* of $\mathcal{D}_X$-modules necessary to define the category of holonomic modules.

Some other references for this material include Bernstein's lecture notes "Algebraic theory of $\mathcal{D}$-modules", chapters VI and VII of the book "Algebraic $\mathcal{D}$-modules" by Borel et al., and Jonathan Wang's Cambridge Part III essay, "Introduction to $\mathcal{D}$-modules and representation theory". There is a riotous surfeit of differing notations for the basic functors: when the standard references disagree, which is nearly always, we will use the notation in HTT.

Week 2: basic definitions; coherence and quasi-coherence

Let k be an algebraically closed field of characteristic 0, and let X be a smooth scheme of pure dimension n over k, with structure sheaf $\mathcal{O}_X$. The *cotangent sheaf* Ω^1_X is a locally free $\mathcal{O}_X$-module of rank n, and hence so is its $\mathcal{O}_X$-dual, the *tangent sheaf* $\Theta_X = \mathcal{H}om_{\mathcal{O}_X}(\Omega^1_X, \mathcal{O}_X)$. (In other sources, Θ_X is frequently denoted $\mathcal{T}_X$.) By the universal property of Ω^1_X as a sheaf of Kähler differentials, Θ_X is identified with the sheaf of k-linear *derivations* (or "vector fields") on X:

$$\Theta_X \simeq \mathcal{D}er_k(\mathcal{O}_X) = \{\delta \in \mathcal{E}nd_k(\mathcal{O}_X) \,|\, \delta(fg) = \delta(f)g + f\delta(g) \,\forall f, g \in \mathcal{O}_X\}$$

where we abuse notation by writing $f \in \mathcal{O}_X$ when f is a local section of $\mathcal{O}_X$. Both $\mathcal{O}_X$ (as "multiplications") and Θ_X (as derivations) are subsheaves of $\mathcal{E}nd_k(\mathcal{O}_X)$.

Definition. *$\mathcal{D}_X$, the sheaf of k-*linear differential operators *on X, is the k-subalgebra of $\mathcal{E}nd_k(\mathcal{O}_X)$ generated by $\mathcal{O}_X$ and Θ_X.*

Let $x \in X$ be given. Since X is smooth, there is an open affine neighborhood $x \in U \subset X$ and sections $x_1, \ldots, x_n \in \mathcal{O}_X(U)$, $\partial_1, \ldots, \partial_n \in \Theta_X(U)$ such that

- the ∂_i commute with each other (of course, the x_i always commute with each other);
- $\partial_i(x_j) = \delta_{i,j}$ (Kronecker delta) for all i, j;
- the ∂_i generate $\Theta_X(U)$ over $\mathcal{O}_X(U)$.

The x_i (resp. ∂_i) can be chosen to be lifts of a regular system of parameters for the regular local ring $\mathcal{O}_{X,x}$ (resp. lifts of the dual basis to the differentials $dx_1, \ldots, dx_n \in \Omega^1_{X,x}$). We call $\{x_i, \partial_i\}_{i=1}^n$ a *(local) coordinate system* on U. The x_i are often called *étale coordinates*, since they define an étale morphism $U \to \mathbb{A}^n_k$. Over U, the sheaf $\mathcal{D}_X$ takes the form of a Weyl algebra with respect to the x_i and ∂_i: that is, we have

$$\mathcal{D}_X|_U = \oplus_{\alpha_1,\ldots,\alpha_n \geq 0} \mathcal{O}_U \cdot \partial_1^{\alpha_1} \cdots \partial_n^{\alpha_n}.$$

A *$\mathcal{D}_X$-module* is just a sheaf of *left* modules over the sheaf of non-commutative rings $\mathcal{D}_X$. For example, $\mathcal{O}_X$ is a $\mathcal{D}_X$-module in an obvious way. Observe that $\mathcal{D}_X$ is locally free, and therefore quasi-coherent, as an $\mathcal{O}_X$-module. Often, but not always, we will restrict attention to $\mathcal{D}_X$-modules that are quasi-coherent as $\mathcal{O}_X$-modules. We will work with the following categories:

- $\mathrm{Mod}(\mathcal{D}_X)$, the category of all left $\mathcal{D}_X$-modules;
- $\mathrm{Mod}_{qc}(\mathcal{D}_X)$, the category of all left $\mathcal{D}_X$-modules that are quasi-coherent as $\mathcal{O}_X$-modules;
- $\mathrm{Mod}_c(\mathcal{D}_X)$, the category of all left $\mathcal{D}_X$-modules that are coherent as $\mathcal{D}_X$-modules (that is, quasi-coherent as $\mathcal{O}_X$-modules and locally finitely generated over $\mathcal{D}_X$);
- $\mathrm{Mod}(\mathcal{D}_X^{\mathrm{op}})$, $\mathrm{Mod}_{qc}(\mathcal{D}_X^{\mathrm{op}})$, and $\mathrm{Mod}_c(\mathcal{D}_X^{\mathrm{op}})$, the analogues of the above for *right* $\mathcal{D}_X$-modules.

Since $\mathcal{D}_X$ is quasi-coherent (but not coherent) over $\mathcal{O}_X$, a quasi-coherent $\mathcal{D}_X$-module is just a $\mathcal{D}_X$-module that is quasi-coherent over $\mathcal{O}_X$. On the other hand, a coherent $\mathcal{D}_X$-module need not be coherent over $\mathcal{O}_X$. If it is, then it must be a vector bundle:

Proposition. *If M is a $\mathcal{D}_X$-module that is coherent over $\mathcal{O}_X$, then M is locally free over $\mathcal{O}_X$.*

Proof. Since M is coherent, it suffices to check that M_x is a free $\mathcal{O}_{X,x}$-module for all $x \in X$. Fix $x \in X$ and local coordinates $\{x_i, \partial_i\}$ on an open affine neighborhood U of x, where the x_i are lifts of a set of generators of the maximal ideal $\mathfrak{m}_x \subset \mathcal{O}_{X,x}$. By Nakayama's lemma, we can find generators $s_1, \dots, s_m$ for M_x over $\mathcal{O}_{X,x}$ whose images $\overline{s}_1, \dots, \overline{s}_m$ in $M_x / \sum_i x_i M_x$ form a basis for this vector space over $\mathcal{O}_{X,x}/\mathfrak{m}_x = k$. We claim that $s_1, \dots, s_m$ are *free* generators for M_x. Suppose, for contradiction, that there is a nontrivial dependence relation $\sum_{i=1}^m f_i s_i = 0$ ($f_i \in \mathcal{O}_{X,x}$) such that the minimum of the orders of the f_i is as small as possible (here the *order* of $f \in \mathcal{O}_{X,x}$ is $\max\{l \mid f \in \mathfrak{m}_x^l\}$). Observe that this minimal order must be positive, since if any f_i were a unit, we could reduce the given dependence relation modulo $\mathfrak{m}_x$ to obtain a nontrivial k-linear dependence relation among the $\overline{s}_i$. Relabeling if necessary, we may assume f_1 realizes the minimal order. Choose j such that the order of $\partial_j(f_1)$ is strictly less than that of f_1 (in down-to-earth terms, choose a parameter x_j that occurs in the lowest-degree term of f_1, and differentiate with respect to it). Then we have

$$0 = \partial_j(0) = \partial_j(\sum_{i=1}^m f_i s_i) = \sum_{i=1}^m \partial_j(f_i)s_i + f_i \partial_j(s_i),$$

where the rightmost expression can be expanded as an $\mathcal{O}_{X,x}$-linear combination of $s_1, \dots, s_m$ whose coefficient of minimal order is of strictly smaller order than f_1. This contradiction finishes the proof. □

There is another, more general, definition of $\mathcal{D}_X$, in which we construct $\mathcal{D}_X$ recursively as a *filtered* sheaf of rings. Let $F_0\mathcal{D}_X = \mathcal{O}_X$, and for all $l > 0$, define

$$F_l\mathcal{D}_X = \{\delta \in \mathcal{E}nd_k(\mathcal{O}_X) \mid [\delta, f] \in F_{l-1}\mathcal{D}_X \, \forall f \in \mathcal{O}_X\};$$

finally, set $\mathcal{D}_X = \cup_{l \geq 0} F_l\mathcal{D}_X$. (Here, $[\delta, f]$ denotes the *commutator* $\delta f - f\delta$.) By induction on $l + m$, it is easy to prove that, for all $l, m \geq 0$, we have

- $F_l\mathcal{D}_X \cdot F_m\mathcal{D}_X \subset F_{l+m}\mathcal{D}_X$ (key formula: $[\delta\delta', f] = \delta[\delta', f] + [\delta, f]\delta'$), and
- $[F_l\mathcal{D}_X, F_m\mathcal{D}_X] \subset F_{l+m-1}\mathcal{D}_X$ (key formula: $[[\delta, \delta'], f] = [[\delta, f], \delta'] + [\delta, [\delta', f]]$).

This definition does not require us to make any assumptions (*e.g.* smoothness) on X. When X is smooth, the sheaf $\mathcal{D}_X$ constructed above agrees with our earlier definition. In general, we always have $F_1\mathcal{D}_X = \mathcal{O}_X \oplus \Theta_X$ (given $\delta \in F_1\mathcal{D}_X$, we associate with it the pair $(\delta(1) \in \mathcal{O}_X, \delta - \delta(1) \in \Theta_X)$). Over an open affine $U \subset X$ with a coordinate system $\{x_i, \partial_i\}$, we have

$$F_l\mathcal{D}_U = \oplus_{\alpha_1 + \cdots + \alpha_n \leq l} \mathcal{O}_U \cdot \partial_1^{\alpha_1} \cdots \partial_n^{\alpha_n},$$

so each $F_l\mathcal{D}_X$ is a locally free $\mathcal{O}_X$-module of finite rank. The filtration $\{F_l\mathcal{D}_X\}_{l \geq 0}$ is called the *order filtration* (sometimes "degree filtration") on $\mathcal{D}_X$, and the elements of $F_l\mathcal{D}_X$ are *differential operators of order* $\leq l$.

Finally, observe that if we pass to the *associated graded* sheaf of rings

$$\mathrm{gr}^F\mathcal{D}_X = \oplus_{l=0}^\infty F_l\mathcal{D}_X / F_{l-1}\mathcal{D}_X$$

($F_{-1}\mathcal{D}_X = 0$), we obtain a sheaf of *commutative* rings due to the relation $[F_l\mathcal{D}_X, F_m\mathcal{D}_X] \subset F_{l+m-1}\mathcal{D}_X$. (In fact, the sheaf $\mathrm{gr}^F\mathcal{D}_X$ can be identified with the symmetric algebra of the tangent sheaf Θ_X.) We will discuss $\mathrm{gr}^F\mathcal{D}_X$ in more detail later; it serves a useful role in allowing us to apply techniques of commutative algebra to obtain results about the non-commutative $\mathcal{D}_X$.

Week 3: side-changing operations

Like last week, we let X be a smooth scheme of pure dimension n over an algebraically closed field k of characteristic 0. We let $\mathcal{O}_X$, Θ_X, and Ω^1_X be the structure, tangent, and cotangent sheaves of X respectively, and we let $\mathcal{D}_X$ be the k-subalgebra of $\mathcal{E}nd_k(\mathcal{O}_X)$ generated by $\mathcal{O}_X$ and Θ_X. *For the rest of this mini-course, these notations and hypotheses are fixed*, although we may occasionally need to impose additional conditions on X.

Suppose M is a $\mathcal{O}_X$-module. Since Θ_X generates $\mathcal{D}_X$ over $\mathcal{O}_X$, in order to give M a structure of left or right $\mathcal{D}_X$-module, it is enough to specify how the derivations $\delta \in \Theta_X$ act on M, as long as the relations in and between $\mathcal{O}_X$ and Θ_X are respected. Recall that Θ_X is a sheaf of Lie algebras: the commutator $[\delta, \delta'] = \delta\delta' - \delta'\delta$ of two derivations is again a derivation, and this operation satisfies the Lie algebra axioms. Of course, if $\delta \in \Theta_X$ is a derivation and $f \in \mathcal{O}_X$, then $f\delta$ is again a derivation. We also have the relations $[\delta, f] = \delta(f)$ for all $\delta \in \Theta_X$ and $f \in \mathcal{O}_X$. If we specify elements $\delta \cdot m \in M$ for all $\delta \in \Theta_X$ and $m \in M$ in such a way that, for all $\delta, \delta' \in \Theta_X$, $f \in \mathcal{O}_X$, and $m \in M$, we have

- $[\delta_1, \delta_2] \cdot m = \delta_1 \cdot (\delta_2 \cdot m) - \delta_2 \cdot (\delta_1 \cdot m)$,
- $(f\delta) \cdot m = f(\delta \cdot m)$, and
- $(\delta f) \cdot m = f(\delta \cdot m) + \delta(f)m$,

then we obtain a structure of left $\mathcal{D}_X$-module on M. To obtain a structure of right $\mathcal{D}_X$-module on M in a similar way, we begin by specifying elements $m \cdot \delta$ for all $\delta \in \Theta_X$ and $m \in M$ such that the obvious right-to-left analogues of the first two conditions are satisfied; the replacement for the third condition is

$$m \cdot (f\delta) = f(m \cdot \delta) - \delta(f)m,$$

since $f\delta = \delta f - \delta(f)$.

The above recipe for imposing $\mathcal{D}_X$-structures on $\mathcal{O}_X$-modules can be used to build new $\mathcal{D}_X$-modules from old ones using the tensor and $\mathcal{H}om$ operations over $\mathcal{O}_X$. For example, suppose that M' (resp. N) is a right (resp. left) $\mathcal{D}_X$-module. Then the $\mathcal{O}_X$-module $M' \otimes_{\mathcal{O}_X} N$ becomes a right $\mathcal{D}_X$-module using the formula

$$(m' \otimes n) \cdot \delta = m' \cdot \delta \otimes n - m' \otimes \delta \cdot n$$

for all $\delta \in \Theta_X$, $m' \in M'$, and $n \in N$. To be fully honest, every time a $\mathcal{D}_X$-structure is defined by specifying how the derivations act, we need to check all the relations as above. As a sample calculation, we check the third condition:

$$\begin{aligned}(m' \otimes n) \cdot (f\delta) &= m' \cdot (f\delta) \otimes n - m' \otimes (f\delta) \cdot n \\ &= (f(m' \cdot \delta) - \delta(f)m') \otimes n - m' \otimes f(\delta \cdot n) \\ &= f(m' \cdot \delta \otimes n - m' \otimes \delta \cdot n) - \delta(f)(m' \otimes n) \\ &= (m' \otimes n) \cdot (\delta f) - \delta(f)(m' \otimes n).\end{aligned}$$

As another example, if N' is a right $\mathcal{D}_X$-module, then the $\mathcal{O}_X$-module $\mathcal{H}om_{\mathcal{O}_X}(M', N')$ becomes a left $\mathcal{D}_X$-module using the formula

$$(\delta \cdot \varphi)(m') = \varphi(m' \cdot \delta) - \varphi(m') \cdot \delta$$

for $\varphi \in \mathcal{H}om_{\mathcal{O}_X}(M', N')$, $m' \in M'$, and $\delta \in \Theta_X$.

Many, but not all, combinations of left and right $\mathcal{D}_X$-modules constructed using tensor and $\mathcal{H}om$ over $\mathcal{O}_X$ can be given either a left or right $\mathcal{D}_X$-structure. A way to remember this is to use *Oda's rule*, which says that if X is a smooth curve of genus g and $\mathcal{L}$ is a line bundle on X, then $\mathcal{L}$ can be given a left (resp. right) $\mathcal{D}_X$-module structure if and only if the degree of $\mathcal{L}$ is 0 (resp. $2g-2$). An even simpler mnemonic is "left = 0, right = 1, $\otimes$ = +, and $\mathcal{H}om$ = target minus source", where addition is *not* understood modulo 2: the result of the addition or subtraction must be 0 or 1 if the resulting module is to support a left or right $\mathcal{D}_X$-module structure. Therefore, if M and N are left $\mathcal{D}_X$-modules and M' and N' are right $\mathcal{D}_X$-modules, $M \otimes_{\mathcal{O}_X} N$ can be given a structure of left $\mathcal{D}_X$-module ("$0+0=0$"), whereas $M' \otimes_{\mathcal{O}_X} N'$ cannot be given a structure of either left or right $\mathcal{D}_X$-module ("$1+1=2$"), and neither can $\mathcal{H}om_{\mathcal{O}_X}(M', N)$ ("$0-1=-1$").

The standard example of a left $\mathcal{D}_X$-module is the structure sheaf $\mathcal{O}_X$. Its counterpart, the standard example of a *right* $\mathcal{D}_X$-module, is the *canonical sheaf* $\omega_X = \bigwedge^n \Omega^1_X$, an invertible (locally free of rank 1) $\mathcal{O}_X$-module. This sheaf is denoted Ω_X in HTT; this will be one of our few deviations from the notation of that book. We are going to use the canonical sheaf to set up the *side-changing operations*, which are quasi-inverse functors defining an equivalence of categories between $\mathrm{Mod}(\mathcal{D}_X)$ and $\mathrm{Mod}(\mathcal{D}_X^{\mathrm{op}})$. The side-changing operations have simple descriptions in local coordinates. However, it is useful to define them first using global, sheaf-theoretic constructions and only then to calculate what they do in coordinates, rather than giving an *a priori* coordinate-dependent definition and then providing an independence proof. (We will continue to develop pieces of $\mathcal{D}_X$-module theory in this order.)

The right $\mathcal{D}_X$-module structure on ω_X is defined by means of the *Lie derivative*. We have an isomorphism

$$\omega_X = \bigwedge^n \Omega^1_X = \bigwedge^n \mathcal{H}om_{\mathcal{O}_X}(\Theta_X, \mathcal{O}_X) \simeq \mathcal{H}om_{\mathcal{O}_X}(\bigwedge^n \Theta_X, \mathcal{O}_X)$$

of $\mathcal{O}_X$-modules. Let $\delta \in \Theta_X$ and $\omega \in \omega_X$ be given, and identify ω with an $\mathcal{O}_X$-linear homomorphism $\bigwedge^n \Theta_X \to \mathcal{O}_X$. Then the *Lie derivative* $\mathrm{Lie}_\delta(\omega)$ of the form ω along the derivation δ is the $\mathcal{O}_X$-linear homomorphism $\bigwedge^n \Theta_X \to \mathcal{O}_X$ (that is, element of ω_X) defined by

$$\mathrm{Lie}_\delta(\omega)(\delta_1 \wedge \cdots \wedge \delta_n) = \delta(\omega(\delta_1 \wedge \cdots \wedge \delta_n)) - \sum_{i=1}^n \omega(\delta_1 \wedge \cdots \wedge [\delta, \delta_i] \wedge \cdots \wedge \delta_n)$$

where $\delta_1, \ldots, \delta_n \in \Theta_X$. If we set $\omega \cdot \delta = -\mathrm{Lie}_\delta(\omega)$ for $\omega \in \omega_X$ and $\delta \in \Theta_X$, the axioms above for a right $\mathcal{D}_X$-module structure on ω_X are satisfied.

If M is a left $\mathcal{D}_X$-module, then by the formula given earlier, $\omega_X \otimes_{\mathcal{O}_X} M$ is a right $\mathcal{D}_X$-module. This is the left-to-right side-changing operation. To construct its quasi-inverse, consider the dual sheaf $\omega_X^{\otimes -1} = \mathcal{H}om_{\mathcal{O}_X}(\omega_X, \mathcal{O}_X)$. Let M be a right $\mathcal{D}_X$-module. We have an $\mathcal{O}_X$-module isomorphism

$$\omega_X^{\otimes -1} \otimes_{\mathcal{O}_X} M = \mathcal{H}om_{\mathcal{O}_X}(\omega_X, \mathcal{O}_X) \otimes_{\mathcal{O}_X} M \simeq \mathcal{H}om_{\mathcal{O}_X}(\omega_X, M),$$

and we know by another formula given earlier that any $\mathcal{H}om$ between two right $\mathcal{D}_X$-modules is a left $\mathcal{D}_X$-module. This is the right-to-left side-changing operation, and in fact we have

$$\omega_X^{\otimes -1} \otimes_{\mathcal{O}_X} (\omega_X \otimes_{\mathcal{O}_X} M) \simeq M$$

as *left* $\mathcal{D}_X$*-modules* (and an analogous statement for right modules), that is, the two side-changing operations are quasi-inverse functors.

Finally, we describe the effect of the side-changing operation in local coordinates. Suppose that X is affine with coordinates $\{x_i, \partial_i\}$. In this case, the canonical sheaf ω_X is *globally* trivial: $\mathcal{O}_X \xrightarrow{\sim} \omega_X$ via $1 \mapsto dx_1 \wedge \cdots \wedge dx_n$. (We simply write dx for the top form $dx_1 \wedge \cdots \wedge dx_n$.) Therefore, given any left $\mathcal{D}_X$-module M, the underlying $\mathcal{O}_X$-modules M and $\omega_X \otimes_{\mathcal{O}_X} M$ are isomorphic. To describe the right $\mathcal{D}_X$-action on $\omega_X \otimes_{\mathcal{O}_X} M$, it suffices to specify how the derivations ∂_i act. The key observation here is that $\mathrm{Lie}_{\partial_i}(dx) = 0$ for all i. Indeed, since dx is the dual basis element to $\partial_1 \wedge \cdots \wedge \partial_n \in \bigwedge^n \Theta_X$, the first term in the Lie derivative, $\partial_i(dx(\partial_1 \wedge \cdots \wedge \partial_n))$, is ∂_i applied to a constant and hence vanishes, and the remaining terms vanish because the ∂_j all commute with ∂_i. Therefore, if $m \in M$, we have (by our rule for the right $\mathcal{D}_X$-action on the tensor product of the right $\mathcal{D}_X$-module ω_X with the left $\mathcal{D}_X$-module M)

$$(dx \otimes m) \cdot \partial_i = -\mathrm{Lie}_{\partial_i}(dx) \otimes m - dx \otimes \partial_i \cdot m = -dx \otimes \partial_i \cdot m,$$

so if $f \in \mathcal{O}_X$, we have

$$(dx \otimes m) \cdot (f\partial_i) = (dx \otimes fm) \cdot \partial_i = -dx \otimes \partial_i \cdot (fm).$$

From this calculation, it is easy to see how any element of $\mathcal{D}_X$ acts on M on the right. Under the isomorphism $M \simeq \omega_X \otimes_{\mathcal{O}_X} M$, the element m corresponds to $dx \otimes m$, and using this identification, we see that the right action of $f\partial_i$ on m is the same as the left action of $-\partial_i f$ on m. In general, we can define the right $\mathcal{D}_X$-action on M by $m \cdot \delta = \delta^t \cdot m$, where for any differential operator $\delta \in \mathcal{D}_X$ (not just derivations), we define the "transpose" (or *formal adjoint*) δ^t of δ by setting

$$(x_1^{\alpha_1} \cdots x_n^{\alpha_n} \partial_1^{\beta_1} \cdots \partial_n^{\beta_n})^t = (-1)^{\beta_1 + \cdots \beta_n} \partial_1^{\beta_1} \cdots \partial_n^{\beta_n} x_1^{\alpha_1} \cdots x_n^{\alpha_n}$$

and extending by linearity. Since $\mathcal{D}_X$ is a Weyl algebra with respect to the x_i and ∂_i, this defines the formal adjoint for all differential operators: the derivations are moved to the left past the variables, and each derivation contributes a sign.

Week 4: naive pullback and pushforward

Let $f : X \to Y$ be a morphism of smooth schemes over k. We are going to describe the "naive" (non-derived) inverse and direct image functors for $\mathcal{D}_X$- and $\mathcal{D}_Y$-modules. The story is simpler in the case of the inverse image. If M is a left $\mathcal{D}_Y$-module, its $\mathcal{O}$-module inverse image, $f^*M = \mathcal{O}_X \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}M$, can be given a structure of left $\mathcal{D}_X$-module (here $f^{-1}M$ is the sheaf-theoretic inverse image), as follows. Corresponding to the scheme morphism f, we have a map $f^*\Omega^1_Y \to \Omega^1_X$ of $\mathcal{O}_X$-modules. Taking the $\mathcal{O}_X$-dual, we obtain a map $\Theta_X \to f^*\Theta_Y$, which we denote df and refer to as the *differential* of f. Given $\delta \in \Theta_X$ and $p \otimes m \in f^*M$ (where, abusively, we write m for an element of $f^{-1}M$), we define

$$\delta \cdot (p \otimes m) = \delta(p) \otimes m + p \cdot df(\delta)(1 \otimes m),$$

a "chain-rule-type" action, which makes sense because $df(\delta) \in f^*\Theta_Y$ acts on f^*M.

As an example, suppose $X = \mathbb{A}^n$ and $Y = \mathbb{A}^m$ are affine spaces over k, and let $f : X \to Y$ be a morphism, which must come from a ring map $f^\# : k[y_1, \ldots, y_m] \to k[x_1, \ldots, x_n]$. Write $k[\vec{y}]$ (resp. $k[\vec{x}]$) for these coordinate rings, and let $F_j = f^\#(y_j)$ for $j = 1, \ldots, m$. If M is a left $\mathcal{D}_Y$-module, the action of $\partial_{x_i} \in \mathcal{D}_X$ on $p \otimes m \in f^*M = k[\vec{x}] \otimes_{k[\vec{y}]} M$ defined in the previous paragraph becomes

$$\partial_{x_i} \cdot (p \otimes m) = \partial_{x_i}(p) \otimes m + \sum_{j=1}^{m} p \cdot \frac{\partial F_j}{\partial x_i} \otimes \partial_{y_j} \cdot m,$$

and it is perhaps easier to see here the resemblance to the chain rule. (In Coutinho's book, there is a careful proof that this action respects the relations $[\partial_{x_i}, x_j] = \delta_{i,j}$ in $\mathcal{D}_X$.) More generally, if $f : X \to Y$ is an arbitrary morphism of smooth schemes, $\dim Y = m$, and $\{y_j, \partial_j\}$ are local coordinates on Y, then we have

$$\delta \cdot (p \otimes m) = \delta(p) \otimes m + p \cdot \sum_{j=1}^{m} \delta(y_j \circ f) \otimes \partial_j \cdot m$$

in these coordinates. Here $y_j \circ f$ makes literal sense as a regular function on X if X and Y are, for example, affine varieties; but in the general case of an abstract morphism of schemes, we must use the sheaf map $\mathcal{O}_Y \to f_*\mathcal{O}_X$ to make sense of it. (Borel's book takes the above as the *definition* of the inverse image operation on $\mathcal{D}_Y$-modules, and then sketches a proof that the action so defined is independent of the chosen local coordinates on Y.)

If we apply the inverse image operation to the left $\mathcal{D}_Y$-module $\mathcal{D}_Y$ itself, the resulting object, $f^*\mathcal{D}_Y = \mathcal{O}_X \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}\mathcal{D}_Y$, is a $(\mathcal{D}_X, f^{-1}\mathcal{D}_Y)$-bimodule: the left $\mathcal{D}_X$-action comes from the "chain rule" as above, and the right $f^{-1}\mathcal{D}_Y$-action is just right multiplication on the right tensor factor. We denote this bimodule by $\mathcal{D}_{X\to Y}$. Observe that by the associativity of tensor products, we have

$$f^*M = \mathcal{O}_X \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}M \simeq \mathcal{O}_X \otimes_{f^{-1}\mathcal{O}_Y} (f^{-1}\mathcal{D}_Y \otimes_{f^{-1}\mathcal{D}_Y} f^{-1}M) \simeq \mathcal{D}_{X\to Y} \otimes_{f^{-1}\mathcal{D}_Y} f^{-1}M$$

as left $\mathcal{D}_X$-modules. We may therefore express the inverse image operation as $M \mapsto \mathcal{D}_{X\to Y} \otimes_{f^{-1}\mathcal{D}_Y} f^{-1}M$, which makes it clear that this operation is a functor $f^* : \mathrm{Mod}(\mathcal{D}_Y) \to \mathrm{Mod}(\mathcal{D}_X)$ that is right-exact and preserves quasi-coherence. The complicated nature of the left $\mathcal{D}_X$-action on f^*M is "quarantined" in the first tensor factor.

Consider the following simple example: let $X = \mathbb{A}^n$ and $Y = \mathbb{A}^{n+1}$, and let $i : X \hookrightarrow Y$ be the closed immersion defined by the surjective ring map $i^\# : k[\vec{x}, y] \to k[\vec{x}]$ that sends y to 0. The bimodule $\mathcal{D}_{X\to Y}$ is, by definition, $k[\vec{x}] \otimes_{k[\vec{x},y]} \mathcal{D}_Y$, which is isomorphic to the quotient $\mathcal{D}_Y / y \cdot \mathcal{D}_Y$ of $\mathcal{D}_Y$ by its *right* ideal $y \cdot \mathcal{D}_Y$. Another way to describe this bimodule is as the tensor product $\mathcal{D}_X \otimes_k k[\partial_y]$, which decomposes (as a left $\mathcal{D}_X$-module) into a direct sum of infinitely many copies of $\mathcal{D}_X$, indexed by the powers of ∂_y. It follows that, for any left $\mathcal{D}_Y$-module M, $i^*M \simeq M/y \cdot M$ as left $\mathcal{D}_X$-modules. More generally, if $i : X \hookrightarrow Y$ is any closed immersion between smooth schemes over k, we can choose local coordinates $\{y_i, \partial_i\}_{i=1}^n$ on Y such that $y_{n-c+1} = \cdots = y_n = 0$ are defining equations for X as a closed subscheme of Y, where c is the codimension of X in Y. With respect to these coordinates,

$$\mathcal{D}_{X\to Y} \simeq \mathcal{D}_X \otimes_k k[\partial_{n-c+1}, \ldots, \partial_n]$$

as left $\mathcal{D}_X$-modules. In particular, if $c > 0$, $\mathcal{D}_{X\to Y}$ is a free left $\mathcal{D}_X$-module of infinite rank.

The previous example shows that the inverse image functor does *not*, in general, preserve coherence, because tensoring with an infinite-rank $\mathcal{D}_X$-module does not produce a coherent $\mathcal{D}_X$-module in general. The example also shows that, in contrast to the case of inverse images, we cannot define a direct image functor for $\mathcal{D}$-modules that agrees with the usual f_* on the underlying $\mathcal{O}$-modules. To see this, consider again the surjection $i^\# : k[\vec{x}, y] \to k[\vec{x}]$ defining a closed immersion of affine spaces $X \hookrightarrow Y$. Let M be a left $\mathcal{D}_X$-module. The functor i_* corresponds to restriction of scalars, so y acts as 0 on the $\mathcal{O}_Y$-module i_*M. In order to make i_*M a left $\mathcal{D}_Y$-module, we would need to define the action of ∂_y on M in a manner respecting the relations in $\mathcal{D}_Y$. This is impossible: in $\mathcal{D}_Y$, we have the relation $\partial_y \cdot y - y \cdot \partial_y = 1$, but if y acts as 0 on M, so must $\partial_y \cdot y - y \cdot \partial_y$.

Given a morphism $f : X \to Y$ and a left $\mathcal{D}_X$-module M, our goal is to build a left $\mathcal{D}_Y$-module out of M and f_* in some way; the previous paragraph shows that we cannot simply take f_*M. It turns out to be easier to see what to do if we begin with a *right* $\mathcal{D}_X$-module M. Recall that $\mathcal{D}_{X\to Y}$ is a $(\mathcal{D}_X, f^{-1}\mathcal{D}_Y)$-bimodule. Since M is a right $\mathcal{D}_X$-module, we can form the tensor product $M \otimes_{\mathcal{D}_X} \mathcal{D}_{X\to Y}$. Right multiplication on the second tensor factor gives this product a right $f^{-1}\mathcal{D}_Y$-module structure. If we then apply f_*, we get a right $f_*f^{-1}\mathcal{D}_Y$-module $f_*(M \otimes_{\mathcal{D}_X} \mathcal{D}_{X\to Y})$, which becomes a right $\mathcal{D}_Y$-module via the adjunction unit $\mathcal{D}_Y \to f_*f^{-1}\mathcal{D}_Y$. By using the side-changing operations, we can define a similar operation for a *left* $\mathcal{D}_X$-module M: the sequence of operations

$$M \mapsto \omega_X \otimes_{\mathcal{O}_X} M \mapsto f_*((\omega_X \otimes_{\mathcal{O}_X} M) \otimes_{\mathcal{D}_X} \mathcal{D}_{X\to Y}) \mapsto \omega_Y^{\otimes -1} \otimes_{\mathcal{O}_Y} f_*((\omega_X \otimes_{\mathcal{O}_X} M) \otimes_{\mathcal{D}_X} \mathcal{D}_{X\to Y})$$

produces a left $\mathcal{D}_X$-module. Now recall that the *projection formula* says that if $\mathcal{F}$ is any $\mathcal{O}_X$-module and $\mathcal{L}$ is a line bundle (or any vector bundle) on Y, we have

$$f_*\mathcal{F} \otimes_{\mathcal{O}_Y} \mathcal{L} \simeq f_*(\mathcal{F} \otimes_{\mathcal{O}_X} f^*\mathcal{L}) = f_*(\mathcal{F} \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}\mathcal{L})$$

as $\mathcal{O}_Y$-modules. If we apply the projection formula with $\mathcal{L} = \omega_Y^{\otimes -1}$ and $\mathcal{F} = (\omega_X \otimes_{\mathcal{O}_X} M) \otimes_{\mathcal{D}_X} \mathcal{D}_{X\to Y}$, and use the commutativity and associativity of tensor products, we obtain an isomorphism

$$\omega_Y^{\otimes -1} \otimes_{\mathcal{O}_Y} f_*((\omega_X \otimes_{\mathcal{O}_X} M) \otimes_{\mathcal{D}_X} \mathcal{D}_{X\to Y}) \simeq f_*((\omega_X \otimes_{\mathcal{O}_X} \mathcal{D}_{X\to Y} \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}\omega_Y^{\otimes -1}) \otimes_{\mathcal{D}_X} M)$$

of left $\mathcal{D}_X$-modules. Therefore we can write the direct image operation as a functor $\mathrm{Mod}(\mathcal{D}_X) \to \mathrm{Mod}(\mathcal{D}_Y)$ defined by $M \mapsto f_*(\mathcal{D}_{Y\leftarrow X} \otimes_{\mathcal{D}_X} M)$, where $\mathcal{D}_{Y\leftarrow X}$ is the $(f^{-1}\mathcal{D}_Y, \mathcal{D}_X)$-bimodule

$$\mathcal{D}_{Y\leftarrow X} = \omega_X \otimes_{\mathcal{O}_X} \mathcal{D}_{X\to Y} \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}\omega_Y^{\otimes -1}.$$

We call $\mathcal{D}_{X\to Y}$ and $\mathcal{D}_{Y\leftarrow X}$ the *transfer bimodules* associated with f. Because the naive candidate for a direct image functor just defined mixes a left exact functor (f_*) with a right exact functor ($\otimes$), we will only ever work with its derived version, to avoid difficulties with homological algebra and properties such as the composition rule.

Week 5: good filtrations; structure of $\mathcal{D}_X$ and $\mathrm{gr}\,\mathcal{D}_X$

Recall that $\mathcal{D}_X = \cup_l F_l\mathcal{D}_X$ is a *filtered* sheaf of rings on X, via the order (or degree) filtration. Since $[F_l\mathcal{D}_X, F_m\mathcal{D}_X] \subset F_{l+m-1}\mathcal{D}_X$ for all l and m, the associated graded sheaf, $\mathrm{gr}^F\mathcal{D}_X = \oplus_{l=0}^{\infty} F_l\mathcal{D}_X/F_{l-1}\mathcal{D}_X$, is a sheaf of commutative rings on X. If $\{x_i, \partial_i\}$ are local coordinates on an open affine $U \subset X$, let ξ_i be the image of ∂_i in $F_1\mathcal{D}_U/F_0\mathcal{D}_U \subset \mathrm{gr}^F\mathcal{D}_U$ (called the *principal symbol* of ∂_i); then $\mathrm{gr}^F\mathcal{D}_U \simeq \mathcal{O}_U[\xi_1, \ldots, \xi_n]$.

We are going to discuss filtrations on left $\mathcal{D}_X$-modules. For now, we assume that X is affine, and let D_X be the filtered ring $\mathcal{D}_X(X)$. We simply write $\mathrm{gr}\,D_X$ for $\mathrm{gr}^F D_X$. A left D_X-module M is called a *filtered D_X-module* if it is provided with an increasing, exhaustive filtration F by additive subgroups, $\{F_pM\}_{p=0}^{\infty}$ (that is, we assume that $F_pM \subset F_{p+1}M$ for all p, and that $\cup_p F_pM = M$), such that $F_lD_X \cdot F_pM \subset F_{l+p}M$ for all l and p. The following theory, of course, works for more general filtered rings (not just D_X), as well as for more general filtered modules where we allow nonzero F_pM for $p < 0$ (however, F_pM must be zero for sufficiently negative p).

If M is a filtered D_X-module, its associated graded module, $\mathrm{gr}^F M = \oplus_{p=0}^{\infty} F_pM/F_{p-1}M$, is a module over $\mathrm{gr}\,D_X$. If $\mathrm{gr}^F M$ is a *finitely generated* $\mathrm{gr}\,D_X$-module, we say that F is a *good filtration* on M.

Proposition. *There exists a good filtration F on a left D_X-module M if and only if M is finitely generated over D_X.*

Proof. The "if" direction is easy: given a finite set of generators of M over D_X, it is clear how to define a filtration such that the classes of these generators in $F_0M = F_0M/F_{-1}M \subset \mathrm{gr}^F M$ generate $\mathrm{gr}^F M$ over $\mathrm{gr}\,D_X$. For the "only if" direction, let $m_1 \in F_{p_1}M, \ldots, m_k \in F_{p_k}M$ be such that the classes $\overline{m_i} \in F_{p_i}M/F_{p_i-1}M$ generate $\mathrm{gr}^F M$ over $\mathrm{gr}\,D_X$ (just pick any finite set of generators for $\mathrm{gr}^F M$ and split each generator up into its homogeneous components). We claim that $m_1, \ldots, m_k$ generate M over D_X. It suffices to show that $F_pM \subset \sum_i D_X \cdot m_i$ for all p. We use induction on p; since $F_pM = 0$ for negative p, the base case is obvious. Let p and $m \in F_pM \setminus F_{p-1}M$ be given, and assume the statement for smaller values of p. By assumption, the class $\overline{m} \in F_pM/F_{p-1}M$ can be written $\overline{m} = \sum_i \overline{\delta_i}\,\overline{m_i}$ where $\delta_i \in \mathrm{gr}\,D_X$. This equality still holds if we replace $\overline{\delta_i}$ by its homogeneous component of degree $p - p_i$. After doing so, choose lifts $\delta_i \in D_X$ of $\overline{\delta_i}$, and apply the induction hypothesis to $m - \sum_i \delta_i m_i \in F_{p-1}M$. □

The proof of the "only if" direction shows that $F_pM = \sum_{p\ge p_i}(F_{p-p_i}D_X)\cdot m_i$ for all p. That is, all good filtrations arise from shifts of the filtration on $\mathcal{D}_X$ after choosing generators for M. This fact can be used to compare two good filtrations on the same module.

Proposition. *Let M be a finitely generated left D_X-module. Let F and G be good filtrations on M. There exists an integer a such that $F_{p-a}M \subset G_pM \subset F_{p+a}M$ for all p.*

Proof. By symmetry, it suffices to assume only that F is good and to show the first containment. By the previous result, there exist $m_1, \dots, m_k \in M$ and $p_1, \dots, p_k \geq 0$ such that $F_pM = \sum_{p \geq p_i}(F_{p-p_i}D_X) \cdot m_i$ for all p. Since $\cup_q G_qM = M$, for all i we can choose q_i such that $m_i \in G_{q_i}M$. Let $a = \max\{q_i - p_i\}$. Then we have

$$F_pM = \sum_{p \geq p_i}(F_{p-p_i}D_X) \cdot m_i \subset \sum_{p \geq p_i} F_{p-p_i}D_X \cdot G_{q_i}M \subset \sum_{p \geq p_i} G_{p-p_i+q_i}M \subset G_{p+a}M,$$

as claimed (up to a shift). □

We say that the filtrations F and G are *neighboring* if the integer a in the proposition can be taken to be 1. The proposition implies that any two good filtrations on a left D_X-module can be connected by a chain of pairs of neighboring filtrations, which will be useful later in proofs.

The fact that $\operatorname{gr} D_X$ is a commutative ring can be used to reduce proofs of properties of the non-commutative ring D_X to proofs involving its commutative approximation $\operatorname{gr} D_X$. Perhaps the easiest example of this strategy is the following:

Proposition. *The ring D_X is left and right Noetherian.*

Proof. Let $I \subset D_X$ be a left ideal. Make I into a filtered D_X-module by setting $F_lI = I \cap F_lD_X$ for all l. Then $\operatorname{gr} I$ is an ideal in the Noetherian commutative ring $\operatorname{gr} D_X$, hence is finitely generated; it follows that I is finitely generated over D_X. The proof for right ideals is exactly the same. □

Another result (whose proof is more involved) about D_X that is proved by reducing everything to the setting of the commutative $\operatorname{gr} D_X$ involves global dimensions. Recall that the *left (resp. right) global dimension* of a ring A is the supremum of the set of projective dimensions of left (resp. right) A-modules. If A is left and right Noetherian, its left and right global dimensions coincide, and we speak simply of its global dimension. The global dimension of the commutative ring $\operatorname{gr} D_X$ is $2n$.

Proposition. *The global dimension of D_X is $\leq 2n$.*

(In fact, it is exactly $2n$, but this requires much more work to prove.) The idea behind the proof of this weaker result is the following. It suffices to prove that $\operatorname{Ext}^{2n+1}_{D_X}(M, N) = 0$ for all finitely generated left D_X-modules M and N. Fix good filtrations on M and N. We know that $\operatorname{Ext}^{2n+1}_{\operatorname{gr} D_X}(\operatorname{gr} M, \operatorname{gr} N) = 0$, because the global dimension of $\operatorname{gr} D_X$ is $2n$. There is a filtration on $\operatorname{Ext}^{2n+1}_{D_X}(M, N)$ such that $\operatorname{gr} \operatorname{Ext}^{2n+1}_{D_X}(M, N)$ is isomorphic to a subquotient of $\operatorname{Ext}^{2n+1}_{\operatorname{gr} D_X}(\operatorname{gr} M, \operatorname{gr} N) = 0$, which completes the proof. The key to this last step is showing that for any good filtered D_X-module M, there is a resolution $F_\bullet \to M$ by filtered finite free D_X-modules that descends to a resolution $\operatorname{gr} F_\bullet \to \operatorname{gr} M$.

Week 6: resolutions; derived pullback and pushforward

Beginning now, we add to our list of permanent assumptions about our scheme X (and other schemes Y, Z that we will map to and from X) that it be *separated* and *finite type* over k. In particular, X is now quasi-compact. What is more, X has the *resolution property*: any coherent $\mathcal{O}_X$-module is a quotient of a locally free $\mathcal{O}_X$-module. Totaro proved that for a smooth scheme X of finite type over k, the resolution property is equivalent to X having affine diagonal (a weaker condition than separated). HTT assume that X is quasi-projective, in which case the resolution property is easy to see.

If $M \in \operatorname{Mod}(\mathcal{D}_X)$, M has an injective resolution $M \to I^\bullet$ and a flat resolution $F^\bullet \to M$ by left $\mathcal{D}_X$-modules: this is a general fact for left modules over any sheaf of rings. Last week, we sketched a proof that the ring of sections $\mathcal{D}_X(U)$ has global dimension $\leq 2n$ for any open $U \subset X$.

It follows that M has *bounded* injective and flat resolutions. Suppose furthermore that M is a quasi-coherent $\mathcal{D}_X$-module. If $F \twoheadrightarrow M$ is an $\mathcal{O}_X$-linear surjection from a locally free $\mathcal{O}_X$-module, then $\mathcal{D}_X \otimes_{\mathcal{O}_X} F \twoheadrightarrow \mathcal{D}_X \otimes_{\mathcal{O}_X} M \twoheadrightarrow M$ is a $\mathcal{D}_X$-linear surjection, and $\mathcal{D}_X \otimes_{\mathcal{O}_X} F$ is a locally free $\mathcal{D}_X$-module. It follows that every $M \in \mathrm{Mod}_{qc}(\mathcal{D}_X)$ has a resolution by locally free $\mathcal{D}_X$-modules, and therefore a bounded resolution by locally projective $\mathcal{D}_X$-modules (using the finiteness of global dimension).

If M is *coherent* as a $\mathcal{D}_X$-module, it has a resolution by locally free $\mathcal{D}_X$-modules of finite rank. To see this, we must replace M in the proof above by a coherent $\mathcal{O}_X$-submodule $M' \subset M$ that generates M (globally) over $\mathcal{D}_X$. Such a thing exists by the following argument: take a finite open affine cover $\{U_i\}$ of X such that $M|_{U_i}$ is generated over $\mathcal{D}_{U_i}$ by a coherent $\mathcal{O}_{U_i}$-submodule M_i'' (such a cover exists by the definition of coherence), extend each M_i'' to a coherent $\mathcal{O}_X$-submodule M_i' of M, then simply take $M' = \sum_i M_i'$.

We now introduce derived categories of $\mathcal{D}_X$-modules, the correct setting for the inverse and direct image operations. Recall that the derived category $D(\mathcal{D}_X) = D(\mathrm{Mod}(\mathcal{D}_X))$ is obtained by taking the Abelian category of complexes of left $\mathcal{D}_X$-modules, forming its quotient by chain homotopy equivalences, and finally inverting all quasi-isomorphisms (maps of complexes inducing isomorphisms on all cohomology objects). A single $\mathcal{D}_X$-module M is viewed as an object in this category by considering the complex which is M in degree zero and 0 elsewhere (its sole nonzero cohomology object is $H^0(M) = M$). The category $D(\mathcal{D}_X)$ is no longer Abelian; its consolation prize is a *triangulated category* structure. A morphism $M^\bullet \to N^\bullet$ in $D(\mathcal{D}_X)$ need not be induced by a single map of complexes: instead, such morphisms are equivalence classes of "roofs" $M^\bullet \leftarrow P^\bullet \to N^\bullet$ where the map $P^\bullet \to M^\bullet$ (but not necessarily the other map) is a quasi-isomorphism. Standard references on derived categories are chapter 1 of Hartshorne's *Residues and Duality* and chapter 10 of Weibel's *Introduction to Homological Algebra*. Appendices B and C of HTT include an excellent summary of the theory, but with most proofs omitted.

Variants on the basic derived category $D(\mathcal{D}_X)$ include $D^b(\mathcal{D}_X)$ (resp. $D^+(\mathcal{D}_X), D^-(\mathcal{D}_X)$), where only bounded (resp. bounded below, bounded above) complexes are considered, and $D^*_{qc}(\mathcal{D}_X)$ (resp. $D^*_c(\mathcal{D}_X)$), where only complexes with quasi-coherent (resp. coherent) cohomology objects are considered (here $*$ stands for b, $+$, $-$, or no superscript). It follows from the discussion above and standard derived category techniques that every object of $D^b(\mathcal{D}_X)$ is represented by a bounded complex of flats and a bounded complex of injectives, and every object of $D^b_{qc}(\mathcal{D}_X)$ is represented by a bounded complex of locally-projectives.

Let $f : X \to Y$ be any morphism. Recall that we associated with f a pair of *transfer bimodules*:

$$\mathcal{D}_{X\to Y} = f^*\mathcal{D}_Y = \mathcal{O}_X \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}\mathcal{D}_Y$$

is a $(\mathcal{D}_X, f^{-1}\mathcal{D}_Y)$-bimodule, where the left $\mathcal{D}_X$-action is the "chain-rule-type" action defined during Week 4, and the right $f^{-1}\mathcal{D}_Y$-action is just right multiplication on the second tensor factor; by applying side-changing operations, we obtain

$$\mathcal{D}_{Y\leftarrow X} = \omega_X \otimes_{\mathcal{O}_X} \mathcal{D}_{X\to Y} \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}\omega_Y^{\otimes -1},$$

which is a $(f^{-1}\mathcal{D}_Y, \mathcal{D}_X)$-bimodule.

Definition. *The* inverse image functor $\mathbb{L}f^* : D^b(\mathcal{D}_Y) \to D^b(\mathcal{D}_X)$ *is defined by* $\mathbb{L}f^*(M^\bullet) = \mathcal{D}_{X\to Y} \otimes^{\mathbb{L}}_{f^{-1}\mathcal{D}_Y} f^{-1}M^\bullet$*, and the* direct image functor $\int_f : D^b(\mathcal{D}_X) \to D^b(\mathcal{D}_Y)$ *is defined by* $\int_f M^\bullet = \mathbb{R}f_*(\mathcal{D}_{Y\leftarrow X} \otimes^{\mathbb{L}}_{\mathcal{D}_X} M^\bullet)$.

Observe that these are the same as the "naive" operations defined in Week 4, except that all occurrences of tensor product or pushforward have been replaced with their derived versions. (There is a variant of $\int_f$ defined for derived categories of *right* modules.) The functor f^{-1} is exact, so does not need to be derived. To compute $\mathbb{L}f^*(M^\bullet)$, we replace $M^\bullet$ with a bounded flat resolution $F^\bullet$ and then form the tensor product $\mathcal{D}_{X\to Y} \otimes_{f^{-1}\mathcal{D}_Y} f^{-1}F^\bullet$. (Alternatively, we could replace $\mathcal{D}_{X\to Y}$ with a resolution by flat *right* $f^{-1}\mathcal{D}_Y$-modules.) Since the direct image functor is a composition of a left derived functor and a right derived functor, we would in general need first to replace $M^\bullet$ with a bounded flat resolution $F^\bullet$, and then replace $\mathcal{D}_{Y\leftarrow X} \otimes_{\mathcal{D}_X} F^\bullet$ with a bounded injective resolution before applying f_*.

As a complex of $\mathcal{O}_X$-modules, $\mathbb{L}f^*M^\bullet$ is naturally isomorphic to $\mathcal{O}_X \otimes^{\mathbb{L}}_{f^{-1}\mathcal{O}_Y} f^{-1}M^\bullet$, using associativity of the derived tensor product. Since quasi-coherence of a $\mathcal{D}_X$-module simply means quasi-coherence of the underlying $\mathcal{O}_X$-module, it is fairly straightforward to see that $\mathbb{L}f^*$ preserves quasi-coherence, that is, restricts to a functor $D^b_{qc}(\mathcal{D}_Y) \to D^b_{qc}(\mathcal{D}_X)$. However, $\mathbb{L}f^*$ does not, in general, preserve coherence: we have $\mathbb{L}f^*\mathcal{D}_Y = \mathcal{D}_{X\to Y}$, and we have seen before that if f is a nontrivial closed immersion, $\mathcal{D}_{X\to Y}$ is a locally free left $\mathcal{D}_X$-module of *infinite* rank. Finally, we have a *composition rule* for the inverse image: if $X \xrightarrow{f} Y \xrightarrow{g} Z$, then the functors $\mathbb{L}(g\circ f)^*$ and $\mathbb{L}f^*\circ\mathbb{L}g^*$ from $D^b(\mathcal{D}_Z)$ to $D^b(\mathcal{D}_X)$ are naturally isomorphic. This is easy to prove using associativity of the derived tensor product and the fact that sheaf-theoretic inverse image (f^{-1}) commutes with tensor product (note that since $\mathcal{D}_Y$ is locally free over $\mathcal{O}_Y$, $f^{-1}\mathcal{D}_Y$ is certainly flat over $f^{-1}\mathcal{O}_Y$, and so $\mathcal{D}_{X\to Y} = \mathcal{O}_X \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}\mathcal{D}_Y = \mathcal{O}_X \otimes^{\mathbb{L}}_{f^{-1}\mathcal{O}_Y} f^{-1}\mathcal{D}_Y$).

The analogous statements for the direct image $\int_f$ are all true, but with different and more involved proofs. There is a composition rule ($\int_{g\circ f}$ and $\int_g \circ \int_f$ are naturally isomorphic) for which it is necessary to use the derived version of the direct image, whereas the composition rule for inverse images is even true for the "naive" version. The direct image does not preserve coherence, even for open immersions; however, it preserves coherence if f is proper, in particular for closed immersions. Finally, direct image does preserve quasi-coherence: $\int_f$ restricts to a functor $D^b_{qc}(\mathcal{D}_X) \to D^b_{qc}(\mathcal{D}_Y)$. However, it is not clear that $\mathcal{D}_{Y\leftarrow X} \otimes^{\mathbb{L}}_{\mathcal{D}_X} M^\bullet$ has an $\mathcal{O}_X$-module structure, and so this proof does not immediately reduce to a proof for the $\mathcal{O}$-module categories. Instead, we will factor a general morphism f into manageable pieces.

In fact, this strategy is how we will understand direct and inverse images more generally. The preceding discussion is about as far as we will go with arbitrary morphisms f. Given such a morphism $f : X \to Y$, we can factor it as

$$X \xrightarrow{\Gamma_f} X \times Y \xrightarrow{p_2} Y,$$

where the first map is the *graph* $\Gamma_f = (\mathrm{id}_X, f)$ of f, and the second map is projection on the second factor. Since X and Y are smooth, p_2 is smooth. Since X and Y are separated, Γ_f, which is a base change of the diagonal Δ_f, is a closed immersion. We therefore focus our attention on the special cases of closed immersions and projections (more generally, smooth morphisms). Closed immersions will be given pride of place because of their importance for Kashiwara's theorem. We remark that, by the composition rule, it will suffice to show that $\int_f$ preserves quasi-coherence in case f is a closed immersion or a projection in order to conclude it for arbitrary f.

We first indicate what happens in the easiest case of all: when $j : U \hookrightarrow X$ is an open immersion. In this case, $j^{-1}\mathcal{D}_X$ is just $\mathcal{D}_U$, from which it follows that $\mathcal{D}_{U\to X}$ and $\mathcal{D}_{X\leftarrow U}$ are both simply $\mathcal{D}_U$. Therefore $\mathbb{L}j^*$ is just j^{-1}, and $\int_j$ is just $\mathbb{R}j_*$.

Next we consider closed immersions, which will occupy us for some time. Let $i : X \hookrightarrow Y$ be a closed immersion, where $\dim X = r$ and $\dim Y = n$. As we saw in Week 4, we can choose local coordinates $\{y_j, \partial_{y_j}\}_{j=1}^n$ on Y such that $y_{r+1} = \cdots = y_n = 0$ are local defining equations for the immersion i. Write $x_j = y_j \circ i$ for $j = 1, \ldots, r$: then $\{x_j, \partial_{x_j}\}_{j=1}^r$ are local coordinates on X. In these coordinates, $\mathcal{D}_{X\to Y} \simeq \mathcal{D}_X \otimes_k k[\partial_{r+1}, \ldots, \partial_n]$ as left $\mathcal{D}_X$-modules. Recall that $\mathcal{D}_{X\to Y}$

is a $(\mathcal{D}_X, i^{-1}\mathcal{D}_Y)$-bimodule. As sheaves of k-spaces, $\mathcal{D}_{X\to Y} \simeq \mathcal{D}_{Y\leftarrow X}$ (using the simultaneous trivializations of ω_X and ω_Y by $dx_1 \wedge \cdots \wedge dx_r$ and $dy_1 \wedge \cdots \wedge dy_n$); the right $\mathcal{D}_X$- and left $i^{-1}\mathcal{D}_Y$-actions on $\mathcal{D}_{Y\leftarrow X}$ are the transposes (formal adjoints) of those on $\mathcal{D}_{X\to Y}$.

In particular, we observe that $\mathcal{D}_{Y\leftarrow X}$ is a locally free *right* $\mathcal{D}_X$-module. The definition of direct image $\int_i$ involves a tensor product over $\mathcal{D}_X$ with $\mathcal{D}_{Y\leftarrow X}$. In contrast, the definition of inverse image $\mathbb{L}i^*$ involves a tensor product over $i^{-1}\mathcal{D}_Y$ with $\mathcal{D}_{X\to Y}$, and the latter is *not* a locally free left $i^{-1}\mathcal{D}_Y$-module. Therefore the *direct* image is actually simpler in this case. Since i is affine, i_* is exact, and therefore both derived functors occurring in the definition of $\int_i$ can be replaced with their non-derived versions. If $M \in \mathrm{Mod}(\mathcal{D}_X)$ (a single $\mathcal{D}_X$-module), the complex $\int_i M$ therefore has no cohomology in nonzero degrees, and the functor

$$\int_i^0 M = H^0(\int_i M) = i_*(\mathcal{D}_{Y\leftarrow X} \otimes_{\mathcal{D}_X} M)$$

is exact. Observe that this functor clearly preserves quasi-coherence, since $\mathcal{D}_{Y\leftarrow X}$ is locally free as a right $\mathcal{D}_X$-module. On the other hand, $\mathbb{L}i^*M$ may have nontrivial cohomology in nonzero degrees.

Week 7: pushforward for closed immersions; summary for smooth morphisms

Last time, we saw that if $i : X \hookrightarrow Y$ is a closed immersion, then the functor $\int_i^0 : \mathrm{Mod}(\mathcal{D}_X) \to \mathrm{Mod}(\mathcal{D}_Y)$ defined by $\int_i^0 M = i_*(\mathcal{D}_{Y\leftarrow X} \otimes_{\mathcal{D}_X} M)$ is exact (and preserves quasi-coherence). We write $\mathrm{Mod}^X_{qc}(\mathcal{D}_Y)$ for the category of all quasi-coherent left $\mathcal{D}_Y$-modules supported on X: clearly $\int_i^0 M \in \mathrm{Mod}^X_{qc}(\mathcal{D}_Y)$.

Theorem. *(Kashiwara) The functor $\int_i^0 : \mathrm{Mod}_{qc}(\mathcal{D}_X) \to \mathrm{Mod}^X_{qc}(\mathcal{D}_Y)$ is an equivalence of categories. Indeed, $\int_i^0$ possesses a right adjoint $i^\natural$ such that $(\int_i^0, i^\natural)$ is an adjoint equivalence.*

We will construct the right adjoint (and prove it is a right adjoint), leaving for next week the proof of equivalence. In fact, $i^\natural N$ will be defined for *any* left $\mathcal{D}_Y$-module N, and the two will be *adjoint* as functors $\mathrm{Mod}(\mathcal{D}_X) \leftrightarrow \mathrm{Mod}(\mathcal{D}_Y)$. Given any such N, we define

$$i^\natural N = \mathcal{H}om_{i^{-1}\mathcal{D}_Y}(\mathcal{D}_{Y\leftarrow X}, i^{-1}N),$$

which is naturally a left $i^{-1}\mathcal{D}_Y$-module and can be viewed as a left $\mathcal{D}_X$-module using the right $\mathcal{D}_X$-structure on $\mathcal{D}_{Y\leftarrow X}$. The functor $i^\natural : \mathrm{Mod}(\mathcal{D}_Y) \to \mathrm{Mod}(\mathcal{D}_X)$, being the composition of a sheaf $\mathcal{H}om$ and the exact functor i^{-1}, is left exact. A more sophisticated version of Kashiwara's equivalence (stated in terms of a triangulated equivalence between derived categories) uses the right derived functor $\mathbb{R}i^\natural$ of this left exact functor. We remark here that if $N^\bullet \in D^b(\mathcal{D}_Y)$, we have $\mathbb{R}i^\natural N^\bullet \simeq \mathbb{L}i^* N^\bullet[\dim X - \dim Y]$ in $D^b(\mathcal{D}_X)$. (The proof of this fact is an explicit calculation using a locally free left $i^{-1}\mathcal{D}_Y$-resolution of $\mathcal{D}_{Y\leftarrow X}$.)

Suppose $\psi : \mathcal{D}_{Y\leftarrow X} \to i^{-1}N$ is an $i^{-1}\mathcal{D}_Y$-linear map. In local coordinates $\{y_i, \partial_i\}$ on Y, we have, as discussed last time, an isomorphism $\mathcal{D}_{Y\leftarrow X} \simeq \mathcal{D}_X \otimes_k k[\partial_{r+1}, \ldots, \partial_n]$ as k-spaces (here $r = \dim X$ and $n = \dim Y$). The left $i^{-1}\mathcal{D}_Y$- and right $\mathcal{D}_X$-actions are transposes of those on $\mathcal{D}_{X\to Y}$. As a left $i^{-1}\mathcal{D}_Y$-module, $\mathcal{D}_{Y\leftarrow X}$ is generated by $1 \otimes 1$. Let $\mathcal{I} \subset \mathcal{O}_Y$ be the defining ideal sheaf of $i : X \hookrightarrow Y$ (locally generated by $y_{r+1}, \ldots, y_n$). Then $i^{-1}\mathcal{I}$ annihilates $1 \otimes 1$, because it annihilates the left tensor factor ($y_{r+1}, \ldots, y_n$ all act as zero on $\mathcal{D}_X$). Since ψ is $i^{-1}\mathcal{D}_Y$-linear, we have $i^{-1}\mathcal{I} \cdot \psi(1 \otimes 1) = 0$. It follows that the image of ψ lies in $i^{-1}\underline{\Gamma}_X N$, where $\underline{\Gamma}_X N$ is the subsheaf of sections of N supported on X.

This observation is crucial for the proof that $i^\natural$ is right adjoint to $\int_i^0$, which we give now. Let $M \in \mathrm{Mod}(\mathcal{D}_X)$ and $N \in \mathrm{Mod}(\mathcal{D}_Y)$ be given. We have functorial $\mathcal{D}_Y$-module isomorphisms

$$\begin{aligned} i_*\mathcal{H}om_{\mathcal{D}_X}(M, \mathcal{H}om_{i^{-1}\mathcal{D}_Y}(\mathcal{D}_{Y\leftarrow X}, i^{-1}N)) &\simeq i_*\mathcal{H}om_{\mathcal{D}_X}(M, \mathcal{H}om_{i^{-1}\mathcal{D}_Y}(\mathcal{D}_{Y\leftarrow X}, i^{-1}\underline{\Gamma}_X N)) \\ &\simeq i_*\mathcal{H}om_{i^{-1}\mathcal{D}_Y}(\mathcal{D}_{Y\leftarrow X} \otimes_{\mathcal{D}_X} M, i^{-1}\underline{\Gamma}_X N) \\ &\simeq \mathcal{H}om_{\mathcal{D}_Y}(i_*(\mathcal{D}_{Y\leftarrow X} \otimes_{\mathcal{D}_X} M), \underline{\Gamma}_X N) \\ &\simeq \mathcal{H}om_{\mathcal{D}_Y}(i_*(\mathcal{D}_{Y\leftarrow X} \otimes_{\mathcal{D}_X} M), N), \end{aligned}$$

where the first and fourth isomorphisms follow from the previous paragraph, the second isomorphism is a form of $\otimes - \mathcal{H}om$ adjunction, and the third isomorphism uses the full faithfulness of i_* as well as the identification $i_*i^{-1}\underline{\Gamma}_X N \simeq \underline{\Gamma}_X N$ (if we begin with a sheaf supported on X, pulling back to X and then pushing forward to Y changes nothing). If we take global sections of both sides, we obtain functorial bijective correspondences

$$\begin{aligned} \mathrm{Hom}_{\mathcal{D}_X}(M, i^\natural N) &= \Gamma(X, \mathcal{H}om_{\mathcal{D}_X}(M, \mathcal{H}om_{i^{-1}\mathcal{D}_Y}(\mathcal{D}_{Y\leftarrow X}, i^{-1}N))) \\ &= \Gamma(Y, i_*\mathcal{H}om_{\mathcal{D}_X}(M, \mathcal{H}om_{i^{-1}\mathcal{D}_Y}(\mathcal{D}_{Y\leftarrow X}, i^{-1}N))) \\ &\simeq \Gamma(Y, \mathcal{H}om_{\mathcal{D}_Y}(i_*(\mathcal{D}_{Y\leftarrow X} \otimes_{\mathcal{D}_X} M), N)) \\ &= \mathrm{Hom}_{\mathcal{D}_Y}(\int_i^0 M, N), \end{aligned}$$

so that $(\int_i^0, i^\natural)$ form an adjoint pair, as claimed.

Recall that our strategy for studying the inverse and direct image functors along general morphisms was to factor such morphisms into closed immersions followed by smooth morphisms (specifically, projections) using the graph. Before proving Kashiwara's theorem and discussing its consequences next time, we briefly sketch what happens in the smooth case.

Let $f : X \to Y$ be a smooth morphism. If $M \in \mathrm{Mod}(\mathcal{D}_Y)$, then $\mathbb{L}f^*M \simeq \mathcal{O}_X \otimes^{\mathbb{L}}_{f^{-1}\mathcal{O}_Y} f^{-1}M$ in $D^b(\mathcal{O}_X)$. Since f is smooth, it is in particular flat, so $\mathcal{O}_X$ is flat over $f^{-1}\mathcal{O}_Y$. It follows that $\mathbb{L}f^*M$ has cohomology only in degree zero (that is, up to identifying a left $\mathcal{D}_X$-module with a complex concentrated in degree zero, we simply have $\mathbb{L}f^* = f^*$).

To see what happens for the direct image, we use de Rham complexes. Recall that the (absolute) *de Rham complex* $\Omega_X^\bullet$ on X takes the form

$$0 \to \mathcal{O}_X \xrightarrow{d} \Omega_X^1 \to \cdots \to \Omega_X^i \xrightarrow{d^i} \Omega_X^{i+1} \to \cdots \to \Omega_X^n \to 0,$$

where $\Omega_X^i = \bigwedge^i \Omega_X^1$, the map d is the universal derivation, and the maps d^i, called *exterior derivatives*, are induced by d. The objects in this complex are coherent $\mathcal{O}_X$-modules, but the maps are merely k-linear. There is a relative version of this complex: if r is the relative dimension $\dim X - \dim Y$ of the smooth morphism f, then by replacing $\Omega_X^1 = \Omega_{X/k}^1$ with $\Omega_{X/Y}^1$, we obtain a complex $\Omega_{X/Y}^\bullet$ of length r whose objects are coherent $\mathcal{O}_X$-modules but whose maps are $f^{-1}\mathcal{O}_Y$-linear.

If M is an $\mathcal{O}_X$-module, a *connection* on M is a k-linear map $\nabla : M \to \Omega_X^1 \otimes_{\mathcal{O}_X} M$ such that $\nabla(fm) = df \otimes m + f\nabla(m)$ for all $m \in M$ and $f \in \mathcal{O}_X$. The datum of a connection on M is equivalent to that of a k-linear map $\Theta_X \to \mathcal{E}nd_k(M)$ ($\delta \mapsto \nabla_\delta$) such that $f\nabla_\delta = \nabla_{f\delta}$ and $\nabla_\delta f = f\nabla_\delta + \delta(f)$. Recall from Week 3 that this means M is "two-thirds of the way to being a $\mathcal{D}_X$-module": if we want to extend the $\mathcal{O}_X$-structure on M to a $\mathcal{D}_X$-structure by specifying how the derivations $\delta \in \Theta_X$ act, these are two of the three required properties. If the map $\Theta_X \to \mathcal{E}nd_k(M)$ satisfies the third property (that $\nabla_{[\delta_1,\delta_2]} = [\nabla_{\delta_1}, \nabla_{\delta_2}]$ for all $\delta_1, \delta_2 \in \Theta_X$), the connection ∇ is called *integrable* (or *flat*), and we see that an integrable connection on M and a left $\mathcal{D}_X$-module structure on M amount to the same thing. Given a connection ∇ on M, the maps $\nabla^i : \Omega^i_X \otimes_{\mathcal{O}_X} M \to \Omega^{i+1}_X \otimes_{\mathcal{O}_X} M$ defined by $\nabla^i(\omega \otimes m) = d^i\omega \otimes m - \omega \wedge \nabla(m)$ form a complex

$$\mathrm{DR}_X(M) = (0 \to M \xrightarrow{\nabla} \Omega^1_X \otimes_{\mathcal{O}_X} M \to \cdots \to \Omega^n_X \otimes_{\mathcal{O}_X} M \to 0)$$

if and only if ∇ is integrable. Therefore, for any left $\mathcal{D}_X$-module M, we can build the complex $\mathrm{DR}_X(M)$, which is called the *de Rham complex of* M. Likewise, we can define the relative version $\mathrm{DR}_{X/Y}(M)$. The complex $\mathrm{DR}_X(\mathcal{D}_X)[-n]$ is a locally free resolution of ω_X as a right $\mathcal{D}_X$-module, and the complex $\mathrm{DR}_{X/Y}(\mathcal{D}_X)[-r]$ is a locally free resolution of $\mathcal{D}_{Y\leftarrow X}$ as a left $f^{-1}\mathcal{D}_Y$-module. Therefore, if $M \in \mathrm{Mod}(\mathcal{D}_X)$, we have $\int_f M = \mathbb{R}f_*(\mathrm{DR}_{X/Y}(M)[-r])$ in $D^b(\mathcal{D}_Y)$. From this, we see at once that $\int_f$ preserves quasi-coherence (the objects in the complex $\mathrm{DR}_{X/Y}(M)$ are quasi-coherent if M is), so by our factoring argument, $\int_f$ descends to a functor $D^b_{qc}(\mathcal{D}_X) \to D^b_{qc}(\mathcal{D}_Y)$ for *any* morphism f.

Week 8: Kashiwara's theorem

Recall that if $i : X \hookrightarrow Y$ is a closed immersion, we have the exact functor $\int_i^0 : \mathrm{Mod}(\mathcal{D}_X) \to \mathrm{Mod}(\mathcal{D}_Y)$ defined by $\int_i^0 M = i_*(\mathcal{D}_{Y\leftarrow X} \otimes_{\mathcal{D}_X} M)$, and the left exact functor $i^\natural : \mathrm{Mod}(\mathcal{D}_Y) \to \mathrm{Mod}(\mathcal{D}_X)$ defined by $i^\natural N = \mathcal{H}om_{i^{-1}\mathcal{D}_Y}(\mathcal{D}_{Y\leftarrow X}, i^{-1}N)$. We saw last time that $(\int_i^0, i^\natural)$ form an adjoint pair. Today, we will prove *Kashiwara's theorem*, stated last time: if we restrict these functors to

$$\int_i^0 : \mathrm{Mod}_{qc}(\mathcal{D}_X) \to \mathrm{Mod}^X_{qc}(\mathcal{D}_Y), \qquad i^\natural : \mathrm{Mod}^X_{qc}(\mathcal{D}_Y) \to \mathrm{Mod}_{qc}(\mathcal{D}_X)$$

(where the superscript X means those $\mathcal{D}_Y$-modules supported on X), we obtain an (adjoint) equivalence of categories. The same is true with *qc* replaced by *c*; one must simply check that $\int_i^0$ and $i^\natural$ preserve coherence (we already know they preserve quasi-coherence).

Proof. Since the functors form an adjoint pair, it suffices to show that the *unit* $M \to i^\natural \int_i^0 M$ and *counit* $\int_i^0 i^\natural N \to N$ are isomorphisms for all $M \in \mathrm{Mod}_{qc}(\mathcal{D}_X)$ and $N \in \mathrm{Mod}^X_{qc}(\mathcal{D}_Y)$. By the composition rule, we may factor i into a sequence of codimension-one closed immersions and thus assume that i itself is of codimension one. To prove that two sheaves are isomorphic is a local question, so we may assume that we have coordinates $\{y_i, \partial_i\}_{i=1}^n$ on Y such that $y_n = 0$ is a defining equation for X. We write y for y_n and ∂ for ∂_n. In these coordinates, the transfer bimodule $\mathcal{D}_{Y\leftarrow X}$ takes the form $k[\partial] \otimes_k \mathcal{D}_X$. The canonical bundles of X and Y have been simultaneously trivialized by these coordinates, so $\mathcal{D}_{Y\leftarrow X} = \mathcal{O}_X \otimes_{i^{-1}\mathcal{O}_Y} i^{-1}\mathcal{D}_Y$, with left $i^{-1}\mathcal{D}_Y$-action given by the transpose of the obvious right action.

We first consider $i^\natural N$ for arbitrary $N \in \mathrm{Mod}(\mathcal{D}_Y)$. Since $\mathcal{O}_X$ is just $i^{-1}\mathcal{O}_Y/(y)$, we see that $i^\natural N = \mathcal{H}om_{i^{-1}\mathcal{D}_Y}(\mathcal{O}_X \otimes_{i^{-1}\mathcal{O}_Y} i^{-1}\mathcal{D}_Y, i^{-1}N)$ is just the kernel of $y \in i^{-1}\mathcal{D}_Y$ acting on $i^{-1}N$. Let $M \in \mathrm{Mod}_{qc}(\mathcal{D}_X)$ be given. Then

$$\int_i^0 M = i_*(\mathcal{D}_{Y\leftarrow X} \otimes_{D_X} M) \simeq i_*((k[\partial] \otimes_k \mathcal{D}_X) \otimes_{\mathcal{D}_X} M) \simeq i_*(k[\partial] \otimes_k M) = k[\partial] \otimes_k i_*M,$$

so $i^\natural \int_i^0 M$ is the kernel of y acting on $i^{-1}(k[\partial] \otimes_k i_*M) = k[\partial] \otimes_k M$. It is easy to see that the unit $M \to i^\natural \int_i^0 M$ sends m to $1 \otimes m$ under this identification. We claim that the kernel is precisely $1 \otimes_k M$, from which it will follow that the unit is an isomorphism. Write the left $i^{-1}\mathcal{D}_Y$-module $k[\partial] \otimes_k M$ as $\oplus_{j\geq 0}\partial^j \cdot M$, where the powers of ∂ appear on the left because the action is transposed. For any $j \geq 0$ and any $m \in M$, we have

$$y\partial^j m = \partial^j ym - j\partial^{j-1}m = -j\partial^{j-1}m;$$

that is, $ym = 0$, but powers of ∂ can "absorb" powers of y before they reach m. It follows that $y\partial^j m = 0$ if and only if $j = 0$, that is, the kernel of y is $1 \otimes_k M \subset k[\partial] \otimes_k M$. (*Caveat lector*: in Week 6, I stated incorrectly in the seminar that y annihilates all of $k[\partial] \otimes_k M$! This has now been fixed in the posted notes.)

We now consider the counit. Let $N \in \mathrm{Mod}_{qc}^X(\mathcal{D}_Y)$. Let θ be the linear operator $y\partial : N \to N$, and let $N^j = \{n \in N | \theta n = jn\}$, for $j \in \mathbb{Z}$, be the j-eigenspace of this operator. Note that $\theta : N^j \to N^j$ is an isomorphism (multiplication by j) if $j \neq 0$, and $\partial y = \theta + 1 : N^j \to N^j$ is an isomorphism (multiplication by $j+1$) if $j \neq -1$. Let $n \in N^j$, for any j, be given. The calculation

$$\theta(yn) = y\partial yn = y(y\partial + 1)n = y\theta n + yn = y(jn) + yn = (j+1)yn$$

shows that $yn \in N^{j+1}$, that is, $yN^j \subset N^{j+1}$. A similar calculation shows that $\partial N^j \subset N^{j-1}$. Therefore, if $j < -1$, the isomorphisms $\theta : N^{j+1} \to N^{j+1}$ and $\theta + 1 : N^j \to N^j$ can be factored $N^{j+1} \xrightarrow{\partial} N^j \xrightarrow{y} N^{j+1}$ and $N^j \xrightarrow{y} N^{j+1} \xrightarrow{\partial} N^j$, and so $\partial : N^{j+1} \to N^j$ and $y : N^j \to N^{j+1}$ are *both* isomorphisms for such j.

We claim that $N = \oplus_{j>0} N^{-j}$ as k-spaces. Before proving this claim, we show how the remaining part of Kashiwara's theorem follows from it. Assuming that $N = \oplus_{j>0} N^{-j}$, we have $N = k[\partial] \otimes_k N^{-1}$ (since $\partial : N^{-j} \to N^{-j-1}$ is an isomorphism for all $j > 0$) and $i^\natural N$, the kernel of y acting on $i^{-1}N = \oplus_{j>0} i^{-1}N^{-j}$, is $i^{-1}N^{-1}$ (since $y : i^{-1}N^{-j} \to i^{-1}N^{-j+1}$ is an isomorphism for $j > 1$). Therefore

$$\int_i^0 i^\natural N \simeq \int_i^0 i^{-1}N^{-1} \simeq i_*((k[\partial] \otimes_k \mathcal{D}_X) \otimes_{\mathcal{D}_X} i^{-1}N^{-1}) \simeq i_*(k[\partial] \otimes_k i^{-1}N^{-1}) \simeq k[\partial] \otimes_k N^{-1} \simeq N,$$

which completes the proof of Kashiwara's theorem modulo the (omitted) verification that the composite isomorphism above coincides with the counit.

Finally, we prove the claimed direct sum decomposition. This step is where the quasi-coherence of N is essential: since N is supported on X, this quasi-coherence implies that every $n \in N$ is annihilated by some power of the defining ideal (y) of X. Therefore it will be enough to show that

$$\ker(N \xrightarrow{y^k} N) \subset \oplus_{j=1}^k N^{-j}$$

for all $k \geq 1$ (then just take the ascending union of both sides). We prove this last statement by induction on k. For the base case, if $n \in N$ is such that $yn = 0$, then $\theta n = y\partial n = \partial yn - n = -n$, so $n \in N^{-1}$. Now assume that $y^k n = 0$ for some $k > 1$. Since $y^{k-1}(yn) = 0$, the induction hypothesis implies that $yn \in \oplus_{j=1}^{k-1} N^{-j}$. Applying ∂ (which drops the eigenvalue by 1), we find $\theta n + n = \partial yn \in \oplus_{j=2}^k N^{-j}$. On the other hand, the element $\theta n + kn$ is also annihilated by y^{k-1}: we have $y^{k-1}(\theta n + kn) = y^k \partial n + ky^{k-1}n = \partial(y^k n) = \partial(0)$. So the induction hypothesis also gives $\theta n + kn \in \oplus_{j=1}^{k-1} N^{-j}$, from which it follows that the difference $(k-1)n = (\theta n + kn) - (\theta n + n)$ belongs to $\oplus_{j=1}^{k-1} N^{-j} - \oplus_{j=2}^k N^{-j} \subset \oplus_{j=1}^k N^{-j}$. Since $k > 1$, $k - 1$ is invertible, and the proof is complete. □

The version of Kashiwara's theorem just proved can be viewed as the base case of a proof by induction (on *cohomological length*) for a more general derived category version of the theorem: namely, $\int_i : D^b_{qc}(\mathcal{D}_X) \to D^{b,X}_{qc}(\mathcal{D}_Y)$ is an equivalence of triangulated categories with quasi-inverse $\mathbb{R}i^\natural \simeq \mathbb{L}i^*[\dim X - \dim Y]$, where the target category is the bounded derived category of $\mathcal{D}_Y$-modules whose cohomology sheaves are both quasi-coherent and supported on X. (Again, the same is true with qc replaced by c.) We also remark that if $Z \hookrightarrow X$ is a closed immersion where X is smooth but Z is not, we can *define* (inspired by Kashiwara's theorem) $\mathrm{Mod}_{qc}(\mathcal{D}_Z)$ to be the subcategory $\mathrm{Mod}^Z_{qc}(\mathcal{D}_X)$. In Dennis Gaitsgory's notes on geometric representation theory, there is a proof that the resulting category is independent of the choice of X and embedding $Z \hookrightarrow X$.

Week 9: characteristic varieties

Last week we proved Kashiwara's theorem: if $i : X \hookrightarrow Y$ is a closed immersion, then $\int_i^0 : \mathrm{Mod}_{qc}(\mathcal{D}_X) \to \mathrm{Mod}_{qc}^X(\mathcal{D}_Y)$ is an equivalence of categories. Suppose that X is a single closed point $P \in Y$. The category $\mathrm{Mod}_{qc}(\mathcal{D}_P)$ is simply the category of k-spaces. Therefore, by Kashiwara's theorem, every object in the equivalent category $\mathrm{Mod}_{qc}^P(\mathcal{D}_Y)$ is isomorphic to a direct sum of copies of $\int_i^0 k$, and this last object is isomorphic to $i_*(k[\partial_1, \dots, \partial_n])$ where $\{y_i, \partial_i\}$ are local coordinates on an open affine U containing P. The left action of $\mathcal{D}_Y$ on the skyscraper sheaf $i_*(k[\partial_1, \dots, \partial_n])$ was described last week: the ∂_i act in the obvious way, and the y_i annihilate any element of $i_*(k[\partial_1, \dots, \partial_n])$ that lacks a ∂_i term to "absorb" them. If $R = k[y_1, \dots, y_n]$ and $Y = \operatorname{Spec} R$, then $k[\partial_1, \dots, \partial_n]$ with this action is isomorphic as a $\mathcal{D}(R, k)$-module to the top local cohomology module $H^n_{\mathfrak{m}}(R)$ supported at the irrelevant maximal ideal $\mathfrak{m} = (y_1, \dots, y_n)$. The module $H^n_{\mathfrak{m}}(R)$ is often denoted E (it is an injective hull of k over R). Concretely, E is usually described as the module of "inverse polynomials" $\oplus k \cdot y_1^{-\alpha_1} \cdots y_n^{-\alpha_n}$ where $(\alpha_1, \dots, \alpha_n)$ runs over all n-tuples of strictly positive integers, and the action of R on E sends any non-negative power of a variable to zero. The map $\mathcal{D}_Y / \mathcal{D}_Y \cdot \mathfrak{m} = k[\partial_1, \dots, \partial_n] \to E$ that sends $\partial_i^{\alpha_i}$ to $y_i^{-\alpha_i - 1}$ becomes an isomorphism once the images $y_i^{-\alpha_i - 1}$ are multiplied by appropriate constants determined by the quotient rule action of $\mathcal{D}(R, k)$ on E.

We briefly describe one application of the preceding description of quasi-coherent $\mathcal{D}$-modules supported at a single point. Recall that if X is an affine scheme, the global section functor $\Gamma(X, -) : \mathrm{Mod}_{qc}(\mathcal{O}_X) \to \mathrm{Mod}(\Gamma(X, \mathcal{O}_X))$ is exact. We define a smooth scheme X to be $\mathcal{D}$-*affine* if $\Gamma(X, -) : \mathrm{Mod}_{qc}(\mathcal{D}_X) \to \mathrm{Mod}(\Gamma(X, \mathcal{D}_X))$ is exact, and if whenever $\Gamma(X, M) = 0$ for a quasi-coherent $\mathcal{D}_X$-module M, then $M = 0$. (From these two conditions, it follows that every $M \in \mathrm{Mod}_{qc}(\mathcal{D}_X)$ is generated over $\mathcal{D}_X$ by its global sections, and that the functor $\Gamma(X, -)$ as above is an equivalence of categories.) Obviously, if X is smooth and affine, then X is $\mathcal{D}$-affine. However, there are other examples of $\mathcal{D}$-affine schemes:

Theorem. *(Beilinson-Bernstein) Projective space $\mathbb{P}^n_k$ is $\mathcal{D}$-affine.*

Sketch of proof. We use the following notation: $V = k^{n+1}$ with coordinates $x_0, \dots, x_n$; $V^\circ = V \setminus \{0\}$; $X = \mathbb{P}^n_k$; $\pi : V^\circ \to X$ is the projection; and $j : V^\circ \hookrightarrow V$ is the inclusion. Let $0 \to M' \to M \to M'' \to 0$ be a short exact sequence of quasi-coherent $\mathcal{D}_X$-modules. The functor π^* is exact (since π is smooth) and j_* is left exact, so we obtain a long exact sequence

$$0 \to j_*\pi^*M' \to j_*\pi^*M \to j_*\pi^*M'' \to \mathbb{R}^1 j_*\pi^*M' \to \cdots,$$

where the term $\mathbb{R}^1 j_*\pi^*M'$ and all later terms are quasi-coherent $\mathcal{D}_V$-modules supported at $V \setminus V^\circ = \{0\} \subset V$. By Kashiwara's theorem, all of these sheaves are direct sums of copies of the skyscraper $k[\partial_0, \dots, \partial_n]$ at 0. Consider the *Euler operator* $\theta = \sum_{i=0}^n x_i \partial_i$. By an explicit calculation, the eigenvalues of θ acting on $k[\partial_0, \dots, \partial_n]$ are strictly negative integers, while the 0-eigenspace of θ acting on $\Gamma(V, j_*\pi^*N) = \Gamma(V^\circ, \pi^*N)$ is exactly $\Gamma(X, N)$ for *any* $\mathcal{D}_X$-module N. By first taking global sections (an exact functor on the affine space V) and then passing to the 0-eigenspaces of θ, we obtain the desired *short* exact sequence $0 \to \Gamma(X, M') \to \Gamma(X, M) \to \Gamma(X, M'') \to 0$, since the global sections of the later terms in the long exact sequence have trivial 0-eigenspaces. □

We turn next to the problem of assigning a *dimension* to every coherent $\mathcal{D}_X$-module. Recall that the associated graded sheaf $\operatorname{gr} \mathcal{D}_X$ (with respect to the order filtration) is isomorphic to the symmetric algebra $\operatorname{Sym} \Theta_X$ and locally takes the form $\mathcal{O}_U[\xi_1, \dots, \xi_n]$, where $\{x_i, \partial_i\}$ are coordinates on an open affine $U \subset X$ and ξ_i, the *(principal) symbol* of ∂_i, is the class of ∂_i in $F_1\mathcal{D}_X / F_0\mathcal{D}_X \subset \operatorname{gr} \mathcal{D}_X$. We can glue the spectra of the rings $\mathcal{O}_U[\xi_1, \dots, \xi_n]$ as U varies, obtaining a smooth scheme $\mathbf{Spec}_X \operatorname{gr} \mathcal{D}_X$ of dimension $2n$ together with an affine projection map $\pi : \mathbf{Spec}_X \operatorname{gr} \mathcal{D}_X \to X$ such that $\pi_* \mathcal{O}_{\mathbf{Spec}_X \operatorname{gr} \mathcal{D}_X} = \operatorname{gr} \mathcal{D}_X$. Since $\operatorname{gr} \mathcal{D}_X \simeq \operatorname{Sym} \Theta_X$, the scheme $\mathbf{Spec}_X \operatorname{gr} \mathcal{D}_X \simeq \mathbf{Spec}_X \operatorname{Sym} \Theta_X$ together with the affine morphism π is the *geometric vector bundle* over X whose sheaf of sections is the $\mathcal{O}_X$-dual of Θ_X, namely the cotangent sheaf Ω^1_X (see exercise II.5.18 in Hartshorne). We therefore refer to $\mathbf{Spec}_X \operatorname{gr} \mathcal{D}_X$ as the *cotangent bundle* over X, denoted T^*X. Locally, π takes the form of a projection $\pi|_U : T^*U = U \times k^n \to U$, where the ξ_i are coordinates for k^n.

Let M be a coherent $\mathcal{D}_X$-module. We saw in Week 6 that there exists a coherent $\mathcal{O}_X$-submodule M_0 of M that generates M over $\mathcal{D}_X$. If we set $F_i M = (F_i \mathcal{D}_X) \cdot M_0 \subset M$ for all $i \geq 0$, we obtain a *global* good filtration on M: the associated graded sheaf of modules $\operatorname{gr}^F M$ is coherent over $\operatorname{gr} \mathcal{D}_X = \pi_* \mathcal{O}_{T^*X}$. (Note that the existence of M_0 means that we do not need to worry about patching good filtrations on $M|_U$ as U varies.) By pulling this coherent sheaf back to the cotangent bundle, we obtain a coherent $\mathcal{O}_{T^*X}$-module

$$\widetilde{\operatorname{gr}^F M} = \mathcal{O}_{T^*X} \otimes_{\pi^{-1}\pi_* \mathcal{O}_{T^*X}} \pi^{-1}(\operatorname{gr}^F M)$$

such that $\pi_* \widetilde{\operatorname{gr}^F M} = \operatorname{gr}^F M$.

Definition. *The support* $\operatorname{Supp} \widetilde{\operatorname{gr}^F M}$ *of* $\widetilde{\operatorname{gr}^F M}$*, a closed subset of* T^*X *which we endow with the reduced subscheme structure, is called the* characteristic variety *of* M *and denoted* $\operatorname{Ch}(M)$.

The closed subscheme $\operatorname{Ch}(M)$ is defined by the *characteristic ideal* $\sqrt{\operatorname{Ann}_{\operatorname{gr} \mathcal{D}_X} \operatorname{gr}^F M}$. Since $\operatorname{gr}^F M$ is a graded module over $\operatorname{gr} \mathcal{D}_X$, and the grading on the latter comes (locally) from the ξ_i, the graded ideal $\sqrt{\operatorname{Ann}_{\operatorname{gr} \mathcal{D}_X} \operatorname{gr}^F M}$ is homogeneous with respect to the ξ_i, from which it follows that $\operatorname{Ch}(M)$ is *conic* (closed under scalar multiplication on the fibers of π). In other words, if $(\vec{x}, \vec{\xi}) \in \operatorname{Ch}(M)$, then $(\vec{x}, \lambda\vec{\xi}) \in \operatorname{Ch}(M)$ for all $\lambda \in k$.

In order for the characteristic variety to be a well-defined construction, we must show that $\operatorname{Supp} \widetilde{\operatorname{gr}^F M}$ is independent of the choice of good filtration F on M. For this, we may assume X is affine and M is a finitely generated left module over the ring $D_X = \mathcal{D}_X(X)$.

Proposition. *Let* M *be a finitely generated left module over* D_X*, and let* F *and* G *be good filtrations on* M*. Then* $\operatorname{Supp} \operatorname{gr}^F M = \operatorname{Supp} \operatorname{gr}^G M$ *as subsets of* $\operatorname{gr} D_X$.

Proof. We may assume that $F_p M \subset G_p M \subset F_{p+1} M$ for all p, since by a result from Week 5, any two good filtrations can be linked by a finite chain of such pairs. Fix p and consider the natural map $\varphi_p : F_p M / F_{p-1} M \to G_p M / G_{p-1} M$ induced by the inclusions. The kernel of φ_p is $G_{p-1} M / F_{p-1} M$, and using the fact that $G_{p-1} M \subset F_p M$, we see that the cokernel of φ_p is $(G_p M / G_{p-1} M)/(F_p M / G_{p-1} M) \simeq G_p M / F_p M$. That is, we have $\ker \varphi_p \simeq \operatorname{coker} \varphi_{p-1}$ for all p. Passing to the direct sum over all $p \geq 0$, the map $\varphi = \oplus \varphi_p : \operatorname{gr}^F M \to \operatorname{gr}^G M$ has $\ker \varphi \simeq \operatorname{coker} \varphi$ (disregarding the gradings). By considering the short exact sequences

$$0 \to \ker \varphi \to \operatorname{gr}^F M \to \operatorname{im} \varphi \to 0, \; 0 \to \operatorname{im} \varphi \to \operatorname{gr}^G M \to \operatorname{coker} \varphi \to 0$$

we have equalities

$$\operatorname{Supp} \operatorname{gr}^F M = \operatorname{Supp} \ker \varphi \cup \operatorname{Supp} \operatorname{im} \varphi = \operatorname{Supp} \operatorname{coker} \varphi \cup \operatorname{Supp} \operatorname{im} \varphi = \operatorname{Supp} \operatorname{gr}^G M,$$

where the second equality uses the isomorphism $\ker \varphi \simeq \operatorname{coker} \varphi$. The proof is complete. □

If $M \in \operatorname{Mod}_c(\mathcal{D}_X)$, the preceding proposition shows that $\operatorname{Ch}(M)$ is a well-defined invariant of M, and therefore so is its dimension, $d(M) = \dim \operatorname{Ch}(M)$, which we refer to as the *dimension* of M. Given a short exact sequence $0 \to M' \to M \to M'' \to 0$ in $\operatorname{Mod}_c(\mathcal{D}_X)$, we have $\operatorname{Ch}(M) = \operatorname{Ch}(M') \cup \operatorname{Ch}(M'')$ (a good filtration on M induces good filtrations on M' and M'', then we consider the supports of the terms of the short exact sequence obtained upon passing to associated graded objects); and therefore $d(M) = \max\{d(M'), d(M'')\}$. *A priori*, the dimension of M is an integer

$d(M)$ such that $0 \leq d(M) \leq 2n$ (since $\mathrm{Ch}(M) \subset T^*X$). Next week, we will prove *Bernstein's inequality*, a much stronger lower bound for $d(M)$.

Week 10: Bernstein's inequality; holonomic modules

Recall from last week how we assign a *dimension* $d(M)$ to every coherent $\mathcal{D}_X$-module M: we choose a global good filtration F on M, form the associated graded $\operatorname{gr}\mathcal{D}_X$-module $\operatorname{gr}^F M$, and then pull this module back along the cotangent bundle $\pi : \mathbf{Spec}_X \operatorname{gr}\mathcal{D}_X = T^*X \to X$, obtaining a coherent $\mathcal{O}_{T^*X}$-module $\widetilde{\operatorname{gr}^F M}$. The support $\operatorname{Supp}\widetilde{\operatorname{gr}^F M}$, a closed subset of T^*X which we endow with the reduced subscheme structure, does not depend on the choice of good filtration F on M: it is called the *characteristic variety* $\mathrm{Ch}(M)$ of M, and we take for $d(M)$ its dimension $\dim(\mathrm{Ch}(M))$. Because $\mathrm{Ch}(M) \subset T^*X$, we have $0 \leq d(M) \leq 2n$.

Theorem. *(Bernstein's inequality) Let $M \neq 0$ be a coherent $\mathcal{D}_X$-module. Then we have $d(M) \geq n$.*

Proof. We will use induction on $\dim(X)$, the base case (where $0 = n = 2n$) being obvious. We observe first that the support of M as a sheaf on X is precisely $\pi(\mathrm{Ch}(M))$. To see this, note that if M_0 is a coherent $\mathcal{O}_X$-submodule of M that generates M over $\mathcal{D}_X$, then M and M_0 have the same support on X: clearly $\operatorname{Supp} M_0 \subset \operatorname{Supp} M$, and if $x \notin \operatorname{Supp} M_0$, then there is an open affine neighborhood U of x such that $(M_0)|_U$, and consequently $M(U) = \mathcal{D}_X(U) \cdot M_0(U)$, vanishes. If we consider the good filtration F on M defined by $F_iM = (F_i\mathcal{D}_X) \cdot M_0$, then $\operatorname{gr}^F M$ is generated over $\operatorname{gr}\mathcal{D}_X$ by M_0, from which it follows that $\pi(\mathrm{Ch}(M)) = \operatorname{Supp} M_0 = \operatorname{Supp} M$. Note that this subset is closed, since M_0 is coherent over $\mathcal{O}_X$.

Suppose for contradiction that $\dim(\mathrm{Ch}(M)) < n = \dim(X)$. Then $S = \pi(\mathrm{Ch}(M))$ is a *proper* closed subscheme of X. By the generic smoothness theorem (Corollary III.10.7 in Hartshorne), after replacing X with a dense open subset (which must have the same dimension), we may assume that S is smooth. Let $i : S \hookrightarrow X$ be the inclusion. By the coherent version of Kashiwara's theorem, $M \simeq \int_i^0 N$ for some coherent $\mathcal{D}_S$-module N, so we simply replace M with $\int_i^0 N$. By induction on dimension, we may assume that $d(N) = \dim(\mathrm{Ch}(N)) \geq \dim(S)$. We claim that $\dim(\mathrm{Ch}(N)) + r = \dim(\mathrm{Ch}(M))$, where r is the codimension of S in X (this claim will imply that $\dim(\mathrm{Ch}(M)) \geq \dim(S) + r = \dim(X)$, a contradiction that completes the proof).

By the composition rule for direct images, it suffices to assume that $r = 1$. We also assume that we have coordinates $\{x_i, \partial_i\}$ on X such that $x_n = 0$ is a defining equation for S. We write x for x_n, ∂ for ∂_n, and ξ for the principal symbol ξ_n. We are simply going to prove here that $\dim(\mathrm{Ch}(N)) + 1 = \dim(\mathrm{Ch}(M))$. A stronger statement is true: if ρ is the natural map $S \times_X T^*X \to T^*S$, which is smooth of relative dimension 1, then $\rho^{-1}(\mathrm{Ch}(N)) = \mathrm{Ch}(M) \subset S \times_X T^*X \subset T^*X$. (This statement is local on S and would be needed for our passage to local coordinates to be rigorous.)

In coordinates, we know that $M \simeq k[\partial] \otimes_k i_*N \simeq \oplus_{j \geq 0} \partial^j \cdot N$. We will show that if we choose our good filtrations carefully, then $\operatorname{gr} M$ is obtained from $\operatorname{gr} N$ by adjoining a coordinate (namely ξ). Let G be a good filtration on N; after shifting, we may assume $G_{-1}N = 0$. Define a good filtration F on M by setting

$$F_jM = \partial^j \cdot G_0N + \partial^{j-1} \cdot G_1N + \cdots + G_jN;$$

then we have

$$\begin{aligned}\frac{F_jM}{F_{j-1}M} &= \frac{\partial^j \cdot G_0N + \partial^{j-1} \cdot G_1N + \cdots + \partial \cdot G_{j-1}N + G_jN}{\partial^{j-1} \cdot G_0N + \partial^{j-2} \cdot G_1N + \cdots + G_{j-1}N} \\ &\simeq \frac{G_jN}{G_{j-1}N} \oplus \xi \cdot \frac{G_{j-1}N}{G_{j-2}N} \oplus \cdots \oplus \xi^j \cdot G_0N,\end{aligned}$$

so that

$$\begin{aligned}\operatorname{gr}^F M &= \oplus_j F_j M / F_{j-1} M \\ &\simeq \bigoplus_j \left(\frac{G_j N}{G_{j-1} N} \oplus \xi \cdot \frac{G_{j-1} N}{G_{j-2} N} \oplus \cdots \oplus \xi^j \cdot G_0 N \right) \\ &\simeq \oplus_j \xi_j \cdot \operatorname{gr}^G N \\ &\simeq k[\xi] \otimes_k \operatorname{gr}^G N,\end{aligned}$$

from which it is clear that $\dim \operatorname{Supp} \operatorname{gr}^G N + 1 = \dim \operatorname{Supp} \operatorname{gr}^F M$, as desired. □

If M is any coherent $\mathcal{D}_X$-module, there exists (see Appendix D in HTT) a filtration

$$0 = C^{2n+1} M \subset C^{2n} M \subset \cdots \subset C^1 M \subset C^0 M = M$$

by coherent $\mathcal{D}_X$-submodules such that, for all j such that $C^j M / C^{j+1} M \neq 0$, every irreducible component of $\operatorname{Ch}(C^j M / C^{j+1} M)$ has dimension $2n - j$. Note that $C^j M / C^{j+1} M$ is a coherent $\mathcal{D}_X$-module. It follows from Bernstein's inequality that we must have $C^{n+1} M = C^{n+2} M = \cdots = C^{2n+1} M = 0$, and therefore

Corollary. *Let $M \neq 0$ be a coherent $\mathcal{D}_X$-module. Then every irreducible component of $\operatorname{Ch}(M)$ has dimension at least $n = \dim(X)$.*

The coherent $\mathcal{D}_X$-modules that are "as small as possible" (either zero, or of dimension exactly n) form a privileged class with especially nice properties.

Definition. *Let M be a coherent $\mathcal{D}_X$-module. We say that M is* holonomic *if $M = 0$ or $d(M) = n$.*

It follows from the previous corollary that if $M \neq 0$ is holonomic, every irreducible component of $\operatorname{Ch}(M)$ has dimension exactly n (that is, $\operatorname{Ch}(M)$ is equidimensional). Recall from last week that if $0 \to M' \to M \to M'' \to 0$ is a short exact sequence in $\operatorname{Mod}_c(\mathcal{D}_X)$, then $\operatorname{Ch}(M)$ is the union of $\operatorname{Ch}(M')$ and $\operatorname{Ch}(M'')$, and consequently $d(M)$ is the maximum of $d(M')$ and $d(M'')$. It follows that M is holonomic if and only if M' and M'' are holonomic. Therefore, the full subcategory $\operatorname{Mod}_h(\mathcal{D}_X) \subset \operatorname{Mod}_c(\mathcal{D}_X)$ consisting of holonomic modules is a Serre subcategory; since the category $\operatorname{Mod}_c(\mathcal{D}_X)$ is Abelian, so also is $\operatorname{Mod}_h(\mathcal{D}_X)$.

If $M \in \operatorname{Mod}_c(\mathcal{D}_X)$ is actually coherent as an $\mathcal{O}_X$-module, we saw in Week 2 that M must be locally free over $\mathcal{O}_X$. Given such an M, we can define a filtration F on M by setting $F_i M = 0$ if $i < 0$ and $F_i M = M$ if $i \geq 0$. Since every $F_i M$ is coherent over $\mathcal{O}_X$, F is a good filtration, with $\operatorname{gr}^F M \simeq M$. Note that, for all i, $\xi_i \in \operatorname{gr} \mathcal{D}_X$ annihilates $\operatorname{gr}^F M$ (because it increases degree by 1). Therefore $\operatorname{Ch}(M)$ is only the zero section $T^*_X X \subset T^* X$ (locally, $T^*_U U = U \times \{0\} \subset U \times k^n = T^* U$). Since $\dim T^*_X X = \dim X = n$, M is holonomic. In fact, if $M \neq 0$, $\operatorname{Ch}(M) = T^*_X X$ if and only if M is coherent over $\mathcal{O}_X$ (if and only if M is locally free over $\mathcal{O}_X$). From this characterization, we can see that every holonomic $\mathcal{D}_X$-module is generically of this form:

Proposition. *Let M be a holonomic $\mathcal{D}_X$-module. There exists an open dense subset $U \subset X$ such that $M|_U$ is coherent as an $\mathcal{O}_U$-module.*

Proof. Let $S = \operatorname{Ch}(M) \setminus T^*_X X$. If S is empty, then either $M = 0$ or $\operatorname{Ch}(M) = T^*_X X$; in the latter case, M is already coherent over $\mathcal{O}_X$ and so we can take X for U. Otherwise, since $\operatorname{Ch}(M)$ is conic, the fibers of $\pi|_S : S \to \pi(S)$ are at least one-dimensional, so $\dim \pi(S) < \dim S = \dim X$, and there is an open dense subset U of X that lies within the complement $X \setminus \pi(S)$. But then $\operatorname{Ch}(M|_U)$ is contained in $T^*_U U$, so $M|_U$ is coherent over $\mathcal{O}_U$. □

We remark that in fact $\pi(S)$ in the preceding proof is already closed, so we can simply take $X \setminus \pi(S)$ for U. To see this, note that the restriction of π to $T^*X \setminus T^*_X X$ factors through the projectivized cotangent bundle $\mathbf{Proj}_X \operatorname{Sym} \Theta_X = \mathbb{P}(\Theta_X) \to X$, and this last map is proper; since S is conic, its image in X must be closed. We also remark that if M is merely coherent over $\mathcal{D}_X$ (not necessarily holonomic), there exists an open dense subset V (not the same as the U above) such that $M|_V$ is *projective* over $\mathcal{O}_V$.

Week 11: more on holonomic modules; sins of omission

Last week we defined the category $\mathrm{Mod}_h(\mathcal{D}_X)$ of *holonomic* $\mathcal{D}_X$-modules: the Serre subcategory of $\mathrm{Mod}_c(\mathcal{D}_X)$ consisting of 0 together with all coherent $\mathcal{D}_X$-modules M such that $d(M) = n$ (by Bernstein's inequality, this is the smallest possible dimension). In order to prove our next structural result for $\mathrm{Mod}_h(\mathcal{D}_X)$, we need to introduce multiplicities. If M is any coherent $\mathcal{D}_X$-module with a chosen good filtration F, and $C \subset T^*X$ is an irreducible component of $\mathrm{Ch}(M)$, the *multiplicity* of $\widetilde{\mathrm{gr}^F M}$ along C is the length of the stalk of $\widetilde{\mathrm{gr}^F M}$ at the generic point η of C, viewed as a module (which must be of finite length) over the local ring $\mathcal{O}_{T^*X,\eta}$. Let $m_{d(M)}(M)$ be the sum of the multiplicities of $\widetilde{\mathrm{gr}^F M}$ along the $d(M)$-dimensional irreducible components of $\mathrm{Ch}(M)$. This sum is independent of the choice of good filtration F (the proof of this fact is similar to the proof that $\mathrm{Ch}(M)$ is independent: prove it first for neighboring filtrations), and since length is additive in short exact sequences, we have $m_{d(M)}(M) = m_{d(M)}(M') + m_{d(M)}(M'')$ for every short exact sequence $0 \to M' \to M \to M'' \to 0$ in $\mathrm{Mod}_c(\mathcal{D}_X)$.

Theorem. *Every holonomic $\mathcal{D}_X$-module M is of finite length.*

Proof. Every irreducible component of $\mathrm{Ch}(M)$, as well as of $\mathrm{Ch}(M')$ for any $\mathcal{D}_X$-submodule $M' \subset M$ or of $\mathrm{Ch}(M'')$ for any $\mathcal{D}_X$-module quotient $M \twoheadrightarrow M''$, has dimension exactly n. Consider the multiplicity $m(M) = m_n(M)$. If $m(M) = 0$, then $\mathrm{Ch}(M)$ is empty, and therefore $M = 0$. Therefore, if $M' \subset M$ is a proper $\mathcal{D}_X$-submodule (so that $M'' = M/M' \neq 0$ and thus $m(M'') > 0$), we have $m(M) = m(M') + m(M'') > m(M')$. The finite length of M follows by induction on $m(M)$. ☐

Recall that if $f : X \to Y$ is a morphism, the direct image $\int_f$ and (shifted) inverse image $f^\dagger = \mathbb{L}f^*[\dim X - \dim Y]$ preserve quasi-coherence, but do not preserve coherence in general (if f is proper, $\int_f$ preserves coherence). Remarkably, both $\int_f$ and $f^\dagger$ preserve *holonomy* for general f:

Theorem. *(Preservation of holonomy) Let $f : X \to Y$ be a morphism. If $M^\bullet \in D^b_h(\mathcal{D}_X)$, then $\int_f M^\bullet \in D^b_h(\mathcal{D}_Y)$, and if $M^\bullet \in D^b_h(\mathcal{D}_Y)$, then $f^\dagger M^\bullet \in D^b_h(\mathcal{D}_X)$, where the subscript h indicates that we are considering the derived category of complexes whose cohomology objects are all holonomic.*

Next week (the final seminar), we will give a proof of the first part of this theorem. In fact the most difficult step is proving that $\int_f$ preserves holonomy when f is the projection $k^n \to k^{n-1}$; the general statement reduces without too much trouble to this one, and the statement for $f^\dagger$ reduces to the statement for $\int_f$ using some standard distinguished triangles in $D^b(\mathcal{D}_X)$. We consider an application. Suppose that $Y = P$ is a point. Since X is smooth, the projection $p : X \to P$ is a smooth morphism. Therefore, we can compute the complex $\int_p M \in D^b(\mathcal{D}_P)$ using the de Rham complex: $\int_p M$ is represented by $\mathbb{R}p_*(\mathrm{DR}_X)$, since the de Rham complex of X relative to the point P is just DR_X. The functor p_* is simply the global section functor on X, and so $\mathbb{R}p_*$ is the hypercohomology functor. That is, $\int_p M$ is represented by $\mathbb{R}\Gamma(X, \mathrm{DR}_X)$. The cohomology objects (k-spaces) of the complex $\mathbb{R}\Gamma(X, \mathrm{DR}_X)$ are called the *de Rham cohomology spaces* $H^*_{dR}(M)$ of M. If M is holonomic, then by the theorem above, all $H^*_{dR}(M)$ are holonomic $\mathcal{D}_P$-modules, that is, finite dimensional k-spaces (a $\mathcal{D}_P$-module is coherent if and only if it is holonomic). We have therefore proved

Corollary. *If M is a holonomic $\mathcal{D}_X$-module, the de Rham cohomology spaces of M are finite-dimensional.*

Finally, we are going to summarize a couple of the important pieces of the algebraic theory (the first three chapters of HTT) that we did not cover in detail: the "sins of omission" of this mini-course. One such omission is the behavior of direct and inverse images for smooth morphisms, which was summarized in Week 7. We discuss two more here: *duality* and *minimal (or intermediate) extensions.*

Duality. We define a contravariant *dual* functor $\mathbb{D} : D^b(\mathcal{D}_X) \to D^b(\mathcal{D}_X)$ by

$$\mathbb{D}M^\bullet = \mathbb{R}\mathcal{H}om_{\mathcal{D}_X}(M^\bullet, \mathcal{D}_X) \otimes_{\mathcal{O}_X} \omega_X^{\otimes -1}[n],$$

a derived category version of the usual dual module construction over a non-commutative ring. Here $\mathbb{R}\mathcal{H}om_{\mathcal{D}_X}(M^\bullet, \mathcal{D}_X)$ becomes a complex of right $\mathcal{D}_X$-modules by the right action of $\mathcal{D}_X$ on itself, and so the side-changing operation is necessary to recover a complex of left $\mathcal{D}_X$-modules. The functor $\mathbb{D}$ preserves coherence, and if $M^\bullet \in D^b_c(\mathcal{D}_X)$, then $M^\bullet \xrightarrow{\sim} \mathbb{D}\mathbb{D}M^\bullet$ (the indicated map always exists; the coherence is necessary for the map to be an isomorphism).

Now we consider the case where $M^\bullet$ is concentrated in degree zero, that is, the case of a single coherent $\mathcal{D}_X$-module M. Consider the cohomology objects $\mathcal{E}xt^i_{\mathcal{D}_X}(M, \mathcal{D}_X)$ of the complex $\mathbb{R}\mathcal{H}om_{\mathcal{D}_X}(M, \mathcal{D}_X)$. It is true in general that $\mathcal{E}xt^i_{\mathcal{D}_X}(M, \mathcal{D}_X) = 0$ for $i < 2n - d(M)$, and since $\mathcal{D}_X$ has global homological dimension $\leq n$, we also know that $\mathcal{E}xt^i_{\mathcal{D}_X}(M, \mathcal{D}_X) = 0$ for $i > n$. Therefore, if M is holonomic ($d(M) = n$), there is exactly one nonvanishing $\mathcal{E}xt$, namely $\mathcal{E}xt^n_{\mathcal{D}_X}(M, \mathcal{D}_X)$. In this case, $\mathbb{D}M$ is isomorphic in $D^b(\mathcal{D}_X)$ to a single $\mathcal{D}_X$-module concentrated in degree *zero* (because of the degree shift in the definition of $\mathbb{D}$) and so we identify $\mathbb{D}M$ with this module, namely the left $\mathcal{D}_X$-module $\mathcal{E}xt^n_{\mathcal{D}_X}(M, \mathcal{D}_X) \otimes_{\mathcal{O}_X} \omega_X^{\otimes -1}$. (We abuse notation by writing $\mathbb{D}M$ instead of $\mathbb{D}^0 M$.) If M is holonomic, so again is $\mathbb{D}M$; in fact, their characteristic varieties coincide. As a special case, if M is locally free as an $\mathcal{O}_X$-module, $\mathbb{D}M \simeq \mathcal{H}om_{\mathcal{O}_X}(M, \mathcal{O}_X)$ as left $\mathcal{D}_X$-modules (recall from Week 2 that we can define a left $\mathcal{D}_X$-module structure on the right-hand side by specifying how the derivations act). Taking $M = \mathcal{O}_X$, we see that $\mathcal{E}xt^n_{\mathcal{D}_X}(\mathcal{O}_X, \mathcal{D}_X) \simeq \omega_X$, which is nonzero when restricted to any open affine $U \subset X$; from this it follows that the ring $\mathcal{D}_X(U)$ has global dimension *exactly* n.

Returning to the general framework of derived categories with coherent cohomology, $\mathbb{D}$ commutes with $\int_f$ when $f : X \to Y$ is proper; the commutativity of $\mathbb{D}$ with inverse images is a more complicated story.

Minimal extensions. Let M be a holonomic $\mathcal{D}_X$-module. Since M has finite length, there exists a composition series $0 = M_{l+1} \subset M_l \subset \cdots \subset M_1 \subset M_0 = M$ of (holonomic) $\mathcal{D}_X$-submodules such that M_i/M_{i+1} is *simple* (contains no nontrivial $\mathcal{D}_X$-submodule) for all i. The theory of minimal extensions provides a classification of all simple holonomic $\mathcal{D}_X$-modules. We remark that there exist simple $\mathcal{D}_X$-modules that are *not* holonomic. If $X = \mathbb{A}^2 = \operatorname{Spec} k[x, y]$ and we let

$$\delta = x + y + \partial_x + \partial_y + y\partial_x\partial_y \in \mathcal{D}_X(X) = A_2$$

(the second Weyl algebra), then δ generates a maximal left ideal in A_2. The quotient $A_2/A_2 \cdot \delta$ is a simple left A_2-module that is not holonomic (its dimension is 3). This example is due to Toby Stafford.

If $f : X \to Y$ is a morphism, we define a functor $\int_{f!} = \mathbb{D}_Y \circ \int_f \circ \mathbb{D}_X : D^b_h(\mathcal{D}_X) \to D^b_h(\mathcal{D}_Y)$ (we are using the theorem on preservation of holonomy). There is a natural transformation $\int_{f!} \to \int_f$ which is an isomorphism when f is proper (because in that case $\int_f$ commutes with the dual functors). Now suppose that $Y \subset X$ is a locally closed smooth subscheme such that the inclusion map $i : Y \hookrightarrow X$ is *affine*. Then $\mathbb{R}i_* \simeq i_*$ (i_* is exact), and $\mathcal{D}_{Y\leftarrow X}$ is a locally free right $\mathcal{D}_Y$-module: if we factor i as $Y \xrightarrow{j} W \xrightarrow{\sigma} X$ where j is an open immersion and σ a closed immersion, then $\mathcal{D}_{W\leftarrow Y} = j^{-1}\mathcal{D}_W (= \mathcal{D}_Y)$ and so

$$\mathcal{D}_{X\leftarrow Y} \simeq j^{-1}\mathcal{D}_{X\leftarrow W} \otimes_{j^{-1}\mathcal{D}_W} \mathcal{D}_{W\leftarrow Y} \simeq j^{-1}\mathcal{D}_{X\leftarrow W},$$

and since σ is a closed immersion, $\mathcal{D}_{X\leftarrow W}$ is a locally free right $\mathcal{D}_W$-module. Therefore, if M is a holonomic $\mathcal{D}_Y$-module, we can think of $\int_i M$ and $\int_{i!} M$ as single $\mathcal{D}_X$-modules $\int_i^0 M$ and $\int_{i!}^0 M$ (all derived functors involved in their definitions become exact), and both of these $\mathcal{D}_X$-modules are holonomic. We have a morphism $\int_{i!}^0 M \to \int_i^0 M$ in $\mathrm{Mod}_h(\mathcal{D}_X)$ as above.

Definition. *Let $Y \subset X$ be a locally closed smooth subscheme such that the inclusion map $i : Y \hookrightarrow X$ is affine. Let M be a holonomic $\mathcal{D}_Y$-module. The image of the morphism $\int_{i!}^0 M \to \int_i^0 M$ is called the* minimal *(or* intermediate*) extension of M, and denoted $\mathcal{L}(Y, M)$.*

Observe that if Y is actually a *closed* subscheme, then i is proper, in which case $\mathcal{L}(Y, M)$ is simply $\int_i^0 M$. The following is the main classification theorem for simple holonomic $\mathcal{D}_X$-modules using minimal extensions:

Theorem. *Let Y, X, and i be as in the definition above. If M is a simple holonomic $\mathcal{D}_Y$-module, then $\mathcal{L}(Y, M)$ is a simple holonomic $\mathcal{D}_X$-module; indeed, it is the* unique *simple $\mathcal{D}_X$-submodule of $\int_i^0 M$. Conversely, given X, any simple holonomic $\mathcal{D}_X$-module is of the form $\mathcal{L}(Y, M)$ for some locally closed smooth subscheme Y of X with affine inclusion map and some simple $\mathcal{D}_Y$-module that is* coherent *as an $\mathcal{O}_Y$-module.*

Week 12: preservation of holonomy

We begin by revisiting the example from last week of a simple, non-holonomic module. Let $X = \mathbb{A}^n = \operatorname{Spec} k[x_1, \ldots, x_n]$, so that $D_X = \mathcal{D}_X(X)$ is the nth Weyl algebra A_n. Let $\delta \in D_X$ be given. We are going to compute the characteristic variety of $D_X/(D_X \cdot \delta)$. Suppose that $\delta \in F_l D_X \setminus F_{l-1} D_X$, and let $\xi = \sigma_l(\delta) \neq 0$ be the principal symbol of δ in $F_l D_X / F_{l-1} D_X \subset \operatorname{gr} D_X$. The order filtration F on D_X induces filtrations F' on the left ideal $D_X \cdot \delta$ and F'' on the quotient $D_X / D_X \cdot \delta$. We have $F'_p(D_X \cdot \delta) = (F_{p-l} D_X) \cdot \delta$ for all p, and so $\operatorname{gr}^{F'}(D_X \cdot \delta) = (\operatorname{gr} D_X) \cdot \xi$. The short exact sequence $0 \to D_X \cdot \delta \to D_X \to D_X/(D_X \cdot \delta) \to 0$ induces a short exact sequence of associated graded objects (since F' and F'' are induced from F), from which we conclude that

$$\operatorname{gr}^{F''}(D_X/(D_X \cdot \delta)) \simeq \operatorname{gr} D_X / (\operatorname{gr} D_X \cdot \xi),$$

a hypersurface ring. Recall that $\operatorname{gr} D_X$ is simply the polynomial ring $k[x_1, \ldots, x_n, \xi_1, \ldots, \xi_n]$. The characteristic variety $\operatorname{Ch}(D_X/(D_X \cdot \delta))$ is the closed set $V(\xi) \subset \operatorname{Spec} k[x_1, \ldots, x_n, \xi_1, \ldots, \xi_n]$. In particular, its dimension (the dimension of the D_X-module $D_X/(D_X \cdot \delta)$) is $2n - 1$. In our example from last week, $n = 2$ and $\delta = x + y + \partial_x + \partial_y + y\partial_x\partial_y$. Since $\delta \in F_2 D_X \setminus F_1 D_X$, its principal symbol in $k[x, y, \xi_x, \xi_y]$ is $y\xi_x\xi_y$, and its characteristic variety is $V(y\xi_x\xi_y) \subset \operatorname{Spec} k[x, y, \xi_x, \xi_y]$, a three-dimensional closed subscheme (the union of three coordinate hyperplanes in affine 4-space).

Before turning to the proof of the theorem on preservation of holonomy, we need to discuss a certain distinguished triangle in the derived category $D^b_{qc}(\mathcal{D}_X)$. Suppose first that X is merely a topological space and $Z \subset X$ is a closed subset. Write i for the closed immersion $Z \hookrightarrow X$ and j for the complementary open immersion $U = X \setminus Z \hookrightarrow X$. If $\mathcal{I}$ is an *injective* object in the category of Abelian sheaves on X, there is a well-known short exact sequence

$$0 \to \underline{\Gamma}_Z(\mathcal{I}) \to \mathcal{I} \to j_* j^{-1} \mathcal{I} \to 0$$

where $\underline{\Gamma}_Z$ denotes the subsheaf of sections supported on Z. (The kernel of $\mathcal{F} \to j_* j^{-1}\mathcal{F}$ is $\underline{\Gamma}_Z(\mathcal{F})$ for *any* Abelian sheaf $\mathcal{F}$ on X, more or less by definition, but if $\mathcal{F}$ is not an injective sheaf, then $\mathcal{F} \to j_* j^{-1}\mathcal{F}$ may have a nontrivial cokernel.) This short exact sequence for injective sheaves induces a distinguished triangle

$$\mathbb{R}\underline{\Gamma}_Z(\mathcal{F}^\bullet) \to \mathcal{F}^\bullet \to \mathbb{R}j_* j^{-1}\mathcal{F}^\bullet \xrightarrow{+1}$$

in the bounded derived category of Abelian sheaves on X. Now suppose X is a smooth scheme (satisfying our current laundry list of hypotheses) and $M^\bullet \in D^b_{qc}(\mathcal{D}_X)$. Then $\mathbb{R}j_* j^{-1}M^\bullet = \int_j j^{-1}M^\bullet = \int_j j^\dagger M^\bullet$, where $j^\dagger$ is the shifted inverse image functor (there is no degree shift in this case, because U is a dense open subset). If we assume furthermore that Z is smooth, then $\mathbb{R}\underline{\Gamma}_Z(M^\bullet) \simeq \int_i i^\dagger M^\bullet$ (this is an application of the derived version of Kashiwara's theorem; here the quasi-coherence is necessary). Therefore, if Z is a smooth closed subscheme of a smooth scheme X, we have a distinguished triangle

$$\int_i i^\dagger M^\bullet \to M^\bullet \to \int_j j^\dagger M^\bullet \xrightarrow{+1}$$

in $D^b_{qc}(\mathcal{D}_X)$ whenever $M^\bullet$ belongs to this derived category.

We now sketch a proof of the "forward" direction of the theorem on preservation of holonomy, which we restate below. (The "backward" direction, for the inverse image functor $f^\dagger$, can be reduced to the "forward" direction using the distinguished triangle above.)

Theorem. *Let $f : X \to Y$ be any morphism between smooth schemes. If $M^\bullet \in D^b_h(\mathcal{D}_X)$, then $\int_f M^\bullet \in D^b_h(\mathcal{D}_Y)$.*

Sketch of proof. By induction on the cohomological length of $M^\bullet$, it suffices to prove the theorem in the case where $M \in \mathrm{Mod}_h(\mathcal{D}_X)$ is a single holonomic $\mathcal{D}_X$-module concentrated in degree zero. Using the graph of f, we can factor f into a closed immersion followed by a projection, so it suffices to prove the theorem in case f is a morphism of one of these two types. If $f = i$ is a closed immersion $X \hookrightarrow Y$, then $\int_i^0$ is exact, and a coherent $\mathcal{D}_X$-module M is holonomic *if and only if* $\int_i^0 M$ is a holonomic $\mathcal{D}_Y$-module, as we saw in the proof of Bernstein's inequality (the functor $\int_i^0$ increases the dimension of $\mathrm{Ch}(M)$ by exactly $\dim Y - \dim X$). The fact that this is an *if and only if* statement in this case will be important below.

Now suppose $f = p$ is a projection $X = Y \times Z \to Y$. We need only to show that if $M \in \mathrm{Mod}_h(\mathcal{D}_{Y\times Z})$, then $\int_p M \in D^b_h(\mathcal{D}_Y)$. We may assume that both X and Y are affine: the problem is local on Y, so we may immediately reduce to the case of affine Y, and as for X, we use the Čech complex. We can find a finite affine open covering $\{U_i = Y \times V_i\}$ of X. Since we have assumed X is separated, M is isomorphic in $D^b(\mathcal{D}_X)$ to the Čech complex of M with respect to the cover $\{U_i\}$. If U is the intersection of some subcollection of the U_i and $j : U \hookrightarrow X$ is the corresponding (affine) open immersion, then $j_*(M|_U) = \int_j j^* M$ (since j is affine and hence $\mathbb{R}j_* = j_*$). The Čech complex of M is a direct sum of sheaves of the form $j_*(M|_U)$. Therefore, to show that $\int_p M \in D^b_h(\mathcal{D}_Y)$, it is enough (since finite direct sums of holonomic modules are again holonomic) to show that $\int_p \int_j j^* M = \int_{p\circ j} j^* M \in D^b_h(\mathcal{D}_Y)$, and therefore we may assume $X = U$ is affine as well. Fix embeddings $\alpha : X \hookrightarrow k^n$, $\beta : Y \hookrightarrow k^m$ of X and Y into affine spaces, and consider the graph morphism $\Gamma_p : X \to X \times Y$. The composite

$$X \xrightarrow{\Gamma_p} X \times Y \xrightarrow{\alpha\times\beta} k^{n+m} = k^n \times k^m \xrightarrow{\pi} k^m,$$

where π is the projection, coincides with the composite

$$X \xrightarrow{p} Y \xrightarrow{\beta} k^m$$

by definition of the graph. Since β is a closed immersion, $\int_p M \in D_h^b(\mathcal{D}_Y)$ if and only if $\int_\beta \circ \int_p M = \int_{\beta\circ p} M \in D_h^b(\mathcal{D}_{k^m})$. Since $\beta \circ p = \pi \circ (\alpha \times \beta) \circ \Gamma_p$, it is enough to check that $\int_{\pi\circ(\alpha\times\beta)\circ\Gamma_p} M = \int_\pi \circ \int_{(\alpha\times\beta)} \circ \int_{\Gamma_p} M \in D_h^b(\mathcal{D}_{k^m})$. But both Γ_p and $\alpha \times \beta$ are closed immersions, so we already know $\int_{(\alpha\times\beta)} \circ \int_{\Gamma_p} M \in D_h^b(\mathcal{D}_{k^{n+m}})$. We are therefore reduced to the case where p is the projection $\pi : k^{n+m} \to k^m$. By factoring this projection into n projections of relative dimension one, we may assume finally that $X = k^n = k \times k^{n-1}$, that $Y = k^{n-1}$, and that $p : X \to Y$ is the projection onto the final $n-1$ coordinates. Beginning now, we will not distinguish between the sheaf $\mathcal{D}_{k^n}$ and its ring of global sections D_{k^n} (the nth Weyl algebra); likewise for $\mathcal{D}_{k^{n-1}}$ and for modules over these two rings.

We write i for the closed immersion $\{0\} \times k^{n-1} = k^{n-1} \hookrightarrow k^n$ and j for the complementary open immersion $k^* \times k^{n-1} \hookrightarrow k^n$. What follows is a summary of the proof that the projection p preserves holonomy. We have a distinguished triangle

$$\int_i i^\dagger M \to M \to \int_j j^\dagger M \xrightarrow{+1}$$

which, upon taking cohomology, induces an exact sequence

$$0 \to h^0(\int_i i^\dagger M) \to M \to h^0(\int_j j^\dagger M) \to h^1(\int_i i^\dagger M) \to 0$$

since M is concentrated in degree 0. We know M is holonomic, so if $h^0(\int_j j^\dagger M)$ is holonomic, then $h^l(\int_i i^\dagger M)$ is holonomic for $l = 0, 1$ by the exact sequence above. As $\int_i i^\dagger M$ has nontrivial cohomology only in degrees $0, 1$ ($\int_i$ is exact, and $i^\dagger$ has cohomological dimension equal to the codimension of i), we have $\int_i i^\dagger M \in D_h^b(\mathcal{D}_{k^n})$. Since i is a closed immersion, this implies $i^\dagger M \in D_h^b(\mathcal{D}_{k^{n-1}})$.

For the proof outline above to be complete, we would need to fill the following gaps:

(a) if M is holonomic, $h^0(\int_j j^\dagger M)$ is holonomic (here $h^0(\int_j j^\dagger M)$ is just the localization M_{x_1}, but proving this is holonomic takes some work);
(b) if $i^\dagger M \in D_h^b(\mathcal{D}_{k^{n-1}})$, then $\int_p M \in D_h^b(\mathcal{D}_{k^{n-1}})$ as well.

The key to filling in the gaps above is the Fourier transform operation. Given a $\mathcal{D}_{k^n}$-module M, its *Fourier transform* $\widehat{M}$ is the $\mathcal{D}_{k^n}$-module whose underlying Abelian group is the same but whose $\mathcal{D}$-action is defined by $x_i * m = -\partial_i \cdot m$, $\partial_i * m = x_i \cdot m$. Here $\{x_i, \partial_i\}_{i=1}^n$ is a coordinate system for k^n, $*$ defines the $\mathcal{D}$-action on $\widehat{M}$, and the minus sign is necessary for the new action to respect the relations in $\mathcal{D}$. It is a general fact that a coherent $\mathcal{D}_{k^n}$-module M is holonomic if and only if $\widehat{M}$ is as well. We can pass from the cohomology of $\int_p M$ to the cohomology of $i^\dagger M$ by a sequence of Fourier transforms and degree shifts, which is how (b) is proved.

In order to prove (a) and the fact that Fourier transform respects holonomy, one uses the *Bernstein filtration*, where both x_i and ∂_i have degree 1, rather than the usual order filtration. If F is a good filtration on a coherent $\mathcal{D}_{k^n}$-module M with respect to the Bernstein filtration, then $F_l M$ is a finite-dimensional k-space for all l, which allows us to use Hilbert function techniques to study such filtered modules (there is a criterion for holonomy of M based on the asymptotic growth of the k-dimensions of $F_l M$). □