D-modules

Christopher Elliott

April 27th, 2010

Acknowledgements. I would like to thank my supervisor Prof. Grojnowski for many pieces of helpful advice. I would also like to thank S. Glasman for a number of useful conversations, and A. Shannon for encouragement, and suggestions on the example in Chapter 7.

Contents

1 Introduction

The origin of D-modules is in the work of the Japanese school of Mikio Sato in the mid-twentieth century on *algebraic analysis*. The aim of this program was to understand systems of linear partial differential equations on manifolds, and their generalisations, using the techniques of algebraic geometry and sheaf theory. One approach that proved very fruitful not only for the study of analysis, but also for algebraic geometry, Hodge theory and representation theory was the notion of a D-module. Essentially, we define a sheaf of *differential operators* on a complex manifold or smooth algebraic variety, and consider modules over that sheaf. Systems of linear PDEs arise as an example when the variety is affine n-space.

The derived category of D-modules on a smooth algebraic variety has a rich structure. If $f\colon X \to Y$ is a morphism of smooth varieties, we can push forward and pull back a D-module along the morphism. There are also a number of other natural functors between categories of D-modules which together make up a version of the so-called "six operations formalism" of Grothendieck.

The *Riemann-Hilbert correspondence* is the most important result presented in this essay. As the 21^{st} of his famous problems, Hilbert proposed the classical question "Can one always produce a Fuchsian differential operator with prescribed singularities and monodromy?" The Riemann-Hilbert correspondence as proved by Kashiwara and Mebkhout uses the concept of D-modules to answer and vastly generalise this. Indeed, we first introduce the category of *holonomic* D-modules. Although the solution space of a system of ordinary equations is always finite dimensional, when we pass to higher dimensions this is no longer necessarily true. Holonomic D-modules give the right generalisation of this notion of finite dimensional solution space.

With an extra *regularity* condition governing the growth behaviour near singularities, we produce the full subcategory of the derived category consisting of complexes of D-modules with *regular holonomic cohomology*. The Riemann-Hilbert correspondence says that this category is equivalent to the derived category of complexes of $\mathbb{C}$-modules with *constructible cohomology*, in a way that is compatible with many natural functors we can define on the two categories. This gives us a relationship between the geometric category of D-modules and the topological category of constructible sheaves. In particular, there is a subcategory of the constructible derived category consisting of so-called *perverse sheaves*, which are extremely interesting when understanding the topology of a space. The Riemann-Hilbert correspondence gives us a deep relationship between perverse sheaves and regular holonomic D-modules.

The main application of D-modules explained in this essay is to Hodge theory. There is a result of Griffiths which says that if we have a one-parameter family of smooth varieties $f : X \to \mathbb{C} \setminus \Sigma$ for Σ a finite set, then the cohomology of the fibres $H^w(X_\lambda)$ forms firstly a local system on $\mathbb{C} \setminus \Sigma$, but even more so, a *variation of Hodge structure*, which roughly means a family of Hodge structures on the points, varying continuously and compatibly with the usual Hodge structures on each $H^w(X_\lambda)$. The theory of D-modules can be used to help us generalise this to explain what happens at *singular fibres*. We can introduce the notion of the *Kashiwara-Malgrange filtration* on a D-module along a subvariety. This is a way of understanding this singular behaviour, and is intimately related to the *vanishing cycles* of Deligne. This helps us understand the topology, but to understand the Hodge theory along the singular fibres an even more sophisticated method had to be developed: the *mixed Hodge modules* invented and studied by Saito. The definition of these objects is rather elaborate, and will only be touched on in this essay.

Another important application of this theory is to the representation theory of semisimple Lie algebras. Although this will not be covered in this essay, I give a brief summary here. Let $\mathfrak{g}$ be a semisimple Lie algebra. Let $M(\lambda)$ be the *Verma module* for $\mathfrak{g}$ of weight λ, and let $L(\lambda)$ be its unique irreducible proper quotient. What is the character of $L(\lambda)$? Alternatively, what is the multiplicity $[M(\lambda) \colon L(\lambda)]$?

A conjectural answer to these question was given by Kazhdan and Lusztig, but it was only after work of Beilinson and Bernstein that progress could be made. They proved that there exist certain regular holonomic D-modules whose rings of global sections (given natural $\mathfrak{g}$-module structures) were isomorphic to $M(\lambda)$ and $L(\lambda)$. Not only that, but under the Riemann-Hilbert correspondence, these D-modules correspond to particularly nice perverse sheaves, understandable and amenable to computation. With this, one can convert difficult questions about infinite-dimensional representations of Lie algebras to much more tractable questions about perverse sheaves. It was using this machinery that the Kazhdan-Lusztig conjecture was proved. For details on this, see [16], [12] and the second half of [8].

For the first part of this essay, I essentially follow the methods of the recent book [8] by Hotta, Takeuchi and Tanisaki, with frequent reference to the unpublished notes [3] of Bernstein. When considering the Kashiwara-Malgrange filtration and applications to Hodge theory, I referred extensively to the book [15] of Peters and Steenbrink, and the notes [17] of Sabbah.

2 D-module fundamentals

In this section we will present the elementary definitions of the sheaf D_X and modules over it, and the elementary operations that can be performed, namely inverse and direct images and tensor product. The key theorem of this section is Kashiwara's theorem (2.26), which explains how the category of coherent D-modules on a variety behaves under closed embeddings. All algebraic varieties are over the field of complex numbers. Although in this chapter the results will also be true for a general algebraically closed field of characteristic zero, later on we will use the fact that smooth complex varieties can be given the structure of complex manifolds. Although there is a rich theory of D-modules in characteristic p, we will not touch on this at all in this essay.

2.1 D_X and D_X-modules

Definition 2.1. Let X be a smooth algebraic variety (over $\mathbb{C}$). We define several sheaves of a natural geometric origin that will be important in what follows. Firstly, let $\mathcal{E}nd(\mathcal{O}_X)$ denote the sheaf of endomorphisms of $\mathcal{O}_X$. We can naturally identify the structure sheaf $\mathcal{O}_X$ with a subsheaf of $\mathcal{E}nd(\mathcal{O}_X)$ by

$$f \in \mathcal{O}_X \rightsquigarrow (g \mapsto fg).$$

Define the *(pre)sheaf of derivations* of X on an open set U by

$$\mathcal{E}nd(\mathcal{O}_X)(U) \supseteq \mathcal{D}er(\mathcal{O}_X)(U) = \{\theta \in \mathcal{E}nd(\mathcal{O}_X) \,:\, \theta(fg) = f\theta(g) + \theta(f)g \;\forall f, g \in \mathcal{O}_X(U)\}.$$

In fact, this presheaf is a coherent sheaf of $\mathcal{O}_X$-modules, which we more commonly denote Θ_X, the *tangent sheaf* or *sheaf of vector fields* on X. Dually, we define the *cotangent sheaf* or *sheaf of 1-forms* to be the $\mathcal{O}_X$-module $\mathcal{H}om_{\mathcal{O}_X}(\Theta_X, \mathcal{O}_X)$, denoted Ω^1_X. We will often need to refer to the top exterior power of Ω^1_X, i.e. $\bigwedge^n \Omega^1_X$ where n is the dimension of X. We call this the *canonical sheaf*, and denote it ω_X.

Definition 2.2. Let X be a smooth algebraic variety. The *sheaf of differential operators* on X, denoted D_X, is defined to be the subsheaf of $\mathcal{E}nd(\mathcal{O}_X)$ generated by $\mathcal{O}_X$ and Θ_X as a $\mathbb{C}$-algebra.

We will often refer to the following motivating example when seeking geometric intuition:

Example 2.3. Let $X = \mathbb{C}^n$, affine n-space. What does D_X look like in this case? $\mathcal{O}_X = \mathbb{C}[x_1, \ldots, x_n]$, the polynomial algebra in n variables, and the derivations of this algebra are precisely those generated by the formal differential operators

$$\partial_i(x_1^{a_1} \cdots x_n^{a_n}) = a_i x_1^{a_1} \cdots x_i^{a_i - 1} \cdots x_n^{a_n} \quad \text{for } i = 1, \ldots n$$

extended linearly to endomorphisms of all of $\mathbb{C}[x_1 \dots, x_n]$. Thus the ring D_X is the ring

$$\mathbb{C}[x_1, \dots x_n; \partial_1, \dots, \partial_n],$$

a subring of the endomorphism ring. We can describe this formally as the ring generated over $\mathbb{C}$ by elements $x_1, \dots, x_n$ and $\partial_1, \dots, \partial_n$ with commutation relations

$$[x_i, x_j] = 0 = [\partial_i, \partial_j], \quad [x_i, \partial_j] = \delta_{ij}.$$

The algebra $D_{\mathbb{C}^n}$ is usually written D_n, and called the nth *Weyl algebra.*

This allows us to describe D_X locally for a general smooth complex variety X. Indeed, consider a point $x \in X$. Then we can always find *some* affine open neighbourhood U on which we have *local coordinates* $x_1, \dots, x_n \in \mathcal{O}_X(U)$ and $\partial_1, \dots, \partial_n \in \Theta_X(U)$. That is, sections such that

$$[x_i, x_j] = [\partial_i, \partial_j] = 0 \text{ and } \partial_i(x_j) = \delta_{ij}.$$

Furthermore, we can choose them so that the ∂_i generate $\Theta_X(U)$ and the x_i generate the maximal ideal $\mathfrak{m}_x$ of $\mathcal{O}_{X,x}$ ([8] A.5.2).

Remark 2.4. This description leads to a natural filtration F on the sheaf D_X. On a neighbourhood V with local coordinates $\{x_i, \partial_j\}$ as described above one can write

$$D_V = \bigoplus_{\alpha \in \mathbb{N}^n} \mathcal{O}_V \partial^\alpha$$

where $\partial^\alpha = \partial_1^{\alpha_1} \cdots \partial_n^{\alpha_n}$ is multi-index notation. Let U be an open subset of X. Then define

$$F_i D_X(U) = \left\{ P \in D_X(U) \, : \, P \in \bigoplus_{|\alpha| \leq i} \mathcal{O}_V \partial^\alpha \text{ for all } V \subseteq U \text{ with local coordinates} \right\}.$$

In the affine case this is just the *order filtration* by order of a differential operator. F is an increasing exhaustive filtration (i.e. $F_i D_X \subseteq F_{i+1} D_X$, and $\bigcup F_i D_X = D_X$), and $(F_i D_X)(F_j D_X) \subseteq (F_{i+j} D_X)$.

Remark 2.5. We can form the *associated graded module* $\mathrm{gr}^F(D_X)$ with respect to this filtration:

$$\mathrm{gr}^F(D_X) = \bigoplus_{i=1}^{\infty} F_i D_X / F_{i-1} D_X.$$

This is naturally isomorphic to the sheaf $\pi_* \mathcal{O}_{T^*X}$, where $\pi \colon T^*X \to X$ is the cotangent bundle. To see this we simply look in local co-ordinates.

Example 2.6. $\mathrm{gr}^F(D_n) \cong \mathbb{C}[\xi_1, \dots, \xi_{2n}]$, the polynomial ring in $2n$ variables.

Definition 2.7. A (left) *D-module* M on a smooth complex variety X (our fundamental object of study) is a sheaf of (left) D_X-modules.

Examples 2.8. Let $X = \mathbb{C}^n$, so $D_X = D_n$. Then we have the following examples:

1. $M = \mathcal{O}_X$ with the natural action of x_i, and trivial action of ∂_i.

2. $M = D_X$ with the natural action of both x_i and ∂_i.

3. On $X = \{pt\}$, $D_{\{pt\}} \cong \mathbb{C}$, so the D-modules on $\{pt\}$ are simply the complex vector spaces. Similarly, if X is now a general variety, $p \in X$, and V is a complex vector space, we can consider the *skyscraper sheaf* of V at p as a D_X-module with the trivial action. We will see in section 2.3 that these are the direct image of $D_{\{pt\}}$-modules under the inclusion $pt \mapsto p$.

4. Let $P_1 \dots P_m$ be a set of differential operators in D_n (i.e. a system of linear PDEs in n variables). Then the quotient
$$D_n / \sum_{i=1}^{m} D_n P_m$$
is naturally a D_n-module. We will see when we study the solution functor that this is the module we must consider when applying the theory of D-modules to the classical theory of linear PDEs.

A very natural way of interpreting D-modules is as $\mathcal{O}$-modules with additional structure. Written out explicitly this has the following form:

Proposition 2.9. Let M be a sheaf of $\mathcal{O}_X$-modules. Giving M the structure of a D_X-module is equivalent to describing a $\mathbb{C}$-linear morphism $\nabla\colon \Theta_X \to \mathcal{E}nd(M)$ such that

1. $\nabla_{f\theta}(m) = f\nabla_\theta(m)$
2. $\nabla_\theta(fm) = \theta(f)m + f\nabla_\theta(m)$
3. $(\nabla_{[\theta_1,\theta_2]} - [\nabla_{\theta_1}, \nabla_{\theta_2}])(m) = 0$

for all $f \in \mathcal{O}_X$, θ and $\theta_i \in \Theta_X$, and $m \in M$

Proof. A D_X-module structure on M is an action of Θ_X on M compatible with the $\mathcal{O}_X$-action. Let U be an affine open subset. On U, $D_X(U)$ is generated by $\mathcal{O}_X(U)$ and $\Theta_X(U)$ with the relation $[\theta, f] = \theta(f)$ for $\theta \in \Theta_X(U)$, $f \in \mathcal{O}_X(U)$. But this is precisely the condition (2) above. Interpreting $\nabla_\theta(m)$ as the action θm, the result is then trivial. □

Remark 2.10. The choice of notation above is deliberately evocative. Indeed, consider the case when M is locally free of finite rank as an $\mathcal{O}_X$-module, i.e. equivalent to the sheaf of sections of a *vector bundle* on X. Then the proposition tells us that a D-module structure on M is precisely a *flat* or *integrable connection* on the bundle.

We could equally well have chosen to consider *right* D_X-modules. There is an important example of a right D_X-module to consider.

Example 2.11. The canonical sheaf ω_X has a well-known right D_X-module structure, via the *Lie derivative*:
$$\mathrm{Lie}_\theta(\eta)(\theta_1, \dots, \theta_n) = \theta(\eta(\theta_1, \dots, \theta_n)) - \sum_{i=1}^{n} \eta(\theta_1, \dots, [\theta, \theta_i], \dots, \theta_n).$$
That this defines a right D_X-module is a simple check by the obvious right module analogue of proposition 2.9.

As the following construction shows, we can pass between the settings of left and right D_X-modules freely:

Proposition 2.12. Let X be a smooth variety. Then there is an equivalence of categories
$$\mathrm{Mod}(D_X) \to \mathrm{Mod}(D_X^{op})$$
(where $\mathrm{Mod}(D_X^{op})$ is naturally identified with the category of *right* D_X modules), given by
$$\omega_X \otimes_{\mathcal{O}_X} (-)$$
with quasi-inverse
$$\omega_X^\vee \otimes_{\mathcal{O}_X} (-).$$
Here we use the notation $\omega_X^\vee$ to denote the dual $\mathcal{O}_X$-module $\mathcal{H}om_{\mathcal{O}_X}(\omega_X, \mathcal{O}_X)$.

Remark 2.13. In a neighbourhood U with local co-ordinates, this is just the fact that specifying a left D_U-module structure on some $\mathcal{O}_U$-module M is equivalent to specifying a right D_U-module structure via the *formal adjoint*. Explicitly, if $P \in D_U$ is a differential operator, we form its formal adjoint via $x_i \mapsto x_i, \partial_i \mapsto -\partial_i$, so

$$\sum (f_\alpha(x_1, \ldots, x_n)\partial^\alpha) \mapsto \sum (-\partial)^\alpha f_\alpha(x_1, \ldots x_n).$$

Then it is easy to check that P acting on the left is equivalent to the formal adjoint of P acting on the right.

The reason the global version of this statement needs to appeal to the canonical sheaf is that in general the formal adjoint of a differential operator acts naturally on ω_X.

Proof. Let M be a left D_X-module and N be a right D_X-module. Observe first that $\omega_X \otimes_{\mathcal{O}_X} M$ is a right D_X-module via the action

$$(\eta \otimes m)\theta = \mathrm{Lie}_\theta\, \eta \otimes m - \eta \otimes \theta m \quad \theta \in \Theta_X.$$

A similar result holds for $\omega_X^\vee \otimes_{\mathcal{O}_X} N \cong \mathcal{H}om_{\mathcal{O}_X}(\omega_X, N)$, i.e. it has left D_X-module structure via

$$\theta(\phi)(\eta) = \phi(\eta)\theta + \phi(\mathrm{Lie}_\theta\, \eta) \quad \theta \in D_X,\ \phi \in \mathcal{H}om_{\mathcal{O}_X}(\omega_X, N),\, \eta \in \omega_X.$$

Now, by associativity of the tensor product, there is an isomorphism

$$\omega_X^\vee \otimes_{\mathcal{O}_X} \omega_X \otimes_{\mathcal{O}_X} M \cong M$$

as required. □

2.2 Inverse images

Let $f \colon X \to Y$ be a morphism of smooth varieties, and let M be a D_Y-module. We'd like to pull M back under f to a D_X -module f^*M, analogously to the inverse image of $\mathcal{O}$-modules. For inverse images this will work rather simply: we can give a D-module structure to the $\mathcal{O}$-module inverse image.

Define f^*M to be the $\mathcal{O}$-module inverse image

$$\mathcal{O}_X \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}M.$$

We must introduce a D-module structure to this $\mathcal{O}$-module. First, we describe the action of some $\theta \in \Theta_X$ on a typical element $\phi \otimes f^{-1}(m)$ for $\phi \in \mathcal{O}_X$, $m \in M$ in a set of local co-ordinates (y_i, ∂_i) for Y. Put

$$\theta(\phi \otimes f^{-1}(m)) = \theta(\phi) \otimes f^{-1}(m) + \phi \sum_{i=1}^{n} \theta(f^{-1}(y_i) \otimes \partial_i m).$$

Indeed, by 2.9 this defines a local D_X-module structure. Essentially, if we stare at this we can convince ourselves that this is just a form of the chain rule. It suffices to describe the action of $\theta \in \Theta_X$ locally to give a D_X-module structure, so we have produced the D-module inverse image.

Define the *transfer module* $D_{X\to Y}$ to be the module

$$f^*D_Y = \mathcal{O}_X \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}D_Y.$$

This is a left D_X module as described above, but also a right $f^{-1}D_Y$-module via right multiplication on the second multiplicand. This gives an alternative description of the inverse image: namely

$$\begin{aligned} f^*M &= \mathcal{O}_X \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}M \\ &= \mathcal{O}_X \otimes_{f^{-1}\mathcal{O}_Y} (f^{-1}D_Y \otimes_{f^{-1}D_Y} f^{-1}M) \\ &= D_{X\to Y} \otimes_{f^{-1}D_Y} f^{-1}M. \end{aligned}$$

Thus we have defined a right exact functor

$$f^*\colon \operatorname{Mod}(D_Y) \to \operatorname{Mod}(D_X).$$

Remark 2.14. Notationally I am not distinguishing between the inverse image functor for $\mathcal{O}$-modules and the inverse image functor for D-modules, but this will not cause any ambiguity.

Lemma 2.15. As in the $\mathcal{O}$-module case, if $f\colon X \to Y$ and $g\colon Y \to Z$ are morphisms of smooth varieties, then $(fg)^* = g^* f^*$.

The proof is an easy manipulation from the definitions, and is omitted.

Example 2.16. We compute the transfer module $D_{X\to Y}$ locally in the case of $\iota : X \to Y$ a closed embedding. By taking local co-ordinates (y_i, ∂_i) on Y, we must calculate the transfer module under the standard embedding $\iota : \mathbb{C}^n \to \mathbb{C}^N$. Split up the ring D_N as

$$\begin{aligned} D_N &\cong \bigoplus_{m_1,\dots,m_n} \mathcal{O}_{\mathbb{C}^n} \partial_1^{m_1} \cdots \partial_n^{m_n} \otimes_{\mathbb{C}} \mathbb{C}[\partial_{n+1}, \dots, \partial_N] \\ &= D' \otimes_{\mathbb{C}} \mathbb{C}[\partial_{n+1}, \dots, \partial_N] \end{aligned}$$

where the isomorphism is in the category of left D'-modules. Then

$$\begin{aligned} D_{n\to N} &= \mathcal{O}_n \otimes_{\iota^{-1}\mathcal{O}_N} \iota^{-1} D_N \\ &\cong \mathcal{O}_n \otimes_{\iota^{-1}\mathcal{O}_N} \iota^{-1}(D' \otimes_{\mathbb{C}} \mathbb{C}[\partial_{n+1}, \dots, \partial_N] \\ &\cong (\mathcal{O}_n \otimes_{\iota^{-1}\mathcal{O}_N} \iota^{-1} D') \otimes_{\mathbb{C}} \mathbb{C}[\partial_{n+1}, \dots, \partial_N] \\ &\cong D_n \otimes_{\mathbb{C}} \mathbb{C}[\partial_{n+1}, \dots, \partial_N] \end{aligned}$$

where $D_{n\to N}$ denotes $D_{\mathbb{C}^n \to \mathbb{C}^N}$ and $\mathcal{O}_n$ denotes $\mathcal{O}_{\mathbb{C}^n}$. The left module structure is given by multiplication on the left component $D_n \cong \mathcal{O}_n \otimes_{\iota^{-1}\mathcal{O}_N} \iota^{-1} D'$. The right module structure is defined by considering $P \in \iota^{-1} D_N$ as living in $\iota^{-1}(D') \otimes_{\mathbb{C}} \mathbb{C}[\partial_{n+1}, \dots, \partial_N]$.

2.3 Direct images

Let $f\colon X \to Y$ be a morphism of smooth varieties as before. Motivated by the above, we might first try to define the D-module direct image f_* in the following way: let M be a *right* D_X-module. The transfer module $D_{X\to Y}$ is a right $f^{-1}D_Y$-module, so we can produce a right D_Y-module via

$$f_*(M \otimes_{D_X} D_{X\to Y}),$$

where f_* here is the sheaf theoretic direct image. Using the side-changing operations defined in 2.12 we can even convert this into a functor from left D_X-modules to left D_Y-modules. Indeed, define the *transfer module*

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

a $(f^{-1}D_Y, D_X)$-bimodule. Then we produce a left D_Y-module via

$$f_*(D_{Y\leftarrow X} \otimes_{D_X} M).$$

However, this definition is not the right one. In order to do homological algebra with our direct image we'd like a left exact functor. However, due to the mixture between the left exact functor f_* and the right exact functor $\otimes_{D_X}$, the naïve direct image defined above is not exact on either side. An example of the resulting bad behaviour is the failure of the analogue of lemma 2.15 for this functor.

To remedy this, we work in the *derived category* of chain complexes of D_X-modules, which we will denote $D(D_X)$. Usually we will be interested in the bounded derived category $D^b(D_X)$, or the bounded derived category with coherent cohomology $D^b_{coh}(D_X)$. Recall that the bounded derived category is the category obtained from the homotopy category $K^b(D_X)$ of bounded complexes of D_X-modules by formally inverting the quasi-isomorphisms. We quickly review these notions, which are explained fully in [11].

Definition 2.17. Let X be a smooth variety, and let A be a sheaf of rings on X, usually D_X or $\mathbb{C}_X$. Define $K(A)$ to be the homotopy category of complexes of A-modules. It has full subcategories $K^+(A)$, $K^-(A)$ and $K^b(X)$ of complexes bounded below, bounded above, and bounded on both sides. We can form the *derived category* of any of these categories by formally localising with respect to the system of quasi-isomorphisms, as described in [11] sections 1.6 and 1.7. We denote these derived categories $D(A), D^+(A), D^-(A)$ and $D^b(A)$ respectively.

Note that all these categories possess the structure of a *triangulated category.* From now on I will assume familiarity with the language of triangulated categories, as in [11] 1.5.

We will also refer to the full subcategories $D^\sharp_{qcoh}(A)$ and $D^\sharp_{coh}(A)$ of complexes with *quasicoherent* or *coherent* cohomology respectively, where $\sharp$ is any of $+, -, b$ or nothing.

Definition 2.18. We define *truncation functors* $\tau^{\leq k}, \tau^{\geq k}$ for each k by

$$\tau^{\leq k}(\cdots \to F^{k-1} \to F^k \overset{\partial^k}{\to} F^{k+1} \to \cdots) = (\cdots \to F^{k-1} \to F^k \to \operatorname{im}(\partial^k) \to 0)$$
$$\tau^{\geq k}(\cdots \to F^{k-1} \overset{\partial^{k-1}}{\to} F^k \to F^{k+1} \to \cdots) = (0 \to \ker(\partial^{k-1}) \to F^k \to F^{k+1} \to \cdots).$$

We may also refer to $\tau^{<k} = \tau^{\leq k-1}$ or $\tau^{>k} = \tau^{\geq k+1}$. The truncation functors $\tau^{\leq 0}, \tau^{\geq 0}$ define a *t-structure* on the triangulated category $D^b_{coh}(D_X)$. See for instance [8] chapter 8 or [11] chapter 10 for details.

Definition 2.19. For $f\colon X \to Y$ a morphism of smooth varieties, the D-module *direct image* functor

$$f_*\colon D^b(D_X) \to D^b(D_Y)$$

is defined by

$$f_* M^\bullet = Rf_*(D_{Y\leftarrow X} \otimes^L_{D_X} M^\bullet)$$

for $M^\bullet \in D^b(D_X)$.

We can also produce derived versions of the inverse image functor, namely

$$f^* N^\bullet = D_{X\to Y} \otimes^L_{f^{-1}D_Y} f^{-1}N^\bullet$$

and the shifted version

$$f^\dagger N^\bullet = f^* N^\bullet[\dim X - \dim Y]$$

which will often simplify notation. I am deliberately overloading the notation f_* and f^*. From now on, unless otherwise specified, these notations will refer to the functors between *derived* categories just defined.

To produce actual D-modules rather than complexes, we can take the cohomology of these complexes.

Example 2.20. Consider the closed embedding $\iota\colon \mathbb{C}^n \to \mathbb{C}^N$. In example 2.16, we computed the transfer module $D_{n\to N}$ to be $D_n \otimes_{\mathbb{C}} \mathbb{C}[\partial_{n+1}, \ldots, \partial_N]$. Similarly, $D_{N\leftarrow n}$ can be calculated by an analogous procedure to give

$$D_{N\leftarrow n} \cong \mathbb{C}[\partial_{n+1}, \ldots, \partial_N] \otimes_{\mathbb{C}} D_n$$

with $(\iota^{-1}D_N, D_n)$-bimodule structure given by right multiplication on the right component and left action via the isomorphism

$$D_{N\leftarrow n} \cong \iota^{-1}(D_N \otimes_{\mathcal{O}_N} \omega^\vee_{\mathbb{C}^N}) \otimes_{\iota^{-1}\mathcal{O}_N} \omega_{\mathbb{C}^n}$$

which holds more generally than this specific case. This gives a local description of a general closed embedding $\iota\colon X \to Y$. Indeed, we can go further. Let M be a D_X-module. We can describe $H^0\iota_*M$ locally using this:

$$\begin{aligned} H^0\iota_*M &\cong \iota_*(\mathbb{C}[\partial_n+1,\dots,\partial_N]\otimes_{\mathbb{C}} D_n \otimes_{D_n} M) \\ &\cong \mathbb{C}[\partial_n+1,\dots,\partial_N]\otimes_{\mathbb{C}} \iota_*(M). \end{aligned}$$

For $j > 0$, $H^j\iota_*M = 0$ locally, and hence globally.

Recall that a sheaf of $\mathcal{O}_X$-modules F on X is called *coherent* if it is locally finitely generated, and for any open $U \subseteq X$, and any locally finitely generated submodule of $F|_U$, is locally finitely presented.

Example 2.21. In general coherence is *not* preserved under either direct or inverse images of D-modules. Indeed consider the closed embedding $\iota\colon \{pt\} \to \mathbb{C}$. We calculated that as a $\mathbb{C}$-algebra,

$$\iota^*D_1 = D_{0\to 1} \cong \mathbb{C}[x]\otimes_{\mathbb{C}}\mathbb{C}[y]$$

which is not finitely generated. For direct images, even coherence as $\mathcal{O}$-modules is not in general preserved, as (for instance) there exist varieties whose structure sheaves have infinitely generated rings of global sections [20].

Remark 2.22. For certain classes of morphisms the subcategory of coherent D-modules *is* preserved under direct and inverse images. For example:

- Suppose $f\colon X \to Y$ is a proper morphism of smooth varieties, and suppose $M^\bullet$ is a coherent complex of D_X-modules. Then $f_*M^\bullet$ is a coherent complex of D_Y-modules. Essentially, this is proved in the quasi-projective case by factoring f as the composite of a closed embedding and a projection, and checking the result locally in both cases (see [8] theorem 2.5.1). In particular, for $\iota\colon X \to Y$ a closed embedding, ι_* preserves coherence.
- Now, suppose $f\colon X \to Y$ is a smooth morphism of smooth varieties. Then the functor f^* preserves coherence.

2.4 Tensor products

We will still need a few more functors on categories of D-modules. Here we introduce the notion of a tensor product of D-modules.

Recall the *exterior tensor product* of $\mathcal{O}$-modules. This is a bifunctor

$$\bullet\boxtimes\bullet\colon \operatorname{Mod}(\mathcal{O}_X)\times\operatorname{Mod}(\mathcal{O}_Y)\to\operatorname{Mod}(\mathcal{O}_{X\times Y})$$

defined by

$$M\boxtimes N = \mathcal{O}_{X\times Y}\otimes_{\pi_1^{-1}\mathcal{O}_X\otimes\pi_2^{-1}\mathcal{O}_Y}(\pi_1^{-1}M\otimes_{\mathbb{C}}\pi_2^{-1}N),$$

where π_1 and π_2 are the projections from $X\times Y$ to its two factors. We'd like to define a similar notion for a tensor product of D-modules.

For M a D_X-module, N a D_Y-module, we can form the *exterior tensor product* of D-modules analogously:

$$M\boxtimes N = D_{X\times Y}\otimes_{\pi_1^{-1}D_X\otimes\pi_2^{-1}D_Y}(\pi_1^{-1}M\otimes_{\mathbb{C}}\pi_2^{-1}N).$$

Indeed, as an $\mathcal{O}_{X\times Y}$-module this is isomorphic to the exterior tensor product of $\mathcal{O}$-modules defined above, but it comes equipped with a natural $D_{X\times Y}$-module structure via $D_{X\times Y}\cong D_X\boxtimes D_Y$. Thus we have defined a functor

$$\bullet\boxtimes\bullet\colon \operatorname{Mod}(D_X)\times\operatorname{Mod}(D_Y)\to\operatorname{Mod}(D_{X\times Y}).$$

In fact, this is exact on both factors ([8] pp38-9), so extends to a functor on derived categories

$$\bullet\boxtimes\bullet\colon D^b(D_X)\times D^b(D_Y)\to D^b(D_{X\times Y}).$$

2.5 Coherent D_X-modules

For most of the rest of this report we will study the full subcategory of *coherent* D_X-modules. From this point on, we assume for simplicity that all varieties are quasi-projective.

Proposition 2.23. D_X is coherent as a sheaf of rings.

Proof. We will actually prove a stronger statement

Claim. If M is a locally finitely generated D_X-module which is quasicoherent over $\mathcal{O}_X$ then it is coherent over D_X.

To see this, let U be an affine open subset of X. Let $f\colon D_U^m \to M|_U$ be any homomorphism. We must prove that $K = \ker f$ is finitely generated, i.e. there exists an exact sequence

$$D_U^n \to D_U^m \xrightarrow{f} M|_U.$$

But the functor $\Gamma(U, -)$ is exact, so this exact sequence exists if and only if an exact sequence of the form

$$D_X(U)^n \to D_X(U)^m \xrightarrow{f(U)} M(U)$$

exists, with $f(U)$ the induced map. This however is immediate, as $D_X(U)$ is a (left) Noetherian ring, so the kernel of $f(U)$ is finitely generated. □

We already encountered in 2.10 a special subcategory of $\operatorname{Mod} D_X$, namely the full subcategory of *integrable connections*, denoted $\operatorname{Conn}(X)$. These are the D_X modules with are locally free of finite rank as $\mathcal{O}_X$-modules, i.e. essentially vector bundles on X equipped with flat connections. We can now give an alternative characterisation of this subcategory:

Proposition 2.24. A D_X-module M is coherent over $\mathcal{O}_X$ if and only if it is an integrable connection.

Proof (following [8] 1.4.10). Clearly integrable connections are coherent, so let M be an $\mathcal{O}$-coherent D_X-module. Take a point $x \in X$. It suffices to prove that for every such x, the stalk M_x is free over the local ring $\mathcal{O}_{X,x}$, as then M is locally free over $\mathcal{O}_X$ as required. Let $\mathfrak{m}$ be the maximal ideal of $\mathcal{O} = \mathcal{O}_{X,x}$. We can find generators $P_1, \dots, P_k$ for M_x over $\mathcal{O}$ such that their images $\overline{P_i}$ in the quotient $M_x/\mathfrak{m}M_x$ generate the quotient over $\mathbb{C} = \mathcal{O}/\mathfrak{m}$. Indeed, take any basis for the $\mathbb{C}$-vector space $M_x/\mathfrak{m}M_x$ and apply Nakayama's lemma to it.

Suppose for contradiction that $P_1, \dots P_k$ are linearly dependent. We will produce a non-trivial dependence relation among the $\overline{P_i}$, which is a contradiction. For a set $F = \{f_1, \dots f_k\} \in \mathcal{O}^k$, define

$$\operatorname{ord}(F) = \min_i (\max\{p : f_i \in \mathfrak{m}^p\})$$

so if $\operatorname{ord}(F) = 0$ then some f_i is not in $\mathfrak{m}$. Take a dependence relation

$$\sum_{i=1}^k f_i P_i = 0$$

and put $F = \{f_1, \dots f_k\}$. If we can always reduce $\operatorname{ord}(F)$ then we're done, as we can reduce $\operatorname{ord}(F)$ down to zero, and then reducing mod $\mathfrak{m}$ gives us a non-trivial dependence relation in the $\overline{P_i}$ as required. We do this by applying the operators ∂_j to get

$$\sum_{i=1}^k (\partial_j f_i) P_i + f_i(\partial_j P_i) = 0$$

which is a new dependence relation. We can always pick some j such that $\operatorname{ord}(F)$ decreases, because $\operatorname{ord}(f+g) = \min\{\operatorname{ord}(f), \operatorname{ord}(g)\}$, and there is always some ∂_j that reduces the order of a given single $f \in \mathcal{O}$. □

Remark 2.25. This immediately implies that $\operatorname{Conn}(X)$ is an abelian category, since kernels and cokernels of morphisms of D-modules coherent over $\mathcal{O}$ are themselves coherent over $\mathcal{O}$.

2.6 Kashiwara's theorem

One can check (e.g. [8] 1.5.25) that for $\iota\colon X \to Y$, the functor $\iota^\dagger$ is right adjoint to the functor ι_*. Obviously the pair do not in general describe an equivalence of categories, but in fact if we restrict to an appropriate subcategory of $D^b_{coh}(D_Y)$ then the functors *do* restrict to an equivalence. This result of Kashiwara could fairly be described as the first major result of the theory of D-modules.

We will use the notation $\mathrm{Mod}^X(D_Y)$ for the full subcategory of D_Y-modules whose support is contained in the image of X. Likewise, we have a derived version, $D^X(D_Y)$.

Theorem 2.26 (Kashiwara)**.** Let $\iota : X \to Y$ be a closed embedding of varieties. Then there is an equivalence of categories:

$$\iota_*\colon D^b_{coh}(D_X) \rightleftarrows D^{b,X}_{coh}(D_Y)\ : \iota^\dagger$$

and hence in particular

$$H^0\iota_*\colon \mathrm{Mod}_{coh}(D_X) \rightleftarrows \mathrm{Mod}^X_{coh}(D_Y)\ : H^0\iota^\dagger$$

Remark 2.27. Firstly, note that the functors above do indeed preserve coherence, by the remark 2.22. We should note that this result is somewhat surprising. Indeed, the analogous result for $\mathcal{O}$-modules is completely false. Take for example the closed embedding $\iota\colon \{pt\} \to \mathbb{C}$. Then the $\mathcal{O}$-modules on $\{pt\}$ are just complex vector spaces, and their direct images are just skyscraper sheaves of such vector spaces. However, there are many more $\mathcal{O}_\mathbb{C}$ modules supported only at the origin, for example

$$\mathbb{C}[x]/x^k \quad \text{for any } k \in \mathbb{N}.$$

Of course, in general these are *not* D-modules under the standard action of ∂.

Proof. We prove the latter statement first, and then induct on *cohomological length*, i.e.

$$\ell(M^\bullet) = \max\{i \in \mathbb{Z} : H^i(M^\bullet) \neq 0\} - \min\{i \in \mathbb{Z} : H^i(M^\bullet) \neq 0\}.$$

To prove this latter statement, we show that the unit and counit of the adjunction are isomorphisms. Since the problem is local, and by induction on the codimension of X in Y it suffices to prove the statement for the standard embedding of a hypersurface

$$\iota\colon \mathbb{C}^{n-1} \to \mathbb{C}^n \qquad \text{via } (x_1, \dots, x_n) \mapsto (x_1, \dots, x_n, 0).$$

That is, we must prove that the morphisms

$$M \mapsto H^0\iota^\dagger(\mathbb{C}[\partial] \otimes_\mathbb{C} \iota_* M) \quad \text{and} \quad \mathbb{C}[\partial] \otimes_\mathbb{C} \iota_* H^0\iota^\dagger N \mapsto N$$

are isomorphisms for $M \in \mathrm{Mod}_{coh}(D_n), N \in \mathrm{Mod}^{\mathbb{C}^n}_{coh}(D_{n+1})$, where ∂ denotes ∂_{n+1}.

Define the operator $\xi = y\partial$, where $y = y_{n+1}, \partial = \partial_{n+1}$. For $N \in \mathrm{Mod}^{\mathbb{C}^n}_{coh}(D_{n+1})$ we consider the eigenspaces

$$N^j = \{m \in N : \xi m = jm\}.$$

We'd like to show that

$$N = \bigoplus_{j \in \mathbb{N}} N^{-j} \cong \mathbb{C}[\partial] \otimes_\mathbb{C} N^{-1}$$

since then clearly $H^0\iota^\dagger N = \ker(i^{-1}N \xrightarrow{y} i^{-1}N) = i^{-1}N^{-1}$, which implies the result we want. To prove the direct sum decomposition we use an induction argument on k: firstly, notice that it suffices to prove that for each k

$$\ker(N \xrightarrow{y^k} N) \subseteq N^{-1} \oplus \cdots \oplus N^{-k}$$

because N is coherent and supported in $\mathbb{C}^n$, so any $m \in N$ is annihilated by some y^k. If $ym = 0$ then $\xi m = \partial y m - m = -m$, so $m \in N^{-1}$. For the induction step,

$$\begin{aligned} y^k m = 0 &\implies ym \in N^{-1} \oplus \cdots \oplus N^{-(k-1)} \\ &\implies \partial y m = \xi m + m \in N^{-1} \oplus \cdots \oplus N^{-k} \end{aligned}$$

$$\begin{aligned} \text{but} \quad & y^{k-1}(\xi m + km) = y^k \partial m + k y^{k-1} m = 0 \\ \text{so} \quad & \xi m + km \in N^{-1} \oplus \cdots \oplus N^{-(k-1)} \end{aligned}$$

$$\begin{aligned} &\implies km - m \in N^{-1} \oplus \cdots \oplus N^{-k} \\ &\implies m \in N^{-1} \oplus \cdots \oplus N^{-k} \quad \text{for } k > 1 \end{aligned}$$

which is what we wanted.

The induction on $\ell(M)$ step is now rather easy. We have proven the case $\ell(M^\bullet) = 0$, so suppose $\ell(M^\bullet) > 0$. We then split $M^\bullet$ into two pieces; that is, we pick k such that the truncation functors $\tau^{\leq k}$ and $\tau^{>k}$ both reduce cohomological length. Now consider the diagram

$$\begin{tikzcd} \tau^{\leq k}(M^\bullet) \arrow[r] \arrow[d, "\alpha_1"'] & M^\bullet \arrow[r] \arrow[d, "\alpha_2"'] & \tau^{>k}(M^\bullet) \arrow[r, "+1"] \arrow[d, "\alpha_3"'] & {} \\ \iota^\dagger \iota_* \tau^{\leq k}(M^\bullet) \arrow[r] & \iota^\dagger \iota_* M^\bullet \arrow[r] & \iota^\dagger \iota_* \tau^{>k}(M^\bullet) \arrow[r, "+1"] & {} \end{tikzcd}$$

whose rows are distinguished triangles. The induction hypothesis implies α_1 and α_3 are isomorphisms, hence so is α_2. Thus the unit is an isomorphism. An analogous argument shows the counit is also an isomorphism. □

Remark 2.28. This allows us to make sense of the notion of 'D-modules on a *general* variety X, not necessarily smooth'. Indeed, we can realise X as a closed subvariety of a smooth variety X', and define $\mathrm{Mod}(D_X)$ to be $\mathrm{Mod}^X(D_{X'})$. Then by Kashiwara's theorem we feel reassured that this is the right generalisation of the smooth case.

3 Holonomic D-modules

For a coherent D_X-module M we can introduce an associated geometric invariant called the *characteristic variety* associated to M. This is inspired by the classical case in the following way. If P is a linear differential operator on $\mathbb{C}$, we can consider the graded module $\mathrm{gr}^F(D_1/D_1P)$ with respect to $F^\bullet$ a so-called *good filtration*. The annihilator of this module cuts out a subvariety in $\mathbb{C}^2$, called the characteristic variety of the module D_1/D_1P. The important issue is whether this variety is one-dimensional, or all of $\mathbb{C}^2$. If it is one-dimensional, we say the D-module is *holonomic*, which in this classical case essentially corresponds to the differential equation being maximally overdetermined. Holonomic D-modules in general form a particularly well-behaved subcategory of coherent D-modules.

We will define good filtrations on D-modules and the characteristic variety, and prove it is independent of the choice of good filtration. We compute the characteristic variety in a number of examples. The main result of this section is that holonomicity of D-modules is invariant under all the natural operations so far defined, but to prove this we will need to understand the machinery of duality of D-modules. On the way we prove the key lemma of Bernstein on b-functions, which is of considerable independent interest.

3.1 Good filtrations

Let M be a quasi-coherent D_X-module. This aim of this section is to define a filtration on M that has some of the nice properties of the order filtration on D_X itself defined in 2.4. Given an exhaustive increasing filtration $F_\bullet$ on M, we can form the associated graded module

$$\mathrm{gr}^F(M) = \bigoplus_{i=1}^{\infty} F_i M/F_{i-1}M.$$

This is a module over the sheaf $\mathrm{gr}^F(D_X) \cong \pi_* \mathcal{O}_{T^*X}$ as noted in 2.5.

Definition 3.1. Let $(F_i)_{i\in\mathbb{N}}$ on M be an exhaustive increasing filtration on M such that

$$(F_i D_X)(F_j M) \subseteq F_{i+j}M.$$

$F_\bullet$ is called *good* if the associated graded module $\mathrm{gr}^F(M)$ is coherent.

Proposition 3.2. A quasi-coherent D_X-module M if coherent if and only if it admits a good filtration $F_\bullet$.

Proof. First, suppose M is coherent. Then M is generated by a submodule N which is coherent over $\mathcal{O}_X$ (this follows from standard algebro-geometric results, e.g. [8] 1.4.17). Define an exhaustive increasing filtration by

$$G^i(M) = (F^i D_X)(N)$$

where F^i is the standard filtration. This filtration is clearly good.

Conversely, suppose M is quasi-coherent, but not coherent. Find an open set U such that the ring $M|_U$ is not finitely-generated over $D_X|_U$. Let $F^\bullet$ be any increasing exhaustive filtration on $M|_U$. Suppose $\mathrm{gr}^F(M|_U)$ is finitely generated. Then there exists a finite set $\{m_i \in F^i M|_U \ : i = 0, \ldots, k\}$ such that $\{m_i/F^{i-1}M|_U\}$ generates $\mathrm{gr}^F(M|_U)$. But then, by induction on j,

$$F^j M|_U = \sum_{i\leq j} F^{j-1}M|_U m_i$$

which is a contradiction of $M|_U$ finitely generated. □

3.2 Characteristic Varieties

Let X be a smooth variety. Let M be a coherent D_X-module equipped with a good filtration F. We can associate to M an algebraic subset of the cotangent bundle T^*X called the *characteristic variety* as follows. Let

$$\mathrm{gr}^F M = \bigoplus_{i\in\mathbb{Z}} F^i M/F^{i-1}M$$

be the graded module over $\mathrm{gr}^F D_X = \pi_* \mathcal{O}_{T^*X}$, where $\pi\colon T^*X \to X$ is the canonical projection map. We can produce a coherent $\mathcal{O}_{T^*X}$-module by

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

Then the characteristic variety $\mathrm{Ch}(M)$ is the support of $\widetilde{\mathrm{gr}^F(M)}$, sometimes also called the *singular support* of M.

Remark 3.3. For affine D-modules, we can compute the characteristic variety in a more concrete way. If $M \in \mathrm{Mod}_{coh}(D_n)$, then its graded version is a module over the polynomial ring $\mathbb{C}[x_1, \ldots, x_{2n}]$. To calculate its characteristic variety we take the annihilator of $\mathrm{gr}(M)$, which is an ideal in the polynomial ring, and compute the variety in $\mathbb{C}^{2n}$ defined by this ideal.

Lemma 3.4. The characteristic variety of M is independent of the choice of good filtration F.

Proof. Following the method in [8] D.3.1, we prove this in two steps. Let F_i and F_i' be two good filtrations on M. We say F and F' are *adjacent* if for all i,

$$F_iM \subseteq F_i'M \subseteq F_{i+1}M.$$

We connect any two good filtrations by a finite chain of adjacent pairs of good filtrations, then prove the assertion for an adjacent pair.

Define a chain of good filtrations $F^{(k)}$ by

$$F_i^{(k)}M = F_iM + F_{i+k}'M.$$

Then for k large and negative $F^{(k)} = F$, and for k large and positive $F^{(k)} = F'$ up to a shift. Also, $F^{(k)}$ and $F^{(k+1)}$ are adjacent for every k, (see e.g. [8] D.1.3).

With this out of the way, consider the adjacent case. By definition we have natural homomorphisms $\theta_i \colon F_iM/F_{i-1}M \to F_i'M/F_{i-1}'M$ for each i, and hence $\theta \colon \operatorname{gr}^F M \to gr^{F'}M$. Consider the exact sequence

$$0 \longrightarrow \ker\theta \longrightarrow \operatorname{gr}^F M \xrightarrow{\theta} gr^{F'}M \longrightarrow \operatorname{coker}\theta \longrightarrow 0\,.$$

For each i, $\ker\theta_i \cong F_{i-1}'M/_{i-1}M \cong \operatorname{coker}\theta_{i-1}$ by the definition of adjacency. hence $\ker\theta \cong \operatorname{coker}\theta$, and so taking supports in the exact sequence above,

$$\begin{aligned}\operatorname{supp}(\operatorname{gr}^F M) &= \operatorname{supp}(\ker\theta) \cup \operatorname{supp}(\operatorname{im}\theta)\\ &= \operatorname{supp}(\operatorname{coker}\theta) \cup \operatorname{supp}(\operatorname{im}\theta)\\ &= \operatorname{supp}(\operatorname{gr}^{F'} M)\end{aligned}$$

which implies the characteristic varieties are isomorphic. □

Examples 3.5.
1. $\operatorname{Ch}(\mathcal{O}_X) = X$ (the zero section) for any smooth variety X.
2. $\operatorname{Ch}(D_X) = T^*X$.
3. Let $\iota \colon \mathbb{C}^n \to \mathbb{C}^{n+1}$ be the inclusion map. Let N be a coherent D_n-module, and let M be a coherent D_{n+1}-module. We can investigate how the characteristic varieties of these modules change under application of ι_*. Indeed, take M to be ι_*N. Then we know from example 2.20 that

$$M \cong \mathbb{C}[\partial_{n+1}] \otimes_{\mathbb{C}} \iota_\bullet N$$

where the $\iota_\bullet$ is the $\mathcal{O}$-module direct image. Take a good filtration F_i on N, and produce a new filtration F_i' on M by

$$F_i'M = \sum_{j=0}^{i}\sum_{k=0}^{j} \mathbb{C}\partial_{n+1}^k \otimes_{\mathbb{C}} \iota_\bullet F_{i-j}N.$$

Since F_i good, certainly F_i' is also good. The associated graded module with respect to this filtration is then

$$\operatorname{gr}^{F'}(M) = \mathbb{C}[\partial_{n+1}] \otimes_{\mathbb{C}} \operatorname{gr}^F(N)$$

which has the same annihilator as $\operatorname{gr}^F(N)$ (naturally embedded in $\operatorname{gr}(D_{n+1})$). Thus we can see

$$\operatorname{Ch}(M) = \{(x,\xi) \in T^*\mathbb{C}^{n+1} \,:\, x \in \mathbb{C}^n,\ (x,\pi(\xi)) \in \operatorname{Ch}(N)\}$$

where $\pi \colon T^*\mathbb{C}^{n+1} \to T^*\mathbb{C}^n$ is the projection. In particular, $\dim\operatorname{Ch}(M) = \dim\operatorname{Ch}(N) + 1$.

4. Let $M \in \mathrm{Mod}(D_X)$, $N \in \mathrm{Mod}(D_Y)$. Then we can calculate the characteristic variety of $M \boxtimes N$, namely
$$\mathrm{Ch}(M \boxtimes N) = \mathrm{Ch}(M) \times \mathrm{Ch}(N).$$

5. Let $\pi\colon X \times Y \to Y$ be the projection morphism. Then if $M \in \mathrm{Mod}(D_Y)$ then it is immediate from the definitions that
$$\pi^* M \cong \mathcal{O}_Z \boxtimes M,$$
and hence $\mathrm{Ch}(\pi^* M) = Z \times \mathrm{Ch}(M)$.

We have the following crucial theorem:

Theorem 3.6 (Bernstein's inequality)**.** $\dim \mathrm{Ch}(M) \geq \dim X$ for any D_X-module $M \neq 0$.

Proof. Induct on $\dim X$. $\dim X = 0$ is trivial, so suppose $\dim X > 0$. Without loss of generality we may assume $\operatorname{supp} M$ is contained in a smooth hypersurface $S \xhookrightarrow{\iota} X$. Indeed, if $\operatorname{supp} M = X$ then $\mathrm{Ch}(M)$ is just the zero-section on T^*X, and so has dimension $\dim X$. Pass to $\operatorname{supp} M \subseteq S$ by replacing X with a suitable open subset (to ensure S is non-singular).

Now, Kashiwara's theorem implies that $M = H^0 \iota_* N$ for some coherent D_S-module N. We want to use this to apply this induction hypothesis. But example 3.5 (3) tells us that
$$\dim \mathrm{Ch}(M) = \dim \mathrm{Ch}(N) + 1 \geq \dim(S) + 1 = \dim(X)$$
as required. □

Remark 3.7. There is an alternative proof of this fact due to Kashiwara, Kawai and Sato [19] by proving that for a coherent D_X-module M, the characteristic variety $\mathrm{Ch}(M)$ is *involutive* with respect to the symplectic structure of T^*X. That is, for any $p \in X$, the tangent space T_pX satisfies $T_pX \supseteq (T_pX)^\perp$ in $T_p(T^*X)$ (we say T^pX is an *involutive subspace*). Here $\bullet^\perp$ is defined in terms of the standard symplectic form on T^*X. Since in general
$$\dim(X) = \dim(T_pX) = 2n - \dim((T_pX)^\perp),$$
this immediately implies $\dim(X) \geq n$.

Definition 3.8. In the special case where $\dim \mathrm{Ch}(M) = \dim X$ we say M is *holonomic*. We denote the full subcategory of holonomic D_X-modules $\mathrm{Mod}_h(D_X)$, and the derived category of complexes of D_X-modules with holonomic cohomology $D_h(D_X)$.

Example (4) of 3.5 has the following immediate corollary:

Corollary 3.9. If M and N are holonomic then so is $M \boxtimes N$.

Remark 3.10. In many situations, the category of holonomic D-modules is the right category to think about. In a sense that will be made precise later, the holonomic D-modules correspond to systems of PDEs with *finite-dimensional* spaces of solutions.

Lemma 3.11. Let
$$0 \to N \to M \to L \to 0$$
be a short exact sequence of coherent D_X-modules. Then M is coherent if and only if N and L are both coherent. In particular the full subcategory of holonomic D_X-modules is preserved under subobjects and quotients.

Proof. Pick a good filtration F for M on an open set. Under the induced filtrations for N and L we have a short exact sequence (on an open set)
$$0 \to \mathrm{gr}^F N \to \mathrm{gr}^F M \to \mathrm{gr}^F L \to 0.$$
So locally, and hence globally, $\mathrm{Ch}(M) = \mathrm{Ch}(N) \cup \mathrm{Ch}(L)$. This implies the result immediately. □

We have the following criterion for holonomicity:

Lemma 3.12. Let $M^\bullet \in D^b(D_X)$. Then $M^\bullet$ is holonomic if and only if $M^\bullet$ is coherent, and all the fibres $i_x^*(M^\bullet)$ are finite dimensional. Here i_x is the inclusion of the point $x \in X$.

Proof. We'll prove the direction $\Longleftarrow$, which is the only way we will need. Let M be coherent with finite dimensional fibres. Without loss of generality $\operatorname{supp}(M^\bullet) = X$. We will induct on $\dim(X)$. Take $Y \overset{j}{\hookrightarrow} X$ open, such that $\dim(X \setminus Y) < \dim(X)$. Consider the complex $j^\dagger M^\bullet$. We show this complex is holonomic. Take an open dense subset such that the i^{th} cohomology is locally free over $\mathcal{O}$. The fibres are finite-dimensional since those for $M^\bullet$ are, so $H^i(j^\dagger M^\bullet)$ is coherent over $\mathcal{O}$, and hence holonomic for every i as required.

Now, produce the distinguished triangle

$$N^\bullet \to M^\bullet \to j_* j^\dagger M^\bullet \overset{+1}{\to},$$

so $N^\bullet$ is the cocone. $N^\bullet$ is supported on $X \setminus Y$, is coherent and has finite dimensional fibres, so by the induction hypothesis it is holonomic. as j_* preserves holonomicity [1], $j_* j^\dagger M^\bullet$ is holonomic, so $M^\bullet$ is holonomic as required. $\square$

Let f be a section of $\mathcal{O}_X$. There is an affine open subvariety of X associated to f by taking the complement of the zero locus:

$$U = X_f = \{x \in X \,:\, f(x) \neq 0\}.$$

If u is a section of D_U then the D_U-module generated by u is clearly holonomic. We can actually show something much stronger than this.

Lemma 3.13 (Bernstein's Lemma on b-functions)**.** Let σ be any section of $\mathcal{O}_X$, and u be any section of D_{X_σ}. Then there exist polynomials $p \in D_X(X)[x]$ and $b \neq 0 \in \mathbb{C}[x]$ such that

$$p(n)\sigma^{n+1}u = b(n)\sigma^n u$$

for all $n \in \mathbb{N}$.

This lemma will be crucial when we prove that direct and inverse image functor preserve holonomicity in section 3.23. To prove it however, we will need a little more machinery: namely the notion of *duality* for D-modules.

3.3 Duality for D-modules

Definition 3.14. Define the *duality functor* for D-modules to be

$$\begin{aligned} \mathbb{D}\colon D^b_{coh}(D_X) &\to D^b_{coh}(D_X)^{op} \\ M^\bullet &\mapsto (R\mathcal{H}om_{D_X}(M^\bullet, D_X) \otimes_{\mathcal{O}_X} \omega_X^\vee)[\dim X]. \end{aligned}$$

Lemma 3.15. $\mathbb{D}^2 \cong \mathrm{id}$.

The proof is technical, and omitted. See [3] page 24, or [8] proposition 2.6.5.

Remark 3.16. This definition needs some justification. In particular, why the degree shift by $\dim X$? Suppose M is a holonomic D_X-module. Then the complex $R\mathcal{H}om(M, D_X)$ has non-zero homology precisely in degree $\dim X$. In fact, for coherent modules the converse is also true. This will follow from the following theorem of Roos:

[1] This uses the key result 3.23 for open embeddings below. This is not cheating, as the proof of that specific case does not use this lemma.

Theorem 3.17 (Roos)**.** let $M \in \mathrm{Mod}(D_X)$ be coherent. Then

$$\operatorname{codim} \mathrm{Ch}(\mathcal{E}xt^i_{D_X}(F, D_X \otimes_{\mathcal{O}_X} \omega_X^\vee)) \geq i \quad \forall\ i.$$

In particular we have

$$\mathcal{E}xt^i_{D_X}(F, D_X \otimes_{\mathcal{O}_X} \omega_X^\vee) = 0$$

for all $i > \operatorname{codim}(\mathrm{Ch}(M))$.

We will not prove this here. See [3] for a proof sketch.

Corollary 3.18. Let $M \in \mathrm{Mod}(D_X)$ be coherent. Then M is holonomic if and only if $\mathbb{D}M \in \mathrm{Mod}(D_X)$, i.e. $H^i(\mathbb{D}M) = 0$ for all $i \neq 0$. In particular, $\mathbb{D}$ preserves the full subcategory of holonomic D_X-modules.

Proof. First suppose M is holonomic. Roos's theorem implies that $H^i(\mathbb{D}M) = 0$ for all $i < 0$. But $H^i(\mathbb{D}M) = 0$ for all $i > 0$ always. Conversely, if $\mathbb{D}M$ is a module N say, then $M = \mathbb{D}N$. So

$$M = \mathcal{E}xt^{\dim X}_{D_X}(N, D_X) \otimes_{\mathcal{O}_X} \omega_X^\vee \neq 0,$$

so $\operatorname{codim}(\mathrm{Ch}(M)) \geq \dim X$ by Roos's theorem and M is holonomic. □

Corollary 3.19. Suppose M is a D_X-module, and $U \subseteq X$ is an open subset such that some submodule $N \leq M|_U$ is holonomic as a D_U-module. Then we can find a holonomic D_X-submodule $M' \leq M$ such that $M'|_U = N$.

Proof. Without loss of generality M is coherent, and $N = M|_U$ ([**?**]). Consider the complex $\mathbb{D}M$. Define M' to be $\mathbb{D}H^0(\mathbb{D}M)$. By the previous corollary, to show M' is a holonomic D_X-module we need only check $H^0(\mathbb{D}M)$ is a holonomic D_X-module. But it is certainly non-zero, so by Roos's theorem $\operatorname{codim}(\mathrm{Ch}(H^0(\mathbb{D}M))) \geq \dim X$ and so it is holonomic also. Now we only need to show that $M'|_U = M|_U$. But

$$M'|_U = \mathbb{D}H^0(\mathbb{D}M|_U) = \mathbb{D}\mathbb{D}M|_U \cong M|_U.$$

□

We can use the techniques of duality to neatly prove the following proposition on holonomic D-modules:

Proposition 3.20. Holonomic D-modules have *finite length*, i.e. if M is a holonomic D_X-module, there exists some $n \in \mathbb{N}$ such that any chain of D_X-submodules

$$M_0 \subsetneqq M_1 \subsetneqq \cdots \subsetneqq M_m = M$$

has $m \leq n$. Equivalently, M is both Artinian and Noetherian.

Proof. First note that since $\mathbb{D}\colon \mathrm{Mod}_h(D_X) \to \mathrm{Mod}_h(D_X)^{op}$ gives a contravariant duality, it is only necessary to prove M is Noetherian, (i.e. $\mathbb{D}M$ Noetherian $\implies$ M Artinian). The problem is local since we can cover X by finitely many affine subsets with local co-ordinates. But certainly D_n is a (left) Noetherian ring, and since M is coherent it is locally finitely generated, and hence locally Noetherian. □

Proof of 3.13. Write Y for X_σ. The crucial idea of the proof is to *extend scalars* $\mathbb{C} \to \mathbb{C}(\lambda)$, the field of rational functions. Then we have extended varieties $\iota\colon \hat{Y} \to \hat{X}$. The reason we do this is that because global sections of the extended sheaf $D_{\hat{X}}$ are precisely quotients of polynomials $D_X(X)[\lambda]/\mathbb{C}[\lambda]$. Thus it suffices to produce some $P \in D_{\hat{X}}$ such that

$$P(\sigma^{\lambda+1}u) = \sigma^\lambda u$$

and then put $P = d/b$.

How can we produce such a P? The section u generates a holonomic D_Y-module M, and this extends to a holonomic $D_{\hat{Y}}$-module $\hat{M}$ via the construction

$$\hat{M} = \sigma^\lambda(\mathbb{C}[\lambda] \otimes_{\mathbb{C}} M)$$

with $\theta \in \Theta_{\hat{Y}}$ acting by

$$\theta(\sigma^\lambda m) = \frac{\lambda(\theta(\sigma))}{\sigma}\sigma^\lambda m + \sigma^\lambda \theta(m).$$

We observe that $\hat{M}$ is still holonomic (as the construction clearly does not change the dimension of the characteristic variety), Thus we can apply 3.19 to the direct image $\iota_*(\hat{M})$ to produce a submodule $N \leq \iota_*(\hat{M})$ which restricts to $\hat{M}$ on $\hat{Y}$.

Now, this has a natural interpretation, namely that the quotient $\iota_*(\hat{M})/N$ is zero on the affine piece $Y \subseteq X$. Hence the image of the section u of D_Y is zero in this quotient, i.e. if $\hat{u} = \sigma^\lambda u$ in $\hat{M}$, then by Kashiwara's theorem $\iota_*\hat{u}$ is supported on $\hat{Y}$, so is a section of N. We now use the fact that holonomic D-modules have finite length to produce a section $P \in D_{\hat{X}}$ such that

$$P(\sigma^{n+1}\hat{u}) = \sigma^n \hat{u}.$$

But this is exactly what we need! Simply change variables by $\lambda \mapsto \lambda + n$, and we see this condition on P is precisely the condition

$$P(\sigma^{\lambda+1}u) = \sigma^\lambda u$$

that we were looking for. □

Definition 3.21. For a morphism $f: X \to Y$, we introduce two new functors on D-modules:

$$f_! = \mathbb{D}f_*\mathbb{D} \quad \text{and} \quad f^! = \mathbb{D}f^\dagger\mathbb{D},$$

the exceptional direct and inverse image functors. These are produced by analogy with the functors crucial for Verdier duality, which we will introduce in section 4.5. By a routine check, we can see that $f_!$ is left adjoint to $f^\dagger$, and $f^!$ is left adjoint to f_*. Furthermore, for any f one can produce a canonical natural transformation of functors

$$f_! \to f_*$$

(see [8] 3.2.16 for details). We will use this to classify the simple holonomic D-modules in 5.18.

Remark 3.22. We have now defined six functors on D-modules, namely $f_*, f^\dagger, f_!, f^!, \mathbb{D}, \boxtimes$. This is an example of the *six operations formalism* of Grothendieck. I will deliberately avoid being too explicit about this means, but this terminology is widely used, e.g. in the monograph [14] of Mebkhout.

In section 4 we will introduce an analogous set of six operations on the derived category of complexes of $\mathbb{C}_X$-modules with constructible cohomology, and explain how the two families relate together via the *de Rham functor*.

3.4 Preservation of holonomicity

Theorem 3.23. Let $f: X \to Y$ be a morphism of smooth varieties. Then the functors f_* and f^* preserve the full derived subcategories of holonomic D-modules; i.e.

$$f_* D_h(D_X) \subseteq D_h(D_Y) \quad \text{and} \quad f^* D_h(D_Y) \subseteq D_h(D_X).$$

Remark 3.24.

1. Notice how different this is from the coherent case. We remarked before in 2.22 that only certain special classes of morphisms preserve coherence of D-modules when we take direct and inverse images. Thus this theorem is telling us – in a sense – that the full subcategory of holonomic D-modules is easier to work with than the whole category of coherent D-modules.

2. We've already proven that $\mathbb{D}$ and $\boxtimes$ preserve holonomicity in 3.18 and 3.9 respectively. Thus holonomicity is preserved by all of the six operations we mentioned in remark 3.22.

Proof. We proceed by reducing the problem to a few special cases which will follow from calculations we've already done. Indeed, first we reduce to the direct image case, i.e. show that the theorem for f_* implies the theorem for f^*. We can always decompose f as a product $\pi \circ \iota$ where $\pi\colon Z \times Y \to Y$ is a projection and $\iota\colon X \to W\,(= Z \times Y)$ is a closed embedding, so it suffices to check these two cases (since we're assuming X is quasi-projective throughout). The projection follows from calculations we've already done. Indeed, π^* is exact, so $\pi^* M \in \mathrm{Mod}(D_{Z\times Y})$ for $M \in \mathrm{Mod}_h(D_Y)$. Furthermore, by example 3.5 (5),

$$\begin{aligned}\mathrm{Ch}(\pi^* M) &= \mathrm{Ch}(\mathcal{O}_Z \boxtimes M)\\ &= \mathrm{Ch}(\mathcal{O}_Z) \times \mathrm{Ch}(M)\end{aligned}$$

which has dimension $\dim(Z) + \dim(Y) = \dim(X \times Y)$. Thus $\pi^* M$ is holonomic.

The closed embedding now follows from the direct image case by homological algebra. Take $j : U \to W$ to be the complementary open embedding to ι. There is in general a distinguished triangle ([8] 1.7)

$$\iota_* \iota^\dagger M \to M \to j_* j^\dagger M \xrightarrow{+1} .$$

If $j_* j^\dagger M$ is holonomic then so is $\iota_* \iota^\dagger M$, and hence $\iota^* M$ by the computation in example 3.5. But $j_* j^\dagger M \cong j_* M|_U$ which is holonomic by the direct image case.

Now, to deduce the statement for direct images we again split up f into a product of familiar morphisms. This time let $f = \pi \circ \iota$ with $\pi\colon Y \times Z \to Y$ a projection with Z *complete*, and $\iota\colon X \to W\,(= Z \times Y)$ a *locally* closed embedding. We check the theorem for π first. If $M \in \mathrm{Mod}_{Y\times Z}$, we check the holonomicity of $\pi_* M$ using the criterion 3.12 (thus this case follows from the open embedding case). Certainly $\pi_* M$ is coherent because π is proper (see the remark 2.22). To see it has finite dimensional fibres,

$$i_y^* \pi_* M = \pi_{y*} i_{Z_y}^* M$$

for i_{Z_y} the inclusion of the fibre above $y \in Y$, and π_y the projection from the fibre Z_y to the point y. This is a kind of base change (see [3] 2.3) which is not used anywhere else, so for brevity the proof is omitted. Thus $i_y^* \pi_* M$ is a direct image under a proper map of a coherent complex, hence coherent, i.e. finite dimensional.

Finally, we must check the case of a locally closed embedding. We've already checked closed embeddings in example 3.5, so we must check the case where ι is an open embedding to complete the proof. This will follow from the lemma on b-functions. Indeed, let M be a holonomic D_X-module. To reduce to the case where we can apply the lemma, we need ι to be of the form $Y = X_\sigma \hookrightarrow X$, but we can always do this by covering by finitely many affines, and replacing M by its Čech resolution.

Now, take M a holonomic D_Y-module generated by u. Then $\iota_* M$ is generated by $\sigma^n u$ for $n \in \mathbb{Z}$. By the lemma on b-functions, we actually only need a finite set of such $\sigma^n u$, so $\iota_* M$ is coherent. Furthermore, following through the proof of the lemma, we actually showed that the D-module with extended scalars $\hat{M}$ was holonomic, generated by $\sigma^\lambda u$. We want to pass from this statement for $D_{\hat{X}}$ to one for D_X. Since finitely many $\sigma^n u$ generate M it suffices to prove that almost all $\sigma^n u$ lie in a holonomic submodule. The ideas is to look at $\hat{M}$ and set $\lambda = n$.

More precisely, $\hat{M}$ holonomic means there exist $P_1, \ldots, P_k \in D_{\hat{X}}$ such that their images $Q_1, \ldots Q_k$ in the graded module cut out by a variety of dimension $\leq \dim X$. Set $\lambda = n$ to get operators $P_1', \ldots, P_k' \in D_X$, and for all but finitely many n their images in the graded module still cut out a variety of dimension $\leq \dim X$. Furthermore $P_n(\sigma^n u) = 0$ for all n, so $\sigma^n u$ lies in a holonomic submodule as required.

□

4 The de Rham functor

In this section we introduce the analytic setting for D-modules. Although they often behave very differently, it is useful for our purposes to consider D-modules on complex manifolds, rather than smooth varieties. Every smooth complex variety can be considered as a complex manifold, so no information is lost, however in some settings complex manifolds are easier to deal with. Most crucially, we will define the *de Rham* functor for a D_X-module, which will relate D-modules to topological objects called constructible and perverse sheaves. Precisely how they relate is one of the main topics of interest in the next few chapters.

4.1 Analytic D-modules

Let X be a complex manifold. X comes equipped with a sheaf $\mathcal{O}_X$ of holomorphic functions, and a sheaf Θ_X of holomorphic vector fields, so we can define D-modules on X in a similar way to the algebraic case. We will frequently want to consider the complex manifold topology on smooth algebraic varieties. If we want to make this explicit, we will use the notation X^{an} to denote a smooth variety considered with the induced topology from $\mathbb{C}^n$.

Definition 4.1. Define the sheaf D_X of *holomorphic differential operators* to be the sheaf generated by $\mathcal{O}_X$ and Θ_X in $\mathcal{E}nd(\mathcal{O}_X)$. A *$D$-module* is a (left) D_X-module.

Remark 4.2. Most of the key facts and results about algebraic D-modules also hold in the analytic case. In particular

- The properties of D_X described in section 2.1 still hold as before. We have an analogue of side-changing as in 2.12, and we can define the notion of an *integrable connection* as in 2.9. These form an abelian category.
- A D_X-modules is coherent over $\mathcal{O}_X$ if and only if it is an integrable connection, as in 2.5.
- We define inverse images, direct images and tensor products of D-modules under holomorphic maps in the analytic setting exactly as we did in the algebraic setting, by defining transfer modules.
- Kashiwara's theorem 2.26 still holds for closed embeddings of complex manifolds.
- We can define good filtrations at least locally, and this allows us to define the characteristic variety of a coherent D_X-module M to be the unique closed subvariety $\mathrm{Ch}(M) \subseteq T^*X$ such that for any open set $U \subseteq X$ equipped with a good filtration F_i, we have

  $$T^*U \cap \mathrm{Ch}(M) = \mathrm{supp}(\mathrm{gr}^F(M|_U)).$$

  (See e.g. [8] p100.) Given this, the characteristic variety behaves similarly to the algebraic case. Most importantly, Bernstein's inequality 3.6 still holds, so we can define a *holonomic* D_X-module to be one whose characteristic variety has dimension exactly $\dim X$ as before.
- We can define duality functors $\mathbb{D}_X$ for coherent D-modules as before, and holonomicity is preserved by dualising.

Example 4.3. Although in the analytic setting theorem 3.23 remains true for inverse images, it *fails* for direct images. Indeed, consider the open embedding $j\colon \mathbb{C}^\times \hookrightarrow \mathbb{C}$, considered as a morphism of complex manifolds. Let $\mathcal{O}_{\mathbb{C}^\times}$ be the sheaf of holomorphic functions, a holonomic $D_{\mathbb{C}^\times}$-module. If we apply j_* to this module, we produce the $D_{\mathbb{C}}$-module of functions holomorphic away from 0. This is not coherent, because the behaviour at 0 can be arbitrarily bad: i.e. we may have an essential singularity at 0.

In the algebraic setting, things were far more rigid, which is why this problem didn't arise. The direct image in the algebraic category was simply $\mathcal{O}_{\mathbb{C}}[x^{-1}]$.

Remark 4.4. The analytic techniques we are about to introduce will help us even in the algebraic setting. Given X a smooth quasi-projective variety, over $\mathbb{C}$, we can associate to X a complex manifold X^{an}. This is a consequence of the fact that quasi-projective complex varieties can be naturally topologised by the subspace topology in $\mathbb{P}^n_{\mathbb{C}}$, making them – in the smooth case – into complex manifolds.

4.2 DR and Sol

We remain in the analytic category, i.e. $M^\bullet \in D^b(D_X)$ for X a complex manifold. We define two crucial functors: the *de Rham functor*

$$\begin{aligned} DR_X \colon D^b(D_X) &\to D^b(\mathbb{C}_X) \\ M^\bullet &\mapsto \omega_X \otimes^L_{D_X} M^\bullet \end{aligned}$$

and the *solution functor*

$$\begin{aligned} Sol_X \colon D^b(D_X) &\to D^b(\mathbb{C}_X)^{op} \\ M^\bullet &\mapsto R\mathcal{H}om_{D_X}(M^\bullet, \mathcal{O}_X). \end{aligned}$$

Lemma 4.5. The two functors relate to one another in the following way:

$$DR_X(M^\bullet) \cong Sol_X(\mathbb{D}_X M^\bullet)[\dim(X)],$$

so we only need consider one functor to understand the properties of both.

Proof. First notice that

$$\begin{aligned} Sol_X(\mathbb{D}_X M^\bullet) &\cong R\mathcal{H}om_{D_X}(\mathcal{O}_X, M^\bullet) \\ &\cong R\mathcal{H}om_{D_X}(\mathcal{O}_X, D_X) \otimes^L_{D_X} M^\bullet. \end{aligned}$$

So we simply resolve $\mathcal{O}_X$ as follows:

$$\begin{aligned} R\mathcal{H}om_{D_X}(\mathcal{O}_X, D_X) &\cong \mathcal{H}om_{D_X}(D_X \otimes_{\mathcal{O}_X} \wedge^0 \Theta_X, D_X) \to \cdots \mathcal{H}om_{D_X}(D_X \otimes_{\mathcal{O}_X} \wedge^{\dim X} \Theta_X, D_X) \\ &\cong \mathcal{H}om_{\mathcal{O}_X}(\wedge^0 \Theta_X, D_X) \to \cdots \mathcal{H}om_{\mathcal{O}_X}(\wedge^{\dim X} \Theta_X, D_X) \\ &\cong \Omega^0_X \otimes_{\mathcal{O}_X} D_X \to \cdots \to \Omega^{\dim X}_X \otimes_{\mathcal{O}_X} D_X \\ &\cong \omega_X[-\dim(X)] \end{aligned}$$

in the derived category. Finally, applying $\otimes^L M^\bullet$ gives the result. ☐

Example 4.6. Let us demonstrate the behaviour of these functors in the classical setting of linear PDEs. Take $M = D_n / \sum_{i=1}^N D_n P_i$, where P_i are differential operators in D_n. Then the solution functor is just $\mathrm{Hom}_{D_n}(M, \mathcal{O}_n)$. Hence it fits into an exact sequence

$$0 \to Sol(M) \to \mathcal{O}_n \stackrel{\sum P_i}{\to} \mathcal{O}_n^N$$

justifying its role as the space of solutions of the system of linear differential operators.

Recall that a *local system* on a complex (or smooth) manifold X (or even a more general topological space) is a locally constant sheaf of $\mathbb{C}_X$-modules with finite-dimensional stalks. Assuming X is connected these stalks all have the same dimension, called the *rank* of the local system.

Examples 4.7.

1. The constant sheaf $\mathbb{C}^n_X$ is a rank n local system.
2. If X is a smooth manifold, the *orientation sheaf* is defined to be $\mathcal{L}_{or} = H^n(X, \mathbb{C}_X)^\vee$. This is a rank one local system which is constant if and only if X is orientable.

3. Let $X = \mathbb{C}^{\times}$. There is a bijective correspondence between rank one local systems on X and numbers $0 \neq \lambda \in \mathbb{C}$. Given a rank one local system $\mathcal{L}$, let γ be a loop in X around 0. We can cover the loop by open sets on which $\mathcal{L}$ is trivial, so by composing the isomorphisms on overlaps we produce an isomorphism $\mathbb{C} \xrightarrow{\times\lambda} \mathbb{C}$.

This is an example of the more general phenomenon of *monodromy.* We can interpret the λ above as one-dimensional representations of the fundamental group of $\mathbb{C}^{\times}$. In general, for X a well-behaved topological space (e.g. any smooth manifold), there is an equivalence of categories

$$\left\{\begin{matrix}\text{rank } m \text{ local systems}\\ \text{on X}\end{matrix}\right\} \leftrightarrow \left\{\begin{matrix}m \text{ dimensional representations}\\ \text{of } \pi_1(X)\end{matrix}\right\}.$$

In section 4.3 we will introduce a generalisation of local systems: constructible sheaves. For basic facts on local systems see e.g. [16] or [6].

Theorem 4.8 (de Rham)**.** Let X be a complex manifold of dimension n. The de Rham functor gives an equivalence of categories

$$H^{-n} \circ DR_X \colon \operatorname{Conn}(X) \to \operatorname{Loc}(X).$$

Proof. Let M be a rank m integrable connection on X. Firstly, we must check that $H^{-n}(DR_X(M))$ is a rank m local system on X. In the derived category, we have a nice representative for $DR_X(M)[-n]$, namely

$$\Omega^0_X \otimes_{\mathcal{O}_X} M \to \cdots \to \Omega^n_X \otimes_{\mathcal{O}_X} M$$

with differentials induced from the differentials in the de Rham complex. This is because we can resolve the right D_X-module ω_X ([8] 1.5.27) by

$$0 \to \Omega^0_X \otimes_{\mathcal{O}_X} D_X \to \cdots \to \Omega^n_X \otimes_{\mathcal{O}_X} D_X \to \omega_X \to 0$$

where

$$d(\eta \otimes P) = d\eta \otimes P + \sum_{i=1}^{n} dx_i \wedge \eta \otimes \partial_i P$$

for (x_i, ∂_i) local coordinates. When we look at the zeroth cohomology (since we've shifted by $-n$), we find that it is just

$$\ker(\nabla \colon M \to \Omega^1_X \otimes_{\mathcal{O}_X} M)$$
$$\nabla(m) = \sum_i (x_i \otimes \partial_i m).$$

This is nothing but the ∇ arising when M is thought of as a connection on an $\mathcal{O}$-module. Thus

$$H^0(DR_X(M)) = \{m \in M \ : \nabla m = 0\}$$

is the sheaf of *horizontal sections* of M. To see it is locally free over $\mathbb{C}_X$, we use Frobenius's theorem. See [21] 9.11 and 9.12 for full details.

We can explicitly produce a quasi-inverse for this functor as follows: let $\mathcal{L}$ be a rank m local system. Let $M = \mathcal{O}_X \otimes_{\mathbb{C}_X} \mathcal{L}$, with integrable connection defined by the composite

$$\nabla \colon \ \Omega^0_X \otimes_{\mathcal{O}_X} M \xrightarrow{\ \sim\ } \mathcal{O}_X \otimes_{\mathbb{C}_X} \mathcal{L} \xrightarrow{d\otimes 1_{\mathcal{L}}} \Omega^1_X \otimes_{\mathbb{C}_X} \mathcal{L} \xrightarrow{\ \sim\ } \Omega^1_X \otimes_{\mathcal{O}_X} M\ .$$

To check these functors give an equivalence, we first observe that if M is an integrable connection then $\mathcal{O}_X \otimes_{\mathbb{C}_X} M^{\nabla} \cong M$. Similarly, if $\mathcal{L}$ is a local system, then $(\mathcal{O}_X \otimes_{\mathbb{C}_X} \mathcal{L})^{\nabla} \cong \mathcal{L}$ for the given connection. □

Example 4.9. Let $\mathcal{L}$ be the trivial local system on $\mathbb{C}$. The proof of 4.8 gives us a recipe for computing the integrable connection associated to this sheaf. Indeed, we can see what the result must be in this case: $\mathcal{O}_{\mathbb{C}}$ with the standard connection $\nabla \colon f \mapsto df$. This has horizontal sections precisely the constant functions.

4.3 Constructible sheaves

Essentially, a constructible sheaf is one that can be built up from local systems, in a way we will make precise.

Definition 4.10. Let X be a complex manifold. A *stratification* of X is a decomposition of X as a disjoint union

$$X = \coprod_i X_i$$

of non-empty connected sets (the *strata*) such that for each i, the boundary ∂X_i is a union of strata, and both $\overline{X_i}$ and ∂X_i are complex manifolds.

Definition 4.11. A $\mathbb{C}_X$-module M is called *constructible* if there exists some stratification $\coprod X_i$ of X such that for each stratum X_i, the restriction $M|_{X_i}$ is a local system. In the derived category $D(\mathbb{C}_X)$, we say a complex $M^\bullet$ is *constructible* if all of its cohomology sheaves are constructible. We denote the full subcategory of such complexes $D_c(\mathbb{C}_X)$.

If X is instead a smooth variety, then a $\mathbb{C}_X$-module M is called *constructible* if there exists a stratification of X by algebraic varieties such that M is a constructible $\mathbb{C}_{X^{an}}$-module with respect to this stratification. We have an analogous notion in the derived category to the analytic case.

Remark 4.12. One can define the notion of a stratification in more generality than described above. However, since we are working in a special case, we can actually impose extra conditions on our stratifications. For X a complex manifold (or more generally any 'analytic space'), a *Whitney stratification* is a stratification by submanifolds such that, for any two strata Y and Y', and for any sequences (y_i) in Y and (y'_i) in Y'

- If $y_i \to y' \in Y'$ and the tangent spaces T_i to Y at y_i converge to T, then $T \subseteq T_{y'}Y'$. (Whitney's condition A)
- If (y_i) is as above, and $y'_i \to y'$ also, and the secant lines L_i from y_i to y'_i converge to a line L, then $L \subseteq T$. (Whitney's condition B)

In fact, such stratifications exist for any complex manifold X. Not only that, but given a stratification of X, there always exists a refinement that satisfies the Whitney conditions. For more details, see Appendix 1. of [13].

4.4 Kashiwara's constructibility theorem

In this section we explain the statement in remark 3.10 about holonomic D-modules being those with finite-dimensional solutions. What we will in fact show is that applying the solution functor to a holonomic D-module yields a *constructible* complex, i.e. $Sol_X(M)$ is constructible for holonomic M.

Theorem 4.13. Let $M \in \mathrm{Mod}_h(X)$ for a complex manifold X. Then both $Sol_X(M)$ and $DR_X(M)$ have constructible cohomology, i.e. are objects of $D^b_c(X)$.

Remark 4.14. First, note that by lemma 4.5, we need only show this for DR_X to prove the result in general.

For the proof we'll need the following technical lemma, which I won't prove. See for example [8] 4.4.6. Unfortunately, the proof of the theorem will require the lemma in its full, rather complicated, generality.

Lemma 4.15. Let X be a complex manifold, $x \in X$. Let M be a coherent D_X-module. Let U_t, $t \in (0, 1]$ be an increasing family of open, relatively compact subsets of X, whose boundaries are real smooth hypersurfaces in X. If we also require

- Every U_t is *holomorphically separable and convex* [2]
- $\bigcup_{t<s} U_t = U_s$
- $\bigcap_{t>s} (U_t \setminus U_s) = \partial U_s$, and
- $\mathrm{Ch}(M) \cap T^*_{\partial U_t}(X_{\mathbb{R}}) \subseteq$ the zero section in T^*X

then we conclude that there is are isomorphisms

$$R\Gamma(\bigcup_t U_t, R\mathcal{H}om_{D_X}(M, \mathcal{O}_X)) \xrightarrow{\sim} R\Gamma(U_s, R\mathcal{H}om_{D_X}(M, \mathcal{O}_X)) \quad \forall\ s \in (0,1].$$

In particular, for any $0 < \varepsilon_1 < \varepsilon_2$ small enough, there is an isomorphism

$$R\Gamma(B_{\varepsilon_2}(x), R\mathcal{H}om_{D_X}(M, \mathcal{O}_X)) \xrightarrow{\sim} R\Gamma(B_{\varepsilon_1}(x), R\mathcal{H}om_{D_X}(M, \mathcal{O}_X))$$

where $B_\varepsilon(x)$ denotes the ε-ball about x.

Sketch proof of theorem 4.13. We do this in three steps. First, we construct an appropriate Whitney stratification of X. Secondly, we prove that $H^i(Sol_X(M))$ is locally constant on strata for every i, and finally we prove that these cohomology sheaves have finite-dimensional stalks everywhere.

Step 1: We produce a Whitney stratification of X such that $\mathrm{Ch}(M) \subseteq \coprod_j T^*X_j$ in T^*X. For this, we must use the fact that $\mathrm{Ch}(M)$ is a Lagrangian submanifold of T^*X stable under the action of $\mathbb{C}^\times$ on the fibres (a *conic* Lagrangian submanifold). This diversion into symplectic geometry is outside the scope of this essay, so the interested reader is referred to [8] E.3.6 and E.3.9.

Step 2: This proceeds by induction on i. Write $Sol_X(M) = N^\bullet$, so we are looking at the sheaves $H^i(N^\bullet)$. First consider the case $i = 0$. For a fixed stratum $X_\alpha \subseteq X$, since we are working locally we may assume that

$$X_\alpha = \{x_1 = \cdots = x_k = 0\} \subseteq \mathbb{C}^n = X.$$

Pick some $x \in X_\alpha$. To prove the sheaf is locally constant at x, we must prove that $\exists\ \varepsilon > 0$ such that the restriction map onto the stalk at $x' \in \overline{B_\varepsilon}(x)$ is an isomorphism for all such x'. Now, for this we will use the lemma. Indeed, we take a collection of open sets U_t as in the lemma indexed by $t \in (0,1]$ such that $U_1 = U$ and $\bigcap U_t = \{x\}$. We can satisfy the conditions of the lemma because of the way we chose our stratification . Hence we have an open $U \supseteq \overline{B_\varepsilon}(x)$ such that for fixed x' there is a quasi-isomorphism

$$\begin{aligned}
& R\Gamma(U, N^\bullet) \to N^\bullet x' \\
\implies & H^0(\Gamma(U, N^\bullet)) \cong H^0(N^\bullet)_{x'} \\
\implies & \Gamma(U, H^0(N^\bullet)) \cong H^0(N^\bullet)_{x'} \\
\implies & \varinjlim_{U \supseteq \overline{B_\varepsilon}(x)} \Gamma(U, H^0(N^\bullet)) \cong H^0(N^\bullet)_{x'} \\
\implies & \Gamma(\overline{B_\varepsilon}(x), H^0(N^\bullet)) \cong H^0(N^\bullet)_{x'}
\end{aligned}$$

as required.

For the induction step, we use homological algebra. Indeed, for fixed i we consider the following diagram whose rows are distinguished triangles. Take an ε such that there is a quasi-isomorphism

[2] i.e. for any two points in U_t, there exists a holomorphic function separating them, and U_t is equal to its holomorphic convex hull.

$\phi\colon R\Gamma(\overline{B_\varepsilon}(x), N^\bullet) \to N^\bullet_{x'}$ for all $x' \in B_\varepsilon(x)$. Extend this to the diagram

$$\begin{tikzcd} R\Gamma(\overline{B_\varepsilon}(x), H^i(N^\bullet)) \arrow[r] \arrow[d] & R\Gamma(\overline{B_\varepsilon}(x), N^\bullet) \arrow[r] \arrow[d, "\phi"'] & R\Gamma(\overline{B_\varepsilon}(x), \tau^{>i}(N^\bullet)) \arrow[r, "+1"] \arrow[d] & . \\ H^i(N^\bullet)_{x'} \arrow[r] & N^\bullet_{x'} \arrow[r] & \tau^{>i}N^\bullet_{x'} \arrow[r, "+1"] & \end{tikzcd}$$

Now, if $H^i(N^\bullet)|_{X_\alpha}$ is locally constant then for small enough ε the leftmost arrow is certainly a quasi-isomorphism. Hence the rightmost arrow is also a quasi-isomorphism, and so by an argument similar to the case $i = 0$ above, $H^{i+1}(N^\bullet)|_{X_\alpha}$ is locally constant, completing the step.

Step 3: The proof of this step is essentially another application of the lemma (although this time we can get away with the weaker statement about open balls), but requires some rather technical functional analysis so is omitted. All the details are included in [8] 4.6.2, or in the original paper [10] of Kashiwara (where the slightly different Theorems 1.2 and 1.6 plays the role of our technical lemma above). □

Notice that this is a much stronger result than the condition eluded to in 3.10.

Remark 4.16. A concept whose importance will become clear later is that of a *perverse sheaf*. Recall we mentioned in 2.18 the standard t-structure on a derived category such as $D(\mathrm{Mod}(\mathbb{C}_X))$. There is an alternative t-structure available for $D^b_c(X)$ called the *perverse t-structure*. Indeed, define full subcategories ${}^pD^{\leq 0}_c(X)$ and ${}^pD^{\geq 0}_c(X)$ by

$$F^\bullet \in {}^pD^{\leq 0}_c \text{ if } H^j(\iota^{-1}_{X_i}F^\bullet) = 0 \ \ \forall\ j > \dim(X) \text{ and for all strata}$$
$$F^\bullet \in {}^pD^{\geq 0}_c \text{ if } H^j(\iota^{!}_{X_i}F^\bullet) = 0 \ \ \forall\ j < \dim(X) \text{ and for all strata}$$

where $\coprod X_i$ is a stratification of X, and ι_{X_i} denotes the inclusion of the stratum. Here $\iota^!$ is the derived exceptional inverse image functor, which we will discuss in the section on Verdier duality below. Denote the associated truncation functors by ${}^p\tau^{\leq 0}$ and ${}^p\tau^{\geq 0}$. We will note in the next section that the perverse sheaves comprise the subcategory of $D^b_c(X)$ preserved by Verdier duality.

If $F^\bullet$ is in the heart of this t-structure, we say $F^\bullet$ is a *perverse sheaf*. The category $\mathrm{Perv}(X)$ of such complexes is then an abelian sub-category of the triangulated category $D(\mathrm{Mod}(\mathbb{C}_X))$. Under the Riemann-Hilbert correspondence (5.22), the category of perverse sheaves will correspond to the full subcategory

$$\mathrm{Mod}_{rh}(D_X) \subseteq D^b_{rh}(D_X)$$

of regular holonomic D_X-modules.

For example, if $X = \{pt\}$, the perverse sheaves on X are simply complex vector spaces. Indeed $D^b_c(X) = D^b(X) = K^b(X)$, the category of complexes of vector spaces up to homotopy. Such a complex is in the heart of the perverse t-structure if and only if it has cohomology only in degree 0, i.e. if it is homotopic to a complex supported only in degree 0.

4.5 Verdier duality

The category $D^b(\mathbb{C}_X)$ admits a number of functors analogous to those defined in chapter 1. In this section, we will describe them, and how they relate to one another. When we study the Riemann-Hilbert correspondence in chapter 5, , we will see that these structures are not merely analogous, but actually the same, under an equivalence of categories given by the de Rham functor.

Throughout this section, we will consider certain continuous maps $f\colon X \to Y$ of topological spaces. The theory of Verdier duality works in great generality, but for our purposes we need only know that the results hold for holomorphic maps between complex manifolds, or morphisms between algebraic varieties over $\mathbb{C}$. We will work with sheaves of $\mathbb{C}_X$-modules, but if we were working in full generality we could replace $\mathbb{C}$ with a more general class of rings. For a complete exposition of the theory of Verdier duality, see [6], or [11].

Definition 4.17. Let $f\colon X \to Y$ be a map as above. Then we can define functors on sheaves of $\mathbb{C}$-vector spaces

$$f_*\colon \operatorname{Mod}(\mathbb{C}_X) \to \operatorname{Mod}(\mathbb{C}_Y)$$
$$f^{-1}\colon \operatorname{Mod}(\mathbb{C}_Y) \to \operatorname{Mod}(\mathbb{C}_X)$$

and their derived functors

$$Rf_*\colon D(\mathbb{C}_X) \to D(\mathbb{C}_Y)$$
$$f^{-1}\colon D(\mathbb{C}_Y) \to D(\mathbb{C}_X)$$

Where, of course, $D(\mathbb{C}_X)$ means $D(\operatorname{Mod}(\mathbb{C}_X))$. The functor f^{-1} is exact. These functors send bounded complexes to bounded complexes in the derived category.

Definition 4.18. For spaces X and Y, we can define an *exterior tensor product*

$$\bullet \boxtimes \bullet\colon D(\mathbb{C}_X) \times D(\mathbb{C}_Y) \to D(\mathbb{C}_{X\times Y})$$

analogously to the exterior tensor product of D-modules defined in 2.4 in the obvious way. (As in, e.g. [11] p97).

Definition 4.19. For $f\colon X \to Y$ as above, we define a functor on sheaves of $\mathbb{C}$-vector spaces

$$f_!\colon \operatorname{Mod}(\mathbb{C}_X) \to \operatorname{Mod}(\mathbb{C}_Y),$$

the *direct image with compact supports*, by

$$(f_!F)(U) = \{s \in \Gamma(f^{-1}(U), F)\ :\ f|_{\operatorname{supp}(s)} \text{ is proper }\}.$$

This functor is left exact, so we will also consider its derived functor

$$Rf_!\colon D(\mathbb{C}_X) \to D(\mathbb{C}_Y).$$

The slightly subtle issue is how to define duality. There is an extremely powerful result of Verdier, generalising the classical duality theorems of Poincaré which will show us the way.

Theorem 4.20 (Verdier duality)**.** Let $f\colon X \to Y$. There exists an additive functor

$$f^!\colon D^+(\mathbb{C}_Y) \to D^+(\mathbb{C}_X)$$

of triangulated categories: the *exceptional inverse image*, such that there is a natural isomorphism

$$R\mathcal{H}om(Rf_!F^\bullet, G^\bullet) \cong Rf_*R\mathcal{H}om(F^\bullet f^!G^\bullet)$$

for any complexes $F^\bullet \in D^b(\mathbb{C}_X)$ and $G^\bullet \in D^b(\mathbb{C}_Y)$.

See [6] 2.3.21 or [11] 3.1.5.

Remark 4.21. By applying the 0th hypercohomology $\mathbb{H}^0$ to the above isomorphism, we produce exactly the statement that $f_!$ is right adjoint to $Rf_!$.

Definition 4.22. Let a_X be the unique morphism from X to the one point space $\{pt\}$. We define the *dualising complex* for X to be

$$\Delta_X = a_X^!\mathbb{C}_{\{pt\}} \in D^b(\mathbb{C}_X).$$

Example 4.23. Let X be a complex manifold. Then the dualising complex is simply $\mathbb{C}_X[2\dim(X)]$. Indeed, for general real topological manifolds, the dualising complex has cohomology $H^i(\Delta_X) \cong \mathcal{L}_{or}$ if $i = -\dim X$, and 0 otherwise. Since complex manifolds are orientable the result is a shifted constant sheaf. ([9] VI.2.5 and 3.2).

Definition 4.24. With all this in mind, define the *Verdier dual* of a complex $F^\bullet \in D^b(\mathbb{C}_X)$ to be

$$\boldsymbol{D}_X(F^\bullet) = R\mathcal{H}om(F^\bullet, \Delta_X).$$

One consequence of this definition of the dual of a complex is that, by 4.20, we can relate the standard and exceptional push-forwards and pull-backs by isomorphisms

$$(\boldsymbol{D}_X \circ Rf_*)(F^\bullet) \cong (Rf_! \circ \boldsymbol{D}_X)(F^\bullet) \text{ and } (\boldsymbol{D}_X \circ f^{-1})(G^\bullet) \cong (f^! \circ \boldsymbol{D}_X)(G^\bullet)$$

for any $F^\bullet \in D^b(\mathbb{C}_X)$ and $G^\bullet \in D^b(\mathbb{C}_Y)$. For $f_!, f^!$ for D-modules, this was given as the *definition* of those functors. These relationships and their correspondences under the Riemann-Hilbert correspondence gives a powerful method for understanding the behaviour of D-modules.

Example 4.25. Many classical forms of duality arise as special cases of this, for example, we recover classical Poincaré duality for a topological n-manifold $X$by applying the methods of Verdier duality to the sheaf $\mathcal{A}_X[k]$ of smooth functions on X with a shift in homological degree. See [6] theorem 3.3.1.

Remark 4.26. One of the reasons perverse sheaves were initially introduced is that they comprise the largest subcategory of $D^b_c(\mathbb{C}_X)$ which is preserved by Verdier duality. Essentially, the perverse t-structure is reversed by dualising, so only its heart is preserved. This is an analogue for constructible sheaves of our result 3.18 that duality for D-modules preserves the full subcategory of holonomic complexes of D-modules.

Remark 4.27. Thus, we have produced a collection of functors between derived categories

$$Rf_*,\ Lf^*,\ Rf_!,\ f^!,\ \boxtimes,\ \boldsymbol{D}_X$$

for a morphism $f\colon X \to Y$ of complex manifolds or algebraic varieties. This is very similar to the situation in 3.22 for holonomic D_X-modules. In 5.6, we will see that these so-called "six operations" correspond exactly under an equivalence of categories given by the de Rham functor.

5 Regular Singularities and the Riemann-Hilbert Correspondence

In this section we begin by reviewing the classical theory of regular singularities of systems of differential equations, following the article [7] of Haefliger. The work we've already done allows us to extend this concept to D-modules on general complex manifolds and smooth algebraic varieties, and hence to introduce the category of *regular holonomic D-modules*. Along the way we will study the case of D-modules on the complex line with singularities, as understood by analysts such as Fuchs in the 19$^{\text{th}}$ century. This gives us the motivation to introduce definitions for more general behaviour of D-modules with singularities. The highlight of the chapter is the Riemann-Hilbert correspondence, as requiring holonomicity and regular behaviour on singularities are precisely the conditions we need for an equivalence of categories with the derived category of constructible sheaves on a variety.

5.1 Fuchs Theory

This subsection will essentially consist of a single worked example: linear differential equations on $\mathbb{C}$ whose coefficients have at worst a pole at 0, and are holomorphic elsewhere.

Definition 5.1. Throughout this section, let K denote the field of meromorphic functions holomorphic away from zero, i.e. K is the field of fractions of the stalk $\mathcal{O} = (\mathcal{O}_\mathbb{C})_0$. More concretely, $K \cong$

$\mathbb{C}\{\{x\}\}[x^{-1}]$ where $\mathbb{C}\{\{x\}\}$ denotes the ring of *convergent* power series at $x = 0$. We will use the notation $\widetilde{\mathcal{O}}$ for the ring of stalks of *multivalued* holomorphic functions on $\mathbb{C}$ at 0, i.e.

$$\widetilde{\mathcal{O}} = \varinjlim_{\varepsilon \to 0} (\mathcal{O}(\widetilde{D^\bullet_\varepsilon})$$

where $\mathcal{O}(\widetilde{D^\bullet_\varepsilon})$ is the ring of holomorphic functions on the universal cover of the punctured disc of radius ε. These are, essentially, germs at 0 of holomorphic functions on the universal cover.

Consider a differential equation

$$P = \sum_{i=o}^{n} a_i(x) \left(\frac{d}{dx}\right)^{n-i}$$

for $a_i \in K$, $a_0 \neq 0$. This is equivalent to a system of linear first order equations, which we can express in the form

$$\frac{d}{dx} u_i(x) = \sum_{j=0}^{n} a_{ij}(x) u_j(x),$$

for $i = 1, \dots, n$, or equivalently

$$\frac{d}{dx} u(x) = A(x) u(x) \tag{1}$$

for $A \in \mathrm{Mat}_n(K)$. Two such systems $\frac{d}{dx}u(x) = A_i(x)u(x)$, $i = 1, 2$ are called *equivalent* if there exists a matrix $T \in GL_n(K)$ such that

$$A_1 = T A_2 T^{-1} - T \frac{d}{dx} T^{-1}$$

which is simply a change of variable condition: the result of setting $v(x) = Tu(x)$.

We can view this in several different ways. Firstly, we can characterise such systems by the *monodromy* of their solutions about the origin. Indeed, we consider multivalued solutions of the system (1): solutions on the universal cover $\widetilde{\mathbb{C}^\times}$ of $\mathbb{C}^\times$, where we pass to the universal cover by the change of variables $x = e^{2\pi i t}$. The space of such solutions is an n-dimensional $\mathbb{C}$-vector space S of functions of t, with basis $(v_1(t), \dots, v_n(t))$ say. But clearly $(v_1(t+1), \dots, v_n(t+1))$ are still linearly independent, so $t \mapsto t + 1$ gives an automorphism $S \to S$, i.e. a matrix $C \in GL_n(\mathbb{C})$. This matrix is called the *monodromy matrix* of the system (1). Two systems are equivalent if and only if they have the same monodromy.

Alternatively, let us introduce the following notion, which will be the right concept to generalise to other complex manifolds and smooth varieties:

Definition 5.2. A *meromorphic connection* on $\mathbb{C}$ at 0 is a finite dimensional vector space M over K equipped with a $\mathbb{C}$-linear map $\nabla\colon M \to M$ such that

$$\nabla(fm) = \frac{df}{dx} m + f\nabla(m) \quad \forall\, f \in K,\ \forall\, m \in M.$$

Two such meromorphic connections (M_1, ∇_1) and (M_2, ∇_2) are *isomorphic* if there exists an isomorphism $\phi\colon M_1 \to M_2$ such that $\phi \circ \nabla_1 = \nabla_2 \circ \phi$.

Let us describe a correspondence between systems of linear differential equations of form (1), and meromorphic connections. Let (M, ∇) be a meromorphic connection. Pick a vector space basis $(e_1, \dots, e_n)$ for M, and define a matrix $A = (a_{ij}) \in \mathrm{Mat}_n(K)$ by

$$\nabla e_j = \sum_{i=1}^{n} a_{ij} e_i.$$

Thus, for $u \in M$,

$$\begin{aligned}
\nabla u = 0 &\iff \nabla\left(\sum_i u_i e_i\right) = 0 \\
&\iff \sum_i \left(\frac{du_i}{dx} - \sum_j a_{ij} u_j\right) e_i = 0 \\
&\iff \frac{du_i}{dx} - \sum_j a_{ij} u_j = 0 \quad \forall\, i \\
&\iff \frac{d}{dx} u(x) = A(x) u(x).
\end{aligned}$$

Conversely, given such a system of equations, characterised by $A = (a)_{ij}$, we produce a meromorphic connection by putting

$$\nabla e_j = \sum_{i=1}^{n} a_{ij} e_i$$

where $e_1, \ldots, e_n$ is a basis for an n-dimensional K-vector space. It is clear that under this correspondence, equivalence classes of systems of differential equations correspond bijectively with isomorphism classes of meromorphic connections.
We now introduce the notion of regularity, via *moderate growth* of solutions.

Definition 5.3. Consider a *multivalued solution* $u \in \widetilde{\mathcal{O}}$ of the system $\frac{d}{dx}u(x) = A(x)u(x)$. The solution u is said to have *moderate growth* if for any sector of the form

$$S = \{z = (r, \theta) \,:\, 0 < r < \varepsilon,\, \theta_0 < \theta < \theta_1\},$$

there exists a constant $c > 0$ and an integer $j \in \mathbb{N}$ such that

$$|u(z)| < \frac{c}{|z|^j}$$

for all $z \in S$. If the solutions of the system have moderate growth, we say the system is *regular*, or has regular singularity at 0.

The following is a classical result on regularity.

Theorem 5.4. The following are equivalent:

1. The system

$$\frac{d}{dx} u(x) = A(x) u(x) \tag{2}$$

   is regular.

2. The system (2) is equivalent to one of the form

$$\frac{d}{dx} u(x) = \frac{B(x)}{x} u(x)$$

   where $B(z)$ is a matrix with holomorphic coefficients.

3. The system (2) is equivalent to one of the form

$$\frac{d}{dx} u(x) = \frac{C}{x} u(x)$$

   where C is a matrix with constant coefficients.

For a proof see [7] 1.3.1 or [8] 5.1.4.

Using this, we produce the classical result of Fuchs.

Theorem 5.5. The system $\frac{d}{dx}u(x) = A(x)u(x)$ is equivalent to a single equation $\sum_{i=0}^{n} a_i(x) \left(\frac{d}{dx}\right)^{n-i} u = Pu = 0$. It is regular if and only if $\frac{a_i}{a_0}$ has a pole at 0 of order at most i, for $i = 1, \ldots, n$.

Proof. One direction is now immediate. If the coefficients a_i satisfy the given conditions (the "Fuchs conditions") then the system is regular by 5.4. Conversely, suppose the system is regular. We'd like to proceed by induction on n, the order of P. To do this, first notice that by multiplying through by $\frac{x^n}{a_0(x)}$, the given equation is equivalent to one of the form

$$\theta^n u + \sum_{i=1}^{n} b_i(x)\theta^{n-i}u = Qu = 0$$

where $\theta = x\frac{d}{dx}$, and the Fuchs conditions are equivalent to requiring that all b_i are *holomorphic*. We need to reduce this to an equation of lower degree. We can always find some solution of the form $u(x) = x^\alpha h(x)$, where $h \in K$ and $e^{2\pi i\alpha}$ is an eigenvalue of the monodromy (as remarked in [7]). Take any $v \in \widetilde{\mathcal{O}}$. When do we have

$$Q(uv) = 0?$$

Plugging it in, we see that this holds if and only if θv is a solution of a degree $n-1$ equation of form

$$(\theta^{n-1} + c_1\theta^{n-2} + \cdots + c_{n-1})(\theta v) = 0$$

where the c_i have the form

$$c_i = b_i + f_{i,i-1}b_{i-1} + \cdots + f_{i,1}b_1 \quad f_{i,j} \text{ holomorphic.}$$

Then, by regularity, we observe that since the solution uv has moderate growth, so do the solutions θv of this equation of *lower degree*! Thus we can apply our induction hypothesis to show that all the coefficients c_i must be holomorphic, and hence all the b_i must be holomorphic also as required. $\square$

Let's reinterpret this in terms of meromorphic connections.

Definition 5.6. Let (M, ∇) be a meromorphic connection. We say the connection is *regular* (at 0) if there exists a basis $e_1, \ldots, e_n$ of M over K such that

$$\nabla e_i = -\sum_j \frac{b_{ij}(z)}{z} e_j \quad b_{ij} \in \mathcal{O}.$$

Equivalently, there exists a finitely generated submodule L of M such that $x\nabla L \subseteq L$, generating M over K.

In view of the previous two theorems, we can see that this is equivalent to the regularity of the associated system of linear differential equations. Indeed, the meromorphic connection corresponds to the system

$$\frac{du_i}{dx} = \sum_j \frac{b_{ij}(z)}{z} u_j$$

which, by 5.4, is regular precisely when we can choose $b_{ij} \in \mathcal{O}$, i.e. when there exists a basis of the above form for the associated meromorphic connection.

5.2 Regularity on algebraic curves

Let X be a smooth algebraic curve. Let us generalise the notion of a meromorphic connection, and regularity (5.6) to this setting. Fix a point $x \in X$. We replace $\mathcal{O}$ by the stalk $\mathcal{O}_{X,x}$, and K by its field of fractions $K_{X,x}$. In this setting we can make the following definition:

Definition 5.7. Let X and x be as above. A *meromorphic connection* on X at x is a pair (M, ∇), where M is a finite-dimensional $K_{X,x}$-module, and ∇ is a $\mathbb{C}$-linear map

$$\nabla \colon M \to \Omega^1_{X,x} \otimes_{\mathcal{O}_{X,x}} M$$

such that

$$\nabla(fm) = df \otimes m + f\nabla m \quad \forall\ f \in K_{X,x},\ u \in M.$$

To generalise the notion of regularity of a meromorphic connection, we use the second phrasing given in definition 5.6. This is naturally adaptable to out new situation:

Definition 5.8. Let (M, ∇) be a meromorphic connection on the algebraic curve X at a point x. We say the connection is *regular* if there exists a finitely generated submodule L of M such that$\xi \nabla L \subseteq \Omega^1_{X,x} \otimes_{\mathcal{O}_{X,x}} L$ for ξ a local parameter at x, with

$$K_{X,x} L = M$$

(i.e. L generates M over $K_{X,x}$.)

Now, since X is a curve, there exists a unique (up to isomorphism) smooth completion $\overline{X}$ of X. Take M an integrable connection on X. We can use the notion we just introduced to say what it means for M to be *regular*. Indeed, push forward M to j_*M, where $j \colon X \to \overline{X}$ is the open embedding. At each point $x \in \overline{X} \setminus X$ in turn, consider the stalk $(j_*M)_x$. We can give this the structure of a meromorphic connection at x by

$$\nabla(m) = d\xi \otimes \partial m \quad m \in (j_*M)_x,$$

where (ξ, ∂) are local coordinates for $D_{\overline{X}}$ around the point x. So

Definition 5.9. An integrable connection M is called *regular* (or has *regular singularities*) if the induced meromorphic connection $((j_*M)_x, \nabla)$ is regular at every $x \in \overline{X} \setminus X$.

Examples 5.10.

1. By considering $\mathbb{C}$ as an algebraic curve we recover some of the previous Fuchsian theory. Indeed, consider a D_1-module

   $$M = D_1/D_1P \quad P \text{ a differential operator.}$$

   When is M regular? Certainly M is holonomic, because the annihilator of $\mathrm{gr}(M)$ is non-zero. Consider the embedding $j \colon \mathbb{C} \to \mathbb{P}^1$. M is an integrable connection on $\mathbb{C}$ minus a finite set of points: the singular points of P . To be regular at such a point means precisely for the solutions of P to have moderate growth there, and likewise to be regular at ∞ means to have moderate growth as $|z| \to \infty$. Equivalently, if

   $$P = \sum_{i=0}^{n} a_i(x) \left(\frac{d}{dx}\right)^{n-i}$$

   a_i/a_0 has a pole of order at most i at each singular point in $\mathbb{P}^1$. Such a P is said to be *Fuchsian*.

2. ([8] 5.1.24) This approach does not work as nicely for Riemann surfaces as it does for algebraic curves. The problem arises when we take an integrable connection on a Riemann surface U, and attempt to produce a meromorphic connection at $x \in X \setminus U$ where $X \setminus U$ is a finite set (a

divisor). For example, consider $j\colon \mathbb{C}^\times \to \mathbb{C}$. First consider $\mathbb{C}$ as an algebraic curve. Then take two integrable connections on U

$$\begin{aligned} M &= D_U/D_U\partial \\ \text{and } N &= D_U/D_U(x^2\partial - 1). \end{aligned}$$

Pushing forward under j to $M' = j_*M$, $N' = j_*N$, these are certainly not isomorphic. Even M'^{an} and N'^{an} are not isomorphic since M'^{an} is regular at 0 but N'^{an} is not. However in the analytic category $M^{an} \cong N^{an}$, via

$$P + D\partial \leftrightarrow P\exp(1/x) + D(x^2\partial - 1)$$

since $P \in D\partial \iff P\exp(1/x) \in D(x^2\partial - 1)$. Thus taking two different – naïve algebraic – extensions to meromorphic connections may produce different results.

Proposition 5.11. Let M be a coherent D_X-module for X any algebraic variety (not necessarily a curve). Then M if is holonomic then there exists an open dense subvariety $U \subseteq X$ such that $M|_U$ is an integrable connection. For X a curve, the converse also holds.

Proof. First, notice that M is forced to be an integrable connection if and only if its characteristic variety $\mathrm{Ch}(M)$ is precisely the zero section $T^*_X X$ in T^*X. Indeed, if M is an integrable connection then we easily see this by noticing that

$$F_i M = \begin{cases} M & i \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

defines a good filtration. Conversely, work in local coordinates, and let F be a good filtration. We prove that for large i, $F_iM = F_{i+1}M = M$, and so F_iM coherent over $\mathcal{O}_X \implies M$ coherent over $\mathcal{O}_X$.

To see this, recall that in local coordinates, we can define the ideal $I \subseteq \mathbb{C}[x_1, \ldots, x_{2n}] = R_{2n}$ by

$$I = \sum_{i=1}^{n} R_{2n} x_{n+i}.$$

This ideal is the radical of $\mathrm{Ann}(gr^F(M))$ by the definition of the characteristic variety, since $\mathrm{Ch}(M) = T^*_X X$. It is Noetherian, so for some power m of I, $I^m \subseteq \mathrm{Ann}(gr^F(M))$. Hence for a multi-index $|\alpha| = m$,

$$\partial^\alpha F_j M \subseteq F_{m+j-1}M.$$

Now, this is what we need, because now for big enough j

$$\begin{aligned} F_{(m+j)}M &= (F_m D_X)(F_j M) \qquad \text{as } F \text{ is a good filtration} \\ &= \sum_{|\alpha| \leq m} \mathbb{C}[x_1, \ldots, x_n]\partial^\alpha F_j M \\ &\subseteq F_{(m+j)-1}M \end{aligned}$$

as required.

Given this, consider the set $S = \mathrm{Ch}(M) \cap (T^*X \setminus T^*_X X)$. It suffices to prove that $\dim(\pi(S)) < \dim(X)$, as then we can find an open dense subset contained in $X \setminus \pi(S)$ as required. To see this, we observe that the fibres of π over S have dimension bigger than 0. This is always true for points in the characteristic variety (since the characteristic variety is *conic*. See the proof of 4.13 and [8] appendix E.) Thus the result follows.

Conversely, if M is generically an integrable connection on a set U then it is certainly holonomic by 3.12, as it is forced to have finite-dimensional fibres on $X \setminus U$. $\square$

With this result in mind, we make the following definition:

Definition 5.12. Let X be a smooth algebraic curve. A holonomic D_X-module M is called *regular* if there exists open dense $U \subseteq X$ such that $M|_U$ is a regular integrable connection.

5.3 Regularity on algebraic varieties

Generalising these notions to higher dimensions is not so easy. At this point the algebraic and analytic paths diverge. Although there is an interesting body of theory in the analytic case (as described in chapter 5 of [8]), from now on we will focus attention on the algebraic case. We begin by introducing regularity for integrable connections on a smooth variety X, then digress to a classification theorem for simple holonomic D-modules before returning and using this to define regularity for general holonomic D-modules.

Definition 5.13. Let M be an integrable connection on X. We say M is *regular* if for every morphism $f\colon C \to X$ where C is a smooth algebraic curve, the inverse image f^*M is a regular integrable connection on C.

Remark 5.14. Let Y be a smooth algebraic variety, and let D be a divisor on Y arising as the complement of an open subvariety $j\colon X \hookrightarrow Y$. We define *meromorphic connections along* D to be D_Y-modules which are isomorphic over $\mathcal{O}_Y$ to coherent $j_*\mathcal{O}_X$ modules. There is an equivalence of categories

$$\{\text{meromorphic connections along } D\} \leftrightarrow \mathrm{Conn}(X).$$

We will not use these meromorphic connections, but they can be used to demonstrate the uniqueness of meromorphic extensions in the algebraic setting. That is, nothing like 5.10 example 2. can go wrong here.

5.4 Classification of simple holonomic D-modules

Definition 5.15. Let M be a coherent D_X-module. We say M is *simple* if $M \neq 0$ and M has no non-zero proper D_X-submodules (i.e. the only D_X-submodules are 0 and M).

Proposition 5.16. Let M be a holonomic D_X-module. Then there exists a finite *composition series* of submodules

$$0 = M_0 \leq M_1 \leq \cdots \leq M_r = M$$

such that each M_i is holonomic and M_i/M_{i-1} is simple for every i. We call the quotients M_i/M_{i-1} the *composition factors*.

Proof. Immediate from the fact that submodules and quotients of holonomic D-modules are holonomic, and the fact that holonomic D-modules have finite length (3.20). □

The aim of this section is to classify simple holonomic D-modules by showing they are all isomorphic to *minimal extensions* of the following form:

Definition 5.17. Let $j\colon Y \to X$ be a locally closed affine embedding of smooth varieties. Let M be a holonomic D_Y-module. We already know that j_*M and $j_!M$ are holonomic D_X-modules (the fact that the map is affine ensures that they are bona fide modules, rather than simply complexes of modules). In fact, as we already mentioned (see definition 3.21 and [8] 3.2.16) there is a natural transformation $j_! \to j_*$. Call the image of $j_!M \to j_*M$ the *minimal extension* of $M \in \mathrm{Mod}_h(D_Y)$, denoted $L(Y, M)$.

Theorem 5.18. Let $j\colon Y \to X$ be as above, and $N \in \mathrm{Mod}_h(D_Y)$. If N is simple then so is $L(Y, N)$, and furthermore, any simple holonomic D_X-module M is isomorphic to $L(Y, N)$ for some such Y and N.

Proof. The second assertion follows from the first. Indeed, let M be a simple holonomic D_X-module. By 5.11 there exists an open dense $j\colon Y \hookrightarrow X$ such that $j^\dagger M$ is an integrable connection. The module $j^\dagger M$ is simple by 3.19. If $L(Y, j^\dagger M)$ is simple then we're done, as

$$\mathrm{Hom}_{D_X}(j_! j^\dagger M, M) \cong \mathrm{Hom}_{D_Y}(j^\dagger M, j^\dagger M),$$

so there exists a non-zero surjective morphism $j_! j^\dagger M \to M$, so M is a quotient of $j_! j^\dagger M$, so M is a non-zero submodule of $L(Y, j^\dagger M)$.

Now we prove the first assertion, i.e. that $L(Y, N)$ is simple. First, we notice that $j_* M$ has a unique simple holonomic submodule L. Indeed, suppose there exist two such, $L_1 \neq L_2$. Then

$$j^\dagger (L_1 \oplus L_2) = j^\dagger L_1 \oplus J^\dagger L_2.$$

But this is a contradiction, since for any submodule $N \leq j_* M$, we have $j^\dagger N \cong M$. This is a little bit fiddly to see. Factor j as $j_2 \circ j_1$ for j_1 closed, j_2 open. $j^\dagger$ is exact (this is a consequence of Kashiwara's theorem) so $j^\dagger N \to j^\dagger j_* M \cong M$ is injective. By the adjunction $j_2^! \dashv j_{2*}$, the inclusion $N \to j_* M$ gives a non-zero morphism $j_2^! N \cong j_2^\dagger N \to j_{1*} M$, which is surjective because the righthand side is simple holonomic. Hence we get a surjective morphism

$$j^\dagger N \cong j_1^\dagger j_2^\dagger N \to j_1^\dagger j_{1*} M \cong M.$$

Thus we have the required isomorphism.

Given this L a unique simple submodule, it suffices to factor the canonical morphism $j_! M \to j_* M$ through L. This is now easy by the adjunction $i_! \dashv i^\dagger$:

$$\mathrm{Hom}(i_! M, i_* M) \cong \mathrm{Hom}(M, i^\dagger i_* M) \cong \mathrm{Hom}(M, M) \cong \mathrm{Hom}(M, i^\dagger L) \cong \mathrm{Hom}(i_! M, L).$$

Thus $L \cong L(Y, M)$ is simple as required. □

5.5 Preservation of regularity

Definition 5.19. We can now define what it means for a holonomic D_X-module M to be *regular*. Let M be such a D_X-module. By 5.18 we know that its composition factors F_i are isomorphic to minimal extensions $L(Y_i, N_i)$. We say M is *regular* if each N_i is a regular integrable connection on Y_i.

We denote the full subcategory of $\mathrm{Mod}(D_X)$ consisting of regular holonomic D_X-modules by $\mathrm{Mod}_{rh}(D_X)$, and we denote the full subcategory of $D^b(D_X)$ consisting of complexes with regular holonomic cohomology by $D^b_{rh}(D_X)$.

Remark 5.20. There is an alternative characterisation of regular holonomic D_X-modules, similar to definition 5.13: $M^\bullet \in D^b_h(X)$ is regular if and only if for any morphism $f\colon C \to X$ where C is a smooth algebraic curve, $f^\dagger(M^\bullet)$ has regular cohomology. The equivalence of these two notions is fairly tricky to prove, and I will not do so here. See, for example, [8] 6.1.6 for a full proof.

The following is the crucial theorem of this section: an analogue of 3.23 for $\mathrm{Mod}_{rh}(D_X)$:

Theorem 5.21. Let $f\colon X \to Y$ be a morphism of smooth varieties. Then the functors f_* and f^* preserve the full derived subcategories of regular holonomic D-modules. i.e.

$$f_* D_{rh}(D_X) \subseteq D_{rh}(D_Y) \quad \text{and} \quad f^* D_{rh}(D_Y) \subseteq D_{rh}(D_X).$$

Again, we will not prove this here. There are full accounts in [3] and in [8].

5.6 The Riemann-Hilbert correspondence

As already mentioned, we have now described two categories: $D^b_{rh}(D_X)$ and $D^b_c(\mathbb{C}_X)$ for a smooth variety X. In each, we have a family of operations, enumerated in 3.22 and 4.27 respectively. We also have a functor

$$DR_X \colon D^b_{rh}(D_X) \to D^b_c(\mathbb{C}_X)$$

by Kashiwara's constructibility theorem 4.13. The following result is probably the most important theorem in the classical theory of D-modules, describing how these families relate.

Theorem 5.22. Let X be a smooth algebraic variety. Then there is an equivalence of categories

$$D^b_{rh}(D_X) \xrightarrow{\sim} D^b_c(\mathbb{C}_X)$$

given by the de Rham functor. Furthermore, if $f \colon X \to Y$ is a morphism of smooth varieties, the de Rham functor is compatible with the functors described in 3.22. That is, the following squares commute:

$$\begin{tikzcd} D^b_{rh}(D_X) \arrow[r, "DR_X"] \arrow[d, "f_*"'] & D^b_c(\mathbb{C}_X) \arrow[d, "f_*"'] \\ D^b_{rh}(D_Y) \arrow[r, "DR_Y"] & D^b_c(\mathbb{C}_Y) \end{tikzcd} \qquad \begin{tikzcd} D^b_{rh}(D_X) \arrow[r, "DR_X"] & D^b_c(\mathbb{C}_X) \\ D^b_{rh}(D_Y) \arrow[r, "DR_Y"] \arrow[u, "f^*"] & D^b_c(\mathbb{C}_Y) \arrow[u, "f^*"] \end{tikzcd}$$

for all morphisms $f \colon X \to Y$, and

$$\boldsymbol{D}_X \circ DR_X = DR_X \circ \mathbb{D}_X,$$

which implies similar compatibility relations for $f_!$ and $f^!$.

Remark 5.23. We can immediately see that the de Rham theorem 4.8 is a very special case of this result. If we restrict to constructible sheaves on a single stratum (i.e. local systems), then the corresponding regular holonomic D-modules are simply the integrable connections on X.

Example 5.24. Let's consider the classical problem from which the Riemann-Hilbert correspondence gets its name. Let P be a linear differential operator on $X = \mathbb{C} \setminus S$ for S a finite set, and suppose P is regular at each $s \in S$. That is to say $P \in \mathrm{Mod}_{rh}(D_X)$. Then we can produce a representation of $\pi_1(X, x)$ (the monodromy) by an analogous procedure to the single singularity case described in section 5.1. The twenty-first Hilbert problem asked: given a set S of singularities and a monodromy representation ρ, can one always produce a differential operator P with this monodromy?

By the above theorem, we can answer this question affirmatively. More interestingly, one can view the correspondence as a vast generalisation of this situation to n dimensions and n variables, with the local systems on the punctured plane (or equivalently on a Riemann surface) replaced by perverse sheaves, and the ordinary differential operators replaced with regular holonomic D-modules.

Remark 5.25. The Riemann-Hilbert correspondence gives another natural interpretation of the category of perverse sheaves. Namely, they are the image of the full subcategory $\mathrm{Mod}_{rh}(D_X) \subseteq D^b_{rh}(D_X)$ under the de Rham functor DR_X. To see this corresponds to our previous definition, notice that the de Rham functor is t-exact with respect to the standard t-structure on $D^b_{rh}(D_X)$ and the perverse t-structure on $D^b_c(\mathbb{C}_X)$, and thus sends the heart of the former t-structure (i.e. $\mathrm{Mod}_{rh}(D_X)$) to the heart of the latter (i.e. perverse sheaves). Historically, it was in this context that perverse sheaves were first explicitly considered.

6 Vanishing Cycles and the Kashiwara-Malgrange Filtration

6.1 Nearby and vanishing cycles

The notions of nearby and vanishing cycles functors go back to the classical case of Milnor fibrations. If one wishes to understand a one-parameter degeneration of complex manifolds, one can investigate what happens to the (co)homology of the fibres when passing from the (homotopy equivalent) non-singular fibres to a singular fibre at zero. Deligne came up with a formal algebraic way of formalising and generalising this idea to a wider setting, as described originally in the SGA seminar, and published in [5].

The situation we want to understand is the following. Let $f\colon X \to \mathbb{C}$ be a morphism of smooth varieties. The fibre X_0 at 0 may be singular, but suppose the fibres X_t for t in the punctured disc $D^0_\varepsilon(0)$ are not singular. We aim to understand the singular fibre, first by the behaviour of the fibres ('nearby cycles'), then by analysing what is lost when passing from the nearby fibres to the singular fibre ('vanishing cycles'). More concretely, consider the following example:

Example 6.1. Consider the family of curves

$$X_\lambda\colon y^2 = x(x-1)(x-\lambda)$$

paramterised by $\lambda \in \mathbb{C}$. X_λ is an elliptic curve when $\lambda \neq 0, 1$. Let $f\colon X \to \mathbb{C}$ be the morphism of smooth varieties whose fibre over λ is X_λ. Let $\mathcal{L}$ be a local system on the smooth fibre, characterised by monodromies (g_1, g_2), $g_i \in GL_n(\mathbb{C})$ about the two generators of the fundamental group of X_λ. When the local system degenerates at the singular fibres X_0, X_1, heuristically we "lose" the information about some of the monodromy, but not all of it. This is what we hope to understand via nearby and vanishing cycles. We will consider this example in more detail in the next chapter, in section 7.2.

In the following definition we assume X_0 is the only singular fibre. We can do this by working in a small open disc around 0, and considering the neighbourhood $f^{-1}(D^0_\varepsilon(0))$ of the punctured disc in X. For simplicity of notation this is suppressed in the definitions to follow, where I will simply talk about $X \setminus X_0$.

Definition 6.2. Suppose $f\colon X \to \mathbb{C}$ is as above. We first construct the *nearby cycle functor*

$$\psi\colon D^b_c(\mathbb{C}_X) \to D^b_c(\mathbb{C}_{X_0})$$

as follows. Consider the diagram

$$\begin{tikzcd} E \arrow[r, "\widetilde{\pi}"] \arrow[d] & X \setminus X_0 \arrow[r, "j"] \arrow[d, "f"] & X & X_0 \arrow[l, "\iota"'] \\ \widetilde{\mathbb{C}^\times} \arrow[r, "\pi"] & \mathbb{C}^\times & & \end{tikzcd}$$

where E is the pullback of f by π. Then define ψ to be the composite

$$\psi = \iota^* \circ R(j \circ \widetilde{\pi})_*(j \circ \widetilde{\pi})^*\colon D^b_c(\mathbb{C}_X) \to D^b_c(\mathbb{C}_{X_0}).$$

(Of course, the inverse image functors are exact, so only the direct image needs to be replaced with the derived functor). The idea is to pull back to a *generic fibre* (E, which is homotopic to any non-singular X_λ) and then retract to the singular fibre X_0. See for example the exposition in [15] section 11.2.

Heuristically, when we pass to the singular fibre we should define the vanishing cycles to be "everything that is not a nearby cycle" in some sense. We do this via a distinguished triangle in $D^b_c(X_0)$.

Definition 6.3. Since $(j\circ\widetilde{\pi})^* \dashv R(j\circ\widetilde{\pi})_*$ is an adjunction, the unit of the adjunction gives a morphism

$$\iota^* F^\bullet \to \psi F^\bullet.$$

We define the *vanishing cycle functor*

$$\phi\colon D^b_c(X) \to D^b_c(X_0)$$

via the mapping cone of this morphism:

$$\iota^* F^\bullet \to \psi F^\bullet \overset{can}{\to} \phi F^\bullet \overset{+1}{\to}.$$

The morphism labeled *can* is called the *canonical* morphism.

Remark 6.4. Both $\psi F^\bullet$ and $\phi F^\bullet$ come equipped with natural *monodromy actions*. Indeed, let T be the anticlockwise generator of $\pi_1(\mathbb{C}^\times, *)$. Then T is an automorphism of $E = (X \setminus X_0) \times_{\mathbb{C}^\times} \widetilde{\mathbb{C}^\times}$ satisfying

$$\widetilde{\pi} \circ T = \widetilde{\pi},$$

so applying T to E induces an automorphism of $\psi F^\bullet$.

What about vanishing cycles? Consider now the following commutative diagram with distinguished rows:

$$\begin{tikzcd} \iota^* F^\bullet \arrow[r] \arrow[d] & \psi F^\bullet \arrow[r, "can"] \arrow[d, "T-1"] & \phi F^\bullet \arrow[r, "+1"] \arrow[d, dashed, "var"] & {} \\ 0 \arrow[r] & \psi F^\bullet \arrow[r, equal] & \psi F^\bullet \arrow[r] & 0 \end{tikzcd}$$

The morphism *var* (the *variation* morphism) comes from completing the diagram to a morphism of distinguished triangles. This means that in particular, we can describe a monodromy automorphism on vanishing cycles by

$$T - 1 = can \circ var.$$

The terminology of vanishing cycles comes from the classical theory of vanishing cycles for Milnor fibrations, as described in [15] Appendix C.2.

6.2 Perverse sheaves in 1 dimension

In this section we compute the simplest non-trivial example: namely the case when

$$f = j \colon \mathbb{C}^\times \to \mathbb{C}.$$

In general one can use the notion of nearby and vanishing cycles to classify the perverse sheaves on X a Riemann surface, following [2] and [6] 5.2.26. For simplicity, lets look at $X = \mathbb{C}$, $P^\bullet \in \mathrm{Perv}(\mathbb{C})$. First note the following

Example 6.5. Let $X = \{pt\}$. Then $\mathrm{Perv}(X) = \mathrm{Vect}(\mathbb{C})$. Indeed, sheaves of $\mathbb{C}_X$-modules are simply $\mathbb{C}$-vector spaces, and the derived category of such is simply the homotopy category of complexes of vector spaces. Such a complex $V^\bullet$ is perverse if and only if it has cohomology only in degree 0 as remarked in 4.16, i.e. if and only if it is homotopic to a complex supported only in degree 0.

We will use without proof the following highly non-trivial result of Gabber ([15] proposition 13.29, [6] theorem 5.2.21):

Theorem 6.6. Write ${}^p\psi$ and ${}^p\phi$ for $\psi[-1]$ and $\phi[-1]$ respectively, functors from $D^b_c(X) \to D^b_c(X_0)$. Then both ${}^p\psi$ and ${}^p\phi$ preserve the full subcategory of perverse sheaves, i.e. restrict to functors

$$\mathrm{Perv}(X) \to \mathrm{Perv}(X_0).$$

Thus, if $P^\bullet \in \mathrm{Perv}(\mathbb{C})$, then ${}^p\psi P^\bullet$ and ${}^p\phi P^\bullet$ are in $\mathrm{Perv}(\{pt\}) = \mathrm{Vect}(\mathbb{C})$. So we have a diagram of vector spaces and linear maps

$$\begin{tikzcd} {}^p\psi P^\bullet \arrow[r, shift left, "can"] & {}^p\phi P^\bullet \arrow[l, shift left, "var"] \end{tikzcd},$$

i.e. a representation of the quiver

$$\begin{tikzcd} \bullet \arrow[r, shift left] & \bullet \arrow[l, shift left] \end{tikzcd}. \tag{3}$$

In fact, this is all the information we need to characterise $P^\bullet$. For simplicity of notation, write $\mathrm{Perv}(\mathbb{C})$ for the category of perverse sheaves with respect to the *specific* stratification $\mathbb{C}^\times \cup \{0\}$. Then we have the following result:

Proposition 6.7. There is an equivalence of categories between $\mathrm{Perv}(\mathbb{C})$ and the category of representations of the quiver (3).

Proof. We have already described the functor from perverse sheaves to representations, so we will describe a functor in the other direction and check they are quasi-inverse to one another. Indeed, let

$$\Psi \underset{v}{\overset{c}{\rightleftarrows}} \Phi$$

be a representation of the quiver (3). By calculation, we can see where the irreducible representations must go, namely

$$\begin{array}{lll} \mathbb{C} \rightleftarrows 0 & \mapsto \mathbb{C}_0 & \text{(skyscraper sheaf)} \\ 0 \rightleftarrows \mathbb{C} & \mapsto \mathbb{C}[1] & \text{(constant sheaf)} \\ \mathbb{C} \underset{\times\lambda}{\overset{\times 1}{\rightleftarrows}} \mathbb{C} & \mapsto j_*\mathcal{L}_{\lambda+1}[1] & \end{array}$$

where $j\colon \mathbb{C}^\times \to \mathbb{C}$ and $\mathcal{L}_\lambda$ is the rank one local system on $\mathbb{C}^\times$ with monodromy λ.

In general, we produce a perverse sheaf

$$P = \left(0 \to P^{-1} \xrightarrow{0} P^0 \to 0\right)$$

as follows. Define P^{-1} with respect to the given stratification by $P_0^{-1} = \ker(c)$, and $P^{-1}|_{\mathbb{C}^\times} = \Psi$ with monodromy $v \circ c + 1$. Define $P_0^0 = \mathrm{coker}(c)$, $P^0|_{\mathbb{C}^\times} = 0$. These are both certainly constructible, and the whole complex P is perverse because one can show (see [6] example 5.2.23) that the perverse complexes with respect to this stratification are precisely those of form

$$0 \to F^{-1} \to F^0 \to 0$$

with $H^0(F^\bullet)$ supported on $\{0\}$ and $\Gamma_{\{0\}}(H^{-1}(F^\bullet)) = 0$.

We must check that these functors define an equivalence. Firstly it is clear that for a representation of the quiver, applying the two functors gives back precisely the original representation, because the nearby and vanishing cycle complexes of these perverse sheaves are easy to calculate. We must check the other direction, so let $P^\bullet$ be a perverse sheaf on $\mathbb{C}$ with respect to the given stratification, without loss of generality of the form

$$P = 0 \to P^{-1} \xrightarrow{0} P^0 \to 0.$$

Then by computing the vanishing cycles of this complex the result follows readily. □

6.3 The Kashiwara-Malgrange filtration

At this point, it is natural to ask how we might interpret these vanishing cycles in the category of regular holonomic D-modules, under the Riemann-Hilbert correspondence. The answer is to introduce the *Kashiwara-Malgrange filtration* on such a D-module. This filtration V^q will relate to the vanishing cycle functor according to the following result ([18], 3.4.12).

Recall that for any invertible operator T on a finite-dimensional complex vector space, we can decompose T uniquely as a product $T = T_sT_u$ where T_s is semisimple, T_u is unipotent, and $T_uT_s = T_sT_u$. By the *semisimple part* of the monodromy operator, we mean the semisimple T_s in this Jordan decomposition.

Theorem 6.8. let $f\colon X \to \mathbb{C}$ be a regular function. Let M be a regular holonomic D_X-module, and let $\widetilde{M}$ be its pushforward under the graph $(1,f)\colon M \to M \times \mathbb{C}$. If the filtration V^q exists, then there are canonical isomorphisms

$$DR_{X_0}(gr_q^V \widetilde{M}) \cong {}^p\phi_{\exp(2\pi i q)}(DR_X(M)) \quad q \in (-1,0]$$

where $\mathrm{gr}_q^V = V^q/V^{q+1}$, and ${}^p\phi_{\exp(2\pi i q)}$ means the $\exp(2\pi i q)$-eigenspace of ${}^p\phi$ with respect to the action of the semisimple part of the monodromy T described in 6.14.

A good introduction to the Kashiwara-Malgrange filtration is [4]. There are also accounts in [2], [15] and [14].

Definition 6.9. Let X be a smooth algebraic variety, $Y \subseteq X$ a codimension one subvariety in X. We first describe a filtration called the V-filtration on D_X associated to Y. Then the Kashiwara-Malgrange filtration on a holonomic D_X-module M will be a filtration on M that is compatible with the V-filtration in a natural way.

Let $\mathcal{I}$ be the sheaf of ideals corresponding to Y. Define an exhaustive decreasing filtration on $\mathcal{O}_X$ first by

$$V^k\mathcal{O}_X = \begin{cases} \mathcal{I}^k & \text{if } k > 0 \\ \mathcal{O}_X & \text{if } k \leq 0 \end{cases}.$$

This then defines a filtration on D_X by putting

$$V^k D_X = \{P \in D_X \,:\, P\mathcal{I}^i \subseteq \mathcal{I}^{i+k} \,\forall\, i \in \mathbb{Z}\}.$$

This agrees with the notion just given on the subsheaf $\mathcal{O}_X$, and if locally X has coordinates $x_1, \ldots, x_l, y_1, \ldots, y_m$, and $Y = \{y_1 = \cdots = y_m = 0\}$ we can view the filtration as

$$V^k D_X = \sum_{|\beta|-|\gamma| \geq k} f_{\alpha,\beta,\gamma}(x)\partial_x^\alpha y^\beta \partial_y^\gamma$$

with the natural multi-index notation.

Definition 6.10. Let M be a holonomic D_X-module, and let V^i be the V-filtration associated to a codimension one subvariety Y in X. A *V-filtration on M* is a decreasing exhaustive *rational* filtration $(V^qM)_{q\in\mathbb{Q}}$ by coherent V^0D_X-modules such that

1. $(V^kD_X)(V^qM) \subseteq V^{q+k}M$. This is what is meant by compatibility with the filtration of 6.9.
2. If $q > 0$, then the inclusion $(V^1D_X)(V^qM) \hookrightarrow V^{q+1}M$ is an equality.
3. The action of $y\partial - q$ on $\mathrm{gr}_q^V M$ is nilpotent.

Lemma 6.11. If such a V-filtration exists, it is unique, and is called the *Kashiwara-Malgrange filtration.*

Proof. Suppose V_1 and V_2 were two filtrations satisfying the conditions on a given $M \in \mathrm{Mod}_h(D_X)$. Then first of all, we can find rational numbers a and b such that

$$V_1^{q+a}M \supseteq V_2^qM \supseteq V_1^{q+b}M \quad \forall\, q \in \mathbb{Q}.$$

Suppose for contradiction that we can pick such a, b such that $b - a > 0$ is minimal. Nilpotence of $y\partial - q$ on $\mathrm{gr}_q^{V_1} M$ implies that

$$\begin{aligned} &\mathrm{gr}_q^{V_2}\mathrm{gr}_{q'}^{V_1} M = 0 \quad \forall\, q \neq q' \\ \Longrightarrow\ &V_2^q\,\mathrm{gr}_{q+b}^{V_1} M = 0 \quad \forall\, b > 0 \\ \Longrightarrow\ &V_2^qM \subseteq V_1^{q+b-\varepsilon} \quad \text{for some } \varepsilon > 0 \end{aligned}$$

which is a contradiction of the minimality of $b - a$, so in fact $V_1 = V_2$. $\square$

Definition 6.12. What if Y is a codimension one subvariety which is *not* smooth? In this case, we can still define the V-filtration along Y. Indeed, we have a morphism of locally constant $\mathcal{O}_X$-modules $\mathcal{O}_X \hookrightarrow \mathcal{O}_X(Y)$, with associated bundle morphism $\sigma\colon X \to E$. Let E_0 be the zero section of E. We say that a rational filtration $V^q M$ is a V-filtration on M associated to $Y \subseteq X$ if $\sigma_*(V^q M)$ is a V-filtration on $\sigma_* M$ associated to $E_0 \subseteq E$ in the above sense.

Remark 6.13. For Y smooth, one can check that this definition agrees with the original definition, so this does make sense.

Suppose X is fibred by $f\colon X \to \mathbb{C}$, and $Y = X_0$ is a possibly singular fibre. Under what conditions does the Kashiwara-Malgrange filtration exist for this Y? For regular holonomic D_X-modules M we have a strong condition for the existence of the filtration, proven by Malgrange. ([15] Theorem 14.23)

Theorem 6.14. We say $M \in \mathrm{Mod}_{rh}(D_X)$ has *quasi-unipotent monodromy* along the fibre X_0 if the monodromy action T on $\psi(DR_X(M))$ described in remark is *quasi-unipotent*: that is, if there exists some $n \in \mathbb{N}$ such that $T^n - 1$ is nilpotent. If M has quasi-unipotent monodromy then then Kashiwara-Malgrange filtration exists on M.

7 Applications to Hodge Theory

We conclude the essay by giving some applications of the theory of D-modules to problems of Hodge theory. We introduce the classical language of Hodge theory: Hodge structures on vector spaces. These can be viewed as local systems on a point, and so we can use the techniques introduced in chapter 4 to give an integrable connection analogue on a more general variety: a *variation of Hodge structure*. We then use the Kashiwara-Malgrange filtration to describe how one might understand degenerations of Hodge structures near a singular fibre of a one-parameter family of curves. Finally we remark on Saito's work on Mixed Hodge Modules, which is a vast generalisation of the ideas in this chapter.

The main sources for this material are the book [15] of Peters and Steenbrink, and the lecture notes [17] of Sabbah.

7.1 Hodge structures and variations of Hodge structure

The classical result that inspired the development of Hodge theory was the following *Hodge decomposition* for de Rham cohomology of a compact Kähler manifold:

Theorem 7.1. Let X be a compact Kähler manifold. Let $H^{p,q}(X)$ be the spaces of cohomology classes of type (p,q). Then there is a direct sum decomposition

$$H^m(X, \mathbb{C}_X) \cong \bigoplus_{p+q=m} H^{p,q}(X)$$

and $H^{p,q}(X) \cong \overline{H^{p,q}(X)}$.

The idea is to consider in more generality modules with an analogous Hodge decomposition, defined as follows:

Definition 7.2. Let $V_{\mathbb{Q}}$ be a $\mathbb{Q}$ vector space. A *Hodge structure* of weight w on $V_{\mathbb{Q}}$ is a direct sum decomposition

$$V = V_{\mathbb{Q}} \otimes_{\mathbb{Q}} \mathbb{C} \cong \bigoplus_{p+q=w} V^{p,q}$$

such that $V^{p,q} \cong \overline{V^{p,q}}$. Equivalently, we can define a Hodge structure via a filtration

$$F_p V_{\mathbb{C}} = \bigoplus_{p' \geq p} V^{p',w-p'}$$

called the *Hodge filtration*, satisfying the condition that for any $p \in \mathbb{Z}$,

$$F_p V \cap \overline{F_{w-p+1}} = \{0\}.$$

A morphism of Hodge structures is, unsurprisingly, a linear map between vector spaces that preserves the Hodge filtration.

Remark 7.3. Here, we are considering $\mathbb{Q}$-Hodge structures. We could equally well have taken $V_{\mathbb{Z}}$ to be a $\mathbb{Z}$-module and made the same definition. Indeed this is the approach taken in [15]. However, when we go on to talk about variations of Hodge structures, we'd like to use the analogy that a $\mathbb{Q}$ vector space is a $\mathbb{Q}$-local system at a point, for close analogy with our discussion of local systems in chapter 4.

Example 7.4. We've already seen one example of a Hodge structure, namely cohomology groups. Let's introduce another example, via *Tate twists.* First we define the basic Hodge structures $\mathbb{Z}(k)$. If $k = 0$, this is just the Hodge structure $\mathbb{C}$ with the standard underlying $\mathbb{Q}$-vector space $\mathbb{Q}$. If $k = 1$, this is the Hodge structure $\mathbb{C}$ with underlying $\mathbb{Q}$-vector space $2i\pi\mathbb{Q}$. Then for general $k \in \mathbb{Z}$, $\mathbb{Z}(k) = \mathbb{Z}(1)^{\otimes k}$, with weight $-2k$.

If V is a weight w Hodge structure, its k^{th} *Tate twist* is a weight $w - 2k$ Hodge structure with underlying vector space $V \otimes \mathbb{Z}(k)$, and Hodge structure $V(k)^{p,q} = V^{p-k,q-k}$. The name comes from the similar notion of the Tate twist of a Galois module.

Definition 7.5. Let V be a Hodge structure of weight w. A *polarization* of the Hodge structure V is a $(-1)^w$-symmetric morphism of Hodge structures of weight w

$$Q\colon V \otimes V \to \mathbb{Z}(-w)$$

such that the bilinear form

$$Q'(u, v) = (2i\pi)^w Q(Cu, v)$$

is positive definite. Here C is the operator that sends $u \in V^{p,q}$ to $i^{p-q}u$, sometimes called the Weil operator.

We can consider Hodge structures as vector bundles on a point, with some additional structure. The right idea to generalise this to other varieties is the correspondence between integrable connections and local systems (4.8).

Definition 7.6. Let X be a smooth variety. A *variation of Hodge structure* on X of weight w is a local system $V_{\mathbb{Q}}$ of $\mathbb{Q}_X$-modules on X equipped with a finite increasing filtration

$$F_p(V_{\mathbb{Q}} \otimes_{\mathbb{Q}} \mathcal{O}_X)$$

of subbundles (the *Hodge filtration*) satisfying the following conditions:

1. At each point $y \in Y$, the filtered vector space $(V_{\mathbb{Q},y} \otimes_{\mathbb{Q}} \mathbb{C}, F^\bullet)$ is a Hodge structure of weight w on the stalk $V_{\mathbb{Q},y}$.

2. (*Griffiths Transversality*) By theorem 4.8, the local system $V_{\mathbb{Q}}$ induces an integrable connection (V, ∇) on X. This connection should satisfy the condition

$$\nabla(F_i V_{\mathbb{Q}}) \subseteq F_{i-1} V_{\mathbb{Q}} \otimes \Omega^1_X$$

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

Example 7.7. Let V be a Hodge structure of weight w, and let

$$\rho\colon \pi_1(Y,*) \to \mathrm{Aut}(V)$$

be a representation of π_1, with corresponding local system V_ρ. The Hodge structure on V gives a natural Hodge filtration to V_ρ which trivially satisfies Griffiths transversality. Thus ρ makes V into a variation of Hodge structure. Denote $V_{\mathbf{1}}$ by V_Y.

Definition 7.8. let V be a variation of Hodge structure. A *polarization* of V is a morphism of variations of Hodge structure

$$Q\colon V \otimes V \to \mathbb{Z}(-k)_Y$$

such that the induced morphisms on the fibres are themselves polarizations of Hodge structures.

Example 7.9 ([15] 10.7)**.** We have the following crucial example of a variation of Hodge structure: let $f\colon X \to Y$ be a proper smooth morphism of smooth varieties. Then the k^{th} cohomology of the fibres $H^k(X_\lambda, \mathbb{Q})$ defines a local system on S. In fact this local system underlies a variation of Hodge structure of weight k such that the induced Hodge structures on the stalks are just the usual Hodge structures on the cohomology of the fibres $H^k(X_\lambda, \mathbb{C})$. This is because there exists a connection on the *relative de Rham cohomology*

$$H^k_{dR}(X/Y) = R^k f_* \mathbb{C}_X \otimes_{\mathbb{C}} \mathcal{O}_Y$$

called the *Gauss-Manin connection* ∇_{GM}, and a Hodge filtration satisfying Griffiths transversality. Such examples are sometimes called *geometric* variations of Hodge structure.

Example 7.10. Let $f\colon X \to \mathbb{C}$ have smooth fibres away from a finite set $S \subseteq \mathbb{C}$. Then, as in the above example, the cohomology groups $H^k(X_\lambda)$ form a local system on $\mathbb{C} \setminus S$, which can naturally be viewed as the underlying local system of a variation of Hodge structure.

The above example illustrates how the theory of D-modules that we have developed may be helpful. Suppose that rather than simply being interested in a local system on $\mathbb{C} \setminus S$ like the cohomology of the smooth fibres, we want to understand a sheaf that also takes into account data on the singular fibres. Such an object will then be a constructible complex with respect to the stratification $(\mathbb{C} \setminus S) \cup S$: i.e. an object of $D^b_c(\mathbb{C})$. By the Riemann-Hilbert correspondence we may equivalently view this as a complex of regular holonomic D-modules. In the previous chapter, we introduced a method for understanding the behaviour of such objects, namely the Kashiwara-Malgrange filtration. We will explain how this works with an extended example.

7.2 A one-parameter degeneration of elliptic curves

In this section we will study the example described in 6.1. That is, the smooth variety $f\colon X \to \mathbb{C}$ whose fibres X_λ are the curves

$$y^2 = x(x-1)(x-\lambda).$$

The fibres are smooth with the exception of the two singular fibres X_0 and X_1. We will consider degenerations of Hodge structures on X at these singular fibres. Due to time constraints I will unfortunately only lay out how the calculation might go, and will not actually perform the calculation.

Consider the sheaf on $\mathbb{C}$ whose stalks are the k^{th} cohomology groups of the fibre X_λ, i.e. $H^k(X_\lambda, \mathbb{Q})$. This is constructible with respect to the obvious stratification $(\mathbb{C} \setminus \{0,1\}) \cup \{0,1\}$, as in example 7.9, so corresponds to a complex $M^\bullet \in D^b_{rh}(D_1)$ by the Riemann-Hilbert correspondence. We would like to compute the Kashiwara-Malgrange filtration on $M^\bullet$ at the two singular fibres.

Equivalently, by 6.8, if we understand the graded structure of $M^\bullet$ under the Kashiwara-Malgrange filtration, then we know about the eigenspaces of the monodromy operator on vanishing cycles of this cohomology sheaf at the singular fibres.

How might we go about computing this filtration? Using the techniques of [14], we can reduce the problem to another more tractable one. Indeed, if we have a holomorphic D_X-module, by Kashiwara's theorem we can consider it as a holonomic D_3 module with support contained in X. This allows us to consider a local neighbourhood about the singular point of X at 0, with coordinates x, y and t. Take $M \in \mathrm{Mod}_h^X(D_3)$, and let m be a local generator at this point. By proposition 4.2.1 in [14], the existence of a V-filtration is equivalent to the existence of a b-function for m, i.e. a non-zero $b \in \mathbb{C}[s]$ such that

$$b(x\partial_x + y\partial_y + t\partial_t)m \in (V^{-1}D_3)m$$

where the V-filtration is with respect to the hyperplane $t = 0$ say.

Take the constant D_X-module, and let M be its direct image under the embedding $\iota\colon X \to \mathbb{C}^3$. Let m be a local generator near the singularity at 0. We could aim to compute the b-function for m, as M should correspond, under the de Rham functor and direct image along f, to the cohomology sheaf we want to understand.

Thus we have produced a recipe for computing the vanishing cycles near the singularity, namely consider the constant D-module on the total space of the family as a D_3-module, compute the b-function of a local generator near the singularity, compute the Kashiwara-Malgrange filtration by techniques in [14], push down along f and apply the de Rham functor.

7.3 Remarks on mixed Hodge modules

I will now describe some of the generalisations of this theory developed by Saito, described in [15]. Let X be a *general* algebraic variety, not necessarily smooth. One can produce an analogue of the Riemann-Hilbert correspondence in this setting. Essentially, Saito's idea was to replace the filtered integrable connection of a variation of Hodge structure with a D-module equipped with a good filtration. This is, then, an analogue of a variation of Hodge structure where singularities are allowed to arise (i.e. we're working with regular holonomic D-modules rather than integrable connections). These objects are called *Hodge modules*, but the precise definition is rather involved. In this chapter we will at most sketch out some of the ideas and consequences of the theory.

Since we are now thinking about Hodge structures on constructible sheaves, the behaviour along singularities becomes crucial. Luckily, we have introduced the dual perspectives of vanishing cycles and the Kashiwara-Malgrange filtration that allow us to analyse such things. Essentially, a Hodge module $M$of weight w is a filtered regular holonomic D_X-module with an underlying rational structure which is a *rational perverse sheaf.* Furthermore, we require the existence of certain filtrations on the nearby and vanishing cycle complexes of M, whose k^{th} graded parts are *themselves* weight k Hodge modules supported on the singular loci. Thus in a sense we define Hodge modules inductively, with variations of Hodge structure playing the role of the base case.

The following theorem of Saito is a generalisation of the similar property of variations of Hodge structure:

Theorem 7.11 (Hodge-Saito)**.** Let X be a projective variety, and let M be a polarized Hodge module of weight w on X. Then there is the natural structure of a polarized Hodge structure on the hypercohomology

$$\mathbb{H}^k(X, DR_X(M)[\dim(X)])$$

arising from the Hodge module structure on M.

The real objects Saito aimed to study were *mixed* analogues of these Hodge modules. The notion of a mixed Hodge *structure* is classical:

Definition 7.12. A *mixed Hodge structure* V is a $\mathbb{Q}$-vector space equipped with *two* filtrations: a decreasing Hodge filtration $F^\bullet$ as before, and also an increasing filtration $W^\bullet$ called the *weight filtration*. We require that $F^\bullet$ induces a (pure) Hodge structure on each graded piece

$$\mathrm{gr}_W^k(V) = W^k V / W^{k-1} V$$

of weight k.

Such structures were first invented to define Hodge structures on intersection cohomology groups analogous to the classical Hodge structures on (singular) cohomology. [3] This allows one to prove results like the two Lefschetz theorems in the setting of intersection cohomology. By generalising these notions, one can prove powerful results like the intersection homology analogue of the *decomposition theorem*:

Theorem 7.13 (Decomposition theorem for mixed Hodge modules)**.** Let $f\colon X \to Y$ be a proper map of varieties. Denote the category of mixed Hodge modules on X by $MHM(X)$. Then if $M \in D^b(MHM(X)$ is semisimple and pure (i.e. actually just a complex of polarized Hodge modules) then

$$f_*M \cong \bigoplus_{i\in\mathbb{Z}} \mathcal{H}^i(f_*M)[-i]$$

and all the $\mathcal{H}^i(f_*M)$ are themselves semisimple and pure.

From this, it is possible to prove

Theorem 7.14 (Decomposition theorem for intersection complexes)**.** Let $f\colon X \to Y$ be a proper map of varieties. Denote the intersection complex of X by IC_X. Then there is an isomorphism

$$Rf_*IC_X \cong \bigoplus_{i\in\mathbb{Z}} {}^p\mathcal{H}^i(Rf_*IC_X)[-i]$$

and all the perverse sheaves ${}^p\mathcal{H}^i(Rf_*IC_X)$ are semisimple. Here ${}^p\mathcal{H}^i$ are the so-called *perverse cohomology sheaves.*

For a good exposition of these ideas and more, see the article [1] of Andrea, De Cataldo and Migliorini.

References

[1] M. Andrea, A. De Cataldo, and L. Migliorini. The decomposition theorem, perverse sheaves and the topology of algebraic maps. *Bull. Amer. Math. Soc*, 46(4):535–633, 2009.

[2] D. Arapura. Notes on D-modules and connections with Hodge theory. Unpublished notes available online at http://www.math.purdue.edu/ dvb/preprints/dmod.pdf.

[3] J. Bernstein. Algebraic theory of D-modules. Unpublished notes available online at http://www.math.uchicago.edu/∼arinkin/langlands/Bernstein/Bernstein-dmod.pdf.

[4] N. Budur. On the V-filtration of D-modules. In F. Bogomolov and Y. Tschinkel, editors, *Geometric methods in algebra and number theory*, volume 235 of *Progress In Mathematics*. Birkhäuser, 2005.

[5] P. Deligne and N. Katz. *Groups de Monodromie en Géométrie Algébrique. SGA 7.II*, volume 340 of *Lecture Notes In Mathematics*. Springer-Verlag, 1973.

[6] A. Dimca. *Sheaves in Topology*. Springer-Verlag, 2004.

[3] See [12] for a general introduction to intersection (co)homology.

[7] A. Haefliger. Local theory of meromorphic connections in dimension one (Fuchs theory). In A. Borel, editor, *Algebraic D-modules*, chapter 3, pages 151–172. Academic Press, Inc., 1987.

[8] R. Hotta, K. Takeushi, and T. Tanisaki. *D-Modules, Perverse Sheaves and Representation Theory*, volume 236 of *Progress In Mathematics.* Birkhäuser, 2008.

[9] B. Iversen. *Cohomology of Sheaves.* Springer-Verlag, 1986.

[10] M. Kashiwara. On the maximally overdetermined system of linear differential equations I. *Publ. Res. Inst. Math. Sci.*, 10(2):563–579, 1974.

[11] M. Kashiwara and P. Schapira. *Sheaves on Manifolds*, volume 292 of *A Series of Comprehensive Studies in Mathematics.* Springer-Verlag, 1990.

[12] F. Kirwan and J. Woolfe. *An Introduction to Intersection Homology Theory.* Chapman & Hall, 2^{nd} edition, 2006.

[13] R. MacPherson. Intersection homology and perverse sheaves. Unpublished notes available online at http://faculty.tcu.edu/gfriedman/notes/ih.pdf, 1990.

[14] Z. Mebkhout. *Le Formalise des Six Opérations de Grothendieck pour les D_X-Modules Cohérents*, volume 35 of *Travaux en Cours.* Hermann, 1989.

[15] C. Peters and J. Steenbrink. *Mixed Hodge Structures*, volume 52 of *A Series of Modern Surveys In Mathematics.* Springer-Verlag, 2008.

[16] K. Rietsch. An introduction to perverse sheaves. In *Representations of Finite Dimensional Algebras and Related Topics in Lie Theory and Geometry*, volume 40 of *Fields Institute Communications.* AMS, 2004.

[17] C. Sabbah. Hodge theory, singularities and D-modules. Unpublished notes available online at http://www.math.polytechnique.fr/cmat/sabbah/sabbah_luminy07.pdf, 2007.

[18] M. Saito. Modules de Hodge polarisables. *Publ. Res. Inst. Math. Sci. Kyoto Univ.*, 24(6):849–995, 1988.

[19] M. Sato, T. Kawai, and M. Kashiwara. Microfunctions and psuedo-differential equations. In *Hyperfunctions and Pseudo-differential Equations (Katana 1971)*, volume 287 of *Lecture Notes In Mathematics.* Springer-Verlag, 1973.

[20] R. Vakil. An example of a nice variety whose ring of global sections is not finitely generated. Unpublished notes available online at http://math.stanford.edu/ vakil/files/nonfg.pdf, 2000.

[21] C. Voisin. *Hodge Theory and Complex Algebraic Geometry I.* Cambridge University Press, 2002.