

Cohomology

This lecture offers a first glimpse at sheaf theory, which involves an interplay of algebra, geometry, and topology. The goal is to introduce sheaves through examples, and to motivate their study by pointing out connections to analysis and to topology. In the process, there is little proof and much hand-waving. The reader may wish to revisit this lecture when equipped with the algebraic tools necessary to appreciate the principles outlined here.

The point of view adopted here is that sheaves are spaces of functions. This is the approach pioneered by the French school around Cartan and Leray. Functions are typically defined by local conditions, and it is of interest to determine global functions with prescribed local properties. For example, the Mittag-Leffler problem specifies the principal parts of a holomorphic function on a Riemann surface at finitely many points, and asks for the existence of a global holomorphic function with the given principal parts. Such local-to-global problems are often nontrivial to solve, and lead to the notion of sheaf cohomology.

There are many sources for reading on sheaves and their applications. For the algebraic geometry approach, we refer to Hartshorne's book [61]; for the differential geometric aspects, one can consult Godement [47] and Griffiths and Harris [51]. The required homological background can be found in Weibel [161], while Gelfand and Manin [46] and Iversen [83] show the workings of homological algebra in the context of sheaf theory. For the link between calculus and cohomology, we recommend [14, 110], and [29, 88] for connections with D -modules and singularity theory.

1. Sheaves

We assume that the reader is familiar with basic concepts of point set topology. Let us fix a space X with topology \mathcal{T}_X .

Example 2.1. While open and closed sets abound in some familiar spaces such as \mathbb{R}^n , a topology may be quite sparse. At one extreme case, the only sets in \mathcal{T}_X are X and the empty set \emptyset . This scenario is known as the *trivial topology*. At the other extreme, all subsets of X are open, and hence all are closed as well. In that case, X is said to have the *discrete topology*.

Most topologies are somewhere between the two extremes considered in the preceding example. The case of interest to us is the spectrum of a commutative ring, with the Zariski topology, and typically this has fewer open sets than one is accustomed to in \mathbb{R}^n .

Example 2.2. Let (X, \mathcal{T}_X) be the spectrum of the ring $R = \mathbb{C}[x]$. The points of X are prime ideals of R , which are the ideals $\{(x - c)\}_{c \in \mathbb{C}}$ together with the ideal (0) . The closed sets in \mathcal{T}_X are finite collections of ideals of the type $(x - c)$, and the set X . This topology is fairly coarse: for instance, it is not Hausdorff. (Prove this.)

Let F be a topological space. We attach to each open set U of X , the space of all continuous functions $C(U, F)$ from U to F . The following is the running example for this lecture.

Example 2.3. Let X be the unit circle \mathbb{S}^1 with topology inherited from the embedding in \mathbb{R}^2 . Any proper open subset of \mathbb{S}^1 is the disjoint union of open connected arcs. Let F be \mathbb{Z} with the discrete topology.

Let U be an open set of \mathbb{S}^1 , and pick an element $f \in C(U, \mathbb{Z})$. For each $z \in \mathbb{Z}$, the preimages $f^{-1}(z)$ and $f^{-1}(\mathbb{Z} \setminus \{z\})$ form a decomposition of U into disjoint open sets. Therefore, if U is connected, we deduce that $f(U) = \{z\}$ for some z , and hence $C(U, \mathbb{Z}) = \mathbb{Z}$.

In order to prepare for the definition to come, we revise our point of view. Consider the space $\mathbb{Z} \times X$ with the natural projection $\pi: \mathbb{Z} \times X \rightarrow X$. For each open set U in X , elements of $C(U, \mathbb{Z})$ can be identified with continuous functions $f: U \rightarrow \mathbb{Z} \times X$ such that $\pi \circ f$ is the identity on U . In this way, the elements of $C(U, \mathbb{Z})$ become *sections*, that is, continuous lifts for the projection π .

Definition 2.4. Let X be a topological space. A *sheaf* \mathcal{F} on X is a topological space F , called the *sheaf space* or *espace étalé*, together with a surjective map $\pi_{\mathcal{F}}: F \rightarrow X$ which is locally a homeomorphism.

The sheaf \mathcal{F} takes an open set $U \subseteq X$ to the set $\mathcal{F}(U) = C(U, F)$ of continuous functions $f: U \rightarrow F$ for which $\pi_{\mathcal{F}} \circ f$ is the identity on U . Each

inclusion $U' \subseteq U$ of open sets of X yields a *restriction map*

$$\rho_{U,U'}: \mathcal{F}(U) \longrightarrow \mathcal{F}(U').$$

If $U'' \subseteq U' \subseteq U$ are open subsets of X , then

$$\rho_{U',U''} \circ \rho_{U,U'} = \rho_{U,U''}.$$

The elements of $\mathcal{F}(U)$ are *sections* of \mathcal{F} over U ; if $U = X$ these are known as *global sections*.

Remark 2.5. The sheaf space F of X is locally homeomorphic to X in the sense that for every point $f \in F$ there is an open neighborhood $V \subseteq F$ of f such that $\pi_{\mathcal{F}}: V \longrightarrow \pi_{\mathcal{F}}(V)$ is a homeomorphism.

This does not mean that F is a covering space of X . Indeed, let $X = \mathbb{R}^1$ and S a discrete space with more than one point. Construct the sheaf space F of \mathcal{F} as the quotient space of $X \times S$ obtained by identifying (x, s) with (x, s') for $x \neq 0$ and $s, s' \in S$. Set $\pi_{\mathcal{F}}(x, s) = x$. Then \mathcal{F} is a *skyscraper sheaf* with fiber S at $x = 0$. Evidently, $\pi_{\mathcal{F}}$ is not a covering map:




Figure 2.1. A skyscraper sheaf over the real line

An important class of sheaves, including Example 2.3, is the class of *constant sheaves*. These arise when the sheaf space F is the product of X with a space F_0 equipped with the discrete topology. In this case, sections of \mathcal{F} on an open set U are identified with continuous maps from U to F_0 .

Remark 2.6. We will see in Lecture 12 that one may specify a sheaf without knowing the sheaf space F . Namely, one may prescribe the sections $\mathcal{F}(U)$, so long as they satisfy certain compatibility properties. This is useful, since for many important sheaves, particularly those that arise in algebraic geometry, the sheaf space F is obscure, and its topology \mathcal{T}_F complicated.

In general, the sections $\mathcal{F}(U)$ of a sheaf form a set without further structure. There are, however, more special kinds of sheaves, for example, sheaves of Abelian groups, sheaves of rings, etc., where, for each open set U , the set

$\mathcal{F}(U)$ has an appropriate algebraic structure, and the restriction maps are morphisms in the corresponding category.

Example 2.7 (The constant sheaf \mathcal{Z}). With $X = \mathbb{S}^1$ as in Example 2.3, let \mathcal{Z} be the sheaf associated to the canonical projection $\pi_{\mathcal{Z}}: X \times \mathbb{Z} \rightarrow X$ where \mathbb{Z} carries the discrete topology. Since \mathbb{Z} is an Abelian group, $\mathcal{Z}(U)$ is an Abelian group as well under pointwise addition of maps. The restriction maps $\mathcal{Z}(U) \rightarrow \mathcal{Z}(U')$, for $U' \subseteq U$, are homomorphisms of Abelian groups.

The sheaf \mathcal{Z} can be used to show that the circle is not contractible. The remainder of this lecture is devoted to two constructions which relate \mathcal{Z} to the topology of \mathbb{S}^1 : Čech complexes and derived functors.


2. Čech cohomology

Consider the unit circle \mathbb{S}^1 embedded in \mathbb{R}^2 . The sets $U_1 = \mathbb{S}^1 \setminus \{-1\}$ and $U_2 = \mathbb{S}^1 \setminus \{1\}$ constitute an open cover of \mathbb{S}^1 . Set $U_{1,2} = U_1 \cap U_2$. Since U_1 and U_2 are connected, $\mathcal{Z}(U_i) = \mathbb{Z}$ for $i = 1, 2$. The restriction maps $\mathcal{Z}(\mathbb{S}^1) \rightarrow \mathcal{Z}(U_i)$ are isomorphisms, since each of \mathbb{S}^1 , U_1 , and U_2 is a connected set. With $U_{1,2}$, however, the story is different: its connected components may be mapped to distinct integers, and hence $\mathcal{Z}(U_{1,2}) \cong \mathbb{Z} \times \mathbb{Z}$. More generally, for an open set U , we have


$$\mathcal{Z}(U) = \prod \mathbb{Z},$$

where the product ranges over the connected components of U .

One has a diagram of inclusions



which gives rise to a commutative diagram of Abelian groups



Consider the maps $\mathcal{Z}(U_i) \rightarrow \mathcal{Z}(U_{1,2})$ for $i = 1, 2$. An element of $\mathcal{Z}(U_{1,2})$ is given by a pair of integers (a, b) . If this element lies in the image of $\mathcal{Z}(U_i)$, then we must have $a = b$, since U_i is connected. It follows that the image of $\mathcal{Z}(U_i) \rightarrow \mathcal{Z}(U_{1,2})$ is $\mathbb{Z} \cdot (1, 1)$ for $i = 1, 2$.

In a sense, the quotient of $\mathcal{Z}(U_{1,2})$ by the images of $\mathcal{Z}(U_i)$ measures the insufficiency of knowing the value of a section at one point in order to determine the entire section. In more fancy terms, it describes the possible \mathbb{Z} -bundles over \mathbb{S}^1 . Since U_1 and U_2 are contractible, any bundle on them is given by the product of U_i with the fiber of the bundle. The question then arises how the sections on the open sets are identified along their intersection $U_{1,2}$. Choose generators for $\mathcal{Z}(U_1) = \mathbb{Z} = \mathcal{Z}(U_2)$, and suppose the two generators agree over one connected component of $U_{1,2}$. On the other connected component, these generators could agree, or could be inverses under the group law. In the former case, one gets the trivial bundle on \mathbb{S}^1 , and in the latter case, the total space of the bundle is a “discrete Möbius band.” The trivial bundle corresponds to the section $(1, 1)$ over $U_{1,2}$, while the Möbius strip is represented by $(1, -1)$. These are displayed in Figure 2.2.




Figure 2.2. The trivial bundle and the Möbius bundle over the circle

In terms of sections, an element $(a, b) \in \mathcal{Z}(U_1) \times \mathcal{Z}(U_2)$ lifts to an element of $\mathcal{Z}(\mathbb{S}^1)$ if and only if $\rho_{U_1, U_{1,2}}(a) = \rho_{U_2, U_{1,2}}(b)$. Algebraically this can be described as follows. Consider the complex

$$0 \longrightarrow \mathcal{Z}(\mathbb{S}^1) \longrightarrow \mathcal{Z}(U_1) \times \mathcal{Z}(U_2) \xrightarrow{d^0} \mathcal{Z}(U_{1,2}) \longrightarrow 0,$$

where $d^0 = \rho_{U_1, U_{1,2}} - \rho_{U_2, U_{1,2}}$.

The negative sign on the second component of d^0 ensures that we have a complex, i.e., that the composition of consecutive maps is zero. The discussion above reveals that the complex is exact on the left and in the middle, and has a free group of dimension one as cohomology on the right, representing the existence of nontrivial \mathbb{Z} -bundles on \mathbb{S}^1 .

It is reasonable to ask what would happen if we covered \mathbb{S}^1 with more than two open sets.

Exercise 2.8. Cover the circle \mathbb{S}^1 with the complements of the third roots of unity; call these U_i with $1 \leq i \leq 3$. The intersections $U_{i,j} = U_i \cap U_j$ are homeomorphic to disjoint pairs of intervals, and the triple intersection $U_{1,2,3}$ is the complement of the roots, hence is homeomorphic to three disjoint intervals. This gives us a diagram of restriction maps:

$$\begin{array}{ccccccc}
 & & \mathcal{Z}(U_1) & \longrightarrow & \mathcal{Z}(U_{1,2}) & & \\
 & \nearrow & & & \searrow & & \\
 \mathcal{Z}(\mathbb{S}^1) & \longrightarrow & \mathcal{Z}(U_2) & & \mathcal{Z}(U_{1,3}) & \longrightarrow & \mathcal{Z}(U_{1,2,3}) \\
 & \searrow & & & \nearrow & & \\
 & & \mathcal{Z}(U_3) & \longrightarrow & \mathcal{Z}(U_{2,3}) & &
 \end{array}$$

which can be used to construct a complex

$$0 \longrightarrow \mathcal{Z}(\mathbb{S}^1) \longrightarrow \prod_{i=1}^3 \mathcal{Z}(U_i) \xrightarrow{d^0} \prod_{1 \leq i < j \leq 3} \mathcal{Z}(U_{i,j}) \xrightarrow{d^1} \mathcal{Z}(U_{1,2,3}) \longrightarrow 0,$$

where the maps d^i are built using the restriction maps, suitably signed to ensure that compositions of maps are zero. Find such a suitable choice of signs, or look ahead at Definition 2.9. Determine the various groups and the matrix representations of the homomorphisms that appear in this complex. Show that the only nonzero cohomology group is $\ker d^1 / \text{image } d^0$, which is isomorphic to \mathbb{Z} .

The preceding examples are intended to suggest that there is a certain invariance to the computations induced by open covers. This is indeed the case, as is discussed later, but we first clarify the algebraic setup.

Definition 2.9. Let \mathcal{F} be a sheaf on a topological space X , and let $\mathfrak{U} = \{U_i\}_{i \in I}$ be an open cover of X , where I is a totally ordered index set. Given a finite subset J of I , let $U_J = \bigcap_{j \in J} U_j$. We construct a complex $\check{C}^\bullet(\mathfrak{U}; \mathcal{F})$ with t -th term

$$\check{C}^t(\mathfrak{U}; \mathcal{F}) = \prod_{|J|=t+1} \mathcal{F}(U_J), \quad \text{where } t \geq 0.$$

Let J be a subset of I with $|J| = t+1$, and let $i \in I \setminus J$. We define $\text{sgn}(J, i)$ to be -1 raised to the number of elements of J that are greater than i . Set

$$d^t: \check{C}^t(\mathfrak{U}; \mathcal{F}) \longrightarrow \check{C}^{t+1}(\mathfrak{U}; \mathcal{F})$$

to be the product of the maps

$$\text{sgn}(J, i) \cdot \rho_{U_J, U_{J \cup \{i\}}}: \mathcal{F}(U_J) \longrightarrow \mathcal{F}(U_{J \cup \{i\}}).$$

The sign choices ensure that $d^{t+1} \circ d^t = 0$, giving us a complex

$$0 \longrightarrow \check{C}^0(\mathfrak{U}; \mathcal{F}) \longrightarrow \check{C}^1(\mathfrak{U}; \mathcal{F}) \longrightarrow \check{C}^2(\mathfrak{U}; \mathcal{F}) \longrightarrow \dots$$

called the *Čech complex* associated to \mathcal{F} and \mathfrak{U} . The cohomology of $\check{C}^\bullet(\mathfrak{U}; \mathcal{F})$ is *Čech cohomology*, $\check{H}^\bullet(\mathfrak{U}; \mathcal{F})$.

Note that $\mathcal{F}(X)$ is not part of the Čech complex. However, as the following exercise shows, $\mathcal{F}(X)$ can be read from $\check{C}^\bullet(\mathfrak{U}; \mathcal{F})$.

Exercise 2.10. Let \mathcal{F} be a sheaf on X , and \mathfrak{U} an open cover of X . Prove that $\check{H}^0(\mathfrak{U}; \mathcal{F}) = \mathcal{F}(X)$.

Next we discuss the dependence of Čech cohomology on covers.

Definition 2.11. An open cover $\mathfrak{V} = \{V_i\}_{i \in I}$ of X is a *refinement* of an open cover $\mathfrak{U} = \{U_i\}_{i \in I'}$ if there is a map $\tau: I \rightarrow I'$ with $V_i \subseteq U_{\tau(i)}$ for each $i \in I$. The map τ need not be injective or surjective.

If \mathfrak{V} refines \mathfrak{U} , one has restriction maps $\mathcal{F}(U_{\tau(i)}) \rightarrow \mathcal{F}(V_i)$. More generally, for each finite subset $J \subseteq I$, there is a map $\mathcal{F}(U_{J'}) \rightarrow \mathcal{F}(V_J)$ where $J' = \{\tau(j) \mid j \in J\}$. These give a morphism of complexes

$$\check{\tau}: \check{C}^\bullet(\mathfrak{U}; \mathcal{F}) \longrightarrow \check{C}^\bullet(\mathfrak{V}; \mathcal{F}).$$

This map depends on τ , but the induced map on cohomology does not. In fact, the morphism $\check{\tau}$ is independent up to homotopy of the choice of τ .

Consider the class of open covers of X with partial order $\mathfrak{V} \geq \mathfrak{U}$ if \mathfrak{V} refines \mathfrak{U} . The class of covers may not be a set, but one can get around this problem; see [155, pp. 142/143] or [47, §5.8]. One can form the direct limit of the corresponding complexes. This is filtered, so Exercise 4.34 implies

$$H^t\left(\varinjlim_{\mathfrak{U}} \check{C}^\bullet(\mathfrak{U}; \mathcal{F})\right) \cong \varinjlim_{\mathfrak{U}} \check{H}^t(\mathfrak{U}; \mathcal{F}).$$

Definition 2.12. Set $\check{H}^t(X; \mathcal{F}) = \varinjlim_{\mathfrak{U}} \check{H}^t(\mathfrak{U}; \mathcal{F})$; this is called the *t-th Čech cohomology group* of \mathcal{F} .

Each open cover \mathfrak{U} gives rise to a natural map

$$\check{H}^t(\mathfrak{U}; \mathcal{F}) \longrightarrow \check{H}^t(X; \mathcal{F}).$$

The next result gives conditions under which this map is an isomorphism.

Theorem 2.13. Let \mathcal{F} be a sheaf on X . Suppose X has a base \mathfrak{V} for its topology that is closed under finite intersections and such that $\check{H}^t(V; \mathcal{F}) = 0$ for each $V \in \mathfrak{V}$ and $t > 0$.

If \mathfrak{U} is an open cover of X such that $\check{H}^t(U; \mathcal{F}) = 0$ for each $t > 0$ and U is a finite intersection of elements of \mathfrak{U} , then

$$\check{H}^t(\mathfrak{U}; \mathcal{F}) \cong \check{H}^t(X; \mathcal{F}) \quad \text{for all } t \geq 0. \quad \square$$

The theorem follows from Théorèmes 5.4.1 and 5.9.2 in [47]; see also Theorem 2.26. Using this theorem seems difficult as the hypotheses involve \mathfrak{U} and an auxiliary cover \mathfrak{V} . There are important situations where one has a cover \mathfrak{V} with the desired properties. If \mathcal{F} is a locally constant sheaf on a real manifold, for \mathfrak{V} one can take any base such that each V is homeomorphic to a finite number of disjoint copies of \mathbb{R}^n . For \mathcal{F} a quasi-coherent sheaf on a suitable scheme, one can take \mathfrak{V} to be any affine open cover.

We close with a few remarks about cohomology of constant sheaves.

Remark 2.14. Suppose that

$$0 \longrightarrow A' \longrightarrow A \longrightarrow A'' \longrightarrow 0$$

is an exact sequence of Abelian groups. Consider the induced constant sheaves $\mathcal{A}', \mathcal{A}, \mathcal{A}''$ on a space X . Suppose \mathfrak{U} is an open cover of X such that for any finite intersection U_I of sets in \mathfrak{U} , we get an exact sequence

$$0 \longrightarrow C(U_I, A') \longrightarrow C(U_I, A) \longrightarrow C(U_I, A'') \longrightarrow 0.$$

One then obtains an exact sequence of Čech complexes, which, in turn, yields an exact sequence of cohomology groups

$$\dots \longrightarrow \check{H}^t(\mathfrak{U}; \mathcal{A}') \longrightarrow \check{H}^t(\mathfrak{U}; \mathcal{A}) \longrightarrow \check{H}^t(\mathfrak{U}; \mathcal{A}'') \longrightarrow \check{H}^{t+1}(\mathfrak{U}; \mathcal{A}') \longrightarrow \dots$$

Not all open covers of X produce long exact sequences: consider Example 2.22 with the open cover $\mathfrak{U} = \{X\}$. In other words, $\check{H}^\bullet(\mathfrak{U}; -)$ need not be a δ -functor. On the other hand, $\check{H}^\bullet(X; -)$ is a δ -functor in favorable cases, such as on paracompact spaces, or for quasi-coherent sheaves on schemes.

Let \mathbb{K} be a field of characteristic zero endowed with the discrete topology, and let \mathcal{K} be the sheaf on \mathbb{S}^1 that assigns to each open U the set of continuous functions from U to \mathbb{K} . Then

$$\check{H}^\bullet(\mathbb{S}^1; \mathcal{K}) \cong \check{H}^\bullet(\mathbb{S}^1; \mathcal{Z}) \otimes_{\mathbb{Z}} \mathbb{K},$$

since tensoring with a field of characteristic zero is an exact functor.

The situation is different if \mathbb{K} has positive characteristic.

Exercise 2.15. Recall that the real projective plane \mathbb{RP}^2 is obtained by identifying antipodal points on the 2-sphere \mathbb{S}^2 . Cover \mathbb{RP}^2 with three open hemispheres whose pairwise and triple intersections are unions of disjoint contractible sets. Use this cover to prove that $\check{H}^1(\mathbb{RP}^2; \mathcal{Z}) = 0$.

Let \mathcal{Z}_2 be the sheaf for which $\mathcal{Z}_2(U)$ is the set of continuous functions from U to $\mathbb{Z}/2\mathbb{Z}$, where $\mathbb{Z}/2\mathbb{Z}$ has the discrete topology. Prove that $\check{H}^1(\mathbb{RP}^2; \mathcal{Z}_2) \cong \mathbb{Z}/2\mathbb{Z}$. In particular, the map $\check{H}^1(\mathbb{RP}^2; \mathcal{Z}) \longrightarrow \check{H}^1(\mathbb{RP}^2; \mathcal{Z}_2)$ induced by the natural surjection $\mathbb{Z} \longrightarrow \mathbb{Z}/2\mathbb{Z}$ is not surjective. This reflects the fact that \mathbb{RP}^2 is non-orientable.

3. Calculus versus topology

We now switch gears and investigate what happens if instead of making the cover increasingly fine, we replace \mathcal{F} by a resolution of more “flexible” sheaves. Specifically, let \mathcal{D} be the sheaf on \mathbb{S}^1 that attaches to each open set U , the ring of real-valued smooth functions. In this case, $\mathcal{D}(U)$ is a vector space over \mathbb{R} .

Example 2.16. Consider the cover of \mathbb{S}^1 by the open sets $U_1 = \mathbb{S}^1 \setminus \{-1\}$ and $U_2 = \mathbb{S}^1 \setminus \{1\}$. As before, we can construct a complex of the form

$$0 \longrightarrow \mathcal{D}(\mathbb{S}^1) \longrightarrow \mathcal{D}(U_1) \times \mathcal{D}(U_2) \longrightarrow \mathcal{D}(U_{1,2}) \longrightarrow 0.$$

If $(f_1, f_2) \in \mathcal{D}(U_1) \times \mathcal{D}(U_2)$ maps to zero in the above complex, then the functions f_1 and f_2 agree on $U_{1,2}$. It follows that f_1 and f_2 have (at worst) a removable singularity at -1 and 1 respectively, since f_2 has no singularity at -1 and f_1 has no singularity at 1 , and they are smooth outside these points. In particular, there is a function $f \in \mathcal{D}(\mathbb{S}^1)$ such that each f_i is the restriction of f to U_i , as predicted by Exercise 2.10.

Now consider the cohomology on the right. A function $f \in \mathcal{D}(U_{1,2})$ has no singularities except possibly at 1 and -1 . Let u be a smooth function on $U_{1,2}$ that takes value 0 near -1 , and value 1 near 1 . Of course, $f = (1 - u)f + uf$. Now $(1 - u)f$ can be extended to a smooth function on U_1 , and uf can be extended to a smooth function on U_2 . It follows that f is in the image of $\mathcal{D}(U_1) \times \mathcal{D}(U_2) \longrightarrow \mathcal{D}(U_{1,2})$. Hence the complex

$$0 \longrightarrow \mathcal{D}(U_1) \times \mathcal{D}(U_2) \longrightarrow \mathcal{D}(U_{1,2}) \longrightarrow 0$$

has cohomology groups $\check{H}^0(\mathfrak{U}; \mathcal{D}) = \mathcal{D}(\mathbb{S}^1)$ and $\check{H}^1(\mathfrak{U}; \mathcal{D}) = 0$.

We remark that any other open cover of \mathbb{S}^1 would give us a Čech complex with a unique nonzero cohomology group in degree zero. In terms of the limit over all open covers, this says that the sheaf \mathcal{D} on \mathbb{S}^1 has $\check{H}^0(\mathbb{S}^1; \mathcal{D}) = \mathcal{D}(\mathbb{S}^1)$ and $\check{H}^1(\mathbb{S}^1; \mathcal{D}) = 0$. The key point is the existence of partitions of unity:

Definition 2.17. Let M be a smooth manifold and $\mathfrak{U} = \{U_i\}_{i \in I}$ a locally finite open cover of M , i.e., any point $m \in M$ belongs to only finitely many of the open sets U_i . A *partition of unity subordinate to \mathfrak{U}* is a collection of smooth functions $f_i: M \longrightarrow \mathbb{R}$ for $i \in I$, such that $f_i|_{M \setminus U_i} = 0$ and $\sum f_i = 1$. Note that this is a finite sum at any point of M .

The partition in our case was $u + (1 - u) = 1$, and its existence allowed a section of \mathcal{D} on $U_{1,2}$ to be expressed as a sum of sections over U_1 and U_2 . Partitions of unity make the sheaf \mathcal{D} sufficiently fine and flexible, so that $\check{H}^t(M; \mathcal{D}) = 0$ for all $t \geq 1$.

Next, let \mathcal{R} be the sheaf that sends an open set $U \subseteq X$ to the continuous functions $C(U, \mathbb{R})$, where \mathbb{R} is endowed with the discrete topology. This is

the constant sheaf associated to \mathbb{R} ; see Example 2.7. The elements of $\mathcal{R}(U)$ are smooth functions, hence \mathcal{R} may be viewed as a “subsheaf” of \mathcal{D} ; you can make the meaning of this precise once you have seen Definition 2.19.

Example 2.18. Let U be a proper open subset of \mathbb{S}^1 . We claim there is an exact sequence

$$0 \longrightarrow \mathcal{R}(U) \longrightarrow \mathcal{D}(U) \xrightarrow{\frac{d}{dt}} \mathcal{D}(U) \longrightarrow 0,$$

where the first map is inclusion, while the second map is differentiation by arclength. To see that this sequence is exact, note that U is the disjoint union of open arcs that are diffeomorphic to the real line. On such an open arc, (i) the constants are the only functions that are annihilated by differentiation, and (ii) every smooth function can be integrated to a smooth function. On the other hand, as we will discuss later, the sequence is not exact on the right if U is \mathbb{S}^1 .

For open sets $V \subseteq U$, sequences of the form above, along with appropriate restriction maps, fit together to give a commutative diagram

$$\begin{array}{ccccc} \mathcal{R}(U) & \longrightarrow & \mathcal{D}(U) & \xrightarrow{\frac{d}{dt}} & \mathcal{D}(U) \\ \downarrow & & \downarrow & & \downarrow \\ \mathcal{R}(V) & \longrightarrow & \mathcal{D}(V) & \xrightarrow{\frac{d}{dt}} & \mathcal{D}(V). \end{array}$$

Definition 2.19. A *morphism of sheaves* $\varphi: \mathcal{F} \longrightarrow \mathcal{G}$ on X is a collection of maps $\varphi_U: \mathcal{F}(U) \longrightarrow \mathcal{G}(U)$ for $U \in \mathcal{T}_X$, such that for each inclusion $V \subseteq U$ of open sets, there is a commutative diagram

$$\begin{array}{ccc} \mathcal{F}(U) & \xrightarrow{\varphi_U} & \mathcal{G}(U) \\ \rho_{U,V} \downarrow & & \downarrow \rho'_{U,V} \\ \mathcal{F}(V) & \xrightarrow{\varphi_V} & \mathcal{G}(V) \end{array}$$

where ρ and ρ' are the restriction maps for the sheaves \mathcal{F} and \mathcal{G} respectively.

If \mathcal{F} and \mathcal{G} are sheaves of Abelian groups, rings, etc., a morphism of such sheaves is a collection as above with the additional property that the maps φ_U are morphisms in the appropriate category.

A natural next step would be to discuss the Abelian category structure on the category of sheaves. This then leads to a notion of a cohomology of a complex of sheaves, and hence to exactness. We take a shorter route.

Definition 2.20. Let \mathcal{F} be a sheaf on a space X , and consider a point $x \in X$. The class of open sets U of X containing x forms a filtered direct

system, and its limit

$$\mathcal{F}_x = \varinjlim_{U \ni x} \mathcal{F}(U)$$

is the *stalk* of \mathcal{F} at x . Any additional structure \mathcal{F} has is usually inherited by the stalks; for instance, if \mathcal{F} is a sheaf of Abelian groups or rings, then \mathcal{F}_x is an Abelian group or a ring, respectively.

It follows from standard properties of direct limits that a morphism $\varphi: \mathcal{F} \rightarrow \mathcal{G}$ of sheaves on X induces, for each point $x \in X$, a morphism $\varphi_x: \mathcal{F}_x \rightarrow \mathcal{G}_x$ in the appropriate category; passage to a stalk is a functor.

Definition 2.21. A *complex* \mathcal{F}^\bullet of sheaves of Abelian groups is a sequence

$$\dots \longrightarrow \mathcal{F}^{t-1} \xrightarrow{d^{t-1}} \mathcal{F}^t \xrightarrow{d^t} \mathcal{F}^{t+1} \xrightarrow{d^{t+1}} \dots$$

of morphisms of sheaves where $d^{t+1} \circ d^t = 0$ for all t . Such a complex \mathcal{F}^\bullet is *exact* if the induced complex of Abelian groups

$$\dots \longrightarrow \mathcal{F}_x^{t-1} \longrightarrow \mathcal{F}_x^t \longrightarrow \mathcal{F}_x^{t+1} \longrightarrow \dots$$

is exact for each $x \in X$.

For instance, if X has a base \mathfrak{U} of open sets such that for each U in \mathfrak{U} , the sequence of Abelian groups

$$\dots \longrightarrow \mathcal{F}^{t-1}(U) \longrightarrow \mathcal{F}^t(U) \longrightarrow \mathcal{F}^{t+1}(U) \longrightarrow \dots$$

is exact, then the complex of sheaves is exact; this is because computing direct limits commutes with cohomology; see Exercise 4.34. The converse is certainly not true; see the example below.

Example 2.22. It follows from the preceding discussion that the following sequence of sheaves from Example 2.18 is exact:

$$0 \longrightarrow \mathcal{R} \longrightarrow \mathcal{D} \xrightarrow{\frac{d}{dt}} \mathcal{D} \longrightarrow 0.$$

Taking global sections, we get a sequence

$$0 \longrightarrow \mathcal{R}(\mathbb{S}^1) \longrightarrow \mathcal{D}(\mathbb{S}^1) \xrightarrow{\frac{d}{dt}} \mathcal{D}(\mathbb{S}^1) \longrightarrow 0$$

which is exact at the left and in the middle, and we shall see that it is not exact on the right. If a smooth function on \mathbb{S}^1 arises as a derivative, then the fundamental theorem of calculus implies that its average value is 0. Consequently, nonzero constant functions on \mathbb{S}^1 do not arise as derivatives. On the other hand, for any smooth function $f: \mathbb{S}^1 \rightarrow \mathbb{R}$, the function

$$g(t) = f(t) - \frac{1}{2\pi} \int_0^{2\pi} f(u) du$$

integrates to 0 on \mathbb{S}^1 and hence $\int_0^t g(\tau) d\tau$ is a function on \mathbb{S}^1 whose derivative is $g(t)$. It follows that the complex of global sections

$$0 \longrightarrow \mathcal{D}(\mathbb{S}^1) \longrightarrow \mathcal{D}(\mathbb{S}^1) \longrightarrow 0$$

has cohomology in degrees 0 and 1, and that both cohomology groups are isomorphic to the space of constant functions on \mathbb{S}^1 .

This discussion leads to another type of cohomology theory.

Definition 2.23. A sheaf \mathcal{E} of Abelian groups is *injective* if every injective morphism $\mathcal{E} \rightarrow \mathcal{F}$ of sheaves of Abelian groups splits. Every sheaf \mathcal{F} of Abelian groups can be embedded into an injective sheaf; see Remark 12.29.

A complex \mathcal{G}^\bullet of sheaves is a *resolution* of \mathcal{F} if there is an exact sequence

$$0 \longrightarrow \mathcal{F} \longrightarrow \mathcal{G}^0 \longrightarrow \mathcal{G}^1 \longrightarrow \mathcal{G}^2 \longrightarrow \dots .$$

It is an *injective* resolution if each \mathcal{G}^t is injective. The *sheaf cohomology* $H^t(X, \mathcal{F})$ is the t -th cohomology group of the complex

$$0 \longrightarrow \mathcal{G}^0(X) \longrightarrow \mathcal{G}^1(X) \longrightarrow \mathcal{G}^2(X) \longrightarrow \dots .$$

The cohomology groups do not depend on \mathcal{G}^\bullet ; see Lectures 3 and 12.

Definition 2.24. A sheaf \mathcal{F} is *acyclic* if $H^t(X, \mathcal{F}) = 0$ for all $t \geq 1$.

For the proof of the following proposition, see [61, Proposition III.1.2A].

Proposition 2.25 (Acyclicity principle). *If \mathcal{G}^\bullet is a resolution of \mathcal{F} by acyclic sheaves, then $H^t(X, \mathcal{F}) = H^t(\mathcal{G}^\bullet(X))$ for each t .* \square

In Example 2.16 we verified that \mathcal{D} is acyclic on \mathbb{S}^1 . As noted in Example 2.22, the constant sheaf \mathcal{R} on \mathbb{S}^1 has an acyclic resolution

$$0 \longrightarrow \mathcal{D} \xrightarrow{\frac{d}{dt}} \mathcal{D} \longrightarrow 0,$$

so $H^t(\mathbb{S}^1, \mathcal{R}) = 0$ for $t \geq 2$. Moreover, $H^0(\mathbb{S}^1, \mathcal{R}) \cong \mathbb{R} \cong H^1(\mathbb{S}^1, \mathcal{R})$, and these agree with the singular cohomology groups of \mathbb{S}^1 .

4. Čech cohomology and derived functors

We now compare the two approaches we have taken. The constant sheaf \mathcal{R} on \mathbb{S}^1 has Čech cohomology groups $\check{H}^t(\mathbb{S}^1; \mathcal{R}) = 0$ if $t \geq 2$, and

$$\check{H}^0(\mathbb{S}^1; \mathcal{R}) \cong \mathbb{R} \cong \check{H}^1(\mathbb{S}^1; \mathcal{R}).$$

The sheaf cohomology groups turned out to be the same as these. We take an open cover \mathcal{U} of \mathbb{S}^1 , fine enough so that each finite intersection of the open sets is an open arc. By the discussion following Theorem 2.13, one has $\check{H}^\bullet(\mathcal{U}; \mathcal{R}) = \check{H}^\bullet(\mathbb{S}^1; \mathcal{R})$ and $\check{H}^\bullet(\mathcal{U}; \mathcal{D}) = \check{H}^\bullet(\mathbb{S}^1; \mathcal{D})$. Moreover, for each U_I , there is an exact sequence

$$0 \longrightarrow \mathcal{R}(U_I) \longrightarrow \mathcal{D}(U_I) \longrightarrow \mathcal{D}(U_I) \longrightarrow 0,$$

and hence an exact sequence of complexes

$$0 \longrightarrow \check{C}^\bullet(\mathcal{U}; \mathcal{R}) \longrightarrow \check{C}^\bullet(\mathcal{U}; \mathcal{D}) \longrightarrow \check{C}^\bullet(\mathcal{U}; \mathcal{D}) \longrightarrow 0.$$

By Remark 2.14, there exists an exact sequence of cohomology groups

$$0 \longrightarrow \check{H}^0(\mathbb{S}^1; \mathcal{R}) \longrightarrow \check{H}^0(\mathbb{S}^1; \mathcal{D}) \longrightarrow \check{H}^0(\mathbb{S}^1; \mathcal{D}) \longrightarrow \check{H}^1(\mathbb{S}^1; \mathcal{R}) \longrightarrow 0,$$

where the zero on the right follows from the fact that $\check{H}^t(\mathbb{S}^1; \mathcal{D}) = 0$ for all $t \geq 1$; see Example 2.16. We conclude that $\check{H}^t(\mathbb{S}^1; \mathcal{R}) = 0$ for all $t \geq 2$, and that $\check{H}^0(\mathbb{S}^1; \mathcal{R})$ and $\check{H}^1(\mathbb{S}^1; \mathcal{R})$ arise naturally as the kernel and cokernel respectively of the differentiation map

$$\frac{d}{dt}: \mathcal{D}(\mathbb{S}^1) \longrightarrow \mathcal{D}(\mathbb{S}^1).$$

In particular, the Čech cohomology of \mathcal{R} on \mathbb{S}^1 can be read off from the global sections of the morphism $\mathcal{D} \longrightarrow \mathcal{D}$ given by differentiation.

In Lecture 19 we revisit this theme of linking differential calculus with sheaves and topology. The main ideas are as follows: one can get topological information from the Čech approach, since fine open covers turn the computation into a triangulation of the underlying space. On the other hand, one can replace the given sheaf by a suitable complex of acyclic sheaves and consider the cohomology of the resulting complex of global sections.

The following statement summarizes the relationship between Čech cohomology and sheaf cohomology; see [47, Chapitre 5] for proofs.

Theorem 2.26. *Let X , \mathcal{F} , \mathfrak{U} , and \mathfrak{V} be as in Theorem 2.13. Then there are isomorphisms*

$$\check{H}^t(\mathfrak{U}; \mathcal{F}) \cong \check{H}^t(X; \mathcal{F}) \cong H^t(X, \mathcal{F}) \quad \text{for each } t. \quad \square$$

Connectedness

This lecture deals with connections between cohomological dimension and connectedness of varieties. An important ingredient is a local cohomology version of the Mayer-Vietoris theorem encountered in topology.

1. Mayer-Vietoris sequence

Let $\mathfrak{a}' \supseteq \mathfrak{a}$ be ideals in R and M an R -module. The inclusion $\Gamma_{\mathfrak{a}'}(-) \subseteq \Gamma_{\mathfrak{a}}(-)$ induces, for each n , an R -module homomorphism

$$\theta_{\mathfrak{a}',\mathfrak{a}}^n(M): H_{\mathfrak{a}'}^n(M) \longrightarrow H_{\mathfrak{a}}^n(M).$$

This homomorphism is functorial in M , that is to say, $\theta_{\mathfrak{a}',\mathfrak{a}}^n(-)$ is a natural transformation from $H_{\mathfrak{a}'}^n(-)$ to $H_{\mathfrak{a}}^n(-)$. Given ideals \mathfrak{a} and \mathfrak{b} in R , for each integer n we set

$$\begin{aligned} \iota_{\mathfrak{a},\mathfrak{b}}^n(M): H_{\mathfrak{a}+\mathfrak{b}}^n(M) &\longrightarrow H_{\mathfrak{a}}^n(M) \oplus H_{\mathfrak{b}}^n(M), \\ z &\longmapsto (\theta_{\mathfrak{a}+\mathfrak{b},\mathfrak{a}}^n(z), \theta_{\mathfrak{a}+\mathfrak{b},\mathfrak{b}}^n(z)), \quad \text{and} \\ \pi_{\mathfrak{a},\mathfrak{b}}(M): H_{\mathfrak{a}}^n(M) \oplus H_{\mathfrak{b}}^n(M) &\longrightarrow H_{\mathfrak{a}\cap\mathfrak{b}}^n(M), \\ (x, y) &\longmapsto \theta_{\mathfrak{a},\mathfrak{a}\cap\mathfrak{b}}^n(x) - \theta_{\mathfrak{b},\mathfrak{a}\cap\mathfrak{b}}^n(y). \end{aligned}$$

Clearly the homomorphisms $\iota_{\mathfrak{a},\mathfrak{b}}^n(M)$ and $\pi_{\mathfrak{a},\mathfrak{b}}(M)$ are functorial in M .

Theorem 15.1 (Mayer-Vietoris sequence). *Let \mathfrak{a} and \mathfrak{b} be ideals in a Noetherian ring R . For each R -module M , there exists a sequence of R -modules*

$$\begin{aligned} 0 \longrightarrow H_{\mathfrak{a}+\mathfrak{b}}^0(M) &\xrightarrow{\iota_{\mathfrak{a},\mathfrak{b}}^0(M)} H_{\mathfrak{a}}^0(M) \oplus H_{\mathfrak{b}}^0(M) \xrightarrow{\pi_{\mathfrak{a},\mathfrak{b}}^0(M)} H_{\mathfrak{a}\cap\mathfrak{b}}^0(M) \\ &\longrightarrow H_{\mathfrak{a}+\mathfrak{b}}^1(M) \xrightarrow{\iota_{\mathfrak{a},\mathfrak{b}}^1(M)} H_{\mathfrak{a}}^1(M) \oplus H_{\mathfrak{b}}^1(M) \xrightarrow{\pi_{\mathfrak{a},\mathfrak{b}}^1(M)} H_{\mathfrak{a}\cap\mathfrak{b}}^1(M) \longrightarrow \dots, \end{aligned}$$

which is exact, and functorial in M .

Proof. It is an elementary exercise to verify that one has an exact sequence

$$0 \longrightarrow \Gamma_{\mathfrak{a}+\mathfrak{b}}(M) \xrightarrow{\iota_{\mathfrak{a},\mathfrak{b}}^0(M)} \Gamma_{\mathfrak{a}}(M) \oplus \Gamma_{\mathfrak{b}}(M) \xrightarrow{\pi_{\mathfrak{a},\mathfrak{b}}^0(M)} \Gamma_{\mathfrak{a} \cap \mathfrak{b}}(M).$$

We claim that $\pi_{\mathfrak{a},\mathfrak{b}}^0(M)$ is surjective whenever M is an injective R -module. Indeed, since $\Gamma_{\mathfrak{a}'}(-)$ commutes with direct sums for any ideal \mathfrak{a}' , it suffices to consider the case where $M = E_R(R/\mathfrak{p})$ for a prime ideal \mathfrak{p} of R . Then the asserted surjectivity is immediate from Example 7.6.

Let I^\bullet be an injective resolution of M . In view of the preceding discussion, we have an exact sequence of complexes,

$$0 \longrightarrow \Gamma_{\mathfrak{a}+\mathfrak{b}}(I^\bullet) \longrightarrow \Gamma_{\mathfrak{a}}(I^\bullet) \oplus \Gamma_{\mathfrak{b}}(I^\bullet) \longrightarrow \Gamma_{\mathfrak{a} \cap \mathfrak{b}}(I^\bullet) \longrightarrow 0.$$

The homology exact sequence arising from this is the one we seek. The functoriality is a consequence of the functoriality of $\iota_{\mathfrak{a},\mathfrak{b}}^n(-)$ and $\pi_{\mathfrak{a},\mathfrak{b}}^n(-)$ and of the connecting homomorphisms in cohomology long exact sequences. \square

Recall that a topological space is *connected* if it cannot be written as a disjoint union of two proper closed subsets. The Mayer-Vietoris sequence has applications to connectedness properties of algebraic varieties.

Exercise 15.2. Let R be a Noetherian ring. Prove that the support of a finitely generated indecomposable R -module is connected.

2. Punctured spectra

We focus now on punctured spectra of local rings.

Definition 15.3. The *punctured spectrum* of a local ring (R, \mathfrak{m}) is the set

$$\text{Spec}^\circ R = \text{Spec } R \setminus \{\mathfrak{m}\},$$

with topology induced by the Zariski topology on $\text{Spec } R$. Similarly, if R is graded with homogeneous maximal ideal \mathfrak{m} , then its punctured spectrum refers to the topological space $\text{Spec } R \setminus \{\mathfrak{m}\}$.

Connectedness of the punctured spectrum can be interpreted entirely in the language of ideals.

Remark 15.4. Let (R, \mathfrak{m}) be a local ring and \mathfrak{a} an ideal. The punctured spectrum $\text{Spec}^\circ(R/\mathfrak{a})$ of the local ring R/\mathfrak{a} is connected if and only if the following property holds: given ideals \mathfrak{a}' and \mathfrak{a}'' in R with

$$\text{rad}(\mathfrak{a}' \cap \mathfrak{a}'') = \text{rad } \mathfrak{a} \quad \text{and} \quad \text{rad}(\mathfrak{a}' + \mathfrak{a}'') = \mathfrak{m},$$

either $\text{rad } \mathfrak{a}'$ or $\text{rad } \mathfrak{a}''$ equals \mathfrak{m} ; equivalently, $\text{rad } \mathfrak{a}''$ or $\text{rad } \mathfrak{a}'$ equals $\text{rad } \mathfrak{a}$.

Indeed, this is a direct translation of the definition of connectedness, keeping in mind that

$$V(\mathfrak{a}') \cup V(\mathfrak{a}'') = V(\mathfrak{a}' \cap \mathfrak{a}'') \quad \text{and} \quad V(\mathfrak{a}') \cap V(\mathfrak{a}'') = V(\mathfrak{a}' + \mathfrak{a}'').$$

Exercise 15.5. Prove that if R is a local domain, then $\text{Spec}^\circ R$ is connected.

Exercise 15.6. Let $R = \mathbb{R}[x, y, ix, iy]$ where $i^2 = -1$. Note that

$$R \cong \mathbb{R}[x, y, u, v]/(u^2 + x^2, v^2 + y^2, xy + uv, xv - uy).$$

Show that $\text{Spec}^\circ R$ is connected, but $\text{Spec}^\circ(R \otimes_{\mathbb{R}} \mathbb{C})$ is not.

The next few results identify conditions under which the punctured spectrum is connected. The first is a straightforward application of the Mayer-Vietoris sequence, Theorem 15.1.

Proposition 15.7. *If R is local with $\text{depth } R \geq 2$, then $\text{Spec}^\circ R$ is connected.*

Proof. Let \mathfrak{m} be the maximal ideal of R , and let \mathfrak{a}' and \mathfrak{a}'' be ideals with $\text{rad}(\mathfrak{a}' \cap \mathfrak{a}'') = \text{rad}(0)$ and $\text{rad}(\mathfrak{a}' + \mathfrak{a}'') = \mathfrak{m}$. Proposition 7.3(2) and the depth sensitivity of local cohomology, Theorem 9.1, imply that

$$H_{\mathfrak{a}' \cap \mathfrak{a}''}^0(R) = R \quad \text{and} \quad H_{\mathfrak{a}' + \mathfrak{a}''}^0(R) = 0 = H_{\mathfrak{a}' + \mathfrak{a}''}^1(R).$$

The Mayer-Vietoris sequence now yields an isomorphism of R -modules

$$H_{\mathfrak{a}'}^0(R) \oplus H_{\mathfrak{a}''}^0(R) \cong R.$$

But R is indecomposable as a module over itself by Exercise 15.8, so, without loss of generality, we may assume that $H_{\mathfrak{a}'}^0(R) = R$ and $H_{\mathfrak{a}''}^0(R) = 0$. This implies that $\text{rad } \mathfrak{a}' = \text{rad}(0)$, as desired. \square

Exercise 15.8. Show that a local ring is indecomposable as a module.

The hypothesis on depth in Proposition 15.7 is optimal in view of Exercise 15.6 as well as the following example:

Example 15.9. The ring $R = \mathbb{K}[[x, y]]/(xy)$ is local with $\text{depth } R = 1$. Moreover, one has that

$$V(x) \cup V(y) = \text{Spec } R \quad \text{and} \quad V(x) \cap V(y) = \{(x, y)\}.$$

Thus, $\text{Spec}^\circ R$ is not connected.

Here is an amusing application of Proposition 15.7:

Example 15.10. Let $R = \mathbb{K}[[x, y, u, v]]/(x, y) \cap (u, v)$. It is not hard to check that $\dim R = 2$. On the other hand, $\text{Spec}^\circ R$ is disconnected so $\text{depth } R \leq 1$ by Proposition 15.7. In particular, R is not Cohen-Macaulay.

More sophisticated results on connectedness of punctured spectra can be derived from the Hartshorne-Lichtenbaum vanishing theorem. The one below was proved in the equicharacteristic case by Faltings [36, 37] using a different method; the argument presented here is due to Brodmann and Rung [18]. Recall that $\text{ara } \mathfrak{a}$, the arithmetic rank of \mathfrak{a} , is the least number of generators of an ideal with the same radical as \mathfrak{a} .

Theorem 15.11 (Faltings' connectedness theorem). *Let R be a complete local domain. If \mathfrak{a} is an ideal of R with $\text{ara } \mathfrak{a} \leq \dim R - 2$, then $\text{Spec}^\circ(R/\mathfrak{a})$, the punctured spectrum of R/\mathfrak{a} , is connected.*

Proof. Let \mathfrak{m} be the maximal ideal of R and \mathfrak{a}' , \mathfrak{a}'' ideals with $\text{rad}(\mathfrak{a}' \cap \mathfrak{a}'') = \text{rad } \mathfrak{a}$ and $\text{rad}(\mathfrak{a}' + \mathfrak{a}'') = \mathfrak{m}$. We prove that one of $\text{rad } \mathfrak{a}'$ or $\text{rad } \mathfrak{a}''$ equals \mathfrak{m} .

Set $d = \dim R$. Since $\text{ara } \mathfrak{a} \leq d - 2$, we have

$$H_{\mathfrak{a}' \cap \mathfrak{a}''}^n(R) = H_{\mathfrak{a}}^n(R) = 0 \quad \text{for } n = d - 1, d,$$

where the first equality holds by Proposition 7.3(2) and the second by Proposition 9.12. Keeping in mind that $H_{\mathfrak{a}' + \mathfrak{a}''}^n(R) = H_{\mathfrak{m}}^n(R)$ for each n , again by Proposition 7.3(2), the Mayer-Vietoris sequence gives us

$$0 = H_{\mathfrak{a}' \cap \mathfrak{a}''}^{d-1}(R) \longrightarrow H_{\mathfrak{m}}^d(R) \longrightarrow H_{\mathfrak{a}'}^d(R) \oplus H_{\mathfrak{a}''}^d(R) \longrightarrow H_{\mathfrak{a}' \cap \mathfrak{a}''}^d(R) = 0.$$

Theorem 9.3 yields $H_{\mathfrak{m}}^d(R) \neq 0$, so the exact sequence above implies that one of $H_{\mathfrak{a}'}^d(R)$ or $H_{\mathfrak{a}''}^d(R)$ must be nonzero; say $H_{\mathfrak{a}'}^d(R) \neq 0$. This implies $\text{cd}_R(\mathfrak{a}') \geq d$, and hence $\text{rad } \mathfrak{a}' = \mathfrak{m}$ by Theorem 14.1. \square

Hochster and Huneke have obtained generalizations of Faltings' connectedness theorem. One such is [74, Theorem 3.3]:

Theorem 15.12. *Let (R, \mathfrak{m}) be a complete equidimensional ring of dimension d such that $H_{\mathfrak{m}}^d(R)$ is indecomposable as an R -module; equivalently, the canonical module ω_R is indecomposable.*

If \mathfrak{a} is an ideal of R with $\text{ara } \mathfrak{a} \leq d - 2$, then $\text{Spec}^\circ(R/\mathfrak{a})$ is connected. \square

Faltings' theorem leads to another connectedness result, originally due to Fulton and Hansen [43]. An interesting feature of the proof is the use of 'reduction to the diagonal' encountered in the proof of Theorem 1.33.

Theorem 15.13 (Fulton-Hansen theorem). *Let \mathbb{K} be an algebraically closed field, and let X and Y be irreducible sets in $\mathbb{P}_{\mathbb{K}}^n$.*

If $\dim X + \dim Y \geq n + 1$, then $X \cap Y$ is connected.

Proof. Let \mathfrak{p} and \mathfrak{q} be homogeneous prime ideals of $\mathbb{K}[x_0, \dots, x_n]$ such that

$$\mathbb{K}[x_0, \dots, x_n]/\mathfrak{p} \quad \text{and} \quad \mathbb{K}[x_0, \dots, x_n]/\mathfrak{q}$$

are homogeneous coordinate rings of X and Y , respectively. Then

$$\mathbb{K}[x_0, \dots, x_n]/(\mathfrak{p} + \mathfrak{q})$$

is a homogeneous coordinate ring for the intersection $X \cap Y$. If $X \cap Y$ is disconnected, then so are the punctured spectra of the graded rings

$$\frac{\mathbb{K}[x_0, \dots, x_n]}{(\mathfrak{p} + \mathfrak{q})} \cong \frac{\mathbb{K}[x_0, \dots, x_n, y_0, \dots, y_n]}{(\mathfrak{p} + \mathfrak{q}' + \Delta)},$$

where \mathfrak{q}' is the ideal generated by polynomials obtained by substituting y_i for x_i in a set of generators for \mathfrak{q} , and $\Delta = (x_0 - y_0, \dots, x_n - y_n)$ is the ideal defining the diagonal in $\mathbb{P}_{\mathbb{K}}^{2n+1}$. The complete local ring

$$R = \mathbb{K}[[x_0, \dots, x_n, y_0, \dots, y_n]]/(\mathfrak{p} + \mathfrak{q}')$$

is a domain by Exercises 15.15(4) and 15.17, and

$$\dim R = \dim X + 1 + \dim Y + 1 \geq n + 3.$$

Evidently $\dim \Delta \leq n + 1 = (n + 3) - 2$. This, however, contradicts Faltings' connectedness theorem. \square

Exercise 15.14. If \mathbb{K} is a field which is *not* algebraically closed, construct a \mathbb{K} -algebra R such that R is an integral domain, but $R \otimes_{\mathbb{K}} R$ is not.

Exercise 15.15. Let $\mathbb{K} \subseteq \mathbb{L}$ be fields, where \mathbb{K} is algebraically closed.

- (1) If a family of polynomials $\{f_i\}$ in $\mathbb{K}[x_1, \dots, x_n]$ has no common root in \mathbb{K}^n , prove that it has no common root in \mathbb{L}^n .
- (2) If $g \in \mathbb{K}[x_1, \dots, x_n]$ is irreducible, prove that it is also irreducible as an element of $\mathbb{L}[x_1, \dots, x_n]$. Hint: Use (1).
- (3) If R is a finitely generated \mathbb{K} -algebra which is an integral domain, prove that $R \otimes_{\mathbb{K}} \mathbb{L}$ is an integral domain.
- (4) If R and S are finitely generated \mathbb{K} -algebras which are integral domains, prove that $R \otimes_{\mathbb{K}} S$ is an integral domain. Hint: Take \mathbb{L} to be the fraction field of S and use (3).

Exercise 15.16. Give an example of a local domain (R, \mathfrak{m}) whose \mathfrak{m} -adic completion is not a domain.

Exercise 15.17. If (R, \mathfrak{m}) is an \mathbb{N} -graded domain finitely generated over a field R_0 , prove that its \mathfrak{m} -adic completion is also a domain.

D-modules

We introduce rings of differential operators and modules over these rings, i.e., *D*-modules. The focus is on Weyl algebras, with a view towards applications to local cohomology in Lecture 23.

1. Rings of differential operators

Let \mathbb{K} be a field and R a commutative \mathbb{K} -algebra.

The composition of elements P, Q of $\text{Hom}_{\mathbb{K}}(R, R)$ is denoted $P \cdot Q$. With this product, $\text{Hom}_{\mathbb{K}}(R, R)$ is a ring; it is even a \mathbb{K} -algebra since each P is \mathbb{K} -linear. The *commutator* of P and Q is the element

$$[P, Q] = P \cdot Q - Q \cdot P.$$

Since R is commutative, the map $r \mapsto (s \mapsto rs)$ gives an embedding of \mathbb{K} -algebras: $R \subseteq \text{Hom}_{\mathbb{K}}(R, R)$. Note that the natural left-module structure of R over $\text{Hom}_{\mathbb{K}}(R, R)$ extends the one of R over itself.

Definition 17.1. Set $D_0(R; \mathbb{K}) = R$ viewed as a subring of $\text{Hom}_{\mathbb{K}}(R, R)$, and for each $i \geq 0$, let

$$D_{i+1}(R; \mathbb{K}) = \{P \in \text{Hom}_{\mathbb{K}}(R, R) \mid [P, r] \in D_i(R; \mathbb{K}) \text{ for each } r \in R\}.$$

It is easy to verify that if P is in $D_i(R; \mathbb{K})$ and Q is in $D_j(R; \mathbb{K})$, then $P \cdot Q$ is in $D_{i+j}(R; \mathbb{K})$. Thus, one obtains a \mathbb{K} -subalgebra of $\text{Hom}_{\mathbb{K}}(R, R)$,

$$D(R; \mathbb{K}) = \bigcup_{i \geq 0} D_i(R; \mathbb{K}).$$

This is the ring of \mathbb{K} -linear *differential operators* on R . Elements of $D_i(R; \mathbb{K})$ are said to have *order* i . Note that $D_1(R; \mathbb{K})$ is the \mathbb{K} -span of R and the *derivations*, i.e., the maps δ with $\delta(rs) = \delta(r)s + r\delta(s)$ for $r, s \in R$.

The *tensor algebra* in y_1, \dots, y_m over \mathbb{K} is denoted $\mathbb{K}\langle y_1, \dots, y_m \rangle$.

Example 17.2. Let \mathbb{K} be a field of characteristic zero. Set $R = \mathbb{K}[x]$ and $\partial = \partial/\partial x$, the derivative with respect to x ; it is a derivation on R . The ring $D(R; \mathbb{K})$ is $\mathbb{K}\langle x, \partial \rangle$ modulo the two-sided ideal generated by $\partial \cdot x - x \cdot \partial - 1$.

As the example above shows, $D(R; \mathbb{K})$ need not be commutative.

Exercise 17.3. Find the rings of \mathbb{C} -linear differential operators on $\mathbb{C}[[x]]$ and on $\mathbb{C}[x, y]/(xy)$.

Definition 17.4. Set $R = \mathbb{K}[x_1, \dots, x_n]$, a polynomial ring in x_1, \dots, x_n . The n -th *Weyl algebra* over \mathbb{K} is the ring

$$D_n(\mathbb{K}) = \mathbb{K}\langle x_1, \dots, x_n, \partial_1, \dots, \partial_n \rangle / \mathfrak{a},$$

where \mathfrak{a} is the two-sided ideal generated by the elements

$$(17.4.1) \quad x_i \cdot x_j - x_j \cdot x_i, \quad \partial_i \cdot x_j - x_j \cdot \partial_i - \delta_{i,j}, \quad \partial_i \cdot \partial_j - \partial_j \cdot \partial_i,$$

with $\delta_{i,j}$ the Kronecker delta. Observe that R is a subring of $D_n(\mathbb{K})$.

Viewing the element $\partial_i \in D_n(\mathbb{K})$ as partial differentiation with respect to x_i , one can realize $D_n(\mathbb{K})$ as a subring of $D(R; \mathbb{K})$. The derivations are the R -linear combinations of $\partial_1, \dots, \partial_n$. If \mathbb{K} has characteristic zero, then $D_n(\mathbb{K})$ equals $D(R; \mathbb{K})$; see [54, Theorem 16.11.2] or [25, Theorem 2.3].

Exercise 17.5. When \mathbb{K} is of positive characteristic, say p , show that $D(R; \mathbb{K})$ is bigger than the Weyl algebra, $D_n(\mathbb{K})$.

Hint: Consider *divided powers* $\frac{\partial^p}{p! \partial x_i^p}$.

Exercise 17.6. Prove that the only two-sided ideals of the Weyl algebra $D_n(\mathbb{K})$ are 0 and $D_n(\mathbb{K})$ itself; i.e., $D_n(\mathbb{K})$ is simple.

A theorem of Stafford [149] implies that every ideal of $D_n(\mathbb{K})$ can be generated by two elements; see also [12, §1.7]. Algorithms for obtaining the two generators may be found in [66, 99].

Remark 17.7. Let \mathfrak{a} be an ideal in a commutative \mathbb{K} -algebra R , and set $S = R/\mathfrak{a}$. One can identify $D(S; \mathbb{K})$ with the subring of $D(R; \mathbb{K})$ consisting of operators that stabilize \mathfrak{a} , modulo the ideal generated by \mathfrak{a} .

For example, for $S = \mathbb{C}[x, y]/(xy)$, the ring $D(S; \mathbb{C})$ is the S -algebra generated by $x\partial_x^n, y\partial_y^n$ for $n \in \mathbb{N}$, modulo the ideal generated by xy .

Exercise 17.8. Let $S = \mathbb{C}[s^{\pm 1}, t^{\pm 1}]$. Prove that $D(S; \mathbb{C}) = S\langle \partial_s, \partial_t \rangle$.

Let $R = \mathbb{C}[s, st, st^2]$, which is the semigroup ring for the semigroup generated by the columns of the matrix

$$A = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \end{bmatrix}.$$

The ring $D(R; \mathbb{C})$ turns out to be the R -subalgebra of $D(S; \mathbb{C})$ generated by

$$I_{(p,q)}(\theta_s, \theta_t) \cdot s^p t^q, \quad p, q \in \mathbb{Z},$$

where $\theta_s = s \cdot \partial_s$, $\theta_t = t \cdot \partial_t$ and $I_{(p,q)}(\alpha, \beta)$ is the defining ideal of the Zariski closure of the set of points $((p, q) + \mathbb{N}A) \setminus \mathbb{N}A$ in $\text{Spec } \mathbb{C}[\alpha, \beta]$. Check that these operators are indeed in $D(R; \mathbb{C})$.

For example, if $p = q = -1$, then

$$((p, q) + \mathbb{N}A) \setminus \mathbb{N}A = \{(2k - 1, k - 1)\}_{k \geq 1} \cup \{(-1, k)\}_{k \geq -1},$$

so that $I_{(-1,-1)}(\theta_s, \theta_t) = (\theta_t + 1)(2\theta_2 - \theta_t + 1)$.

This exercise is a special case of a result about differential operators on affine semigroup rings; see [137, Theorem 2.1].

Exercise 17.9. Let f be an element in R . Use the order filtration on $D(R; \mathbb{K})$, as in Definition 17.1, to prove that each $P \in D(R; \mathbb{K})$ induces a differential operator on R_f . Thus, $R_f \otimes_R D(R; \mathbb{K})$ is a subset of $D(R_f; \mathbb{K})$.

Exercise 17.10. Check that $R_f \otimes_R D(R; \mathbb{K})$ is a ring. Prove that if R is a finitely generated \mathbb{K} -algebra, then $R_f \otimes_R D(R; \mathbb{K}) = D(R_f; \mathbb{K})$.

Hint: Each $P \in D(R_f; \mathbb{K})$ maps R to R_f . Since R is finitely generated, the image of P lies in $(1/f^k)R$ for some k . Use induction on the order of P .

Exercise 17.11. Show that $D(\mathbb{C}[x]; \mathbb{C})$ has proper left ideals that are not principal. In contrast, every left ideal of $D(\mathbb{C}(x); \mathbb{C})$ is principal.

Remark 17.12. Let \mathbb{K} be a field of characteristic zero, and R a domain finitely generated over \mathbb{K} . When R is regular, [54, Theorem 16.11.2] implies that $D(R; \mathbb{K})$ is the \mathbb{K} -algebra generated by R and the derivations. The converse is *Nakai's conjecture*.

2. The Weyl algebra

For the rest of this lecture, the field \mathbb{K} is of characteristic zero, R is a polynomial ring over \mathbb{K} in $\mathbf{x} = x_1, \dots, x_n$, and $D = D(R; \mathbb{K})$ is the Weyl algebra. We use \bullet to denote the action of D on R ; it is defined by

$$\partial_i \bullet f = \frac{\partial f}{\partial x_i}, \quad x_i \bullet f = x_i f \quad \text{for } f \in R.$$

For example, $\partial_i \bullet x_i = 1$ while $(\partial_i \cdot x_i)(f) = f + x_i \partial_i(f)$ for $f \in R$.

We study filtrations and associated graded rings on D .

Definition 17.13. The commutator relations (17.4.1) imply that each element of D can be written uniquely as

$$\sum_{\alpha, \beta \in \mathbb{N}^n} c_{\alpha, \beta} \mathbf{x}^\alpha \boldsymbol{\partial}^\beta,$$

where all but finitely many $c_{\alpha,\beta} \in \mathbb{K}$ are zero. The monomials $\mathbf{x}^\alpha \boldsymbol{\partial}^\beta$ form a \mathbb{K} -basis for D , which is the *Poincaré-Birkhoff-Witt basis* or *PBW basis*.

This basis can be used to solve Exercise 17.6. In what follows, we use it to define filtrations on D ; the discussion involves notions from Lecture 5.

Definition 17.14. A *weight* on D is a pair of vectors ω_x and ω_∂ in \mathbb{Z}^n , denoting weights on \mathbf{x} and $\boldsymbol{\partial}$, respectively, such that $\omega_x + \omega_\partial \geq 0$ componentwise. Given such an $\omega = (\omega_x, \omega_\partial)$, consider the subspaces of D

$$F_t = \mathbb{K} \cdot \{\mathbf{x}^\alpha \boldsymbol{\partial}^\beta \mid \langle \omega, (\alpha, \beta) \rangle \leq t\}.$$

The nonnegativity condition implies that no term in the expansion of $\boldsymbol{\partial}^\beta \mathbf{x}^\alpha$ in the PBW basis has weight larger than $\langle \omega, (\alpha, \beta) \rangle$. Check that $\bigcup_t F_t = D$ and that $F_s \cdot F_t \subseteq F_{s+t}$.

Strictly speaking, F is not a filtration in the sense of Definition 5.4, since F_t may be nonzero for $t < 0$, as in (3) below. This does not happen if the weight ω is nonnegative, as will be the case for most of this lecture.

- (1) $\omega = (\mathbf{0}, \mathbf{1}) = (0, \dots, 0, 1, \dots, 1)$ yields the *order filtration*: operators in $F_t = \mathbb{K} \cdot \{\mathbf{x}^\alpha \boldsymbol{\partial}^\beta \mid \sum_{j=1}^n \beta_j \leq t\}$ are precisely those with order $\leq t$.
- (2) $\omega = (\mathbf{1}, \mathbf{1}) = (1, \dots, 1, 1, \dots, 1)$ produces the *Bernstein filtration*, for which $F_t = \mathbb{K} \cdot \{\mathbf{x}^\alpha \boldsymbol{\partial}^\beta \mid \sum_{j=1}^n (\alpha_j + \beta_j) \leq t\}$.
- (3) $\omega = (-1, 0, \dots, 0, 1, 0, \dots, 0)$ yields the *V-filtration along $x_1 = 0$* , where $F_t = \mathbb{K} \cdot \{\mathbf{x}^\alpha \boldsymbol{\partial}^\beta \mid \beta_1 - \alpha_1 \leq t\}$.

Remark 17.15. Let ω be a weight vector on D so that F , as in Definition 17.14, is a filtration. Recall that $D = T/\mathfrak{a}$, where T is the tensor algebra on $2n$ variables, and \mathfrak{a} is the ideal generated by elements in (17.4.1).

The weight ω gives a filtration on T with the t -th level the linear combinations of all words that have ω -weight at most t . Then $\text{gr}_\omega(D)$ is the quotient of $\text{gr}_\omega(T)$ by $\text{gr}_\omega(\mathfrak{a})$. Note that $\text{gr}_\omega(\mathfrak{a})$ is generated by

$$\begin{aligned} & \{[x_i, x_j]\}, \quad \{[x_i, \partial_j] \mid i \neq j\}, \quad \{[\partial_i, \partial_j]\}, \\ & \{[x_i, \partial_i] \mid \omega_{x_i} + \omega_{\partial_i} > 0\}, \quad \text{and} \quad \{[x_i, \partial_i] + 1 \mid \omega_{x_i} + \omega_{\partial_i} = 0\}. \end{aligned}$$

It follows that the PBW basis on D gives one on $\text{gr}_\omega(D)$.

We focus on the Bernstein filtration B .

Exercise 17.16. Prove that $\text{gr}_B(D) = \mathbb{K}[\mathbf{x}, \boldsymbol{\partial}]$.

Proposition 17.17. *The Weyl algebra D is left and right Noetherian.*

Proof. Let \mathfrak{b} be a proper left ideal of D . The Bernstein filtration B on D induces an increasing filtration on \mathfrak{b} where

$$\mathfrak{b}_t = \mathfrak{b} \cap B_t \quad \text{for } t \geq 0.$$

Note that $\mathfrak{b}_0 = 0$ and that $\bigcup_t \mathfrak{b}_t = \mathfrak{b}$. It is easily checked that $\text{gr}(\mathfrak{b})$ is an ideal of $\text{gr}(D)$. Since $\text{gr}(D)$ is Noetherian by Exercise 17.16, there exist finitely many elements b_i in \mathfrak{b} such that their images generate the ideal $\text{gr}(\mathfrak{b})$. We claim that $\mathfrak{b} = \sum_i Db_i$, and hence that \mathfrak{b} is finitely generated.

Indeed, if not, then let t be the least integer for which there exists an element b in $\mathfrak{b}_t \setminus \sum_i Db_i$. By the choice of b_i , there exist elements d_i in D such that $b = \sum_i d_i b_i \pmod{\mathfrak{b}_{t-1}}$. But then

$$b = (b - \sum_i d_i b_i) + \sum_i d_i b_i \in \mathfrak{b}_{t-1} + \sum_i Db_i \subseteq \sum_i Db_i,$$

where the inclusion holds by choice of t . This is a contradiction.

At this point we know that D is left Noetherian. The proof that it is right Noetherian is similar. \square

Definition 17.18. Let M be a finitely generated D -module. The Bernstein filtration B on D induces a filtration, say G , on M as in Exercise 5.15, and the $\text{gr}_B(D)$ -module $\text{gr}_G(M)$ is Noetherian.

We write $\dim_B M$ and $e_B(M)$ for the *dimension* and *multiplicity* of M with respect to B ; see Definition 5.19.

Example 17.19. One has $\dim_B D = 2n$ and $e_B(D) = 1$ since $\text{gr}_B(D)$ is a polynomial ring in $2n$ variables.

One might expect the dimension of D -modules to take all values between 0 and $2n$. The next result, due to Bernstein, thus comes as a surprise.

Theorem 17.20. *If M is a finitely generated nonzero D -module, then one has $n \leq \dim_B M \leq 2n$.*

Proof. Let B be the Bernstein filtration on D , and G an induced filtration on M . We claim that the map

$$B_t \longrightarrow \text{Hom}_{\mathbb{K}}(G_t, G_{2t}) \quad \text{where } P \longmapsto (u \longmapsto Pu)$$

is injective. Indeed, this is trivially true for $t \leq 0$. Suppose $P \in B_t$ is nonzero and the claim holds for smaller t . If x_i , respectively, ∂_i , occurs in a term of P , then it is not hard to verify using the PBW basis that $[P, \partial_i]$, respectively, $[P, x_i]$, is in B_{t-1} and nonzero. Since the corresponding term does not annihilate G_{t-1} , one deduces that P cannot annihilate G_t . This settles the claim. It then follows that

$$\text{rank}_{\mathbb{K}}(B_t) \leq \text{rank}_{\mathbb{K}} \text{Hom}_{\mathbb{K}}(G_t, G_{2t}) = (\text{rank}_{\mathbb{K}} G_t)(\text{rank}_{\mathbb{K}} G_{2t}).$$

Therefore $\dim_B D \leq 2 \dim_B M$; now recall Example 17.19. \square

Exercise 17.21. Let B be the Bernstein filtration on D , and P an element in B_t . Prove the following statements.

- (1) $[P, x_i]$ and $[P, \partial_i]$ are in B_{t-1} .
- (2) $[P, x_i] \neq 0$, respectively, $[P, \partial_i] \neq 0$, if and only if ∂_i , respectively, x_i , occurs in the representation of P in terms of the PBW basis.

Extend this result to a general commutator $[P, Q]$.

We record some statements about general filtrations.

Remark 17.22. Let ω be a nonnegative weight on D . A spectral sequence argument as in [13, A:IV] shows that

$$\dim_{\omega} M = 2n - g(M) \quad \text{for } g(M) = \min\{i \mid \text{Ext}_D^i(M, D) \neq 0\}.$$

Thus, the dimension is independent of the nonnegative filtration. For example, the B -dimension and the dimension induced by the order filtration coincide. For the order filtration the analogue of Theorem 17.20 is the *weak fundamental theorem of algebraic analysis*; see [11] or [25].

Let ω be a weight with $\omega_x + \omega_{\partial} > 0$ componentwise. In this case, $\text{gr}_{\omega}(D) = \mathbb{K}[\mathbf{x}, \boldsymbol{\partial}]$. Let M be a finitely generated D -module. Theorem 17.20 remains true, with the caveat that there may be nonzero modules M with $\text{gr}_{\omega}(M) = 0$; consider $n = 1$, $M = D/D(x - 1)$ and $\omega = (-1, 2)$. See [148].

The annihilator of the $\text{gr}_{\omega}(D)$ -module $\text{gr}_{\omega}(M)$ is the *characteristic ideal* of M , with respect to ω . The associated subvariety of $\mathbb{A}_{\mathbb{K}}^{2n}$ is the *characteristic variety* of M . The weak fundamental theorem says that the characteristic variety with respect to the order filtration has dimension at least n . A stronger result is proved in [135]:

Theorem 17.23. *Let M be a nonzero finitely generated D -module. Each component V of its characteristic variety with respect to the order filtration satisfies $n \leq \dim V \leq 2n$. \square*

3. Holonomic modules

Recall that R is a polynomial ring in n variables over a field of characteristic zero, and that D is the corresponding Weyl algebra.

We single out a class of D -modules which enjoy special properties:

Definition 17.24. A finitely generated D -module M is *holonomic* when $\dim_B M \leq n$; equivalently, $\dim_B M = n$ or $M = 0$.

It follows from Remark 17.22 that a finitely generated D -module M is holonomic if and only if $\text{Ext}_D^i(M, D) = 0$ for $i < n$.

Exercise 17.25. Set $n = 1$. Prove that D/I is holonomic when $I \neq 0$.

Exercise 17.26. Prove that the D -module R is holonomic.

Exercise 17.27. Let $\mathfrak{m} = (\mathbf{x})$ in R . Verify that the image of $1/(x_1 \cdots x_n)$ generates $H_{\mathfrak{m}}^n(R)$ as a D -module. Prove that $H_{\mathfrak{m}}^n(R)$ is holonomic, and that it is isomorphic to $D/D\mathbf{x}$.

Remark 17.28. Consider an exact sequence of D -modules

$$0 \longrightarrow M' \longrightarrow M \longrightarrow M'' \longrightarrow 0.$$

Since D is Noetherian, M is finitely generated if and only if M' and M'' are finitely generated. When this holds, it follows from Exercise 5.20 and Theorem 17.20 that M is holonomic if and only if M' and M'' are holonomic, and also that, in this case, one has an equality

$$e_B(M) = e_B(M') + e_B(M'').$$

The definition of a holonomic module does not prepare us for the following remarkable result. Here B is the Bernstein filtration.

Theorem 17.29. *If M is a holonomic D -module, $\ell_D(M) \leq e_B(M) < \infty$.*

Proof. Let $0 \subset M_1 \subset \cdots \subset M_{l-1} \subset M_l = M$ be a strictly increasing filtration by D -submodules. It follows from Remark 17.28 that each M_i is holonomic. Moreover, the exact sequences of D -modules

$$0 \longrightarrow M_{i-1} \longrightarrow M_i \longrightarrow M_i/M_{i-1} \longrightarrow 0$$

imply that $e_B(M_1) < e_B(M_2) < \cdots < e_B(M)$, since $M_i/M_{i-1} \neq 0$. Therefore one obtains $l \leq e_B(M)$. This implies the desired result. \square

Another noteworthy property of holonomic modules is that they are cyclic; see [12, Theorem 1.8.18].

Exercise 17.30. Suppose $n = 1$. The D -modules $R \oplus R$ and $R \oplus D/Dx$ are holonomic. Verify that they are cyclic.

4. Gröbner bases

Gröbner basis techniques can be transplanted to Weyl algebras. To begin with, fix a weight filtration on D . A Gröbner basis for an ideal \mathfrak{a} is a system of generators for \mathfrak{a} whose images in $\text{gr}(D)$ generate $\text{gr}(\mathfrak{a})$.

Exercise 17.31. Let ω be a weight vector that is componentwise positive, and let \leq be the induced weight order. Show that \leq can be refined to a term order \leq_{lex} using a lexicographic order on the variables.

Show that if G is a \leq_{lex} -Gröbner basis for an ideal \mathfrak{a} , then it is also a Gröbner basis with respect to \leq .

For \leq a term order, Buchberger's algorithm applied to a generating system for \mathfrak{a} yields a Gröbner basis. If ω does not satisfy $\omega \geq 0$ componentwise, as in the case of the V -filtration, it cannot be refined to a term order. An approach to constructing Gröbner bases for such orders is outlined below.

Definition 17.32. The *homogenized Weyl algebra*, denoted $D^{(h)}$, is the ring $\mathbb{K}[h]\langle x_1, \dots, x_n, \partial_1, \dots, \partial_n \rangle$ modulo the two-sided ideal generated by

$$x_i \cdot x_j - x_j \cdot x_i, \quad \partial_i \cdot x_j - x_j \cdot \partial_i - \delta_{i,j} h^2, \quad \partial_i \cdot \partial_j - \partial_j \cdot \partial_i.$$

The *dehomogenization map* $D^{(h)} \rightarrow D$ is the surjective homomorphism induced by $h \mapsto 1$.

If a term order \prec on monomials of $D^{(h)}$ is such that $h^2 \prec x_i \partial_i$ for all i , then Algorithm 5.46 for the normal form terminates provided that G consists of homogeneous elements in $D^{(h)}$. Thus, Buchberger's algorithm terminates on homogeneous inputs in $D^{(h)}$; see [136, Proposition 1.2.2]. It remains to reduce Gröbner computations in D to homogeneous computations in $D^{(h)}$.

Exercise 17.33. Starting with the weight ω for D with $\omega_x + \omega_\partial \geq 0$ componentwise, let $t \geq 0$ be such that $2t \leq \omega_{x_i} + \omega_{\partial_i}$ for each i . Show that the weight (t, ω) for $D^{(h)}$ defines an order which, when refined with any term order for which $h < x_i$ and $h < \partial_i$ for each i , is an order \prec as discussed in the previous paragraph.

Let $\mathfrak{a} = (P_1, \dots, P_k)$ be an ideal in D . Show that the dehomogenization of a \prec -Gröbner basis for $P_1^{(h)}, \dots, P_k^{(h)}$ gives an ω -Gröbner basis for \mathfrak{a} .

Macaulay 2 has a package [98] devoted to D -modules.