The structure of Coh($\mathbb{P}^1$)

Team Edward, Talks 1 & 2

Some References:

Explicit Methods for Derived Categories of Sheaves, Alastair Craw
Algebraic Geometry, Robin Hartshorne - Chapters II and III

1 Coherent sheaves

1.1 Some preliminary comments

(We assume a basic familiarity with sheaves and affine/projective schemes, but review some of the relevant concepts here. We assume the reader has seen the Spec and Proj constructions, as well as the definition of a sheaf.)

Definition 1.1. *Suppose X is a noetherian scheme, and $\mathcal{F}$ is a sheaf of modules over $\mathcal{O}_X$. We call $\mathcal{F}$* **quasicoherent** *if for every affine open set $U = Spec(A)$ of X, $\mathcal{F}|_U \simeq \widehat{M}$, where M is a module over A. $\mathcal{F}$ is called* **coherent** *if these modules are moreover required to be finitely presented.*

For example, if $X = \mathrm{Spec}A$ is an affine scheme, then the functors

$$\Gamma(X, -) : \mathrm{QCoh}(X) \to \mathrm{Mod}(A)$$

$$\Gamma(X, -) : \mathrm{Coh}(X) \to \mathrm{fpMod}(A)$$

are equivalences. In both cases, the inverse equivalence is the functor $M \mapsto \widehat{M}$, sending a module M to the sheaf defined to equal to $\widehat{M_f}$ on the distinguished affine corresponding to the element f, for every nonzero f (and thus, defined inductively, since X is noetherian).

In general, we should think of *quasicoherent sheaves on a scheme* as being like a generalization of the notion of *modules on a ring* (after all, a quasicoherent sheaf is just given by locally patching together modules). It is then not hard to believe that

Proposition 1.2. *The categories $QCoh(X)$ and $Coh(X)$ are abelian.*

Almost all of the required constructions carry over from the corresponding constructions on modules by just doing the same thing on sections of sheaves, and checking compatibility with restriction maps. The zero sheaf functions as a zero object, and direct sums and kernels can be constructed sectionwise. But things are a little different for coker: if we try to construct the cokernel of a morphism sectionwise, we don't get a sheaf! This is because the functor $\mathcal{F} \mapsto \mathcal{F}(U)$ is not an exact functor on the category of quasicoherent sheaves (and so the sheaf condition, which consists of a limit diagram, no longer holds). An attempt to preserve exactness is precisely what gives rise to *sheaf cohomology* (which we'll see later on). We actually have to *sheafify* the presheaf we get from this construction, and this works as the cokernel: I won't go into details here.

There's also a tensor product structure (that's symmetric up to isomorphism). As with the cokernel, we need to sheafify if we aren't working over an affine scheme. If X is affine, then the equivalence $\Gamma(X, -)$ preserves this tensor product.

It is also worth mentioning that if $X = \mathrm{Proj}(S)$ where S is a graded ring (for example, $\mathbb{P}^1_k = \mathrm{Proj}(k[x_0, x_1])$), then coherent sheaves on X arise via the $M \mapsto \widehat{M}$ construction on *graded* modules M over S, like in the affine case. Given a graded module M over S, we can construct a sheaf $\widehat{M}$, defined by the property that the stalk at a point p consists of the degree 0 elements of M_p. Any graded module gives rise to a sheaf in this way, every coherent sheaf arises this way, and two modules M and M' gives rise to the same sheaf iff, for n sufficiently large, $M_n = M'_n$.

1.2 Locally free sheaves, and the Serre twisting sheaf

Definition 1.3. *A sheaf $\mathcal{F}$ on X is called* **locally free** *(or a* **vector bundles***) if there is an open affine cover $\{U_i\}_{i \in I}$ of X such that $\mathcal{F}|_U \simeq \mathcal{O}_X^{\oplus m}|_U$ for some m. Equivalently, $\mathcal{F}$ is locally free iff, for every $p \in X$, $\mathcal{F}_p \simeq \mathcal{O}_{X,p}^m$ for some m. We call $\mathcal{F}$ a* **line bundle** *in the case $m = 1$.*

It's worth noting that these are, in some sense, the 'generators' for $\mathrm{Coh}(X)$. Remember that if R is a ring of dimension n, then any module M

over R has a free resolution

$$0 \to F_n \to F_{n-1} \to \cdots \to F_1 \to M \to 0$$

(called a *syzygy*). Similarly, any coherent sheaves over $\mathbb{P}^n_k$ has a resolution by sheaves which are sums of line bundles, where the length of the resolution will be $n+1$.

So exactly what data goes into a line bundle over $\mathbb{P}^1$? Recall that $\mathbb{P}^1$ consists of two copies of $\mathbb{A}^1$ (call them $U_x = k[x]$ and $U_y = k[y]$), glued along the affine subscheme $U_x \cap U_y = k[x,y]/(xy-1) \simeq k[s,s^{-1}]$. Here, we should think of $x = x_0/x_1$ whenever $x_1 \neq 0$, and $y = x_1/x_0$ whenever $x_0 \neq 0$, where x_0, x_1 are the homogeneous coordinates. A coherent sheaf $\mathcal{F}$ consists of a module $\mathcal{F}(U_x)$ over $k[x]$ and a module $\mathcal{F}(U_y)$ over $k[y]$, along with some sort of clutching function in the middle to glue them together. If $\mathcal{F} = \mathcal{L}$ is a line bundle, then these are required to be locally free of rank 1.

Theorem 1.4. *(Serre-Swan) A finitely generated module over a commutative ring is locally free if and only if it is projective.*

Theorem 1.5. *Finitely generated projective modules over a PID are free.*

In particular, since $k[x]$ is a PID, this implies that locally free modules over $\mathbb{A}^1$ are always free. So it follows that $\mathcal{L}(U_x)$ and $\mathcal{L}(U_y)$ are, in fact, free over $k[x]$ and $k[y]$, respectively. Now consider the diagram

$$\begin{array}{ccc} \mathcal{L}(U_x) & & \mathcal{L}(U_y) \\ & \searrow \quad \swarrow & \\ & \mathcal{L}(U_x \cap U_y) & \end{array}$$

where these are rank one free modules over $k[x]$, $k[y]$, and $k[s,s^{-1}]$, respectively. The first map is a $k[x]$-module homomorphism (where x acts by s) and the second is a $k[y]$-module homomorphism (where y acts by s^{-1}). The two maps send the generators of $\mathcal{L}(U_x)$ and $\mathcal{L}(U_y)$ to two generators of $\mathcal{L}(U_x \cap U_y)$. The second can therefore be written as the first times an element $A(s,s^{-1})$ of $k[s,s^{-1}]$.

The line bundle $\mathcal{L}$ is entirely determined by this Laurent polynomial, called the *clutching function*. This is precisely the matrix $A(s, s^{-1})$ we got in the previous section. We know this Laurent polynomial is, after an appropriate automorphism of the structure sheaf, an integer power of s. So we have

Proposition 1.6. *For any n, there is exactly one line bundle $\mathcal{O}(d)$ (up to isomorphism) on $\mathbb{P}^n$ for every $d \in \mathbb{Z}$. $\mathcal{O}(d)$ is defined by the clutching function $A(s, s^{-1}) = s^{-d}$.*

Let's explicitly compute the sections of this line bundle. Any open set other than the entirety of $\mathbb{P}^1$ is contained within an affine, and so we know what sections look like. A global section of the line bundle $\mathcal{O}(d)$ consists of an element of $\mathcal{L}(U_x)$ and an element of $\mathcal{L}(U_y)$ which restrict to the same element of $\mathcal{L}(U_x \cap U_y$. Let g_x denote the generator of $\mathcal{L}(U_x)$, and let g_y denote the generator of $\mathcal{L}(U_y)$. Then we can check that if the clutching function is s^{-d}, then we have $d+1$ linearly independent pairs of sections that glue: $(g_x, g_y y^d), (g_x x, g_y y^{d-1}), \ldots, (g_x x^d, g_y)$. Hence, we have, by an obvious bijection,

Proposition 1.7. *The global sections of $\mathcal{O}(d)$ correspond to homogeneous degree d elements of $k[x_0, x_1]$.*

For $\mathbb{P}^n_k$, the story of these line bundles is much the same. $\mathbb{P}^n$ consists of $n+1$ copies of $\mathbb{A}^n$, called $U_0, \ldots, U_n$, where $U_i = \mathrm{Spec}(k[x_{0/i}, x_{1/i}, \ldots, x_{n/i}])$, glued in the natural way. We can construct a line bundle on $\mathbb{P}^n$ by picking transition functions along $U_i \cap U_j$ which are compatible on the intersections $U_i \cap U_j \cap U_k$. For the line bundle $\mathcal{O}(d)$, the transition function from U_i to U_j should involve multiplication by $(x_i/x_j)^d$.

It turns out that $\mathcal{O}(d)$ is an important line bundle to consider on a general projective scheme $\mathrm{Proj}(S)$ (for example, a projective variety), and is defined as follows.

Definition 1.8. *(Serre twisting sheaf)*

Given a graded ring S, define $\mathcal{O}(1)$ (called the twisting sheaf of Serre) as $\widehat{M}$ where for each n, $M_n = S_{n+1}$. Similarly define $\mathcal{O}(d)$ as $\widehat{M}$ where for each n, $M_n = S_{n+d}$.

Corollary 1.9. *The global sections of the sheaf $\mathcal{O}(d)$ on $\mathbb{P}^n_k$ correspond to the homogeneous elements of $k[x_0, \ldots, x_n]$ of degree d, which form a space of dimension $\binom{d+n}{n}$. In particular, when $d > 0$, this is a polynomial of degree n in d, and when $d < 0$, there are no global sections.*

Exercise: Check that this definition of $\mathcal{O}(d)$ agrees in the case of $\mathbb{P}^1$, i.e., the line bundle with clutching function s^d is, in fact, constructed from the graded ring whose degree n part is the degree $n + d$ part of $k[x_0, x_1]$.

It is also worth noting that if X is any projective variety over k (and thus is equipped with a map $f : X \to \mathbb{P}^n_k$ for some n), then we can pull back the line bundles $\mathcal{O}(d)$ by f to get line bundles $\mathcal{O}_X(d) := f^*\mathcal{O}(d)$ on X. The sections of this line bundle correspond to degree d terms in the coordinate ring, as we'd expect.

1.3 Tensor products, dual, and inner hom

Next, we'd like to understand what morphisms are possible between two line bundles $\mathcal{O}(d_1)$ and $\mathcal{O}(d_2)$. We need to carry over some more of the tools and language from the category of modules over a commutative ring in order to do so. The following general statements about sheaves will work over Proj(S) for any graded ring S.

Proposition 1.10. *Let $\widehat{M}$ be any sheaf over $X = Proj(S)$, where M is a graded module over S. Then $\mathcal{O}(1) \otimes \widehat{M} = \widehat{M(1)}$, where $M(1)_n = M_{n+1}$.*

Proof. $M(1) = M \otimes_S S(1)$, and $\mathcal{O}(1) = \widehat{S(1)}$. So the result holds because the construction $M \mapsto \widehat{M}$ commutes with tensor product. □

Corollary 1.11. $\mathcal{O}(m) \otimes \mathcal{O}(n) \simeq \mathcal{O}(m + n)$.

(Again, in the world of K-theory, this tells us that the tensor product on line bundles adds the Chern classes.)

Next, recall that the category of sheaves is equipped with an *inner hom.* Just as $\hom_R(M, N)$ has the structure of an R-module, where R is a commutative ring and M and N are R-modules, we can construct a **presheaf** of modules over $\mathcal{O}_X$, $hom_{\mathcal{O}_X}(\mathcal{F},\mathcal{G})$, defined by the property that

$$hom_{\mathcal{O}_X}(\mathcal{F},\mathcal{G})(U) = \mathrm{Hom}(\mathcal{F}|_U,\mathcal{G}|_U)$$

This is a sheaf (not hard to check) whose *global sections are Hom.* There is also an adjunction

$$hom_{\mathcal{O}_X}(\mathcal{F}\otimes\mathcal{G},\mathcal{H}) \simeq hom_{\mathcal{O}_X}(\mathcal{F},hom_{\mathcal{O}_X}(\mathcal{G},\mathcal{H}))$$

This object leads to some results we might expect from the analogy with R-modules.

Definition 1.12. *For a locally free sheaf $\mathcal{F}$, define its* **dual** $\mathcal{F}^\vee = hom_{\mathcal{O}_X}(\mathcal{F},\mathcal{O}_X)$.

Proposition 1.13. $\mathcal{O}(d)^\vee \simeq \mathcal{O}(-d)$.

Proof. Simply check that the morphism $\mathcal{O}(d)\otimes hom_{\mathcal{O}_X}(\mathcal{O}(d),\mathcal{O}) \to \mathcal{O}$ is an isomorphism on the level of stalks. □

Corollary 1.14.

$$hom(\mathcal{O}(m),\mathcal{O}(n)) \simeq \mathcal{O}(n-m)$$

Proof. Use the previous proposition, along with the adjunction. □

Thus, because $\mathrm{Hom}(\mathcal{F},\mathcal{G}) = \Gamma(X, hom(\mathcal{F},\mathcal{G}))$, we see that $\mathrm{Hom}(\mathcal{O}(m),\mathcal{O}(n)) = \Gamma(X,\mathcal{O}(n-m))$ which is nonzero if and only if $n \geq m$. This illustrates an important principle for line bundles: *there are nontrivial morphisms upwards in degree, but not downwards in degree.*

1.4 Skyscraper sheaves, and a decomposition theorem

We've pretty thoroughly explored the story of line bundles, but obviously, not every coherent sheaf is locally free! Here, we'll briefly explore the structure of skyscraper sheaves (more generally, torsion sheaves, whose stalks are zero-dimensional at all but finitely many points), and then give a general decomposition theorem for coherent sheaves on $\mathbb{P}^1$.

Consider the morphism of line bundles $\mathcal{O}_X \to \mathcal{O}_X(d)$ given by multiplication by some nonzero homogeneous polynomial f of degree d. The kernel is the zero sheaf, because f is nonzero, and $\mathbb{A}^n$ is an integral domain. To compute the cokernel, we'll have to look at the level of stalks. It turns out that at a point p, the cokernel of the map $\mathcal{O}_{X,p} \to \mathcal{O}_{X,p}(1)$ is a vector space over k whose dimension is the *order of* p *as a zero of* f. So at points where f is nonzero, the stalk will be zero, and at points where f has a zero of some nonzero multiplicity m, the stalk will be a vector space of dimension m. Summarizing,

Definition 1.15. *Let $\mathcal{O} \to \mathcal{O}(d)$ be the map given by multiplication by a homogeneous degree d polynomial f. Then the cokernel sheaf is called $\mathcal{O}_f$, and the stalk at any point p is a vector space over k whose dimension is the order of f at p. In the case where f has degree* 1 *(for example, $f = x$), we call $\mathcal{O}_f$ a skyscraper sheaf. We define $\mathcal{O}_f$ to have* degree *equal to deg(f).*

(Intuitively, we should think of torsion sheaves as being very small. For example, one can easily show that $\mathcal{O}_f = \widehat{M}$, where $M = k[x_0, x_1]/(f)$. Note that this does not always decompose into skyscraper sheaves: for example, if $f = (x_0 - x_1)^2$, we get the sheaf associated to $k[x_0, x_1]/((x_0 - x_1)^2)$. This is not the same thing as the sheaf associated to $k[x_0, x_1]/(x_0 - x_1) \oplus k[x_0, x_1]/(x_0 - x_1)$!)

Now we have enough language to understand how arbitrary coherent sheaves on $\mathbb{P}^1$ look.

Theorem 1.16. *Any coherent sheaf on $\mathbb{P}^1_k$ is a finite direct sum of line bundles and degree one skyscraper sheaves.*

The proof is roughly the same idea as the structure theorem for finitely generated modules over a PID: applying linear algebra to get a Jordan-type form.

2 Serre Duality

Serre duality relates the cohomology groups of a locally free sheaf $\mathcal{F}$ on a smooth projective scheme X of dimension n (over base field k). It states that there is a unique (up to isomorphism) line bundle ω_X (called the *dualizing sheaf*), such that for all coherent $\mathcal{F}$, i, there is a natural isomorphism

$$\mathrm{Ext}^i(\mathcal{F}, \omega) \simeq H^{n-i}(X, \mathcal{F})^*$$

If $\mathcal{F}$ is locally free, the left side becomes $H^i(X, \mathcal{F}^\vee \otimes \omega_X)$, so this relates the cohomology groups of $\mathcal{F}$ and $\mathcal{F}^\vee \otimes \omega_X$ (ω_X provides a sort of 'twisting'). For example, in the case of an oriented compact manifold, this is an analogue of Poincare duality (the dualizing sheaf is trivial).

We'll build up some of the machinery necessary to understand and prove this theorem, and also see explicitly that the dualizing sheaf is precisely the *canonical bundle*: the top exterior power of the sheaf of differential 1-forms. Then we'll see various statements of the theorem.

2.1 An algebraic tool with short exact sequences

With functors like H^* and Ext^*, a short exact sequence (or fiber sequence) of objects leads us to a long exact sequence of groups. This makes short exact sequences invaluable in computing things: if we can understand two objects in a short exact sequence, usually we can understand the third. But what about when we have a resolution that's longer than 3 objects? We might as well show how to utilize such sequences here.

Proposition 2.1. *Let*

$$0 \longrightarrow A_n \xrightarrow{\varphi_n} A_{n-1} \xrightarrow{\varphi_{n-1}} \cdots \xrightarrow{\varphi_2} A_1 \xrightarrow{\varphi_1} A_0 \longrightarrow 0$$

be a long exact sequence of objects in an abelian category. Then we have a collection of short exact sequences with the induced maps

$$0 \longrightarrow A_n \longrightarrow A_{n-1} \longrightarrow coker(\varphi_n) \longrightarrow 0$$

$$0 \longrightarrow coker(\varphi_n) \longrightarrow A_{n-2} \longrightarrow coker(\varphi_{n-1}) \longrightarrow 0$$

$$0 \longrightarrow coker(\varphi_{n-1}) \longrightarrow A_{n-3} \longrightarrow coker(\varphi_{n-2}) \longrightarrow 0$$

$$\vdots$$

$$0 \longrightarrow coker(\varphi_3) \longrightarrow A_1 \longrightarrow A_0 \longrightarrow 0$$

The proof is easy. This statement is, however, important. It means that if we have a resolution of a sheaf $\mathcal{F}$ by objects we understand well (for example, locally free sheaves), then we can understand $\mathcal{F}$ as well by understanding all of the successive cokernels. In particular, in $\mathbb{P}^n_k$, any coherent sheaf $\mathcal{F}$ has a resolution by locally free sheaves

$$0 \to \mathcal{F}_n \to \mathcal{F}_{n-1} \to \cdots \to \mathcal{F}_0 \to \mathcal{F}$$

so if we can compute the cohomology of the locally free sheaves, we can compute the cohomology of any coherent sheaf. In the derived category, we can literally just replace any object with one composed of locally free sheaves in this way.

2.2 Cech cohomology, and computation of the cohomology of line bundles on projective space

The cohomology of a sheaf can be defined using injective resolutions. But in practice, this is a poor way of doing computations: injective sheaves are generally huge, and it's not clear how you'd generate a resolution of them! However, a scheme is usually presented to us in terms of affines (i.e., perhaps there's a nice affine cover we can pick). This seems useful, because of the following theorem.

Theorem 2.2. *Let $X = Spec(A)$ be the spectrum of a noetherian ring A. Then for all quasicoherent sheaves $\mathcal{F}$ on X, and all $i > 0$, $H^i(X, \mathcal{F}) = 0$.*

Proof. This is clear, because $\Gamma(X, -) : \mathrm{QCoh}(X) \to \mathrm{Mod}(A)$ is an equivalence of categories (so $\Gamma(X, -)$ is exact if X is affine). □

Exercise: Use this to deduce that skyscraper sheaves have trivial cohomology groups.

So, perhaps there's a better, more obvious choice of "fibrant replacement" we can use to compute cohomology. It turns out that there is.

Definition 2.3. *Let X be a scheme, and $\mathfrak{U} = \{U_i\}$ a cover by open affines, where the index set is ordered. For a coherent sheaf $\mathcal{F}$ on X, define its* Cech resolution

$$\mathcal{C}^p(\mathfrak{U}, \mathcal{F}) = \prod_{|I|=p} (f_I)_*(\mathcal{F}|_{U_I})$$

where the differential map $\mathcal{C}^p \to \mathcal{C}^{p+1}$ is defined on a section s by

$$(ds)_{i_0 \cdots i_{p+1}} = \sum_{k=0}^{p+1} (-1)^k s_{i_0 \cdots \hat{i_k} \cdots i_{p+1}}|_{U_{i_0 \cdots i_{p+1}}}$$

Define the Cech complex $C^*(\mathfrak{U}, \mathcal{F}) = \Gamma(X, \mathcal{C}^*(\mathfrak{U}, \mathcal{F}))$

Proposition 2.4. *If X is additionally a noetherian separated scheme, then $H^i(C^p(\mathfrak{U}, \mathcal{F})) = H^i(X, \mathcal{F})$. ($X$ must be separated so that the intersections U_I are all affine.)*

Roughly, the idea is that this Cech resolution is taking the place of the injective resolution as the fibrant replacement for $\mathcal{F}$. This is immensely useful, because in many cases, the Cech complex is quite easy to compute! The reason this works is, roughly, because we're choosing a different model structure on $\mathrm{Ch}^+(\mathrm{Coh}(X))$: one where the fibrant objects are now complexes of *sheaves which are direct sums of sheaves supported on an affine open.*

Let A be a noetherian ring, and let $S = A[x_0, \ldots, x_n]$ be a graded ring. Let $X = \mathrm{Proj}(S) = \mathbb{P}^n_A$. Denote by $\mathcal{O}_X(1)$ the usual twisting sheaf of Serre.

Theorem 2.5.

1. *$H^i(X, \mathcal{O}_X(d)) = 0$ for $0 < i < n$ and all $d \in \mathbb{Z}$.*
2. *$H^n(X, \mathcal{O}_X(-n-1)) \simeq A$.*
3. *The natural map*

   $$H^0(X, \mathcal{O}_X(d)) \times H^n(X, \mathcal{O}_X(-d-n-1)) \to H^n(X, \mathcal{O}_X(-n-1)) \simeq A$$

   is a perfect pairing of finitely generated free A-modules, for each $d \in \mathbb{Z}$.

Proof. As usual, for $i = 0, 1, \ldots, n$, let $U_i = \mathrm{Spec}(A[x_0/x_i, x_1/x_i, \ldots, x_n/x_i])$ be the i-th distinguished open affine. Not surprisingly, we will use the open cover $\mathfrak{U} = \{U_0, \ldots, U_n\}$ to compute Cech cohomology. However, we will be computing the cohomology of $\mathcal{F} = \bigoplus_{d \in \mathbb{Z}} \mathcal{O}(d)$ and will keep track of the grading by d.

Note that $\mathcal{F}(U_{i_0 i_1 \cdots i_p}) = S_{x_{i_0} \cdots x_{i_p}}$. So the Cech complex is

$$0 \to \prod_{i_0} S_{x_{i_0}} \to \prod_{i_0 < i_1} S_{x_{i_0} x_{i_1}} \to \cdots \to S_{x_0 x_1 \cdots x_n} \to 0$$

where the maps are induced by the localization maps. We claim that this sequence is exact at every stage except for possibly at the beginning and at the end. As an example, we will show it is exact at $\prod_{i_0 < i_1} S_{x_{i_0} x_{i_1}}$. Pick an element s of this group, and let's call the component in $S_{x_i x_j}$ s_{ij}. Suppose it is in the kernel of the differential. Then it must map to zero in every $S_{x_{i_0} x_{i_1} x_{i_2}}$. Therefore, the elements in $s_{i_0 i_1}, s_{i_0 i_2}$, and $s_{i_1 i_2}$ must sum to zero. It follows that $s_{i_0 i_1}$ cannot have any terms with the x_{i_0} and x_{i_1} exponents negative. In general, in s, there are thus no terms with more than one exponent negative. Next, consider the x_{i_0}-exponent-negative terms in s_{i_0} and s_{i_1}. Because $s_{i_0 i_1 i_2} = 0$, these terms must be equal (recall the powers of -1 in the differential of the Cech complex). It follows that we can pick an element of $S_{x_{i_0}}$ which maps to each of these terms. We can do this for every index and we then get that s comes from an element of $\prod_{i_0} S_{x_{i_0}}$, as desired. This prove (1).

Now we must compute the cohomology at either end, namely, the kernel of the first map and the cokernel of the last. We already know H^0: this is just global sections. So it remains to compute H^n. $S_{x_0 \ldots x_n}$ is a free A-module spanned by monomials with exponents arbitrary, while the image of $\prod_i S_{x_0 \cdots \hat{x_i} \cdots x_n}$ is spanned by the monomials with at least one exponent nonnegative. Thus, the cokernel is spanned by the monomials $x_0^{l_0} \cdots x_n^{l_n}$ with $l_0, l_1, \ldots, l_n < 0$. This has dimension 0 in degree $-n$ and greater, and for any degree $d \leq -n-1$, it has the dimension $\binom{-d-1}{n}$, which is the same as the dimension of $H^0(X, \mathcal{O}(-d-n-1))$. In particular, $H^n(X, \mathcal{O}(-n-1)) \simeq A$. This proves (2).

Now, to prove (3), we just have to find a nondegenerate pairing for

$d \geq 0$ (and since the two modules are isomorphic and finite dimensional, this will complete the proof). $H^0(X, \mathcal{O}(d))$ has a basis consisting of the monomials of degree d with all exponents nonnegative, and as we just saw, $H^n(X, \mathcal{O}(-d-n-1))$ has a basis consisting of monomials of degree $-d-n-1$ with all exponents negative. The natural pairing arises from taking the product on these monomials, and we will declare the product to be zero in $H^n(X, \mathcal{O}(-n-1))$ if any of the exponents is nonnegative. It is not hard to check that this is a perfect pairing. $\square$

(**Note:** The dimension of $H^0(X, \mathcal{O}(d))$ is $\binom{d+n}{n}$. For $d \geq -n$, this is equal to the polynomial function $\frac{(d+n)(d+n-1)\cdots(d+1)}{n!}$, but for $d \leq -n-1$, the dimension is zero, whereas this function is not. The H^n terms we see showing up serve to make it so that the *Euler characteristic* χ satisfies this polynomial, i.e.,

$$\chi(X, \mathcal{F}) := \sum_{i \geq 0} (-1)^i \dim_k H^i(X, \mathcal{F}) = \frac{(d+n)(d+n-1)\cdots(d+1)}{n!}$$

You might recognize this as the *Hilbert polynomial* of the graded ring S!)

2.3 The canonical bundle is $\mathcal{O}_X(-n-1)$ on $\mathbb{P}^n_k$

We want to show that, for $X = \mathbb{P}^n_k$, the sheaf $\omega_X = \wedge^n(\Omega_{X/k})$ of differential n-forms is $\mathcal{O}(-n-1)$. We'll show this in two different ways.

First, let's figure out what the local clutching functions should look like. Pick coordinates $x_{10}, x_{20}, \ldots, x_{n0}$ on U_0 and $x_{01}, x_{21}, \ldots, x_{n1}$ on U_1. Consider the basis element $dx_{10} \wedge dx_{20} \wedge \cdots \wedge dx_{n0}$ on U_0. What does this equal in terms of the coordinates on U_1? Rewriting,

$$= d\left(\frac{1}{x_{01}}\right) \wedge d\left(\frac{x_{21}}{x_{01}}\right) \wedge d\left(\frac{x_{31}}{x_{01}}\right) \wedge \cdots \wedge d\left(\frac{x_{n1}}{x_{01}}\right)$$

$$= -\frac{dx_{01}}{x_{01}^2} \wedge \bigwedge_{i=2}^{n} \left(\frac{dx_{i1}}{x_{01}} - \frac{x_{i1}}{x_{01}^2} dx_{01}\right)$$

$$= -\frac{1}{x_{01}^{n+1}} dx_{01} \wedge dx_{21} \wedge \cdots \wedge dx_{n1}$$

That is, the transition function involves multiplying by an element of degree $-n-1$. These are precisely the clutching functions defining $\mathcal{O}(-n-1)$. Let's now show this another way.

Proposition 2.6. *There is a short exact sequence of sheaves over* $X = \mathbb{P}^n$

$$0 \to \Omega_{X/k} \to \mathcal{O}(-1)^{\oplus(n+1)} \to \mathcal{O}_X \to 0$$

Proof. Let $E = S(-1)^{n+1}$ be generated by elements $e_0, e_1, \ldots, e_n$ in degree 1. Then we have a homomorphism $E \to S$ defined by $e_i \mapsto x_i$. Let M be the kernel of this map, so that we have the sequence

$$0 \longrightarrow M \longrightarrow E \longrightarrow S$$

We then have a short exact sequence of sheaves

$$0 \longrightarrow \tilde{M} \longrightarrow \mathcal{O}(-1)^{\oplus(n+1)} \longrightarrow \mathcal{O}_X \longrightarrow 0$$

Why is the last morphism surjective? The map $E \to S$ is not a surjection of modules, but the image is all of S in degree sufficiently large (where sufficiently large here means at least 1). Another way to see it is that, for each x_i, the map of localized modules $E_{x_i} \to S_{x_i}$ is surjective, so the morphism of sheaves is surjective on each distinguished open: hence it is surjective.

Hence, it just suffices to show that $\tilde{M} = \Omega_{X/k}$. We'll show that the modules are equal on the distinguished opens U_i. On U_i, the module of differentials is generated in grading 0 by $dx_{0i}, dx_{1i}, \ldots, d\hat{x}_{ii}, \ldots, dx_{ni}$ as a module over $k[x_{0i}, \ldots, \hat{x_{ii}}, \ldots, x_{ni}]$ (obviously the higher gradings are given by multiplying by appropriate integer powers of x_i). On the other hand, the graded module M consists of $(n+1)$-tuples of homogeneous polynomials $(f_0, \ldots, f_n)$ such that $x_0 f_0 + \ldots + x_n f_n = 0$. Passing to M_{x_i}, we see that

$$-f_i = x_{0i} f_0 + \ldots + x_{ni} f_n$$

(where obviously $x_{ii} f_i$ is not on the right). So M_{x_i} consists, in degree 0 of n-tuples $(f_0, \ldots, \hat{f}_i, \ldots, f_n)$ with no conditions (f_i is determined by the above equation). Again, the higher degrees are given by multiplying by appropriate powers of x_i. Hence, $\widehat{M_{x_i}} = (\Omega_{X/k})|_{U_i}$. It is not hard to check that the transition maps are equal as well. So $\widehat{M} \simeq \Omega_{X/k}$, as desired. $\square$

From here, we can just use the following easy fact.

Proposition 2.7. *Let S be a graded ring, and suppose $0 \to A \to B \to C \to 0$ is a short exact sequence of graded S-modules. $\wedge^n B$ has a filtration*

$$E_0 \hookrightarrow E_1 \hookrightarrow \cdots \hookrightarrow E_n = \wedge^n B$$

where E_i is generated by elements $b_1 \wedge b_2 \wedge \cdots \wedge b_n$ where $n - i$ of the b_j's are in A. Then each quotient E_{i+1}/E_i is isomorphic to $\wedge^{n-i} A \otimes \wedge^i C$.

Corollary 2.8. $\wedge^n \Omega_{X/k} \simeq \mathcal{O}(-n-1)$.

Proof. Recall we have the short exact sequence

$$0 \to \Omega_{X/k} \to \mathcal{O}(-1)^{\oplus(n+1)} \to \mathcal{O}_X \to 0$$

Now look at $\wedge^{n+1}(\mathcal{O}(-1)^{\oplus(n+1)})$. This is the sheaf associated to the module $\wedge^{n+1}(S(-1)^{n+1})$ (where $\wedge$ is taken over S-modules). This is generated by the $(n+1)$-fold wedge tensors where each element is concentrated in a single coordinate. In other words, this is isomorphic to $\bigotimes_{i=0}^{n} S(-1) \simeq S(-n-1)$. Hence, $\wedge^{n+1}(\mathcal{O}(-1)^{\oplus(n+1)}) \simeq \mathcal{O}(-n-1)$.

On the other hand, because $\wedge^{n+1}(\Omega_{X/k}) = 0$ and $\wedge^2 \mathcal{O}_X = 0$, we see that in the filtration from the previous proposition, there is only one nonzero quotient. That is,

$$\wedge^{n+1}(\mathcal{O}(-1)^{\oplus(n+1)}) \simeq \wedge^n \Omega_{X/k} \otimes \mathcal{O}_X \simeq \wedge^n \Omega_{X/k}$$

as desired. □

2.4 Serre Duality

Theorem 2.9. *(Serre duality for $\mathbb{P}^n_k$) Let $X = \mathbb{P}^n_k$, and let $\omega_X = \wedge^n \Omega_{X/k} \simeq \mathcal{O}(-n-1)$. Then*

1. *$H^n(X, \omega_X) \simeq k$. Fix one isomorphism.*
2. *For any coherent sheaf $\mathcal{F}$ over X, the natural pairing*

$$Hom^i_{\mathcal{O}_X}(\mathcal{F}, \omega_X) \times H^n(X, \mathcal{F}) \to H^n(X, \omega_X) \simeq k$$

is a perfect pairing of finite-dimensional vector spaces over k

3. For every i, there's an isomorphism

$$\mathit{Ext}^i(\mathcal{F}, \omega_X) \to H^{n-i}(X, \mathcal{F})^*$$

which, for $i = 0$, is induced by the map form 2.

Proof. (1) holds from Theorem 5.5. For (2), recall that if $0 \to \mathcal{H} \to \mathcal{G} \to \mathcal{F} \to 0$ is a short exact sequence of sheaves, then we have long exact sequences

$$0 \to \mathrm{Hom}(\mathcal{F}, \omega_X) \to \mathrm{Hom}(\mathcal{G}, \omega_X) \to \mathrm{Hom}(\mathcal{H}, \omega_X) \to \mathrm{Ext}^1(\mathcal{F}, \omega_X) \to$$

and

$$\cdots \to H^1(X, \mathcal{H})^* \to H^0(X, \mathcal{F})^* \to H^0(X, \mathcal{G})^* \to H^0(X, \mathcal{H})^* \to 0$$

We have maps connecting these two sequences in parallel (like a ladder). Recall that any coherent sheaf $\mathcal{F}$ can be resolved by sheaves that are sums of line bundles. Such a resolution can be decomposed into n short exact sequences. Hence, if we can prove the desired statement for line bundles, then by applying the 5-lemma n times in ladders like those above, it will follow for any coherent sheaf. This is just Theorem 5.5, though.

To show (3), we simply need to construct such a map, and then can again apply the 5-lemma trick (because the statement holds for line bundles). This follows from some stuff about universal δ-functors. □

Note that when $\mathcal{F}$ is locally free, $\mathrm{Ext}^i(\mathcal{F}, \omega_X) \simeq \mathrm{Ext}^i(X, \mathcal{F}^\vee \otimes \omega_X) \simeq H^i(X, \mathcal{F}^\vee \otimes \omega_X)$, giving the statement described at the beginning of the section.

We can, in fact, generalize the conditions of the theorem. The same statement holds whenever X is a projective scheme that is Cohen-Macaulay and all irreducible components have the same dimension. In particular, it holds for X a smooth projective variety.

2.5 In the Derived Category

The notion of a *Serre duality functor* can be defined in a more general triangulated category (see notes from lecture 3). We'll reproduce the definition here.

Definition 2.10. *Let $\mathcal{D}$ be a k-linear triangulated category with finite dimensional Hom's (so that taking the dual twice gives back the Hom space), where k is an algebraically closed field. A Serre functor $S : \mathcal{D} \to \mathcal{D}$ is an additive equivalence of categories, with bi-functorial isomorphisms*

$$Hom_{\mathcal{D}}(A, B) \to Hom_{\mathcal{D}}(B, S(A))^{\vee}$$

for any $A, B \in \mathcal{D}$. Such a functor also commutes with the translation functor, and the composite

$$Hom(A, B) \to Hom(B, S(A))^{\vee} \to Hom(S(A), S(B))$$

is the isomorphism given by S.

Proposition 2.11. *Let X be a smooth projective scheme of dimension n. Then $- \otimes \omega_X[n]$ is a Serre duality functor on $D^b Coh(X)$.*

Proof. (Sketch) Serre duality implies that if A, B are locally free sheaves on X, then $\mathrm{Ext}^i(A, B) \simeq \mathrm{Ext}^{n-i}(B, A \otimes \omega_X)^{\vee}$, simply because we can freely move locally free sheaves back and forth over Ext's, as long as we replace them by their duals. Next, remember that in $D^b\mathrm{Coh}(X)$, Hom comes equipped with all of the Ext's as homotopy groups. Now we just carefully count gradings: the i-th homotopy group of $\mathrm{Hom}(B, A \otimes \omega_X[n])^{\vee}$ equals the $(i-n)$-th homotopy group of $\mathrm{Hom}(B, A \otimes \omega_X)^{\vee}$, which is $\mathrm{Ext}^{n-i}(B, A \otimes \omega_X)$, so the map between the Hom spaces is an equivalence. □

3 Understanding $D^b\mathbf{Coh}(X)$

3.1 An exercise from early in the semester

You might remember an exercise from the second lecture revolving around the relation between the algebra $\mathrm{Ext}^\bullet(\mathcal{O}\oplus\mathcal{O}(1), \mathcal{O}\oplus\mathcal{O}(1))$ and the Kronecker quiver. Well now, we can go ahead and complete this exercise.

Proposition 3.1. *Let $X = \mathbb{P}^1_k$, and let $\mathcal{M} = \mathcal{O} \oplus \mathcal{O}(1)$. Then the graded algebra Ext$^\bullet(\mathcal{M}, \mathcal{M})$ is concentrated in degree 0, and is the path algebra of the Kronecker quiver.*

Proof. First, note that $\mathrm{Ext}^i(\mathcal{O}(l), \mathcal{O}(k)) = \mathrm{Ext}^i(\mathcal{O}, \mathcal{O}k - l) = H^i(X, \mathcal{O}(k - l))$. We've seen that for $i > 0$, this is zero unless $i = 1$ and $k - l = -2$. So if $k, l \in \{0, 1\}$, this is zero. Hence, the higher Ext groups vanish.

To compute $\mathrm{Ext}^0 = \mathrm{Hom}$, we recall that $\mathrm{Hom}(\mathcal{O}(l), \mathcal{O}(k))$ is a vector space of dimension $k - l + 1$ given by the homogeneous degree $k - l$ polynomials in two variables. It is easy to see that we get the path algebra for the Kronecker quiver: the trivial paths both correspond to multiplication by 1, and the two nontrivial paths are the morphisms from $\mathcal{O}$ to $\mathcal{O}(1)$ corresponding to multiplication by x_0 and x_1, respectively. □

Additionally, it turns out that the functor

$$D^b\mathrm{Coh}(\mathbb{P}^1) \to D^b\mathrm{Rep}(\cdot \rightrightarrows \cdot)$$

$$E \mapsto \mathrm{Ext}^\bullet(\mathcal{O} \oplus \mathcal{O}(1), E)$$

is an equivalence of categories. The utility of this statement is that now, instead of thinking about the derived category of sheaves on $\mathbb{P}^1$, we can instead just think of the derived category of modules over this algebra, and modules over an algebra are usually simpler to think about than the more general sheaves over a scheme.

$\mathcal{O}\oplus\mathcal{O}(1)$ is an example of what is called a *tilting sheaf* on $\mathbb{P}^1$. The derived endomorphism algebra of this sheaf essentially contains all of the complexity of the derived category of coherent sheaves on $\mathbb{P}^1$. We can generalize this theory substantially, and will now do so.

3.2 Tilting Sheaves

Let X be a smooth projective scheme over an algebraically closed field k, and let $\mathcal{D} = D^b\mathrm{Coh}(X)$, as usual.

Definition 3.2. *A coherent sheaf $\mathcal{T}$ on X is called a* **tilting sheaf** *if*

1. *(T1) $A := End_{\mathcal{O}_X}(\mathcal{T})$ has finite global dimension. That is, any module over A has a finite projective resolution.*
2. *(T2) $Ext^i_{\mathcal{O}_X}(\mathcal{T}, \mathcal{T}) = 0$ for $i > 0$.*
3. *(T3) $\mathcal{T}$ classically generates $\mathcal{D}$. That is, $\mathcal{D}$ has no nontrivial triangulated subcategory containing $\mathcal{T}$ that is closed under isomorphisms, shifts, taking cones of morphisms, and taking direct summands.*

Here, (T3) is the relevant property that $\mathcal{T}$ is 'complicated enough' that its endomorphism algebra contains all of the necessary information to understand coherent sheaves on X (in the derived world, at least).

It's also worth noting here that, in $D^b\mathrm{Coh}(X)$, sheaves supported on an affine scheme (hence, rank zero sheaves) are trivial. So it's enough to consider only *tilting bundles*, i.e., locally free tilting sheaves.

Theorem 3.3. *Let $\mathcal{T}$ be a tilting sheaf on a smooth projective scheme X, with associated tilting algebra $A = End_{\mathcal{O}_X}(\mathcal{T})$. Then the functors*

$$F(-) := Hom_{\mathcal{O}_X}(\mathcal{T}, -) : Coh(X) \to mod(A^{op})$$

and

$$G(-) := - \otimes_A \mathcal{T} : mod(A^{op}) \to Coh(X)$$

induce equivalences of triangulated categories

$$\mathbf{R}F : D^b Coh(X) \to D^b mod(A^{op})$$

$$\mathbf{L}G : D^b mod(A^{op}) \to D^b Coh(X)$$

that are quasi-inverse to each other.

Proof. (Sketch) It is clear that if $\mathcal{F}$ is a coherent sheaf on X, then $\mathrm{Hom}_{\mathcal{O}_X}(\mathcal{T},\mathcal{F})$ is a finitely generated right A-module (and thus a left A^{op}-module). On the other hand, if M is a finitely generated right A-module, we can construct $M \otimes_A \mathcal{T}$ as the sheaf whose stalk at p is $M \otimes_A \mathcal{T}_p$, and this is a coherent sheaf on X.

In order to prove the statement, we do in fact need to check that the induced functors land in the *bounded* derived categories (so that the derived functors above are well-defined). For $\mathbf{R}F$, this is where the smoothness of X comes in, because we need all but finitely many of the Ext groups to be zero. For $\mathbf{L}G$, this is where condition (T1) comes in, because we need A-modules to have finite projective resolutions.

Not surprisingly, the claim that these two functors are quasi-inverse is where we have to use conditions (T2) and (T3). Roughly, (T2) tells us that $\mathbf{R}F \circ \mathbf{L}G(A) = \mathbf{R}F(\mathcal{T}) = A$ (because the higher Ext groups from T to itself vanish), which tells us that $\mathbf{R}F \circ \mathbf{L}G$ is the identity on A. Hence, it is on finitely generated projective A-modules as well, and then because any coherent A-module has a finite projective resolution, it is the identity on all of $D^b\mathrm{mod}(A^{\mathrm{op}})$. In order to get that the other composite is quasi-isomorphic to the identity, we use (T3) to see that the image of $\mathbf{L}G$, which is the triangulated subcategory generated by $\mathcal{T}$, is all of $D^b\mathrm{Coh}(X)$. □

Corollary 3.4. *If $\mathcal{T}$ is a coherent sheaf on X satisfying (T1) and (T2), then $\mathcal{T}$ satisfies (T3) iff $\mathbf{L}G \circ \mathbf{R}F(\mathcal{F}) \simeq \mathcal{F}$ for all $\mathcal{F} \in D^b Coh(X)$.*

Example: If X is affine, $\mathcal{O}_X$ is a tilting sheaf. (If you don't see why this is true, take a minute to think about it: you should be able to say why in one line!)

Theorem 3.5.

$$\mathcal{O} \oplus \mathcal{O}(1) \oplus \cdots \oplus \mathcal{O}(n)$$

is a tilting sheaf on $X = \mathbb{P}^n$.

Proof. If we have $\mathcal{O} \oplus \cdots \oplus \mathcal{O}(n)$, then by taking direct summands, we get $\mathcal{O}, \ldots, \mathcal{O}(n)$. There is an exact sequence

$$0 \to \mathcal{O}^{\binom{n+1}{0}} \to \mathcal{O}(1)^{\binom{n+1}{1}} \to \cdots \to \mathcal{O}(n+1)^{\binom{n+1}{n+1}} \to 0$$

Hence, by taking cones, we can get $\mathcal{O}(n+1)$. Line bundles are flat, so we can tensor the above exact sequence by any $\mathcal{O}(d)$, and inductively, by taking cones, we can get any $\mathcal{O}(d)$ for $d \geq n+1$. Similarly, we can get the negative degree line bundles by taking cones and shifting. So we can get all line bundles. Any coherent sheaf can be resolved by sheaves that are sums of line bundles, so we can get all coherent sheaves, as desired. □

Exercise: Which quiver (with relations) has path algebra equal to the endomorphism algebra of this tilting sheaf?

3.3 Mukai Pairing

As we have seen, when a smooth projective scheme X has a tilting sheaf, $D^b\mathrm{Coh}(X)$ is equivalent to $D^b\mathrm{mod}(A^{\mathrm{op}})$, where A is the algebra of endomorphisms of the tilting sheaf. So now the question is, which X's have tilting sheaves? We'll provide a statement towards this end, by imposing additional structure on our categories and seeing that the equivalence we provided further preserves this structure. Moreover, we'll see that this allows us to understand $K(\mathrm{Coh}(X))$.

First, we define a sort of enriched version of the Euler characteristic.

Definition 3.6. *Let $\mathcal{A}$ be an abelian category. Then define a map*

$$[-] : D^b(\mathcal{A}) \to K(\mathcal{A})$$

by sending an object E to $[E] = \sum_{i\in\mathbb{Z}} (-1)^i [E^i]$.

Exercise: Check that this map is well-defined, and compatible with the additive structure.

Now note that $K(\mathrm{Coh}(X))$ (if X is smooth and projective) admits a bilinear form

$$\langle [E], [F] \rangle := \sum_i (-1)^i \dim_k \mathrm{Ext}^i_{\mathcal{O}_X}(E, F)$$

called the *Mukai pairing.* Since Ext^i is just H^i of the $\mathbf{R}\mathrm{Hom}$ object (or π_i of the mapping space, if you like), you can literally think of this as the Euler characteristic of the derived Hom complex. Similarly, since $A =$

$\mathrm{Hom}_{\mathcal{O}_X}(\mathcal{T},\mathcal{T})$ has finite global dimension (so that any pair of A-modules has a finite number of Ext groups), we have an inner product on $K(\mathrm{mod}(A^{\mathrm{op}}))$

$$\langle [M],[N]\rangle := \sum_i (-1)^i \dim_k \mathrm{Ext}^i_A(M,N)$$

Both of these inner products make more sense when imagined in the derived category, since they are the Euler characteristic of the respective $\mathbf{R}\mathrm{Hom}$ complexes (of k-vector spaces). Recall $\mathbf{R}F : D^b\mathrm{Coh}(X) \to D^b\mathrm{mod}(A)$ is an equivalence of categories, with inverse $\mathbf{L}G$. So if $A, B \in D^b\mathrm{Coh}(X)$, we get maps between $\mathrm{Hom}_{D^b\mathrm{Coh}(X)}(A,B)$ and $\mathrm{Hom}_{D^b\mathrm{mod}(A^{\mathrm{op}})}(\mathbf{R}F(A),\mathbf{R}F(B))$ which are inverse to one another (homotopy equivalences in the world of ∞-categories, quasi-isomorphisms in the world of derived categories). Hence, they must preserve the inner product on these derived categories. So, roughly, we should expect

Theorem 3.7. *Let $\mathcal{T}$ be a tilting sheaf on a smooth projective X, and let $A = End_{\mathcal{O}_X}(\mathcal{T})$. The homomorphism*

$$[\mathbf{R}F(-)] : K(Coh(X)) \to K(mod(A^{op}))$$

$$[\mathbf{R}F(\mathcal{F})] = \sum_i (-1)^i Ext^i_{\mathcal{O}_X}(\mathcal{T},\mathcal{F})$$

is an isomorphism of abelian groups that preserves the bilinear forms on either side.

Corollary 3.8. *If X smooth and projective admits a tilting sheaf, then its Grothendieck group is finitely generated and free.*

In particular, this tells us that any Bridgeland stability condition on $K(\mathrm{Coh}(X))$ factors through the isomorphism $K(\mathrm{Coh}(X)) \to \Gamma$, where Γ is a finite-dimensional lattice!

Interesting fact: $K(\mathrm{Coh}(\mathbb{P}^n)) \simeq \mathbb{Z}^{n+1}$. One way to see this is as follows: there is a homomorphism

$$K(\mathrm{Coh}(\mathbb{P}^n)) \to \frac{1}{n!}\mathbb{Z}[x]/(x^{n+1})$$

defined by sending a sheaf to its **Hilbert polynomial**. Indeed, if two (virtual) sheaves are equal in the Grothendieck group, they should have the same Hilbert polynomial (namely, the polynomial $p(N) = \dim(M_N)$ for N sufficiently large). It's also clear that this is surjective onto a lattice of full rank: we can generated sheaves with $n+1$ linearly independent Hilbert polynomials, and in fact, $\mathcal{O}, \mathcal{O}(1), \ldots, \mathcal{O}(n)$ work. The same argument we saw before shows that these $n+1$ line bundles generate $K(\mathrm{Coh}(\mathbb{P}^n))$, which implies that this map is a monomorphism. Hence, it is an isomorphism onto a lattice of full rank.

The upshot: any virtual sheaf in the Grothendieck group is uniquely defined by the associated Hilbert polynomial. In particular, any element of $K(\mathrm{Coh}(\mathbb{P}^1))$ is defined by **rank** and **degree**, each of which can be any integer.

3.4 Exceptional sequences

Exceptional sequences appear in algebraic geometry. We'll see that if a tilting sheaf has a decomposition as a sum of line bundles, then those line bundles form an exceptional sequence, and in this case, the endomorphism algebra is the path algebra of an easy to describe quiver with relations, where the vertices of this quiver correspond to the line bundles of the exceptional sequence. This then gives us a simple description of $K(\mathrm{Coh}(X))$.

Definition 3.9. *An object E in a triangulated k-linear category $\mathcal{D}$ is* **exceptional** *if*

$$Hom_{\mathcal{D}}(E, E[n]) \simeq \begin{cases} k & \text{if } n = 0 \\ 0 & \text{otherwise} \end{cases}$$

A sequence $(E_0, \ldots, E_m)$ of exceptional objects is called **exceptional** *if*

$$\mathbf{R}Hom_{\mathcal{D}}(E_i, E_j) = 0 \quad \text{if} \quad i > j$$

and **strongly exceptional** *if, additionally,*

$$Hom_{\mathcal{D}}(E_i, E_j[n]) = 0 \quad \text{if} \quad i < j \quad \text{and} \quad n \neq 0$$

Moreover, the sequence is called **full** *(or* **complete** *if $E_0, \ldots, E_m$ generate $\mathcal{D}$ classically.*

Proposition 3.10. *Let $\mathcal{T}$ be a locally free sheaf on X, and $T = \bigoplus_{i=0}^{m} \mathcal{E}_i$ a decomposition into distinct locally free sheaves with $Hom_{\mathcal{O}_X}(\mathcal{E}_i, \mathcal{E}_i) \simeq k$ (i.e., line bundles). Then*

1. *If $\mathcal{T}$ satisfies (T1) and (T2), then after reordering appropriately, $(\mathcal{E}_0, \ldots, \mathcal{E}_m)$ forms a strongly exceptional sequence.*
2. *If $\mathcal{T}$ additionally satisfies (T3), then $(\mathcal{E}_0, \ldots, \mathcal{E}_m)$ is a fully strongly exceptional sequence.*

Conversely, if we have any full strongly exceptional sequence $(\mathcal{E}_0, \ldots, \mathcal{E}_m)$, then $\mathcal{T} = \bigoplus_{i=0}^{m} \mathcal{E}_i$ is a tilting sheaf.

Proof. (Sketch) In order to prove the forwards direction, first note that

$$0 = \mathrm{Ext}^i_{\mathcal{O}_X}(\mathcal{T}, \mathcal{T}) \simeq \bigoplus_{i,j} \mathrm{Ext}^i_{\mathcal{O}_X}(\mathcal{E}_i, \mathcal{E}_j)$$

so all higher Ext spaces vanish. Each $\mathcal{E}_i$ is exceptional, because $\mathrm{Hom}_{\mathcal{O}_X}(\mathcal{E}_i, \mathcal{E}_i) \simeq k$. So we just need to show that we can rearrange the $\mathcal{E}_i$ so that there are no nontrivial morphisms backwards in the sequence. For every two indices i, j, one of the spaces $\mathrm{Hom}_{\mathcal{O}_X}(\mathcal{E}_i, \mathcal{E}_j)$ and $\mathrm{Hom}_{\mathcal{O}_X}(\mathcal{E}_j, \mathcal{E}_i)$ is trivial (because otherwise, every morphism $\mathcal{E}_i \to \mathcal{E}_i$ factors through $\mathcal{E}_j$, by dimensionality, and similarly for $\mathcal{E}_j$, which implies that the two are isomorphic). We can then accordingly order the two in the way that we want. Careful checking will show that this choice can be made consistently over all pairs.

In order to check the backwards direction (namely, that every full strongly exceptional sequence defines a tilting sheaf), the only nontrivial thing to check is that $A = \mathrm{End}_{\mathcal{O}_X}(\bigoplus_{i=0}^{m} \mathcal{E}_i)$ has finite global dimension. Because there are no morphisms backwards in the exceptional sequence, this algebra is isomorphism to an algebra of lower triangular matrices, which implies that A has finite global dimension. □

Proposition 3.11. *Let $(\mathcal{E}_0, \ldots, \mathcal{E}_m)$ be a collection of line bundles, and let $\mathcal{T} = \bigoplus_{i=0}^{m} \mathcal{E}_i$. Then $End_{\mathcal{O}_X}(T)$ is the path algebra of a quiver (Q, R), where*

1. *The vertices of Q correspond to $\mathcal{E}_0, \ldots, \mathcal{E}_m$.*

2. *The edges from $\mathcal{E}_i$ to $\mathcal{E}_j$ correspond to a basis for the vector space $Hom_{\mathcal{O}_X}(\mathcal{E}_i, \mathcal{E}_j)$.*

3. *Two paths are declared to be equal (i.e., their difference is in R) if the corresponding morphisms are equal.*

Exercise: Prove it! Note that this is not the only construction of the quiver: we could have more edges and more relations, so as to get the same path algebra.

Corollary 3.12. *If $(\mathcal{E}_0, \dots, \mathcal{E}_m)$ is a full strongly exceptional sequence for X, then $K(Coh(X))$ is freely generated by $\mathcal{E}_0, \dots, \mathcal{E}_m$.*

In particular, $(\mathcal{O}, \mathcal{O}(1), \dots, \mathcal{O}(n))$ form a full strongly exceptional sequence on $\mathbb{P}^n$!