18.726: Algebraic Geometry (K.S. Kedlaya, MIT, Spring 2009)
Dualizing sheaves and Riemann-Roch (updated 6 May 09)

In this lecture, we introduce dualizing sheaves for projective schemes over a field, then use them to derive the Riemann-Roch theorem for curves. Throughout, let k be a field (not necessarily algebraically closed), let $j : X \to P = \mathbb{P}^N_k$ be a closed immersion with X of dimension n, and let $\mathcal{O}_X(1)$ be the corresponding twisting sheaf.

1 Dualizing sheaves

For V a k-vector space, let $V^\vee$ denote the dual space $\mathrm{Hom}_k(V, k)$. A *dualizing sheaf* for X is a coherent sheaf ω_X° equipped with a *trace morphism* $t : H^n(X, \omega_X^\circ) \to k$, such that for all coherent sheaves $\mathcal{F}$ on X, the composition

$$\mathrm{Hom}_X(\mathcal{F}, \omega_X^\circ) \times H^n(X, \mathcal{F}) \to H^n(X, \omega_X^\circ) \xrightarrow{t} k$$

of the natural pairing with the trace morphism induces an isomorphism

$$\mathrm{Hom}_X(\mathcal{F}, \omega_X^\circ) \cong H^n(X, \mathcal{F})^\vee.$$

By interpreting this in terms of representing a certain functor, we see that a dualizing sheaf is unique up to unique isomorphism if it exists.

Theorem. *There exists a dualizing sheaf for X.*

This also holds when X is proper, but I won't give the proof in this course (see the references at the end of Hartshorne III.7).

The previous theorem is not so useful unless one can identify the dualizing sheaf explicitly. This is tricky in general, but can be done well in the smooth case.

Theorem. *Suppose that X is smooth and irreducible over k. Then the canonical sheaf ω_X is a dualizing sheaf.*

2 Application to Riemann-Roch

Modulo the previous two theorems, we can already deduce Riemann-Roch for curves. Suppose in this section that k is algebraically closed, and that X is smooth over k, irreducible, and of dimension 1.

For any divisor D on X, the identification of the canonical sheaf $\omega_X \cong \Omega_{X/k}$ with the dualizing sheaf ω_X° gives us an isomorphism

$$\begin{aligned} H^0(X, \omega_X \otimes \mathcal{L}(-D)) &\cong \mathrm{Hom}_X(\mathcal{L}(D), \omega_X) \\ &\cong \mathrm{Hom}_X(\mathcal{L}(D), \omega_X^\circ) \\ &\cong H^1(X, \mathcal{L}(D))^\vee. \end{aligned}$$

This tells us several useful things. First, the genus $g = g(X)$, which is typically defined as $\dim_k H^0(X, \omega_X)$, is also equal to $\dim_K H^1(X, \mathcal{O})$. Second, the desired statement of Riemann-Roch is now

$$\begin{aligned}\deg(D) + 1 - g &\stackrel{?}{=} \dim_k H^0(X, \mathcal{L}(D)) - \dim_k H^0(X, \omega_X \otimes \mathcal{L}(-D)) \\ &= \dim_k H^0(X, \mathcal{L}(D)) - \dim_k H^1(X, \mathcal{L}(D)) \\ &= \chi(X, \mathcal{L}(D)).\end{aligned}$$

Third, Riemann-Roch does indeed hold for $D = 0$ (by the previous two assertions).

To finish the proof, it is enough to show that the Riemann-Roch equality for a given divisor D is equivalent to its truth for the divisor $D + (Q)$ for any closed point $Q \in X(k)$. (With that in hand, we can walk from 0 to any other divisor by adding or subtracting points.) So let us see how much both sides of the Riemann-Roch equality change when we add the point Q. On one hand, obviously

$$(\deg(D + (Q)) + 1 - g) - (\deg(D) + 1 - g) = 1.$$

On the other hand, we have a short exact sequence

$$0 \to \mathcal{L}(D) \to \mathcal{L}(D + (Q)) \to \mathcal{O}_Q \to 0$$

where $\mathcal{O}_Q$ denotes the skyscraper sheaf k at the point Q. Since Euler characteristics add in short exact sequences,

$$\chi(X, \mathcal{L}(D + (Q))) - \chi(X, \mathcal{L}(D)) = \chi(X, \mathcal{O}_Q) = 1.$$

Hence Riemann-Roch for D is equivalent to Riemann-Roch for $D + (Q)$.

3 Construction of the dualizing sheaf

We now go back and construct dualizing sheaves following the argument in Hartshorne (but fleshing out some details which he leaves opaque). Recall that we already know the duality theorem for $X = P$, with the dualizing sheaf being the canonical sheaf ω_P. The plan is to manufacture a dualizing sheaf on X out of ω_P, using Serre duality for P. That tells us that if we fix an isomorphism $H^N(P, \omega_P) \cong k$ of k-vector spaces, then for any coherent sheaf $\mathcal{F}$ on X,

$$H^n(X, \mathcal{F}) = H^n(P, j_*\mathcal{F}) \cong \mathrm{Ext}_P^{N-n}(j_*\mathcal{F}, \omega_P)^\vee.$$

So we are reduced to finding a sheaf ω_X° on X for which there is a functorial isomorphism

$$\mathrm{Hom}_X(\mathcal{F}, \omega_X^\circ) \cong \mathrm{Ext}_P^{N-n}(j_*\mathcal{F}, \omega_P).$$

(We then get the required trace map $H^n(X, \omega_X^\circ) \to k$ by tracing the identity element of $\mathrm{Hom}_X(\omega_X^\circ, \omega_X^\circ)$ through the identifications.)

One might imagine that this isomorphism comes from an isomorphism of sheaves

$$\mathscr{H}om_X(\mathcal{F}, \omega_X^\circ) \overset{?}{\cong} \mathscr{E}xt_P^{N-n}(j_*\mathcal{F}, \omega_P)$$

by taking global sections. Taking $\mathcal{F} = \mathcal{O}_X$ in this hypothetical isomorphism suggests the right definition:

$$\omega_X^\circ = j^*\mathscr{E}xt_P^{N-n}(j_*\mathcal{O}_X, \omega_P).$$

We would like to get back from this to the claimed isomorphism

$$\mathrm{Hom}_X(\mathcal{F}, \omega_X^\circ) \cong \mathrm{Ext}_P^{N-n}(j_*\mathcal{F}, \omega_P).$$

by first forming the canonical identification

$$\mathscr{H}om_X(\mathcal{F}, j^*\mathscr{H}om_P(j_*\mathcal{O}_X, \cdot)) \cong \mathscr{H}om_P(j_*\mathcal{F}, \cdot)$$

(local version: for A a ring, I an ideal, $M \in \underline{\mathrm{Mod}}_{A/I}$, $N \in \underline{\mathrm{Mod}}_A$, we identify $\mathrm{Hom}_A(M, N) \cong \mathrm{Hom}_{A/I}(M, \mathrm{Hom}_A(A/I, N))$), then evaluating the resulting derived functors at ω_P, then taking global sections. This is complicated by the fact that in the second step, $\mathscr{H}om_X(\mathcal{F}, \cdot)$ is not exact, and in the third step, taking global sections is not exact.

To straighten these things out, we need to know more about the relationship between the sheaf $\mathscr{E}xt$ and the global Ext. For starters, here is one thing I can say using Serre vanishing. (See Hartshorne Proposition III.6.9.)

Proposition. *Let $\mathcal{F}$ and $\mathcal{G}$ be coherent sheaves on X. Then there exists an integer q_0 depending on $\mathcal{F}$ and $\mathcal{G}$, such that for every $q \geq q_0$, we have*

$$\mathrm{Ext}_X^i(\mathcal{F}, \mathcal{G}(q)) \cong \Gamma(X, \mathscr{E}xt_X^i(\mathcal{F}, \mathcal{G})(q)).$$

Proof. This holds for $i = 0$ without any restriction on q. For $\mathcal{F}$ locally free, the right side is zero for $i > 0$, whereas the left side vanishes for n large enough by Serre's vanishing theorem. The general case then follows by a dimension shifting argument; see Hartshorne Proposition III.6.9. □

Next, I must recall a theorem which I skipped over earlier.

Theorem (Grothendieck)**.** *For any $\mathcal{F} \in \underline{\mathrm{Sh}}_{\underline{\mathrm{Ab}}}(X)$, $H^i(X, \mathcal{F}) = 0$ for $i > n$.*

Proof. This holds with X replaced by any noetherian topological space of dimension n. The argument is a rather elaborate dimension-shifting argument; see Hartshorne Theorem III.2.7. (See also Hartshorne exercise III.4.8(d), which is enough for this discussion.) □

Corollary. *For any coherent sheaf $\mathcal{F}$ on X, we have $\mathscr{E}xt_P^i(j_*\mathcal{F}, \omega_P) = 0$ for $i < N - n$.*

Proof. Put $\mathcal{F}_i = \mathscr{E}xt_P^i(j_*\mathcal{F}, \omega_P)$. On one hand, for q large,

$$\Gamma(P, \mathcal{F}_i(q)) = \mathrm{Ext}_P^i(j_*\mathcal{F}, \omega_P(q)) \cong H^{N-i}(P, j_*\mathcal{F}(-q))^\vee$$

by Serre duality for P. For $i < N - n$, $H^{N-i}(P, j_*\mathcal{F}(-q)) = 0$ by the theorem. Hence $\Gamma(P, \mathcal{F}_i(q)) = 0$ for q large. On the other hand, since $\mathcal{F}_i$ is coherent, for q large, $\mathcal{F}_i(q)$ is generated by global sections. This forces $\mathcal{F}_i(q) = 0$ for q large, whence $\mathcal{F}_i = 0$. □

At this point, we can finish with the following argument; compare Hartshorne Lemma III.7.4. (Once again, there is a spectral sequence hiding behind this, but never mind.) Take an injective resolution $\mathcal{I}^{\cdot}$ of ω_P, so we can compute $\mathscr{E}xt^{\cdot}(j_*\mathcal{F}, \omega_P)$ as the cohomology of $\mathscr{H}om_P(j_*\mathcal{F}, \mathcal{I}^{\cdot})$, and similarly for Ext and Hom. But using the canonical identification from earlier, if we write $\mathcal{J}^{\cdot} = j^*\mathscr{H}om_P(j_*\mathcal{O}_X, \mathcal{I}^{\cdot})$, we can also compute $\mathscr{E}xt^{\cdot}(j_*\mathcal{F}, \omega_P)$ as the cohomology of $\mathscr{H}om_X(\mathcal{F}, \mathcal{J}^{\cdot})$, and similarly for Ext and Hom. So now what we need to know is that

$$\mathscr{H}om_X(\mathcal{F}, \omega_X^\circ) \overset{?}{\cong} h^{N-n}(\mathscr{H}om_X(\mathcal{F}, \mathcal{J}^{\cdot}))$$

and similarly with straight Homs.

However, the sheaves $\mathcal{J}^{\cdot}$ are *injective* $\mathcal{O}_X$-modules. (Local version: if A is a ring, I an ideal, and $I \in \underline{\mathrm{Mod}}_A$ is injective, then $\mathrm{Hom}_A(A/I, M)$ is an injective A/I-module; this uses the previous local identification.) By the previous corollary, the complex $\mathcal{J}^{\cdot}$ (whose cohomology computes $\mathscr{E}xt^{\cdot}(j_*\mathcal{O}_X, \omega_P)$) is acyclic in degrees up to $N-n-1$. We can then split it into *two* complexes of injectives $\mathcal{J}_1^{\cdot}, \mathcal{J}_2^{\cdot}$, where $\mathcal{J}_1^{\cdot}$ is exact and only has terms in degrees up to $N-n$, and $\mathcal{J}_2^{\cdot}$ only has terms in degrees at least $N-n$ (exercise).

Since $\mathcal{J}_1^{\cdot}$ is a bounded complex of injectives, it splits into a series of split short exact sequences; thus it stays exact no matter what left exact functors you apply to it. So we can replace $\mathcal{J}$ by $\mathcal{J}_2^{\cdot}$ for the purposes of computing derived functors, i.e., what we need to prove is reduced to

$$\mathscr{H}om_X(\mathcal{F}, \omega_X^\circ) \overset{?}{\cong} h^{N-n}(\mathscr{H}om_X(\mathcal{F}, \mathcal{J}_2^{\cdot}))$$

and similarly for straight Hom. But $\mathcal{J}_2^{\cdot}$ only starts in degree $N-n$, and Hom and $\mathscr{H}om$ are left exact, so we have

$$\begin{aligned}
\mathscr{E}xt_P^{N-n}(j_*\mathcal{F}, \omega_P) &\cong h^{N-n}(\mathscr{H}om_X(\mathcal{F}, \mathcal{J}_2^{\cdot})) \\
&\cong \mathscr{H}om_X(\mathcal{F}, h^{N-n}(\mathcal{J}_2^{\cdot})) \\
&\cong \mathscr{H}om_X(\mathcal{F}, h^{N-n}(\mathscr{H}om_X(\mathcal{O}_X, \mathcal{J}_2^{\cdot}))) \\
&\cong \mathscr{H}om_X(\mathcal{F}, \mathscr{E}xt^{N-n}(j_*\mathcal{O}_X, \omega_P)) \\
&\cong \mathscr{H}om_X(\mathcal{F}, \omega_X^\circ)
\end{aligned}$$

and similarly

$$\mathrm{Ext}_P^{N-n}(j_*\mathcal{F}, \omega_P) \cong h^{N-n}(\mathrm{Hom}_X(\mathcal{F}, \mathcal{J}_2^{\cdot})) \cong \mathrm{Hom}_X(\mathcal{F}, \omega_X^\circ).$$

That completes the proof that

$$\omega_X^\circ = \mathscr{E}xt_P^{N-n}(j_*\mathcal{O}_X, \omega_P)$$

is a dualizing sheaf for X.

4 Calculation of the dualizing sheaf for smooth schemes

To finish the proof of Riemann-Roch, we must still show that we can take $\omega_X^\circ = \omega_X$ when X is smooth over k. Fortunately, this is a local problem.

Theorem. *Suppose that X is a local complete intersection in P. Let $\mathcal{I}$ be the ideal sheaf of X. Then there is a canonical isomorphism*

$$\mathscr{E}xt^r_P(j_*\mathcal{O}_X, \omega_P) \cong \omega_P \otimes j_*\mathcal{O}_X \otimes \wedge^r(\mathcal{I}/\mathcal{I}^2)^\vee.$$

The local complete intersection condition asserts that $\mathcal{I}$ is locally generated by $N - n$ elements; this is true for X smooth basically by the Jacobian criterion. See Hartshorne Theorem II.8.17. The fact that the right side gives ω_X comes from the exact sequence

$$0 \to \mathcal{I}/\mathcal{I}^2 \to \Omega_{P/k} \otimes j_*\mathcal{O}_Y \to j_*\Omega_{Y/k}$$

by taking exterior powers; see Hartshorne Proposition II.8.20. The stated theorem itself is proved by computing in local coordinates; see Hartshorne Theorem III.7.11