BABY ALGEBRAIC GEOMETRY SEMINAR: AN ALGEBRAIC PROOF OF RIEMANN-ROCH

RAVI VAKIL

Contents

1. Introduction

I'm going to present an algebraic proof of Riemann-Roch. This is a hefty task, especially as I want to say enough that you can genuinely fill in all the details yourself. There's no way I can finish in an hour without making this less self-contained, so what I'll do is explain a bit about cohomology, and reduce Riemann-Roch to Serre duality. That should be do-able within an hour. Then I'll give people a chance to leave, and after that I'll prove Serre duality in the following half-hour.

Also, these notes should help. And you should *definitely* stop me and ask questions. For example, if I mention something I defined last semester in my class, and you'd like me to refresh your memory as to what the definition was, please ask.

The proof I'll present is from Serre's *Groupes algébriques et corps de classes*, Ch. 2, [S]. I found out this past Tuesday that this proof is originally due to Weil.

Throughout, C is a non-singular projective algebraic curve over an algebraically closed field $\overline{k}$.

Date: Friday, February 11, 2000.

2. Cohomology of sheaves

Definition. If S is a sheaf on C, define $H^0(C,S)$ to be the set of global sections of S over C. If S is an $\mathcal{O}_C$-module, then $H^0(C,S)$ has the structure of a $\overline{k}$-vector space. In this situation, define $h^0(C,S)$ to be the dimension of $H^0(X,S)$ as a vector space. (I made this definition in the last class of last semester.)

I'll now define $H^1(C,S)$ (Cech cohomology), a little loosely. Here again, S is an $\mathcal{O}_C$-module. Elements of $H^1(C,S)$ are given by the following data. Let $U_1, \dots,$ U_n be an open cover of C. Let $U_{ij} = U_i \cap U_j$, $U_{ijk} = U_i \cap U_j \cap U_k$ for convenience. Then the data $(f_{ij} \in H^0(U_{ij}, S))$ satisfying $f_{ij} - f_{jk} + f_{ik} = 0$ in $H^0(U_{ijk}, S)$ (called a *cocycle*) gives an element of $H^1(C,S)$.

This element is declared to be 0 if there are sections $g_i \in H^0(U_i, S)$ such that $f_{ij} = g_i - g_j$ in $H^0(U_{ij}, S)$, and in general 2 cocyles give the same element of H^1 if their difference is 0.

If you take a finer partition than U_i, and take the corresponding cocycle, then this is declared to be the same element of cohomology.

You can clearly add cocyles given using the same covering (describe); and any two coverings have a common refinement, so you can add *any* two cocycles.

The proper way to define $H^1(C,S)$ is to take the direct limit over all coverings.

You can actually define H^2, H^3, etc. in an analogous way, but we won't need them here. It is a fact (due to Grothendieck, see [H] Theorem III.2.7 for the pretty proof) that $H^i(C,S) = 0$ for all $i > 1$ (and more generally if X is a noetherian topological space of dimension n, then for all sheaves of abelian groups S on X and all $i > n$, $H^i(X,F) = 0$).

Define $\chi(C,S) = h^0(C,S) - h^1(C,S)$. (It would be better to define $\chi(C,S) = \sum_{i \geq 0} h^i(C,S)$, but that would make this proof longer.)

Example of a sheaf with no higher cohomology: a constant sheaf. Let G be an abelian group. Suppose $\underline{G}$ is a *constant sheaf* over C, whose sections over any open set of C is G, with the restriction map being the identity.

Easy Exercise. Check from the definition that $H^1(C,G) = 0$.

Another example: a skyscraper sheaf. If P is a given point of C, define a sheaf $\overline{k}_P$ as follows. The sections of $\overline{k}_P$ over U (i.e. $\overline{k}_P(U)$) are 0 if $P \notin U$, and $\overline{k}$ if $P \in U$. This is clearly a sheaf. We make it an $\mathcal{O}_C$-module in the only reasonable way we could: any section of $\mathcal{O}_C(U)$ (i.e. function on U) has a value at P, so the action of this section on the vector space $\overline{k} = \overline{k}_P(U)$ is multiplication by this value.

We immediately have $H^0(C, \overline{k}_P) = \overline{k}$.

Easy Exercise. $H^1(C,S)=0$.

Hence $\chi(\overline{k}_P)=h^0(C,\overline{k}_P)-h^1(C,\overline{k}_P)=1-0=1$.

Short exact sequences of sheaves. Suppose F, G, H are sheaves of $\mathcal{O}_C$-modules on C. A morphism of sheaves $F\to G$ is given by morphisms $F(U)\to G(U)$ that agree with restriction maps:

$$\begin{array}{ccc} F(U) & \rightarrow & G(U) \\ \downarrow & & \downarrow \\ F(V) & \rightarrow & G(V) \end{array}$$

commutes. It is easy to check that a morphism of sheaves induces morphisms of all the stalks. Then we say that $F\to G\to H$ is *exact* at G if the morphism of stalks is exact. (This isn't the best way to say it, but it'll work.)

Lemma/Exercise. Suppose $0\to F\to G\to H\to 0$ is a short exact sequence of sheaves of $\mathcal{O}_C$-modules on C. Then the following sequence is exact.

$$0\to H^0(C,F)\to H^0(C,G)\to H^0(C,H)\to H^1(C,F)\to H^1(C,G)\to H^1(C,H).$$

This just involves diagram chasing. As you have probably guessed, this sequence continues with H^2's etc.

Example: an important short exact sequence of sheaves.

Let P be some point on C. Let $\mathcal{O}_C(-P)$ be the sheaf of functions vanishing on P; in other words, sections of $\mathcal{O}_C(-P)$ over an open set U are those functions on U vanishing on P.

Then there is a short exact sequence of sheaves

$$0\to\mathcal{O}_C(-P)\to\mathcal{O}_C\to\overline{k}_P\to 0.$$

(Describe the morphisms.)

Similarly, if $\mathcal{L}$ is any invertible sheaf, there is a short exact sequence of sheaves

$$0\to\mathcal{L}(-P)\to\mathcal{L}\to\overline{k}_P\to 0.$$

Again, sections of $\mathcal{L}(-P)$ over an open set U are sections of $\mathcal{L}$ that vanish at P.

We played around with these objects a fair bit in class, and we saw for example that $\mathcal{L}(-P)$ was an invertible sheaf too.

Taking the long exact sequence associated to that short exact sequence, we get:

$$0\to H^0(C,\mathcal{L}(-P))\to H^0(C,\mathcal{L})\to H^0(C,\overline{k}_P)$$

$$\to H^1(C,\mathcal{L}(-P))\to H^1(C,\mathcal{L})\to H^1(C,\overline{k}_P)=0.$$

In particular, $\chi(C,\mathcal{L}(-P))=\chi(C,\mathcal{L})-\chi(C,\overline{k}_P)=\chi(C,\mathcal{L})-1$.

Lemma (Cheap Riemann-Roch). $\chi(C,\mathcal{L}) = \chi(C,\mathcal{O}_C) + \deg(\mathcal{L})$. Written explicitly (with $\mathcal{L} = \mathcal{O}_C(D)$):

$$h^0(C,\mathcal{O}(D)) - h^1(C,\mathcal{O}(D)) = d + 1 - h^1(C,\mathcal{O}_C).$$

We'll use this later in the proof of Serre duality.

Proof. Remember that every invertible sheaf was of the form $\mathcal{O}_C(p_1 + \cdots + p_a - q_1 - \cdots - q_b)$ for some points $p_1, \ldots, q_b$, where $a - b = \deg \mathcal{L}$. Then just do it by induction on $a + b$. (Do the first step for them.) □

That's it for the background.

3. Statements of Riemann-Roch and Serre Duality; Riemann-Roch from Serre Duality

Recall the invertible sheaf of differentials Ω^1.

The Riemann-Roch Theorem (for nonsingular projective algebraic curves over an algebraically closed field). $\mathcal{L}$ a degree d invertible sheaf on C. Then $h^0(C,\mathcal{L}) - h^0(C,\Omega^1 \otimes \mathcal{L}^\vee) = d - g + 1$.

I'll assume that you have some idea as to why this is such an incredibly important result. It is actually just the smallest case of important Riemann-Roch-type theorems. The hard part of the proof requires:

Serre duality (curve case). There is a natural perfect pairing $H^0(C,\Omega^1 \otimes \mathcal{L}^\vee) \times H^1(C,\mathcal{L}) \to \overline{k}$. Hence $h^1(C,\mathcal{L}) = h^0(C,\Omega^1 \otimes \mathcal{L}^\vee)$.

The form we will prove is: there is a natural perfect pairing $H^0(C,\Omega^1_C(-D)) \times H^1(C,\mathcal{O}_C(D)) \to \overline{k}$ (leave on board).

Proving Riemann-Roch using Serre duality. Recall that genus g of C was defined as $g = h^0(C,\Omega^1_C)$.

$$\begin{aligned}
h^0(C,\mathcal{L}) - h^0(C,\Omega^1 \otimes \mathcal{L}^\vee) &= h^0(C,\mathcal{L}) - h^1(C,\mathcal{L}) \\
&= \chi(C,\mathcal{L}) \\
&= d + \chi(C,\mathcal{O}_C) \\
&= d + h^0(C,\mathcal{O}_C) - h^1(C,\mathcal{O}_C) \\
&= d + 1 - h^0(C,\Omega^1_C) \\
&= d + 1 - g.
\end{aligned}$$

□

So we're left with proving Serre duality.

Let people leave!

4. Proof of Serre duality

4.1. **Repartitions, $I(D)$, and $J(D)$.** I've divided the proof into several parts. I think this first section is the hardest, so if it gets heavy, don't lose heart! The main idea of this proof is to interpret $H^1(C, \mathcal{O}(D))$ in the language of *repartitions*, or *adeles*. These are quite strange objects the first time you see them!

A repartition is an indexed set $\{r_P\}_{P \in C}$ where r_P is an element of $k(C)$, and $r_P \in \mathcal{O}_P$ for all but finitely many P. This is a (hideously huge) ring; call it R. Notice that $k(C)$ is naturally a subring of R (recall that an element of $k(C)$ is regular at all but finitely many points of C). R is also a $k(C)$-algebra.

If D is a divisor on C, define $R(D)$ as the (additive) subgroup of repartitions consisting of $\{r_P\}_{P \in C}$ where $v_P(r_P) + v_P(D) \geq 0$. (This is analogous to $\mathcal{O}_C(D)$ being the sheaf of rational functions where, at each point, the valuation of the function plus the valuation of D is non-negative.) Note that if $D' \geq D$ then that $R(D) \subset R(D')$; informally, "the bigger D is, the bigger $R(D)$ is".

(Pause for questions.)

Proposition. $I(D) := H^1(C, \mathcal{O}_C(D)) \cong \frac{R}{R(D)+k(C)}$ (canonically). (Write $I(D)$ with Serre duality statement.)

Proof. Let $\underline{k(C)}$ be the constant sheaf whose sections are $k(C)$. There's a natural injection $\mathcal{O}_C(D) \hookrightarrow \underline{k(C)}$; let S be the cokernel.

We have a short exact sequence

$$0 \to \mathcal{O}_C(D) \to \underline{k(C)} \to S \to 0$$

so we have the long exact sequence in cohomology, which gives us:

$$k(C) \to H^0(C, S) \to H^1(C, \mathcal{O}_C(D)) \to H^1(C, \underline{k(C)}) = 0$$

To complete the proof of the proposition, we just need to check that the global sections of S are $R/R(D)$.

To do this, we'll need to know what a quotient sheaf really is; I didn't do it in my class, but I think those who are taking Johan's followup class may have seen it.

The stalk S_P of S at a point P is the quotient of the stalks of $\underline{k(C)}$ by stalks of $\mathcal{O}_C(D)$, which is $k(C)$ modulo those functions with valuation at P at least $-v_P(D)$. This is the "P-part" of $R/R(D)$, and $R/R(D) = \oplus_P S_P$. Let T_P be the skyscraper sheaf at P, with values in this stalk S_P. We'll soon see that $S = \oplus_P S_P$. In other words, a section of S is a selection of values of S_P over all P, where almost all choices are 0. Then you can check that the global sections of S are indeed $\oplus_P S_P = R/R(D)$.

The way we'll check that S really is the direct sum of skyscraper sheaves, is by showing that if you're given a point $P \in C$, and a local section s of S defined on a neighborhood U, there is a smaller neighborhood U' such that the section s vanishes on $U' \setminus P$. This proof is getting long, so perhaps this last statement is best left as an exercise: given a section s in the stalk of S_P, it has a lift (in some neighbourhood) to a section s' of $\underline{k(C)}$, and this lift is an element of $k(C)$. Let U' be a smaller neighbourhood of P away from $\{\operatorname{supp} D \setminus P\}$ and also $\{$ the poles of $s' \setminus P\}$. □

That was the first of two tricky parts.

Define $J(D) := I(D)^* = (R/(R(D)+k(C)))^*$. (Add to Serre duality statement.) An element of $J(D)$ is a linear form on R, that is 0 on $R(D)$ and $k(C)$. Now if $D' \geq D$, $R(D) \subset R(D')$, so $J(D') \subset J(D)$. Let $J := \cup_D J(D)$. (Elements of J are linear functionals on R that vanish on $k(C)$, and also on *some* $R(D)$.)

We next show that J is a $k(C)$-vector space. Suppose $f \in k(C)$, $\alpha \in J$. Consider $f\alpha : R \to \overline{k}$, $r \mapsto \langle \alpha, fr \rangle$. This is a linear functional on R, vanishing on $k(C)$. If $\alpha \in J(D)$, and $(f) = \Delta$ then $f\alpha$ vanishes on $R(D-\Delta)$ ($r \in R(D-\Delta) \Rightarrow fr \in R(D) \Rightarrow \langle \alpha, fr \rangle = 0$), so it belongs to $J(D-\Delta)$, and hence it belongs to J. Hence $(f, \alpha) \to f\alpha$ gives J the structure of a $k(C)$-vector space.

Proposition. $\dim_{k(C)} J \leq 1$.

Proof. (I'll skip this argument if I'm short on time, which I likely will be.) Otherwise, suppose α and α' are 2 elements of J that are linearly independent (over $k(C)$). Suppose $\alpha, \alpha' \in J(D)$, and let $d = \deg(D)$. Suppose Δ_n is any divisor of degree n (with n to be declared later, large enough to give a contradiction). Then for any $f \in H^0(C, \mathcal{O}(\Delta_n))$, $f\alpha \in J(D-\Delta_n)$, by the argument above, and ditto with f replaced by g and α replaced by α'. Now as α and α' are linearly independent over $k(C)$, we know that $f\alpha + g\alpha' = 0 \Rightarrow f = g = 0$. Thus the map

$$(f, g) \mapsto f\alpha + g\alpha'$$

is an injection of the direct sum $H^0(C, \mathcal{O}(\Delta_n)) \oplus H^0(C, \mathcal{O}(\Delta_n))$ into $J(D-\Delta_n)$, so we have the inequality $\dim_{\overline{k}} J(D-\Delta_n) \geq 2h^0(C, \mathcal{O}(\Delta_n))$.

Now let's estimate both sides.

The left side is $\dim I(D-\Delta_n) = H^1(C, \mathcal{O}_C(D-\Delta_n))$. By "cheap" Riemann-Roch, this is $H^0(C, \mathcal{O}_C(D-\Delta_n)) - (d-n) + \text{constant} = n + \text{constant}$ if we pick n big enough so the degree $d-n$ of $D-\Delta_n$ is negative. (By constant, I mean that it doesn't depend on n or Δ_n.)

By cheap Riemann-Roch, the right side is twice at least $\deg(\Delta_n) + \text{constant}$, which is at least $2n + \text{constant}$. When n is chosen to be huge, there's no way in which the left side can be bigger than the right. □

That was the second tricky point.

4.2. **Differentials enter the picture.** Recall the invertible sheaf of differentials Ω^1_C; also the invertible sheaf $\Omega^1_C(D)$. Let M be the set of meromorphic differentials; note that it is a one-dimensional $k(C)$-vector space. Recall that one can define a residue map $\text{res}_P : M \to \overline{k}$. It vanishes on any differential that has no pole at P.

Definitions / Proofs that res_P *is well-defined.* 1) If $\overline{k} = \mathbb{C}$, you can use the complex analytic definition and proof. 2) Otherwise, you can write it locally as $(a_{-n}/t^n + \cdots + a_{-1}/t)dt+$ something regular, where t is a uniformizing parameter, define the residue as a_{-1}, and show that this definition is independent of t. If the characteristic is 0, this latter step isn't too hard, but if the characteristic is positive, it can be tricky; see [S] for a proof, and [T] for a nicer proof.

Kiran also told me a slick proof that a_{-1} is indepenendent of choice of uniformizer t that I'll put in the notes: the fact that $\sum_{i=-n}^{-1} a_n t^n dt$ has the same (-1)-coefficient as $\sum_{i=-n}^{-1} b_n u^n du$ for $t = u + \sum_{k=2}^{\infty} c_k u^k$ is (for a fixed choice of the pole order of the differential in question) a polynomial identity in the a's and b's with coefficients in $\mathbb{Z}$. It holds identically over $\mathbb{C}$ by the analytic proof 1), so it holds in every field.

Residue Theorem. For every meromorphic differential $\omega \in M$, $\sum_{P \in C} \text{res}_P(\omega) = 0$.

Proofs. 1) If $\overline{k} = \mathbb{C}$, Stokes' theorem works. 2) In characteristic 0, it isn't hard. 3) In positive characteristic, see [S] or [T]. □

4.3. **Proving duality.** For every meromorphic differential $\omega \in M$, define the divisor $(\omega) = \sum_{P \in C} v_P(\omega) P$, so $\Omega^1(-D)$ is "the sheaf of differentials satisfying $(\omega) \geq D$". Next define a pairing $\langle \omega, r \rangle$ between meromorphic differentials ω and repartitions r given by $\langle \omega, r \rangle = \sum_{P \in C} \text{res}_P(r_P \omega)$. (This is well-defined — only a finite number of terms in the sum are non-zero.)

Note that:

a) $\langle \omega, r \rangle = 0$ if $r \in k(C)$ (Residue theorem).
b) $\langle \omega, r \rangle = 0$ if $r \in R(D)$ and $\omega \in H^0(C, \Omega^1(-D))$ (as then $r_P \omega$ has no pole at P for all $P \in C$).
c) If $f \in k(C)$, then $\langle f\omega, r \rangle = \langle \omega, fr \rangle$. (Both are $\sum_{P \in C} \text{res}_P(f \omega r)$.)

Hence for each differential ω in $H^0(C, \Omega^1(-D))$, this defines a linear functional $\theta(\omega)$ on $R/(R(D) + k(C))$. $\theta : H^0(C, \Omega^1(-D)) \to J(D)$. (Also $\theta : M \to J$.)

Lemma. If ω is a meromorphic differential such that $\theta(\omega) \in J(D)$, then $\omega \in \Omega^1(-D)$.

Proof. Otherwise, assume $\omega \notin \Omega^1(D)$. We'll get a contradiction, and find an element of $R(D)$ that $\theta(\omega)$ doesn't kill.

There's some point $P \in C$ such that $v_P(\omega) < v_P(-D)$ (i.e. ω has a bigger pole than allowed by D). Define a repartition $r \in R(D)$ by $r_Q = 0$ for $Q \neq P$, and $r_P = 1/t^{v_P(\omega)+1}$ (where t is some uniformizer). Then as $v_P(r_P\omega) = -1$, $\langle \omega, r \rangle = \sum_Q \operatorname{res}(r_Q\omega) \neq 0$, proving the lemma. □

Theorem (Serre duality). θ gives an isomorphism from $H^0(C, \Omega^1(-D))$ to $J(D) = H^1(C, \mathcal{O}_C(D))^*$.

Proof. First, θ is injective. Reason: if $\theta(\omega) = 0$, then by the lemma, $\omega \in \Omega(-\Delta)$ for every Δ, so $\omega = 0$. (Explain.)

Next, θ is surjective. By c), θ is a $k(C)$-linear map from M to J. As M has dimension 1 (as a $k(C)$-vector space), and J has dimension at most 1 (earlier Proposition), we get surjectivity. □

That's it! The proof at first appears to be sleight of hand, but there's a lot going on under the surface that one eventually finds quite enlightening.

I hope you're not too shelshocked; if there are any parts of the argument I can elaborate on, please let me know. Thank you!

References

[F] W. Fulton, *Algebraic curves*. He gives a low-tech proof, but I didn't find it enlightening. (The first five chapters are great.)

[H] R. Hartshorne, *Algebraic geometry*. His proof of duality is quite general, but you don't need that big machinery if you're just dealing with curves.

[S] J.-P. Serre, *Groupes algébriques et corps de classes*, see esp. Ch. II (on algebraic curves), pp. 17–35, and 76–81. Serre is, of course, a god. Caution: his $\Omega^1(D)$ is what people would now call $\Omega^1(-D)$.

[T] J. Tate, *Residues of differentials on curves*, Ann. Scient. Éc. Norm. Sup., 4e série, t. 1, 1968, 149–159. He gives a nice characteristic-free definition of residues of differentials on curves. His proof of invariance of the definition (in positive characteristic) is an improvement on [S].