The Canonical Sheaf

Stefano Filipazzi

September 14, 2015

These notes are supposed to be a handout for the student seminar in algebraic geometry at the University of Utah.

In this seminar, we will go over some of the main and astonishing properties of the canonical sheaf of an algebraic variety.

The genus

It is well known that for a closed and orientable topological surface S, the genus represents what one intuitively would call "the number of handles". This concept has many interpretations and ways to be defined. From a topological perspective, we can make this concept more formal by talking about the rank of the free abelian group $H^1(S, \mathbb{Z})$, which is twice of the genus of S. Alternatively, one has that the genus g can be expressed in terms of the Euler characteristic $\chi(S)$ as $\chi(S) = 2-2g$. Another equivalent way can be found looking at the holomorphic (algebraic) structure such a surface can be endowed with.

Definition. Given a Riemann surface C, we call the abelian differential a global $(1,0)$-form on C, i.e. a differential one-form that locally can be written as $f(z)dz$, with f a holomorphic function. We call geometric genus, denoted by p_g, the dimension of the $\mathbb{C}$-vector space given by the abelian differentials.

The aim of this section is to show that the intuitive concept of geometric genus and the algebraic one agree. First, we compute a couple of examples.

The Riemann sphere

The Riemann sphere is obtained from two copies of $\mathbb{C}$, say $(\mathbb{C}, z)$ and $(\mathbb{C}, w)$, with the identification $z = w^{-1}$ on $\mathbb{C} \setminus \{0\}$. Assume ω is an abelian differential. On the chart $(\mathbb{C}, z)$ has the form $a(z)dz$, with

$$a(z) = \sum_{i \geq 0} a_i z^i; \tag{1}$$

on the other hand, on the chart $(\mathbb{C}, w)$, it has to look like

$$\omega = \sum_{i \geq 0} a_i w^{-i} d(w^{-1}) = \sum_{i \geq 0} -a_i w^{-i-2}. \tag{2}$$

For ω not to have a pole at ∞, i.e. at $w = 0$, we need $a_i = 0$ for all i. This shows that there are no nontrivial abelian differentials on the Riemann sphere.

Elliptic curves

An elliptic curve E is given by the quotient of $\mathbb{C}$ by a full-rank lattice Λ, i.e. the group generated over $\mathbb{Z}$ by two vectors linearly independent over $\mathbb{R}$. Since dz is invariant under the action of the lattice (indeed dz is translation invariant), it descends to an abelian differential on E. By abuse of notation, we will denote it as dz. Now, assume ω is another abelian differential on E; since dz is nowhere vanishing, ω/dz defines a holomorphic function on E. By the maximum principle, it is constant. Thus, we have just proved that dz spans the space of abelian differentials on E.

Higher genus: the hyperelliptic case

Before considering the next example, note that the idea of "quotient of two abelian differentials" sketched in the elliptic case can be still useful. Given a Riemann surface C and two abelian differentials α and β, the quotient α/β is a well-defined meromorphic function. Indeed, on a chart we have $\alpha = a(z)dz$ and $\beta = b(z)dz$, so $\alpha/\beta = a(z)/b(z)$; furthermore, the coefficients that pop out in front of dz in a change of chart cancel out, resulting in the local descriptions of α/β gluing together. Now, the residue theorem and Stokes' theorem tell us that on a Riemann surface we have

$$\sum_{P\in C} res_P\omega = 0 \tag{3}$$

for any meromorphic $(1,0)$-form on C. In particular, if we consider the logarithmic differential $d\log(f) = \frac{df}{f}$ of a meromorphic function, it tells us that the number of zeroes (with multiplicity) has to conincide with the number of poles (with multiplicity). Since in α/β the zeroes come from the ones of α and the poles come from the zeroes of β, we have that two abelian differentials on a Riemann surface C have the same number of zeroes (with multiplicity).

Now, we will sketch a way to construct a hyperelliptic curve of a given genus (in the sense of "number of handles"). Given a polynomial $F(y)$ with $2n$ distinct roots, consider the locus $Z \subset \mathbb{C}^2$ defined by $x^2 = F(y)$. Since the roots of $F(y)$ are all distinct, Z is smooth (it follows from the fact that on Z the partials of $x^2 - F(y)$ cannot vanish simultaneously). Say $F(y) = \prod_{i=1}^{2n}(y - \xi_i)$. Now, in the y-plane, partition the roots of F in pairs and join them with a segment (a path, in general). For simplicity, say that we pair them in pairs of the form ξ_{2j}, ξ_{2j+1}.

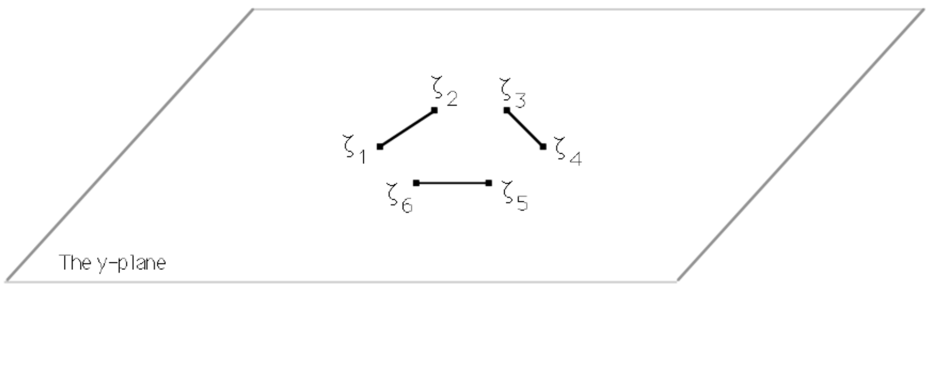

Now, the function $(y - \xi_i)(y - \xi_{i+1})$ has two holomorphic square roots on the complement of the segment $\overline{\xi_i \xi_{i+1}}$ (it follows from the fact that each $y - \xi_i$ has a square root in the complement of a ray starting at ξ_i, and the discrepancy between the two sheets along the cut is $e^{\pi\sqrt{-1}}$; multiplying the two together, where allowed, cancels the discrepancy); fix one and call it f_i. Then the function $f = \prod_{i=1}^{n} f_{2i-1}$ is defined on the complement of the segments $\overline{\xi_i \xi_{i+1}}$ in the y-plane.

Now, call $G(x, y) = x^2 - F(y)$. Note that, on Z, $\frac{\partial G}{\partial x}(x, y)$ does not vanish if $x \neq 0$, and $\frac{\partial G}{\partial y}(x, y)$ does not vanish when $x = 0$. Thus, the implicit function theorem allows us to use y as coordinate on Z away from $x = 0$, and x as coordinate around $x = 0$. If we consider the y-plane as $\{x = 0\} \subset \mathbb{C}^2$, we can study the portion of Z that lies above the y-plane with the segments $\overline{\xi_i \xi_{i+1}}$ removed (i.e. the points of Z of the form (x, y) with $x \neq 0$ and y not along any of those segments). Since away from $x = 0$, y is a coordinate for Z, and we have a well-defined square root f of F defined over the slit plane (the plane with the segments missing), we get exactly the two sheets

$$x = f(y), \quad x = -f(y), \tag{4}$$

each of which is isomorphic to the slit plane. Now, we want to put these two sheets together in order to get a Riemann surface with $n - 1$ handles. First, we compactify the two sheets as $y \to \infty$. To this aim, we can glue two small disks $\{w||w| < \epsilon\}$ and $\{t||t| < \epsilon\}$ as follows:

$$\begin{aligned} w &\mapsto (x, y), \quad \text{where} \quad x = f(y),\ y = \frac{1}{w}; \\ t &\mapsto (x, y), \quad \text{where} \quad x = -f(y),\ y = \frac{1}{t}. \end{aligned} \tag{5}$$

To get the Riemann surface C from the two completed copies of the slit plane, we have to glue them according to the following prescription: we attach one side of each slit of the "lower" sheet to the opposite side of the corresponding slit on the "upper" sheet.

The following pictures show how to get the Riemann surface from the sheets. Here the two sheets are indentified.

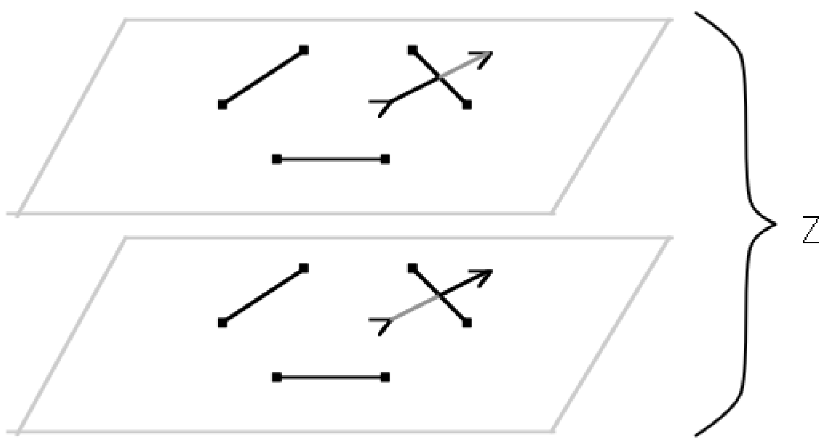

Here there is a more topological way to see the identification.

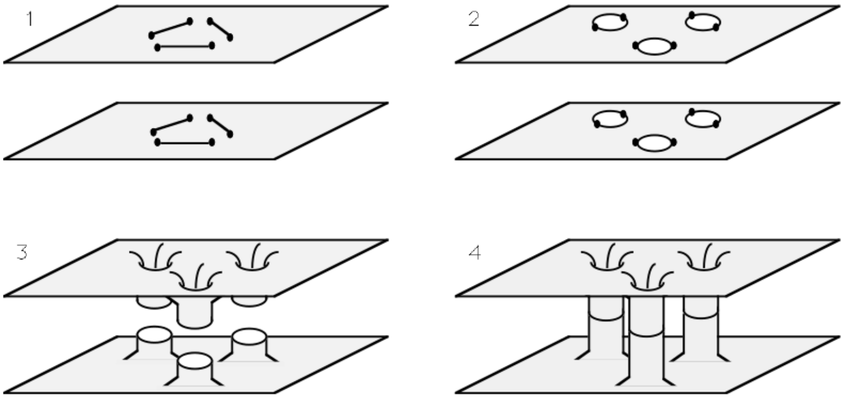

Here is the resulting surface, after the compactification.

Now, for computational reasons, assume $F(y) = y^6 - 1$. Call ω the differential defined as

$$\omega = \frac{dx}{\frac{\partial F}{\partial y}} = \frac{dy}{\frac{\partial F}{\partial x}}, \tag{6}$$

where the first expression is used around $x = 0$ (in this case around the sixth roots of unity), and the second one where $x \neq 0$. So, we get a nowhere-vanishing holomorphic one-form on Z. If we consider the points at infinity, say around $w = 0$, we get

$$\omega = \frac{w^3}{2\sqrt{1-w^6}} d\left(\frac{1}{w}\right) = -\frac{w}{2\sqrt{1-w^6}} dw, \tag{7}$$

and analogously around $t = 0$. In the general case, at the two points at infinity we get zeroes of order $n-2$. Thus, ω induces an abelian differential on C with two zeroes, each of them of order $n-2$. Now, we have observed that the genus of C is $n-1$, so we will denote it by g and start using it instead of n. Since the function

y is meromorphic on C with two poles of order one at $w = 0$ and $t = 0$, we have g linearly independent abelian differentials on C, namely

$$\omega,\ y\omega, \dots,\ y^{g-1}\omega. \tag{8}$$

We claim that this is all, i.e. any other abelian differential is a $\mathbb{C}$-linear combination of these ones. On C we have a natural involution, say ι, induced by switching the sign of x in $\mathbb{C}^2$. Thus, we can write any abelian differential α as sum of its ι-invariant and ι-anti-invariant parts,

$$\beta = \frac{\alpha + \iota^*(\alpha)}{2}, \quad \gamma = \frac{\alpha - \iota^*(\alpha)}{2} \tag{9}$$

respectively. Since β is ι-invariant, for a fixed y in the slit plane, the value on $(f(y), y)$ and $(-f(y), y)$ agree: thus β descends to a holomorphic one-form on the slit y-plane. Now, we show we can extend it on the whole y-plane. Around the slits we can use x as coordinate, and write

$$\beta = \sum_{i \geq 0} b_i x^i. \tag{10}$$

Since β is ι-invariant, all $b_i = 0$ for any even i must vanish (note we gain one (-1) from the differential of ι!). Thus we can write

$$\beta = 2h(x^2)xdx = 6h(F(y))y^5dy \tag{11}$$

for some holomorphic function h. Thus we have a holomorphic one-form on the y-plane; since β is defined on the whole C, and in particular at $w = 0$ and $t = 0$, we can extend the form we got on the y-plane to the Riemann sphere. By the previous example, we know it must be zero, and thus $\beta = 0$. Thus we have that any abelian differential on C is anti-invariant with respect to the hyperelliptic involution ι. Thus, for any α we can consider the meromorphic function α/ω; since they are both anti-invariant, their ratio is invariant under the action of ι. Also, since ω does not vanish away from $t = 0$ and $w = 0$, the function is holomorphic on the two slit planes $x = f(y)$ and $x = -f(y)$. Since we have a pole, and not an essential singularity at the respective points at infinity $w = 0$ and $t = 0$, we have that the function is polynomial on both the slit planes. Since it is ι-invariant, it will depend just on y. Thus, we get $\alpha/\omega = P(y)$ on C. Since ω has poles of order $g-1$ at $w = 0$ and $t = 0$, for α not to have poles, $P(y)$ has to have degree at most $g - 1$. Hence, we have that α is a $\mathbb{C}$-linear combination of the distinguished abelian differentials $\omega,\ y\omega, \dots,\ y^{g-1}\omega$. Also, notice that this implies that any abelian differential on C has exactly $2g - 2$ zeroes.

How to infer the statement for any Riemann surface

The three examples above show that we can count the genus in three different ways:

(i) the topological one (i.e. the number of handles);

(ii) the number of linearly independent abelian differentials;

(iii) $\frac{k+2}{2}$, where k is the number of zeroes of an abelian differential.

The natural question is whether it holds true for any Riemann surface C. At this point we covered one case per topological genus, but we do not know how varying the complex structure could affect the space of abelian differentials for $g \geq 2$. The key point will be to relate algebraic (complex) invariants of a Riemann surface to topological ones.

Given a Riemann surface C, consider the exact sequence of sheaves of abelian groups

$$0 \to \mathbb{C} \to \mathcal{O}_C \xrightarrow{d} \Omega_C \to 0 \tag{12}$$

where $\mathbb{C}$ denotes the constant sheaf $\mathbb{C}$, $\mathcal{O}_C$ the sheaf of regular functions, and Ω_C the sheaf of abelian differentials. The sequence is exact, since on a small disk (it is enough it is simply connected) all 1-forms are exact (remember that for a sequence of sheaves, it is enough to check exactness locally on stalks). Thus, we get the associated long exact sequence in cohomology

$$\begin{aligned} 0 \to &H^0(C,\mathbb{C}) \to H^0(C,\mathcal{O}_C) \to H^0(C,\Omega_C) \\ \to &H^1(C,\mathbb{C}) \to H^1(C,\mathcal{O}_C) \to H^1(C,\Omega_C) \\ \to &H^2(C,\mathbb{C}) \to H^2(C,\mathcal{O}_C) = 0. \end{aligned} \tag{13}$$

Here we are deeply relying on the equivalence between the complex analytic and the algebraic perspectives. Namely, we use the fact that any Riemann surface is a projective algebraic variety (i.e. it can be realized in some $\mathbb{CP}^n$ as the vanishing locus of polynomial equations). Also, Serre's GAGA guarantees that for a complex projective variety X, there is an equivalence of categories between the "algebraic" coherent sheaves on X and the "analytic" coherent sheaves. Thus, we can view $\mathcal{O}_C$ and Ω_C from two points of view (note: $\mathbb{C}$ is not a coherent sheaf!). In particular, we can apply Grothendieck's vanishing to $\mathcal{O}_C$ and get $H^2(C,\mathcal{O}_C) = 0$ (it guarantees that if X is an algebraic variety and $\dim X = n$, then $H^i(X,\mathcal{F}) = 0$ for any sheaf of abelian groups and $i \geq n+1$).

Also, we know that any global regular function on C is constant; thus, the map $H^0(C,\mathbb{C}) \to H^0(C,\mathcal{O}_C)$ is an isomorphism. Hence, by exactness of the sequence, we get that the following map factors through zero, guaranteeing the exactness of the following sequence

$$0 \to H^0(C,\Omega_C) \to H^1(C,\mathbb{C}) \to H^1(C,\mathcal{O}_C) \to H^1(C,\Omega_C) \to H^2(C,\mathbb{C}) \to 0. \tag{14}$$

Now, we know that C is a complex manifold of complex dimension one; hence it is an orientable real manifold of dimension two. This guarantees us $H^2(C,\mathbb{C}) = \mathbb{C}$.

Now a good place to mention an important statement has arrived...

Theorem (Serre duality)**.** Let X be a projective nonsingular variety of dimension n over an algebraically closed field. Let Ω_X be the sheaf of algebraic one-forms on X, and ω_X the n-th exterior power of Ω_X. We call ω_X the canonical sheaf of X, and we denote by K_X the divisor associated to it. Then for any invertible sheaf $\mathcal{L}$ on X we have a functorial isomorphism between $H^i(X,\mathcal{L}^\vee \otimes \omega_X)$ and $H^{n-i}(X,\mathcal{L})^\vee$ for all $i \geq 0$, where $\mathcal{L}^\vee = \mathcal{H}om(\mathcal{L},\mathcal{O}_X)$.

Now, since C has dimension one as algebraic variety, we have $\Omega_X = \omega_X$. Thus, we can apply Serre duality to $H^0(C, \Omega_C)$, and get that it is dual exactly to $H^1(C, \mathcal{O}_C)$. In the same way, we have that $H^1(C, \Omega_C)$ is dual to $H^0(C, \mathcal{O}_C)$; thus, we get that $H^1(C, \Omega_C)$ has dimension one, since $H^0(C, \mathcal{O}_C) \cong \mathbb{C}$. Since by exactness we have that $H^1(C, \Omega_C)$ surjects onto $H^2(C, \mathbb{C})$, and they are both one-dimensional, this map has to be an isomorphism. Thus, again by exactness, we get a short exact sequence

$$0 \to H^0(C, \Omega_C) \to H^1(C, \mathbb{C}) \to H^1(C, \mathcal{O}_C) \to 0. \tag{15}$$

Since $H^0(C, \Omega_C)$ and $H^1(C, \mathcal{O}_C)$ are dual to each other, we get

$$\dim H^1(C, \mathbb{C}) = \frac{\dim H^0(C, \Omega_C)}{2} = \frac{\dim H^1(C, \mathcal{O}_C)}{2}. \tag{16}$$

But $H^1(C, \mathbb{C})$ is a topological invariant, whose dimension is exactly $2g$ (where g denotes the topological genus), while $H^0(C, \Omega_C)$ is an algebraic (or analytic) invariant. Thus we have that the dimension of the space of abelian differentials is a topological invariant. Since we checked for one curve per topological genus that the number of handles agrees with the dimension of such space, we have that the two concepts of genus agree.

Observation. Also, the description in terms of number of zeroes of an abelian differential holds true for any Riemann surface. One way to show it is via the Riemann-Roch theorem.

Observation. The exactness of the sequences used and the dimension count could be done via a purely analytic approach, using Laplace kind of equations and the interpretation of cohomology via the Dolbeaut isomorphism.

The canonical sheaf of $\mathbb{P}^n$

As stated in the Serre duality theorem, for a non-singular variety, the canonical sheaf is the sheaf of top forms. It is interesting and useful to compute such a sheaf in the case of projective space. Thus, we will consider $\mathbb{CP}^n$ and try to explicitly compute what the transition functions of such a line bundle should be.

We will use the standard cover of $\mathbb{CP}^n$ with $U_0, \ldots, U_n$, where $U_i = \{[x_0 : \ldots : x_n] \in \mathbb{CP}^n | x_i \neq 0\}$. Now, for simplicity, we will focus on the change of charts between U_0 and U_1. We will write $(z_1, \ldots, z_n)$ and $(w_1, \ldots, w_n)$ for the natural coordinates on U_0 and U_1, respectively. The restriction to U_i of $\omega_{\mathbb{CP}^n}$ gives ω_{U_i}, which is the trivial line bundle of differential forms on the affine n-space. Thus, ω_{U_0} is generated by the section $dz_1 \wedge \ldots \wedge dz_n$, and ω_{U_1} is generated by the section $dw_1 \wedge \ldots \wedge dw_n$. Hence, in order to understand what kind of sections and transition functions $\omega_{\mathbb{CP}^n}$, it will be enough to understand how to express $dz_1 \wedge \ldots \wedge dz_n$ in terms of $dw_1 \wedge \ldots \wedge dw_n$. On $U_{01} = U_0 \cap U_1$ both x_0 and x_1 are invertible functions. Thus, for $i \geq 2$, we have

$$dz_i = d\left(\frac{x_i}{x_0}\right) = d\left(\frac{x_i}{x_1} \cdot \frac{x_1}{x_0}\right) = w_1^{-1} dw_i + w_i dz_1. \tag{17}$$

On the other hand, for $i = 1$ we have

$$dz_1 = d(w_1^{-1}) = -w_1^{-2} dw_1. \tag{18}$$

These facts all together amount to say that

$$dz_1 \wedge \ldots \wedge dz_n = -w_1^{-n-1} dw_1 \wedge \ldots \wedge dw_n. \tag{19}$$

Thus, the standard section on U_0 gains a pole of order $n+1$ along the hyperplane at infinity. Modulo a negative sign (in the change of charts from U_i to U_j you have to be careful about the ordering of i and j), this is the change of chart that occurs for any pair i and j, if we pick the standard trivialization of ω_{U_i} and ω_{U_j}. In particular, there are no global sections.

On the other hand, the negative sign in front of w_1^{-n-1} (and analogously the one that might appear in the other changes of charts) is not relevant, since we could trivialize ω_{U_0} choosing the section $-dz_1 \wedge \ldots \wedge dz_n$ instead. Thus, we have that the the canonical line bundle on $\mathbb{CP}^n$ is isomorphic to the line bundle $\mathcal{L}$, given by the following data:

(i) $\mathcal{L}$ is trivial on each U_i, i.e. $\mathcal{L}_{|U_i} \cong U_i \times \mathbb{C}$;

(ii) the change of chart from U_i to U_j is given by $f_{ij}(\underline{z})(v) = z_i^{-n-1} v$, where $\underline{z}$ are the coordinate functions on U_i and v is in the stalk over the point with coordinates $\underline{z}$.

Thus, we notice that all the data we need is the number $-n-1$. It also agrees with the divisor we might want to associate to this line bundle. Indeed, the section $dz_1 \wedge \ldots \wedge dz_n$ is a rational section of $\omega_{\mathbb{CP}^n}$ (i.e. it has some poles); in particular, it never vanishes, and has a pole of order $n+1$ along the hyperplane $H = \{x_0 = 0\}$. Then it follows that $\omega_{\mathbb{CP}^n} \cong \mathcal{O}_{\mathbb{CP}^n}((-n-1)H)$.

Observation. What we have just showed, namely that the relevant piece of data is the integer $-n-1$, is a fact that holds true more generally. On $\mathbb{P}^n_k$, the projective space over a field k, isomorphism classes of line bundles are parametrized by $\mathbb{Z}$. To denote a line bundle, we just write $\mathcal{O}_{\mathbb{P}^n_k}(l)$ for some integer l. Also, if we try to work out the transition functions for such a line bundle on $\mathbb{CP}^n$, we get $f_{ij}(\underline{z})(v) = z_i^l v$. Also, the fact that in computing the transition functions we could neglect the -1 can be made more formal, looking at the definition of Cartier divisor.

Observation. In general, if we consider $\mathbb{P}^n_k$, then the canonical sheaf $\omega_{\mathbb{P}^n_l}$ is isomorphic to $\mathcal{O}_{\mathbb{P}^n_k}(-n-1)$. It can be shown via the Euler sequence involving the sheaf of relative differentials $\Omega_{\mathbb{P}^n_k|k}$.