BABY'S FIRST RIEMANN-ROCH

ROK GREGORIC

This note summarizes a rather elegant and intuitive "proof" of the Riemann-Roch Theorem, in an elementary fashion and entirely from first principles (of course, in the complex analytic setting). I have learned it from the lovely notes on curves by Harris [Harr], but it is very similar to Riemann's original consideration [Riem].

0.1. **Contents.** We lay this, Riemann's contribution, out in Chapter 1. Then we discuss Roch's contribution and generalization to arbitrary divisors in Chapter 2. In Chapter 3, we briefly discuss how the theorem applies to line bundles in Chapter 3. Pursuing this perspective on the Riemann-Roch story, we find in Chapter 4 how it leads us naturally to consider sheaf cohomology. Finally in Chapter 5, we sketch the usual cohomological proof of the Riemann-Roch Theorem, and compare it to the explicit proof sketched in Chapter 2.

The reader who wishes only to walk away with intuition about Riemann-Roch, and understanding of what it says, may safely skip Capters 3, 4, and 5.

0.2. **What we hope to do.** Our primary goal is to convince the reader that, while perhaps foreboding at first with all the terminology that goes into it, Riemann-Roch is actually not a scary theorem. In face, remembering some basic ideas from your complex analysis course, you would very likely stubmle very close to it yourself, if somebody asked you to ponder what we call in this note the *Riemann-Roch Problem*.

0.3. **What we do not do.** Please note that we do not actually *prove* the theorem, as we leave a key step unverified in section 1.10. If you wish to see how this can be resolved, and a real proof obtained, please consult [Harr].

1. Fun & easy complex anlysis game

1.1. **Holomorphic functions.** Let X be a compact connected Riemann surface, fixed throughout for the rest of this note. The most basic type of complex analytic functions on X that we could consider are holomorphic functions.

Alas, we know from basic complex analysis that non-constant holomorphic functions are incapable of attaining extrema. That is rather problematic, since the compactness of X implies by basic point-set topology that any continuous function on it will have to have maxima and minima. Consequently, the only holomorphic functions on X are the constant ones. In symbols, we have $\mathscr{O}(X) = \mathbf{C}$.

1.2. **Meromorphic functions.** From our basic course in complex analysis, we recall another type of functions: meromorphic. They are functions almost everywhere holomorphic, but which are allowed to have isolated poles at certain points.

We know from section 1.1. that, if we do not allow for any poles, we get no non-constant functions on X. How about if we allow a single pole? Or two? Or three? Etc. This is the basic question in the Riemann-Roch circle of ideas:

Date: February 4, 2020.
University of Texas at Austin.

The Riemann-Roch Problem (1st version). *Determine the number of meromorphic functions on X with prescribed poles.*

1.3. **Vector space $L(D)$.** Thus let us organize ourselves. Fix some arbitrary points $x_1, \dots, x_d \in X$, and consider the for notational convenience $D = x_1 + \dots + x_d$, really just a formal sum of those points. Then we set $L(D)$ to consisting of functions which are holomorphic away from the x_j, and are allowed to have at most a simple pole at each point x_j. We also allow the constant function 0 in $L(D)$, thus making it into a complex vector space. Since this is a vector space, the "number" of functions, alluded to in the previous function, should be interpreted as its dimesnion $\ell(D) = \dim L(D)$.

1.4. **The Riemann-Roch problem.** The Riemann-Roch problem may now be expressed as

The Riemann-Roch Problem (2nd version). *Determine $\ell(D)$ in terms of data of the points themselves, and the topology of X.*

Since we feel that the choice of points should not effect the question too much, it seems reasonable to expect their contribution to be in terms of their number d. On the other hand, the topology of X is entirely determined, among topological surfaces, by the non-negative integer g, its genus. Thus we wish to describe $\ell(D)$ in terms of d and g.

1.5. **Residues.** Well, let us think about functions in $f \in L(D)$ then! Since we do not know much about them, there is very little we can do.

If you know that a meromorphic function f on X has a pole at a point $x \in X$, then you you know that in terms of a small enough local holomorphic coordinate chart[1] (z, U) around x, you'll be able to write

$$f|_U = \sum_{j=-1}^{\infty} a_j z^j$$

for some constants $a_j \in \mathbf{C}$. This is the Laurent series expansion of f around x. The interesting part, i.e. the part pertaining to the pole, is concentrated in $a_{-1} = \mathrm{Res}_x(f)$, the residue - thus is the same definition as in basic complex analysis.

The assignment $f \mapsto (\mathrm{Res}_{x_0}(f), \dots, \mathrm{Res}_{x_d}(f))$ defines[2] a clearly $\mathbf{C}$-linear map $\mathrm{Res} : L(D) \to \mathbf{C}^d$. Its kernel consists of those functions $f \in L(D)$ whose residues at all the points x_j are zero, i.e. which are holomorphic at all those points. Since a function in $L(D)$ is already holomorphic everywhere else, the kernel of Res is $\mathcal{O}(X) \simeq \mathbf{C}$. At this point we have the exact sequence

$$0 \to \mathbf{C} \to L(D) \xrightarrow{\mathrm{Res}} \mathbf{C}^d,$$

which however is not necessarily exact on the right - the map Res need not be surjective. That is to say, we have

$$\ell(D) = \dim \mathbf{C} + \dim \mathrm{Res}(L(D)) = 1 + \dim \mathrm{Res}(L(D)). \tag{1}$$

[1] That is, an open subset $U \subseteq X$ together with a holomorphic embedding $z : U \hookrightarrow \mathbf{C}$. Any Riemann surface may by definition be covered by such charts, and so local questions on it translate purely into ones of complex analysis in one complex variable.

[2] We should point out that this map *does depend* on the choice of coordinates (z_j, U_j) around each point x_j. There is nothing wrong with that *per se*, though it might feel a little dirty. In fact, we will return to this after a fashion in section 1.9.

We have made good progress towards Riemann-Roch here! Now the problem has become determining the image $\mathrm{Res}(L(D))$ inside $\mathbf{C}^d$. We approach this by finding a system of linear constraints that its elements must satisfy.

1.6. **The Residue Theorem.** We take inspiration from our course in complex analysis. The one thing we might remember having learned there about residues, is the *Residue Theorem*. It says that if f is a meromorphic function on (neighborhood of) a relatively compact open domain $U \subseteq \mathbf{C}$, with poles $z_1, \dots, z_n \in U$, then we may express the path integral of f around the boundry of U as

$$\int_{\partial U} f(z)\, dz = 2\pi i \sum_{j=1}^{n} \mathrm{Res}_{z_j}(f)$$

We might imagine doing this on the Riemann surface X as well, perhaps by cutting it into coordinate chart pieces, and then canceling the integral contributions along the cuts (as they would enter the total sum of the integral once in one direction, and once in the opposite, hence with twice with opposite signs).

But since the Riemann surface in question has no boundary, the LHS integral would evaluate to 0, leaving us with the claim that the sum of residues is zero. That is precisely the sort of thing we are looking for: a *linear relation among residues*!

1.7. **A problem with this idea.** That is a great idea, but it has one fatal flaw: it is *not possible* to consistantly define an integral of a holomorphic function along a curve on a Riemann surface. Roughly, that is because they are missing the dz term. That must come from a holomorphic coordinate z, and on $\mathbf{C}$ (where your complex analysis likely took place), one may choose such a coordinate once and for all and express everthing in terms of it (again, almost *certainly* how you did it in your class). Then it doesn't matter if we are thinking of $f(z)$ or $f(z)\, dz$, as we can always multiply or divide with the chosen dz to pass from one to another. The issue is that, on a general Riemann surface X, it will generally not be possible to define such a *globally* well-defined dz.

1.8. **Differential forms.** This is not hard to fix though: instead of considering differential functions, i.e. functions $f : X \to \mathbf{C}$ which look in terms of every holomorphic chart $(z, U))$ as a holomorphic function $f|_U = f(z)$, we may instead consider *holomorphic 1-forms.*

A holomorphic 1-form ω is something that looks on every holomorphic chart (z, U) as $\omega|_U = g(z)\, dz$ for some holomorphic function $g(z)$ on the complex plane. As such, it is *precisely* the kind of thing that may be integrated along paths on X! And while a *globally* well-defined dz might not exist, holomorphic 1-forms, which must only look so *locally* very likely will.

Replacing the word "holomorphic" with "meromorphic" everywhere in the previous paragraph, we also get meromorphic 1-forms. Those are the ones which may posses residues, and so they are the ones to which the Residue Theorem will apply, guaranteeing that all their residues together sum to 0.

1.9. **Linear constraints on** $\mathrm{Res}(L(D))$**.** Thus to estimate the image of Res, following the program of section 1.6., let's fix an arbitrary non-zero holomorphic 1-form ω on X. For any non-zero $f \in L(D)$, the product $f\omega$ is then a meromorphic 1-form on X with poles only possible at the points x_i, and so the Residue Theorem shows that

$$\mathrm{Res}_{x_1}(f\omega) + \cdots + \mathrm{Res}_{x_d}(f\omega) = 0.$$

In holomorphic local coordinates (z_j, U_j) around each point x_j, suppose you write $\omega|_{U_j} = g_j\, dz_j$ for a holomorphic function $g_j \in \mathscr{O}(U_j)$. Then you have

$$\mathrm{Res}_{x_j}(f\omega) = \mathrm{Res}_{x_j}(f) g_j(x_j)$$

and consequently the above Residue Theorem relation becomes

$$\sum_{j=1}^{d} \mathrm{Res}_{x_j}(f) g_j(x_j) = 0.$$

Thus if $(a_1, \ldots, a_n) \in \mathbf{C}^d$ belongs to the image of the map Res, it must satisfy the linear relation

$$\sum_{j=1}^{d} a_j(f) g_j(x_j) = 0.$$

This was for a fixed holomorphic 1-form ω, and since the relation is clearly linear in ω, we get $\dim \Omega^1(X)$-many linear relations on the image of Res.

Here we have used $\Omega^1(X)$ to denote the vector space of holomorphic 1-forms on X. It is a basic truth of Riemann suface theory that $\dim \Omega^1(X) = g$, the genus of X. Indeed, that may (and in algebro-geometric contexts often is) taken as the definition of the genus.

1.10. **Dependence among the linear constraints.** In the previous section, we found a system of linear constraints on the elements of $\mathrm{Res}(L(D)) \subseteq \mathbf{C}^2$, one for each element $\omega \in \Omega^1(X)$.

But not all of these constraints will be independent: indeed, if $g_j(x_j) = 0$ for all $1 \le j \le d$, i.e. if ω vanishes at all the points x_j , then the relation imposed by ω is void. Let $\Omega^1_D(X)$ denote the vector subspace of $\Omega^1(X)$ consisting of all holomorphic 1-forms which vanish at all points x_j. Then the non-trivial linear constraints obtained in 8. are given by $\Omega^1(X)/\Omega^1_D(X)$.

These linear constraints are also sufficient[3], i.e. they precisely cut out the subset $\mathrm{Res}(L(D)) \subset \mathbf{C}^d$. It follows that

$$\begin{aligned} \dim \mathrm{Res}(L(D)) &= \dim \mathbf{C}^d - \#\text{ constraints} \\ &= d - g + \dim \Omega^1_D(X). \end{aligned}$$

Combining this with the formula (1) we got before for $\ell(D)$ in terms of $\mathrm{Res}(L(D))$, we get the formula

$$\ell(D) = 1 + d - g + \dim \Omega^1_D(X). \tag{2}$$

This is the Riemann-Roch formula! In many ways, it is the solution to the Riemann-Roch Problem, as set out it section 1.4. - it expresses $\ell(D)$ in terms of d and g. We could perfectly well stop our discussion here.

2. Divisors & stuff

But let us instead try to relate equation (2) to the usual formulation of the Riemann-Roch theorem you might have come across.

You might justifiably object to calling equation (2) the Riemann-Roch theorem. For instance, you might object that it is not actually a good solution to the Riemann-Roch Problem, as set out in section 1.4., since it includes the term $\dim \Omega^1_D(X)$.

[3]This is the one point in this argument where we merely assert something, and not prove it. That is fine though, as shown in [Harr], it is not all that hard to make this into a real proof.

2.1. **Riemann's Inequality.** Indeed, Riemann himself had come only this far in [Riem], though like us in 1.10., he didn't really prove that the system of constraints is sufficient. Instead he settled for the inequality

$$\ell(D) \geq 1 + d - g.$$

From the perspective of the 1st version the Riemann-Roch Problem from section 1.2., this is quite reasonable. If you are trying to figure out how many poles you must allow meromorphic functions to have until some non-constant ones actually start existing, this inequality will get you there!

For instance, Riemann's goal in [Riem] was to prove that every compact Riemann surface is algebraic. To do so, he required a good supply of linearly independent meromorphic functions (from which to build the embedding $X \hookrightarrow \mathbf{P}^n$, the image of which he then proved to be cut out by polynomial equations). His inequality kindly provided him with information about poles at (at least) how many points were neccessary for that.

2.2. **Roch's contribution.** It was Riemann's student Roch, who in [Roch] completed the inequality to the equality. First he showed that the system of constraints is indeed unique, and secondly, he recognized $\Omega^1_D(X)$ as something else. To understand what this something else is, it will be useful to first discuss how RR can be generalized, as it will make us organically stuble upon the ingredients.

2.3. **Higher order poles.** Instead of only considering meromorphic functions with simple poles, as we have up to now, we might have allowed higher poles too.

Indeed, imitating section 1.3., we might have fixed points $x_1, \dots, x_k \in X$ together with degrees $n_1, \dots, n_k \in \mathbf{Z}_{\geq 0}$, and considered meromorphic functions on X which are holomorphic away from the points x_j, and may have a pole of order up to n_j at x_j, for every $j = 1, \dots, k$. Collecting this data into the form of a *divisor* $D = n_1 x_1 + \dots + n_k x_k$, the associated meromorphic functions form the vector space $L(D)$.

Then we may imitate the steps of sections 1.4. – 1.10. with only the discussion in section 1.5. needing to be seriously modified. Indeed, if a function f is known to have a pole of degree up to n at a point x, then in terms of a local holomorphic coordinate (z, U) around x, its Laurent expansion will be

$$f|_U = \sum_{j=-n}^{\infty} a_j z^j.$$

The "interesting" information, i.e. pertaining to the pole, now resides in the n-tuple of coefficients $a_{-n}, \dots, a_{-1} \in \mathbf{C}$.

Thinking about sections 1.6. – 1.10. we see that this means the role of d in them will be played by the number

$$\deg D = n_1 + \dots + n_k,$$

the *degree of the divisor* D. In particular, in place of (2), the Riemann-Roch formula becomes

$$\ell(D) = 1 + \deg D - g + \dim \Omega^1_D(X). \tag{3}$$

We still haven't come to identifying the $\dim \Omega^1_D(X)$ term with anything else, but we might have come to appreciate the role of divisors in the Riemann-Roch story better. At this point, it may occur to us that this whole story has secretly been about divisors all along, so it is worth our while to consider them slightly more closely.

2.4. **Divisors.** It pays to generalize the notion of a divisor slightly, by allowing it to be $D = n_1x_1 + \cdots + n_kx_k$ for any points $x_j \in X$ and any degrees $n_j \in \mathbf{Z}$. To differentiate them, divisors with all the coefficients $n_j \geq 0$, such as we considered in the previous paragraph, are called *effective divisors*.

This way, given a meromorphic function f on X, we can associate to it the *principal divisor*

$$\operatorname{div}(f) = \{\text{zeros of } f\} - \{\text{poles of } f\}.$$

Both zeros and poles here are counted with multiplicity, which amounts for the degree coefficients in the divisor. These may be viewed as the prime examples of divisors.

Note that the points x_j appearing in the divisor $D = n_1x_1 + \cdots + n_kx_k$ are somewhat artificial: if $x_{k+1} \in X$ is another point we would like to add, we may add it as $D = n_1x_1 + \cdots + n_kx_k + 0x_{k+1}$. In general, we could view any divisor as a sum $D = \sum_{x \in X} n_x x$, but where we require the degrees $n_x \in \mathbf{Z}$ to vanish for all but finitely many points $x \in X$.

From this, it is obvious that we can sum and substract divisors, and that they form an abelian group. Indeed, it is the free abelian group spanned by the points of X. But we can also put a compatible partial ordering on divisors: if $D = \sum_{x \in X} n_x x$ and $D' = \sum_{x \in X} n'_x x$, then we say that $D \leq D'$ if and only if $n_x \leq n'_x$ for all $x \in X$.

2.5. **The Riemann-Roch Problem in terms of divisors.** Using the technology of divisors from the last section, we can rephrase the definition of the vector space $L(D)$, associated in section 2.3. to an effective divisor $D = n_1x_1 + \cdots n_kx_k$, in terms of it.

Indeed, for a meromorphic function f on X, note that the condition that is holomorphic away from the points x_j and allowed to have a pole of order up to n_j at each x_j, is equivalent to demanding that its divisor satisfy $\operatorname{div}(f) + D \geq 0$.

This works, at least as a definition, for any divisor D:

$$L(D) = \{f \in \mathscr{M}(X) : \operatorname{div}(f) + D \geq 0\} \cup \{0\}.$$

That is, if $D = n_1x_1 + \cdots + n_kx_k$, a function $f \in L(D)$ is

- holomorphic away from the points x_j, or at x_j if $n_j = 0$,
- *allowed* to have a pole of oder *up to* n_j at the point x_j if $n_j > 0$,
- *must* have a zero of order *at least* n_j at the point x_j if $n_j < 0$.

The Riemann-Roch problem, as always, is to determine $\ell(D)$ in terms of the genus g of X and the degree of the divisor $\deg D = n_1 + \ldots + n_k$. In fact, formula (3) is still valid, so long as we interpret $\Omega^1_D(X)$ as conmsisting of those meromorphic 1-forms which are

- holomorphic away from the points x_j, or at x_j if $n_j = 0$,
- *must* have a zero of order *at least* n_j at the point x_j if $n_j > 0$,
- *allowed* to have a pole of oder *up to* n_j at the point x_j if $n_j < 0$.

Note that this is precisely opposite to the requirements for a meromorphic function to be in $L(D)$, suggesting that perhaps $\Omega^1_D(X)$ has to do with $-D$ instead of D.

2.6. **The canonical divisor.** Finally we come to Roch's contribution to the Riemann-Roch story, alluded to in section 2.2. First, he introduced a special kind of divisor on X, called the *canonical divisor* and denoted by K. This divisor has the special property that

$$L(K) \simeq \Omega^1(X).$$

We may define it by taking an arbitrary non-zero meromorphic 1-form ω, and associating to a divisor $\operatorname{div}(\omega)$ of its zeros minus its poles, just as we did for functions in section 2.4. This gives rise to the canonical divisor $K = -\operatorname{div}(\omega)$.

The isomorphism $L(K) \simeq \Omega^1(X)$ is now given by sending $f \mapsto f\omega$, for the chosen meromorphic 1-form ω. Since the points at which ω has a pole of some order, f must have a zero of at least that order, the obtained 1-form $f\omega$ is indeed homolorphic. Conversely, for an arbitrary $\omega' \in \Omega^1(X)$, a holomorphic chart (z, U) gives rise to local expressions $\omega|_U = g(z)\,dz$ and $\omega' = g'(z)\,dz$ for a holomorphic function g and a meromorphic function g'. Then

$$\left.\frac{\omega'}{\omega}\right|_U = \frac{g'(z)\,dz}{g(z)\,dz} = \frac{g'(z)}{g(z)} \in \mathscr{M}(U), \tag{4}$$

and so $f = \frac{\omega'}{\omega}$ is a meromorphic function on X, which must have zeros where ω has poles, and may have poles where ω has zeros. I.e. we have found an element $f \in L(K)$, such that $\omega' = f\omega$, establishing the desired isomorphism.

2.7. **The Riemann-Roch formula.** The argument of the previous section is easy to modify in the presence of a divisor D. It allows us to entirely analogously obtain a vector space isomorphism

$$L(K - D) \simeq \Omega^1_D(X), \tag{5}$$

realizing the criptic remark at the end of section 2.5.

Now we can deliver on the promise in 2.2. that we would come to Roch's contribution to the Riemann-Roch story. Using this identification, the Riemann-Roch formula (3) may be rewritten in the today-standard form

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

As we mentioned in section 2.5., this works for an arbitrary divisor D on X, and this is the celebrated *Riemann-Roch Theorem* in its definitive form.

3. Where line bundles fit into all this

Let us give a few hints and remarks as to how the Riemann-Roch story ties into the theory of holomorphic line bundles.

3.1. **Divisors and line bundles.** In section 2.6 we found a divisor K, for which $\Omega^1(X) \simeq L(K)$. Similarly, it is clear that $\mathscr{O}(X) \simeq L(0)$. This might lead us to ask: what sorts of functions or function-like things on the Riemann surface X can be expressed as $L(D)$ for some divisor D.

The answer turns out to be: all sections of holomorphic line bundles. Given a holomorphic line bundle $\mathscr{L}$ on X, there always exists a divisor D on X for which there is an isomorphism $L(D) \simeq \Gamma(X; \mathscr{L})$, identifying the its global (holomorphic) sections.

Indeed, we may find it by following the same recpie we used for Ω^1_X in section 2.6.: let $s \in \Gamma(X; \mathscr{L} \otimes_{\mathscr{O}_X} \mathscr{M}_X)$ be a non-zero meromorphic section of the line bundle $\mathscr{L}$. To obtain the desired divisor D, we take its divisor of poles minus its divisor of zeros. Then an argument entirely analogous to the one in section 2.6. shows that the map $L(D) \to \Gamma(X; \mathscr{L})$, sending $f \mapsto fs$, is an isomorphism of vector spaces.

3.2. **Divisor class group.** There is some ambiguity in the process outlined in the previous though: different choices of the section s may lead to different divisors D. Alas, since $\mathscr{L}$ is a line bundle, an argument along the lines of (4) from section 2.6., shows that any two non-zero meromorphic sections s and s' of $\mathscr{L}$ must be related by $s = fs'$ for some meromorphic function f on X. It follows that the two divisors obtained from s and s' will differ by a principal divisor $\mathrm{div}(f)$.

Thus while the divisor D associated to the line bundle $\mathscr{L}$ may not be well-defined as an element of the divisor group $\mathrm{Div}(X)$, it *is* well defined as an element of the *divisor class group* $\mathrm{Cl}(X) = \mathrm{Div}(X)/\mathrm{div}(\mathscr{M}(X)^\times)$.

3.3. **The first Chern class.** The divisor D, associated as described in section 3.1. to the line bundle $\mathscr{L}$, is called the *first Chern class of* $\mathscr{L}$, in which capacity it is denoted $c_1(\mathscr{L})$. Using it, we may define the *degree of a line bundle* in terms of the degree of a divisor as $\mathscr{L}$ as $\deg \mathscr{L} = \deg c_1(\mathscr{L})$. If you are concerned about the uniqueness of this definition, see Appendix A.

In light of section 3.2., we may view the first Chern class as a map

$$c_1 : \mathrm{LinBun}(X) \to \mathrm{Cl}(X)$$

from (holomorphic) line bundles on X (modulo isomorphisms) to the divisor class group. Furthermore, it is clear from its description that

$$c_1(\mathscr{L} \otimes_{\mathscr{O}_X} \mathscr{L}') = c_1(\mathscr{L}) + c_1(\mathscr{L}'),$$

making it a group[4] homomorphism. In fact it is even an *isomorphism.* That is, line bundles and divisor classes encode the same data.

3.4. **Equivalence between line bundles and divisors.** To convince ourselves of this claim, we construct and inverse to c_1, imitating (or more precisely, localizing) the construction of the vector space $L(D)$ from a divisor D that we had given in section 2.5.

Given a divisor D, we define the holomorphic line bundle $\mathscr{L}(D)$ so that its sections[5] over any open $U \subseteq X$ are given by

$$\Gamma(U; \mathscr{L}(D)) = \{f \in \mathscr{M}(U) : \mathrm{div}(f|_U) + D|_U \geq 0\} \cup \{0\}.$$

Here $D|_U$ is obtained from the divisor $D = \sum_{x \in X} n_x x$ by throwing away all the points not in U, which is to say, $D|_U = \sum_{x \in U} n_x x$. It is quite easy to see that $c_1(\mathscr{L}(D)) = D$ (at least, up to a principal divisor), while an argument analogous to the one we made in section 2.6. for $\mathscr{L} = \Omega^1_X$, shows that $\mathscr{L}(c_1(\mathscr{L})) \simeq \mathscr{L}$.

3.5. **Riemann-Roch for line bundles.** If $\mathscr{L}$ is a line bundle, then $\mathscr{L} = \mathscr{L}(D)$ for some divisor D. Since $\mathscr{L}(D) \otimes \mathscr{L}(-D) = \mathscr{L}(0) = \mathscr{O}_X$, we may identify the dual line bundle to $\mathscr{L}$ with $\mathscr{L}^\vee = \mathscr{L}(-D)$. In light of this, the isomorphism (5) of section 2.7. may be rephrased in terms of an arbitrary line bundle $\mathscr{L}$ on X as the isomorphism

$$\Gamma(X; \mathscr{L}^\vee \otimes_{\mathscr{O}_X} \Omega^1_X) \simeq \Omega^1_{c_1(\mathscr{L})}(X)$$

[4]The group of holomorphic line bundles on X up to isomorphism, equipped with the operation of tensor product, is called the *Picard group* of X.

[5]Thus we are actually defining $\mathscr{L}(D)$ as a sheaf, and claim that it is in fact the sheaf of sections of a holomorphic line bundle. This kind of identification of bundles and their sheaves of sections is hihgly standard in algebraic geometry.

of complex vector spaces.Thus the Riemann-Roch formula of section 2.7. may be expressed in terms of line bundles as

$$\dim \Gamma(X; \mathscr{L}) - \dim \Gamma(X; \mathscr{L}^\vee \otimes_{\mathscr{O}_X} \Omega^1_X) = \deg \mathscr{L} + 1 - g. \tag{6}$$

4. Cohomology and all that bizz

This section is a little harder than the previous ones, and we hope that this it will not spoil the fun-and-easy flavor of what came so far. Perhaps everything discussed from here on might be better saved for baby's second Riemann-Roch.

4.1. **What are we even doing?** When we set off on our Riemann-Roch adventure, we found in section 1.4. that we were interested in the vector space $L(D)$ for a (general, since section 2.5.) divisor D on X. In section 3.4. we learned that this vector space $L(D) = \Gamma(X; \mathscr{L}(D))$ may be viewed as the global sections of a holomorphic line bundle $\mathscr{L}(D)$ on X, and that divisors and line bundles are in bijection through this construction.

Thus, instead of the divisor D, we might have started of with a holomorphic line bundle $\mathscr{L}$, and ask ourselves the following question:

The Riemann-Roch Problem (3rd version)**.** *Determine* $\Gamma(X; \mathscr{L})$ *in terms of data about the line bundle and the topology of* X*.*

With that as our problem, surely the Riemann-Roch formula (6) of section 3.5. is the answer. But there is nonetheless something new to be learned: our focus has shifted to global sections of line bundles. So let us talk a little about those.

4.2. **Global sections.** We want to think about the process $\mathscr{L} \mapsto \Gamma(X; \mathscr{L})$, of sending a line bundle to its vector space of global sections.

Our argument to establish Riemann-Roch used in 1.5. the technology of exact sequences, which is to say homological aglera, and it seems convenient to ask for it in this setting too. Sadly, line bundles are ill-equipped to handle such things: they can not accomodate the most basic operation of homological algebra: taking direct sums. This issue is of course that, if $\mathscr{L}$ and $\mathscr{L}'$ are line bundles, then their direct sum $\mathscr{L} \oplus \mathscr{L}'$, while still a vector bundle, will have rank 2 instead of 1.

4.3. **So, vector bundles then?** Based on that observation, we should extend our scope to (holomorphic) vector bundles $\mathscr{E}$ on X, for which global sections $\mathscr{E} \mapsto \Gamma(X; \mathscr{E})$ work just as well. But we are not out of the woods yet! While vector bundles allow us perform many operations of homological algebra, i.e. direct sums and kernels, they fail to allow cokernels.

The issue is simple to illustrate: consider for a fixed point $x \in X$ the inclusion $i : \mathscr{O}_X \hookrightarrow \mathscr{L}(x)$, where $\mathscr{L}(x)$ is the line bundle associated to the divisor x. Following the explicit description of the latter from section 3.4., we see that i is the inclusion of holomorphic functions into functions which are holomorphic away from x, and may have at most a simple pole at x. The cokernel of this map $\operatorname{coker}(i) = \mathscr{L}(x)/\mathscr{O}_X$ should thus has vanishing fibers at all the points $y \in X$, since the two bundles agree precisely on any small enough neighborhood of y.

To understand what is happening at the distinguished point x, we recall the Laurent series expansion from 1.5. In a holomorphic chart (z, U) around x, the inclusion $i : \mathscr{O}_X \hookrightarrow \mathscr{L}(x)$ is given, in terms of the Laurent expansions, as the inclusion of (convergent) series

$$\left\{ \sum_{j=0}^{\infty} a_j z^j : a_j \mathbf{C} \right\} \hookrightarrow \left\{ \sum_{j=-1}^{\infty} a_j z^j : a_j \in \mathbf{C} \right\}.$$

To obtain the cokernel, we extract the coefficients $a_{-1} \in \mathbf{C}$, or as we recalled in 1.5., compute the residue at x. It follows that the fiber of $\operatorname{coker}(i)$ at the point x must be $\mathbf{C}$.

Alas, if $\operatorname{coker}(i)$ were to be a vector bundle, such as $\mathscr{L}(x)$ and $\mathscr{O}_X$ are, then all of its fibers would have to be isomorphic as complex vector spaces. We conclude that, in order to be able to form cokernels, holomorphic vector bundles are not enough.

4.4. **Coherent sheaves.** Once again, we expand our objects of interest, this time explicitly by building in cokernels.

Indeed, recall (as you might know, e.g. from a course in algebraic geometry) that a vector bundle on[6] X may equivalently be defined as a sheaf $\mathscr{E}$ of $\mathscr{O}_X$-modules, such that there exists for every point a neighborhood $U \subseteq X$ and an isomorphism of $\mathscr{O}_U$-modules

$$\mathscr{E}|_U \simeq \mathscr{O}_U^{\oplus r}.$$

That is, *locally* $\mathscr{E}$ is isomorphic to the direct sum $\mathscr{O}_X^{\oplus r}$, where r is its rank, but not necessarily *globally.*

Well, since we saw in the last section that vector bundles are not closed under cokernels, we define a *coherent sheaf*[7] to be a sheaf $\mathscr{F}$ of $\mathscr{O}_X$-modules, such that for every point there exists a neighborhood $U \subseteq X$ and an exact sequence of $\mathscr{O}_X$-modules

$$\mathscr{O}_U^{\oplus p} \to \mathscr{O}_U^{\oplus q} \to \mathscr{F}|_U \to 0.$$

Thus a coherent sheaf is locally a cokernel of a vector bundle map. Finally coherent sheaves $\operatorname{Coh}(X)$ turn out to be a convenient place to do homological algebra. In fancy words, we can say that they form an abelian category.

But before we delve into the homological algebra we wish to perform in this abelian category, let us examine some of its objects, to get a bit of a feeling for coherent sheaves, and hopefully dispel any air of mystery surrounding them.

4.5. **Skyscrapers.** Coherent sheaves on X by design contain all vector bundles, along with a bit more. The archetypical example is the cokernel from section 4.3. that we could not form inside vector bundles.

That is, for any point $x \in X$, the cokernel of the inclusion $\mathscr{O}_X \to \mathscr{L}(x)$ is denoted $x_*(\mathbf{C})$, and called the *skyscraper sheaf at x*. Equivalently (and somewhat more standardly), it is the cokernel of the inclusion $\mathscr{L}(-x) \to \mathscr{O}_X$, of holomorphic functions vanishing at x into all holomorphic functions. An easy generalization is the skyscraper sheaf $x_*(V)$, where V is any finite dimensional complex vector space. It may be formed as a direct sum of copies of $x_*(\mathbf{C})$.

It follows from our discussion in section 4.3. that $x_*(V)$ has the fiber V at the point x, and 0 at al other points in X. In this sense, its support is the 0-dimensional

[6]This would work just as well if X were not just a Riemann surface, but also a complex manifold, complex space, algebraic variety, scheme, etc.

[7]Technically, what we are defining here is called a *finitely presented $\mathscr{O}_X$-module.* Under nice circumstances, such as that when X is a compact Riemann surface (or, if you prefer, smooth algebraic curve), the two agree.

subspace $\{x\} \subseteq X$, so unlike line bundles and vector bundles on X, which have to do with its 1-dimensional geometry, skyscraper sheaves are inherently 0-dimensional.

4.6. **Coherent sheaves = vector bundles + skyscrapers.** Vector bundles and skyscrapers are the two extreme cases for how coherent sheaves can look, and a general coherent sheaf will be something in the middle. Let us elaborate on this a little. As the example of the skyscraper sheaf illustrates, the dimension of the fiber of a coherent sheaf need, unlike that of a vector bundle, not be locally constant. Instead, it may exhibits jumps, though as it turns out, at most at finitely many points[8]. In this sense, coherent sheaves look like a vector bundle almost everywhere, with a finite number of skyscraper protrusions.

4.7. **Taking global sections is not exact.** When working in the context of abelian categories, the functors that we like best are exact ones. That is to say, those which preserve short exact sequences.

The abelian categories of interest to us here is $\mathrm{Coh}(X)$ and $\mathrm{Vect}_{\mathbf{C}}$, and the functor is $\Gamma : \mathrm{Coh}(X) \to \mathrm{Vect}_{\mathbf{C}}$, given by $\mathscr{F} \mapsto \Gamma(X; \mathscr{F})$, which sends a coherent sheaf to its global sections. Alas, this functor is not exact.

Indeed, consider the short exact sequence of coherent sheaves

$$0 \to \mathscr{O}_X \to \mathscr{L}(x) \to x_*(\mathbf{C}) \to 0,$$

which comes about from the *definition* of the skyscraper sheaf in section 4.5 as a cokernel. Passing to global sections, we obtain the exact sequence

$$0 \to \mathbf{C} \to L(x) \to \mathbf{C}, \tag{7}$$

the map $\mathbf{C} \to L(x)$ including the constant functions into functions on X that have at most a simple pole at x, and $L(x) \to \mathbf{C}$ computing the residue at x. We encountered this sequence in section 1.5. already, and indeed, we are disussing the same difficulty as there: that the right-most arrow Res_x is not necessarily surjective. *If it were*, that is, if the sequence *were short exact*, we would find that

$$\dim L(x) = \dim \mathbf{C} + \dim \mathbf{C} = 2. \tag{8}$$

That is true for the Riemann sphere, which is to say when $X = \mathbf{P}^1$, but not for a general compact Riemann surface. To convince ourselves of this, we work it out for one particular family of Riemann surfaces X.

4.8. **Working this out on an elliptic curve.** Consider the case of the elliptic curve $X = \mathbf{C}/(\mathbf{Z} + i\mathbf{Z})$, which is topologically a torus.

Indeed, recall that the torus admits a continuum of complex structures, parametrized by $\tau \in \mathbf{C}$ with $\mathrm{Re}(z) > 0$. For every such τ, the quotient $\mathbf{C}/(\mathbf{Z}+\tau\mathbf{Z})$ equips the torus with a non-isomorphic Riemann surface structure. The argument below would work for an arbitrary such τ, but for simplicity we stick with $\tau = i$.

A holomorphic (resp. meromorphic) function on X may be identified with a holomorphic (resp. meromorphic) function $f(z)$ on $\mathbf{C}$, which is *doubly periodic* in the sense that

$$f(z) = f(z+1) = f(z+i).$$

[8] In the case of compact Riemann surfaces = algebraic curves that we are considering, or along codimension ≥ 1 subschemes/complex subspaces in general.

Let $x \in X$ be the image of $z_0 = \frac{1}{\sqrt{2}}(1+i) \in \mathbf{C}$, the "center of the fundamental square". In order for f to define a function in $L(x)$, it must be holomorphic away from z_0, and may have a simple pole there.

The value of the residue $\mathrm{Res}_x(f)$ may be computed by integrating[9] $f(z)\,dz$ around z_0. The choice of the counter of integration is arbitrarym, so long as it does not cross any poles of f, (that is the content of the Cauchy Integral Theorem, another staple from your basic complex analysis course,) and so we may "drag it out" (homotope it) to the "boundy of the fundamental square", i.e. the square in $\mathbf{C}$ spanned by the points $0, 1, i$, and $1+i$. The required periodicity of $f(z)$ shows that the integrals over opposite sides of the square, which must be taken with opposite orientation, cancel with each other, and we find that $\mathrm{Res}_x(f) = 0$.

In conclusion, a meromorphic function on X that is not holomorphic must have at least two poles. Consequently $L(x) = \mathscr{O}(X) = \mathbf{C}$, and $\dim L(x) = 1$, violating the dimensional equality (8) and consequently showing that the sequence (7) can not be short exact.

4.9. **Cohomology to the rescue.** We learned from section 4.7. that the global sections functor $\Gamma : \mathrm{Coh}(X) \to \mathrm{Vect}_{\mathbf{C}}$ is not exact. That hinders our ability to apply the methods of homological algebra to its study. We could try to works around it, and live with this failure for the rest of our lives. Instead, holding our head up, we get to work fixing it!

The solution will turn out to consider not only the functor Γ, but a whole family of functors H^i for $i \geq 0$, with $\mathrm{H}^0 = \Gamma$. This is *sheaf cohomology*, and the reason it appears in the usual treatment of the Riemann-Roch Theorem in the literature.

4.10. **Derived functors.** The method of resolving the issue of functors not being exact is a general one, and though we usually choose specific over general in this note, let us once for a change discuss the general picture. Suppose we are given a functor between abelian categories $F : \mathcal{A} \to \mathcal{B}$, which is left exact, that is, it preserves left exact sequences, but not necessarily right exact ones. Then there exists a canonical family of functors $\mathrm{R}^i F : \mathcal{A} \to \mathcal{B}$ for $i \geq 1$, called the *right derived functors of F*, such that for any *short exact* sequence

$$0 \to A' \to A \to A'' \to 0$$

in $\mathcal{A}$, the *left exact* sequence in $\mathcal{B}$, obtained by applying F to this short exact sequence, canonically extends to a *long exact* sequence

$$\begin{array}{l} 0 \longrightarrow F(A') \longrightarrow F(A) \longrightarrow F(A'') \\ \hookrightarrow \mathrm{R}^1F(A') \longrightarrow \mathrm{R}^1F(A) \longrightarrow \mathrm{R}^1F(A'') \\ \hookrightarrow \mathrm{R}^2F(A') \longrightarrow \mathrm{R}^2F(A) \longrightarrow \mathrm{R}^2F(A'') \longrightarrow \cdots. \end{array}$$

In light of this, we often write $F = \mathrm{R}^0 F$ in analogy with $\mathrm{R}^i F$.

[9]It turns out that for elliptic curves, the issue of computing residues, that we discussed in section 1.7., is avoided. That has to do with a canonical choice of a holomorphic 1-form, possible because an elliptic curve is a so-called *Calabi-Yau manifold*. If you are around algebraic geometers, or mirror symmetrists more generally, you have likely heard that name before.

If F is already exact, then $\mathrm{R}^i F = 0$ for all $i \geq 1$. If not, the derived functors $\mathrm{R}^i F$ provide a quantitative measure to how far the functor F is from being exact.

4.11. **Computing derived functors by resolutions*.** *This section may safely be skipped.* The property of turning short exact sequences into long exact ones, that we talked about in the previous section, is really all you need know about derived functors. But for what it's worth, we briefly describe a method of computing them explicitly. Perhaps this might explain the relationship with cohomology.

Given an object $A \in \mathcal{A}$, pick a *resolution*, which is to say[10], an exact sequence

$$0 \to A \to A^0 \to A^1 \to A^2 \to \cdots$$

in $\mathcal{A}$. Since F is not right exact applying it to this sequence will destroy its exactness. None the less, the sequence

$$0 \to F(A^0) \to F(A^1) \to F(A^2) \to \cdots$$

is still a cochain complex, in the sense that each two subsequent arrows compose to zero. Thus we may compute its *cohomology*, i.e. quotienting out the image of one arrow inside the kernel of the next. This gives rise to the derived functors as

$$\mathrm{R}^i F(A) = \mathrm{H}^i(F(A^\bullet)).$$

4.12. **Sheaf cohomology.** Returning from the abstract digression of the previous two sections to the setting of section 4.9. let us see what our new-found knowledge bares for the global section functor $\Gamma : \mathrm{Coh}(X) \to \mathrm{Vect}_{\mathbf{C}}$. For a coherent sheaf $\mathscr{F}$ on X, we are led to define *i-th sheaf cohomology of $\mathscr{F}$* as

$$\mathrm{H}^i(X; \mathscr{F}) := \mathrm{R}^i\Gamma(\mathscr{F}).$$

In particular, $\mathrm{H}^0(X;\mathscr{F}) = \Gamma(X;\mathscr{F})$ recovers global sections, and for any short exact sequence of coherent sheaves

$$0 \to \mathscr{F}' \to \mathscr{F} \to \mathscr{F}'' \to 0,$$

there is a canonical long exact sequence of vector spaces

$$\begin{aligned}0 \longrightarrow \mathrm{H}^0(X;\mathscr{F}') \longrightarrow \mathrm{H}^0(X;\mathscr{F}) \longrightarrow \mathrm{H}^0(X;\mathscr{F}'') \\ \longrightarrow \mathrm{H}^1(X;\mathscr{F}') \longrightarrow \mathrm{H}^1(X;\mathscr{F}) \longrightarrow \mathrm{H}^1(X;\mathscr{F}'') \\ \longrightarrow \mathrm{H}^2(X;\mathscr{F}') \longrightarrow \mathrm{H}^2(X;\mathscr{F}) \longrightarrow \mathrm{H}^2(X;\mathscr{F}'') \longrightarrow \cdots.\end{aligned}$$

Of course the construction of sheaf cohomology works entirely analogously[11] when X is not a compact Riemann surface, i.e. algebraic curve, but an arbitrary complex space, or scheme, etc.

[10]Furthermore, the resolution must be *F-acyclic*, which means that $\mathrm{R}^i F(A^j) = 0$ for all $i \geq 1$ and $j \geq 0$. Thus, in order for this process of computing derived functors to have any chance of being useful, we must first understand F-acyclic objects in $\mathcal{A}$.

[11]Furthermore, we need not have worked in the context of coherent sheaves. Sheaves of $\mathscr{O}_X$-modules would be just as good, though then the vanishing results of section 4.13. would not automatically apply. Similarly, when working in the algebro-geometric context, we could work with the abelian category of quasi-coherent sheaves, which is larger and often a little more convenient than the category of coherent ones.

4.13. **Vanishing above the dimension.** But the situation is even a little nicer than the general story would let us believe! If X is a d-dimensional complex manifold/smooth scheme, then sheaf cohomology vanishes above degree d. That is, if $\mathscr{F}$ is a coherent sheaf on X, then $\mathrm{H}^i(X;\mathscr{F}) = 0$ for all $i > d$.

In the case of Riemann surfaces that we are concerned with in this note, we have $d = 1$, and so the only non-vanishing sheaf cohomology groups are H^0 and H^1.

5. Cohomological aspects of Riemann-Roch

Now equipped with the technology of sheaf cohomology, we sketch the standard cohomological approach to proving the Riemann-Roch theorem, that one may find in most textbooks, e.g. [Fors] or [Hart].

5.1. **The Euler characteristic in topology.** You likely recall the Euler characteristic from your topology class, but you may have encountered it earlier.

Its origin is in the study of polyhedra. For a polyhedron with v vertices, e edges, and f faces, its Euler characteristic is defined as $\chi = v - e + f$. This is useful for many applications, for instance in the classical theory of surfaces, and in graph theory.

For topological purposes, it may be generalized to any simplicial complex, as an alternating sum of the number of the number of similices it possesses in each dimension. Even more generally, for an d-dimensional manifold M, its Euler characteristic may be defined as

$$\chi(M) = \sum_{i=0}^{d} (-1)^i \dim \mathrm{H}^i(M;\mathbf{C}).$$

For a polyhedron or a simiplicial complex, this recovers the older notion, but it has the advantage of being phrased entirely in terms of cohomology, which makes sense for any space, and not any extra structure, such as a triangulation.

5.2. **The Euler characteristic of a sheaf.** We imitate this definition is sheaf cohomology. Returning to the context of section 4.13., let X be a d-dimensional complex manifold/smooth scheme and $\mathscr{F}$ a coherent sheaf on it. We set its *Euler characteristic* to be

$$\chi(\mathscr{F}) = \sum_{i=0}^{d} (-1)^i \dim \mathrm{H}^i(X;\mathscr{F})$$

In analogy with the topological Euler characteristic of section 5.1., we might better use the notation $\chi(X;\mathscr{F})$, but we are mostly sloppy and supress the underlying complex manifold/smooth scheme X, which is presumed to be understood from context.

Specializing even further to the context of a compact Riemann surface/algebraic curve, that we are chiefly interested in here, the Euler characteristic is just

$$\chi(\mathscr{F}) = \dim \mathrm{H}^0(X;\mathscr{F}) - \dim \mathrm{H}^1(X;\mathscr{F}).$$

5.3. **Additivity of Euler characteristic.** What makes Euler characteristic a useful way of packaging the information of sheaf cohomology, is its additivity. We explain what this means in this section. Let

$$0 \to \mathscr{F}' \to \mathscr{F} \to \mathscr{F}'' \to 0$$

be a short exact sequence of coherent sheaves on X a d-dimensional complex manifold/smooth scheme. Then we know from section 4.12. that sheaf cohomology gives rise to a long exact sequence, which in light of the vansihing result of section 4.13. has the form

$$
\begin{array}{l}
0 \longrightarrow \mathrm{H}^0(X;\mathscr{F}') \longrightarrow \mathrm{H}^0(X;\mathscr{F}) \longrightarrow \mathrm{H}^0(X;\mathscr{F}'') \\
\hookrightarrow \mathrm{H}^1(X;\mathscr{F}') \longrightarrow \mathrm{H}^1(X;\mathscr{F}) \longrightarrow \mathrm{H}^1(X;\mathscr{F}'') \cdots \\
\cdots \to \mathrm{H}^d(X;\mathscr{F}') \longrightarrow \mathrm{H}^d(X;\mathscr{F}) \longrightarrow \mathrm{H}^d(X;\mathscr{F}'') \longrightarrow 0.
\end{array}
$$

Since any exact sequence of vector spaces satisfies the property that the alternating sum of the dimensions of its terms is zero (fun and easy-enough exercise in linear algebra!), we find by "summing each of the three columns" that

$$\chi(\mathscr{F}') - \chi(\mathscr{F}) + \chi(\mathscr{F}'') = 0.$$

This is the additivity of the Euler characteristic. Not only when $\mathscr{F} = \mathscr{F}' \oplus \mathscr{F}''$, but also when $\mathscr{F}$ just fits into the middle of a short exact sequence with $\mathscr{F}'$ and $\mathscr{F}''$ on the sides, the Euler characteristic pretends as if it were a sum.

5.4. **Corrupted beyond all recognition.** Let us now firmly return to the contex we are actually considering in this note: let X be a compact Riemann surface.

In section 4.1. we rephrased the Riemann-Roch Problem as asking about the global sections $\Gamma(X;\mathscr{L})$ of a holomorphic line bundle $\mathscr{L}$, which is to say about $\dim \Gamma(X;\mathscr{L})$, in terms of the topology of X.

Alas, perhaps section 4.7. convinced us that global sections are not the most well-behaved thing to look at. Instead, motivated by the nice behavior we saw it exhibit in section 5.3., we might satisfy ourselves with asking about the Euler characteristic.

The Riemann-Roch Problem (4th version)**.** *Determine $\chi(\mathscr{L})$ in terms of data about the line bundle and the topology of X.*

Accepting this Riemann-Roch Problem into our life amounts to accepting that we will only get an answer up to an "error term" $\mathrm{H}^1(X;\mathscr{L})$.

That is not even all that radical: all the Riemann-Roch statements we saw so for, for instance (6) of section 3.5., had error terms in them as well! Those error terms were phrased in terms of holomorphic 1-forms, but an association game with de Rham cohomology might lead us to believe that H^1 might not be that far off. We will later see this articulated precisely in section 5.9. in the form of Serre duality.

5.5. **The key short exact sequence.** Thus we wish to understand $\chi(\mathscr{L})$ for a given (holomorphic) line bundle $\mathscr{L}$. How to approach this? Had we not discussed the Riemann-Roch theorem without the cohomological language previously, we might be out of ideas. Instead, we harken to section 1.5. for inspiration.

Recall the short exact sequence of coherent sheaves

$$0 \to \mathscr{O}_X \to \mathscr{L}(x) \to x_*(\mathbf{C}) \to 0,$$

that we already saw in Section 4.7, and whose global sections featured prominently in section 1.5. Since $\mathscr{L}$ is a line bundle, tensoring with it preserves short exact sequences, and so we obtain another short exact sequence

$$0 \to \mathscr{L} \to \mathscr{L} \otimes_{\mathscr{O}_X} \mathscr{L}(x) \to \mathscr{L} \otimes_{\mathscr{O}_X} x_*(\mathbf{C}) \to 0.$$

It is clear from looking at fibers that $\mathscr{L} \otimes_{\mathscr{O}_X} x_*(\mathbf{C})$ is still a skyscraper sheaf at x, whose fiber is the fiber of $\mathscr{L}$ at that point. Of course, $\mathscr{L}$ is a line bundle, and so the fiber in question is $\mathbf{C}$. We know from section 3.4. that $\mathscr{L} \simeq \mathscr{L}(D)$ for some divisor D on X, and so putting all this together, we find ourselves before the short exact sequence of coherent sheaves on X

$$0 \to \mathscr{L}(D) \to \mathscr{L}(D+x) \to x_*(\mathbf{C}) \to 0.$$

Using the additivity if the Euler characteristic from section 5.5., we get

$$\chi(\mathscr{L}(D+x)) = \chi(\mathscr{L}(D)) + \chi(x_*(\mathbf{C})). \tag{9}$$

To be able to make use of that, we must compute

5.6. **The Euler characteristic of the skyscraper.** Let us determine the Euler characteristic of the skyscraper sheaf $x_*(\mathbf{C})$. Clearly

$$\mathrm{H}^0(X; x_*(\mathbf{C})) = \Gamma(X; x_*(\mathbf{C})) = \mathbf{C},$$

while the vanishing property of section 4.13. may be used[12] with $d = 0$ to see that $\mathrm{H}^1(X; x_*(\mathbf{C})) = 0$. Indeed, while $x_*(\mathbf{C})$ is defined on X, it is supported entirely on the 0-dimensional subspace $\{x\} \subset X$. It may furthermore be shown that sheaf cohomology computed on this subspace $\{x\}$ will agree with the one computed on X. This justifies applying the vanishing property of section 4.13. Altogether, the Euler characterstic of the skyscraper sheaf is

$$\chi(x_*(\mathbf{C})) = 1 - 0 = 1.$$

5.7. **Easy cohomological Riemann-Roch.** Combining the computation from the previous section and the Euler sequence equality (9), we find that

$$\chi(\mathscr{L}(D+x)) = \chi(\mathscr{L}(D)) + 1.$$

Observe that both the divisor D and the point x were entirely arbitrary in the above discussion. Since any divisor may be obtained from the zero divisor D by adding/substracting $(\deg D)$-many points, it follows inductively that

$$\chi(\mathscr{L}(D)) = \chi(\mathscr{O}_X) + \deg D.$$

To make this into an solution of the Riemann-Roch Problem of section 5.5., it remains to determine $\chi(\mathscr{O}_X)$. We already know since section 1.1. that $\mathrm{H}^0(X; \mathscr{O}_X) = \mathscr{O}(X) = \mathbf{C}$, while we may just as well define the genus of X to be $g = \dim \mathrm{H}^1(X; \mathscr{O}_X)$.

Putting all of this together, we get the formula

$$\chi(\mathscr{L}(D)) = \deg D + 1 - g,$$

which is one incarnation of the Riemann-Roch theorem.

5.8. **The genus revisited.** Above in section 5.7. we elected to take the genus of a compact Riemann surface X to be $g = \mathrm{H}^1(X; \mathscr{O}_X)$.

This might not seem unreasonable, seeing how we might remember from our topology class that the genus of a topological surface X is precisely $g = \frac{1}{2} \dim \mathrm{H}^1(X; \mathbf{C})$. Indeed, you might recall the relationship $\chi(X) = 2 - 2g$ between the topological Euler characteristic and the genus. So in particular, there is a relationship between the genus and some kind of 1st cohomology.

[12] This also follows particularly easily from the Čech approach to computing sheaf cohomology.

The alert-minded reader may not be so trusting though. In section 1.9. mentioned that the genus may also be defined as $g = \dim \Omega^1(X)$. Why should these two different proceedures output the same number?

5.9. **Serre duality.** The answer to the question with which we concluded the last section comes in the form of *Serre duality*.

It says that, if $\mathscr{E}$ is a rank r holomorphic line bundle on a compact d-dimensional complex manifold/proper smooth scheme X, then there is a canonical isomorphism

$$\mathrm{H}^i(X;\mathscr{E}) \simeq \mathrm{H}^{n-i}(X;\mathscr{E}^\vee \otimes_{\mathscr{O}_X} \Omega^d_X)^\vee$$

for all integers $i \geq 0$.

This is a reasonably involved theorem, and we do not prove it here. The standard algebro-geometric approach to proving it reduces it to the case of $X = \mathbf{P}^d$, where vector bundles are easy to understand. When working in the complex analytic setting, a reasonably simple proof via Hodge theory is also available.

In any case, we apply Serre duality to the case when X is a compact Riemann surface, so that $d = 1$, and when $\mathscr{E} = \mathscr{L}$ is a line bundle, so that $r = 1$. The statement of Serre duality for $i = 1$ becomes

$$\mathrm{H}^1(X;\mathscr{L}) \simeq \mathrm{H}^0(X;\mathscr{L}^\vee \otimes_{\mathscr{O}_X} \Omega^1_X)^\vee.$$

Writing the bundle in terms of a divisor as $\mathscr{L} = \mathscr{L}(D)$, and recalling from the arguments in section 2.6. that $\Omega^1_X \simeq \mathscr{L}(K)$ (though we only proved this there on the level of global sections, as we noted in Chapter 3, the proof may be repeated over an arbitrary open, and so identifies the sheaves), this further simplifies to

$$\mathrm{H}^1(X;\mathscr{L}(D)) = L(K-D)^\vee.$$

Consequently the Euler characteristic, as computed by the easy cohomological Riemann-Roch formula of section 5.7., is nothing but

$$\begin{aligned}
\chi(\mathscr{L}(D)) &= \dim \mathrm{H}^0(X;\mathscr{L}(D)) - \dim \mathrm{H}^1(X;\mathscr{L}(D)) \\
&= \dim \Gamma(X;\mathscr{L}(D)) - \dim \Gamma(X;\mathscr{L}(K-D))^\vee \\
&= \ell(D) - \ell(K-D),
\end{aligned}$$

recovering the "full" Riemann-Roch formula from section 2.7. This concludes the usual cohomological proof of the Riemann-Roch Theorem. □

5.10. **Remarks about this proof.** This cohomological approach to proving the Riemann-Roch has become standard in the literature, no doubt because Serre duality, encountered along the way, is such an important result in its own right. This proof seems to be due to Serre, first appearing in [Serre], alongside the articulation and proof of Serre dulity.

This proof is sometimes, most infamously in the introduction to the chapter on curves in [Hart], criticised as unenlightening. We respectfully disagree. As we tried to make abundantly clear in section 5.5., the key idea of the proof is still the same residue trick we first encountered in section 1.5., albeit disguised in the language of cohomology. The use of Serre duality is also merely a profound elaboration on Roch's contribution to the Riemann-Roch story from section 2.2.

Nonetheless, as shown by Serre, Hirzebruch, and Grothendieck, the cohomological interpretation of the theorem is indispensible in understanding how to generalize the Riemann-Roch theorem from Riemann surfaces to higher dimensions.

Appendix

We collect some remarks that do not fit neately into the above discussion, but that might nonetheless be helpful in dispelling confusion.

A. **Degree of principal divisors.** It may seem to one at first glance at the divisor class group in section 3.2. that it should be trivial. Indeed, given any finite collection of points in the complex plane $\mathbf{C}$, one may find a holomorphic function which vanishes there (or has poles, with arbitrary specified degree, etc.). Alas, we are working on a *compact* Riemann surface X, and this makes all the difference.

Indeed, to convince ourselves that many divisors are not principal, note that the degree of any principal divisor is zero. For a non-zero meromorphic function f on X, we have

$$\deg \operatorname{div}(f) = \#\text{ zeros of } f - \#\text{ poles of } f,$$

counted with multiplicity. Recall now the *Argument Principle* from your complex analysis class: if f is meromorphic on a (neighborhood of a) relatively compact domain $U \subseteq \mathbf{C}$, so that all its poles and zeros are contained inside U, then

$$\#\text{ zeros of } f - \#\text{ poles of } f = \frac{1}{2\pi i}\int_{\partial U} \frac{f'(z)}{f(z)}\, dz,$$

where the left-hand side must be counted with multiplicites. The integrand on the right-hand side may be rewritten as $d(\log f)$, showing that the issue of section 1.7. does not apply. Indeed, covering X by such holomorphic local charts U, we obtain a global result valid on X. But since X does not have any boundary, which is to say, because the contributions from different ∂U cancel out, we conclude that $\#$ zeros of $f - \#$ poles of $f = 0$ for any non-zero meromorphic function f. Thus $\deg \operatorname{div}(f) = 0$.

It follows that the notion of the degree $\deg D$ of a divisor D indeed depends only on its divisor class. This was used implicitly in section 3.3. to make sense of the degree of a line bundle.

B. **What Chern classes have to do with Chern classes.** Use of the term "first Chern class" in section 3.3. may seem rather vexing to readers familiar with Chern classes from topology.

In algebraic topology, a rank r complex vector bundle E on a topological space X get Chern classes associated to it as certain characteristic classes $c_i^{\mathrm{top}}(E) \in \mathrm{H}^*(X;\mathbf{Z})$ for integers $0 \le i \le r$. In particular, when L is a complex line bundle on a topological surface X, its topological first Chern class is an element $c_1^{\mathrm{top}}(L) \in \mathrm{H}^2(X;\mathbf{Z})$. As with the algebraic first Chern class of section 3.3. the topological first Chern class determines the isomorphism class of a complex line bundle. The difference is thus in whether we are considering the complex or the holomorphic vector bundle. Let us explain the relationship more precisely.

Consider the inclusion $\mathbf{Z} \hookrightarrow \mathscr{O}(X)$ of contant $\mathbf{Z}$-valued functions on a Riemann surface X into its holomorphic functions. This gives rise to the short exact sequence of sheaves of abelian groups on X

$$0 \to \mathbf{Z} \to \mathscr{O}_X \xrightarrow{\exp} \mathscr{O}_X^\times \to 1,$$

where exp really stands for $f \mapsto e^{2\pi i f}$. This is called the *exponential sequence*, and it encodes much about the complex geometry of X. Zooming in on the connecting map in the induced long exact sequence on sheaf cohomology, we obtain a group homomorphism

$$\mathrm{H}^1(X;\mathscr{O}_X^\times) \to \mathrm{H}^2(X;\mathbf{Z}).$$

The left-hand side may be identified with the Picard group of holomorphic line bundles on X up to isomorphism.

Indeed, suppose we are given a holomorphic line bundle $\mathscr{L}$ on X, and a covering of X by holomorphic charts (z_i, U_i) over which $\mathscr{L}$ trivializes. That is to say, we are given isomorphisms $\mathscr{L}|_{U_i} \simeq \mathscr{O}_{U_i}$, and composeng these for each pair of indices i, j, we may extract its *transition maps* $f_{ij} \in \Gamma(U_{ij};\mathscr{O}_{U_{ij}}^\times)$. Together, $\{f_{ij}\}_{i,j}$ is the *gluing cocycle* for $\mathscr{L}$, and determines the holomorphic line bundle. On the other hand, through the mehcanism of Čech cohomology, $\{f_{ij}\}$ represents an element in $\mathrm{H}^1(X;\mathscr{O}_X^\times)$. It can be shown that this proceedure yields the desired isomorphism.

Combing this with the isomorphism of section 3.3., we obtain a homomorphism

$$\mathrm{Cl}(X) \simeq \mathrm{H}^1(X;\mathscr{O}_X^\times) \to \mathrm{H}^2(X;\mathbf{Z}),$$

and this is the map under which $c_1(\mathscr{L}) \mapsto c_1^{\mathrm{top}}(\mathscr{L})$ for any holomorphic line bundle $\mathscr{L}$ on X, relating the algebraic and topological first Chern class.

Here $\mathrm{H}^2(X;\mathbf{Z})$ may similarly be identified with the group of topological complex line bundles on X up to isomorphism. This may be shown by an analogous gluing cocyle construction as for holomorphic line bundles above. But for those familiar with classifying spaces, it may also be shown nicely as

$$\mathrm{H}^2(X;\mathbf{Z}) \simeq \pi_0\mathrm{Map}(X,\mathrm{B}^2\mathbf{Z}) \simeq \pi_0\,\mathrm{Map}(X,\mathrm{BU}(1)) \simeq \pi_0\mathrm{Bun}_{\mathrm{U}(1)}(X),$$

where second isomorphism is because $\mathrm{U}(1) \simeq \mathrm{B}\mathbf{Z}$ and the third one because the classifying space $\mathrm{B}G$ classifies G-bundles. Finally we may use the standard representation of $\mathrm{U}(1)$ on $\mathbf{C}$ to identify principal $\mathrm{U}(1)$-bundles and complex line bundles. It follows that we have spent this section discussing the commutative square

$$\begin{array}{ccc} \mathrm{LinBun}(X) & \longrightarrow & \mathrm{LinBun}^{\mathrm{top}}(X) \\ {\scriptstyle c_1}\downarrow{\scriptstyle \simeq} & & {\scriptstyle c_1^{\mathrm{top}}}\downarrow{\scriptstyle \simeq} \\ \mathrm{Cl}(X) & \longrightarrow & \mathrm{H}^2(X;\mathbf{Z}) \end{array}$$

whose top horizontal map discards the holomorphic structure of a holomorphic line bundle and only remembers the underlying complex line bundle.

References

[Fors] O. Forster, *Lectures on Riemann Surfaces*, Graduate Texts in Mathematics **81**, Springer-Verlag, New York, Berlin 1991.

[Harr] Michael Harris. *Geometry of Algebraic Curves, Fall 2011, Course taught by Joe Harris, Notes by Atanas Atanasov.* Lecture notes, available at http://staff.math.su.se/shapiro/UIUC/curvesHarris.pdf

[Hart] Robin Hartshorne. *Algebraic Geometry.* Graduate Texts in Mathematics **52**, Springer-Verlag, New York, Heidelberg 1977.

[Riem] Bernhard Riemann. *Theorie der Abel'schen Functionen.* Journal für die reine und angewandte Mathematik. 54: 115–155, (1971).

[Roch] Gustav Roch. *Ueber die Anzahl der willkurlichen Constanten in algebraischen Functionen.* Journal für die reine und angewandte Mathematik. 64: 372–376, (1865).

[Serre] J.-P. Serre, *Un théorème de dualité*, Comment. Math. Helv. **29**, 1955, str. 9–26.