DIVISORS AND THE ABEL-JACOBI MAP

SAYANTAN KHAN

Contents

A divisor on a compact Riemann surface is just a labelling of finitely many points with an integer. This may not seem like a lot of data, but a divisor manages to encode a fair amount of information. In this expository article, we'll see how a divisor is really the same thing as a map into a projective space, and a complex line bundle on a surface. This will let us translate properties of line bundles like the existence of enough sections into statements about divisors, which are easier to verify. We will also construct a complex manifold called the *Jacobian* of the Riemann surface, and show that it is isomorphic to the collection of divisors (modulo some relations) via the *Abel-Jacobi map*. The Jacobian turns out to be a manifold which is easy to understand and provides some insight into the group of divisors.

Many of the results that we show here will also hold for non-singular curves over any algebraically closed field, not just $\mathbb{C}$. Some of the proof techniques however don't go through, especially if they involve integration of forms. We'll point out whenever something like this does happen and outline alternative approaches to proofs that work over other fields.

1. Divisors

We shall denote a compact Riemann surface by the symbol X in this and the subsequent sections.

Definition 1.1 (Divisor)**.** A divisor on a compact Riemann surface X is a finitely supported function $D : X \to \mathbb{Z}$. The collection of points p where $D(p) \neq 0$ is called the *support* of the divisor D

The collection of all divisors $\text{Div}(X)$ forms an abelian group under pointwise addition. To emphasize the fact that divisors form a group, we shall usually denote a divisor as a finite $\mathbb{Z}$-linear combination of points in X.

$$D := \sum_{p \in X} D(p) \cdot p$$

The most important examples of divisors come from meromorphic functions (i.e. holomorphic maps to $\mathbb{CP}^1$) and meromorphic 1-forms on X.

Definition 1.2 (Principal divisor)**.** Let f be a non-zero meromorphic function on X. Since X is compact, f has finitely many zeroes and poles. Let $\{z_1, \ldots, z_k\}$ be its zeroes and $\{p_1, \ldots, p_l\}$ be its poles, both counted with multiplicities. Then we can define a divisor (f), called the *principal divisor* of f.

$$(f) := \sum_{i=1}^{k} z_i - \sum_{j=1}^{l} p_j$$

Example 1.1. Let $X = \mathbb{CP}^1$, let f be the function $\dfrac{z-a_1}{(z-a_2)^2}$. This function has a zero of order 1 at a_1 and another zero of order 1 at ∞. It has a pole of order 2 at a_2. The principal divisor (f) associated to this function is $a_1 + \infty - 2 \cdot a_2$.

In a similar manner, we can associate a divisor to a meromorphic 1-form.

Definition 1.3 (Canonical divisor)**.** Let ω be a meromorphic 1-form on X. Let $\{z_1, \ldots, z_k\}$ be its zeroes and $\{p_1, \ldots, p_l\}$ be its poles, both counted with multiplicities. Then we can define a divisor (ω), called the *canonical divisor* of ω.

$$(\omega) := \sum_{i=1}^{k} z_i - \sum_{j=1}^{l} p_j$$

Example 1.2. Consider $\mathbb{CP}^1$ with two coordinates z and w around 0 and ∞ respectively, related by $z = \frac{1}{w}$. We can define a form dz around 0, which transforms to $\frac{1}{w^2}dw$. That means this form has no zeroes, and a pole of order 2 at ∞. The canonical divisor (ω) is given by $-2 \cdot \infty$.

Example 1.3. Let X be a torus, i.e. a quotient of $\mathbb{C}$ by a lattice $(1, \tau)$, where $\tau \in \mathbb{H}^+$. The form dz on $\mathbb{C}$ descends to a form dz on the torus. This form has no zeroes or poles anywhere, which means the associated canonical divisor is just the empty divisor, i.e. no point of X is in the support.

For any divisor, we can take the coefficients of each point of X and add them up. Since the support of a divisor is always finite, this gives us an integer called the *degree* of the divisor.

Definition 1.4 (Degree of a divisor)**.** Let D be a divisor on X. Then the degree of the divisor is defined in the following manner.

$$\deg(D) := \sum_{p \in X} D(p)$$

The subgroup of $\mathrm{Div}(X)$ consisting of degree 0 divisors is denoted by $\mathrm{Div}_0(X)$. These two groups fit into the following short exact sequence.

$$0 \longrightarrow \mathrm{Div}_0(X) \longrightarrow \mathrm{Div}(X) \xrightarrow{\deg} \mathbb{Z} \longrightarrow 0$$

If we pick a basepoint x_0 of X, then there's a map $\mathrm{Div}(X)$ to $\mathrm{Div}_0(X)$ given by $D \mapsto D - \deg(D) \cdot x_0$. This map is a splitting of the short exact sequence. We get a decomposition of $\mathrm{Div}(X)$ as $\mathrm{Div}_0(X) \oplus \mathbb{Z}$. This means if we want to understand $\mathrm{Div}(X)$, it suffices to understand $\mathrm{Div}_0(X)$.

The divisor we constructed in Example 2.2 and 1.3 have degree 0, whereas the one in Example 1.2 has degree equal to -2. Notice that the principal divisor we constructed turned out to have degree 0. This is not a coincidence: any other principal divisor will also turn out to have degree 0.

Lemma 1. *Let (f) be a principal divisor on any compact Riemann surface. Then the degree of (f) is 0.*

Proof. Recall that the principal divisor (f) is $\sum_i z_i - \sum_j p_j$, where z_i and p_j are zeroes and poles counted with multiplicities. To show that the degree is 0, we'll need to argue that the number of zeroes equals the number of poles. In fact, a slightly stronger result is true: the function ϕ that sends $p \in \mathbb{CP}^1$ to the number of points in $f^{-1}(p)$ (counted with multiplicities) is a constant function. To show this, it will suffice to show that ϕ is locally constant, since $\mathbb{CP}^1$ is connected.

To show ϕ is locally constant, pick any point $p \in \mathbb{CP}^1$ and let $p_1, \ldots, p_k$ be its preimages, counted *without* multiplicities, i.e. just the set $f^{-1}(p)$. Let the multiplicity of the preimage p_i be d_i. The value of the function ϕ at p is thus $\sum_i d_i$. Around each of these points, we can pick a coordinate chart such that the map f looks like $z \mapsto z^{d_i}$ in the local coordinates. If we perturb p by a small amount to p', all preimages of p' will also be close to p. In fact, something better happens: in a chart where the map f looked like $z \mapsto z^{d_i}$, there will be d_i many preimages of p', and all of them will have multiplicity 1. This follows from a result in complex analysis which states that for $p' \neq 0$, the function $z^d - p'$ has d simple zeroes. What this means for us is that the number of preimages of p' are $\sum_i d_i$, which was the same as the number of preimages of p counted with multiplicities. We have shown that for a p' close enough to p, $\phi(p) = \phi(p')$, which means ϕ is locally constant. □

The above proof can also be made completely algebraic. The key idea stays the same, i.e. the number of preimages of any point is a constant function. In fact, that constant number turns out to be $[\mathcal{M}(X) : \mathcal{M}(\mathbb{CP}^1)]$, where $\mathcal{M}(X)$ and $\mathcal{M}(\mathbb{CP}^2)$ are the field of meromorphic functions on these spaces. Theorem 1.5 in [7] fills in more details.

Principal divisors also behave nicely with respect to multiplication with functions, i.e. they satisfy the following identities for any meromorphic functions f and g.

$$(fg) = (f) + (g)$$

In particular, this shows the set of principal divisors $\mathrm{PDiv}(X)$ forms a subgroup of $\mathrm{Div}(X)$ and also of $\mathrm{Div}_0(X)$ (by Lemma 1). To understand these groups better, it's useful to look at the quotient groups $\mathrm{Pic}(X) := \mathrm{Div}(X)/\mathrm{PDiv}(X)$ and $\mathrm{Pic}_0(X) := \mathrm{Div}(X)/\mathrm{PDiv}(X)$. These groups are called the *Picard groups* of the Riemann surface X. As it turns out, the quotient group $\mathrm{Pic}(X)$ and the principal divisors $\mathrm{PDiv}(X)$ are far easier to understand than the infinitely generated group $\mathrm{Div}(X)$. Another reason why the Picard group is interesting is because a lot of divisors that show up naturally are only well defined up to a principal divisor. In that case, it makes sense to think of these naturally occurring divisors as living in $\mathrm{Pic}(X)$, rather than $\mathrm{Div}(X)$.

The next thing to consider are spaces of functions associated with divisors. In particular, we'll look at the space $L(D)$ for a given divisor D.

Definition 1.5. The space $L(D)$ is the vector space consisting of all meromorphic functions f such that $(f)+D \geq 0$, i.e. the divisor $(f) + D$ has a non-negative value for all points $p \in X$.

Now suppose we pick another divisor $D+(g)$. It then turns out that the space $L(D)$ and $L(D+(g))$ are isomorphic vector spaces, where the isomorphism is given by $f \mapsto fg^{-1}$. This is an example of a natural construction which descends to the Picard group.

The space $L(D)$ will crop up again in the future, and we'll need the following lemma.

Lemma 2. *If D is a divisor of negative degree, then the space $L(D)$ consists only of the function 0, i.e. it's trivial.*

Proof. Suppose there was some meromorphic function g such that $(g)+D \geq 0$. That would mean $\deg((g))+\deg(D) \geq 0$. But the degree of a principal divisor is 0, and $\deg(D)$ is negative, which leads to a contradiction.[1] □

Another natural construction that descends to the Picard group is the collection of canonical divisors. Let ω and γ be two meromorphic 1-forms on a Riemann surface X. Then there exists a meromorphic function f such that $\gamma = f\omega$. This follows from the fact that in one local chart where $\omega = H(z)dz$ and $\gamma = G(z)dz$, $\omega = \frac{H(z)}{G(z)}\gamma$. On some other coordinates w, where $z = c(w)$, the forms become $\omega = H(c(w))c'(w)dw$ and $\gamma = G(c(w))c'(w)dw$. In this coordinate chart, $\omega = \frac{H(c(w))}{G(c(w))}$, but this function equals $\frac{H(z)}{G(z)}$ on the overlap of the two charts. Any two meromorphic functions that agree on an open set must agree everywhere, which means $\omega = \frac{H}{G}\gamma$ globally. What this means in particular is that the canonical divisors (ω) and (γ) go to the same element in $\mathrm{Pic}(X)$. This is why a canonical divisor is often called *the* canonical divisor, and denoted by K.

We can now state the Riemann-Roch theorem, but we will not prove it.

Theorem 3 (Riemann-Roch)**.** *If X is a compact Riemann surface with genus g, and D is a divisor on X, then the dimensions of $L(D)$ and $L(K-D)$ are related in the following manner.*

$$\dim(L(D)) - \dim(L(K-D)) = \deg(D) - g + 1$$

This theorem has several consequences; one consequence will be especially important when analyzing the Abel-Jacobi map later in the article, and that is the following.

Corollary 4. *The space of holomorphic 1-forms $\Omega^1(X)$ on a Riemann surface X of genus g is g-dimensional.*

Proof. If D is the empty divisor, then the space $L(D)$ is just the space of all holomorphic functions (because the empty divisor forces it to have no poles). But the space of holomorphic functions on X is just $\mathbb{C}$, hence 1-dimensional. Plugging in the empty divisor into the Riemann-Roch formula, and using the fact we just mentioned, we get the following.

$$\dim(L(K)) = g$$

Suppose now that $K = (\omega)$ for some meromorphic 1-form ω. Then for any $f \in L(K)$, $f\omega$ will not have any poles, because if it did, $(f) + (\omega)$ would not be non-negative everywhere. That means $f\omega$ is a holomorphic 1-form. Conversely, suppose γ is a holomorphic 1-form. It can be written as $g\omega$, where g is some meromorphic function. Since $(g) + (\omega) = (\gamma) \geq 0$, $g \in L(K)$. This shows that the space of holomorphic 1-forms is isomorphic to $L(K)$ and thus has dimension g. □

[1] In all earnestness, we haven't defined what the principal divisor associated to the 0 function means. One can define it abstractly to be the 0 vector in the space $L(D)$. A more intuitive, but informal way of thinking about it is as a divisor supported everywhere on X with its value at every point being $+\infty$. That way $((0)) + D$ is always non-negative, no matter what D is.

These results hold for algebraic curves over any closed fields. When working specifically with Riemann surfaces, it's possible to bypass Riemann-Roch entirely, and get this result via Hodge theory. Hodge theory gives a decomposition of $H^1_{\mathrm{dR}}(X)$ as $H^{1,0}(X) \oplus H^{0,1}(X)$, where $H^{1,0}(X)$ is the space of holomorphic forms, and $H^{0,1}(X)$ the space of anti-holomorphic forms. Since $H^1_{\mathrm{dR}}(X)$ is $2g$-dimensional, the space of holomorphic 1-forms turns out to be g dimensional.

2. Maps into Projective Spaces

We have already seen examples of maps from Riemann surfaces to projective spaces, namely meromorphic maps, which are holomorphic maps from the Riemann surface to $\mathbb{CP}^1$. In terms of coordinates on $\mathbb{CP}^1$, a map like this can be thought of as sending z to $[1 : f(z)]$. This makes sense at all points except at the poles of f, where the coordinates end up looking like $[1 : \infty]$. But this function can be extended holomorphically over the poles. Around each pole, choose local coordinates such that $f(z) = \frac{1}{z^d}$. For all non-zero z in this chart, $[1 : f(z)] = \left[1 : \frac{1}{z^d}\right] = [z^d : 1]$. The last expression is valid even at the point $z = 0$, and its clear that this is a holomorphic extension.

This suggests what a holomorphic map into $\mathbb{CP}^n$ should look like: it should locally like a holomorphic map into a domain in $\mathbb{C}^n$.

Definition 2.1 (Holomorphic maps to $\mathbb{CP}^n$)**.** A map f from a compact Riemann surface X to $\mathbb{CP}^n$ is holomorphic if around each point $p \in X$ we can find a coordinate chart such that f in that coordinate chart looks like the following.

$$f(z) = [f_0(z) : f_1(z) : \cdots : f_n(z)]$$

Furthermore, one of the f_i, say f_k should not vanish anywhere on the coordinate chart, and the functions $\frac{f_i}{f_k}$ should be holomorphic.

The first thing we need to verify is that under this definition of a holomorphic mapping into $\mathbb{CP}^n$, there actually exist some non-trivial functions.

Example 2.1. Let $\{f_0, \ldots, f_n\}$ be meromorphic functions on X. Then the map given by $f(z) = [f_0(z), \ldots, f_n(z)]$ is a well defined map everywhere except at the common zeroes of f_i, and anywhere where one of the f_i has a pole. It's clear that outside of such points, the map f is holomorphic. We can extend it holomorphically across the problematic points as well. Suppose $p \in X$ is a common zero of all the f_is. We can pick local coordinates around p such that $f_i(z) = z^{d_i}(g_i(z))$, where $g_i(z)$ doesn't vanish at p. Then the map around p looks like the following.

$$z \mapsto \left[z^{d_0} g_0(z) : \cdots : z^{d_n} g_n(z)\right]$$

Suppose d_k is the minimum of the d_is. Then the holomorphic extension across p is given by the following.

$$z \mapsto \left[z^{d_0 - d_k} g_0(z) : \cdots : g_k(z) : \cdots : z^{d_n - d_k} g_n(z)\right]$$

The function can be similarly extended across the poles. This gives an example of a non-constant holomorphic function into $\mathbb{CP}^n$.

We now have a large class of holomorphic maps into $\mathbb{CP}^n$. However, it turns out that every holomorphic function into $\mathbb{CP}^n$ is of the form we just described. Furthermore, its expression as ratios of $n+1$ meromorphic functions is unique up to multiplication by a meromorphic function.

Proposition 5 (Proposition 4.3 from [6])**.** *Let $f : X \to \mathbb{CP}^n$ be a holomorphic map. Then there exist $n+1$ meromorphic functions $\{f_0, \ldots, f_n\}$ such that $f(z) = [f_0(z) : \cdots : f_n(z)]$. If $\{g_0, \ldots, g_n\}$ are another collection such that $f(z) = [g_0(z) : \cdots : g_n(z)]$, then there exists a meromorphic function m such that $f = mg$.*

Proof. The first step in the proof is to actually find the candidates for the meromorphic functions f_i. To do that, we can assume without loss of generality that for some $p \in X$, the x_0-coordinate of $\mathbb{CP}^n$ is non-zero, otherwise we could rearrange the coordinates. We have $n+1$ functions $\{\pi_0, \pi_1, \ldots, \pi_n\}$ from $\mathbb{CP}^n$ to $\mathbb{CP}^1$, i.e. a point $[x_0 : x_1 : \cdots : x_n]$ gets mapped to $\frac{x_i}{x_0}$ by the function π_i. We compose these functions with f to get $f_i := \pi_i \circ f$ which is a map from X to $\mathbb{CP}^1$. We claim that these functions f_i are meromorphic, and are our candidate functions, i.e. $f(z) = [f_0(z) : \cdots : f_n(z)]$. The latter claim is easier to verify: since f_i is defined $\frac{f(z)_i}{f(z)_0}$, where $f(z)_i$ is the i^{th} coordinate of f, we can multiply out by $f(z)_0$ to get f, and doing that doesn't affect the projective coordinates.

What we do need to verify is that f_is are meromorphic functions. This property can be checked locally, and this is where we use the fact that maps to $\mathbb{CP}^n$ locally look like $[h_0(z) : \cdots : h_n(z)]$, where the h_i are holomorphic. In that case $f_i(z) = \frac{h_i(z)}{h_0(z)}$, i.e. a ratio of two holomorphic functions, which is clearly meromorphic.

Suppose we have another collection of functions g_i such that $[f_0(z) : \cdots f_n(z)] = [g_0(z) : \cdots : g_n(z)]$ for all $z \in X$. Locally, we can look at the ratio $m(z) = \frac{f_0(z)}{g_0(z)}$, which will be a meromorphic function. Since $[f_0(z) : \cdots f_n(z)] = [g_0(z) : \cdots : g_n(z)]$, we must have that for any i, $\frac{f_i(z)}{g_i(z)} = m(z)$. This proves uniqueness modulo $\mathcal{M}(X)$. □

Since any map ϕ from X to $\mathbb{CP}^n$ is given by $n+1$ meromorphic functions, we can define what is called the *linear system* of the map ϕ.

Definition 2.2 (Linear system of a holomorphic map)**.** If $\phi : X \to \mathbb{CP}^n$ is given by $z \mapsto [f_0(z) : \cdots : f_n(z)]$, then let D be the divisor such that $-D(p) = \min((f_i)(p))$ for all $p \in X$. Let V_ϕ be the linear span of the meromorphic functions $\{f_0, \ldots, f_n\}$. Then the linear system $|\phi|$ defined by the map ϕ is the following set of divisors.

$$|\phi| = \{(g) + D \mid g \in V_\phi\}$$

It's not quite clear at this point that this definition is useful, or even that the linear system $|\phi|$ is well defined, because the definition involves choosing $\{f_0, \ldots, f_n\}$. This is where Proposition 5 comes in. Any other choice of meromorphic functions would differ from $\{f_0, \ldots, f_n\}$ by multiplication with a meromorphic function, i.e. it would look something like $\{mf_0, \ldots, mf_n\}$. Let $-D = \min((f_i))$ and let $-D' = \min(mf_i)$. Then $D' = D - (m)$. Any non-negative divisor of the form $(g')+D'$ where g' is in the span of $\{mf_0, \ldots, mf_n\}$ will become $(m)+(g)+D'-(m)$, where g is now in the span of $\{f_0, \ldots, f_n\}$. The (m) cancels out the negative (m) giving the same divisor, and shows that $|\phi|$ is well defined.

We also need to understand why this is a useful construction. Paraphrasing [4], a linear system $|\phi|$ encodes an external construction like a map into $\mathbb{CP}^n$ into something intrinsic to the surface X, i.e. a collection of divisors on it.

One can also construct linear systems purely in terms of some divisor D. Consider a vector space V of meromorphic functions such that for any $g \in V$, $(g) + D \geq 0$. In the case of $|\phi|$, the vector space we consider is the vector space V_ϕ for some fixed representation of ϕ as $[f_0 : \cdots : f_n]$. The question we should ask ourselves now is that when does an abstract linear system arise as a linear system of some map. The answer to that is when it is *base point free.*

Definition 2.3 (Base point free linear systems)**.** A linear system Q is said to be base point free if for all $p \in X$, there is some divisor $E \in Q$ such that $E(p) = 0$.

For a linear system arising from a map into $\mathbb{CP}^n$, it's not too hard to see that it's base point free. That's because for any point $p \in X$, we can pick $\{f_0, \ldots, f_n\}$ such that at least one of $f_i(p)$ is not equal to 0. The converse is not too hard to show.

Proposition 6 (Proposition 4.15 from [6])**.** *If Q is a base point free linear system, then $Q = |\phi|$ for some holomorphic map $\phi : X \to \mathbb{CP}^n$. Furthermore, ϕ is unique up to a choice of coordinates in $\mathbb{CP}^n$.*

Proof. Since Q is a linear system, its elements are of the form $(g) + D$ for some g coming from a vector space V. Let $\{f_0, \ldots, f_n\}$ be a basis of V. Define a map $\phi : z \mapsto [f_0(z) : \cdots : f_n(z)]$. We claim that $Q = |\phi|$. It's clear that $Q \subseteq |\phi|$. To show the inclusion the other way, we use the fact that Q is base point free. The linear system Q being base point free implies that $-D = \min((f_i))$. If this weren't the case, then for some p, $-D(p) < (f_i)(p)$ and hence $(g)(p) + D(p) > 0$ Since $-D = \min((f_i))$, Q must equal $|\phi|$.

Suppose Q was also equal to some other $|\phi'|$. We could then write the generators of Q as $(f_i) + D$ as well as $(f_i') + D'$. Here f_i is the basis for the vector space V of functions (f) such that $(f) + D \geq 0$, and f_i' is the analogous construction for D' coming from ϕ'. We could change coordinates on the ϕ' map to get $(f_i) + D = (f_i') + D'$ for all i. That means $(f_i) - (f_i')$ is a fixed divisor for all i, which means $\frac{f_i}{f_i'}$ is a fixed meromorphic function, and ϕ and ϕ' define the same map. □

This is the first real result we have managed to obtain from the theory of divisors. We have effectively parameterized the collection of maps into $\mathbb{CP}^n$ as a something intrinsic to X, i.e. base point free linear systems on X.

For any given divisor D, we can look at the linear system $|D|$ which is the set of all non-negative divisors of the form $(g) + D$. Then there's a simple condition D must satisfy for $|D|$ to be base point free.

Proposition 7 (Lemma 3.15 from [6])**.** *The linear system $|D|$ is base point free iff the following equality holds for all points $p \in X$.*

$$\dim(L(D - p)) = \dim(L(D)) - 1$$

Proof. Pick a local coordinate system around p. We can write out the Laurent series for every function in $L(D)$ as $cz^n + \Omega(z^{n+1})$, where $n = -D(p)$. Then the space $L(D-p)$ can be thought of as the kernel of the functional that sends functions in $L(D)$ to the coefficient of the z^n, i.e. c. This functional is non-zero since $|D|$ is base point free, because we can actually find an f such that $(f)(p) + D(p) = 0$, which means the Laurent series of f will have a non-zero c. Since the functional is non-zero, its kernel will have codimension 1. Conversely, if $L(D-p)$ has codimension 1, we know the functional is non-zero, i.e. for some f, $(f)(p) + D(p) = 0$. □

Given a base point free $|D|$, we can ask ourselves under what conditions is the associated map into $\mathbb{CP}^n$ is a nice map, i.e. an embedding.

Proposition 8 (Proposition 4.20 from [6])**.** *Let D be a divisor such that $|D|$ is base point free. Then the associated map $\phi_D : X \to \mathbb{CP}^n$ is an embedding iff the following identity holds for all points p and q in X.*

$$\dim(L(D-p-q)) = \dim(L(D)) - 2$$

Proof. We'll first show that ϕ_D is injective iff for distinct p and q, $\dim(L(D-p-q)) = \dim(L(D)) - 2$. Suppose $\phi_D(p) = \phi_D(q)$. We pick a basis of $L(D-p)$, $\{f_1, \ldots, f_n\}$ and extend it to a basis of $L(D)$, $\{f_0, \ldots, f_n\}$. Then the map given by $z \mapsto [f_0(z) : \cdots : f_n(z)]$ sends p to $[1 : 0 : \cdots : 0]$. This happens because at p, $(f_i)(p) > -D(p)$ for $i > 0$, but $(f_0)(p) = -D(p)$. Since $\phi_D(p) = \phi_D(q)$, the image of q must also be $[1 : 0 : \cdots : 0]$. The latter equality implies that $(f_i)(q) > (f_0)(q)$ for $i > 0$. That means the f_i for $i > 0$ generate $L(D-q)$, and $L(D-q) = L(D-p)$. Conversely, if $L(D-p) = L(D-q)$, then $\phi_D(p) = \phi_D(q)$. This means every function f such that $(f)(p) > -D(p)$ also satisfies $(f)(q) > -D(q)$, which in particular means $L(D-p) \subseteq L(D-p-q)$, which in turn implies $L(D-p-q) = L(D-p)$. Conversely, if $L(D-p) = L(D-p-q)$, then since $L(D-p-q)$ is codimension 1, it equals $L(D-q)$, which by transitivity, equals $L(D-p)$. We have thus shown that $\phi_D(p) = \phi_D(q)$ iff $L(D-p) = L(D-p-q)$. If we take the negation of these statements, we see that $\phi_D(p) \neq \phi_D(q)$ iff $L(D-p) \neq L(D-p-q)$. This happens only if $\dim(L(D-p-q)) = \dim(L(D)) - 2$, which proves the injectivity result.

To check whether an injective map from a compact Riemann surface is an embedding, it suffices to check whether the map is an immersion, i.e. the map on tangent spaces is an injection. To check this at any point $p \in X$, we can assume without loss of generality that $\phi_D(p) = [1 : 0 : \cdots : 0]$ via some change of coordinates on $\mathbb{CP}^n$. The coordinates on the image are given by $\frac{f_i(z)}{f_0(z)}$. We need to verify at least one of these coordinate functions has non-zero derivative at p, which will only happen if $\frac{f_i(p)}{f_0(p)}$ has a simple zero at p. In terms of divisors, that translates to $(f_i)(p) = (f_0)(p) + 1$. Notice that $L(D-p)$ is generated by $\{f_1, \ldots, f_n\}$. If some (f_i) was equal to $(f_0) + 1 \cdot p$, it wouldn't be in $L(D-2p)$, i.e. $L(D-2p)$ would not equal $L(D-p)$. Conversely, if $L(D-2p) \neq L(D-p)$, we could find an f_i that's in $L(D-p)$ but not in $L(D-2p)$. Then $\frac{f_i}{f_0}$ would be simple zero. We have thus shown that the condition for the map ϕ_D to be an immersion is that for all points p, we must have$L(D-2p) \neq L(D-p)$, i.e. $\dim(L(D-2p)) = \dim(L(D)) - 2$.

Thus for a map to be both injective and an immersion, we must have for all points p and q in X, $\dim(L(D-p-q)) = \dim(L(D)) - 2$. □

A divisor D for which $\dim(L(D-p-q)) = \dim(L(D)) - 2$ for all p and q is called a *very ample divisor*, and if some positive multiple of D is very ample, D is called ample. A very ample divisor corresponds to an embedding of the surface in some $\mathbb{CP}^n$, which means in most cases it suffices to find a very ample divisor if one is trying to embed a Riemann surface into $\mathbb{CP}^n$. Checking whether a divisor is very ample can be often done easily, via the Riemann-Roch theorem.

Example 2.2 (Very ample divisors)**.** Let X be a compact Riemann surface of genus g and let D be a divisor of degree greater than or equal to $2g+1$. It then turns out that D is a very ample divisor. Plug in the divisors D and $D-p-q$ into the Riemann-Roch formula to get the following equations.

$$\dim(L(D)) - \dim(L(K-D)) = \deg(D) - g + 1 \tag{1}$$

$$\dim(L(D-p-q)) - \dim(L(K-D+p+q)) = \deg(D-p-q) - g + 1 \tag{2}$$

If we subtract (2) from (1), we get the following.

$$\dim(L(D)) - \dim(L(D-p-q)) = 2 - (\dim(L(K-D+p+q)) - \dim(L(K-D)))$$

To proceed from here, we need to figure out the degree of the canonical divisor K. The degree of the canonical divisor on a surface of genus g turns out to be $2g-2$. We won't prove it here, but it's an easy consequence of the Riemann-Hurwitz formula. In that case, the divisor $K-D+p+q$ has degree -1, and the divisor $K-D$ has degree -3. By Lemma 2, both the dimension terms on the right hand side become 0, giving us the exact condition for D to be a very ample divisor.

3. Line bundles

So far, we have seen two kinds of functions on Riemann surfaces: the first kind are the holomorphic functions, which output an element of $\mathbb{C}$ for each point p in the Riemann surface, and the second kind are the holomorphic 1-forms, which output an element which locally looks like a multiple of dz, for some dz. It's not quite clear what sort of space the latter output lives in. What we do know is that the outputs of both of these kind of functions can be treated as elements of a vector space, i.e. they can be added and multiplied by a scalar globally. The notion of a holomorphic line bundle is a formalization of this idea, i.e. a collection of maps from a Riemann surface whose outputs form a 1-dimensional vector space. The line bundle is a space E with a map to the Riemann surface X, and maps from X to V satisfying certain conditions are elements of a vector space. Here's a more formal definition.

Definition 3.1 (Holomorphic line bundle)**.** A holomorphic line bundle (E, π) over X is a topological space E along with a map $\pi : E \to X$ which satisfies the following conditions.

(i) For any point $x \in X$, there exists an open set U (called the *locally trivial neighbourhood*) around x such that $\pi^{-1}(U)$ is isomorphic to $U \times \mathbb{C}$ via a map α (called the *local trivialization*) in a manner that makes the diagram in Figure 1 commute.

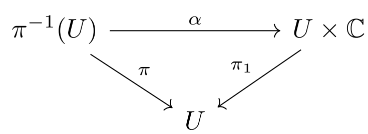

Figure 1. The local trivialization α commutes with the projections π and π_1.

(ii) For any two local trivializations α and β from $\pi^{-1}(U)$ to $U \times \mathbb{C}$, the map $\alpha \circ \beta^{-1}$ is $(\mathrm{id}, C(x))$, where $C(x)$ is a holomorphic map from U to $\mathrm{GL}(1, \mathbb{C})$. This map is called the transition map from (U, β) to (U, α).

Definition 3.2 (Sections of a line bundle)**.** If (E, π) is a holomorphic line bundle on X, a holomorphic section is a map s from X to E such that $\pi \circ s = \mathrm{id}$, and in any locally trivial neighbourhood U, with local trivialization α, $\pi_2 \circ \alpha \circ s$ is a holomorphic map from U to $\mathbb{C}$. A meromorphic (or rational) section of E is a section defined on all but finitely many points of X such that it locally looks like a meromorphic function.

We have already seen two examples of holomorphic line bundles and sections of them: the first one is $X \times \mathbb{C}$, whose sections are just the regular holomorphic functions, and the second is the cotangent bundle T^*X, whose sections are the holomorphic 1-forms.

The next thing one should define now are maps from line bundles.

Definition 3.3 (Line bundle morphisms)**.** If $\mathcal{L}$ and $\mathcal{K}$ are two line bundles over X, then a line bundle morphism from $\mathcal{L}$ to $\mathcal{K}$ is a map f which makes the following diagram commute.

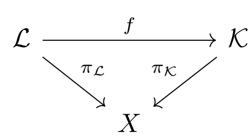

Furthermore, the map f should be linear on the fibre over every point of X. This map is an isomorphism if there's a map g in the opposite direction such that $f \circ g$ and $g \circ f$ are identity on each fibre.

Analogous to dual of a vector space, we can also construct the dual of a line bundle $\mathcal{L}$. Abstractly, it's given by a line bundle map from $\mathcal{L}$ to the trivial line bundle $X \times \mathbb{C}$, which mirrors the construction of the dual of a vector space, which is given by a map from V to the base field k. Concretely, it's constructed by taking the data of the transition maps of the line bundle $\mathcal{L}$, and replacing them with their conjugate. It's an easy exercise to see that these two constructions really give the same line bundle.

The last construction we need on line bundles is the notion of a tensor product. Abstractly, $\mathcal{L} \otimes \mathcal{K}$ is the set of all bilinear bundle maps from $\mathcal{L} \times \mathcal{K}$ to $X \times \mathbb{C}$. Concretely, it's constructed by taking the transition maps of $\mathcal{L}$ and $\mathcal{K}$, and tensoring them.

The collection of line bundles over X modulo isomorphism is denoted by the $\mathrm{LB}(X)$. The operation of taking a dual and tensor product factor onto this set in a well defined manner, and they make $\mathrm{LB}(X)$ into an abelian group, with multiplication given by $\mathcal{L} \times \mathcal{K} := \mathcal{L} \otimes \mathcal{K}$, inverse given by $\mathcal{L}^{-1} := \mathcal{L}^*$, and the identity element is the class of the trivial bundle $X \times \mathbb{C}$.

This is where divisors enter the picture. It turns out that the group $\mathrm{LB}(X)$ is isomorphic to the group $\mathrm{Pic}(X)$. To see this, we must first describe a map from $\mathrm{Pic}(X)$ to $\mathrm{LB}(X)$, and then a map in the other direction.

Suppose $D = p \cdot x$ is a divisor with just one point x in its support. To construct a line bundle from this divisor, pick a small coordinate chart (U, α) around p, and let V be the open set $X \setminus \{x\}$. Let these be the locally trivial neighbourhood. The only thing we need to specify is how does a point in the fibre over some point in $U \cap V$ change as we change from the local trivialization in U to the one in V. That will determined by the integer p. Recall that $U \cap V$ is biholomorphic to the open disc minus a point in $\mathbb{C}$. We need a holomorphic function from this set into $\mathrm{GL}(1, \mathbb{C})$ to define the transition map. We'll define the map to be $z \mapsto z^p$. This gives us a line bundle over X. For divisors of the form $\sum_i p_i \cdot x_i$, we define the line bundle associated to these as the $\bigotimes_i \mathcal{L}_i$, where $\mathcal{L}_i$ is the line bundle associated to the divisor $p_i \cdot x_i$. We only need to check one thing: this descends to a map from $\mathrm{Pic}(X)$ to $\mathrm{LB}(X)$, i.e. doing this construction for any principal divisor should give a line bundle isomorphic to $X \times \mathbb{C}$. This will show it's a well-defined map from $\mathrm{Pic}(X)$ to $\mathrm{LB}(X)$ and that it's injective.

Showing that the map is surjective is slightly trickier. The way one does it is the following: start with any rational section of the line bundle $\mathcal{L}$. The rational section may vanish at a few places, and may have poles somewhere, so it defines a divisor D. It turns out that is the divisor associated to the line bundle, i.e. if one performs the construction in the previous paragraph with the divisor D, one gets a line bundle $\mathcal{L}_D$ which is isomorphic to $\mathcal{L}$. The thing that needs to be checked here is that the divisor's class is independent of the choice of the rational function. We have already seen something like this in the case of the cotangent bundle, where the divisor associated to any meromorphic 1-form (i.e. rational section of the cotangent bundle) descended to the same element in $\mathrm{Pic}(X)$, i.e. the canonical class. We have thus proved the following result, modulo a lot of details.

Theorem 9. *The groups* $\mathrm{Pic}(X)$ *and* $\mathrm{LB}(X)$ *are isomorphic.*

This correspondence lets us a conclude a lot of things about line bundles from the corresponding results about divisors. We first need the following lemma.

Lemma 10. *Let D be a divisor, and $\mathcal{L}$ the line bundle associated to the divisor. Then the space of holomorphic sections of $\mathcal{L}$ is isomorphic to $L(D)$.*

Proof. For clarity's sake, we'll prove this when D just has one point in its support, i.e. $D = p \cdot x$. Recall the locally trivial neighbourhoods of $\mathcal{L}$. One of them was a small open set U around x, and the other was $V = X \setminus \{x\}$. Any section holomorphic on V must be a rational function f on X. To check its actually a holomorphic section, we need to check its also holomorphic on U. Since the transition function is $z \mapsto z^p$, the section looks like $z^p f$ in the locally trivial neighbourhood U. This will be holomorphic if f had a pole of order less than or equal to p at x. But that means f belongs to $L(D)$. Conversely, any function in $L(D)$ will have a pole of order less than or equal to p at x will define a holomorphic section of $\mathcal{L}$ by the argument above. The same proof works for divisors with multiple points in the support. □

One consequence of this lemma and the Riemann-Roch theorem is that for any holomorphic line bundle, the space of its holomorphic sections is finite dimensional.

One can now also talk about the ampleness of line bundles, by talking about the ampleness of their associated divisors. If a line bundle is very ample, a basis consisting of holomorphic sections will give a map of the base space X into a projective space $\mathbb{CP}^n$. Furthermore, line bundles can be enriched with additional geometric structure, like *connections*, which allows us to talk about the *curvature* of a line bundle. The line bundle is said to be positive if its curvature form satisfies certain positivity conditions. The *Kodaira embedding theorem* shows that the notions of ampleness and positivity are equivalent, thus allowing the use of techniques from the differential geometry to prove ampleness, and techniques from algebraic geometry to prove positivity results. Chapter 4 of [2] has more details on the interaction between algebraic and differential geometry arising from the Kodaira embedding theorem.

4. The Abel-Jacobi Map

We have seen that the group $\mathrm{Pic}(X)$ classifies line bundles on X up to isomorphism. Consequently, it might be useful to understand what the group itself looks like. We already have some idea of the structure of the group: recall that after picking a base point, the group $\mathrm{Div}(X)$ factors as $\mathrm{Div}_0(X) \oplus \mathbb{Z}$. It's not too hard to see that this splitting factors onto the group $\mathrm{Pic}(X)$. We end up with a similar splitting after picking a basepoint $x_0 \in X$.

$$\mathrm{Pic}(X) = \mathrm{Pic}_0(X) \oplus \mathbb{Z}$$

Let's focus our attention to $\mathrm{Pic}_0(X)$. To understand this group, we need to be able to determine when a degree 0 divisor is principal. This is where the Abel-Jacobi map comes in. We will construct a complex torus $J(X)$, and a map from $\mathrm{Div}_0(X)$ to $J(X)$ whose kernel is exactly the group of principal divisors. We will also show that the map is surjective, which will prove that $\mathrm{Pic}_0(X)$ is isomorphic as a group to $J(X)$.

4.1. **Construction of the Jacobian.** Let $\Omega^1(X)^*$ be the dual of the space of holomorphic 1-forms $\Omega^1(X)$. We know what some elements of $\Omega^1(X)^*$ look like: they're given by integrating the given holomorphic 1-form along some fixed curve γ. Suppose γ is closed curve and γ' another closed curve which differs from γ by a boundary, i.e. $\gamma = \gamma' + \partial B$, for some singular 2-chain B. We can integrate any holomorphic 1-form ω along γ.

$$\begin{aligned}\int_\gamma \omega &= \int_{\gamma'} \omega + \int_{\partial B} \omega \\ &= \int_{\gamma'} \omega + \int_B d\omega\end{aligned}$$

Since ω is a holomorphic 1-form, it locally looks like $f(z)dz$ for some holomorphic f. Its exterior derivative will be $\partial_z f dz \wedge dz + \partial_{\overline{z}} f dz \wedge d\overline{z} = 0$[2]. What this means is that the second term in the right hand side of the above expression vanishes. We thus get the following identity for homologous chains γ and γ' for any holomorphic 1-form ω.

$$\int_\gamma \omega = \int_{\gamma'} \omega$$

We get a well defined map from $H_1(X;\mathbb{Z})$ to $\Omega^1(X)^*$, given by the following expression.

$$[\gamma] \mapsto \int_\gamma$$

The image of this map forms a subgroup of $\Omega^1(X)^*$, which we'll call *periods*, and denote by Λ. The Jacobian $J(X)$ is the space $\Omega^1(X)^*/\Lambda$.

We can explicitly see what the space $J(X)$ looks like. By Poincaré duality, $H_1(X;\mathbb{Z})$ is isomorphic to $H^1(X;\mathbb{Z})$. We also know that $H^1(X;\mathbb{Z})$ forms a lattice in $H^1(X;\mathbb{R})$. A result in Chapter 8 of [1] says that $H^1(X;\mathbb{R})$ is isomorphic (as an $\mathbb{R}$-vector space) to $\Omega^1(X)$, which, after picking a basis, is isomorphic to $\Omega^1(X)^*$. That means Λ forms a lattice inside $\Omega^1(X)^*$, and $J(X)$ is a complex g-dimensional torus.

4.2. **Construction of the Abel-Jacobi map.** To construct the Abel-Jacobi map, we'll need to pick a basepoint, say p; though we'll show later it's not really required. We first define the Abel-Jacobi map A from X to $J(X)$.

$$A(x) = \left[\int_{\gamma(p,x)}\right]$$

Here, $\gamma(p,x)$ is a path from p to x, and $\left[\int_{\gamma(p,x)}\right]$ denotes the equivalence class of $\int_{\gamma(p,x)}$ in $J(X)$. This is a well defined map, since any path $\gamma'(p,x)$ will differ from $\gamma(p,x)$ by a closed loop, which will go to 0 in $J(X)$. We can extend this map by linearity to $\mathrm{Div}(X)$.

$$A\left(\sum_i c_i \cdot x_i\right) = \sum_i \left(\left[c_i \int_{\gamma(p,x_i)}\right]\right)$$

If we restrict the map A to divisors of degree 0, it turns out to be independent of the choice of base point p.

Proposition 11. *Let A be the Abel-Jacobi map with respect to the basepoint p, and A' be the Abel-Jacobi map with respect to the basepoint p'. If D is a divisor of degree 0, then $A(D) = A'(D)$.*

Proof. Every divisor of degree 0 can be written as a sum of divisors of the form $a - b$, where a and b are points on X. It suffices to show $A(a-b) = A'(a-b)$. Let γ_1 be a path from p to a, and γ_2 be a path from p to b. Let γ_1' and γ_2' be the corresponding paths from p'. If we compute $A(a-b) - A'(a-b)$, we get the following integral.

$$\begin{aligned}A(a-b) - A'(a-b) &= \left[\left(\int_{\gamma_1} - \int_{\gamma_2}\right) - \left(\int_{\gamma_1'} - \int_{\gamma_2'}\right)\right] \\ &= \left[\int_{\gamma_1 - \gamma_1' + \gamma_2' - \gamma_2}\right]\end{aligned}$$

From Figure 2 it's clear that the path $\gamma_1 - \gamma_1' + \gamma_2' - \gamma_2$ is a closed path, which means that integral is 0 in $J(X)$, and $A(a-b) = A'(a-b)$. □

[2] We have shown here that holomorphic 1-forms are closed. A similar proof also works for anti-holomorphic 1-forms.

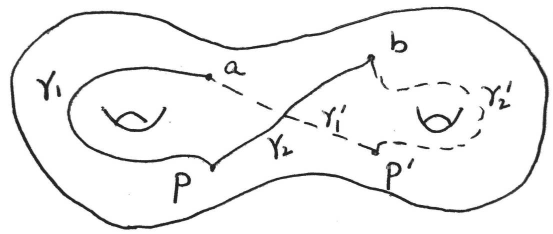

FIGURE 2. The Abel-Jacobi map is independent of base point for the divisor $a - b$

What we have gotten so far is an explicit description of the Abel-Jacobi map from $\mathrm{Div}_0(X)$ to $J(X)$. The next couple of things we need to show about this map are that its kernel is exactly the group of principal divisors $\mathrm{PDiv}(X)$, and that the map is surjective. The former result is called Abel's theorem, and the latter is called the Jacobi Inversion problem. We will prove both these results in the next two sections.

4.3. **Abel's Theorem.** We shall replicate the proof of Abel's theorem presented in Chapter VIII of [6]. We first prove one direction of the result, the direction which states that the image of a principal divisor under the Abel-Jacobi map is 0. To prove this, we'll need the notion of the *trace* of meromorphic 1-form. The operation of taking a trace is in some sense dual to pulling back a 1-form along a map f. Given a non-constant holomorphic map f between compact Riemann surfaces X and Y, the trace of a meromorphic 1-form ω on X is an associated 1-form $\mathrm{Tr}(\omega)$ on Y, which acts like the pushforward of ω along the map f. We will make it clearer in what sense it is a pushforward, but before that, we'll define the trace.

Since f is a non-constant holomorphic map from X to Y, it must be a ramified cover of some degree d, because it's surjective (its image is compact and open by the Open Mapping theorem), and is singular at only finitely many points. At an unramified point q in Y, the map f must be a local isomorphism, i.e. there exists an open set V around q such that $f^{-1}(V)$ is a disjoint union of d open sets $\{U_1, \ldots, U_d\}$ in X, all of which are biholomorphic to V. Let ϕ_i be the biholomorphism from V to U_i. We define $\mathrm{Tr}(\omega)$ at these points in the following manner.

$$\mathrm{Tr}(\omega)(q) = \sum_{p \in f^{-1}(q)} \phi_i^*(\omega)(p)$$

It's clear that in the neighbourhood V, $\mathrm{Tr}(\omega)$ is a meromorphic 1-form, since it's locally a sum of meromorphic 1-forms. The only difficulty lies in extending this 1-form on the points $q \in Y$, where f is ramified, because it won't be a local biholomorphism at such a point, and there will be no map going in the other direction. We can still define the trace though. First, assume that q has just one preimage p, and pick local coordinates z around p and around q such that the map f sends z to z^d. Suppose the form ω locally looks like $h(z)dz$. Then we define its trace by the following formula.

$$\mathrm{Tr}(\omega) = \sum_{i=0}^{d-1} \frac{h\left(\zeta^i z\right)}{d\left(\zeta^i z\right)^{d-1}} dz \tag{3}$$

This formula seems to come out of nowhere, but we have already seen this before. For a non-ramified point p, the value of $\mathrm{Tr}(\omega)$ at p was given by summing up $\phi_i^*(\omega)$ over all its preimages. This is also what is happening here: for any non-zero z, the map is unramified, and the above formula is doing precisely what it should, i.e. summing the pullback over all preimages. The symbol ζ denotes a d^{th} root of 1, i.e. $\exp\left(\frac{2\pi\sqrt{-1}}{d}\right)$. The expression given by (3) defines a meromorphic 1-form in the neighbourhood of 0. If we write out the Laurent series of $h(z)dz$, which we'll denote by $\sum_n c_n z^n dz$ and plug that into the above formula and simplify, we get the Laurent series for $\mathrm{Tr}(\omega)$.

$$\mathrm{Tr}(\omega) = \sum_n c_{nd-1} z^{n-1} dz \tag{4}$$

For an arbitrary point q with preimages $\{q_1, \ldots, q_k\}$, where the ramification degree is d_i, one does the above construction for each preimage, and adds up all the traces. This defines the trace of a meromorphic 1-form. From the formula of the trace, it's clear that the trace of a holomorphic 1-form is holomorphic.

Given a path $\gamma \in Y$, we can pull it back to d paths in X by using the path lifting properties of covers, as long as the path doesn't go through any ramification points. Even if it does, we can just delete those points, lift the path up, and take the closure, which will give us d paths. We denote the pullback of a path γ defined in this manner by

$f^*(\gamma) = \gamma_1 + \cdots + \gamma_d$, where γ_i's are the paths that γ lifts to. We can define the pullback of any singular 1-chain in this manner. This lets us state the following lemma.

Lemma 12 (Integration of trace)**.** *Let $f : X \to Y$ be a non-constant holomorphic function, ω a holomorphic 1-form on X, and γ a chain on Y. We then have the following equation.*

$$\int_{f^*\gamma} \omega = \int_\gamma \mathrm{Tr}(\omega)$$

Proof sketch. Since the ramification points are finite, and have measure 0, they do not affect this integral. We might as well assume that the paths don't go through the ramification points. In that case, the left hand side is just an integration along the d lifts of γ. If one recalls the definition of trace of a 1-form, it's really the sum of the 1-forms at the preimage points, which means the right hand side is also the sum of the integral of ω along the lifted paths $f^*\gamma$. □

We can now prove one of the directions of Abel's theorem.

Proposition 13. *Let (f) be a principal divisor on X. Then the Abel-Jacobi map sends (f) to 0 in $J(X)$.*

Proof. The divisor (f) comes from a holomorphic map $f : X \to \mathbb{CP}^1$. Let $\{z_1, \ldots, z_d\}$ be the preimages of 0 counted with multiplicities, and let $\{p_1, \ldots, p_d\}$ be the preimages of ∞, counted with multiplicities. The divisor (f) can be written in the following manner.

$$(f) = \sum_{i=1}^{d} (z_i - p_i)$$

Pick a base point x_0, and paths α_i and β_i from x_0 to z_i and p_i respectively. Then the action of the Abel-Jacobi map on (f) is given by the following expression.

$$A((f)) = \sum_{i=1}^{d} \left[\int_{\alpha_i} - \int_{\beta_i} \right] \tag{5}$$

Pick a path γ in $\mathbb{CP}^1$ from 0 to ∞ that doesn't pass through any ramification points. This path will pullback to d paths $\{\gamma_1, \ldots, \gamma_d\}$ in X such that γ_i goes from z_i to p_i. Then for each i, $\alpha_i + \gamma_i - \beta_i$ is a closed loop. Using this fact along with (5) gives us the following.

$$A((f)) = -\sum_{i=1}^{d} \left[\int_{\gamma_i} \right]$$

One of the preimages of $A((f))$ in $\Omega^1(X)^*$ is $-\sum_{i=1}^{d} \int_{\gamma_i}$. We will show that this functional is evaluates to 0 on every holomorphic 1-form ω, which will show the image of the functional in $J(X)$ is 0, and consequently, $A((f)) = 0$.

If we evaluate the functional on any holomorphic 1-form ω, we get the following.

$$-\sum_{i=1}^{d} \int_{\gamma_i} \omega = -\int_{f^*\gamma} \omega$$

The right hand side becomes $-\int_\gamma \mathrm{Tr}(\omega)$ by Lemma 12. Furthermore, since ω was holomorphic, so is its trace. But there are no non-zero holomorphic 1-forms on $\mathbb{CP}^1$ since its genus is 0. This proves the result. □

Proving the other part of Abel's theorem is slightly more involved. We'll need a few lemmas before we can prove the result. First, consider the surface X as the quotient of a $4g$-sided polygon $\mathcal{P}$, with the sides given by $\bigcup_{i=1}^{g} \{a_i, b_i, a_i', b_i'\}$, a_i is identified with a_i' with orientation reversed, and b_i is identified with b_i' with orientation reversed. Note that the homology $H_1(X; \mathbb{Z})$ is generated by the loops $\{a_i\}$ and $\{b_i\}$. For any 1-form σ, its integral around the loop a_i is denoted by $A_i(\sigma)$, and its integral around the loop b_i is denoted by $B_i(\sigma)$.

Pick a point x_0 in the interior of the polygon $\mathcal{P}$, and a smooth form σ. We can define the following function on $\mathcal{P}$.

$$f_\sigma(x) = \int_{\gamma(x_0, x)} \sigma$$

Here, $\gamma(a, b)$ is a path from a to b in the polygon $\mathcal{P}$. Since the polygon is simply connected, any two paths are homotopic, and the function is well defined. Furthermore, if σ is a holomorphic 1-form, then f_σ is a holomorphic function. We can now state our first few lemmas.

Lemma 14. *Let σ and τ be closed forms on X. Then we have the following equation.*

$$\int_{\partial\mathcal{P}} f_\sigma \tau = \sum_{i=1}^{g} A_i(\sigma)B_i(\tau) - A_i(\tau)B_i(\sigma)$$

Proof. We can explicitly write down what the given integral will be.

$$\begin{aligned}\int_{\partial\mathcal{P}} f_\sigma \tau &= \int_{\sum_{i=1}^{g} a_i + b_i - a_i' - b_i'} f_\sigma \tau \\ &= \sum_{i=1}^{g} \left(\int_{p\in a_i} (f_\sigma(p) - f_\sigma(p'))\tau + \int_{q \in b_i} (f_\sigma(q) - f_\sigma(q'))\tau \right)\end{aligned}$$

In the above equation, we're using the fact that a_i and a_i' get identified, and b_i and b_i' get identified, and if p and p' are corresponding points on a_i and a_i', the value of τ at both these points will be the same.

We now need to figure out what $f_\sigma(p) - f_\sigma(p')$ is. That is given by the following integral.

$$\begin{aligned}f_\sigma(p) - f_\sigma(p') &= \int_{\gamma(x_0,p)} \sigma - \int_{\gamma(x_0,p')} \sigma \\ &= \int_{\gamma(x_0,p)+\gamma(p,p')-\gamma(x_0,p')} \sigma - \int_{\gamma(p,p')} \sigma\end{aligned}$$

The first term in the above expression is 0, since the associated path is contractible. The second path $\gamma(p,p')$ is homotopic to the path b_i on $\mathcal{P}$, so we end up getting $-B_i(\sigma)$. Similarly, the expression $f_\sigma(q) - f_\sigma(q')$ becomes $A_i(\sigma)$, which gives us the result we wanted. ☐

Lemma 15. *If ω is a non-zero holomorphic 1-form on X, we have the following inequality.*

$$\operatorname{Im} \sum_{i=1}^{g} A_i(\omega)\overline{B_i(\omega)} < 0$$

Proof. We use Lemma 14 with the forms ω and $\overline{\omega}$. We get the following.

$$\int_{\partial\mathcal{P}} f_\omega \overline{\omega} = \sum_{i=1}^{g} A_i(\omega)B_i(\overline{\omega}) - A_i(\overline{\omega})B_i(\omega)$$

It's not too hard to check that $A_i(\overline{\omega}) = \overline{A_i(\omega)}$ and $B_i(\overline{\omega}) = \overline{B_i(\omega)}$. Substituting these relations back into the above equation, we get the following.

$$\int_{\partial\mathcal{P}} f_\omega \overline{\omega} = 2\sqrt{-1} \cdot \operatorname{Im}(A_i(\omega)\overline{B_i(\omega)})$$

We can compute the left hand side term using Stoke's theorem.

$$\begin{aligned}\int_{\partial\mathcal{P}} f_\omega \overline{\omega} &= \int_{\mathcal{P}} d(f_\omega \overline{\omega}) \\ &= \int_{\mathcal{P}} df_\omega \wedge \overline{\omega} + f_\omega d\overline{\omega} \\ &= \int_{\mathcal{P}} \omega \wedge \overline{\omega} + 0\end{aligned}$$

The term df_ω becomes ω simply as a consequence of the fundamental theorem of calculus, and $\overline{\omega}$ is a closed form because anti-holomorphic 1-forms are closed (just like we showed holomorphic 1-forms are closed).

Locally, ω looks like $f(z)dz$, so $\omega \wedge \overline{\omega}$ will look like $|f(z)|^2 dz \wedge d\overline{z}$. Changing coordinates to $dx \wedge dy$, this becomes $-2\sqrt{-1}|f|^2 dx \wedge dy$. The integral of this will be purely imaginary, and will have negative imaginary part. This proves that $\frac{\int_{\mathcal{P}} f_\omega \overline{\omega}}{2i} < 0$, which is what we wanted to show. ☐

Lemma 16. *If ω is a holomorphic 1-form such that $A_i(\omega) = 0$ for all i, then $\omega = 0$. Then same holds if the A_i are replaced by B_i.*

Proof. This follows obviously from Lemma 15. ☐

Since the space of holomorphic 1-forms is g-dimensional, we can encapsulate the data of the values of $A_i(\omega)$ and $B_i(\omega)$ in the form of two $g \times g$ matrices, which we'll denote by $\mathbf{A}$ and $\mathbf{B}$. We start by picking a basis $\{\omega_1, \ldots, \omega_g\}$ of $\Omega^1(X)$. The ij^{th} entry of $\mathbf{A}$ is $A_i(\omega_j)$, and similarly, the ij^{th} entry of $\mathbf{B}$ is $B_i(\omega_j)$. These matrices are called the *period matrices* with respect the given basis of $\Omega^1(X)$ and $H_1(X;\mathbb{Z})$. They have a few useful properties.

Lemma 17. *The matrices* $\mathbf{A}$ *and* $\mathbf{B}$ *are non-singular.*

Proof. Let ω be a holomorphic 1-form in the kernel of $\mathbf{A}$. That means $A_i(\omega) = 0$ for all i. But by Lemma 16, ω must be 0, hence $\mathbf{A}$ must be non-singular. The same proof works for $\mathbf{B}$. □

Lemma 18. *The matrices* $\mathbf{A}$ *and* $\mathbf{B}$ *satisfy the following relation.*

$$\mathbf{A}^\top \mathbf{B} = \mathbf{B}^\top \mathbf{A}$$

Proof. Apply Lemma 14 to the form ω_i and ω_j. The left hand side becomes the following.

$$\begin{aligned}\int_{\partial \mathcal{P}} f_{\omega_i}\omega_j &= \int_{\mathcal{P}} d(f_{\omega_i}\omega_j) \\ &= \int_{\mathcal{P}} \omega_i \wedge \omega_j + f_{\omega_i} d\omega_j \\ &= 0 + 0\end{aligned}$$

The first term vanishes because ω_i locally looks like $f_i(z)dz$ and ω_j looks like $f_j(z)dz$, and their wedge product is 0. The second term vanishes because holomorphic 1-forms are closed. We can rearrange the right hand side of Lemma 14 to get the following equality.

$$\sum_{n=1}^{g} A_n(\omega_i) B_n(\omega_j) = \sum_{n=1}^{g} B_n(\omega_i) A_n(\omega_j)$$

The left hand side is the ij^{th} entry of $\mathbf{A}^\top \mathbf{B}$ and the right hand side is the ij^{th} entry of $\mathbf{B}^\top \mathbf{A}$. This proves the lemma. □

We just need to prove one last lemma before we can prove the other direction of Abel's theorem.

Lemma 19. *Let D be a degree 0 divisor on X such that $A(D)$ equals 0 in $J(X)$. Then there exists a meromorphic 1-form ω satisfying the following conditions.*

(i) ω has simple poles on the support of D and is holomorphic outside the support of D.
(ii) $\operatorname{Res}_p(\omega) = D(p)$ for all points $p \in X$.
(iii) $A_i(\omega)$ and $B_i(\omega)$ are integral multiples of $2\pi\sqrt{-1}$ for all i.

Proof. By Proposition 1.15 of Chapter VII in [6], there exists a meromorphic 1-form τ that satisfies the first two conditions. Then $\tau - \omega$, where ω is a holomorphic 1-form also satisfies the first two conditions. If $\{\omega_1, \dots, \omega_g\}$ is a basis of holomorphic forms, then the problem boils down to finding coefficients $\{c_1, \dots, c_g\}$ such that $\tau - \sum_{i=1}^{g} c_i\omega_i$ satisfies the third condition.

We can assume without loss of generality that the curves a_i and b_i don't pass through the poles of τ, otherwise we could always shift them to avoid the pole. Define the numbers ρ_i in the following manner.

$$\begin{aligned}\rho_i &= \frac{1}{2\pi\sqrt{-1}} \int_{\partial\mathcal{P}} f_{\omega_i}\tau \\ &= \sum_{p \in X} \operatorname{Res}_p(f_{\omega_i}\tau)\end{aligned}$$

The second equality follows from the residue theorem. Since f_{ω_i} is holomorphic, and the residue of τ at p is $D(p)$, we get the following.

$$\begin{aligned}\rho_i &= \sum_{p \in X} f_{\omega_i}(p) \cdot D(p) \\ &= \sum_{p \in X} D(p) \int_{x_0}^{p} \omega_k\end{aligned}$$

Recall that x_0 was the basepoint we picked in X. Notice that this is precisely the functional $A(D)$ acting upon ω_i. But recall that $A(D)$ was 0 in $J(X)$. That means the functional $\sum_{p\in X} D(p) \int_{x_0}^{p}$ is of the form $\int_\gamma$, for some closed loop γ. We can write γ in terms of the generators of $H_1(X;\mathbb{Z})$ as the following.

$$[\gamma] = \sum_{i=1}^{g} m_i[a_i] - n_i[b_i]$$

Applying this functional to ω_k, we see that it evaluates to the following.

$$\rho_k = \sum_{i=1}^{g} A_i(\omega_k) - B_i(\omega_k) \tag{6}$$

But Lemma 14 gives us that ρ_k also equals the following expression.

$$\rho_k = \frac{1}{2\pi\sqrt{-1}} \sum_{i=1}^{g} B_i(\tau)A_i(\omega_k) - A_i(\tau)B_i(\omega_k) \tag{7}$$

If we equate (6) and (7), we get the following expression.

$$\sum_{i=1} (B_i(\tau) - 2m_i\pi\sqrt{-1})A_i(\omega_k) = \sum_{i=1}^{g} (A_i(\tau) - 2n_i\pi\sqrt{-1})B_i(\omega_k)$$

The above equation is really an explicit way of writing the following equation.

$$\mathbf{A}^\top \mathbf{b} = \mathbf{B}^\top \mathbf{a}$$

Here, $\mathbf{b}$ is the vector whose k^{th} coordinate is $B_k(\tau) - 2m_i\pi\sqrt{-1}$, and $\mathbf{a}$ is the corresponding vector for A_k.

Consider the following maps.

$$\begin{aligned}
&\alpha : \mathbb{C}^g \to \mathbb{C}^{2g} \\
&\alpha : v \mapsto \begin{pmatrix} \mathbf{A}v \\ \mathbf{B}v \end{pmatrix} \\
&\beta : \mathbb{C}^{2g} \to \mathbb{C}^g \\
&\beta : \begin{pmatrix} v \\ w \end{pmatrix} \mapsto \mathbf{B}^\top v - \mathbf{A}^\top w
\end{aligned}$$

By Lemma 18, $\beta \circ \alpha = 0$, and because $\mathbf{A}$ and $\mathbf{B}$ are full rank, we must have $\ker(\beta) = \text{Image}(\alpha)$. Notice that the vector $\begin{pmatrix} \mathbf{a} \\ \mathbf{b} \end{pmatrix}$ is in $\ker(\beta)$, which means there's a $\mathbf{c} \in \mathbb{C}^g$ such that $\mathbf{a} = \mathbf{A}\mathbf{c}$ and $\mathbf{b} = \mathbf{B}\mathbf{c}$. This vector $\mathbf{c}$ will tell us how to modify τ to get the form we want. Let ω be the following 1-form.

$$\omega = \tau - \sum_{i=1}^{g} c_i\omega_i$$

If we compute $A_i(\omega)$, and $B_i(\omega)$, we get the following.

$$\begin{aligned}
A_i(\omega) &= 2n_i\pi\sqrt{-1} \\
B_i(\omega) &= 2m_i\pi\sqrt{-1}
\end{aligned}$$

This proves the lemma. □

We can now prove Abel's theorem.

Theorem 20 (Abel's Theorem). *A degree* 0 *divisor* D *is principal iff* $A(D)$ *goes to* 0 *in* $J(X)$.

Proof. If D is principal, $A(D) = 0$ by Proposition 13. Now suppose $A(D) = 0$. By Lemma 19, we can find a meromorphic 1-form that satisfies the three conditions in Lemma 19. Define a function f in the following manner.

$$f(x) = \exp\left(\int_{x_0}^{x} \omega\right)$$

Since the integral of ω over any closed loop is an integral multiple of $2\pi\sqrt{-1}$, the above function is well defined. Furthermore, the function f is holomorphic wherever ω is. In particular, that means it has only finitely many singularities. We need to show that the singularities are all poles, and $D = (f)$.

Pick a coordinate neighbourhood around a pole p of ω. Since ω has a simple pole of order $D(p)$ there, locally it looks like the following.

$$\omega = \left(\frac{D(p)}{z} + g(z)\right) dz$$

That means the integral $\int_{x_0}^z$ locally looks like the following.

$$\int_{x_0}^{z} \omega = D(p)\ln(z) + h(z)$$

And if we take the exponential, we get the following.

$$f(z) = z^{D(p)} e^{h(z)}$$

This is clearly meromorphic, and the order of f at p is $D(p)$. This shows $(f) = D$ and proves the result. ☐

Abel's theorem implies that the Abel-Jacobi map factors through the group of principal divisors, and induces the following injection.

$$A : \mathrm{Pic}_0(X) \hookrightarrow J(X)$$

4.4. **Jacobi Inversion.** Our next goal is to show the map from $\mathrm{Pic}_0(X)$ to $J(X)$ is also surjective. That's equivalent to showing that the Abel-Jacobi map A into $J(X)$ is surjective. To show that the map A is surjective, it will suffice to show that its image in $J(X)$ contains an open subset of $J(X)$. Since $J(X)$ is a connected topological group, the result will follow from the following easy theorem (which we won't prove).

Theorem 21. *Let G be a connected topological group, and H a subgroup that contains an open subset of G. Then $H = G$.*

Pick a base point x_0 in X. Consider the following map I from X^g to $J(X)$, where g is the genus of X.

$$I((x_1,\ldots,x_g)) = A\left(\sum_{i=1}^{g} x_i - x_0\right)$$

To show that the image of A contains an open set, it suffices to show that the image of I contains an open set. Because both X^g and $J(X)$ g-dimensional complex varieties, one approach would be to show that there is some point $(x_1,\ldots,x_g) \in X^g$ where the derivative of I is non-singular. The inverse function theorem will then tell us that the image contains an open set.

Proposition 22. *There exists a point $(x_1,\ldots,x_g) \in X^g$ such that the derivative of I at that point is an invertible $g \times g$ matrix.*

Our proof for Proposition 22 is based on the proof of Theorem 5.2 in [5].

Proof. Pick points $\{x_1,\ldots,x_g\} \in X$, and local coordinates $\{z_1,\ldots,z_g\}$ around them. We'll specify later in the proof what these points need to be. This gives a coordinate chart around the point $(x_1,\ldots,x_g) \in X^g$ given by $(z_1,\ldots,z_g)$. Pick paths $\{\gamma_1,\ldots,\gamma_g\}$ going from x_0 to x_i respectively. We also need pick a basis of $\Omega^1(X)$, say $\{\omega_1,\ldots,\omega_g\}$, which gives us coordinates on $J(X)$ from the coefficients of the dual basis $\{\omega_1^*,\ldots,\omega_g^*\}$. In terms of these coefficients, the map I looks like the following.

$$I((x_1,\ldots,x_g)) = \begin{pmatrix} \sum_{i=1}^{g}\int_{\gamma_i}\omega_1 \\ \sum_{i=1}^{g}\int_{\gamma_i}\omega_2 \\ \vdots \\ \sum_{i=1}^{g}\int_{\gamma_i}\omega_g \end{pmatrix}$$

We can now compute the derivative matrix of this map. If we denote the j^{th} row of the output of I by I_j, then $\frac{\partial I_j}{\partial x_k}$ is given by the following expression.

$$\frac{\partial I_j}{\partial x_k} = \frac{\omega_j}{dz_k}$$

The above expression needs an explanation. The expression for I_j is given by $\sum_{i=1}^{g}\int_{\gamma_i}\omega_j$, and if we differentiate this with respect to the x_k coordinate, all the terms except the k^{th} term will vanish, since they don't depend on x_k. The derivative of the k^{th} term with respect to the coordinate x_k can be computed by writing ω_j locally in the z_k coordinates as $\omega_j = h_{jk}dz_k$. In that case, the expression for the derivative becomes the following.

$$\frac{\partial I_j}{\partial x_k} = \frac{\partial}{\partial x_k}\int_{x_0}^{x_k} h_{jk}dz_k$$

The fundamental theorem of calculus tells us that the above expression equals h_{jk}. We can denote it succinctly by $\frac{\omega_j}{d_{zk}}$. Thus the derivative matrix at $(x_1, \ldots, x_g)$ looks like the following.

$$DI((x_1, \ldots, x_g)) = \begin{pmatrix} \frac{\omega_1(x_1)}{dz_1} & \cdots & \frac{\omega_1(x_g)}{dz_g} \\ \vdots & \ddots & \vdots \\ \frac{\omega_g(x_1)}{dz_1} & \cdots & \frac{\omega_g(x_g)}{dz_g} \end{pmatrix}$$

We need to show that this matrix is non-singular for some point $(x_1, \ldots, x_g) \in X^g$.

Pick a point x_1 where the form ω_1 does not vanish. And then modify the basis of $\Omega^1(X)$ in the following manner: replace ω_i for $i \geq 2$ with the ω_i', where ω_i' is given by the following expression.

$$\omega_i' = \omega_i - \frac{\omega_i(x_1)}{\omega_1(x_1)}\omega_1$$

This is well defined, since $\omega_1(x)$ is not 0. Furthermore, we can locally divide forms, since they're sections of a line bundle. With respect to the new basis of $\Omega^1(X)$, the expression for DI at any point with the first coordinate x_1 becomes the following.

$$DI((x_1, \ldots, x_g)) = \begin{pmatrix} \frac{\omega_1(x_1)}{dz_1} & \frac{\omega_1(x_1)}{dz_2} & \cdots & \frac{\omega_1(x_g)}{dz_g} \\ 0 & \frac{\omega_2'(x_2)}{dz_2} & \cdots & \frac{\omega_2'(x_g)}{dz_g} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \frac{\omega_g'(x_2)}{dz_2} & \cdots & \frac{\omega_g'(x_g)}{dz_g} \end{pmatrix}$$

We can repeat this process again, by finding a point $x_2 \in X$ where ω_2' does not vanish, and modifying the basis elements ω_i' for $i \geq 3$. If we do this g times, we find g points $\{x_1, \ldots, x_g\}$ in X such that the derivative of I at $(x_1, \ldots, x_g)$ is given by an upper triangular matrix with non-zero coefficients on the diagonal. This shows that at that point, the derivative matrix is invertible. The proves the result. □

We have shown that the Abel-Jacobi map is surjective, and the induced map from $\mathrm{Pic}_0(X)$ to $J(X)$ is an isomorphism. We have thus managed to answer the question we asked ourselves at the beginning of the section, i.e. what does the group $\mathrm{Pic}(X)$ look like. The answer is that it looks like $J(X) \oplus \mathbb{Z}$, i.e. the product of a torus and $\mathbb{Z}$.

Our construction of the Jacobian, as well as the results we proved about it, i.e. the Jacobian is isomorphic to $\mathrm{Pic}_0(X)$, relied heavily on tools available only over $\mathbb{C}$ like integration of differential forms and singular homology. Over an arbitrary field, the Jacobian of a curve X is an abelian variety (a variety which is also an algebraic group) $J(X)$ such that $X \hookrightarrow J(X)$, and the extending the map linearly to divisors gives an isomorphism from $\mathrm{Pic}_0(X)$ and $J(X)$, along with a couple of other conditions. This works over any algebraically closed field, but it comes at a cost of being harder to describe, or do computations with. Chapter A.8 in [3] has more details on the properties and constructions of the Jacobian over arbitrary fields.

References

[1] Simon Donaldson. *Riemann surfaces*. Oxford University Press, 2011.
[2] Phillip Griffiths and Joseph Harris. *Principles of algebraic geometry*. John Wiley & Sons, 2014.
[3] Marc Hindry and Joseph H Silverman. *Diophantine geometry: an introduction*, volume 201. Springer Science & Business Media, 2013.
[4] Rene Schipperus (https://math.stackexchange.com/users/149912/rene schipperus). What's the intuition behind divisors? Mathematics Stack Exchange. URL:`https://math.stackexchange.com/q/2771427` (version: 2018-05-07).
[5] SJ Kleinerman. The jacobian, the abel-jacobi map, and abel's theorem.
[6] Rick Miranda. *Algebraic curves and Riemann surfaces*, volume 5. American Mathematical Soc., 1995.
[7] Brian Osserman. Divisors on nonsingular curves. URL: `https://www.math.ucdavis.edu/~osserman/classes/248A-F09/divisors.pdf`. Last visited on 2018/11/25.