Lecture 14: Homolomorphic one-forms and vector fields

June 24, 2015

Let X be a connected Riemann surface with basepoint x. In the previous lecture, we introduced the universal cover of X: this is a Riemann surface $\widetilde{X}$ mapping down to X, such that X identifies with the quotient of $\widetilde{X}$ by a certain action of the fundamental group $\pi_X(x, x)$. In this lecture we will address the question of what structure on X would be needed to identify $\widetilde{X}$ with the complex numbers $\mathbb{C}$, and hence to uniformize X by $\mathbb{C}$. The answer will be a *nowhere-vanishing holomorphic one-form*, satisfying a certain technical condition which is irrelevant when X is compact (our case of interest).

In the end, we will get the following theorem:

Theorem 0.1. *Let X be a compact connected Riemann surface, and suppose that there exists a nowhere vanishing holomorphic one-form on X. Then:*

1. *X is isomorphic to $\mathbb{C}/\Gamma$ for some full lattice $\Gamma \subset \mathbb{C}$;*
2. *The fundamental group $\pi_X(x, x)$ at any basepoint $x \in X$ is isomorphic to $\mathbb{Z} \oplus \mathbb{Z}$.*

First we have to say what a holomoprhic one-form is. Essentially, it is something which holomorphically measures first-order change on a Riemann surface. For example, a holomorphic function f defines a holomorphic one-form df, its *differential*, which measures the first-order change as seen by f. But globally, there can be more one-forms than these, essentially because to glue df to dg we don't need to have $f = g$ exactly: they can differ by an overall constant.

From another perspective, these df's are things which make rigorous Leibniz's notation

$$\frac{df}{dz} = f'(z).$$

Thus, the derivative f' is the quotient of the first-order change of f by the first order change of z. Using one forms we can also make more sense of the integral notation

$$\int_\gamma f(z)dz,$$

since now the expression $f(z)dz$ itself has invariant meaning — it is a one-form.

But we're getting ahead of ourselves. First, recall the notion of tangent space: to every Riemann surface X and point $p \in X$, we can attach a one-dimensional complex vector space

$$T_p(X)$$

of "tangent vectors to X at p". We never gave a precise definition of this space, but here are two possibilities:

Definition 0.2. *The space $T_p(X)$ is defined as either:*

1. *The set of equivalence classes of smooth paths $\gamma : (-\epsilon, \epsilon) \to X$ such that $\gamma(0) = p$, modulo the equivalence relation which identifies two paths which agree to first order at 0. (More precisely, this means that the derivatives of the paths, calculated in a chart for X around x, should agree at 0).*

2. *The set of associations $f \mapsto \partial(f) \in \mathbb{C}$, where f is a holomorphic function defined in a neighborhood of p normalized so that $f(p) = 0$, satisfying the following conditions:*

   - *$\partial(f + g) = \partial(f) + \partial(g)$ and $\partial(c \cdot f) = c \cdot \partial(f)$ for $c \in \mathbb{C}$, i.e. ∂ is a linear map;*
   - *$\partial(f) = 0$ if f vanishes to order ≥ 2 at p.*

The intuitive content of the first definition is hopefully clear: a tangent vector is the first-order information of a path. As for the second definition, it thinks of a tangent vector as a direction along which one can differentiate functions; thus the tangent vector is encoded as the differentiation operator ∂, whose properties are axiomatized.

To see that the definitions agree and give a one-dimensional complex vector space, one notes that both are invariant under isomorphisms and under shrinking X to a neighborhood of p, so it suffices to do just the calculation of $T_p(\mathbb{C})$. There one sees that every element of $T_p(\mathbb{C})$ is of the form

$$\lambda \cdot \frac{\partial}{\partial z}$$

for some unique $\lambda \in \mathbb{C}$. Here $\frac{\partial}{\partial z}$ is notation for the standard unit vector 1 based at p in $\mathbb{C}$. Thus, $\lambda \cdot \frac{\partial}{\partial z} \in T_p(\mathbb{C})$ is the vector from p to $p + \lambda$. More explicitly, in the first definition, $\lambda \cdot \frac{\partial}{\partial z}$ is represented by the path $\gamma(t) = p + \lambda \cdot t$, and in the second definition, $\lambda \cdot \frac{\partial}{\partial z}$ is represented by the differentiation operator

$$\partial(f) = \lambda \cdot f'(p).$$

(To connect the two representations, note that we can also write $\partial(f) = \frac{d}{dt}(f \circ \gamma) \mid_{t=0}$.)

Thus, if we choose a local isomorphism of (X, p) with $(\mathbb{C}, 0)$, then a tangent vector is just represented by a complex number, thought of as a vector in $\mathbb{C}$. If you change the local isomorphism by a local biholomorphism φ of $(\mathbb{C}, 0)$, then the complex number changes according to the rule $\lambda \mapsto \lambda \cdot \varphi'(0)$.

From either of the above definitions it should be clear that a holomorphic map $f : X \to Y$ gives a linear map

$$df_x : T_x(X) \to T_y(Y)$$

for any $x \in X$ with image $y = f(x)$ in Y. In local coordinates, this becomes a linear map from $\mathbb{C} \to \mathbb{C}$, i.e.a one-by-one matrix; the entry of this matrix is just the usual derivative of f at x. Also, the chain rule $d(f \circ g) = df \circ dg$ holds.

Now on to business.

Definition 0.3. *Let X be a Riemann surface. A* vector field *on X is an assignment of, to every point $x \in X$, an element $v_x \in T_x(X)$. A* one form *on X is an assignment of, to every point $x \in X$, a $\mathbb{C}$-linear map $\omega_x : T_x(X) \to \mathbb{C}$.*

A vector field is called holomorphic *if, in local charts, it can be written as $v_z = f(z) \cdot \frac{\partial}{\partial z}$, where f is a holomorphic function. A one form is called* holomorphic *if, in local charts, $\omega_x(\frac{\partial}{\partial z})$ is a holomorphic function of x.*

The set of holomorphic vector fields on X is denoted $T(X)$; the set of holomorphic one forms is denoted $\Omega(X)$.

Thus, a vector field attaches a first-order change to every point of X; and a one-form assigns numbers to every possible first-order change at every possible point of X. It may seem from this description that specifying a one-form involves specifying more information than a vector field, but this is not so: because $T_x(X)$ is one-dimensional as a complex vector space, to specify ω_x one needs only specify the value of ω_x at any nonzero vector of $T_x(X)$. It follows that, locally, both vector fields and one forms amount to just functions f. Namely, on an open subset U of $\mathbb{C}$, every vector field is of the form

$$v_x = f(x) \cdot \frac{\partial}{\partial z}$$

for a unique function f, and every holomorphic one-form is of the form

$$\omega_x(c \cdot \frac{\partial}{\partial z}) = c \cdot f(x)$$

for a unique function f. Furthermore, the vector field (resp. the one-form) is holomorphic if and only if the function f is. Note, however, that since the standard unit vector field $\frac{\partial}{\partial z}$ is not fixed by most coordinate changes, this identification of vector fields and one-forms with functions does not generally hold globally.

We will focus more on one-forms than vector fields, because they are easier to manipulate algebraically when it comes to elliptic curves. But we gave the definition of vector field to illustrate the duality between the two concepts, and since a vector field is more geometrically accessible (it's just a bunch of vectors combing your space). A manifestation of this duality is the difference in functoriality: for a holomorphic map $f : X \to Y$, we get natural maps

$$f_* : T(X) \to T(Y)$$

and

$$f^* : \Omega(Y) \to \Omega(X)$$

in different directions.

The most basic examples of holomorphic one forms are df, where f is a holomorphic function. This is defined as follows: since $f : X \to \mathbb{C}$ is holomorphic, there is an induced map

$$df_x : T_x(X) \to T_{f(x)}\mathbb{C} = \mathbb{C}$$

for all $x \in X$, thus a one-form. It is holomorphic, since in local charts we have $df_x(\frac{\partial}{\partial z}) = f'(x)$, and f' is holomorphic when f is.

The very simplest case of this occurs when $X = \mathbb{C}$ and f is the identity map. This map is sometimes also just denoted "z", so the one-form in this case is denoted dz. It is uniquely characterized by the fact that, at every point,

$$dz(\frac{\partial}{\partial z}) = 1.$$

In particular, dz is an example of a *nowhere vanishing one-form*: at every point, the linear map dz is not the zero map.

Starting from this dz we can produce other examples. Exercise: if $f : X \to \mathbb{C}$ is holomorphic, then the pullback f^*dz equals df. Exercise: the zeroes of df are exactly the points where f has multiplicity

≥ 2. Exercise: dz doesn't change under pullback by translations of $\mathbb{C}$; thus it *descends* to a well-defined holomorphic one-form on $\mathbb{C}/\Gamma$ for any group Γ acting properly freely on $\mathbb{C}$, meaning there is a unique $\omega \in \Omega(\mathbb{C}/\Gamma)$ such that $p^*\omega = dz$, where $p : \mathbb{C} \to \mathbb{C}/\Gamma$ is the natural projection. So non-vanishing holomoprhic one-forms exist on any Riemann surface uniformized by $\mathbb{C}$; we're aiming for the converse.

Finally, one last exercise: both holomorphic vector fields and holomorphic one-forms can be multiplied by arbitrary holomorphic functions. E.g., if $\omega \in \Omega(X)$ and f is holomorphic, define

$$(f \cdot \omega)_x(\xi) = f(x) \cdot \omega(\xi).$$

Furthermore, the product rule $d(f \cdot g) = f \cdot dg + g \cdot df$ holds.

Now, note that, by definition, holomorphic one-forms are local objects: to specify $\omega \in \Omega(X)$, it suffices to specify ω on some open cover of X, provided there is agreement on overlaps. Thus, to understand the possibilities for arbitrary holomorphic one-forms, it suffices to understand holomorphic one forms on the open unit disk $\mathbb{D}$. As remarked after the definition of one-form, these can all be put in the form $f(z) \cdot dz$ for a unique holomorphic function f. In the following lemma we give another, different description of all holomorphic one-forms on $\mathbb{D}$:

Lemma 0.4. *Every holomorphic one form ω on $\mathbb{D}$ can be written as*

$$\omega = dg$$

for some holomorphic function g, which is unique up to adding constants.

Proof. Write $\omega = f(z) \cdot dz$. Then $\omega = dg$ if and only if $g'(z) = f(z)$, as follows straight from the definitions. Thus we need to show that every holomorphic function on $\mathbb{D}$ has an antiderivative, unique up to constants. This is standard complex function theory; it can be done either by integration or by finding the correct power series expansion for the antiderivative. □

Now we can turn to the main point. By definition, one-forms assign numbers to tangent vectors, which are first-order germs of paths. We want to show that they actually assign numbers to paths, by integrating the assignments at all the tangent vectors along the path. Instead of explicitly defining the integral, however, we will cheat by using the previous lemma in the construction.

Proposition 0.5. *Let X be a Riemann surface and ω a holomorphic one-form on X. There is a unique way to assign a number $\int_\gamma \omega \in \mathbb{C}$ to every path $\gamma : [0,1] \to X$, in such a way that the following properties are satisfied:*

1. *If $\omega = df$ in an open neighborhood of the path, then $\int_\gamma \omega = f(\gamma(1)) - f(\gamma(0))$ (first-order change of f integrates to global change of f).*
2. *The value $\int_\gamma \omega$ only depends on the homotopy class $[\gamma]$ of γ, so it can be written $\int_{[\gamma]} \omega$.*
3. *The integral is additive under composition of paths:*

$$\int_{[\gamma_1]\cdot[\gamma_2]} \omega = \int_{[\gamma_1]} \omega + \int_{[\gamma_2]} \omega.$$

Proof. First we prove the proposition when X is isomorphic to the unit disk $\mathbb{D}$. There, by the lemma, we can write $\omega = df$, where f is unique up to adding a constant. Then we define

$$\int_\gamma \omega = f(\gamma(1)) - f(\gamma(0)).$$

This is independent of f, because the constants will cancel. The verification of the three properties is obvious, as is the uniqueness.

It follows that, returning to case of general X, we get a well-defined integral $\int_\gamma \omega$ whenever γ lies entirely in some open subset isomorphic to a disk. It doesn't matter which open subset we choose, because the two potentially different f's as above will have to differ by an overall constant all along the path, by analytic continuation.

Now we treat a general path $\gamma : [0,1] \to X$. We can break γ up into finitely many pieces, each of which lives in some open subset of X isomorphic to a disk. Then we can define $\int_\gamma \omega$ by additivity: it is the sum of the integrals over the various pieces. Since any two ways of breaking up an interval admit a common refinement, and the additivity property 3 is valid locally, this definition is independent of the way we break up the interval. Thus it is well-defined.

To prove the first property, note that if $\omega = df$, then we can use this same f locally all along the path in the above definition, which gives the result by a telescoping sum.

To prove the second property, note that, just as an path can be broken up into pieces each of which lives in a disk, so can a homotopy. But in a disk we know the homotopy invariance, so we deduce it in general.

The third property follows immediately from the definition: if we have a way of breaking up two paths, we have a way of breaking up their composite which makes the claim obvious. ☐

Exercise: let f be a nonzero meromorphic function on a Riemann surface X, and $x \in X$. Then f is nonzero and holomorphic in a small punctured neighborhood of x, so that df/f is a well-defined holomorphic one-form in a small punctured neighborhood of x, and

$$\int_\gamma \frac{df}{f} = 2\pi i \cdot v_x(f),$$

where γ is a small loop running once counterclockwise around x, and $v_x(f)$ is the valuation of f at x.

Corollary 0.6. *Let X be a Riemann surface and $\omega \in \Omega(X)$. Let x be a point of X, with corresponding universal cover $p : \widetilde{X} \to X$ and its basepoint $\widetilde{x}$, given by the identity path at x. There is a unique holomorphic function $f : \widetilde{X} \to \mathbb{C}$ such that $f(\widetilde{x}) = 0$ and $p^*\omega = df$.*

This function f is defined by

$$f(\gamma) = \int_\gamma \omega.$$

Proof. Exericse. ☐

This corollary is very useful, since it shows how to produce explicit functions on the universal cover, which itself was defined rather abstractly. For example, if $X = \mathbb{C} \setminus \{0\}$, $x = 1$, and $\omega = \frac{dz}{z}$, then we get a well-defined holomorphic function $f : \widetilde{X} \to \mathbb{C}$, which truly deserves to be called log. In fact, $\widetilde{X}$ is the Riemann surface associated to any choice of logarithm function defined locally on X, in the sense we discussed at the beginning of this course. It is where log naturally lives. The fact that the fibers of

$p : \widetilde{X} \to X$ are labelled by the integers corresonds to the fact that log is only well-defined up to adding elements of $2\pi i \cdot \mathbb{Z}$.

So we have found a way to map out of the universal cover. But if we want to identify the universal cover, we also need to find a way to map in to it, which means we need to produce paths instead of function on paths. This is the dual problem, so we should look at the dual to the notion of one form, which is vector field.

This turns out to be a very natural idea. If v is a smooth vector field on a Riemann surface X, then a *solution* to v is a path

$$\gamma : (a, b) \to X$$

for some open interval $(a, b) \subset \mathbb{R}$, such that $\gamma'(t) = v(\gamma(t))$ for all $t \in (a, b)$. Thus, a solution is a path which "follows the vectors" in the vector field.

It should be intuitively clear that there exists a solution passing through any point. Namely, just start out at that point, then drive your car on the surface by taking your velocity at a given point to be the vector sticking out at that point. The mathematical justification for this comes from the theory of ordinary differential equations, which gives:

Theorem 0.7. *Let X be a Riemann surface, v a smooth vector field on X, and $x \in X$ a point. Then there is an interval (a, b) containing 0 and a solution*

$$\gamma : (a, b) \to X$$

to v such that $\gamma(0) = x$. This solution is unique in the sense that any two such agree on their common domain of definition.

If we take the union over all possible (a, b) as in the statement, we get a maximal interval of definition for the solution γ with $\gamma(0) = x$. When this maximal interval of definition equals all of $\mathbb{R}$, we say the solution *exists for all time*. Not every solution exists for all time: if the vectors grow large enough, they can shuffle a solution out to ∞ in finite time. But this problem doesn't exist when X is compact, and so:

Theorem 0.8. *Let X be a compact Riemann surface, v a smooth vector field on X, and $x \in X$. Then there is a unique smooth map*

$$\gamma : \mathbb{R} \to X$$

such that $\gamma(0) = x$ and $\gamma'(t) = v(\gamma(t))$ for all $t \in \mathbb{R}$.

When the vector field is *holomorphic*, such a solution can be extended to "complex time" as well, giving a holomorphic map $\gamma : \mathbb{C} \to X$. This can be proved by following the proof in the real case, and noting that it works just fine with $\mathbb{C}$ replacing $\mathbb{R}$. Alternately, we can reduce the complex case to the real one by the following trick: for $\lambda \in \mathbb{C}$, consider the scaled vector field $\lambda \cdot v$, and write its real-time solution $\gamma_\lambda : \mathbb{R} \to X$. Then the value of our desired solution γ at complex time λ will be $\gamma_\lambda(1)$ (Exercise).

Corollary 0.9. *Let X be a compact Riemann surface, v a holomorphic vector field on X, and $x \in X$ with associated universal cover $p : \widetilde{X} \to X$. Then there is a unique holomorphic map*

$$g : \mathbb{C} \to \widetilde{X}$$

such that $g(0) = \widetilde{x}$ and $(p \circ g)_* \frac{d}{dt} = v$.

Proof. If $\gamma : \mathbb{C} \to X$ is the complex-time solution to v, we can define $g(\lambda)$ to be the path $[0,1] \to X$ defined by

$$t \mapsto \gamma(t \cdot \lambda).$$

We leave the verification of the required properties as an exercise. $\square$

Let us give an example of such a map g. Take $X = \mathbb{C} \setminus \{0\}$, and the vector field $v \in T(X)$ defined by

$$v_z = z \cdot \frac{\partial}{\partial z}.$$

Although X is not compact, this actually won't matter, since (as we'll soon see explicitly) all solutions to this vector field exist for all time anyway. This is due to a very careful balance between the vector field and the non-compactness of X: for example it would fail if we used either $\frac{\partial}{\partial z}$ or $z^2 \cdot \frac{\partial}{\partial z}$ (why?).

Explicitly, a solution γ to v is a function $\gamma : \mathbb{C} \to \mathbb{C} \setminus \{0\}$ such that

$$\gamma'(t) = \gamma(t).$$

We all know how to produce solutions to this differential equation: the unique solution with $\gamma(0) = 1$ is given by

$$\gamma(t) = exp(t).$$

Exercise: sketch the vector field v and convince yourself graphically that this is the solution, based on properties of the complex exponential function.

Thus, in this case, the map $g : \mathbb{C} \to \widetilde{X}$ is given by

$$g(\lambda) = \text{the path } t \mapsto exp(\lambda \cdot t) \text{ for } t \in [0,1].$$

Note that the composition $p \circ g : \mathbb{C} \to X \setminus \{0\}$ is the exponential map $\lambda \mapsto exp(\lambda)$.

Alright, we're almost done. We have both a way to map in to $\widetilde{X}$, and a way to map out of it. Putting things together, we get the following:

Proposition 0.10. *Let X be a compact Riemann surface and $x \in X$ a point, with corresponding universal cover $(\widetilde{X}, \widetilde{x})$. To every nowhere-vanishing holomorphic one-form ω on X there corresponds a unique isomorphism $f : \widetilde{X} \simeq \mathbb{C}$ sending $\widetilde{x}$ to 0 and satisfying $f^*(dz) = p^*\omega$.*

This proposition implies the theorem stated at the beginning of this lecture, for the following reason: the proposition implies that such an X is uniformized by $\mathbb{C}$. But we studied all Riemann surfaces uniformized by $\mathbb{C}$, and the only compact ones were of the form described in part 1 of that theorem. As for part 2, it follows because in the previous lecture we saw that the uniformizing group Γ identifies with the fundamental group of X. On the other hand, a full lattice in $\mathbb{C}$ is isomorphic to $\mathbb{Z} \oplus \mathbb{Z}$ by choosing generators.

Now let's prove the proposition.

Proof. By one of our results above, integrating the one-form ω gives a unique map

$$f : \widetilde{X} \to \mathbb{C}.$$

sending $\widetilde{x}$ to 0 and satisfying $f^*dz = p^*\omega$. We need to see that f is an isomorphism. However, we can use ω to produce a *dual* holomorphic vector field v, defined by

$$v_x = \text{ the unique tangent vector satisfying } \omega_x(v_x) = 1.$$

Such a v_x exists and is unique because ω_x is a nonzero map from a one-dimensional vector space to the complex numbers, hence is an isomorphism. In local coordinates, if $\omega = f(z) \cdot dz$, then $v = \frac{1}{f(z)} \cdot \frac{\partial}{\partial z}$, which proves holomorphicity.

Thus we also get a map

$$g : \mathbb{C} \to \widetilde{X},$$

by solving the differential equation corresponding to the vector field v as described above. By construction, the composition $f \circ g : \mathbb{C} \to \mathbb{C}$ sends 0 to 0 and has derivative equal to 1 everywhere. Thus it is the identity map $z \mapsto z$. We conclude that the map f has a *section*, meaning there is a map g backwards with $f \circ g = id$. Then to finish, we need only apply the following lemma. □

Lemma 0.11. *Let $f : X \to Y$ be a map of connected Riemann surfaces. If f has a section, then f is an isomorphism.*

Proof. Let g be a section. Then g is certainly injective, since $f \circ g = id$. If we prove surjectivity, then g, and hence f, will be an isomorphism, as desired.

Since g is injective, it is non-constant; thus from the local structure theorem for holomorphic maps it follows that $g(Y)$ is an open subset of X. On the other hand, I claim $g(Y)$ is also closed. Indeed, suppose $g(y_i)$ is a sequence of points in $g(Y)$ tending to a point $x \in X$. Applying f, we find that the y_i tend to some point $y = f(x)$. Applying g again, we find that the $g(y_i)$ tend to $g(y)$. But they also tend to x. Thus, since X is Hausdorff, we must have $x = g(y)$, so $x \in g(Y)$, proving closedness.

Hence $g(Y)$ is both open and closed. But X is connected and Y is nonempty, so this implies $g(Y) = X$, meaning g is surjective, as desired. □

We will investigate this theorem more explicitly in the next lecture. But for now let us briefly illustrate it in the simpler analogous case of $X = \mathbb{C} \setminus \{0\}$ with the non-vanishing holomorphic one-form

$$\omega = \frac{dz}{z}$$

and the basepoint $1 \in X$.

Here the function $f : \widetilde{X} \to \mathbb{C}$ is the logarithm; we are claiming it is an isomorphism. The dual vector field is

$$v = z \cdot \frac{\partial}{\partial z},$$

which gives the above-discussed map $g : \mathbb{C} \to \widetilde{X}$, essentially the exponential map. We indeed have $f \circ g = id$, and hence both are isomorphisms. The conclusion is that $\mathbb{C} \setminus \{0\} \simeq \mathbb{C}/2\pi i \cdot \mathbb{Z}$ via the exponential map, and that the fundamental group $\pi_X(1, 1)$ is $\mathbb{Z}$.