Lecture 11: Riemann surfaces uniformized by $\mathbb{P}^1$ and $\mathbb{C}$

October 14, 2014

Today we'll start to explore some consequences of the uniformization theorem.

Recall that the uniformization theorem says that every connected Riemann surface is isomorphic to some $\mathbb{M}/\Gamma$, where $\mathbb{M}$ is either $\mathbb{P}^1, \mathbb{C}$, or $\mathbb{D}$, and Γ is a subgroup of $Aut(\mathbb{M})$ acting properly freelyon $\mathbb{M}$.

Recall also that "properly freely" means the following:

1. For every point $z_0 \in \mathbb{M}$, there is an open neighborhood U of z_0 such that the open sets $\{\gamma(U)\}_{\gamma\in\Gamma}$ are disjoint from one another.

Note that a special case is that no $\gamma \neq id$ in Γ can have a fixed point. However, in our analysis of the automorphisms of these three Riemann surfaces $\mathbb{M}$, we've seen (by case-by-case analysis) that every γ which has no fixed point must actually be an oriented isometry of $\mathbb{M}$ in its canonical geometry. So, in the statement of the uniformization theorem, we could just as well say that Γ needs to be a subgroup of $Isom^+(\mathbb{M})$ acting freely and properly on $\mathbb{M}$. Then we can also rephrase the above condition in metrical terms:

1. For every point $z_0 \in \mathbb{M}$, the minimal distance $\inf_{\gamma\neq id} d(z_0, \gamma(z_0))$ is attained by some $\gamma \neq id$.

Now, our goal this week is to understand all the possible Γ's and $\mathbb{M}/\Gamma$'s in the two cases $\mathbb{M} = \mathbb{P}^1$ and $\mathbb{M} = \mathbb{C}$. We'll see that the list of possibilities is fairly small, so that actually, most Riemann surfaces must arise from the case $\mathbb{M} = \mathbb{D}$.

We start with $\mathbb{M} = \mathbb{P}^1$. We've already seen that every automorphism of $\mathbb{P}^1$ must have a fixed point. Thus, because of the freeness condition, the only possibility is $\Gamma = \{id\}$, which gives $\mathbb{M}/\Gamma = \mathbb{P}^1$. So the only Riemann surface uniformized by $\mathbb{P}^1$ is $\mathbb{P}^1$ itself. We've already made a detailed study of $\mathbb{P}^1$, so we can consider this "case closed".

So we move on to $\mathbb{M} = \mathbb{C}$. Thus, let Γ be a group of oriented isometries acting properly and freely on $\mathbb{C}$ (in the Euclidean metric). Because every non-translation isometry of $\mathbb{C}$ has a fixed point, Γ must consist entirely of translations. Then we will see that there are three general possibilities for such a Γ, acting properly freely on $\mathbb{C}$:

1. $\Gamma = \{id\}$;

2. Γ consists of translations by integer multiples of some nonzero $v \in \mathbb{C}$;

3. Γ consists of translations by $\mathbb{Z}$-linear combinations of two $\mathbb{R}$-linearly independent $v_1, v_2 \in \mathbb{C}$.

To prove this, we identify $\mathbb{C}$ with $\mathbb{R}^2$, and we identify Γ with an additive subgroup of $\mathbb{R}^2$ by identifying a translation map with the vector which it is a translation by. Then the freeness condition translates into $\Gamma \subset \mathbb{R}^2$ being *discrete*, i.e. every point of Γ is isolated, meaning has an open neighborhood which contains no other points of Γ.

In fact a nice remark is that it's equivalent to say that just the point $0 \in \Gamma$ should be isolated. Indeed, if U isolates 0, then $\gamma + U$ isolates γ.

Anyway, in this context we have:

Proposition 0.1. *Let Γ be a discrete subgroup of $\mathbb{R}^n$. Then there is an $\mathbb{R}$-linearly independent subset $S \subset \mathbb{R}^n$ such that Γ is equal to the $\mathbb{Z}$-span of S.*

In the case $n = 2$, this proposition implies our desired classification. Indeed, a linearly independent subset of a two dimensional vector space must have either 0, 1, or 2 elements, and these correspond to the three possibilities listed above. So let us prove the proposition.

Proof. We induct on n, the case $n = 0$ being trivial.

Now take an n, and suppose the claim known for $n-1$. If $\Gamma = \{0\}$, then the claim is obvious. Otherwise, we can choose a nonzero $\gamma \in \Gamma$ of minimal distance to the origin. (Proof: first choose an arbitrary nonzero element of Γ, and say it has norm R. Then there are only finitely many points of Γ in $\overline{D}(0;R)$, since this $\overline{D}(0;R)$ is compact and Γ is discrete. Thus we can minimize the distance to the origin among these finitely many points, and hence among all points, since all the others have distance $> R$.)

Now consider the orthogonal projection $\Gamma \to (\mathbb{R}\cdot\gamma)^\perp$ to the orthogonal complement of γ in $\mathbb{R}^n$ (with the standard inner product). Certainly γ maps to zero in this orthogonal projection, so there is an induced map of quotients

$$\Gamma/(\mathbb{Z}\cdot\gamma) \to (\mathbb{R}\cdot\gamma)^\perp.$$

I claim this map is injective, and that it exhibits $\Gamma/(\mathbb{Z}\cdot\gamma)$ as a discrete subgroup of $(\mathbb{R}\cdot\gamma)^\perp$.

We will get both claims at once if we show that there is a neighborhood U of 0 in $(\mathbb{R}\cdot\gamma)^\perp$ which meets the projection of no point of $\Gamma/(\mathbb{Z}\cdot\gamma)$ besides 0. In other words, we need an $\epsilon > 0$ such that the tube of radius ϵ around $\mathbb{R}\cdot\gamma$ contains no points of Γ besides the integer multiples of γ.

But because γ was assumed minimal, the ball of radius $R = |\gamma|$ around 0 contains no points of Γ besides 0. Thus for all integers k, the ball of radius R around $k\cdot\gamma$ contains no points of Γ besides $k\cdot\gamma$. Thus the union of all these balls contains no points of Γ besides those in $\mathbb{Z}\cdot\gamma$. However, a simple matter of Euclidean geometry shows that this union of balls contains a tube of radius $\epsilon = \frac{\sqrt{3}}{2}\cdot R$ around $\mathbb{R}\cdot\gamma$. That proves the claim.

Thus $\Gamma/(\mathbb{Z}\cdot\gamma) \to (\mathbb{R}\cdot\gamma)^\perp$ gives a discrete subgroup of the $n-1$ dimensional $(\mathbb{R}\cdot\gamma)^\perp$. By the inductive hypothesis, then, there is a set $\overline{S}$ of $\mathbb{R}$-linearly independent vectors of $\mathbb{R}\cdot\gamma)^\perp$ such that $\Gamma/(\mathbb{Z}\cdot\gamma)$ consists of the $\mathbb{Z}$-span of the elements of $\overline{S}$.

Now, define $S \subset \Gamma$ consist of an arbitrarily chosen lift of every element of $\overline{S}$, together with the single element γ. Then all our desired conditions are satisfied, as an easy check shows. This finishes the proof. $\square$

Now we can analyze the possibilities for the quotient Riemann surface $\mathbb{C}/\Gamma$. In the first case, $\Gamma = \{id\}$, so we get $\mathbb{C}$ back again. That's not so interesting.

Now consider the second case, $\mathbb{C}/\mathbb{Z}\cdot w$ with $w \neq 0$. I claim that the answer is independent of w, in the sense that the resulting Riemann surfaces are all isomorphic. Indeed, suppose given another $w' \neq 0$. Let $\lambda = w'/w$, and consider the map

$$\mathbb{C} \xrightarrow{\cdot\lambda} \mathbb{C}.$$

This is obviously an isomorphism of Riemann surfaces. But moreover it converts translation by w into translation by w'. Thus it induces an isomorphism

$$\mathbb{C}/(\mathbb{Z}\cdot w) \xrightarrow{\cdot\lambda} \mathbb{C}/(\mathbb{Z}\cdot w'),$$

whence the claim.

On the other hand, we know already that when $w = 2\pi i$, the exponential map identifies this quotient with $\mathbb{C}\setminus\{0\}$. Thus the only Riemann surface we get from $\mathbb{C}/\Gamma$ in the second case is $\mathbb{C}\setminus\{0\}$.

So far it's been pretty boring: after all this work, we've only recovered three Riemann surfaces: $\mathbb{P}^1$, $\mathbb{C}$, and $\mathbb{C}\setminus\{0\}$. But it's about to get more interesting, when we turn to the third case, where Γ is generated by two linearly independent vectors $v_0, v_1 \in \mathbb{C}$. It turns out that in this case:

1. There are many different possible quotients, even up to isomorphism;
2. The quotients are all compact Riemann surfaces;
3. We can identify the quotients (after removing a point) with the Riemann surface associated to a cubic equation of the form $y^2 = f(x)$ where $f(x)$ is a complex cubic polynomial with distinct roots.

First let us see what $\mathbb{C}/\Gamma$ "looks like" in this case. Let

$$\overline{\mathcal{P}} = \{\lambda_1 \cdot v_1 + \lambda_2 \cdot v_2 \mid -\frac{1}{2} \leq \lambda_1, \lambda_2 \leq \frac{1}{2}\},$$

a *fundamental parallelogram* cenetered at 0.

Then the restriction of the projection $\mathbb{C} \to \mathbb{C}/\Gamma$ to $\overline{\mathcal{P}}$ is surjective:

$$\overline{\mathcal{P}} \twoheadrightarrow \mathbb{C}/\Gamma.$$

Moreover it's almost injective: if $z, w \in \overline{\mathcal{P}}$ have the same image in the quotient, then either $z = w$ or z and w lie on opposite sides of the boundary of the parallelogram, and differ by translation by v_1 or v_2.

Thus we can picture $\mathbb{C}/\Gamma$ as being obtained by taking this parallelogram $\overline{\mathcal{P}}$ and gluing opposite sides together by translation. This is, topologically speaking, a *torus*.

A corollary is that $\mathbb{C}/\Gamma$ is compact. Indeed, $\overline{\mathcal{P}}$ is closed and bounded, hence compact; and the continuous image of a compact space is compact.

Given this compactness, and recalling our discussion of the algebraization theorem, it's reasonable to ask what the field of meromorphic functions on $\mathbb{C}/\Gamma$ is. The answer is surprisingly simple:

Theorem 0.2. *There exist two meromorphic functions $\wp, \wp' \in \mathcal{M}(\mathbb{C}/\Gamma)$ with the following properties:*

1. *Every element of $\mathcal{M}(\mathbb{C}/\Gamma)$ can be written uniquely in the form*

$$A(\wp) + B(\wp) \cdot \wp',$$

   where A and B are rational functions (quotients of polynomials).

2. *There is an equation of the form*

$$(\wp')^2 = 4 \cdot (\wp - e_1)(\wp - e_2)(\wp - e_3)$$

   for some distinct $e_1, e_2, e_3 \in \mathbb{C}$.

Note that this theorem tells you exactly what the field $\mathcal{M}(\mathbb{C}/\Gamma)$ is: it's in bijection with pairs (A, B) of rational functions in $\mathbb{C}(z) \times \mathbb{C}(z)$; the addition is component-wise; and the multiplication is

$$(A, B) \cdot (A', B') = (AA' + BB' \cdot 4 \cdot (z - e_1)(z - e_2)(z - e_3), AB' + A'B).$$

In particular the structure of the field $\mathcal{M}(\mathbb{C}/\Gamma)$ is determined by these three constants e_1, e_2, and e_3, which in turn depend on the lattice Γ (we'll see exactly how). By the way, a related theorem, which we'll also prove in the next lecture, is the following:

Theorem 0.3. *The functions $\wp$ and $\wp'$ induce an isomorphism of Riemann surfaces*

$$(\mathbb{C}/\Gamma) \setminus \{0\} \xrightarrow{\sim} \{y^2 = 4(x - e_1)(x - e_2)(x - e_3)\} \subset \mathbb{C} \times \mathbb{C}.$$

One way of reading this is that these functions $\wp$ and $\wp'$ parametrize all solutions to this particular cubic equation. Actually it turns out that, perhaps after some coordinate change, all equations of the form $y^2 = f(x)$, for any cubic polynomial f with distinct roots, arise in this way. This is a very useful fact with some nice classical consequences, since such equations $y^2 = f(x)$ tend to arise in various places. For example they arise when calculating the arc length of an ellipse. So from the above theorem one learns that the arc length of an ellipse can be expressed in terms of these functions $\wp$ and $\wp'$ (which we'll write down explicitly). A perhaps more interesting example is that they arise from the differential equations describe the motion of a simple pendulum. Thus the trajectory of a pendulum is also controlled by these same functions.

Without further ado, let's get to work on the first theorem, describing $\mathcal{M}(\mathbb{C}/\Gamma)$. First, some general nonsense. Let Γ be any group acting freely and properly on a Riemann surface X, and consider the quotient X/Γ. Then giving a holomorphic map from X/Γ to another Riemann surface Y is the same as giving a Γ-invariant holomorphic map from X to Y, i.e. a holomorphic map $f : X \to Y$ with $f \circ \gamma = f$ for all $\gamma \in \Gamma$.

Indeed, on the level of set theory this is clear, since the quotient $p : X \to X/\Gamma$ has the property that $p(x) = p(y)$ if and only if x and y are related by the action of some γ. Thus we need only see that a map $X/\Gamma \to Y$ is holomorphic if and only if its composition with p is holomorphic. That follows because p is (by definition) a local isomoprhism.

Thus, taking $Y = \mathbb{P}^1$, we find that our exotic-looking field $\mathcal{M}(\mathbb{C}/\Gamma)$ can be identified with something one could study in classical complex analysis: it is the field of meromorphic functions on $\mathbb{C}$ which are invariant under translation by Γ.

Using this description, we can produce our functions $\wp$ and $\wp'$. Actually, as the notation suggests, $\wp'$ will be the derivative of $\wp$ (which is called the *Weierstrass p-function*, by the way). Here is the proposition:

Proposition 0.4. *There exists a unique Γ-invariant meromophic function $\wp$ on $\mathbb{C}$ with the following two properties:*

1. *The set of poles of $\wp$ is exactly $\Gamma \subset \mathbb{C}$;*
2. *The Laurent series expansion of $\wp$ takes the form*

$$\wp(z) = \frac{1}{z^2} + (\textit{terms of degree} \geq 1).$$

Note that if we consider $\wp$ as a meromorphic function on $\mathbb{C}/\Gamma$, then properties 1 and 2 in particular say that $\wp$ has a single pole at 0, of multiplicity 2.

In the proof we will construct an explicit such function $\wp$; for example we will see how to calculate its entire Laurent series expansion at 0. However, for the proof of the theorems stated above we will only need to invoke the above two properties of $\wp$. So in some sense one can forget the formula for $\wp$.

Proof. First we show uniqueness. Suppose $\wp$ and p are two such functions, and consider

$$\wp - p.$$

This is a meromorphic function on $\mathbb{C}/\Gamma$ with no poles outside 0, since both $\wp$ and p are. But it also has no pole at 0, since the $\frac{1}{z^2}$'s in the Laurent series expansions cancel. Thus $\wp - p$ is a holomorphic map $\mathbb{C}/\Gamma \to \mathbb{P}^1$ between compact connected Riemann surfaces, such that its fiber at ∞ is empty. Thus, by invariance of degree, if $\wp - p$ is nonconstant then the fiber above every point must be empty. But this is absurd, so $\wp - p$ must be constant. Then since $\wp - p$ has constant term 0 at $z = 0$, this constant must be 0. So $\wp = p$, as claimed.

Now we show existence. The idea is the following. A first naive guess would be to take the basic function $\frac{1}{z^2}$ and then translate it by all the elements of Γ, and sum up, that is:

$$\wp(z) = \sum_{\gamma \in \Gamma} \frac{1}{(z-\gamma)^2}.$$

Unfortunately, this sum does not converge, so that won't be our definition. However, we have better luck with what would be the derivative $\wp'(z)$, and we can take this as a definition: let

$$\wp'(z) := \sum_{\gamma \in \Gamma} \frac{-2}{(z-\gamma)^3}.$$

One can see that this converges uniformly on closed disks, and hence limits to a meromorphic function on $\mathbb{C}$, by comparison with the integral

$$\int_{v \in \mathbb{R}^2} \frac{1}{|v|^3} |dv|,$$

which can be evaluated by switching to polar coordinates.

Noe that the poles of $\wp'$ are exactly the points of Γ, and that its Laurent series at 0 is of the form

$$\wp'(z) = \frac{-2}{z^3} + (\text{terms of degree } \geq 1),$$

where the terms of degree ≥ 1 occur only in odd degrees. (The coefficients can be calculated by successively differentiating and plugging in $z = 0$; the even terms must vanish because $\wp'$ is an *odd* function: $\wp'(-z) = -\wp'(z)$.)

In the interior of the fundamental parallelogram $\mathcal{P} \subset \overline{\mathcal{P}}$, the only pole of $\wp'$ is $z = 0$. Thus there is a unique antiderivative $\wp$ of $\wp'$ in this region whose constant term at 0 is 0. This antiderivative being determined by integrating the Laurent series expansion term-by-term, and so it has the proper expansion at $z = 0$ to satisfy condition 2.

To finish the proof, we need to see that $\wp$ extends to a Γ-invariant function on all of $\mathbb{C}$.

It suffices to see that $\wp(z + v_1) = \wp(z)$ and $\wp(z + v_2) = \wp(z)$ when all of the points in question lie in $\mathcal{P}$, because granting this we can inductively extend $\wp$ to all of $\mathbb{C}$ by using translation to move to neighboring parallelograms.

But, consider $\wp(z + v_1) - \wp(z)$. This function is constant $= C$, since its derivative vanishes and the region where it's defined is connected. On the other hand $\wp$ is even, since $\wp'$ is odd. Thus

$$C = \wp(z + v_1) - \wp(z) = \wp(-z - v_1) - \wp(-z) = \wp(w) - \wp(w + v_1) = -C,$$

where $w = -z - v_1$. Hence $C = 0$, as claimed. The same argument for v_2 gives the claim. $\square$

We don't have time today to prove the main theorems, so instead let me indicate the overall strategy. We view our function $\wp$ as a holomorphic map

$$\wp : \mathbb{C}/\Gamma \to \mathbb{P}^1.$$

Then $\wp$ has degree two, since it has only one pole of order 2. Thus, generically $\wp$ is two-to-one. In fact we can be more precise: $\wp(z) = \wp(w)$ if and only if $z = \pm w$. There are four "branch points" of $\wp$ (points of multiplicity > 1) of multiplicity 2, corresponding to the four points on $\mathbb{C}/\Gamma$ which satisfy $z = -z$. One of these is $z = 0$, which maps to $\infty \in \mathbb{P}^1$; the other three will map to the points we call $e_1, e_2,$ and e_3 in $\mathbb{C} \subset \mathbb{P}^1$.

This fact that $\mathbb{C}/\Gamma$ is degree two over $\mathbb{P}^1$ is the key to the first claim above, that $\mathcal{M}(\mathbb{C}/\Gamma)$ is built out of two copies of $\mathcal{M}(\mathbb{P}^1) = \mathbb{C}(z)$. The idea behind the second claim is that this map $\wp$ makes $\mathbb{C}/\Gamma$ look a lot like the Riemann surface of $y^2 = p(x)$ with $p(x)$ a polynomial of degree 3 with roots e_1, e_2, e_3. Indeed, that Riemann surface is also a degree two cover of $\mathbb{C}$, via projection to the x-axis (given x, there are two choices for y, corresponding to the two square roots of $p(x)$). And the branch points are the same too, since the square root is branched exactly when $p(x) = 0$.

In total, we can say that once we have this one meromorphic function $\wp$, we can use it to understand completely the structure of meromorphic functions on $\mathbb{C}/\Gamma$. This is also the prototype for algebraization: we first show that there is a non-constant meromorphic function on any compact Riemann surface X, then we use the geometry of the resulting map $X \to \mathbb{P}^1$ to completely understand $\mathcal{M}(X)$ in terms of $\mathbb{C}(z)$.