Randomness in Groups

Motivating question: What does a random 2×2 matrix look like?

Step 0 is simply figuring out what the question means. How do you pick a random matrix? And what do we mean by what it "looks like"? Let's consider the latter question first

Let's begin by restricting ourselves to $SL(2,\mathbb{Z})$. Being discrete, this makes things slightly nicer. Given a matrix $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in SL(2,\mathbb{Z})$, we can think of it as a Möbius transformation of the upper half plane $z \mapsto \frac{az+b}{cz+d}$. We can easily imagine such maps giving a clear meaning to thinking about what it looks like. There are 3 distinct cases: 1) A has 2 distinct real eigenvalues, 2) A has a repeated (real) eigenvalue and 3) A has a complex conjugate pair of eigenvalues.

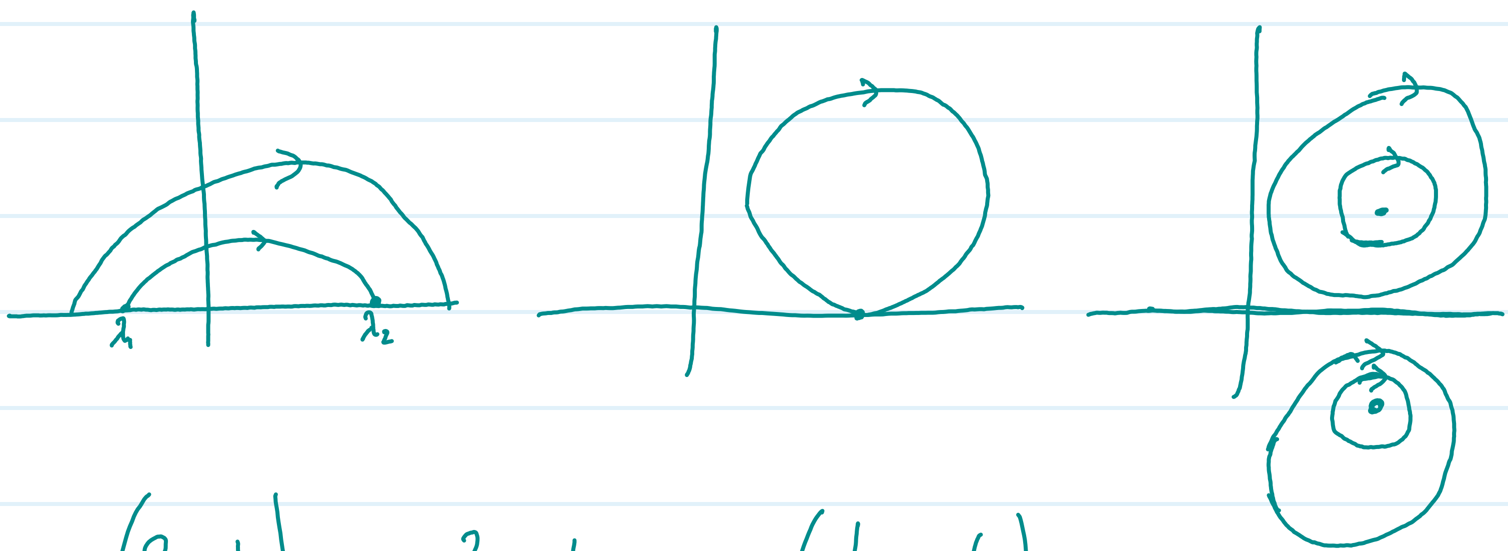

$$\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \qquad \frac{2z+1}{z+1} \qquad \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

$$(\lambda-1)^2 = 0$$

$$z+1$$

Introduction to Teichmuller Theory

Let X be a Riemann surface (i.e. a 2-dim manifold with a complex structure).

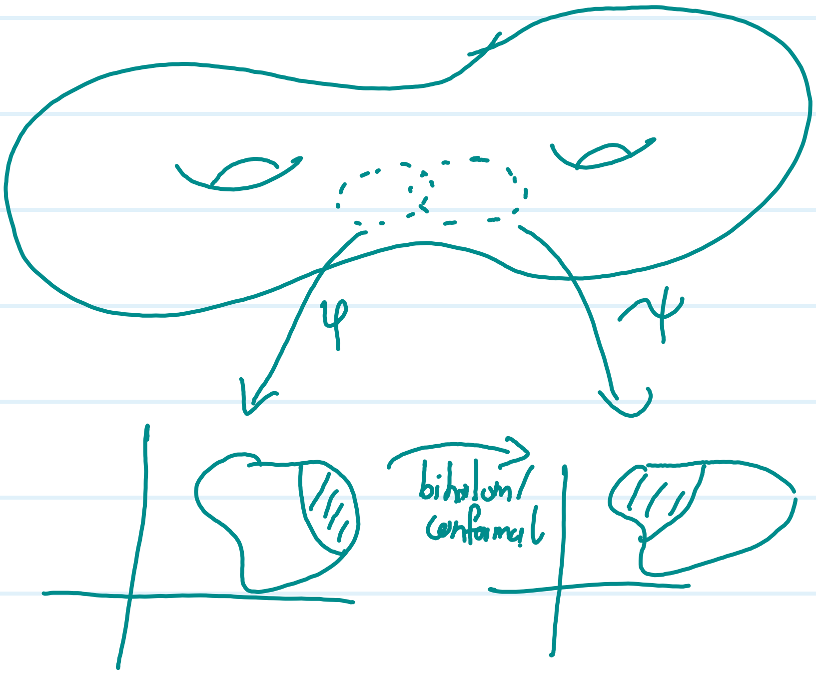

In fact since X is a surface, we know there is a (topological) surface S which X is homeomorphic to. A homeomorphism $f: X \to S$ induces an isomorphism $f_*: \pi_1(X) \to \pi_1(S)$. In fact given an isomorphism $\pi_1(X) \to \pi_1(S)$ we can reconstruct the homeomorphism f (up to isotopy) through something called the Alexander trick. ← look into later

A marked Riemann surface is a Riemann surface X along with a marking map (i.e. homeomorphism) $f: S \to X$ where S is simply a topological surface of genus g. Then Teichmuller space is basically the space of all marked Riemann surfaces, i.e.

top. surface of genus g ↓   Riem surface ↙

$$\mathcal{T}(S_g) = \{f: S_g \to X \text{ homeo}\}/\sim \qquad f_1 \sim f_2 \text{ if } \exists \varphi \quad \begin{tikzcd} & X_1 \arrow[d, "\varphi"] \\ S_g \arrow[ur, "f_1"] \arrow[r, "f_2"'] & X_2 \end{tikzcd} \quad \text{s.t. } \varphi \circ f_1 \text{ isotopic to } f_2$$

We would like to put a topology on $\mathcal{T}(S)$. There are many diff ways of doing this. One way to do this is through representation theory

Recall that the uniformisation theorem that the only simply connected Riemann surfaces are the upper half plane (equiv disk), the complex plane and the Riemann sphere. Given a marked Riemann surface $S \to X$, we can lift the homeomorphism to a homeomorphism of their universal covers and for any surface besides the sphere and torus, the universal cover is going to be $\mathbb{H}^2$.

$$\begin{tikzcd} \tilde{S} \arrow[r] \arrow[d] & \tilde{X} \arrow[r] \arrow[d] & \mathbb{H}^2 \arrow[d, dashed] \\ S \arrow[r] & X & \mathbb{H}^2/\Gamma \arrow[l, dashed] \end{tikzcd}$$

We know that $\pi_1(S)$ acts on the universal cover $\tilde{S}$ and hence by the above identification $\pi_1(S)$ acts on $\mathbb{H}^2$. In fact it must act conformally (since $\pi_1(X)$ acts conformally on $\tilde{X}$) and the conformal maps on $\mathbb{H}^2$ are exactly the isometries of $\mathbb{H}^2$ Long story short then, we can think of a marked Riemann surface as a group homomorphism $\pi_1(S) \to \Gamma \subset \mathrm{Isom}(\mathbb{H}^2)$ In fact provided that Γ act discretely and freely on $\mathbb{H}^2$, $\mathbb{H}^2/\Gamma$ is exactly X Hence we can think of Teichmuller space as

$$\mathcal{T}(S) = \left\{ \rho : \pi_1(S) \to \mathrm{Isom}(\mathbb{H}^2) \;\middle|\; \begin{array}{l} \rho \text{ is 1-1 group homo} \\ \text{and } \rho(\pi_1(S)) \text{ acts} \\ \text{freely and discretely} \end{array} \right\} \Big/ \text{conjugation}$$

b/c you can 'move' all of $\mathbb{H}^2$ around by a conformal map and get same guy

We can give a topology to this more easily. Suppose $\pi_1(S)$ is generated by $a_1, \dots, a_n$. Then a representation ρ_1 is near ρ_2 if

$$\begin{matrix} \rho_1(a_1) & & \rho_2(a_1) \\ \rho_1(a_2) & \text{are all near} & \vdots \\ \vdots & & \\ \rho_1(a_n) & & \rho_2(a_n) \end{matrix}$$

these are just (conjugacy classes of) 2x2 real matrices!

There is a slightly different way of measuring the 'distance' between 2 conformal structures. Suppose X and Y are 2 Riemann surfaces. If there is a conformal map between them then they are the same. If there is a nearly conformal map between them then they are nearly the same. A nearly conformal map is called a quasi conformal map.

Recall that conformal maps are maps which have a non-zero complex derivative everywhere. This means that conformal maps take (small) circles to circles

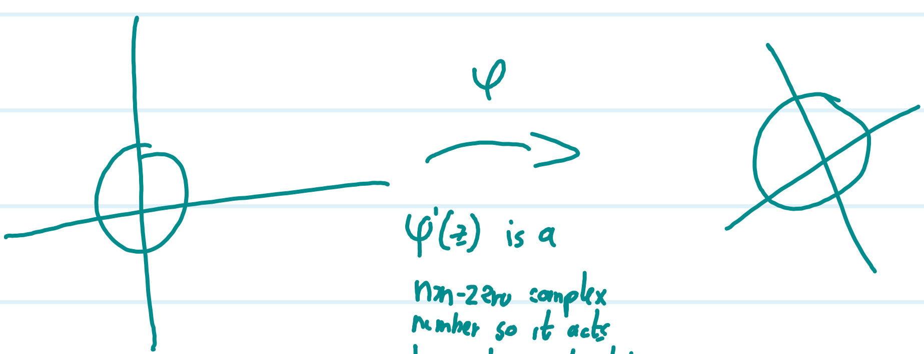

$\varphi'(z)$ is a non-zero complex number so it acts by scaling and rotation

A quasiconformal map, then, can be understood as one that takes circles to ellipses. We can make this statement more precise.

we secretly think of this as $\mathbb{C}$

We know that if $f: \mathbb{R}^2 \to \mathbb{R}^2$ is differentiable, then $f'(z)$ is a linear map (from $T_z\mathbb{R}^2 \cong \mathbb{R}^2$ to $T_{f(z)}\mathbb{R}^2 \cong \mathbb{R}^2$)

Any linear map (from $\mathbb{R}^2$ to $\mathbb{R}^2$) can be written as the sum of a holomorphic map (also known as homothety in this case) and an antiholomorphic map. In other words for every $A: \mathbb{R}^2 \to \mathbb{R}^2$ linear, there exist unique $\alpha, \beta \in \mathbb{C}$ such that

obstruction to being conformal

$$Az = \alpha z + \beta \bar{z}$$

This β then is exactly the obstruction to being conformal. Since we don't really care about scaling we often consider the quantity β/α as the degree of obstruction. In the case that we have a differentiable function f this quantity is $\frac{\partial f/\partial \bar{z}}{\partial f/\partial z}$ (Recall by the Cauchy-Riemann eqns that f is conformal/holomorphic exactly when $\partial f/\partial \bar{z} = 0$).

The form $d\bar{z}/dz$ is called the Beltrami differential. In this way we get a 1-1 correspondence between small deformations of conformal structure (roughly speaking a small open nbhd in $\mathcal{T}(S)$ around some X) and Beltrami differentials. In other words, the tangent space of $\mathcal{T}(S)$ at some point X can be identified with the space of Beltrami differential and the cotangent space (i.e. the dual to the space of Beltrami differentials) is the space of quadratic differentials.

Quadratic differentials

Given a Riemann surface X, a quadratic differential q on it is a form which locally (i.e. when viewed on a chart) looks $q(z)dz^2$ where $q(z)$ is a holomorphic function.

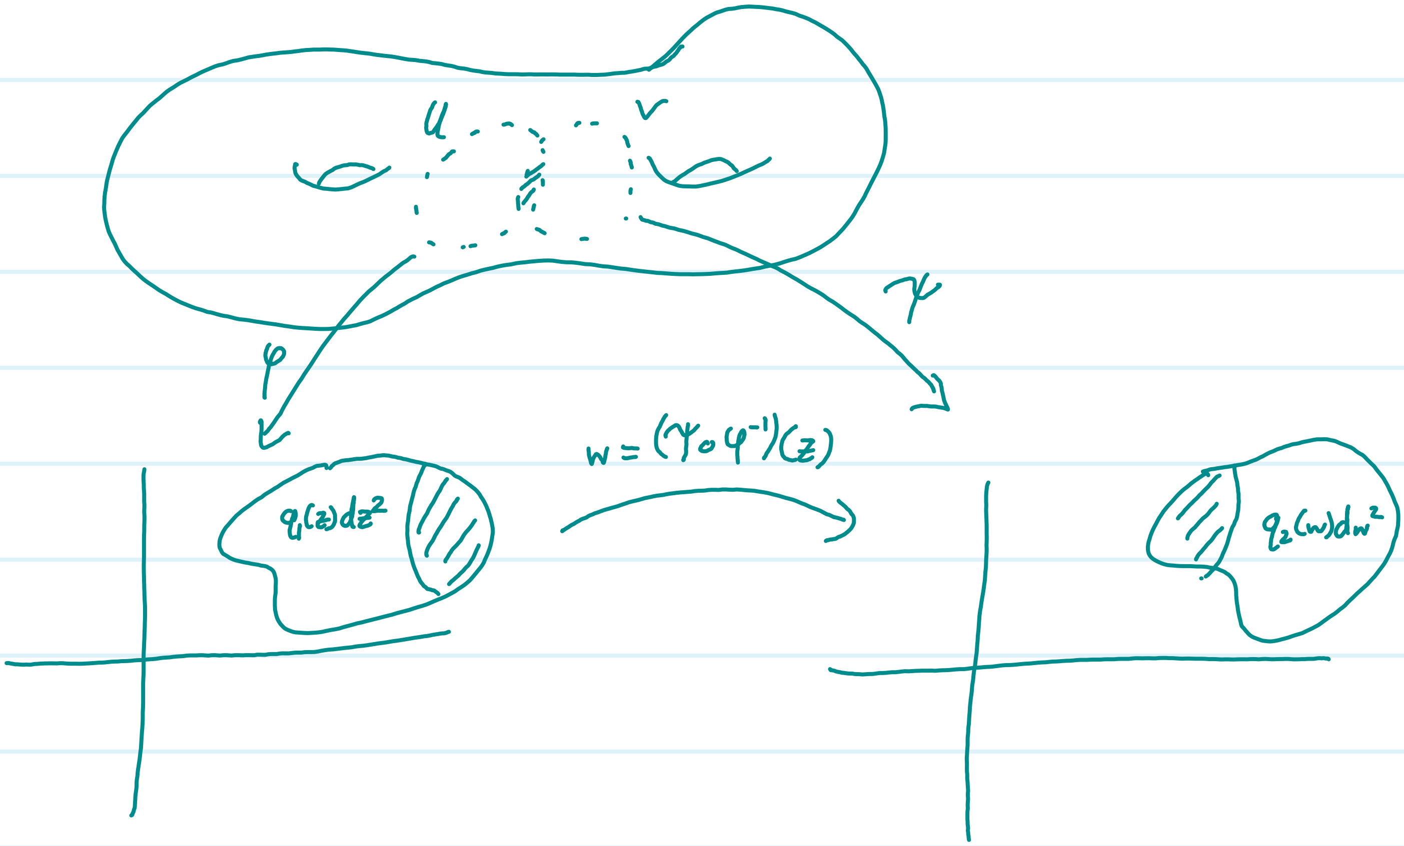

A quadratic differential is really just a holomorphic function on every chart. As always the real interesting question is how are these glued together. What the dz^2 symbol tells us is that the transition maps are given, not by derivatives like dw/dz as would be the case for 1-forms, but rather by the square of the derivative

$$q_1(z)dz^2 = q_2(w)dw^2$$

square of derivative

$$q_1(z) = q_2(w)\left(\frac{dw}{dz}\right)^2$$

change of coord map!

Now the natural question is why do we care about quadratic differentials? What do they give us? For one thing they give us a Euclidean metric (away from certain bad points) We will see that away from the bad points (which specifically are the zeroes of the differential) we can find charts so the differential looks like dz^2 (the holomorphic function in front is just 1)

Notice the form dz^2 has the nice property that if we look at its absolute value we have

actual multiplication (not wedge product)

$$|dz^2| = |dz|^2 = dz\, d\bar{z} = (dx + i\,dy)(dx - i\,dy) = dx^2 + dy^2$$

Euclidean metric!

How can we choose a nice chart like this where the quadratic differential is given by just dz^2? Suppose locally our quadratic differential is given by $q(z)\,dz^2$ where q is non-zero. Then define

$$w(z) = \int_{z_0}^{z} \sqrt{q(z)}\, dz \quad \text{so} \quad \frac{dw}{dz} = \sqrt{q(z)}$$

With this we get

$$q(z) = \tilde{q}(w)\left(\frac{dw}{dz}\right)^2$$

$$q(z) = \tilde{q}(w)\; q(z)$$

$$1 = \tilde{q}(w)$$

zeroes of q (the quad diff) are well-defined and discrete

Notice in order to take the square root we need $q(z)$ to be non-zero. Moreover we have 2 choices of square roots ($+$ or $-$) and we can translate w by a constant (the $+c$ that comes with integrals). But really these are the only choices we have. What this means is that if we pick good charts then the transition maps are of the form $w \longmapsto \pm w + c$. On any given chart we know what the horizontal and vertical directions are. In general these are not well-defined across the manifold, however, as an arbitrary holomorphic function (which is what the typical transition function looks like) can send horizontal and vertical lines to all sorts of crazy things. However if the transition functions are all given by translations and reflections then horizontal and vertical directions do make sense In summary then, if we have a Riemann surface X with a quadratic differential q, we can use q to find charts on X such that the charts are of the form $w \longmapsto \pm w + c$ In this way, we get a Euclidean structure on X In fact we get slight more we know what the horizontal and vertical directions are (if, for example, we allowed rotations this would no longer be true) As a reminder all of this can only be done away from the zeroes of q. Surfaces where the transition maps are given by $w \mapsto \pm w + c$ are called half-translation surfaces In fact we can go the other way as well Given a half-translation surface, the form dz^2 is preserved under the transition maps so is well-defined on the surface Let's look at an example

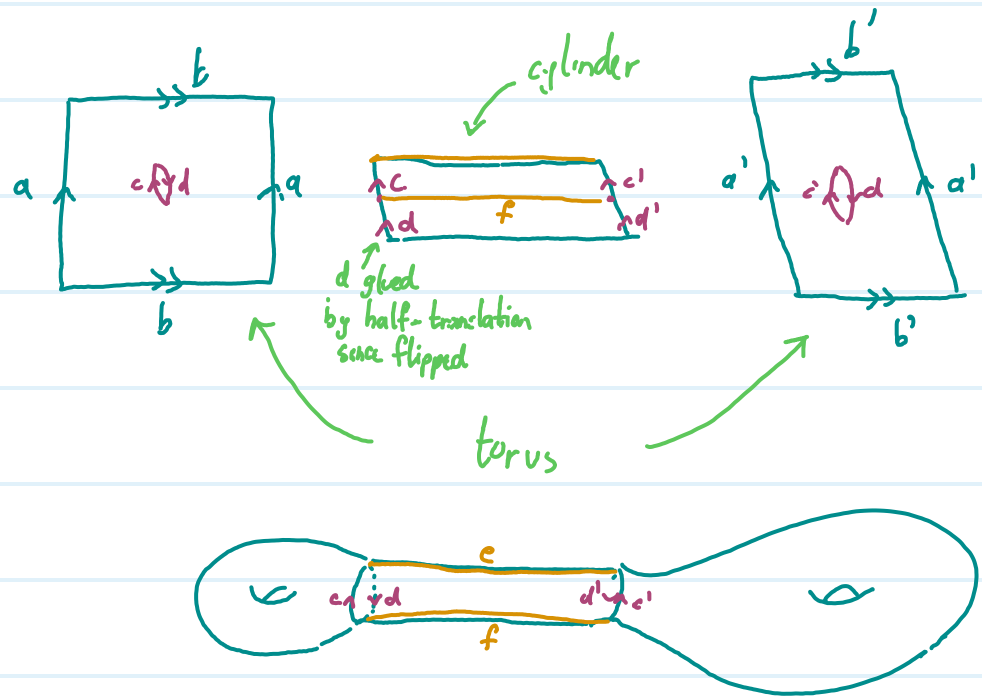

We can form a genus 2 surface as a half translation surface by taking 2 tori (which are translation surfaces) and gluing them via a cylinder (which is also a translation surface It's gluing of the cylinder to the tori which requires the half translation) by cutting a slit on each torus Thus we get an almost Euclidean structure on a genus 2-surface, in particular away from finitely many points As one might guess the problem points occur exactly at the half translations

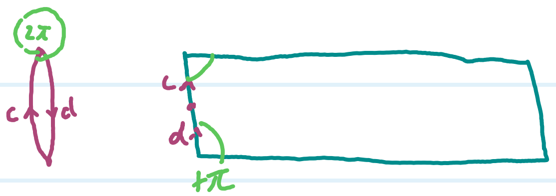

Total angle $= 3\pi$ so could not be Euclidean at that point

Let's look at the case of the torus in a bit more detail.

Let S be the topological torus. Let $\Omega(S)$ be the space of 1-forms on the torus. Of course forms require a conformal structure so really these forms come with a marking map. In other words then we have

Riemann surface, holom 1-form, $X \sim X'$ if $[X]=[X']$ in $\mathcal{T}(S)$

$$\Omega(S) = \{ (X, \omega) : S \to X \text{ homeo} \}/\sim \qquad \omega \not\equiv 0$$

Since S is a torus we can describe this space much more simply. Every Riemann surface X which is homeomorphic to a torus is given by $\mathbb{C}/\Lambda$ where $\Lambda \subset \mathbb{C}$ is a lattice. Such a lattice is determined by 2 $\mathbb{R}$-linearly independent complex numbers u, v. Now suppose ω is a 1-form on $X = \mathbb{C}/\Lambda$. This lifts to a 1-form on the universal cover $\mathbb{C}$ which is invariant under Λ. A 1-form on $\mathbb{C}$ is just $f(z)\,dz$ and since the 1-form is invariant Λ, $f(z)$ must be doubly periodic ($f(z+u) = f(z)$ and $f(z+v) = f(z)$). Thus $f(z)$ is a holomorphic function that is bounded on $\mathbb{C}$ and hence must be a constant. Hence $\omega = c\,dz$ for some $c \in \mathbb{C}$. Notice in fact what we get then is that $\Omega(S)$ forms a bundle over $\mathcal{T}(S) \cong \mathbb{H}^2$ (in the case of S = torus) with fiber given by the different 1-forms one can have on a given Riemann surface X which in this case is given by $\mathbb{C}^*$

There is a different way for us to understand $\Omega(S)$, as sitting inside $\underset{\mathbb{C}^2}{\underset{\shortparallel}{H^1(S,\mathbb{C})}} = \{\text{Linear } H_1(\underset{\mathbb{Z}^2}{\underset{\shortparallel}{S,\mathbb{Z}}}) \to \mathbb{C}\}$. Indeed, let α, β form a basis for $H_1(S,\mathbb{Z})$. Any map in $H^1(S,\mathbb{C})$ is completely determined by where α, β are sent and given some $(X,\omega) \in \Omega(S)$ there are 'obvious' complex numbers that α, β can be sent to, namely

$$\alpha \longmapsto \int_\alpha \omega \quad , \quad \beta \longmapsto \int_\beta \omega$$

This way we get a map $\Omega(S) \longrightarrow H^1(S,\mathbb{C})$. This map is injective since, by the Alexander method, an isomorphism $\pi_1(S) \longrightarrow \Lambda \subset \mathbb{C}^2$ (where Λ is a lattice) let's us recover the marking map. However this map is not onto (for example points on diagonal of $\mathbb{C}^2$ would produce a degenerate lattice which does not give us a Riemann surface) In this way we get a coordinate system for $\Omega(S)$ called period coordinates

We can also define coordinates on $\Omega(S)$ using measured foliations. Suppose we are given $(X,\omega) \in \Omega(S)$. We know that $X = \mathbb{C}/\langle u,v\rangle$ where $u = \int_\alpha \omega$, $v = \int_\beta \omega$ and $\omega = dz$ (or some complex multiple of this but that is not so important. There is an 'obvious' foliation of $\mathbb{C}$ that descends to X. A foliation is a way to decompose our space into a union of lower dimensional manifolds in such a way that locally things look like a product of a manifold (these are the leaves of the foliation) and an open interval. This is much easier to understand pictorially

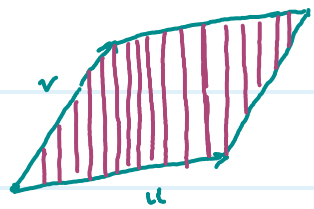

Example of foliation

The 'obvious' foliation on $\mathbb{C}$ then is to use the fact that $\mathbb{C} \cong \mathbb{R}\times\mathbb{R}$ has a product structure (so any 'nice' quotient will have a local product structure). In fact there are 2 obvious foliations on $\mathbb{C}$, namely the horizontal one and the vertical one and both of these will descend to X It would be nice if we could find these foliations without having to go to $\mathbb{C}$ and just use the data given Indeed we can – through ω. Recall we are taking $\omega = dz$. Then if we write $z = x + iy$ we have $Re(dz) = dx$ and $Im(dz) = dy$ The vertical foliation is given by $\ker(dx)$ (for any vertical tangent vector dx is 0 This defines a vector field which is integrable and hence defines the required submanifolds) We will denote this vertical foliation by μ_- We suggestively denote the foliation using the notation of measures because the foliation comes equipped with a transverse measure namely dx itself. Indeed if γ is an arc transverse to the foliation then $\left|\int_\gamma dx\right|$ gives the transverse length of γ. Notice this measure is invariant under isotopy along leaves

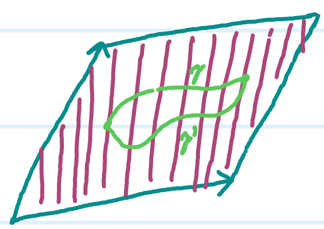

$\left|\int_\gamma dx\right|$ gives transverse length of γ, which is same as tranverse length of γ'

Completely analogously of course $dy = Im(\omega)$ defines a horizontal foliation on X which we call μ_+.

Now we want some coordinates for this space We can encode the topological information of a foliation on a torus via its slope The slope of a foliation is the slope of the line of any lift of a leaf. Thus we can topologically identify a foliation by a point in $\mathbb{RP}^1 \cong S^1$. Since we want measured foliations, we need additionally a non-negative real number c which gives the measure

$c\,dx$ (or $c\,dy$ if we prefer). It is a fact that the only measures on foliations of a torus are scalings of the Lebesgue measure. Thus the space looks like $\mathbb{R}^2 \setminus \{0\} / \pm 1$ (a point $[(x,y)]$ determines a slope (notice (x,y) and $(-x,-y)$ give the same slope.) and the length $\sqrt{x^2+y^2}$ gives the measure) Notice this is homeomorphic to $\mathbb{R}^2 \setminus \{0\}$ (by the map $z \mapsto z^2$) The only measure foliation we are missing now is the 'zero foliation' which assigns length 0 to all curves. This measure foliation does not depend on the underlying topological foliation and hence naturally sits at the origin (imagine taking a measured foliation and making the measure smaller. Via the map this traces out a path towards 0 and naturally then the limit would be 0) In other word then $MF(S)$ the space of measured foliations on a torus is homeomorphic to $\mathbb{R}^2$

Finally, it is useful to know that we can determine the slope of a foliation without having to look at lifts Indeed if μ_- is determined by dx (as above) we have

$$\text{slope}(\mu_-) = \frac{\int_\beta dx}{\int_\alpha dx}$$

And similarly of course we could find the slope of the horizontal foliation μ_+ by using dy instead. In summary then, we have a nice map

$$\Omega(S) \longrightarrow (MF(S) \times MF(S)) \setminus \Delta \quad \text{diagonal}$$

$$(X,\omega) \longmapsto (\mu_+, \mu_-)$$

In fact this map is 1-1 and onto (we need to remove the diagonal of course since the vertical and horizontal foliations cannot be the same) Suppose we are given a pair (μ_+, μ_-) Then we can recover the Riemann surface and 1-form as follows The 1-form is easy to recover by $\omega = \mu_- + i\mu_+$. (recall the measures are simply scalings of the Lebesgue measure so the sum here is going to be holomorphic For the Riemann surface X, we define it to be $\mathbb{C}/\langle u,v\rangle$ by $Re(u) = \int_\alpha \mu_-$, $Re(v) = \int_\beta \mu_-$

$I_m(u) = \int_\alpha \mu_+$ and $I_m(v) = \int_\beta \mu_+$ In this way we get a map $\Omega(S) \to MF(S) \times MF(S) - \Delta \to \mathbb{R}^4$ which gives yet another way of giving coordinates to $\Omega(S)$.

We would like to define a measure on MF(S), the space of measured foliations on a torus There is a nice way to this if our torus comes equipped with a length function on simple closed curves Recall on a torus the set of simple closed curves (up to isotopy) can be identified with $\mathbb{Q} \cup \{\infty\}$. So let $\rho : \mathbb{Q} \to \mathbb{R}_+$ be our length function We can extend this to a length function on all of $\mathbb{R}^2$ by requiring ρ to be homogeneous and continuous (a point in $\mathbb{Q}$ is then represented by a pair (p,q) of coprime integers, the set of all rational slopes is dense so indeed we can extend ρ to $\mathbb{R}^2$) Then we can for example define

we often quotient ± 1 since don't care abt orientation

$$\mu_{Th}(\{(x,y) : \rho(x,y) \le 1\}) = \lim_{L \to \infty} \frac{|\{\gamma_{p,q} : \rho(\gamma_{p,q}) \le L\}|}{L^2}$$

num simple closed curves of length $\le L$

numerator should grow quadratically (since function of area) so divide by L^2

Consider the example where $\rho(x,y) = \sqrt{x^2+y^2}$ (the usual Euclidean metric)

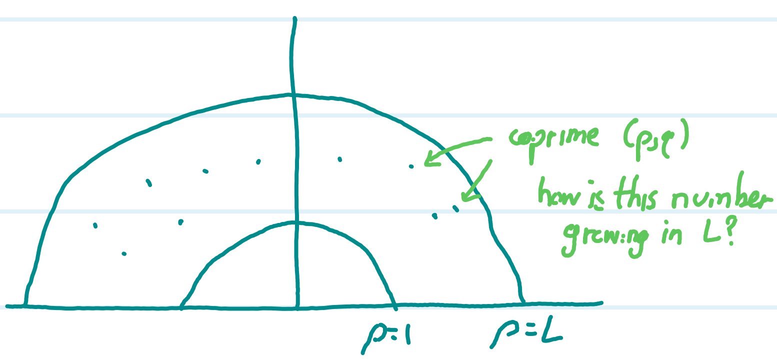

It is known that $\lim_L \frac{\{(p,q) \in \mathbb{Z}^2 : \rho(p,q) \le L\}}{\pi L^2} = 1$ (the number of lattice points in the disk is exactly proportional to the area)

Now we want to determine the density of $\mathbb{Q} \subset \mathbb{Z}^2$ (thought of as coprime pairs). Notice we can write

$$\mathbb{Z}^2 = \underbrace{\mathbb{Q}}_{\gcd=1} \cup \underbrace{2\mathbb{Q}}_{\gcd=2} \cup \underbrace{3\mathbb{Q}}_{\gcd=3} \cup \cdots$$

Suppose $\mathbb{Q}$ has density d. Then $2\mathbb{Q}$ has density $d/2^2$ (expanding the set reduces the density, in proportion to the area since we are on the plane) Similarly $3\mathbb{Q}$ has density $d/3^2$ and so on So equating the densities on both sides we get

$$1 = d + \frac{d}{2^2} + \frac{d}{3^2} + \cdots \Rightarrow d = \frac{1}{\zeta(2)}$$

(Yes this is an imprecise, heuristic argument but the statement is true nevertheless)

In summary then we have two ways of giving coordinates to $\Omega(S)$ We have the period coordinates given by $\Omega(S) \to H^1(S,\mathbb{C})$ (by integrating against ω) and we have $\Omega(S) \to MF(S) \times MF(S)$ By choosing a basis for homology we can identify $H^1(S,\mathbb{C})$ with $\mathbb{C}^2$ The usual volume form on $\mathbb{C}^2$ defines a measure on $\Omega(S)$ called the Masur-Veech measure Alternatively, if we have a length function on S we can define a measure on $MF(S)$, which gives another way of getting a measure on $\Omega(S)$ Different length functions will give different measures (in fact there is one choice of ρ which will give the Masur-Veech volume)

We would like to do something similar for $QD(S)$, the space of quadratic differentials on a surface S (here we genuinely mean an arbitrary surface S and not just the torus) Period coordinates are not so easy to define for quadratic differentials (although we will do it!) However quadratic differentials do define (singular) measured foliations so we should be able to use that theory to define a measure on $QD(S)$.