Counting Problems in Teichmüller Space II

Talk by Alex Eskin

August 24, 2007

Let Q_g denote the space of quadratic differentials. $q \in Q_g$ naturally defines a pair of measured foliations which are the horizontal and vertical foliations α_h and α_v in the flat structure associated to q. This gives a map $Q_g \to \mathcal{MF} \times \mathcal{MF}$ which sends $q \mapsto (\alpha_h, \alpha_v)$.

While the set of all measured foliations on a surface is very large, one can put an equivalence relation on this set which makes $\mathcal{MF}$, the space of measured foliations, into a finite dimensional manifold having half the dimension of the space of quadratic differentials. Another nice property of this space is that it admits a measure on $\mathcal{MF}$ called the Thurston measure.

The following theorem is due to Hubbard and Masur. For a point X in Teichmüller space, let $Q_g(X)$ denote the space of quadratic differentials holomorphic at X.

THEOREM: (Hubbard-Masur) The map $\Phi : Q_g(X) \to \mathcal{MF}$ given by sending q to its horizontal foliation is a homeomorphism such that $\Phi(\{q \in Q_g(X) \mid \text{Area}(q) \leq 1\})$ is the set $\{\alpha \in \mathcal{MF} \mid \text{Ext}_\alpha(X) \leq 1\}$

The extremal length of a measured foliation α, $\text{Ext}_\alpha(X)$, can be defined as $\lim_{i\to\infty} \text{Ext}_{\alpha_i}(X)$, where α_i is a sequence of simple closed curves converging to α.

This theorem allows us to define an interesting function. At each point X, we can look at the subset of $Q_g(X)$ consisting of area one quadratic differentials. This is analogous to a unit ball in this space. The Hubbard-Masur function then measures the area of the image of this ball under Φ, i.e. $\Lambda : \mathcal{T}(S) \to \mathbb{R}$ is given by

$$\Lambda(X) = \mu_{Th}\{\alpha \in \mathcal{MF} \mid \text{Ext}_\alpha(X) \leq 1\}$$

This function ends up being useful in answering many counting problems. We will mention three of these results in this talk.

The first question we will look at involves maping class orbits of points in Teichmüller space. Let $\Gamma = Mod_g$, the mapping class group of a genus g surface. The orbit of a point

X under the action of Mod_g is a discrete set in Teichmüller space. For any $X, Y \in \mathcal{T}(\mathcal{S})$, we would like to understand how the number of mapping class translates of X in the the ball of radius R around Y grows as a function of R.

The following theorems are due to Arethreya, Bufetov, Eskin and Mirzakhani. The constant h is equal to $3g-3$ in each.

THEOREM 2:

$$|\Gamma X \cap B_R(Y)| \sim \frac{1}{h}\Lambda(X)\Lambda(Y)e^hR$$

The next theorem (which, loosely speaking, can still be considered a "counting problem") shows how the measure of the ball of radius R grows as R increases.

THEOREM 3: $m(B_R(X)) \sim \frac{1}{h}\Lambda(X)e^{hR}\left(\int_{\mathcal{M}_g}\Lambda(y)dy\right)$, where $\mathcal{M}_g$ denotes the moduli space of a surface of genus g.

THEOREM: The number of multicurves on X of extremal length less than L is asymptotic to $\Lambda(X)L^h$.

Roughly speaking, multicurves form a lattice in the space of measured foliations, so the above theorem fits into a family of results involving counting lattice points in manifolds. The last theorem above can be taken as an exercise in the theory of measured laminations.

The outlines of the proofs all of the above theorems are similar: follow the methods from Margulis' thesis in the compact part of moduli space and then deal with the non-compactness appropriately. We will not discuss the proofs of these theorems here, but will instead return to the problem of counting closed geodesics in moduli space. Recall the following theorem, which was stated in the previous lecture:

THEOREM 1: (Mirzakhani-Eskin, Veech, Hamenstädt, Rafi, Bufetov)

$$N(R) \sim \frac{e^{hR}}{hR}, \text{ where } h = 6g-6$$

We won't be able to give this proof in detail, but we will be able to sketch some of the ideas involved. To start off, we will sketch an elegant argument of Margulis' that describes how to count geodesics in a compact negatively curved manifold M. This argument doesn't work for moduli space, as moduli space is not compact. We will come back to the problem of noncompactness in a moment, but for the time being lets pretend that $\mathcal{M}_g$ is a compact manifold and see what we can do.

Fake Proof. Let U be a small box in the tangent space $T\mathcal{M}_g$, and let μ be a measure on $T\mathcal{M}_g$ that is invariant under the geodesic flow ϕ_t, normalized so that $\mu(T\mathcal{M}_g) = 1$.

The first fact we will need is that the geodesic flow on moduli space is mixing, i.e. for $A, B \subset (T\mathcal{M})$, $\mu(\phi_t(A) \cap B) \approx \mu(A)\mu(B)$ for large values of t.

We now want to consider what happens to U under the action of the flow. We can think of the action of the geodesic flow as stretching U in one direction and contracting U in another. For simplicity, we'll finish the argument in as though we were working in two dimensions, which gives a small cross section of the actual picture. Because of the mixing property of the flow, we see the image of U under the flow intersecting U many times.

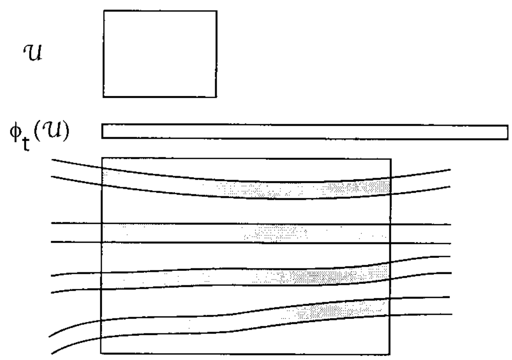

Intersection of $\mathcal{U}$ and $\phi_R(\mathcal{U})$ for large R

Let R be large enough that $\phi_R(U) \cap U$ has more than one component. One can show that each component of the intersection $\phi_R(U) \cap U$ contains exactly one closed geodesic of length between R and $R + \epsilon$, where ϵ is the diameter of U. This gives that counting connected components of this intersection gives the number of closed geodesics of length around R. From this we see that the number of closed geodesics of length between R and $R + \epsilon$ is given by

$$\frac{\mu(U \cap \phi_t(U))}{\mu(A)}$$

where A is a single connected component. The flow contracts in one direction by e^{-R}, so the area of each component of the intersection has area about $\mu(U)e^{-R}$. By the mixing property, we know that the numerator in the preceding equation is about $\mu(U)^2$, so we get that the number of geodesics of length about R is $\mu(U)e^R$.

Of course not all geodesics will pass through U. To catch all of them, we create a tiling of moduli space by boxes U_i. It would be nice to argue that the number of geodesics of length R passing through U_i is about $\mu(U_i)e^R$, so the total number of geodesics of length R is $\sum \mu(U_i)e^R = e^R \sum \mu(U_i) = e^R$. This is wrong, of course, because a geodesic will pass through many boxes. If the width of the boxes is ϵ, and the length of the geodesic is about R we get that each of our paths passed through about R/ϵ boxes. This shows that our counting was off by a factor of R/ϵ, so our final count for the number of geodesics of length between R and $R+\epsilon$ is $e^R/(R/\epsilon) = \epsilon \cdot e^R/R$.

Had we worked in the correct dimension, there would be a factor of h in the above estimate, giving that the number of closed geodesics between R and $R+\epsilon$ is about $\epsilon \cdot e^{hR}/R$. Thus to find the total number of geodesics of length less than R, we must estimate the integral of the function e^{hx}/x. As we are only interested in the growth of this function, we don't care about the behavior of this function for small x. Thus up to a constant, the function of R we are interested in is

$$\int_1^R \frac{e^{hx}}{x} dx$$

This is a hard integral to take. Luckily we don't need to compute it but only need to show that

$$\lim_{R\to\infty} \frac{\int_1^R e^{hx}/x\, dx}{e^{hR}/hR} = 1$$

Repeated application of integration by parts gives that the above integral equals

$$\left[\frac{e^{hx}}{hx} + \frac{e^{hx}}{(hx)^2}\right]_1^R + 2\int_1^R \frac{e^{hx}}{h^2x^3} dx \sim \frac{e^{hR}}{hR} + \frac{e^{hR}}{(hR)^2} + 2\int_1^R \frac{e^{hx}}{h^2x^3} dx$$

Clearly the ratio of the second term and $e^{hR}/(hR$ goes to zero as R goes to ∞, so it remains to show that this is also true for the remaining integral. The integral is clearly smaller than $R \cdot e^{hR}/(h^2R^3)$ for large values of R, and

$$\lim_{R\to\infty} \frac{R \cdot e^{hR}/(h^2R^3)}{e^{hR}/(hR)} = 0$$

This shows that $\int_1^R \frac{e^{hx}}{x} dx \sim e^{hR}/(hR)$, concluding the fake proof that the number of geodesics of length less or equal to R is asymptotic to $e^{hR}/(hR)$.

The argument sketched above made some implicit assumptions about moduli space, some of which are false. The most difficult to get around is the compactness assumption. There may be closed geodesics far out in the thin part of moduli space, that are therefore able to stay outside of compact sets. For the remainder of this lecture we will describe how to deal with non-compactness.

The following lemma helps deal with the thin part by guaranteeing that short geodesics in $\mathcal{M}_g$ stay close to the thick part.

LEMMA: (Veech) There exists a constant $C = C(g)$ such that if γ is a closed geodesic in $\mathcal{M}_g$ of length R, then γ is contained in a cR neighborhood of the thick part of moduli space. Here the distance is given by the Teichmüller metric projected to moduli space.

Proof. We will need a few facts to get us started. The first is that for curves with small extremal length, extremal length is close to hyperbolic length. Thus in discussing short curves, we don't need to make any distinction between extremal length and hyperbolic length.

The second fact is that if we move a distance τ in Teichmüller space, the extremal length of any curve can change by a factor of at most $e^{2\tau}$. This follows from Kerckhoff's formula:

$$d_{\mathcal{T}}(X,Y) = \sup_{\alpha} \frac{1}{2}\left(\frac{\mathrm{Ext}_{\alpha}(X)}{\mathrm{Ext}_{\alpha}(Y)}\right)$$

Let ϵ be a number such that any pair of curves on a hyperbolic surface of length less than ϵ have disjoin collars. By a "short" curve, we mean a curve of length less than ϵ. Note that the set of short curves is finite; indeed, it is less than $3g-3$, the size of a pants decomposition on S.

Now we can proceed with the proof. Suppose that Veech's lemma does not hold. Then there is a closed geodesic γ in moduli space, and some point X on γ which is very far from the thick part of moduli space, in which no curve is shorter than ϵ. Kerckhoff's formula gives that this is equivalent to the extremal length of some simple closed curve being very short on X. By the first fact above, this is equivalent to the hyperbolic length of some closed curve on X being very short.

As was mentioned in the last lecture, each closed geodesic in moduli space is associated to a conjugacy class of pseudo-Anosovs, and each individual pseudo-Anosov acts on Teichmüller space via a deck transformation corresponding to $[\gamma] \in \pi_1(\mathcal{M}_g)$.

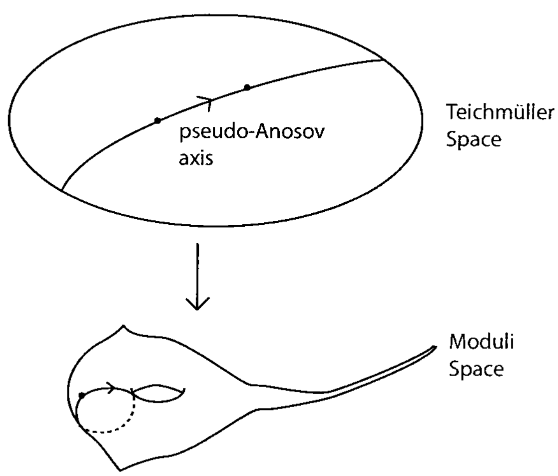

Let f be one of these pseudo-Anosovs. The translation length of f is exactly R, the length of γ, so f moves the lifts of points on γ exactly R. As hyperbolic structures, $f(X) = X$, but f may change the name of α in terms of the original marking on X. However, the second fact above tells us that f changes the extremal length of α by a factor of at most e^{2R}. We can then apply the first fact to deduce that the hyperbolic length of α has changed by about a factor of e^{2R} as well. Repeating this argument, we can see that if α is of length less than $\epsilon \cdot e^{-2R(3g-3)}$, then $f^k(\alpha)$ is short for $0 \leq k \leq 3g-3$. As $f^k(\alpha)$ is short curve on X and there are at most $3g-3$ short curves on X, so $f^k(\alpha) = f^j(\alpha)$ for some k and j between 0 and $3g-3$, thus some power of f fixes a curve. This is a contradiction, as f was assumed to be a pseudo-Anosov.

This shows that $l(\alpha) > \epsilon \cdot e^{-2R(3g-3)}$, so for Y in the thick part of moduli space,

$$\frac{\text{Ext}_\alpha(X)}{\text{Ext}_\alpha(Y)} > \frac{\epsilon \cdot e^{-2R(3g-3)}}{\epsilon} = e^{-2R(3g-3)}$$

and by Kirckhoff's formula $d_{\mathcal{T}}(X,Y) \leq R(3g-3)$. □

In light of this lemma, one may wonder whether there are any geodesics in the thin part. The following theorem shows that there are.

THEOREM: (Rafi, Hamenstädt) Fix any compact set K in moduli space. Let $N_K(R)$ be the number of closed geodesics of length at most R which do not intersect K. For all K, $N_K(R) \sim Ce^{(h-1)R}$.

A complementary result is given by the following:

THEOREM: (Eskin-Mirzakhani) There exists a compact set K, which is the δ-thick part of moduli space for some δ such that

$$N_K(R) < C_\epsilon e^{(h-1+\epsilon)R}$$

We would like to outline some of the ideas in the proof of this result. To do so, we first pick any net N in Teichmüller space, i.e. a discrete set of points such that every point in Teichmüller space is close to one of the points in the net. Fix $\tau >> 0$. We now define a random walk on the net by stipulating that each step in the walk moves from a net point to another net point τ close to it.

We now can define a map Φ from the space of geodesics on moduli space to the set of random walk trajectories. To do this we discretize our geodesic by taking evenly spaced points along it, and then picking the nearest net point to each of these.

We want to estimate the number of geodesics that stay in the thin part, and to do this we count the number of random walk trajectories that stay in the thin part. This may seem like a hopeless way of counting geodesics, as the map Φ is clearly not surjective. There are many random walks which do not look like nearest point projections from geodesics.

Such a walk may move much less efficiently than a geodesic, for example by backtracking. Negative curvature saves the day however, as random walks in a negative curvature are very unlikely to backtrack. In the end, the approximation is weakened by counting random walks rather than geodesics, but the price to pay can be shown to be relatively small.

Another problem is that Φ is not injective. In fact, Φ even fails to be bounded-to-one. The reason for this is that geodesics in Teichmüller space can track each other for a very long period of time. While in Teichmüller space these geodesics will eventually diverge, they may descend to curves in moduli space that stay very close together. If these curves are bunched close enough together, their projections to the net points can be identical.

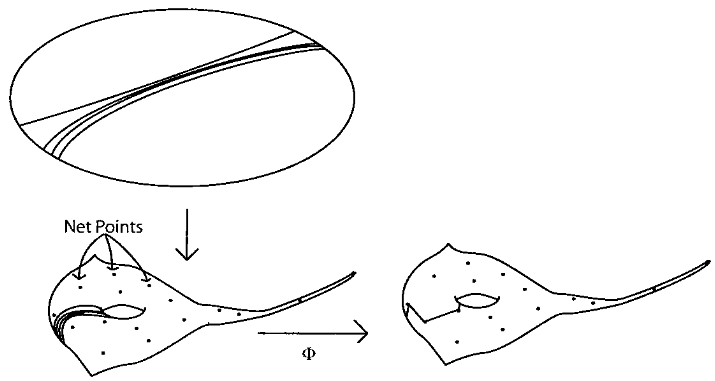

The following argument shows, however, that Φ can only be badly non-injective very far out in the thin part, i.e. for all ϵ there exists K such that in the ϵ-thick part of moduli space, Φ is at most K to 1. To see this, suppose there are many closed geodesics in moduli space which track each other. Then some pair γ_1 and γ_2 will run ϵ close together. Let $\widetilde{\gamma_1}$ and $\widetilde{\gamma_2}$ be ϵ close lifts of γ_1 and γ_2 to Teichmüller space. Let ϕ be the mapping class that translates along $\widetilde{\gamma_1}$ and λ the mapping class that translates along $\widetilde{\gamma_2}$. If X is a point on $\widetilde{\gamma_1}$ that is within ϵ of γ_2, then $d(\lambda^{-1}\phi(X), X) \sim \epsilon$. But the only mapping class elements that have very small translation lengths are given by Dehn twists about short curves. This shows that such bunching can only occur in the presence of short curves, i.e. in the thin part of moduli space. Reasoning appropriately, one can deduce that the multiplicity of the map Φ is approximated by $\prod_\alpha \frac{1}{l_\alpha(X)}$, where the product is taken over all short curves α.

$u(X) = \sum \epsilon_j f_j(X)$, where $f_j(X) = \prod_{1 \leq i \leq j} \frac{1}{l_j(X)^{1/2}}$, $l_j(X)$ is the length of the j-th shortest curve on X and ϵ_j is an appropriately chosen constant. This function u is large on the thin part of moduli space, and small on the thick part of moduli space. It is then possible to derive the following inequality:

$$(A_\tau u)(X) \leq C\tau e^{-N(X)\tau} u(X)$$

where $(A_\tau u)(X)$ denotes the average of u over the ball of radius τ about X, C is a constant and $N(X)$ is the number of short curves on X. For very large values of τ, $C\tau e^{-N(X)\tau} << 1$ whenever any curve is short, so the average of $u(X)$ over the ball of radius τ is less than $u(x)$. This shows that at each step of the random walk, there is a very high probability that $u(X)$ will decrease. However, from the definition of $u(X)$, one sees that $u(X)$ decreases exactly when we head toward the thick part of moduli space, so we see that with a very high probability a random walk will go back to the thick part of moduli space.

The estimate $(A_\tau u)(X) \leq C\tau e^{-N(X)\tau}u(X)$ comes from Minsky's product formula. In doing this calculation, we may assume that we are in the thin part of moduli space (if our random walk stays in the thick part then we can't contribute to $N_K(R)$, as K contains the thick part of moduli space). Recall that Minsky's product formula says that the thin part of moduli space can be coarsely parameterized by a product of hyperbolic planes corresponding to the short curves and the Teichmüller space of the cusped surface given by pinching all short curves to cusps. Recall also that the hyperbolic planes in the product formula have y coordinate given by the reciprocal of the length of a short curve, so $u(X)$ can be estimated using the y coordinates of the hyperbolic plane factors.

In the hyperbolic plane, we can compute the average of the function $f(x, y) = y^{1/2}$ over a ball of radius τ and we find $(A_\tau(f))(x, y) \leq Ce^{-\tau}f(x, y)$. Using the fact that with $N(X)$ short curves we are in a product of $N(X)$ hyperbolic planes with the sup metric (the other component of the product formula is negligable) and that $u(X) \sim \prod_{i=1}^{i=N(X)} y_i^{1/2}$, we get the inequality $(A_\tau(f))(x, y) \leq Ce^{N(X)-\tau}f(x, y)$.