Model Manifolds for Punctured Torus Groups and Constructions in Higher Complexity

Talk by Yair Minsky

August 22, 2007

For the first part of this lecture, we will finish up the discussion of punctured torus groups from last time. For the remainder we will discuss the curve complex of a general surface, and how it is used to understand surface groups of higher complexity.

Recall the setup from last time: S is a punctured torus, and $\rho : \pi_1(S) \to \mathrm{PSL}_2(\mathbb{C})$ is a discrete faithful quasi-Fuchsian representation with such that for all generators A and B of $\pi_1(S)$, $\rho(ABA^{-1}B^{-1})$ is parabolic.

EXERCISE: Last time the parabolicity condition was stated slightly differently. Our previous stipulation was that $[\gamma]$ must be sent to a parabolic, where γ is the curve homotopic to the boundary. Show that a pair of curves on the punctured torus are a pair of generators for the fundamental group if and only if they come from curves that intersect exactly once, and deduce the equivalence of the two phrasings of the parabolicity condition.

In the previous lecture, we showed how the geometry manifolds coming from quasi-Fuchsian punctured torus groups $\Gamma(X, Y)$ can be deduces from the positions of X and Y in Teichmüller space with respect to the Farey graph.

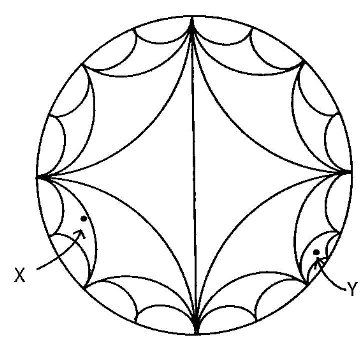

Recall the following theorem from the last lecture:

THEOREM: If $\{\alpha_n\}$ is the pivot sequence from X to Y, then there exists a constant L_0 such that for all n $l(\rho(\alpha_n)) \leqslant L_0$.

We will now describe how to use this information to build a coarse model of what the manifoldnlooks like geometrically. The first step is to describe how to find 1-Lipschitz surfaces in the manifold based on this information.

PROPOSITION: There exists L_1 such that if A and B are a pair of generators for the fundamental group of S with $l(\rho(A)), l(\rho(B)) < L_0$, then there exists $(X, f) \in LIP(M)$ such that the length of both A and B in X are at most L_1.

The following figure illustrates this proposition, where A^* and B^* denote the geodesic representatives of the curves corresponding to $\rho(A)$ and $\rho(B)$.

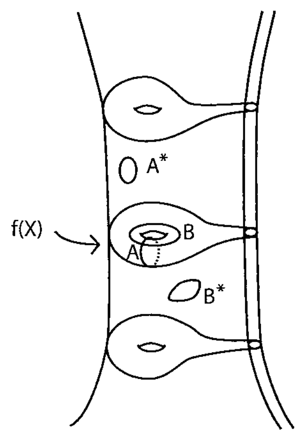

The proposition implies the claim from the last lecture that if $\rho(A)$ and $\rho(B)$ are both bounded then their product is bounded. To see this, notice that if such a surface exists, then the product of $\rho(A)$ and $\rho(B)$ is bounded by $2l(A)+2l(B)$ as a curve on X representing the product of A and B traces along A twice and B twice.

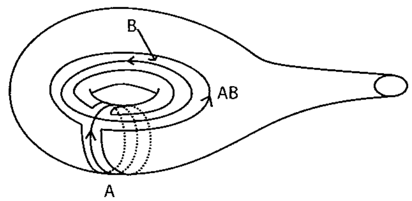

As we will see below, it turns out that this claim is essentially what is used to prove the proposition, so the claim and the above proposition are essentially equivalent.

Sketch Proof of Proposition: To simplify the argument assume $\rho(A)$ and $\rho(B)$ both have translation length 1. Looking in $\mathbb{H}^3$, we can drop a perpendicular from the axis of $\rho(A)$ to the axis of $\rho(B)$. $\rho(A)$ takes some point x_1 on one side of this axis to a point y_1 a distance 1 away from it on the other side of it. Likewise there are points x_2 and y_2 on the axis of $\rho(B)$ on either side of the perpendicular such that $\rho(B)x_2 = y_2$. We may draw hyperplanes through each of the points x_1, y_1, x_2 and y_2. If the length L of the perpendicular is long, then these hyperplanes are disjoint, as shown below.

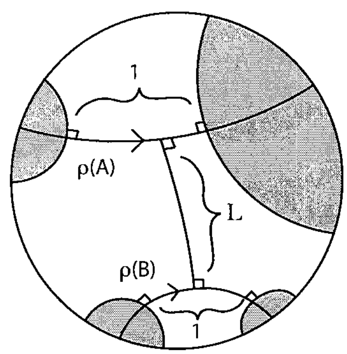

The following exercise shows that this is precluded by the parabolicity condition.

EXERCISE: Show that if the hyperplanes shown above are disjoint, then $\rho(A)$ and $\rho(B)$ generate a Schottky group, and hence consists entirely of loxodromic elements.

As $\rho([A, B])$ must be a parabolic, this is impossible so the axes of $\rho(A)$ and $\rho(B)$ must be close together. At the midpoint of the perpendicular, both $\rho(A)$ and $\rho(B)$ both have short translation length, hence at the image of this point in the quotient manifold there are loops homotopic to A and B that are simultaneously short. We then see a pair of loops in the manifold through a single point, which can be completed to a triangulation and used to construct a 1-Lipschitz surface.

$\square$

As was mentioned in the last lecture the proof of the previous theorem can be done rigorously using trace identities, but the above pictures give some of the intuition behind this result.

We can now describe how to put all this information together to build a geometric model of the manifold. The pivot sequence gives a list of curves $\{\alpha_n\}$ which must have bounded length in the 3-manifold. The curves α_n and α_{n+1} are neighbors in the Farey graph, which means that they intersect exactly once, and hence that they give a pair of generators of the punctured torus group. We can then apply the above proposition to find a 1-Lipschitz surface on which both curves sit simultaneously. So for each n, we have $(X_n, f_n) \in LIP(M)$ such that α_n and α_{n+1} are simultaneously bounded. The following picture suggestively shows how these surfaces sit in the manifold. Note that these surfaces are not necessarily embedded, so this picture is schematic.

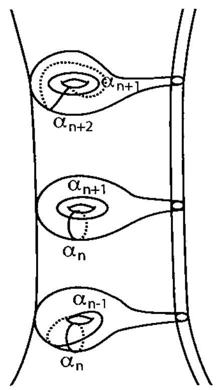

We have not proven that these surfaces show up in the correct order, but this is in fact the case. As shown in the picture, the short curves on the surfaces are given by the pivots. While in the above picture α_{n-1} and α_{n+1} look like the same curve on the surface they sit on, these curves are not homotopic in the manifold as they differ by Dehn twists (exactly w_n of them) about the curve α_n.

We would like to describe how two successive surfaces are related. If we look at f_n and f_{n-1}, the image of α_n under both maps has controlled length. If we look at the complement of an annulus around α_n in X_n and X_{n+1}, then we see two three holed spheres.

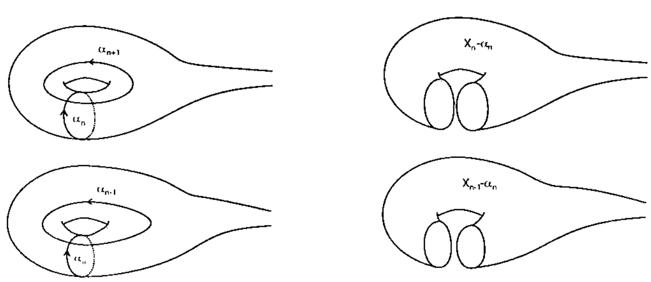

Because α_n has roughly the same length in X_n and X_{n-1}, these two three holed spheres have boundary lengths which are roughly the same. Assume for simplicity that the boundary lengths are exactly the same. The following exercise shows that this information is enough to allow us to assume that the two 3-holed spheres have the exact same hyperbolic structures.

EXERCISE: Show that a hyperbolic structure on a three holed sphere is determined by the lengths of its boundary components.

The above exercise gives that we can assume that $X_n - \mathcal{N}(\alpha_n)$ and $X_{n-1} - \mathcal{N}(\alpha_n)$ are isometric to a single surface X, so we may assume that the restrictions of f_n and f_{n-1} to the complement of α_n have the same domain X. The is useful, because now we can claim that $f_n|_X$ and $f_{n-1}|_X$ are homotopic maps into the three manifold with a short homotopy between them. That these two 3-holed spheres are homotopic follows from the fact that the two copies of α_n are homotopic in the three manifold. The following picture can be used to show that this homotopy cannot be too long.

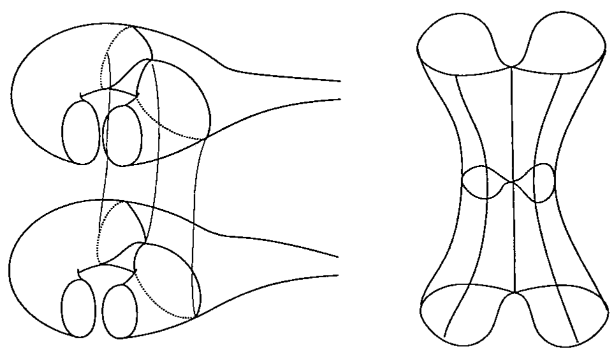

Applying some elementary hyperbolic geometry, one can show that because the loops on either side have bounded lendth, if the two sides of the homotopy are very far apart then somewhere along this homotopy we must see two very short loops based at the same point. Margulis' lemma gives that if two essential closed curves based at the same point in a hyperbolic manifold are very short, then they must generate an abelian group. But these two loops generate a free group, so this is impossible. This gives a bound on how far apart the two surfaces can be. There is another possibility in the above discussion, which is that the two surfaces are close together but any homotopy between them takes a very long time. We won't discuss it here, but this difficulty can be dealt with.

Now we have that for any pair of successive surfaces, the surfaces given by removing the curve that they have in common are connected by a bounded homotopy. Therefore if we take the 3-holed sphere X crossed with an interval, this chunk maps into the surface in a controlled way via $H(x,t) = h_t(x)$, where h_t is the homotopy between $f_{n-1}|_X$ and $f_n|_X$. This map gives a Lipschitz map from $X \times I$ to the three manifold. We can now stack these blocks on top of each other as shown below and piece these maps together to get a Lipschitz map from a model we have constructed to the 3-manifold.

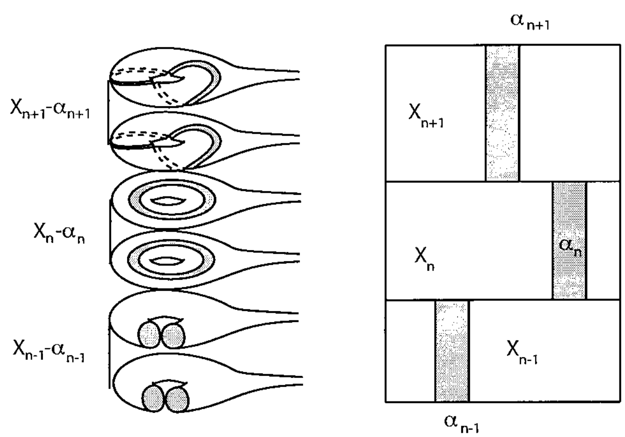

Schematic Picture

Each of the maps $f_n|_X$ misses and annulus crossed with an interval, i.e. a solid torus corresponding to a neighborhood of the pivot curve α_n, and has a corresponding hole in its domain. To complete the model, one has to fill in these holes with the appropriate Margulis tubes, but we won't give the details here.

We now have completed our model of the manifold and out map f, and it remains to prove that map we have constructed is nice. We first show that the map is proper. Given this fact, we can show that the map is degree 1 by compactifying the model and extending f to a map from the compactified model to the hyperbolic manifold together with its conformal boundaries. We then show that the map is degree 1 on the boundary, and use the fact that given a proper map between compact manifolds the degree of the map is the degree of the restriction of the map to the boundary. Then comes the task of showing that the map is in fact bi-Lipschitz, but we won't say anything about that here.

Model Manifolds and the Ending Lamination Conjecture-

Coming up with a bi-Lipschitz model of a manifold may seem interesting enough in its own right, but we will now show how this can be used to prove the ending lamination conjecture.

We commented last time that the discussion for quasi-Fuchsian punctured torus groups $\Gamma(X,Y)$ with X and Y in the interior of Teichmüller space can be repeated in the case where

X and Y are points on the boundary of Teichmüller space. If X and Y are distinct irrational points on the boundary (i.e. the images under stereographic projection of irrational points on the real line to a circle of radius 1 tangent to the real axis at 0) then $\Gamma(X, Y)$ will be what is called a degenerate group. An example of such a group is given by considering the cyclic cover of a hyperbolic three manifold that fibers over the circle. For such a group, the limit set is the whole sphere and the domain of discontinuity has disappeared, so such groups do not have conformal boundaries. The points X and Y no longer correspond to the conformal structure at infinity, but to a different kind of data called ending laminations. Work of Bonahon and Thurston showed how to associate such data to a group, but it was not known until recently whether this ending data determined the group. Thurston's ending lamination conjecture, now a theorem, asserts that this is the case.

We will now show how the model can be used to prove this claim for punctured torus groups. Suppose ρ_1 and ρ_2 are two punctured torus groups with the same pair (X, Y) of points on the boundary of Teichmüller space. We claim that the quotient manifolds for ρ_1 and ρ_2 are isometric.

Given the pair (X, Y) we can produce a model M_0 as above independent of the representations ρ_1 and ρ_2. We can then produce a bi-Lipschitz map $F_1 : M_0 \to N_1 = \mathbb{H}^3/\rho_1(\pi_1(S))$ as above, and a bi-Lipschitz map $F_2 : M_0 \to N_2 = \mathbb{H}^3/\rho_2(\pi_1(S))$. $\Phi = F_2 \circ F_1^{-1} : N_1 \to N_2$ is a bi-Lipschitz map as well, showing these two manifolds are Lipschitz equivalent. We then use some rigidity results to finish the argument. Using Mostow's argument, we can show that Φ lifts to a map $\widetilde{\Phi}$ that conjugates the actions of ρ_1 and ρ_2 and extends to a map whose restriction to the boundary $\widehat{\Phi} : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ is quasi-conformal. We then apply Sullivan's rigidity theorem, which gives that because the limit set is the whole sphere $\widehat{\Phi}$ is actually conformal. Since a conformal map from the Riemann sphere to the Riemann sphere is a Möbius transformation and Möbius transformations extend to isometries of $\mathbb{H}^3$, we end up with an isometry that conjugates the actions of ρ_1 and ρ_2. This shows that the quotient manifolds are isometric.

Constructions for more complicated surface groups-

We would like to repeat this discussion for Kleinian surface groups other than the punctured torus. The first thing to replace is the Farey-Graph. The analogous object in higher complexity is $\mathcal{C}(S)$, the complex of curves for the surface S. For simplicity we will assume that S is not equal to an annulus, a thrice punctured sphere, a once punctured torus or a four times punctured sphere, as for all of these surfaces $\mathcal{C}(S)$ must be defined differently (e.g. for the once punctured torus, $\mathcal{C}(S)$ is identified with the Farey tesselation as above). We also don't define the curve complex for the sphere, disk, or torus.

Definition For a compact orientable surface S other than those listed above, $\mathcal{C}(S)$ is a simplicial complex with a vertex for every essential (i.e. non-peripheral, non-trivial) simple closed curve, and vertices $\{\alpha_1, \alpha_2, ..., \alpha_n\}$ span an n-simplex if their associated curves can be realized disjointly.

As an example, the following curves on a genus two surface span a 2-simplex.

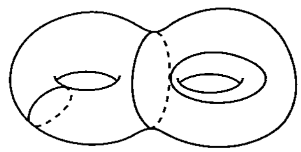

It is clear that this definition will not yield the Farey graph for the punctured torus. On the punctured torus, any two essential curves that do not intersect must be homotopic, so instead we alter the definition so that curves are connected by an edge if they intersect exactly once. Similar alterations are necessary in the other low complexity cases, for example in the case of the 4-holed sphere v_1 and v_2 are connected by an edge if they intersect twice.

To motivate studying this object, we remark that if $[\alpha, \beta]$ is an edge in the curve complex of $\mathcal{C}(S)$, and $\rho : \pi_1(S) \to \mathrm{PSL}_2(\mathbb{C})$ is a surface group, then there exists $(X, f) \in LIP(\mathbb{H}/\rho(\pi_1(S)))$ such that f takes α and β to their geodesic representatives. This remark follows from the techniques related to those described in Ken Bromberg's lectures, as given two disjoint curves you can add a vertex on each curve and extend to a two-vertex triangulation. This shows that if α and β are neighbors in the curve complex then they have a common 1-Lipschitz surface to live on.

By finding surfaces on which the geodesic representative of different curves live, we can build ladders of surfaces through the manifold.

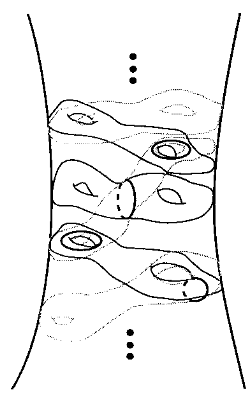

This suggests that the combinatorics of the curve complex can be very useful in understanding the geometry of the manifold.

Recall that the identification of the short curves in the punctured torus case relied on the fact that the edges in the Farey graph separate the disk. This is a tree-like property of the Farey graph, in that to travel between two points one cannot avoid crossing any edge in between them

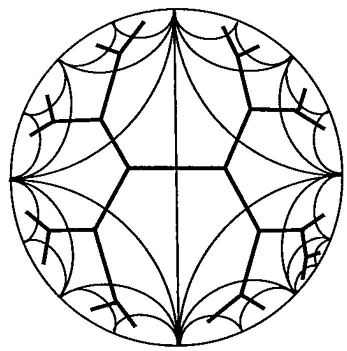

In higher complexity the curve complex is not tree-like, and we need a different way to overcome this difficulty. It turns out that higher complexity curve complexes have a property that compensates for this.

DEFINITION: A geodesic metric space is δ-hyperbolic if for all geodesic triangles $[X, Y, Z]$, $[X, Y] \subset \mathcal{N}_\delta([X, Z] \cup [Y, Z])$.

THEOREM: (Masur-Minsky) $\mathcal{C}(S)$ is δ-hyperbolic.

In a δ-hyperbolic space geodesic triangles are thin. An important consequence of this is the stability of quasi-geodesics, i.e. there is a unique geodesic that stays a bounded distance away from it. This is not the case in spaces that are not δ-hyperbolic. In the Euclidean plane, for example, a logarithmic spiral is a quasi-geodesic which does not stay a bounded distance from any straight line. In a δ-hyperbolic space, quasi-geodesics cannot behave too badly, they must do roughly what the geodesic that tracks them does.

Therefore if a geodesic which tracks a quasi-geodesic in a δ-hyperbolic space passes through a certain set of points, then though the quasi-geodesic need not pass through them precisely, it must at least come close. This is the property of the curve complex that stands in for the separation property of the Farey graph. This will be described in more detail in later lectures.

As a final topic, we will describe the analog of the rotation numbers w_n that appeared in the discussion for the Farey graph. Loosely speaking, these numbers indicate how twisted

X is about α relative to Y. If $\{...\alpha_{-1}, \alpha_0, \alpha_1, ...\}$ are the curves in the pivot sequence, for each α_n we see a fan of w_n triangles. The curve α_{n+1} is given by doing w_n Dehn twists about α to α_{n-1}, so the number w_n give exactly how much twisting about α_n is necessary to get between these two curves. For short curve γ on X, we must cross this fan to from γ to Y, and hence γ must be twisted w_n times around α if it wants to become a short curve on Y.

There is a direct analogy in higher genus, which is that for any curves α, β and γ, we can defined the twisting of β about α relative to γ. It turns out that there is another way of defining something analogous to rotation number which is more inductive in nature and ends up being more useful in the construction of the model manifolds. One looks at the subsurface $Z = S - \gamma$ and the sets $\alpha_Z = \alpha \cap Z$ and $\beta_Z = \beta \cap Z$. α_Z and β_Z will both be a collection of arcs. We can then take a single strand α'' from α_Z and produce a new curve α' by taking an essential component of the boundary of a closed regular neighborhood of α''.

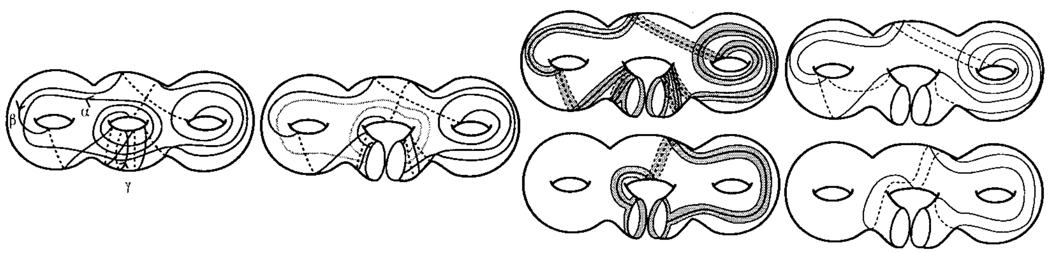

After doing the same to β to get a curve β', we can measure how far apart α' and β' are in the curve complex of Z. This is called a subsurface coefficient, and these play the same role in the general theory as the rotation numbers played in the discussion for the punctured torus.

This exhibits an essential feature of the discussion for higher genus. Not only does one have to study the curve complex of the surface itself, but one must also understand the curve complex of all its subsurfaces and how they fit together.