The curve complex and its relatives

Talk by Moon Duchin

August 16, 2007

•How to think about Teichmüller space (and why to think about the curve complex)

As we have seen in the previous lecture, Teichmüller space $\mathcal{T}(S)$ for a surface of finite type $S = S_{g,n}$ (genus g and n punctures or boundary components) is a parameter space for various types of structures that can be imposed on the same topological surface S. I'll adopt the view that $\mathcal{T}(S)$ is the space of metrics on S, marked with a set of generators for the fundamental group, up to (measurable) conformal equivalence: in that way, each point represents a conformal class of metrics. Let us use $\xi = \xi(S) = 3g - 3 + n$ as a measure of the topological *complexity* of S; then when $\xi \geq 1$, there are hyperbolic metrics on S. Every conformal class contains many singular flat metrics and, by uniformization, exactly one *Poincaré metric* (constant curvature -1). Thus we can also identify points of $\mathcal{T}(S)$ with their hyperbolic representatives.

The Teichmüller metric, as we saw before, is defined by the minimum quasi-conformal constants for maps between two such structures. (Recall that a *conformal* map takes circles to circles, infinitesimally. A *quasi-conformal* map takes circles to ellipses with a bounded ratio of axes.) This gives us a very straightforward description of geodesics in terms of flat structures. For instance, consider the flat structure obtained when opposite sides of a regular octagon are identified by translation: it is a metric on $S_{2,0}$ which is isometrically Euclidean except in a neighborhood of its unique singular point (a cone point with excess angle). With respect to the quasi-conformal constants, an efficient way to deform this octagon metric is to stretch in the horizontal direction, while compressing in the vertical direction by the same multiplicative factor. This gives a family of different metrics along a geodesic in $\mathcal{T}(S_{2,0})$, but because uniformization is something of a black box, we lack a correspondingly nice description of how the hyperbolic metrics vary along this trajectory. So a host of questions about Teichmüller geometry can be interpreted in this way: what is the long-range behavior of the hyperbolic metrics along this trajectory? Given two conformally equivalent flat structures, if you apply this stretch-and-shrink geodesic flow to each, how quickly do they diverge in Teichmüller distance? And so on.

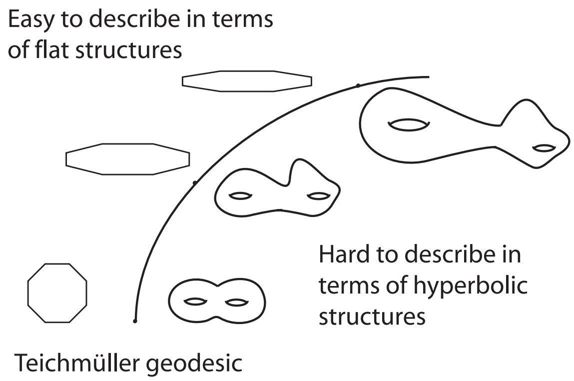

As we'll see, a good way to access these kinds of questions is by careful bookkeeping of $\mathcal{S}$, the set of curves on S. Recall from the last lecture that the Thurston compactification of $\mathcal{T}(S)$ is given by compatible embeddings $\mathcal{T} \hookrightarrow \mathbb{PR}^{\mathcal{S}}$ (sending curves to their hyperbolic lengths in the Poincaré metric) and $\mathcal{PMF} \hookrightarrow \mathbb{PR}^{\mathcal{S}}$ (sending curves to their weights in the measured foliation) so that $\mathcal{PMF}$ compactifies $\mathcal{T}$. So basically, we know where we are in the Teichmüller space, or even in its boundary at infinity, if we have kept track of the length/weight status of all curves. Below, we will build a combinatorial gadget $\mathcal{C}$ to do just that. We'll make repeated use of the fundamental fact that two intersecting curves can't be short at the same time in a metric of negative curvature: the crucial data in our object $\mathcal{C}$ will therefore be the intersection patterns of curves on the surface.[1]

• Coordinates on Teichmüller space

A maximal collection of disjoint, essential, non-peripheral simple closed curves on a surface S, no two of which are freely homotopic, is called a *pants decomposition.* There are $\xi(S)$ curves in a pants decomposition.

It is a fact of hyperbolic geometry (easy to check via right-angled hexagons in $\mathbb{H}$) that a Poincaré metric on a pair of pants is determined by the lengths of its three boundary curves. So given a pants decomposition of a surface of genus $g \geq 2$, in order to specify a hyperbolic metric on it we need only to give the lengths of the boundary components of the pairs of pants, along with twisting parameters to show how they are glued together at the "cuffs." This information gives the Fenchel-Nielsen coordinates for the hyperbolic structure on the surface. Teichmüller space $\mathcal{T}(S)$ therefore has real dimension 2ξ.

As an alternative to specifying twist parameters, given a pants decomposition $\{\alpha_1, \alpha_2, ..., \alpha_\xi\}$,

[1] I've motivated the curve complex as a way to study Teichmüller geometry, but it is also a way to *replace* Teichmüller geometry with combinatorics in the study of 3-manifolds via Kleinian groups or Heegaard splittings.

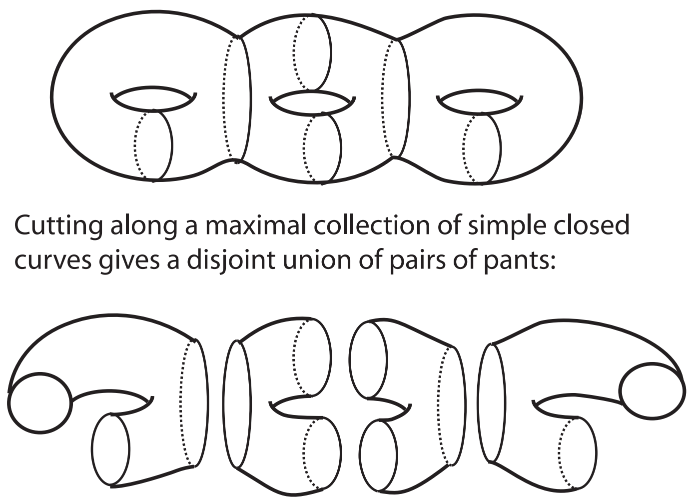

we may choose appropriate curves $\{\beta_1, \beta_2, ..., \beta_\xi\}$, with β_i transverse to α_i, (intersecting no other pants curves, and intersecting other transversals minimally) and specify the lengths of each β_i. This also determines the hyperbolic structure on the surface. The collection $\{\alpha_1, \beta_1, ..., \alpha_\xi, \beta_\xi\}$ is called a *marking* (a complete, clean marking in the sense of Masur-Minsky).

• Notions of coarse geometry and the definition of the curve complex

A *quasi-isometry* between metric spaces is a map $f : (X, d_1) \to (Y, d_2)$ such that for some constants K and c and all $x_1, x_2 \in X$,

$$\frac{1}{K} d_1(x_1, x_2) - c \leq d_2(f(x_1), f(x_2)) \leq K d_1(x_1, x_2) + c,$$

and whose image is metrically almost everything. A *coarse Lipschitz* map is one that satisfies the right-hand side of the above inequality (so that distance can decrease under f but not increase by much).

Let $\mathcal{S}$ be the set of free homotopy classes of essential, non-peripheral simple closed curves on a given surface. The curve complex $\mathcal{C}$, introduced by Harvey in 1978, is a simplicial complex that encodes intersections in $\mathcal{S}$. The vertices of $\mathcal{C}$ are elements of $\mathcal{S}$, and $\gamma_1, \gamma_2, ..., \gamma_n$ span an n-simplex if they can be realized disjointly. A maximal dimensional simplex in the curve complex corresponds to a pants decomposition. The combinatorics of the 1-skeleton gives $\mathcal{C}$ the structure of a metric space: the distance $d_{\mathcal{C}}(\alpha, \beta)$ of two curves is the smallest number of edges that must be traversed in $\mathcal{C}$ to connect them.

Maximal collection of simple closed curves:

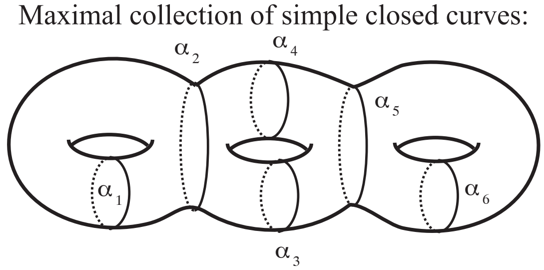

Collection of transversals for the above curves:

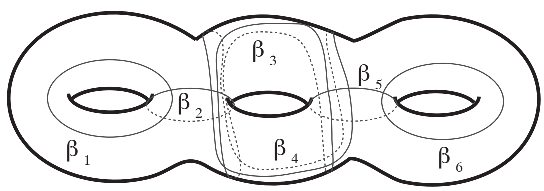

Marking of the surface:

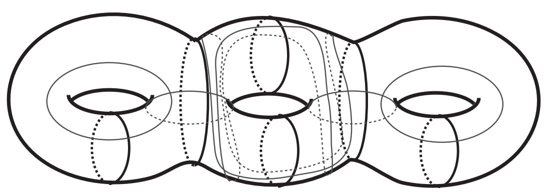

In low complexity, $\mathcal{C}$ is trivial, so one typically uses a modified complex $\mathcal{C}'$, which I will describe in certain cases later. First, it's useful to consider a second complex $\mathcal{CA}$, the complex of curves and arcs, which has as vertices not only curves but also (properly embedded) arcs with endpoints on the boundary of the surface; edges still connect vertices with disjoint representatives. In high complexity ($\xi \geq 3$), the curve complex and the complex of curves and arcs are quasi-isometric, so have the same coarse geometry.

As an example, we can work out what $\mathcal{CA}$ looks like for the punctured torus. Each rational point on the unit circle corresponds to a simple closed curve on the torus, which can be represented in a given homology basis by a pair (p, q) for relatively prime (p, q). Two curves $\alpha = (p, q)$ and $\beta = (r, s)$ intersect once iff $\begin{pmatrix} p & r \\ q & s \end{pmatrix}$ has determinant ± 1. Notice that on the punctured torus, for every simple closed curve there is exactly one simple arc disjoint from it. Let α' be the arc disjoint from α, β' the arc disjoint from β. It is easy to verify that α' and β' are disjoint exactly when α and β intersect once. So $\mathcal{CA}(S_{1,1})$ is coarsely equal to the Farey graph, where every Farey edge is decorated $\alpha - \alpha' - \beta' - \beta$.

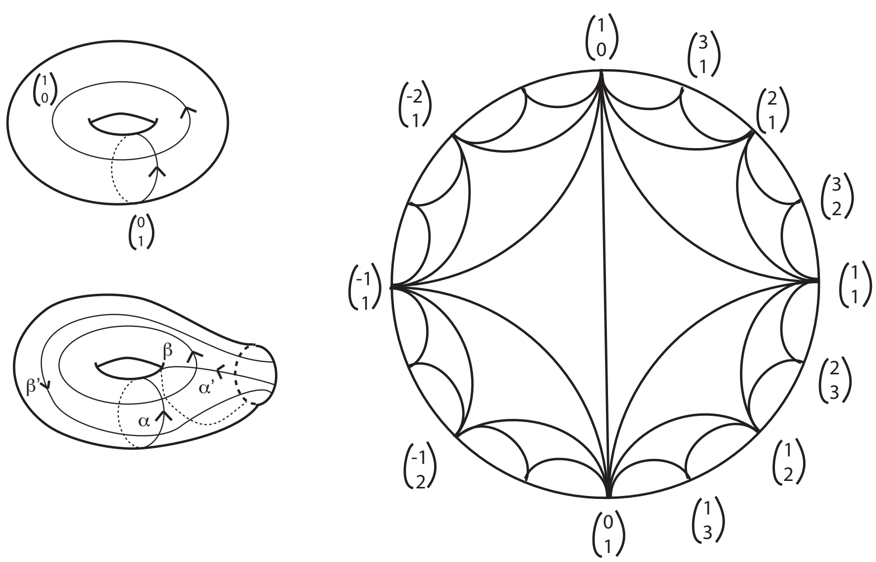

In very low complexity, we modify the definition. For the annulus, we let the vertices of $\mathcal{C}'$ be arcs considered rel boundary (otherwise, Dehn twists would not be recorded). To compute $\mathcal{C}'$, observe that picking a point on each boundary circle and a number of times to wrap around the annulus specifies an element. Setwise this looks like $S^1 \times S^1 \times \mathbb{Z} = T^2 \times \mathbb{Z}$ along with some identifications and some adjacencies between neighboring copies of the torus. Thus $\mathcal{C}'$ is quasi-isometric to the real line.[2]

[2] So, for those keeping track of all the cases: the curve complex of the annulus is coarsely $\mathbb{R}$ and for the

Note that the curve complex is (hideously) locally infinite: many isotopy classes of curves can be realized disjointly from a given curve (already true of the Farey graph), and there are in general infinitely many ways to get between two curves at distance two (in higher complexity).

• THE MAPPING CLASS GROUP AND BASIC PROPERTIES OF THE CURVE COMPLEX

The mapping class group, $Mod(S) = \text{Diff}^+(S)/\text{Diff}_0(S)$, should be thought of as the discrete part of the group of diffeomorphisms; in the torus case, it is $SL_2(\mathbb{Z})$. $Mod(S)$ can be shown to be generated by Dehn twists (to do a *Dehn twist* about α, cut the surface along α, twist by 2π, and reglue), plus some half Dehn twists in the boundary case. Thurston gave a classification of the elements of $Mod(S)$ by showing that an element is either finite order, reducible, or so-called "pseudo-Anosov".

Pseudo-Anosovs generalize Anosov maps on the torus to higher complexity. The most basic Anosov map on the torus is given by composing a Dehn twist about the $\binom{1}{0}$ curve with a reverse Dehn twist about the $\binom{0}{1}$ curve, obtaining the transformation $\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$. This map exhibits North-South dynamics on the Poincaré disk, i.e., under iteration it has one attracting fixed point and one repelling fixed point on the boundary, and an axis between them along which elements are translated. I'll define a *pseudo-Anosov* map on S as one which exhibits North-South dynamics on Teichmüller space and (consequently) on the curve complex. Thurston constructed examples of such maps just as above, by taking $T_\alpha^{-1}T_\beta$, where T_α and T_β are Dehn twists, and α and β are curves that fill the surface (two curves *fill* if every element of $\mathcal{S}$ must intersect at least one of them).[3] (See the figure for an example of filling curves on $S_{2,0}$.) This might seem very special, but in fact it's generic: for many appropriate ways of measuring, almost every mapping class is pseudo-Anosov.

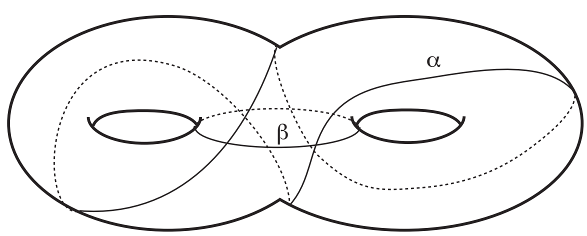

The action of the mapping class group permutes $\mathcal{S}$ and preserves disjointness, so it induces

torus it is declared the same as for the punctured torus (the Farey graph). A pair of pants has no curve complex. When $\xi \geq 1$, use $\mathcal{CA}$. This may seem ad hoc, but all of these definitions fit together correctly when considering projections to subsurfaces, a tool discussed briefly below.

[3]For α and β the longitude and meridian of the punctured torus, T_α alone is reducible (meaning it leaves a curve fixed, namely α), and you can check that the product $T_\alpha T_\beta$ is finite order.

a simplicial action on $\mathcal{C}$, which we can use to start to study the structure of $\mathcal{C}$. For instance, we've got a complex with infinitely many vertices and edges, and we'd like to see it's not just coarsely trivial: we should see why $\mathcal{C}$ has infinite diameter. Most curves we are likely to be able to visualize on a surface lie close together in the curve complex—it even takes a touch of work to find a pair at distance three. Two vertices are connected in $\mathcal{C}$ by a path of length two exactly when there exists a third curve disjoint from both, so in other words

$$\alpha, \beta \text{ fill } \iff d_{\mathcal{C}}(\alpha, \beta) \geq 3.$$

To see why the curve complex has infinite diameter, consider iterations of a pseudo-Anosov mapping class. Because a pseudo-Anosov has North-South dynamics, it has a foliation F as its attracting fixed point and it the curves along its axis in $\mathcal{C}$ are getting more like that complicated foliation (recall that § is dense in $\mathcal{PMF}$, so the curves are approaching the foliation in that topology). This suggests that the curves in this sequence intersect a fixed curve more and more, so that they leave every bounded subset of the curve complex.

We would also like to see that the curve complex is connected. Suppose α and β are curves that intersect many times on the surface. Form an arc α_0 by following α between two successive intersections with β; its endpoints divide β into two arcs β_1 and β_2. Complete α_0 to a simple closed curve α' by choosing the arc β_i that intersects α fewer times. Then $i(\alpha', \beta) \leq \frac{1}{2} i(\alpha, \beta)$, so our surgery on α has reduced the intersection number by half. Running this argument carefully, we can deduce something stronger than the connectedness of $\mathcal{C}$, namely:

$$d_{\mathcal{C}}(\alpha, \beta) \leq 2 \log_2 i(\alpha, \beta) + 2.$$

In their two papers on the geometry of the curve complex, Masur and Minsky introduce a suite of techniques for taming the unruly local infinitude of $\mathcal{C}$. For instance, they define *tight geodesics*, which have the property that there are only finitely many between a given two points. Also, they analyze the *hierarchy* structure given by considering all the subsurfaces of S: let Y be a subsurface of S with α as a boundary curve and observe that curves on Y are disjoint from α, so adjacent to it in $\mathcal{C}(S)$. Thus curve complexes of subsurfaces appear in links of vertices. We can get some control on the geometry of $\mathcal{C}(S)$ by projecting curves and arcs to $\mathcal{C}(Y)$ via intersection with Y. A great deal of useful information appears in these two Masur-Minsky papers, but let me highlight one especially important fact: the curve complex, with its combinatorial metric, is *delta-hyperbolic*, which means that, despite the infinite number of choices, fairly efficient paths in $\mathcal{C}$ between common endpoints can't be that different from each other.

• Relatives of the curve complex

Just as we built a curve complex $\mathcal{C}$ whose vertices are curves, we can build a marking graph $\mathcal{M}$ whose vertices are markings, or a pants complex $\mathcal{P}$ whose vertices are pants decompositions. We'll connect vertices of $\mathcal{M}$ or $\mathcal{P}$ when the objects differ minimally: by an *elementary move* (which we won't describe here).

We get a coarsely commutative diagram featuring the marking graph, Teichmüller space and the curve complex.

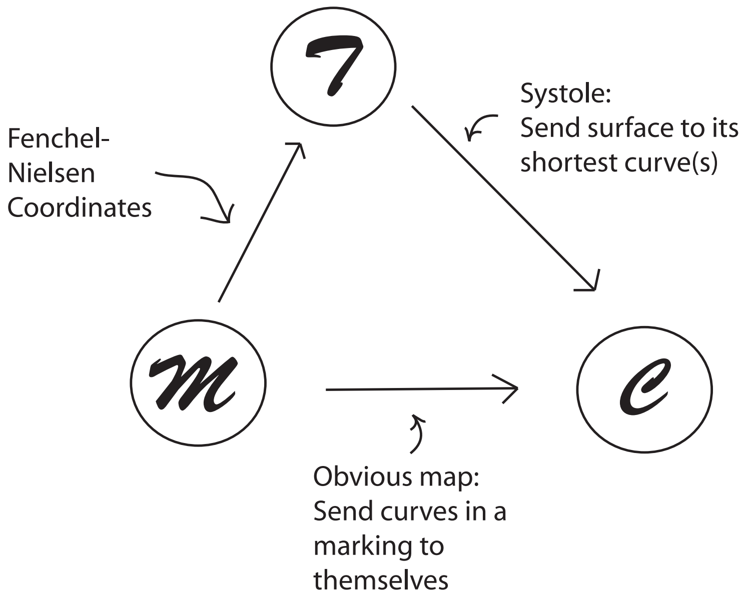

I propose the view that these different combinatorial models describe different parts of Teichmüller space. $\mathcal{M}$ can be thought of as a model for the thick part (those hyperbolic metrics with the injectivity radius bounded below by a constant ϵ), while $\mathcal{C}$ describes the thin part (those hyperbolic metrics with some curve shorter than ϵ). More specifically, $\mathcal{C}$ records the intersection patterns of the thin strata, as intersecting curves cannot be simultaneously short by the collar lemma. (So $Thin_\alpha$ and $Thin_\beta$ overlap if and only if α,β are adjacent in $\mathcal{C}$.) On the other hand, the FN map in the diagram (build a surface from a marking by taking Fenchel-Nielsen coordinates giving all curves the same length) misses almost all of $\mathcal{T}$, but if instead we took a marking to the region in $\mathcal{T}$ where it is shortest, we would cover all of the thick part of $\mathcal{T}$ with sets of bounded size.

The first diagram sits as a piece of the one below, showing how several coarse models of Teichmüller space are related. It includes the facts that the marking graph turns out to be quasi-isometric to the mapping class group with the word metric, and Brock's theorem that Teichmüller space with the Weil-Petersson metric is quasi-isometric to the pants complex. In the diagram, each downward-pointing arrow is a coarse Lipschitz map.

As a closing example, to illustrate these successive collapses between these metrics, consider the effect of a Dehn twist at various levels in this diagram. Iterating a Dehn twist moves

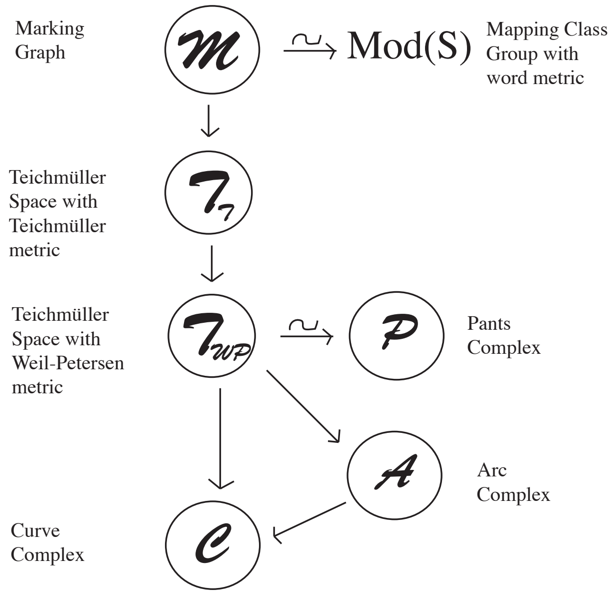

a point linearly in the mapping class group, logarithmically in the Teichmüller metric and not at all in the curve complex. The marking graph (or mapping class group) is the "biggest" metric space; it has many quasi-flats (think of subgroups $\mathbb{Z}^k$ generated by disjoint Dehn twists). At each stage more of the metric space is being collapsed out, until at the bottom we arrive at the curve complex, where all the flats are gone and the metric is finally delta-hyperbolic.

References:

- Howard Masur and Yair Minsky, *Geometry of the complex of curves. I. Hyperbolicity.*
- Howard Masur and Yair Minsky, *Geometry of the complex of curves. II. Hierarchical structure.*
- Saul Schleimer, *Notes on the complex of curves*