Curves, Surfaces, and Hyperbolic 3-Manifolds

Andrew Yarmola

Abstract These notes follow a mini-course by Yair Minsky at the GEAR Junior Retreat in 2014. Our focus is to explore the relationship between the geometry of 3-manifolds and the combinatorics of an embedded surface. The fundamental example will be a 3-manifold fibering over a circle. The geometry of such manifolds is deeply related to the monodromy map of the fiber. We use this example as a stepping stone to work towards a model for the geometry of a 3-manifold mediated by the combinatorics of the mapping class group and the curve complex of an embedded surface.

1 Introduction

Let M be a 3-manifolds that admits a complete hyperbolic structure and Σ be a compact surface of negative Euler characteristic. Recall that a complete hyperbolic structure on M is a complete Riemannian metric of constant sectional curvature -1. In particular, we can identify the universal cover $\tilde{M}$ with $\mathbb{H}^3$, hyperbolic 3-space, and $M \cong \mathbb{H}^3/\Gamma$ for some discrete group $\Gamma \subset \mathrm{Isom}^+(\mathbb{H}^3) \cong \mathrm{PSL}(2,\mathbb{C})$ isomorphic to $\pi_1(M)$. Given a π_1-injective embedding $f : \Sigma \to M$, we would like to study the connection between combinatorics of the surface and the hyperbolic structure on M.

By "combinatoritcs" of a surface, we mean two objects: the mapping class group and the curve graph of Σ. In the first half of these notes, we will work with the mapping class group to build fundamental examples of these manifolds. In the second half, we will study the relationship between laminations and the curve graph on Σ to get a handle on the geometry of M.

Andrew Yarmola
Boston College, Department of Mathematics, Chestnut Hill, MA 02467, e-mail: yarmola@bc.edu

2 Fibering over a circle

Let $M \to S^1$ be an orientable smooth 3-manifold (possibly with boundary) fibered over a circle with fiber Σ, a compact connected surface of negative Euler characteristic. If α generates $\pi_1(S^1)$, we get a monodromy map $f : \Sigma \to \Sigma$ associated to α. One can recover M from this information as the mapping torus $M_f \cong M$ of f.

$$M_f = \Sigma \times \mathbb{R}/(x,t) \sim (f(x), t+1)$$

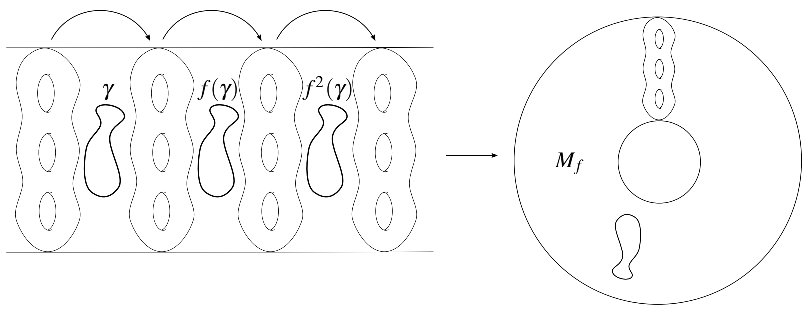

Fig. 1 Cyclic cover of a mapping torus.

This identification of M is not unique. M could fiber in a different way over S^1. Indeed, if $\pi : M_f \to S^1$ is the standard fibering, we can consider the form $\pi^* d\theta$ on M_f, where $d\theta$ is the length element on S^1. Notice that the kernel of $\pi^* d\theta$ is the fiber Σ of π. Now, if we perturb $\pi^* d\theta$, we get a new direction for the kernel, giving us a new fiber. In fact, if $\dim H_1(M, \mathbb{R}) > 1$, then there are infinitely many distinct ways for M to fiber over a circle. In [16], Thurston defines a (pseudo-)norm on $H_2(M, \mathbb{R})$ that allows him to pick out the integral homology classes that realize all the possible fibers.

Remark 1. In [9], Allen Hatcher provides another construction in the case that $M = \Sigma \times S^1$ and $\pi : M \to S^1$ is the standard projection. Pick distinct points $x_i \in S^1$ for $i = 1, \ldots n$ and consider the fibers $F_i = \pi^{-1}(x_i)$. Fix a non-separating simple closed curve C on Σ and consider the torus $T = C \times S^1$. Cut F_i along T and re-glue the resulting surfaces so that F_i connects to F_{i+1} when it crosses T, taking the subscripts mod n. One obtains a connected n-sheeted cover of Σ as a fiber of a new fibering of M over S^1. Notice that the monodromy of this new fibering will be periodic of order n.

We would like to understand what conditions are necessary on the pair (Σ, f) for the interior of M_f to admit a complete hyperbolic structure. To do this, we will fix Σ and understand the properties of f. Since $M_f \cong M_{f'}$ whenever f and f' are isotopic, we will think of f as an element of $\mathrm{Mod}(\Sigma)$, the mapping class group of Σ. Recall the following definitions.

Definition 1. For a smooth surface Σ, we define the *mapping class group*

$$\mathrm{Mod}(\Sigma) = \mathrm{Diff}^+(\Sigma, \partial\Sigma)/\mathrm{Diff}_0(\Sigma, \partial\Sigma)$$

i.e. the set of all orientation preserving diffeomorphism that restrict to the identity on $\partial\Sigma$, up to boundary preserving isotopy.

Definition 2. $f \in \mathrm{Mod}(\Sigma)$ is *reducible* if there exists a finite collection of disjoint, homotopically nontrivial, non-peripheral simple closed curves $\mathcal{C}$ such that $f^n(\mathcal{C}) \sim \mathcal{C}$ for some $n > 0$.

Definition 3. $f \in \mathrm{Mod}(\Sigma)$ is *pseudo-Anosov* if it is irreducible.

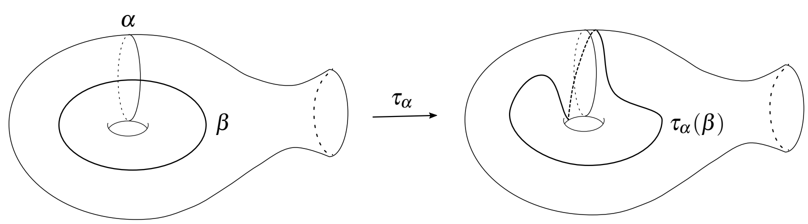

Fig. 2 A Dehn twist locally twists every transverse arc once around α to the right. It is reducible as it fixes the curve α.

With this notation in mind, we obtain a definitive answer to our question.

Theorem 1. *(Thurston, [17]) If Σ is a compact connected surface of negative Euler characteristic and $f \in Mod(\Sigma)$, the interior of M_f admits a complete hyperbolic structure if and only if f is pseudo-Anosov. By Mostow rigidity, this structure is unique up to isometry.*

The fact that f *must* be pseudo-Anosov is clear from the following observation. If f is reducible, then for some $n > 0$, f^n fixes (up to isotopy) a non-peripheral simple closed curve C. The cylinder $C \times \mathbb{R} \subset \Sigma \times \mathbb{R}$ projects to an essential (i.e. π_1-injective) torus in M_f, which is not possible if the interior of M_f admits a complete hyperbolic structure.

Exercise 1. Show that a closed hyperbolic 3-manifold cannot contain an essential torus.

Another remark worth making is that one can build a singular Euclidean metric on M_f using the following fact about pseudo-Anosov maps.

Theorem 2. *(Nielsen-Thurston classification of pseudo-Anosov maps, [6]) Given a pseudo-Anosov $f \in Mod(\Sigma)$, $g \geq 2$, there exist transverse singular measured foliations (F_s, μ_s) and (F_u, μ_u) on Σ, and a real number $\lambda > 1$ so that*

$$f \cdot (F_u, \mu_u) = (F_u, \lambda \mu_u) \quad \textit{and} \quad f \cdot (F_s, \mu_s) = (F_s, \lambda^{-1} \mu_s)$$

We refer the reader to [6] for a complete definition of a measured foliation. Below is a local picture to keep in mind, where f acts by "stretching" (F_u, μ_u) and "contracting" (F_s, μ_s). The foliations are called unstable and stable, respectively.

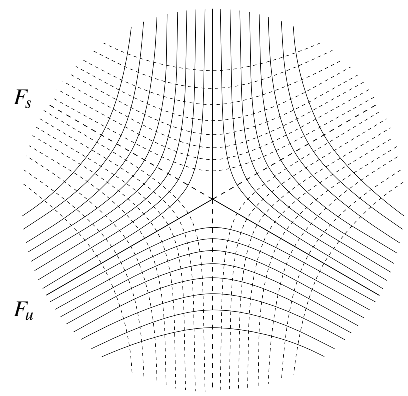

Fig. 3 Local picture of an order 3 singularity of two transverse foliations.

These traverse singular measured foliations give rise to a singular Euclidean metric on Σ, where the cone points in the interior of Σ have angle bigger than 2π. The leaves of the foliations are geodesic in this metric. Since f topologically fixes the foliations, one gets a singular Euclidean metric on M_f with all the negative curvature concentrated at the singularities. The product of the leaves with S^1 gives geodesic subsurfaces. Unfortunately, this construction does not guarantee the existence of a hyperbolic metric.

The actual proof of Theorem 1 uses a limiting procedure by iterating f to build a representation $\rho : \pi_1(\Sigma) \to \mathrm{PSL}(2,\mathbb{C})$ on which f acts by isometry. We explore a similar picture in the next section.

Our last remark on this fundamental example is that every finite volume hyperbolic 3-manifold has a finite index cover isometric to some M_f. This is the Virtually Fibered Theorem, a deep result that follows from the work of Agol, Kahn, Markovic, Wise, and others, see [1].

3 Iterating mapping class group elements

In this section we would like to understand what M_{f^n} and $M_{f\circ g^n}$ look like as $n \to \infty$ for $f, g \in \mathrm{Mod}(\Sigma)$. Let us assume that f is pseudo-Anosov and let g be a Dehn twist or another nice reducible element of $\mathrm{Mod}(\Sigma)$. We will also assume that Σ is *closed*. The reader may replace mapping tori with their interiors to recover the compact surface case.

3.1 Iterating psuedo-Anosovs

For our first example, we have an n-fold cover $M_{f^n} \to M_f$ on which f acts as an order n isometry. The hyperbolic structure on these covers comes from lifting the unique hyperbolic structure on M_f. We can think of the $\mathbb{Z}$-cover $\hat{M}_f \cong \Sigma \times \mathbb{R}$ as a limit of M_{f^n} as $n \to \infty$.

The $\mathbb{Z}$-cover $\hat{M}_f$ carries more structure than just its homeomorphism type, in particular, f acts on it by isometry. This means that for a simple closed geodesic $\gamma \in \hat{M}_f$, the simple closed geodesics $f^n(\gamma)$ leave every compact set as $n \to \pm\infty$. In fact, the curves $f^n(\gamma)$ as $n \to \pm\infty$ limit to two laminations (see below), those associated to the stable and unstable foliations of f. These type of laminations are called ending laminations.

$$
\begin{array}{ccccc}
\hat{M}_f & \longrightarrow & M_{f^n} & \longrightarrow & M_f \\
\downarrow & & \downarrow & & \downarrow \\
\mathbb{R} & \longrightarrow & S^1 & \xrightarrow{\cdot n} & S^1
\end{array}
$$

3.2 Geometric Limits

The notion of a limit above can be made precise. For this, we need to consider hyperbolic manifolds with baseframes (M_n, ω_n), where ω_n is an orthonormal frame at some point $p_n \in M_n$.

Definition 4. We say that $\{(M_n, \omega_n)\}$ converges *geometrically* to (M, ω) if one can find smooth embeddings $\phi_n : U_n \to M_n$ with $p \in U_n \subset M$ and $d\phi_n(\omega) = \omega_n$, with the property that for any $\varepsilon > 0$ and $r > 0$, there exists L such that for all $n > L$, $B_r(p) \subset U_n$ and ϕ_n is ε-close to an isometry in C^2. The maps ϕ_n are called *comparison maps*.

From the heuristic image on $\hat{M}_f$ above, it is easy to imagine choosing appropriate base frames on M_{f^n} and $\hat{M}_f$ along with the sets $U_n \in \hat{M}_f$ and maps $\phi_n : U_n \to M_{f^n}$. One can show that in this case, the choice of convergent baseframes does not matter.

Fix the pair $(\mathbb{H}^3, \eta)$ by choosing some orthonormal frame η. Then for (M_n, ω_n) there is a unique covering map $(\mathbb{H}^3, \eta) \to (M_n, \omega_n)$ sending η to ω_n. This gives us a representation $\rho_n : \pi_1(M_n, p_n) \to \mathrm{PSL}(2, \mathbb{C})$ with image $\Gamma_n \leq \mathrm{PSL}(2, \mathbb{C})$. Taking geometric limits as defined above is the same as taking limits of discrete subgroups Γ_n in the Hausdorff/Chabauty topology on closed subsets of $\mathrm{PSL}(2, \mathbb{C})$, see [5]. This, however, is different from taking algebraic limits in the space of representations, see [3] for some examples.

3.3 Iterating Dehn twists

Our next example are the manifolds $M_{f\circ\tau^n}$ where f is pseudo-Anosov and τ is a Dehn twist around an essential curve γ on Σ. See Section 2 for an image of the action of τ. Since a Dehn twist is a local operation, we pick a normal neighborhood $N_\gamma \subset \Sigma$ of γ. We visualize $M_{f\circ\tau^n}$ as $\Sigma\times[0,1]\setminus N_\gamma\times[1/3,2/3]$ where we glue in the solid torus $N_\gamma\times[1/3,2/3]$ by τ^n on $N\times\{2/3\}$ and identity everywhere else. We then identify $\Sigma\times\{0\}$ with $\Sigma\times\{1\}$ via f to get $M_{f\circ\tau^n}$.

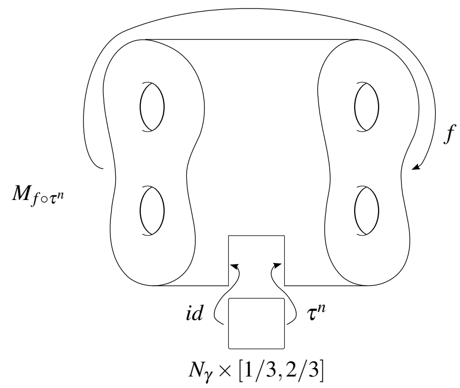

Fig. 4 The gluing description of $M_{f\circ\tau^n}$.

To think about the geometric limit, we want to confirm that $M_{f\circ\tau^n}$ admits a complete hyperbolic metric for all but finitely many n. It is a theorem of Fathi [7] that $f\circ\tau^n$ is pseudo-Anosov except for at most 7 consecutive values of n. The geometric limit $N_{f,\tau}$ of $M_{f\circ\tau^n}$ is then a finite volume hyperbolic manifold homeomorphic to $M_f\setminus\gamma\times\{1/2\}$. $N_{f,\tau}$ has exactly one (rank 2) cusp corresponding to the "drilled" curve γ.

Note that we can recover $M_{f\circ\tau^n}$ from $N_{f,\tau}$ by the process of Dehn filling. Let N be a finite volume hyperbolic manifold with one cusp. We can topologically compactify N so that its boundary is a torus. Then, there is a neighborhood $V\cong(D_2\setminus D_1)\times S^1$ of of the boundary, where $D_1\subset D_2$ are disks. The curves $\lambda=\{1\}\times S^1$ and $\mu=\partial D_2\times\{1\}$ are the *longitude* and *meridian*, respectively. A p/q Dehn filling $N_{p/q}$ of N corresponds to gluing in a solid torus so that its meridian attaches to the curve $p\mu+q\gamma$.

Theorem 3 (Thurson, [17]). *Given a finite volume hyperbolic manifold N with one cusp, for all but finitely many $p/q\in\mathbb{Q}\cup\{\infty\}$ there exists a hyperbolic structure on $N_{p/q}$*

In our case, $M_{f\circ\tau^n}$ corresponds to an $1/n$ Dehn filling of $N_{f,\tau}$. In particular, for Theorem 3 above, we see that $N_{p/q}\to N$ geometrically as $p/q\to 0$.

Exercise 2. Prove that $M_f\setminus\gamma\times\{1/2\}$ is Haken.

Exercise 3. Show that $f\circ\tau^n$ is pseudo-Anosov for all but finitely many n.

3.4 Iterating partial pseudo-Anosovs

Let $g \in \mathrm{Mod}(\Sigma)$ be reducible such that there is a compact subsurface $W \subset \Sigma$ with $g|_W$ pseudo-Anosov and $g|_{\Sigma \setminus W} = id$. One way to construct such mapping classes is take two simple closed curves α, β that fill a subsurface W and let $g = \tau_\alpha \circ \tau_\beta$. Then g is pseudo-Anosov on W and identity outside. We turn our attention to the manifolds $M_{f \circ g^n}$.

As before, our topological picture of $M_{f \circ g^n}$ is obtained by starting with $\Sigma \times [0,1] \setminus W \times [1/3, 2/3]$ and gluing in $W \times [1/3, 2/3]$ via g^n on $W \times \{2/3\}$ and identity outside. We then identify $\Sigma \times \{0\}$ with $\Sigma \times \{1\}$ via f to get $M_{f \circ g^n}$.

Exercise 4. Show that for all but finitely many n, $f \circ g^n$ is pseudo-Ansov.

We can understand the geometry of the limit $N_{f,g}$ of $M_{f \circ g^n}$ by looking into the region $W \times [1/3, 2/3]$. In that region the geometry looks more and more like that of the mapping torus W_{g^n} as $n \to \infty$. In particular, as we go from left to right through $W \times [1/3, 2/3]$, the region looks more and more like the "right" end of $\hat{W}_g$. That is, geodesic curves are carried out by g^n for large $n > 0$. As we go from right to left, the picture looks more like the "left" end of $\hat{W}_g$. The curves on ∂W become rank 1 cusps in $N_{f,g}$. See Figure 5 for for the picture and [3] for more details on this example.

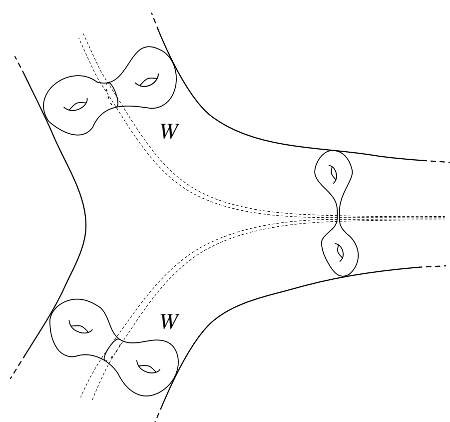

Fig. 5 The geometry of $N_{f,g}$.

Remark 2. Given a factorization of a mapping class $\psi = f \circ g_1^{n_1} \circ \ldots \circ g_k^{n_k}$, the above examples provide a picture of M_ψ localized to parts on $\Sigma \times [0,1]$. Since we can partially recover M_ψ by surgeries or quotients on the cusps or ends, we may hope that there is some relationship between factorization and the geometric decomposition of M_ψ. Unfortunately, factoring mapping class group elements in not simple.

4 Pleated surfaces

Our next goal is to find geometrically interesting embeddings in a homotopy class of π_1-injective maps $[h : \Sigma \to N]$, where N is a complete hyperbolic 3-manifold. For simplicity, we assume in this section that Σ is closed. Keep in mind that there is a lift $\hat{h} : \Sigma \to \hat{N}$ of h such that $\hat{h}$ is a homotopy equivalence. As we saw earlier, the geometry of a mapping torus M_f as better understood by such a cover.

One approach to pick out interesting members of $[h]$ is to use dynamics. Fix a Riemannian structure on S and let $h_E : \Sigma \to N$ be the energy minimizing map. Here, *energy* is the integral of $|dh|^2$. In fact, the energy minimizing map is attained and harmonic, see [12]. We will, however, focus on a more geometric construction.

4.1 Almost pleated surfaces

Fix a point $x \in \Sigma$ and a loop α based at x. Let T be a triangulation of Σ whose only vertex is x and α an edge. Notice that every edge is a loop based at x. See Figure 7.

Exercise 5. (a) How many edges does T have? (b) How many triangles are in T?

Given a class $[h]$, we construct an embedding $h_T : \Sigma \to N$. Let α^* be the totally geodesic representative of $h(\alpha)$ and x^* a point on α^*. For every other edge $e \in T$, there is a geodesic arc e^* homotopic to $h(e)$ such that e^* is a loop at x^*. Notice that e^* might have a corner at x^*. For any three edges e_1, e_2, e_3 that bound a triangle, there is a unique hyperbolic triangle in N with edges e_1^*, e_2^*, e_3^*. The map h_T takes $x \mapsto x^*$, every edge $e \mapsto e^*$ and every triangle to its corresponding hyperbolic one.

The image $h_T(\Sigma)$ is totally geodesic away from the 1-skeleton. Further, we can define a path metric to Σ by pulling back via h_T. This metric will be hyperbolic away from x. At x, we will have a cone angle of $\geq 2\pi$. Indeed, since α^* is totally geodesic, we get antipodal points at x^* and so the total angle at x has to be $\geq 2\pi$.

Exercise 6. Show that $\text{area}(h_T(\Sigma)) \leq 2\pi|\chi(\Sigma)|$.

Since the area of the image is controlled by the topology of the surface, we can see that the length of the shortest curve on $h_T(\Sigma)$ is bounded by some constant $L(\chi(\Sigma))$. Indeed, if $p \in h_T(\Sigma)$ is a point of the shortest curve, then balls centered at p in N must intersect for large radii because $h_T(\Sigma)$ has finite area.

If T, T' are two triangulations of the above type sharing the same vertex x, then we can *interpolate* between the (almost) pleated surfaces $h_T, h_{T'}$.

Theorem 4 (Hatcher, [8]). *The triangulations T and T' of the type above for a closed surface Σ are connected by elementary moves of the form seen in Figure 6.*

The theorem gives a sequence of triangulations $T = T_0, \ldots, T_k = T'$ where each step differs by an elementary move. Such a move defines a hyperbolic tetrahedron in N. By a careful analysis of this situation, Canary show that there is a homotopy H_t between h_{T_i} and $h_{T_{i+1}}$, such that H_t is an (almost) pleated surface with at most two vertices for each t. See [4] for details.

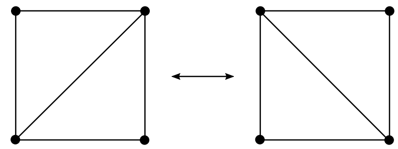

Fig. 6 An elementary move for triangulations.

4.2 Pleated surfaces

To get a pleated surface from the above construction, we would like to force the cone angle at x^* to be 2π. As before, we want some map h_λ in the homotopy class $[h]$ whose image is totally geodesic outside of some codimenension 1 locus. Instead of the 1-skeleton of a triangulation, we will want to use a *lamination*.

Definition 5. A *lamination* on a smooth closed surface Σ is a closed subset of Σ foliated by smooth simple curves. These simple curves are called *leaves*.

Definition 6. A *pleated surface* is a continuous map $h : \Sigma \to N$ whose image is totally geodesic in the complement of a lamination on Σ and the leaves of the lamination are mapped to geodesics. We also require that every smooth arc map to a rectifiable one.

How do we construct such a map? Well, we can modify the triangulation T in a controlled manner. Let τ_α be a Dehn twist around α and define $T_n = \tau_\alpha^n(T)$. To picture this, think of an edge going towards α and spiraling to the right n times around α before hitting x. Let λ be the image of the edges in $\lim_{n\to\infty} T_n$. Notice that the edges become simple curves that spiral on α.

Another variation of this construction is to drag x around α many times and bring the edges along with it, see Figure 7.

The sequence of maps h_{T_n} is anchored at α^*, so the images are all contained in some compact set $K \subset N$ by the area bound. It follows that the injectivity radius of all the $h_{T_n}(\Sigma)$ is bounded below by some $\varepsilon > 0$ and the Arezlá-Ascoli theorem gives a convergent subsequence of $h_{T_n} \to h_\lambda$, a pleated surface. The cone angle at x^* vanishes and surface $h_\lambda(\Sigma)$ is totally geodesic outside of the image of λ. Further, the path metric to Σ inherited from the metric on N by f_λ is hyperbolic.

4.3 Uniform Injectivity

Given a pleated surface $f : \Sigma \to N$, we endow Σ with the induced path metric σ. It follows that f maps every geodesic segment in (Σ, σ) to a rectifiable arc in N which has the same length. In particular, the associated lamination is foliated by geodesics on (Σ, σ). To better understand pleated surfaces $f : \Sigma \to N$, we would like to know how distances of points on (Σ, σ) differ from distances in N.

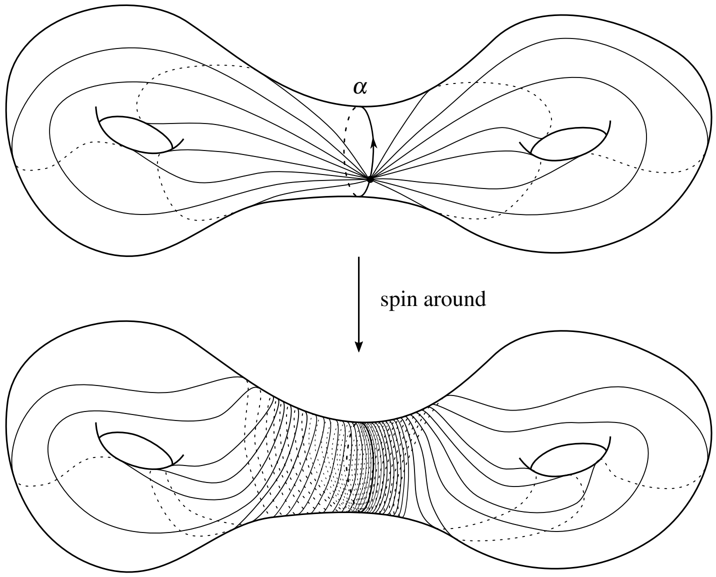

Fig. 7 Constructing a lamination by dragging the unique vertex of a triangulation around a simple curve α.

Definition 7. Given $\varepsilon > 0$ and a hyperbolic manifold M, the ε-thin part is $M_{<\varepsilon} = \{p \in M \mid \mathrm{inj}(p) \leq \varepsilon\}$, where $\mathrm{inj}(p)$ is the injectivity radius at p. The ε-tick part $M_{>\varepsilon}$ is just the complement.

Let $\mathbb{P}U(N)$ the the projectivized unit tangent bundle of N. Let λ be the lamination of f, define $t_f : \lambda \to \mathbb{P}U(N)$ by mapping $x \in \lambda$ to the unique (projective) tangent vector along λ. Thurston proves the following theorem.

Theorem 5 (Thurston, [15]). *Fix $\overline{\varepsilon} > 0$ and Σ, a smooth closed surface of genus at most $g \geq 2$. Suppose $f : \Sigma \to N$ is a pleated surface with lamination λ* **and** *f is a homotopy equivalence. Then for every $\varepsilon > 0$ there exists $\delta > 0$, depending only on $\varepsilon, \overline{\varepsilon}$ and g, such that for any $x, y \in \lambda$ in the $\overline{\varepsilon}$-think part of (Σ, σ)*

$$d_{\mathbb{P}U(N)}(t_f(x), t_f(y)) < \delta \implies d_\sigma(x, y) < \varepsilon$$

Proof. By contradiction, suppose there exists $\varepsilon > 0$ and a sequence $f_n : \Sigma \to N_n$ with points $x_n, y_n \in \lambda_n \cap (\Sigma, \sigma_n)_{>\overline{\varepsilon}}$ such that

$$d_{\mathbb{P}U(N_n)}(t_{f_n}(x_n), t_{f_n}(y_n)) < 1/n \quad \text{but} \quad d_{\sigma_n}(x_n, y_n) > \varepsilon$$

Our goal is to take a limit and extract a topological contradiction. Let $\Sigma_n = (\Sigma, \sigma_n)$ be the induced (marked) hyperbolic structures. There is a notion of geometric convergence in the setting of pleated surfaces: with frames chosen at x_n (or y_n), we want the comparison maps for the limits $\lim_n \Sigma_n$ and $\lim_n N_n$ to commute with f_n.

A sequence Σ_n of (marked) hyperbolic structures either converges or short curves develop. However, since x_n is in the $\overline{\varepsilon}$-thick part of Σ_n, one can show that the injectivity radius of $f_n(x_n)$ in N_n must be greater than some $\overline{\varepsilon}' > 0$ that depends only on $\overline{\varepsilon}$ and g. From this, we deduce that the space of possible surfaces and 3-manifolds with baseframes in the thick part is compact in the geometric convergence topology. Since f_n is 1-Lipschitz for all n, the Arezlá-Ascoli theorem gives a convergent subsequence.

Assume we converge to $f_\infty : \Sigma_\infty \to N_\infty$. Passing to further subsequences, we can assume that $x_n \to x_\infty, y_n \to y_\infty$ and the laminations of f_n converge so some λ_∞. Note, that since f_n are homotopy equivalences, f_∞ is at least π_1-injective. See [15] for details.

From the choice of sequence f_n, it follow that $t_{f_\infty}(x_\infty) = t_{f_\infty}(y_\infty)$ but $d_{\sigma_\infty}(x_\infty, y_\infty) > 0$. This means that the distinct leaves of λ_∞ at x_∞ and y_∞ are mapped to the same geodesic in N_∞. This implies that f_∞ is no longer π_1-injective, a contradiction. In fact, f_∞ factors through a non-trivial covering map. □

In [18], Thuston proves a more general version of the theorem.

5 Curve graph

Using pleated surfaces and uniform injectivity, we can try to model complete hyperbolic 3-manifolds N homotopy equivalent to a surface Σ. It is a theorem of Bonahon [2] that all such N are homeomorphic to $\Sigma \times \mathbb{R}$. A fundamental example of these the manifolds is $\hat{M}_f$, the $\mathbb{Z}$-cover the mapping torus of a pseudo-Anosov. Recall that $\hat{M}_f$ has infinitely many simple geodesic curves $f^n(\gamma)$ of the same length. Our goal in this section will be to study the geodesic curves in N of length $< L$.

Definition 8. Given a surface Σ, the *curve graph* $\mathcal{C}(\Sigma)$ is a graph where each simple closed curve on Σ is a vertex and $[\alpha, \beta]$ is an edge if α, β can be simultaneously realized as disjoint curves. $\mathcal{C}(\Sigma)$ carries the path metric where edges have length 1.

Let $h : \Sigma \to N$ be a homotopy equivalence and let $[\alpha, \beta]$ be an edge in $\mathcal{C}(\Sigma)$. Since α and β are disjoint on Σ, we can build a two vertex triangulation of Σ with the vertices on α and β. Taking the limit of almost pleated surfaces by Dehn twisting about α and β, we get a pleated surface that has both α^* and β^* as totally geodesic in N. From this, we see that paths in $\mathcal{C}(\Sigma)$ correspond to sequences of (interpolated) pleated surfaces.

Theorem 6 (Masur-Minsky, [11]). $\mathcal{C}(\Sigma)$ *is* δ*-hyperbolic.*

Definition 9. The space of ending laminations $\mathcal{EL}(\Sigma)$ is the collection of laminations whose complementary regions in Σ are ideal polygons or once-punctured ideal polygons.

Theorem 7 (Klarreich, [10]). *The Gromov boundary* $\partial_\infty \mathcal{C}(\Sigma)$ *is homeomorphic to* $\mathcal{EL}(\Sigma)$.

Given a homotopy equivalence $f : \Sigma \to N$, let

$$l_N(\alpha) = \text{ length of the geodesic homotopic to } f(\alpha) \text{ in } N$$

Define the set

$$\mathcal{C}_{f,L} = \{\alpha \in \mathcal{C} : l_N(\alpha) < L\}$$

This set will turn out to be quasiconvex. Recall that

Definition 10. A subset A of a geodesic metric space X is *K-quasiconvex* if for any $x, y \in A$, the geodesic between x, y lies in a K-neighborhood of A.

Theorem 8 (Minsky, [14]). *Let Σ be a closed smooth surface of genus g and $f : \Sigma \to N$ a homotopy equivalence. There exits a constant L_0, depending only on g, such that for all $L > L_0$ we can find a K so that $\mathcal{C}_{f,L}$ is K-quasiconvex.*

Proof. In a δ-hyperbolic space, a subset is quasiconvex if it admits a Lipschitz coarse retraction from the ambient space, see [13]. For $\alpha \in \mathcal{C}(\Sigma)$, let

$$\mathbf{pleat}(\alpha) = \{h : \Sigma \to N : h \sim f \text{ is a pleated surface with } h(\alpha) \text{ geodesic in } N\}$$

From the previous section, we know that this set is non empty. Our retraction map will be $\pi : \mathcal{C}(\Sigma) \to \mathcal{C}_{f,L}$ given by

$$\pi(\alpha) = \text{ any choice of shortest geodesic (systole) on any } h \in \mathbf{pleat}(\alpha)$$

We take L_0 to be the Bers constant, which guarantees that the shorted geodesic has length less than $L_0 < L$.

Before we argue that π is a Lipschitz coarse retraction, we first would like to show that all the possible choices for $\pi(\alpha)$ lie in a bounded set in $\mathcal{C}(\Sigma)$ where the bound is uniform for fixed genus. If we pick $h \in \mathbf{pleat}(\alpha)$, then a surgery argument shows that systoles of (Σ, σ_h), must be *simultaneously disjoint* on the surface, so they lie in a ball of diameter 1 in $\mathcal{C}(\Sigma)$.

Exercise 7. Prove that the systoles on hyperbolic surface can be realized simultaneously disjoint.

Fix $h \in \mathbf{pleat}(\alpha)$ as a reference pleated surface. We will argue that there exists a curve γ that has simultaneously bounded length in $h \in \mathbf{pleat}(\alpha)$ and any other $h' \in \mathbf{pleat}(\alpha)$. A priori, we have no control of the σ_h-length of α, so it is worthwhile to think of α as being very long in (Σ, σ_h). Since Σ is closed, the injectivity radius of (Σ, σ_h) is greater than some $\varepsilon > 0$. It follows that there exists some $s > 0$ such that $s \cdot \varepsilon > \text{area}(\Sigma)$. By taking an $\varepsilon/2$ neighborhood of α in (Σ, σ_h), we can find a simple geodesic arc $v \subset \alpha$ of σ_h-length at most s with endpoints x, y such that $d_N(h(x), h(y)) \leq d_{\sigma_h}(x, y) \leq \varepsilon$. Let γ be the curve obtained by closing off v by the geodesic shortcut between x, y on (Σ, σ_h), see Figure 8. A little care will guarantee that γ is simple and essential. Notice that the larger the σ_h-length of α, the more parallel the geodesic representative of α at x and y.

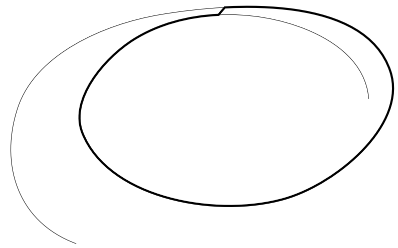

Fig. 8 Building the curve γ.

For any other $h' \in \mathbf{pleat}(\alpha)$, $h'(\alpha) = h(\alpha)$, so to bound the length of γ in $(\Sigma, \sigma_{h'})$, we have to control the length of the geodesic arc between x, y in $(\Sigma, \sigma_{h'})$. This is precisely where Thurston's Uniform Injectivity gives us our bound. When $h(\alpha)$ is long, $h(\alpha)$ has to be almost parallel at $h(x)$ and $h(y)$. By Theorem 5, it follows that the distance on any pleated surface between $h(x), h(y)$ is bounded, see Figure 9. Therefore, γ has uniformly bounded length for any pleated surface in $\mathbf{pleat}(\alpha)$.

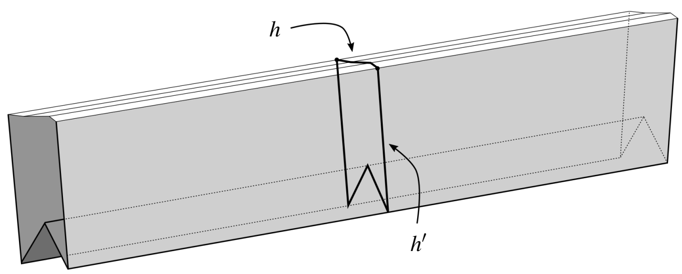

Fig. 9 The geodesic path between $h(x), h(y)$ for h and h'.

Curves of bounded length on a hyperbolic surface have bounded pairwise intersection number and so they lie a bounded distance from any systole in the curve graph. In particular, any systole of an element in $\mathbf{pleat}(\alpha)$ is a bounded distance from γ. Therefore, the set of all choices for $\pi(\alpha)$ will have bounded diameter in the curve graph.

Exercise 8. Prove that curves of bounded length on a hyperbolic surface have bounded pairwise intersection number.

Exercise 9. Prove that if the intersection number of $\alpha, \beta \in \mathcal{C}(\Sigma)$ is bounded, then their distance in the curve complex is also bounded.

To show that π is a Lipschitz coarse retraction, we have to show that if $d_{\mathcal{C}}(\alpha, \beta) \leq 1$ then $d_{\mathcal{C}}(\pi(\alpha), \pi(\beta))$ is bounded. Since we know that $\mathbf{pleat}(\alpha) \cap \mathbf{pleat}(\beta)$ is non-empty by a previous construction, the assertion follows from the bounded diameter result. $\square$

5.1 Towards the bi-Lipschitz model

Let $\Sigma = \Sigma_{0,5}$ be the sphere with 5 holes and $f : \Sigma \to N$ a homotopy equivalence. If N has simple geodesic curves of bounded length going out to the ends, then one can obtain ending lamination ν_-, ν_+. The points $\nu_\pm$ are on the boundary of the curve graph and it is possible to choose a bi-infinite geodesic in $\mathcal{C}(\Sigma)$ with endpoints $\nu_\pm$. Let v_i for $i \in \mathbb{Z}$ be the vertices of this geodesic. Notice that v_i is disjoint from v_{i-1}, v_{i+1}, so we can cut the surface Σ along v_i to get a subsurface Σ_i containing v_{i-1} and v_{i+1}. We can then connect v_{i-1} and v_{i+1} in $\mathcal{C}(\Sigma_i)$ by a geodesic path as in Figure 10.

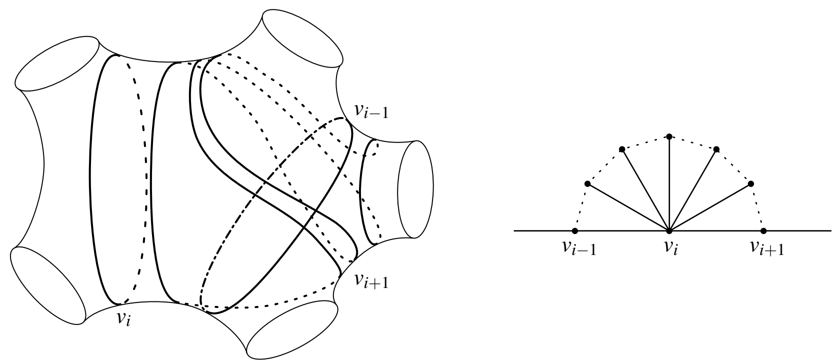

Fig. 10 Three vertices of the bi-infinite geodesic and the neighbouring path in the curve graph.

This data picks out the curves that carry all the important structural information of N. One can build a manifold with a piecewise Riemannian metric that models N via a bi-Lipschitz map to N. This determines the hyperbolic structure of N by a theorem of Sullivan. For the details on the construction for this example, see [14].

References

1. Aschenbrenner, M., Friedl, S., Wilton, H.: 3-manifold groups. ArXiv e-prints (2012)
2. Bonahon, F.: Bouts des variétés hyperboliques de dimension 3. Annals of Mathematics pp. 71–158 (1986)
3. Brock, J.F.: Iteration of mapping classes and limits of hyperbolic 3-manifolds. Inventiones mathematicae **143**(3), 523–570 (2001)
4. Canary, R.D.: A covering theorem for hyperbolic 3-manifolds and its applications. Topology **35**(3), 751–778 (1996)
5. Canary, R.D., Marden, A., Epstein, D.B.A.: Fundamentals of hyperbolic manifolds: Selected expositions. 328. Cambridge University Press (2006)
6. Farb, B., Margalit, D.: A Primer on Mapping Class Groups. Princeton Mathematical Series. Princeton University Press (2012)
7. Fathi, A.: Dehn twists and pseudo-anosov diffeomorphisms. Inventiones mathematicae **87**(1), 129–151 (1987). DOI 10.1007/BF01389156.

8. Hatcher, A.: On triangulations of surfaces. Topology and its Applications **40**(2), 189–194 (1991)
9. Hatcher, A.: Fibers of fibrations of a 3-manifold over S^1.
URL http://mathoverflow.net/a/61014 Cited 1 Aug 2015.
10. Klarreich, E.: The boundary at infinity of the curve complex and the relative Teichmüller space. preprint **226**, 230 (1999)
11. Masur, H.A., Minsky, Y.N.: Geometry of the complex of curves I: Hyperbolicity. Inventiones mathematicae **138**(1), 103–149 (1999)
12. Minsky, Y.N.: Harmonic maps into hyperbolic 3-manifolds. Transactions of the American Mathematical Society **332**(2), 607–632 (1992)
13. Minsky, Y.N.: Bounded geometry for Kleinian groups. Inventiones mathematicae **146**(1), 143–192 (2001)
14. Minsky, Y.N.: Combinatorial and geometrical aspects of hyperbolic 3-manifolds. In: Y. Komori, V. Markovic, S. C. (eds.) Kleinian Groups and Hyperbolic 3-Manifolds, pp. 3–40. Cambridge University Press (2003). Cambridge Books Online
15. Thurston, W.P.: Hyperbolic structures on 3-manifolds I: Deformation of acylindrical manifolds. Annals of Mathematics pp. 203–246 (1986)
16. Thurston, W.P.: A norm for the homology of 3-manifolds. Memoirs of the American Mathematical Society **59**(339), 99–130 (1986)
17. Thurston, W.P.: Hyperbolic Structures on 3-manifolds, II: Surface groups and 3-manifolds which fiber over the circle. ArXiv Mathematics e-prints (1998)
18. Thurston, W.P.: Hyperbolic Structures on 3-manifolds, III: Deformations of 3-manifolds with incompressible boundary. ArXiv Mathematics e-prints (1998)