arXiv:1102.2247v2 [math.DS] 14 Sep 2011

GEOMETRIZATION OF POSTCRITICALLY FINITE BRANCHED COVERINGS

SYLVAIN BONNOT, MICHAEL YAMPOLSKY

Abstract. We study canonical decompositions of postcritically finite branched coverings of the 2-sphere, as defined by K. Pilgrim. We show that every hyperbolic cycle in the decomposition does not have a Thurston obstruction. It is thus Thurston equivalent to a rational map.

Foreword. A proof of the main result was announced by Nikita Selinger at the Workshop "Holomorphic Dynamics around Thurston's Theorem" which took place at Roskilde University from September 27 - October 1, 2010 (a preprint [9] has soon appeared). At the same conference, we independently proposed a different approach to the proof, which is presented here. We are very grateful to Nikita for several useful discussions, and, particularly, for pointing out an error in a previous version of this paper, and helping us to correct it in the proof. Our argument bears a certain ideological similarity to that of [9]. However, it is based on a specific geometric surgery construction, rather than the more abstract concept of augmented moduli space used by Nikita.

1. Introduction and statement of the result

Thurston maps and multicurves. In this section we recall the basic setting of Thurston's characterization of rational functions. Let $f : S^2 \to S^2$ be an orientation-preserving branched covering map of the two-sphere. We define the *postcritical set* P_f by

$$P_f := \bigcup_{n>0} f^{\circ n}(\Omega_f),$$

where Ω_f is the set of critical points of f. When the postcritical set P_f is finite we say that f is a *Thurston mapping*.

Two Thurston maps f and g are *Thurston equivalent* if there are homeomorphisms $\phi_0, \phi_1 : S^2 \to S^2$ such that

(1) the maps ϕ_0, ϕ_1 coincide on P_f, send P_f to P_g and are isotopic rel P_f;

Date: January 13, 2019.

(2) the diagram

$$\begin{array}{ccc} S^2 & \xrightarrow{\phi_1} & S^2 \\ \Big\downarrow f & & \Big\downarrow g \\ S^2 & \xrightarrow{\phi_0} & S^2 \end{array}$$

commutes.

Given a Thurston map $f : S^2 \to S^2$, we define a function $N_f : S^2 \to \mathbb{N} \cup \infty$ as follows:

$$N_f(x) = \begin{cases} 1 & \text{if } x \notin P_f, \\ \infty & \text{if } x \text{ is in a cycle containing a critical point,} \\ \operatorname{lcm}\limits_{f^k(y)=x} \deg_y(f^{\circ k}) & \text{otherwise.} \end{cases}$$

The pair (S^2, N_f) is called the *orbifold of* f. The *signature* of the orbifold (S^2, N_f) is the set $\{N_f(x) \text{ for } x \text{ such that } 1 < N_f(x) < \infty\}$. The *Euler characteristic* of the orbifold is given by

$$\chi(S^2, N_f) := 2 - \sum_{x \in P_f} \left(1 - \frac{1}{N_f(x)}\right).$$

One can prove that $\chi(S^2, N_f) \leq 0$. In the case where $\chi(S^2, N_f) < 0$, we say that the orbifold is *hyperbolic*. Observe that most orbifolds are hyperbolic: indeed, as soon as the cardinality $|P_f| > 4$, the orbifold is hyperbolic.

We recall that a simple closed curve $\gamma \subset S^2 - P_f$ is *non essential* if it bounds a disk, and is *peripheral* if it bounds a punctured disk. We call a homotopy class of simple closed curves $[\gamma]$ *trivial* if it is either non-essential or peripheral.

Definition 1.1. *A multicurve Γ on (S^2, P_f) is a set of disjoint, non-homotopic, essential, nonperipheral simple closed curves on $S^2 - P_f$. A multicurve Γ is f-stable if for every curve $\gamma \in \Gamma$, each component α of $f^{-1}(\gamma)$ is either trivial or homotopic rel P_f to an element of Γ.*

To any f-stable multicurve is associated its *Thurston linear transformation* $f_\Gamma : \mathbb{R}^\Gamma \to \mathbb{R}^\Gamma$, best described by the following transition matrix

$$M_{\gamma\delta} = \sum_\alpha \frac{1}{\deg(f : \alpha \to \delta)}$$

where the sum is taken over all the components α of $f^{-1}(\delta)$ which are isotopic rel P_f to γ. Since this matrix has nonnegative entries, it has a leading eigenvalue $\lambda(\Gamma)$ that is real and nonnegative (by the Perron-Frobenius theorem).

We can now state Thurston's theorem:

Thurston's Theorem. *Let $f : S^2 \to S^2$ be a Thurston map with hyperbolic orbifold. Then f is Thurston equivalent to a rational function g if and only if $\lambda(\Gamma) < 1$ for every f-stable multicurve Γ. The rational function g is unique up to conjugation with an automorphism of $\mathbb{P}^1$.*

When a stable multicurve Γ has a leading eigenvalue $\lambda(\Gamma) \geq 1$, we call it a *Thurston obstruction*.

Pilgrim's canonical obstructions. Below we describe a particular type of Thurston obstructions, which were defined by K. Pilgrim in [7]. Let us assume that a Thurston map f has a hyperbolic orbifold. Let us denote $\mathcal{T}_f$ the Teichmüller space of the punctured sphere $S \equiv S^2 \setminus P_f$, and $d_{\mathcal{T}_f}(\cdot, \cdot)$ the Teichmüller distance; $\mathcal{M}_f$ will denote the moduli space of $S^2 \setminus P_f$; $p_f : \mathcal{T}_f \to \mathcal{M}_f$ will be the covering map. Further, for a choice of the complex structure τ on S, we let ρ_τ denote the hyperbolic metric on the Riemann surface $S_\tau \equiv (S, \tau)$, length_τ the hyperbolic length, and d_τ the hyperbolic distance. Similarly, for a general hyperbolic Riemann surface W we denote ρ_W, d_W, and length_W the hyperbolic metric, distance, and length on W; $\mathcal{T}_W$ the Teichmüller space, etc.

Definition 1.2. *For a non-trivial homotopy class of simple closed curves $[\gamma]$ on S we let $\ell_\tau([\gamma])$ denote the length of the unique geodesic representative of $[\gamma]$ in S_τ.*

The map f induces an analytic mapping on $\mathcal{T}_f$:

$$\sigma_f : \mathcal{T}_f \to \mathcal{T}_f, \text{ where } \sigma_f([\tau]) = [f^*\tau].$$

We note that this mapping projects to a *finite* cover of the moduli space:

Proposition 1.1 (Lemma 5.2 of [3])**.** *There exists a tower*

$$\mathcal{T}_f \xrightarrow{\tilde{p}_f} \widetilde{\mathcal{M}}_f \xrightarrow{\bar{p}_f} \mathcal{M}_f$$

of covering spaces, such that $\bar{p}_f$ is a finite cover, and a map $\tilde{\sigma}_f : \widetilde{\mathcal{M}}_f \to \mathcal{M}_f$, such that the diagram below commutes:

$$\begin{array}{ccc} \mathcal{T}_f & \xrightarrow{\sigma_f} & \mathcal{T}_f \\ \downarrow{\scriptstyle \tilde{p}_f} & & \downarrow{\scriptstyle p_f} \\ \widetilde{\mathcal{M}}_f & \xrightarrow{\tilde{\sigma}_f} & \mathcal{M}_f \end{array}$$

There exists a rational mapping $R : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ which is Thurston equivalent to f if and only if there is a fixed point $[\tau_*] = \sigma_f[\tau_*]$. The map σ_f does not increase Teichmüller distance. In the absence of a Thurston obstruction, σ_f^2 is strictly contracting [3], so for any choice of the starting

point $[\tau_0] \in \mathcal{T}_f$, the iterates $[\tau_n] \equiv \sigma_f^n([\tau_0])$ converge to $[\tau_*]$ geometrically fast.

Pilgrim showed that a presence of an obstruction implies that the sequence $[\tau_n]$ diverges to infinity in $\mathcal{T}_f$ in the following specific sense:

Theorem 1.2 ([7]). *Suppose f is obstructed. Then*

(I) *there exists a class $[\gamma]$ such that*
$$\ell_{\tau_n}([\gamma]) \to 0;$$
(II) *for a non-trivial homotopy class $[\gamma]$ the above property is independent of the starting point $[\tau_0] \in \mathcal{T}_f$;*
(III) *the union of all classes $[\gamma]$ as above forms a Thurston obstruction Γ_c.*

Pilgrim calls Γ_c the *canonical* Thurston obstruction. Thus, the existence of an obstruction implies that the canonical obstruction exists (that is $\Gamma_c \neq \emptyset$).

Pilgrim further showed:

Theorem 1.3 ([7]). *Let $\tau_0 \in \mathcal{T}_f$. There exists a constant $E = E(\tau_0)$ such that for every non-trivial simple closed curve $\gamma \notin \Gamma_c$ we have*
$$\inf \ell_{\sigma_f^n \tau_0}([\gamma]) > E.$$

For future reference, denote

$$D_\tau \equiv d_{\mathcal{T}_f}([\tau], \sigma_f[\tau]). \tag{1.1}$$

Pilgrim's decompositions and combinations of Thurston maps. What follows is a very brief review; the reader is referred to K. Pilgrim's book [8] for details. We adhere to the notation of [8], for ease of reference.

As a motivation, consider that for the canonical Thurston obstruction $\Gamma_c \ni \gamma$, there is a choice of complex structure τ for which $\ell_\tau([\gamma])$ is arbitrarily small, and remains small under pull-back by f. It is thus natural to think of the punctured sphere $S^2 \setminus P_f$ as pinching along the homotopy classes $[\gamma] \in \Gamma_c$; instead of a single sphere we then obtain a collection of spheres, interchanged by a map f.

More specifically, let f be a Thurston map, and $\Gamma = \cup \gamma_j$ an f-stable multicurve. Consider also a finite collection of disjoint annuli $A_{0,j}$ whose core curves are the respective γ_j. For each $A_{0,j}$ consider only non-trivial preimages; these form a collection of annuli $A_{1,k}$ each of which is homotopic to one of the curves in Γ. Pilgrim says that the pair (f, Γ) is in a *standard form* (see Figure 1) if there exists a collection of annuli $A_{0,j}$ as above such that the following properties hold:

(a) for each curve γ_j the annuli $A_{1,k}$ in the same homotopy class are contained inside $A_{0,j}$;

(b) moreover, the two outermost annuli $A_{1,k}$ as above share their outer boundary curves with $A_{0,j}$;
(c) finally, restricted to a boundary curve χ of $A_{0,j}$, the map $f : \chi \to f(\chi)$ is, up to a homeomorphic change of coordinates in the domain and the range, given by $z \mapsto z^d : S^1 \to S^1$, for some $d \geq 1$.

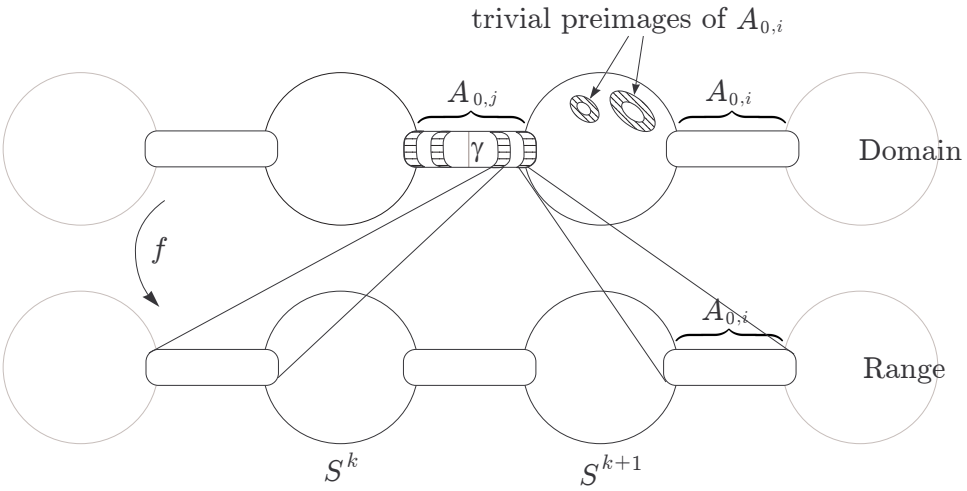


FIGURE 1. Pilgrim's decomposition of a Thurston mapping

A Thurston map with a multicurve in a standard form can be decomposed as follows. First, all annuli $A_{0,j}$ are removed, leaving a collection of spheres with holes, denoted $S_0(j)$. For each j, there exists a unique connected component $S_1(j)$ of $f^{-1}(\cup S_0(j))$ which has the property $\partial S_0(j) \subset \partial S_1(j)$. Any such $S_1(j)$ is a sphere with holes, with boundary curves being of two types: boundaries of removed annuli, or boundaries of trivial preimages of the removed annuli.

The holes in $S_0(j) \subset S^2$ can be filled using the property (c) above. Namely, let χ be a boundary curve of a component D of $S^2 \setminus S_0(j)$. Let $k \in \mathbb{N}$ be the first iterate $f^k : \chi \to \chi$, if it exists. For each $0 \leq i \leq k-1$ the curve $\chi_i \equiv f^i(\chi)$ bounds a component D_i of $S^2 \setminus S_0(m_i)$ for some m_i. Denote d_i the degree of $f : \chi_i \to \chi_{i+1}$. Select homeomorphisms

$$h_i : \bar{D}_i \to \bar{\mathbb{D}} \text{ so that } h_{i+1} \circ f \circ h_i^{-1}(z) = z^{d_i}.$$

Set $\tilde{f} \equiv f$ on $\cup S_0(j)$. Define new punctured spheres $\tilde{S}(j)$ by adjoining cups $h_i^{-1}(\bar{\mathbb{D}} \setminus \{0\})$ to $S_0(j)$. Extend the map $\tilde{f}$ to each D_i by setting

$$\tilde{f}(z) = h_{i+1}^{-1} \circ (h_i(z))^d.$$

We have thus replaced every hole with a cap with a single puncture.

By construction, the map

$$\tilde{f} : \cup\tilde{S}(j) \to \cup\tilde{S}(j)$$

contains a finite number of periodic cycles of punctured spheres. For every periodic cycle of spheres, pick a representative $\tilde{S}(j)$, and denote by $\mathcal{F}$ the first return map $f^{k_j} : \tilde{S}(j) \to \tilde{S}(j)$. This is again a Thurston map. The collection of maps $\mathcal{F}$ and the combinatorial information required to glue the spheres $S_0(j)$ back together is what Pilgrim calls a *decomposition* of f.

Pilgrim shows:

Theorem 1.4. *For every obstructed Thurston map f there exists an equivalent map g, whose canonical obstruction we denote Γ_c^g, such that (g, Γ_c^g) is in a standard form, and thus can be decomposed.*

Statement of the geometrization result. It is natural to ask whether the canonical decomposition described above has the maximality property: that is, whether the restrictions of the return map $\mathcal{F}$ to spheres $\tilde{S}(j)$, which have topological degrees greater than one, and a hyperbolic orbifold, are *unobstructed*. In view of Thurston's theorem, this would imply that for every such $\tilde{S}(j)$ there is a unique, up to a normalization, rational map $R : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ which is equivalent to $\mathcal{F} : \tilde{S}(j) \to \tilde{S}(j)$. This constitutes the main conjecture posed by Pilgrim in [8].

In our main result we prove Pilgrim's conjecture. Thus, Pilgrim's decomposition of an obstructed Thurston mapping canonically breaks it into unobstructed, and thus *geometrizable*, pieces.

Main Theorem. *Let $\mathcal{F} : \cup\tilde{S}(j) \to \cup\tilde{S}(j)$ be given by Pilgrim decomposition of an obstructed Thurston map along the canonical obstruction. For every j such that $\mathcal{F} : \tilde{S}(j) \to \tilde{S}(j)$ has a topological degree greater than 1 and a hyperbolic orbifold, there does not exist a Thurston obstruction in $\tilde{S}(j)$.*

2. Preliminaries

Some notation. Let $\tilde{S}(j)$ be as in the statement of the main theorem. It is obtained by adding caps D_i to a sphere with holes $S_1(j) \subset S$. The restriction of the first return map $\mathcal{F}$ to $\tilde{S}(j)$ is an iterate $\tilde{f}^n$, so that we have a cycle of punctured spheres

$$S^0 \equiv \tilde{S}(j) \xrightarrow{\tilde{f}} S^1 \xrightarrow{\tilde{f}} \cdots \xrightarrow{\tilde{f}} S^{n-1} \xrightarrow{\tilde{f}} S^0.$$

Collar Lemmas. Let us denote by $s(x)$ the function

$$s(x) = \sinh^{-1}(1/\sinh(x/2)).$$

Note that $s(x)$ decreases from ∞ to 0 as x increases from 0 to ∞. The *collar* around a simple closed hyperbolic geodesic γ on a hyperbolic Riemann surface W is the neighborhood

$$\mathcal{C}(\gamma) \equiv \{z \in W \ | d_W(z, \gamma) < s(\text{length}_W(\gamma))\}.$$

The following is known as Collar Lemma (cf. [2]):

Theorem 2.1 (**Collar Lemma for closed geodesics**). *The collar $\mathcal{C}(\gamma)$ is an annulus. Further, if γ and δ are two disjoint simple closed geodesics on a hyperbolic Riemann surface W, then*

$$\mathcal{C}(\gamma) \cap \mathcal{C}(\delta) = \emptyset.$$

We also recall a limiting version of Collar Lemma for cusps (see [5]).

Lemma 2.2 (**Collar Lemma for cusps**). *Let us denote $\mathcal{C}$ the quotient of the region*

$$\{z : \text{Im}(z) > 1\} \subset \mathbb{H}$$

by the translation $z \mapsto z + 2$. For every cusp κ of a hyperbolic surface X there is an isometry

$$\iota_\kappa : \mathcal{C} \to C(\kappa)$$

between $\mathcal{C}$ and a collar neighborhood $C(\kappa) \subset X$. Furthermore, the collars about different cusps are disjoint, and $C(\kappa)$ is disjoint from the collar $C(\gamma)$ about any simple geodesic γ on X.

For $\Delta > 1$ we denote $\mathcal{C}_\Delta \subset \mathcal{C}$ the quotient of $\{z : \text{Im}(z) > \Delta\} \subset \mathbb{H}$ by $z \mapsto z + 2$, and call $\iota_\kappa(\mathcal{C}_\Delta)$ the Δ-*neighborhood of the cusp* κ.

Teichmüller spaces. Here we describe a key technical tool that we will use, which is due to Y. Minsky (see [6]).

Let S be a surface of finite type, and $\Gamma = \{\gamma_1, \ldots, \gamma_k\}$ be a system of disjoint, homotopically distinct, non trivial simple closed curves on S. Then we denote by $\text{Thin}_\epsilon(S, \Gamma)$ the *thin* part of the Teichmüller space of S given by

$$\text{Thin}_\epsilon(S, \Gamma) = \{[\tau] \in \mathcal{T}_s \text{ for which } \sup_i \ell_\tau([\gamma_i]) < \epsilon\}.$$

Denote S_Γ the multiply connected surface obtained from S by decomposing along the collection γ_i (that is, cutting along γ_i, and capping every hole by a puncturd disk). There is a natural homeomorphism

$$\Pi_\Gamma : \mathcal{T}_S \to \mathcal{T}_{S_\Gamma} \times \mathbb{H}_1 \times \ldots \times \mathbb{H}_k$$

defined as follows. Firstly, we complete the family Γ into a larger family of curves $\hat{\Gamma}$ decomposing S into pairs of pants (possibly degenerate). Recall the definition of Fenchen-Nielsen coordinates

$$(\ell_\tau(\gamma), \mathrm{Twist}_\tau(\gamma))_{\gamma\in\hat{\Gamma}} \text{ on } \mathcal{T}_S,$$

given by the homotopy classes of curves in $\hat{\Gamma}$. The projection $\mathcal{T}_S \mapsto \mathcal{T}_{S_\Gamma}$ is obtained by restriction of these coordinates to $\hat{\Gamma} \setminus \Gamma$. For each $\gamma_i \in \Gamma$, the map Π sends

$$(\ell_\tau([\gamma_i]), \mathrm{Twist}_\tau([\gamma_i])) \longrightarrow \left(\frac{1}{\ell_\tau([\gamma_i])}, \mathrm{Twist}_\tau([\gamma_i])\right) \in \mathbb{H}_i.$$

Minsky's result gives a distortion bound on this homeomorphism for small enough values of ϵ:

Theorem 2.3 ([6])**.** *Endow the product $\mathcal{T}_{S_\Gamma} \times \mathbb{H}_1 \times \ldots \times \mathbb{H}_k$ with the maximum of the Teichmüller distances on the connected components of $\mathcal{T}_{S_\gamma}$, and the hyperbolic distances on each copy of $\mathbb{H}$. Then, for all $\epsilon > 0$ sufficiently small, the homeomorphism Π_Γ, restricted to $\mathrm{Thin}_\epsilon(S, \gamma)$ distorts distances by a bounded additive amount.*

Decompositions. To fix the ideas, for the remainder of this paper, the decomposition of a Riemann surface S with a canonical multicurve Γ_c on the level of Teichmüller spaces will be carried out by restricting the Fenchen-Nielsen projection

$$\Pi_\Gamma : \mathcal{T}_S \to \mathcal{T}_{S_{\Gamma_c}} \times \mathbb{H}_1 \times \ldots \times \mathbb{H}_k \tag{2.1}$$

to the Teichmüller space $\mathcal{T}_{S_{\Gamma_c}}$.

Geometrically, this is achieved by cutting and capping the Riemann surface S *preserving the hyperbolic structure.* Note, that the gluing of two hyperbolic pieces along (piecewise) geodesic boundaries is possible whenever the angles at the corners of the boundaries add up to 2π and the lengths of the boundaries agree. We cannot decompose S along geodesic loops in Γ_c, as there are no closed geodesics surrounding a cusp. Instead, we will use *geodesic bigons* bounded by geodesic arcs meeting at the same angles α at the two vertices. An easy application of Gauss-Bonnet Theorem together with the Collar Lemmas shows that:

Proposition 2.4. *There exists $\delta_0 > 0$ such that, assuming that*

$$\max_{\gamma\in\Gamma} \ell_\tau([\gamma]) < \delta_0,$$

the surface (S, τ) can be decomposed into $\cup\tilde{S}(j)$ in such a way that:

- *the inclusion of the surface with holes $S_1(j) \subset \tilde{S}(j)$ into (S, τ) is a hyperbolic isometry;*

- *the caps D_i are bounded by geodesic bigons.*

Furthermore, for every $\epsilon > 0$ there exists $\delta > 0$ such that if

$$\max_{\gamma \in \Gamma} \ell_\tau([\gamma]) < \delta,$$

then the surface can be decomposed so that every cap $D_i \subset S_0(j) \setminus S_1(j)$ has spherical diameter less than ϵ in (S, τ).

Combinations of punctured surfaces. Let us briefly describe the inverse of the above decomposition procedure, for future reference. Pilgrim's *combination* applied to $\cup \tilde{S}(j)$. involves connecting the punctured spheres by annuli to reverse the decomposition construction. The end result is a branched covering which is equivalent to $f : S^2 \to S^2$. Let $S_1 \equiv \tilde{S}(j_1)$ and $S_2 \equiv \tilde{S}(j_2)$ be two punctured spheres to be combined at the pair of punctures $p_1 \in S_1$ and $p_2 \in S_2$. We will glue in an annulus by removing a small topological disk $D_i \subset S_i$ around each of the punctures p_i, and attaching the boundary curves of the annulus to the boundary circles ∂D_i. We again would like to perform the gluing in a hyperbolically rigid fashion.

Let $W_1, \ldots, W_N$ be a finite collection of punctured spheres, each one being a hyperbolic Riemann surface. For each puncture $p_j \in W_1 \sqcup \ldots \sqcup W_N$ we choose a disk around it D_{p_j} that contains only the puncture p_j.

Definition 2.1 (Admissible pairing of caps)**.** *We say that a finite family of pairs (p_i, p_j) is an admissible pairing of punctures if*

1. *for each pair (p_i, p_j) the punctures p_i and p_j belong to distinct spheres;*
2. *the connected sum $W_1 \# \cdots \# W_N$ obtained by gluing along the pairs (p_i, p_j) is a punctured sphere with at least three punctures.*

We can now describe combining S_1 and S_2.

Lemma 2.5. *Let us consider two hyperbolic punctured spheres S_1, S_2 each one with a distinguished puncture $p_1 \in S_1$ and $p_2 \in S_2$. There exists an $\epsilon_0 > 0$ such that for any $\epsilon \leq \epsilon_0$ and every $\Delta > 1$ the following is true:*

1. *for an angle $\alpha = \alpha(\epsilon, \Delta)$ there exist topological disks $D_1 \ni p_1$, $D_2 \ni p_2$ in the respective Δ-neighborhoods of p_1, p_2 respectively, which are isotopic to the collar neighborhoods of the cusps (relative to the complements of the neighborhoods), and such that the boundary circle $C_i \equiv \partial D_i$ is made of two geodesic arcs meeting at two points with the angle α;*
2. *there exists a hyperbolic annulus A such that the core curve of A has length ϵ, and A is bounded by two topological circles C_1' and C_2', each one of which is made of two geodesic arcs meeting at two*

points at angle $2\pi-\alpha$, and having the same lengths as the geodesics forming C_1 and C_2 respectively.

Proof. Using Gauss-Bonnet Theorem on topological annuli bounded by piecewise geodesic arcs, one can find arbitrarily thin annuli bounded by pairs of geodesic arcs meeting at an angle arbitrarily close to π. Using the same theorem, one can find circles made of two geodesic arcs meeting at an angle arbitrarily close to $2\pi - \pi = \pi$ near a given cusp.

□

By applying the above statement repeatedly, we obtain the following lemma.

Lemma 2.6. *Let $S_1, \dots, S_N$ be a collection of hyperbolic punctured spheres together with an admissible pairing of punctures $\{(p_{i_1}, p_{j_1}), \dots (p_{i_K}, p_{j_K})\}$. There exists $\epsilon_0 > 0$ so that for any $\epsilon < \epsilon_0$ and $\Delta > 1$ one can perform a gluing on all the pairs (p_i, p_j) by using hyperbolic annuli A_{p_i,p_j} so that the resulting hyperbolic Riemann surface S has the following properties:*

1. *each S_j is isometrically embedded into S with complementary annuli A_{p_i,p_j};*
2. *the boundary circles of the annuli are geodesic bigons contained in the Δ-neighborhoods of the cusps;*
3. *the length of the core curve of each A_{p_i,p_j} is equal to ϵ.*

More on decompositions. The following facts will help us in bounding the changes to hyperbolic lengths, and complex structures, when decomposing Riemann surfaces. First, we note a standard consequence of the comparison of the hyperbolic metric on a connected subdomain Ω of $\hat{\mathbb{C}}$ with $\frac{1}{d(z,\partial\Omega)}|dz|$:

Lemma 2.7. *Let Ω denote the Riemann sphere $\hat{\mathbb{C}}$ with p punctures ($p \geq 3$) with its Poincaré metric ρ_Ω, and let γ be a non peripheral curve. For any constant $\epsilon > 0$ there exists $\delta > 0$ so that the hyperbolic length of γ does not change by more than ϵ when removing disks around the punctures with spherical radii less than δ.*

Next, we quote a lemma due to Haïssinsky [4]:

Lemma 2.8 (Haïssinsky)**.** *Let X be a finite type hyperbolic surface and $x \in X$. For any neighborhood of $([id], X)$ in Teich(X), there exists $r_0 > 0$ such that, for any $0 < r < r_0$ the deformations restricted to the disk centered at x and with hyperbolic radius r are contained in that neighborhood.*

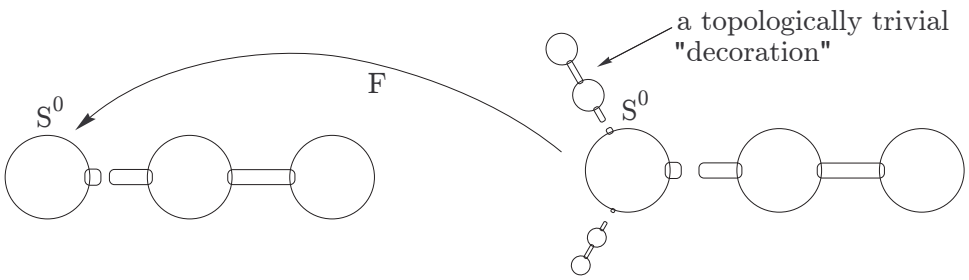


FIGURE 2. Pull-back of a decomposition

Idea of the proof. When we pull-back a complex structure τ_0 by f, only the geodesics of the canonical restriction become arbitrarily short. If we restrict to a subsurface $S_1(j)$, then, with respect to the structure

$$\tau_n \equiv \sigma_f^n \tau_0,$$

all simple closed curves have hyperbolic length bounded from below independently of n (Theorem 1.3). However, if we restrict τ_0 to a complex structure μ_0 on $\tilde{S}_0(j)$, and pull back by $\sigma_{\mathcal{F}}$, the result may *a priori* be different. As illustrated in Figure 2, the restrictions of τ_n to $S_0(j)$ will include "decorations": trivial preimages of the other subsurfaces. They are attached by very thin tubes, and will not create much difference for the first few pull-backs. However, without some uniform distortion bound on the change in the complex structure caused by restriction to the subsurface, we will not be able to control their impact on the length of geodesics. Here Minsky's Theorem 2.3 will be a key tool for us, as it produces a uniform bound on the change in complex structure caused by a decomposition.

3. Proof of the Main Theorem

Let us start by fixing $\tau_0 \in \mathcal{T}_f$, and setting $\tau_n \equiv \sigma_f \tau_0$. Let $E = E(\tau_0)$ be given by Theorem 1.3. Let

$$\mathcal{T}_{\mathcal{F}} \xrightarrow{\tilde{p}_{\mathcal{F}}} \widetilde{\mathcal{M}_{\mathcal{F}}} \xrightarrow{\bar{p}_{\mathcal{F}}} \mathcal{M}_{\mathcal{F}}$$

be as in Proposition 1.1. We first prove:

Lemma 3.1. *There exists a compact set $K \Subset \widetilde{\mathcal{M}_{\mathcal{F}}}$ and a sequence $n_j \to \infty$ with the following properties. Consider the decomposition of (S, τ_n) given by*

$$\mu_n \equiv \pi_{\mathcal{T}_{\mathcal{F}}} \circ \Pi_{\Gamma_c}(\tau_n) \in \mathcal{T}_{\mathcal{F}}$$

where $\pi_{\mathcal{T}_{\mathcal{F}}}$ denotes the restriction of the Fenchen-Nielsen projection (2.1) to the first coordinate. Then we have:

- *for all* $n \in \mathbb{N}$ *we have the projections* $\tilde{\mu}_n \equiv \tilde{p}_{\mathcal{F}}(\mu_n) \in K$*;*
- *for each* $k \in \mathbb{N}$ *we have* $\tilde{\mu}_{n_j+k} \to \tilde{\alpha}_k$*.*

Proof. As before, let Γ_c stand for the canonical obstruction of f. By Theorem 1.2, for every $\epsilon > 0$ there exists n_0 such that for all $n > n_0$ we have

$$[\tau_n] \in \mathrm{Thin}_\epsilon(S, \Gamma_c).$$

By the Collar Lemma, and Lemma 2.7, for every $\epsilon_1 > 0$ there exists $n_1 \in \mathbb{N}$ such that for every essential curve $\gamma \subset S_0(j)$

$$||\ell_{\tau_n}([\gamma]) - \ell_{\tau_n|S_0(j)}([\gamma])|| < \epsilon_1.$$

By Theorem 1.3, for every simple closed curve $\gamma \subset \cup S_0(j)$ such that:

(a) $[\gamma]$ is non-trivial in S, and
(b) $[\gamma] \notin \Gamma_c$,

we have $\ell_{\tau_n}([\gamma]) > E$. Thus, by Mumford Compactness Theorem, the projections

$$p_{\mathcal{F}} \circ \pi_{\mathcal{T}_{\mathcal{F}}} \circ \Pi_{\Gamma_c}(\tau_n) = \tilde{p}_{\mathcal{F}}(\mu_n)$$

lie in a compact subset K of $\mathcal{M}_{\mathcal{F}}$. The next statement follows from compactness of K in a standard fashion.

□

We prove the following:

Proposition 3.2. *Let* N *be the least common multiple of the periods of the cycles of decomposed surfaces. For every* $k, m \in \mathbb{N}$

$$d_{k,m,l} \equiv d_{\mathcal{T}_{\mathcal{F}}}([\sigma^m_{\mathcal{F}}(\mu_k)], [\pi_{\mathcal{T}_{\mathcal{F}}} \circ \Pi_{\Gamma_c}(\sigma^m_{f^N}(\tau_{n_l+k}))]) \underset{l\to\infty}{\longrightarrow} 0.$$

Proof. Recall that the surface S^0 is obtained by capping the surface with holes $S_1(j) \subset S_0(j)$ as described in Proposition 2.4. Let us *combine* the truncated surface

$$S' \equiv (S \setminus S_0(j), \sigma^m_{f^N}(\tau_{n_l+k}))$$

with the surface

$$(S^0, \sigma^m_{\mathcal{F}}(\mu_k)).$$

Using Lemma 2.5 we can combine the two surfaces into a surface $(S, \lambda_{k,m})$ of the same topological type as S, with the only difference of the complex structure on the component $S_0(j)$, which is now given by $\sigma^m_{\mathcal{F}}(\mu_k)$.

By Minsky's Theorem, up to a small error, the Teichmüller distance between $(S, \lambda_{k,m})$ and $(S, \sigma^m_{f^N}(\tau_{n_l+k}))$ is given by the distance $d_{k,m,l}$. We need to see the effect of adding "decorations" to the complex structure $\sigma^m_{cF}(\mu_k)$ (cf. Figure 2).

By the Collar Lemma, for n_l large enough, the value of Δ in Lemma 2.5 can be chosen arbitrarily large. By Haïssinsky's result 2.8,

$$d_{\mathcal{T}_f}([(S, \lambda_{k,m})], [(S, \sigma_{f^N}^m(\tau_{n_l+k}))]) \underset{l\to\infty}{\longrightarrow} 0.$$

By Minsky's Theorem 2.3, and Lemma 3.1

$$d_{k,m,l} \underset{l\to\infty}{\longrightarrow} 0.$$

□

We can now complete the proof of the Main Theorem as in [9]. Note that by Proposition 1.1 and Proposition 3.2, the projections

$$\tilde{\sigma}_{\mathcal{F}}(\tilde{\alpha}_k) = \hat{\alpha}_{k+1} \equiv \bar{p}_{\mathcal{F}}(\tilde{\alpha}_{k+1}).$$

Since $\bar{p}_{\mathcal{F}}$ is a finite cover, the points $(\bar{p}_{\mathcal{F}})^{-1}(\hat{\alpha}_{k+1})$ form a discrete set. Hence, Proposition 3.2 implies that

$$\tilde{p}_{\mathcal{F}}(\sigma_{\mathcal{F}}(\mu_{n_j+k})) \to \tilde{\alpha}_{k+1}.$$

The pull-back $\sigma_{\mathcal{F}}$ is a strict contraction on every component of the preimage $(\tilde{p}_{\mathcal{F}})^{-1}(K)$. We thus have a limit point $\tilde{\lambda} \in K = \lim \tilde{\alpha}_{m_j}$ such that

$$\lim \tilde{p}_{\mathcal{F}}(\sigma_{\mathcal{F}}(\alpha_{m_j})) = \tilde{\lambda}.$$

Hence, $\mathcal{F}$ is Thurston equivalent to a rational mapping, and thus, not obstructed.

References

[1] S. Bonnot, M. Braverman, and M. Yampolsky, *Thurston equivalence is decidable*, e-print Arxiv.org **1009.5713**.

[2] P. Buser, *Geometry and spectra of compact Riemann surfaces*, Progress in Mathematics, vol. 106, Birkhäuser Boston Inc., Boston, MA, 1992.

[3] A. Douady and J.H. Hubbard, *A proof of Thurston's topological characterization of rational functions*, Acta Math. **171** (1993), 263–297.

[4] P. Haïssinsky, *Déformation localisée de surfaces de Riemann*, Publ. Mat. **49** (2005), no. 1, 249–255.

[5] C.T. McMullen, *Complex dynamics and renormalization*, Annals of Math. Studies, vol. 135, Princeton University Press, Princeton, NJ, 1994.

[6] Y. Minsky, *Extremal length estimates and product regions in Teichmuller space*, Duke Math. J. **83** (1996), 249–286.

[7] K. Pilgrim, *Canonical Thurston obstructions*, Adv. Math. **158** (2001), no. 2, 154–168.

[8] ———, *Combinations of complex dynamical systems*, Lecture Notes in Mathematics, vol. 1827, Springer, 2003.

[9] N. Selinger, *Thurston's pullback map on the augmented Teichmüller space and applications*, e-print at Arxiv.org **1010.1690** (October 11, 2010).