

CENTRALIZERS IN MAPPING CLASS GROUP AND DECIDABILITY OF THURSTON EQUIVALENCE

KASRA RAFI, NIKITA SELINGER, MICHAEL YAMPOLSKY

ABSTRACT. We find a constructive bound for the word length of a generating set for the centralizer of an element of the Mapping Class Group. As a consequence, we show that it is algorithmically decidable whether two postcritically finite branched coverings of the sphere are Thurston equivalent.

1. INTRODUCTION

In [SY15], it has been shown that there exists an algorithm that finds a canonical decomposition of an obstructed Thurston map as well as geometrization of all cycles in that decomposition (see Theorem 7.1 for a precise statement). Our motivation for the present article is to build an algorithm that can check whether pairwise equivalence of the pieces of canonical decompositions of two Thurston maps can be glued together into a global equivalence between the two.

A prototype of the algorithm has been already presented in [SY15] where the same result has been shown for a subclass of Thurston maps that are only allowed to have hyperbolic cycles in their canonical decompositions. This restriction significantly simplifies the problem as the group of self-equivalences of a Thurston map with hyperbolic orbifold is trivial (which follows from the fact that an equivalence between two Thurston maps with hyperbolic orbifolds is unique [DH93]). We study the self-equivalence groups in the other cases. By constructively characterizing the generators of all of the groups involved, we reduce a countable search to solving a finite number of linear problems. A different approach to the problem of algorithmically verifying Thurston equivalence has been studied in [BD17].

Our key result is a complexity bound on self-equivalences in the case when a first return map on a component in the canonical decomposition of a Thurston map is a homeomorphism, that is, a bound on centralizers of elements of the Mapping Class Group. To accomplish this, we prove the following theorem:

Theorem 1.1. *For every element ϕ of the Mapping Class Group, the centralizer of ϕ has a generating set where every element has a word length that is bounded by a uniform multiple of the word length of ϕ .*

Armed with this statement, we obtain:

Theorem 1.2. *There exists an algorithm which for any two Thurston maps f and g outputs an equivalence ϕ if f and g are equivalent, and outputs **maps are not equivalent** otherwise.*

2. BACKGROUND

2.1. Thurston maps. Let $f : S^2 \rightarrow S^2$ be an orientation-preserving branched covering self-map of the two-dimensional topological 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 *postcritically finite*.

A (*marked*) *Thurston map* is a pair (f, Q_f) where $f : S^2 \rightarrow S^2$ is a postcritically finite ramified covering of degree at least 2 and Q_f is a finite collection of marked points $Q_f \subset S^2$ which contains P_f and is f -invariant: $f(Q_f) \subset Q_f$. In particular, all elements of Q_f are pre-periodic for f .

Thurston equivalence. Two marked Thurston maps (f, Q_f) and (g, Q_g) are *Thurston (or combinatorially) equivalent* if there are homeomorphisms $\phi_0, \phi_1 : S^2 \rightarrow S^2$ such that

- (1) the maps ϕ_0, ϕ_1 coincide on Q_f , send Q_f to Q_g and are isotopic rel Q_f ;
- (2) the diagram

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

commutes.

We will call (ϕ_0, ϕ_1) an *equivalence pair*.

Let Q be a finite collection of points in S^2 . Recall that a simple closed curve $\gamma \subset S^2 - Q$ is *essential* if it does not bound a disk, is *non-peripheral* if it does not bound a punctured disk.

Definition 2.1. A *multicurve* Γ on (S^2, Q) is a set of disjoint, nonhomotopic, essential, non-peripheral simple closed curves on $S^2 \setminus Q$. Let (f, Q_f) be a Thurston map, and set $Q = Q_f$. A multicurve Γ on $S \setminus Q$ is *f -stable* if for every curve $\gamma \in \Gamma$, each component α of $f^{-1}(\gamma)$ is either trivial (meaning inessential or peripheral) or homotopic rel Q to an element of Γ .

Definition 2.2. A *Levy cycle* is a multicurve

$$\Gamma = \{\gamma_0, \gamma_1, \dots, \gamma_{n-1}\}$$

such that each γ_i has a nontrivial preimage γ'_i , where the topological degree of f restricted to γ'_i is 1 and γ'_i is homotopic to $\gamma_{(i-1) \bmod n}$ rel Q . A Levy cycle is *degenerate* if each γ_i has a preimage γ'_i as above such that γ'_i bounds a disk D_i and the restriction of f to D_i is a homeomorphism and $f(D_i)$ is homotopic to $D_{(i+1) \bmod n}$ rel Q .

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

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

where the sum is taken over all the components α of $f^{-1}(\delta)$ which are isotopic rel Q 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).

The celebrated Thurston's Theorem [DH93] is the following:

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

In view of this, an f -stable multicurve Γ with $\lambda(\Gamma) \geq 1$ is called a *Thurston obstruction*.

In [SY15], the second and third authors obtained a similar statement for Thurston maps with parabolic orbifolds:

Theorem 2.3. *Let f be a Thurston map with postcritical set P and marked set $Q \supset P$ such that the associated orbifold is parabolic and the associated matrix is hyperbolic. Then either f is equivalent to a quotient of an affine map or f admits a degenerate Levy cycle.*

Furthermore, in the former case the affine map is defined uniquely up to affine conjugacy.

3. CENTRALIZERS OF ELEMENTS IN THE MAPPING CLASS GROUP

Let ϕ be an element of the mapping class group $\text{Mod}(S)$ of a surface S of finite type. Let

$$C(\phi) = \left\{ \psi \in \text{Mod}(S) \mid \psi \phi = \phi \psi \right\}$$

be the centralizer of ϕ in $\text{Mod}(S)$. Fix a generating set \mathcal{G} for $\text{Mod}(S)$ and let $\|\cdot\|_{\mathcal{G}}$ denote the word length with respect to this generating set. For $M > 0$ define

$$C(\phi, M) = \left\{ \psi \in C(\phi) \mid \|\psi\|_{\mathcal{G}} \leq M \right\}.$$

In this section, we prove the following version of Theorem 1.1:

Theorem 3.1. *There is a constant M_0 , depending on S and the generating set \mathcal{G} , so that for every $\phi \in \text{Mod}(S)$, $C(\phi, M_0 \|\phi\|_{\mathcal{G}})$ generates $C(\phi)$.*

A computational consequence of the above theorem is the following:

Corollary 3.2. *There is an algorithm which, given $\phi \in \text{Mod}(S)$, outputs a set of generators of $C(\phi)$.*

3.1. Some tools. Our main tool is the following theorem of J. Tao.

Theorem 3.3 ([Tao13]). *For any fixed generating set \mathcal{G} for $\text{Mod}(S)$, there exists a constant K , such that if $\phi, \phi' \in \text{Mod}(S)$ are conjugate, then there is a conjugating element η with*

$$\|\eta\|_{\mathcal{G}} \leq K \max(\|\phi\|_{\mathcal{G}} + \|\phi'\|_{\mathcal{G}}).$$

Let us introduce the following notations: $a \asymp_C b$, will mean $a < Cb + C$ and $b < Ca + C$, and $a \stackrel{*}{\prec} b$ will mean $a < Nb$ for some fixed N , and $a \stackrel{*}{\succ} b$ will mean $a \stackrel{*}{\prec} b$ and $b \stackrel{*}{\prec} a$.

We will also need the Masur-Minsky distance formula [MM00]. For every subsurface $R \subset S$, they define a measure of complexity between two curve systems $d_R(\cdot, \cdot)$ called the *subsurface projection distance* (see [MM00] for more details). Recall that a marking on a surface S is a pants decomposition and a set of dual curves for each pants curve (again, see [MM00] for precise definition and details).

Theorem 3.4. *For any generating set \mathcal{G} , any marking μ_0 , and any threshold k that is sufficiently large, there is a uniform constant C so that, for any $\eta \in \text{Mod}(S)$, we have*

$$\|\eta\|_{\mathcal{G}} \asymp_C \sum_{R \subset S} \left[d_R(\mu_0, \eta(\mu_0)) \right]_k.$$

Here the sum is over all subsurfaces R of S , and the function $[\cdot]_k$ is a truncation function with $[x]_k = x$ when $x \geq k$ and 0 otherwise.

3.2. Special cases.

Proposition 3.5. *Theorem 3.1 holds if ϕ is finite order.*

Proof. There are finitely many conjugacy classes of finite order elements in $\text{Mod}(S)$ (see for example [FM12, Theorem 7.13]). By Theorem 3.3, it is sufficient to show that, for each finite order element ϕ , $C(\phi)$ is finitely generated. Indeed, consider a set \mathcal{F} of finite order elements by picking one representative from every conjugacy class. If each $C(\xi)$, $\xi \in \mathcal{F}$ is finitely generated, then there is a uniform upper-bound M_1 for the word length of all elements in any such generating set. If ϕ is conjugate to $\phi' \in \mathcal{F}$, there is a conjugating element η , where $\|\eta\|_{\mathcal{G}} \leq K \max(\|\phi\|_{\mathcal{G}} + \|\phi'\|_{\mathcal{G}})$. Then, we can find a generating set for $C(\phi)$ by conjugating a generating set for $C(\phi')$. But $\|\phi'\|_{\mathcal{G}}$ is uniformly bounded (\mathcal{F} is finite). Hence, there is M_0 where the word length of this generating set for $C(\phi)$ are bounded by $M_0\|\phi\|_{\mathcal{G}}$.

Now let ϕ be any finite order element. To see that $C(\phi)$ is finitely generated, let Σ be the orbifold quotient of S by ϕ and $\text{Mod}^{\circ}(\Sigma)$ be the orbifold mapping class group of Σ . Then, $\text{Mod}^{\circ}(\Sigma)$ is a finite index subgroup of $\text{Mod}(\Sigma)$ and hence (say, using Schreier's lemma) is finitely generated. There is a finite index sub-group of $\text{Mod}^{\circ}(\Sigma)$ that lifts to sub-group C_{Σ} of $\text{Mod}(S)$ (see MacLachlan and Harvey [MH75, Theorem 10]) which is also finitely generated. Finally, $C(\phi)$ is a finite extension of C_{Σ} and hence is also finitely generated. \square

Proposition 3.6. *Theorem 3.1 holds if ϕ is a pseudo-Anosov element.*

Proof. By [McC94] $C(\phi)$ is a virtually cyclic where the degree of the extension is uniformly bounded, in particular $C(\phi)$ is finitely generated. In fact, if F_- and F_+ are the stable and unstable measured foliations associated to ϕ , then any $\psi \in C(\phi)$ preserves the pair (F_-, F_+) as a set.

To prove the theorem, it is sufficient to show that, for any $\psi \in C(\phi)$ there is a power m so that $\|\psi\phi^m\|_{\mathcal{G}} \prec^* \|\phi\|_{\mathcal{G}}$. Indeed, this shows that $C(\phi)$ is generated by ϕ and elements in $C(\phi)$ whose word length is less than a multiple of $\|\phi\|_{\mathcal{G}}$.

We use Theorem 3.4 to find such a bound. First, we claim that there exists an integer m so that

$$(1) \quad d_S(\mu_0, \psi\phi^m(\mu_0)) \prec^* d_S(\mu_0, \phi(\mu_0)).$$

Let $\mathcal{B} = \mathcal{B}_{\phi}$ be the quasi-axis of ϕ in the curve graph of S , that is a geodesic in the curve graph that is preserved by a power $\phi^{m'}$ of ϕ (see [Bow08]). Then \mathcal{B} limits to F_{\pm} in the boundary of the curve graph. And assuming \mathcal{B} is tight, there are only finitely many such quasi-axes and ϕ permutes them (again see [Bow08]). Hence, for some power m'' , $\psi\phi^{m''}$ also preserves \mathcal{B} . Choose $m = m'' + p \cdot m'$ so that the translation of length $\psi\phi^m$ along \mathcal{B} is less than or equal to that of $\phi^{m'}$. Both the distance from μ_0 to \mathcal{B} and the translation distance of $\phi^{m'}$ along \mathcal{B} are bounded by the word length of ϕ . Hence, the claims follows.

Choose such m' so that the translation length of $\phi^{m'}$ is large enough to ensure that the geodesic in the curve graph S connecting μ_0 to $\phi^{m'}(\mu_0)$ passes near \mathcal{B} (the curve graph is gromov hyperbolic). Then for every subsurface $R \subsetneq S$, if $d_R(\mu_0, \phi^{m'}(\mu_0))$ is large, then either

$$d_R(\mu_0, F_+) \quad \text{or} \quad d_R(F_-, \phi^{m'}(\mu_0))$$

is large. That is, there is a constant k so that

$$d_R(\mu_0, \phi^{m'}(\mu_0)) \geq \min(d_R(\mu_0, F_+), d_R(F_-, \phi^{m'}(\mu_0))) - k.$$

Using $d_R(F_-, \phi^{m'}(\mu_0)) = d_R(F_-, \mu_0)$ we get

$$\left[d_R(\mu_0, \phi^{m'}(\mu_0)) \right]_{2k} \stackrel{*}{\succ} \left[d_R(\mu_0, F_+) \right]_k + \left[d_R(F_-, \mu_0) \right]_k.$$

Hence,

$$(2) \quad \|\phi\|_{\mathcal{G}} \stackrel{*}{\succ} \|\phi^{m'}\|_{\mathcal{G}} \prec_C \sum_{R \subsetneq S} \left[d_R(\mu_0, \eta(\mu_0)) \right]_{2k} \\ \stackrel{*}{\succ} d_S(\mu_0, \phi^{m'}(\mu_0)) + \sum_{R \subsetneq S} \left[d_R(\mu_0, F_-) \right]_k + \sum_{R \subsetneq S} \left[d_R(\mu_0, F_+) \right]_k.$$

Now, let m be as in (1) and let $\xi = \psi\phi^m$. Further assume k is large enough so that $d_R(F_-, F_+) < k$. Then

$$\|\xi\|_{\mathcal{G}} \stackrel{*}{\succ} d_S(\mu_0, \xi(\mu_0)) + \sum_{R \subsetneq S} \left[d_R(\mu_0, \xi(\mu_0)) \right]_{2k} \\ \stackrel{*}{\succ} d_S(\mu_0, \xi(\mu_0)) + \sum_{R \subsetneq S} \left[d_R(\mu_0, F_+) \right]_k + \left[d_R(\xi(\mu_0), F_+) \right]_k$$

If $\xi(F_+) = F_+$.

$$\sum_{R \subsetneq S} \left[d_R(\xi(\mu_0), F_+) \right]_k = \sum_{R \subsetneq S} \left[d_R(\mu_0, F_+) \right]_k$$

and if $\xi(F_+) = F_-$.

$$\sum_{R \subsetneq S} \left[d_R(\xi(\mu_0), F_+) \right]_k = \sum_{R \subsetneq S} \left[d_R(\mu_0, F_-) \right]_k$$

In either case, the last two terms in estimate above given for $\|\xi\|_{\mathcal{G}}$ are less than the lower-bound Equation (2) given for $\|\phi\|_{\mathcal{G}}$. We also know from Equation (1) that the first term is bounded above by $d_S(\mu_0, \phi(\mu_0))$ which is also bounded by above by a multiple of $\|\phi\|_{\mathcal{G}}$. The Theorem follows. \square

3.3. The general case. Recall from the Nielsen-Thurston classification of surface homeomorphisms [Thu88, FLP79] that there is a normal form for any homeomorphism ϕ of a surface S of finite type. That is,

- (1) There is a multicurve Γ_ϕ that is preserved by ϕ , called the *canonical reducing system*, defined as follows: consider the set \mathcal{A}_ϕ consisting of all curves α so that $\phi^k(\alpha) = \alpha$ up to isotopy, for some $k > 0$ and let Γ_ϕ be the boundary of the subsurface of S that is filled with curves in \mathcal{A}_ϕ . The curve system Γ_ϕ is empty if ϕ is pseudo-Anosov or has finite order.
- (2) The components of $S - \Gamma_\phi$ are decomposed into ϕ -orbits $\{V_1, \dots, V_\ell\}$ where

$$\phi(V_i) = V_{i+1} \quad \text{for} \quad i \in \mathbb{Z}/\ell\mathbb{Z}.$$

- (3) For every every ϕ -orbit $\mathcal{V} = \{V_1, \dots, V_\ell\}$ the first return map $\phi^\ell: V_1 \rightarrow V_1$ is either finite order or pseudo-Anosov.

It is convenient to fix a topological surface δ that is homeomorphic to every V_i . Choosing a homeomorphism $\delta \rightarrow V_i$, the map

$$\phi^\ell|_{V_i}: V_i \rightarrow V_i$$

defines a conjugacy class in $\text{Mod}(\delta)$ that is independent on i or the homeomorphism from δ to V_i . That is, it depends only on the ϕ -orbit \mathcal{V} . We denote this conjugacy class by $[\phi_\mathcal{V}]$. We say the ϕ -orbit \mathcal{V} is *of type δ with the first return map $[\phi_\mathcal{V}]$* .

We start by modifying the generating set \mathcal{G} and conjugating ϕ so that they are compatible with each other.

If we choose $\varphi \in [\phi]$ with $\|\varphi\|_{\mathcal{G}} \stackrel{*}{\prec} \|\phi\|_{\mathcal{G}}$ then, by Theorem 3.3, the conjugating element η satisfies $\|\eta\|_{\mathcal{G}} \stackrel{*}{\prec} \|\phi\|_{\mathcal{G}}$. But η conjugates a generating set for $C(\varphi)$ to a generating set for $C(\phi)$ which means it would be enough to prove the theorem for φ . Our goal is to find a representative of the conjugacy class $[\phi]$ of ϕ which has (as much as possible) a standard form.

There are finitely many topological types possible for subsurfaces of S . Let Δ be the set of surfaces that can be a subsurface of S . That is, for every subsurface V of S , there is a (unique) surface $\delta \in \Delta$ that is homeomorphic to V . We fix a generating set \mathcal{G}_{δ} for every surface $\delta \in \Delta$. In fact, we assume \mathcal{G}_{δ} consists of Dehn twists around a finite set of curves μ_{δ} . Curves in μ_{δ} fill the surface δ , that is every curve in δ intersects a curve in μ_{δ} .

Also, up to a homeomorphism, there are finitely many multicurves on a surface S . Let Λ be a fixed set consisting of a representative for every homeomorphism type of a multicurve in S . For any simple closed curve γ in S , let D_{γ} denote the Dehn twist around γ . For each $\Gamma \in \Lambda$ let μ_{Γ} be a set of curves on S with the following properties.

- (1) $\Gamma \subset \mu_{\Gamma}$.
- (2) The set $\mathcal{G}_{\Gamma} = \{D_{\gamma} \mid \gamma \in \mu_{\Gamma}\}$ generates $\text{Mod}(S)$.
- (3) for every subsurface V that is a component of $S - \Gamma$ that is homeomorphic to δ , there is a homeomorphism $m_V: \delta \rightarrow V$ so that $m_V(\mu_{\delta})$ is exactly the set of curves in μ_{Γ} that are contained in V . In particular, $m_V(\mathcal{G}_{\delta}) \subset \mathcal{G}_{\Gamma}$ generates $\text{Mod}(V)$.

Note that, μ_{Γ} fills S . For the rest of this article, we assume

$$\mathcal{G}_0 = \bigcup_{\Gamma \in \Lambda} \mathcal{G}_{\Gamma} \quad \text{and} \quad \mu_0 = \bigcup_{\Gamma \in \Lambda} \mu_{\Gamma}.$$

Note that $\|\cdot\|_{\mathcal{G}}$ differs from $\|\cdot\|_{\mathcal{G}_0}$ by a uniform multiplicative amount.

We are now ready to construct φ . Let $\Gamma \in \Lambda$ be the curve system in Λ that has the same homeomorphism type as Γ_{ϕ} . Conjugate ϕ to ϕ' by a homeomorphism that sends Γ_{ϕ} to Γ . Then, ϕ' partitions the components of $S - \Gamma$ to ϕ' -orbits similar to ϕ . We then further modify ϕ' to φ whose orbits are the same as the orbits of ϕ' so that, if $\mathcal{V}' = \{V'_1, \dots, V'_\ell\}$ is a ϕ' -orbit of size ℓ associated to the ϕ -orbit \mathcal{V} , then

- (1) for $i = 1, \dots, \ell - 1$, we have

$$m_{V_{i+1}}^{-1} \varphi m_{V_1}: \delta \rightarrow \delta \quad \text{is the identity map}$$

- (2) The map

$$m_{V_1}^{-1} \varphi m_{V_{\ell}}: \delta \rightarrow \delta$$

is the representative of $[\phi_{\mathcal{V}}]$ that has the shortest word length with respect to \mathcal{G}_{δ} .

We can make this canonical by choosing, ahead of time, a representative for every conjugacy class in $\text{Mod}(\delta)$.

Proposition 3.7. *For φ constructed as above, we have*

$$\|\varphi\|_{\mathcal{G}_0} \stackrel{*}{\prec} \|\phi\|_{\mathcal{G}}$$

Proof. The proposition follows from the Masur-Minsky distance formula (Theorem 3.4), which we apply to φ . Since $\varphi(\Gamma) = \Gamma$, for every R that intersects Γ , we have

$$d_R(\mu, \varphi(\mu)) \stackrel{\pm}{\asymp} d_R(\Gamma, \varphi(\Gamma)) = O(1).$$

Hence, choosing k large enough, these terms disappear from the distance formula.

For a surface R , let $\text{Mark}(R)$ be the set of markings on the surface R . For every ϕ -orbit, \mathcal{V}' associated to the ϕ -orbit \mathcal{V} , we have,

$$\begin{aligned}
 \|\varphi^\ell|_{\mathcal{V}}\|_{\mathcal{G}_\delta} &\stackrel{*}{\asymp} \sum_{\rho \subset \delta} \left[d_\rho(\mu_\delta, \phi_{\mathcal{V}}(\mu_\delta)) \right]_k \\
 &\stackrel{*}{\asymp} \min_{\mu \in \text{Mark}(\delta)} \sum_{\rho \subset \delta} \left[d_\rho(\mu, \phi_{\mathcal{V}}(\mu)) \right]_k \\
 &= \min_{\mu \in \text{Mark}(V_1)} \sum_{R \subset V_1} \left[d_R(\mu, \phi^\ell(\mu)) \right]_k \\
 &\stackrel{*}{\prec} \sum_{R \subset V_1} \left[d_R(\mu_0, \phi^\ell(\mu_0)) \right]_k \\
 &\stackrel{*}{\prec} \sum_{R \subset S} \left[d_R(\mu_0, \phi^\ell(\mu_0)) \right]_k \stackrel{*}{\asymp} \|\phi^\ell\|_{\mathcal{G}_0} \leq \ell \|\phi\|_{\mathcal{G}_0}
 \end{aligned}$$

Now, let μ_S be a marking for S associated to the generating set \mathcal{G}_S and let μ_0 be the marking in δ that is the image of the projection of μ_S to V_1 under $m_{V_1}^{-1}$. That is, for every sub-surface R of δ , we have

$$d_{m_{V_1}(R)}(\mu_S, m_{V_1}(\mu_0)) = O(1).$$

We can now compare the word length of $\phi_{\mathcal{V}}$ with that of ϕ^ℓ which send V_1 to itself.

$$\begin{aligned}
 \|\phi_{\mathcal{V}}\|_{\mathcal{G}_\delta} &\prec \sum_{R \subset \delta} \left[d_R(\mu_0, \phi_{\mathcal{V}}(\mu_0)) \right]_k \\
 &\stackrel{*}{\asymp} \sum_{R \subset V_1} \left[d_R(\mu_S, \phi_{\mathcal{V}}(\mu_S)) \right]_k \\
 &\prec \sum_{R \subset S} \left[d_R(\mu_S, \phi^\ell(\mu_S)) \right]_k \asymp \|\phi^\ell\|_{\mathcal{G}_S}
 \end{aligned}$$

where the last inequality is the distance formula in the surface S . The lemma now follows since $\|\phi^\ell\|_{\mathcal{G}_S} \leq \ell \|\phi\|_{\mathcal{G}_S}$ and ℓ and all other related constant are independent of ϕ . \square

Proposition 3.8. *There is a constant M_δ , depending only on the generating sets \mathcal{G}_δ and \mathcal{G} , so that for every $\phi \in \text{Mod}(S)$ that has a ϕ orbit $\mathcal{V} = \{V_1, \dots, V_\ell\}$ of type δ , with the first return map $\phi_{\mathcal{V}}$, we have*

$$\|\phi_{\mathcal{V}}\|_{\mathcal{G}_\delta} \leq M_\delta \|\phi\|_{\mathcal{G}}.$$

Proof. If $\phi_{\mathcal{V}}$ is finite order, then the statement is clear since there are only finitely many conjugacy classes of finite order elements and $\phi_{\mathcal{V}}$ is one of the finitely many fixed representatives of these classes. Hence, $\|\phi_{\mathcal{V}}\|_{\mathcal{G}_\delta}$ is uniformly bounded. However, we give a general argument that works in both cases using the Masur-Minsky distance formula Theorem 3.4.

Note that the changing η by a conjugation is the same as changing the marking μ_δ . Since we have chosen $\phi_{\mathcal{V}}$ to be the representative $[\phi_{\mathcal{V}}]$ with the smallest word length, we have,

$$\|\phi_{\mathcal{V}}\|_{\mathcal{G}_\delta} \prec \sum_{R \subset \delta} \left[d_R(\mu, \phi_{\mathcal{V}}(\mu)) \right]_k \quad \text{for any marking } \mu.$$

Here, \prec means less than up a uniform multiplicative and additive error.

Now, let μ_S be a marking for S associated to the generating set \mathcal{G}_S and let μ_0 be the marking in δ that is the image of the projection of μ_S to V_1 under $m_{V_1}^{-1}$. That is, for every

sub-surface R of δ , we have

$$d_{m_{V_1}(R)}(\mu_S, m_{V_1}(\mu_0)) = O(1).$$

We can now compare the word length of $\phi_{\mathcal{V}}$ with that of ϕ^ℓ which send V_1 to itself.

$$\begin{aligned} \|\phi_{\mathcal{V}}\|_{\mathcal{G}_\delta} &\prec \sum_{R \subset \delta} \left[d_R(\mu_0, \phi_{\mathcal{V}}(\mu_0)) \right]_k \\ &\stackrel{*}{\asymp} \sum_{R \subset V_1} \left[d_R(\mu_S, \phi_{\mathcal{V}}(\mu_S)) \right]_k \\ &\prec \sum_{R \subset S} \left[d_R(\mu_S, \phi^\ell(\mu_S)) \right]_k \asymp \|\phi^\ell\|_{\mathcal{G}_S} \end{aligned}$$

where the last inequality is the distance formula in the surface S . The lemma now follows since $\|\phi^\ell\|_{\mathcal{G}_S} \leq \ell \|\phi\|_{\mathcal{G}_S}$ and ℓ and all other related constant are independent of ϕ . \square

Now, consider an element $\psi \in C(\phi)$. First notice that if $\alpha \in \mathcal{A}_\phi$ then

$$\phi^k(\psi(\alpha)) = \psi(\phi^k(\alpha)) = \psi(\alpha),$$

which means $\psi(\mathcal{A}_\phi) = \mathcal{A}_\phi$. Hence, ψ also preserves the subsurface that is filled with the curves in \mathcal{A}_ϕ . Therefore, $\psi(\Gamma_\phi) = \Gamma_\phi$ and ψ permutes the components of $S - \Gamma_\phi$.

Assume, $\psi(V_1) = W_1$ where V_1 is in a ϕ -orbit \mathcal{V} and W_1 is in a ϕ -orbit \mathcal{W} . We observe that \mathcal{V} and \mathcal{W} are of the same type δ . Also, for every i ,

$$\psi(V_i) = \psi(\phi^i(V_1)) = \phi^i(\psi(V_1)) = \phi^i(W_1) = W_i.$$

Hence, the orbit \mathcal{V} is mapped to the \mathcal{W} , which in particular implies $|\mathcal{V}| = |\mathcal{W}| = \ell$. We further have

$$\psi|_{V_1} \phi|_{V_1}^\ell \psi|_{V_1}^{-1} = \phi|_{W_1}^\ell,$$

where $\psi|_{V_1}$ is the restriction of ψ to V_1 (similarly, $\phi|_{V_1}$ and $\phi|_{W_1}$ are restrictions of ϕ to V_1 and W_1 respectively). That is, $[\phi_{\mathcal{V}}] = [\phi_{\mathcal{W}}]$ which implies $\phi_{\mathcal{V}} = \phi_{\mathcal{W}}$.

To sum up, ψ induces a permutation of ϕ -orbits, however, it can only send an orbit to another orbit if the orbits have the same size, same topological type and if the associated first return maps are the same. Also, ψ has to send adjacent components of $S - \Gamma_\phi$ to adjacent components. To keep track of this information, we consider the a decorated dual graph defined as follows. Let G be a graph whose vertices are components of $S - \Gamma_\phi$ and edges are pairs of adjacent components. We decorate a vertex $V \in G$ (which is a component of $S - \Gamma_\phi$) with the name \mathcal{V} of the associated ϕ -orbit, the topological type δ and the first return map $\phi_{\mathcal{V}} \in \text{Mod}(\delta)$. We say a map $f: G \rightarrow G$ is an automorphism of the decorated graph if

- (1) f is a graph automorphism.
- (2) There is a permutation σ of the ϕ -orbits so that, if

$$f(V, \mathcal{V}, \delta, \phi_{\mathcal{V}}) = (W, \mathcal{W}, \delta', \phi_{\mathcal{W}})$$

then $\delta = \delta'$, $|\mathcal{V}| = |\mathcal{W}|$, $\mathcal{W} = \sigma(\mathcal{V})$ and $\phi_{\mathcal{V}} = \phi_{\mathcal{W}}$.

Note that, the set of automorphisms of the decorated graph G form a group which we denote by $\text{Aut}(G)$. For a given $\psi \in C(\phi)$, we denoted the induced graph map by f_ψ and the induced permutation of ϕ -orbits by σ_ψ . We have a homomorphism

$$\pi_G: C(\phi) \rightarrow \text{Aut}(G)$$

projecting ψ to the induced action f_ψ on the decorated graph G .

For each orbit \mathcal{V} , let

$$C_{\mathcal{V}} \subset \ker(\pi_G)$$

be the set of elements of $C(\phi)$ that fix every subsurface in $S - \Gamma_{\phi}$ whose restriction to any subsurface that is not in \mathcal{V} is identity. Then the $\cup_{\mathcal{V}} C_{\mathcal{V}}$ generates $\ker(\pi_G)$ and the intersection $\cap_{\mathcal{V}} C_{\mathcal{V}}$ is the set of multi-twists around the curves Γ . Also, for $\psi \in C_{\mathcal{V}}$ where $\mathcal{V} = \{V_1, \dots, V_{\ell}\}$,

$$\psi|_{V_{i+1}} = \phi^{-i} \psi|_{V_1} \phi^i.$$

That is, the restriction $\psi|_{V_1}$ determines the restriction of ϕ to every other subsurface in the ϕ -orbit \mathcal{V} . Therefore, considering the homeomorphism

$$\pi_{\mathcal{V}}: C_{\mathcal{V}} \rightarrow \text{Mod}(\delta), \quad \psi \rightarrow m_{V_1}^{-1} \psi|_{V_1} m_{V_1}$$

we have that any $\psi \in C_{\mathcal{V}}$ is determined, up to possibly a multi-twist around Γ , by its projection $\pi_{\mathcal{V}}(\psi)$ to $\text{Mod}(\delta)$ which lies in $C(\phi_{\mathcal{V}})$.

We have shown that every element of $C(\phi)$ is determined, up to a multi-twist around Γ , by its projection to $\text{Aut}(G)$ and to $C(\phi_{\mathcal{V}})$. We now examine which multi-twists around curves in Γ lie in $C(\phi)$. Consider the action of ϕ on Γ . It decomposes Γ into orbits $\bar{\gamma} = \{\gamma_1, \dots, \gamma_j\}$ where $\phi(\gamma_i) = \gamma_{i+1}$ for $i \in \mathbb{Z}/j\mathbb{Z}$. We call such orbit $\bar{\gamma}$ an admissible multicurve if ϕ^j sends γ_1 to γ_1 preserving the orientation. For any admissible multicurve $\bar{\gamma}$, define

$$D_{\bar{\gamma}} = D_{\gamma_1} \dots D_{\gamma_j},$$

where D_{γ_i} is a Dehn twist if γ_i is non-separating and a half-twist if γ_i is separating. That is, $D_{\bar{\gamma}}$ is the product of Dehn twists (or half-twists) around the curves in $\bar{\gamma}$. The set of multi-twists around the curve in Γ that commute with ϕ is generated by $\{D_{\bar{\gamma}}\}_{\bar{\gamma}}$. Note that this maybe an empty set. This is because, if an element $\psi \in C(\phi)$ twists around γ is also has to twist by the same amount around $\phi(\gamma)$. However, if ϕ^j sends γ to itself reserving the orientation, then ϕ^j conjugates D_{γ} to D_{γ}^{-1} . Hence, D_{γ} does not commute with ϕ and no Dehn twists around such γ is possible.

There is no homomorphism back from $\text{Aut}(G)$ or $C(\phi_{\mathcal{V}})$ to $\text{Mod}(S)$. But to find a generating set for $C(\phi)$ it is enough to choose a section. To summarize the above discussion, we have shown:

Summary. Let $\mathcal{G}_{\mathcal{V}}$ be a generating set for $C(\phi_{\mathcal{V}})$ and consider arbitrary sections

$$\pi_G^{-1}: \text{Aut}(G) \rightsquigarrow C(\phi) \quad \text{and} \quad \pi_{\mathcal{V}}^{-1}: C(\phi_{\mathcal{V}}) \rightsquigarrow C_{\mathcal{V}}.$$

Then $C(\phi)$ is generated by the union of the following sets:

- (1) The set $\{D_{\bar{\gamma}}, \gamma \in \Gamma\}$, where $\bar{\gamma} \subset \Gamma$ is an admissible multicurve.
- (2) The image of $\text{Aut}(G)$ under π_G^{-1}
- (3) The images of $\mathcal{G}_{\mathcal{V}}$ under maps $\pi_{\mathcal{V}}^{-1}$.

What remains is to bound the word length of the elements of this generating set. An element in (1) is a product of Dehn twists around a uniformly bounded number of curves and these Dehn twists are already in our generating set. For (2), we build the section to be as close to the identity as possible. Namely, for any $f \in \text{Aut}(G)$ and induced permutation σ , let $\sigma(V_i) = W_j$ where V_i in the ϕ -orbit \mathcal{V} and W_j in a ϕ -orbit \mathcal{W} . We define ψ to be the map that also sends V_i to W_j and so that $m_{W_j}^{-1} \psi m_{V_i}$ is the identity. Then ψ is clearly in $C(\phi)$. For (3), given an element $g \in \mathcal{G}_{\mathcal{V}}$ associated to a ϕ -orbit $\mathcal{V} = \{V_1, \dots, V_{\ell}\}$ there is mapping class ψ , that acts on subsurfaces of $S - \mathcal{A}_{\phi}$ the same way as ϕ , its restriction to \mathcal{V} is the same as ϕ and is the identity on every other orbit. Again, ψ clearly commutes with ϕ . The desired upper-bound for the word length of ψ follows from Proposition 3.8.

4. SELF-EQUIVALENCES OF THURSTON MAPS

The results of the previous section imply the following theorem.

Theorem 4.1. *Let f be either a Thurston map with empty canonical obstruction or a homeomorphism. Then the group $C(f)$ of all self-equivalences of f is finitely generated. Moreover, there is an algorithm that finds a generating set for $C(f)$.*

Proof. I. If f is an unobstructed Thurston map with hyperbolic orbifold then it is equivalent to a rational map (possibly with extra marking) and $C(f)$ is trivial (cf. [DH93, BGL14]). The same argument applies in the case of a Thurston map with parabolic orbifold unless P_f contains exactly 4 points, and f is equivalent to a quotient of an affine map $Ax + b$ with hyperbolic associated matrix A (see [SY15]).

II. Let f be a Thurston map with parabolic orbifold such that $C(f)$ is non-trivial. If $Q_f = P_f$ contains exactly 4 points, then the pure mapping class group of (S^2, Q_f) is isomorphic to the modular group $\Lambda = \mathrm{PGL}(2, \mathbb{Z})/\mathrm{PGL}(2, \mathbb{Z}/2\mathbb{Z})$. In this case, $C(f)$ is the subgroup of Λ of all matrices that commute with A . It consists of the matrices which diagonalize simultaneously with A , and thus its generating set can be easily computed.

If Q_f has more than 4 points, let us denote $C(f, P_f)$ the group of self-equivalences of f with only the points in P_f marked. Clearly, $C(f)$ is isomorphic to a finite index subgroup of $C(f, P_f)$. Indeed, if a self-equivalence ϕ is homotopic to the identity in (S^2, Q_f) it will also be homotopic to the identity in (S^2, P_f) . Therefore every self-equivalence can be represented by an affine homeomorphism. Some elements of $C(f, P_f)$, however, may have affine representatives that do not fix points in Q_f but instead send them to different pre-periodic orbits. Determining which subgroup of $C(f, P_f)$ fixes points in Q_f is a straightforward exercise in linear algebra.

III. If f is a homeomorphism then Corollary 3.2 can be applied. □

5. HURWITZ CLASSIFICATION OF BRANCHED COVERS

Let X and Y be two finite type Riemann surfaces. We recall that two finite degree branched covers ϕ and ψ of Y by X are *equivalent in the sense of Hurwitz* if there exist homeomorphisms $h_0 : Y \rightarrow Y, h_1 : X \rightarrow X$ such that

$$h_0 \circ \phi = \psi \circ h_1.$$

An equivalence class of branched covers is known as a *Hurwitz class*. Enumerating all Hurwitz classes with a given ramification data is a version of the *Hurwitz Problem*. The classical paper of Hurwitz [Hur91] gives an elegant and explicit solution of the problem for the case $X = \hat{\mathbb{C}}$.

We will need the following narrow consequence of Hurwitz's work (for a modern treatment, see [BBGvBK13]):

Theorem 5.1. *There exists an algorithm \mathcal{A} which, given PL branched covers ϕ and ψ of PL spheres and a PL homeomorphism h_0 mapping the critical values of ϕ to those of ψ , does the following:*

- (1) *decides whether ϕ and ψ belong to the same Hurwitz class or not;*
- (2) *if the answer to (1) is affirmative, decides whether there exists a homeomorphism h_1 such that $h_0 \circ \phi = \psi \circ h_1$.*

6. EQUIVALENCE ON THICK PARTS

6.1. Canonical obstructions and thin-thick decompositions of Thurston maps.

Let f be a Thurston map, and $\Gamma = \{\gamma_j\}$ an f -stable multicurve. Consider a finite collection

of disjoint closed annuli $A_{0,j}$ which are homotopic to 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 . Following Pilgrim, we say that the pair (f, Γ) is in a *standard form* (see Figure 1) if there exists a collection of annuli $A_{0,j}$, which we call *decomposition annuli*, 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 $A_{0,j}$.

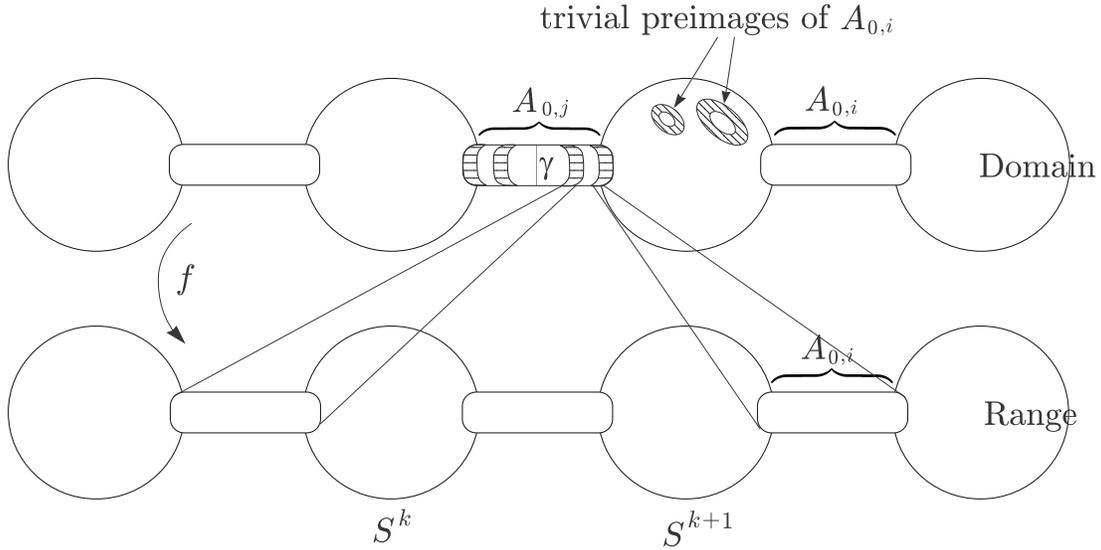


FIGURE 1. Pilgrim's decomposition of a Thurston map

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 component $S_1(j)$ is a sphere with holes, with boundary curves being of two types: boundaries of the removed annuli, or boundaries of trivial preimages of the removed annuli.

The holes in $S_0(j) \subset S^2$ can be filled as follows. Let χ be a boundary curve of a component D of $S^2 \setminus S_0(j)$. Let $k \in \mathcal{N}$ be the first iterate $f^k : \chi \rightarrow \chi$, 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 $f : \chi_i \rightarrow \chi_{i+1}$. Select homeomorphisms

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

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

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

We have thus replaced every hole with a cap with a single puncture. We call such a procedure *patching* a component.

By construction, the map

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

contains a finite number of periodic cycles of punctured spheres. For every periodic sphere $\tilde{S}(j)$ denote by \mathcal{F} the first return map $\tilde{f}^{k_j} : \tilde{S}(j) \rightarrow \tilde{S}(j)$. This is again a Thurston map or a homeomorphism. The collection of maps \mathcal{F} and the combinatorial information required to glue the spheres $S_0(j)$ back together is what Pilgrim called a *decomposition* of f along Γ ; we will denote it \mathcal{S}_Γ .

Pilgrim showed:

Theorem 6.1. *For every obstructed marked Thurston map f with an obstruction Γ there exists an equivalent map g such that (g, Γ) is in a standard form, and thus can be decomposed.*

Pilgrim [Pil01] defined a *canonical decomposition* of a Thurston map f based on his definition of a canonical Thurston obstruction. His original definition was framed in the language of iteration on a Teichmüller space; we will give an equivalent definition discovered by the second author [Sel13]:

Theorem 6.2. *Suppose f is an obstructed Thurston mapping. Then there exists a unique minimal (with respect to inclusion) obstruction Γ_f , which is the canonical obstruction in the sense of [Pil01], with the following properties.*

- *If a first-return map \mathcal{F} of a cycle of components in \mathcal{S}_{Γ_f} is a $(2, 2, 2, 2)$ -map, then every curve of every simple Thurston obstruction for \mathcal{F} has two postcritical points of f in each complementary component and the two eigenvalues of $\tilde{\mathcal{F}}_*$ are equal or non-integer.*
- *If the first-return map \mathcal{F} of a cycle of components in \mathcal{S}_{Γ_f} is not a $(2, 2, 2, 2)$ -map nor a homeomorphism, then there exists no Thurston obstruction for \mathcal{F} .*

Definition 6.3. If Γ_f is the canonical obstruction, then the decomposition \mathcal{S}_{Γ_f} is the *canonical decomposition* of f . In this case, we call the components of the complement of the decomposition annuli the *thick parts*, and the decomposition annuli themselves the *thin parts*.

From this point, only canonical decompositions of Thurston maps will be considered.

Definition 6.4. By *equivalence on thick parts* ϕ between f and g we mean a homeomorphism defined on the union of patched thick parts of f onto the union of patched thick parts of g such that the following holds:

- Denote ϕ_W the restriction of ϕ to any patched thick component W . If X is a periodic patched thick component of f then $Y = \phi(X)$ is periodic for \tilde{g} with the same period. If $\mathcal{F}_X, \mathcal{G}_Y$ denote the first return maps of X and Y respectively, then ϕ_X is an equivalence of \mathcal{F}_X and \mathcal{G}_Y .
- Let X be a periodic patched thick component and let W be a preimage of X so that

$$\tilde{f}^n(W) = X.$$

Denote $Y = \phi_X(X)$ and $Z = \phi(W)$. Then $\tilde{g}^n(Z) = Y$ and ϕ_W is a lift of ϕ_X through actions of \tilde{f} and \tilde{g} :

$$\phi_X \circ \tilde{f}^n = \tilde{g}^n \circ \phi_W.$$

6.2. Centralizer on thick parts. For each periodic patched thick component $X = \mathcal{F}(X)$ of the canonical decomposition of f , denote $C_X(f) \subset \text{PMCG}(X)$ the group of self-equivalences of the first return Thurston mapping $\mathcal{F}|_X : X \rightarrow X$.

Definition 6.5. We define the *centralizer on thick parts* $C_{\text{thick}}(f)$ of f to be the group of all self-equivalences of f on thick parts. By the previous definition, $C_{\text{thick}}(f)$ is isomorphic to the subgroup of the free abelian product

$$C_{\text{periodic}}(f) \equiv \coprod_{\text{periodic components } X} C_X(f)$$

consisting of all elements ϕ such that for every thick patched component W with $\tilde{f}^n(W) = X$, one can define ϕ_W so that $\phi_X \circ \tilde{f}^n = \tilde{f}^n \circ \phi_W$ (that is, ϕ can be lifted via the action of \tilde{f} to all strictly pre-periodic preimages of X).

Note that since all C_X are finitely generated (Theorem 4.1) and $C_{\text{thick}}(f)$ is a subgroup of finite index, $C_{\text{thick}}(f)$ is also finitely generated. Furthermore,

Lemma 6.6. *A generating set of $C_{\text{thick}}(f)$ can be computed explicitly.*

Proof. By Theorem 4.1, for each periodic component X , a generating set A of $C_{\text{periodic}}(f)$ can be computed explicitly. Given the topological complexity of the covering maps $\tilde{f}^n : W \rightarrow X$ for all thick preimages of periodic components, it is straightforward to obtain an upper bound on the word length (in terms of the elements of A) of the generating set of $C_{\text{thick}}(f)$. By Theorem 5.1, we can verify algorithmically, which of the words, whose length is under this bound, correspond to elements of $C_{\text{thick}}(f)$. \square

Consider two equivalences on thick parts ϕ and ψ between two Thurston maps f and g . Then $\phi^{-1} \circ \psi$ is a self-equivalence of f . This yields the following.

Lemma 6.7. *Let ϕ be an equivalence on thick parts between two Thurston maps f and g . Then any other equivalence can be written $\phi \circ l$ where $l \in C_{\text{thick}}(f)$.*

7. ALGORITHMIC GEOMETRIZATION OF THICK PARTS

The second and third authors proved the following [SY15, Theorem 6.1]:

Theorem 7.1 (Canonical geometrization). *There exists an algorithm which for any Thurston map f finds its canonical obstruction Γ_f .*

Furthermore, let \mathcal{F} denote the collection of the first return maps of the canonical decomposition of f along Γ_f . Then the algorithm outputs the following information:

- *for every first return map with a hyperbolic orbifold, the unique (up to Möbius conjugacy) marked rational map equivalent to it;*
- *for every first return map of type $(2, 2, 2, 2)$ the unique (up to affine conjugacy) affine map of the form $z \mapsto Az + b$ where $A \in \text{SL}_2(\mathbb{Z})$ and $b \in \frac{1}{2}\mathbb{Z}^2$ with marked points which is equivalent to f after quotient by the orbifold group G ;*
- *for every first return map which has a parabolic orbifold not of type $(2, 2, 2, 2)$ the unique (up to Möbius conjugacy) marked rational map map equivalent to it, which is a quotient of a complex affine map by the orbifold group.*

8. EXTENDING EQUIVALENCE FROM THICK TO THIN PARTS

The following is standard (see e.g. [FM]):

Proposition 8.1. *For every Thurston obstruction $\Gamma = \{\alpha_1, \dots, \alpha_n\}$, the Dehn twists T_{α_j} , $j = 1 \dots n$ generate a free Abelian subgroup of $\text{PMCG}(S \setminus Q_f)$.*

We write $\mathbb{Z}^\Gamma \simeq \mathbb{Z}^n$ to denote the subgroup generated by T_{α_j} . We will need the following straightforward generalization of [SY15, Proposition 7.7]:

Proposition 8.2. *Let f, g be equivalent Thurston maps. Let the pair (ϕ_1, ϕ_2) realize the equivalence of the thick components of f and g . Extend ϕ_1 to a homeomorphism of the whole sphere $S^2 \setminus Q_f$, defining it on the thin parts in an arbitrary fashion. Then there exist $m \in \mathbb{Z}^\Gamma$, $\psi \in C_{thick}$, and an equivalence pair (h_1, h_2) for f, g such that $h_1 = \phi_1 \circ \psi \circ m$.*

Notice that if $h_1 \circ f = g \circ h_2$ where $h_1 = \phi_1 \circ m_1$ for some $m_1 \in \mathbb{Z}^\Gamma$, then h_2 is homotopic to $\phi_1 \circ m_2$ for some other $m_2 \in \mathbb{Z}^\Gamma$. If $m_1 = m_2$ then h_1 is homotopic to h_2 and these two homeomorphisms realize an equivalence between f and g . Since we cannot check whether this happens for all elements of \mathbb{Z}^Γ we will require the following proposition [SY15, Proposition 7.8]:

Proposition 8.3. *There exists explicitly computable $N \in \mathbb{N}$ such that if $n \in \mathbb{Z}^\Gamma$ where all coordinates of n are divisible by N , then*

$$(\phi_1 \circ (m_1 + n)) \circ f = g \circ (\phi_2 \circ (m_2 + M_\Gamma n)),$$

whenever

$$(\phi_1 \circ m_1) \circ f = g \circ (\phi_2 \circ m_2).$$

9. CHECKING THURSTON EQUIVALENCE.

We are now ready to present an algorithm which checks whether two Thurston maps f and g are equivalent or not.

Algorithm.

- (1) Find the canonical obstructions $\Gamma_f = \{\alpha_1, \dots, \alpha_n\}$ and $\Gamma_g = \{\beta_1, \dots, \beta_n\}$ (Theorem 7.1).
- (2) Check whether the cardinality of the canonical obstructions $\Gamma_f = \{\alpha_1, \dots, \alpha_n\}$ and $\Gamma_g = \{\beta_1, \dots, \beta_n\}$ is the same, and whether the corresponding Thurston matrices coincide. If not, output **maps are not equivalent** and halt.
- (3) Denote the thin parts (decomposition annuli) of f and g by A_i and B_i respectively. Construct the first return maps \mathcal{F} and \mathcal{G} of the periodic patched thick parts for f and g and geometrize them (Theorem 7.1). Are the geometrizations of \mathcal{F} and \mathcal{G} the same up to reordering of the components of the first return map? If not, output **maps are not equivalent** and halt.
- (4) **for all** permutations $\sigma \in S_n$ **do**
 - (5) Is there a homeomorphism

$$h_\sigma : S^2 \setminus Q_f \rightarrow S^2 \setminus Q_g$$

sending $A_i \rightarrow B_{\sigma(i)}$? If not, **continue**.

- (6) Is it true that for every periodic patched thick component X of f the geometrization of $\mathcal{F}|_X$ is the same as the geometrization of $\mathcal{G}_{h_\sigma(X)}$? If not, **continue**.
- (7) For all thick components C_j^f check whether the Hurwitz classes of the patched coverings

$$\tilde{f} : \widetilde{C_j^f} \rightarrow \widetilde{f(C_j^f)} \text{ and } \tilde{g} : \widetilde{h_\sigma(C_j^f)} \rightarrow \widetilde{g(h_\sigma(C_j^f))}$$

are the same (Theorem 5.1). If not, **continue**.

- (8) Construct equivalence pairs (η_0^X, η_1^X) between first return maps \mathcal{F}_X and $\mathcal{G}_{h_\sigma(X)}$ of periodic patched thick components corresponding by h_σ and the group $C_{\text{periodic}}(g)$ of self-equivalences of \mathcal{G} . If the maps of some pair are not equivalent, **continue**.
- (9) Find an equivalence between first return maps \mathcal{F}_X and $\mathcal{G}_{h_\sigma(X)}$ in the form $\phi_X = \psi \circ \eta_0^X$ with $\psi \in C_{\text{periodic}}(g)$ that can be lifted via branched covers \tilde{f} and \tilde{g} to every preimage of every thick component and preserves the set of marked points. Since C_{thick} is a finite index subgroup of $C_{\text{periodic}}(g)$, this is a finite check (for representatives of each coset), which can be carried out algorithmically by (Theorem 5.1 and Lemma 6.7). If not possible, **continue**.
- (10) Lift the equivalences, to obtain a homeomorphism ϕ_1 defined on all thick parts.
- (11) Compute $C_{\text{thick}}(g)$ (Lemma 6.7).
- (12) Pick some initial homomorphisms $a_i : A_i \rightarrow B_{\sigma(i)}$ so that the boundary values agree with ϕ_1 . This defines ϕ_1 on the whole sphere.
- (13) Find the set of vectors $m_1 \in \mathbb{Z}^\Gamma$ with coordinates between 0 and $N - 1$, where N is as in Proposition 8.3 such that $h_1 = \phi_1 \circ m_1$ lifts through f and g so that

$$(\phi_1 \circ m_1) \circ f = g \circ h_2.$$

For all vectors m_1 in this set **do**

- (14) By the discussion above $h_2 = m_2 \circ \phi_2$ with $m \in \mathbb{Z}^\Gamma$. Compute m_2 .
- (15) Find the finite index subgroup G_1 of $C_{\text{thick}}(f)$ of all elements ψ such that $\psi \circ h_1$ lifts through f and g (Lemma 6.6).
- (16) For every $\psi \in G_1$ we have

$$\psi \circ h_1 \circ f = g \circ n_\psi \circ \psi \circ h_2$$

where $n_\psi \in \mathbb{Z}^\Gamma$. The map $\psi \mapsto n_\psi - m_2$ is a homomorphism.

- (17) Similarly, find the finite index subgroup G_2 of \mathbb{Z}^Γ of all elements k such that $k \circ h_1$ lifts through f and g . For every $k \in G_1$ we have

$$k \circ h_1 \circ f = g \circ n_k \circ h_2$$

where $n_k \in \mathbb{Z}^\Gamma$. The map $k \mapsto n_k - m_2$ is also a homomorphism (linear).

- (18) Using generators of $C_{\text{thick}}(g)$ construct ψ and k such that $k + m_1 = n_k + n_\psi + m_2$. If $m_1 - m_2$ is not in the image of $n_k + n_\psi - k$, **continue**.
- (19) Output **maps are equivalent** and $\psi \circ k \circ h_1$; halt.

(20) **end do**

(21) **end do**

(22) output **maps are not equivalent** and halt.

If the algorithm exits on step 17, then $\phi \circ h_0$ realizes the equivalence between f and g , by construction. Otherwise, no such equivalence exists, by Proposition 8.2, and thus the above algorithm satisfies the conditions of our main theorem.

REFERENCES

- [BBGvBK13] L. Bartholdi, X. Buff, H.-C. Graf von Bothmer, and J. Kröcker, *Algorithmic construction of Hurwitz maps*, e-print arXiv:1303.1579 (2013).

- [BD17] Laurent Bartholdi and Dzmityr Dudko, *Algorithmic aspects of branched coverings*, Ann. Fac. Sci. Toulouse Math. (6) **26** (2017), no. 5, 1219–1296. MR 3746628
- [BGL14] X. Buff, C. Guizhen, and T. Lei, *Teichmüller spaces and holomorphic dynamics*, Handbook of Teichmüller theory. Volume IV, IRMA Lect. Math. Theor. Phys., vol. 19, Eur. Math. Soc., Zürich, 2014, pp. 717–756.
- [Bow08] B.H. Bowditch, *Tight geodesics in the curve complex*, Invent. Math. **171** (2008), no. 2, 281–300. MR 2367021
- [DH93] A. Douady and J.H. Hubbard, *A proof of Thurston’s topological characterization of rational functions*, Acta Math. **171** (1993), 263–297.
- [FLP79] A. Fathi, F. Laudenbach, and V. Poénaru, *Travaux de Thurston sur les surfaces*, Astérisque, vol. 66-67, Société Mathématique de France, 1979.
- [FM] B. Farb and D. Margalit, *A Primer on Mapping Class Groups*, Princeton University Press.
- [FM12] ———, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012. MR 2850125 (2012h:57032)
- [Hur91] A. Hurwitz, *Ueber Riemann’sche Flächen mit gegebenen Verzweigungspunkten*, Math. Ann. **39** (1891), 1–60.
- [McC94] J. McCarthy, *Normalizers and centralizers of pseudo-anosov mapping classes*, preprint, 1994.
- [MH75] C. Maclachlan and W. J. Harvey, *On mapping-class groups and Teichmüller spaces*, Proc. London Math. Soc. (3) **30** (1975), no. part 4, 496–512. MR 0374414
- [MM00] H.A. Masur and Y. N. Minsky, *Geometry of the complex of curves. II. Hierarchical structure*, Geom. Funct. Anal. **10** (2000), no. 4, 902–974.
- [Pil01] K. Pilgrim, *Canonical Thurston obstructions*, Adv. Math. **158** (2001), no. 2, 154–168.
- [Sel13] N. Selinger, *Topological characterization of canonical Thurston obstructions*, Journal of Modern Dynamics **7** (2013), 99 – 117.
- [SY15] N. Selinger and M. Yampolsky, *Constructive geometrization of Thurston maps and decidability of Thurston equivalence*, Arnold Math. J. **1** (2015), 361–402.
- [Tao13] J. Tao, *Linearly bounded conjugator property for mapping class groups*, Geom. Funct. Anal. **23** (2013), no. 1, 415–466. MR 3037904
- [Thu88] W. P. Thurston, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. (N.S.) **19** (1988), no. 2, 417–431.