THURSTON EQUIVALENCE TO A RATIONAL MAP IS DECIDABLE

SYLVAIN BONNOT, MARK BRAVERMAN, MICHAEL YAMPOLSKY

ABSTRACT. We demonstrate that the question whether or not a given topological ramified covering map of the 2-sphere is Thurston equivalent to a rational map is algorithmically decidable.

1. INTRODUCTION

This paper solves a long-standing open problem in one-dimensional Complex Dynamics: we show that Thurston's equivalence to a post-critically finite rational map is algorithmically decidable. Thurston's theorem [3] is central to the subject. In the case when a rational mapping exists, it is essentially unique, and the proof of the theorem [3] supplies an iterative algorithm for computing its coefficients. When, for instance, the rational mapping is a quadratic polynomial, the Spider algorithm of Hubbard and Schleicher [6], computes the coefficients starting from a convenient combinatorial description of the branched covering. However, the algorithm will go astray if the branched covering data cannot be realized by a polynomial. Thus the question we answer in this work is both natural and important.

Here is the Equivalence problem we consider in this note:

Problem: Equivalence to a Rational Map. *Given a piecewise linear post-critically finite ramified covering map f from $\hat{\mathbb{C}}$ to itself, determine whether or not it is Thurston equivalent to a rational map. If it is equivalent, give an algorithm to compute the coefficients of the corresponding rational map (defined up to a conjugacy with a Möbius map).*

Our main result is:

Theorem 1.1. *The problem of equivalence to a rational map is algorithmically solvable.*

Date: May 9, 2010.

2. Thurston mappings

In this section we recall the basic setting of Thurston's characterization of rational functions.

2.1. **Ramified covering maps.** 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.*

Thurston equivalence. 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;

(2) the diagram

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

commutes.

Orbifold of a Thurston map. 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:

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

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.

Thurston linear transformation. We recall that a simple closed curve $\gamma \subset S^2 - P_f$ 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, 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 (meaning inessential or peripheral) 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 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*.

Several examples. Let us first give an example of a quadratic rational map f with an f-stable multicurve Γ. The map is given by the formula:

$$f(z) = \frac{z^2 + c}{z^2 - 1}, \ \text{ with } c = \frac{1 + i\sqrt{3}}{2}.$$

The picture of its Julia set is seen in Figure 1, it is popularized as the cover art of the Stony Brook preprint series. The map f is known as the *mating* of two quadratic Julia sets: Douady's rabbit and the basilica (see e.g. [12]).

The two critical points of f are $a_1 = 0$ and $b_1 = \infty$. Both of them are periodic:

$$a_1 = 0 \overset{f}{\mapsto} a_2 \overset{f}{\mapsto} a_3 \overset{f}{\mapsto} a_1,$$

$$b_1 = \infty \overset{f}{\mapsto} b_2 = 1 \overset{f}{\mapsto} b_1.$$
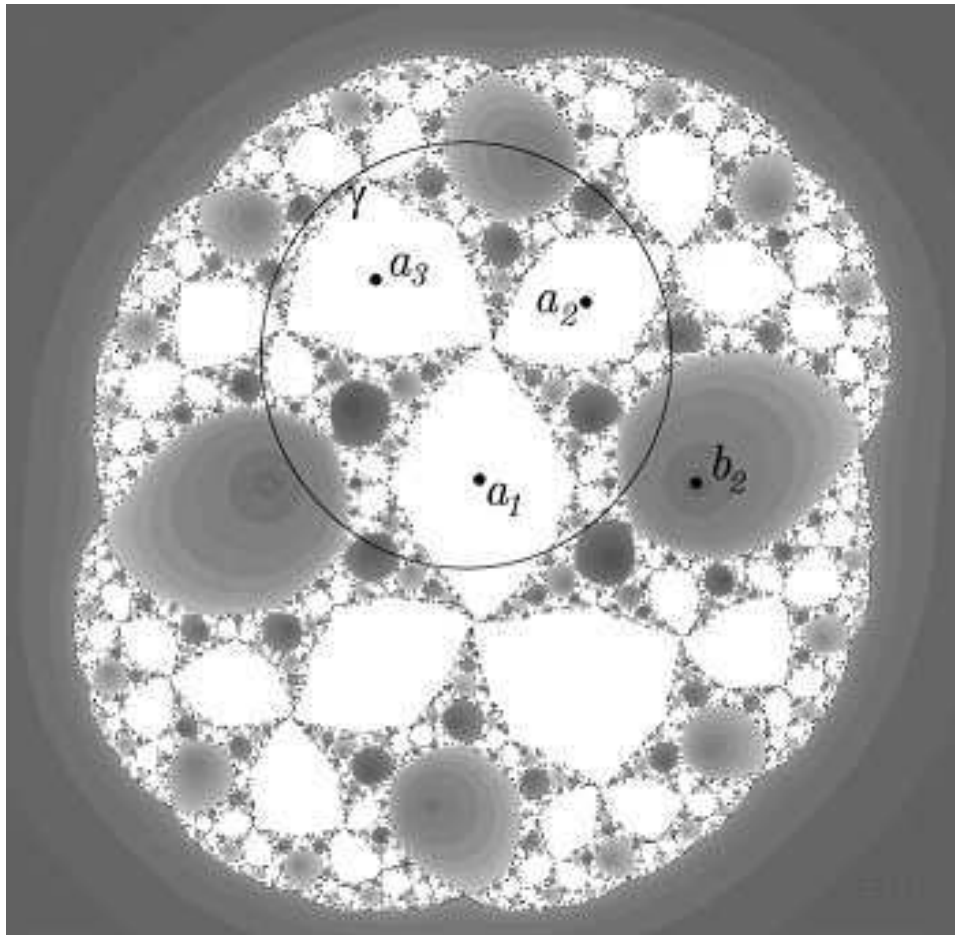
FIGURE 1. A stable multicurve of a quadratic rational map.

Our stable multicurve Γ consists of a single simple closed curve γ which separates a_i's from b_i's. It is easy to see that the corresponding transition matrix consists of a single entry $1/2$. Thus, $\lambda(\Gamma) = 1/2$.

To give an example of a Thurston obstruction, we will need to work a little harder. We again use the procedure known as *mating*. Let us again start with Douady's rabbit polynomial, $f_c(z) = z^2 + c$ which is the unique quadratic polynomial with $\mathrm{Im}\, c > 0$ such that the critical point 0 is peridic with period 3. Thus, the postcritical set of f_c is $\{w_0 = 0, w_1 = c, w_2 = c^2 + c, w_3 = \infty\}$. Consider also the complex conjugate, the polynomial $f_{\bar{c}}$ whose postcritical set we denote $\{w'_0 = 0, w'_1 = \bar{c}, w'_2 = \bar{c}^2 + \bar{c}, w'_3 = \infty\}$. The *formal mating* of these two polynomials is the branched covering mapping of S^2 which is obtained as follows. We first compactify the complex plane by adjoining a circle of directions
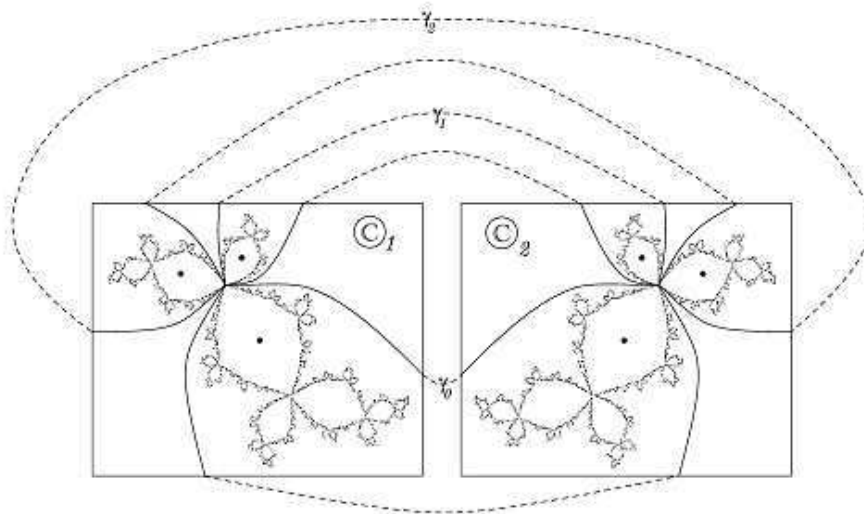
FIGURE 2. An example of Thurston obstruction.

at infinity, $\{\infty \cdot e^{2\pi i\theta}|\ \theta \in \mathbb{R}/\mathbb{Z}\}$. We denote such compactification with the natural topology by $\overline{\mathbb{C}}$. Let us now glue two copies $\overline{\mathbb{C}}_1$, $\overline{\mathbb{C}}_2$ along the circles at infinity using the equivalence relation $\sim_\infty$ given by

$$\infty \cdot e^{2\pi i\theta} \in \overline{\mathbb{C}}_1 \sim_\infty \infty \cdot e^{-2\pi i\theta} \in \overline{\mathbb{C}}_2.$$

Evidently,

$$S = \overline{\mathbb{C}}_1 \sqcup \overline{\mathbb{C}}_2 / \sim_\infty \simeq S^2.$$

The formal mating of f_c and $f_{\bar{c}}$ is the well-defined branched covering map F of the 2-sphere S which is given by f_c on $\overline{\mathbb{C}}_1$ and $f_{\bar{c}}$ on $\overline{\mathbb{C}}_2$. By construction, this map has an invariant equator (the two circles at ∞ glued together), and its postcritical set is the union

$$\{w_0, w_1, w_2\} \cup \{w_0', w_1', w_2'\}.$$

An obstruction for this mapping is given by a multicurve Γ consisting of three loops γ_i separating w_i, w_i' from the rest of the postcritical set (see Figure 2). It is easy to see that Γ is an F-stable multicurve, with the associated transition matrix

$$\begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}$$

so that $\lambda(\Gamma) = 1$.

2.2. A piecewise-linear Thurston mapping.

For the purposes of an algorithmic analysis, we will require a finite description of a branched covering $f : S^2 \to S^2$.

Since we will work mainly in the piecewise linear category, it is convenient to recall here some definitions.

Simplicial complexes. Following [17] (chapter 3.2 and 3.9) we call *a simplicial complex* any locally finite collection Σ of simplices satisfying the following two conditions:

- a face of a simplex in Σ is also in Σ, and
- the intersection of any two simplices in Σ is either empty or a face of both.

The union of all simplices in Σ is the *polyhedron* of Σ (written $|\Sigma|$).

Piecewise linear maps. A map $f : M \to N$ from a subset of an affine space into another affine space is *piecewise linear (PL)* if it is the restriction of a simplicial map defined on the polyhedron of some simplicial complex.

We also define *piecewise linear (PL) manifolds* as manifolds having an atlas where the transition maps between overlapping charts are piecewise linear homeomorphisms between open subsets of $\mathbb{R}^n$. It is well known that any piecewise linear manifold has a triangulation: there is a simplicial complex Σ together with a homeomorphism $|\Sigma| \to X$ which is assumed a PL map (see [17], proof of theorem 3.10.2).

One example of such a manifold is the standard piecewise linear (PL) 2-sphere, which is nicely described in [17] as follows: pick any convex 3-dimensional polyhedron $K \subset \mathbb{R}^3$, and consider the charts corresponding to all the possible orthogonal projections of the boundary (topological) sphere ∂K onto hyperplanes in $\mathbb{R}^3$. The manifold thus obtained is the *standard piecewise linear 2-sphere.* One can prove that another choice of polyhedron would lead to an isomorphic object (see exercise 3.9.5 in [17]).

It is known that in dimension three or lower, every topological manifold has a PL structure, and any two such structures are PL equivalent (in dimension 2, see [14], for the dimension 3 consult [1]).

Piecewise linear branched covers. We begin by formulating the following proposition which describes how to lift a triangulation by a PL branched cover (see [4],section 6.5.4):

Proposition 2.1 (**Lifting a triangulation**). *Let B be a compact topological surface, $\pi : X \to B$ a finite ramified cover of B. Let Δ be the set of branch points of π, and let $\mathcal{T}$ be a triangulation of B such that Δ is a subset of vertices of $\mathcal{T}$ ($\Delta \subset K_0(\mathcal{T})$ in the established notation). Then there exists a triangulation $\mathcal{T}'$ of X, unique up to a bijective*

change of indices, so that the branched covering map $\pi : X \to B$ sends vertices to vertices, edges to edges and faces to faces. Moreover, if $X = B$ is a standard PL 2-sphere and π is PL, then $\mathcal{T}'$ can be produced constructively given a description of $\mathcal{T}$.

We consider PL maps of the standard PL 2-sphere which are topological branched coverings with a finite number of branch points. We call such a map a *piecewise linear Thurston mapping*.

Remark 2.1. Note that any such covering may be realized as a piecewise-linear branched covering map of a triangulation of $\mathbb{C}$ with rational vertices. An algorithmic description of a PL branched covering could thus either be given by the combinatorial data describing the simplicial map, or as a collection of affine maps of triangles in $\hat{\mathbb{C}}$ with rational vertices. We will alternate between these descriptions as convenient.

We note:

Proposition 2.2. *Every Thurston mapping f is Thurston equivalent to a PL Thurston mapping.*

Before proving the above Proposition, let us formulate a basic topological fact, known as Alexander's trick:

Alexander's trick. *Two homeomorphisms of the closed n-dimensional ball, which are isotopic on the boundary, are isotopic.*

Proof of Proposition 2.2. We may start with a triangulation $\mathcal{T}_1$ of S^2 whose vertices include the postcritical set P_f. Refining the triangulation $\mathcal{T}_1$ to $\mathcal{T}_2$, if necessary, we isotope f to a map which leaves the vertices and the edges of $\mathcal{T}_2$ invariant. Finally, every topological map from a triangle to a triangle can be isotoped into a simplicial map using Alexander's trick. We can thus further isotope our map to a PL Thurston mapping with triangulation $\mathcal{T}_2$. □

3. OUTLINE OF THE PROOF OF THEOREM 1.1

The proof will rely on a construction of two explicit algorithms, A_1 and A_2, which, given a postcritically finite piecewise linear branched covering $f : S^2 \to S^2$ with a hyperbolic orbifold, perform the following tasks:

A_1 If f has a Thurston obstruction, the algorithm A_1 will terminate and output the obstruction. It will not terminate otherwise.

A_2 If f is Thurston equivalent to a rational mapping R, then the algorithm A_2 will terminate. It will identify the rational mapping by outputting a ball in an appropriate parameter space

of rational maps which isolates the rational mapping R from postcritically finite mappings of the same degree and with the same size of the postcritical set. If f is not equivalent to any rational mapping, then the algorithm A_2 will not terminate.

We further will use a polynomial root-finding algorithm A_3 which finds an isolated root $\bar{x}_* \in \mathbb{R}^m$ of a system of polynomial equations $\{P_i(\bar{x}) = 0\}$.

A_3 the input of the algorithm is: a system of polynomial equations $\{P_i(\bar{x}) = 0\}$ for $\bar{x} \in \mathbb{C}^m$ (the coefficients of the polynomials P_i are either given through an oracle, or computed with an arbitrarily high precision via a given algorithm); a rational ball $B(\bar{w}, r) \subset \mathbb{C}^m$ which contains $\bar{x}_*$ and such that $B(\bar{w}, 2r)$ does not contain any other roots; a natural number n. The output is $\bar{d}_n \in \mathbb{C}^m$ with the property $||\bar{x}_* - \bar{d}_n|| < 2^{-n}$.

Proof of Theorem 1.1 assuming the existence of A_1 and A_2. Given a post-critically finite piecewise linear map f with a hyperbolic orbifold, we will run the two algorithms A_1 and A_2 in parallel. One and only one of them will terminate. If it is A_1, then we conclude that f is not equivalent to any rational map. If it is A_2 then we know that a Thurston equivalent rational map R exists, and we are given an isolating neighborhood for it in the parameter space. The root-finding algorithm A_3 can then be employed to find the coefficients of R with any given precision. $\square$

4. SOME TOPOLOGICAL PRELIMINARIES

Simple closed curves. Recall that two curves are in a *minimal position* if they realize the minimal number of intersections in their homotopy classes.

Lemma 4.1 (The Bigon Criterion). *Two transverse simple closed curves on a surface S are in a minimal position if and only if the two arcs between any pair of intersection points do not bound an embedded disk in S.*

Let us also formulate an elementary fact:

Lemma 4.2. *Two simple closed curves on a surface S are homotopic if and only if they can be isotoped to boundary curves of an annulus.*

We now prove:

Proposition 4.3. *There exists an algorithm to check whether two simple closed polygonal curves on a triangulated surface S are homotopic.*

Proof. The algorithm works as follows:

(I) If necessary, isotope the curves so that all the intersections are transverse.
(II) While there exists a pair of intersection points which bounds a disk **do**:
        push one of the curves through the disk to remove the two intersection points. **end do**
(III) Does there exist an intersection point? If yes, output **the curves are not homotopic** and halt. If no, proceed to step (IV).
(IV) Do the two curves bound an annulus? If no, output **the curves are not homotopic** and halt. If yes, output **the curves are homotopic** and halt.

To verify the algorithm, we note that the Bigon Criterion implies that step (II) can be performed until the curves are in a minimal position. The correctness of the algorithm now follows by Lemma 4.2.

$\square$

Maps isotopic to the identity. The following theorem of Ladegaillerie [8] will be useful to us in what follows:

Theorem 4.4. *Let K be a compact topological 1-complex, X an oriented compact surface with boundary, i_0, i_1 two embeddings of K into the interior of X. There is an equivalence between the two following properties:*

(1) *i_0 and i_1 are isotopic by an ambient isotopy of X (fixed on ∂X)*
(2) *i_0 and i_1 are homotopic and there is an orientation preserving homeomorphism $h : X \to X$ such that $h \circ i_0 = i_1$.*

We formulate the following corollary:

Proposition 4.5. *There exists an algorithm A which does the following. Given a triangulated sphere with a finite number of punctures $S = S^2 - Z$ and a triangulated homeomorphism $h : S \to S$, the algorithm identifies whether h is isotopic to the identity.*

Proof. Let $x \notin Z$ be a vertex in the triangulation $\mathcal{T}$. Consider a collection of closed loops γ_i in $\partial \mathcal{T}$ passing through the basepoint x such that $\{\gamma_i\}$ forms a basis of $\pi_1(S)$ (refine the triangulation, if necessary). By Theorem 4.4, it is sufficient to verify that $h(\gamma_i)$ is homotopic to γ_i for all i. Indeed, this is equivalent to the existence of a global isotopy of S which moves $h(\gamma_i)$ to γ_i. By the Alexander's trick, the latter statement means that h is isotopic to the identity.
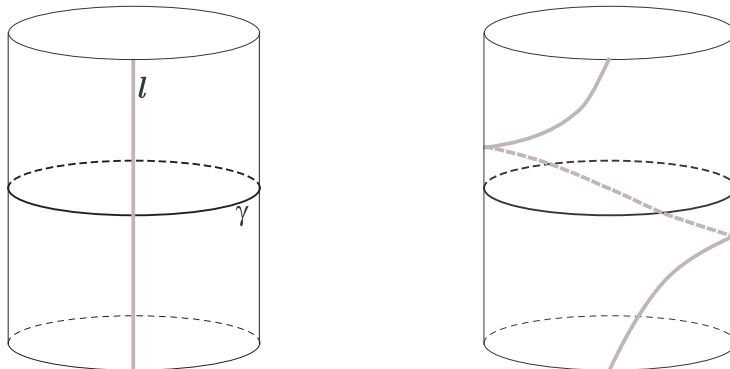
$\square$

FIGURE 3. Dehn twist of a cylinder about an equatorial
curve γ: a vertical line l before (left) and after (right)
the twist is shown.

Dehn twists. Recall the definition of a Dehn twist. Let γ be a simple
closed curve on a surface S, and let A be a tubular neighborhood of γ.
Choose a homeomorphism

$$h : S^1 \times [0,1] \to A,$$

which endows the annulus A with a coordinate system (θ, r) where
$\theta \in \mathbb{R}/\mathrm{mod}\, 2\pi\mathbb{Z}$ is the angular coordinate in S^1, and $r \in [0,1]$. A
Dehn twist about γ is the homeomorphism

$$f : S \to S$$

which is identical outside A, and is given by

$$f : h(\theta, r) \mapsto h(\theta + 2\pi r, r).$$

Mapping class group. Since Thurston equivalence involves isotopies
preserving pointwise the points of P_f, we are led to consider the *pure
mapping class group* $\mathrm{PMod}(S^2 - P_f)$. It is the group of homeomor-
phisms of $S^2 - P_f$ fixing P_f pointwise, modulo isotopies fixing P_f
pointwise. The mapping class group acts on the set of isotopy classes
of simple closed curves.

We use the following fact:

Proposition 4.6. *The group $PMod(S^2 - P_f)$ is generated by a finite
number of explicit Dehn twists.*

The finiteness of the number of generating twists is a classical result
of Dehn; Lickorish [9] has made the construction explicit. See, for
example, [5] for an exposition.

5. Algorithm A_1: detecting an obstruction

Enumeration of the multicurves. We first prove the following proposition:

Proposition 5.1. *There exists an algorithm A which enumerates all non-peripheral multicurves on $S^2 - P_f$.*

For ease of reference let us state the following elementary fact:

Proposition 5.2. *Let S^* denote S^2 with a finite number of punctures. Consider two simple closed curves γ_1 and γ_2 in S^*. Assume that a component of $S^* - \gamma_1$ contains the same number of punctures as some component of $S^* - \gamma_2$. Then there exists a self-homeomorphism of S^* which sends γ_1 to γ_2.*

We fix a finite collection of simple closed curves $c_1, \ldots, c_M$ so that the Dehn twists $T_1, \ldots, T_M$ around those curves generate the mapping class group $\mathrm{PMod}(S^2 - P_f)$. This construction can be performed algorithmically by Proposition 4.6. We further refine our initial triangulation of the sphere so that these Dehn twists can be considered as piecewise linear maps relatively to the refined triangulation.

For every set of j punctures with $j \in \{2, \ldots, |P_f| - 2\}$ we choose one polygonal simple closed curve which separates them from the rest of P_f. Denote these curves γ_k (a simple count shows that $k = 1 \ldots 2^{|P_f| - 1} - |P_f| - 1$.

Proof of Proposition 5.1. To enumerate all multicurves, we proceed inductively as follows. At step 0, our collection of multicurves consists of all finite subsets of the set $\{\gamma_k\}$.

At step $N \in \mathbb{N}$, we generate all the possible images of the curves γ_k by reduced words in the T_i, T_i^{-1} of length less than N. Using Proposition 4.3, we remove all duplications from this finite collection. All inessential or peripheral curves are likewise removed.

We then consider all finite subsets of this collection. Using Proposition 4.3 again, we remove all subsets which have been generated previously, at steps $0, 1, \ldots, N - 1$.

By Proposition 5.2 and Proposition 4.6, every multicurve is thus generated. $\qquad\square$

Construction of the algorithm A_1. Denote $A(n)$ the algorithm of Proposition 5.1, which generates the exhaustive sequence of multicurves Γ_n. Set $n = 1$.

    (I) use Proposition 4.3 to check whether Γ_n is invariant. If not, proceed to step (V).

(II) Compute the transition matrix $M_{\gamma\delta}$ of the associated Thurston linear transformation f_{Γ_n}. Denote $P_n(\lambda)$ the characteristic polynomial of $M_{\gamma\delta}$.

(III) Is 1 a root of P_n? If yes, go to step (VI).

(IV) For $1 \leq j \leq n$ **do**
  − Use A_3 to query whether P_j has a root λ with $B(\lambda, 2^{-3i}) \subset [1 + 2^{-i}, \infty)$. If yes, go to step (VI).
  end do

(V) $n \mapsto n + 1$. Return to step (I).

(VI) Return **there exists a Thurston obstruction** and halt.

6. ALGORITHM A_2: FINDING AN EQUIVALENT RATIONAL MAP

6.1. **Moduli space of rational maps.** Let Rat_d denote the space of all holomorphic maps of degree $d \geq 2$ to itself. These maps can be written as fractions $\frac{p(z)}{q(z)}$, where the polynomials p, q are relatively prime and $d = \max(\deg p, \deg q)$. It can be shown that Rat_d is a connected complex-analytic manifold of dimension $2d+1$. Denoting $\mathrm{Res}(P, Q)$ the resultant of P and Q, one can represent Rat_d as the open set $\mathbb{P}^{2d+1}/V$, where $V = \{P, Q : \mathrm{Res}(P, Q) = 0\}$. (See for example [16], page 169).

Since we are interested in equivalence classes of rational maps under conjugation by Möbius maps, we are led to consider the moduli space $\mathcal{M}_d = \mathrm{Rat}_d/\mathrm{PSL}_2(\mathbb{C})$ (observe that it is the same as $\mathrm{Rat}_d/\mathrm{PGL}_2(\mathbb{C})$).

We note the following easy fact:

Proposition 6.1. *The moduli space $\mathcal{M}_d$ has the structure of a complex orbifold of dimension $2d - 2$.*

Proof. The stabilizer of a rational map $f \neq \mathrm{Id}$ under the action of $\mathrm{PSL}_2(\mathbb{C})$ is the subgroup $\mathcal{S}(f)$ consisting of Möbius maps which commute with f. There exists $n \in \mathbb{N}$ such that the set P_n consisting of periodic points of f with periods less than n has at least three points. Since every $M \in \mathcal{S}(f)$ must permute the points in P_n, the stabilizer $\mathcal{S}(f)$ is necessarily finite. □

As an example of a rational map with a non-trivial stabilizer, consider $f(z) = z^d$ for $d > 2$, in which case,

$$\mathcal{S}(f) = \{z \mapsto \lambda z | \text{ where } \lambda^{d-1} = 1\}.$$

We will now require a more computation-friendly description of $\mathcal{M}_d$. There are several similar approaches to this in the existing literature; we use the work [2]. As a first step we note the following standard fact:

Proposition 6.2. *Suppose R is not conjugate to a map of the form $z \mapsto z^{\pm d}$. Then the union of the critical and the postcritical sets $C_f \cup P_f$ contains at least three points.*

Mapping scheme. A *mapping scheme* of degree d is a triple (N, τ, ω), where $N \in \mathbb{N}$ and $N \geq 3$; τ is a *dynamics function*

$$\{1, \ldots, N\} \xrightarrow{\tau} \{1, \ldots, N\},$$

and ω is a *local degree function*

$$\{1, \ldots, N\} \xrightarrow{\omega} \mathbb{N}.$$

For a postcritically finite branched covering map $f : S^2 \to S^2$ denote Z_f is the union $C_f \cup P_f$ of the critical set and the postcritical set of f. We say that f *realizes* $X = (N, \tau, \omega)$ if we can choose a bijection $\psi : \{1, \ldots, N\} \to Z_f$ such that:

- $\psi(\tau(x)) = f(\psi(x))$;
- the local degree of f at $\psi(x)$ is equal to $\omega(x)$.

Following [2], a *normalization* of a mapping scheme $X = (N, \tau, \omega)$ is an injection

$$\alpha : \{0, 1, \infty\} \to \{1, \ldots, N\}.$$

We will denote a pair (X, α) by X_α, and refer to it as a *marked mapping scheme*.

We say that a rational map $f : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ realizes a marked mapping scheme X_α if the bijection ψ as defined above has the additional property:

- $\{0, 1, \infty\} \subset Z_f$ and $\psi^{-1}(i) = \alpha(i)$.

The set of all rational realizations of a specific marked mapping scheme $X^* = X_\alpha$ will be written as $\text{Rat}^\times(X^*)$.

Viewing the points in $\{1, \ldots, N\}$ as vertices of a weighted directed graph with arrows connecting x with $\tau(x)$ having weights $\omega(x)$, we define a *signature* $\mathcal{N}_X$ as follows. The signature is the set

$$\mathcal{N}_X = \{N(x), \ x \in \tau(\{1, \ldots, N\})\},$$

where $N(x)$ is defined as the least common multiple, over all directed paths of any length joining y to x, of the product of the weights of edges along this path. The significance of this definition lies in the following:

Proposition 6.3 ([3])**.** *If f is a rational mapping which realizes the mapping scheme X, then $\mathcal{N}_X$ is the signature of the orbifold of f.*

We can now state:

Theorem 6.4 ([2])**.** *Given a mapping scheme X whose signature is not $(2, 2, 2, 2)$ and a normalized marking scheme $X^* = (X, \tau, \omega)$, there is an injection $\iota : Rat^\times(X^*) \to \mathbb{C}^N$ such that the image is a zero-dimensional affine variety $V(\mathcal{I})$ determined by an ideal $\mathcal{I} = \mathcal{I}_{X^*}$, where $\mathcal{I}$ is defined over the rationals. In particular, $Rat^\times(X^*)$ is finite. Furthermore, a basis for $\mathcal{I}$ can be algorithmically computed.*

Let us give an indication of how the injection ι may be defined. Enumerate the elements of Z_f as $\{z_i\}$ so that $\psi(i) = z_i$. Isolate the elements sent respectively to zero and the infinity as follows:

$$FZ = \{n | z_n \in Z_f - \{\infty\} \text{ and } f(z_n) = 0\},$$

$$\text{and } FP = \{m | z_m \in Z_f - \{\infty\} \text{ and} f(z_m) = \infty\}.$$

Then any normalized rational realization f can be written *uniquely* as

$$(6.1) \qquad f(z) = \frac{a_0 + \ldots + a_r z^r}{b_0 + b_1 z + \ldots + 1 \cdot z^s} \cdot \frac{\prod_{n \in FZ}(z - z_n)^{d_n}}{\prod_{m \in FP}(z - z_m)^{d_m}}$$

Thus, the coordinates a_j, b_j, and z_j, specify the normalized rational map as a point in $\mathbb{C}^N$.

6.2. **The algorithm A_2.** We start with a piecewise-linear Thurston map $f : S^2 \to S^2$ with triangulation $\mathcal{T}$. Let X be a mapping scheme which is realized by f, and let α be a marking of X. Set $X^* = X_\alpha$. Let N and $\mathcal{I}_{X^*}$ be as in Theorem 6.4. Compute the finite set $g = (g_1, \ldots g_r)$ of polynomials with rational coefficients, which generates the ideal I_{X^*}. Denote A the algorithm of Proposition 5.1. Let $\{R_1, \ldots, R_m\}$ be the finite collection of rational maps which realize the marked mapping scheme X^*.

The algorithm A_2 works as follows:

  (I) use A to enumerate as D_i^k the representatives of the mapping class groups $\text{PMod}(S^2 - P_{R^k})$;
  (II) for every $1 \le k \le m$ and every pair $(i, j) \subset \mathbb{N} \times \mathbb{N}$ **do**
    – discretize the map $(D_j)^{-1} \circ R_k \circ D_j$ to a piecewise linear Thurston mapping M_j^k with triangulation $\mathcal{T}_1$. Note that at this stage the algorithm A_3 may need to be invoked to better estimate the coefficients of R_k;
    – refine the triangulations $\mathcal{T}$ and $\mathcal{T}_1$ to obtain a triangulation $\mathcal{T}_2$ on which both f and $M_{i,j}^k$ are defined;
    – identify all triangulated orientation preserving homeomorphisms h_j of $\mathcal{T}_2$; use the algorithm of Proposition 4.5 to list all $w_k = h_{j_k}$ which are isotopic to the identity;

– perform a finite check to determine whether there exists a pair w_i, w_j such that

$$w_i \circ M_{i,j}^k = w_j \circ f.$$

If yes, go to (III). **End do**

(III) output f **is Thurston equivalent to a rational map**, output the isolating neighborhood for R_k, and exit the algorithm.

7. CONCLUDING REMARKS

Let us note an easy corollary of our main result. Consider the following decidability problem:

Problem (a): *Given two piecewise linear Thurston mappings f and g with hyperbolic orbifolds and without Thurston obstructions, are f and g Thurston equivalent?*

Theorem 7.1. *Problem (a) is algorithmically decidable.*

Proof. Denote R_f and R_g the rational maps equivalent to f and g respectively. The existence of such maps is guaranteed by Thurston's theorem. They are defined up to a Möbius conjugacy. The algorithm works as follows:

(I) Check if f and g have identical mapping schemes $X = (N, \tau, \omega)$. If not, output **the maps are not Thurston equivalent** and exit.

(II) Run the algorithm A_2 to find isolating 2^{-n}-neighborhoods U_f and U_g of R_f and R_g respectively in the parameter space of normal forms (6.1).

(III) For each ordered triple of distinct natural numbers (a, b, c) between 1 and N **do**
    – Normalize the mapping scheme X by $\alpha : (0, 1, \infty) \mapsto (a, b, c)$. Calculate a point of the parameter space $\bar{w}_{(a,b,c)}$ representing the normal form (6.1) of R_f corresponding to (X, α) with precision $2^{-(n+2)}$.
    – If $\bar{w}_{(a,b,c)} \in U_g$ then output **the maps are Thurston equivalent** and exit.

(IV) Output **the maps are not Thurston equivalent** and exit.

$\square$

Consider the following natural generalization of Problem (a), suggested to us by M. Lyubich:

Problem (b): *Given two piecewise linear Thurston mappings f and g with hyperbolic orbifolds, are f and g Thurston equivalent?*

Algorithmic decidability of Problem (b) presents an interesting direction of further study.

We find that recognizability problems of combinatorial equivalence of piecewise-linear maps are analogous to recognizability problems of piecewise-linear manifolds, which are solvable in dimensions 3 or less, and become algorithmically intractable in dimensions greater than 4 (and possibly 4 as well) – see, for example, the book of Weinberger [18]. In conclusion, we speculate that natural notions of equivalence of maps in higher dimensions will lead to algorithmically unsolvable problems – and a new interplay between Dynamics and Computability.

Appendix A. A root-finding algorithm.

The existence of a root-finding algorithm A_3 is a classical result of H. Weyl [19]. Consider a system of analytic functions $g = \{g_i(z), \ i = 1, \ldots, n\}$ defined in a box $D \subset \mathbb{C}^n$. Denote $N(d, D)$ the number of common zeroes of g in D, and assume that there are no common zeroes on the boundary. Then $N(d, D)$ can be computed using the following multidimensional residue formula (see [15] page 324):

$$N(g, D) = \frac{(n-1)!}{(2\pi i)^n} \int_{\partial D} \frac{1}{|g|^{2n}} \cdot \sum_{j=1}^{n} \overline{g_j} dg_j \wedge d\overline{g_j} \wedge dg_1 \wedge d\overline{g_1} \ldots \widehat{[j]} \ldots d\overline{g_r} \wedge dg_r,$$

where $\widehat{[j]}$ means that we omit the term $dg_j \wedge d\overline{g_j}$.

Observe that our space $Rat^{\times}(X^{\star})$ is indeed given as the zeroes of N polynomials in $\mathbb{C}^N$.

Weyl's algorithm to locate all roots of $g = 0$ in an isolating neighborhood G works as follows. Begin by setting $j = 0$ covering G by a cubic grid of size $2^{-j} = 1$. In each cube C_k of the grid, use the residue formula to check whether there are any zeros in it. Since we cannot catch zeros on the boundary of a cube, perform the check for a cube of twice the size – it is guaranteed to catch any zeros in the closure of C_k.

Throw away all cubes without any zeros. Increment $j \mapsto j + 1$ and divide the remaining cubes C_k into cubes with side 2^{-j}. Repeat the process, until all zeros are identified with the desired precision.

References

[1] R. H. Bing, *An alternative proof that 3-manifolds can be triangulated*, Annals of Math. (2) **69** (1959), 37–65.

[2] E. Brezin, R. Byrne, J. Levy, K. Pilgrim, and K. Plummer, *A census of rational maps*, Conformal Geometry and Dynamics **4** (2000), 35–74.

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

[4] R. Douady and A. Douady, *Algèbre et théories galoisiennes*, Cassini, 2005.

[5] B. Farb and D. Margalit, *A Primer on mapping class groups*.

[6] J.H. Hubbard and D. Schleicher, *The Spider Algorithm*.

[7] A. Kameyama, *The Thurston equivalence for postcritically finite branched coverings*, Osaka J. Math. **38** (2001), 565–610.

[8] Y. Ladegaillerie, *Classes d'isotopie de plongements de 1-complexes dans les surfaces*, Topology **23** (1984), 303–311.

[9] W. B. R. Lickorish, *A finite set of generators for the homeotopy group of a 2-manifold*, Proc. Cambridge Philos. Soc. **60** (1964), 769–778.

[10] J. Milnor and W. Thurston, *On iterated maps of the interval*.

[11] V. Nekrashevych, *Self-similar groups*, Vol. 117, American Mathematical Society, 2005.

[12] K. Pilgrim, *Combinations of complex dynamical systems*, Lecture Notes in Mathematics, vol. 1827, Springer, 2003.

[13] ———, *An algebraic formulation of Thurston's combinatorial equivalence*, Proc. Amer. Math. Soc. **131** (2003), 3527–3534.

[14] T. Radó, *Uber den Begriff der Riemannschen Flächen.*, Acta Litt. Sci. Szeged (1925), 101–121.

[15] B. Shabat, *Introduction à l'analyse complexe*, Mir, 1990.

[16] J. H. Silverman, *The arithmetic of dynamical systems*, Vol. 241, Springer, 2007.

[17] W. P. Thurston, *Three-dimensional Geometry and Topology*, Princeton University Press, 1997.

[18] S. Weinberger, *Computers, rigidity, and moduli. The large-scale fractal geometry of Riemannian moduli space. M. B. Porter Lectures*, Princeton University Press, Princeton, NJ, 2005.

[19] H. Weyl, *Randbemerkungen zu Hauptproblemen der Mathematik, II, Fundamentalsatz der Algebra and Grundlagen der Mathematik*, Math. Z. **20** (1924), 131–151.