

ALGORITHMIC CONSTRUCTION OF HURWITZ MAPS

LAURENT BARTHOLDI, XAVIER BUFF, HANS-CHRISTIAN GRAF VON BOTHMER,
AND JAKOB KRÖKER

ABSTRACT. We describe an algorithm that, given a k -tuple of permutations representing the monodromy of a rational map, constructs an arbitrarily precise floating-point complex approximation of that map.

We then explain how it has been used to study a problem in dynamical systems raised by Cui.

1. INTRODUCTION

Let \mathbb{S} be a topological oriented 2-sphere. The branched coverings $\mathbb{S} \rightarrow \mathbb{S}$ considered in this article are all orientation preserving. Let $Q := \{Q_i\}_{i \in I}$ be a finite subset of \mathbb{S} with $I = \mathbb{Z}/k\mathbb{Z}$. In [9], Hurwitz describes an elegant classification of branched coverings $\mathbb{S} \rightarrow \mathbb{S}$ with branch value set contained in Q in terms of *admissible* k -tuples of permutations $(\sigma_i \in \mathfrak{S}_d)_{i \in I}$. A k -tuple is admissible if:

- the permutations $(\sigma_i)_{i \in I}$ generate a transitive subgroup of \mathfrak{S}_d ,
- $\sigma_1 \cdot \sigma_2 \cdots \sigma_k = \text{id}$ and
- the cycle lengths satisfy the condition

$$(1) \quad \sum_{i \in I} \sum_{\substack{c \text{ cycle} \\ \text{of } \sigma_i}} (\text{length}(c) - 1) = 2d - 2.$$

See §1.1 for more details regarding the classification.

It is easy, using a computer algebra system such as GAP [7], to enumerate all admissible k -tuples of permutations; it is an altogether different problem to *construct* an analytic model of a covering associated to a given admissible k -tuple of permutations. The purpose of this note is to describe such an algorithm, its implementation, and an application to complex dynamics.

1.1. Hurwitz's classification. Two branched coverings $f : \mathbb{S} \rightarrow \mathbb{S}$ and $g : \mathbb{S} \rightarrow \mathbb{S}$ are equivalent if there is an orientation preserving homeomorphism $h : \mathbb{S} \rightarrow \mathbb{S}$ such that $g = f \circ h$. Hurwitz's result is a classification of equivalence classes of coverings in this sense.

Choose a basepoint $* \in \mathbb{S} \setminus Q$. For each $i \in I$, choose a path γ_i joining $*$ to Q_i in $\mathbb{S} \setminus Q$, in such a way that

- the paths γ_i intersect only at $*$,
- the paths $(\gamma_1, \dots, \gamma_k)$ are ordered cyclically counterclockwise around $*$.

The fundamental group $G = \pi_1(\mathbb{S} \setminus Q, *)$ is generated by paths $\hat{\gamma}_i$ that follow γ_i , wind once counterclockwise around Q_i , and return to $*$ along γ_i . It has the presentation

$$G = \langle \hat{\gamma}_i, i \in I \mid \hat{\gamma}_1 \cdot \hat{\gamma}_2 \cdots \hat{\gamma}_k = \text{id} \rangle.$$

Let $f : (\mathbb{S}, C) \rightarrow (\mathbb{S}, Q)$ be a covering branched over Q . Number $\{*_1, \dots, *_d\}$ the f -preimages of $*$. Then, for each $i \in I$ and each $m \in \{1, \dots, d\}$, the path $\hat{\gamma}_i$ lifts to a path starting at $*_m$ and ending at $*_n$ for some $n =: \sigma_i(m)$. This defines a permutation σ_i for each $i \in I$. Note that the k -tuple $(\sigma_i)_{i \in I}$ is admissible:

- since $\mathbb{S} \setminus C$ is connected, the group $\langle \sigma_i \rangle$ is transitive on $\{1, \dots, d\}$;
- since $\hat{\gamma}_1 \cdot \hat{\gamma}_2 \cdots \hat{\gamma}_k = \text{id}$, we have that $\sigma_1 \cdot \sigma_2 \cdots \sigma_k = \text{id}$;
- computing the Euler characteristic of $\mathbb{S} \setminus C$ via the Riemann-Hurwitz formula yields (1).

Conversely, let $(\sigma_i)_{i \in I}$ be an admissible k -tuple of permutations. Define a branched covering as follows: start with d disjoint copies of \mathbb{S} , cut open along the paths γ_i . If $\sigma_i(m) = n$, glue the right boundary of γ_i on m -th sphere to the left boundary of γ_i on the n -th sphere. This defines a covering with branch value set contained in Q . It is connected because $\langle \sigma_i \rangle$ is transitive on $\{1, \dots, d\}$. The Euler characteristic of the cover is 2, because of (1) and the Riemann-Hurwitz formula; so it is a sphere.

The k -tuple $(\sigma_i)_{i \in I}$ must be considered up to diagonal conjugation by \mathfrak{S}_d , which amounts to numbering the spheres differently. The constructions above then define a bijection between equivalence classes of branched coverings and equivalence classes of appropriate k -tuples of permutations.

A coarser equivalence relation on coverings has also been considered, but is not the main focus of this article: two coverings $f, g : \mathbb{S} \rightarrow \mathbb{S}$ are *Hurwitz equivalent* if there exist homeomorphisms $h_0, h_1 : \mathbb{S} \rightarrow \mathbb{S}$ with $f \circ h_1 = h_0 \circ g$. Hurwitz classes of coverings may also be classified by k -tuples of permutations; namely, by the orbits on appropriate k -tuples of the symmetric group \mathfrak{S}_d (acting as above) and the pure braid group on k strings. The latter group's generators act by conjugating, for any two consecutive points Q_i, Q_{i+1} in Q , the permutations σ_i and σ_{i+1} by $\sigma_i \sigma_{i+1}$. This amounts to changing the “spider” $\bigcup_{i \in I} \gamma_i$ by twisting the legs γ_i and γ_{i+1} around each other.

1.2. Analytic models. Assume now $Q \subset \mathbb{P}^1(\mathbb{C})$ and that $f : \mathbb{S} \setminus C \rightarrow \mathbb{P}^1(\mathbb{C}) \setminus Q$ is a covering map. Then, f defines holomorphic charts on $\mathbb{S} \setminus C$ and it is not difficult to see that the points in C are removable singularities: we denote by \mathbb{S}_f the corresponding Riemann surface. By the Uniformization Theorem, there is a conformal homeomorphism $\phi_f : \mathbb{S}_f \rightarrow \mathbb{P}^1(\mathbb{C})$. The map $F := f \circ \phi_f^{-1} : \mathbb{P}^1(\mathbb{C}) \rightarrow \mathbb{P}^1(\mathbb{C})$ is a holomorphic branched covering, i.e., a rational map. Assume $g = f \circ h : \mathbb{S} \rightarrow \mathbb{P}^1(\mathbb{C})$ for some homeomorphism $h : \mathbb{S} \rightarrow \mathbb{S}$. Let $\phi_g : \mathbb{S}_g \rightarrow \mathbb{P}^1(\mathbb{C})$ be a conformal homeomorphism and set $G := g \circ \phi_g^{-1} : \mathbb{P}^1(\mathbb{C}) \rightarrow \mathbb{P}^1(\mathbb{C})$ be the corresponding rational map. Then, $H = \phi_f \circ \phi_g^{-1} : \mathbb{P}^1(\mathbb{C}) \rightarrow \mathbb{P}^1(\mathbb{C})$ is a Möbius transformation and $F = G \circ H$.

Therefore, up to precomposition by a Möbius transformation, the rational map F only depends on the equivalence class of covering $f : (\mathbb{S}, C) \rightarrow (\mathbb{P}^1(\mathbb{C}), Q)$. We say that f is an *analytic model*.

1.3. Dynamics. If $\#Q = 3$, then we may assume $Q_1 = \infty$, $Q_2 = 0$ and $Q_3 = 1$ within $\mathbb{P}^1(\mathbb{C})$. Furthermore, precomposing f by an appropriate Möbius transformation, we may also assume that $Q \subset f^{-1}(Q)$. Then $f : \mathbb{P}^1(\mathbb{C}) \rightarrow \mathbb{P}^1(\mathbb{C})$ becomes a *post-critically finite map*: a self-map of the sphere all of whose critical points have finite orbits. Pilgrim linked in [13] the “dessin d’enfant” (the full preimage of the segment $[0, 1]$) of f with a dynamical invariant, its “Hubbard tree”.

Post-critically finite maps may be considered as purely topological objects, namely branched self-coverings of the topological sphere \mathbb{S} ; a fundamental theorem of Thurston (see [6]) proves (except in few well-understood, low-complexity cases) that such a branch covering has at most one holomorphic realization, up to conjugation by a Möbius transformation; furthermore, if $\#Q = 3$ as above, then it has precisely one holomorphic realization.

We describe in Section 8 a question by Cui in the theory of holomorphic dynamical systems, and give an explicit holomorphic realization of a topological map he constructed.

This will also be our running example in the text. With $Q_1 = \infty$, $Q_2 = 0$ and $Q_3 = 1$, the permutations representing the map are

$$\begin{aligned} \sigma_1 &= (1, 7, 11, 2)(3, 8)(4, 5)(6, 10)(9, 12, 13), \\ \sigma_2 &= (1, 3, 12, 4)(5, 9)(6, 7)(10, 13, 11)(2, 8), \\ \sigma_3 &= (1, 5, 13, 6)(7, 10)(2, 3)(8, 11, 12)(4, 9). \end{aligned} \tag{2}$$

We seek a degree-13 rational map f .

In this specific example, the search can be made more feasible as follows. Setting all critical points as unknowns and eliminating is out of the question. With a little faith that the symmetry between $\infty, 0, 1$ translates to f , let $\rho(z) = 1/(1-z)$ be the rotation permuting $\infty, 0, 1$, and note that $\mathbb{P}^1(\mathbb{C})/\rho$ is a sphere, branched at the two fixed points of ρ . If f descends to a map g on $\mathbb{P}^1(\mathbb{C})/\rho$, then (after change of variables) it has the form $g(z) = z(p(z)/q(z))^3$ for degree-4 polynomials, such that $g(z) = 1 + \mathcal{O}((z-1)^4)$ at $z = 1$, and such that 1 is the image of four other points with local degrees 3, 2, 2, 2 respectively. We are grateful to Noam Elkies and Curt McMullen for having pointed out to us the feasibility of this approach.

Nevertheless, we will show that our algorithm is strong enough to produce a solution even without exploiting the symmetry of the Hurwitz data.

1.4. Simple cases. If $\#Q = 2$, then there is a unique solution represented, up to diagonal conjugation, by the pair of permutations

$$\sigma_1 = (1, 2, \dots, d) \quad \text{and} \quad \sigma_2 = (d, \dots, 2, 1).$$

If $Q_1 = \infty$ and $Q_2 = 0$, an analytic model is $f(z) = z^d$.

However, the case $\#Q = 3$ seems already as complicated as the general case, and has only been addressed in the literature for small d . Such maps are often called “dessins d’enfant”, see [8]; the corresponding combinatorial objects for the modular surface $\mathfrak{h}/\mathrm{PSL}_2(\mathbb{Z})$ are called “Conway diagrams”, see [1, §3.4]. Methods of constructing them are addressed, *inter alia*, in [3–5].

In this section, we consider the case $d \leq 3$ which can completely be solved. If $\#Q = 2$, then as we said above we may choose $Q = \{\infty, 0\}$ and $f(z) = z^d$. If $d = \#Q = 3$ then we may choose $Q = \{\infty, 0, 1\}$. Without loss of generality, we may assume that all points of Q are branched values, since otherwise we are reduced to the case $\#Q = 2$. Up to permutation of the points in Q and the indices, the only possible triple of permutations is

$$\sigma_1 = (1, 2, 3), \quad \sigma_2 = (1, 2) \quad \text{and} \quad \sigma_3 = (2, 3).$$

To find an analytic model, we seek a rational map f of degree 3 such that

$$\infty \xrightarrow{3:1} \infty, \quad 0 \xrightarrow{2:1} 0 \quad \text{and} \quad 1 \xrightarrow{2:1} 1.$$

This implies $f(z) = 3z^2 - 2z^3$ as the only realization.

The next case we consider is $d = 3$ and $\#Q = 4$. Using Möbius transformations, we may normalise Q to be $\{\infty, 0, 1, w\}$. Up to conjugation in \mathfrak{S}_3 we may take the first permutations to be $\sigma_1 = \dots = \sigma_i = (1, 2)$. The condition that the permutations generate a transitive group imply that one of them is not $(1, 2)$. Up to conjugation, we may assume that the first permutation which is not $(1, 2)$ is $\sigma_{i+1} = (2, 3)$. Since $\sigma_1 \cdot \sigma_2 \cdot \sigma_3 \cdot \sigma_4 = \text{id}$, this gives four possibilities, namely, writing $\sigma = (\sigma_1, \sigma_2, \sigma_3, \sigma_4)$,

$$\begin{aligned} \sigma &= ((1, 2), (1, 2), (2, 3), (2, 3)), & \sigma &= ((1, 2), (2, 3), (1, 2), (1, 3)), \\ \sigma &= ((1, 2), (2, 3), (1, 3), (2, 3)), & \sigma &= ((1, 2), (2, 3), (2, 3), (1, 2)). \end{aligned}$$

To find the corresponding f , assume without loss of generality that f maps $\infty \mapsto \infty$, $0 \mapsto 0$, $1 \mapsto 1$ and $v \mapsto w$. This forces the map f to have the form

$$f_a(z) = z^2 \frac{a(z-1) + 1}{(a+2)(z-1) + 1},$$

for some parameter a subject to $(a+1)(a-1)^3 = wa(a+2)^3$; then $(a+1)(a-1) = va(a+2)$. Since $w \neq 0, 1$, the equation defining a in terms of w has four distinct roots, leading to four candidate maps f_a . There is a bijection between the maps f and the triples of permutations above, but no canonical one — it will depend on the specific choice of $\#Q$ generators $\hat{\gamma}_i$ of $\pi_1(\mathbb{P}^1(\mathbb{C}) \setminus Q, *)$. Note also that these four solutions are part of a single Hurwitz class.

2. OVERVIEW OF THE ALGORITHM

We are given a list $\sigma_1, \dots, \sigma_k$ of permutations in \mathfrak{S}_d with product $\sigma_1 \cdots \sigma_k = 1$, and points $Q_1, \dots, Q_k \in \mathbb{P}^1(\mathbb{C})$.

Let $\alpha_i = (\alpha_{i,1}, \dots, \alpha_{i,\ell_i})$ be the cycle lengths of σ_i ; we have $\sum_j \alpha_{i,j} = d$ for all i , and $\sum_{i,j} (\alpha_{i,j} - 1) = 2d - 2$. In the first part of the algorithm, we enumerate all rational maps with critical values Q_1, \dots, Q_k such that the multiplicities of the preimages of Q_i are $\alpha_{i,1}, \dots, \alpha_{i,\ell_i}$. In the second part, we select the appropriate rational map among these candidates.

The approach in the first part of the algorithm seems to originate in Malle [11]; see also [12].

For the sake of describing its workflow more clearly, the actual algorithm (described in the remainder of the text) has been slightly simplified.

Normalization: Without loss of generality, we assume $Q_1 = \infty$, $Q_2 = 0$ and $Q_3 = 1$. We approximate the other Q_i by $\tilde{Q}_i \in \mathbb{P}^1(\overline{\mathbb{Q}})$. The rational map we seek will leave Q_1, Q_2, Q_3 fixed.

Finite field solution: We pick a prime p , such that the points \tilde{Q}_i have realizations $\overline{Q}_i \in \mathbb{P}^1(\mathbb{F}_p)$. We then list all degree- d rational maps \overline{F} over \mathbb{F}_p with poles and zeroes of multiplicities α_1 and α_2 respectively, and by brute force check for each \overline{F} whether $\overline{F} - \overline{Q}_i$ has zeroes of multiplicities α_i for all $i \geq 3$. Note that the rational map \overline{F} is a solution to our original problem over \mathbb{F}_p . (If there are no solutions, we restart with a different prime p).

p -adic solution: Write $\overline{F} = \overline{W}_2/\overline{W}_1$ with \overline{W}_2 monic of degree d , and \overline{W}_1 of degree less than d . (In fact, we later write the denominator as λW_1 with W_1 monic. The present discussion uses a simplified notation.) For $i \geq 3$, let $\overline{W}_i = \overline{W}_2 - \overline{Q}_i \overline{W}_1$ be the numerator of $\overline{F} - \overline{Q}_i$. We compute high-precision p -adic approximations \hat{Q}_i of the \tilde{Q}_i , and lift each \overline{W}_i from

\mathbb{F}_p to a high-precision polynomial \hat{W}_i over \mathbb{Z}_p , in such a manner that we have $\hat{W}_i = \hat{W}_2 - \hat{Q}_i \hat{W}_1 + O(p^N)$ for large N . This lifting can be done by Hensel's lemma, because by Corollary 3.3, the Jacobian of the system $\{W_i = W_2 - Q_i W_1\}$ is invertible at a solution for almost every prime p . (If $D\bar{F}$ happens not to be invertible, we restart with a different prime).

Algebraic solution: Using the lattice-reduction algorithm LLL [10], we find polynomials W_i over $\bar{\mathbb{Q}}$, with coefficients of small height (small degree and coefficients of minimal polynomial) that are close to \hat{W}_i obtained at the previous step. Using exact arithmetic over $\bar{\mathbb{Q}}$, we check that the solution W_2/W_1 is correct. (If not, we either compute a finer p -adic approximation, or higher-degree algebraic number approximations, or we restart altogether with a larger prime).

Complex solution: For each root $c_{i,j} \in \bar{\mathbb{Q}}$ of W_i , given by its minimal polynomial over \mathbb{Q} , we compute (to user-specified precision) all the complex roots $\tilde{c}_{i,j,s}$ of its minimal polynomial. We still need to find out which of the Galois conjugates $\tilde{c}_{i,j,s}$ belong to the same solution. For this, we also choose small, random integers m, n , compute the minimal polynomial of $mc_{i,j} + nc_{i',j'}$, and find complex approximations $\tilde{d}_{i,j,i',j',t}$ of its roots. We then pair together those roots $(\tilde{c}_{i,j,s}, \tilde{c}_{i',j',s'})$ for which $\tilde{d}_{i,j,i',j',t} \approx m\tilde{c}_{i,j,s} + n\tilde{c}_{i',j',s'}$ for some t . By considering enough of these pairs we can stitch together a collection of compatible coordinate approximations $\tilde{c}_{i,j,s}$ embracing all i, j .

We call \tilde{C} the collection of all roots $\tilde{c}_{i,j,s}$, and note that the rational map is determined by its zeroes and poles, and the normalization that $Q_1 = 1$ is fixed. Since $\{0, \infty\} \subset Q$, these zeroes and poles are in \tilde{C} .

The second step of the algorithm checks, by path lifting, that the monodromy around Q_i is correct. For each of the Galois conjugate solutions (\tilde{C}, \tilde{f}) obtained in the first step, we do the following:

Triangulate: We are given a floating-point approximation \tilde{Q} of Q . We compute a triangulation \mathcal{Q} of $\mathbb{P}^1(\mathbb{C})$ whose vertex set contains \tilde{Q} , and a triangulation \mathcal{C} of $\mathbb{P}^1(\mathbb{C})$ whose vertex set contains \tilde{C} . For efficiency reasons, we use *Delaunay triangulations*, see §7. We compute the dual triangulation \mathcal{Q}^\perp ; it has one vertex per face of \mathcal{Q} , and edges transverse to those of \mathcal{Q} . We fix a vertex $* \in \mathcal{Q}^\perp$ as our basepoint.

Lift the triangulation: Let W denote the vertices of \mathcal{Q}^\perp . For each $w \in W$, we number arbitrarily w_1, \dots, w_d the \tilde{f} -preimages of w .

For each edge $\varepsilon \in \mathcal{Q}^\perp$ going from w' to w'' , we compute a permutation $\varsigma_\varepsilon \in \mathfrak{S}_d$ such that the \tilde{f} -lift of ε starting at w'_i ends at $w''_{\varsigma_\varepsilon(i)}$. There are two strategies for this, one is by subdividing appropriately the path ε and playing “connect-the-dots”, the other uses more efficiently the triangulation \mathcal{C} .

Read permutations: For each critical value $Q_i \in Q$, let $\varepsilon(1), \dots, \varepsilon(n)$ be the sequences of edges traversed by a path in \mathcal{Q}^\perp that starts and ends in the basepoint $*$, and surrounds once counterclockwise the point Q_i and no other vertex of Q . Compute the permutation $\sigma'_i = \varsigma_{\varepsilon(1)} \cdots \varsigma_{\varepsilon(n)}$.

Check: The data (C, f) are a valid solution to the Hurwitz problem if and only if there exists a permutation $\tau \in \mathfrak{S}_d$ such that $\sigma'_i = (\sigma_i)^\tau$.

2.1. Implementation. The fourth-named author has implemented the first part of the algorithm, mainly in C, and the second-named author has implemented the second part of the algorithm, mainly in GAP [7]. By far the most time-consuming part of the procedure is the search for a solution over a finite field. Example 4.9 required approximately 15 Minutes on a desktop, 30-SPECint2006 computer. The code is maintained by the fourth-named author, and is available at

<https://github.com/jakobkroeker/HMAC>

3. THE SPACE OF RATIONAL MAPS

We show, in this section, that (as soon as the prime p is sufficiently large) we may lift every \mathbb{F}_p -solution to \mathbb{Z}_p . This follows from the well known fact that the Hurwitz spaces are smooth. We could not find the precise statement we need in the literature, so we give a complete proof.

Let $d \geq 2$ be an integer and denote by Rat_d the space of rational maps of degree d , which may be identified with a Zariski open subset of $\mathbb{P}^{2d+1}(\mathbb{C})$.

Let $k \geq 3$ be an integer and let $F : \mathbb{S} \rightarrow \mathbb{S}$ be a ramified covering branched over $Q = \{Q_1, \dots, Q_k\}$. Note that, according to the Riemann-Hurwitz Formula, $C := F^{-1}(Q)$ contains exactly $(k-2)d+2$ points. We write $C = \bigcup_i C_i$ with $C_i := F^{-1}\{Q_i\} = \{C_{i,1}, \dots, C_{i,\ell_i}\}$, and for each $j \in \{1, \dots, \ell_i\}$ we let $\alpha_{i,j}$ be the local degree of F at $C_{i,j}$.

Let \mathfrak{Q} be the smooth quasiprojective variety of injective maps $\mathfrak{q} : Q \rightarrow \mathbb{P}^1(\mathbb{C})$. For $\mathfrak{q} \in \mathfrak{Q}$, we use the notation $q_i := \mathfrak{q}(Q_i)$. Similarly, let \mathfrak{C} be the smooth quasiprojective variety of injective maps $\mathfrak{c} : C \rightarrow \mathbb{P}^1(\mathbb{C})$. For $\mathfrak{c} \in \mathfrak{C}$, we use the notation $c_{i,j} := \mathfrak{c}(C_{i,j})$. The quasiprojective variety $\mathfrak{Y} := \mathfrak{C} \times \mathfrak{Q}$ is smooth. We shall prove that the subvariety

$$\mathfrak{X} := \{(\mathfrak{c}, \mathfrak{q}) \in \mathfrak{Y} \mid (\exists f \in \text{Rat}_d) (\forall i, j) f(c_{i,j}) = q_i \text{ and } \deg_{c_{i,j}} f = \alpha_{i,j}\}$$

is also smooth, and regularly parametrised:

Proposition 3.1. *The variety \mathfrak{X} is smooth of dimension $k+3$, locally regularly parametrised by $(q_1, \dots, q_k, c_{1,1}, c_{2,1}, c_{3,1})$.*

Observe that, for $(\mathfrak{c}, \mathfrak{q}) \in \mathfrak{X}$, there is a unique rational map $f \in \text{Rat}_d$ such that $f(c_{i,j}) = q_i$ and $\deg_{c_{i,j}} f = \alpha_{i,j}$ for all (i, j) . Indeed, knowing a rational map above three points completely determines the rational map (it is even enough to know the full preimage of two points plus one preimage of a third point).

Note that the group of Möbius transformations acts on \mathfrak{C} and \mathfrak{Q} by postcomposition:

$$(M, N) \cdot (\mathfrak{c}, \mathfrak{q}) := (M \circ \mathfrak{c}, N \circ \mathfrak{q}).$$

The quotient space may be identified with $\mathfrak{Y}_0 := \mathfrak{C}_0 \times \mathfrak{Q}_0$ with

$$\mathfrak{C}_0 := \{\mathfrak{c} \in \mathfrak{C} \mid c_{1,1} = \infty, c_{2,1} = 0 \text{ and } c_{3,1} = 1\}$$

and

$$\mathfrak{Q}_0 := \{\mathfrak{q} \in \mathfrak{Q} \mid q_1 = \infty, q_2 = 0 \text{ and } q_3 = 1\}.$$

The projection $\mathfrak{Y} \rightarrow \mathfrak{Y}/(\text{PSL}_2(\mathbb{C}) \times \text{PSL}_2(\mathbb{C})) \cong \mathfrak{Y}_0$ is a submersion.

The action preserves \mathfrak{X} as indicated on the following commutative diagram:

$$\begin{array}{ccccc} \mathbb{C} & \xrightarrow{c} & \mathbb{P}^1(\mathbb{C}) & \xrightarrow{M} & \mathbb{P}^1(\mathbb{C}) \\ F \downarrow & & \downarrow f & & \downarrow N \circ f \circ M^{-1} \\ \mathbb{Q} & \xrightarrow{q} & \mathbb{P}^1(\mathbb{C}) & \xrightarrow{N} & \mathbb{P}^1(\mathbb{C}). \end{array}$$

It is therefore enough to show that $\mathfrak{X}_0 := \mathfrak{X} \cap \mathfrak{Y}_0$ is a smooth subvariety of \mathfrak{Y}_0 locally regularly parametrised by (q_4, \dots, q_k) .

We first write equations for \mathfrak{X}_0 . To each $(\mathbf{c}, \mathbf{q}, \lambda) \in \mathfrak{C}_0 \times \mathfrak{Q}_0 \times \mathbb{C}^*$, we associate a collection of monic polynomials $(W_i)_{i \in \{1, \dots, k\}}$ defined by

$$W_1(z) := \prod_{j=2}^{\ell_1} (z - c_{1,j})^{\alpha_{1,j}} \quad \text{and for } i \geq 2 \quad W_i(z) := \prod_{j=1}^{\ell_i} (z - c_{i,j})^{\alpha_{i,j}}$$

and a collection of rational maps $(f_i)_{i \in \{2, \dots, k\}}$ defined by

$$f_i := \frac{W_i}{\lambda W_1} + q_i.$$

Note that these are degree- d rational maps with poles of order $\alpha_{1,j}$ at $c_{1,j}$. In addition, f_i maps $c_{i,j}$ to q_i with local degree $\alpha_{i,j}$. It follows that $(\mathbf{c}, \mathbf{q}) \in \mathfrak{X}_0$ if and only if there is a $\lambda \in \mathbb{C}^*$ such that $f_i = f_2$ for all $i \in \{3, \dots, k\}$, that is, $F_i = 0$ with

$$(3) \quad F_i := W_i + \lambda q_i W_1 - W_2.$$

In that case, we use the notation

$$f_{(\mathbf{c}, \mathbf{q}, \lambda)} := f_2 = f_3 = \dots = f_k.$$

In other words, consider the map

$$\mathcal{F} := (F_3, \dots, F_k) : \mathfrak{Y}_0 \times \mathbb{C} \rightarrow (\mathbb{C}[z]_{\deg < d})^{k-2}.$$

Then, $(\mathbf{c}, \mathbf{q}) \in \mathfrak{X}_0$ if and only if there is a $\lambda \in \mathbb{C}^*$ such that $\mathcal{F}(\mathbf{c}, \mathbf{q}, \lambda) = 0$.

According to the following Lemma and the Implicit Function Theorem, the subvariety of $\mathfrak{Y}_0 \times \mathbb{C}$ defined by the equation $\mathcal{F} = 0$ is smooth of dimension $k - 3$, locally regularly parametrised by (q_4, \dots, q_k) . It follows that its projection to the \mathfrak{Y}_0 component, namely \mathfrak{X}_0 , is also smooth of dimension $k - 3$, locally regularly parametrised by (q_4, \dots, q_k) .

Lemma 3.2. *If $\mathcal{F}(\mathbf{c}, \mathbf{q}, \lambda) = 0$, then the derivative $D_{(\mathbf{c}, \mathbf{q}, \lambda)} \mathcal{F}$ restricts to an isomorphism $T_{\mathbf{c}} \mathfrak{C}_0 \times \{0\} \times T_{\lambda} \mathbb{C} \rightarrow T_0(\mathbb{C}[z]_{\deg < d})^{k-2}$.*

We postpone the proof of the lemma to Section 3.2 and mention immediately a corollary that we shall use later. For $\mathbf{q} \in \mathfrak{Q}$, let $\mathcal{F}_{\mathbf{q}} : \mathfrak{C}_0 \times \mathbb{C}^* \rightarrow (\mathbb{C}[z]_{\deg < d})^{k-2}$ be defined by

$$\mathcal{F}_{\mathbf{q}}(\mathbf{c}, \lambda) := \mathcal{F}(\mathbf{c}, \mathbf{q}, \lambda).$$

Corollary 3.3. *Assume that $(\mathbf{c}, \mathbf{q}, \lambda)$ is defined over $\overline{\mathbb{Q}}$ with $\mathcal{F}(\mathbf{c}, \mathbf{q}, \lambda) = 0$. Then, for almost every prime p , the derivative $D\mathcal{F}_{\mathbf{q}}$ at (\mathbf{c}, λ) is invertible mod p .*

Proof. Since the point $(\mathbf{c}, \mathbf{q}, \lambda)$ is defined over $\overline{\mathbb{Q}}$, its coördinates may be written using integers, and reduced mod p . For all except finitely many values of p , the resulting reduction gives a genuine point, namely where the reductions of (\mathbf{c}, \mathbf{q}) are injective and the reduction of $f_{(\mathbf{c}, \mathbf{q}, \lambda)}$ has degree d . Since $D\mathcal{F}$ is invertible over $\overline{\mathbb{Q}}$,

it may be written as a/N for an algebraic integer a and $N \in \mathbb{N}$; then the reduction modulo p of $D\mathcal{F}$ is invertible for all primes not dividing N . \square

Before embarking in the proof of the Lemma, we first build up a description of the tangent space of Rat_d .

3.1. The tangent space to rational maps. Consider a rational map $f \in \text{Rat}_d$. A tangent vector to f is

$$\dot{f} := \left. \frac{df_t}{dt} \right|_{t=0}$$

for a holomorphic family of rational maps $(f_t)_{|t|<\epsilon}$ with $f_0 = f$.

For every $z \in \mathbb{P}^1(\mathbb{C})$, the vector $\dot{f}(z)$ is a tangent vector in $\mathbb{P}^1(\mathbb{C})$ at $f(z)$; in other words, \dot{f} is a section of the pullback bundle $f^*(T\mathbb{P}^1(\mathbb{C}))$. It can be pulled back to a vector field on $\mathbb{P}^1(\mathbb{C})$, as

$$\eta(z) = -(D_z f)^{-1}(\dot{f}(z)),$$

or, in coördinates, $\eta(z) = -\dot{f}(z)/f'(z)$. Therefore, $\eta(z)$ is a meromorphic vector field on $\mathbb{P}^1(\mathbb{C})$, holomorphic away from critical points of f , and with a pole of order at most m at critical points of multiplicity m .

Geometrically, $\eta(z)$ is the movement at time $t = 0$ of the point $z_t = f_t^{-1}(f(z))$. This point f_t can be followed away from critical points, by the Implicit Function Theorem.

We consider now local perturbations of f at a critical point, i.e. we assume that the vector field \dot{f} is given by a path $f_t = \phi_t \circ f \circ \psi_t^{-1}$ with ϕ_t, ψ_t analytic perturbations of the identity at the critical value and point v, c respectively of f .

Let c_t denote the critical point of f_t and let v_t denote its critical value; then $c_t = \psi_t(c)$ and $v_t = \phi_t(v)$. Let \dot{c} denote the motion vector of c_t , and let \dot{v} denote the motion vector of v_t . Then $\dot{c} = \dot{\psi}(c)$ and $\dot{v} = \dot{\phi}(v)$. Now $\dot{f} = \dot{\phi} \circ f - Df \circ \dot{\psi}$, because $\phi_0 = \psi_0 = \text{id}$. Therefore,

$$\eta + f^* \dot{\phi} = \dot{\psi}.$$

At v , the vector field $\dot{\phi}$ takes value \dot{v} , the vector field $\eta + f^* \dot{\phi}$ is holomorphic at c , and its constant term is \dot{c} . If $\dot{v} = 0$, then $f^* \dot{\phi}$ is holomorphic near c and vanishes at c .

Therefore, whenever we have a family of rational maps (f_t) for which we can follow a critical point c_t and its associated critical value v_t with $\dot{v} = 0$, the vector field η is holomorphic near c and coincides with \dot{c} at c .

3.2. Proof of Lemma 3.2. For $i \in \{3, \dots, k\}$, we have

$$f_i - f_2 = \frac{F_i}{\lambda W_1}.$$

Recall that if $\mathcal{F}(c, q, \lambda) = 0$, then $f_i = f_2$ for all $i \in \{3, \dots, k\}$. We denote by f this common rational map of degree d . If in addition $(\dot{c}, \dot{q}, \dot{\lambda})$ belongs to the Kernel of $D\mathcal{F}$ at (c, q, λ) , then

$$\dot{f}_i - \dot{f}_2 = \frac{\dot{F}_i}{\lambda W_1} - \frac{F_i \cdot (\dot{\lambda} W_1 + \lambda \dot{W}_1)}{(\lambda W_1)^2} = 0,$$

so $\dot{f}_i = \dot{f}_2$ for all $i \in \{3, \dots, k\}$. We denote by \dot{f} this common tangent vector to Rat_d at f and by η the corresponding meromorphic vector field on $\mathbb{P}^1(\mathbb{C})$.

As $(\mathbf{c}, \mathbf{q}, \lambda)$ varies in $\mathfrak{Y}_0 \times \mathbb{C}^*$, the f_i -preimages of the points $q_1 = \infty$ and q_i vary holomorphically: they are the points $c_{1,j}$ and $c_{i,j}$. According to the previous remark, we see that if $\dot{\mathbf{q}} = 0$ then, for all $i \in \{2, \dots, k\}$, the meromorphic vector field η is holomorphic near $c_{1,j}$, coincides with $\dot{c}_{1,j}$ at $c_{1,j}$, and furthermore is holomorphic near $c_{i,j}$ and coincides with $\dot{c}_{i,j}$ at $c_{i,j}$.

Thus, η is a holomorphic vector field on the whole sphere $\mathbb{P}^1(\mathbb{C})$ and coincides with $\dot{c}_{i,j}$ at $c_{i,j}$. In particular, it vanishes at $c_{1,1} = \infty$, $c_{2,1} = 0$ and $c_{3,1} = 1$. A holomorphic vector field with at least 3 zeroes globally vanishes. Therefore, $\eta = 0$ and $\dot{c}_{i,j} = 0$ for all (i, j) . In addition, for all $i \in \{1, \dots, k\}$, we have that $\dot{W}_i = 0$ and for all $i \in \{3, \dots, k\}$, we have that $0 = \dot{F}_i = \dot{\lambda} q_i W_1$. This shows that $\dot{\lambda} = 0$.

Let us summarize: if $\mathcal{F}(\mathbf{c}, \mathbf{q}, \lambda) = 0$ and if $(\dot{\mathbf{c}}, 0, \dot{\lambda})$ belongs to the Kernel of $D\mathcal{F}$ at $(\mathbf{c}, \mathbf{q}, \lambda)$, then $\dot{\mathbf{c}} = 0$ and $\dot{\lambda} = 0$. So, the restriction of $D_{(\mathbf{c}, \mathbf{q}, \lambda)}\mathcal{F}$ to $T_{\mathbf{c}}\mathcal{C}_0 \times \{0\} \times T_{\lambda}\mathbb{C}$ is injective. Since $T_{\mathbf{c}}\mathcal{C}_0 \times \{0\} \times T_{\lambda}\mathbb{C}$ and $T_0(\mathbb{C}[z]_{\deg < d})^{k-2}$ have the same dimension, that is $(k-2)d$, this restriction is an isomorphism as required.

4. FINDING A SOLUTION IN A FINITE FIELD

We describe in this section an efficient method of finding a rational function over a finite field with prescribed critical values and multiplicities.

We start by recalling some facts about univariate polynomials over non-algebraically-closed fields \mathbb{k} of arbitrary characteristic. For this we need some notation:

Notation 4.1. An ordered sequence $\alpha = (\alpha_1, \dots, \alpha_k)$ with $\alpha_1 \geq \dots \geq \alpha_k > 0$ and $\sum a_i = d$ is called a *partition* of d . With the shorthand notation

$$\beta^\mu := (\underbrace{\beta, \dots, \beta}_{\mu \text{ times}})$$

We can always write $\alpha = (\alpha_1, \dots, \alpha_k) = (\beta_1^{\mu_1}, \dots, \beta_n^{\mu_n})$ with $\beta_1 > \dots > \beta_n$ and appropriate μ_i . For example, $13 = 4 + 3 + 2 + 2 + 2$ is written as $\alpha = (4, 3, 2, 2, 2) = (4^1, 3^1, 2^3)$.

The partition α^* defined by $\alpha_j^* := \#\{i: \alpha_i \geq j\}$ is called the *dual partition* of α . For example, $\alpha^* = (5, 5, 2, 1)$.

Let $f \in \mathbb{k}[x]$ be a degree- d polynomial, let $l_i = (x - C_i) \in \overline{\mathbb{k}}[x]$ be its distinct linear factors over the algebraic closure of \mathbb{k} , and let α_i be their multiplicities, so that

$$f = \prod_{i=1}^k l_i^{\alpha_i}.$$

Without restriction we can assume $\alpha_1 \geq \dots \geq \alpha_k$ and $\alpha = (\alpha_1, \dots, \alpha_k)$ is a partition of d . In this situation we say that f is of *shape* α .

If we write $\alpha = (\alpha_1, \dots, \alpha_k) = (\beta_1^{\mu_1}, \dots, \beta_n^{\mu_n})$ as above, we can write

$$f = \prod_{i=1}^n f_i^{\beta_i}$$

with $\deg f_i = \mu_i$ and f_i the product of those linear forms that have multiplicity β_i . In this situation the f_i are coprime. \triangle

Lemma 4.2. Let \mathbb{k} be a field of characteristic p , let $f \in \mathbb{k}[x]$ be a univariate polynomial of shape α and write $f = \prod_{i=1}^k l_i^{\alpha_i}$ with $l_i \in \overline{\mathbb{k}}[x]$. If $\alpha_i < p$ for all i

then

$$\gcd(f, f') = \prod_{i=1}^k l_i^{\alpha_i - 1}.$$

Proof. We have

$$f' = \sum_{i=1}^k \alpha_i \frac{f}{l_i} l'_i = \left(\prod_{i=1}^k l_i^{\alpha_i - 1} \right) \sum_{i=1}^k \alpha_i \left(\prod_{j \neq i} l_j \right) l'_i.$$

This shows that $\prod_{i=1}^k l_i^{\alpha_i - 1}$ divides the gcd. Assume now that there is another linear factor l in the gcd. Since the gcd divides f there exists an index s with $l = l_s$. Since the gcd divides f' we have that l divides

$$\sum_{i=1}^k \alpha_i \left(\prod_{j \neq i} l_j \right) l'_i.$$

Now all summands except for $\alpha_s \left(\prod_{j \neq s} l_j \right) l'_s$ are divisible by l . Since α_s and l'_s are nonzero in \mathbb{k} and $\overline{\mathbb{k}}[x]$ respectively, it follows that l must divide $\prod_{j \neq s} l_j$. This is impossible since the l_s are pairwise coprime. \square

Corollary 4.3. *With the notations of the previous Lemma we have*

$$\gcd(f, f', \dots, f^{(e)}) = \prod_{i: \alpha_i > e} l_i^{\alpha_i - e}.$$

Proof. Lemma 4.2 and induction. \square

Corollary 4.4. *With the notations above let α^* be the dual partition of α . Then*

$$\alpha_e^* = \deg \gcd(f, f', \dots, f^{(e-1)}) - \deg \gcd(f, f', \dots, f^{(e)}).$$

Proof. We have

$$g_e := \frac{\gcd(f, f', \dots, f^{e-1})}{\gcd(f, f', \dots, f^e)} = \frac{\prod_{i: \alpha_i > e-1} l_i^{\alpha_i - e + 1}}{\prod_{i: \alpha_i > e} l_i^{\alpha_i - e}} = \prod_{i: \alpha_i \geq e} l_i.$$

It follows that

$$\deg g_e = \deg \prod_{i: \alpha_i \geq e} l_i = \#\{i: \alpha_i \geq e\} = \alpha_e^*. \quad \square$$

Algorithm 4.5 (Compute the shape of a polynomial).

Given: a polynomial $f \in \mathbb{k}[x]$

Return: the shape $(\alpha_1, \dots, \alpha_k)$ of f .

Write $d = \deg(f)$. For each $e = 0, \dots, d$, compute $g_e = \gcd(f, f', \dots, f^{(e)})$. For each $e = 1, \dots, d$ define then $\alpha_e^* = g_{e-1} - g_e$. Return the dual of the partition $(\alpha_1^*, \dots, \alpha_d^*)$.

Proof of validity. This directly follows from Corollary 4.4. \square

We may collect linear factors of the same multiplicity, so as to avoid field extensions:

Corollary 4.6. *With the notation of Lemma 4.2 choose $\beta_1 > \dots > \beta_n$ among the α_i such that $f = \prod_{i=1}^n f_i^{\beta_i}$ with $f_j = \prod_{i: \alpha_i = \beta_j} l_i$. Then the f_i are defined over \mathbb{k} .*

Proof. Since the calculation of a gcd does not require field extensions, we have

$$\gcd(f, f', \dots, f^{(e)}) \in \mathbb{k}[x].$$

By Corollary 4.3 we then have

$$g_j := \frac{\gcd(f, f', \dots, f^{(\beta_j-1)})}{\gcd(f, f', \dots, f^{(\beta_j)})} = \frac{\prod_{i: \beta_i > \beta_j-1} f_i^{\beta_i - \beta_j + 1}}{\prod_{i: \beta_i > \beta_j} f_i^{\beta_i - \beta_j}} = \prod_{i=1}^j f_i \in \mathbb{k}[x]$$

so

$$f_j = \frac{g_j}{g_{j+1}} \in \overline{\mathbb{k}}[x] \cap \mathbb{k}(x) = \mathbb{k}[x]. \quad \square$$

We are now ready to describe our algorithm searching for rational maps over \mathbb{F}_p .

Algorithm 4.7 (Compute all rational maps over \mathbb{k} with given shape above given points).

Given: a finite field \mathbb{k} , a list of points $\overline{Q}_1, \dots, \overline{Q}_k \in \mathbb{P}^1(\mathbb{k})$, an integer d , and a list of partitions $(\alpha_1, \dots, \alpha_k)$ of d with $\alpha_i = (\alpha_{i,1}, \dots, \alpha_{i,\ell_i})$ satisfying $\sum_i \sum_j (\alpha_{i,j} - 1) = 2d - 2$

Return: all rational maps over \mathbb{k} of degree d such that every \overline{Q}_i has ℓ_i preimages with local degrees $\alpha_{i,1}, \dots, \alpha_{i,\ell_i}$ respectively.

We choose a Möbius transformation $M \in \mathrm{PSL}_2(\mathbb{k})$ sending \overline{Q}_1 to ∞ and \overline{Q}_2 to 0.

We write each partition α_i in compacted form as $\alpha_i = (\beta_{i,1}^{\mu_{i,1}}, \dots, \beta_{i,n_i}^{\mu_{i,n_i}})$.

We enumerate all ℓ_1 -tuples of monic polynomials (f_1, \dots, f_{n_1}) with $\deg(f_j) = \mu_{1,j}$, and all ℓ_2 -tuples of monic polynomials (g_1, \dots, g_{n_2}) with $\deg(g_j) = \mu_{2,j}$.

For each such pair of tuples, we compute

$$W_1 = \prod_{j=1}^{n_1} f_j^{\beta_{1,j}} \quad \text{and} \quad W_2 = \prod_{j=1}^{n_2} g_j^{\beta_{2,j}}.$$

Using Algorithm 4.5, we filter those (W_1, W_2) such that the shape of W_1 is α_1 and the shape of W_2 is α_2 (this fails only if a pair f_i, g_j is not coprime).

By computing their g.c.d., we filter those (W_1, W_2) such that W_1 and W_2 are coprime.

For each $i = 3, \dots, k$, let $\Lambda_i \subset \mathbb{k}$ be the set of $\lambda \in \mathbb{k}$ such that the shape (computed using Algorithm 4.5) of $W_i := W_2 - \lambda M(\overline{Q}_i)W_1$ is α_i . We filter those rational maps for which $\bigcap_{i=3}^k \Lambda_i$ is non-empty.

We return all the rational maps $M^{-1} \circ (W_2/\lambda W_1)$, for all $\lambda \in \bigcap_{i=3}^k \Lambda_i$, that survived the filtering.

Proof of validity. Let first $f := M^{-1} \circ (W_2/\lambda W_1)$ be a rational map returned by the algorithm. For $i = 2, \dots, k$, consider the rational map $f_i = M^{-1} \circ (W_i/\lambda W_1)$. By the very definition of W_i (compare with (3)), we have $f = f_2 = \dots = f_k$. On the other hand, the f_i -preimages of \overline{Q}_i are the zeroes of W_i , so they have multiplicities α_i .

On the other hand, let f be a rational map such that every \overline{Q}_i has ℓ_i preimages with local degrees $\alpha_{i,1}, \dots, \alpha_{i,\ell_i}$ respectively. Then, for every Möbius transformation M sending \overline{Q}_1 to ∞ and \overline{Q}_2 to 0, the rational map $M \circ f$ will be of the form $W_2/\lambda W_1$, for monic polynomials W_1, W_2 of respective shapes α_1, α_2 and a scalar $\lambda \in \mathbb{k}$. Furthermore, by Corollary 4.6, both W_1 and W_2 factor over $\mathbb{k}[x]$ into polynomials of degrees $\mu_{1,j}$ and $\mu_{2,j}$ respectively. The rational map $M \circ f - \overline{Q}_i$ will be of

the form $W_i/\lambda W_1$ with W_i of shape α_i ; therefore, that solution f will be returned by the algorithm. \square

Remark 4.8. Algorithm 4.7 is the most computationally-intensive part of our procedure. Its performance is improved in the following ways:

- (1) If $\mu_{i,j} = 1$ for some i, j , then we may assume, after permuting the shapes α_i , that $i = 1$ so that W_1 contains a power of a linear factor $f_i^{\beta_{i,j}}$. Fixing the corresponding preimage of ∞ to be ∞ amounts to the choice $f_i = 1$, so that the degree of W_1 is actually $d - \beta_{i,j}$. This speeds up the search by a factor p .

Similarly, if up to permutation of the indices there are more $\mu_{i,j} = 1$, with $i \in \{1, 2\}$, then the corresponding factors may be assumed to be x and $x - 1$.

On the other hand, if all $\mu_{i,j} \geq 2$, then no normalization of the critical points may be assumed, and in particular ∞ should not be assumed to be a preimage of some \overline{Q}_i .

- (2) When using Corollary 4.4 one can detect a wrong shape already if $\gcd(f, f')$ or for that matter any $\gcd(f, \dots, f^{(e)})$ with $e < \alpha_1$ has the wrong degree. We stop the calculation of gcd's as soon as this happens. This speeds up the process by a factor of about α_1 . Similarly, as soon as the intersection of the Λ_i already computed is empty, the pair (W_1, W_2) should be discarded.
- (3) If the largest $\mu_{i,j}$ is small enough (e.g. 7 or 8 in \mathbb{F}_{11}) we can enumerate all monic irreducible homogeneous polynomials of degree $\leq \mu_{i,j}$ and build the f_i and g_i out of them, while taking care that no irreducible piece is used twice. We can then omit checking shape and coprimeness of W_1 and W_2 as these conditions are then automatically satisfied.

Example 4.9. Over \mathbb{F}_{11} we searched for a rational map of shape $(4, 3, 2, 2, 2)$, $(4, 3, 2, 2, 2)$ and $(4, 3, 2, 2, 2)$; we chose $\overline{Q}_1 = \infty$ and $\overline{Q}_2 = 0$, and didn't specify \overline{Q}_3 , letting on the contrary the algorithm determine choose it for us. We found the solution

$$\frac{W_2}{W_1} = \frac{x^4(x+3)^3(x^3-3x-5)^2}{(x-5)^3(x^3+3x^2+2x+3)^2}.$$

It has indeed the desired shape as we have the following factorisation

$$W_2 + 4W_1 = (x-1)^4(x-3)^3(x^3-2x-3)^2,$$

which implies, for the choice $\overline{Q}_3 = 1$, the value $\lambda = -4$.

5. LIFTING A SOLUTION FROM \mathbb{F}_p TO \mathbb{Z}_p

The lift from \mathbb{F}_p to \mathbb{Z}_p is done using Hensel's lemma (namely, Newton's method in positive characteristic):

Proposition 5.1 (Hensel's Lemma). *Let $\mathcal{F} = (F_1, \dots, F_m)$ be a vector of polynomials, with $F_i \in \mathbb{Z}[x_1, \dots, x_m]$, and let $J = (\frac{dF_i}{dx_j})$ be the Jacobian matrix of \mathcal{F} . Assume that $a = (a_1, \dots, a_m) \in \mathbb{Z}^m$ satisfies*

$$\mathcal{F}(a) \equiv 0 \pmod{p^N},$$

that $J(a)$ is invertible modulo p^N , and let $J^{-1}(a)$ be an inverse modulo p^N . Then

$$\mathcal{F}(a + bp^N) \equiv 0 \pmod{p^{2N}},$$

for

$$b := -\frac{F(a)}{p^N} J^{-1}(a).$$

Furthermore $J(a + bp^N)$ is invertible modulo p^{2N} .

Proof. $F(a)$ is divisible by p^N since $F(a) \equiv 0 \pmod{p^N}$. Therefore b is well defined. We have

$$\begin{aligned} F(a + p^N b) &\equiv F(a) + bJ(a)p^N \pmod{p^{2N}} \\ &\equiv F(a) - \frac{F(a)}{p^N} J^{-1}(a)J(a)p^N \pmod{p^{2N}} \\ &\equiv 0 \pmod{p^{2N}}. \end{aligned}$$

The invertibility holds more generally. Let A and B be matrices with $AB \equiv \mathbb{1} \pmod{p^N}$. We can then write

$$AB \equiv \mathbb{1} + p^N C \pmod{p^{2N}}.$$

In this situation we have

$$\begin{aligned} A(B - p^N BC) &\equiv \mathbb{1} + p^N C - p^N ABC \pmod{p^{2N}} \\ &\equiv \mathbb{1} \pmod{p^{2N}} \end{aligned}$$

since $AB \equiv \mathbb{1} \pmod{p^N}$; so $B' = B - p^N BC$ is an inverse to A modulo p^{2N} . \square

Consider the following data: a ring \mathbb{k} ; a family of polynomials $W_i \in \mathbb{k}[x]$ of degree at most d , for $i = 1, \dots, k$, with factorisations $W_i = \prod_{j=1}^{n_j} W_{i,j}^{\beta_{i,j}}$; a parameter $\lambda \in \mathbb{k}^\times$; and a sequence of points $Q_1 = \infty, Q_2, \dots, Q_k \in \mathbb{P}^1(\mathbb{k})$. We say that they are *coherent* if $W_i/\lambda W_1 + Q_i$ is independent of $i = 2, \dots, k$.

We say that they are *normalised* if the following holds: the first three values Q_1, Q_2, Q_3 are $\infty, 0, 1$ respectively; and the first three preimages $C_{1,1}, C_{2,1}, C_{3,1}$ are also respectively $\infty, 0, 1$. This means that we assume that W_1 has degree $d - \alpha_{1,1}$, that $W_{2,1} = x^{\alpha_{2,1}}$, and that $W_{3,1} = (x - 1)^{\alpha_{3,1}}$.

Note that this assumption is not innocuous: it may well be that no critical point $C_{i,j}$ is defined over \mathbb{k} . The normalization may be imposed at no cost if (after permutation of the indices) $\mu_{1,1} = \mu_{2,1} = \mu_{3,1} = 1$.

We are now ready to detail the lifting algorithm. Out of coherent data in \mathbb{F}_p and a parameter N , it computes a p^{-N} -approximation of the corresponding coherent data in \mathbb{Z}_p , in the form of an approximation in \mathbb{Z}/p^N .

Algorithm 5.2 (Lift a solution p -adically).

Given: coherent data $\overline{W}_i = \prod_{j=1}^{n_j} \overline{W}_{i,j}^{\beta_{i,j}} \in \mathbb{F}_p[x]$, $\overline{\lambda} \in \mathbb{F}_p^\times$, and $\overline{Q}_i \in \mathbb{P}^1(\mathbb{F}_p)$; a parameter $N \in \mathbb{N}$; and lifts $Q_i \in \mathbb{P}^1(\mathbb{Z}/p^N)$ of the points \overline{Q}_i

Return for infinitely many p : coherent data $W_i = \prod_{j=1}^{n_j} W_{i,j}^{\beta_{i,j}} \in (\mathbb{Z}/p^N)[x]$ and $\lambda \in (\mathbb{Z}/p^N)^\times$ that reduce mod p to \overline{W}_i .

First, we assume that the data may be normalised. This amounts to requiring at least three of the $W_{i,j}$, for distinct i 's, to have a linear factor. This holds for a positive proportion of primes p . If no such three factors exist, the algorithm aborts. Otherwise, we silently replace the three corresponding $\beta_{i,j}^{\mu_{i,j}}$ by $(\beta_{i,j}^1, \beta_{i,j}^{\mu_{i,j}-1})$ in the shapes so as to create a term with $\mu_{i,j} = 1$.

We write now each $W_{i,j}$ in the form

$$(4) \quad W_{i,j} = x^{\mu_{i,j}} + \sum_{s=1}^{\mu_{i,j}} w_{i,j,s} x^{\mu_{i,j}-s},$$

for unknowns $w_{i,j,s} \in \mathbb{Z}/p^N$.

Recall from (3) the expressions $F_i = W_i + \lambda q_i W_1 - W_2$ and $\mathcal{F} = (F_3, \dots, F_k)$. The F_i are polynomials in the variables $\{w_{i,j,s} : (i,j) \neq (1,1), (2,1), (3,1)\} \cup \{\lambda\}$. We lift the coefficients of the coherent data $(\overline{W}_i, \overline{\lambda})$ to \mathbb{Z} , to obtain an initial parameter $a^0 = (w_{1,2,1}^0, \dots, w_{k,n_k,\mu_k,n_k}^0, \lambda^0)$. Since the original data is coherent, we have $\mathcal{F}(a^0) \equiv 0 \pmod{p}$. For almost all p , the Jacobian $D\mathcal{F}$ is invertible by Corollary 3.3; if $D\mathcal{F}$ is not invertible at a^0 , then we abort the algorithm. Otherwise, we apply repeatedly Hensel's Lemma 5.1 to obtain a solution a to $\mathcal{F}(a) \equiv 0 \pmod{p^N}$.

Finally, we reconstruct the polynomials $W_{i,j}$ out of their coefficients (which are just coordinates of a).

Proof of validity. The invertibility of the Jacobian was expressed in Corollary 3.3 in terms of the variables $C_{i,j}$. This does not make any difference: here we express them in terms of the $w_{i,j,s}$, which are elementary symmetric functions of the $C_{i,j}$. \square

Example 5.3. Consider the shapes $\alpha_1 = \alpha_2 = \alpha_3 = (4^1, 3^1, 2^3)$. Our example

$$\begin{aligned} W_1 &= (x-5)^3(x^3+3x^2+2x+3)^2 \\ W_2 &= x^4(x+3)^3(x^3-3x-5)^2 \\ W_3 &= W_2 + 4W_1 = (x-1)^4(x-3)^3(x^3-2x-3)^2 \end{aligned}$$

gives a vector of coefficients

$$\begin{aligned} a^0 &= (w_{1,2,1}, w_{1,3,1}, w_{1,3,2}, w_{1,3,3}, \\ &\quad w_{2,2,1}, w_{2,3,1}, w_{2,3,2}, w_{2,3,3}, \\ &\quad w_{3,2,1}, w_{3,3,1}, w_{3,3,2}, w_{3,3,3}, \lambda) \\ &= (-5, 3, 2, 3, 3, -3, 0, -5, -3, -2, 0, -3, -4) \end{aligned}$$

with $F(w_0) \equiv 0 \pmod{11}$. The lift is

$$a^1 = (50, -41, 13, 25, -19, -33, 19, -60, -47, 11, -46, -58, 51)$$

with $F(w_1) \equiv 0 \pmod{11^2}$. We can continue this process inductively. Notice that the precision doubles in every step.

6. PROMOTING A SOLUTION FROM \mathbb{Z}_p TO A NUMBER FIELD \mathbb{K}

If the Hurwitz problem has a solution over \mathbb{Z} that reduces to a given solution over \mathbb{F}_p , then Hensel lifting will find it after a finite number of steps. Unfortunately the solutions usually involve fractional coefficients, and are usually defined over a finite extension \mathbb{K} of \mathbb{Q} . Our first goal will therefore be to determine this extension.

Consider a degree- e extension \mathbb{K} of the rationals, and $a \in \mathbb{K}$. Then $1, a, \dots, a^e$ are linearly dependent over \mathbb{Q} , and therefore also over \mathbb{Z} , i.e there exists a polynomial

$$P = p_0 + p_1 t + \dots + p_e t^e$$

with all $p_i \in \mathbb{Z}$ and $P(a) = 0$. Let now $p \in \mathbb{N}$ be a prime such that P splits over \mathbb{Z}_p ; so that we may view $\mathbb{Q}(a)$ as a subfield of \mathbb{Q}_p . Assume also that a is invertible modulo p , so that we may consider a as an element of \mathbb{Z}_p .

Consider now $\tilde{a} \equiv a \pmod{p^N}$; then we have the equation

$$P(\tilde{a}) = p_{e+1}p^N$$

which is linear in $\{p_0, \dots, p_{e+1}\}$. We use the LLL algorithm [10] to find small integer solutions to this linear equation. The default implementation uses a simple heuristic to guess the correct precision N and the correct extension degree: for a initial precision we start with extension degree $e = 1$ and increase e until a solution is found (i.e. $F(\tilde{a}) = 0 \pmod{p^{2N}}$) or the computed shortest lattice basis vector norm is the same for e and $e - 1$. If the computed vector norm did not change, we increase the p -adic precision. If we have a-priori knowledge about the minimum or maximum expected extension degree, then it can be passed to the algorithm, which is more likely to find quickly a solution.

The following algorithm is described as a process that, receiving as input an infinite feed of ever-more-precise approximations of a p -adic number that is known to be algebraic, produces an infinite stream of ever-more-likely minimal polynomials of that p -adic number.

Algorithm 6.1 (Convert a p -adic number to an algebraic number).

Given: approximations, to arbitrary precision, of an algebraic number $a \in \mathbb{Z}_p$

Return: polynomials $P(t) \in \mathbb{Z}[t]$ whose likelihood converges to 1 of being the minimal polynomial of a , as the precision of a improves.

Assume that, for each $N \in \mathbb{N}$, the algorithm may receive an approximation a_N , to N base- p digits, of a . The element a_N is represented as an integer in $\{0, \dots, p^N - 1\}$.

Start with $d = 1$ and $N = 1$. Then, repeat the following. Consider the lattice in \mathbb{R}^{d+1} generated by the columns of the matrix

$$M = \begin{pmatrix} p^N & -a_N & -a_N^2 & \dots & -a_N^d \\ 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \end{pmatrix}.$$

Using the LLL algorithm, find a vector (P_0, P_1, \dots, P_d) in the lattice, of small norm $\theta_{N,d}$. Form the polynomial $P(t) = \sum_{i=0}^d P_i t^i$.

If $P(a_N) \equiv 0 \pmod{p^{2N}}$, or if $d > 1$ and $\theta_{N,d} = \theta_{N,d-1}$, then output P as a candidate polynomial. Repeat then, after having incremented d in the first case holds, and doubled N otherwise.

Proof of validity. The algorithm repeatedly increases N and d . Note that the polynomials returned may have degree $< d$, so increasing d is harmless, and the precision is increased as soon as increase in maximal degree does not improve the solution.

Let (P_0, \dots, P_d) be a short lattice vector. Then this vector is

$$M \cdot {}^t(P(a_N)/p^N, P_1, \dots, P_d)$$

and in particular $P(a_N) \equiv 0 \pmod{p^N}$. On the other hand, the coefficients P_i are small, so P is likely to be the minimal polynomial of a . \square

Example 6.2. Considering our example $\alpha_1 = \alpha_2 = \alpha_3 = (4^1, 3^1, 2^3)$ we found

$$w_{1,2,1} = 1400834756308742009361916361765119584358776523123371526525883115012 \in \mathbb{Z}/11^{26}$$

and $P_{1,2,1}(w_{1,2,1}) \equiv 0 \pmod{11^{2^6}}$ for

$$P_{1,2,1}(t) = 39t^6 + 117t^5 + 195t^4 + 195t^3 + 141t^2 + 63t + 16.$$

Higher precision values of $w_{1,2,1}$ are also zeroes of the same polynomial $P_{1,2,1}$. We take this as a hint that $P_{1,2,1}$ is indeed the minimal polynomial of the coördinate $w_{1,2,1}$ in the lift of our finite field solution.

Having found the minimal polynomials $P_{i,j,s} \in \mathbb{Z}[t]$ for all coördinates $w_{i,j,s}$ of our solution vector, we determine the field \mathbb{K} on which they are all defined, as the compositum of all field extensions defined by the $P_{i,j,s}$. If these field extensions were independent, then we should just consider all zeroes in \mathbb{C} of the $P_{i,j,s}$ and return the corresponding rational functions.

However, in general, the field extensions will be highly dependent. To simplify notation, let us assume that all $P_{i,j,s}$ are of degree e , and that \mathbb{K} itself is a degree- e extension. Then there are e possible values for each coördinate. At this stage we do not know how to combine these single coördinate solutions to a solution vector (there are $e^{d(k-2)}$ possible combinations). To solve this problem we use the following method.

To illustrate our method, consider $a, b \in \mathbb{K}$ and let $P_a, P_b, P_{a+b} \in \mathbb{Z}[x]$ be minimal polynomials of $a, b, a+b$ respectively. Assume furthermore that P_a, P_b and P_{a+b} are of degree e , and let $a_1, \dots, a_e, b_1, \dots, b_e$ and c_1, \dots, c_e respectively be approximations over \mathbb{C} of the zeroes of P_a, P_b and P_{a+b} . Consider the $e \times e$ “root compatibility matrix” $M = (M_{i,j})$ defined by

$$M_{i,j} = \begin{cases} 1 & \text{if there exists } k \text{ with } a_i + b_j \approx c_k, \\ 0 & \text{otherwise.} \end{cases}$$

If M is a permutation matrix, then it describes which root b_j should be paired with a_i , namely it is characterised by $M_{i,i} = 1$. In this manner, all other coördinates are chosen, dependent on the first choice of a root of $P_{1,2,1}$.

Example 6.3. Tentative minimal polynomials of $w_{1,3,1}$ and $a = w_{1,2,1} + w_{1,3,1}$ are

$$\begin{aligned} P_{1,3,1} &= 28431t^6 + 255879t^5 + 982449t^4 + 2056509t^3 + 2465721t^2 + 1597239t + 439138 \\ P_a &= 28431t^6 + 341172t^5 + 1844856t^4 + 5660928t^3 + 10384524t^2 + 10807344t + 5068144 \end{aligned}$$

We obtain the following approximate zeroes over \mathbb{C} using Brent’s method, implemented in PARI [2]; we preserve the ordering in which the roots were returned.

$w_{1,2,1}$	$w_{1,3,1}$	$w_{1,2,1} + w_{1,3,1}$
$-.150 + .807i$	$-1.5 - 1.02i$	$-2.161 - 1.184i$
$-.150 - .807i$	$-1.5 + 1.02i$	$-2.162 + 1.184i$
$-.5 + .440i$	$-2.012 - .377i$	$-2 - 1.462i$
$-.5 - .440i$	$-2.012 + .377i$	$-2 + 1.462i$
$-.850 + .807i$	$-.988 - .377i$	$-1.839 - 1.184i$
$-.850 - .807i$	$-.988 + .377i$	$-1.839 + 1.184i$

Now consider the compatibility matrix M . It is

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

i.e. $M_{i,j}$ is 1 precisely when $(w_{1,2,1})_i + (w_{1,3,1})_j$ approximates one of the values in the last column of the table above. We note that M is a permutation matrix. Applying the permutation M to the list of roots $(w_{1,3,1})_i$ of $P_{1,3,1}$ leads to a valid coordinate pairing: now $(Mw_{1,3,1})_i$ corresponds to the $(w_{1,2,1})_i$.

Example 6.4. There are situations in which the matrix M is not a permutation matrix but nevertheless contains useful information. Consider the finite algebraic set

$$S = \left\{ (i, \sqrt{2}), (-i, \sqrt{2}), (i, -\sqrt{2}), (-i, -\sqrt{2}) \right\}.$$

In this case the minimal polynomials of the coordinates are $x^2 + 1$ and $y^2 - 2$ and the root compatibility matrix is $M = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$. All 4 possible pairings lead to correct solutions.

There are also situations in which the matrix M differs from the permutation matrix giving the correct identification of coordinates, even with exact arithmetic. For example, if ξ denotes a fifth root of unity, and

$$S = \left\{ (\xi, \xi^2), (\xi^2, \xi^4), (\xi^3, \xi), (\xi^4, \xi^3) \right\}$$

then the coordinates have respectively $x^4 + x^3 + x^2 + x + 1$ and $y^4 + y^3 + y^2 + y + 1$ as their minimal polynomial, and the root compatibility matrix is insufficient to recover S . Indeed the sum of the coordinates $z = x + y$, which has minimal polynomial $z^4 + 2z^3 + 4z^2 + 3z + 1$, does not distinguish solutions in S from the non-solutions

$$S' = \left\{ (\xi, \xi^3), (\xi^2, \xi), (\xi^3, \xi^4), (\xi^4, \xi^2) \right\}$$

Note that S and S' are distinguished by the equation $x^2 - y$ which holds in S but not in S' . A linear form $ax + by$ with $0 \neq a \neq b \neq 0$ would also do.

There are also situations with more than two variables in which all compatibility matrices between two variables lead to possible pairings, but their combined information is not enough to identify the correct solutions. For example, consider

$$S = \left\{ (\varepsilon_1\sqrt{2}, \varepsilon_2\sqrt{3}, \varepsilon_3\sqrt{5}) : \varepsilon_i \in \{\pm 1\}, \varepsilon_1\varepsilon_2\varepsilon_3 = 1 \right\}.$$

All pairings between two coordinates are allowed, but there are 4 solutions in total, and not 8.

Above we have used the linear forms $x_i + x_j$ to determine the compatibility. Our algorithm uses random linear forms to avoid these problems, or at least make them less probable.

We are now ready to explain the algorithm computing the solutions in number fields that reduce modulo p^N to given approximate solutions in \mathbb{Z}_p . That problem is in fact an instance of the following, more general problem.

The following algorithm is described as a process that, receiving as input an algebraic system over \mathbb{Z} and an infinite feed of ever-more-precise p -adic approximations of a solution, produces a stream of algebraic solutions (in the form of a minimal polynomials over \mathbb{Z} and a complex numbers singling out roots of the minimal polynomials). The stream eventually exhausts all solutions conjugate to the p -adic solution.

Algorithm 6.5 (Solve 0-dimensional algebraic systems).

Given: a polynomial system of equations $\mathcal{F} = (F_1, \dots, F_m)$ in variables x_1, \dots, x_m , having a finite number of solutions; and approximations, to arbitrary precision, of a solution $(\widehat{x}_1, \dots, \widehat{x}_m)$ in \mathbb{Z}_p

Return: a number field $\mathbb{K} \subset \mathbb{C}$, and exact solutions (x_1^i, \dots, x_m^i) in \mathbb{K} , for $i = 1, \dots, s$; each element of \mathbb{K} is given by its minimal polynomial and an approximation in \mathbb{C} of a particular root.

We construct the solutions (x_1^i, \dots, x_m^i) iteratively, entry by entry, by constructing partial tables $\{(x_1^i, \dots, x_k^i) : i = 1, \dots, s\}$. We start by an empty table, with $s = 1$ and $k = 0$, and let ϵ be a small number.

Then, for each $k = 1, \dots, m$, we do the following. Using Algorithm 6.1, we compute a likely minimal polynomial P_k of \widehat{x}_k , say of degree e . We compute, to precision better than ϵ , approximate roots $r_1, \dots, r_e \in \mathbb{C}$ of P_k . Set

$$\delta_1 = \min\{|r_i - r_j| : 1 \leq i \neq j \leq e\}.$$

If $\delta_1 < \epsilon$, we halve ϵ and restart all over.

We next choose randomly a linear form $L = l_1x_1 + \dots + l_kx_k$ with $l_k \neq 0$. Again using Algorithm 6.1, we compute a minimal polynomial P_L for $L(\widehat{x}_1, \dots, \widehat{x}_k)$. If the degree of P_L is not divisible by s , we choose a different linear form L and repeat the above. Otherwise, we let δ_1 be the minimal distance between roots of P_L . If $\delta_1 < \epsilon$, we halve ϵ and restart all over.

We then compute the $s \times e$ matrix $M = (M_{i,j})$ with

$$M_{i,j} = \begin{cases} 1 & \text{if } |P_L(L(x_1^i, \dots, x_{k-1}^i, r_j))| < \frac{1}{3} \min\{\delta_1, \delta_2\} \|L\|_1, \\ 0 & \text{otherwise.} \end{cases}$$

Let the degree of P_L be st . If M contains t ones per row and one one per column, then we replace s by st and replace each row $(x_1^i, \dots, x_{k-1}^i)$ is the partial table by t rows $(x_1^i, \dots, x_{k-1}^i, r_j)$ for all j such that $M_{i,j} = 1$. Otherwise, we repeat the step with a different linear form or, if that failed more than ten times in a row, we simply skip the iteration.

When the iteration finished with $k = m$, we have obtained s candidate solutions, which we check algebraically by evaluating \mathcal{F} on them. We output all those that are certifiably valid solutions, and restart the algorithm with better approximations of the $\widehat{x}_1, \dots, \widehat{x}_m$.

Proof of validity. First, all the solutions returned are valid, since they were checked (using exact algebra) by evaluating \mathcal{F} on them.

Let now (x_1, \dots, x_m) be a solution that is conjugate to $(\widehat{x}_1, \dots, \widehat{x}_m)$. In particular, the minimal polynomials of the x_i and \widehat{x}_i are the same, so they will eventually be found by Algorithm 6.1. Similarly, for every linear form L with integer coefficients, $L(x_1, \dots, x_m)$ and $L(\widehat{x}_1, \dots, \widehat{x}_m)$ also have the same minimal polynomial, so it will also be eventually found by Algorithm 6.1. \square

We apply this algorithm to the same polynomial equations (3). The variables x_i are a relabeling of the $w_{i,j,s}$ from (4).

7. COMPUTING THE MONODROMY

In this section, we detail the second part of the algorithm sketched in §2.

We are given an approximation of a degree- d rational map $f \in \mathbb{C}(z)$, as well as an approximation of the critical values $Q \subset \mathbb{C}$, and the local degrees $\alpha_{i,1}, \dots, \alpha_{i,\ell_i}$ above each critical value Q_i . We are asked to compute the monodromy of the covering induced by f .

The first step is to compute a triangulation \mathcal{Q} of $\mathbb{P}^1(\mathbb{C})$ by arcs of circle, and containing Q among its vertices. A particularly efficient triangulation is the *Delaunay triangulation*. This is a decomposition of $\mathbb{P}^1(\mathbb{C})$ into triangles, such that, for any two triangles with a common edge, the sum of their opposite angles is $> \pi$. Such a triangulation always exists; is essentially unique; and may be computed e.g. using [14].

For performance reasons, we refine the triangulation by adding vertices to it: whenever we encounter a triangle whose ratio “circumradius / shortest side” is larger than 1000, we add the circumcenter to the triangulation. This process converges, and gives a reasonably good triangulation in that its triangles are not too acute; see [15].

The dual decomposition \mathcal{Q}^\perp of the sphere is the associated *Voronoi diagram*. It has one vertex, called a *dual vertex*, per Delaunay triangle and one edge, called *dual edge*, across every Delaunay edge. Each of its edges ε is parametrised as the preimage, under a Möbius transformation μ_ε , of the arc $[0, 1]$. We denote by W the vertex set of \mathcal{Q}^\perp , and choose a basepoint $* \in W$. For each $w \in W$, we number arbitrarily the elements of the fibre $f^{-1}(w)$ as $\{w_1, \dots, w_d\}$. Because W is far from Q , there are d preimages of each $w \in W$, and their computation is numerically stable.

There are now two strategies, which have both been tested and implemented. The first one is a bit simpler, but the second one performs better in practice. Both associate a permutation ς_ε of $\{1, \dots, d\}$ with each edge $\varepsilon \in \mathcal{Q}^\perp$ from $w' \in W$ to $w'' \in W$, in such a way that the the f -lift of ε that starts at w'_j ends at $w''_{\varsigma_\varepsilon(j)}$. Both are explained below; assuming them, we finish the description of the algorithm.

Let $(\hat{\gamma}_i)_{i=1, \dots, k}$ be non-crossing (but possibly overlapping) paths in \mathcal{Q}^\perp that start and end in $*$, cyclically ordered around $*$, such that $\hat{\gamma}_i$ surrounds once counterclockwise the point Q_i and no other vertex of Q . These paths may be selected as follows: choose first the path $\hat{\gamma}_1$ arbitrarily, and mark its edges. Then, for $i = 2, \dots, k$, choose the path $\hat{\gamma}_i$ in such a manner that it does not cross the previously chosen paths (i.e., it may follow a marked path, but must depart from it on the same side as it joined it), and starts at $*$ in counterclockwise order between the paths $\hat{\gamma}_{i-1}$ and $\hat{\gamma}_1$.

These paths form a basis for the fundamental group $\pi_1(\mathbb{P}^1(\mathbb{C}) \setminus Q, *)$, compatible with the description from §1.1. Let $(\varepsilon_{i,1}, \dots, \varepsilon_{i,n_i})$ be the edges along $\hat{\gamma}_i$, and compute the permutation $\sigma_i = \varsigma_{\varepsilon_{i,1}} \cdots \varsigma_{\varepsilon_{i,n_i}}$. Then the monodromy representation of f is given by the family $(\sigma_i)_{i=1, \dots, k}$.

7.1. Connect-the-dots. For each dual edge $\varepsilon \in \mathcal{Q}^\perp$, going from w' to w'' , we do the following. Knowing the spherical distance from w' to w'' and using coarse estimates on $|f'(z)|$, we have an upper bound on the length of each of the d preimages

of ε . We attempt to match each w'_j with a $w''_{\varsigma(j)}$ for some permutation $\varsigma \in \mathfrak{S}_d$, by matching each w'_j to the closest $w''_{\varsigma(j)}$. If more than one match is compatible with the upper bound on the length of an arc above the arc from w' to w'' , we subdivide the edge ε .

Algorithm 7.1 (Lifting edges by connect-the-dots).

Given: a rational map $f \in \mathbb{C}(z)$, an edge $\varepsilon = \mu_\varepsilon^{-1}[0, 1]$ from $w' = \mu_\varepsilon^{-1}(0)$ to $w'' = \varepsilon^{-1}(1)$ and orderings w'_1, \dots, w'_d and w''_1, \dots, w''_d of the preimages of w', w'' respectively

Return: a permutation $\varsigma \in \mathfrak{S}_d$ such that the f -lift of ε starting at w'_j ends at $w''_{\varsigma(j)}$

More generally, the algorithm computes, for arbitrary $0 \leq s < t \leq 1$, the matching between f -preimages of $\mu_\varepsilon^{-1}(s)$ and of $\mu_\varepsilon^{-1}(t)$; the solution is provided by the matching for $s = 0$ and $t = 1$. We write $\{w_j(t)\}_{j=1, \dots, d}$ for the f -preimages of $\mu_\varepsilon^{-1}(t)$.

If there is only one possible match between the sets $\{w_j(s)\}$ and $\{w_j(t)\}$, given by a permutation ς , then the algorithm returns that permutation. Otherwise, set $u = (s+t)/2$, compute the preimages $\{w_j(u)\}_{j=1, \dots, d}$ of $\mu_\varepsilon^{-1}(u)$, recursively compute the matching between the $w_j(s)$ and $w_j(u)$ and between the $w_j(u)$ and $w_j(t)$, and return the product of the corresponding permutations.

7.2. Using two triangulations. The second algorithm is more efficient, and uses fundamentally the fact that the arcs in the triangulation \mathcal{Q}^\perp and in its f -preimage are given by algebraic curves.

We initially compute the Delaunay triangulation \mathcal{C} on $f^{-1}(Q)$, and parametrise its edges e via Möbius transformations ν_e such that $e = \nu_e^{-1}([0, 1])$.

It is straightforward to lift \mathcal{Q}^\perp through f : its edges are all the curves defined by equations $(\mu_\varepsilon \circ f)(z) \in [0, 1]$.

Algorithm 7.2 (Lifting edges using two triangulations).

Given: a rational map $f \in \mathbb{C}(z)$, an edge $\varepsilon = \mu_\varepsilon^{-1}[0, 1]$ from $w' = \mu_\varepsilon^{-1}(0)$ to $w'' = \varepsilon^{-1}(1)$ and orderings w'_1, \dots, w'_d and w''_1, \dots, w''_d of the preimages of w', w'' respectively

Return: a permutation $\varsigma \in \mathfrak{S}_d$ such that the f -lift of ε starting at w'_j ends at $w''_{\varsigma(j)}$

For each $i = 1, \dots, d$, we seek the $j \in \{1, \dots, d\}$ such that the lift of ε starting at w'_i ends at w''_j . The permutation to return is then the map $i \mapsto j$.

We first determine in which triangle T of \mathcal{C} the lift w'_i lies. Then we compute whether the lift $\tilde{\varepsilon}$ of ε starting at w'_i leaves T . If this happens, then it must cross an edge e of T , namely, we have $\mu_\varepsilon \circ f \circ \nu_e^{-1}([0, 1]) \in [0, 1]$. This entails, firstly, that the imaginary part of $\mu_\varepsilon \circ f \circ \nu_e^{-1}$ vanishes, and secondly that its real part belongs to $[0, 1]$. Both are polynomial conditions imposed on real-valued polynomials, and are efficiently computable numerically. We also keep track of the point of intersection \tilde{w}_e of $\tilde{\varepsilon}$ and e .

In that case, we move to the neighbouring triangle T' of T along edge e , and continue. When we do not detect more intersections with edges of \mathcal{C} , we know in which triangle of \mathcal{C} the vertex w''_j lies.

It may happen that two or more vertices w''_j belong to the same triangle T that we have found in the previous paragraph. In that case, we let \tilde{w} denote the last point on $\tilde{\varepsilon}$ that was computed — possibly w'_i ; it also belongs to T . We consider in turn all candidates w''_j , and compute the straight path δ_j from \tilde{w} to w''_j and its image $f(\delta_j)$. If there exists a unique j such that $f(\delta_j)$ lies in the two triangles of \mathcal{Q}

to which ε belongs, then we have found the desired j . If there are no such j , then we interpolate. Let $t_0 \in [0, 1]$ be such that $f(\tilde{w}) = \varepsilon(t_0)$. Consider $t \in (t_0, 1)$, and those lifts $w \in f^{-1}(\varepsilon(t))$ that belong to T ; we then consider the paths δ_j from w to w'_j as before, and continue with increasing t .

Additional care must be taken for tangent crossings of edges (when the imaginary part of $\mu_\varepsilon \circ f \circ \mu_\varepsilon^{-1}$ has a multiple zero), and when vertices of \mathcal{C} lie on edges of \mathcal{Q}^\perp or conversely; these are treated as special cases. However, because of the necessary crudeness of norm estimates on $|f'(z)|$ in the first method, this second method is preferable.

8. AN APPLICATION TO DYNAMICAL SYSTEMS

The *post-critical set* of a branched self-covering $f : \mathbb{S} \rightarrow \mathbb{S}$ with branch value set Q_f is

$$P_f := \bigcup_{n \geq 0} f^{\circ n}(Q_f).$$

We are interested in the case where P_f is finite and we consider f up to isotopy rel P_f ; namely, $f \sim g$ if there exists a path of branched self-coverings from f to g whose post-critical set moves smoothly. We say that f is *equivalent* to a rational map F if there are homeomorphisms $\phi : \mathbb{S} \rightarrow \mathbb{P}^1(\mathbb{C})$ and $\psi : \mathbb{S} \rightarrow \mathbb{P}^1(\mathbb{C})$ such that $F \circ \phi = \psi \circ f$ and ϕ is isotopic to ψ rel P_f .

On the one hand, many examples of branched self-coverings can be constructed combinatorially, via triangulations; for these, it is natural to consider the maps up to isotopy. On the other hand, a fundamental theorem by Thurston claims


Theorem 8.1 (Thurston; see [6]). *Let $f : \mathbb{S} \rightarrow \mathbb{S}$ be branched self-covering with $\#P_f \geq 3$. Then f is equivalent to a rational map if and only if f admits no “Thurston obstruction”, namely, if and only if, for every collection \mathcal{C} of isotopy classes of non-peripheral curves on $\mathbb{S} \setminus P_f$, the $\mathbb{Q}\mathcal{C}$ -endomorphism*

$$\mathcal{C} \ni c \mapsto \sum_{d \in f^{-1}(c) \cap \mathcal{C}} \frac{1}{\deg(f : d \rightarrow c)} \cdot d \in \mathbb{Q}\mathcal{C}$$

has spectral radius < 1 .

Furthermore, in that case, the rational map is unique up to conjugation by a Möbius transformation.

8.1. Cui’s Problem. Cui Guizhen suggested in 2010 that if f is a “Sierpiński map”, namely a rational map whose Julia set is a Sierpiński carpet, then there should exist an essential, non-peripheral, simple curve γ such that $f^{-n}(\gamma)$ contains at least two components homotopic to γ rel P_f , for some n large enough. He then found a counterexample to his suggestion, given combinatorially as follows:



In that case, $P_f = Q := \{Q_1, Q_2, Q_3\}$ with $Q_1 := \infty$, $Q_2 := 0$ and $Q_3 := 1$, and f fixes P_f pointwise. Since f has only three post-critical points, all curves are peripheral. On the other hand, the Julia set of f is a Sierpiński carpet. First, the Julia set of f is connected and locally connected. Imagine that there are two Fatou components of f whose closure intersect. Consider the curve between their centres and consisting of internal rays. Iterating forward sufficiently many times, one obtains a curve γ connecting two points (say Q_1, Q_2) of Q , such that $f(\gamma)$ is simple. This is impossible, because every simple curve connecting Q_1 to Q_2 and mapping bijectively under f must pass through a preimage of Q_3 .

From this picture, it is easy to compute the monodromy action about Q : the permutations are those given in (2), namely

$$\begin{aligned}\sigma_1 &= (1, 7, 11, 2)(3, 8)(4, 5)(6, 10)(9, 12, 13), \\ \sigma_2 &= (1, 3, 12, 4)(5, 9)(6, 7)(10, 13, 11)(2, 8), \\ \sigma_3 &= (1, 5, 13, 6)(7, 10)(2, 3)(8, 11, 12)(4, 9).\end{aligned}$$

Because $\#Q = 3$, all curves on $\mathbb{P}^1(\mathbb{C}) \setminus Q$ are peripheral. By Thurston's rigidity theorem, there is a unique map with monodromy $(\sigma_1, \sigma_2, \sigma_3)$ and fixing $\infty, 0, 1$ with local degree 2, so it is the map computed by our algorithm (perhaps after precomposition with a Möbius transformation so that $\infty, 0, 1$ are critical of order 2).

Our algorithm searched in fact for a map with $\infty, 0, 1$ of order 4. This is an improvement to searching immediately for the correct map, because there are three points of order 2 above each of $\infty, 0, 1$, and they may lie in a strict field extension. Needless to say, it suffices, once such a map is found, to precompose it with a finite number of possible Möbius transformations.

Our algorithm found a solution (mod 11) of the defining equations for a map; then lifted them (mod 11^{2^6}) and finally obtained six Galois conjugate solutions. The correct one (with correct choice of point of order 2 above $\infty, 0, 1$) was then found. The original map is of the form

$$f(z/w) = \lambda \frac{(z - b_4w)^4(z - b_3w)^3(z - b_{2,1}w)^2(z - b_{2,2}w)^2(z - b_{2,3}w)^2}{(z - a_4w)^4(z - a_3w)^3(z - a_{2,1}w)^2(z - a_{2,2}w)^2(z - a_{2,3}w)^2};$$

here are the preimages a_i, b_i, c_i of $\infty, 0, 1$ respectively:

$$\begin{aligned}
 a_4 &= \infty \text{ (meaning the term } z - a_4w \text{ should be replaced by 1),} \\
 a_3 &\approx 0.50000000000000000000000000000000 - 0.439846359796987134487167714627i, \\
 a_{2,1} &\approx 1.61268567872451072013417667720 - 0.490182463946729812334860743821i, \\
 a_{2,2} &\approx 0.50000000000000000000000000000000 - 0.0415300696430258467988035191529i, \\
 a_{2,3} &\approx -0.612685678724510720134176677204 - 0.490182463946729812334860743821i, \\
 b_4 &= 0, \\
 b_3 &\approx 1.12748515145901194873474709466 - 0.991840479188802206853242764751i, \\
 b_{2,1} &\approx 1.98629656633071582984701575517 - 0.164982069462835473582606346591i, \\
 b_{2,2} &\approx 0.567640411622375679553529964298 - 0.172536644477962176299255320022i, \\
 b_{2,3} &\approx -0.995164705141609432502666361446 - 0.796186860797306011242450678339i, \\
 c_4 &= 1, \\
 c_3 &\approx -0.127485151459011948734747094655 - 0.991840479188802206853242764751i, \\
 c_{2,1} &\approx 0.432359588377624320446470035702 - 0.172536644477962176299255320022i, \\
 c_{2,2} &\approx -0.986296566330715829847015755165 - 0.164982069462835473582606346591i, \\
 c_{2,3} &\approx 1.99516470514160943250266636145 - 0.796186860797306011242450678339i, \\
 \lambda &\approx 0.130027094895701439414281708196i.
 \end{aligned}$$

The required Möbius transformation μ maps $(a_{2,2}, b_{2,3}, c_{2,2})$ to $(\infty, 0, 1)$ respectively, so

$$\mu(z) = \frac{z - b_{2,3}}{z - a_{2,2}} \cdot \frac{c_{2,2} - a_{2,2}}{c_{2,2} - b_{2,3}}$$

and the required solution is $f \circ \mu$. Its Julia set is displayed in Figure 1.




FIGURE 1. Julia set of Cui's map

Acknowledgments. We are grateful to Kevin Pilgrim for valuable remarks on a preliminary version of the text, and for more examples of Sierpiński maps with three post-critical points.

REFERENCES

- [1] A. Oliver L. Atkin and H. Peter F. Swinnerton-Dyer, *Modular forms on noncongruence subgroups*, Combinatorics (Proc. Sympos. Pure Math., Vol. XIX, Univ. California, Los Angeles, Calif., 1968), Amer. Math. Soc., Providence, R.I., 1971, pp. 1–25. MR0337781 (49 #2550)
- [2] *PARI/GP, version 2.5.0*, The PARI Group, Bordeaux, 2011.
- [3] Philip L. Bowers and Kenneth Stephenson, *Uniformizing dessins and Belyĭ maps via circle packing*, Mem. Amer. Math. Soc. **170** (2004), no. 805, xii+97. MR2053391 (2005a:30068)
- [4] Jean-Marc Couveignes and Louis Granboulan, *Dessins from a geometric point of view*, The Grothendieck theory of dessins d’enfants (Luminy, 1993), London Math. Soc. Lecture Note Ser., vol. 200, Cambridge Univ. Press, Cambridge, 1994, pp. 79–113. MR1305394 (96b:14015)
- [5] Jean-Marc Couveignes, *Quelques revêtements définis sur \mathbb{Q}* , Manuscripta Math. **92** (1997), no. 4, 409–445, DOI 10.1007/BF02678203 (French, with French summary). MR1441485 (98c:14021)
- [6] Adrien Douady and John H. Hubbard, *A proof of Thurston’s topological characterization of rational functions*, Acta Math. **171** (1993), no. 2, 263–297. MR1251582 (94j:58143)
- [7] The GAP Group, *GAP — Groups, Algorithms, and Programming, Version 4.4.10*, 2008.
- [8] Alexandre Grothendieck, *Esquisse d’un programme*, Geometric Galois actions, 1, London Math. Soc. Lecture Note Ser., vol. 242, Cambridge Univ. Press, Cambridge, 1997, pp. 5–48 (French, with French summary). With an English translation on pp. 243–283. MR1483107 (99c:14034)
- [9] Adolf Hurwitz, *Ueber Riemann’sche Flächen mit gegebenen Verzweigungspunkten*, Math. Ann. **39** (1891), no. 1, 1–60, DOI 10.1007/BF01199469 (German). MR1510692
- [10] Arjen K. Lenstra, Hendrik W. Lenstra Jr., and László Lovász, *Factoring polynomials with rational coefficients*, Math. Ann. **261** (1982), no. 4, 515–534, DOI 10.1007/BF01457454. MR682664 (84a:12002)
- [11] Gunter Malle, *Polynomials for primitive nonsolvable permutation groups of degree $d \leq 15$* , J. Symbolic Comput. **4** (1987), no. 1, 83–92, DOI 10.1016/S0747-7171(87)80056-1. MR908415 (89b:12007)
- [12] Gunter Malle and Bernd H. Matzat, *Realisierung von Gruppen $\mathrm{PSL}_2(\mathbb{F}_p)$ als Galoisgruppen über \mathbb{Q}* , Math. Ann. **272** (1985), no. 4, 549–565, DOI 10.1007/BF01455866 (German). MR807290 (87e:12002)
- [13] Kevin M. Pilgrim, *Dessins d’enfants and Hubbard trees*, Ann. Sci. École Norm. Sup. (4) **33** (2000), no. 5, 671–693, DOI 10.1016/S0012-9593(00)01050-8 (English, with English and French summaries). MR1834499 (2002m:37062)
- [14] Robert J. Renka, *Algorithm 772: STRIPACK: Delaunay triangulation and Voronoi diagram on the surface of a sphere*, ACM Trans. Math. Software **23** (1997), no. 3, 416–434, DOI 10.1145/275323.275329. MR1672176
- [15] Jonathan R. Shewchuk, *Delaunay refinement algorithms for triangular mesh generation*, Comput. Geom. **22** (2002), no. 1-3, 21–74, DOI 10.1016/S0925-7721(01)00047-5. 16th ACM Symposium on Computational Geometry (Hong Kong, 2000). MR1893652 (2003b:65131)

L.B.: MATHEMATISCHES INSTITUT, GEORG-AUGUST UNIVERSITÄT ZU GÖTTINGEN

X.B.: INSTITUT DE MATHÉMATIQUES DE TOULOUSE, UNIVERSITÉ PAUL SABATIER, TOULOUSE

H.-C.G.v.B., J.K.: COURANT RESEARCH CENTRE “HIGHER ORDER STRUCTURES”, GEORG-AUGUST UNIVERSITÄT ZU GÖTTINGEN