Another proof of convergence of Thurston's iteration on Teichmüller space for postcritically finite rational maps

Jeremy Kahn and Eduardo Sodré

December 19, 2024

Let $f\colon S^2 \to S^2$ be a postcritically finite branched cover of the 2-sphere, C_f be its critical points, and $P_f = \bigcup_{n\geq 1}^{\infty} f^n(C_f)$ be its postcritical set. We want to prove that if f has no *obstruction* and *hyperbolic orbifold*, then it is *Thurston equivalent* to a rational map: there exists homeomorphisms $h, \widetilde{h} : S^2 \to \hat{\mathbb{C}}$ isotopic rel P_f such that the diagram below commutes. Essentially, f has the same combinatorics as that of some rational map.

$$\begin{array}{ccc} (S^2, P_f) & \xrightarrow{\widetilde{h}} & \hat{\mathbb{C}} \\ \Big\downarrow f & & \Big\downarrow g \\ (S^2, P_f) & \xrightarrow{h} & \hat{\mathbb{C}} \end{array}$$

An obstruction for f is a collection Γ of simple, closed, disjoint and non-trivial (that is, not null-homotopic and non-peripheral) curves in $S^2 \setminus P_f$, with the following properties. It is f*-stable*, meaning that for $\gamma \in \Gamma$, all preimages of γ under f are either trivial or homotopic to some curve in Γ. Being f-stable, we may define the *Thurston linear transformation* $f_\Gamma : \mathbb{R}^\Gamma \to \mathbb{R}^\Gamma$ given by

$$f_\Gamma(\gamma_j) = \sum_i \sum_\alpha \frac{1}{d_{j,i,\alpha}} \gamma_i,$$

where α ranges over all components of $f^{-1}(\gamma_j)$ which are homotopic to γ_i, and $d_{j,i,\alpha}$ is the degree of the map $f|_{\gamma_\alpha} : \gamma_\alpha \to \gamma_j$. By Perron-Frobenius, f_Γ

has a largest non-negative real eigenvalue $\lambda(f, \Gamma)$. Then Γ is an obstruction if $\lambda(f, \Gamma) \geq 1$.

The map f induces an iteration $\sigma_f\colon \operatorname{Teich}(S^2, P_f) \to \operatorname{Teich}(S^2, P_f)$ given by pulling back complex structures by f. If we view elements of $\operatorname{Teich}(S^2, P_f)$ as diffeomorphisms $\varphi : (S^2, P_f) \to \hat{\mathbb{C}}$, where $\varphi_1 \sim \varphi_2$ if $\varphi_2 \circ \varphi_1^{-1}$ is isotopic to a conformal map rel $\varphi_1(P_f)$, we have a projection $\pi\colon \operatorname{Teich}(S^2, P_f) \to M(S^2, P_f)$ to the Moduli space of (S^2, P_f), viewed as the injections $\iota : P_f \to \hat{\mathbb{C}}$ equivalent under post-composition with Möbius transformations. The projection is then just the restriction of φ to P_f. In our approach, we avoid the use of deformation theory on $\operatorname{Teich}(S^2, P_f)$ as in [1] by considering an intermediary quotient space between $\operatorname{Teich}(S^2, P_f)$ and $M(S^2, P_f)$, which we call the Rees space, on which the iteration σ_f is well-defined and behaves nicely. The most technical tool utilized is then the existence and uniqueness theorem of Teichmüller maps.

Given $\tau_1, \tau_2 \in \operatorname{Teich}(S^2, P_f)$, we recall the definition of the Teichmüller distance $d(\tau_1, \tau_2)$. If μ_1, μ_2 are representative conformal structures, it is equal to $\frac{1}{2} \inf_\psi \log \operatorname{Dil}(\psi)$, where the infimum is over all quasiconformal maps $\psi : (S^2, \mu_1) \to (S^2, \mu_2)$ isotopic to the identity rel P_f, and $\operatorname{Dil}(\psi)$ is its (maximal) dilatation. Due to Teichmüller's existence and uniqueness theorem, this infimum is realized uniquely by a Teichmüller map q, which stretches the complex structures along some quadratic differentials in (S^2, μ_1) and (S^2, μ_2). For any $\tau \in \operatorname{Teich}(S^2, P_f)$, we let $d(\tau) = d_{\text{Teich}}(\tau, \sigma_f \tau)$. By lifting q to a quasiconformal map $\widetilde{q}$ with same dilatation via the diagram

$$\begin{array}{ccc} (S^2, f^* f^* \mu) & \xrightarrow{\widetilde{q}} & (S^2, f^* \mu) \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle f} \\ (S^2, f^* \mu) & \xrightarrow[q]{} & (S^2, \mu), \end{array}$$

we observe that

$$d(\sigma_f \tau) \leq d(\tau), \tag{1}$$

and more generally, distances are non-increasing with respect to σ_f:

$$d(\sigma_f \tau_1, \sigma_f \tau_2) \leq d(\tau_1, \tau_2).$$

Given our hypothesis of f, we have reduced the problem to proving that σ_f has a fixed point.

We now describe the *Rees space* $R(F)$ of a portrait, and the associated Rees space $R(f)$ of f. Given abstract not necessarily disjoint finite sets C and P, a *portrait* F is a map $F : C \cup P \to P$ and a function $\deg : C \to \{2, 3, \dots\}$, meant to record the "0-dimensional data" of a postcritically finite branched cover. The Rees space $R(F)$ consists of tuples (e, G, μ), where μ is a complex structure on S^2, $e : C \cup P \to (S^2, P_f)$ is an injection, and $G : (S^2, \mu) \to (S^2, \mu)$ is a topological branched cover with $C_g = e(C)$, $P_g = e(P)$ such that $G \circ e = e \circ F$ and $\deg G|_{e(c)} = \deg(c)$ for all $c \in C$. Two tuples (e_i, G_i, μ_i) for $i = 1, 2$ are equivalent if the corresponding self-maps of S^2 are Thurston equivalent by a conformal h, that is, so that $h^*\mu_2 = \mu_1$ and $h \circ e_1|_P = e_2|_P$:

$$\begin{array}{ccc}(S^2, \mu_1) & \xrightarrow{\widetilde{h}} & (S^2, \mu_2) \\ {\scriptstyle G_1}\big\downarrow & & \big\downarrow{\scriptstyle G_2} \\ (S^2, \mu_1) & \xrightarrow{h} & (S^2, \mu_2).\end{array}$$

Moreover, it is possible to normalize the complex structure by taking an equivalent tuple to assume that the μ is the standard complex structure on $\hat{\mathbb{C}}$, so that $G : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ is a topological branched cover.

An equivalent description of the Rees space which will be important for us is as tuples (e, e', g, q), where $e, e' : C \cup P \to \hat{\mathbb{C}}$ are injections, $g : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ is holomorphic and $q : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ is a homeomorphism such that:

(i) $g \circ e' = e \circ F$ on $C \cup P$;

(ii) $\deg g_{e'(c)} = \deg(c)$, for all $c \in C$;

(iii) $q \circ e' = e$ on P.

$$\begin{array}{ccc}C \cup P & \xrightarrow{e'} & \hat{\mathbb{C}} \\ {\scriptstyle F}\big\downarrow & & \big\downarrow{\scriptstyle g} \\ P & \xrightarrow{e|_P} & \hat{\mathbb{C}}\end{array} \qquad \begin{array}{ccc} & & \hat{\mathbb{C}} \\ & {\scriptstyle e'|_P}\nearrow & \big\downarrow{\scriptstyle q} \\ P & \xrightarrow{e|_P} & \hat{\mathbb{C}}\end{array}$$

Two tuples are equivalent if there are conformal isomorphisms $h, h' : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$, that is, Möbius transformations, such the first diagram below commutes and the second commutes up to isotopy rel $e_1'(P)$ with respect to the q_i. We recover the first description by considering the branched cover $g \circ q^{-1}$ on

$\hat{\mathbb{C}} = (S^2, \hat{\mu})$, where $\hat{\mu}$ is the standard complex structure on $\hat{\mathbb{C}}$.

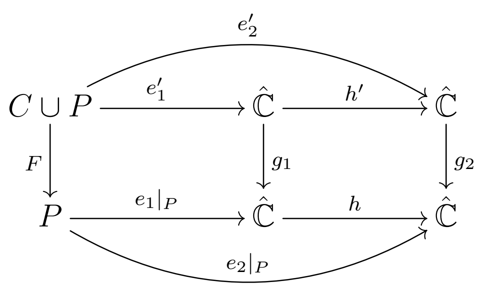

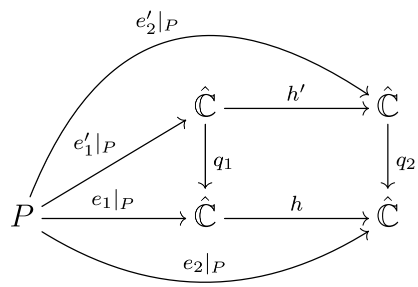

There are natural maps $\mathrm{Teich}(S^2, P_f) \to R(F) \to M(S^2, P_f)$, where the image of $\mathrm{Teich}(S^2, P_f)$ in $R(F)$ is denoted by $R(f)$ and corresponds to the data of all maps Thurston equivalent to f. Hence elements of $R(f)$ may be viewed as conformal structures μ where $\mu_1 \sim \mu_2$ if there exists a homeomorphism h such that $h^*\mu_2 = \mu_1$ and h is a self-Thurston equivalence of f (note that any map isotopic to a Thurston equivalence is itself a Thurston equivalence). Equivalently, we may define a subgroup of $\mathrm{MCG}(S^2, P_f)$ called the *special liftable* mapping classes ($\mathrm{LS}(f)$), namely those h which form a self-equivalence of f, so $h \in \mathrm{LS}(f)$ if there exists $h \sim \tilde{h}$ (rel P_f) such that $h \circ f = f \circ \tilde{h}$. Then $R(f) = \mathrm{Teich}(S^2, P_f)/\,\mathrm{LS}(f)$, and we imbue the Rees space with the quotient topology.

The iteration σ_f descends continuously to the quotient, inducing an iteration $\sigma_f : R(f) \to R(f)$. Moreover, since the action of $\mathrm{MCG}(S^2, P_f)$ on $\mathrm{Teich}(S^2, P_f)$ is isometric, the map $d : \mathrm{Teich}(S^2, P_f) \to [0, +\infty)$ also descends continuously to $R(f)$: if $h \in \mathrm{LS}(f)$,

$$d(\mu, f^*\mu) = d(h^*\mu, h^*f^*\mu) = d(h^*\mu, \widetilde{h}^*f^*\mu) = d(h^*\mu, f^*h^*\mu),$$

since $h \sim \widetilde{h}$ rel P_f. In fact, these properties motivate the definition of $\mathrm{LS}(f)$ and $R(f)$ in the first place, since for those mapping classes $h \in \mathrm{LS}(f)$ we have $f^*h^*\mu = h^*f^*\mu$ for all μ, where the equality is to be understood as for elements of $\mathrm{Teich}(S^2, P_f)$.

Note that $d(\alpha) = \frac{1}{2}\log \mathrm{Dil}(q)$ is not necessarily the distance between α and $\sigma_f(\alpha)$ in $R(f)$, but if for some $\alpha \in R(f)$ we have $d(\alpha) = 0$, then f is Thurston equivalent to a rational map. With this in mind, we search for global minima of d on an adequately chosen compact subset of $R(f)$.

Lemma 1. *The map* $(\pi_M, d) : R(f) \to M(S^2, P_f) \times [0, +\infty)$ *is proper.*

Proof. We may assume $|P_f| \geq 3$; let $K \subset M(S^2, P_f)$ be compact and $d \geq 0$. We want to show that $R_{K,d} \coloneqq \{\alpha \in R(f) \mid \pi_M(\alpha) \in K,\ d(\alpha) \leq d\}$ is compact. Let $(\alpha_i)_{i\geq 0}$ be a sequence in $R(f)$, where $\alpha_i = [e_i, e_i', g_i, q_i]$ and the q_i are Teichmüller maps. Given Möbius transformations M_i and N_i, we have

$$\alpha = [e_i, e_i', g_i, q_i] = [M_i e_i, N_i e_i', M_i g_i N_i^{-1}, M_i q_i N_i^{-1}],$$

so that we may assume there exists $p_1, p_2, p_3 \in P$ which are sent to $0, 1, \infty$ by all the e_i by choosing appropriate M_i, and that these points are fixed by q_i by choosing appropriate N_i. Under this normalization, and due to compactness in the moduli space, we may pass to a subsequence and assume that $e_i|_P$ converges to an injection $\iota : P \to \hat{\mathbb{C}}$. Moreover, since the q_i are normalized e^{2d}-quasiconformal maps, we may also take a subsequence and assume that $q_i \to q$ uniformly, so that $e_i'|_P \to q^{-1}\iota$.

We consider diffeomorphisms $h_i : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ such that $h_i \circ e_i|_P = \iota$, "pushing" the points in $e_i(P)$ towards their limit $\iota(P)$, where the h_i converge uniformly to the identity on $\hat{\mathbb{C}}$. In fact, we may assume that the h_i are quasiconformal and, for sufficiently large i, they are the identity outside of some small neighborhood of $\iota(P)$, shrinking with respect to i, and $\mathrm{Dil}(h_i) \to 1$. Let $V = F(C) \subseteq P$, representing the abstract critical values, and $V_i = e_i(V) = g_i(e_i'(C))$ the critical values of g_i. The composition $h_i g_i : \hat{\mathbb{C}} \setminus g^{-1}(V_i) \to \hat{\mathbb{C}} \setminus \iota(V)$ is a covering map of degree d, inducing a subgroup $H_i < \pi_1(\hat{\mathbb{C}} \setminus \iota(V))$ of index d. There are only finitely many subgroups of index d, so by passing to a subsequence we may assume all of the H_i are equal to some fixed $H_\infty < \pi_1(\hat{\mathbb{C}} \setminus \iota(V))$. There exists an associated covering space $g : E \to \hat{\mathbb{C}} \setminus \iota(V)$ and diffeomorphisms $k_i : \hat{\mathbb{C}} \setminus g_i^{-1}(V_i) \to E$ such that the diagram below commutes:

$$\begin{array}{ccc} \hat{\mathbb{C}} \setminus g_i^{-1}(V_i) & \xrightarrow{k_i} & E \\ \downarrow{\scriptstyle g_i} & & \downarrow{\scriptstyle g} \\ \hat{\mathbb{C}} \setminus V_i & \xrightarrow[h_i]{} & \hat{\mathbb{C}} \setminus \iota(V) \end{array}$$

Topologically, E is a 2-sphere with some finite number of punctures, and if we pullback the conformal structure of $\hat{\mathbb{C}} \setminus \iota(V)$ by g we may assume that $E = \hat{\mathbb{C}} \setminus L$. Consequently, the k_i are quasiconformal diffeomorphisms such that $\mathrm{Dil}(k_i) \leq \mathrm{Dil}(h_i) \to 1$, and each extends to $g_i^{-1}(V_i)$, mapping it bijectively onto L.

$$\begin{array}{ccccc} C \cup P & \xrightarrow{e_i'} & \hat{\mathbb{C}} & \xrightarrow{k_i} & \hat{\mathbb{C}} \\ \downarrow{\scriptstyle F} & & \downarrow{\scriptstyle g_i} & & \downarrow{\scriptstyle g} \\ P & \xrightarrow{e_i|_P} & \hat{\mathbb{C}} & \xrightarrow{h_i} & \hat{\mathbb{C}} \end{array}$$

The maps k_i send $0, 1, \infty$ into points in $g^{-1}(\iota(F(\{p_1, p_2, p_3\}))$, so by passing to a subsequence we may assume that the k_i are normalized to send $0, 1, \infty$ into three specific points independently of i. Moreover, by taking an isomorphic covering, we may assume that these three points are themselves $0, 1, \infty$. With this normalization, k_i converges uniformly to the identity, and by continuity of composition under uniform convergence,

$$g_i = h_i^{-1}(h_i g_i) = h_i(g k_i) \to g.$$

For $p \in P$, we see that $g k_i e_i'(p) = \iota(p)$, so that by passing to a subsequence $k_i(e_i(p))$ is a fixed element $\widetilde{p}$ in $g^{-1}(\iota(p))$. Then

$$e_i'(p) = k_i^{-1}(\widetilde{p}) \to \widetilde{p},$$

but also $e_i'(p) \to q^{-1}\iota(p)$, so that $\widetilde{p} = q^{-1}\iota(p)$. In particular

$$\widetilde{p} = q^{-1}\iota(p) \implies g(\widetilde{p}) = \iota(p) = g q^{-1}(\iota(p)).$$

This shows that g is the desired rational map and q the quasiconformal map whose composition forms a postcritically finite branched cover. We also get the Thurston equivalences

$$\begin{array}{ccc} \hat{\mathbb{C}} & \xrightarrow{q k_i q_i^{-1}} & \hat{\mathbb{C}} \\ {\scriptstyle g_i q_i^{-1}}\downarrow & & \downarrow{\scriptstyle g q^{-1}} \\ \hat{\mathbb{C}} & \xrightarrow{h_i} & \hat{\mathbb{C}} \end{array}$$

where $h_i \sim q k_i q_i^{-1}$ rel $e_i(P)$ since both are sufficiently close to the identity, and $\log \mathrm{Dil}(h_i) \to 0$, showing convergence of the $(\alpha_i)_i$. □

More is true; if we imbue every "component" of $R(F)$, corresponding to post-critically finite branched covers in distinct Thurston equivalence classes, with the quotient topologies from the corresponding Teichmüller spaces, we get that (π_M, d) is proper in the entire domain $R(F)$.

We also have the following result, effectively proven in [1]:

Lemma 2. *Suppose there exists τ for which $d(\tau) \leq D$. Then there is an integer $m \geq 1$ and a non-empty compact set $K \subset M(S^2, P_f)$ such that, if $\pi(\tau) \in K$ and $d(\tau) \leq D$, then $\pi(\sigma_f^m \tau) \in K$.*

Proof. For $\tau \in \text{Teich}(S^2, P_f)$, define

$$\omega(\tau) \coloneqq \sup_{\gamma}\{-\log l_\tau(\gamma)\},$$

where the supremum ranges over all non-trivial simple closed curves γ in $S^2 \setminus P_f$ and $l_\tau(\gamma)$ is the length of the unique geodesic of τ homotopic to γ. Proposition 7.3 in [1] states that ω is 2-Lipschitz and that, for all M, $\{\tau \in \text{Teich}(S^2, P_f) \colon \omega(\tau) \leq M\}$ is the preimage of a compact set K in $M(S^2, P_f)$, so that a sequence in $\text{Teich}(S^2, P_f)$ can only go to infinity in $M(S^2, P_f)$ if the length of some geodesic is going to 0.

Given that f is unobstructed, Proposition 8.2 in [1] guarantees that there exists an integer $m \geq 1$ and $C > 0$ such that, if $\omega(\tau) > C$ and $d(\tau) \leq D$, then $\omega(\sigma_f^m \tau) < \omega(\tau)$. Suppose then that $d(\tau) \leq D$ and $\omega(\tau) \leq C + 2mD$. We divide into two cases; if $\omega(\tau) \leq C$, then because ω is 2-Lipschitz we get $\omega(\sigma_f^m \tau) \leq C + 2mD$. If $\omega(\tau) > C$, then by Proposition 8.2 we get that $\omega(\sigma_f^m \tau) < \omega(\tau) \leq C + 2mD$. This concludes that $K = \{\tau \colon \omega(\tau) \leq C + 2mD\}$ is our desired compact set. □

Letting

$$E_{K,D} = \{\tau \in \text{Teich}(S^2, P_f) \mid \pi(\tau) \in K \text{ and } d(f) \leq D\} = \pi_R^{-1}(R_{K,D}),$$

we see that $E_{K,D}$ is σ_f^m-invariant, for K and D given in Lemma 2. Consequently $\sigma_f^m(R_{K,D}) \subseteq R_{K,D}$, and it is compact. We then require only one more lemma:

Theorem 3. *For all τ there exists n such that $d(\sigma_f^n \tau) < d(\tau)$.*

Assuming Theorem 3, we can then prove that σ_f has a fixed point. As $R_{K,D}$ is non-empty and compact, we may find $\alpha \in R_{K,d_1}$ for which $d(\alpha)$ is minimal. If $d(\alpha) = 0$, we're done. Otherwise, $\sigma_f^{mn}(\alpha) \in R_{K,D}$ and $d(\sigma_f^{mn}(\alpha)) \leq d(\sigma_f^n(\alpha)) < d(\alpha)$, a contradiction.

Let's now prove Theorem 3. We need the following combinatorial lemma:

Lemma 4. *If f has a hyperbolic orbifold and ψ is a quadratic differential for (S^2, P_f), then there exists n such that $(f^n)^*\psi$ has a pole outside P_f.*

Proof. We only need combinatorial properties of quadratic differentials for this result. Recall that, counting multiplicity, $Z - P = -4$, where Z is the number of zeros of ψ and P is the number of poles. We assume that the set of poles of ψ is contained in P_f. Moreover, if w is a zero of order $m \in \{-1, 0, \ldots\}$ for ψ, and z is such that it maps to w with local degree k, then z is a zero of order $k(m+2)-2$ for $f^*\psi$. This is because a zero of order m corresponds to a $(m+2)$-pronged singularity, which is pulled back to a $k(m+2)$-pronged singularity by f.

Since poles are always simple for integrable holomorphic quadratic differentials, it is sufficient to show that the number of zeros eventually increases under iterated pullbacks. If w is a zero of order $m \geq 1$ for ψ, then each $z \in f^{-1}(w)$ is a zero of order $k_z(m+2) - 2 \geq 1$. Hence $f^*\psi$ doesn't have more zeros than ψ only when ψ itself has no zeros, and only 4 simple poles p_1, p_2, p_3, p_4 contained in P_f.

If $f^*\psi$ has no zeros, then all points mapping to some p_i must map with local degree $k = 1$ or 2, and if w is not a pole and $z \in f^{-1}(w)$, then z must map to w with local degree 1. This in particular implies that P_f is exactly the set of 4 poles. If some pole is also a critical point, it would be a regular point for $f^*\psi$. But as $f^*\psi$ has at least 4 poles, there would some pole for $f^*\psi$ outside of P_f. Therefore $C_f \cap P_f = \emptyset$, and the orbifold is euclidean of type $(2, 2, 2, 2)$, contradicting our assumptions. □

In fact, Lemma 2 in [1] sharpens this result to show that we can always take $n = 2$.

Now, given τ, we can form the Teichmuller map from τ to $\sigma_f\tau$, where we stretch along some quadratic differential ψ. This Teichmuller map lifts to a quasiconformal map from $\sigma_f\tau$ to $\sigma_f^2\tau$:

$$\begin{array}{ccc} (S^2, f^*f^*\mu) & \xrightarrow{\widetilde{q}} & (S^2, f^*\mu) \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle f} \\ (S^2, f^*\mu) & \xrightarrow[q]{} & (S^2, \mu). \end{array}$$

We may take these iterated lifts, and as soon as $f^{n*}(\psi)$ has a pole outside of P_f, the associated map from $\sigma_f^n\tau$ to $\sigma_f^{n+1}\tau$ will not be a Teichmüller map. This then implies that $d(\sigma_f^n\tau) < d(\tau)$, and concludes the final argument needed for the result.

References

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