Annals of Mathematics **192** (2020), 1–46
https://doi.org/10.4007/annals.2020.192.1.1

A positive characterization of rational maps

By Dylan P. Thurston

This paper is dedicated to William Thurston, 1946–2012, *and Tan Lei*, 1963–2016.
I will miss the joy they brought to the subject.

Abstract

When is a topological branched self-cover of the sphere equivalent to a post-critically finite rational map on $\mathbb{CP}^1$? William Thurston gave one answer in 1982, giving a negative criterion (an obstruction to a map being rational). We give a complementary, positive criterion: the branched self-cover is equivalent to a rational map if and only if there is an elastic graph spine for the complement of the post-critical set that gets "looser" under backwards iteration.

Contents

1. Introduction

In this paper, we complete the program laid out in earlier work [Thu16], and give a positive characterization of post-critically finite rational maps among branched self-covers of the sphere.

Definition 1.1. A *topological branched self-cover* of the sphere is a finite set of points $P \subset S^2$ and a map $f\colon (S^2, P) \to (S^2, P)$, also written $f\colon (S^2, P) \circlearrowleft$, so that f is an orientation-preserving covering map (with degree greater than 1)

Keywords: Complex dynamics, rational maps, elastic graphs, quasiconformal surgery, extremal length, Thurston obstruction

AMS Classification: Primary: 37F10; Secondary: 37E25, 37F31.

when restricted to a map from $S^2 \backslash f^{-1}(P)$ to $S^2 \backslash P$. That is, f is a branched cover so that $f(P) \subset P$ and P contains the critical values. (As a result, P contains the post-critical set of f.) Two branched self-covers are *equivalent* if they are related by conjugacy of S^2 (possibly changing the set P) and homotopy relative to P.

One source of topological branched self-covers is post-critically finite rational maps. Let $\widehat{\mathbb{C}} = \mathbb{CP}^1$, and suppose $f(z) = P(z)/Q(z)$ is a rational map with a finite, forward-invariant set P that contains all critical values. Then, if we forget the conformal structure, $f\colon (\widehat{\mathbb{C}}, P) \circlearrowleft$ is a topological branched self-cover.

Question 1.2. When is a topological branched self-cover equivalent to a post-critically finite rational map?

One answer to Question 1.2 was given by W. Thurston 30 years ago [DH93], recalled as Theorem 7.4. He proved a negative characterization: there is a certain combinatorial object (an *annular obstruction*) that exists exactly when $f\colon (S^2, P) \circlearrowleft$ is *not* equivalent to a rational map. In this paper, we give a complementary, positive, characterization: a combinatorial object that exists exactly when f *is* equivalent to a rational map.

Before stating the main theorem, we give a combinatorial description of topological branched self-covers $f\colon (S^2, P) \circlearrowleft$ in terms of graph maps.[1] Pick a graph spine Γ_0 for $S^2 \backslash P$ (a deformation retract of $S^2 \backslash P$, i.e., an embedded graph so each complementary region is a punctured disk) and consider its inverse image $\Gamma_1 = f^{-1}(\Gamma_0) \subset S^2 \backslash f^{-1}(P)$. There are two natural homotopy classes of maps from Γ_1 to Γ_0.

- A covering map π_Γ commuting with the action of f:

$$\begin{tikzcd} \Gamma_1 \arrow[r, "\pi_\Gamma"] \arrow[d, hook] & \Gamma_0 \arrow[d, hook] \\ S^2 \backslash f^{-1}(P) \arrow[r, "f"] & S^2 \backslash P. \end{tikzcd}$$

- A map φ_Γ commuting up to homotopy with the inclusion of $S^2 \backslash f^{-1}(P)$ in $S^2 \backslash P$. The homotopy class $[\varphi_\Gamma]$ is unique, since Γ_0 is a deformation retract of $S^2 \backslash P$:

$$\begin{tikzcd} \Gamma_1 \arrow[r, "\varphi_\Gamma"] \arrow[d, hook] \arrow[dr, phantom, "\sim"] & \Gamma_0 \arrow[d, hook] \\ S^2 \backslash f^{-1}(P) \arrow[r, hook, "\text{incl.}"] & S^2 \backslash P. \end{tikzcd}$$

---

[1] In this paper, a graph map is a continuous map, not necessarily taking vertices to vertices.

This data $\pi_\Gamma, \varphi_\Gamma \colon \Gamma_1 \rightrightarrows \Gamma_0$ is a *virtual endomorphism* of Γ_0. It, together with a ribbon graph structure on Γ_0, determines f up to equivalence. (See Section 2 for examples. This is proved in Theorem 3 in Section 2. Example 2.13 shows that the extra ribbon structure is necessary to determine f.)

For our characterization of rational maps, we also need an *elastic structure* on $\Gamma = \Gamma_0$, by which we mean an elastic length $\alpha(e) \in \mathbb{R}_{>0}$ on each edge e of Γ. (We treat α as an ordinary length for purposes of differentiation, etc.) An *elastic graph* $G = (\Gamma, \alpha)$ is a graph Γ and an elastic structure α on Γ. For a piecewise-linear (PL) map $\psi \colon G_1 \to G_2$ between elastic graphs, the *embedding energy* is

$$\operatorname{Emb}(\psi) \coloneqq \operatorname*{ess\,sup}_{y \in G_2} \sum_{x \in \psi^{-1}(y)} |\psi'(x)|. \tag{1.3}$$

(We identify each edge e with an interval of length $\alpha(e)$ to compute derivatives.) The essential supremum ignores sets of measure zero, which for PL maps amounts to ignoring vertices of G_2 and images of vertices of G_1. On a homotopy class, $\operatorname{Emb}[\psi]$ is defined to be the infimum of $\operatorname{Emb}(\varphi)$ for $\varphi \in [\psi]$. $\operatorname{Emb}[\psi]$ is realized and controls whether G_1 is “looser” than G_2 as an elastic graph [Thu19, Th. 1].

In the context of a branched self-cover $f \colon (S^2, P) \circlearrowleft$, if $G = G_0$ is a spine for $S^2 \setminus P$ with an elastic structure, we get a virtual endomorphism $\pi_G, \varphi_G \colon G_1 \rightrightarrows G_0$, where G_1 inherits an elastic structure by pulling back via π_G. We can then consider $\operatorname{Emb}[\varphi_G]$, the embedding energy of $\varphi_G \colon G_1 \to G_0$, or the iterated version $\operatorname{Emb}[\varphi_G^n]$. (See Section 5 for iteration.)

In a mild generalization, we also consider disconnected surfaces.

Definition 1.4. A *branched self-cover* $f \colon (\Sigma, P) \circlearrowleft$ is a (possibly disconnected) compact closed surface Σ, a finite subset $P \subset \Sigma$, and a map $f \colon \Sigma \to \Sigma$ that

- is a branched covering map with branch values contained in P,
- has constant degree greater than 1,
- maps P to P, and
- is a bijection on components of Σ.

(The restriction to π_0-bijective maps avoids dynamically uninteresting cases.) A standard Euler characteristic argument shows that each component of Σ is either a sphere or a torus, and that in the torus case there is no branching.

Definition 1.5. For the purposes of this paper, a *Riemann surface* $S = (\Sigma, \omega)$ is a topological surface Σ, possibly disconnected or with boundary, and a conformal structure ω on Σ. A *rational map* is a closed Riemann surface S and a conformal, π_0-bijective map $f \colon S \circlearrowleft$.

Definition 1.6. A branched self-cover $f\colon (\Sigma, P) \circlearrowleft$ is of *non-compact type* if each component of Σ contains a point of P that eventually falls into a cycle with a branch point (under forward iteration of f). It is of *hyperbolic type* if each component of Σ contains a point of P and each cycle of P contains a branch point.

In either case, the branching is non-trivial, so Σ is a union of spheres. If f is a rational map, it is of non-compact type if and only if the Julia set does not contain any component of Σ and it is of hyperbolic type if and only if the dynamics on the Julia set is hyperbolic. Thus the term "non-compact type" refers to the orbifold of f, while the term "hyperbolic type" refers to dynamics of f and not to the orbifold.

THEOREM 1. *Let $f\colon (\Sigma, P) \circlearrowleft$ be a branched self-cover of hyperbolic type. Then the following conditions are equivalent:*

(1) *The branched self-cover f is equivalent to a rational map.*
(2) *There is an elastic graph spine G for $\Sigma \backslash P$ and an integer $n > 0$ so that* $\operatorname{Emb}[\varphi_G^n] < 1$.
(3) *For every elastic graph spine G for $\Sigma \backslash P$ and for every sufficiently large n (depending on f and G), we have* $\operatorname{Emb}[\varphi_G^n] < 1$.

Loosely speaking, Theorem 1 says that f is equivalent to a rational map if and only if elastic graph spines get looser under repeated backwards iteration. As compared to the earlier Theorem 7.4, Theorem 1 makes it easier to prove a map is rational: you can just exhibit an elastic graph spine G and a suitable map in the homotopy class of φ_G^n. (In practice, $n = 1$ often suffices; see Section 1.2.) We prove Theorem 1 as Theorem 1′ in Section 5, including some additional equivalent conditions.

Theorem 1 and the older Theorem 7.4 are complementary; neither one implies the other, and the proofs are largely independent. It is easy to see one implication: if there is an elastic graph spine that gets looser under backwards iteration in the sense of Theorem 1, then there is no annular obstruction in the sense of Theorem 7.4. See [Thu16, §8.4] for the argument, or Section 7 of this paper for generalizations.

There are several ingredients to prove Theorem 1, as outlined in Figure 1 and summarized below. Much of this has appeared in other papers; the main new contributions of this paper are in Ingredients 4 and 5, in Sections 4 and 5, respectively. In the proof as a whole, the hardest parts are Ingredient 1, which has been known for some time; Ingredient 2 relating conformal embeddings to extremal length, particularly the behavior under covers; and Ingredient 3, which involves a combinatorial understanding of the embedding energy in equation (1.3).

* * *

$$\begin{tikzcd}
f \text{ rational} \arrow[d, Leftrightarrow, "\text{Ingdt. 1: Th. 3.3}"'] & \\
\begin{array}{c}\exists S\colon f^{-1}(S) \text{ strictly}\\ \text{conformally embeds in } S\end{array} \arrow[d, Leftrightarrow, "\text{Ingdt. 2: Th. 1.9}"'] & \exists G,n\colon \operatorname{Emb}[\varphi_S^n] < 1 \arrow[d, Leftrightarrow, "\text{Ingdt. 3: Th. 1.14}"] \\
\exists S\colon \operatorname{SF}[\varphi_S] < 1 \arrow[r, Leftrightarrow, dashed, "\text{Ingdt. 4: Th. 2}"] \arrow[d, Leftrightarrow, "\text{Prop. 5.6}"'] \arrow[dr, phantom, "\text{Ingdt. 5: \S5}"] & \exists G,n\colon \operatorname{SF}[\varphi_G^n] < 1 \arrow[d, Leftrightarrow, "\text{Prop. 5.6}"] \\
\overline{\operatorname{SF}}[\pi_S,\varphi_S] < 1 \arrow[r, Leftrightarrow, "\text{Prop. 5.14}"'] & \overline{\operatorname{SF}}[\pi_G,\varphi_G] < 1
\end{tikzcd}$$

Figure 1. An outline of the equivalences used in proving Theorem 1, for a fixed branched self-cover $f\colon (\Sigma, P) \circlearrowleft$

The zeroth ingredient is the graphical description of topological branched self-covers in terms of spines, crucial to our entire approach. This is essentially a graphical version of Nekrashevych's automata for iterated monodromy groups [Nek05]. It is described in Section 2, culminating in Theorem 3, giving a graphical model for branched self-covers.

* * *

The first ingredient is a characterization of rational maps in terms of conformal embeddings of Riemann surfaces, a surface version of the graph criterion in Theorem 1. This has been folklore in the community for some time and is recalled as Theorem 3.3 in Section 3.

* * *

The second ingredient is a characterization of conformal embeddings of Riemann surfaces in terms of extremal length of multi-curves on the surface. This appeared in earlier work with Kahn and Pilgrim [KPT15], as we now summarize. Recall that the *extremal length* of a simple multi-curve c on a Riemann surface measures the maximum thickness of a collection of annuli around c.

Definition 1.7. For $f\colon R \hookrightarrow S$ a topological embedding of Riemann surfaces, the (extremal length) *stretch factor* of f is the maximal ratio of extremal lengths between the two surfaces:

$$\mathrm{SF}[f] \coloneqq \sup_{c\colon C\to R} \frac{\mathrm{EL}_S[f\circ c]}{\mathrm{EL}_R[c]},$$

where the supremum runs over all simple multi-curves c on R with $\mathrm{EL}_R[c] \neq 0$.

Definition 1.8. An *annular extension* of a Riemann surface R is any surface obtained by attaching a conformal annulus to each boundary component of R, and a conformal embedding $f\colon R \hookrightarrow S$ between Riemann surfaces is *annular* if it extends to a conformal embedding of an annular extension of R into S.

THEOREM 1.9 (Kahn-Pilgrim-Thurston [KPT15]). *Let R and S be Riemann surfaces, and let $f\colon R \hookrightarrow S$ be a topological embedding so that no component of $f(R)$ is contained in a disk or a once-punctured disk. Then f is homotopic to a conformal embedding if and only if* $\mathrm{SF}[f] \leqslant 1$. *Furthermore, f is homotopic to an annular conformal embedding if and only if* $\mathrm{SF}[f] < 1$.

We also use Theorem 5.11, a strengthening of Theorem 1.9 that behaves well under covers.

* * *

The third ingredient is a relation between the embedding energy of equation (1.3) to a stretch factor of maps between graphs (rather than surfaces) [Thu19].

Definition 1.10. Let $G = (\Gamma, \alpha)$ be an elastic graph. A *multi-curve* on G is a (not necessarily connected) 1-manifold C and a PL map $c\colon C \to G$. It is (strictly) *reduced* if c is locally injective. The *extremal length* of (C, c) is

$$\mathrm{EL}(c) \coloneqq \int_{y\in G} n_c(y)^2 \, d\alpha(y), \tag{1.11}$$

where $n_c(y)$ is the number of elements in $c^{-1}(y)$. (See [Thu16, §5.2] for motivation on why this is called extremal length.) If c is reduced, then $n_c(y)$ depends only on the edge containing y, and equation (1.11) reduces to

$$\mathrm{EL}(c) = \sum_{e\in\mathrm{Edge}(G)} \alpha(e) n_c(e)^2. \tag{1.12}$$

$\mathrm{EL}[c]$ is the extremal length of any reduced representative of $[c]$.

For $[\varphi]\colon G \to H$ a homotopy class of maps between elastic graphs, the (extremal length) *stretch factor* is the maximum ratio of extremal lengths,

$$\mathrm{SF}[\varphi] \coloneqq \sup_{c\colon C\to G} \frac{\mathrm{EL}_H[\varphi\circ c]}{\mathrm{EL}_G[c]}, \tag{1.13}$$

where the supremum runs over all non-trivial multi-curves on G.

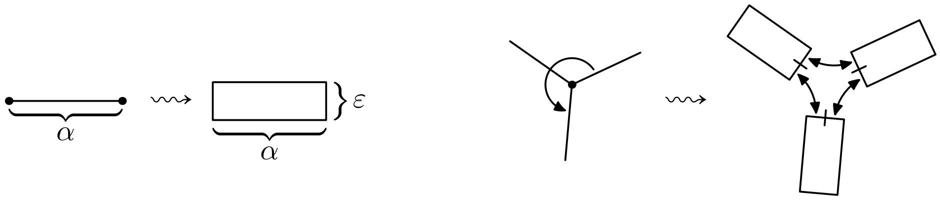

Figure 2. Geometrically thickening an elastic ribbon graph. Left: An edge of elastic length α is thickened to an $\alpha \times \varepsilon$ rectangle. Right: At a vertex, glue each half of the end of each rectangle to one of the neighbors according to the ribbon structure.

Theorem 1.14 ([Thu19, Th. 1]). *For* $[\varphi]\colon G \to H$ *a homotopy class of maps between elastic graphs,*

$$\mathrm{Emb}[\varphi] = \mathrm{SF}[\varphi].$$

The above two quantities are also equal to the maximum ratio of Dirichlet energies between the two graphs. This arises naturally in the proof of Theorem 1.14 and justifies the terminology of "loosening" of elastic graphs. This fact is not used in the present paper, so we will not develop it further here.

* * *

The fourth ingredient is a relation between extremal lengths on a graph and on a certain degenerating family of surfaces. Suppose that G is an elastic ribbon graph, where a *ribbon graph* has a specified counterclockwise cyclic order on edges incident to each vertex (Definition 2.5). Its *ε-thickening* $N_\varepsilon G$ is the conformal surface obtained by replacing each edge e of G by a rectangle of size $\alpha(e) \times \varepsilon$ and gluing the rectangles at the vertices using the given cyclic order, as shown in Figure 2. A *ribbon map* $\varphi\colon G_1 \to G_2$ between ribbon graphs is a map that lifts to a topological embedding $N_\varepsilon\varphi\colon N_\varepsilon G_1 \hookrightarrow N_\varepsilon G_2$ (Definition 2.7).

Theorem 2. *Let G and H be two elastic ribbon graphs with only trivalent vertices, and let $\varphi\colon G \to H$ be a ribbon map between them. Let m be the minimal value of $\alpha(e)$ for e an edge in G or H. Then, for $\varepsilon < m/2$,*

$$\mathrm{SF}[\varphi]/(1 + 8\varepsilon/m) \leqslant \mathrm{SF}[N_\varepsilon\varphi] \leqslant \mathrm{SF}[\varphi] \cdot (1 + 8\varepsilon/m).$$

Theorem 2 is proved in Section 4. We can get some intuition for Theorem 2 from a corollary, which motivates the term "embedding energy" but is not otherwise used.

Corollary 1.15. *Let G_1 and G_2 be two elastic ribbon graphs with trivalent vertices, and let $\varphi\colon G_1 \to G_2$ be a ribbon map between them. Then if* $\mathrm{Emb}[\varphi] < 1$, *for all sufficiently small ε, there is a conformal embedding in $[N_\varepsilon \varphi]$. On the other hand, if for some sufficiently small ε there is a conformal embedding in $[N_\varepsilon \varphi]$, then* $\mathrm{Emb}[\varphi] \leqslant 1$.

Proof. Immediate from Theorem 2, Theorem 1.14, Proposition 4.12, and Theorem 1.9. ☐

* * *

The fifth and final ingredient is a study of the behavior of the stretch factors and embedding energy under iteration. Let $f\colon (S^2, P) \circlearrowleft$ be a branched self-cover. Then we have a virtual endomorphism $\pi, \varphi\colon X_1 \rightrightarrows X_0$, where the X_i are either Riemann surfaces or elastic graphs. By iterating, we get a sequence of virtual endomorphisms $\pi^n, \varphi^n\colon X_n \rightrightarrows X_0$, each with its own stretch factor $\mathrm{SF}[\varphi^n]$ (Definition 5.2). From general principles (Proposition 5.6), it is not hard to prove that the stretch factor grows or shrinks exponentially. That is,

$$\overline{\mathrm{SF}}[\pi, \varphi] \coloneqq \lim_{n\to\infty} \sqrt[n]{\mathrm{SF}[\varphi^n]} \tag{1.16}$$

exists, whether we are working with elastic graphs or with Riemann surfaces. We call this limit the *asymptotic stretch factor*. General principles also show that $\overline{\mathrm{SF}}[\pi, \varphi]$ does not depend on the particular conformal or elastic structure we start with (Proposition 5.7), and furthermore Theorem 2 implies that $\overline{\mathrm{SF}}[\pi, \varphi]$ is the same in these two cases (Proposition 5.14). Theorem 3.3 and Theorem 5.11 (a strengthening of Theorem 1.9) then show that $\overline{\mathrm{SF}}[\pi, \varphi] < 1$ if and only if f is equivalent to a rational map. This is then translated to a proof of Theorem 1, as explained in Section 5.

* * *

The last two sections have material complementary to the main proof. In Section 6, we explain how a virtual endomorphism (π, φ) gives an asymptotic energy $\overline{E}^p_p[\pi, \varphi]$ for each p, with $\overline{E}^2_2$ essentially agreeing with $\overline{\mathrm{SF}}[\pi, \varphi]$. We give bounds on how $\overline{E}^p_p$ varies as a function of p and a few comments on what this might mean.

In Section 7 we compare Theorem 1 with the original characterization, Theorem 7.4.

1.1. *Prior work.* Kahn's work on degenerating surfaces [Kah06, §3] has close ties to this work. In particular, his work is quite similar in spirit to Corollary 1.15. His notion of *domination* of weighted arc diagrams is equivalent to embedding energy being less than one, in the following dualizing sense. Given a elastic ribbon graph $G = (\Gamma, \alpha)$, each edge e of Γ has a dual arc A_e on $N\Gamma$, the arc between boundary components that meets e in one point. We can then consider the dual weighted arc diagram

$$W_G := \sum_{e \in \mathrm{Edge}(\Gamma)} \frac{A_e}{\alpha(e)}.$$

PROPOSITION 1.17. *If* $\varphi\colon G_1 \to G_2$ *is a ribbon map of elastic ribbon graphs, then* $\mathrm{Emb}[\varphi] \leqslant 1$ *if and only if* W_{G_1} *dominates* W_{G_2}.

The proof follows from tracing through the definitions of both notions.

The notion of weighted arc diagrams is a little more general than elastic ribbon graphs, as only weighted arc diagrams for which the set of arcs is filling can be written as W_G for some G. See also [Thu16, §8.5].

The graphical description of branched self-covers in Section 2 is closely related to the description in terms of automata and bisets by Nekrashevych. More specifically, these graphical descriptions are examples of his combinatorial models for expanding dynamical systems [Nek14]. (Nekrashevych also allows models with higher-dimensional cells. See also Section 1.2.)

Theorem 3.3 has been circulating in the community. The non-trivial direction is written as [CPT16, Th. 5.2] or [Wan14, Th. 7.1].

The overall plan of this argument was summarized earlier [Thu16], and some of the arguments were sketched there. For completeness, the logically necessary arguments are reproduced and expanded here.

1.2. *Future directions.* Theorem 1 raises several questions and future directions. First, there is the question of the iteration (n in the theorem statement), and whether it is strictly necessary to iterate to give a positive certificate for any given rational map. It is first of all clear that the necessary value of n depends on the elastic graph spine G (including a dependence on the elastic constants). See [Thu16, Fig. 2] for one concrete example. So then the question becomes whether we can always find an elastic spine G so that $n = 1$ suffices. There are some rational where $n = 1$ does *not* suffice for any G by a case analysis. These include the barycentric subdivision rational map [CFKP03, Fig. 25],

$$f(z) = \frac{4}{27}\frac{(z^2 - z + 1)^3}{(z(z-1))^2},$$

and rational map #3.2 in the census of quadratic rational maps with at most four post-critical points [BBL$^+$00],

$$f(z) \approx \frac{13.531903z - 13.531903}{(z + 2.3829758)^2}.$$

Analysis of the first case is made easier using the symmetries of the map. Analysis of the second case is more involved, and we do not give the details. In both these cases, the Julia set is a Sierpiński carpet. (The second example is the unique quadratic map in that census with a Sierpiński carpet Julia set.)

For many other examples, it seems that $n = 1$ does suffice. In particular, a crucial role appears to be played by the *crochet maps*, rational maps where two Fatou components can be connected by a path that intersects the Julia set in only countably many points. (The terminology was introduced by Dudko, Hlushchanka, and Schleicher, in ongoing work.) They are also conjecturally the rational maps for which the Ahlfors regular conformal dimension is equal to 1.

CONJECTURE 1.18. *For any crochet post-critically finite rational map,* $n = 1$ *suffices in Theorem 1: There is an elastic* (*orbi*)*graph spine* G *for which* $\mathrm{Emb}_G[\varphi] < 1$ *without the need to iterate.*

Question 1.19. For crochet rational maps, is there a preferred "best" spine for $\widehat{\mathbb{C}}\setminus P$? For polynomials, there is a slight modification of the Hubbard tree that serves this purpose [Thu16, Def. 8.8], and in many cases there appear to be good candidates.

Beyond this, one might hope for several extensions of Theorem 1, in a variety of different directions. We briefly survey some of them.

First, it is possible to relax the restriction to maps of hyperbolic type in Theorem 1, to allow maps of non-compact type. As stated, the theorem is false in this generality. Consider, for example, the dendritic polynomial $z^2 + i$, with graph virtual endomorphism $\pi, \varphi \colon \Gamma_1 \rightrightarrows \Gamma_0$ $\pi, \varphi \colon \Gamma_1 \rightrightarrows \Gamma_0$

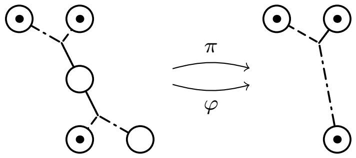

.

The black loops around the marked points are mapped to themselves by degree 1. This persists under covers, so $\overline{\mathrm{SF}}[\pi, \varphi] = 1$ and we do not have the strict inequality we need for Theorem 1. Essentially, with our definition of spines, the loops around these marked points count as "Levy cycles," obstructing our criterion.

A proper treatment of this family of maps uses *orbigraphs*, spaces locally modeled on a graph modulo a finite group, so that the graph virtual endomorphism in this example becomes

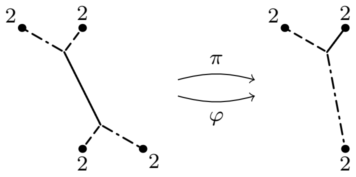

,

where the "2" mark an orbifold point of order 2, the quotient of an edge by an involution. Details will appear in a future paper; see also [Thu16, Prob. 8.22].

On the other hand, any graph-based criterion is unlikely to work when the branched self-cover f has *no* cycles with branch points. For rational maps, these are the cases when the Julia set is the whole Riemann sphere $\widehat{\mathbb{C}}$. The issue is that for a graph virtual endomorphism that is *contracting* (i.e., in the language of Section 6, $\overline{E}_{\infty}^{\infty}[\pi, \varphi] < 1$), by a result of Nekrashevych [Nek14] the Julia set is homeomorphic to the inverse limit of graphs Γ_n with respect to the maps φ_{n-1}^{n}. (See Definition 5.1 for terminology.) But any inverse limit of graphs has topological dimension 1, and so cannot be homeomorphic to $\widehat{\mathbb{C}}$.

One might also ask for a generalization of Theorem 1 to allow the post-critical set P to be infinite, probably with some other restrictions. For W. Thurston's original theory, this a fruitful area of research, with a series of papers [CJS03a], [CJS03b], [ZJ09], [CT11], [CT18] leading ultimately to an obstruction theorem for maps where the accumulation set of P is finite (the *geometrically finite maps*).

In another direction, exponential maps and other transcendental maps $\mathbb{C} \to \mathbb{C}$ have been a fruitful area of research, with substantial more complexity than the rational map case. In particular, the natural analog of Question 1.2 is almost entirely open. This deals with topological maps of transcendental type that are post-singularly finite (to include the forward orbits of asymptotic values). For the obstruction criterion, there is a result for the exponential family [HSS09], but this is a special case and it is not clear what to expect in general. It is natural to ask whether there is any analogue of the positive criterion presented in this paper.

In both these cases (geometrically finite or transcendental maps), any graphical model needs to allow for infinite graphs: either because the post-critical set is infinite (and we want some sort of spine for its complement), or because the degree of the cover is infinite. Some of the ingredients in this paper carry over without issue. For instance, Propositions 4.8 and 4.9, relating extremal length on graphs and their thickening to a surface, hold for infinite graphs. (Indeed, Theorem 2 can be generalized considerably, to allow grafting along arbitrary embedded arcs and/or circles, with some weakening of the conclusion.) Likewise the quasi-conformal surgery techniques recalled in Section 3 have been well studied in these more general contexts.

Other ingredients appear harder to generalize to the setting of infinite graphs (or infinitely-generated π_1). For instance, in earlier work, we gave several equivalent conditions for conformal embeddings of Riemann surfaces of finite type with some "space" around them [KPT15]. These conditions are unlikely to be equivalent in the setting of arbitrary Riemann surfaces.

Another possible generalization is to higher-dimensional maps, for instance studying maps of higher-dimensional manifolds that are post-critically finite in a suitable sense. It is not entirely clear what the right questions or conjectures should be, but one can often find suitable combinatorial models for such maps as virtual endomorphisms of CW complexes. Nekrashevych has both a general theory [Nek14] and concrete examples [Nek16]. The major obstacle to developing a theory similar to the one in this paper in the higher-dimensional setting is finding the right analogue of the energies E^p_p recalled in Section 6. The definitions rely heavily on the underlying objects being 1-dimensional. (See [Thu19, Def. A.24] for a possible direction towards a generalization.) Furthermore, the proof of existence of minimizers from earlier work is combinatorial and produces piecewise-linear minimizing maps; this approach will not work in higher dimensions. This remains work in progress.

Another direction is investigating the meaning of the additional energies $\overline{E}^p_p$ in Section 6. In work in progress with Kevin Pilgrim, we relate these energies to the Ahlfors regular conformal dimension [PT].

Acknowledgements. This project grew out of joint work with Kevin Pilgrim and owes a great deal to conversations with him and with Jeremy Kahn. In addition, there were many helpful conversations and comments from Maxime Fortier-Bourque, Frederick Gardiner, Mikhail Hlushchanka, Volodymyr Nekrashevych, Tan Lei, and the anonymous referee.

This material is based upon work supported by the National Science Foundation under Grant Number DMS-1507244.

2. Spines for branched self-covers

Definition 2.1. A *virtual endomorphism* of a group G is a finite-index subgroup $H \subset G$ and a homomorphism $\varphi\colon H \to G$.

A *virtual endomorphism* of a topological space X consists of a space Y and a pair of maps

$$\pi, \varphi\colon Y \rightrightarrows X,$$

where π is a covering map of constant, finite degree and φ is considered up to homotopy.

A virtual endomorphism of spaces gives a virtual endomorphism of groups, as follows. Suppose X and Y are connected and locally connected and $x_0 \in X$ is a basepoint. If we pick $y_0 \in \pi^{-1}(x_0)$, then $\pi_1(Y, y_0)$ is naturally a subgroup of $\pi_1(X, x_0)$. If we homotope φ so that $\varphi(y_0) = x_0$, then φ_* gives a group homomorphism from $\pi_1(Y, y_0)$ to $\pi_1(X, x_0)$, i.e., a virtual endomorphism of $\pi_1(X, x_0)$.

Virtual endomorphisms of topological (orbi)spaces are also called *topological automata* by Nekrashevych [Nek14]. If you drop the condition that π be a covering map, the same structures were called *topological graphs* or *topological correspondences* by Katsura [Kat04] and *multi-valued dynamical systems* by Ishii and Smillie [IS10].

Definition 2.2. A *homotopy morphism* between two virtual endomorphisms, from $\pi_X, \varphi_X \colon X_1 \rightrightarrows X_0$ to $\pi_Y, \varphi_Y \colon Y_1 \rightrightarrows Y_0$, is a pair of maps $f_0 \colon X_0 \to Y_0$ and $f_1 \colon X_1 \to Y_1$ so that

$$\begin{aligned} f_0 \circ \pi_X &= \pi_Y \circ f_1, \\ f_0 \circ \varphi_X &\sim \varphi_Y \circ f_1, \end{aligned}$$

where $\sim$ means homotopy of maps.

A *homotopy equivalence* between (π_X, φ_X) and (π_Y, φ_Y) is a pair of homotopy morphisms (f_0, f_1) from X to Y and (g_0, g_1) from Y to X, so that $f_0 \circ g_0 \sim \mathrm{id}_{Y_0}$ and $g_0 \circ f_0 \sim \mathrm{id}_{X_0}$. This implies that $f_1 \circ g_1 \sim \mathrm{id}_{Y_1}$ and $g_1 \circ f_1 \sim \mathrm{id}_{X_1}$, as shown below:

$$\begin{tikzcd}
X_1 \arrow[r, bend left, "\pi_X"] \arrow[r, bend right, "\varphi_X"'] \arrow[d, bend left, "f_1"] \arrow[d, phantom, "\sim"] & X_0 \arrow[d, bend left, "f_0"] \arrow[d, phantom, "\sim"] \\
Y_1 \arrow[u, bend left, "g_1"] \arrow[r, bend left, "\pi_Y"] \arrow[r, bend right, "\varphi_Y"'] & Y_0 \arrow[u, bend left, "g_0"]
\end{tikzcd}$$

If $f \colon (\Sigma, P) \circlearrowleft$ is a branched self-cover of a surface, let $\Sigma_0 = \Sigma \setminus P$ and $\Sigma_1 = \Sigma \setminus f^{-1}(P)$. The restriction of f gives a covering map $\pi_\Sigma \colon \Sigma_1 \to \Sigma_0$, and the natural inclusion of surfaces gives a map $\varphi_\Sigma \colon \Sigma_1 \to \Sigma_0$, together forming a surface virtual endomorphism

$$\pi_\Sigma, \varphi_\Sigma \colon \Sigma_1 \rightrightarrows \Sigma_0. \tag{2.3}$$

A *spine* of Σ_0 is a graph $\Gamma_0 \subset \Sigma_0$ that is a deformation retract of Σ_0. If we replace Σ_0 in (2.3) by a spine Γ_0, we get spaces and maps

- $\Gamma_1 = f^{-1}(\Gamma_0) \subset \Sigma_1$;
- deformation retractions $\kappa_i \colon \Sigma_i \to \Gamma_i$;
- the restriction of f to a covering of graphs $\pi_\Gamma \colon \Gamma_1 \to \Gamma_0$; and
- $\varphi_\Gamma = \kappa_0 \circ \varphi_\Sigma \colon \Gamma_1 \to \Gamma_0$.

These form a graph virtual endomorphism

$$\pi_\Gamma, \varphi_\Gamma \colon \Gamma_1 \rightrightarrows \Gamma_0. \tag{2.4}$$

Since the κ_i are homotopy equivalences, $[\varphi_\Gamma]$ is determined by $[\varphi_\Sigma]$. While φ_Σ is a topological inclusion, φ_Γ is just a continuous map of graphs. We say $(\pi_\Gamma, \varphi_\Gamma)$ is *compatible* with the branched self-cover f. Since any two spines for Σ_0 are homotopy equivalent, the homotopy equivalence class $[\pi_\Gamma, \varphi_\Gamma]$ is determined by f.

To go the other direction and recover the branched self-cover from the graph virtual endomorphism $\pi_\Gamma, \varphi_\Gamma \colon \Gamma_1 \rightrightarrows \Gamma_0$, we need some more data.

Definition 2.5. A *ribbon structure* on a graph Γ is, for each vertex v of Γ, a cyclic ordering on the ends of edges incident to v, thought of as the counterclockwise ordering. A ribbon structure gives a canonical thickening of Γ into an oriented surface with boundary $N\Gamma$, the underlying topological surface of the Riemann surface $N_\varepsilon\Gamma$ from Figure 2. There is a natural inclusion $i_{N\Gamma} \colon \Gamma \hookrightarrow N\Gamma$ and projection $\pi_{N\Gamma} \colon N\Gamma \to \Gamma$.

We will prove that a virtual endomorphism of a ribbon graph is compatible with at most one branched self-cover.

For an example of what this data looks like, consider the rational map

$$f(z) = (1+z^2)/(1-z^2),$$

with critical portrait

$$\begin{tikzcd} 0 \arrow[r, "(2)"] & 1 \arrow[r] & \infty \arrow[r, bend left, "(2)"] & -1. \arrow[l, bend left] \end{tikzcd} \tag{2.6}$$

We take P to be the post-critical set $\{-1, 1, \infty\}$. We can take Γ_0 to be a Θ-graph embedded in $\Sigma_0 = S^2 \setminus P$ and take Γ_1 to be $f^{-1}(\Gamma_0)$, as indicated in Figure 3a. The map π is the covering map that preserves labels and orientations on the edges. The map φ might, for instance, be chosen so that

- the two a edges of Γ_1 map to the a and b edges of Γ_0;
- the two b edges of Γ_1 map to the c edge of Γ_0; and
- the two c edges of Γ_1 map with a constant map to the two vertices of Γ_0.

To read off the critical portrait, first recall that from a connected ribbon graph embedded in the plane, the complementary regions are intrinsically determined by following the boundary of the ribbon surface. Thus we can talk about the regions of Γ_0 and Γ_1. Then, for instance, the point "1" is in the region of Γ_1 surrounded by an a edge and a b edge, so must map by f to the point "∞," which is in the exterior region of Γ_0, also surrounded by an a edge and a b edge. On the other hand, "∞" in the exterior region of Γ_1 is surrounded by an a, c, a, and c edge, and so maps with double branching to "-1," in the region of Γ_0 surrounded by a and c. Proceeding in this way, we recover the critical portrait (2.6).

This data is essentially equivalent to an automaton in the style of Nekrashevych [Nek05]. To construct the automaton, first choose a spanning tree T_0 inside Γ_0 and collapse it to get a rose graph spine R_0 for Σ_0. (A *rose graph* is a graph with one vertex.) If we collapse $\pi^{-1}(T_0)$ inside Γ_1, we get R_1, which is

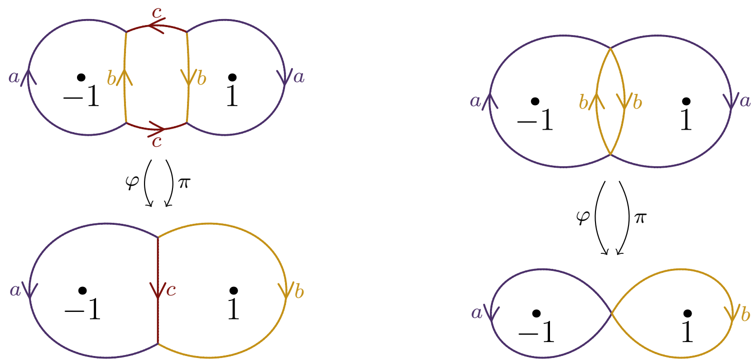

(a) Virtual endomorphism of a spine for $S^2 \backslash \{-1, 1, \infty\}$. The marked point ∞ is at infinity.

(b) Virtual endomorphism of a rose graph.

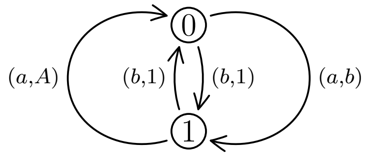

(c) Dual Moore diagram. 1 is the identity element in the group $F_2 = \langle a, b \rangle$, and capital letters denote inverses.

$$\begin{aligned} a(0 \cdot w) &= 1 \cdot b(w) \\ a(1 \cdot w) &= 0 \cdot A(w) \\ b(0 \cdot w) &= 1 \cdot w \\ b(1 \cdot w) &= 0 \cdot w \end{aligned}$$

(d) Textual description of automaton.

Figure 3. Representations of the rational map $z \mapsto (1 + z^2)/(1 - z^2)$.

likewise a spine for Σ_1. (R_1 is not itself a rose graph.) Since R_0 is also a spine for $S^2 \backslash P$, there is a virtual endomorphism $\pi_R, \varphi_R \colon R_1 \rightrightarrows R_0$. In the running example, if we take the spanning tree of Γ_0 to be edge c in Figure 3a, we get the graphs R_0 and R_1 in Figure 3b.

The graph R_1 constructed above is quite close to the dual Moore diagram for the corresponding automaton, shown in Figure 3c for the running example. To get from Figure 3b to Figure 3c, perform the following steps:

(1) Homotop the graph map $\varphi \colon R_1 \to R_0$ so that it sends vertices to the vertex of R_0. In the example, the two b edges of R_1 get mapped to points.
(2) As a graph, the dual Moore diagram D is R_1, with vertices numbered arbitrarily.
(3) Label each edge e of D by, first, the label of e in R_1 and, second, the element of $\pi_1(R_0)$ represented by $\varphi(e)$.

The dual Moore diagram encodes an automaton, which in the example is given textually in Figure 3d.

Returning to the general theory, not all combinations of a graph virtual endomorphism $\pi, \varphi\colon \Gamma_1 \rightrightarrows \Gamma_0$ and a ribbon graph structure on Γ_0 are compatible with a branched self-cover.

Definition 2.7. If Γ and Γ' are ribbon graphs, a *ribbon map* $\varphi\colon \Gamma \to \Gamma'$ is a map that lifts to an orientation-preserving topological embedding $N\varphi\colon N\Gamma \hookrightarrow N\Gamma'$, in the sense that $\varphi = \pi_{N\Gamma'} \circ N\varphi \circ i_{N\Gamma}$.

Lemma 2.8. *If $\varphi\colon \Gamma \to \Gamma'$ is a map between ribbon graphs and Γ and Γ' are connected, then up to isotopy there is at most one orientation-preserving lift $N\varphi\colon N\Gamma \to N\Gamma'$.*

Proof. This follows from the fact that any two orientation-preserving homotopic embeddings from one connected surface to another are isotopic, which in turn follows from work of Epstein [Eps66] by looking at the boundary curves [Put16]. It is also proved as a side effect of work of Fortier Bourque on conformal embeddings [FB18]. ☐

Definition 2.9. Suppose that $\pi, \varphi\colon \Gamma_1 \rightrightarrows \Gamma_0$ is a graph virtual endomorphism where Γ_0 has a ribbon graph structure. We can use the covering map π to pull back the ribbon structure on Γ_0 to a ribbon structure on Γ_1. Then we say the data form a *ribbon virtual endomorphism* if φ is a ribbon map. A *ribbon homotopy morphism* between two ribbon virtual endomorphisms is a homotopy morphism as in Definition 2.2 so that f_0 and f_1 are ribbon maps.

Remark 2.10. It is not immediately clear how to give an efficient algorithm to check whether a topological map $\varphi\colon \Gamma \to \Gamma'$ between ribbon graphs is a ribbon map, but we can give an inefficient algorithm. If we specify, for each regular point $y \in \Gamma'$, the order in which the points in $\varphi^{-1}(y)$ appear on the corresponding cross-section of $N\Gamma'$, it is easy to check locally whether there is an embedded lift. Since there are only finitely many choices of orders, this can be checked algorithmically.

Definition 2.11. A map $\varphi\colon X \to Y$ between locally path-connected topological spaces is π_0*-bijective* if it gives a bijection from the connected components of X to the connected components of Y. (Recall that branched self-covers are assumed to be π_0-bijective.) The map φ is π_1*-surjective* if, for each $x \in X$, the induced map $\varphi_*\colon \pi_1(X, x) \to \pi_1(Y, \varphi(x))$ is surjective.

Theorem 3. *Branched self-covers of surfaces $f\colon (\Sigma, P) \circlearrowleft$, up to equivalence, are in bijection with ribbon virtual endomorphisms $\pi, \varphi\colon \Gamma_1 \rightrightarrows \Gamma_0$ so that φ is π_0-bijective and π_1-surjective, up to ribbon homotopy equivalence.*

Remark 2.12. Theorem 3 does not assume *a priori* that the surface Σ is a (union of) spheres or that the Γ_i are planar. Once the theorem is proved then the usual Euler characteristic arguments imply that each component of Σ is a sphere or a torus, with sphere being by far the more interesting case.

Proof. If we are given a branched self-cover $f\colon (\Sigma, P) \circlearrowleft$, we have already seen how to pick a compatible spine $\Gamma_0 \subset \Sigma \backslash P$ and construct a ribbon virtual endomorphism $\pi, \varphi\colon \Gamma_1 \rightrightarrows \Gamma_0$, unique up to ribbon homotopy equivalence. It is immediate that φ is π_0-bijective and π_1-surjective.

It remains to check the other direction. Suppose we have a ribbon virtual endomorphism as in the statement. Let $\Sigma_0 = N\Gamma_0$ and $\Sigma_1 = N\Gamma_1$. Since φ is π_0-bijective, Lemma 2.8 tells us the lift $N\varphi\colon \Sigma_1 \hookrightarrow \Sigma_0$ is unique. Let $\widehat{\Sigma}_0$ be the marked surface obtained by attaching a disk with a marked point in the center to each boundary component of Σ_0. Let $P_0 \subset \widehat{\Sigma}_0$ be the set of marked points.

Recall that a simple closed curve C on a closed surface is separating if and only if it is homologically trivial, and that it is non-separating if and only if there is a "dual" simple closed curve C' that intersects C transversally in one point.

Let C_1 be a boundary component of Σ_1, and consider the simple curve $C_0 = N\varphi(C_1) \subset \Sigma_0 \subset \widehat{\Sigma}_0$. If C_0 is non-separating, a dual curve cannot be homotoped to lie in the image of $N\varphi$, contradicting π_1-surjectivity. So C_0 is separating and divides $\widehat{\Sigma}_0$ into two components, with one component containing the image of $N\varphi$. If the other component is not a disk with 0 or 1 marked points, then again φ is not π_1-surjective. If $N\varphi(C_1)$ bounds a disk with no marked points in $\widehat{\Sigma}_0$ (so bounds a disk in Σ_0), say that C_1 is *collapsed*.

Now construct $\widehat{\Sigma}_1$ by attaching a disk D_i to each boundary component C_i of Σ_1. Mark the center of D_i if C_i is not collapsed. Let $P_1 \subset \widehat{\Sigma}_1$ be the set of marked points. By the choices made in the construction, $N\varphi$ extends to a homeomorphism $g\colon \widehat{\Sigma}_1 \to \widehat{\Sigma}_0$ inducing a bijection from P_1 to P_0.

Since the ribbon structure on Γ_1 is the pull-back of the ribbon structure on Γ_0, the covering map π extends to a covering of surfaces $N\pi\colon \Sigma_1 \to \Sigma_0$. Since $N\pi$ restricts to a covering map from $\partial N\Gamma_1$ to $\partial N\Gamma_0$, we can extend $N\pi$ to a branched cover $h\colon \widehat{\Sigma}_1 \to \widehat{\Sigma}_0$ with $h(P_1) \subset P_0$ and branch values contained in P_0.

The desired branched self-covering is then $f = h \circ g^{-1}\colon (\widehat{\Sigma}_0, P_0) \circlearrowleft$. The original virtual endomorphism $\pi, \varphi\colon \Gamma_1 \rightrightarrows \Gamma_0$ is compatible with f. ☐

Example 2.13. To see that the ribbon structure is necessary in Theorem 3, consider the 1/5 and 2/5 rabbit (i.e., the centers of the 1/5 and 2/5 bulb of the Mandelbrot set), with compatible graph virtual endomorphisms shown in Figure 4. The two branched self-covers are different, but the graph virtual endomorphisms are the same except for the ribbon structure.

Remark 2.14. The relationship with Nekrashevych's automata can be used to give another proof of the uniqueness part of Theorem 3 [Nek05, Th. 6.5.2].

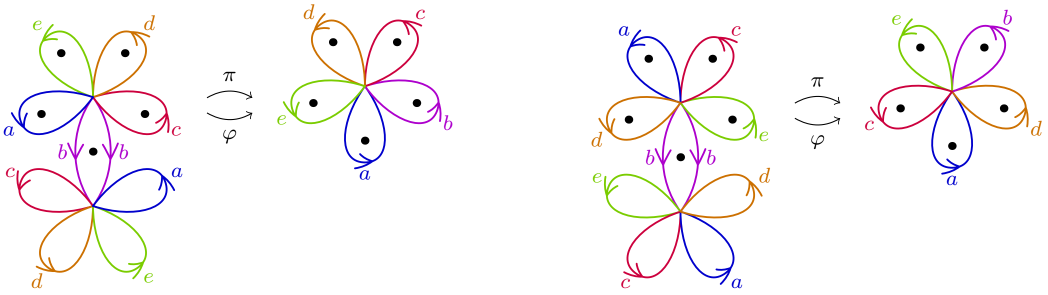

Figure 4. Spines for the 1/5 rabbit (left) and 2/5 rabbit (right). There are extra marked points at infinity. The map π is the cover that preserves labels and the map φ is determined by the deformation retraction.

3. Quasi-conformal surgery

We now turn to the (standard) characterization of rational maps in terms of conformal embeddings of surfaces. For this section, we generalize to maps of non-compact type, since we can do it with little extra work.

Let $f\colon (\Sigma, P) \circlearrowleft$ be a branched self-cover of non-compact type. Let P_F (with F for Fatou) be the set of points in P whose forward orbit under f lands in a cycle with a branch point, and let P_J (with J for Julia) be $P \setminus P_F$. Let Σ_0 be the complement of a disk neighborhood of P_F in Σ, with marked subset $P_{0J} \coloneqq P_J$. Parallel to (2.3), there is a *branched* virtual endomorphism

$$\pi, \varphi\colon (\Sigma_1, P_{1J}) \rightrightarrows (\Sigma_0, P_{0J}),$$

where $\Sigma_1 \coloneqq f^{-1}(\Sigma_0)$, P_{1J} is P_J as a subset of Σ_1, π is a branched cover with branch values contained in P_{0J}, and φ induces a bijection between P_{1J} and P_{0J}. We consider φ up to homotopy relative to P_{1J}.

Given a branched virtual endomorphism of surfaces $\pi, \varphi\colon (\Sigma_1, P_{1J}) \rightrightarrows (\Sigma_0, P_{0J})$, if there is a conformal structure ω_0 on Σ_0, there is a pull-back conformal structure $\omega_1 \coloneqq \pi^*\omega_0$ on Σ_1. Then (π, φ) is said to be *conformal* with respect to ω_0 if φ is homotopic (rel P_{1J}) to an annular conformal embedding from (Σ_1, ω_1) to (Σ_0, ω_0).

Definition 3.1. The *Teichmüller space* of a branched virtual endomorphism of surfaces $\pi, \varphi\colon (\Sigma_1, P_{1J}) \rightrightarrows (\Sigma_0, P_{0J})$ is the space $\mathrm{Teich}(\pi, \varphi)$ of isotopy classes of complex structures ω_0 on Σ_0 so that (π, φ) is conformal with respect to ω_0.

Remark 3.2. In the literature, the condition that the embedding φ be annular is often omitted.

Theorem 3.3. *Let* $f\colon (\Sigma, P) \circlearrowleft$ *be a branched self-cover of non-compact type, and let* $\pi, \varphi\colon (\Sigma_1, P_{1J}) \rightrightarrows (\Sigma_0, P_{0J})$ *be the associated branched virtual endomorphism. Then* f *is equivalent to a rational map if and only if* $\operatorname{Teich}(\pi, \varphi)$ *is non-empty.*

Proof. We start with the easy direction. If f is equivalent to a rational map, replace it with its rational map $f\colon (\widehat{\mathbb{C}}, P) \circlearrowleft$. Let $J \subset \widehat{\mathbb{C}}$ be the Julia set of f. Then $P \cap J = P_J$. Set S_0 to be a suitable open neighborhood of J, chosen so that $\overline{f^{-1}(S_0)} \subset S_0$. Then we can take $S_1 = f^{-1}(S_0)$.

To be concrete about the "suitable open neighborhood," let F be the Fatou set of f, and choose a Green's function on F, by which we mean a harmonic function $G\colon F \to (0, \infty]$ (taking the value ∞ only at discrete points) so that

$$G(f(z)) = \lambda_z G(z),$$

where $\lambda_z > 1$ is a locally constant function on F so that $\lambda_{f(z)} = \lambda_z$. (This implies G goes to zero near ∂F.) Concretely, if z is attracted to a cycle in P_F of period p and total branching d, then $\lambda_z = d^{1/p}$. On a basin B_i of F with Böttcher coordinate $\varphi_i\colon B_i \to \mathbb{D}$, we can take $G(z) = -K_i \log|\varphi_i|$ on B_i for some constant K_i.

Extend G to all of $\widehat{\mathbb{C}}$ by setting $G(z) = 0$ for $z \in J$. Pick $\varepsilon > 0$, and take S_0 to be the union of $G^{-1}([0, \varepsilon])$ and all basins in F that do not contain a point of P_F. We then have a conformal branched virtual automorphism $\pi, \varphi\colon (f^{-1}(S_0), P_J) \rightrightarrows (S_0, P_J)$.

The other direction is a special case of [CPT16, Th. 5.2] or [Wan14, Th. 7.1]. The technique is due to Douady and Hubbard [DH85]. We sketch the proof here.

Suppose ω_0 is a point in $\operatorname{Teich}(\pi, \varphi)$, and consider the corresponding conformal maps $\pi, \varphi\colon (S_1, P_{1J}) \rightrightarrows (S_0, P_{0J})$. We extend π and φ to maps $\widehat{\pi}, \widehat{\varphi}\colon (\widehat{S}_1, P_1) \rightrightarrows (\widehat{S}_0, P_0)$ between closed surfaces as in the proof of Theorem 3, but with attention to keeping the maps conformal except in controlled ways.

- Let $\widehat{S}_0$ be the compact Riemann surface without boundary obtained by attaching disks to the boundary components of S_0. Let D_0 be the union of the new disks, and let $V_0 \subset D_0$ be a union of concentrically contained smaller disks. Let P_{0F} be the set of points in the center of each component of V_0 (and D_0). Let $P_0 = P_{0J} \sqcup P_{0F}$.
- Let $\widehat{S}_1$ be the corresponding conformal branched cover of $\widehat{S}_0$, branched at points in P_0, with $\widehat{\pi}\colon \widehat{S}_1 \to \widehat{S}_0$ extending π. Let D_1 be $\widehat{\pi}^{-1}(D_0)$. Define $P_{1F} \subset \widehat{S}_1$ by picking those points in $\widehat{\pi}^{-1}(P_{0F})$ that are in the center of non-collapsed boundary components of S_1, as in the proof of Theorem 3. Let $V_1 \subset D_1$ be the union of those components of $\widehat{\pi}^{-1}(V_0)$ that contain a point in P_{1F}, and let $P_1 = P_{1J} \sqcup P_{1F}$.

Next we will define a diffeomorphism $\widehat{\varphi}\colon (\widehat{S}_1, P_1) \to (\widehat{S}_0, P_0)$ in stages.

- On S_1, the map $\widehat{\varphi}$ agrees with φ.
- Each component of V_1 contains a point $p_1 \in P_{1F}$, which corresponds to a point $p_0 \in P_{0F}$. On this component, $\widehat{\varphi}$ is a conformal identification of the disk of V_1 to the disk of D_0 containing p_0, mapping p_1 to p_0.
- It remains to define $\widehat{\varphi}$ on $D_1 \backslash V_1$, which is a union of annuli and unmarked disks. Define $\widehat{\varphi}$ to be an arbitrary diffeomorphism extending the maps defined so far. This is possible since we have not changed the isotopy class from the homeomorphism from Theorem 3.

Observe that $\widehat{\varphi}$ and $\widehat{\varphi}^{-1}$ are K-quasi-conformal for some $K \geqslant 1$, since they are diffeomorphisms on compact manifolds. The map $\widehat{\varphi}$ takes the pieces of $\widehat{S}_1$ to the pieces of $\widehat{S}_0$:

$$\begin{tikzcd}
S_1 \arrow[dr, "\text{conf.}"'] & \sqcup & D_1\backslash V_1 \arrow[dl, "\text{q.c.}"] & \sqcup & V_1 \arrow[dl, "\text{conf.}"] \\
& S_0 & \sqcup & D_0. &
\end{tikzcd}$$

Now let $f\colon (\widehat{S}_0, P_0) \circlearrowleft$ be $\widehat{\pi} \circ \widehat{\varphi}^{-1}$. Since $\widehat{\pi}$ is conformal, f is K-quasi-regular, f is conformal on $\varphi(S_1) \sqcup D_0$, and $f(D_0) = V_0 \subset D_0$. If f is not conformal at $x \in \widehat{S}_1$, then $f(x) \in D_0$, so f is conformal at $f(x)$ and at all further forward iterates of $f(x)$. That is, in the forward orbit of any $x \in \widehat{S}^1$, there is at most one n so that f is not conformal at $f^{\circ n}(x)$. Thus $f^{\circ n}$ is also K-quasi-regular with the same value of K. We can therefore apply Sullivan's Averaging Principle to find an invariant measurable complex structure on $\widehat{S}_1$, which by the Measurable Riemann Mapping Theorem can be straightened to give an honest conformal structure and a post-critically finite map [Sul81]. ☐

Remark 3.4. Theorem 3.3 is also true if we take P_F to be an invariant subset of the points of P whose forward orbit lands in a branched cycle, as long as there is at least one point of P_F in each component of Σ. For instance, for polynomials we may take $P_F = \{\infty\}$; this is essentially the setting of Douady-Hubbard's original paper.

4. **Extremal length on thickened surfaces**

Our next goal is to relate extremal length on elastic graphs and on Riemann surfaces. We first recall different definitions of extremal length on surfaces.

Definition 4.1. Let A be a conformal annulus. Then the *extremal length* $\mathrm{EL}(A)$ is defined by the following equivalent definitions (equivalence is standard):

1. There are real numbers $s, t \in \mathbb{R}_{>0}$ so that A is conformally equivalent to the quotient of the rectangle $[0, s] \times [0, t]$ by identifying $(0, y)$ with (s, y) for $y \in [0, t]$. Then $\mathrm{EL}(A)$ is s/t, the circumference divided by the height.
2. Pick a Riemannian metric g in the conformal class of A. For $\rho\colon A \to \mathbb{R}_{\geqslant 0}$ a suitable (Borel-measurable) scaling function, let $\ell_\circ(\rho g)$ be the minimal length, with respect to the pseudo-metric ρg, of any curve homotopic to the core of A, and let $\mathrm{Area}(\rho g)$ be the area of A with respect to ρg. Then
$$\mathrm{EL}(A) = \mathrm{EL}_\circ(A) = \sup_\rho \frac{\ell_\circ(\rho g)^2}{\mathrm{Area}(\rho g)}. \tag{4.2}$$
3. For g and ρ as above, let $\ell_\perp(\rho g)$ be the minimal length with respect to ρg of any curve running from one boundary component of A to the other. Then
$$\mathrm{EL}_\perp(A) = \sup_\rho \frac{\ell_\perp(\rho g)^2}{\mathrm{Area}(\rho g)}, \tag{4.3}$$
$$\mathrm{EL}(A) = 1/\,\mathrm{EL}_\perp(A). \tag{4.4}$$

Definition 4.5. Let S be a general Riemann surface, and let (C, c) be a simple multi-curve on S (a union of non-intersecting simple closed curves), with components $c_i\colon C_i \to S$. Then $\mathrm{EL}_S[c]$ is defined in the following equivalent ways:

1. If g is a Riemannian metric on S in the distinguished conformal class, then
$$\mathrm{EL}_S[c] = \sup_\rho \frac{\ell_{\rho g}[c]^2}{\mathrm{Area}_{\rho g}(S)}, \tag{4.6}$$
where again $\rho\colon S \to \mathbb{R}_{\geqslant 0}$ runs over all Borel-measurable scaling factors and $\ell_{\rho g}[c]$ is the minimal length of any multi-curve in $[c]$ with respect to ρg.
2. Extremal length may be defined by finding the "fattest" set of annuli around $[c]$, as follows. For $i = 1, \dots, k$, let A_i be a (topological) annulus. Then
$$\mathrm{EL}[c] = \inf_{\omega, f} \sum_{i=1}^{k} \mathrm{EL}_\omega(A_i), \tag{4.7}$$
where the infimum runs over all conformal structures ω on the A_i (i.e., a choice of modulus) and over all embeddings $f\colon \bigsqcup_i A_i \hookrightarrow S$ that are conformal with respect to ω and so that f restricted to the core curve of A_i is isotopic to c_i.

These two definitions are equivalent [KPT15, Prop. 3.7].

Also define $\mathrm{EL}_\perp(\bigsqcup A_i)$, the perpendicular extremal length of a union of annuli, as the extremal length of the union of path families running between boundary components:

$$\mathrm{EL}_\perp\Bigl(\bigsqcup_i A_i\Bigr) := \sup_\rho \frac{\bigl(\min_i \ell_\perp(A_i;\rho g)\bigr)^2}{\mathrm{Area}(\rho g)}.$$

With this definition, it is easy to verify that $\sum_i \mathrm{EL}(A_i) = 1/\mathrm{EL}_\perp\bigl(\bigsqcup_i A_i\bigr)$.[2]

For a non-simple homotopy class of multi-curves on S, use equation (4.6) as the definition of extremal length.

To prove Theorem 2, we need to estimate extremal length on $N_\varepsilon G$ from below and above. We prove two propositions for the two directions.

PROPOSITION 4.8. *Let G be an elastic ribbon graph. Then for $[c]$ any homotopy class of multi-curve on NG, we have*

$$\mathrm{EL}_G[c] \leqslant \varepsilon\,\mathrm{EL}_{N_\varepsilon G}[c].$$

(We use $[c]$ for the homotopy class on both G and on $N_\varepsilon G$.)

Proof. We use equation (4.6) to estimate $\mathrm{EL}_{N_\varepsilon G}[c]$ from below. Take as the base metric g the standard piecewise-Euclidean metric in which an edge e gives an $\alpha(e) \times \varepsilon$ rectangle $N_\varepsilon e$.

The test function ρ is the piecewise-constant function which is $n_c(e)$ on $N_\varepsilon e$. (Recall that $n_c(e)$ is the number of times c runs over e.) Then the shortest representative of $[c]$ will run down the center of each rectangle, so

$$\ell_{\rho g}[c] = \sum_{e\in\mathrm{Edge}(\Gamma)} (n_c(e))^2\alpha(e).$$

On the other hand, the area is

$$\mathrm{Area}(\rho g) = \sum_{e\in\mathrm{Edge}(G)} (\alpha(e)n_c(e))\cdot(\varepsilon n_c(e)),$$

so

$$\varepsilon\,\mathrm{EL}_{N_\varepsilon G}[c] \geqslant \frac{\varepsilon\ell_{\rho g}[c]^2}{\mathrm{Area}(\rho g)} = \frac{\varepsilon\Bigl(\sum n_c(e)^2\alpha(e)\Bigr)^2}{\varepsilon\sum n_c(e)^2\alpha(e)} = \mathrm{EL}_G[c]. \qquad \square$$

PROPOSITION 4.9. *Let $G = (\Gamma,\alpha)$ be an elastic ribbon graph with trivalent vertices, and let $m = \min\{\,\alpha(e) \mid e \in \mathrm{Edge}(\Gamma)\,\}$ be the lowest weight of any edge in G. Then, for $\varepsilon < m/2$ and c any simple multi-curve on NG, we have*

$$\varepsilon\,\mathrm{EL}_{N_\varepsilon G}[c] < \mathrm{EL}_G[c]\cdot(1 + 8\varepsilon/m).$$

---

[2] Conceptually, in the standard definition of extremal length with families of paths, there are two ways to combine path families. You can take the *union* of the two families, as in the definition of $\mathrm{EL}_\perp(\sqcup A_i)$; this decreases the extremal length. Alternately, you can *join* the two families, where a valid path consists of one from each family, as implicitly happens in $\sum \mathrm{EL}(A_i)$.

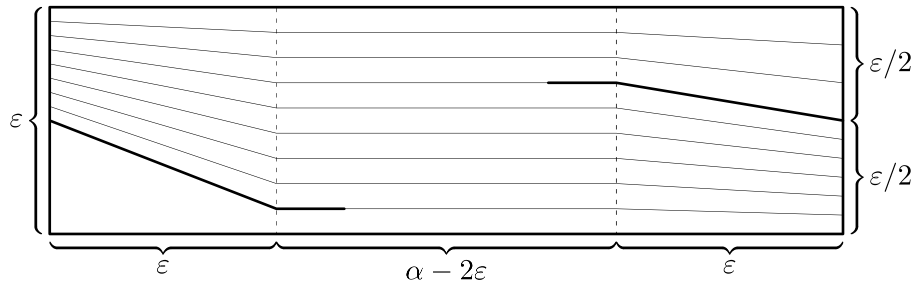

Figure 5. The annuli for Proposition 4.9 inside the rectangle corresponding to an edge of elastic weight α. Each portion of an annulus lies within a strip bounded by horizontal segments in the middle and diagonal segments near the ends.

Proof. We use equation (4.7) to estimate $\mathrm{EL}_{N_\varepsilon G}[c]$ from above.

First find suitable embedded annuli. We have $n_c(e)$ sections of annuli running over $N_\varepsilon e$, which we divide into pieces corresponding to these different annuli. Divide up the central portion of $N_\varepsilon e$ into n horizontal strips of equal height ε/n. Inside an $\varepsilon \times \varepsilon$ square near each end, make adjustments so the annuli will glue together well at the vertices. (These squares do not overlap since $\varepsilon < m/2$.) Specifically, near one end of e, n_1 of the annulus sections will continue to the left-hand neighbor of e at the corresponding vertex and n_2 will continue to the right-hand neighbor, with $n_1 + n_2 = n_c(e)$. Divide the interval $[0, \varepsilon/2]$ into n_1 equal sections, divide the interval $[\varepsilon/2, \varepsilon]$ into n_2 equal sections, and connect the corresponding endpoints by diagonal lines. Do the same near the other end of e. Let A be the resulting union of conformal annuli, as shown in Figure 5.

To give an upper bound on the total extremal length of annuli in A, we will give a lower bound on $\mathrm{EL}_\perp(A)$. We do this by considering the restriction to the annuli of a suitable test metric ρg, where g the standard piecewise-Euclidean metric from the proof of Proposition 4.8.

With this setup, take ρ to be $n_c(e)$ on the central section of $N_\varepsilon e$ and $\sqrt{5}n_c(e)$ on the squares at the ends of $N_\varepsilon e$. In the standard metric g, the vertical height of the annuli is ε/n_c in the center section and at least $\varepsilon/2n_c$ in the end squares. In the metric ρg, the vertical height is at least ε in the center section and $\varepsilon \cdot \sqrt{5}/2$ in the squares. In the squares, since the edges of the annuli are sloped, the actual distance $\ell_\perp(A)$ between the boundaries may be less than the vertical height; but since the slope of the edges dividing different pieces of A is in $[-1/2, 1/2]$, we have

$$\ell_{\rho g,\perp}(A) \geqslant \varepsilon \cdot \sqrt{5}/2 \cdot \cos(\tan^{-1}(1/2)) = \varepsilon.$$

Thus we have

$$\begin{aligned}
\operatorname{Area}(\rho g) &= \sum_{e\in \operatorname{Edge}(\Gamma)} \left[n_c(e)^2(\alpha(e)-2\varepsilon)\varepsilon + 2\cdot 5n_c(e)^2\varepsilon^2\right]\\
&\leqslant \varepsilon \operatorname{EL}_G[c]\cdot(1+8\varepsilon/m)\\
\operatorname{EL}_\perp(A) &\geqslant \frac{\ell_{\rho g,\perp}(A)^2}{\operatorname{Area}(\rho g)} \geqslant \frac{\varepsilon}{\operatorname{EL}_G[c]\cdot(1+8\varepsilon/m)}\\
\varepsilon \operatorname{EL}_{N_\varepsilon G}[c] &\leqslant \varepsilon/\operatorname{EL}_\perp(A) \leqslant \operatorname{EL}_G[c]\cdot(1+8\varepsilon/m).
\end{aligned}$$

The test function ρ is never continuous, so there is a strict inequality. □

The constant in Proposition 4.9 depends only on the local geometry of G and is thus unchanged under covers.

Remark 4.10. The restriction to trivalent graphs in Proposition 4.9 can presumably be removed. Since every graph is homotopy-equivalent to a trivalent graph, it is not necessary for our applications.

We have to do a little more work to deduce Theorem 2 from Propositions 4.8 and 4.9: stretch factor for graphs is defined with respect to *all* multi-curves, while for surfaces we restrict attention to *simple* multi-curves. We must check that the difference between two notions of stretch factor does not matter.

Definition 4.11. Let S_1 and S_2 be Riemann surfaces. For $\varphi\colon S_1 \hookrightarrow S_2$ a topological embedding, the *simple stretch factor* $\operatorname{SF}_{\mathrm{simp}}[\varphi]$ is the stretch factor from Definition 1.7. For $\varphi\colon S_1 \to S_2$ a continuous map, the *general stretch factor* is

$$\operatorname{SF}_{\mathrm{gen}}[\varphi] \coloneqq \sup_{[c]\colon C\to S_1} \frac{\operatorname{EL}[\varphi\circ c]}{\operatorname{EL}[c]},$$

where the supremum runs over *all* multi-curves on S_1 (not necessarily simple).

Let G_1 and G_2 be elastic graphs. For $\varphi\colon G_1 \to G_2$ a map between them, the *general stretch factor* $\operatorname{SF}_{\mathrm{gen}}[\varphi]$ is the stretch factor from Definition 1.10.

Now suppose $\varphi\colon G_1 \to G_2$ is a ribbon map between ribbon graphs. A *simple multi-curve* on G_1 is a multi-curve $c\colon C \to G_1$ that lifts to a simple multi-curve on NG_1 (i.e., so that c is a ribbon map). Then the *simple stretch factor* is

$$\operatorname{SF}_{\mathrm{simp}}[\varphi] \coloneqq \sup_{c \text{ simple on } G_1} \frac{\operatorname{EL}[\varphi\circ c]}{\operatorname{EL}[c]},$$

where the supremum runs over all homotopy classes of simple multi-curves $c\colon C \to G_1$ on G_1. Observe that if c is a simple multi-curve, then $\varphi\circ c$ is also a simple multi-curve, since $N\varphi$ is an embedding.

Proposition 4.12. *Let $\varphi\colon G_1 \to G_2$ be a ribbon map between elastic ribbon graphs. Then*

$$\mathrm{SF}_{\mathrm{gen}}[\varphi] = \mathrm{SF}_{\mathrm{simp}}[\varphi] = \mathrm{Emb}[\varphi].$$

To prove Proposition 4.12, we use train tracks.

Definition 4.13 ([Thu19, Def. 3.14]). A *train track* T is a graph in which the edges incident to each vertex are partitioned into equivalence classes, called *gates*, with at least two gates at each vertex. A *train-track (multi-)curve* on T is a (multi-)curve that enters and leaves by different gates each time it passes through a vertex. A *weighted train track* is a train track T with a weight $w(e)$ for each edge e of T, satisfying a triangle inequality at each gate g of each vertex v:

$$w(g) \leqslant \sum_{\substack{g' \text{ gate at } v \\ g' \neq g}} w(g'), \tag{4.14}$$

where $w(g)$ is the sum of the weights of all edges in g. (If there are only two gates at v, this inequality is necessarily an equality.)

Lemma 4.15. *Let (T, w) be a weighted train track with a ribbon structure. Then there is a sequence of simple train-track multi-curves (C_i, c_i) and positive weighting factors k_i so that $k_i c_i$ approximates w, in the sense that*

$$\lim_{i\to\infty} k_i n_{c_i} = w.$$

This lemma is close to standard facts in the theory of train tracks. Note there is no assumption that the train track structure and the ribbon structure are compatible. (This avoids assuming that the optimizer for $\mathrm{Emb}[\varphi]$ is a ribbon map.) Compare Lemma 4.15 to [Thu19, Prop. 6.13], which gives the exact weights (without approximation), but yields multi-curves that are not simple.

Proof. We first prove that if w is integer-valued and has even total weight at each vertex, then there is a simple train-track multi-curve (C, c) so that $n_c = w$.

On each edge e of T, take $w(e)$ parallel strands on Ne. We must show how to stitch together these strands at the vertices without crossing strands or making illegal train-track turns. Focus on a vertex v. If one of the incident edges has zero weight, delete it. If one of the train-track triangle inequalities is an equality, smooth the vertex (in the sense of [Thu19, Def. 3.14]) so that there are only two gates at v. After this, if there are at least three gates at v, then all inequalities are strict and equation (4.14) is strengthened to

$$w(g) \leqslant -2 + \sum_{\substack{g' \text{ gate at } v \\ g' \neq g}} w(g') \tag{4.16}$$

by the parity condition.

In either the two-gate or more-gate case, find any two edges at v that are adjacent in the ribbon structure and belong to different gates. Join adjacent outermost strands from these two edges. We are left with a smaller problem, where the weights on these two strands are reduced by 1. The train-track inequalities are still satisfied, using equation (4.16) when there are three or more gates. By induction we can join up all the strands to a multi-curve C. Because we always join strands that are adjacent among the remaining strands in the ribbon structure, C is simple.

For general weights, find a sequence of even integer weights w_i on T and factors k_i so that $\lim_{i\to\infty} k_i w_i = w$ and the w_i satisfy the train-track inequalities for T. The above argument gives a simple multi-curve (C_i, c_i) for each w_i, as desired. ☐

Proof of Proposition 4.12. We already know that $\mathrm{Emb}[\varphi] = \mathrm{SF}_{\mathrm{gen}}[\varphi]$. From the definition, it is clear that $\mathrm{SF}_{\mathrm{simp}}[\varphi] \leqslant \mathrm{SF}_{\mathrm{gen}}[\varphi]$. It remains to prove that $\mathrm{Emb}[\varphi] \leqslant \mathrm{SF}_{\mathrm{simp}}[\varphi]$. By [Thu19, Prop. 6.13], there is a weighted train-track T that fits into a tight sequence

$$T \xrightarrow{t} G_1 \xrightarrow{\psi} G_2,$$

where t is the inclusion of a subgraph, $\psi \in [\varphi]$ and $\psi \circ t$ is a train-track map. ("Tight" means that the energies are multiplicative, which in this case means that $\mathrm{EL}[\psi \circ t] = \mathrm{EL}[t]\,\mathrm{Emb}[\psi]$, and furthermore all three maps t, ψ, and $\psi \circ t$ are minimizers in their homotopy classes.) T inherits a ribbon structure as a subgraph of G_1. By Lemma 4.15, we can find a sequence of simple multi-curves (c_i, C_i) on T so that

$$C_i \xrightarrow{c_i} T \xrightarrow{t} G_1 \xrightarrow{\psi} G_2$$

approaches a tight sequence, in the sense that $t \circ c_i$ and $\psi \circ t \circ c_i$ are both reduced and

$$\lim_{i\to\infty} \frac{\mathrm{EL}[\psi \circ t \circ c_i]}{\mathrm{EL}[t \circ c_i]} = \frac{\mathrm{EL}[\psi \circ t]}{\mathrm{EL}[t]} = \mathrm{Emb}[\varphi].$$

Since t is an inclusion, the sequence of weighted multi-curves $t \circ c_i$ is simple, as desired. ☐

Proof of Theorem 2. Immediate from Propositions 4.8, 4.9, and 4.12. ☐

Question 4.17. For $\varphi\colon S_1 \hookrightarrow S_2$ a topological embedding of Riemann surfaces, how does $\mathrm{SF}_{\mathrm{gen}}[\varphi]$ behave? By considering quadratic differentials, it is

$$\begin{tikzcd}
 & & & & X_n \arrow[dlll, dashed] \arrow[dl, dashed] \arrow[dr, dashed] \arrow[drrr, dashed] \arrow[ddllll, dashed, bend right, "\varphi^n"'] \arrow[ddrrrr, dashed, bend left, "\pi^n"] & & & & \\
 & X_1 \arrow[dl, "\varphi"'] \arrow[dr, "\pi"] & & X_1 \arrow[dl, "\varphi"'] \arrow[dr, dash] & \cdots & X_1 \arrow[dl, dash] \arrow[dr, "\pi"] & & X_1 \arrow[dl, "\varphi"'] \arrow[dr, "\pi"] & \\
X_0 & & X_0 & & \cdots & & X_0 & & X_0
\end{tikzcd}$$

$n+1$ copies of X_0, n copies of X_1

Figure 6. The pullback construction of X_n.

not hard to see that if φ is not homotopic to an annular conformal embedding, then

$$\mathrm{SF}_{\mathrm{simp}}[\varphi] = \mathrm{SF}_{\mathrm{gen}}[\varphi] \geqslant 1.$$

On the other hand, if φ is homotopic to a conformal embedding, then

$$\mathrm{SF}_{\mathrm{simp}}[\varphi] \leqslant \mathrm{SF}_{\mathrm{gen}}[\varphi] \leqslant 1.$$

But this leaves many questions open.

5. Iteration and asymptotic stretch factor

5.1. *General theory.* To complete the proof of Theorem 1, we turn to the behavior of energies under iteration. Recall first that if $\varphi\colon X_1 \to X_0$ is any continuous map and $\pi_0\colon Y_0 \to X_0$ is a covering map, we can form the pull-back

$$\begin{tikzcd}
Y_1 \arrow[r, dashed, "\widetilde{\varphi}"] \arrow[d, dashed, "\pi_1"] & Y_0 \arrow[d, "\pi_0"] \\
X_1 \arrow[r, "\varphi"] & X_0.
\end{tikzcd}$$

Then π_1 is also a covering map. In this setting, we call $\widetilde{\varphi}$ a *cover* of φ.

Definition 5.1. A *topological correspondence* is a pair of topological spaces X_1 and X_0 and a pair of maps between them: $\pi, \varphi\colon X_1 \rightrightarrows X_0$.

For $n \geqslant 0$, the n'th *orbit space* X_n of a topological correspondence is the n-fold product of X_1 with itself over X_0 using π and φ, i.e., the pull-back X_n in Figure 6. Concretely, X_n is the set of tuples

$$(y_0, x_1, y_1, x_2, y_2, \ldots, y_{n-1}, x_n, y_n) \in X_0 \times (X_1 \times X_0)^n$$

so that $\varphi(x_i) = y_{i-1}$ and $\pi(x_i) = y_i$ for $1 \leqslant i \leqslant n$, with the subspace topology. Of the natural maps from X_n to X_0, we distinguish

$$\begin{aligned}\varphi^n(y_0, x_1, \dots, x_n, y_n) &= y_0,\\ \pi^n(y_0, x_1, \dots, x_n, y_n) &= y_n.\end{aligned}$$

The corresponding *iterate* of (π, φ) is the virtual endomorphism $\pi^n, \varphi^n\colon X_n \rightrightarrows X_0$. If (π, φ) is a virtual endomorphism, so is (π^n, φ^n). We also distinguish intermediate maps $\pi^n_k, \varphi^n_k\colon X_n \rightrightarrows X_k$ for $k \leqslant n$ by

$$\begin{aligned}\varphi^n_k(y_0, x_1, \dots, x_n, y_n) &= (y_0, x_1, \dots, x_k, y_k)\\ \pi^n_k(y_0, x_1, \dots, x_n, y_n) &= (y_{n-k}, x_{n-k+1}, \dots, x_n, y_n).\end{aligned}$$

(The convention is that superscripts refer to the domain and subscripts refer to the range. This fits well with seeing φ^n as an "iterate" of φ.)

If we have two topological correspondences $\pi_X, \varphi_X\colon X_1 \rightrightarrows X_0$ and $\pi_Y, \varphi_Y\colon Y_1 \rightrightarrows Y_0$ and a morphism (f_0, f_1) from (π_X, φ_X) to (π_Y, φ_Y) (in the sense of Definition 2.2), then we can also use the pull-back property to iterate the morphism, getting a map $f_n\colon X_n \to Y_n$. Concretely,

$$f_n(y_0, x_1, y_1, \dots, x_n, y_n) = (f_0(y_0), f_1(x_1), f_0(y_1), \dots, f_1(x_n), f_0(y_n)).$$

If φ is injective (as for surface virtual endomorphisms from branched self-covers), $\pi \circ \varphi^{-1}$ is a partially-defined map on X_0. Then φ^n is also injective and $\pi^n \circ (\varphi^n)^{-1}$, where defined, is the n-fold iterate $(\pi \circ \varphi^{-1})^{\circ n}$. This justifies the term "iteration."

Definition 5.2. Consider a category of spaces with a structure that can be lifted to covers (like an elastic structure on graphs or a conformal structure on surfaces). Suppose that we have a non-negative *energy* E defined for suitable maps $\varphi\colon X \to Y$ that is *sub-multiplicative*, in the sense that

$$E(\psi \circ \varphi) \leqslant E(\psi)E(\varphi), \tag{5.3}$$

and *invariant under covers*, in the sense that if $\widetilde{\varphi}\colon \widetilde{X} \to \widetilde{Y}$ is a cover of φ, then

$$E(\widetilde{\varphi}) = E(\varphi). \tag{5.4}$$

Then for $\pi, \varphi\colon X_1 \rightrightarrows X_0$ a virtual endomorphism between such spaces, where the structure on X_1 is lifted from the structure on X_0 via π, the *asymptotic energy* is

$$\overline{E}(\pi, \varphi) := \lim_{n\to\infty} E(\varphi^n)^{1/n}. \tag{5.5}$$

Proposition 5.6. *Let $\pi, \varphi\colon X_1 \rightrightarrows X_0$ be a virtual endomorphism and let E be an energy that is sub-multiplicative and invariant under covers. Then the limit defining the asymptotic energy converges and is equal to the infimum of the terms. In particular, $\overline{E}(\pi, \varphi) \leqslant E(\varphi)$.*

Proof. Note that $\varphi^n = \varphi^k \circ \varphi^n_k$. An examination of the diagrams reveals that φ^n_k is a cover of φ^{n-k}, so $E(\varphi^n_k) = E(\varphi^{n-k})$. We therefore have $E(\varphi^n) \leqslant E(\varphi^k)E(\varphi^{n-k})$ by equation (5.3). Then the sequence $\log(E(\varphi^n))$ is sub-additive, and Fekete's Lemma gives the result.[3] ☐

If an energy $E(\varphi)$ is invariant under homotopy of φ, then we will write it as $E[\varphi]$.

Proposition 5.7. *Let $\pi, \varphi\colon X_1 \rightrightarrows X_0$ be a virtual endomorphism, and let E be an energy that is sub-multiplicative, invariant under covers, and invariant under homotopy. Then $\overline{E}(\pi, \varphi)$ is invariant under homotopy equivalence of (π, φ).*

Proof. Let $\sigma, \psi\colon Y_1 \rightrightarrows Y_0$ be a virtual endomorphism homotopy equivalent to (π, φ), with homotopy equivalences given by $f_i\colon X_i \to Y_i$ and $g_i\colon Y_i \to X_i$ for $i = 0, 1$. We need to compare $E[\varphi^n]$ and $E[\psi^n]$. Let $f_n\colon X_n \to Y_n$ be the iterate of (f_0, f_1). Since f_n is a cover of f_0, we have $E[f_n] = E[f_0]$. Furthermore, φ^n is homotopic to $g_0 \circ \psi^n \circ f_n$. Then

$$\begin{aligned} E[\varphi^n] &\leqslant E[g_0]E[\psi^n]E[f_n] = E[\psi^n]\big(E[f_0]E[g_0]\big), \\ E[\varphi^n]^{1/n} &\leqslant E[\psi^n]^{1/n}\big(E[f_0]E[g_0]\big)^{1/n}. \end{aligned}$$

Passing to the limit on both sides, we have

$$\overline{E}(\pi, \varphi) \leqslant \overline{E}(\sigma, \psi).$$

By the same reasoning in the other direction, $\overline{E}(\sigma, \psi) \leqslant \overline{E}(\pi, \varphi)$. ☐

If E is sub-multiplicative, invariant under covers, and invariant under homotopy, we will write $\overline{E}[\pi, \varphi]$ to indicate that the asymptotic energy is independent of homotopy equivalence.

5.2. *Specific energies.* We now turn to the specific energies of interest on elastic graphs or conformal surfaces. There are three relevant energies:

- the stretch factor SF for conformal surfaces;
- the stretch factor SF for elastic graphs; and
- the embedding energy Emb for elastic graphs.

The last two are equal by Theorem 1.14, although we sometimes distinguish when we need to use that theorem. All three are invariant under homotopy by definition.

For φ a map between elastic graphs, the embedding energy $\operatorname{Emb}[\varphi]$ is sub-multiplicative by [Thu19, Prop. 2.15].

---

[3] Fekete's Lemma: If $(a_n)_{n=1}^{\infty}$ is sub-additive, then $\lim_{n\to\infty} a_n/n$ exists and is equal to the infimum of the terms.

LEMMA 5.8. *Let* $\varphi\colon X \to Y$ *and* $\psi\colon Y \to Z$ *be either topological embeddings of conformal surfaces or maps between elastic graphs. Then stretch factor is sub-multiplicative*:

$$\mathrm{SF}[\psi \circ \varphi] \leqslant \mathrm{SF}[\varphi]\,\mathrm{SF}[\psi].$$

Proof. In either case, the stretch factor is a supremum over multi-curves (simple multi-curves for maps between surfaces). For c any suitable multi-curve on X, if $\mathrm{EL}_Y[\varphi \circ c] \neq 0$, we have

$$\frac{\mathrm{EL}_Z[\psi \circ \varphi \circ c]}{\mathrm{EL}_X[c]} = \frac{\mathrm{EL}_Z[\psi \circ \varphi \circ c]}{\mathrm{EL}_Y[\varphi \circ c]}\frac{\mathrm{EL}_Y[\varphi \circ c]}{\mathrm{EL}_X[c]} \leqslant \mathrm{SF}[\psi]\,\mathrm{SF}[\varphi].$$

If $\mathrm{EL}_Y[\varphi \circ c] = 0$, then also $\mathrm{EL}_Z[\psi \circ \varphi \circ c] = 0$ and we get the same inequality. Since $\mathrm{SF}[\psi \circ \varphi]$ is the supremum of the left-hand side over all c, we get the desired result. ☐

PROPOSITION 5.9. *Stretch factor for maps between elastic graphs is invariant under covers.*

Proof. Let $\varphi\colon G \to H$ be a map between elastic graphs, and let $\widetilde{\varphi}\colon \widetilde{G} \to \widetilde{H}$ be a cover of φ of degree d. First note that we can pull-back a multi-curve (C, c) on G to a multi-curve $(\widetilde{C}, \widetilde{c})$ on $\widetilde{G}$, with $\mathrm{EL}[\widetilde{c}] = d\,\mathrm{EL}[c]$ and $\mathrm{EL}[\widetilde{\varphi} \circ \widetilde{c}] = d\,\mathrm{EL}[\varphi \circ c]$. It follows that $\mathrm{SF}[\widetilde{\varphi}] \geqslant \mathrm{SF}[\varphi]$.

For the other inequality, we use Theorem 1.14. Let $\psi \in [\varphi]$ be a map with $\mathrm{Emb}(\psi) = \mathrm{Emb}[\varphi]$, and let $\widetilde{\psi}$ be the corresponding lift. Then $\mathrm{Emb}[\widetilde{\varphi}] \leqslant \mathrm{Emb}(\widetilde{\psi}) = \mathrm{Emb}(\psi) = \mathrm{Emb}[\varphi]$. ☐

For graphs, embedding energy/stretch factor fits nicely into the general theory laid out in Section 5.1. For surfaces, SF is not invariant under covers [KPT15, Ex. 6.6]. We therefore modify the definition.

Definition 5.10. For $\varphi\colon R \hookrightarrow S$ a topological embedding of conformal surfaces, the *lifted stretch factor* $\widetilde{\mathrm{SF}}[\varphi]$ is

$$\widetilde{\mathrm{SF}}[\varphi] \coloneqq \sup_{\substack{\widetilde{\varphi}\text{ finite}\\ \text{cover of }\varphi}} \mathrm{SF}[\widetilde{\varphi}].$$

THEOREM 5.11 (Kahn-Pilgrim-Thurston [KPT15, Th. 3]). *Let* $\varphi\colon R \hookrightarrow S$ *be a topological embedding of Riemann surfaces. If* $\mathrm{SF}[\varphi] \geqslant 1$, *then* $\widetilde{\mathrm{SF}}[\varphi] = \mathrm{SF}[\varphi]$. *If* $\mathrm{SF}[\varphi] < 1$, *then*

$$\mathrm{SF}[\varphi] \leqslant \widetilde{\mathrm{SF}}[\varphi] < 1.$$

LEMMA 5.12. $\widetilde{\mathrm{SF}}$ *is sub-multiplicative.*

Proof. Any cover of a composition $\varphi \circ \psi$ factors as a composition $\widetilde{\varphi} \circ \widetilde{\psi}$ of covers of φ and ψ. ☐

Lemma 5.13. $\widetilde{\mathrm{SF}}$ *is invariant under covers.*

Proof. Any two finite covers of a map have a common finite cover. □

Let us recap what we have so far. For $\pi_G, \varphi_G \colon G_1 \rightrightarrows G_0$ a virtual endomorphism of elastic graphs, we have an asymptotic energy $\overline{\mathrm{SF}}[\pi_G, \varphi_G]$, invariant under homotopy equivalence. In particular, $\overline{\mathrm{SF}}$ is independent of the elastic structure on G_0, since the identity is a homotopy equivalence..

For $\pi_S, \varphi_S \colon S_1 \rightrightarrows S_0$ a virtual endomorphism of conformal surfaces, we have an asymptotic energy $\overline{\widetilde{\mathrm{SF}}}[\pi_S, \varphi_S]$, which we will also write $\overline{\mathrm{SF}}[\pi_S, \varphi_S]$. (See Corollary 5.15 below.) This is invariant under quasi-conformal homotopy equivalences and therefore is independent of the conformal structure, as long as we do not change a puncture to a boundary component or vice versa.

In particular, if $\pi, \varphi \colon G_1 \rightrightarrows G_0$ is a ribbon virtual endomorphism of elastic ribbon graphs, then the asymptotic energy of the induced virtual endomorphism $N_\varepsilon\pi, N_\varepsilon\varphi \colon N_\varepsilon G_1 \rightrightarrows N_\varepsilon G_0$ is independent of ε. Thus we will drop ε from the notation.

Proposition 5.14. *Let* $\pi, \varphi \colon \Gamma_1 \rightrightarrows \Gamma_0$ *be a ribbon graph virtual endomorphism. Then*

$$\overline{\mathrm{SF}}[\pi, \varphi] = \overline{\mathrm{SF}}[N\pi, N\varphi] = \lim_{n\to\infty} \sqrt[n]{\mathrm{SF}[N\varphi^n]}.$$

Proof. By Proposition 5.7, we can replace the virtual endomorphism with a ribbon homotopy equivalent one without changing $\overline{\mathrm{SF}}$. So we may assume that G_0 and G_1 are trivalent, with some minimum elasticity m on any edge. Pick $\varepsilon < m/2$.

Theorem 2 says that $\mathrm{SF}[N_\varepsilon\varphi^n]$ is within a factor of $1 + 8\varepsilon/m$ of $\mathrm{SF}[\varphi^n]$ for all n. Similarly, since SF for graphs is invariant under covers and the estimates depend only on the local geometry, $\widetilde{\mathrm{SF}}[N_\varepsilon\varphi^n]$ is within a factor of $1 + 8\varepsilon/m$ of $\widetilde{\mathrm{SF}}[\varphi^n]$. When we take the n'th root in limit for the three terms in the statement, this factor disappears, as in Proposition 5.7. □

Corollary 5.15. *For any virtual surface endomorphism* $\pi, \varphi \colon S_1 \rightrightarrows S_0$ *where* S_0 *and* S_1 *have no punctures,*

$$\overline{\mathrm{SF}}[\pi, \varphi] = \lim_{n\to\infty} \sqrt[n]{\mathrm{SF}[\varphi^n]}.$$

Proof. If S_0 and S_1 are closed surfaces, then $\mathrm{SF}[\varphi] \geqslant 1$ (since there is never an annular conformal embedding) so SF is invariant under covers and the statement is trivial.

Otherwise, we will see that $\lim \sqrt[n]{\mathrm{SF}[\varphi^n]}$ is independent of the conformal structure on S_0 (since the general theory of asymptotic energies does not apply), and then replace S_0 by an ε-thickening of a graph and apply Proposition 5.14.

For that purpose, let $\pi, \varphi\colon S_1 \rightrightarrows S_0$ and $\pi', \varphi'\colon S_1' \rightrightarrows S_0'$ be two virtual endomorphisms of conformal surfaces, and let $f_i\colon S_i \to S_i'$ be one direction of a homotopy equivalence between them. If $\mathrm{SF}[f_0] \geqslant 1$, then $\mathrm{SF}[f_n] = \mathrm{SF}[f_0]$ for all n by the first part of Theorem 5.11. If $\mathrm{SF}[f_0] < 1$, then $\mathrm{SF}[f_n] < 1$ by the second part of Theorem 5.11. Either way, we have the inequalities we need to deduce that $\mathrm{SF}[\varphi^n]$ is within a uniform constant factor of $\mathrm{SF}[\varphi'^n]$, as in the proof of Proposition 5.7. ☐

5.3. *Proof of Theorem 1.* We are now ready to prove Theorem 1. We expand the statement to include asymptotic energies.

THEOREM 1′. *Let $f\colon (\Sigma, P) \circlearrowleft$ be a branched self-cover of hyperbolic type with associated surface virtual endomorphism (π_S, φ_S). Then the following conditions are equivalent:*

1. *f is equivalent to a rational map;*
2. *there is an elastic graph spine G for $\Sigma \backslash P$ and an integer $n > 0$ so that* $\mathrm{Emb}[\varphi_G^n] < 1$;
3. *for every elastic graph spine G for $\Sigma \backslash P$ and for every sufficiently large n (depending on f and G), we have* $\mathrm{Emb}[\varphi_G^n] < 1$;
4. $\overline{\mathrm{SF}}[\pi_S, \varphi_S] < 1$; *and*
5. $\overline{\mathrm{SF}}[\pi_G, \varphi_G] < 1$.

Proof. Conditions (4) and (5) are equivalent by Proposition 5.14. Conditions (2) and (3) are equivalent to Condition (5) by Proposition 5.6 and Theorem 1.14.

To see that Condition (1) implies Condition (4), suppose f is equivalent to a rational map. Then by Theorem 3.3, there is a conformal virtual endomorphism $\pi_S, \varphi_S\colon S_1 \rightrightarrows S_0$ compatible with f. Since φ_S is annular, by Theorem 5.11, $\widetilde{\mathrm{SF}}[\varphi_S] < 1$, so by Proposition 5.6, $\overline{\mathrm{SF}}[\pi_S, \varphi_S] < 1$.

Conversely, to see Condition (4) implies Condition (1), suppose $\overline{\mathrm{SF}}[\pi_S, \varphi_S] < 1$ with respect to any conformal structure S_0. By Proposition 5.6, there is some n so that $\mathrm{SF}[\varphi_S^n] \leqslant \widetilde{\mathrm{SF}}[\varphi_S^n] < 1$, so by Theorem 1.9, φ_S^n is homotopic to an annular conformal embedding. Then by Theorem 3.3, the n-fold composition $f^{\circ n}$ is equivalent to a rational map, which implies that f itself is equivalent to a rational map. ☐

Remark 5.16. The last step in the proof follows from W. Thurston's Obstruction Theorem (Theorem 7.4 below), since an obstruction for f is also an obstruction for $f^{\circ n}$, but does not use the full strength of that theorem. It suffices, for instance, to know that some power of the pull-back map on Teichmüller space is contracting [BCT14, §2.5].

6. Asymptotics of other energies

The theory of asymptotic energies developed at the beginning of Section 5 applies to any energy that is sub-multiplicative and invariant under homotopy and covers. In particular, it applies to any of the p-conformal energies E_p^p defined in [Thu19, App. A]. These energies $E_p^p[\varphi]$ are a simultaneous generalization of best Lipschitz constant $\operatorname{Lip}[\varphi] = E_\infty^\infty[\varphi]$, and the embedding energy $\operatorname{Emb}[\varphi] = \bigl(E_2^2[\varphi]\bigr)^2$.

Recall that a p-conformal graph, for $1 < p \leqslant \infty$, is a graph with a p-length $\alpha(e)$ on each edge e, which we will treat as a metric. A 1-conformal graph is instead a weighted graph, with a weight $w(e) > 0$ on each edge. For $\varphi\colon G_1 \to G_2$ a PL map between p-conformal graphs, $E_p^p[\varphi]$ is defined by

$$\begin{aligned} E_p^p(\varphi) &:= \begin{cases} \operatorname*{ess\,sup}_{y\in G_2} \dfrac{n_{\varphi,w}(y)}{w(y)} & p=1,\\ \operatorname*{ess\,sup}_{y\in G_2} \Biggl(\displaystyle\sum_{x\in\varphi^{-1}(y)} \lvert\varphi'(x)\rvert^{p-1}\Biggr)^{1/p} & 1<p<\infty,\\ \operatorname{Lip}(\varphi) & p=\infty, \end{cases} \\ E_p^p[\varphi] &:= \inf_{\psi\in[\varphi]} E_p^p(\psi). \end{aligned} \tag{6.1}$$

(In the $p = 1$ case, we count weighted preimages: $n_{\varphi,w}(y) = \sum_{x\in\varphi^{-1}(y)} w(x)$.) Like Emb, the energy E_p^p is sub-multiplicative and invariant under covers, whether or not we take homotopy classes. For a graph virtual endomorphism $\pi, \varphi\colon \Gamma_1 \rightrightarrows \Gamma_0$, we thus have asymptotic energies $\overline{E}_p^p[\pi,\varphi]$. By Proposition 5.7, $\overline{E}_p^p$ is invariant under homotopy equivalence.

$E_p^p[\varphi]$ can be characterized as a stretch factor for energies of maps to length graphs [Thu19, Th. 6]. For $1 \leqslant p \leqslant \infty$ and $f\colon G \to K$ a PL map from a p-conformal graph to a length graph, there are energies

$$\begin{aligned} E_\infty^p(f) &:= \lVert f'\rVert_p, \\ E_\infty^p[f] &:= \inf_{g\in[f]} E_\infty^p(g). \end{aligned} \tag{6.2}$$

Then, for $\varphi\colon G_1 \to G_2$ a PL map between p-conformal graphs,

$$E_p^p[\varphi] = \sup_{[f]\colon G_2\to K} \frac{E_\infty^p[f\circ\varphi]}{E_\infty^p[f]}, \tag{6.3}$$

where the supremum runs over all length graphs K and homotopy classes $[f]$ of maps.

To study the behavior of $\overline{E}_p^p$ as p varies, we compare p-conformal energies and q-conformal energies for $p \leqslant q$.

Definition 6.4. For a metric graph G, let

$$m(G) \coloneqq \min_{e \in \mathrm{Edge}(G)} \alpha(e),$$

$$M(G) \coloneqq \sum_{e \in \mathrm{Edge}(G)} \alpha(e).$$

Lemma 6.5. *For $1 \leqslant p \leqslant q \leqslant \infty$ and $f\colon G \to K$ a constant-derivative map from a metric graph to a length graph,*

$$E_\infty^q(f) \leqslant m(G)^{-\frac{1}{p}+\frac{1}{q}} E_\infty^p(f).$$

The same inequality is true when we minimize over the homotopy class.

Proof. In general, there is an inequality

$$\Bigl(\sum_i {x_i}^q\Bigr)^{1/q} \leqslant \Bigl(\sum_i {x_i}^p\Bigr)^{1/p}. \tag{6.6}$$

With positive weights w_i with $m = \min_i w_i$, this becomes

$$\Bigl(\sum_i w_i {x_i}^q\Bigr)^{1/q} \leqslant m^{-\frac{1}{p}+\frac{1}{q}} \cdot \Bigl(\sum_i w_i {x_i}^p\Bigr)^{1/p}. \tag{6.7}$$

(Apply equation (6.6) to the sequence ${w_i}^{1/q} x_i$.) Apply equation (6.7) to the definition of E_∞^p.

To get the statement for homotopy classes, apply this inequality to a map f that minimizes $E_\infty^p(f)$ in its homotopy class. □

Lemma 6.8. *For $1 \leqslant p \leqslant q \leqslant \infty$ and $f\colon G \to K$ a PL map from a metric graph to a length graph,*

$$E_\infty^p(f) \leqslant M(G)^{\frac{1}{p}-\frac{1}{q}} \cdot E_\infty^q(f).$$

The same inequality is true when we minimize over the homotopy class.

Proof. Use sub-multiplicativity of the energies [Thu19, Prop. A.12]:

$$E_\infty^p(f) \leqslant E_q^p(\mathrm{id}) E_\infty^q(f).$$

From the definition of E_q^p [Thu19, eq. (A.6)], we see that $E_q^p(\mathrm{id}) = M(G)^{\frac{q-p}{pq}}$. (Alternatively, apply Hölder's inequality to equation (6.2).)

To get the statement for homotopy classes, apply this inequality to a map f that minimizes $E_\infty^q(f)$ in its homotopy class. □

We can now see that these energies give nothing new for ordinary (non-virtual) graph endomorphisms (i.e., outer automorphisms of the free group). Define the asymptotic energy of an ordinary endomorphism φ to be $\overline{E}[\varphi] \coloneqq \overline{E}[\mathrm{id}, \varphi] = \lim_{n\to\infty} \sqrt[n]{E[\varphi^{\circ n}]}$.

Proposition 6.9. *For $[\varphi]\colon \Gamma \to \Gamma$ an endomorphism of a graph,*

$$\overline{E}^p_p[\varphi] = \overline{\mathrm{Lip}}[\varphi].$$

Proof. By Lemmas 6.5 and 6.8 there is a constant $C \geqslant 1$ so that

$$\frac{1}{C}\frac{\mathrm{Lip}[f \circ \varphi^{\circ n}]}{\mathrm{Lip}[f]} \leqslant \frac{E^p_\infty[f \circ \varphi^{\circ n}]}{E^p_\infty[f]} \leqslant C\frac{\mathrm{Lip}[f \circ \varphi^{\circ n}]}{\mathrm{Lip}[f]}.$$

Equation (6.3) then shows that $E^p_p[\varphi^{\circ n}]$ is within a factor of C of $\mathrm{Lip}[\varphi^{\circ n}]$. The constant factor disappears in the limit defining $\overline{E}^p_p[\varphi]$. □

For virtual endomorphisms, the situation is more interesting.

Proposition 6.10. *For $\pi, \varphi\colon \Gamma_1 \rightrightarrows \Gamma_0$ a virtual endomorphism of graphs, $\overline{E}^p_p[\pi, \varphi]$ is a non-increasing function of p: if $1 \leqslant p \leqslant q \leqslant \infty$,*

$$\overline{E}^q_q[\pi, \varphi] \leqslant \overline{E}^p_p[\pi, \varphi].$$

Proof. Pick a metric structure G_0 on Γ_0 and lift it to get a series of metric graphs G_n as usual. Then for any $f\colon G_0 \to K$ a map to a length graph,

$$\frac{E^q_\infty[f \circ \varphi^n]}{E^q_\infty[f]} \leqslant \frac{1}{C}\frac{E^p_\infty[f \circ \varphi^n]}{E^p_\infty[f]}$$

for some constant C, since $m(G_n) = m(G_0)$. Now $E^q_q[\varphi^n] \leqslant E^p_p[\varphi^n]/C$, and the constant factor disappears in the limit as usual. □

Proposition 6.11. *For $\pi, \varphi\colon \Gamma_1 \rightrightarrows \Gamma_0$ a virtual endomorphism of graphs with π a covering of degree d, if $1 \leqslant p \leqslant q \leqslant \infty$,*

$$\overline{E}^q_q[\pi, \varphi] \geqslant d^{-\frac{1}{p}+\frac{1}{q}} \cdot \overline{E}^p_p[\pi, \varphi].$$

Proof. Pick a metric structure G_0 on Γ_0, and lift it as in Proposition 6.10. Observe that $M(G_n) = d^n M(G_0)$. Then by Lemmas 6.5 and 6.8, for any length graph K and $[f]\colon G_0 \to K$,

$$\frac{E^p_\infty[f \circ \varphi^n]}{E^p_\infty[f]} \leqslant \left(\frac{M(G_n)}{m(G_0)}\right)^{\frac{1}{p}-\frac{1}{q}} \frac{E^q_\infty[f \circ \varphi^n]}{E^q_\infty[f]} = C \cdot d^{n\left(\frac{1}{p}-\frac{1}{q}\right)} \cdot E^q_q[\varphi^n]$$

for some constant C. Then taking the supremum over K and $[f]$, taking the n'th root, and passing to the limit shows that $\overline{E}^p_p[\pi, \varphi] \leqslant d^{1/p-1/q} \cdot \overline{E}^q_q[\pi, \varphi]$. □

Corollary 6.12. *$\overline{E}^p_p[\pi, \varphi]$ is a continuous function of p.*

Question 6.13. What more can be said about $\overline{E}^p_p[\pi, \varphi]$ as p varies? For instance, an examination of equation (6.1) shows that for any map $\varphi\colon G_1 \to G_2$ between metric graphs and $1 \leqslant p \leqslant q \leqslant \infty$,

$$\left(E^q_q(\varphi)\right)^{q/(q-1)} \leqslant \left(E^p_p(\varphi)\right)^{p/(p-1)}$$

and so for a virtual endomorphism

$$, \left(\overline{E}_q^q[\pi,\varphi]\right)^{q/(q-1)} \leqslant \left(\overline{E}_p^p[\pi,\varphi]\right)^{p/(p-1)}. \tag{6.14}$$

When $\overline{E}_p^p[\pi,\varphi] < 1$, this is stronger than Proposition 6.10. Is there more? For instance, is $\overline{E}_p^p[\pi,\varphi]$ a convex function of p in some sense?

The cases $p = 1$, 2, or ∞ of the asymptotic energy are of particular interest.

- The most important case is $\overline{E}_\infty^\infty[\pi,\varphi]$, with $\overline{E}_\infty^\infty < 1$ if and only if $\pi,\varphi\colon G_1 \rightrightarrows G_0$ is a combinatorial model for an *expanding* dynamical system in the sense of Nekrashevych [Nek14].[4] This notion of expanding is quite important. In the expanding case, there is a well-defined (ordinary) dynamical system on an inverse limit Julia set, independent of the details of the combinatorial model. Furthermore, the iterated monodromy groups of an expanding dynamical system are well behaved [Nek05], having, for instance, solvable word problem, while still allowing for many interesting examples (e.g., groups of intermediate growth).
- For $p = 2$, Theorem 1 relates $\overline{E}_2^2[\pi,\varphi] < 1$ to rational maps.
- The other natural special case is $p = 1$. If the weights are all 1, $E_1^1[\varphi] < 1$ implies that φ is null-homotopic, so in non-trivial cases $\overline{E}_1^1[\pi,\varphi] \geqslant 1$. If $\overline{E}_\infty^\infty[\pi,\varphi] < 1$ (so there is a Julia set), it appears that $\overline{E}_1^1[\pi,\varphi] > 1$ when the Julia set has Sierpiński-carpet-like behavior, and that $\overline{E}_1^1[\pi,\varphi] = 1$ when the Julia set has many local cut points in the sense of Carrasco Piaggio [CP14].

In forthcoming joint work with Kevin Pilgrim [PT], we will show that, if p^* is the Ahlfors regular conformal dimension of the Julia set of a virtual endomorphism (π,φ), then

$$\overline{E}_{p^*}^{p^*}[\pi,\varphi] = 1.$$

Combined with the bounds on how $\overline{E}_p^p$ varies as a function of p, this allows us to give concrete bounds on the Ahlfors regular conformal dimension.

7. Obstructions

7.1. *Obstructions for rational maps.* In this section we relate Theorem 1 to W. Thurston's Obstruction Theorem. With an eye to generalizations, we rephrase it in terms of elastic multi-curves without the assumption of hyperbolic type.

Definition 7.1. An *elastic multi-curve* $A = (C,\alpha,c)$ on a surface Σ is a multi-curve $c\colon C \hookrightarrow \Sigma$, together with an elastic structure (i.e., metric) α on C. The *support* of A is the underlying multi-curve (C,c).

---

[4] Nekrashevych's combinatorial models are more general, allowing higher-dimensional cells.

There are two natural operations on elastic multi-curves. First, if $\pi\colon \widetilde{\Sigma} \to \Sigma$ is a covering map and A is an elastic multi-curve on Σ, there is a multi-curve $\widetilde{A} = \pi^{-1}(A)$ on $\widetilde{\Sigma}$ obtained by pull-back in the usual way.

Second, if $A = (C, \alpha, c)$ is an elastic multi-curve on Σ, then the *join* $\mathrm{Join}(A)$ is the elastic multi-curve obtained by

- deleting all components of C whose images are null-homotopic or bound a punctured disk; and
- replacing any components $(C_1, \alpha_1), \dots, (C_k, \alpha_k)$ of A whose images are parallel with a single component (C_0, α_0), with elastic length obtained by the harmonic sum:

$$\alpha_0 = \alpha_1 \oplus \cdots \oplus \alpha_k = \frac{1}{\frac{1}{\alpha_1} + \cdots + \frac{1}{\alpha_k}}.$$

(The harmonic sum comes from the parallel law for resistors or for springs.)

Definition 7.2. An *obstruction* for f is

- an elastic multi-curve A on $\Sigma \setminus P$ and
- a map $\psi\colon A \to \mathrm{Join}(f^{-1}(A))$

so that

- ψ commutes up to homotopy with the maps to $\Sigma \setminus P$ and
- $\mathrm{Emb}(\psi) \leqslant 1$.

Remark 7.3. Contrast the obstruction map ψ with the map $\varphi\colon f^{-1}(G) \to G$ in the statement of Theorem 1: The maps are going the opposite direction.

THEOREM 7.4 (W. Thurston, Douady-Hubbard [DH93]). *Let $f\colon (\Sigma, P) \circlearrowleft$ be a topological branched self-cover so that the first return map is not a Lattés map on any component. Then f is equivalent to a rational map if and only if there is no obstruction for f.*

The usual formulation of Theorem 7.4 refers to the maximum eigenvalue of a matrix constructed out of the multi-curves underlying A. The above formulation is equivalent by Perron-Frobenius theory, as we spell out in Proposition 7.14 below. Intuitively, Theorem 7.4 says that f is rational if and only if there is no conformal collection of annuli that gets (weakly) wider under backwards iteration.

7.2. *Obstructions for virtual endomorphisms.* We now turn to obstructions in the more general setting of the asymptotic p-conformal energies from Section 6. We also switch to virtual endomorphisms of topological spaces (e.g., graphs) or orbifolds. (In the context of branched self-covers $f\colon (\Sigma, P) \circlearrowleft$, we should consider the orbifold of f.)

Definition 7.5. For $1 < p < \infty$ and $\alpha_1, \alpha_2 \in \mathbb{R}_{>0}$, the *p-harmonic sum* of α_1 and α_2 is

$$\alpha_1 \oplus_p \alpha_2 \coloneqq \bigl((\alpha_1)^{1-p} + (\alpha_2)^{1-p}\bigr)^{1/(1-p)}. \tag{7.6}$$

For $p = \infty$, set

$$\alpha_1 \oplus_\infty \alpha_2 \coloneqq \min(\alpha_1, \alpha_2).$$

This definition is chosen so that the p-energies satisfy a parallel law.

PROPOSITION 7.7. *For $1 < p \leqslant q \leqslant \infty$, let $[\varphi]\colon G^p \to H^q$ be a homotopy class of maps from a p-conformal graph to a q-conformal graph. Suppose that G has two parallel edges e_1 and e_2 that are mapped to homotopic paths by φ. Let G_3 be the p-conformal graph G with e_1 and e_2 replaced by a single edge e_3 with $\alpha(e_3) = \alpha(e_1) \oplus_p \alpha(e_2)$, and let $[\varphi_3]\colon G_3 \to H$ be the natural homotopy class. Then*

$$E^p_q[\varphi_3] = E^p_q[\varphi].$$

Proof. If $q = \infty$, then the optimal maps in φ and in φ' will be constant-derivative. The result follows by examining the energy and comparing the derivatives. The general statement follows from the $q = \infty$ case by equation (6.3). ☐

For $p = 1$, we do not have a p-length in the same way; instead, a 1-conformal graph is a weighted graph, and when joining them in parallel we add the weights.

Definition 7.8. A *p-conformal multi-curve* $A = (C, \alpha, c)$ on a space X is a multi-curve $c\colon C \to \Gamma$, together with a p-conformal structure α on C. (We will sometimes allow punning and write the domain of the map c as A.) The *join* $\mathrm{Join}_p(A)$ is obtained by deleting components of C whose image is null-homotopic or torsion in $\pi_1(X)$ and replacing components of C whose images are parallel (up to homotopy) by a single component so that

- for $p > 1$, the new p-length is the p-harmonic sum of the constituent p-lengths; and
- for $p = 1$, the new weight is the (ordinary) sum of the constituent weights.

If $\pi\colon X_1 \to X_0$ is a covering map and A is a p-conformal multi-curve on X_0, we have the usual *pull-back* $\pi^* A$, a p-conformal multi-curve on X_1. If $\varphi\colon X_1 \to X_0$ is a map and $A = (C, \alpha, c)$ is a p-conformal multi-curve on X_1, the *push-forward* $\varphi_* A$ is $\mathrm{Join}_p(C, \alpha, \varphi \circ c)$.

If $\pi, \varphi\colon X_1 \rightrightarrows X_0$ is a virtual endomorphism, a *p-obstruction* for (π, φ) is

- a p-conformal multi-curve $A = (C, \alpha, c)$ on Γ_0 and
- a map $\psi\colon A \to \varphi_* \pi^* A$

so that

- ψ commutes up to homotopy with the maps to X_0 and
- $E_p^p(\psi) \leqslant 1$.

PROPOSITION 7.9. *If $1 \leqslant p \leqslant q \leqslant \infty$, G is a q-conformal graph, and A is a p-conformal multi-curve on G, then $E_q^p[A] = E_q^p[\mathrm{Join}_p(A)]$.*

Proof. Parallel to Proposition 7.7. ☐

PROPOSITION 7.10. *Let $\pi, \varphi \colon G_1 \rightrightarrows G_0$ be a virtual endomorphism of p-conformal graphs. If there is a p-obstruction for (π, φ), then $E_p^p[\varphi] \geqslant 1$. Likewise, if $f\colon (\Sigma, P) \circlearrowleft$ is a topological branched self-cover of hyperbolic type compatible with (π, φ) and there is a p-obstruction for f, then $E_p^p[\varphi] \geqslant 1$.*

Remark 7.11. In the branched self-cover case, if f is not of hyperbolic type and we do not use orbifolds, then considering boundary curves shows that we always have $E_p^p[\varphi] \geqslant 1$. See Section 1.2.

Proof. Let (A, ψ) be a p-obstruction for(π, φ). We have a diagram of maps

$$\begin{tikzcd}
\pi^*A \arrow[r, "\widetilde{c}"'] \arrow[d] & G_1 \arrow[dd, "\varphi"] \\
\varphi_*\pi^*A \arrow[dr] & \\
A \arrow[u, "\psi"] \arrow[r, "c"'] & G_0,
\end{tikzcd}$$

commuting up to homotopy. Since E_p^p is invariant under covers, $E_p^p[\widetilde{c}] = E_p^p[c]$. Proposition 7.9 guarantees that

$$E_p^p[\varphi_*\pi^*A \to G_0] = E_p^p[\widetilde{c} \circ \varphi].$$

Then by sub-multiplicativity and the assumptions, we have

$$E_p^p[c] \leqslant E_p^p[\psi] E_p^p[\widetilde{c}] E_p^p[\varphi] \leqslant \cdot E_p^p[c] \cdot E_p^p[\varphi],$$

so $E_p^p[\varphi] \geqslant 1$. The statement for branched self-covers follows immediately. ☐

See Corollary 7.15 for a strengthening of Proposition 7.10.

7.3. *Duality.* There is a notion of *p-conductance*, dual to p-length. Recall that we can describe electrical networks either in terms of resistances or in terms of conductances. Resistances add in series and change by a harmonic sum in parallel. Conductances add in parallel and change by a harmonic sum in series. Alternately, we can think about the relation between extremal length and modulus of conformal annuli.

There is a similar story for general p-conformal graphs when $1 < p < \infty$. Recall [Thu19, Def. A.17] that we can think of an edge e in a p-conformal graph as an equivalence class of rectangles of length ℓ and height w, with the p-length α given by

$$\alpha = \frac{\ell}{w^{1/(p-1)}}.$$

Now consider a dual view, interchanging the role of length and height. The *p-conductance* is

$$\gamma := \alpha^{1-p} = \frac{w}{\ell^{p-1}} = \frac{w}{\ell^{1/(p^\vee-1)}},$$

where $p^\vee = p/(p-1)$ is the Hölder conjugate of p.

Propositions 7.7 and 7.9 say that p-conductances add in parallel.

In checking whether a p-conformal multi-curve A is a p-obstruction, there are two basic operations. Let us recall what happens to the p-lengths.

- We pass to a cover by taking $f^{-1}(A)$ (i.e., the pullback by the covering map). A connected cover of a circle is necessarily a (longer) circle, which is series composition. Thus if a component of $f^{-1}(A)$ covers a component of A by a degree d map, the p-length gets multiplied by d.
- We merge parallel components in the Join_p operation. The p-weights change by the p-harmonic sum (equation (7.6)).

If we work with the dual p-conductances instead, the two operations switch in complexity.

- If a circle of p-conductance γ is covered by a degree d map, the pull-back p-conductance is $d^{1-p}\gamma$.
- The parallel composition in Join_p becomes simpler: add the constituent p-conductances.

7.4. *Obstruction matrices and invariant multi-curves.* We now investigate when we can choose p-lengths on a given curve to make it into a p-obstruction. As a result of the previous section, if we use p-conductances to search for p-obstructions, we get linear inequalities and can construct a matrix.

Definition 7.12. Let $\pi, \varphi\colon X_1 \rightrightarrows X_0$ be a virtual endomorphism, let $c\colon C \to X_0$ be a multi-curve, and let $1 \leqslant p < \infty$. Then the *p-obstruction matrix $M_{C,p}$* of C is the square matrix with rows and columns indexed by the components of C, with the (C_i, C_j) entry given by

$$\sum_{\substack{D \in \pi^* C_i \\ \varphi_* D \sim C_j}} \bigl(\deg(D \xrightarrow{\pi} C_i)\bigr)^{1-p}.$$

In other words, consider all components D of the multi-curve $\pi^* C_i$ on X_1 that push forward to C_j, and sum a power of the degree that D covers C_i. There may be components of $\pi^* C_i$ that do not push forward to any C_j; these components are ignored.

The matrix $M_{C,p}$ is designed to mimic the action of $\varphi_*\pi^*$ on p-conformal multi-curves. Suppose C has n components, and let $\gamma = (\gamma_i)_{i=1}^n \in \mathbb{R}^n_{\geqslant 0}$ be a non-negative vector. Then define $A(\gamma)$ to be the p-conformal multi-curve in which a component C_i of C is given p-conductance γ_i, or is dropped if $\gamma_i = 0$.

Lemma 7.13. *For γ as above,*

$$\varphi_*\pi^*A(\gamma) = A(M_{C,p}\gamma) + A',$$

where A' is a p-conformal multi-curve whose support is disjoint from C.

Proof. Immediate from Proposition 7.9. ☐

Observe that $M_{C,p}$ has non-negative entries. Therefore, by the Perron-Frobenius Theorem, it has a positive eigenvalue of maximum absolute value with a non-negative eigenvector. (Our assumptions do not guarantee that $M_{C,p}$ is irreducible, so the eigenvector might not be unique.) Let $\lambda(M_{C,p})$ be this positive eigenvalue.

Proposition 7.14. *Let $\pi, \varphi\colon X_1 \rightrightarrows X_0$ be a virtual endomorphism, and let (C, c) be a multi-curve on X_0. Then $\lambda(M_{C,p}) \geqslant 1$ if and only if there is a p-obstruction whose support is a sub-multi-curve of C.*

Proof. If $\lambda = \lambda(M_{C,p}) \geqslant 1$, let γ be the corresponding non-negative eigenvector. Then by Lemma 7.13,

$$\varphi_*\pi^*A(\gamma) = A(\lambda\gamma) + A'.$$

The p-conductances on the components of $A(\gamma)$ are multiplied by λ in $\varphi_*\pi^*A(\gamma)$ and the p-lengths are multiplied by $\lambda^{-1/(p-1)}$. Thus, tracing through the definition of E^p_p from equation (6.1), we have $E^p_p\bigl[A(\gamma) \to \varphi_*\pi^*A(\gamma)\bigr] = \lambda^{-1/p} \leqslant 1$, so $A(\gamma)$ is an obstruction.

Conversely, suppose that we have a p-obstruction A whose support is a sub-multi-curve of C. Form a vector γ from the p-conductances of A, extended by 0 for components of C that are not in A. Then Lemma 7.13 and the assumption that A is a p-obstruction say that each component of $M_{C,p}\gamma$ is greater than or equal to the corresponding component of γ. By the Collatz-Wielandt formula, this implies that $\lambda(M_{C,p}) \geqslant 1$. ☐

Corollary 7.15. *If $\pi, \varphi\colon \Gamma_1 \rightrightarrows \Gamma_0$ has a p-obstruction, then so do the iterates $\pi^n, \varphi^n\colon \Gamma_n \rightrightarrows \Gamma_0$. In particular, if there is a p-obstruction for (π, φ), then $\overline{E}^p_p[\pi, \varphi] \geqslant 1$.*

Proof. For any $n \geqslant 1$, a matrix M has an eigenvalue of absolute value greater than 1 if and only if M^n does. Apply Propositions 7.14 and 7.10. ☐

Let us investigate what the support of a p-obstruction can look like.

Definition 7.16. Let $\pi, \varphi\colon X_1 \rightrightarrows X_0$ be a virtual endomorphism, and let (C, c) be a multi-curve on X_0. Then C is *forwards-invariant* if each component of C is homotopic to a component of $\varphi_*\pi^*C$. (Here, φ_* is defined as in Definition 7.8, but without any α-lengths.) C is *irreducible* if, for any two components C_i and C_j of C, there is some n so that C_j appears in $(\varphi_*\pi^*)^n C_i$. (An irreducible curve is necessarily forwards-invariant.) C is *back-invariant* if each component of $\varphi_*\pi^*C$ is homotopic to a component of C, and C is *totally invariant* if the components of C are in bijection with the components of $\varphi_*\pi^*C$.

To explain the terminology, consider a branched self-cover R $f\colon (S^2, P) \circlearrowleft$. Let C be a multi-curve on $S^2 \setminus P$.

- C is back-invariant if and only if, up to homotopy in $S^2 \setminus P$, we have $C \subset f^{-1}(C)$.
- C is forwards-invariant if and only if C is homotopic in $S^2 \setminus P$ to a multi-curve C_1 on $S^2 \setminus f^{-1}(P)$ with $f(C_1) \subset C$.[5]

If A is a p-obstruction for (π, φ), then the underlying multi-curve C of A must be forwards-invariant. (Otherwise, no map $\psi\colon A \to \varphi_*\pi^*A$ is possible.) The matrix $M_{C,p}$ is irreducible in the Perron-Frobenius sense if and only if C is irreducible as a multi-curve. On the other hand, in the context of rational maps it is more traditional to look at back-invariant multi-curves.

We can often switch between back-invariant and forward-invariant multi-curves as follows. First, some graph-theory terminology. In a directed graph, a *strongly-connected component* (SCC) is a maximal set S of vertices so that every ordered pair of vertices in S can be connected by a directed edge-path. Every directed graph is a disjoint union of its SCCs. A *strict* SCC is an SCC in which every pair of vertices can be connected by a non-trivial directed edge-path. A non-strict SCC is a single vertex with no self-loop.

Given a virtual endomorphism $\pi, \varphi\colon X_1 \rightrightarrows X_0$ and a multi-curve C on X_0, form the directed graph $\Gamma(C)$ whose vertices are the components of C, with an arrow from C_i to C_j if C_j appears as a component of $\varphi_*\pi^*C_i$. A strict SCC of $\Gamma(C)$ gives a forward-invariant multi-curve, although not all forward-invariant multi-curves arise in this way.

Proposition 7.17. *Let $\pi, \varphi\colon X_1 \rightrightarrows X_0$ be a virtual endomorphism, and let C be a multi-curve on X_0. Then there is an irreducible forwards-invariant sub-multi-curve $C_0 \subset C$ with $\lambda(M_{C_0,p}) = \lambda(M_{C,p})$.*

Proof. $M_{C,p}$ is block triangular with respect to the partial order on the SCCs of $\Gamma(C)$, so its maximum eigenvalue will be equal to the Perron-Frobenius

---

[5] Recall that the forward image of a multi-curve in $S^2 \setminus P$ is not well defined.

eigenvalue of a diagonal block corresponding to an SCC S. If S is a single vertex with no self-loop, then $\lambda(M_{C,p}) = 0$ and we can take C_0 to be empty. Otherwise, take C_0 to be the union of multi-curves in S. □

PROPOSITION 7.18. *Let* $\pi, \varphi \colon \Sigma_1 \rightrightarrows \Sigma_0$ *be a surface virtual endomorphism* (*with* φ *a surface embedding*), *and let* C *be a simple forward-invariant multi-curve on* Σ_0. *Then there is a simple back-invariant multi-curve* $C_\infty \supset C$.

Proof. For $i \geqslant 1$, let $C_i = \varphi_* \pi^* C_{i-1}$ by induction. Then each C_i is a simple multi-curve and $C_{i-1} \subset C_i$. Since there is a bound on how many components a simple multi-curve on a surface of finite type can have, the C_i eventually stabilize into a back-invariant multi-curve. □

Although back-invariant multi-curves are more traditional, forwards-invariant multi-curves appear to be more generally useful. In a non-surface setting, or if $p < 2$, obstructions need not be simple and there is no obvious analogue of Proposition 7.18.

7.5. *Annular obstructions and asymptotic energy.* Let $\pi, \varphi \colon \Gamma_1 \rightrightarrows \Gamma_0$ be a virtual endomorphism. For any forwards-invariant multi-curve C on Γ_0, there is unique value of p so that $\lambda(M_{C,p}) = 1$, which we denote $Q(C)$ [HP08, Lemma A.2]. Define $Q(\pi, \varphi)$ to be the maximum of $Q(C)$ over all forwards-invariant multi-curves C.

PROPOSITION 7.19. *Let* $\pi, \varphi \colon \Gamma_1 \rightrightarrows \Gamma_0$ *be a graph virtual endomorphism. Then if* $\overline{E}_p^p[\pi, \varphi] \geqslant 1$, *we have* $Q(\pi, \varphi) \leqslant p$.

Proof. Immediate from Corollary 7.15 and Propositions 7.14 and 6.10. □

Compare Proposition 7.19 to the following result of Haïssinsky and Pilgrim.

THEOREM 7.20 (Haïssinsky–Pilgrim [HP08]). *Suppose* $f \colon S^2 \to S^2$ *is topologically cxc. Then* $Q(f) \leqslant \operatorname{confdim}_{AR}(f)$.

Here, $Q(f)$ is the version of $Q(\pi, \varphi)$ for branched self-covers. *Topologically cxc* is a topological notion of expanding branched self-covers and, in particular, implies that there are no cycles with branch points in P (the opposite of the hyperbolic case). The *Ahlfors regular conformal dimension* $\operatorname{confdim}_{AR}(f)$ is an analytically defined quantity, the minimal Hausdorff dimension of any Ahlfors regular metric in a certain quasi-symmetry class of expanding metrics canonically associated to f.

Theorem 7.20, like Proposition 7.19, gives upper bounds on $Q(f)$. However, it gives bounds in terms of the purely analytic Ahlfors regular conformal dimension, rather than the asymptotic energy. In addition, Theorem 7.20 applies to maps with *no* branched cycles in P, while Proposition 7.19 is vacuous unless *every* cycle in P is branched. (If there is an unbranched cycle in P, then $\overline{E}^p_p[\varphi] \geqslant 1$ for every $p \in [1, \infty]$.)

Question 7.21. Suppose $f\colon (\Sigma, P) \circlearrowleft$ is a branched self-cover of hyperbolic type, with compatible virtual endomorphism (π, φ). Is it true that $\overline{E}^p_p[\pi, \varphi] < 1$ if and only if there is no p-obstruction?

Theorems 7.4 and 1 combine to say that the answer to Question 7.21 is positive for $p = 2$. However, the proof is quite roundabout, needing the full strength of both theorems. One could hope for a more direct proof of equivalence of the two criteria and a generalization to other values of p. (This might also give another proof of Theorem 7.4.)

References

[BBL$^+$00] E. Brezin, R. Byrne, J. Levy, K. Pilgrim, and K. Plummer, A census of rational maps, *Conform. Geom. Dyn.* **4** (2000), 35–74. MR `1749249`. Zbl `0989.37038`. https://doi.org/10.1090/S1088-4173-00-00050-3.

[BCT14] X. Buff, G. Cui, and L. Tan, Teichmüller spaces and holomorphic dynamics, in *Handbook of Teichmüller Theory. Vol. IV*, *IRMA Lect. Math. Theor. Phys.* **19**, Eur. Math. Soc., Zürich, 2014, pp. 717–756. MR `3289714`. Zbl `1314.30079`. https://doi.org/10.4171/117-1/17.

[CFKP03] J. W. Cannon, W. J. Floyd, R. Kenyon, and W. R. Parry, Constructing rational maps from subdivision rules, *Conform. Geom. Dyn.* **7** (2003), 76–102. MR `1992038`. Zbl `1077.37037`. https://doi.org/10.1090/S1088-4173-03-00082-1.

[CP14] M. Carrasco Piaggio, Conformal dimension and canonical splittings of hyperbolic groups, *Geom. Funct. Anal.* **24** no. 3 (2014), 922–945. MR `3213834`. Zbl `1300.30100`. https://doi.org/10.1007/s00039-014-0282-7.

[CJS03a] G. Cui, Y. Jiang, and D. Sullivan, On geometrically finite branched coverings. I. Locally combinatorial attracting, in *Complex Dynamics and Related Topics: Lectures from the Morningside Center of Mathematics*, *New Stud. Adv. Math.* **5**, Int. Press, Somerville, MA, 2003, pp. 1–14. MR `2504307`. Zbl `1201.37068`.

[CJS03b] G. Cui, Y. Jiang, and D. Sullivan, On geometrically finite branched coverings. II. Realization of rational maps, in *Complex Dynamics and Related Topics: Lectures from the Morningside Center of Mathematics*, *New Stud. Adv. Math.* **5**, Int. Press, Somerville, MA, 2003, pp. 15–29. MR `2504308`. Zbl `1201.37067`.

[CPT16] G. Cui, W. Peng, and L. Tan, Renormalizations and wandering Jordan curves of rational maps, *Comm. Math. Phys.* **344** no. 1 (2016), 67–115. MR `3493138`. Zbl `1362.37095`. https://doi.org/10.1007/s00220-016-2623-x.

[CT11] G. Cui and L. Tan, A characterization of hyperbolic rational maps, *Invent. Math.* **183** no. 3 (2011), 451–516. MR `2772086`. Zbl `1230.37052`. https://doi.org/10.1007/s00222-010-0281-8.

[CT18] G. Cui and L. Tan, Hyperbolic-parabolic deformations of rational maps, *Sci. China Math.* **61** no. 12 (2018), 2157–2220. MR `3881954`. Zbl `1404.37049`. https://doi.org/10.1007/s11425-018-9426-4.

[DH85] A. Douady and J. H. Hubbard, On the dynamics of polynomial-like mappings, *Ann. Sci. École Norm. Sup.* (4) **18** no. 2 (1985), 287–343. MR `0816367`. Zbl `0587.30028`. https://doi.org/10.24033/asens.1491.

[DH93] A. Douady and J. H. Hubbard, A proof of Thurston's topological characterization of rational functions, *Acta Math.* **171** no. 2 (1993), 263–297. MR `1251582`. Zbl `0806.30027`. https://doi.org/10.1007/BF02392534.

[Eps66] D. B. A. Epstein, Curves on 2-manifolds and isotopies, *Acta Math.* **115** (1966), 83–107. MR `0214087`. Zbl `0136.44605`. https://doi.org/10.1007/BF02392203.

[FB18] M. Fortier Bourque, The holomorphic couch theorem, *Invent. Math.* **212** no. 2 (2018), 319–406. MR `3787830`. Zbl `1394.30028`. https://doi.org/10.1007/s00222-017-0769-6.

[HP08] P. Haïssinsky and K. M. Pilgrim, Thurston obstructions and Ahlfors regular conformal dimension, *J. Math. Pures Appl.* (9) **90** no. 3 (2008), 229–241. MR `2446078`. Zbl `1213.30048`. https://doi.org/10.1016/j.matpur.2008.04.006.

[HSS09] J. Hubbard, D. Schleicher, and M. Shishikura, Exponential Thurston maps and limits of quadratic differentials, *J. Amer. Math. Soc.* **22** no. 1 (2009), 77–117. MR `2449055`. Zbl `1206.37026`. https://doi.org/10.1090/S0894-0347-08-00609-7.

[IS10] Y. Ishii and J. Smillie, Homotopy shadowing, *Amer. J. Math.* **132** no. 4 (2010), 987–1029. MR `2663646`. Zbl `1207.37013`. https://doi.org/10.1353/ajm.0.0126.

[Kah06] J. Kahn, A priori bounds for some infinitely renormalizable quadratics. I. Bounded primitive combinatorics, 2006, preprint ims06-05, Stony Brook IMS. arXiv `math/0609045v2`.

[KPT15] J. Kahn, K. M. Pilgrim, and D. P. Thurston, Conformal surface embeddings and extremal length, 2015. arXiv `1507.05294`.

[Kat04] T. Katsura, A class of C^*-algebras generalizing both graph algebras and homeomorphism C^*-algebras. I. Fundamental results, *Trans. Amer. Math. Soc.* **356** no. 11 (2004), 4287–4322. MR `2067120`. Zbl `1049.46039`. https://doi.org/10.1090/S0002-9947-04-03636-0.

[Nek05] V. Nekrashevych, *Self-Similar Groups*, *Math. Surveys Monogr.* **117**, Amer. Math. Soc., Providence, RI, 2005. MR `2162164`. Zbl `1087.20032`. https://doi.org/10.1090/surv/117.

[Nek14] V. Nekrashevych, Combinatorial models of expanding dynamical systems, *Ergodic Theory Dynam. Systems* **34** no. 3 (2014), 938–985. MR `3199801`. Zbl `1350.37034`. https://doi.org/10.1017/etds.2012.163.

[Nek16] V. Nekrashevych, Mating, paper folding, and an endomorphism of $\mathbb{PC}^2$, *Conform. Geom. Dyn.* **20** (2016), 303–358. MR `3574443`. Zbl `1408.37082`. https://doi.org/10.1090/ecgd/302.

[PT] K. Pilgrim and D. P. Thurston, Graph energies and Ahlfors-regular conformal dimension, in preparation.

[Put16] A. Putman, Answer to Who proved that two homotopic embeddings of one surface in another are isotopic?, 2016, MathOverflow. Available at http://mathoverflow.net/a/248866.

[Sul81] D. Sullivan, On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions, in *Riemann Surfaces and Related Topics*: *Proceedings of the* 1978 *Stony Brook Conference* (State Univ. New York, Stony Brook, N.Y., 1978), *Ann. of Math. Stud.* **97**, Princeton Univ. Press, Princeton, N.J., 1981, pp. 465–496. MR `0624833`. Zbl `0567.58015`. https://doi.org/10.1515/9781400881550-035.

[Thu16] D. P. Thurston, From rubber bands to rational maps: a research report, *Res. Math. Sci.* **3** (2016), Paper No. 15, 49. MR `3500499`. Zbl `1360.37119`. https://doi.org/10.1186/s40687-015-0039-4.

[Thu19] D. P. Thurston, Elastic graphs, *Forum Math. Sigma* **7** (2019), e24, 84. MR `3993808`. Zbl `07094012`. https://doi.org/10.1017/fms.2019.4.

[Wan14] X. Wang, A decomposition theorem for Herman maps, *Adv. Math.* **267** (2014), 307–359. MR `3269181`. Zbl `1418.37078`. https://doi.org/10.1016/j.aim.2014.09.004.

[ZJ09] G. Zhang and Y. Jiang, Combinatorial characterization of sub-hyperbolic rational maps, *Adv. Math.* **221** no. 6 (2009), 1990–2018. MR `2522834`. Zbl `1190.37051`. https://doi.org/10.1016/j.aim.2009.03.009.

(Received: December 14, 2016)
(Revised: May 1, 2020)

Indiana University, Bloomington, IN
E-mail: dpthurst@indiana.edu