Groups Geom. Dyn. 20 (2026), 971–1050
DOI 10.4171/GGD/914

Ahlfors-regular conformal dimension and energies of graph maps

Kevin M. Pilgrim and Dylan P. Thurston

Abstract. For a hyperbolic rational map f with connected Julia set, we give upper and lower bounds on the Ahlfors-regular conformal dimension of its Julia set J_f from a family of energies of associated graph maps. Concretely, the dynamics of f is faithfully encoded by a pair of maps $\pi, \phi: G_1 \rightrightarrows G_0$ between finite graphs that satisfies a natural expanding condition. Associated to this combinatorial data, for each $q \geq 1$, is a numerical invariant $\overline{E}^q[\pi, \phi]$, its asymptotic q-conformal energy. We show that the Ahlfors-regular conformal dimension of J_f is contained in the interval where $\overline{E}^q = 1$. Among other applications, we give two families of quartic rational maps with Ahlfors-regular conformal dimension approaching 1 and 2, respectively.

Contents

1. Introduction

1.1. Motivation

The iterates of a rational function f define a holomorphic dynamical system on the Riemann sphere $\widehat{\mathbb{C}}$. Its Julia set, typically fractal, may be defined as the smallest set J_f satisfying $J_f = f^{-1}(J_f) = f(J_f)$ and $\#J_f \geq 3$.

For instance, Figure 1 shows the Julia set of the rational function $f(z) = \frac{4}{27}\frac{(z^2-z+1)^3}{(z(z-1))^2}$. It turns out that as a topological space, this J_f is a *Sierpiński carpet*—the complement in

Mathematics Subject Classification 2020: 37F10 (primary); 20E08, 30L10 (secondary).
Keywords: conformal dimension, dynamics of rational maps, quasi-conformal geometry, Ahlfors-regular conformal dimension.

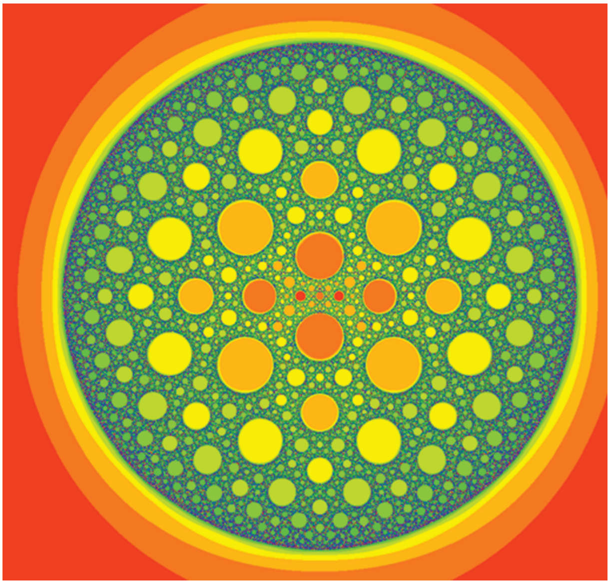

Figure 1. The Julia set of f.

the sphere of a countable collection of Jordan domains whose closures are disjoint and whose diameters tend to zero. With the spherical metric d_{sph} inherited from the round metric on $\widehat{\mathbb{C}}$, the Julia set J_f becomes a metric space. As a dynamical system, the map f is *hyperbolic*—each critical point converges to an attracting cycle—and *critically finite*—the orbits of the critical points are finite. Hyperbolicity is equivalent to the condition that the restriction $f\colon J_f \to J_f$ is an expanding self-covering map. In this setting, hyperbolicity is a dynamical regularity condition that leads to strong metric consequences. As is visually evident, J_f is *approximately self-similar* (see Definition 3.4). A key invariant of J_f is its Hausdorff dimension $\mathrm{hdim}(J_f)$.

D. Sullivan [48, Theorem 4] showed that for any hyperbolic rational map f, upon setting $q := \mathrm{hdim}(J_f)$ and $\mathcal{H}^q$ the corresponding Hausdorff measure on J_f, one has $0 < q < 2$ and $0 < \mathcal{H}^q(J_f) < \infty$; see also [43, Theorem 9.1.6, Corollary 9.1.7] for more general statements. In particular, for any ball $B(x, r)$ with $x \in J_f$ and any $r \le \mathrm{diam}(J_f)$, we have $\mathcal{H}^q(B(x, r)) \asymp r^q$, with implicit constants independent of x and r. This latter condition is known as *Ahlfors q-regularity*. If in addition $J(f)$ is a carpet, we have $1 < \mathrm{hdim}(J_f) < 2$; see [28, 44] for these lower and upper bounds, respectively.

A homeomorphism between metric spaces is *quasi-symmetric* (qs) if it does not distort the roundness of balls too much; see Section 3.3. The *Ahlfors-regular conformal gauge* $\mathcal{G}$ of J_f is the set of all metric measure spaces (X, d, μ) such that there exists a quasi-symmetric homeomorphism $(J_f, d_{\mathrm{sph}}) \to (X, d)$ and, for some $q > 0$, the measure μ is q-Ahlfors regular with respect to d; see [24, 29]. The regularity assumption on μ implies that μ is comparable to the q-dimensional Hausdorff measure $\mathcal{H}^q$ on X. The

Ahlfors-regular conformal dimension of J_f is the infimum over such exponents q, that is,

$$\mathrm{ARCdim}(J_f) := \inf\{\mathrm{hdim}(X) \mid (X, d, \mu) \in \mathcal{G}(J_f)\};$$

see [35] for an introduction. For approximately self-similar carpets X, such as the Julia set J_f of Figure 1, we know $1 < \mathrm{ARCdim}(X) < 2$; see [34] and [43, Corollary 9.17]. Interest in conformal dimension stems in part from the following. Limit sets of Kleinian groups acting on the Riemann sphere and, more generally, boundaries of hyperbolic groups are another source of approximately self-similar spaces. In that setting, the conformal dimension, analogously defined, carries significant information about the group; see [32].

Hyperbolicity and the critically finite property imply, according to a rigidity result of W. Thurston [18], that the geometry and dynamics on J_f are determined by topological data: the conjugacy-up-to-isotopy class of f, relative to its post-critical set. More precisely, if g is another rational map, and if there are orientation-preserving homeomorphisms $\phi_0, \phi_1 : (\widehat{\mathbb{C}}, P_f) \to (\widehat{\mathbb{C}}, P_g)$ such that $\phi_0 \circ f = g \circ \phi_1$ on P_f and ϕ_0 is isotopic to ϕ_1 through homeomorphisms agreeing on P_f, then we may take $\phi_0 = \phi_1$ to be a Möbius transformation. Hence the invariant $\mathrm{ARCdim}(J_f)$ is determined from discrete data.

For general hyperbolic critically finite rational maps with connected Julia set, our main result, Theorem A, implies an estimate for $\mathrm{ARCdim}(J_f)$ in terms of combinatorial data. In concrete cases, by-hand computations with these data can yield nontrivial rigorous upper and lower bounds. For the carpet Julia set of Figure 1, such computations yield $1.6309 \approx \frac{1}{1-\log_6 2} \le \mathrm{ARCdim}(J_f) \le \frac{2}{1-\log_6(10/13)} \approx 1.7445$; see Section 7 for details.

1.2. Combinatorial encoding

Our methods rely on a particular method of combinatorial encoding of rational maps [52]. We choose a finite graph G_0, called a *spine*, onto which $\widehat{\mathbb{C}} - P_f$ deformation retracts. The homotopy type of G_0 depends only on $\#P_f$. Letting $G_1 = f^{-1}(G_0) \subset \widehat{\mathbb{C}} - P_f$, we obtain two graph maps $\pi, \phi : G_1 \rightrightarrows G_0$, where π and ϕ are, respectively, the restrictions of f and of the deformation retraction. The data (π, ϕ) are a *virtual endomorphism of graphs* (see Definition 2.20) and are well defined up to a notion of homotopy equivalence; see [52, Definition 2.2]. We denote by $[\pi, \phi]$ the homotopy class of (π, ϕ). Figure 2 illustrates the data for the map f. See also Figure 11.

For any iterate $n \in \mathbb{N}$, the Julia set of f is the same as that of f^n. It follows from the expanding nature of the dynamics of f that upon replacing f by a suitable iterate, we may assume the virtual endomorphism (π, ϕ) constructed in the previous paragraph is *forward expanding* or, synonymously, *backward contracting*; see Definition 2.22. The critically finite property implies that G_1 and G_0 are connected and ϕ is surjective on the fundamental group. This is a property we call *recurrence*; see Definition 2.23. To summarize, to the dynamics of a critically finite hyperbolic rational map, we associate a forward-expanding recurrent virtual graph endomorphism (π, ϕ).

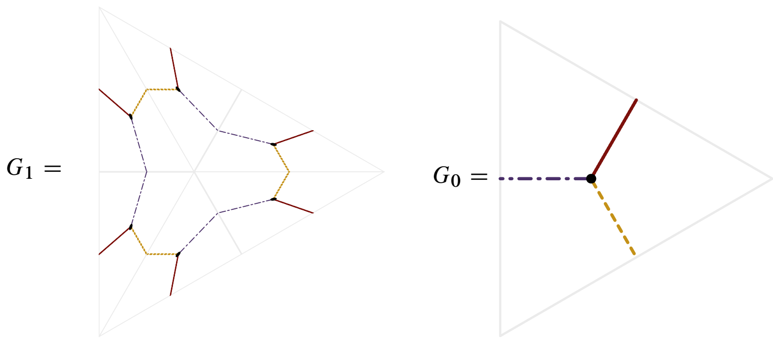

Figure 2. Spines for the map f, up to conjugacy up to isotopy. Doubling the large Euclidean equilateral triangles over their boundaries gives two Riemann surfaces, each isomorphic to $\widehat{\mathbb{C}}$. The map f sends each small triangle at left conformally to one of the two large triangles at right and implements barycentric subdivision. The set of vertices of the large triangles is P_f. Half of a spine G_0 of $\widehat{\mathbb{C}} - P_f$ is shown at right, and half its preimage G_1 under the piecewise-affine map homotopic to f is shown at left. The map π is the covering map preserving colors, while the map ϕ (defined up to homotopy) is induced by the deformation retraction.

It turns out (see Section 2) that any forward-expanding recurrent virtual graph endomorphism (π, ϕ) determines, via now-standard constructions, a dynamical system given by an expanding topological self-cover on a compact metrizable locally connected space $f: \mathcal{J} \to \mathcal{J}$ (Theorem F). The topological conjugacy class of f depends only on the homotopy class $[\pi, \phi]$; see [30, Theorem 4.2].

Via other now-standard constructions, there is an associated conformal gauge of Ahlfors-regular metrics $\mathcal{G}(\mathcal{J}[\pi, \phi])$ associated to $[\pi, \phi]$; see Section 3. If $[\pi, \phi]$ arises from a hyperbolic rational map f, then the spherical metric on J_f belongs to $\mathcal{G}(\mathcal{J}[\pi, \phi])$ (Proposition 3.3).

1.3. Asymptotic q-conformal energies

A virtual endomorphism of graphs (π, ϕ) has, for each $q \in [1, \infty]$, an associated *asymptotic q-conformal graph energy* $\overline{E}^q(\pi, \phi)$, introduced by the second author [51]. We summarize some key points; see Section 4 or the references for more. First, $\overline{E}^q(\pi, \phi)$ depends only on the homotopy class $[\pi, \phi]$ so that these analytic quantities depend only on combinatorial data. If (π, ϕ) arises from a rational map f, different choices of spine lead to homotopic graph endomorphisms so that we may write unambiguously $\overline{E}^q(f)$. In the general case, we will also write $\overline{E}^q[\pi, \phi]$ to indicate that the asymptotic energy depends only on the homotopy class. The inequality $\overline{E}^\infty[\pi, \phi] < 1$ holds if and only if some iterate of (π, ϕ) is homotopic to a backward-contracting virtual graph endomorphism. If (π, ϕ) arises from a hyperbolic critically finite rational map f, then $\overline{E}^2[\pi, \phi] < 1$, and this property characterizes such rational maps among the wider class of their topological counterparts, specifically, critically finite self-branched coverings of the sphere for which each cycle in the post-critical set contains a critical point [52]. As a function of q, the asymptotic energy $\overline{E}^q[\pi, \phi]$ is continuous and non-increasing so that the level set $\{q \mid \overline{E}^q[\pi, \phi] = 1\}$ is an interval $[q_*[\pi, \phi], q^*[\pi, \phi]]$.

1.4. Main result

Our main result relates Ahlfors-regular conformal dimension to these critical exponents and implies that the invariant $\overline{E}^q[\pi, \phi]$ contains useful information for other values of q.

Theorem A. *For any recurrent, forward-expanding virtual graph endomorphism* (π, ϕ)*,*

$$q_*[\pi, \phi] \leq \mathrm{ARCdim}(\mathcal{J}[\pi, \phi]) \leq q^*[\pi, \phi].$$

Equivalently, for $q = \mathrm{ARCdim}(\mathcal{J}[\pi, \phi])$*, we have* $\overline{E}^q[\pi, \phi] = 1$.

In fact, we expect that $q_* = q^*$.

Conjecture 1.1. *For any recurrent virtual graph endomorphism* $[\pi, \phi]$*, the function* $q \mapsto \overline{E}^q[\pi, \phi]$ *is either constant or strictly decreasing.*

In particular, if we suppose that $[\pi, \phi]$ is forward expanding, then (since $\overline{E}^1[\pi, \phi] \geq 1$ and $\overline{E}^\infty[\pi, \phi] < 1$) the conjecture would imply that

$$q_*[\pi, \phi] = q^*[\pi, \phi],$$

and Theorem A characterizes ARCdim.

One might expect more to be true in Conjecture 1.1, in particular some type of convexity of the function $q \mapsto \overline{E}^q[\pi, \phi]$ (after some reparameterization of the domain and/or range). More generally, it would be interesting to know the relationship between our constructions and more classical constructions in thermodynamic formalism. In this vein, we remark that Das, Przytycki, Tiozzo, and Urbański [15] have developed the thermodynamic formalism in a setting which includes the topologically coarse expanding conformal maps $f: \mathcal{J} \to \mathcal{J}$ considered here; see [24].

1.5. Outline

The bulk of the paper is devoted to developing the technology to prove Theorem A.

Topological dynamics (Section 2). We begin with an arbitrary forward-expanding recurrent virtual endomorphism of graphs $(\pi, \phi: G_1 \rightrightarrows G_0)$. Iteration, suitably defined, gives rise to

(1) a sequence of virtual endomorphisms $(\pi^n_{n-1}, \phi^n_{n-1}: G_n \rightrightarrows G_{n-1}), n \in \mathbb{N}$;

(2) a connected, locally connected, compact space $\mathcal{J}$, the inverse limit of

$$\cdots \xrightarrow{\phi^{n+1}_n} G_n \xrightarrow{\phi^n_{n-1}} G_{n-1} \xrightarrow{\phi^{n-1}_{n-2}} \cdots \xrightarrow{\phi^2_1} G_1 \xrightarrow{\phi^1_0} G_0;$$

(3) a positively expansive self-covering map $f\colon \mathcal{J} \to \mathcal{J}$ of degree $d := \deg(\pi)$;

(4) the topological conjugacy class of $f\colon \mathcal{J} \to \mathcal{J}$, depending only on the homotopy class $[\pi, \phi]$.

Our development in this section is quite general. We consider pairs of maps $\pi, \phi\colon X_1 \rightrightarrows X_0$ between finite CW complexes equipped with length metrics and satisfying natural expansion conditions and establish properties of the dynamics on the limit space. The main result, Theorem F, shows that under these conditions, the construction of the conformal gauges given in the next section applies. We also give a general result, Theorem G, which promotes a family of homotopy pseudo-orbits to a map to the limit space in a natural way. This allows us to relate features of the pair (π, ϕ) to features of the limit space.

The conformal gauge (Section 3). We apply a construction in [24] to put a nice metric d_ε on $\mathcal{J}$, called a *visual metric*. It depends on a suitably small but arbitrary parameter ε and on the data of a finite cover $\mathcal{U}_0$ of $\mathcal{J}$ by open, connected sets. Changing these data changes the metric by a special type of quasi-symmetric map called a *snowflake map*. Equipped with a visual metric, f is positively expansive. Even better, any ball of sufficiently small radius is sent homeomorphically and homothetically onto its image, with expansion constant e^ε. The metric space $(\mathcal{J}, d_\varepsilon, \mathcal{H}^q)$ is Ahlfors regular of exponent $q := \log(d)/\varepsilon = \operatorname{hdim}(\mathcal{J}, d_\varepsilon)$. Therefore, the invariant $\operatorname{ARCdim}(\mathcal{J}[\pi, \phi])$ is well defined; see Theorem 3.2.

Energies of graph maps (Section 4). When equipped with natural length metrics, the maps $\phi^n := \phi_0^n\colon G_n \to G_0$ have, for each $q \in [1, \infty]$, a q-conformal energy $E_q^q[\phi^n]$ in the sense introduced by the second author [51]. The growth rate of this energy as n tends to infinity, namely $\overline{E}^q[\pi, \phi] := \lim E_q^q[\phi^n]^{1/n}$, is called the *asymptotic q-conformal energy*. It depends only on the homotopy class $[\pi, \phi]$ and is continuous and non-increasing in q (Proposition 4.13). The expansion hypothesis implies $\overline{E}^\infty[\pi, \phi] < 1$, giving the interval in the statement of Theorem A.

Combinatorial modulus (Section 5). In this section, we recall (and extend slightly) results on a combinatorial version of modulus in a fairly general setting and how it is related to Ahlfors-regular conformal dimension. Although the limit space $\mathcal{J}$ hardly appears in this section, the ultimate motivation is of course to estimate its conformal dimension. In more detail, fix $n \in \mathbb{N}$. There is a natural projection $\phi_n^\infty\colon \mathcal{J} \to G_n$. The collection $\mathcal{V}_n$ of fibers $(\phi_n^\infty)^{-1}(e)$ of closed edges $e \in E(G_n)$ gives a covering of $\mathcal{J}$.[1] Given a family of curves Γ in $\mathcal{J}$ and an exponent $q \geq 1$, we get a numerical invariant, $\operatorname{mod}_q(\Gamma, \mathcal{V}_n)$, the combinatorial modulus of this family.

Sections 5.1–5.4 develop general properties of combinatorial modulus. We need to consider a mild generalization: we define combinatorial modulus for families of *weighted* curves.

[1]For technical reasons, we actually work with a slightly different cover $\mathcal{V}_n$ given by inverse images of a slightly larger set $\hat{e}$ called the *star* of e; see Definition 5.8.

Section 5.5 continues by relating combinatorial modulus to energies of graph maps. For this, it is technically convenient to work with the reciprocal of modulus, namely extremal length. The relation arises via the characterization of graph map energy in terms of maximum distortion of extremal length; see Theorem 4.11, which requires a formulation of extremal length in terms of weighted curves.

Section 5.6 recalls a result of Carrasco [12, Theorem 1.3], which was independently proved by Keith–Kleiner (unpublished). This result states that for a suitably self-similar space Z and a suitable family of coverings $(\mathcal{S}_n)_n$ indexed by $\mathbb{N}$, there is a *critical exponent* q. For a reasonably natural curve family Γ in the space, this critical exponent distinguishes between $\mathrm{mod}_p(\Gamma, \mathcal{S}_n) \to \infty$ if $p < q$ and $\mathrm{mod}_p(\Gamma, \mathcal{S}_n) \to 0$ if $p > q$; see Theorem 5.11.

Sandwiching the dimension (Section 6). The proof of Theorem A applies the developments in Sections 4 and 5. We relate combinatorial modulus of curve families in the limit space to combinatorial modulus of curve families on the graphs G_n and then to energies of graph maps.

Here is a brief summary of the proof. The collection $\{\mathcal{V}_n\}_n$ above forms a family of *snapshots* of the limit space, equipped with a visual metric; see Section 6.1. We show a curve $\gamma\colon C \to G_n$ can be approximately lifted under ϕ_n^∞ to a curve $\gamma'\colon C \to \mathcal{J}$ such that the composition $\gamma'' := \phi_n^\infty \circ \gamma'\colon C \to G_n$ is homotopic to γ, with traces of size uniformly bounded independent of n. This implies that combinatorial moduli for γ and γ'' on G_n are comparable to that of γ' on $\mathcal{J}$ (Lemma 6.4).

With this setup, the upper bound on conformal dimension is straightforward to verify; see Section 6.3. The lower bound is more involved and uses in an essential way the existence of a curve $\gamma\colon C \to G_n$ which is extremal for the distortion of extremal length in the homotopy class of $\phi_0^n\colon G_n \to G_0$. When projected to G_0, the strands of $\phi_0^n \circ \gamma$ are very long and cross edges of G_0 many times. We decompose γ into a family of subcurves ζ, each of which projects to an edge of G_0, and make the needed estimates; see Section 6.4.

1.6. Applications

In Section 7, we give several applications of our methods.

1.6.1. Techniques for estimates. Theorem A yields practical methods for estimating the Ahlfors-regular conformal dimension. If specific ϕ and q are given with $E_q^q(\phi) < 1$, the sub-multiplicativity of energy under composition yields $\overline{E}^q[\pi, \phi] < 1$ and thus, by Theorem A, $\mathrm{ARCdim}(\mathcal{J}[\pi, \phi]) < q$. Furthermore, there are bounds on how quickly $\overline{E}^q[\pi, \phi]$ can decrease as a function of q [52, Proposition 6.11], so if we know $\overline{E}^q[\pi, \phi] < 1$, we get an upper bound on q^* that is smaller than q; see Proposition 7.1.

For lower bounds, we have the following. Set $\overline{N}[\pi, \phi] = \overline{E}^1[\pi, \phi]$. One way to view this quantity is as the asymptotic growth rate of the maximal number of edge-disjoint curves in G_n, each representing a nontrivial loop in G_0. The bounds mentioned above on how fast $\overline{E}^q$ decreases as a function of q yield the following.

Theorem B. *For any recurrent forward-expanding virtual graph automorphism* $[\pi, \phi]$ *where* $\deg(\pi) = d$*, we have*

$$\mathrm{ARCdim}(\mathcal{J}[\pi, \phi]) \geq \frac{1}{1 - \log_d \overline{N}[\pi, \phi]}.$$

For the virtual endomorphism of f from Figures 1 and 2, we determine by hand that $\overline{N}(f) = 2$ and $\overline{E}^2[\pi, \phi] < \sqrt{10/13}$, so $1.6309 \approx \frac{1}{1-\log_6 2} < \mathrm{ARCdim}(J_f) < \frac{2}{1-\log_6(10/13)} \approx 1.7445$; see Section 7.2 for details.

If f is a hyperbolic rational map and $[\pi, \phi]$ an associated virtual graph endomorphism, the quantity $\overline{N}(f)$ seems to be closely related to the topological and metric structure of the Julia set. For example, we show the following.

Theorem C. *Suppose* f *is a critically finite hyperbolic rational map. If* J_f *is a Sierpiński carpet, then* $\overline{N}(f) > 1$.

Examples show that the converse need not hold; see Section 7.1.

1.6.2. When the conformal dimension equals 1. M. Carrasco [13, Theorem 1.2] gives a metric condition, *uniformly well-spread cut points* (UWSCP; Definition 7.15), on a compact, doubling metric space X which guarantees $\mathrm{ARCdim}(X) = 1$. All hyperbolic polynomials and rational maps with "gasket-type" Julia sets satisfy the UWSCP condition. Combining his observation with our Theorem B, we obtain the following.

Theorem D. *Suppose* f *is a hyperbolic rational map. If* J_f *satisfies the UWSCP condition, then* $\overline{N}(f) = 1$.

We also show that Carrasco's criterion for ARCdim $= 1$ is not necessary.

Proposition 1.2. *Let* R *be the rational map obtained by mating the Douady rabbit quadratic polynomial with the basilica polynomial* $z^2 - 1$*. Then* $\mathrm{ARCdim}(J_R) = 1$ *but* J_R *does not satisfy UWSCP.*

This example is shown in Figure 3. More generally, applying our methods, Insung Park [40] has proved the following generalization: *a hyperbolic critically finite rational map* f *satisfies* $\mathrm{ARCdim}(J_f) = 1$ *if and only if* f *is a crochet map, if and only if* $\overline{N}(f) = 1$. Here a *crochet map* is one in which any pair of points in the Fatou set is joined by a path which meets the Julia set in a countable set of points.

1.6.3. Variation in a family. R. Devaney et al. studied the family $f_\lambda(z) = z^2 + \lambda/z^2$ for $\lambda \in \mathbb{C} - \{0\}$ [16]. Figure 4 shows the bifurcation locus in the parameter plane for this family. (The four critical points at the fourth roots of λ end up in the same orbit for this family, making it a critical orbit variety [2] and simplifying the analysis; see equation (7.7).) Parameters taken from the prominent "holes" along the real axis have carpet

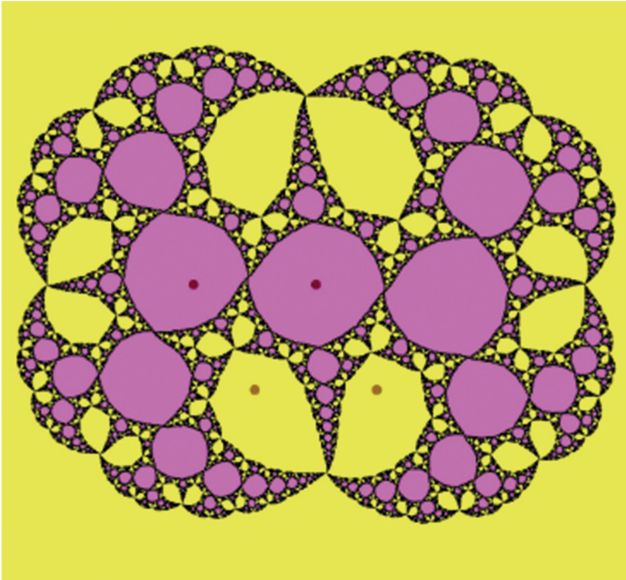

Figure 3. The mating of the rabbit (yellow) and the basilica (pink) polynomials.

Julia sets. In Sections 7.3 and 7.4, we present two sequences $\lambda_n^{\text{skinny}}, \lambda_n^{\text{fat}}$, $n \in \mathbb{N}$, whose values converge, respectively, to the left-most real parameter (at $\lambda = -\frac{3}{16} - \frac{\sqrt{2}}{8}$) and to the origin. Figure 5 illustrates two examples. Note the difference in the apparent "thickness" of the carpets. Though each Julia set is a Sierpiński carpet, that of $\lambda_n^{\text{skinny}}$ visually becomes "skinnier" as $n \to \infty$, while that of λ_n^{fat} visually becomes "fatter" as $n \to \infty$. The former rate seems to be rather gradual, while the latter rate seems to be very fast.

M. Bonk, M. Lyubich, and S. Merenkov showed that a quasi-symmetric map between hyperbolic carpet Julia sets extends to a Möbius transformation [3]. This easily implies that the three carpets shown in Figures 1 and 5 are pairwise quasi-symmetrically inequivalent. Our techniques allow us to quantify this distinction.

Theorem E. *We have*

$$1 < \operatorname{ARCdim}(J_{\lambda_n^{\text{skinny}}}) < 1 + \frac{1}{\log_2(2n+3)},$$

and

$$\frac{2}{1+2^{-n}} \leq \operatorname{ARCdim}(J_{\lambda_n^{\text{fat}}}) < 2.$$

In particular, within this family of fixed degree, there are hyperbolic carpet maps with conformal dimension tending to 1 and to 2. This latter result answers a question of the first author and P. Haïssinsky [26].

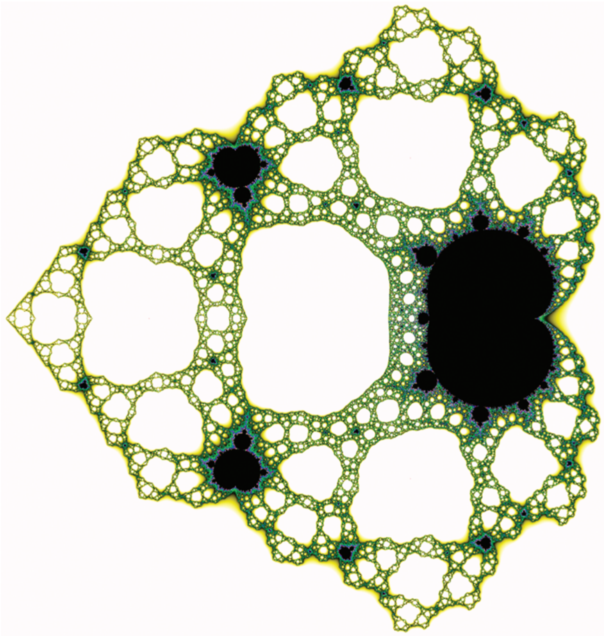

Figure 4. Parameter plane for $f_\lambda(z) = z^2 + \lambda/z^2$. The white regions are where the orbits of the critical points $\lambda^{1/4} \mapsto \pm 2\sqrt{\lambda} \mapsto 4\lambda + 1/4 \mapsto \cdots$ escape to the attracting fixed point at ∞. The small black Mandelbrot sets correspond to cases where that orbit is attracted to a cycle other than the fixed point at infinity.

1.7. Note on notions of conformal dimension

The *conformal Hausdorff dimension* of an Ahlfors-regular metric space X is defined as the infimum of the Hausdorff dimensions of metric spaces quasi-symmetrically homeomorphic to X. *A priori* it could be strictly smaller than $\mathrm{ARCdim}(X)$.

After submission of this article, we learned of a result of S. Eriksson-Bique [20, Theorem 1.6], which implies the following.

Theorem 1.3 (Eriksson-Bique). *Let X be a compact, connected, locally connected, quasi-self-similar metric space. Then the conformal Hausdorff dimension and Ahlfors-regular conformal dimension coincide.*

The property of being quasi-self-similar means arbitrarily small balls in X are quasi-symmetrically equivalent to open sets in X of definite size, where the distortion function in the definition of quasi-symmetry is independent of the location and radius of the chosen ball. This property is invariant under quasi-symmetric maps. The visual metrics we construct in Section 3 are quasi-self-similar. Therefore, all of our main results hold with "Ahlfors-regular conformal dimension" in the conclusion replaced with "conformal Hausdorff dimension".

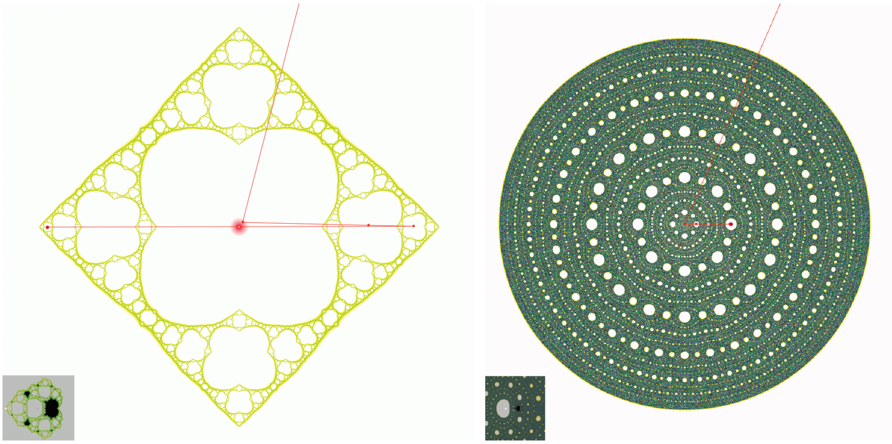

Figure 5. Two Sierpiński carpet Julia sets in the Devaney family corresponding to parameters $\lambda_2^{\text{skinny}} \approx -0.35891$ (left) and $\lambda_2^{\text{fat}} \approx -1.4996769 \times 10^{-5}$ (right). The red points are the common second iterates x_λ of the finite nonzero critical points; the white points show the orbit of x_λ. The insets at lower left are local pictures in the parameter plane. Figures created with FractalStream [39].

1.8. Related work

In a study [33], Kwapisz employs a similar approach to estimating the conformal dimension for Sierpiński carpets. His analysis considers only the standard square carpet but uses very similar notions, with the slight variation that his q-resistance $r(e)$ is related to our q-length $\alpha(e)$ by

$$r(e) = \alpha(e)^{q-1}. \tag{1.4}$$

(In particular, as in this paper, there are quantities associated to the edges rather than the vertices, as contrasted with the more common notions of combinatorial q-modulus in the literature.) With this correspondence, Kwapisz's formulas match ours; for instance, his $P(\mathcal{J})$ [33, equation (1.2)] agrees with our $\left(E_q^1(\phi)\right)^q$ in equation (4.4). One difference between our approaches is that he deals with signed flows as in electrical networks, while we deal with more general unsigned tensions related to elasticity; see the discussion in [51, Appendix B].

A more substantive difference to be noted is that our q-conformal energies enjoy exact sub-multiplicativity [51, Proposition A.12], while Kwapisz only proves weak sub-multiplicativity, up to a constant [33, Theorem 1.3]. On the other hand, by only asking for weak sub-multiplicativity, he is able to get bounds on q-resistance for both the graph analogous to (but more general than) the ones we consider and its dual. Correspondingly, he gets numerical lower bounds, in addition to numerical upper bounds analogous to the ones we get. (Our numerical lower bounds rely on Theorem B, which is unlikely to be sharp in general.)

In another direction, in the non-dynamical setting of Gromov hyperbolic complexes, Bourdon and Kleiner [6] consider families of analytic invariants such as ℓ^p cohomology and separation properties of certain associated function spaces. It would be interesting to know if there is a connection to this work in our setting.

1.9. Notation

We denote the unit interval by $I := [0,1] \subset \mathbb{R}$.

Non-dynamical functions, that is, where the domain and codomain are not identified, are typically denoted using lowercase Greek letters.

If K is a constant, or set of constants, the notation $A \lesssim_K B$ means that $A/B \leq C(K)$ for some constant C depending only on K. The notation $A \asymp_K B$ means $A \lesssim_K B$ and $B \lesssim_K A$.

The study of dynamical systems includes "multi-valued" ones, which are formally pairs of maps between different spaces. We denote these spaces with subscripts, for example, $X_1 \to X_0$ in the case of general spaces or $G_1 \to G_0$ when the spaces are graphs. Such systems generate sequences of spaces and graphs which we denote by X_n and G_n, respectively, for $n = 0, 1, 2, \ldots$ In this context, we decorate maps with a superscript that corresponds to the domain and the subscript that corresponds to the codomain, for example, we have structure maps $\phi_m^n \colon G_n \to G_m$ for $n > m$ (Section 2.5). Similarly, when looking at graph energies $E_q^p(\phi)$, the domain has a p-conformal structure and the codomain has a q-conformal structure (Section 4).

We also often work with maps which are not part of the data of a dynamical system. For such maps, we typically distinguish domain and codomain by using different letters as opposed to subscripts.

2. Topological dynamics

The main result of this section is the following theorem, giving good dynamical properties of the action on a limit space.

Theorem F. *Suppose X_0, X_1 are finite CW complexes equipped with complete length metrics. Suppose $\pi, \phi \colon X_1 \rightrightarrows X_0$ is a λ-backward-contracting and recurrent virtual endomorphism, and let $f \colon \mathcal{J} \to \mathcal{J}$ be the induced dynamics on its limit space. Then*

(1) *the space $\mathcal{J}$ is connected and locally connected;*

(2) *the map $f \colon \mathcal{J} \to \mathcal{J}$ is a positively expansive covering map of degree* $\deg(\pi)$*;*

(3) *the dynamical system $f \colon \mathcal{J} \to \mathcal{J}$ is topologically coarse expanding conformal (cxc), in the sense of Haïssinsky and Pilgrim* [24].

The terminology is defined in Sections 2.1 and 2.5. One main result of [24] is that topologically cxc systems have a canonically associated nonempty set of special Ahlfors-regular metrics, called *visual metrics*. In Section 3, we will describe the geometry of $\mathcal{J}$ when equipped with such metrics.

Theorem F could be deduced as follows. Nekrashevych [38, Theorem 5.10] associates to such a virtual endomorphism a self-similar recurrent contracting group action. Such an object also has an associated limit space homeomorphic to $\mathcal{J}$ that is connected and locally connected [37, Theorem 3.5.1], establishing (1). Moreover, there is an induced dynamics on this limit space naturally conjugate to $f: \mathcal{J} \to \mathcal{J}$. Conclusions (2) and (3) then follow from [25, Theorem 6.15]. To keep this work self-contained, and because we have a later need of certain other related technical facts, we give more direct arguments.

To prove Theorem F, we present the limit space as a subspace of the infinite product $(X_1)^\infty$, following Ishii and Smillie [30]. We generalize the theory of homotopy pseudo-orbits developed there to *families* of homotopy pseudo-orbits parameterized by an auxiliary space A and show, by generalizing their results on homotopy shadowing, that such a family determines a map $A \to \mathcal{J}$ (Theorem G). We prove that the limit space is locally path-connected by applying this generalization to the case when $A = I$ is an interval (Section 2.6). To obtain the needed family of homotopy pseudo-orbits in this setting, we develop in Section 2.4 a general notion of approximate path lifting.

For technical reasons, we need to control the geometry of non-rectifiable paths. To this end, we introduce the *size* of a path, defined to be the diameter of the lift to a universal cover; see Section 2.3.

We also present the associated limit space $\mathcal{J}$ in an equivalent, but more convenient, way as an inverse limit of spaces with increasingly large diameter; see Theorem G′. To find path families, we also need to find "homotopy sections" of projections from $\mathcal{J}$. We therefore need results on homotopy shadowing and homotopy sections in those settings; see Section 2.7.

2.1. Topologically cxc systems

We first recall the notion of topologically cxc systems. To streamline our presentation, we specialize the setup of [24] to the case of self-covers. We next define positively expansive systems. Finally, we show that conclusions (1) and (2) of Theorem F imply conclusion (3).

Suppose $\mathcal{J}$ is a compact, connected, and locally connected topological space and $f: \mathcal{J} \to \mathcal{J}$ is a self-covering map. A finite open cover $\mathcal{U}_0$ of $\mathcal{J}$ by connected sets inductively generates a sequence of coverings $\mathcal{U}_n, n \in \mathbb{N}$, via the recipe

$$\mathcal{U}_{n+1} := \{\widetilde{U} \mid U \in \mathcal{U}_n, \widetilde{U} \text{ is a component of } f^{-1}(U)\}. \tag{2.1}$$

Also set

$$\mathbf{U} := \bigcup_n \mathcal{U}_n.$$

The *mesh* of a covering of a metric space is the supremum of the diameters of its elements.

Definition 2.2. The pair of the dynamical system $f: \mathcal{J} \to \mathcal{J}$ and cover $\mathcal{U}_0$ is said to satisfy Axiom [Expansion] if, for some (equivalently, any) metric on $\mathcal{J}$ compatible with the topology, the mesh of $\mathcal{U}_n$ from equation (2.1) tends to zero as $n \to \infty$. The map f satisfies [Expansion] if there is some $\mathcal{U}_0$ satisfying this condition.

Lemma 2.3. *Suppose $\mathcal{J}$ is compact, connected, and locally connected, $f: \mathcal{J} \to \mathcal{J}$ is a covering map, and $\mathcal{U}_0$ is a covering so $(f, \mathcal{U}_0)$ satisfies Axiom [Expansion]. Then this dynamical system satisfies the following two additional axioms:*

- [Degree] *For fixed n, there is a uniform upper bound on the cardinality of fibers of restrictions $f^n: \widetilde{U}_n \to U_0$ for all $U_0 \in \mathcal{U}_0$ and $\widetilde{U}_n \in \mathcal{U}_n$.*
- [Irreducibility] *For any nonempty open set $W \subset \mathcal{J}$, there is an integer N for which $f^N(W) = \mathcal{J}$.*

On the conditions of this lemma, f is *topologically cxc*. The general definition of topologically cxc is that there exists a finite open cover $\mathcal{U}_0$ such that all three axioms [Expansion], [Degree], and [Irreducibility] hold.

Proof. For a self-cover satisfying Axiom [Expansion], Axiom [Degree] clearly holds, since for large enough n the elements of $\mathcal{U}_n$ will be evenly covered.

To show [Irreducibility], fix arbitrarily a metric on $\mathcal{J}$ compatible with its topology. For $x \in \mathcal{J}$, define

$$\begin{aligned} Y(x) &:= \{y \in X \mid \text{the backward orbit of } y \text{ accumulates at } x\} \\ &= \Big\{y \in X \mid \text{there is a sequence } \tilde{y}_{n_k} \in f^{-n_k}(y) \text{ s.t. } \lim_{k\to\infty} \tilde{y}_{n_k} = x\Big\}. \end{aligned}$$

We first claim $Y(x)$ is nonempty for each x. To see this, fix $x \in \mathcal{J}$, and let y be any accumulation point of the forward orbit of x, that is, $\lim_{k\to\infty} f^{n_k}(x) = y$. Let $U \in \mathcal{U}_0$ contain y. For each sufficiently large k, let $\widetilde{U}_{n_k}$ be the unique component of $f^{-n_k}(U)$ containing x, and pick $\tilde{y}_{n_k} \in \widetilde{U}_{n_k} \cap f^{-n_k}(y)$. Then $\tilde{y}_{n_k} \to x$ since $\operatorname{diam} \widetilde{U}_{n_k} \to 0$. Hence $y \in Y(x)$.

Now suppose $y \in Y(x)$ is arbitrary and choose $U \in \mathcal{U}_0$ with $y \in U$. The same reasoning with $\operatorname{diam} \widetilde{U}_{n_k}$ shows that $U \subset Y(x)$. Since X is connected and covered by $\mathcal{U}_0$, we conclude $Y(x) = X$. Since x is arbitrary, we conclude the set of backward orbits under f of each point is dense in X.

Now fix an arbitrary nonempty open set W as in the statement of Axiom [Irreducible]. By [Expansion], there exists $U \in \mathbf{U}$ with $U \subset W$. Pick $y \in U$. By the previous paragraph, there exists a backward orbit of y accumulating at y. Pulling back U along this backward orbit, [Expansion] implies that there exists some $\widetilde{U} \in \mathbf{U}$ so that $\widetilde{U} \subset U$ and $f^m: \widetilde{U} \to U$ is a covering map for some $m > 0$. The previous paragraph implies $\bigcup_{n=0}^{\infty} f^{nm}(\widetilde{U}) = X$. By our choice of m and $\widetilde{U}$, this is an increasing union. Since X is compact, $f^{nm}(\widetilde{U}) = X$ for some n. Then $N = nm$ suffices for the statement, since $\widetilde{U} \subset W$. ■

Definition 2.4 (Positively expansive). Let $(\mathcal{J}, d)$ be a metric space. A continuous surjection $f: \mathcal{J} \to \mathcal{J}$ is *positively expansive* if there is a constant $e > 0$ so that for any $x \neq y \in \mathcal{J}$ there is an $n \geq 0$ so that $d(f^n(x), f^n(y)) > e$.

Note that a positively expansive map is locally injective.

In the case $\mathcal{J}$ is compact, positively expansive is equivalent to the following condition: there exists a neighborhood $N \supset \{(x, x) \mid x \in \mathcal{J}\}$ of the diagonal such that if $x, y \in \mathcal{J}$ satisfy $(f^i(x), f^i(y)) \in N$ for all $i \geq 0$, then $x = y$. (In particular, positively expansive is independent of the metric.) In addition, by a theorem of Reddy [1, Theorem 2.2.10] there exists a compatible metric D on $\mathcal{J}$, called an *adapted metric*, and *expansion constants* $\delta > 0$ and $0 < \lambda < 1$ such that for any $x, y \in \mathcal{J}$,

$$D(x, y) \leq \delta \implies D(f(x), f(y)) \geq \lambda^{-1} D(x, y).$$

Note that this implies f is a homeomorphism on δ-balls in the D-metric.

Proposition 2.5 (Positively expansive implies [Expansion]). *Suppose $\mathcal{J}$ is compact and locally connected, and $f: \mathcal{J} \to \mathcal{J}$ is a positively expansive covering map. Then f satisfies Axiom [Expansion].*

To prove this, we will need the following result.

Theorem 2.6 (Eilenberg constants [1, Theorem 2.1.1]). *Let X be compact and $f: X \to Y$ be a continuous surjective local homeomorphism. Then there exist two positive numbers τ and μ such that for each subset U of Y with diameter less than τ, there is a decomposition of the set $f^{-1}(U) = U_1 \cup \cdots \cup U_d$ with the following properties:*

(1) *$f: U_i \to U$ is a homeomorphism.*

(2) *For $i \neq j$, no point of U_i is closer than 2μ to a point of U_j.*

(3) *For each $\eta > 0$, there exists $0 < \varepsilon < \tau$ such that if* $\operatorname{diam} U < \varepsilon$, *then for all j,* $\operatorname{diam} U_j < \eta$.

Proof of Proposition 2.5. Equip $\mathcal{J}$ with an adapted metric. We apply Theorem 2.6 with $X = Y = \mathcal{J}$ and obtain the constant τ. Let δ be as in the definition of adapted metric, and take the constant η in Theorem 2.6 so that $\eta < \delta$; we obtain a constant ε. In summary, any open connected set U of diameter at most ε is evenly covered by f and has pre-images $\widetilde{U}_1, \ldots, \widetilde{U}_d$ of diameter at most δ. The definition of adapted metric then implies that $f: \widetilde{U}_j \to U$ expands distances by at least the factor λ^{-1}, and so the inverse branches $f_j^{-1}: U \to \widetilde{U}_j$ contract distances by at least the factor λ.

Since $\mathcal{J}$ is locally connected and compact, there is a finite open cover $\mathcal{U}_0$ by connected sets of mesh ε. The previous paragraph implies that the covering $\mathcal{U}_1$ has mesh at most $\lambda\varepsilon < \varepsilon$. Induction shows $\operatorname{mesh}(\mathcal{U}_n) < \lambda^n \varepsilon \to 0$ as required in Axiom [Expansion]. ■

2.2. Length spaces

In this subsection, we prepare for the proof of local connectivity by collecting some technical results related to covering maps and length spaces. Here, X denotes a finite, hence compact, connected CW complex, equipped with a compatible length metric. The Hopf–Rinow theorem [7, Proposition I.3.7] implies that X is a geodesic metric space. While balls might not be simply connected, they are path-connected. The *systole* is

$$\operatorname{systole}(X) := \inf\{\ell(\gamma) \mid \gamma \text{ is an essential loop in } X\};$$

if there are no essential loops in X, we set $\operatorname{systole}(X) := +\infty$. The systole is positive, since X is compact. Since X is a length space, any cover $\widetilde{X}$ inherits a lifted metric by lengths of paths. A ball of radius less than $\frac{1}{2}\operatorname{systole}(X)$ in X has no essential loops and thus is evenly covered.

Lemma 2.7. *Let X be a finite connected CW complex with a length metric. Suppose $p\colon \widetilde{X} \to X$ is a covering map, $\widetilde{X}$ is equipped with the lifted metric, and $r < \frac{1}{4}\operatorname{systole}(X)$. Then for any $\tilde{x} \in \widetilde{X}$ with $p(\tilde{x}) = x$, the restriction $p\colon B(\tilde{x}, r) \to B(x, r)$ is an isometry.*

This is standard, but we provide a proof for completeness.

Proof. The definition of the lifted metric says that p preserves the length of paths and is therefore 1-Lipschitz. Thus $p(B(\tilde{x}, r)) \subset B(x, r)$. We have $p(B(\tilde{x}, r)) = B(x, r)$ since a geodesic joining x to $y \in B(x, r)$ lifts to a path of the same length joining $\tilde{x}$ to some point $\tilde{y}$ which therefore lies in $B(\tilde{x}, r)$. We now claim that $p\colon B(\tilde{x}, r) \to B(x, r)$ is an isometry. Suppose $\tilde{a}, \tilde{b} \in B(\tilde{x}, r)$, and put $a = p(\tilde{a})$ and $b = p(\tilde{b})$. Then $a, b \in B(x, r)$. Consider the piecewise geodesic path γ comprised of 3 length-minimizing segments which runs from x to a, then from a to b, and then from b to x. This loop may not lie in $B(x, r)$. However, $\ell(\gamma) < 4r$ and both endpoints are at x, so $\gamma \subset B(x, 2r)$, which is evenly covered by the choice of r. It follows that the middle segment lifts to a segment joining $\tilde{a}$ to $\tilde{b}$ of length equal to $d(a, b)$. Hence $d(\tilde{a}, \tilde{b}) \leq d(a, b)$ and the result is proved. ■

If $\phi\colon X \to Y$ is a map between metric spaces, we say a non-decreasing function $\omega_\phi\colon [0, \operatorname{diam}(X)] \to [0, \operatorname{diam}(Y)]$ is a *modulus of continuity* if, for all $E \subset X$, $\operatorname{diam}\phi(E) \leq \omega_\phi(\operatorname{diam} E)$.

Lemma 2.8. *Let X, Y be compact, connected, CW complexes equipped with length metrics, and let $\phi\colon X \to Y$ be a continuous map, with modulus of continuity ω_ϕ. Let $p_Y\colon \widetilde{Y} \to Y$ be any covering map, and equip $\widetilde{Y}$ with the lifted metric. Let $p_X\colon \widetilde{X} \to X$ and $\widetilde{\phi}\colon \widetilde{X} \to \widetilde{Y}$ be the maps induced by pullback, and equip $\widetilde{X}$ with the lifted metric from X. Then $\widetilde{\phi}$ is uniformly continuous, with modulus of continuity $\omega_{\widetilde{\phi}}$ independent of the cover p_Y. Indeed, there exists $r_0 > 0$ depending only on $\phi\colon X \to Y$ so that we can take*

$$\omega_{\widetilde{\phi}}(\delta) = \begin{cases} \omega_\phi(\delta) & \text{if } \delta \leq r_0, \\ \omega_\phi(r_0)\left\lceil \dfrac{d(a,b)}{r_0} \right\rceil & \text{if } \delta > r_0. \end{cases}$$

So $\widetilde{\phi}$ behaves just like ϕ at small scales and is Lipschitz at large scales.

Proof. Put $s_0 := \frac{1}{4}\,\mathrm{systole}(Y)$. By the uniform continuity of ϕ, there exists $0 < r_0 < \frac{1}{4}\,\mathrm{systole}(X)$ such that for each $x \in X$ and $y = \phi(x)$, we have $\phi(B(x, r_0)) \subset B(y, s_0)$. Fix now $\widetilde{x} \in \widetilde{X}$ and put $x = p_X(\widetilde{x})$, $\widetilde{y} = \widetilde{\phi}(\widetilde{x})$, and $y = p_Y(\widetilde{y}) = \phi(x)$. By Lemma 2.7,

$$p_Y^{-1}\colon B(y, s_0) \to B(\widetilde{y}, s_0), \quad p_X\colon B(\widetilde{x}, r_0) \to B(x, r_0)$$

are isometric homeomorphisms. Now fix $0 < r \le r_0$. Then

$$\begin{aligned}\widetilde{\phi}(B(\widetilde{x}, r)) &= (p_Y^{-1} \circ \phi \circ p_X)(B(\widetilde{x}, r)) \\ &= p_Y^{-1}(\phi(B(x, r))) \subset p_Y^{-1}(B(y, \omega_\phi(r))) \subset B(\widetilde{y}, \omega_\phi(r)),\end{aligned}$$

establishing the estimate in the case $\delta \le r_0$. For the other case, suppose $a, b \in \widetilde{X}$ are at distance $R \ge r_0$, and let γ be a geodesic joining a to b. Divide γ into sub-segments $\gamma = \gamma_2 * \gamma_1 * \cdots * \gamma_n$ with $\ell(\gamma_i) \le r_0$ for $i = 1, \ldots, n$ so that $n = \lceil \ell(\gamma)/r_0 \rceil$. Then

$$d(\widetilde{\phi}(a), \widetilde{\phi}(b)) \le \sum_{i=0}^{n} \mathrm{diam}\,\phi(\gamma_i) \le \omega_\phi(r_0) \left\lceil \frac{d(a,b)}{r_0} \right\rceil. \qquad \blacksquare$$

2.3. Sizes of paths and traces of homotopies

It would be nice to always work with the length of paths, but it turns out that not all the paths we consider are rectifiable. (In particular, we consider paths in the Julia set $\mathcal{J}$ and their projections to the finite approximations G_n; these projections are usually not rectifiable.) We could consider the diameter of paths, but we also need to lift paths to covers. We work instead with a hybrid.

Convention 2.9. For paths $\gamma\colon I \to X$, the path $\overline{\gamma}$ is the reversed path, and $\gamma_1 * \gamma_2$ denotes composition of paths, defined when $\gamma_1(1) = \gamma_2(0)$. For homotopies $H\colon I \times A \to X$, we will more generally use the same notations $\overline{H}$ and $H_1 * H_2$, always operating on the first input (which is an interval).

Definition 2.10. For X a locally simply connected length space, A a simply connected auxiliary space (usually the interval), and $\gamma\colon A \to X$ a continuous map, there are lifts $\widetilde{\gamma}\colon A \to \widetilde{X}$ of γ to the universal cover of X. The *size* of γ is the diameter of $\widetilde{\gamma}$ with respect to the lifted metric on $\widetilde{X}$:

$$\mathrm{size}(\gamma) := \max_{s,t \in I} d_{\widetilde{X}}(\widetilde{\gamma}(s), \widetilde{\gamma}(t)).$$

The proof of the following lemma is straightforward.

Lemma 2.11 (Properties of size for paths). *The notion of size for paths (with A the interval) satisfies the following properties:*

(0) *Well defined:* size(γ) *is independent of which lift of* γ *to* $\widetilde{X}$ *you take.*

(1) *Bounded by length: when* $\gamma\colon I \to X$ *is rectifiable,* $\operatorname{size}(\gamma) \le \ell(\gamma)$.

(2) *Invariance under lifts: if* $p\colon X \to Y$ *is a covering map,* Y *is a length space,* X *is equipped with the lifted metric,* $\gamma\colon I \to Y$ *is a path, and* $\widetilde{\gamma}\colon I \to X$ *is a lift of* γ *under* p*, then* $\operatorname{size}(\widetilde{\gamma}) = \operatorname{size}(\gamma)$.

(3) *Sub-additive under path composition:* $\operatorname{size}(\gamma_1 * \gamma_2) \le \operatorname{size}(\gamma_1) + \operatorname{size}(\gamma_2)$.

(4) *Shortening: if* $f\colon X \to Y$ *is* λ*-Lipschitz and* $\gamma\colon I \to X$ *is a path in* X*, then* $\operatorname{size}(f \circ \gamma) \le \lambda \cdot \operatorname{size}(\gamma)$.

As a result of point (1), we will prefer to give statements with hypotheses on the *size* of paths and construct paths with bounds on *length*, even if we do not necessarily need the length bounds for our applications.

Another central feature of our development is the following.

Definition 2.12. For X a locally simply connected length space, A an auxiliary space—now not necessarily simply connected—and $H\colon I \times A \to X$ a homotopy of maps from A to X, a *trace* of H is a path of the form $t \mapsto H(t,a)$ for fixed $a \in A$. The *trace size* of H is the maximum size of a trace:

$$\operatorname{tracesize}(H) := \sup_{a\in A} \operatorname{size}(H(\cdot,a)).$$

If two maps $f, g\colon A \to X$ are homotopic by a homotopy of trace size at most K, then we write $f \sim_K g$.

For a homotopy $H\colon I \times I \to X$ between two paths γ_0 and γ_1, be careful to distinguish between its *size* and *trace size*. For instance, if H has bounded size, then the γ_i must also have bounded size, while two paths that are very long can still have a homotopy of bounded trace size.

The trace size of a homotopy between two paths is sensitive to the parameterization of the domain of the two paths, which in turn is sensitive to details like exactly how one defines the concatenation operation on paths. We will specify the parameterization when necessary.

One thing to note now is that, if β is a path of size R and any $\varepsilon > 0$, then, for any concatenatable γ_1, γ_2 and suitable parameterization of the domain,

$$\gamma_1 * \beta * \overline{\beta} * \gamma_2 \sim_{R+\varepsilon} \gamma_1 * \gamma_2.$$

(The parameterization to make this work uses a very small interval in the domain for $\beta * \overline{\beta}$ on the left-hand side.)

Lemma 2.13. *Suppose* X, Y, Z *are length spaces. If* $f\colon X \to Y$ *and* $g_0, g_1\colon Y \to Z$ *are maps with* $g_0 \sim_K g_1$ *via the homotopy* $\alpha\colon I \times Y \to Z$*, then* $g_0 \circ f \sim_K g_1 \circ f$ *via the homotopy* $f^*\alpha\colon I \times X \to Z$ *given by* $(t,x) \mapsto \alpha(t, f(x))$.

Proof. The only nontrivial point is the bound on the trace size, which follows since every trace of $f^*\alpha$ is a trace of α. ∎

2.4. Approximate path lifting

Suppose X, Y are length spaces, A is a topological space, and suppose we are given maps $\phi\colon X \to Y$ and $g\colon A \to Y$. An *approximate lift* of g under ϕ with constant K is a map $g'\colon A \to X$ such that $\phi \circ g' \sim_K g$.

Now suppose further that X, Y are finite connected CW complexes, equipped with complete length metrics, and $\phi\colon X \to Y$ is continuous and surjective on the fundamental group. We consider the problem of approximately lifting paths in Y under ϕ to paths in X.

In this subsection, the constants appearing in the conclusions depend on $\phi\colon X \to Y$ and on the other constants appearing in the statements. We suppress their dependence on ϕ.

Proposition 2.14 (Controlled approximate path lifting). *Suppose X and Y are finite connected CW complexes equipped with length metrics, $\phi\colon X \to Y$ is continuous, and $\phi_*\colon \pi_1(X) \to \pi_1(Y)$ is surjective. Then there exists $K > 0$ so that for any path $\gamma\colon I \to Y$ joining endpoints y_0 to y_1, and any preimages $x_i \in \phi^{-1}(y_i)$ for $i = 0, 1$ of these endpoints, there exists an approximate lift $\gamma'\colon I \to X$ with $\gamma'(i) = x_i$ for $i = 0, 1$ with homotopy $H\colon I \times I \to Y$ between γ and $\phi \circ \gamma'$ of trace size at most K. Concretely,*

- $H(0,t) = \gamma(t)$ *and* $H(1,t) = \phi \circ \gamma'(t)$ *for* $0 \le t \le 1$*;*
- $H(s,i) = y_i$ *for* $i = 0, 1$ *and* $0 \le s \le 1$*;*
- $\operatorname{size}(s \mapsto H(s,t)) \le K$ *for each* $t \in [0,1]$.

There exist constants C_0 and C_1 so that if γ is rectifiable, then so is γ', and $\ell(\gamma') < C_0 + C_1\ell(\gamma)$.

In other words, approximate lifting increases lengths by controlled amounts, and the failure of a path to lift is measured by a homotopy whose traces are uniformly bounded in size, independent of the path γ; see Figure 6. (We do not use the fact that lengths are increased by controlled amounts in this paper, but it helps add motivation.) We will also say that $\phi\colon X \to Y$ satisfies the *K-APL* condition.

Before proving the general statement, we first prove it for loops γ of bounded length.

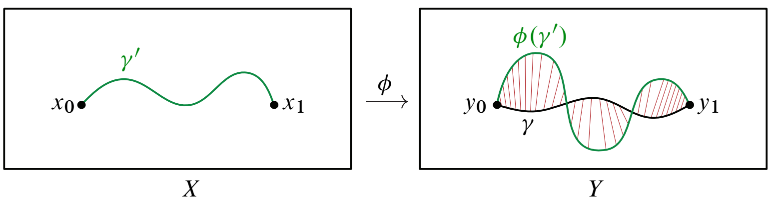

Figure 6. Approximate path lifting. The traces of the homotopies are the red lines on the right; these must be of bounded size.

Lemma 2.15 (Controlled approximate loop lifting). *In the setup of Proposition 2.14, fix $C > 0$. Then there exist constants L and K with the following property. For any $x \in X$, $y = \phi(x)$, and loop $\lambda\colon I \to Y$ based at y with* $\operatorname{size}(\lambda) \le C$, *there is a rectifiable loop $\lambda'\colon I \to X$ based at x with $\ell(\lambda') < L$ so that $\phi \circ \lambda' \sim_K \lambda$.*

Proof. First, fix $x \in X$ and $y = \phi(x)$. The length metrics on X and Y induce norms on the fundamental groups $\pi_1(X, x)$ and $\pi_1(Y, y)$. Since $\operatorname{size}(\lambda) \le C$, the norm of $[\lambda] \in \pi_1(Y, y)$ is also bounded by C. Since $\phi_*\colon \pi_1(X, x) \to \pi_1(Y, y)$ is surjective, there exists $L(x)$ such that the image of the ball of radius $L(x)$ in $\pi_1(X, x)$ contains the ball of radius C in $\pi_1(Y, y)$. Now vary $x \in X$. Since the induced norms on $\pi_1(X, x)$ and $\pi_1(Y, \phi(x))$ vary continuously as a function of x, we can take L to be continuous, and so, by compactness, $\sup_{x \in X} L(x)$ is finite; call this supremum L. Thus there exists a loop λ' based at x of length at most L for which $\phi \circ \lambda' \sim \lambda$.

We must bound the trace size of the homotopy; in fact, we bound its size. For any λ and λ' as above, we can lift $\phi \circ \lambda'$ and λ to paths in the universal cover $\widetilde{Y}$. Fix lifts $\widetilde{\phi \circ \lambda'}$ and $\widetilde{\lambda}$, respectively, joining common endpoints. Then the concatenation $\widetilde{\alpha} := \widetilde{\phi \circ \lambda'} * \overline{\widetilde{\lambda}}$ is a loop in the simply connected length CW complex $\widetilde{Y}$, so, since $\widetilde{\phi}$ is uniformly continuous,

$$\operatorname{diam}(\widetilde{\alpha}) \le \operatorname{diam}(\widetilde{\phi \circ \lambda'}) + \operatorname{diam}(\widetilde{\lambda}) \le \omega_{\widetilde{\phi}}(L) + C =: D_1.$$

By [19, Lemma 9.51], $\widetilde{Y}$ is uniformly simply connected. This implies that $\widetilde{\alpha}$ is homotopic to a constant map via a homotopy whose image has diameter $K := K(D_1)$, as desired. ∎

Proof of Proposition 2.14. Pick (arbitrarily) a basepoint $x_* \in X$, set $y_* := \phi(x_*)$, and pick an arbitrary constant $C > 0$. Divide γ into sub-paths

$$\gamma = \gamma_0 * \cdots * \gamma_{n-1}$$

with $\operatorname{size}(\gamma_i) \le C$. Let the endpoints of γ_i be $z_i, z_{i+1} \in Y$, and for each i pick a path ρ_i from y_* to z_i, of length less than $\operatorname{diam}(Y)$. Pick also paths ρ'_0, ρ'_n from x_* to x_0, x_1, respectively, of length less than $\operatorname{diam}(X)$. Then $\lambda_i := \rho_i * \gamma_i * \overline{\rho_{i+1}}$ is a loop based at y_* of size less than $2\operatorname{diam}(X) + C$. By Lemma 2.15, there are constants L_1, K_1 so that, for each i, there is a loop λ'_i in X based at x_* with $\ell(\lambda'_i) < L_1$ and $\phi \circ \lambda'_i \sim_{K_1} \lambda_i$. In addition, $\sigma_0 := (\phi \circ \rho'_0) * \overline{\rho_0}$ is a loop of size less than $\omega_{\widetilde{\phi}}(\operatorname{diam} Y) + \operatorname{diam} X$, so there are constants L_2, K_2 and a loop σ'_0 based at x_* of length less than L_2 so that $\phi \circ \sigma'_0 \sim_{K_2} \sigma_0$. Similarly, pick an approximate lift σ'_n of $\sigma_n := \rho_n * \overline{(\phi \circ \rho'_n)}$. Now set

$$\begin{aligned} K_3 &:= \max\{K_1, K_2\} \\ D &:= \max\{\operatorname{diam}(X), \omega_{\widetilde{\phi}}(\operatorname{diam}(Y))\} \\ \gamma' &:= \overline{\rho'_0} * \sigma'_0 * \lambda'_1 * \lambda'_2 * \cdots * \lambda'_{n-1} * \sigma'_n * \rho'_n \end{aligned}$$

(with suitable parameterization of the domain for γ') so that

$$\begin{aligned} \phi \circ \gamma' &\sim_{K_3} \overline{(\phi \circ \rho'_0)} * \sigma_0 * \lambda_1 * \lambda_2 * \cdots * \lambda_{n-1} * \sigma_n * (\phi \circ \rho'_n) \\ &\sim_{D+\varepsilon} \gamma_0 * \gamma_1 * \cdots * \gamma_{n-1} = \gamma, \end{aligned}$$

as desired.

To get the bounds on length in the case that γ is rectifiable, choose the initial decomposition of γ into sub-paths γ_i so that, for $i > 0$, $\ell(\gamma_i) \geq C$; then we have $n \leq 1 + \ell(\gamma)/C$. We can choose the paths λ_i', ρ_i', and σ_i' all to have length bounded by a constant, which then gives the desired bound on $\ell(\gamma)$. ∎

Remark 2.16. Lemma 2.15 gives a bound on the overall size of the homotopy, but Proposition 2.14 only gives a bound on the trace sizes.

2.5. Multi-valued dynamical systems

We turn our attention back to dynamics. We think of two spaces and a pair of maps between them

$$X_0 \xleftarrow{\;\phi\;} X_1 \xrightarrow{\;\pi\;} X_0$$

as a *multi-valued dynamical system*. We introduce an associated limit space and describe it in two different ways, as in [30], but adopting slightly different notation. Here is how to translate between their (IS) and our (PT) notation:

$$\begin{aligned} X^i_{\mathrm{IS}} &= X_{i,\mathrm{PT}}; \\ (\iota,\sigma)_{\mathrm{IS}} = (\iota, f)_{\mathrm{IS}} &= (\phi,\pi)_{\mathrm{PT}}; \\ X^{+\infty}_{\mathrm{IS}} &= \mathcal{J}_{\mathrm{PT}}. \end{aligned}$$

(There are also minor differences in the indexing.)

Suppose X_1 and X_0 are compact topological spaces and $\pi, \phi\colon X_1 \rightrightarrows X_0$ are two continuous maps. An *orbit* is a sequence $x = (x_1, x_2, \ldots)$ of points $x_i \in X_1$ with $\pi(x_i) = \phi(x_{i+1})$ when both sides are defined. If x_i is defined for $1 \leq i \leq n$ for some $n \geq 1$, we get the space X_n of orbits of length n. If x_i is defined for all $1 \leq i$, we get the *limit space* $\mathcal{J} \subset X_1^{\mathbb{N}}$ of one-sided infinite orbits. Note the typographical distinction between the abstract limit space $\mathcal{J}$ of a virtual endomorphism and the concrete limit space $J_R \subset \widehat{\mathbb{C}}$ of a rational map R.

With this setup, there are two families of canonical maps

$$\begin{aligned} \phi^{n+1}_n, \pi^{n+1}_n &\colon X_{n+1} \to X_n, \\ \phi^{n+1}_n(x_1, x_2, \ldots, x_n, x_{n+1}) &:= (x_1, x_2, \ldots, x_n), \\ \pi^{n+1}_n(x_1, x_2, \ldots, x_n, x_{n+1}) &:= (x_2, \ldots, x_n, x_{n+1}). \end{aligned}$$

We also set $\phi^1_0 := \phi$ and $\pi^1_0 := \pi$. We can compose these to get maps $\phi^n_k, \pi^n_k\colon X_n \to X_k$ for $0 \leq k < n$. We recall a convention from Section 1.9.

Convention 2.17. In our indexing convention, the index of the domain appears as a superscript and the index of the codomain appears as a subscript. This way, composition corresponds to "contraction" of indices, as is conventional in tensor notation.

We can present the space $\mathcal{J}$ as an inverse limit of the sequence ϕ_n^{n+1}:

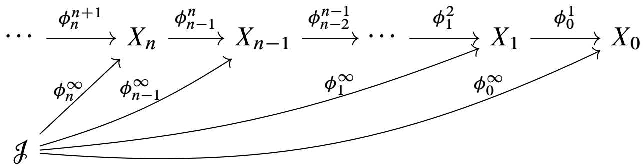

where the $\phi_n^\infty\colon \mathcal{J} \to X_n$ are analogues of ϕ_k^n:

$$\phi_n^\infty(x_1, x_2, \ldots) := (x_1, x_2, \ldots, x_n).$$

There is also a canonical map $f\colon \mathcal{J} \to \mathcal{J}$ induced by the one-sided shift, a kind of analogue of π_n^{n+1}:

$$f(x_1, x_2, x_3, \ldots) := (x_2, x_3, \ldots).$$

There are two natural modifications of a pair of maps (π, ϕ): we can *iterate* it, replacing the pair by

$$\pi_0^n, \phi_0^n\colon X_n \rightrightarrows X_0, \tag{2.18}$$

or we can *reindex* it, replacing the pair by

$$\pi_{n-1}^n, \phi_{n-1}^n\colon X_n \rightrightarrows X_{n-1}. \tag{2.19}$$

Neither of these operations changes the limit space $\mathcal{J}$, but iterating replaces the dynamics of f on $\mathcal{J}$ by f^n, while reindexing does not change f.

We now restrict attention to expanding systems, as in [30], but adopting terminology of Nekrashevych [38] and the second author [52]. We continue with some definitions.

Definition 2.20 (Virtual endomorphisms). A pair of continuous maps $\pi, \phi\colon X_1 \rightrightarrows X_0$ between topological spaces is a *virtual endomorphism* if π is a covering map of finite degree.

Convention 2.21. In this section, we are exclusively concerned with virtual endomorphisms, where X_0, X_1 are finite connected CW complexes, X_0 is equipped with a complete length metric d_0, and X_1 is equipped with the length metric d_1 induced by the covering π, that is, the length metric on X_1 so that π is a local isometry.

Definition 2.22. Suppose $0 < \lambda < 1$. The virtual endomorphism $\pi, \phi\colon X_1 \rightrightarrows X_0$ is λ-*backward-contracting* if $d_0(\phi(a), \phi(b)) \le \lambda d_1(a, b)$ for all $a, b \in X_1$, that is, ϕ is a uniform contraction. Equivalently, we say the virtual endomorphism is λ^{-1}-*forward-expanding*. A virtual endomorphism is backward contracting (equivalently, forward expanding) if it is λ-backward-contracting for some $0 < \lambda < 1$.

Definition 2.23. The virtual endomorphism $\pi, \phi\colon X_1 \rightrightarrows X_0$ is *recurrent* if X_0 and X_1 are connected and $\phi_*\colon \pi_1(X_1) \to \pi_1(X_0)$ is surjective.

If $\pi, \phi\colon X_1 \rightrightarrows X_0$ is recurrent, then X_n is connected for each n.

In the setting of a virtual endomorphism, the map $\pi_n^{n+1}\colon X_{n+1} \to X_n$ defined above is also a covering map. We will use the fact that X_{n+1} is a pullback, which concretely gives the following lemma, among other pullback diagrams.

Lemma 2.24. *Given maps $F\colon A \to X_1$ and $G\colon A \to X_n$ for which $\phi \circ F = \pi_0^n \circ G$, there exists a unique map $A \to X_{n+1}$ such that the following diagram commutes:*

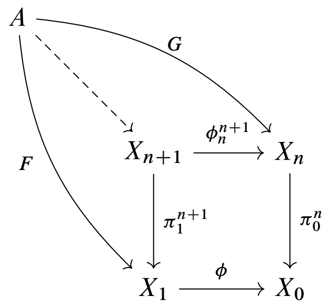

Via pullback of length metrics under π_0^n, each space X_n inherits a length metric d_n. The degree of π_0^n is d^n, and the complexes X_n are locally finite uniformly in n, so $\operatorname{diam} X_n \to \infty$ as $n \to \infty$.

Remark 2.25. From a virtual endomorphism $\pi, \phi\colon X_1 \rightrightarrows X_0$, we will not see every cover of X_0 among the covers $\pi_0^n\colon X_n \to X_0$ (or their normal closures). Studying the exact covers that appear leads to a very interesting group, the *iterated monodromy group* $\mathrm{IMG}(\pi, \phi)$, which is a natural quotient of $\pi_1(X_0)$. This group is the deck group of a Galois cover $\widetilde{X}_0'$ of X_0 usually different from the universal cover. Instead of measuring the size by lifting paths to the universal cover as in Definition 2.10, we could get a different (smaller) notion by lifting to $\widetilde{X}_0'$ instead. The difference is inessential for this paper.

Definition 2.26. A *homotopy pseudo-orbit* $(x, \alpha) = (x_i, \alpha_i)_{i \geq 1}$ is a sequence of points $x_i \in X_1$ and of paths $\alpha_i\colon [0,1] \to X_0$ such that

- $\alpha_i(0) = \pi(x_i)$, $\alpha_i(1) = \phi(x_{i+1})$;
- $\ell(\alpha_i) \leq C$ for some $C \geq 0$ independent of i.

Two homotopy pseudo-orbits $(x, \alpha) = (x_i, \alpha_i)_{i \geq 1}$ and $(x', \alpha') = (x_i', \alpha_i')_{i \geq 1}$ are *homotopic* if there is a sequence $\beta = (\beta_i)_{i \geq 1}$ of paths $\beta_i\colon [0,1] \to X_1$ with

- $\beta_i(0) = x_i$ and $\beta_i(1) = x_i'$;
- for $i \geq 1$, $\alpha_i * (\phi \circ \beta_{i+1}) \sim (\pi \circ \beta_i) * \alpha_i'$;
- $\ell(\beta_i) < B$ for some $B > 0$ independent of i.

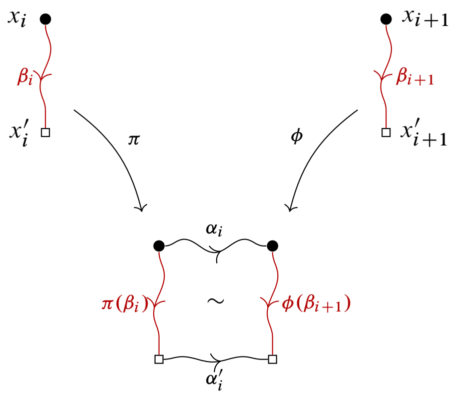

Figure 7. In homotopic pseudo-orbits, $\alpha_i * \phi(\beta_{i+1}) \sim \pi(\beta_i) \cdot \alpha'_i$.

See Figure 7 and [30, Definition 6.4, Figure 5a]. Following Ishii and Smillie, we use lengths in the condition on α_i and β_i, but size would work equally well, since these are paths in a CW complex.

The following result appears in [30, §7].

Theorem 2.27 (Homotopy shadowing). *Suppose $\pi, \phi\colon X_1 \rightrightarrows X_0$ is forward expanding. Then every homotopy pseudo-orbit is homotopic to an orbit, and this orbit is unique.*

We will need a generalization of the homotopy shadowing theorem to a setting, where the orbit depends on a parameter a lying in a space A. We develop this notion in close parallel to the above notions and to those of Ishii and Smillie.

Definition 2.28. For $\pi, \phi\colon X_1 \rightrightarrows X_0$ a virtual endomorphism between locally simply connected length spaces and A an auxiliary space, a *family (x, α) of homotopy pseudo-orbits parameterized by A of trace size at most K* is a sequence $x := (x_i\colon A \to X_1)_{i \geq 1}$ of maps and a sequence $\alpha := (\alpha_i\colon I \times A \to X_0)_{i \geq 1}$ of homotopies, so that

(1) α_i is a homotopy from $\pi \circ x_i$ to $\phi \circ x_{i+1}$, in the sense that $\alpha_i(0, \cdot) = \pi \circ x_i$ and $\alpha_i(1, \cdot) = \phi \circ x_{i+1}$;

(2) there exists a constant $K < \infty$ so that for each $i \geq 1$, we have

$$\operatorname{tracesize}(\alpha_i) \leq K.$$

See Figure 8.

Definition 2.29. Two families (x, α) and (x', α') of homotopy pseudo-orbits parameterized by A are *homotopic* if there exist a constant B and a sequence $\beta = (\beta_i)_{i \geq 1}$ of homotopies $\beta_i\colon I \times A \to X_1$ such that

(1) β_i is a homotopy from x_i to x'_i, in the sense that $\beta_i(0, \cdot) = x_i$ and $\beta_i(1, \cdot) = x'_i$;

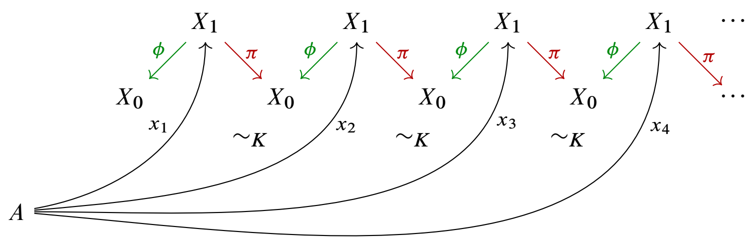

Figure 8. A homotopy pseudo-orbit of maps of A.

(2) for each $i \geq 1$, we have $\operatorname{tracesize}(\beta_i) \leq B$;

(3) for each $i \geq 1$, the map $\alpha_i * (\phi \circ \beta_{i+1})$ is homotopic to $(\pi \circ \beta_i) * \alpha_i'$ in the following sense. There is a map $H_i\colon I \times I \times A \to X_0$ such that $H_i(0,\cdot,\cdot) = \alpha_i * (\phi \circ \beta_{i+1})$, $H_i(1,\cdot,\cdot) = (\pi \circ \beta_i) * \alpha_i'$, and for all $0 \leq t \leq 1$, $H_i(t,0,\cdot) = \pi \circ x_i$ and $H_i(t,1,\cdot) = \phi \circ x_{i+1}'$.

These conditions in (3) guarantee that the homotopic squares appearing in Figure 7 remain squares throughout the interpolating maps $H_i(\cdot,\cdot,a)$. We do not require a size bound on the homotopies H_i in condition (3).

Here is our generalization of Ishii–Smillie's homotopy shadowing result.

Theorem G. *Suppose $\pi, \phi\colon X_1 \rightrightarrows X_0$ is a λ-backward-contracting virtual endomorphism of CW complexes with induced dynamics $f\colon \mathcal{J} \to \mathcal{J}$ on its limit space. Let $(x,\alpha) = (x_i,\alpha_i)_{i\geq 1}$ be a family of homotopy pseudo-orbits parameterized by A of trace size at most K.*

(1) *There exists*

  (a) *a map $x^\infty\colon A \to \mathcal{J}$, that is, a family of orbits of f parameterized by A;*

  (b) *a homotopy $\beta = (\beta_i^\infty)_{i\geq 1}$ from (x,α) to x^∞.*

(2) *We have* $\operatorname{tracesize}(\beta_i^\infty) \leq K' := K/(1-\lambda)$ *for all $i \geq 1$.*

Furthermore, x^∞ is unique among all such maps with uniformly bounded $\operatorname{tracesize}(\beta_i^\infty)$.

To make sense of part (1) (b) of the statement, we regard an orbit as a family of homotopy pseudo-orbits parameterized by A, namely $(x_i^\infty, \alpha_i^\infty)_{i\geq 1}$, where each α_i^∞ is a constant homotopy.

Proof. This is a straightforward modification of the proof of Theorem 7.1 of [30]. Since we will need the notation later, we copy their proof, more or less word for word, with slight adjustments to indexing. We let the pseudo-orbit x now depend on a parameter $a \in A$ so that $x := (x_i\colon A \to X_1)_{i\geq 1}$, and we denote the collection of homotopies by $\alpha := (\alpha_i\colon I \times A \to X_0)_{i\geq 1}$. To ease notation, we think of our homotopies α_i and β_j below as paths in the space of continuous maps from A to X_0 and X_1, respectively.

We inductively define a sequence of families of homotopy pseudo-orbits with successively smaller traces as follows. Set $x_i^0 := x_i$ and $\alpha_i^0 := \alpha_i$. Suppose that a family of homotopy pseudo-orbits $(x_i^n, \alpha_i^n)_{i\geq 1}$ is defined. Then, since π is a covering and $\alpha_i^n(0,\cdot) = \pi \circ x_i^n$, there exists a unique lift $\beta_i^n\colon I \times A \to X^1$ of α_i^n by π so that $\beta_i^n(0,\cdot) = x_i^n$, by the homotopy lifting property of covering maps. Put $\alpha_i^{n+1} := \phi \circ \beta_{i+1}^n$ and $x_i^{n+1} := \beta_i^n(1)$. Then, we have $\pi \circ x_i^{n+1} = \pi \circ \beta_i^n(1) = \alpha_i^n(1) = \phi(x_{i+1}^n) = \phi(\beta_{i+1}^n(0)) = \alpha_i^{n+1}(0)$ and $\phi \circ x_{i+1}^{n+1} = \phi \circ \beta_{i+1}^n(1) = \alpha_i^{n+1}(1)$. This means that, once we verify a trace size bound, $(x_i^{n+1}, \alpha_i^{n+1})_{i\geq 1} =: (x^{n+1}, \alpha^{n+1})$ is a family of homotopy pseudo-orbits.

Contraction implies that the length of the traces of the homotopies α_i^n is bounded by $K\lambda^n$ for $n \geq 1$; this is [30, Lemma 7.2]. Concatenating the homotopies α_i^n for $n = 1, 2, \ldots$ and scaling the time parameters in the homotopy to consecutive intervals in $[0, 1)$ as in their proof, we obtain, for $i \geq 1$, a sequence of maps $\alpha_i^\infty\colon [0,1) \times A \to X_0$ of trace size $K' := K/(1-\lambda)$.

To get a map defined on $[0, 1] \times A$, we need to say a little more. First, as in their proof, for fixed $a \in A$, the path $\alpha_i^\infty(t, a)$ is Cauchy as $t \to 1$ in the sense that, for any $\varepsilon > 0$, there is $\delta < 1$ so that for $t_0, t_1 > \delta$, we have $d_{X_0}(\alpha_i^\infty(t_0, a), \alpha_i^\infty(t_1, a)) < \varepsilon$. Furthermore, these paths are uniformly Cauchy as a varies. There is therefore a well-defined limit $\alpha_i^\infty(1, a)$, and the continuous functions $\alpha_i^\infty(t,\cdot)\colon A \to X_0$ converge uniformly to $\alpha_i^\infty(1,\cdot)$. By the uniform limit theorem, the limiting function α_i^∞ restricted to $\{1\} \times A$ is therefore continuous. The standard proof of the uniform limit theorem shows that, in fact, we get a continuous function $\alpha_i^\infty\colon [0,1] \times A \to X_0$, as desired.

We also have sequences of maps $\beta_i^n\colon I \times A \to X_1$ and concatenating the β_i^n's for $n = 1, 2, 3, \ldots$ and extending by the uniform limit theorem yield β_i^∞, a lift of α_i^∞ under π. We put $x_i^\infty := \beta_i^\infty(\cdot, 1)$ and note that x^∞ defines a family of orbits, that is, a map $x^\infty\colon A \to \mathcal{J}$; the associated homotopies in Definition 2.28 are constant. By construction, β^∞ gives a homotopy between the family of homotopy pseudo-orbits (x, α) and the family of orbits x^∞. Since π is an isometry, the trace sizes of the β_i^∞ are bounded by K', as required. ∎

In our later applications, it will be useful to restate Theorem G in terms of maps into the X_n. For its proof, we will need the following lemma, illustrated in Figure 9.

Lemma 2.30. *Suppose (x, α) is a family of homotopy pseudo-orbits parameterized by A, with trace size at most K. Then for each $n \geq 1$, there are unique lifted maps $\widetilde{x}_n\colon A \to X_n$ and homotopies $\widetilde{\alpha}_n\colon I \times A \to X_n$ so that $\pi_1^n \circ \widetilde{x}_n = x_n$, $\pi_0^n \circ \widetilde{\alpha}_n = \alpha_n$, and $\widetilde{\alpha}_n$ is a homotopy from $\widetilde{x}_n$ to $\phi_n^{n+1} \circ \widetilde{x}_{n+1}$ of trace size at most K.*

Proof. We proceed by induction on n, first constructing $\widetilde{\alpha}_n$ and then $\widetilde{x}_{n+1}$. Recall that $\alpha_n(0) = \pi \circ x_n$ and $\alpha_n(1) = \phi \circ x_{n+1}$. We start by setting $\widetilde{x}_1 := x_1$. If we have constructed $\widetilde{x}_n$, then by unique homotopy lifting applied to the covering map π_0^n, there is a unique function $\widetilde{\alpha}_n$ that is a lift of α_n with starting point $\widetilde{x}_n$; since $\widetilde{\alpha}_n$ is a lift, the trace size is still K, as desired. We next construct $\widetilde{x}_{n+1}$. Let $\widetilde{x}_{n+1,n}\colon A \to X_n$ be $\widetilde{\alpha}_n(1)$. Then, by the defining property of α_n, we have $\pi_0^n \circ \widetilde{x}_{n+1,n} = \phi \circ x_{n+1}$. Since we have a pullback square (Lemma 2.24), there is a unique map $\widetilde{x}_{n+1}\colon A \to X_{n+1}$ compatible with these two projections. ∎

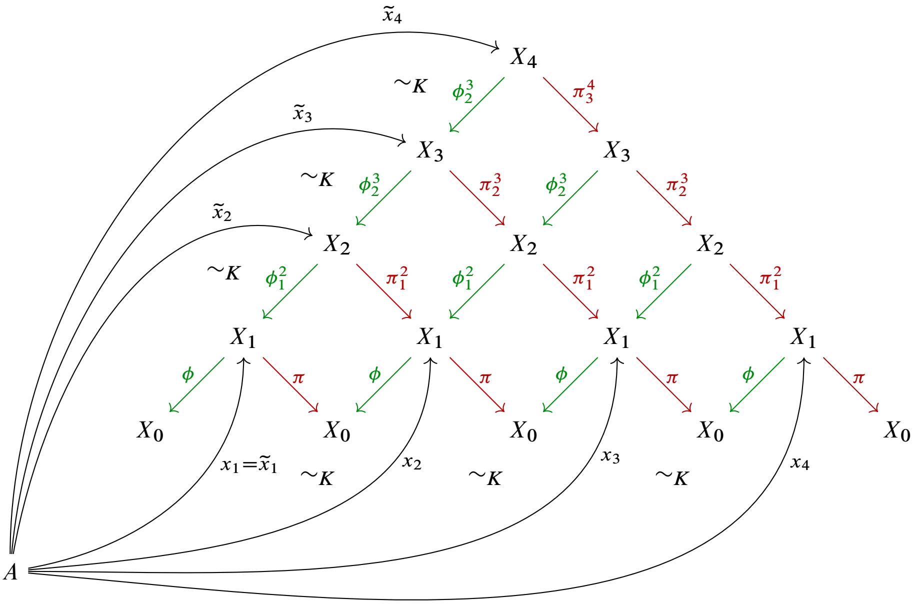

Figure 9. Lifting a family of homotopy pseudo-orbits parameterized by A to covering spaces.

Theorem G′. *Suppose $\pi, \phi\colon X_1 \rightrightarrows X_0$ is a λ-backward-contracting virtual endomorphism of CW complexes with limit space $\mathcal{J}$, and $(x, \alpha) = (x_i, \alpha_i)_{i \geq 1}$ is a family of homotopy pseudo-orbits parameterized by A of trace size at most K. Let x^∞ and β be the families of orbits parameterized by A and homotopies given by Theorem G, and let $\widetilde{x} := (\widetilde{x}_n\colon A \to X_n)_{n\geq 1}$ and $\widetilde{x}^\infty = (\widetilde{x}_n^\infty\colon A \to X_n)_{n\geq 1}$ be the families of maps given by Lemma 2.30 from x and x^∞. Then for each $n \geq 0$, we have $\widetilde{x}_n^\infty = \phi_n^{n+1} \circ \widetilde{x}_{n+1}^\infty$ and $\widetilde{x}_n \sim_{K'} \widetilde{x}_n^\infty$, where $K' := K/(1-\lambda)$.*

Proof. Immediate from Theorem G and Lemma 2.30. ∎

2.6. Proof of Theorem F

An inverse limit of connected spaces is connected, so $\mathcal{J}$ is connected.

We now show that f is a covering map. Let $x^\infty = (x_0, x_1, \ldots) \in X_0 \times X_1 \times \cdots$ represent an element of $\mathcal{J}$. A finite CW complex is locally contractible [27, Proposition A.4], so there is a connected neighborhood U of $x_0 \in X_0$ which is contained in a contractible set. The set U is then evenly covered by π; let $\widetilde{U}_i$, for $i = 1, \ldots, \deg(\pi)$, be the components of its preimages in X_1. The induced map $f\colon \mathcal{J} \to \mathcal{J}$ is the pullback of the covering map π under $\phi_0^\infty\colon \mathcal{J} \to X_0$. This implies that the neighborhood $(\phi_0^\infty)^{-1}(U)$ of x^∞ is evenly covered by the neighborhoods $(\phi_1^\infty)^{-1}(\widetilde{U}_i)$ under f.

We now show that $\mathcal{J}$ is locally path-connected and hence locally connected. In fact, we will show that $\mathcal{J}$ is *weakly locally path-connected* at every point: for all $x \in \mathcal{J}$ and every open neighborhood $U \ni x$, there is a smaller open neighborhood $x \in V \subset U$ so that any two points in V can be connected by a path in U. This is enough, since weak local path connectivity implies that path components of open sets are open.

We switch to thinking of $\mathcal{J}$ as a subset of $(X_1)^\infty$. Let d_1 denote the length metric on X_1. Fix $p^\infty = (p_1, p_2, \dots) \in \mathcal{J}$ and $\varepsilon_0 > 0$. For $m \ge 1$ and $\varepsilon < \varepsilon_0$, let $U_{m,\varepsilon}$ be those points $q^\infty = (q_1, q_2, \dots) \in \mathcal{J}$ for which $d_1(p_i, q_i) < \varepsilon$ for $i = 1, \dots, m$. The definition of the product topology says that for any $\varepsilon_0 > 0$, the $U_{m,\varepsilon}$ are a neighborhood basis of p^∞. Using local contractibility and Lemma 2.7, choose ε_0 so that balls in X_0 of radius ε_0 or smaller are contained in a contractible neighborhood and hence evenly covered by π and so that π is an isometry on these balls. Fix m and $\varepsilon < \varepsilon_0$ and focus attention on $U := U_{m,\varepsilon}$.

Let K be the constant for approximate path lifting given by Proposition 2.14 for the map $\phi\colon X_1 \to X_0$, and choose $n \ge m$ large enough that $\lambda^{n-m} K/(1-\lambda) < \varepsilon/3$. Let $V \subset U$ be $U_{n,\varepsilon/3}$, and fix $q^\infty \in V$. Using Theorem G, we are going to show $\mathcal{J}$ is weakly locally path-connected by constructing a path $x^\infty\colon I \to \mathcal{J}$ from p^∞ to q^∞ which is contained in U.

To construct x^∞, we will first construct a family $x = (x_i\colon I \to X_1)_{i \ge 1}$ of homotopy pseudo-orbits of paths joining p^∞ to q^∞ parameterized by $A = I$, the unit interval; as a path, each x_i joins p_i to q_i. We first construct the x_i for $1 \le i \le n$, where n is the integer from the previous paragraph; in this range, x_i will be an actual orbit (i.e., for $i < n$, the homotopies α_i are constant). We begin by choosing the path x_n. By construction, $d_1(p_n, q_n) < \varepsilon/3$; let $x_n\colon A \to X_1$ be a path exhibiting this. By decreasing induction, for $1 \le i < n$ define x_i to be the lift of $\phi \circ x_{i+1}$ under π starting at p_i. Since ϕ is a contraction and π is an isometry on balls of radius $\varepsilon/3$, the image of x_i is contained in the $\varepsilon/3$-ball about p_i. In particular, q_i is the only element of $\pi^{-1}(\phi(q_{i+1}))$ in this ball, so x_i ends at q_i.

We complete the construction of x by constructing x_i for $i > n$, by increasing induction starting with x_n. For $i \ge n$, supposing we have defined x_i, let x_{i+1} be an approximate lift under ϕ of $\pi \circ x_i$ joining p_{i+1} to q_{i+1}. By Proposition 2.14, $\pi \circ x_i$ and $\phi \circ x_{i+1}$ are joined by a homotopy α_i with trace size bounded by K. This completes the construction of x.

By Theorem G, the family of homotopy pseudo-orbits $(x, \alpha) = (x_i, \alpha_i)_{i \ge 1}$ defines

(i) a family of orbits parameterized by A, $x = (x_i^\infty\colon A \to X_1)_{i \ge 1}$, yielding a map $x^\infty\colon A \to \mathcal{J}$;

(ii) a sequence of homotopies $\beta = (\beta_i^\infty\colon I \times A \to X_1)_{i \ge 1}$ with β_i^∞ joining x_i to x_i^∞.

The bounds on the trace size of the homotopies β_i^∞ from x_i to x_i^∞ given in Theorem G are not enough for our purposes; to make sure the path x^∞ remains within U, we need to make sure that β_i^∞ has trace size less than ε for $1 \le i \le m$, while Theorem G gives a constant trace size $K/(1-\lambda)$. Thus we consider the proof of the theorem, which expresses each element β_i^∞ as a concatenation $\beta_i^1 * \beta_i^2 * \cdots * \beta_i^k * \cdots$, where each β_i^k is obtained from α_{i+k} by repeatedly lifting by π (a total of $k+1$ times) and composing with ϕ (a total of k times). Thus β_i^k is constant (trace size 0) for $i + k < n$, and otherwise has trace size bounded by $\lambda^k K$. In particular, for $i \le m$, we have

$$\operatorname{tracesize}(\beta_i^\infty) \le \sum_{k=n-i}^{\infty} \operatorname{tracesize}(\beta_i^k) \le \frac{\lambda^{n-i}}{1-\lambda} K \le \frac{\lambda^{n-m}}{1-\lambda} K < \frac{\varepsilon}{3}.$$

Finally, we estimate the size of x_i^∞ for $1 \le i \le m$ to check that it remains within U. We have

$$\operatorname{diam}(x_i^\infty) \le \operatorname{tracesize}(\beta_i^\infty) + \operatorname{diam}(x_i) + \operatorname{tracesize}(\beta_i^\infty) < \frac{\varepsilon}{3} + \frac{\varepsilon}{3} + \frac{\varepsilon}{3} = \varepsilon.$$

This concludes the proof of weak local path connectivity, and hence local connectivity.

It remains to prove that $f\colon \mathcal{J} \to \mathcal{J}$ is positively expansive. Let ε_0 be the parameter chosen above so that d_0-balls of radius ε_0 are contained in contractible sets. Consider the neighborhood of the diagonal $N := U_{1,\varepsilon_0} = \{(x, y) \in \mathcal{J} \times \mathcal{J} \mid d_1(x_1, y_1) < \varepsilon_0\}$. To prove that f is positively expansive, it suffices to show that, if $(f^i x, f^i y) \in N$ for each $i \ge 0$, then $x = y$, so let us suppose the iterates remain in N. If we think of $\mathcal{J}$ as a subset of $X_1^{\mathbb{N}}$, the map f is given by the left shift. Thus $d_1(x_i, y_i) < \varepsilon_0$ for each $i \ge 0$.

Since d_1 is a length metric, for each such i there exists a path $\beta_i\colon [0, 1] \to X_1$ with $\ell(\beta_i) < \varepsilon_0$ joining x_i to y_i. By construction, the paths $\pi(\beta_i)$ and $\phi(\beta_{i+1})$ join the same endpoints for each $i \ge 1$. We have $\operatorname{diam}(\pi(\beta_i)) < \varepsilon_0$ by construction, while $\operatorname{diam}(\phi(\beta_{i+1})) < \lambda\varepsilon_0$; hence the union of these paths lies in an ε_0-ball in X_0. Since the ball is contractible, the two paths are homotopic. The collection $\beta := (\beta_i)_{i \ge 1}$ is therefore a homotopy between the orbits $x = (x_i)_i$ and $y = (y_i)_i$. By Theorem 2.27, we then have $x = y$, completing the proof that f is positively expansive. ∎

2.7. Homotopy sections

Suppose $\phi\colon X \to Y$ is a continuous map between topological spaces. A *homotopy section* of ϕ is a continuous map $\sigma\colon Y \to X$ such that $\phi \circ \sigma \sim_K \mathrm{id}_Y$ for some constant K. The main result of this section is the existence of homotopy sections to $\mathcal{J}$.

Proposition 2.31. *Suppose $\phi, \pi\colon X_1 \rightrightarrows X_0$ is a λ-backward-contracting and recurrent virtual endomorphism of finite CW complexes. Suppose $\sigma\colon X_0 \to X_1$ is a homotopy section of ϕ, with $\phi \circ \sigma \sim_K \mathrm{id}_{X_0}$. Then for each $n \in \mathbb{N}$, we have the following:*

(1) *There is a canonically associated family*

$$(x^n, \alpha^n) = (x_i^n\colon X_n \to X_1, \alpha_i^n\colon I \times X_n \to X_0)_{i \ge 1}$$

of homotopy pseudo-orbits parameterized by X_n with trace sizes at most K. When the space X_n is identified with the set of orbits of length n, the map $X_n \to X_1 \times \cdots \times X_1$ determined by the first n coordinates $(x_1^n, x_2^n, \ldots, x_n^n)$ induces the identity map on X_n. Concretely, we require that $x_i = \pi_1^i \circ \phi_i^n$ for $i \le n$ and α_i is constant for $i < n$.

(2) *There exists $\widetilde{x}_\infty^n\colon X_n \to \mathcal{J}$ with $\phi_n^\infty \circ \widetilde{x}_\infty^n \sim_{K'} \mathrm{id}_{X_n}$, where $K' := K/(1-\lambda)$.*

Observe that in the lifted length metrics, the diameters of the X_n typically tend to infinity exponentially fast. Conclusion (2), however, says that the compositions $\phi_n^\infty \circ \widetilde{x}_\infty^n$ are at uniformly bounded distance from the identity.

Proof. We will construct a family of homotopy pseudo-orbits $x^n = (x_i : X_n \to X_1)_{i \geq 1}$. Fix a homotopy $\alpha : I \times X_0 \to X_0$ joining the identity to $\phi \circ \sigma$ with trace size K.

For $1 \leq i \leq n$, let $x_i^n : X_n \to X_1$ be the natural maps as defined in the statement; likewise, for $1 \leq i < n$, let $\alpha_i^n : I \times X_n \to X_0$ be the constant homotopies.

For $i \geq n$, set by induction $x_{i+1} := \sigma \circ \pi \circ x_i$. Then

$$\phi \circ x_{i+1} = \phi \circ \sigma \circ \pi \circ x_i \sim_K \pi \circ x_i$$

with homotopy $\alpha_i = (\pi \circ x_i)^*(\alpha)$ (see Lemma 2.13). This gives the desired family of homotopy pseudo-orbits.

The second assertion follows immediately from Theorem G′, applied to the family (x, α) of homotopy pseudo-orbits parameterized by X_n constructed above. ■

To get started applying this result, we have the following lemma.

Lemma 2.32. *Suppose G, H are finite connected graphs and $\phi : G \to H$ is surjective on the fundamental group. Then there exists a homotopy section $\sigma : H \to G$ of ϕ.*

Proof. The statement is invariant under homotopy equivalence, so we may assume H is a rose of k circles with basepoint y. Let $x \in \phi^{-1}(y)$ be a basepoint for G. Fix one of the k circles, say $\beta \subset H$. The assumption that ϕ induces a surjection between fundamental groups implies there is a loop α based at x for which $\phi(\alpha) \sim_{K(\beta)} \beta$ relative to y. We set $\sigma|_\beta = \alpha$. Doing this for each circle and putting $K := \max_\beta K(\beta)$ proves the claim. ■

As a corollary of Proposition 2.31, we then have the following.

Corollary 2.33. *If $\pi, \phi : G_1 \rightrightarrows G_0$ is a backward-contracting recurrent virtual endomorphism of graphs with limit space $\mathcal{J}$, then there exists a constant $K > 0$ and a family of maps $\sigma_\infty^n : G_n \to \mathcal{J}$ for $n \in \mathbb{N}$ such that $\phi_n^\infty \circ \sigma_\infty^n \sim_K \mathrm{id}_{G_n}$.*

3. The conformal gauge

We recall here from [24] the construction of two natural classes of metrics, one larger than the other, associated to certain expanding dynamical systems.

3.1. Convention

Throughout this section, we suppose $\mathcal{J}$ is compact, connected, and locally connected, and $f : \mathcal{J} \to \mathcal{J}$ is a positively expansive self-cover of degree $d \geq 2$. Proposition 2.5 and

Lemma 2.3 imply that the dynamics of $f\colon \mathcal{J} \to \mathcal{J}$ is topologically cxc. It follows that there exists a finite open cover $\mathcal{U}_0$ by connected sets such that, as in the notation from Section 2.1, the mesh of the coverings $\mathcal{U}_n$ tends to zero as $n \to \infty$, and in addition, for all $\widetilde{U} \in \mathcal{U}_n$ with $f^n\colon \widetilde{U} \to U \in \mathcal{U}_0$, we have $\deg(f^n\colon \widetilde{U} \to U) = 1$, that is, each such U is evenly covered by each iterate.

3.2. Visual metrics

The metrics we construct are most conveniently defined coarse-geometrically as visual metrics on the boundary of a certain rooted Gromov hyperbolic 1-complex. Before launching into technicalities, we quickly summarize the development. The *visual metrics* on $\mathcal{J}$ have the properties that there exists a constant $0 < \theta < 1$ such that for any $n \in \mathbb{N}$ and any $U \in \mathcal{U}_n$, we have $\operatorname{diam} U \asymp \theta^n$, and these U are uniformly nearly round; see Theorem 3.2 for the precise statements. The *snowflake gauge* is the set of metrics bi-Lipschitz equivalent to some power of a visual metric. The snowflake gauge is an invariant of the topological dynamics. Visual metrics are Ahlfors regular with respect to the Hausdorff measure in their Hausdorff dimension, and this measure is comparable to the measure of maximal entropy. The *Ahlfors-regular gauge* of metrics is the larger set of all Ahlfors-regular spaces quasi-symmetrically equivalent to a visual metric; it too is an invariant of the topological dynamics.

1-complex. From coverings as above, we define a rooted hyperbolic 1-complex Σ to get the visual metrics. In addition to the $\mathcal{U}_n$ for $n \geq 0$, let $\mathcal{U}_{-1}$ be the trivial covering $\{\mathcal{J}\}$. The vertices of Σ are the elements of the $\mathcal{U}_n$ for $n \geq -1$, with root $\mathcal{J} \in \mathcal{U}_{-1}$. If $U \in \mathcal{U}_n$, the *level* of U is $|U| := n$. Edges of Σ are of two types:

- *Horizontal* edges join $U, U' \in \mathcal{U}_n$ if $U \cap U' \neq \emptyset$.
- *Vertical* edges join U, U' if $||U| - |U'|| = 1$ and $U \cap U' \neq \emptyset$.

We equip Σ with the word length metric in which each edge has length 1. The *level* of an edge is the maximum of the levels of its endpoints.

Compactification. We compactify Σ in the spirit of W. Floyd rather than of M. Gromov, as follows. Let $\varepsilon > 0$ be a parameter, and let d_ε be the length metric on Σ obtained by scaling, so edges at level n have length θ^n where $\theta = e^{-\varepsilon}$. The metric space (Σ, d_ε) is not complete. Its completion $\overline{\Sigma}_\varepsilon$, sometimes called the *Floyd completion*, adjoins the corresponding boundary $\partial_\varepsilon \Sigma := \overline{\Sigma}_\varepsilon - \Sigma$. The extension of d_ε to a metric on the boundary $\partial_\varepsilon \Sigma$ is called a *visual metric* on the boundary.

Snapshots. To set up the statement of the next theorem, we need a definition.

Definition 3.1. Suppose X is a metric space and $0 < \theta < 1$. A sequence $(\mathcal{S}_n)_n$ of finite coverings of X is called a *sequence of snapshots of X with scale parameter θ* if there exists a constant $C > 1$ such that

(1) (scale and roundness) for all n and all $s \in \mathcal{S}_n$, there exists $x_s \in s$ with

$$B(x_s, C^{-1}\theta^n) \subset s \subset B(x_s, C\theta^n);$$

(2) (nearly disjoint) for all n, the collection of pairs $\{(x_s, s) \mid s \in \mathcal{S}_n\}$ may be chosen so that in addition

$$B(x_s, C^{-1}\theta^n) \cap B(x_{s'}, C^{-1}\theta^n) \neq \emptyset$$

whenever s and s' are distinct elements of $\mathcal{S}_n$.

The elements $s \in \mathcal{S}_n$ need not be open nor connected.

Properties of visual metrics. Theorem 3.2 summarizes highlights of [24, Chapter 3], specialized to the conventions in Section 3.1.

Theorem 3.2 (Visual metrics). *Suppose the topological dynamical system $f\colon \mathcal{J} \to \mathcal{J}$ and sequence of open covers $(\mathcal{U}_n)_n$ satisfy the assumptions in Section 3.1. Let Σ be the associated 1-complex.*

There exists $\varepsilon_0 > 0$ such that for all $0 < \varepsilon < \varepsilon_0$, the boundary $\partial_\varepsilon\Sigma$ equipped with the visual metric d_ε is naturally homeomorphic to $\mathcal{J}$. Moreover, for each ε in this range there exists $C > 1$ such that the following hold. Let $\theta = e^{-\varepsilon} < 1$; we denote by B_ε an open ball with respect to d_ε.

(1) *(Snapshot property) The sequence $(\mathcal{U}_n)_n$ is a sequence of snapshots of $\mathcal{J}$ with scale parameter θ.*

(2) *For $q := \frac{\log d}{\varepsilon}$, for all $x \in \mathcal{J}$ and all $r <$ $\operatorname{diam}(\mathcal{J})$, the Hausdorff q-dimensional measure satisfies*

$$C^{-1}r^q < \mathcal{H}^q(B_\varepsilon(x,r)) < Cr^q.$$

(3) *There exists a unique f-invariant probability measure μ_f of maximal entropy $\log d$ supported on $\mathcal{J}$. The support of μ_f is equal to all of $\mathcal{J}$, and for all Borel sets E, we have $C^{-1} < \mu_f(E)/\mathcal{H}^q(E) < C$.*

(4) *There exists $r_0 < \operatorname{diam}(\mathcal{J})$ so that for all $r < r_0$, we have $f(B_\varepsilon(x,r)) = B_\varepsilon(f(x), \theta^{-1}r)$, and the restriction $f\colon B_\varepsilon(x,r) \to B_\varepsilon(f(x), \theta^{-1}r)$ scales distances by the factor θ^{-1}.*

(5) *If two different parameters ε, ε' and two different open covers $\mathcal{U}_0$, $\mathcal{U}_0'$ are employed in the construction, the resulting metrics d, d' are snowflake equivalent.*

Snowflake equivalence. Two metrics d, d' are *snowflake equivalent* if there exist parameters $\beta, \beta' > 0$ such that the ratio $d^{\beta}/(d')^{\beta'}$ is bounded away from zero and infinity. Put another way, snowflake equivalent means bi-Lipschitz, after raising the metrics to appropriate powers. The *snowflake gauge* of $(f: \mathcal{J} \to \mathcal{J})$ is the snowflake equivalence class of some (and by Theorem 3.2 (5), equivalently, of any) visual metric on $\mathcal{J}$. Thus the snowflake gauge of the dynamical system $f: \mathcal{J} \to \mathcal{J}$ depends only on the topological dynamics and not on choices.

A snowflake equivalence $d \mapsto d'$ preserves the property of being a sequence of snapshots, though the scale parameter and constant C may change. However, if $(\mathcal{S}_n)_n$ is a sequence of snapshots in a metric space X, if Y is another metric space, and if $h: X \to Y$ is only a quasi-symmetry, then the transported sequence $(\mathcal{T}_n)_n$ comprised of images of elements of $\mathcal{S}_n$ under h need not be a sequence of snapshots: the scale condition in the definition of sequence of snapshots (Definition 3.1 (1)) can fail; see Section 3.3 for a definition of quasi-symmetry.

Conformal elevator and naturality of gauges. Suppose now $f: \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ is a hyperbolic rational function. Its Julia set $\mathcal{J}$ can now be equipped with at least two natural metrics: the round spherical metric and a visual metric. The Koebe distortion principles imply that small spherical balls can be blown up via iterates of f to balls of definite size, with uniformly bounded distortion. The same is true for the visual metric by Theorem 3.2 (4). Combining these observations in a technique known as the *conformal elevator* implies the following.

Proposition 3.3. *On Julia sets of hyperbolic rational functions, the spherical and visual metrics are quasi-symmetrically equivalent.*

This is true much more generally [24, Theorems 2.8.2 and 4.2.4]. In all but the most restricted cases, however, the spherical and visual metrics are not snowflake equivalent. In the visual metric, the map f is locally a homothety with a constant factor which is the same at all points. In the spherical metric, the image of a ball under f need not be a ball, and what is more, typically, the magnitude of the derivative $|f'(z)|$ varies as z varies in J_f. The exceptions include maps such as $f(z) = z^d$.

3.3. Ahlfors-regular spaces

We collect several facts here; see [29]. A metric space is *doubling* if there is an integer N such that any ball of radius r is covered by at most N balls of radius $r/2$. Doubling is a finite-dimensionality condition: the Assouad embedding theorem asserts that a doubling metric space is snowflake equivalent to a subset of a finite-dimensional spherical space.

Ahlfors regularity is a homogeneity condition that implies doubling. A space with both a metric and a measure (Z, d, μ) is *Ahlfors regular* of exponent q if $\mu(B(z, r)) \asymp r^q$ for each $r < \operatorname{diam}(Z)$, where the implicit constant is independent of z and of r. The Hausdorff dimension of such a space is necessarily equal to q, and in fact, the given measure μ is comparable to the q-dimensional Hausdorff measure. A metric space is *Ahlfors regular* if its Hausdorff measure in its Hausdorff dimension is Ahlfors regular.

Suppose (Z, d, μ) and (Z', d', μ') are connected, compact, doubling metric spaces. A homeomorphism $h\colon Z \to Z'$ is a *quasi-symmetry* (qs) if there is a constant $H \geq 1$ such that for each $r < \operatorname{diam}(Z)$ and each $z \in Z$, there is an $s > 0$ such that

$$B_{d'}(h(z), s) \subset h(B_d(z, r)) \subset B_{d'}(h(z), Hs).$$

In other words, the image of a round ball is nearly round.[2] A quasi-symmetric map does not, in general, preserve the property of Ahlfors regularity, though it does preserve the property of being doubling.

The *Ahlfors-regular conformal gauge* $\mathcal{G}(Z)$ of a metric space Z is the set of all Ahlfors-regular metric spaces qs equivalent to Z; it may be empty. Let $f\colon \mathcal{J} \to \mathcal{J}$ be a positively expansive self-cover as in the conventions of Section 3.1, and let d_ε be a visual metric from Theorem 3.2. Conclusion (2) of that theorem shows that visual metrics are Ahlfors regular, and so $\mathcal{G}(f\colon \mathcal{J} \to \mathcal{J}) := \mathcal{G}(\mathcal{J}, d_\varepsilon)$ is nonempty. For example, if f is a hyperbolic rational function with Julia set $\mathcal{J}$, the spherical metric d_{sph} on $\mathcal{J}$ is Ahlfors regular, with exponent $q = \operatorname{hdim}(\mathcal{J}, d_{\mathrm{sph}})$; see [43, Corollary 9.1.7]. Proposition 3.3 then implies that the spherical and visual metrics on $\mathcal{J}$ both belong to the gauge $\mathcal{G}(f\colon \mathcal{J} \to \mathcal{J})$.

The *Ahlfors-regular conformal dimension* $\operatorname{ARCdim}(f\colon \mathcal{J} \to \mathcal{J})$ is the infimum of the Hausdorff dimensions of the metrics in the Ahlfors-regular conformal gauge. It is thus another numerical invariant of the topological dynamics. Note that by definition, for any Ahlfors-regular metric d in $\mathcal{G}(f\colon \mathcal{J} \to \mathcal{J})$, we have

$$\operatorname{ARCdim}(f\colon \mathcal{J} \to \mathcal{J}) \leq \operatorname{hdim}(\mathcal{J}, d).$$

3.4. Approximately self-similar spaces

Our proof of Theorem A uses a result of Carrasco and Keith–Kleiner that says that for certain classes of spaces, the Ahlfors-regular conformal dimension is a critical exponent of combinatorial modulus (see Theorem 5.11). Here, we introduce that class of spaces.

The following definition appears in [32, Section 3].

Definition 3.4 (Approximately self-similar). A compact metric space (Z, d) is *approximately self-similar* if there exists a constant $L \geq 1$ such that for every ball $B(z, r) \subset Z$ with radius $0 < r < \operatorname{diam}(Z)$, there exists an open set $U \subset Z$ which is L-bi-Lipschitz equivalent to the rescaled ball $(B(z, r), \frac{1}{r} d)$.

An immediate consequence of Theorem 3.2 is the following.

[2]Strictly speaking, this is the definition of a *weakly qs* map; this is equivalent to the standard but less intuitive definition of a qs map in our setting; see [29].

Corollary 3.5 (Visual metrics are self-similar). *In the setting of Theorem 3.2, visual metrics are approximately self-similar.*

Proof. Let ε, r_0, and θ be as in Theorem 3.2. Set $L := \operatorname{diam}(Z)/(\theta r_0)$. Choose any $0 < r < \operatorname{diam}(\mathcal{J})$. Let $B_\varepsilon(x, r)$ be any ball. If $r > \theta r_0$, we take $U := B_\varepsilon(x, r)$ and note that the metric space $(B_\varepsilon(x, r), d_\varepsilon)$ is L-bi-Lipschitz equivalent to its rescaling $\bigl(B_\varepsilon(x, r), \frac{1}{r} d_\varepsilon\bigr)$. If $r < \theta r_0$, let n be the unique integer for which $\theta r_0 < \theta^{-n} r \le r_0$. Set $U := B_\varepsilon(f^n x, \theta^{-n} r)$. By Theorem 3.2 (4), the open set U is isometric to the rescaled metric space $(B_\varepsilon(x, r), \theta^{-n} d_\varepsilon) = \bigl(B_\varepsilon(x, r), \frac{1}{r} \cdot r\theta^{-n} d_\varepsilon\bigr)$ which is in turn L-bi-Lipschitz equivalent to $\bigl(B_\varepsilon(x, r), \frac{1}{r} d_\varepsilon\bigr)$ by our choice of n. ■

4. Energies of graph maps

In this section, we introduce asymptotic q-conformal energies associated to virtual endomorphisms of graphs and related analytical notions of extremal length. This section is a review of concepts from [51], which contains further motivation, especially in its Appendix A.

In this section, all graphs are assumed to be finite.

4.1. Weighted and q-conformal graphs

Definition 4.1. Suppose G is a graph, and $1 \le q \le \infty$.

(1) For $1 < q \le \infty$, a *q-conformal structure* on G is a positive *q-length* $\alpha(e)$ on each edge e, giving a length metric $|dx|$ in which e has length $\alpha(e)$. For $q = \infty$, this is also called a *length graph*.

(2) For $q = 1$, a 1-*conformal structure* on G is a positive weight $w(e)$ on each edge e. These weights do not determine a length structure. (For instance, in cases where we take derivatives for maps from a 1-conformal graph, the length structure is arbitrary.) A 1-conformal graph is also called a *weighted graph*.

We will write G^q for a graph together with a choice of q-conformal structure on it and denote its underlying graph by G.

Remark 4.2. Although q-conformal structures for $q \in (1, \infty]$ are all formally the same data, we distinguish the value of q for two reasons. First, it helps us keep track of which energies E_q^p to consider. Second, from another point of view a q-conformal structure is naturally thought of as an equivalence class of pairs (ℓ, w) of a length structure ℓ and weights w under a rescaling depending on q [51, Definition A.17], and from that point of view, there are two natural length-like structures in the picture.

The distinctions between q-conformal graphs as q varies arise from how various related analytical quantities are defined and scale under changes of weights. Imagine each edge as "thickened" with an extra $(q - 1)$-dimensional space to obtain a "rectangle" equipped with q-dimensional Hausdorff measure $\mathcal{H}^q$. Scaling the length of an edge by a factor λ changes the total $\mathcal{H}^q$-measure by the factor λ^q and therefore scales in the imaginary direction "orthogonal" to the edge by the factor λ^{q-1}.

An important special case of weighted graphs (i.e., $q = 1$) is when the underlying space is a 1-manifold $\bigcup_i J_i$, where each J_i is an interval or the circle. We may regard each J_i as a graph, by adding the endpoints of the interval, or an arbitrary vertex on the circle. Up to isomorphism, the result is unique. A formal sum $C = \sum w_i J_i$ with $w_i > 0$ then determines a weighted graph in a unique way, by setting the weight of the unique edge in J_i to be w_i. Ignoring the graph structure, we also call C a weighted 1-manifold.

Definition 4.3. A *curve* on a space X is a connected 1-manifold C (either I or S^1) together with a map $\gamma: C \to X$. We refer to the curve by just the map γ to be short, but the underlying domain C is part of the data. (Note the distinction with a *path*, where, for us, the domain is a fixed interval of real numbers.) For a *multi-curve*, we drop the restriction that C be connected, giving $C = \bigcup_i J_i$ where the $J_i \in \{I, S^1\}$ are the connected components and the map γ is determined by $\gamma_i: J_i \to X$. A *strand* of a multi-curve is one of its component curves γ_i. A *weighted multi-curve* $\gamma: C \to X$ is similar, but where $C = \sum w_i J_i$ has the structure of a weighted graph; we also denote this by $\gamma = \sum w_i \gamma_i$. The homotopy class of a curve or multi-curve γ is denoted $[\gamma]$.

4.2. Energies and extremal length

For any $1 \le p \le q \le \infty$ and piecewise linear (PL) map $\phi: G^p \to H^q$ from a p-conformal graph to a q-conformal graph, there is an energy $E_q^p(\phi)$ with well-behaved properties. For any of these energies and a homotopy class $[\phi]$ of maps as above, we denote by $E_q^p[\phi]$ the infimum of the energy over maps in the class. In this paper, we only need a few special cases.

Case $p = 1$ and $1 < q < \infty$. We switch names to consider a PL map $\phi: W^1 \to G^q$ from a weighted (or 1-conformal) graph W^1 to a q-conformal graph G^q, where $1 < q < \infty$. (In this paper, W^1 will typically be a 1-manifold, so this is the structure of a weighted curve on G^q.) Let $q^\vee = q/(q-1)$ be the Hölder conjugate of q.

Define

$$E_q^1(\phi) := \|n_\phi\|_{q^\vee, G^q}.$$

We now explain the notation. For $y \in G^q$, the value

$$n_\phi(y) := \sum_{\phi(x)=y} w(e(x))$$

is the weighted number of preimages of y, where for $x \in W^1$, the quantity $e(x)$ denotes an edge containing x. The edge weights α on G^q, when interpreted as lengths, define a length measure $|dy|$ on the underlying set of G^q such that $\int_e |dy| = \alpha(e)$ for each edge e of G^q. The quantity $n_\phi(y)$ may be infinite (if, e.g., ϕ collapses an edge to the point y) or undefined (if, e.g., $x \in \phi^{-1}(y)$ is a vertex incident to edges with different weights). However, the assumption that ϕ is PL and our convention that the graphs are finite imply that the set of such y is finite, hence of $|dy|$-measure zero. The quantity $\|n_\phi\|_{q^\vee,G^q}$ is then the usual $L^{q^\vee}$ norm of the function n_ϕ with respect to the measure $|dy|$. For instance, $E^1_\infty(\phi)$ is the weighted total length of the image of ϕ.

If n_ϕ is constant on edges of G^q—automatic if ϕ minimizes energy in its homotopy class—the formula for the energy is concretely given by

$$E^1_q(\phi) = \Big(\sum_{e\in\mathrm{Edge}(G)} \alpha(e) n_\phi(e)^{q^\vee}\Big)^{1/q^\vee}. \tag{4.4}$$

Definition 4.5. A weighted multi-curve $\gamma\colon C \to G$ on a graph G is *reduced* if the restriction to each strand is either constant or has arbitrarily small perturbations that are locally injective.

Thus the images of the strands of a reduced multi-curve have no backtracking. Reduced curves ϕ minimize E^1_q in their homotopy class [51, Proposition 3.8 and Lemma 3.10].

We also define the *q-extremal-length* of a homotopy class of maps $\phi\colon W^1 \to G^q$ by

$$\mathrm{EL}_q[\phi] := (E^1_q[\phi])^q. \tag{4.6}$$

To justify the terminology, note that as explained in [50, Section 5.2], the minimizer of extremal length over a suitable homotopy class of maps exists and is realized by a map with nice properties, and the minimum value, when formulated as an extremal problem, mimics the usual definition of extremal length; see also Section 5.5.

Case $1 < p < \infty$ and $q = \infty$. The *p-harmonic energy* of a PL map $\phi\colon G^p \to K^\infty$ from a p-conformal graph to a length graph is

$$E^p_\infty(\phi) := \||\phi'|\|_{p,G^p}.$$

Here $|\phi'|$ denotes the size of the derivative: since $p > 1$, the conformal graph G^p has a length structure, so we can differentiate. If the derivative of ϕ is constant on the edges of G^p—automatic if ϕ minimizes energy in its homotopy class—this is

$$E^p_\infty(\phi) = \Big(\sum_{e\in\mathrm{Edge}(G)} \alpha(e)^{1-p}\ell(\phi(e))^p\Big)^{1/p},$$

where $\ell(\phi(e))$ is the total length of the image of e.

Case $p = 1$ and $q = \infty$. This is the common limit of the above two cases, slightly modified since 1-conformal graphs have a weight w instead of a p-length α. For a PL map $\phi\colon W^1 \to K^\infty$ from a weighted graph to a length graph, set

$$E^1_\infty(\phi) := \int_{x\in W} w(x)|\phi'(x)|\,dx = \int_{y\in K} n_\phi(y)\,dy.$$

This is the weighted length of the image of ϕ.

Case $1 < p = q < \infty$. A PL map $\phi\colon G^q \to H^q$ between q-conformal graphs has *filling function* $\mathrm{Fill}^q(\phi)\colon H^q \to \mathbb{R}$ given at generic points by

$$\mathrm{Fill}^q(\phi)(y) := \sum_{\phi(x)=y} |\phi'(x)|^{q-1} \tag{4.7}$$

and a q-conformal *energy* given by

$$E_q^q(\phi) := (\|\mathrm{Fill}^q(\phi)\|_{\infty,H})^{1/q}. \tag{4.8}$$

If we interpret edges of conformal graphs as "thickened to rectangles" as described above, the filling function sums the "thicknesses" of the "rectangles" over a fiber.

Case $p = q = 1$. A PL map $\phi\colon G^1 \to H^1$ has an energy E_1^1 that is again a limit of the above case, modified to account for weights rather than q-lengths:

$$N(\phi) = E_1^1(\phi) := \operatorname*{ess\,sup}_{y \in H} \frac{n_\phi(y)}{w(y)}.$$

We will apply this in cases where the weights are all 1, in which case this amounts to counting the essential maximum of the number of preimages of any point. ("Essential" as usual means that we can ignore sets of linear measure zero, and in particular, we can ignore images of edges of G^1 that map to a single point in H^1.) We will also call this quantity $N(\phi)$.

Properties of energies. It is not too hard to see that these energies above are all *sub-multiplicative*: if $\phi\colon G^p \to H^q$ and $\psi\colon H^q \to K^r$ are maps from a p-conformal graph to a q-conformal graph to an r-conformal graph (where the energies are defined), then $E_r^p(\psi \circ \phi) \le E_q^p(\phi) E_r^q(\psi)$ [51, Proposition A.12].

Given $\phi\colon G^q \to H^q$, we denote by $[\phi]$ its homotopy class. It is natural to consider minimizers of E_q^q over the class $[\phi]$. The fact that minimizers of these energies exist is not obvious, but more is true. We first give some general context.

Definition 4.9 ([51, Definition 1.33]). Given maps $\phi\colon G^p \to H^q$ and $\psi\colon H^q \to K^r$ between graphs G^p, H^q, and K^r with the respective conformal structures, we say that the sequence $G^p \xrightarrow{\phi} H^q \xrightarrow{\psi} K^r$ is *tight* if

$$E_r^p[\psi \circ \phi] = E_r^p(\psi \circ \phi) = E_q^p(\phi) E_r^q(\psi).$$

Together with sub-multiplicativity of energy, the existence of a tight sequence implies that ϕ and ψ both minimize energy in their respective homotopy classes, and furthermore, the sub-multiplicativity inequalities are sharp. Then [51, Theorem 6, Appendix A] asserts that *any* map $\phi\colon G^q \to H^r$ can be homotoped to be part of a tight sequence (on either side). We will state and use two special cases, starting with the easier one.

Proposition 4.10. *Pick $1 \le q \le \infty$ and let G^q be a q-conformal graph. For any reduced multi-curve $\gamma\colon C^1 \to G^q$ on G, there is a length graph K^∞ and map $\psi\colon G^q \to K^\infty$ so that*

$$E_q^1(\gamma) = E_q^1[\gamma] = \frac{E_\infty^1(\psi \circ \gamma)}{E_\infty^q(\psi)},$$

and the maps ψ and $\psi \circ \gamma$ minimize the energies E_∞^q and E_∞^1, respectively, in their homotopy classes. Furthermore, we can take K^∞ to have the same underlying graph as G^q but with different edge lengths and can take ψ to be the identity.

This is the case $p = 1$ of [51, Theorem 6]; see also [51, Proposition A.10]. It can be viewed as the duality of L^q and $L^{q^\vee}$ norms. The case $q = 2$ appears as the duality map in [51, Definition 6.21]. Since the proof of this case was omitted in the cited paper, we include it here.

Proof. Define K^∞ to have underlying graph G with the length of an edge given by $\ell(e) := \alpha(e) n_\gamma(e)^{1/(q-1)}$. Set ψ to be the identity as a map from G to itself. Then $\psi \circ \gamma\colon C^1 \to K^\infty$ is reduced. We thus have

$$E_\infty^1[\psi \circ \gamma] = E_\infty^1(\psi \circ \gamma) = \sum_e n_\gamma(e) \cdot \ell(e) = E_q^1(\gamma) E_\infty^q(\psi),$$

where last equality is the equality case of the Hölder inequality $\|n_\gamma \cdot (\ell/\alpha)\|_{1,G^q} \le \|n_\gamma\|_{q^\vee,G^q} \cdot \|\ell/\alpha\|_{q,G^q}$ and is immediate if you expand the definitions. We therefore have a tight sequence $C^1 \xrightarrow{\gamma} G^q \xrightarrow{\psi} K^\infty$, yielding all the conclusions in the statement by [51, Lemma 1.34]. ∎

For the other version, we define another quantity, the *stretch factor*. Define

$$\overleftarrow{\mathrm{SF}}_q[\phi] := \sup_{[\gamma]: C \to G} \frac{E_q^1[\phi \circ \gamma]}{E_q^1[\gamma]},$$

where the supremum runs over all nontrivial homotopy classes of weighted multi-curves $\gamma\colon C^1 \to G^q$ (which we may take to be reduced). In other words, by equation (4.6), the q-stretch factor is a power of the maximum ratio of distortion of q-extremal-length of homotopy classes of maps of weighted multi-curves. (It is easy to see that the supremum is the same if we take it over unweighted multi-curves or over unweighted curves, but in either of these other cases, the supremum is not always realized.)

Theorem 4.11. *For $1 \le q \le \infty$ and $[\phi]\colon G^q \to H^q$ a homotopy class of maps between q-conformal graphs, there is a map $\psi \in [\phi]$, a weighted 1-manifold $C^1 = \sum_i w_i J_i$, and a map $\gamma\colon C^1 \to H^q$ fitting into a tight sequence*

$$C^1 \xrightarrow{\gamma} G^q \xrightarrow{\psi} H^q.$$

Explicitly, ψ *minimizes* E_q^q *in* $[\phi]$ *and*

$$E_q^q(\psi) = E_q^q[\phi] = \frac{E_q^1(\psi \circ \gamma)}{E_q^1(\gamma)} = \overleftarrow{\mathrm{SF}}_q[\phi].$$

The multi-curves $\gamma\colon C^1 \to G^q$ *and* $\psi \circ \gamma\colon C^1 \to H^q$ *are reduced and thus minimize* E_q^1 *in their homotopy classes. Furthermore, we can choose the curve so that for each edge* e *of* G^q *and each strand* J_i *of* C^1*, the image of the restriction* $\gamma|J_i$ *meets* e *at most twice.*

Most of this is the case $p = q$ of [51, Theorem 6]. The last conclusion in Theorem 4.11 (on C^1 meeting each edge of G^q at most twice) is a consequence of [51, Proposition 3.19].

4.3. Asymptotic energies of virtual graph endomorphisms

In this subsection, we summarize some results from [52].

We now assume $\pi, \phi\colon G_1 \rightrightarrows G_0$ is a virtual endomorphism of graphs. Following the notation from Section 2.5, for $n = 1, 2, \dots$, let $\phi_0^n\colon G_n \to G_0$ and $\pi_0^n\colon G_n \to G_0$ be the induced maps. For ease of reading, we write $\phi^n := \phi_0^n$ and $\pi^n := \pi_0^n$.

Fix $1 \le q \le \infty$. Fix a q-conformal structure on G_0: for $q > 1$, pick q-lengths α_0, or for $q = 1$, pick weights w_0. This q-conformal structure can be lifted under the coverings $\pi^n\colon G_n \to G_0$ to yield a q-conformal structure α_n or w_n on G_n that we denote G_n^q.

Definition 4.12 (Asymptotic q-conformal energy). Suppose $1 \le q \le \infty$. The *asymptotic* *q-conformal energy* of a virtual endomorphism (π, ϕ) is

$$\overline{E}^q(\pi, \phi) := \lim_{n\to\infty} E_q^q[\phi^n]^{1/n},$$

where $\phi^n\colon G_n^q \to G_0^q$, and we use the lifted q-conformal structure as above.

The limit exists and is equal to the infimum of the terms, since the energies are invariant under passing to covers and are sub-multiplicative; see [52, Proposition 5.6]. Since we take homotopy classes of ϕ^n on the right-hand side, it follows easily [52, Proposition 5.7] that $\overline{E}^q(\pi, \phi)$ only depends on the homotopy class of (π, ϕ) as defined in [52, Definition 2.2]. In particular, it is independent of the choice of q-conformal structure on G_0. So we will henceforth write the asymptotic energy as $\overline{E}^q[\pi, \phi]$. In addition, the asymptotic energy is continuous and non-increasing in the exponent q [52, Proposition 6.10, Corollary 6.12]. We summarize these facts as follows.

Proposition 4.13 ([52, Proposition 6.10]). *Suppose* $1 \le q \le \infty$ *and* $\pi, \phi\colon G_1 \rightrightarrows G_0$ *is a virtual endomorphism of graphs. Then the asymptotic* q*-conformal energy* $\overline{E}^q[\pi, \phi]$ *depends only on the homotopy class of* $[\pi, \phi]$ *and is continuous and non-increasing as a function of* q.

We will assume that $\overline{E}^\infty[\pi, \phi] < 1$; equivalently, after passing to an iterate, there is a metric on G_0 such that $\phi\colon G_1 \to G_0$ is contracting [52, Section 6, p. 36]. This places us in the setup of Section 2, so we have, by Theorem F, a locally connected limit space $\mathcal{J}$ and a positively expansive self-cover $f\colon \mathcal{J} \to \mathcal{J}$. The asymptotic energies $\overline{E}^q[\pi, \phi]$ become numerical invariants of our presentation of the dynamical system $f\colon \mathcal{J} \to \mathcal{J}$.

Question 4.14. Is $\overline{E}^q[\pi, \phi]$ an invariant of the topological conjugacy class of $f\colon \mathcal{J} \to \mathcal{J}$? That is, if two different contracting graph virtual endomorphisms $\pi, \phi\colon G_1 \rightrightarrows G_0$ give homeomorphic limit spaces $\mathcal{J}$ and topologically conjugate maps f on them, do the energies $\overline{E}^q$ coincide? Is this the case if they are combinatorially equivalent in the sense of [38]?

Remark 4.15. We expand on Question 4.14. If we start with a hyperbolic pcf rational map $f\colon \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$, then we can take G_0 to be a spine of $\widehat{\mathbb{C}}$ minus the post-critical set P_f of f and $G_1 := f^{-1}(G_0)$. Since any two spines of the complement of P_f are homotopy equivalent, any two such choices will give homotopy-equivalent virtual endomorphisms and the same energies $\overline{E}^q$ and critical exponents q^* and q_*. But there are other virtual endomorphisms that give the same dynamics on the limit space. For instance, we could reindex, considering $\pi^n_{n-1}, \phi^n_{n-1}\colon G_n \rightrightarrows G_{n-1}$ as in Section 2.5, without changing the dynamics of f on $\mathcal{J}$. (Note, by contrast, that iterating to $\pi^n_0, \phi^n_0\colon G_n \rightrightarrows G_0$ does not change $\mathcal{J}$, but does change f and $\overline{E}^q$ in a predictable way.) More generally, for a hyperbolic pcf rational map, one could look at a forward-invariant set $P' \supsetneq P_f$ of Fatou points, adding some preimages of points in P_f, and construct a virtual endomorphism from that; for $P' = f^{-n}(P_f)$, we get $(\pi^{n+1}_n, \phi^{n+1}_n)$, but there are many other possibilities, yielding virtual endomorphisms that are not homotopy equivalent in the sense of [52, Definition 2.2] (since the graphs G_0 have different ranks) but give the same limiting dynamics on $\mathcal{J}$.

For general dynamics $f\colon \mathcal{J} \to \mathcal{J}$, it is unclear how to parameterize all the different ways to see f as the limit of a graph virtual endomorphism, and thus Question 4.14 is open. In the special case of rational maps $f\colon \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ (or expanding branched self-covers), all spines of $\widehat{\mathbb{C}} - P_f$ are homotopy equivalent, and so we can write $\overline{E}(f)$ unambiguously. Another question for future research is to consider cases where $\mathcal{J}$ is not topologically 1 dimensional and thus not a limit of graph virtual endomorphisms at all: there should be some suitable replacement for the graph energies E^p_q.

5. Combinatorial modulus

In this section, we first recall the definition of the combinatorial q-modulus (or, equivalently, its inverse, the combinatorial q-extremal-length) associated to a family Γ of curves γ on a topological space X equipped with a finite cover $\mathcal{U}$. In fact, there are several such notions, treating families of curves, of multi-curves, and of weighted multi-curves, as defined in Section 4.1. We next show their coarse equivalence. Suppose Γ is a family of curves on X.

- There is a canonically associated family Γ^w of weighted multi-curves $\sum w_i\gamma_i$, $\gamma_i \in \Gamma$, normalized so $\sum_i w_i = 1$. We extend the definition of combinatorial modulus to weighted families so that the moduli of Γ and Γ^w coincide; see Lemma 5.3.
- The family Γ may be chopped up, or *subdivided*, into a family of weighted multi-curves whose strands are shorter curves. Under natural conditions, the moduli of Γ and of the resulting subdivided family are comparable, with control; see Lemma 5.4.
- If $\mathcal{U}$ and $\mathcal{V}$ are two covers with controlled overlap, then the moduli of the family Γ with respect to $\mathcal{U}$ and to $\mathcal{V}$ are comparable, with control; see Lemma 5.5. As an application, two curves that are close, with control, have comparable modulus; see Lemma 5.6.
- If $X = G^q$ is a q-conformal graph, C is a weighted 1-manifold as in Section 4.1, and Γ is the set of maps $\gamma: C \to G^q$ from C to G^q in a fixed homotopy class $[\gamma]$, then the combinatorial q-extremal-length of Γ with respect to the covering by closed edges of G^q is comparable to a power of the q-conformal energy, $(E_q^1[\gamma])^q$, with control; see Proposition 5.10.

We conclude with a result, Theorem 5.11, which says the Ahlfors-regular conformal dimension coincides with a critical exponent for combinatorial modulus of weighted curve families whose elements have diameters bounded from below. This follows from a result of Carrasco [12, Corollary 1.4] and Keith–Kleiner, which holds for unweighted curves, and Lemma 5.3.

5.1. Combinatorial modulus

Let $\mathcal{S}$ be a finite covering of a topological space X, by which we mean a multi-set of subsets s of X whose union is X. Note that we allow repetitions and do not assume the s's are open or connected. For $K \subset X$, denote by $\mathcal{S}(K)$ the set of elements of $\mathcal{S}$ which intersect K, and let $\bigcup \mathcal{S}(K)$ be the union of these elements of $\mathcal{S}$; we may think of $\bigcup \mathcal{S}(K)$ as the "$\mathcal{S}$-neighborhood" of K.

Let $q \geq 1$. A *test metric* is a function $\rho: \mathcal{S} \to [0, \infty)$. The *$q$-volume of X with respect to ρ* is

$$V_q(\rho, \mathcal{S}) := \sum_{s \in \mathcal{S}} \rho(s)^q .$$

For $K \subset X$, the *ρ-length* of K is

$$\ell_\rho(K, \mathcal{S}) := \sum_{s \in \mathcal{S}(K)} \rho(s).$$

If $\gamma: C \to X$ is a curve, we define the ρ-length of its image in X by

$$\ell_\rho(\gamma, \mathcal{S}) := \sum_{s \in \mathcal{S}(\gamma(C))} \rho(s),$$

that is, identifying γ with its image. Since C is implicitly part of the data defining γ, we adopt the briefer notation $\mathcal{S}(\gamma)$ for $\mathcal{S}(\gamma(C))$.

We pause to clarify some subtle points. First, each sum is a sum over a multi-set. So, if an element s of $\mathcal{S}$ is repeated, say twice, then the sum for $V_q(\rho, \mathcal{S})$ has two corresponding terms. Next, recall that the domain of a curve is connected. The ρ-length of a curve γ depends only on its image $\gamma(C)$ and not on its parameterization. So, for example, if γ parameterizes a simple loop, then $\ell_\rho(\gamma, \mathcal{S}) = \ell_\rho(\gamma^n, \mathcal{S})$ for all $n \geq 1$. And finally, more generally, the computation of $\ell_\rho(\gamma, \mathcal{S})$ does not depend on the number of times γ meets an element $s \in \mathcal{S}$; see Section 5.3 for alternatives.

If Γ is a family of curves in X, we make the following definitions leading up to the *combinatorial modulus of* Γ *with respect to* $\mathcal{S}$, denoted $\operatorname{mod}_q(\Gamma, \mathcal{S})$:

$$\begin{aligned} L_\rho(\Gamma, \mathcal{S}) &:= \inf_{\gamma\in\Gamma} \ell_\rho(\gamma, \mathcal{S}) \\ \operatorname{mod}_q(\Gamma, \rho, \mathcal{S}) &:= \frac{V_q(\rho, \mathcal{S})}{L_\rho(\Gamma, \mathcal{S})^q} \\ \operatorname{mod}_q(\Gamma, \mathcal{S}) &:= \inf_\rho \operatorname{mod}_q(\Gamma, \rho, \mathcal{S}). \end{aligned} \tag{5.1}$$

In the last infimum, we restrict to test metrics ρ for which $L_\rho(\Gamma, \mathcal{S}) \neq 0$. We often use an equivalent formulation of $\operatorname{mod}_q(\Gamma, \mathcal{S})$ more familiar to analysts. A test metric is *admissible* for Γ if $\ell_\rho(\gamma, \mathcal{S}) \geq 1$ for each $\gamma \in \Gamma$. Then it is easy to see that

$$\operatorname{mod}_q(\Gamma, \mathcal{S}) = \inf\{V_q(\rho, \mathcal{S}) \mid \rho \text{ admissible for } \Gamma\}.$$

The combinatorial modulus of a nonempty family of curves is always finite and positive.

The *combinatorial extremal length* of a curve family Γ is the reciprocal of the combinatorial modulus:

$$\operatorname{EL}_q(\Gamma, \mathcal{S}) = \sup_\rho \frac{L_\rho(\Gamma, \mathcal{S})^q}{V_q(\rho, \mathcal{S})}.$$

We will relate this to $E_q^1[\gamma]$ in the sense of Section 4; see Section 5.5.

Remark 5.2. In this finite combinatorial world, it makes sense to take $\Gamma = \{\gamma\}$, a single curve. That is, for any $q \geq 1$, any finite cover $\mathcal{S}$, and any curve γ, we have a finite and nonzero $\operatorname{mod}_q(\gamma, \mathcal{S})$. Essentially, discretizing via the coverings thickens a single curve so that it behaves like a family.

5.2. Weighted multi-curves

Recall that a *weighted multi-curve* is a finite formal sum $\gamma = \sum w_i \gamma_i$, where $w_i > 0$ and the γ_i's are curves on X. It is *normalized* if $\sum w_i = 1$.

We will require two sorts of families of weighted multi-curves on a space X.

- Fix a multi-curve $\gamma\colon C \to X$ with strands $\gamma_i\colon J_i \to X$, $J_i \in \{I, S^1\}$, and fix corresponding positive weights w_i. Then we may consider the family whose elements are weighted multi-curves $\gamma = \sum w_i \gamma_i$. The number of strands and the weights are fixed within the family.

- Fix a set Γ of unweighted curves on X. Then we may consider the family whose elements are weighted multi-curves $\gamma = \sum_i w_i \gamma_i$, where $\gamma_i \in \Gamma$. Here, both the number of strands and the weights may vary within the family.

To be concrete about the second type of family, suppose Γ is a family of unweighted curves on X. Define

$$\Gamma^w := \left\{ \sum w_i \gamma_i \;\middle|\; \gamma_i \in \Gamma, \sum w_i = 1 \right\}.$$

This is a family of normalized weighted curves canonically associated to Γ. The assignment $\gamma \mapsto 1 \cdot \gamma$ yields an inclusion $\Gamma \hookrightarrow \Gamma^w$. So, as families of weighted curves, we have $\Gamma \subset \Gamma^w$.

We next extend the definition of combinatorial modulus to families of weighted multi-curves. Let $\mathcal{S}$ be a finite covering of X, and let $q \geq 1$. If $\rho \colon \mathcal{S} \to [0, \infty)$ is a test metric, the *ρ-length* of a weighted multi-curve $\gamma = \sum w_i \gamma_i$ is given by

$$\ell_\rho(\gamma, \mathcal{S}) := \sum w_i \ell_\rho(\gamma_i).$$

We then define the combinatorial modulus $\mathrm{mod}_q(\Gamma^w, \mathcal{S})$ using equation (5.1).

Lemma 5.3 (Moduli of weighted curves). *Suppose Γ is a family of unweighted curves, and Γ^w is the corresponding family of normalized weighted curves. For any $q \geq 1$, we have* $\mathrm{mod}_q(\Gamma^w, \mathcal{S}) = \mathrm{mod}_q(\Gamma, \mathcal{S})$.

Proof. The inequality $\mathrm{mod}_q(\Gamma^w, \mathcal{S}) \geq \mathrm{mod}_q(\Gamma, \mathcal{S})$ holds since $\Gamma^w \supset \Gamma$. Now suppose ρ satisfies $\ell_\rho(\gamma, \mathcal{S}) \geq 1$ for each $\gamma \in \Gamma$, and suppose $c = \sum_i w_i \gamma_i \in \Gamma^w$. Then

$$\ell_\rho(c, \mathcal{S}) = \sum_i w_i \left(\sum_{s \cap \gamma_i \neq \emptyset} \rho(s) \right) = \sum_i w_i \ell_\rho(\gamma_i, \mathcal{S}) \geq \sum_i w_i = 1.$$

Thus each admissible metric for Γ is also admissible for Γ^w, so $\mathrm{mod}_q(\Gamma^w, \mathcal{S}) \leq \mathrm{mod}_q(\Gamma, \mathcal{S})$. ∎

5.3. Subdivision

We now turn to subdividing weighted multi-curves. We only need this construction for the case $\Gamma = \{\gamma\}$, a single weighted multi-curve $\gamma = \sum_i w_i \gamma_i$, and not for curve families.

Suppose $\gamma = \sum w_i \gamma_i$ is a weighted multi-curve corresponding to a map $\gamma \colon C \to X$, where $C = \bigcup_i w_i J_i$ and $\gamma_i \colon J_i \to X$, where each J_i is either an interval or a circle. Suppose that for each i we are given a finite set of maps $\alpha_{i,k} \colon I_{i,k} \to J_i$, where each $I_{i,k}$ is a copy of the unit interval and the images of the $\alpha_{i,k}$ cover J_i without gaps or overlaps, except for singleton points on the boundary. Then the corresponding *subdivision* of γ is the weighted multi-curve on X

$$\sum_i \sum_k w_i (\gamma_i \circ \alpha_{i,k}),$$

corresponding to the map

$$\zeta\colon D = \bigcup_{i,k} w_i I_{i,k} \to X$$

with component maps $\zeta_{i,k} = \gamma_i \circ \alpha_{i,k}$. The subdivision of a normalized weighted multi-curve is not usually normalized.

Lemma 5.4 (Subdivision). *Suppose $\gamma = \sum_i w_i \gamma_i$ is a weighted multi-curve on X, and $\mathcal{S}$ is a finite cover of X. Let $q \geq 1$, and let ζ be a subdivision of γ. Then $\operatorname{mod}_q(\zeta, \mathcal{S}) \leq \operatorname{mod}_q(\gamma, \mathcal{S})$.*

Furthermore, if no $\zeta_{i,k}$ has image contained in a single element of $\mathcal{S}$ and there is a constant K so that, for each $s \in \mathcal{S}$ and each strand i, the number of connected components of $\gamma_i^{-1}(s)$ is bounded by K, then $\operatorname{mod}_q(\gamma, \mathcal{S}) \lesssim_{K,q} \operatorname{mod}_q(\zeta, \mathcal{S})$.

Proof. Suppose ρ is any test metric. For the first assertion, we have $\ell_\rho(\zeta, \mathcal{S}) \geq \ell_\rho(\gamma, \mathcal{S})$. The inequality for q-modulus then follows from the definitions.

For the second assertion, the given conditions imply that if a given element $s \in \mathcal{S}$ meets some γ_i, it then meets at most $2K$ of the restricted curves $\zeta_{i,k}$. Thus $\ell_\rho(\zeta, \mathcal{S}) \leq 2K \cdot \ell_\rho(\gamma, \mathcal{S})$. It follows that $\operatorname{mod}_q(\zeta, \mathcal{S}) \geq (2K)^{-q} \operatorname{mod}_q(\gamma, \mathcal{S})$. ■

5.4. When moduli are comparable

Several different coverings naturally arise in our setting on graphs. We show that given a curve family, the corresponding q-moduli are coarsely equivalent when certain natural regularity conditions hold.

Two coverings $\mathcal{U}, \mathcal{V}$ of a topological space X have *K-bounded overlap* if, for each $u \in \mathcal{U}$, we have $1 \leq \#\mathcal{V}(u) \leq K$, and vice versa for each $v \in \mathcal{V}$ we have $1 \leq \#\mathcal{U}(v) \leq K$. Here, the notation follows that from Section 5.1, so that $\mathcal{V}(u) = \{v \in \mathcal{V} \mid v \cap u \neq \emptyset\}$ and $\mathcal{U}(v) = \{u \in \mathcal{U} \mid u \cap v \neq \emptyset\}$. A covering $\mathcal{U}$ is *K-bounded* if it has K-bounded overlap with itself.

Lemma 5.5 (Bounded overlap). *Suppose $\mathcal{U}, \mathcal{V}$ are two finite coverings of a topological space X with K-bounded overlap. Let Γ be a family of weighted multi-curves on X. Then, for any $q \geq 1$,*

$$\operatorname{mod}_q(\Gamma, \mathcal{U}) \asymp_{K,q} \operatorname{mod}_q(\Gamma, \mathcal{V})$$

or, more precisely,

$$\operatorname{mod}_q(\Gamma, \mathcal{U}) \in [K^{-q-1}, K^{q+1}] \cdot \operatorname{mod}_q(\Gamma, \mathcal{V}).$$

Proof. We imitate the argument in [11, Theorem 4.3.1]. Suppose $\sigma\colon \mathcal{V} \to \mathbb{R}$ is a test metric for $\mathcal{V}$. Define a test metric $\rho\colon \mathcal{U} \to \mathbb{R}$ and a function $f\colon \mathcal{U} \to \mathcal{V}$ by setting

$$\rho(u) = \sigma(f(u)) := \max\{\sigma(v) \mid u \cap v \neq \emptyset\},$$

that is, $f(u) \in \mathcal{V}$ is any element that realizes the maximum.

For any strand γ of an element in Γ, we have

$$\ell_\sigma(\gamma, \mathcal{V}) = \sum_{\substack{v \in \mathcal{V} \\ v \cap \gamma \neq \emptyset}} \sigma(v) \leq \sum_{\substack{u \in \mathcal{U} \\ u \cap \gamma \neq \emptyset}} \Bigg(\sum_{\substack{v \in \mathcal{V} \\ v \cap u \neq \emptyset}} \sigma(v) \Bigg) \leq \sum_{\substack{u \in \mathcal{U} \\ u \cap \gamma \neq \emptyset}} K\rho(u) = K\ell_\rho(\gamma, \mathcal{U}).$$

So for any weighted multi-curve $\gamma = \sum w_i \gamma_i$ in Γ, we therefore have by linearity that

$$\ell_\sigma(\gamma, \mathcal{V}) \leq K\ell_\rho(\gamma, \mathcal{U}).$$

Taking the infimum over the set Γ, we conclude

$$L_\rho(\Gamma, \mathcal{U}) \geq \frac{1}{K} \cdot L_\sigma(\Gamma, \mathcal{V}).$$

From the definition and K-bounded overlap, it is straightforward to see that

$$V_q(\rho, \mathcal{U}) \leq K \cdot V_q(\sigma, \mathcal{V}).$$

We conclude from the definition of modulus as the infimum of the ratio $V_q(\cdot)/(L(\cdot))^q$ that

$$\operatorname{mod}_q(\Gamma, \mathcal{U}) \leq K^{q+1} \operatorname{mod}_q(\Gamma, \mathcal{V}).$$

The other bound follows by symmetry. ∎

We need additional notation to set up the statement of the next lemma. For $\mathcal{S}$ a cover of X, let $\mathcal{S}^2$ be the cover whose elements are $\bigcup \mathcal{S}(s)$ for $s \in \mathcal{S}$. (That is, we take the union of all elements of $\mathcal{S}$ that intersect s.) Inductively define $\mathcal{S}^N$ to be the cover whose elements are $\bigcup \mathcal{S}(t)$ for $t \in \mathcal{S}^{N-1}$. The elements of $\mathcal{S}^N$ are in natural bijection with those of $\mathcal{S}$ but are larger "combinatorial $\mathcal{S}$-neighborhoods".

Lemma 5.6 (Fellow travelers). *Let X be a topological space equipped with a K-bounded finite cover $\mathcal{S}$. Fix a compact 1-manifold $C = \bigcup_i J_i$ and corresponding positive weights w_i, and suppose $\gamma^j = \sum_i w_i \gamma_i^j$, $j = 1, 2$ are two normalized weighted multi-curves given by maps $\gamma^j : C \to X$. Suppose there is an $N \in \mathbb{N}$ so that each strand γ_i^1 of γ^1 is contained in $\bigcup \mathcal{S}^N(\gamma_i^2)$, and similarly γ_i^2 is contained in $\bigcup \mathcal{S}^N(\gamma_i^1)$. Then for all $q \geq 1$, we have*

$$\operatorname{mod}_q(\gamma^1, \mathcal{S}) \asymp_{K,N,q} \operatorname{mod}_q(\gamma^2, \mathcal{S}).$$

Note that the statement does not concern *families* of multi-curves—just single multi-curves. The content of the lemma is that the bound depends only on the constants K, N, q.

Proof. Fix $q \geq 1$. Fix temporarily a test metric $\rho : \mathcal{S} \to [0, \infty)$. For $s \in \mathcal{S}$, we denote by $\widehat{s}$ the corresponding element of $\mathcal{S}^N$. We denote by $\widehat{\rho} : \mathcal{S}^N \to [0, \infty)$ the test metric obtained by setting $\widehat{\rho}(\widehat{s}) := \rho(s)$.

Fix a strand γ_i^2 of γ^2. By assumption, $\gamma_i^2 \subset \mathcal{S}^N(\gamma_i^1)$. This implies that if $s \in \mathcal{S}(\gamma_i^2)$, then $\widehat{s} \in \mathcal{S}^N(\gamma_i^1)$. Thus $\ell_{\widehat{\rho}}(\gamma_i^1, \mathcal{S}^N) \geq \ell_\rho(\gamma_i^2, \mathcal{S})$, and so via linear combinations and the definition, we have $\ell_{\widehat{\rho}}(\gamma^1, \mathcal{S}^N) \geq \ell_\rho(\gamma^2, \mathcal{S})$. Since γ^1, γ^2 are weighted curves, as opposed to families, we have $L_{\widehat{\rho}}(\gamma^1, \mathcal{S}^N) \geq L_\rho(\gamma^2, \mathcal{S})$. Since $V_q(\widehat{\rho}, \mathcal{S}^N) = V_q(\rho, \mathcal{S})$, we conclude

$$\frac{V_q(\widehat{\rho}, \mathcal{S}^N)}{L_{\widehat{\rho}}(\gamma^1, \mathcal{S}^N)^q} \leq \frac{V_q(\rho, \mathcal{S})}{L_\rho(\gamma^2, \mathcal{S})^q}.$$

Taking ρ to realize the infimum of the right-hand ratio, we conclude

$$\operatorname{mod}_q(\gamma^1, \mathcal{S}^N) \leq \operatorname{mod}_q(\gamma^2, \mathcal{S}).$$

The covering $\mathcal{S}^N$ has K^N-bounded overlap with the covering $\mathcal{S}$. By Lemma 5.5, we conclude

$$\operatorname{mod}_q(\gamma^1, \mathcal{S}) \lesssim_{K,N,q} \operatorname{mod}_q(\gamma^1, \mathcal{S}^N) \leq_{K,N,q} \operatorname{mod}_q(\gamma^2, \mathcal{S}).$$

The other bound follows by symmetry. ∎

The next two lemmas are technically convenient.

A graph G has a natural covering $\mathcal{E}$ whose elements are the closed edges of G. There is also a *partial covering* $\mathcal{E}^\circ$ whose elements are the interiors of the edges of G. Since $\mathcal{E}^\circ$ does not cover all of G, there are certain curve families (constant curves at vertices) with length 0 and undefined modulus. However, any non-constant curve intersects at least one element of $\mathcal{E}^\circ$, so if we restrict ourselves to families of non-constant curves, we do not run into trouble defining q-modulus with respect to $\mathcal{E}^\circ$.

Lemma 5.7. *If G is a graph of valence bounded by M and Γ is any family of weighted multi-curves on G, not necessarily normalized, each of whose strands is non-constant, then for any $1 \leq q \leq \infty$,*

$$\operatorname{mod}_q(\Gamma, \mathcal{E}) \asymp_M \operatorname{mod}_q(\Gamma, \mathcal{E}^\circ).$$

Proof. This is similar to Lemma 5.5, but that lemma does not apply directly, since $\mathcal{E}^\circ$ is not a covering. But the same techniques work. Note that there is a canonical bijection $\mathcal{E} \leftrightarrow \mathcal{E}^\circ$, and so a test metric ρ on one determines uniquely a test metric on the other that we denote with the same symbol.

For any test metric ρ on $\mathcal{E}^\circ$, we have $L_\rho(\Gamma, \mathcal{E}^\circ) \leq L_\rho(\Gamma, \mathcal{E})$. Since $V_q(\rho, \mathcal{E}) = V_q(\rho, \mathcal{E}^\circ)$, it follows that

$$\operatorname{mod}_q(\Gamma, \mathcal{E}) \leq \operatorname{mod}_q(\Gamma, \mathcal{E}^\circ).$$

Conversely, given a test metric ρ on $\mathcal{E}$, define a test metric σ on $\mathcal{E}^\circ$ by setting $\sigma(e)$ to be the maximum of $\rho(e')$ for e' equal to e or any of its neighbors. In a graph of valence at most M, a closed edge meets itself and at most $2(M-1)$ other closed edges. Thus $L_\rho(\Gamma, \mathcal{E}) \leq L_\sigma(\Gamma, \mathcal{E}^\circ)$ and $(2M-1)V_q(\rho, \mathcal{E}) \geq V_q(\sigma, \mathcal{E}^\circ)$, so

$$\operatorname{mod}_q(\Gamma, \mathcal{E}) \geq \frac{\operatorname{mod}_q(\Gamma, \mathcal{E}^\circ)}{2M-1}.$$ ∎

We use Lemma 5.7 to make cleaner estimates in the proof of Proposition 5.10, since the elements of the collection $\mathcal{E}^\circ$ do not overlap.

Definition 5.8 (Stars). Suppose G is a length graph with each edge of length 1, and e is an edge of G. The *open star* $\hat{e}$ of e is the open-$1/3$-neighborhood of e:

$$\hat{e} := \bigcup_{x \in \overline{e}} B\left(x, \frac{1}{3}\right).$$

We denote by $\widehat{\mathcal{E}}$ the covering of G by (open) stars of edges.

Lemma 5.9. *If G is a graph of valence bounded by M and Γ is any family of weighted multi-curves on G, not necessarily normalized, then for any $1 \leq q \leq \infty$,*

$$\mathrm{mod}_q(\Gamma, \mathcal{E}) \asymp_M \mathrm{mod}_q(\Gamma, \widehat{\mathcal{E}}).$$

Proof. The coverings $\mathcal{E}$ and $\widehat{\mathcal{E}}$ have M-bounded overlap. Apply Lemma 5.5. ∎

5.5. Combinatorial modulus and graph energies

Our next task is to relate the q-combinatorial modulus of the family comprised of a homotopy class $[\gamma]$ of weighted multi-curve γ with the graphical energy $E^1_q[\gamma]$ from Section 4.

Suppose $\gamma\colon C = \bigcup_i w_i J_i \to G^q$ is a weighted multi-curve on a q-conformal graph G^q, where each $J_i = S^1$. We have defined two quantities. From Section 4, we have the graphical energy $E^1_q[\gamma]$. We also have the combinatorial modulus $\mathrm{mod}_q([\gamma], \mathcal{E})$. The main result of this section is that these two quantities are comparable.

Proposition 5.10. *Suppose $q \geq 1$ and $m \geq 1$. Let G^q be a q-conformal graph with valence bounded by M and q-lengths $\alpha(e) \in [1/m, m]$ bounded above and below. Let $\mathcal{E}$ denote the covering of G^q by closed edges. Let $\gamma\colon C \to G$ be a nonempty weighted multi-curve on G^q with no null-homotopic components. Suppose N is an upper bound on the number of times a reduced representative in $[\gamma]$ passes over each edge of G^q. Then*

$$(E^1_q[\gamma])^q =_{\mathrm{def}} \mathrm{EL}_q([\gamma], \alpha) \asymp_{m,M,N,q} \mathrm{EL}_q([\gamma], \mathcal{E}) =_{\mathrm{def}} \frac{1}{\mathrm{mod}_q([\gamma], \mathcal{E})}.$$

Proof. Up to reparameterization, there is a unique reduced representative in $[\gamma]$; we assume that $\gamma = \sum_i w_i \gamma_i$ is this representative. We will then use Lemma 5.7 to work with the partial covering $\mathcal{E}^\circ$ rather than $\mathcal{E}$.

We then have a quantity $\mathrm{mod}_q(\gamma, \mathcal{E}^\circ)$. Since $\{\gamma\} \subset [\gamma]$, we have $\mathrm{mod}_q(\gamma, \mathcal{E}^\circ) \leq \mathrm{mod}_q([\gamma], \mathcal{E}^\circ)$. Since the underlying set of any strand of a curve in $[\gamma]$ contains that of the corresponding strand in the reduced representative γ, we have $\mathrm{mod}_q(\gamma, \mathcal{E}^\circ) \geq \mathrm{mod}_q([\gamma], \mathcal{E}^\circ)$. Thus $\mathrm{mod}_q(\gamma, \mathcal{E}^\circ) = \mathrm{mod}_q([\gamma], \mathcal{E}^\circ)$. Taking reciprocals, we conclude $\mathrm{EL}_q(\gamma, \mathcal{E}^\circ) = \mathrm{EL}_q([\gamma], \mathcal{E}^\circ)$.

Let $\rho\colon \mathcal{E}^\circ \to \mathbb{R}$ be a test metric; we will assume $\rho(e) > 0$ for each edge. (This does not change the relevant infima/suprema, although it does mean they will not always be realized.) By assumption, $n_\gamma(e) \le N$. We have then

$$L_\rho(\gamma, \mathcal{E}^\circ) = \ell_\rho(\gamma, \mathcal{E}^\circ) = \sum_i w_i \ell_\rho(\gamma_i, \mathcal{E}^\circ) = \sum_i w_i \sum_{e\subset\gamma_i} \rho(e) \asymp_N \sum_e n_\gamma(e)\rho(e);$$

the first equality coming from the fact that γ is a single multi-curve and not a family. For fixed nonzero ρ, consider the length graph K^∞_ρ with underlying graph G but lengths ρ on the edges, and let $\psi_\rho\colon G^q \to K^\infty_\rho$ be the identity map (which we introduce to keep track of which lengths to use). Then the right-hand side above is $E^1_\infty(\psi_\rho \circ \gamma)$. We conclude

$$L_\rho(\gamma, \mathcal{E}^\circ) \asymp_N E^1_\infty(\psi_\rho \circ \gamma).$$

In addition,

$$V_q(\rho, \mathcal{E}^\circ) = \sum_e \rho(e)^q \asymp_{m,q} \sum_e \alpha(e)^{1-q}\rho(e)^q = (E^q_\infty(\psi_\rho))^q,$$

so

$$\mathrm{EL}_q(\gamma, \mathcal{E}^\circ) \overset{\text{def}}{=} \sup_\rho \frac{L_\rho(\gamma, \mathcal{E}^\circ)^q}{V_q(\rho, \mathcal{E}^\circ)} \asymp_{m,N,q} \sup_\rho \left(\frac{E^1_\infty(\psi_\rho \circ \gamma)}{E^q_\infty(\psi_\rho)}\right)^q =_{\text{Proposition 4.10}} \mathrm{EL}_q([\gamma], \alpha),$$

since maximizing over ρ is equivalent to maximizing over metric graphs K^∞. ∎

5.6. Conformal dimension and combinatorial modulus

In this subsection, we give the relationship between critical exponents for combinatorial modulus and conformal dimension used in our main result.

Theorem 5.11 ([12, Corollary 1.4]). *Let X be a connected, locally connected, approximately self-similar metric space. For $\delta > 0$, denote by Γ_δ the family of curves of diameter bounded below by δ and by Γ^w_δ the corresponding family of normalized weighted multi-curves from Section 5.2. Suppose $(\mathcal{S}_n)_{n=0}^\infty$ is a sequence of snapshots at some scale parameter.*

Then the Ahlfors-regular conformal gauge of X is nonempty, and for some $\delta_0 > 0$ and all $0 < \delta < \delta_0$,

$$\mathrm{ARCdim}(X) = \inf\{q \mid \mathrm{mod}_q(\Gamma^w_\delta, \mathcal{S}_n) \to 0 \text{ when } n \to +\infty\}.$$

This theorem in the case of unweighted curves was proved by Carrasco and, in unpublished work, Keith and Kleiner. The statement as given above follows from their result and Lemma 5.3.

The results above let us characterize the conformal dimension in terms of curve families. Consider the limit dynamical system $f\colon \mathcal{J} \to \mathcal{J}$ associated to a forward-expanding recurrent virtual graph endomorphism $\pi, \phi\colon G_1 \rightrightarrows G_0$. We have a covering $\mathcal{U}_0$ of $\mathcal{J}$ by connected open sets, which by Theorem F we can choose to be arbitrarily fine. Via iterated pullback, we get a sequence $\mathcal{U}_n$ of finite covers that form a sequence of snapshots at some scale parameter θ with respect to a visual metric d_{vis}. We then have the following.

Corollary 5.12. *For all sufficiently small $\delta > 0$,*

$$\begin{aligned}\mathrm{ARCdim}(\mathcal{J}, d_{\mathrm{vis}}) &= \inf\{q \mid \mathrm{mod}_q(\Gamma^w_\delta, \mathcal{U}_n) \to 0 \text{ when } n \to +\infty\}\\ &= \sup\{q \mid \mathrm{mod}_q(\Gamma^w_\delta, \mathcal{U}_n) \to \infty \text{ when } n \to +\infty\}.\end{aligned}$$

At the critical exponent, the combinatorial modulus is bounded from below.

Proof. The first equality holds by Theorem 5.11; the second equality and the last assertion follow from [5, Corollary 3.7]. ■

6. Sandwiching ARCdim

To prove Theorem A, we begin by supposing $\pi, \phi\colon G_1 \rightrightarrows G_0$ is a λ-backward-contracting ($\lambda < 1$), recurrent, virtual endomorphism of graphs.

In fact, we will prove something slightly stronger, namely, that the conclusion holds under the weaker assumption that $\overline{E}^\infty[\pi, \phi] < \lambda < 1$. (The earlier assumption was that $E^\infty(\phi) < 1$.) We begin by making some basic simplifications.

First, we may assume G_0 has no vertices of valence 1. Next, a map $\phi\colon G_1 \to G_0$ between connected graphs without leaves which induces a surjection on the fundamental group is necessarily surjective. Finally, the pullback of a surjective map under a connected cover is again surjective. In summary, we may assume G_0, hence G_1 and indeed all the G_n have no leaves, and that for all n, each of ϕ_n^{n+1}, ϕ_0^n, and ϕ_n^∞ is surjective.

There is a uniform upper bound, say M, of the valences of the graphs G_n, $n \in \mathbb{N}$.

For technical reasons, we subdivide any loops of G_0 by adding a vertex on the interior. Since asymptotic energies are independent of the chosen metric, for later convenience, we choose the metric so that each edge of G_0 has length 1. Together, these conventions serve to make the star (recall Definition 5.8) of any edge at any level contractible.

The assumption $\overline{E}^\infty[\pi, \phi] < \lambda < 1$ implies there exists $N \geq 1$ such that $\phi_0^N\colon G_N \to G_0$ is contracting, say with some other constant $\lambda' < 1$. We iterate the pair (π, ϕ) so that the new G_1 is the old G_N, and reset $\lambda := \lambda'$. This modification does not alter the critical exponents q_*, q^*.

We apply the construction in Section 2 to present the limit space $\mathcal{J}$ as an inverse limit; we let G_n, ϕ_n^∞, π_n^∞, etc., be the corresponding spaces and maps as in that section. Theorem F implies that the limit dynamics $f\colon \mathcal{J} \to \mathcal{J}$ associated to (π, ϕ) is topologically cxc with respect to a finite open cover $\mathcal{U}_0$ by open connected sets such that the mesh of $\mathcal{U}_n$ tends to zero as $n \to \infty$. We may also take the mesh of $\mathcal{U}_0$ to be as small as we like.

We equip the limit space $\mathcal{J}$ with a visual metric d_{vis} with parameter $\theta < 1$ as in Theorem 3.2. The collection of coverings $(\mathcal{U}_n)_n$ is then a sequence of snapshots with parameter θ (Section 3.2). Inconveniently, the elements of $\mathcal{U}_n$ are not clearly related to structures in the graphs G_n. We therefore define another collection of coverings of $\mathcal{J}$ as follows. Given $n \in \mathbb{N}$ and $e \in E(G_n)$, we denote by

$$V(e) := (\phi_n^{\infty})^{-1}(\hat{e}) \subset \mathcal{J}$$

the open subset of $\mathcal{J}$ lying over the star of e in G_n; $\mathcal{V}_n$ is the covering of $\mathcal{J}$ by these sets. Unlike $\mathcal{U}_n$, the sets in $\mathcal{V}_n$ are not necessarily connected.

In Section 6.1, we show that the coverings $\mathcal{U}_n$ and $\mathcal{V}_n$ are quantitatively comparable. As a consequence, the sequence $\mathcal{V}_n$ is also a sequence of snapshots. In what follows, we work exclusively with the sequence of coverings $(\mathcal{V}_n)_n$.

In Section 6.2, we relate weighted multi-curves on G_n to weighted multi-curves on $\mathcal{J}$ and show comparability of moduli with respect to natural coverings. We do this by applying the constructions in Section 2.7.

Sections 6.3 and 6.4 give the proofs of the upper and lower bounds on $\mathrm{ARCdim}(\mathcal{J})$, respectively.

6.1. Snapshots from stars

Lemma 6.1. *We have the following properties of the $\mathcal{V}_n$:*

(1) *For $0 \leq k < n$, e an edge of G_k, and $\tilde{e}$ an edge of G_n with $\pi_n^k(\tilde{e}) = e$, the map $f^{n-k}\colon V(\tilde{e}) \to V(e)$ is a homeomorphism.*

(2) *With respect to the visual metric,* $\mathrm{mesh}(\mathcal{V}_n) \to 0$.

Proof. The first assertion follows from the fact that stars are contractible, ϕ_n^{∞} is the pullback of ϕ_k^{∞} under π_k^n, the presentation of $\mathcal{J}$ as the inverse limit of the sequence $(\phi_n^{n+1})_n$, and the definition of the dynamics on the limit space $f\colon \mathcal{J} \to \mathcal{J}$ as induced by π.

For the second, we argue directly, using the definition of the product topology. For $x \in \mathcal{J}$, write $x = (x_0, x_1, x_2, \ldots) \in \prod_{i=0}^{\infty} G_i$. Then there is a neighborhood basis of x in the product topology on $\mathcal{J}$ consisting of sets of the form

$$U(x; \varepsilon, M) := \left(\prod_{i=1}^{M} B_{G_i}(x_i, \varepsilon) \times \prod_{i=M+1}^{\infty} G_i \right) \cap \mathcal{J}$$

for $\varepsilon > 0$ and $M \in \mathbb{N}$. Fix such a neighborhood $U(x; \varepsilon, M)$. Choose $N > M$ large enough that $2\lambda^{N-M} < \varepsilon$, and suppose $n > N$, $e \in E(G_n)$, and $x \in V(e)$. Since $\mathrm{diam}_{G_n}(\hat{e}) < 2$, we have for $i \leq M$ that

$$\mathrm{diam}_{G_i}(\phi_i^n(\hat{e})) < 2\lambda^{n-i} \leq 2\lambda^{N-M} < \varepsilon.$$

Hence for each $0 \leq i \leq M$, we have $\phi_i^{\infty}(V(e)) \subset B_{G_i}(x_i, \varepsilon)$, and so $V(e) \subset U(x; \varepsilon, M)$. The conclusion follows since $\mathcal{J}$ is compact, and the topology determined by the visual metric is the same as that induced from the product topology. ∎

Lemma 6.2 (The $\mathcal{U}_n$ are comparable to the $\mathcal{V}_n$). *There exists $N_0, N_1 \in \mathbb{N}$ with the following property.*

For each $n \geq N_0$ and each $V \in \mathcal{V}_n$, there exist $U \in \mathcal{U}_{n-N_0}$ and $U' \in \mathcal{U}_{N_1+n-N_0}$ such that

$$U' \subset V \subset U.$$

Recall that *Lebesgue number* of a finite open cover $\mathcal{U}$ is the largest $\delta > 0$ such that any subset of diameter less than δ is contained in an element of $\mathcal{U}$.

Proof. Lemma 6.1 implies there exists N_0 such that the mesh of $\mathcal{V}_{N_0}$ is smaller than the Lebesgue number of $\mathcal{U}_0$.

We first show the conclusion holds for $n = N_0$. Fix $V \in \mathcal{V}_{N_0}$ so that $V = V(e)$, where $e \in E(G_{N_0})$. The choice of N_0 implies that $V \subset U$ for some $U \in \mathcal{U}_0$. To find U', let y be the midpoint of e. By the uniform continuity of $\phi^{\infty}_{N_0}$, there exists $\delta > 0$ such that the image of any visual δ-ball under $\phi^{\infty}_{N_0}$ has diameter at most $1/4$. Theorem 3.2 implies there exists $N_1 > N_0$ such that $\operatorname{diam} U' < \delta$ for each $U' \in \mathcal{U}_{N_1}$. Pick any $x \in \mathcal{J}$ with $\phi^{\infty}_{N_0}(x) = y$, and pick any $U' \in \mathcal{U}_{N_1}$ containing y. Then $\phi^{\infty}_{N_0}(U') \subset B(x_{N_0}(e), 1/4) \subset e \subset \widehat{e}$, and so $U' \in V(e)$.

We now prove the conclusion holds for general $n \geq N_0$. Given such $V \in \mathcal{V}_n$, rename it $\widetilde{V}$, put $V := f^{n-N_0}(\widetilde{V})$, and apply the argument above to obtain U' and U. Now let $\widetilde{U}'$ and $\widetilde{U}$ be connected components of the preimages of U' and U, respectively, under f^{n-N_0} meeting $\widetilde{V}$. Recalling the construction in Section 3.1, each maps homeomorphically to its image. Renaming $\widetilde{U}'$ to U' and $\widetilde{U}$ to U yields the result. ■

Proposition 6.3. *There exists $m \in \mathbb{N}$ such that the family $(\mathcal{V}_n)_{n \geq m}$ is a family of snapshots with parameter θ.*

Proof. Let r_0 be the constant from Theorem 3.2 (4). Lemma 6.1 (2) implies there is $m \in \mathbb{N}$ such that for $n \geq m$, each $V \in \mathcal{V}_n$ has diameter at most r_0, and hence the restriction $f\colon V \to f(V)$ is a homeomorphism which scales distances by the factor θ^{-1}.

Recall the graphs G_n are equipped with length metrics in which edges are isometric to the unit interval. For each $V = V(e) \in \mathcal{V}_m$, let $x_V \in \mathcal{J}$ be a point which projects under ϕ^{∞}_m to the midpoint of the edge e. Since there are finitely many edges at level m, we can find $s > 0$ such that for each $V \in \mathcal{V}_m$, $\operatorname{diam}(\phi^{\infty}_m(B(x_V, s))) < 1/2$. This choice of s guarantees that the visual balls $B(x_V, s)$ for $V \in \mathcal{V}_m$ are pairwise disjoint. Let $D = \max\{\operatorname{diam} V \mid V \in \mathcal{V}_m\}$. For each $V \in \mathcal{V}_m$, we have

$$B(x_V, s) \subset V \subset B(x_V, D)$$

so that if we put $C := \max\{D/\theta^m, \theta^m/s\}$, then $C > 1$, and

$$B(x_V, C^{-1}\theta^m) \subset V \subset B(x_V, C\theta^m).$$

The finite collection of pointed sets $\{(V, x_V) \mid V \in \mathcal{V}_m\}$ then satisfies condition (1) in Definition 3.1 on sequences of snapshots. By construction, $C^{-1}\theta^m < s$ and it follows that condition (2) holds as well.

Now suppose $n > m$ and $\widetilde{V} \in \mathcal{V}_n$. The restriction $f^{n-m}|\widetilde{V}$ is an expanding homothety with factor θ^{m-n} onto say $V \in \mathcal{V}_m$. Put $x_{\widetilde{V}} := (f^{n-m}|\widetilde{V})^{-1}(x_V)$. Using the same constant C constructed in the previous paragraph, it follows that the sequence $(\mathcal{V}_n)_{n\geq m}$ is a sequence of snapshots of $\mathcal{J}$ with parameter θ. ∎

For convenience, we reindex our virtual endomorphism and sequence of covers $(\mathcal{V}_n)_n$ so that $m = 0$, as in equation (2.19).

6.2. Comparing modulus of curves on G_n and on $\mathcal{J}$

Suppose $\sigma^n_\infty \colon G_n \to \mathcal{J}$ is the homotopy section of ϕ^∞_n given by Corollary 2.33, and suppose $\gamma \colon C \to G_n$ is a weighted multi-curve on G_n. Define the curve $\gamma' := \sigma^n_\infty \circ \gamma$ on $\mathcal{J}$, and the curve $\gamma'' := \phi^\infty_n \circ \gamma'$ on G_n. By Corollary 2.33, there exists $K > 0$ so that

$$\gamma \sim_K \gamma''.$$

It is important that K is independent of n. Though γ, γ', and γ'' are single curves, they nevertheless have a meaningful combinatorial modulus when regarded as a family; see Remark 5.2.

Recall from Section 5.4 that, on G_n, there are coverings $\mathcal{E}_n$ by closed edges, and $\widehat{\mathcal{E}}_n$ by open stars. Also, recall that M is a uniform upper bound on the degrees of the graphs G_n, $n \in \mathbb{N}$.

Lemma 6.4. *For all $q \geq 1$ and any weighted multi-curve γ on G_n, with γ' and γ'' as above, we have*

$$\operatorname{mod}_q(\gamma', \mathcal{V}_n) = \operatorname{mod}_q(\gamma'', \widehat{\mathcal{E}}_n) \asymp_M \operatorname{mod}_q(\gamma'', \mathcal{E}_n) \asymp_{K,M,q} \operatorname{mod}_q(\gamma, \mathcal{E}_n).$$

In particular, the implicit constants are independent of n and γ.

Tracing through the dependencies, note that the constant K depends on the Lipschitz constant $\lambda < 1$ of ϕ.

Proof. By definition, the map ϕ^∞_n sends γ' to γ'' and induces an isomorphism between the nerve of $\mathcal{V}_n$ and the nerve of $\widehat{\mathcal{E}}_n$. The elements $V(e)$ are given by $(\phi^\infty_n)^{-1}(\widehat{e})$; in particular, they are saturated with respect to the fibers of ϕ^n_∞. Thus $\operatorname{mod}_q(\gamma', \mathcal{V}_n) = \operatorname{mod}_q(\gamma'', \widehat{\mathcal{E}}_n)$. The middle estimate is the content of Lemma 5.9. For the last estimate, γ and γ'' are in uniform combinatorial $\mathcal{E}_n$-neighborhoods of each other since $\gamma \sim_K \gamma''$. Then Lemma 5.6 (with its $K := M$ and its $N := K$) implies that $\operatorname{mod}_q(\gamma'', \mathcal{E}_n) \asymp_{K,M,q} \operatorname{mod}_q(\gamma, \mathcal{E}_n)$. ∎

6.3. Upper bound on ARCdim

Let δ_0 be as in Theorem 5.11, and fix $0 < \delta < \delta_0$.

Proposition 6.5. *Suppose that $q > 1$ and $\overline{E}^q[\pi, \phi] < 1$. Let Γ_δ be the family of curves in $\mathcal{J}$ whose diameters are bounded below by δ. Then $\operatorname{mod}_q(\Gamma_\delta, \mathcal{V}_n) \to 0$ as $n \to \infty$.*

Proof. We begin with the following claim. There is an integer N such that for each $\gamma \in \Gamma_\delta$, there is $e \in E(G_N)$ so that $e \subset \phi_N^\infty(\gamma)$. To see this, we argue as follows. Using Lemma 6.1, choose N so that the mesh of $\mathcal{V}_N$ is at most $\delta/10$. Pick $x^1, x^2 \in \gamma$ with $d_{\text{vis}}(x^1, x^2) > \delta/2$. Choose closed edges $e_1, e_2 \in E(G_N)$ such that $\phi_N^\infty(x_i) \in e_i$. The triangle inequality shows that $e_1 \cap e_2 = \emptyset$. The image $\phi_N^\infty(\gamma)$ is a connected set in G_N which meets two disjoint closed edges, so it must contain an entire edge e.

Let the level N be as in the previous paragraph. For $e \in E(G_N)$, let Γ_e be the family of curves γ in $\mathcal{J}$ so that $e \subset \phi_N^\infty(\gamma)$. Then

$$\Gamma_\delta \subset \bigcup_{e \in E(G_N)} \Gamma_e,$$

and so, for $n \geq N$,

$$\operatorname{mod}_q(\Gamma_\delta, \mathcal{V}_n) \leq \sum_{e \in E(G_N)} \operatorname{mod}_q(\Gamma_e, \mathcal{V}_n).$$

It therefore suffices to show that $\operatorname{mod}_q(\Gamma_e, \mathcal{V}_n) \to 0$ for each $e \in E(G_N)$.

The proposition deals with the asymptotic energy $\overline{E}^q[\pi, \phi]$ and the limit set $\mathcal{J}$. Neither the property that $\overline{E}^q[\pi, \phi] < 1$ nor the conformal gauge of the limit space is changed by reindexing, changing the weights, or iterating. We may therefore choose q-lengths $\alpha \equiv 1$ on G_0 (and thus on G_N) and then reindex to assume $N = 0$. We now have a q-conformal structure G_0^q on G_0. By lifting under the coverings π_0^n, we obtain q-conformal structures G_n^q on G_n for each n. Since $\overline{E}^q[\pi, \phi] < 1$, by iterating we may assume $\phi\colon G_1^q \to G_0^q$ satisfies $E_q^q(\phi) = \lambda < 1$ and so $\phi_0^n\colon G_n^q \to G_0^q$ satisfies $E_q^q(\phi^n) \leq \lambda^n$ for all $n > 0$.

Fix some $e_0 \in E(G_0)$. We must show $\operatorname{mod}_q(\Gamma_{e_0}, \mathcal{V}_n) \to 0$ as $n \to \infty$. Fix $n \in \mathbb{N}$. Recall elements of $\mathcal{V}_n$ are preimages of stars $(\phi_n^\infty)^{-1}(\widehat{s})$. For brevity, we will write $\rho(s)$ instead of $\rho((\phi_n^\infty)^{-1}(\widehat{s}))$. Let x be a local length coordinate on the edge s. Define a test metric $\rho\colon \mathcal{V}_n \to [0, \infty)$ by

$$\rho(s) = \int_{x \in s} |(\phi^n)'(x)||dx|. \tag{6.6}$$

In other words, $\rho(s)$ is the length of $\phi^n(s)$ regarded as a curve on the length graph underlying G_0, which we denote $\ell_{G_0}(\phi^n(s))$. (Note that the image $\phi^n(s)$ may backtrack.)

We now show that this ρ is admissible for Γ_{e_0}. Pick $\gamma \in \Gamma_{e_0}$, and set $\beta := \phi_n^\infty(\gamma)$; it is a curve in G_n. We have

$$e_0 \subset \phi_0^\infty(\gamma) = \phi_0^n(\phi_n^\infty(\gamma)) = \phi_0^n(\beta).$$

Let $s_1, \dots, s_m$ be the edges in $E(G_n)$ met by β. Then

$$\ell_\rho(\gamma, \mathcal{V}_n) = \sum_{j=1}^{m} \rho(s_j) = \sum_{j=1}^{m} \ell_{G_0}(\phi^n(s_j)) \geq \ell_{G_0}(e_0) = 1.$$

The last inequality works even in the presence of backtracking: the total length in G_0 covered by the $\phi^n(s_j)$ is at least as long as e_0, even if β crosses over a given edge s_j multiple times.

Next we estimate the q-volume $V_q(\rho, \mathcal{V}_n)$; for Jensen's inequality, see, for example, [45, Theorem 3.3]. We have

$$\begin{aligned}
V_q(\rho, \mathcal{V}_n) &= \sum_{s \in \mathrm{Edge}(G_n)} \rho(s)^q && \text{Definition of } V_q \\
&= \sum_s \Big(\int_{x \in s} |(\phi^n)'(x)|\, dx \Big)^q && \text{Definition of } \rho \\
&\leq \sum_s \int_{x \in s} |(\phi^n)'(x)|^q\, dx && \text{Jensen's inequality} \\
&= \int_{x \in G_n} |(\phi^n)'(x)|^q\, dx && \\
&= \int_{x \in G_0} \mathrm{Fill}^q(\phi^n)\, dx && \text{Change of variables from } G_n \text{ to } G_0 \\
&\leq \|\mathrm{Fill}^q(\phi^n)\|_{\infty, G_0} \cdot \Big(\int_{x \in G_0} 1\, dx \Big) && \\
&\asymp (E_q^q(\phi^n))^q && \text{Definition of } E_q^q \\
&< \lambda^{nq}.
\end{aligned}$$

We conclude $\mathrm{mod}_q(\Gamma_{e_0}, \mathcal{V}_n) \to 0$, as required. ■

Proposition 6.5 and Theorem 5.11 then imply $\mathrm{ARCdim}(\mathcal{J}) \leq q^*[\pi, \phi]$.

Remark 6.7. The calculations in the proof of Proposition 6.5 can be interpreted as follows. Let K_0^∞ be G_0, considered as a length graph with all edge lengths 1. Let $\psi^n \colon G_n^q \to K_0^\infty$ be ϕ^n, homotoped to be constant derivative on each edge of G_n^q. Then, if you trace through the definitions, $V_q(\rho, \mathcal{V}_n) = \big(E_\infty^q(\psi^n)\big)^q$, and the relevant inequalities are

$$\begin{aligned}
E_\infty^q(\psi^n) \leq E_\infty^q(\phi^n) &\leq E_q^q(\phi^n) \cdot E_\infty^q(\mathrm{id}\colon G_0^q \to K_0^\infty) \\
&< \lambda^n \cdot E_\infty^q(\mathrm{id}\colon G_0^q \to K_0^\infty) \asymp \lambda^n.
\end{aligned}$$

6.4. Lower bound on ARCdim

For this inequality, we will use more of the theory from Section 5. Again let δ_0 be as in Theorem 5.11, and fix $\delta < \delta_0$ that is sufficiently small (to be specified).

Proposition 6.8. *Suppose* $\overline{E}^q[\pi,\phi] > 1$*. Then* $\operatorname{mod}_q(\Gamma^w_\delta, \mathcal{V}_n) \to \infty$.

Proof. Set $\mu = \overline{E}^q[\pi,\phi] > 1$; then $E^q_q[\phi^n] \geq \mu^n$ [52, Proposition 5.6]. We are going to find a normalized weighted multi-curve ζ' on $\mathcal{J}$ in Γ^w_δ such that $\operatorname{mod}_q(\zeta', \mathcal{V}_n) \to \infty$ as $n \to \infty$. This will imply $\operatorname{mod}_q(\Gamma^w_\delta, \mathcal{V}_n) \to \infty$, as required.

Fix $n \in \mathbb{N}$. Recall from Section 5.4 that the graph G_n is naturally covered by its set $\mathcal{E}_n$ of closed edges. By Theorem 4.11, we can find a curve exhibiting $E^q_q[\phi^n]$: there is a reduced weighted multi-curve $\gamma\colon C^1 \to G^q_n$ and a map $\psi\colon G^q_n \to G^q_0$ that minimizes E^q_q in the homotopy class $[\phi^n]$ and fits into a tight sequence

$$C^1 \xrightarrow{\gamma} G^q_n \xrightarrow{\psi} G^q_0. \tag{6.9}$$

(Since the exponent q in G^q_n is fixed, we will suppress it from the notation.) That is,

$$\mu^n \leq E^q_q(\psi) = \frac{E^1_q(\psi \circ \gamma)}{E^1_q(\gamma)} = \frac{E^1_q[\phi^n \circ \gamma]}{E^1_q[\gamma]}. \tag{6.10}$$

This goes to infinity as $n \to \infty$. Furthermore, Theorem 4.11 guarantees that for each strand J_i of C, the restriction $\gamma|J_i$ has image covering a given edge of G_n at most twice so that Proposition 5.10 applies.

The weighted multi-curve $\gamma\colon C^1 \to G_n$ is not unique. In particular, we can and will scale the weights on C^1 so that $E^1_q[\psi \circ \gamma] = 1$. Then $1/E^1_q[\gamma] \geq \mu^n$, and Proposition 5.10 implies that $\operatorname{mod}_q([\gamma], \mathcal{E}_n) \gtrsim \mu^{nq}$.

We next apply the constructions in Section 6.2 to obtain curves $\gamma'\colon C \to \mathcal{J}$ and $\gamma''\colon C \to G_n$ with $\gamma'' := \sigma^n_\infty \circ \gamma'$ and $\gamma \sim_K \gamma''$ for some constant $K > 0$ independent of n. Lemma 6.4 implies

$$\operatorname{mod}_q(\gamma', \mathcal{V}_n) = \operatorname{mod}_q(\gamma'', \mathcal{E}_n) \asymp \operatorname{mod}_q(\gamma, \mathcal{E}_n) \gtrsim \mu^{nq}.$$

Though the diameters of the strands of γ' are bounded from below and the modulus is blowing up, we do not yet have control on the weights of γ'; usually, γ' will not be in Γ^w_δ. (Correspondingly, the strands in γ' are much longer than any constant δ.) We will remedy this by subdividing γ' to obtain a more suitable curve, as in Section 5.3.

The curve $\psi \circ \gamma\colon C \to G_0$ is reduced. The strands $\psi \circ (\gamma \mid J_i)\colon J_i \to G_0$ may run many times over a given edge of G_0. We decompose $\psi \circ \gamma$ into separate pieces, one for each time such a strand runs over an edge of G_0. Formally, for a fixed pair (i, e) consisting of a strand $\psi \circ (\gamma \mid J_i)$ and an edge e of G_0, suppose the restriction $\psi \circ (\gamma \mid J_i)$ runs $N(i, e)$ times over e. We obtain a collection of sub-intervals $I_{i,e,k}$ of J_i for which $\psi \circ \gamma\colon I_{i,e,k} \to e$ is a homeomorphism. (Note the $I_{i,e,k}$ are disjoint except at their endpoints, and since each strand J_i of c is a circle, there are no end issues to worry about.) We identify each $I_{i,e,k}$ with the unit interval and obtain a weighted multi-curve $\zeta\colon D \to G_n$ as follows. With w_i the weight of J_i in C, set

$$D := \bigsqcup_i \bigsqcup_e \bigsqcup_{k=1}^{N(i,e)} w_i I_{i,e,k}.$$

Let $\iota\colon D \to C$ be the natural inclusion of intervals, and define $\zeta := \gamma \circ \iota\colon D^1 \to G_n$. By construction, the function $n_{\psi\circ\zeta}\colon G_0 \to \mathbb{R}^+$ is constant on edges of G_0 and coincides with $n_{\psi\circ\gamma}$.

Recall we have normalized, so $E^1_q[\psi \circ \gamma] = 1$, so with $q^\vee$ the Hölder conjugate of q,

$$1 = E^1_q[\psi \circ \gamma] = \|n_{\psi\circ\gamma}\|_{q^\vee} \asymp_{q,\#E(G_0)} \|n_{\psi\circ\gamma}\|_1 = \|n_{\psi\circ\zeta}\|_1 = \sum_{i,e,k} w_{i,e,k}.$$

At the third step, we use the fact that in $\mathbb{R}^{E(G_0)}$, any two norms are comparable. We conclude that the sum of the weights of D is comparable to 1.

We now decompose the curve $\gamma'\colon C \to \mathcal{J}$ with the same decomposition. Let $\zeta' = \gamma' \circ \iota\colon D \to \mathcal{J}$. We must show that the size of each strand of ζ' has diameter bounded below, independent of n. We focus attention on one component $I_{i,e,k}$ of D and identify that interval with $[0, 1]$. Let $\zeta'' = \phi_n^\infty \circ \zeta'\colon D \to G_n$. Since $\gamma \sim_K \gamma''$ and the decompositions ζ and ζ'' correspond, the endpoints $\zeta''(0)$ and $\zeta''(1)$ are within a uniformly bounded G_n-distance K of $\zeta(0)$ and $\zeta(1)$ (independent of n). Also recall that we assumed that the system is forward expanding, so ϕ is Lipschitz with some constant $\lambda < 1$ and ϕ^n is Lipschitz with constant λ^n. Then

$$\begin{aligned}
|\phi^\infty(\zeta'(0)) - \phi^\infty(\zeta'(1))| &= |\phi^n(\zeta''(0)) - \phi^n(\zeta''(1))| \\
&\geq |\phi^n(\zeta(0)) - \phi^n(\zeta(1))| - |\phi^n(\zeta''(0)) - \phi^n(\zeta(0))| \\
&\qquad - |\phi^n(\zeta(1)) - \phi^n(\zeta''(1))| \\
&\geq 1 - 2\lambda^{-n} K.
\end{aligned}$$

We suppose that n is large enough so that $1 - 2\lambda^{-n}K$ is bigger than $1/2$. We have shown that each strand of ζ', when projected to G_0, has diameter at least $1/2$. Since ϕ^∞ is uniformly continuous (as a function from a compact metric space), it follows that each strand of d' has definite diameter; we choose δ smaller than this diameter. Combining this with the observation that the sum of the weights of D is comparable to 1, we conclude that, after an innocuous rescaling of weights, ζ' lies in Γ^w_δ.

We then have

$$\begin{aligned}
\mu^{nq} &\leq \frac{1}{(E^1_q[\gamma])^q} \\
&\lesssim \operatorname{mod}_q([\gamma], \mathcal{E}_n) && \text{By Proposition 5.10} \\
&= \operatorname{mod}_q(\gamma, \mathcal{E}_n) && \text{Since } \gamma \text{ is reduced} \\
&\asymp \operatorname{mod}_q(\zeta, \mathcal{E}_n) && \text{By Lemma 5.4, subdivision} \\
&\asymp \operatorname{mod}_q(\zeta'', \mathcal{E}_n) && \text{By Lemma 6.4, fellow travelers} \\
&= \operatorname{mod}_q(\zeta', \mathcal{V}_n) && \text{By definition of the cover } \mathcal{V}.
\end{aligned}$$

For the conditions of Lemma 5.4 at the fourth step, note that the G_0-length of each strand of $\phi^n \circ \zeta$ is 1, so by forward expansion the G_n-length of each strand is at least λ^n, and in particular, for large n the image of each strand is not contained in a single edge of G_n.

This completes the proof of Proposition 6.8. ∎

Propositions 6.5 and 6.8 complete the proof of Theorem A.

7. Applications

We turn now to applications. Section 7.1 gives the proof of Theorem C on Sierpiński carpets. Section 7.2 proves the estimates for the barycentric subdivision example mentioned in the introduction. Sections 7.3 and 7.4 give the estimates for the fat and skinny Devaney examples. The brief Section 7.5 introduces Carrasco's UWSCP condition. Section 7.6 applies our methods to examples obtained by the operation of "mating" and concludes with a question about the relationship between the UWSCP condition and other properties.

For some estimates, we rely on explicit bounds on how fast $\overline{E}^q$ can decrease as a function of q [52, Proposition 6.11].

Proposition 7.1. *For $\pi, \phi\colon G_1 \rightrightarrows G_0$ a virtual endomorphism of graphs of degree $d := \deg(\pi)$, if $1 \le p \le q \le \infty$, then*

$$\overline{E}^q[\pi, \phi] \ge d^{-\frac{1}{p}+\frac{1}{q}} \cdot \overline{E}^p[\pi, \phi].$$

As an easy consequence, we have the following theorem announced in the introduction. Recall that the quantity $\overline{N} := \overline{E}_1^1$ (Definition 4.12) counts the asymptotic growth rate, as $n \to \infty$, of the essential number of preimages of ϕ_0^n, minimized among maps within its homotopy class.

Theorem B. *For any recurrent expanding virtual graph endomorphism $[\pi, \phi]$ where $\deg(\pi) = d$, we have*

$$\mathrm{ARCdim}[\pi, \phi] \ge \frac{1}{1 - \log_d \overline{N}[\pi, \phi]}.$$

Proof. By Proposition 7.1 with $p = 1$, for any q we have

$$\overline{E}^q[\pi, \phi] \ge d^{\frac{1}{q}-1}\, \overline{N}[\pi, \phi].$$

If q is less than the quantity in the theorem statement, the right-hand side is greater than 1. The result follows from Theorem A. ∎

7.1. Cases when $\overline{N} > 1$

Our proof of Theorem C is a corollary of a more general result about certain expanding dynamical systems on the sphere.

Theorem C′. *Suppose $g: S^2 \to S^2$ is an expanding Thurston map such that each cycle in its post-critical set P_g contains a critical point. If $\pi, \phi: G_1 \to G_0$ is any virtual graph endomorphism induced by a choice of a spine G_0 for $S^2 - P_g$, then $\overline{N}[\pi, \phi] > 1$.*

A *Thurston map* $g: S^2 \to S^2$ is an orientation-preserving branched self-cover of degree at least 2 such that the post-critical set $P_g := \bigcup_{n>0} g^n(\{\text{branch points}(g)\})$ is finite. The definition of *expanding* Thurston map appearing in the hypothesis of Theorem C′ here is that of Bonk and Meyer [4].

Here is an example of an expanding Thurston map. First, recall from Section 1.1 and Figure 2 that the rational map f there sends the small triangles conformally onto the large triangles. Next, consider the Thurston map g obtained from Figure 2 where now the small triangles are sent not conformally, but Euclidean-planar-affinely, to the large triangles. While the maps f and g are conjugate up to isotopy relative to their post-critical sets, they are not topologically conjugate, since the fixed branch point of g is locally topologically repelling, whereas that of f is attracting. Recall that the Julia set J_f is a Sierpiński carpet. It turns out (see below) that collapsing the Fatou components of f to points yields a map which is topologically conjugate to g. Our proof of Theorem C proceeds by passing to such a quotient and invoking Theorem C′.

The next few paragraphs summarize several results from [4] that we use in the proof.

Fix arbitrarily a metric on S^2 compatible with its topology. Suppose g is an arbitrary Thurston map, and $\mathcal{C} \subset S^2$ is a Jordan curve containing P_g. These data induce a cell structure $\mathcal{T}_0$ on S^2, with two open 2-cells, called *tiles*, given by the components of $S^2 - \mathcal{C}$. Lifting by iterates of g yields a sequence of cell structures $\mathcal{T}_n$ on S^2 for $n = 0, 1, 2, \ldots$; let $\mathrm{mesh}(g, n, \mathcal{C})$ be the maximum diameter of an open 2-cell, or *n-tile*, at level n. The Thurston map g is said to be *expanding* if $\mathrm{mesh}(g, n, \mathcal{C}) \to 0$ as $n \to \infty$; this property is independent of the choice of metric and of $\mathcal{C}$; see [4, Chapter 6]. There exists an iterate k and a Jordan curve $\mathcal{C}'$ isotopic to $\mathcal{C}$ relative to P_g such that $g^k(\mathcal{C}') \subset \mathcal{C}'$; see [4, Theorem 15.1].

In our applications, passing to such an iterate will be innocuous, so in this paragraph we now suppose that g is an expanding Thurston map and that $g(\mathcal{C}) \subset \mathcal{C}$ for some Jordan curve $\mathcal{C} \supset P_g$. Then for each n, the tiling $\mathcal{T}_{n+1}$ refines the tiling $\mathcal{T}_n$ according to a subdivision rule [4, Chapter 12]. A numerical invariant is then the *combinatorial expansion factor*

$$\Lambda_0(g) := \lim_{n\to\infty} D_n(g, \mathcal{C})^{1/n},$$

where $D_n(g, \mathcal{C})$ is, roughly speaking, the minimum number m of closed n-tiles in a chain $t_0, t_1, \ldots, t_m$ with $t_j \cap t_{j+1} \neq \emptyset$, $j = 0, \ldots, m-1$ such that t_0, t_m each contain 1-cells of $\mathcal{T}_0$ whose closures are disjoint, that is, the chain "joins disjoint closed edges" of $\mathcal{C}$. (A slight modification is needed in the case when $\#P_g = 3$ [4, Section 5.7].) We then have that $\Lambda_0(g)$ is independent of $\mathcal{C}$ and is greater than 1 [4, Proposition 16.1]. For each $1 < \theta^{-1} < \Lambda_0(g)$, there exists an analogously defined visual metric on S^2 with expansion factor θ^{-1} [4, Theorem 16.3 (ii)]. Fixing such a choice of metric with its expansion factor θ^{-1}, there is a constant $K > 1$ such that the diameter of each tile t at level n satisfies $\operatorname{diam}(t) \in [\theta^n/K, K\theta^n]$, by [4, Proposition 8.4 (ii)]; compare our Theorem 3.2 (1).

Proof of Theorem C′. Let g be an expanding Thurston map. The property $\overline{N} > 1$ is unchanged under passing to iterates, so we assume there is a Jordan curve $\mathcal{C} \supset P_g$ for which $g(\mathcal{C}) \subset \mathcal{C}$. Let $\Lambda_0(g)$ be the resulting combinatorial expansion factor. We choose a visual metric d_{vis}, and let θ^{-1} be its expansion factor.

In the remainder of the proof, we show $\theta^{-1} \leq \overline{N}$. Since θ^{-1} can be chosen arbitrarily subject to the constraint $\theta^{-1} \leq \Lambda_0(g)$ can be arbitrary, this is enough to conclude that $\overline{N} \geq \Lambda_0(g) > 1$ (see Remark 7.3).

Let $G_0 \subset S^2$ be a realization of the dual of $\mathcal{T}_0$; it is a spine for $S^2 - P_g$ with two vertices. As usual, let $G_n = g^{-n}(G_0)$ so that G_n is the dual of $\mathcal{T}_n$. Note that the mesh of the faces of G_n tends to zero as well with respect to d_{vis}. Fix $p \in P_g$. Let U_0 be the component of the complement of G_0 containing p and let $C_0 = \partial U_0$. Since G_0 is a spine for $S^2 - P_g$, any loop in G_0 that is freely homotopic to a peripheral loop about p contains C_0.

Claim 7.2. There exists $c > 0$ such that for all $n \in \mathbb{N}$, there exist at least $c\theta^{-n}$ pairwise disjoint loops $C_{n,i}$ in G_n freely homotopic to a peripheral loop about p.

Assuming the claim, we show $\theta^{-1} \leq \overline{N}$ as follows. For any $\psi \in [\phi^n]$, the curve $\psi(C_{n,i})$ is freely homotopic to C_0 and so contains C_0. Thus any $y \in C_0$ has, for each n and i, a preimage in $C_{n,i}$, and so $N[\phi^n] \gtrsim \theta^{-n}$, proving $\overline{N}(f) \geq \theta^{-1}$.

We now turn to the proof of Claim 7.2, using the fact that g is expanding on the whole sphere. Let $D = \min\{d_{\mathrm{vis}}(p,q) \mid q \in C_0\}$, let $N_0 \in \mathbb{N}$ be chosen so $D\theta^{-n}/K > 1$ for $n \geq N_0$, and for $n \geq N_0$, let m_n be the greatest positive integer less than or equal to $D\theta^{-n}/K$. Then for all such n, any chain of tiles $t_1, \dots, t_{m_n}$ in $\mathcal{T}_n$ with $p \in t_1$ and $t_i \cap t_{i+1} \neq \emptyset$ for $i = 1, \dots, m_n - 1$ avoids C_0.

For $n \geq N_0$, we will define curves $C_{n,1}, \dots, C_{n,m_n}$. Given $E \subset S^2$ and $n \in \mathbb{N}$, recall that $\mathcal{T}_n(E)$ is the union of the closed tiles at level n meeting E. Let $E_{n,1} = \{p\}$, and for $i = 1, \dots, m_n - 1$, inductively set $E_{n,i+1} := \mathcal{T}_n(E_{n,i})$. Fix one such i. Then $E_{n,i}$ is contained in the interior of $E_{n,i+1}$. The complement $E_{n,i+1} - \operatorname{interior}(E_{n,i})$ is tiled by elements of $\mathcal{T}_n$. Consider the corresponding subgraph of the dual graph G_n. We take $C_{n,i}$ to be a simple cycle in this dual graph that separates C_0 from p. The $C_{n,i}$ are pairwise disjoint, freely homotopic to C_0, and disjoint from C_0, by construction. This proves Claim 7.2 with c approximately D/K. Theorem C′ is proved. ■

Proof of Theorem C. Suppose f is a hyperbolic, critically finite map with carpet Julia set and post-critical set P_f.

By Moore's theorem, since the Julia set is a Sierpiński carpet, the quotient space obtained by collapsing the closures of Fatou components of f to points is a sphere; we denote the resulting projection by $\rho\colon \widehat{\mathbb{C}} \to S^2$. Then ρ gives a semiconjugacy to an induced map $g\colon S^2 \to S^2$. It is shown in [22, Theorem 5.1] that g is an expanding Thurston map.

The quotient map ρ is uniformly approximable by a continuous family of homeomorphisms. Taking any member of this family yields a homeomorphism $h\colon (\widehat{\mathbb{C}}, P_f) \to (S^2, P_g)$, well defined up to isotopy relative to P_f. The map h induces a conjugacy up to isotopy from f to g, so, as in Remark 4.15, $\overline{N}(g) = \overline{N}(f)$. By Theorem C′, we have $\overline{N}(f) = \overline{N}(g) > 1$. ■

Remark 7.3. The major difference between the quantities $\overline{N}(f)$ and $\Lambda_0(f)$ introduced in the proof of Theorem C is that the former represents a maximum growth rate, while the latter represents a minimum growth rate.

The converse to Theorem C need not hold; there are many examples. We sketch two constructions. Begin with the quadratic carpet example of Milnor and Tan [36, Appendix F],

$$f(z) \approx -0.138115091\left(z + \frac{1}{z}\right) - 0.303108805.$$

The two critical points have periods 3 and 4. The thesis of the first author [42, Theorem 7.1] shows that there exists a rational map g which combinatorially is the “tuning” of f and the basilica polynomial along the period 3 critical orbit. The result of tuning replaces each component of the basin of the superattracting 3-cycle with a copy of the basilica Julia set. The Julia set of g is easily seen to have two Fatou components whose closures meet, corresponding to the immediate attracting basins of the basilica, so it is not a carpet. Insung Park [40, Theorem 2] has proved that the energies $\overline{E}^p$ do not decrease under tuning, so in particular $\overline{N}(g) \geq \overline{N}(f) > 1$.

One may easily construct other examples that have not just local cut points, as in the preceding case, but global cut points. The quartic rational map

$$f(z) \approx -\frac{z^3(z+1)}{(z+0.3309124475)^3(z+0.0072626575)}, \tag{7.4}$$

with Julia set shown in Figure 10, has critical points at $p, c, 0, \infty$ with orbits

$$p \xrightarrow{3} \infty \xrightarrow{2} -1 \xrightarrow{1} 0 \overset{3}{\circlearrowleft},$$
$$c \overset{2}{\circlearrowleft}$$

where $-1 < p < c < 0$; the weights on the arrows show the local degree. This example is obtained from taking the torus automorphism $z \mapsto (1+i)z$ on the torus $\mathbb{C}/(\mathbb{Z} \oplus i\mathbb{Z})$ by

- taking its $\mathbb{Z}/4\mathbb{Z}$-quotient to get the Lattés map $g(z) = -z(z+1)/(z+1/2)^2$, with critical orbits $-1/2 \xrightarrow{2} \infty \xrightarrow{2} -1 \to 0 \circlearrowleft$;
- blowing up an arc from the unique fixed post-critical point at 0 to the critical point at $-1/2$ to get a degree-three map h with carpet Julia set; and then
- blowing up an arc from the unique fixed superattracting post-critical point of h to the repelling fixed point of h on the boundary of its basin to get the degree-four map f above.

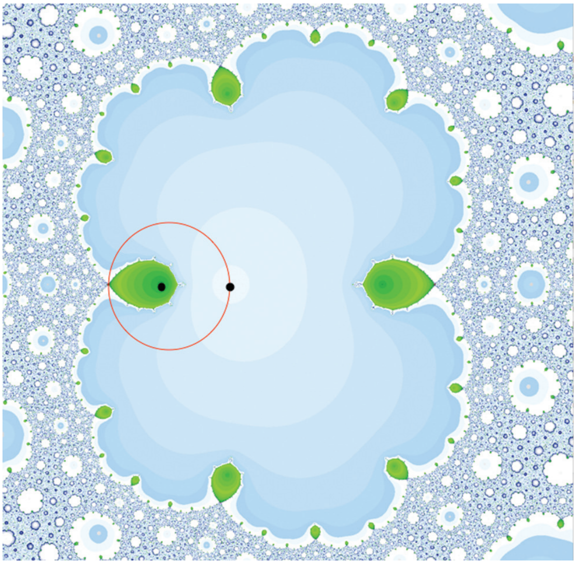

Figure 10. Details of the Julia set for the non-carpet example f in equation (7.4). The post-critical points c and 0 are marked in black, and the curve α is shown in red.

Direct calculation shows that the following condition is satisfied, as shown in Figure 10: there is a curve $\alpha\colon [0,1] \to \widehat{\mathbb{C}}$, with $\alpha(0) = \alpha(1) = 0$, and $\alpha(t) \notin P_f$ for $0 < t < 1$; the image of α is an embedded loop symmetric with respect to the real axis; the bounded component of the complement of the image of α contains c and no other points of P_f; some lift $\tilde{\alpha}$ of α under f is homotopic to α through curves with the same properties. From [42, Theorem 5.14], it follows that the boundary of the immediate basin of the origin is not a Jordan curve, and hence that J_f is not a Sierpiński carpet.

With a bit of work, one can show that $\overline{N}(f) = \sqrt{2} > 1$.

7.2. Barycentric subdivision Julia set

We give details for the estimates for the conformal dimension of the Julia set of the map $f(z) = \frac{4}{27}\frac{(z^2-z+1)^3}{(z(z-1))^2}$ from the introduction, shown combinatorially in Figure 11.

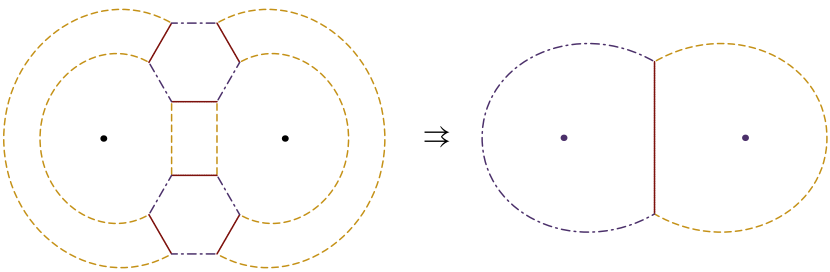

Figure 11. Spines for the barycentric subdivision rational map.

The dynamics on the set of critical points is shown below, with $\omega = \exp(2\pi i/6)$:

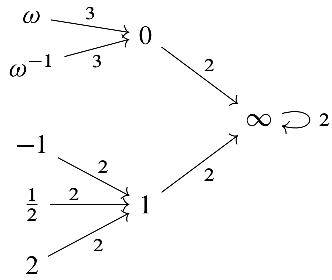

The map f may be described as follows. Consider the spherical "triangle" T_0 (the upper half-plane) with vertices $(\infty, 0, 1)$ and angles (π, π, π). If we give $\widehat{\mathbb{C}}$ the spherical metric to make this triangle equilateral, then the spherical triangle T_1 with vertices $(\infty, \omega, 2)$ has corresponding angles $(\pi/2, \pi/3, \pi/2)$. There is a unique conformal map sending T_1 to T_0 and mapping the vertices $(\infty, \omega, 2) \mapsto (\infty, 0, 1)$. Twelve reflected images of T_1 tile the sphere, and the map f is the unique extension given by Schwarz reflection. The extended real axis is forward invariant, and its preimage under f divides the sphere into 12 small triangles, 6 in each of the upper and lower half-planes. This induces a *finite subdivision rule* on the sphere, in the sense of [10]. We equip the codomain with a cell structure with 0-cells at $0, 1, \infty$ and 1-cells the corresponding segments of the extended real axis. Taking inverse images then refines each 2-cell in a pattern that, combinatorially, effects *barycentric subdivision*; see Figure 2. The function f is a Galois branched covering map, with deck group isomorphic to S_3 and acting by spherical isometries. It gives the j-invariant of an elliptic curve as a function of its λ-invariant; equivalently, it gives the shape invariant of a set of 4 points on $\widehat{\mathbb{C}}$ as a function of the cross-ratio of a corresponding list.

This map was studied by Cannon–Floyd–Parry [10, Example 1.3.1], where they showed that the sequence of tilings $\mathcal{T}_n$ generated by iterated preimages is not conformal in Cannon's sense: there is no metric on the sphere quasi-conformally equivalent to the standard metric in which combinatorial moduli of curve families are comparable to analytic moduli. This is related to the fact that there are fixed critical points, which means the valence of $\mathcal{T}_n$ blows up as $n \to \infty$. Cannon–Floyd–Kenyon–Parry investigated it further [9, Figure 25] and proved that its Julia set is a Sierpiński carpet. Haïssinsky and the second author studied it as well [24, Section 4.6].

To give an upper bound for the conformal dimension for the Julia set of this map, we estimate $\overline{E}^2(f)$, which we know is less than 1 [52, Theorem 1]. To get a concrete estimate, we look at a finite stage and compute $E_2^2[\phi^n]$ for some n. It turns out that $n = 1$ is not enough to get an estimate less than 1, so we consider $\phi^2\colon G_2 \to G_0$, as shown in Figure 12.

Fix $q > 1$. We can take any q-conformal structure G_0^q we like on G_0. It is most convenient to have $\alpha(e) = 1$ for all three edges of G_0. We get a pulled-back structure G_2^q on G_2. In this case, by symmetry, an optimal map will map the complete round dashed central

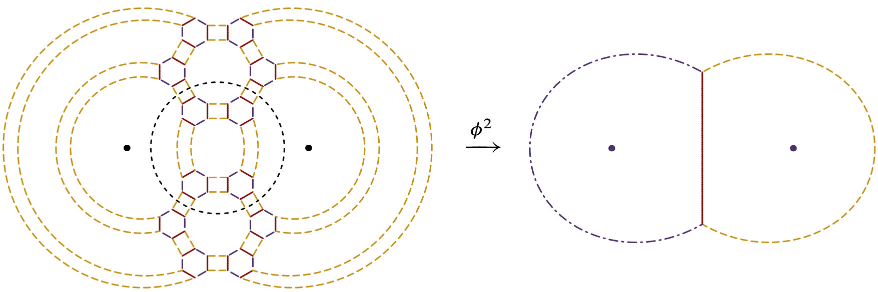

Figure 12. The map $\phi^2\colon G_2 \to G_0$ for the barycentric subdivision map.

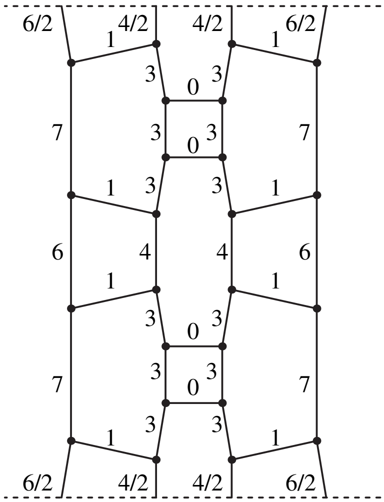

Figure 13. The central portion of the concrete map realizing $E_2^2[\phi^2]$ for the barycentric subdivision map. The edges running off the top and bottom are half-edges in the domain of the map. The numbers are the relative lengths of the images of each edge (or half-edge); to get the actual length, divide by 26. The map itself is projection onto the vertical axis.

circle on the left of Figure 12 to the central edge on the right. (This circle passes through the midpoint of eight edges; the other halves of these edges of G_2 map to different edges of G_0.)

Finding the map that minimizes E_2^2 from a graph to an interval (with specified boundary behavior) is equivalent to finding the resistance of a resistor network (or alternatively a harmonic function on the graph) and can be solved with standard linear algebra techniques. This is easily computed to be $E_2^2[\phi^2] = 10/13$, as shown in Figure 13. (The computation can be further simplified by using the symmetries evident in the figure.) Since the energy of any iterate gives an upper bound for the asymptotic energy [52, Proposition 5.6], we therefore have $\overline{E}^2(f) \le \sqrt{10/13}$.

Proposition 7.5. *For the barycentric carpet map f, we have*

$$\mathrm{ARCdim}(J_f) \le q^* \le \frac{2}{1-\log_6(10/13)} < 1.745.$$

Proof. By definition, we have $\overline{E}^{q^*}(f) = 1$. Apply Proposition 7.1 with $q = 2$ to find

$$\sqrt{\frac{10}{13}} \ge \overline{E}^2(f) \ge 6^{1/2-1/q^*}$$

which simplifies to the desired inequality after taking logs of both sides. ■

Iterating further to better estimate $\overline{E}^2(f)$ will improve the bound in Proposition 7.5, although it probably will not reach the optimal result, since Proposition 7.1 is not, in general, sharp.

To get a lower bound on conformal dimension, we compute $\overline{N}(f)$.

Proposition 7.6. *For the barycentric subdivision rational map, $\overline{N}(f) = 2$.*

An application of Theorem B then yields $\mathrm{ARCdim}(J_f) \ge 1/(1-\log_6(2))$, as claimed at the end of Section 1.1.

Proof. By examining Figure 11, it is easy to find a map in $[\phi]$ for which the inverse image of every generic point is two points. Thus $\overline{N}(f) \le N[\phi] \le 2$.

To get the opposite inequality, we will find 2^n edge-disjoint curves γ_i^n on G_n so that the curves $\phi^n(\gamma_i^n)$ are all homotopic to a simple loop γ_0 on G_0; cf. the proof of Theorem C′. This will immediately imply that $N[\phi^n] \ge 2^n$, as desired.

We find the γ_i^n by inductively finding 2^n edge-disjoint paths connecting any two edges of the dual of the nth barycentric subdivision of a triangle. This is trivial for $n = 0$, and $n = 1$ is shown in Figure 14. This also serves as the inductive step: at level n, in each triangle replace the concrete paths in Figure 14 with the family of 2^{n-1} parallel paths constructed by induction.

The closed curves γ_i^n are obtained by doubling the triangle as usual. ■

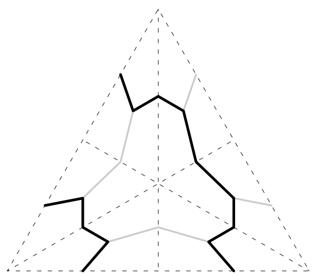

Figure 14. Connecting sides of a triangle in the barycentric subdivision with edge-disjoint paths. This is the case $n = 1$ and the inductive step.

7.3. Fat real Devaney examples

We begin by recalling some specifics concerning the Devaney family $f_\lambda(z) = z^2 + \lambda/z^2$. The points 0 and ∞ are always critical points satisfying the critical orbit relation

$$0 \xrightarrow{2} \infty \overset{2}{\circlearrowleft}.$$

The other critical points have orbits that start

$$\lambda^{1/4} \xrightarrow{2} \pm 2\sqrt{\lambda} \longrightarrow 4\lambda + \frac{1}{4} =: x_\lambda \longrightarrow \cdots. \tag{7.7}$$

A sufficient condition for f_λ to be hyperbolic and have a Sierpiński carpet Julia set is that x_λ eventually iterates into the Fatou component containing the origin. Such parameter values form a countable collection of open disks called *Sierpiński holes*.

As shown by the first author and R. Devaney [17], given any finite word $w = \varepsilon_0\varepsilon_1\cdots\varepsilon_n$ in the alphabet $\{L, R\}$, there exists a unique $\lambda = \lambda_w$, necessarily real and negative, with the following property. For each $i = 0, 1, \ldots, n$, the image $f_\lambda^i(x_\lambda)$ lies to the left (L) or right (R) of the origin, according to the symbol in the ith position of w, and $f_\lambda^{n+1}(x_\lambda) = 0$. Let $\lambda_n^{\mathrm{skinny}}$ and λ_n^{fat} denote the parameter values corresponding to the words LR^n and R^n, respectively.

In this section, we prove the second half of Theorem E, dealing with the fat Devaney family. A virtual endomorphism spine for this rational map is shown in Figure 15 for $n = 3$. As usual, the covering map π preserves the decorations on the edges, and the map ϕ is the deformation retract onto G_0 as a spine for $S^2 - P$, where P consists of the indicated points in the diagram plus an extra point at ∞.

The critical points at the fourth roots of λ are shown schematically on the diagram of G_1 with crosses; they map to the upper and lower critical values on the diagram of G_0, which in turn maps to the sequence of points on the right of G_0, eventually ending at the central critical point at 0, which maps to ∞.

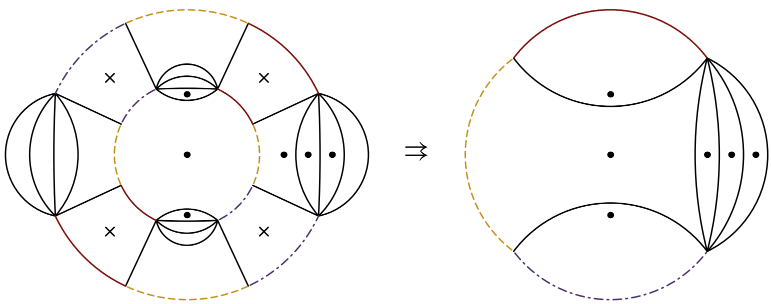

Figure 15. Virtual endomorphism $G_1 \rightrightarrows G_0$ for the fat Devaney family, R^n, shown here with $n = 3$. The post-critical set P is shown with bullets. The crosses on the left are the pre-periodic critical points.

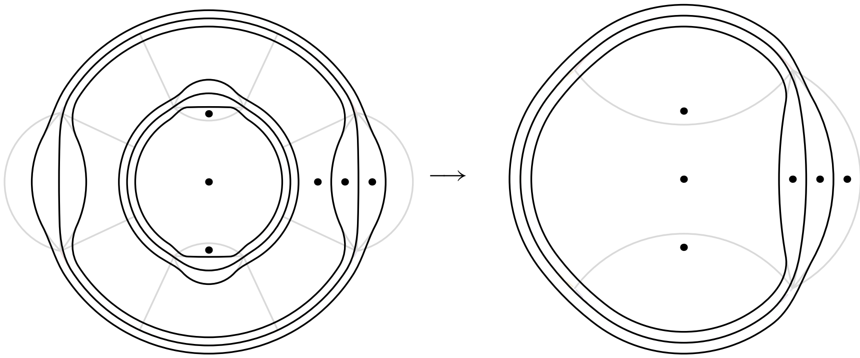

Figure 16. Curves in $S^2 - P$ for the fat Devaney family, and their inverse images. Each inverse image covers the original by degree 2.

To bound the conformal dimension from below, we again find $\overline{N}(f_\lambda)$ and use bounds on how quickly $\overline{E}^q$ decreases as a function of q. To find $\overline{N}$, consider the n-component multi-curve C shown on the right of Figure 16. The curve $f^{-1}(C)$ is shown on the left of Figure 16. (As indicated, it is easy to find $f^{-1}(C)$ by using the fact that G_1 is a cover of G_0.) Each component of $f^{-1}(C)$ covers one of the components of C with degree 2. One of the components of $f^{-1}(C)$ is peripheral (the outer one), and the others are all components of C. Note that we have the following:

- C is *completely invariant* (in the sense of Selinger [46]): $f^{-1}(C) = C$, up to homotopy in $S^2 - P_f$ and dropping inessential or peripheral components.
- C is *Cantor type* (in the sense of Cui–Peng–Tan [14]): for some iterate, each component of C has at least two preimages homotopic to itself.

Lemma 7.8. *For the Devaney family at parameter λ_n^{fat}, let r_n be the largest root of $\lambda^{n+1} - 2\lambda^n + 1$. Then $\overline{N}(f_{\lambda_n^{\mathrm{fat}}}) \geq r_n$.*

Proof. Let $\{\gamma_j\}$ enumerate the components of C. Consider the n-by-n non-negative matrix A whose entry A_{ij} records how many components of $f^{-1}(\gamma_j)$ are homotopic to γ_i (without accounting for the degree of the cover). Concretely, we have

$$A = \begin{pmatrix} 1 & 1 & \cdots & 1 & 1 \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{pmatrix}.$$

Let w be Perron–Frobenius eigenvector of A, with eigenvalue λ; concretely, $w_i = \lambda^{n-i-1}$ and λ is the positive root of

$$\lambda^n = \lambda^{n-1} + \lambda^{n-2} + \cdots + \lambda + 1.$$

Multiplying by $\lambda - 1$ shows that λ is the root r_n given in the lemma statement. Informally, λ is the rate of growth of the number of preimages of any of the curves $[\gamma_i]$.

To see that λ is a lower bound for $\overline{E}^1(f)$, we specialize the discussion in [52, Section 7.5] to the present setting. Let $W_0^1 = \sum w_i[\gamma_i]$ be the weighted multi-curve on G_0 corresponding to w. (Recall that the superscript "1" means we are thinking of this as a 1-conformal graph.) Set $W_1^1 = f^{-1}(W_0^1) = \pi^{-1}(W_0^1)$ be the pullback weighted curve, as a curve on G_1. (Here the pullback is as weighted curves, i.e., not taking into account the degree of the cover as in W. Thurston's obstruction theorem.) By the eigenvalue property, $\phi(W_1^1) = \lambda W_0^1$. We similarly have pullback embeddings $W_n^1 = (\pi^n)^{-1}(W_0^1)$ on G_n, so $\phi^n(W_n^1) = \lambda^n W_0^1$.

Let G_0^1 be G_0 as a weighted graph, with all weights equal to 1. Using the interpretation of E_1^1 via stretch factors—that is, a supremum of ratios of extremal lengths—as in Section 4.2, we have

$$E_1^1[\phi^n] \geq \frac{E_1^1[\phi^n(W_n^1) \to G_0^1]}{E_1^1[W_n^1 \to G_n^1]} = \frac{E_1^1[\lambda^n W_0^1 \to G_0^1]}{E_1^1[W_n^1 \to G_n^1]} = \lambda^n,$$

as desired. ■

Remark 7.9. It is straightforward to find weights on the edges of G_0 to show that the inequality in Lemma 7.8 is an equality, but the lower bound is all we need for our result.

Lemma 7.10. *There exists a unique root r_n of $g_n(\lambda) = \lambda^{n+1} - 2\lambda^n + 1$ in the interval $(1,2)$ and, for $n \geq 1$,*

$$r_n \geq 2 - 2^{-n+1}.$$

Proof. We have $g_n(2) = 1 > 0$. For $n = 1, 2, 3, 4$, we can directly check that $g_n(2 - 2^{-n+1}) \leq 0$. For $n \geq 5$,

$$\begin{aligned} g_n(2 - 2^{-n+1}) &= 2^{n+1}\left(1 - \frac{1}{2^n}\right)^{n+1} - 2^{n+1}\left(1 - \frac{1}{2^n}\right)^n + 1 \\ &\leq 2^{n+1}\left(1 - \frac{n+1}{2^n} + \frac{n(n+1)}{2^{2n+1}}\right) - 2^{n+1}\left(1 - \frac{n}{2^n}\right) + 1 \\ &= -1 + \frac{n(n+1)}{2^n} < 0. \end{aligned}$$

It follows that there is a root of g_n in $[2 - 2^{-n+1}, 2)$. Descartes' rule of signs shows there is a unique root in $(1,2)$. ■

Proof of Theorem E, fat family. By Theorem B, Lemma 7.8, Lemma 7.10, and elementary estimates,

$$\mathrm{ARCdim}(f_{\lambda_n^{\mathrm{fat}}}) \geq \frac{1}{1 - \log_4(\overline{N}(f_{\lambda_n^{\mathrm{fat}}}))} \geq \frac{2}{2 - \log_2 r_n} \geq \frac{2}{1 + 2^{-n}}.$$ ■

We sketch two alternate proofs of the same lower bound on $\mathrm{ARCdim}(J(\lambda_n^{\mathrm{fat}}))$, using the same multi-curve. Fix n; we write $f = f_{\lambda_n^{\mathrm{fat}}}$, and set $J = J(f)$. We equip J with a visual metric, as in Section 3.2; it belongs to the quasi-symmetric gauge of f, by Proposition 3.3.

Our first method is a variant of that employed in [23]. It associates to C a critical exponent for the combinatorial modulus of the family of curves homotopic to C. This exponent is then a lower bound for ARCdim. Abusing notation, let $[C]$ denote the family of curves in J which are homotopic in $\widehat{\mathbb{C}} - P_f$ to some component γ_j of the indicated multi-curve $C = \{\gamma_j\}$. Since each curve in the family is essential and non-peripheral, the diameters of elements of this family are bounded below, say by $\delta > 0$, and so $[C] \subset \Gamma_\delta$. Let $\mathcal{U}_m$, $m \in \mathbb{N}$, denote the sequence of coverings as in Section 6; it is a collection of snapshots, by Proposition 6.3. For $Q > 1$, let A_Q denote the matrix whose (i, j) entry is the sum of terms of the form d_{ijk}^{1-Q} where $d_{ijk} = \deg(\delta_k \to \gamma_j)$ and the curves δ_k range over preimages of γ_j homotopic to γ_i in the complement of the post-critical set. In our case, these covering degrees are all equal to 2, so $A_Q = 2^{1-Q} A$. Thus the Perron–Frobenius eigenvalue λ_Q of A_Q is $2^{1-Q} r_n$, where r_n is the root from Lemma 7.10; note this is strictly decreasing in Q. Setting this equal to 1 and solving for Q yields $Q_* = 1 + \log_2 r_n$ for the critical exponent. Fix now $1 < Q < Q_*$. Then as $m \to \infty$,

$$\mathrm{mod}_Q(\Gamma_\delta, \mathcal{U}_m) \geq \mathrm{mod}_Q([C], \mathcal{U}_m) \gtrsim 1,$$

where the last inequality is the statement of [23, Proposition 5.1]. By Corollary 5.12, we have $\mathrm{ARCdim}(J) > Q$ and hence $\mathrm{ARCdim}(J) \geq Q_*$.

Our second alternate proof we present here as a sketch; the motivation comes from [26]. Associated to the multi-curve C is a holomorphic virtual endomorphism of spaces $\pi_Y, \phi_Y : Y_1 \rightrightarrows Y_0$, where Y_0 is a collection of open Euclidean right annuli of circumference 1 (and geodesic boundary) indexed by the components of C and Y_1 is a collection of pairwise disjoint right Euclidean sub-annuli of Y_0 indexed by the components of $f^{-1}(C)$ homotopic to C. We require that ϕ_Y induces an inclusion $\overline{Y_1} \hookrightarrow Y_0$ conformal on the interior and that $\pi_Y : Y_1 \to Y_0$ is conformal, a local expanding homothety in the Euclidean coordinates with constant factor 2, and with each component mapping by degree 2.

Associated to this conformal expanding dynamical system is a non-escaping set $X \subset Y_0$ and a self-map $g: X \to X$. The set X is isometric to a product $S^1 \times \mathcal{C}$, where S^1 is the Euclidean circle of circumference 1, and $\mathcal{C}$ is a Cantor set associated to a graph-directed iterated function system on a disjoint union of $\#C$ copies of S^1, with contraction maps having factor 2, and where the copies map according to the combinatorics of the map $f^{-1}(C) \to C$. The Hausdorff dimension of $\mathcal{C}$ is equal to $\log_2 r_n$, by a variant of the well-known "pressure formula". It follows from [35, Proposition 4.1.11] that the conformal dimension of X is then equal to $1 + \log_2 r_n = Q_*$. There is a natural, non-surjective semiconjugacy $X \to J$ from g to f.

If this semiconjugacy were a homeomorphism, monotonicity of conformal dimension would imply the desired lower bound on $\mathrm{ARCdim}(J)$; however, it is not. Cui, Peng, and

Tan [14] show that there is a "thick" subset of X—the components living over "buried" points in the Cantor set $\mathcal{C}$—on which the semiconjugacy is injective; this should imply the desired lower bound. For example, passing to some high iterate and deleting the extreme inner and outermost branches of the interval contraction mappings defining the corresponding Cantor set, one obtains a sub-system whose repellor maps injectively to J and whose dimension is close to that of the original system X.

Remark 7.11. It is challenging, using our techniques, to give a concrete upper bound estimate on ARCdim that is less than 2. Although we know that $\overline{E}^2(f) < 2$, at some iterate the actual energy must be less than 2, and at that iterate, we could apply Proposition 7.1 to get an upper bound on ARCdim. But it appears that we have to iterate quite a lot to get to these values and get a good upper bound on ARCdim. In some sense, since the Julia set in these examples is a Sierpiński carpet of Hausdorff dimension close to 2, it is not well approximated by graphs.

7.4. Skinny Devaney examples

We now turn to the other half of Theorem E, dealing with the skinny Devaney family. Suitable spines in this case are shown in Figure 17 for $n = 3$ (kneading sequence LR^3); again, the generalization is evident.

We will find an explicit $q \in (1, 2)$, a q-conformal structure G_0^q given by a metric α on G_0, and map $\phi: G_1^q \to G_0^q$ so that $E_q^q(\phi) = 1$. The metric α is shown on the right of Figure 17, except that we have not yet determined the value x. The map ϕ is indicated schematically on the left of Figure 17: each region surrounded by a green loop is contracted to a point, and ϕ is optimized in the remaining regions.

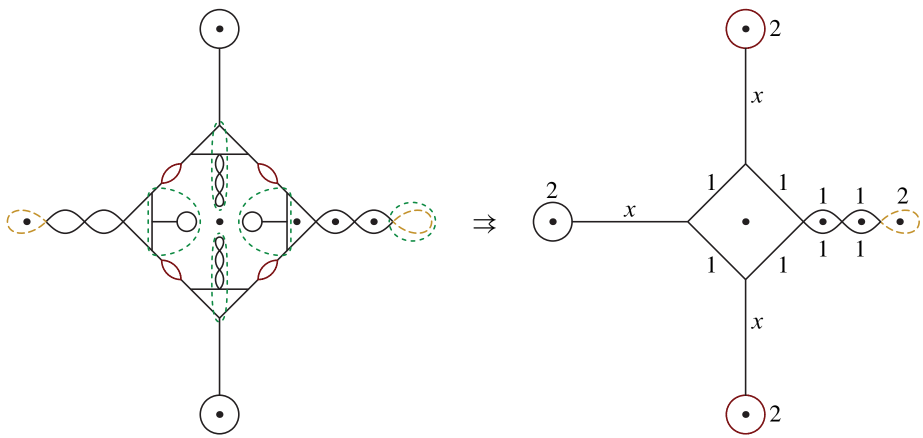

Figure 17. Virtual endomorphism $G_1 \rightrightarrows G_0$ for the skinny Devaney family, LR^n, is shown here with $n = 3$. The map π preserves color and orientation of the plane, as usual. The map ϕ collapses the regions within green dotted circles to points.

Inspection reveals that for most points $y \in G_0^q$, there is a unique point $x \in G_1$ with $\phi(x) = y$ and $|\phi'(x)| = 1$, compatible with $E_q^q(\phi) = 1$. The exceptions are as follows:

(1) On the long left edge, there is a chain in G_1 of n "double edges" (two edges connecting the same vertices), with each edge of length 1:

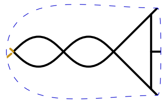

(Note the right vertical edge gets mapped to a point.) This chain maps to an edge in G_0^q of length x.

(2) On each of the four edges of the central square, there is a chain in G_1^q of an edge of length x, two parallel edges of length 2, another edge of length x, and two parallel edges of length 1:

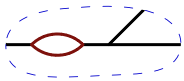

(Note the right-hand endpoints map to the same point.) These chains each map to an edge in G_0 of length 1.

We now solve for x and q to make the supremand in the definition of $E_q^q(\phi)$ equal to 1 on these edges as well. We use the principle that two parallel edges of length a and b in a q-conformal graph can be replaced by a single edge of length

$$a \oplus_q b := (a^{1-q} + b^{1-q})^{1/(1-q)}$$

without changing the optimal E_q^q energy of any map; see [52, Proposition 7.7]. (The case $q = 2$ is the standard parallel law for resistors.)

Thus the n double edges in (1) have an effective q-length of

$$x := n \cdot (1 \oplus_q 1) = n \cdot 2^{1/(1-q)}.$$

We set x to this value to make $E_q^q(\phi) = 1$.

The chain of edges in (2) has an effective q-length of

$$x + (2 \oplus_q 2) + x + (1 \oplus_q 1) = 2x + 3 \cdot 2^{1/(1-q)} = (2n + 3) \cdot 2^{1/(1-q)}.$$

For $q = 1 + 1/\log_2(2n + 3)$, this quantity is equal to 1, which makes $E_q^q(\phi) = 1$, with the supremum in the definition of E_q^q achieved everywhere on G_0.

Proposition 7.12. *For every skinny Devaney example $\lambda_n^{\text{skinny}}$, and for every $q > 1 + 1/\log_2(2n + 3)$, we can find a metric on G_0 in Figure 17 so that $E_q^q(\phi) < 1$. Thus $\overline{E}^q[\pi, \phi] < 1$ and*

$$\operatorname{ARCdim}(J_{\lambda_n^{\text{skinny}}}) \le 1 + \frac{1}{\log_2(2n + 3)}.$$

Proof. For q bigger than $1 + 1/\log_2(2n + 3)$, it is straightforward to modify the metric on G_0 in Figure 17 by adjusting the lengths slightly to make $E_q^q[\phi] < 1$. (For instance, on the chain of loops on the right side, multiply the length around the kth dot in from the end by $(1 + \varepsilon)^k$ for small ε.) Then for these q, we have $\overline{E}^q[\pi, \phi] \leq E_q^q[\phi] < 1$ and $\mathrm{ARCdim} < q$, as desired. ∎

Remark 7.13. One can also give a lower bound on $\mathrm{ARCdim}(J_{\lambda_n^{\text{skinny}}})$ by estimating $\overline{N}[\pi, \phi]$.

7.5. Uniformly well-spread cut points

Definition 7.14. A metric space X is *linearly connected* if there exists a constant $L \geq 1$ such that for each pair of points $x, y \in X$, there is a continuum E containing $\{x, y\}$ such that $\operatorname{diam} E \leq L d(x, y)$.

If X is connected and $f: X \to X$ is metrically cxc, then X is linearly connected [24].

Definition 7.15. A compact, connected metric space X is said to satisfy the UWSCP condition if there exists a constant $C \geq 1$ such that for each $x \in X$ and each $r > 0$, there exists a set $A \subset X$ with $\#A \leq C$ such that no component of $X - A$ meets both $B(x, r/2)$ and $X - \overline{B}(x, r)$.

The following result is [13, Theorem 1.2].

Theorem 7.16. *If X is doubling, compact, connected, linearly connected, and satisfies the UWSCP condition, then* $\mathrm{ARCdim}(X) = 1$.

7.6. Matings

In this section, we give an example showing that among Julia sets of hyperbolic rational functions, the UWSCP condition is sufficient but not necessary for $\mathrm{ARCdim}(J) = 1$. To our knowledge, this is the first result of its kind. We start by recalling generalities on matings.

Formal mating. Here, we denote by $\mathbb{S}^1 := \mathbb{R}/\mathbb{Z}$. For an integer $d \geq 2$, let $\tau_d : \mathbb{S}^1 \to \mathbb{S}^1$ be given by $\tau_d(\theta) = d \cdot \theta$ modulo 1. The action of a degree d monic polynomial on the complex plane may be compactified by extending it to the circle at infinity $\mathbb{S}^1_{f,\infty}$ to give an action on a closed topological 2-disk whose boundary values are given by τ.

The operation of *formal mating* takes as input two monic critically finite polynomials f_1, f_2 of the same degree $d \geq 2$ and returns as output a topological Thurston map $f_1 \amalg f_2 : S^2 \to S^2$. The sphere S^2 on which $f_1 \amalg f_2 : S^2 \to S^2$ acts is obtained by gluing together $\mathbb{S}^1_{f_1,\infty}$ to $\mathbb{S}^1_{f_2,\infty}$ via the map $\theta \mapsto -\theta$. Thus $f_1 \amalg f_2$ preserves a natural "equator circle" on which it acts via τ_d. We refer the reader to [8] for a survey containing facts mentioned below. Our focus here is exclusively on the case when the f_i are hyperbolic; this assumption simplifies the discussion. Abusing terminology, given a hyperbolic critically finite polynomial f, we call a bounded Fatou component a *basin* of f.

Ray equivalence relation and geometric mating. The sphere S^2 on which $f \amalg g$ acts comes equipped with a natural invariant "ray equivalence" relation, $\sim_{\mathrm{ray}}$. If $R\colon \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ is a rational map such that quotienting by the ray equivalence relation yields a continuous semiconjugacy $\rho\colon S^2 \to \widehat{\mathbb{C}}$ from $f \amalg g$ to R that is conformal on the basins of the f_i, we say f_1, f_2 are *geometrically mateable* and that R is the *geometric* mating of f_1 and f_2; see Figure 3.[3] The restriction $\rho\colon \mathbb{S}^1_\infty \to J(R)$ is surjective; its fibers are the intersections of the ray equivalence classes with $\mathbb{S}^1_\infty$. It is interesting to ask how big the fibers can be.

The following is a special case of a general principle, to our knowledge, first articulated by D. Fried [21]: that expanding dynamical systems on reasonable spaces are "finitely presented", in the following precise sense—they are quotients of an expanding subshift of finite type, and there is another subshift which encodes when two points lie in the same fiber. Restricted to the case of matings, we were surprised not to find the following result in the literature; cf. the survey [41, Remark 4.13] where the possibility of ray equivalence classes of unbounded size is entertained. Explicit bounds for the closely related question of the length of ray connections are established in certain cases by W. Jung [31].

Proposition 7.17. *Suppose f_1, f_2 are monic hyperbolic critically finite polynomials of degree $d \geq 2$ that are geometrically mateable, with geometric mating R. Then there exists a constant $C = C(f_1, f_2)$ such that fibers of $\rho\colon \mathbb{S}^1_\infty \to J(R)$ have cardinality at most C.*

For the proof, we apply a general result of Nekrashevych. To motivate the technique, and because we need it anyway, we begin by presenting a construction of a semiconjugacy $\pi\colon \Sigma_d \to \mathbb{S}^1_\infty$ from the full one-sided shift on d symbols to the map τ_d. We then analyze the composition

$$\Sigma_d \xrightarrow{\pi} \mathbb{S}^1_\infty \xrightarrow{\rho} J(R).$$

Proof. We take $x_0 := \infty e^{2\pi i 0}$ as a basepoint and denote by x_k for $k = 0, \ldots, d-1$ its d preimages under the map τ_d. There is a canonical positively oriented (counterclockwise) arc $\alpha_k \subset \mathbb{S}^1_\infty$ joining the basepoint x_0 to each x_k. Fix an iterate $n \geq 1$. Lifting the α_k under $\tau_d, \tau_d^2, \ldots, \tau_d^n$ and concatenating the lifts give an identification of $\tau_d^{-n}(x_0)$ with $\{0, 1, \ldots, d-1\}^n$ given by taking endpoints of iteratively concatenated lifts. The lengths of the concatenations are uniformly bounded, since the nth stage has length bounded by a convergent geometric series with ratio $1/d$, and passing to the limit we obtain the desired semiconjugacy $\pi\colon \Sigma_d \to \mathbb{S}^1_\infty$ from the shift to τ_d. (More simply, we can also see this semiconjugacy by writing an element in $\mathbb{R}/\mathbb{Z}$ in base d.)

[3] According to J. Milnor, geometric mating is "interesting, since it is neither well defined, injective, surjective nor continuous".

Our strategy is to now "push" this construction down to the dynamical plane of R via the semiconjugacy ρ. Denote by P the post-critical set of R. For each $k = 0, 1, \ldots, d-1$, let β_k be a smooth path in $\widehat{\mathbb{C}} - P$ homotopic relative to endpoints to the arc $\rho(\alpha_k)$. Note that since R is hyperbolic, it is expanding outside of a neighborhood of P. Applying the same iteratively lifting and concatenating construction from the previous paragraph using the β_k and R in place of the α_k and τ_d, we get a well-defined composition $\rho \circ \pi\colon \Sigma_d \to J(R)$ induced by taking endpoints of infinitely iterated concatenated lifts of the β_k.

Nekrashevych [37, Proposition 3.6.2] shows that in this setting, there is a finite automaton, called the *nucleus*, whose one-sided infinite paths encode the equivalence relation on Σ_d identifying points in the fiber of the composition $\rho \circ \pi$. It follows that the size of these equivalence classes is uniformly bounded by the size of the nucleus. The fibers of ρ are no larger than those of $\rho \circ \pi$, yielding the result. ■

See also [37, Section 6.13], in which details for a specific example of mating are presented.

Quadratic matings. Theorem 7.18 is due to Tan Lei [49] and M. Shishikura [47] and was proven using ideas of M. Rees.

Theorem 7.18. *Two critically finite quadratic polynomials f_1, f_2 are mateable if and only if they do not lie in conjugate limbs of the Mandelbrot set.*

To a hyperbolic critically finite quadratic polynomial f is associated a unique nontrivial interval $[a, b] \subset \mathbb{R}/\mathbb{Z}$: the external rays of angles a and b land at a common periodic point on the boundary of the immediate basin U containing the critical value of f and separate U from all other periodic attracting basins. For the basilica, $[a, b] = [1/3, 2/3]$, while for the airplane, $[a, b] = [3/7, 4/7]$. The denominators of a and b take the form $2^n - 1$, where n is the period of the finite critical point.

Generalized rabbits. If the hyperbolic component containing f has closure meeting the main cardioid component containing z^2, we call f a *generalized rabbit*. In this case, one may encode f by rational numbers in a different way.

For each $p/q \in \mathbb{Q}/\mathbb{Z} - \{0\}$, there is a unique hyperbolic critically finite quadratic polynomial $f = f_{p/q}$ such that there are q periodic basins meeting at a common repelling fixed point α, and such that the dynamics on the set of these q periodic basins, when equipped with the natural local cyclic ordering near α, is given by a rotation with angle p/q. This fact can be deduced from the classification of critically finite hyperbolic quadratic polynomials via their so-called invariant laminations; see [53]. The basilica polynomial is $f_{1/2}$, while the rabbit polynomial is $f_{1/3}$.

A special feature of generalized rabbit polynomials is the following. If f is a generalized rabbit polynomial and $\theta_1 \sim \theta_2$ is any nontrivial ray-equivalent pair of angles so that the corresponding rays land on a common point z in the Julia set of f, then the point z

is on the boundary of a basin of f. This fact need not hold for other pcf polynomials like the airplane.

Proposition 7.19. *For $i = 1, 2$, suppose $p_i/q_i \in \mathbb{Q}/\mathbb{Z} - \{0\}$ and suppose q_1, q_2 are coprime. Let f_i be the corresponding generalized rabbit quadratic polynomials and R the geometric mating of f_1 and f_2. Then the basins of f_1 and f_2 do not touch in the Julia set of R.*

Note that the mating exists by Theorem 7.18. In the proof below applied to the mating of the basilica and rabbit, the key observation is that the intervals $[1/3, 2/3]$ and $[1-2/7, 1-1/7]$ are disjoint.

Proof. We argue by contradiction. Let U_i be the immediate basin containing the finite critical value of f_i. Suppose $z \in J(R)$ is a point that is simultaneously on the boundary of a basin for f_1 and a basin for f_2. Since, by Proposition 7.17, the ray equivalence classes are finite and the f_i are generalized rabbits, this implies there exists an angle θ such that the ray of angle θ for f_1 lands on a point z_1 on the boundary of a basin of f_1, and the ray of angle $-\theta$ for f_2 lands on a point z_2 on the boundary of a basin of f_2. By iterating this ray pair forward under the formal mating, we may assume that $z \in \partial U_1$. Since q_1 and q_2 are coprime, by passing to a further iterate, we may assume that $z \in \partial U_1 \cap \partial U_2$.

In the dynamical plane of f_i, let (a_i, b_i) be the ray pair of angles landing at the point α_i separating U_i from the remaining immediate attracting basins of f_i, as recalled above. In giving coordinates for the circle at infinity $\mathbb{R}/\mathbb{Z}$, we parameterize to agree with the usual counterclockwise orientation on f_1 and disagree for f_2, so that the coordinates match when we mate. For example, if f_1 is the basilica and f_2 the rabbit, then $(a_1, b_1) = (1/3, 2/3)$ and $(a_2, b_2) = (1-2/7, 1-1/7)$. Then the set of angles of rays landing on ∂U_i is a subset of $[a_i, b_i]$. The condition that q_1 and q_2 are coprime implies that f_1 and f_2 are not in conjugate limbs of the Mandelbrot set, and hence that $[a_1, b_1]$ and $[a_2, b_2]$ are disjoint. But the previous paragraph shows θ belongs to both of these intervals. This is impossible. ∎

Proposition 1.2 is then an immediate consequence of the following two propositions.

Proposition 7.20. *Suppose f_1, f_2, and R are as in Proposition 7.19. Then the Julia set J_R of R does not satisfy UWSCP. In particular, the Julia set of the basilica mated with the rabbit does not satisfy UWSCP.*

Proof. Suppose J_R satisfies UWSCP, and C, x, r, and A be any data as in Definition 7.15. Then there exists a Jordan curve γ surrounding x meeting J_R only in the finite set A. Since the basins of the f_1 and f_2 do not touch, the loop γ lies in the closure of finitely many immediate basins of one of the f_i, say f_1.

Since J_R is the common boundary of the basins of f_1 and of f_2, each component of $\overline{\mathbb{C}} - \gamma$ contains a basin of f_2, say U and V. On the one hand, U and V are joined by a

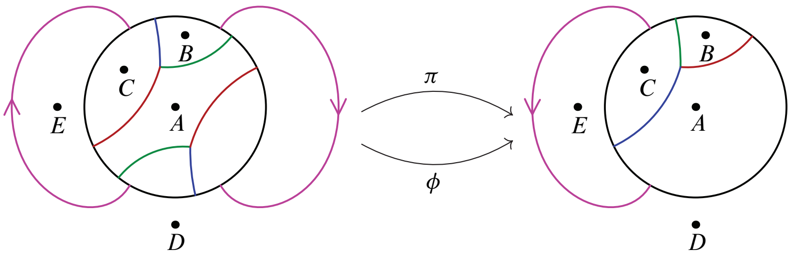

Figure 18. Virtual endomorphism $G_1 \rightrightarrows G_0$ for the mating of the rabbit and the basilica.

path β lying in the closure of finitely many basins of f_2 and meeting $J(R)$ in finitely many points, since this is true in the dynamical plane of f_2. On the other hand, any such path intersects γ and thus passes through a point of A. It follows that the basins of f_1 and f_2 touch in this point of A. By Proposition 7.19, this is impossible. ∎

Proposition 7.21. *The virtual endomorphism F_{RB} of the mating of the rabbit and the basilica satisfies $\overline{N}[F_{RB}] = 1$ and, for $q > 1$, $\overline{E}^q[F_{RB}] < 1$, and so* $\mathrm{ARCdim}(\mathcal{J}_{RB}) = 1$.

Proof. The spines for this mating are shown in Figure 18, with a black equator and colored rays. We may take the map ϕ in its homotopy class to be the map that "pushes" the extra colored edges toward the equator. If we do this suitably, each colored point in G_0 has one (colored) preimage, and each point on the equator has at most three preimages, one on the equator and two colored. It follows by iteration that $N(\phi^n) = 2n + 1$, and so $\overline{N}[F_{RB}] = 1$. For the statement about $\overline{E}^q$, we proceed as in [50, Example 2.4]. Fix $q > 1$, and consider a metric on G_0 where the colored edges have equal length L, and the equator has length 1. Lifting this metric under the covering to a metric on G_1, and lifting the colors as well, each colored edge of G_1 has length L, and the equator has length 2. For the map ϕ described above, $\mathrm{Fill}^q(\phi) = 1$ on the colored edges, and, for L sufficiently large, $\mathrm{Fill}^q(\phi) \approx 2^{1-q} < 1$ by equation (4.7). We can homotop ϕ to make Fill^q less than 1 everywhere by pulling the images of the vertices on the equator very slightly in along the colored edges; this decreases the derivative on the colored edge while increasing the derivative on the equator (but keeping it less than 1). Then $\overline{E}^q[F_{RB}] \le E[\phi] < 1$, as desired. ∎

Remark 7.22. The proof of Proposition 7.21 applies to the mating of any pair of rabbit-type polynomials as in the hypothesis of Proposition 7.19: all such matings have $\overline{N}[\pi, \phi] = 1$. However, if we mate two polynomials that are far out in the limbs, we may end up with a Sierpiński carpet. Figure 19 gives empirical evidence for this assertion. It shows the Julia set of the result of mating of the airplane polynomial f_1 and another polynomial f_2, with corresponding angles $(a_1, b_1) = (3/7, 4/7)$ and $(a_2, b_2) = (3/31, 4/31)$. The mating appears to be a Sierpiński carpet. (This example was found by Insung Park and Caroline Davis.)

Figure 19. The Julia set of the airplane polynomial mated with a Kokopelli-like polynomial. Picture by Arnaud Chéritat.

Conjecture 7.23. *The Julia set of a hyperbolic rational map satisfies UWSCP iff its virtual endomorphism has uniformly bounded* $N[\phi^n]$ *(independent of* n*).*

Acknowledgments. We thank Caroline Davis, InSung Park, and Giulio Tiozzo for useful conversations, and the referee for detailed comments on the first submission.

Funding. The first author was supported by the Simons Foundation under grant numbers 245269 and 615022. The second author was supported by the National Science Foundation under grant numbers DMS-1507244 and DMS-2110143.

References

[1] N. Aoki and K. Hiraide, *Topological theory of dynamical systems*. North-Holland Mathematical Library 52, North-Holland, Amsterdam, 1994 Zbl 0798.54047 MR 1289410

[2] M. Baker and L. De Marco, Special curves and postcritically finite polynomials. *Forum Math. Pi* **1** (2013), article no. e3, 35 pp. Zbl 1320.37022 MR 3141413

[3] M. Bonk, M. Lyubich, and S. Merenkov, Quasisymmetries of Sierpiński carpet Julia sets. *Adv. Math.* **301** (2016), 383–422 Zbl 1358.37083 MR 3539379

[4] M. Bonk and D. Meyer, *Expanding Thurston maps*. Math. Surveys Monogr. 225, American Mathematical Society, Providence, RI, 2017 Zbl 1430.37001 MR 3727134

[5] M. Bourdon and B. Kleiner, Combinatorial modulus, the combinatorial Loewner property, and Coxeter groups. *Groups Geom. Dyn.* **7** (2013), no. 1, 39–107 Zbl 1344.20059 MR 3019076

[6] M. Bourdon and B. Kleiner, Some applications of ℓ_p-cohomology to boundaries of Gromov hyperbolic spaces. *Groups Geom. Dyn.* **9** (2015), no. 2, 435–478 Zbl 1355.20037 MR 3356973

[7] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren Math. Wiss. 319, Springer, Berlin, 1999 Zbl 0988.53001 MR 1744486

[8] X. Buff, A. L. Epstein, S. Koch, D. Meyer, K. Pilgrim, M. Rees, and L. Tan, Questions about polynomial matings. *Ann. Fac. Sci. Toulouse Math. (6)* **21** (2012), no. S5, 1149–1176 Zbl 1364.37099 MR 3088270

[9] 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 Zbl 1077.37037 MR 1992038

[10] J. W. Cannon, W. J. Floyd, and W. R. Parry, Finite subdivision rules. *Conform. Geom. Dyn.* **5** (2001), 153–196 Zbl 1060.20037 MR 1875951

[11] J. W. Cannon and E. L. Swenson, Recognizing constant curvature discrete groups in dimension 3. *Trans. Amer. Math. Soc.* **350** (1998), no. 2, 809–849 Zbl 0910.20024 MR 1458317

[12] M. Carrasco Piaggio, On the conformal gauge of a compact metric space. *Ann. Sci. Éc. Norm. Supér. (4)* **46** (2013), no. 3, 495–548 Zbl 1283.30072 MR 3099984

[13] M. Carrasco Piaggio, Conformal dimension and canonical splittings of hyperbolic groups. *Geom. Funct. Anal.* **24** (2014), no. 3, 922–945 Zbl 1300.30100 MR 3213834

[14] G. Cui, W. Peng, and L. Tan, Renormalizations and wandering Jordan curves of rational maps. *Comm. Math. Phys.* **344** (2016), no. 1, 67–115 Zbl 1362.37095 MR 3493138

[15] T. Das, F. Przytycki, G. Tiozzo, M. Urbański, and A. Zdunik, Thermodynamic formalism for coarse expanding dynamical systems. *Comm. Math. Phys.* **384** (2021), no. 1, 165–199 Zbl 1467.37035 MR 4252875

[16] R. L. Devaney, D. M. Look, and D. Uminsky, The escape trichotomy for singularly perturbed rational maps. *Indiana Univ. Math. J.* **54** (2005), no. 6, 1621–1634 Zbl 1087.37043 MR 2189680

[17] R. L. Devaney and K. M. Pilgrim, Dynamic classification of escape time Sierpiński curve Julia sets. *Fund. Math.* **202** (2009), no. 2, 181–198 Zbl 1160.37362 MR 2506193

[18] A. Douady and J. H. Hubbard, A proof of Thurston's topological characterization of rational functions. *Acta Math.* **171** (1993), no. 2, 263–297 Zbl 0806.30027 MR 1251582

[19] C. Druţu and M. Kapovich, *Geometric group theory*. With an appendix by Bogdan Nica. Amer. Math. Soc. Colloq. Publ. 63, American Mathematical Society, Providence, RI, 2018 Zbl 1447.20001 MR 3753580

[20] S. Eriksson-Bique, Equality of different definitions of conformal dimension for quasiself-similar and CLP spaces. *Ann. Fenn. Math.* **49** (2024), no. 2, 405–436 Zbl 1543.30171 MR 4767001

[21] D. Fried, Finitely presented dynamical systems. *Ergodic Theory Dynam. Systems* **7** (1987), no. 4, 489–507 Zbl 0652.54028 MR 0922362

[22] Y. Gao, P. Haïssinsky, D. Meyer, and J. Zeng, Invariant Jordan curves of Sierpiński carpet rational maps. *Ergodic Theory Dynam. Systems* **38** (2018), no. 2, 583–600 Zbl 1441.37051 MR 3774834

[23] P. Haïssinsky and K. M. Pilgrim, Thurston obstructions and Ahlfors regular conformal dimension. *J. Math. Pures Appl. (9)* **90** (2008), no. 3, 229–241 Zbl 1213.30048 MR 2446078

[24] P. Haïssinsky and K. M. Pilgrim, Coarse expanding conformal dynamics. *Astérisque* **325** (2009), 139 pp. Zbl 1206.37002 MR 2662902

[25] P. Haïssinsky and K. M. Pilgrim, Finite type coarse expanding conformal dynamics. *Groups Geom. Dyn.* **5** (2011), no. 3, 603–661 Zbl 1298.37032 MR 2813529

[26] P. Haïssinsky and K. M. Pilgrim, Examples of coarse expanding conformal maps. *Discrete Contin. Dyn. Syst.* **32** (2012), no. 7, 2403–2416 Zbl 1267.37047 MR 2900552

[27] A. Hatcher, *Algebraic topology*. Cambridge University Press, Cambridge, 2002 Zbl 1044.55001 MR 1867354

[28] S.-M. Heinemann and B. O. Stratmann, Geometric exponents for hyperbolic Julia sets. *Illinois J. Math.* **45** (2001), no. 3, 775–785 Zbl 1055.37056 MR 1879234

[29] J. Heinonen, *Lectures on analysis on metric spaces*. Universitext, Springer, New York, 2001 Zbl 0985.46008 MR 1800917

[30] Y. Ishii and J. Smillie, Homotopy shadowing. *Amer. J. Math.* **132** (2010), no. 4, 987–1029 Zbl 1207.37013 MR 2663646

[31] W. Jung, Quadratic matings and ray connections. 2017, arXiv:1707.00630v1

[32] B. Kleiner, The asymptotic geometry of negatively curved spaces: uniformization, geometrization and rigidity. In *International congress of mathematicians*, pp. 743–768, Invited lectures II, European Mathematical Society, Zürich, 2006 Zbl 1108.30014 MR 2275621

[33] J. Kwapisz, Conformal dimension via p-resistance: Sierpiński carpet. *Ann. Acad. Sci. Fenn. Math.* **45** (2020), no. 1, 3–51 Zbl 1437.28013 MR 4056527

[34] J. M. Mackay, Spaces and groups with conformal dimension greater than one. *Duke Math. J.* **153** (2010), no. 2, 211–227 Zbl 1273.30056 MR 2667133

[35] J. M. Mackay and J. T. Tyson, *Conformal dimension: theory and application*. Univ. Lecture Ser. 54, American Mathematical Society, Providence, RI, 2010 Zbl 1201.30002 MR 2662522

[36] J. Milnor, Geometry and dynamics of quadratic rational maps (with an appendix by Milnor and Tan Lei). *Exp. Math.* **2** (1993), no. 1, 37–83 Zbl 0922.58062 MR 1246482

[37] V. Nekrashevych, *Self-similar groups*. Math. Surveys Monogr. 117, American Mathematical Society, Providence, RI, 2005 Zbl 1087.20032 MR 2162164

[38] V. Nekrashevych, Combinatorial models of expanding dynamical systems. *Ergodic Theory Dynam. Systems* **34** (2014), no. 3, 938–985 Zbl 1350.37034 MR 3199801

[39] M. Noonan, FractalStream. Program for exploring complex dynamical systems. 2008, http://pi.math.cornell.edu/~noonan/fstream.html visited on 14 June 2024

[40] I. Park, *Thurston and Levy obstructions of finite subdivision rules*. Ph.D. thesis, Indiana University, 2021

[41] C. L. Petersen and D. Meyer, On the notions of mating. *Ann. Fac. Sci. Toulouse Math. (6)* **21** (2012), no. 5, 839–876 Zbl 1360.37117 MR 3088260

[42] K. M. Pilgrim, *Cylinders for iterated rational maps*. Ph.D. thesis, University of California, Berkeley, 1994 MR 2691488

[43] F. Przytycki and M. Urbański, *Conformal fractals: ergodic theory methods*. London Math. Soc. Lecture Note Ser. 371, Cambridge University Press, Cambridge, 2010 Zbl 1202.37001 MR 2656475

[44] F. Przytycki, M. Urbański, and A. Zdunik, Harmonic, Gibbs and Hausdorff measures on repellers for holomorphic maps, I. *Ann. of Math. (2)* **130** (1989), no. 1, 1–40 Zbl 0703.58036 MR 1005606

[45] W. Rudin, *Real and complex analysis*. 3rd edn., McGraw-Hill, New York, 1987 Zbl 0925.00005 MR 0924157

[46] N. Selinger, Thurston's pullback map on the augmented Teichmüller space and applications. *Invent. Math.* **189** (2012), no. 1, 111–142 Zbl 1298.37033 MR 2929084

[47] M. Shishikura, On a theorem of M. Rees for the matings of polynomials. In L. Tan (ed.), *The Mandelbrot set, theme and variations*, pp. 289–306, London Math. Soc. Lecture Note Ser. 274, Cambridge University Press, Cambridge, 2000 Zbl 1062.37039 MR 1765095

[48] D. Sullivan, Conformal dynamical systems. In *Geometric dynamics (Rio de Janeiro, 1981)*, pp. 725–752, Lecture Notes in Math. 1007, Springer, Berlin, 1983 Zbl 0524.58024 MR 0730296

[49] L. Tan, Matings of quadratic polynomials. *Ergodic Theory Dynam. Systems* **12** (1992), no. 3, 589–620 Zbl 0756.58024 MR 1182664

[50] D. P. Thurston, From rubber bands to rational maps: a research report. *Res. Math. Sci.* **3** (2016), no. 1, article no. 15, 49 pp. Zbl 1360.37119 MR 3500499

[51] D. P. Thurston, Elastic graphs. *Forum Math. Sigma* **7** (2019), article no. e24, 84 pp. Zbl 1501.37041 MR 3993808

[52] D. P. Thurston, A positive characterization of rational maps. *Ann. of Math. (2)* **192** (2020), no. 1, 1–46 Zbl 1450.37042 MR 4125449

[53] W. P. Thurston, On the geometry and dynamics of iterated rational maps (with an appendix by Dierk Schleicher). In D. Schleicher and N. Selinger (eds.), *Complex dynamics. Families and friends*, pp. 3–137, A K Peters, Wellesley, MA, 2009 Zbl 1185.37111 MR 2508255

Received 20 August 2023.

Kevin M. Pilgrim
Indiana University, 831 E. Third St., Bloomington, IN 47405, USA; pilgrim@iu.edu

Dylan P. Thurston
Boston College, 140 Commonwealth Avenue, Chestnut Hill, MA 02467, USA; thurst@bc.edu