



HAL
open science

An algebraic characterization of expanding Thurston maps

Peter Haïssinsky, Kevin Pilgrim

► **To cite this version:**

Peter Haïssinsky, Kevin Pilgrim. An algebraic characterization of expanding Thurston maps. *Journal of modern dynamics*, 2012, 6 (no. 4), pp.451 – 476. 10.3934/jmd.2012.6.451 . hal-00687764

HAL Id: hal-00687764

<https://hal.science/hal-00687764>

Submitted on 14 Apr 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

An algebraic characterization of expanding Thurston maps

Peter Haïssinsky and Kevin M. Pilgrim

April 14, 2012

Abstract

Let $f : S^2 \rightarrow S^2$ be a postcritically finite branched covering map without periodic branch points. We give necessary and sufficient algebraic conditions for f to be homotopic, relative to its postcritical set, to an expanding map g .

Contents

1	Introduction	2
2	Thurston maps	7
2.1	Curves and arcs	7
2.2	Orbifolds	8
3	Algebraic constructions	9
3.1	The virtual endomorphism	9
3.2	The biset \mathfrak{M}	10
3.3	The selfsimilarity complex Σ	11
4	Dynamical consequences of the hyperbolicity of \mathfrak{M}	12
5	The covering complex Γ	16
6	Coincidence of quasi-isometry types	17
7	Proof of necessity	24
8	Proof of sufficiency	26

1 Introduction

Let $T^2 = \mathbb{R}^2/\mathbb{Z}^2$ denote the usual torus and, using the standard basis, identify T^2 with the quotient $H_1(T^2, \mathbb{R})/H_1(T^2, \mathbb{Z})$. Equip T^2 with the Euclidean metric. The following result is well-known:

Theorem 1.1 *Suppose $f : T^2 \rightarrow T^2$ is a self-covering map of the torus. Then f is homotopic through covering maps to an expanding map $g : T^2 \rightarrow T^2$ if and only if the spectrum of $f_* : H_1(T^2, \mathbb{R}) \rightarrow H_1(T^2, \mathbb{R})$ lies outside the closed unit disk.*

In this setting, one may take g to be the linear map on $T^2 = H_1(T^2, \mathbb{R})/H_1(T^2, \mathbb{Z})$ induced by f_* . One may further arrange so that the map g is a factor of f via a map $\pi : T^2 \rightarrow T^2$ (that is, $\pi \circ f = g \circ \pi$) where π is homotopic to the identity, is monotone, has fibers not disconnecting T^2 , and, if f is itself expanding, is a homeomorphism. The proof is now standard: one lifts the identity map $h_0 : T^2 \rightarrow T^2$ under f and g to obtain $h_1 : T^2 \rightarrow T^2$ and a homotopy $h_t, t \in [0, 1]$ from h_0 to h_1 . By induction, lifting and concatenating the homotopies, one obtains a family $h_t, t \in [0, \infty)$, such that $g \circ h_{t+1} = h_t \circ f$. Expansion of g implies the family h_t converges uniformly to the desired semi-conjugacy. Similar and much more general results hold, cf. [Shu, Theorem 3].

In this work, we formulate and prove an analog of this theorem for a certain class of continuous maps $f : S^2 \rightarrow S^2$ of the sphere to itself, called *Thurston maps*.

A Thurston map $f : S^2 \rightarrow S^2$ is an orientation-preserving branched covering of the two-sphere of degree $d \geq 2$ for which the necessarily nonempty *postcritical set* $P_f = P = \cup_{n>0} f^{on}(\{\text{branch points}\})$ is finite. These were originally introduced by Thurston [DH] as combinatorial objects with which to classify rational functions as dynamical systems on the Riemann sphere $\widehat{\mathbb{C}}$. Two Thurston maps f, g are *combinatorially equivalent* if there exist orientation-preserving homeomorphisms $h_0, h_1 : (S^2, P_f) \rightarrow (S^2, P_g)$ such that $h_0 \circ f = g \circ h_1$ and h_0, h_1 are homotopic relative to P_f . The relation of combinatorial equivalence can be thought of as conjugacy, up to isotopy relative to postcritical sets. Recently, Thurston maps have been investigated in connection with the study of finite subdivision rules, the classification of quasisymmetry classes of metrics on S^2 , and Cannon's conjecture; see [CFP], [BM], [HP1].

A formulation of a variant of Theorem 1.1 for Thurston maps requires the development of a notion of expansion for non-locally injective maps, algebraic/homotopy-theoretic invariants of a self-map of a simply-connected surface, and the synthetic construction of an expanding model map from the algebraic data.

Expansion. Let f be a Thurston map. Since f has branch points, it can never be expanding with respect to a Riemannian metric on S^2 , nor can it be positively expansive. We therefore reformulate the notion of expansion in terms of contraction of inverse branches. Let \mathcal{U}_0 be a finite cover of S^2 by connected open sets, and let \mathcal{U}_n denote the covering whose elements are connected components of $f^{-n}(U)$, $U \in \mathcal{U}_0$. We say that f is *expanding* if the following condition holds.

Definition 1.2 *The map f is expanding if the mesh of the coverings \mathcal{U}_n tends to zero as $n \rightarrow \infty$. That is, for any finite open cover \mathcal{Y} of S^2 by open sets, there exists N such that for all $n \geq N$ and all $U \in \mathcal{U}_n$, there exists $Y \in \mathcal{Y}$ with $U \subset Y$. Equivalently: the diameters of the elements of \mathcal{U}_n , with respect to some (any) metric, tend to zero as $n \rightarrow \infty$.*

Note that this is a purely topological condition. It is at present unknown whether every expanding Thurston map is topologically conjugate to a map which is smooth away from a finite set of points. Perhaps more surprisingly, it is unknown whether every expanding Thurston map is topologically conjugate to a map which preserves some complete length structure on S^2 .

Algebraic invariants. Nekrashevych [Nek1] introduced a collection of algebraic invariants associated to a very wide class of dynamical settings, including that of a Thurston map f . These may be unfamiliar to most readers; see §3 for details. As in [DH], associated to f is an orbifold structure \mathcal{O} on S^2 . In order to have needed finiteness properties, we assume f has no periodic branch points, so that \mathcal{O} is compact. Once a basepoint is fixed, there is a corresponding orbifold fundamental group $\pi_1(\mathcal{O})$, and a so-called *virtual endomorphism* $\phi : \pi_1(\mathcal{O}) \dashrightarrow \pi_1(\mathcal{O})$, well-defined up to postcomposition by inner automorphisms. Roughly speaking, ϕ is the homomorphism induced by the *inverse* of f . The property of ϕ being *contracting* plays the role of the spectral condition on the induced map f_* on homology in Theorem 1.1.

Synthetic model map. When the virtual endomorphism ϕ is contracting, by choosing suitable defining data (basis; generators) one may construct a negatively curved *self-similarity complex* Σ . The construction is similar to that of the boundary of a hyperbolic group from a Cayley graph. The boundary at infinity \mathcal{J} of Σ inherits a self-map $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$. The dynamical system $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$ plays the role of the induced linear map g in Theorem 1.1.

Our main result is

Theorem 1.3 (Characterization of expanding Thurston maps) *Suppose f is a Thurston map without periodic critical points. Then f is homotopic*

(through Thurston maps relative to its postcritical set) to an expanding map g if and only if the virtual endomorphism $\phi : \pi_1(\mathcal{O}) \dashrightarrow \pi_1(\mathcal{O})$ is contracting.

By applying an observation of Rees [Rees, § 1.4], we conclude that a Thurston map without periodic branch points is combinatorially equivalent to an expanding map if and only if its virtual endomorphism on the orbifold fundamental group is contracting.

The hypothesis of no periodic critical points cannot be dropped: the rational function $f(z) = z^2$ acting on the Riemann sphere has contracting virtual endomorphism, but is not homotopic to any expanding map: the nonnegative real axis is an invariant arc joining periodic points in its postcritical set, and by Proposition 2.1 this is an obstruction to expansion.

If the defining data for Σ and an open covering \mathcal{U}_0 of S^2 is chosen in a suitable way, one obtains a semiconjugacy $\pi : S^2 \rightarrow \mathcal{J}$ from $f : S^2 \rightarrow S^2$ to $f_{\Sigma} : \mathcal{J} \rightarrow \mathcal{J}$ playing the role of the map π introduced after Theorem 1.1. In the setting of Theorem 1.3, \mathcal{J} is a sphere, and g is a limit of conjugates of f via a path of conjugacies h_t (with now $t \in [0, 1]$) which gradually collapse the fibers of π to points. Formally, h_t is a *pseudo-isotopy*.

A *finite subdivision rule* \mathcal{R} on the sphere (in the sense of [CFP]) is, roughly, a Thurston map $f : S^2 \rightarrow S^2$ together with a cell structure $S_{\mathcal{R}}$ on S^2 and a refinement $\mathcal{R}(S_{\mathcal{R}})$ of $S_{\mathcal{R}}$ such that $f : \mathcal{R}(S_{\mathcal{R}}) \rightarrow S_{\mathcal{R}}$ is cellular and a homeomorphism on each cell. By taking preimages, one obtains a sequence of refinements $\mathcal{R}^n(S_{\mathcal{R}})$, $n \in \mathbb{N}$. One says that \mathcal{R} has *mesh going to zero* if the diameters of the cells at level n tend to zero as n tends to infinity. It has *mesh going to zero combinatorially* if for some n , (i) every edge of $S_{\mathcal{R}}$ is properly subdivided in $\mathcal{R}^n(S_{\mathcal{R}})$, and (ii) given a pair of disjoint edges e, e' of $S_{\mathcal{R}}$ in the boundary of a tile, no tile t of $\mathcal{R}^n(S_{\mathcal{R}})$ contains edges in both e and in e' . Thus mesh going to zero is clearly an invariant of topological conjugacy, while mesh going to zero combinatorially is a combinatorial property.

By appealing to recent work of Bonk and Meyer [BM], we obtain

Theorem 1.4 (Equivalence of combinatorial expansion) *Suppose f is a Thurston map without periodic branch points. The following conditions are equivalent.*

1. *The map f is equivalent to an expanding map, g .*
2. *There exists m such that f^m is equivalent to a map g which is the subdivision map of a subdivision rule \mathcal{R} with mesh going to zero.*
3. *There exists m such that f^m is equivalent to a map g which is the subdivision map of a subdivision rule \mathcal{R} with mesh going to zero combinatorially.*
4. *The virtual endomorphism ϕ_f is contracting.*

Proof: That (1) implies (2) follows from [BM, Theorem 1.2]; we remark that they do not require the absence of periodic branch points. That (2) implies (1) is a straightforward consequence of the definitions. The equivalence of (2) and (3) is [CFP, Theorem 3.2]. Now suppose (3) holds. By [CFPP, Theorem 1.1], the virtual endomorphism ϕ_g is contracting. This implies (4). That (4) implies (1) is Theorem 1.3. ■

In Theorem 1.4, the topological conjugacy class of the expanding map g is necessarily unique, by e.g. [BM, Theorem 10.4]. A corollary of Theorem 1.4 is thus that the dynamics of an expanding Thurston map is determined by either of two pieces of combinatorial data: a subdivision rule with mesh going to zero combinatorially, or by an expanding virtual endomorphism. If a Thurston map f is already presented in terms of a subdivision rule \mathcal{R} , there is a finite, easily checkable, necessary and sufficient condition for \mathcal{R} to have mesh going to zero combinatorially [CFPP, Theorem 6.1]. If f is not so presented, one can try to verify the contraction of ϕ_f instead. Bartholdi has pointed out that this can be done algorithmically, e.g. by combining a normal forms algorithm for writing elements of the orbifold fundamental group with the FR package [Bar] in the computational algebra program GAP.

In the paragraphs below, we outline in detail the reduction of the proof of Theorem 1.3 to a few key steps.

Necessity. The implications (1) \implies (2) \implies (3) \implies (4) have already been established in the above proof of Theorem 1.4. In §7, however, we provide a new, direct proof. The idea is that the group $\pi_1(\mathcal{O})$ acts on the universal cover $p : \tilde{\mathcal{O}} \rightarrow \mathcal{O}$ and that we may find a lift \tilde{f}_- of “ f^{-1} ” i.e. a map satisfying $f \circ p \circ \tilde{f}_- = p$. The action of the virtual endomorphism on $\pi_1(\mathcal{O})$ can be interpreted in terms of the action of \tilde{f}_- on $\tilde{\mathcal{O}}$. Assuming that f is expanding yields a metric on \mathcal{O} which can be lifted to $\tilde{\mathcal{O}}$ in such a way that (a) $\pi_1(\mathcal{O})$ acts geometrically on $\tilde{\mathcal{O}}$ implying that $\pi_1(\mathcal{O})$ is quasi-isometric to $\tilde{\mathcal{O}}$ and (b) the action of \tilde{f}_- is contracting. Bringing together these two facts leads to the contraction of the virtual endomorphism.

Sufficiency. Suppose the virtual endomorphism ϕ is contracting and f has no periodic branch points. Necessarily $\#P \geq 3$. For technical reasons, it is convenient later to assume $\#P > 3$. When $\#P = 3$, f is equivalent to a rational function which is chaotic on the whole sphere, by Thurston’s characterization [DH], and such a map is expanding.

We first construct the semiconjugacy $\pi : S^2 \rightarrow \mathcal{J}$ from $f : S^2 \rightarrow S^2$ to the model map $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$; recall that \mathcal{J} arises as the boundary of the hyperbolic selfsimilarity complex, Σ . Let \mathcal{U}_0 be an open cover of S^2 by Jordan

domains. By taking inverse images, there is another associated infinite one-complex Γ ; see [HP1] and §5. Using the absence of periodic branch points, we show (Theorem 6.1) that for suitable choices of the covering \mathcal{U}_0 defining Γ and the defining data for Σ , there is a natural quasi-isometry $\Phi : \Sigma \rightarrow \Gamma$. Thus Γ is also hyperbolic. Extending Φ to the boundary by a map of the same name, the composition $\pi := \Phi^{-1} \circ \pi_\Gamma$ is the desired semiconjugacy.

Next, we show that \mathcal{J} is homeomorphic to S^2 , and that the model map f_Σ is a Thurston map. A key role is played by the fact, established in [HP2], that $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$ is a branched covering that, as a topological dynamical system, is expanding and is of so-called *finite type*. In particular, the local degrees of iterates are uniformly bounded. We first develop (§4) some dynamical consequences of the hypothesis that ϕ is contracting, e.g. that there are no so-called Levy cycles. Next, after showing π is monotone (Lemma 5.1), it suffices, by a classical theorem of Moore [Dav, Theorem 25.1], to show that the fibers of π do not disconnect S^2 . To see this, we use a case-by-case analysis of the possibilities for how such fibers separate and intersect P and the absence of Levy cycles and other related obstructions to expansion. We conclude that \mathcal{J} is a sphere, that $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$ is an expanding Thurston map, and that $\pi|_P$ is injective.

It remains to construct an isotopy from f to an expanding map g . By [Dav, Theorems 13.4, 25.1], the decomposition \mathcal{G} of S^2 by the fibers of the semiconjugacy $\pi : S^2 \rightarrow \mathcal{J}$ has the property of being *strongly shrinkable*: there is a one-parameter family of continuous maps $h_t : S^2 \rightarrow S^2, t \in [0, 1]$, called a *pseudoisotopy*, such that $h_0 = \text{id}_{S^2}$, h_t is a homeomorphism for $t \in [0, 1)$, $h_t|_{P_f} = \text{id}_{P_f}$ for all t , and the fibers of h_1 and those of π coincide. The induced homeomorphism $S^2 \rightarrow \mathcal{J}$ conjugates f_Σ to an expanding Thurston map $g : S^2 \rightarrow S^2$, and the family of maps $f_t, t \in [0, 1]$ defined as the unique continuous solution of $h_t \circ f = f_t \circ h_t$, with $f_1 = g$, gives an homotopy through Thurston maps. The desired isotopy follows.

In the remainder of this work, we assume we are given a degree $d \geq 2$ Thurston map $f : S^2 \rightarrow S^2$ with postcritical set P .

Outline. In §2, we collect generalities on Thurston maps and related combinatorial objects, e.g. the pullback relations on curves and on arcs. In §3, we discuss algebraic invariants and define the selfsimilarity complex, Σ . In §4, we derive dynamical consequences from the contraction of the virtual endomorphism. §5 introduces the complex Γ associated to an open cover \mathcal{U}_0 , while §6 gives the proof that Σ and Γ can be chosen so that the natural map between them is a quasi-isometry. §§7 and 8 conclude the proof of Theorem 1.3.

Remark. Nekrashevych [Nek2, Cor. 5.13] has shown that in great generality, a pair of maps $\iota, f : \mathcal{M}_1 \rightarrow \mathcal{M}_0$ satisfying analogous algebraic contraction conditions is combinatorially equivalent to a pair $\iota', f' : \mathcal{M}'_1 \rightarrow \mathcal{M}'_0$ where the \mathcal{M}'_i are simplicial complexes, the maps ι', f' are piecewise affine, and the pair is expanding. While the result applies to our setting, the dimensions of \mathcal{M}'_i can be larger than two, and it seems difficult to arrange for the spaces \mathcal{M}'_i to coincide.

2 Thurston maps

2.1 Curves and arcs

Here, we introduce the pullback relations on curves and arcs defined by a Thurston map, and formulate obstructions to expansion in terms of these relations.

Denote by \mathcal{C} the set of free homotopy classes of essential, *oriented*, simple, closed, nonperipheral (that is, not homotopic into arbitrarily small neighborhoods of elements of P) curves in $S^2 - P$; we use the term *curve* for an element of \mathcal{C} .

Let o denote the union of the homotopy classes in $S^2 - P$ of curves which are either inessential or peripheral; we call such curves *trivial*. The *pullback relation* \longleftarrow on $\mathcal{C} \cup \{o\}$ is defined by setting $o \longleftarrow o$ and $\gamma_1 \longleftarrow \gamma_2$ if and only if γ_2 is homotopic in $S^2 - P$ to a component of $f^{-1}(\gamma_1)$. Thus $\gamma \longleftarrow o$ if and only if some preimage of γ is inessential or peripheral in $S^2 - P$. If $\gamma_1 \longleftarrow \gamma_2$ we write $\gamma_1 \leftrightarrow \gamma_2$ if there is a representative of γ_2 mapping injectively to a representative of γ_1 . The relation defined by \leftrightarrow we refer to as the *univalent pullback relation*.

We next define a similar pullback relation on certain arcs with endpoints in P . Let \mathcal{A} denote the set of isotopy classes (fixing P) of embedded arcs $\alpha \subset S^2$ whose interiors are contained in $S^2 - P$ and whose endpoints lie in P . A *preimage* $\tilde{\alpha}$ of α is, by definition, the closure of a connected component of the inverse image of the interior of α under f . If the endpoints of $\tilde{\alpha}$ lie in P , then $\tilde{\alpha}$ represents an element of \mathcal{A} . Adjoining a symbol o to the set \mathcal{A} to stand for the collection of arcs with one or both endpoints outside P , we obtain similarly a pullback relation on the set $\mathcal{A} \cup \{o\}$.

We recall that a *multicurve* is a finite subset of \mathcal{C} represented by disjoint curves. A *Levy cycle* is a sequence $(\gamma_0, \gamma_1, \dots, \gamma_{p-1})$ with $\gamma_i \leftrightarrow \gamma_{i+1 \bmod p}$ whose elements comprise a multicurve.

The result below is not needed for the proofs of the theorems, since it could be derived from Theorems 1.3 and Theorem 4.1 below. We include the

statement and proof since they are simple and help build intuition. It applies to maps which may have periodic branch points.

Proposition 2.1 *Suppose f is expanding. Given any curve C_0 or arc α_0 , there exists an integer N such that if $C_0 \leftarrow \dots \leftarrow C_n$ or $\alpha_0 \leftarrow \dots \leftarrow \alpha_n$ is an orbit with C_n, α_n nontrivial, then $n \leq N$. In particular, the univalent pullback relation on curves has no (Levy) cycles or wandering curves, and the pullback relation on arcs has no cycles and no wandering arcs.*

Abusing terminology somewhat, we summarize the conclusion by saying that the univalent pullback relation on curves and pullback relation on arcs have *no elements with arbitrarily long nontrivial iterates*.

Proof: If f were expanding with respect to a length structure, the conclusion would follow immediately, since lengths would decrease exponentially as one pulls back. Without this additional structure, we use instead coverings by small open balls. Let $\mathcal{U}_0, \mathcal{U}_i, i \in \mathbb{N}$ be as in the definition of expanding. Choose m sufficiently large so that C_0, α_0 are covered by some number, say L , of elements of \mathcal{U}_m and the union of these elements are contained in a regular neighborhood of C_0, α_0 which meets P in the same manner as does C_0, α_0 (that is, either not at all, or only at the endpoints of α_0) and which is homotopic, relative to P , to C_0, α_0 . It follows that for each n , C_n, α_n are covered by the same number L of elements of \mathcal{U}_{m+n} . Hence if the diameters of the elements of the \mathcal{U}_i tend to zero with i , the integer n cannot be arbitrarily large. ■

Since the pullback relations are natural with respect to combinatorial equivalence, Proposition 2.1 gives a necessary condition for a Thurston map f to be equivalent to an expanding map. It is not, however, a sufficient condition; see §4 for examples.

2.2 Orbifolds

We review here the definition of the orbifolds associated to a Thurston map and their fundamental groups.

For $x, y \in S^2$ let $\nu_0(y) = \text{lcm}\{\deg(f^n, x) : f^n(x) = y\}$ and $\nu_1(x) = \nu_0(f(x))/\deg(f, x)$. For $i = 0, 1$ let $\Sigma_i = \{x \in S^2 : \nu_i(x) > 1\}$, and let \mathcal{O}_i be the orbifold whose underlying topological space is $\{x \in S^2 : \nu_i(x) < \infty\}$ and whose weight function is ν_i . The sets Σ_i are called the *singular sets* of \mathcal{O}_i . Note that $\Sigma_0 = P$. There are no singular points of infinite weight if and only if there are no periodic branch points of f . In this case, the orbifolds \mathcal{O}_i

are compact. The singular sets satisfy $f^{-1}(\Sigma_0) \supset \Sigma_1$. Set $U_0 = S^2 - \Sigma_0$ and $U_1 = S^2 - f^{-1}(\Sigma_0)$. Then $U_1 \subset U_0$ and $f : U_1 \rightarrow U_0$ is a covering map.

Let $b_0 \in U_0$ be a basepoint and $b_1 \in f^{-1}(b_0)$ be one of its preimages. For $i = 0, 1$ let N_i denote the normal subgroup of $\pi_1(U_i, b_i)$ generated by the set of elements of the form g^k , where g is represented by a simple closed peripheral loop γ surrounding a puncture x of U_i , and the exponent k is the weight $\nu_i(x) < \infty$; if the weight is infinite, we do not add such a loop as a generator. The *orbifold fundamental groups* $\pi_1(\mathcal{O}_i, b_i)$ are by definition the quotient groups $\pi_1(U_i, b_i)/N_i$.

In what follows, we put $\mathcal{O} = \mathcal{O}_0$ and $G := \pi_1(\mathcal{O}, b_0)$.

As the (orbifold) universal covering deck group, G acts properly discontinuously on the universal covering of \mathcal{O} .

Note that each element of \mathcal{C} corresponds to a conjugacy class in G whose elements are of infinite order.

3 Algebraic constructions

In this section, we briefly summarize constructions and results of Nekrashevych [Nek1], specializing to the case of Thurston maps.

3.1 The virtual endomorphism

In the setup of §2.2, let $f_* : \pi_1(U_1, b_1) \rightarrow \pi_1(U_0, b_0)$ be the injective homomorphism induced by the covering $f : U_1 \rightarrow U_0$. Since f sends peripheral loops to peripheral loops, it follows from the definitions of the weight functions ν_i that $f_* : N_1 \rightarrow N_0$ is an isomorphism. This observation and the ‘‘Five Lemma’’ of homological algebra imply that the homomorphism f_* descends to a well-defined and injective map $\overline{f}_* : \pi_1(\mathcal{O}_1, b_1) \rightarrow \pi_1(\mathcal{O}_0, b_0) = G$. We denote the image group $\overline{f}_*(\pi_1(\mathcal{O}_1, b_1))$ by H ; it has finite index in G .

Let $\alpha : [0, 1] \rightarrow U_0$ be a path joining b_1 to b_0 and $\alpha_* : \pi_1(U_0, b_1) \rightarrow \pi_1(U_0, b_0)$ the induced isomorphism. Let $N'_0 = \alpha_*^{-1}(N_0)$. Since N_0 is normal, the subgroup N'_0 is normal and is independent of the choice of path α . Set $\pi_1(\mathcal{O}_0, b_1) = \pi_1(U_0, b_1)/N'_0$. Again, the map α_* descends to a well-defined isomorphism $\overline{\alpha}_* : \pi_1(\mathcal{O}_0, b_1) \rightarrow \pi_1(\mathcal{O}_0, b_0)$.

Since the inclusion $\iota : U_1 \hookrightarrow U_0$ sends peripheral loops to loops which are either peripheral or trivial, and since $\nu_0(x)$ divides $\nu_1(x)$ for all x , the induced map $\iota_* : \pi_1(U_1, b_1) \rightarrow \pi_1(U_0, b_1)$ is surjective and sends N_1 to N'_0 . It easily follows that the map ι_* also descends to a surjective map $\overline{\iota}_* : \pi_1(\mathcal{O}_1, b_1) \rightarrow \pi_1(\mathcal{O}_0, b_1)$.

Definition 3.1 *The virtual endomorphism induced by f is the homomorphism $\phi : H \rightarrow G$ defined by*

$$\phi = \bar{\alpha}_* \circ \bar{\iota}_* \circ (\bar{f}_*)^{-1}.$$

By construction, the virtual endomorphism ϕ associated to f is surjective.

The virtual endomorphism depends on the choices of the basepoint b_0 , the preimage b_1 , and the homotopy class of the path α . Different choices yield virtual endomorphisms which differ by pre- and/or post-composition by inner automorphisms. Up to this ambiguity, the virtual endomorphism is an invariant of the homotopy class of f relative to P . A combinatorial equivalence between Thurston maps conjugates, up to inner automorphisms, their virtual endomorphisms; the property of being contracting is preserved.

For $n \geq 2$, the n th iterate ϕ^n is the homomorphism whose domain is defined inductively by

$$\text{dom}\phi = H; \quad \text{dom}\phi^n = \{g \in H : \phi(g) \in \text{dom}\phi^{n-1}\}$$

and whose rule is given by iterating ϕ a total of n times.

Suppose S is a finite generating set for G . We denote by $\|g\|$ the word length of g in the generators S .

Definition 3.2 *The virtual endomorphism $\phi : H \rightarrow G$ is called contracting if the contraction ratio*

$$\rho = \limsup_{n \rightarrow \infty} \left(\limsup_{\|g\| \rightarrow \infty} \frac{\|\phi^n(g)\|}{\|g\|} \right)^{1/n} < 1.$$

The contraction ratio of the virtual endomorphism ϕ is independent of the choices used in its construction.

As an algebraic object, the virtual endomorphism is rather straightforward to describe. The next subsection gives a less familiar, but more natural, construction.

3.2 The biset \mathfrak{M}

Suppose $\ell_0, \ell_1 : [0, 1] \rightarrow S^2 - P$ are two arcs joining the basepoint b_0 to a common point in $S^2 - P$. We say ℓ_0, ℓ_1 are *homotopic in \mathcal{O}* if the loop formed by traversing ℓ_0 first and then the reverse of ℓ_1 represents the trivial element of G . We denote by \mathfrak{M} the set of all homotopy classes $[\ell]$ of arcs ℓ in \mathcal{O} joining b_0 to one of the d elements of $f^{-1}(b_0)$. For $[\ell] \in \mathfrak{M}$ we denote by $z_{[\ell]} \in f^{-1}(b_0)$ the

corresponding common endpoint. The group G acts on \mathfrak{M} via two commuting actions as follows. On the right, it acts by pre-concatenation:

$$\ell \cdot g := \ell * g$$

where first g , then ℓ is traversed. Note that the right action is free and has d orbits. The left action is obtained by taking preimages of the loops under f :

$$g \cdot \ell := f^{-1}(g)[z_{[\ell]}] * \ell$$

where first ℓ , then the lift $f^{-1}(g)[z_{[\ell]}]$ of g based at $z_{[\ell]}$, is traversed. The set \mathfrak{M} equipped with these two commuting actions of G defines the G -biset associated to f .

In the next subsection, we show that choosing certain representatives of elements of \mathfrak{M} and representatives of generators of G gives rise to an abstract 1-complex Σ and a projection map $\pi_\Sigma : \Sigma \rightarrow S^2$. Along the way, we reformulate the notion of contraction in terms of an action of G on the set of words in an alphabet X of size d .

3.3 The selfsimilarity complex Σ

Choose a bijection $\Lambda : X \rightarrow f^{-1}(b_0)$, where X is a finite set of cardinality d . For each $x \in X$, let λ_x be an arc in $S^2 - P_f$ joining b_0 to $\Lambda(x)$. The collection $\{[\lambda_x] : x \in X\}$ of elements of \mathfrak{M} is called a *basis* of the biset $\mathfrak{M} = \mathfrak{M}_f$.

For $n \in \mathbb{N}$, consider now the n th iterate f^n of f . The orbifolds $\mathcal{O} = \mathcal{O}_f$ and \mathcal{O}_{f^n} coincide, so $G = G_f = G_{f^n}$. By lifting the arcs λ_x under f^n , we obtain an identification $\Lambda : X^n \rightarrow f^{-n}(b_0) \times \{n\}$ of the set X^n of words of length n in the alphabet X with the fiber $f^{-n}(b_0) \times \{n\}$ and a corresponding basis for the G -biset \mathfrak{M}_{f^n} associated to the n th iterate of f . We denote by $X^* = \cup_n X^n$, with $X_0 = \{\emptyset\}$ consisting of the empty word. We write $|w| = n$ if $w \in X^n$.

Suppose S is a generating set for G . For each $s \in S$, let γ_s be a loop in $S^2 - P$ based at b_0 representing s . By lifting loops under iterates of f , we obtain an action of G on X^* preserving the lengths of words and acting transitively on each X^n .

The basepoint b_0 , arcs $\lambda_x, x \in X$, and loops $\gamma_s, s \in S$, we refer to as *defining data*.

Since the right action is free, for each word $u \in X^*$, and each $g \in G$, there are a unique word $v \in X^n$ and a unique $h \in G$ with $g.u = v.h$; the element h is called the *restriction*¹ of g to u , and is denoted $g|_u$. The interpretation in terms of defining data is as follows. Let λ_u, λ_v denote the concatenations of lifts of

¹sometimes the term “section” is used.

arcs corresponding to u and v , so that the endpoints of λ_u, λ_v correspond to u and v . Then v is the endpoint of $f^{-n}(g)[u]$ and h is represented by the loop which first traverses λ_u , then $f^{-n}(g)[u]$, then the reverse of λ_v .

By [Nek1, Prop. 2.11.11], contraction of the virtual endomorphism is equivalent to the following: *there is a finite set \mathcal{N} such that for each $g \in G$, there is an integer n such that for all $v \in X^*$ satisfying $|v| \geq n$, the restriction $g|_v \in \mathcal{N}$.* In this case, one says the biset \mathfrak{M} is *hyperbolic*. This is an asymptotic condition: \mathfrak{M}_f is hyperbolic if and only if \mathfrak{M}_{f^m} is hyperbolic for some $m \geq 1$, since the contraction coefficients for ϕ and ϕ^m are necessarily simultaneously less than 1.

Suppose now that defining data has been chosen. The *selfsimilarity complex* Σ is the infinite 1-complex whose vertex set is X^* and whose edges, which come in two types, *horizontal* and *vertical*, are defined as follows. Given $w \in X^*$, a horizontal edge, labelled s , joins w to $s.w$; a vertical edge, labelled x , joins w to xw . The selfsimilarity complex, with edges of length 1, is a proper geodesic metric space, and the local valence at each nonempty vertex is the same if S is symmetric.

The right shift $\sigma : X^* \rightarrow X^*$ determines a cellular self-map $f_\Sigma : \Sigma' \rightarrow \Sigma$ where Σ' is the subcomplex determined by nonempty words. When \mathfrak{M} is hyperbolic, Σ is hyperbolic in the sense of Gromov, and the boundary \mathcal{J} is compact. Since the virtual endomorphism ϕ is surjective, the action is recurrent, so \mathcal{J} is also connected and locally connected [Nek1, Theorem 3.5.1]. The map f_Σ extends to a map of the boundary $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$, yielding a dynamical system.

Although Σ is an abstract 1-complex, path-lifting of the defining data determines a map $\pi_\Sigma : \Sigma \rightarrow S^2$; see [Pil] and [Nek1, Theorem 5.5.3]. If f is not expanding with respect to a length structure on S^2 , however, π_Σ need not extend to a map $\mathcal{J} \rightarrow S^2$.

4 Dynamical consequences of the hyperbolicity of \mathfrak{M}

Suppose f has no periodic branch points. Proposition 2.1 gives necessary conditions for f to be expanding. The main result of this section is to derive the same conclusion from the algebraic assumption that ϕ is contracting; equivalently, that \mathfrak{M} is hyperbolic.

Theorem 4.1 *Suppose f has no periodic branch points. If \mathfrak{M} is hyperbolic, the univalent pullback relation on curves and the pullback relation on arcs have no elements with arbitrarily long nontrivial iterates.*

Proof: The central ingredient in the proof is the following Lemma, which will be proved later.

Lemma 4.2 *Suppose $C_0 \leftrightarrow C_1$. Given any $g_0 \in G$ whose conjugacy class corresponds to C_0 , there exists $g_1 \in G$ whose conjugacy class corresponds to C_1 and $x \in X$ with $g_0 \cdot x = x \cdot g_1 \in \mathfrak{M}$.*

To prove the Theorem, suppose $C_0 \leftrightarrow C_1 \leftrightarrow \dots \leftrightarrow C_n$. Induction and Lemma 4.2 implies that there exist $g_i \in G, x_i \in X$ such that

$$g_0 \cdot x_0 x_1 \dots x_i = x_0 x_1 \dots x_i \cdot g_{i+1}, \quad i = 0, \dots, n-1.$$

In particular, upon setting $w = x_0 x_1 \dots x_n \in X^{n+1}$, we have

$$g_0 \cdot w = w \cdot g_n;$$

recall that by definition, $g_n = g_0|_w$. Since \mathfrak{M} is hyperbolic, there is a finite set $\mathcal{N} \subset G$ such that $g_n = g_0|_w \in \mathcal{N}$ whenever $n = |w|$ is sufficiently large. This immediately implies that the number of distinct elements g_n of \mathcal{C} arising in this way is bounded and hence that the collection of curves $\{C_n\}$ arising in this way is also bounded, independent of n .

To rule out cycles, we use topology to simplify the algebra. Consider first the case of a Levy cycle of period one, i.e. a curve C for which $C \leftrightarrow C$. By varying f within its homotopy class, and using the fact that C is oriented, we may assume $f = \text{id}$ on a neighborhood of C . Take the basepoint b_0 for \mathfrak{M} to lie on C ; it becomes a fixed point of f . Let g be represented by C , regarded as a loop based at b_0 . By choosing a basis for \mathfrak{M} with one arc equal to the constant path at b_0 and corresponding to an element x in the alphabet X , we find $g \cdot x = x \cdot g$. Induction shows then that for all $n \in \mathbb{N}$ and $m \in \mathbb{Z}$ we have that if $w = x x \dots x \in X^n$ then

$$g^m \cdot w = w \cdot g^m.$$

But this is impossible if \mathfrak{M} is hyperbolic: g has infinite order, and the identity above shows that for all n , $g^m|_w = g^m$ which cannot lie in \mathcal{N} if m is sufficiently large.

The case of a Levy cycle of period larger than one can be ruled out by passing to an iterate.

Analyzing the pullback relation on arcs is reduced to that of the pullback relation on curves. An arc which does not eventually become trivial must have both endpoints in cycles of P . For such arcs α_0, α_1 , since f has no periodic branch points, whenever $\alpha_0 \leftarrow \alpha_1$, the boundaries of regular neighborhoods of α_0 and α_1 yield curves C_0 and C_1 for which $C_0 \leftrightarrow C_1$.

If $\#P > 3$ then C_0, C_1 must be nontrivial, and the eventual triviality of curve \leftrightarrow -orbits, already established, implies eventual triviality of arc \leftarrow -orbits. So assume $\#P = 3$. By Thurston's characterization [DH], f is equivalent to a rational function without periodic critical points. Such a map is uniformly expanding with respect to a complete length structure, so such arc orbits are always eventually trivial. ■

We now turn to the proof of the Lemma, which is just an exercise in definitions.

Proof: Suppose defining data has been chosen. Choose a representative γ_0 for g , and let $\tilde{\gamma}_0$ be a preimage of γ_0 corresponding to C_1 ; it exists, by homotopy lifting and the fact that C_1 maps to C_0 by degree one. Let \tilde{b} be the unique preimage of b lying on $\tilde{\gamma}_0$, and suppose x corresponds to \tilde{b} , i.e. $\Lambda(x) = \tilde{b}$. The loop given by

$$\gamma_1 := \bar{\lambda}_x * \tilde{\gamma}_0 * \lambda_x$$

(where λ_x is traversed first) represents an element $g_1 \in G$. Moreover, by the definition of the biset,

$$x \cdot g_1 = [\lambda_x * \bar{\lambda}_x * \tilde{\gamma}_0 * \lambda] = g_0 \cdot x.$$

■ Lemma 4.2

Examples showing the conditions in Propositions 2.1 and Theorem 4.1 are necessary but not sufficient can be found among Thurston maps induced by torus endomorphisms that are hyperbolic but not expanding.

Let $T^2 = \mathbb{R}^2/\mathbb{Z}^2$, let A be an integral matrix with determinant larger than one, and let $F : T^2 \rightarrow T^2$ be the noninvertible endomorphism on the torus induced by A . There is a corresponding pullback relation on curves, defined analogously as in §2. In this setting, all preimages of nontrivial curves are nontrivial, are homotopic to each other, and map by the same degree.

Proposition 4.3 *If a curve C has arbitrarily long nontrivial iterates under the univalent pullback relation, then either $C \leftrightarrow C$ or $C \leftrightarrow -C$, where $-C$ denotes the curve C with the opposite orientation. This occurs if and only if the matrix A has an eigenvalue λ equal to $+1$ or -1 .*

Proof: The sufficiency in the second statement is obvious. The proof will show both the first statement, and the necessity; it is a straightforward reinterpretation of the arguments of [Nek1, Prop. 2.9.2]. Use the standard basis to identify

\mathbb{Q}^2 with $H_1(T^2, \mathbb{Q})$ and $\pi_1(T^2)$ with $H_1(T^2, \mathbb{Z}) = \mathbb{Z}^2$, so that A is the matrix for the induced map on rational and integral homology. Let $B = A^{-1}$. Observe that if $C = C_0 \leftrightarrow C_1$ then in fact all preimages of C_0 map by degree one and the monodromy corresponding to C_0 is trivial. Thus if $C_0 \leftrightarrow C_1 \leftrightarrow C_2 \leftrightarrow \dots$ is an infinite sequence of related curves (either wandering, or cycling) then as an element of $\pi_1(T^2)$, the curve C_0 represents a nontrivial element of the kernel $U < \mathbb{Z}^2 < \mathbb{Q}^2$ of the corresponding iterated monodromy group action. The group U is B -invariant, so $U \otimes \mathbb{Q}$ is a B -invariant \mathbb{Q} -linear subspace of \mathbb{Q}^2 ; let $C = B|_{U \otimes \mathbb{Q}}$. Then the characteristic polynomial q of C is monic, has integer coefficients, and is a factor of the characteristic polynomial of B . If $1/\lambda$ is a root of q then λ is an eigenvalue of A . Both λ and $1/\lambda$ are algebraic integers, hence the norm of λ is equal to 1. The condition $\det(A) > 1$ implies the algebraic degree of λ must be equal to 1, so $\lambda = \pm 1$. ■

The endomorphism $F : T^2 \rightarrow T^2$ is *hyperbolic* if A has no eigenvalues of unit modulus. Under the assumption $\det(A) > 1$, the following is easily shown via straightforward computations.

Proposition 4.4 *The endomorphism F is hyperbolic but not expanding if and only if $\operatorname{tr}(A) - \det(A) > 1$.*

The endomorphism $F : T^2 \rightarrow T^2$ commutes with $\pm \operatorname{id}_{T^2}$ and descends to a Thurston map $f : S^2 \rightarrow S^2$; the corresponding orbifold \mathcal{O} has signature $(2, 2, 2, 2)$. The singular points are the images of the points of order at most 2.

Theorem 4.5 *Suppose $\det(A) > 1$ and $\operatorname{tr}(A) - \det(A) > 1$, and let $f : S^2 \rightarrow S^2$ be the induced Thurston map. Then*

1. *f is not equivalent to an expanding map, but*
2. *the univalent pullback relation on curves and the pullback relation on arcs have no elements with arbitrarily long nontrivial iterates.*

Proof: Proposition, 4.4, Proposition 4.3, and Theorem 1.1 imply both conclusion (1) and the first part of (2). Suppose $\alpha \subset S^2$ is an arc. The two preimages of α under the branched covering $T^2 \rightarrow S^2$, when concatenated, yield a nontrivial curve, C . If $\alpha_0 \leftarrow \alpha_1$ then for the corresponding curves in T^2 we have $C_0 \leftrightarrow C_1$. Thus nontrivial arc orbits lift to nontrivial univalent curve orbits. ■

5 The covering complex Γ

In this section, we begin by summarizing constructions and results of [HP1, Chapter 3], specializing to the case of Thurston maps. We associate to f and a suitable open covering \mathcal{U}_0 another infinite one-complex, Γ , this time a graph. When it is hyperbolic, it has a natural boundary, and there is a natural map $\pi_\Gamma : S^2 \rightarrow \partial\Gamma$.

For convenience, equip S^2 with the usual Euclidean length metric. Let \mathcal{U}_0 be a finite open cover of S^2 by small balls B for which $\#B \cap P \leq 1$ with equality only when the corresponding point of P is in the center of B . For $n \in \mathbb{N}$ let \mathcal{U}_n be the open cover whose elements are connected components of sets of the form $f^{-n}(U), U \in \mathcal{U}_0$. It is convenient to set $\mathcal{U}_{-1} := \{S^2\}$. If $U \in \mathcal{U}_n$ we write $|U| = n$. A key observation here is that if f has no periodic branch points, then there is an upper bound on the degrees of restrictions $f^n : \tilde{U} \rightarrow U, \tilde{U} \in \mathcal{U}_n, n \in \mathbb{N}$.

The vertex set of Γ is defined as $\cup_{n \geq -1} \mathcal{U}_n$, i.e. one vertex for each component. We denote the vertex corresponding to the unique element of \mathcal{U}_{-1} by o . Horizontal edges join U_1, U_2 if $|U_1| = |U_2|$ and $U_1 \cap U_2 \neq \emptyset$. Vertical edges join U, V if $||U| - |V|| = 1$ and $U \cap V \neq \emptyset$. Γ is a graph with root o . We equip Γ with the graph path metric, so that each edge has length 1. Then Γ is a proper, geodesic metric space; let d denote the corresponding distance function. If f has no periodic branch points, then the valence of a vertex of Γ is uniformly bounded, independent of the vertex.

Let $\varepsilon > 0$ be small, and let $\varrho_\varepsilon : \Gamma \rightarrow (0, \infty)$ be given by $\varrho_\varepsilon(x) = \exp(-\varepsilon d(x, o))$. Upon redefining the lengths of curves via

$$\ell_\varepsilon(\gamma) := \int_\gamma \varrho_\varepsilon ds$$

we obtain an incomplete length metric space. Its completion compactifies Γ by adding a boundary at infinity $\partial_\varepsilon\Gamma$. Points in its boundary $\partial_\varepsilon\Gamma$ are represented by vertical geodesic rays from the root. Given $a \in S^2$, let $U_n(a), n = -1, 0, 1, 2, \dots$ be a sequence of open sets for which $a \in U_n \in \mathcal{U}_n$. Then $\{U_n\}$ defines a geodesic ray and, therefore, a point of $\partial_\varepsilon\Gamma$. We obtain a well-defined surjective continuous map $\pi_\Gamma : S^2 \rightarrow \partial_\varepsilon\Gamma$. In particular, $\partial_\varepsilon\Gamma$ is connected and locally connected. When Γ is hyperbolic and $\varepsilon > 0$ is chosen small enough, then $\partial_\varepsilon\Gamma$ coincides with its visual boundary $\partial\Gamma$. We will always choose ε as above when Γ is hyperbolic.

Lemma 5.1 *If Γ is hyperbolic, the projection $\pi_\Gamma : S^2 \rightarrow \partial\Gamma$ is monotone, i.e. its fibers are connected.*

Proof: For $z \in S^2$ and $n \geq -1$, pick $U_n(z) \in \mathcal{U}_n$ such that $z \in U_n(z)$. It follows that the sequence $(U_n(z))_n$ defines a ray in Γ . Suppose $\pi_\Gamma(x) = \pi_\Gamma(y) = \xi \in \partial\Gamma$. From the definition of π and the hyperbolicity of Γ , there is a positive integer D such that for every $n \in \mathbb{N}$, there exist a “chain” $U_n^0, U_n^1, \dots, U_n^D$ of elements of \mathcal{U}_n such that $x \in U_n^0$, $y \in U_n^D$, and $U_n^{i-1} \cap U_n^i \neq \emptyset$ for all $i = 1, \dots, D$. For $n \in \mathbb{N}$ let $Y_n = U_n^0 \cup \dots \cup U_n^D$. Then Y_1, Y_2, Y_3, \dots is a sequence of open connected sets containing both x and y . Set

$$K = \bigcap_{k \geq 0} \overline{\bigcup_{n \geq k} Y_n}.$$

It follows that $K \subset S^2$ is a continuum containing both x and y . We will prove that $\pi_\Gamma(K) = \xi$ and, therefore, that $\pi_\Gamma^{-1}(\xi)$ is connected. Fix $z \in K$. For any $k \geq 0$, $U_k(z) \cap K \neq \emptyset$, so we may find $n \geq k$ such that $Y_n \cap U_k(z) \neq \emptyset$. This means that

$$\begin{aligned} d_\varepsilon(\pi_\Gamma(z), \xi) &\leq d_\varepsilon(\pi_\Gamma(z), U_k(z)) + d_\varepsilon(U_k(z), \{U_n^j, 0 \leq j \leq D\}) \\ &\quad + \text{diam}_\varepsilon\{U_n^j, 0 \leq j \leq D\} + d_\varepsilon(\xi, \{U_n^j, 0 \leq j \leq D\}) \\ &\lesssim e^{-\varepsilon k} + e^{-\varepsilon k} + e^{-\varepsilon n} + e^{-\varepsilon n} \lesssim e^{-\varepsilon k}. \end{aligned}$$

Since k is arbitrary, we obtain that $\pi_\Gamma(z) = \xi$. ■

From [Dav, Theorem 5], we have that the decomposition of S^2 by the fibers of π_Γ is upper-semicontinuous, and the quotient space is homeomorphic, via the map induced by π_Γ , to $\partial\Gamma$.

The map f induces a cellular map $f_\Gamma : \Gamma' \rightarrow \Gamma$, where Γ' is the subgraph on vertices at level ≥ 1 . Since it sends vertical rays to vertical rays, it induces a continuous map $f_\Gamma : \partial\Gamma \rightarrow \partial\Gamma$. Since it sends sufficiently small round balls to round balls [HP1, Prop. 3.2.2], it is an open map; it is easily seen to be at most d -to-1, so it is closed as well. In particular, the local degree function

$$\deg(f_\Gamma, \xi) := \inf_W \{\#f_\Gamma^{-1}(\zeta) \cap W : \zeta \in f_\Gamma(W), \xi \in W \text{ open}\}$$

is well-defined.

6 Coincidence of quasi-isometry types

Recall that if X, Y are metric spaces, a function $\Phi : X \rightarrow Y$ is a *quasi-isometry* if there exist constants $\lambda \geq 1, c \geq 0, b \geq 0$ such that

$$\frac{1}{\lambda}|x - x'| - c \leq |\Phi(x) - \Phi(x')| \leq \lambda|x - x'| + c$$

and such that $|y - \Phi(X)| \leq b$ for all $y \in Y$; here $|x - x'| = d_X(x, x')$, etc.. For geodesic metric spaces, hyperbolicity is preserved under quasi-isometries.

Recall that a choice of defining data defines a projection $\pi_\Sigma : \Sigma \rightarrow S^2$ and an identification map $\Lambda : X^* \rightarrow \cup_{n \geq 0} f^{-n}(\{b\})$. The main result of this section is

Theorem 6.1 *Let $f : S^2 \rightarrow S^2$ be a Thurston map without periodic critical points. Then there exist defining data for the associated selfsimilarity complex Σ and the covering graph Γ and a level-preserving quasi-isometry $\Phi : \Sigma \rightarrow \Gamma$ such that, for every vertex $v \in \Sigma^0$, we have $\Lambda(v) \in \Phi(v)$ and $\Phi \circ f_\Sigma = f_\Gamma \circ \Phi$.*

Remark 6.2 *If we admit periodic branch points, then the same proof will show that Σ is quasi-isometric to a covering complex Γ built from a cover of the complement in the sphere of a neighborhood of the set of periodic branch points.*

Proof: We first choose the defining data and list some elementary properties. For convenience, we endow S^2 with its ordinary Euclidean length metric.

Choosing the defining data. Choose the basepoint $b \in S^2 - P$ to be a regular point. We may choose representative loops $\gamma_s, s \in S$, so that each is simple, peripheral, and has the property that for some common $\epsilon > 0$, the closed ϵ -neighborhood $N_\epsilon(\gamma_s)$ is contained in a closed annulus A contained in $S^2 - P$ surrounding on one side at most one point of P . The peripheral condition and the assumption that there are no periodic critical points implies that there exists N_0 such that for all n ,

$$\deg(f^n : \tilde{A} \rightarrow A) \leq N_0$$

for all preimages \tilde{A} of A under f^{-n} . In this situation, if the above degree is say δ , then in the graph Λ , for each vertex v at level n with $\Lambda(v) \in \tilde{A}$, there is a closed horizontal edge- path of length δ whose edges are labelled by s corresponding and which corresponds to the monodromy at level n induced by γ_s .

We may also choose representative arcs $\lambda_x \subset S^2 - P, x \in X$, so that each is simple and has the property that for some common $\epsilon > 0$, the closed ϵ -neighborhood $N_\epsilon(\lambda_x)$ is a closed disk D contained in $S^2 - P$. Thus

$$\deg(f^n : \tilde{D} \rightarrow D) = 1$$

for all preimages \tilde{D} of D under f^{-n} .

For $w \in X^*$, let λ_{xw} be the lift of λ_x under $f^{|w|}$ which starts at $\pi_\Sigma(w)$.

Since f is a branched covering, given $\delta > 0$, there is a constant $r(\delta) > 0$ with $r(\delta) \rightarrow 0$ as $\delta \rightarrow 0$ such that for any ball $B(y, \delta) \subset S^2$, any preimage of B under f has diameter $< r(\delta)$. It follows that we may choose the defining data \mathcal{U}_0 for the covering graph Γ to consist of a finite collection of small open balls satisfying the following properties.

Given a set $E \subset S^2$, we denote by

$$\mathcal{U}_n(E) = \{U \in \mathcal{U}_n : E \cap U \neq \emptyset\}.$$

1. $\forall U \in \mathcal{U}_0$, $P \cap \partial U = \emptyset$, $\#P \cap U \leq 1$, and if equality holds, then the postcritical point in U is at the center of U ;
2. $\forall U \in \mathcal{U}_0$, $\mathcal{U}_0(U) \cap P \leq 1$;
3. $\forall x \in X$, $\overline{\mathcal{U}_0(\lambda_x)} \subset N_\epsilon(\lambda_x)$;
4. $\forall s \in S$, $\overline{\mathcal{U}_0(\gamma_s)} \subset N_\epsilon(\gamma_s)$;
5. $\forall U_0 \in \mathcal{U}_0$ and $U_1 \in \mathcal{U}_1$ with $U_0 \cap U_1 \neq \emptyset$, the union $U_0 \cup U_1$ is contained in a ball B such that the ball $2B$ of twice the radius meets at most one point of P ;
6. $\forall p \in S^2$ there exists an arc α joining the basepoint b to p such that $\mathcal{U}_0(\alpha)$ is contained in a disk which meets P in at most one point.
7. There is exactly one element of \mathcal{U}_0 containing the basepoint b , and this basepoint is the center of the ball.

Basic properties.

(BP1). Condition 1 implies that every element $U \in \mathcal{U}_n$ has a unique preferred point c_U (we call it the center of U) which is the unique preimage of the center c of the ball $f^{on}(U)$ under f^n .

(BP2). Conditions 1 and 2 imply the following. Since there are no periodic branch points, there is a uniform upper bound q on the local degree of all iterates of f . Suppose $f^n : \tilde{U}_i \rightarrow U_i \in \mathcal{U}_0, i = 0, 1$. If $\tilde{U}_0, \tilde{U}_1 \in \mathcal{U}_n$ intersect, then at most one of \tilde{U}_0, \tilde{U}_1 meets a branch point of f^n ; say it is \tilde{U}_0 . Then there are $q' \leq q - 1$ elements $\tilde{U}_2, \dots, \tilde{U}_{q'}$ of \mathcal{U}_n , each a preimage of U_1 , such that $f^{on} : \tilde{U}_0 \cup \tilde{U}_1 \cup \dots \cup \tilde{U}_{q'} \rightarrow U_0 \cup U_1$ is proper and of degree q' .

(BP3). From condition 6, we define A_0 to be the finite set of curves $\alpha_U : [0, 1] \rightarrow S^2$ joining the center $c_U = \alpha_U(0)$ of U to $b = \alpha_U(1)$. Let also A_n denote all the lifts of these curves under f^n : these are curves joining every center of $U \in \mathcal{U}_n$ to a point of $f^{-n}(b)$. Since every α_U is an arc contained in a disk which intersects P in at most one single point, there is some L independent of n and $U \in \mathcal{U}_n$ such that α_U is covered by at most L elements of \mathcal{U}_n .

(BP4). Condition 7 is merely to make the exposition in what follows more convenient, for with this choice we will define $\Phi : \Sigma \rightarrow \Gamma$ in a simple way.

Strategy of the proof. Define $\Phi : V(\Sigma) \rightarrow V(\Gamma)$ by setting $\Phi(w)$ to be the unique set $U \in \mathcal{U}_{|w|}$ which contains $\Lambda(w)$. We extend Φ over the edges to all of Σ as follows. Suppose γ_s represents a horizontal edge at level 0 (a loop) in Σ . In the sphere, as this edge is traversed, it passes through a sequence of elements of \mathcal{U}_0 ; we pick such a sequence, and associate to this a horizontal loop $\Phi(\gamma_s)$ in Γ at level 0, parameterized at constant speed. We define Φ to send the edge of Σ corresponding to γ_s to this loop $\Phi(\gamma_s)$. Similarly, each horizontal edge $e = e_s$ at level n in Σ^1 corresponds to a lift under f^n of a unique edge γ_s , $s \in S$, at level zero; hence by lifting the chosen covering for γ_s we obtain again a curve $\Phi(e)$ in Γ joining the images under Φ of their extremities by definition.

Similarly, for $x \in X$, we associate to the arc λ_x an edge-path $\Phi(\lambda_x)$ in Γ corresponding to a covering of the chosen arc $\alpha_{\Lambda(x)}$ by elements of \mathcal{U}_0 concatenated with the vertical edge of Γ^1 between the elements $U \in \mathcal{U}_0$ and $V \in \mathcal{U}_1$ containing $\Lambda(x)$. We again extend Φ to the other vertical edges equivariantly by path-lifting.

We also define a coarse inverse $\Psi : V(\Gamma) \rightarrow V(\Sigma)$ by setting $\Psi(U) = w$, where $w \in V(\Sigma)$ is the vertex such that $\Lambda(w)$ is the endpoint of α_U . Note that $\Psi \circ \Phi$ is the identity on $V(\Sigma)$ since α_U is trivial for preimages of $U_b \in \mathcal{U}_0$ containing b , and $\Phi \circ \Psi$ maps a vertex $U \in \Gamma^0$ to the endpoint of α_U .

To prove that Φ is a quasi-isometry, we will show that both Φ and Ψ are Lipschitz on the set of vertices and that the image of Φ is cobounded. The conclusion will follow since $\Psi \circ \Phi$ is the identity on Σ^0 . To prove that each map is Lipschitz, we will show that the distance between the images of the two extremities of an edge is uniformly bounded.

Φ is Lipschitz. Let us consider the horizontal case first. Suppose $v, w \in V(\Sigma)$, $|v| = |w| = n$, and v, w are joined by a horizontal edge, so that $w = s.v$ for some $s \in S$. Equivalently, there is a lift $\tilde{\gamma}_s$ of γ_s under f^n joining $\pi_\Sigma(v)$ to $\pi_\Sigma(w)$. Note that $f^n : \tilde{\gamma}_s \rightarrow \gamma_s$ is 1-1 away from the extremities. The curve γ_s is covered by $\#\mathcal{U}_0(\gamma_s)$ elements of \mathcal{U}_0 . The lift $\tilde{\gamma}_s$ is covered by at most $\#\mathcal{U}_0(\gamma_s) + 1$ elements of \mathcal{U}_n . By the definition of the metric in Γ , this implies that $|\Phi(v) - \Phi(w)|_{hor} \leq (\#\mathcal{U}_0(\gamma_s) + 1)$. Note that we may assume that $\#\mathcal{U}_0(\gamma_s) \leq \#\mathcal{U}_0 - 1$. Since $|\Phi(v) - \Phi(w)|_\Gamma \leq |\Phi(v) - \Phi(w)|_{hor}$, we conclude that

$$|\Phi(v) - \Phi(w)|_\Gamma \leq C_{hor} := \#\mathcal{U}_0.$$

Now let us consider the vertical case. Suppose $|v| = n + 1$, $|w| = n$, and v and w are joined by a vertical edge in Σ . By definition, $w = xv$ for some $x \in X$, and the points $\Lambda(w)$ and $\Lambda(v)$ in S^2 are joined by a lift of the arc λ_x under

f^{-n} starting at $\Lambda(v)$. The arc λ_x is covered by $\#\mathcal{U}_0(\lambda_x)$ elements of \mathcal{U}_0 , each contained in the neighborhood $N_\epsilon(\lambda_x)$, which is a disk in $S^2 - P$. It follows that $\tilde{\lambda}_x$ is covered by the same number $\#\mathcal{U}_0(\lambda_x)$ of elements of \mathcal{U}_n . By the definition of the metric in Γ , this implies that $|\Phi(v) - \Phi(w)|_\Gamma \leq \#\mathcal{U}_0(\lambda_x) + 1$. Hence

$$|\Phi(v) - \Phi(w)|_\Gamma \leq C_{ver} := \max\{\#\mathcal{U}_0(\lambda_x) : x \in X\} + 1.$$

We conclude that Φ is $C = \max\{C_{hor}, C_{ver}\}$ -Lipschitz on $V(\Sigma)$.

The image of Φ is cobounded in Σ . The proof of this follows along exactly the same lines. For any $U \in \mathcal{U}_n$, the curve $\alpha_U \in A_n$ joins c_U to $\Phi \circ \Psi(U)$, whose distance is bounded by L according to the basic property (BP3).

Realization of the covering complex Γ . Proving that Ψ is Lipschitz will require some more data. The covering complex Γ comes equipped with a self-map $f_\Gamma : \Gamma \setminus B(X, 1) \rightarrow \Gamma$. We will find a continuous map $\Gamma \rightarrow S^2$ which conjugates f_Γ and f where defined. First, if $e = (U_1, U_2)$ denotes an edge with $U_1, U_2 \in \mathcal{U}_0$, we consider β_e to be the geodesic (in the round metric) joining their centers; recall that U_j are spherical balls at level 0, so that $\beta_e \subset U_1 \cup U_2$. Denote by B_0 this finite set of curves, and let B_n denote the lifts of the elements of B_0 under f^n .

If $e = (U_0, U_1)$ is a vertical edge of Γ with $U_0 \in \mathcal{U}_0$ and $U_1 \in \mathcal{U}_1$, we may find an arc η_e joining their centers inside $U_0 \cup U_1$. Let H_0 denote the sets of arcs defined for every such edge, and for each $n \geq 0$, define H_n to be the set of their lifts under f^n .

$\Psi := \Phi^{-1}$ is Lipschitz. Let us first consider an edge e joining two sets \tilde{U}_1, \tilde{U}_2 at the same level $n \geq 0$. Let $\tilde{\gamma} = \alpha_{\tilde{U}_1} * \beta_e * \overline{\alpha_{\tilde{U}_2}}$. The image $f^n(\tilde{\gamma})$ is a closed curve based at b which is homotopically nontrivial in $S^2 \setminus P$ since $\tilde{\gamma}$ is not closed. This curve is also the concatenation of three curves belonging to $A_0 \cup B_0$ which is a finite set. Let

$$C_{hor} := \max_\gamma \|g(\gamma)\|$$

be the maximum word length of the corresponding group elements $g(\gamma)$, where γ varies over this finite collection of loops just constructed. We conclude that

$$|\Psi(\tilde{U}_1) - \Psi(\tilde{U}_2)|_\Sigma \leq C_{hor}.$$

The case of vertically adjacent vertices proceeds similarly in spirit. Suppose² \tilde{U}_1, \tilde{U}_0 are vertically adjacent in Γ , with $|\tilde{U}_1| = |\tilde{U}_0| + 1 = n + 1$, and

²Here, it is convenient to index things using zero and one, instead of one and two as above.

let $e = (\tilde{U}_1, \tilde{U}_0)$. The curves $\alpha_{\tilde{U}_1}$ and $\alpha_{\tilde{U}_2}$ join their centers to elements $\pi_\Sigma(w)$ and $\pi_\Sigma(v)$, $w, v \in V(\Sigma)$. Set

$$\tilde{\gamma} = \lambda_w * \alpha_{\tilde{U}_1} * \eta_e * \bar{\alpha}_{\tilde{U}_0}$$

which is a curve joining to points from $\pi_\Sigma(V(\Sigma))$ at the same level n . It follows that $\gamma = f^n(\tilde{\gamma})$ is a loop based at b of the form

$$\gamma = \lambda_x * \alpha_1 * \eta * \bar{\alpha}_0.$$

Hence it belongs to a finite set of curves (see Figure 1).




Figure 1: Showing $|\Psi(\tilde{U}_1) - \Psi(\tilde{U}_0)|_\Sigma = O(1)$. The disk $f(U_1)$ is not shown.

We let

$$C_{ver} := \max_{\gamma} \|g(\gamma)\|$$

be the maximum word length of the corresponding group elements $g(\gamma)$, where γ varies over the finite collection of loops just constructed. Thus

$$|\Psi(\tilde{U}_1) - \Psi(\tilde{U}_0)|_\Sigma \leq C_{ver}.$$

We conclude that Ψ is $\max\{C_{hor}, C_{ver}\}$ -Lipschitz, and the proof is complete. ■

Suppose now that f has no periodic critical points, and that the virtual endomorphism ϕ is contracting. Then Σ is hyperbolic; Theorem 6.1 implies that Γ is also hyperbolic. Hence if ϵ is chosen sufficiently small, the Floyd boundary $\partial\Gamma$ coincides with the Gromov boundary, and Φ induces a homeomorphism $\Phi : \mathcal{J} \rightarrow \partial\Gamma$ which conjugates the shift $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$ to the induced map $f_\Gamma : \partial\Gamma \rightarrow \partial\Gamma$. It follows that

$$\pi := \Phi^{-1} \circ \pi_\Gamma : S^2 \rightarrow \mathcal{J}$$

gives a semiconjugacy of branched coverings from $f : S^2 \rightarrow S^2$ to $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$. The remainder of the proof will use finiteness properties of the dynamical system $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$ to analyze the possibilities for the fibers of π .

We conclude this section with a technical result needed for the analysis in the next section.

By [HP2, Thm. 6.13], $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$ is a branched covering of degree d . Given $E \subset S^2$ a continuum and a component \tilde{E} of $f^{-1}(E)$, we define

$$\deg(f : \tilde{E} \rightarrow E) = \min\{\deg(f : \tilde{U} \rightarrow U) : E \subset U\}$$

where U is a connected open neighborhood of E .

Lemma 6.3 *Suppose $f_\Sigma(\tilde{\xi}) = \xi$, and set $E = \pi^{-1}(\xi)$, $\tilde{E} = \pi^{-1}(\tilde{\xi})$. Then $\deg(f : \tilde{E} \rightarrow E) = \deg(f_\Sigma, \tilde{\xi})$. In particular, the inverse image under f of a fiber of π is a disjoint union of fibers of π , and the degrees of f on these fibers sum to d .*

Proof: We first prove $\deg(f : \tilde{E} \rightarrow E) \geq \deg(f_\Sigma, \tilde{\xi})$. Let U be a small open connected neighborhood of E and \tilde{U} the pullback containing \tilde{E} . Since the decomposition of X by the fibers of π is upper-semicontinuous, and f_Σ is a branched covering, we may assume that the images $V = \pi(U)$, $\tilde{V} = \pi(\tilde{U})$ are neighborhoods of $\xi, \tilde{\xi}$, respectively, that $U = \pi^{-1}(V)$, $\tilde{U} = \pi^{-1}(\tilde{V})$, and that $\deg(f_\Sigma, \tilde{\xi}) = \deg(f_\Sigma : \tilde{V} \rightarrow V)$. Since $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$ is a branched covering, the set of branch values is nowhere dense, so there exists $\zeta \in V$ which is not a branch value of f_Σ . Hence $\#f_\Sigma^{-1}(\zeta) \cap \tilde{V} = \deg(f_\Sigma, \tilde{\xi})$. It follows easily that $\deg(f_\Sigma, \tilde{\xi})$ is a lower bound for $\deg(f : \tilde{U} \rightarrow U)$.

The proof of the inequality $\deg(f : \tilde{E} \rightarrow E) \leq \deg(f_\Sigma, \tilde{\xi})$ is similar. Suppose V, \tilde{V} are respectively neighborhoods of $\xi, \tilde{\xi}$ and set $U = \pi^{-1}(V)$, $\tilde{U} = \pi^{-1}(\tilde{V})$. Choose $\zeta \in V$ which is not a branch value of f_Σ . Then $\#f_\Sigma^{-1}(\zeta) = d$,

and since f is a branched covering of degree d , this implies that for each $z \in \pi^{-1}(\zeta)$, $\#f^{-1}(z) = d$ also. Hence no two elements of $f^{-1}(z)$ belong to the same fiber of π . We conclude $\#f_{\Sigma}^{-1}(\zeta) \cap \tilde{V} = \#f^{-1}(z) \cap \tilde{U}$ and so $\deg(f : \tilde{U} \rightarrow U)$ is a lower bound for $\deg(f_{\Sigma}, \tilde{\xi})$.

■

7 Proof of necessity

Proposition 7.1 *Let $f : S^2 \rightarrow S^2$ be a postcritically finite topological cxc map of the sphere. Then its virtual endomorphism is contracting.*

Let \mathcal{O} be the orbifold modeled on (S^2, P) , $\nu : \mathcal{O} \rightarrow \mathbb{N}$ its weight function, and as usual choose a basepoint $b \in S^2 \setminus P$. Taking an iterate if necessary, we may assume that $f(b) = b$. The group $G = \pi_1(\mathcal{O}, b)$ acts naturally as the deck transformation group of a Galois (regular) ramified universal orbifold covering $p : \tilde{\mathcal{O}} \rightarrow S^2$ such that $\deg(p, \tilde{x}) = \nu(p(\tilde{x}))$ for all $\tilde{x} \in \tilde{\mathcal{O}}$.

It follows that there exists a lift of “ f^{-1} ”, that is, a map $\tilde{f}_- : \tilde{\mathcal{O}} \rightarrow \tilde{\mathcal{O}}$ such that $f \circ p \circ \tilde{f}_- = p$. Since $f(b) = b$, we may choose \tilde{f}_- so that a given $\tilde{b} \in p^{-1}(b)$ is fixed by \tilde{f}_- :

$$\begin{array}{ccc} (\tilde{\mathcal{O}}, \tilde{b}) & \xleftarrow{\tilde{f}_-} & (\tilde{\mathcal{O}}, \tilde{b}) \\ p \downarrow & & \downarrow p \\ (\mathcal{O}, b) & \xrightarrow{f} & (\mathcal{O}, b) \end{array}$$

To see this: at lower-left, replace \mathcal{O} with \mathcal{O}_1 , so that the bottom arrow becomes an orbifold covering. Lift the orbifold structure on \mathcal{O}_1 at lower-left under p to obtain a new orbifold structure $\tilde{\mathcal{O}}_1$ on the underlying space of $\tilde{\mathcal{O}}$ so that $p : \tilde{\mathcal{O}}_1 \rightarrow \mathcal{O}_1$ is an orbifold covering. The map \tilde{f}_- is then a universal orbifold covering of $\tilde{\mathcal{O}}_1$.

In this paragraph, we define a metric \tilde{d} on \mathcal{O} with respect to which \tilde{f}_- becomes a uniform contraction. Let \mathcal{U}_0 be a covering as in §5, and endow S^2 with the metric d_{ε} described in §5. Choose $M > 0$ so small that (i) any d_{ε} -ball on S^2 of radius $2M$ is contained in some element of \mathcal{U}_0 ; from the choice of \mathcal{U}_0 it follows that (ii) any such ball meets P in at most one point. Let \mathcal{D} denote the set of Jordan domains U of diameter at most M . Condition (ii) implies that if $U \in \mathcal{D}$, then the inverse image $p^{-1}(U)$ is a collection $\{\tilde{U}_k\}$ of pairwise disjoint Jordan domains in $\tilde{\mathcal{O}}$ invariant under the covering group action of G (their closures might meet, but we don’t care), and the restriction of p to any domain \tilde{U}_k is proper. Let \tilde{U} be one such component and set, for $x, y \in \tilde{U}$,

$$q_{\tilde{U}}(x, y) = \inf \text{diam}_{\varepsilon} p(K)$$

where the infimum is over all continua $K \subset \overline{\tilde{U}}$ joining x to y . Now suppose x, y in $\tilde{\mathcal{O}}$. We define the metric \tilde{d} on $\tilde{\mathcal{O}}$ by

$$\tilde{d}(x, y) = \inf \sum_j q_{\tilde{U}_j}(x_j, x_{j+1})$$

where the infimum is taken over every chain $x = x_1, x_2, \dots, x_\ell = y$ joining x to y , with the property that for each j , there exists $U_j \in \mathcal{D}$ and a component \tilde{U}_j of $p^{-1}(U_j)$ such that $\{x_j, x_{j+1}\} \subset \tilde{U}_j$. We note that if $\tilde{d}(x, y) \leq M$, then we may find $\tilde{U} \in \mathcal{D}$ such that $\tilde{d}(x, y) = q_{\tilde{U}}(x, y)$. It follows that the infimum in the definition of \tilde{d} is actually attained.

The function \tilde{d} defines a metric and the group G acts by isometries on $(\tilde{\mathcal{O}}, \tilde{d})$.

Lemma 7.2 *The map $\Phi : g \mapsto g(\tilde{b})$ defines a quasi-isometry from G to $(\tilde{\mathcal{O}}, \tilde{d})$.*

Proof: We already know that G is finitely generated and that it acts on $\tilde{\mathcal{O}}$ by isometries, properly discontinuously and cocompactly.

We mimic the proof of the Svarc-Milnor lemma. Let $C \subset \tilde{\mathcal{O}}$ be a compact set containing \tilde{b} such that $G.C = \tilde{\mathcal{O}}$ and fix $D \geq 8M$ such that $C \subset B(\tilde{b}, D/4)$ and let $S = \{g \in G, \tilde{d}(g(\tilde{b}), \tilde{b}) \leq D\}$.

Then $\Phi(G)$ is cobounded, S is a finite generating set and there are constants $\lambda, c > 0$ such that $\tilde{d}(\Phi(g_1), \Phi(g_2)) \leq \lambda \|g_1 g_2^{-1}\|_S + c$ for all $g_1, g_2 \in G$, see [BH, Prop. 8.19]. The only point to prove is the reverse inequality.

Fix $g \in G$ and let us consider x_0, \dots, x_n with $x_0 = \tilde{b}$ and $x_n = g(\tilde{b})$ and let $\tilde{U}_1, \dots, \tilde{U}_n$ be such that

$$\tilde{d}(\tilde{b}, g(\tilde{b})) = \sum q_{\tilde{U}_j}(x_{j-1}, x_j).$$

Let N be the integer part of $4\tilde{d}(\tilde{b}, g(\tilde{b}))/D$. For any $1 \leq k \leq N$, let y_k be the first x_i such that

$$\sum_{j=1}^i q_{\tilde{U}_j}(x_{j-1}, x_j) \geq kD/4.$$

It follows that $\tilde{d}(y_k, y_{k+1}) \leq D/2$ since the points (x_j) are at most M -separated and $M \leq D/8$. Choose now $g_k \in G$ such that $\tilde{d}(g_k(\tilde{b}), y_k) \leq D/4$; then $\tilde{d}(g_k(\tilde{b}), g_{k+1}(\tilde{b})) \leq D$, so that $g_k \circ g_{k+1}^{-1} \in S$ and $\|g\|_S \leq N + 1$. Therefore, $\|g\|_S \leq (4/D)(\tilde{d}(\tilde{b}, g(\tilde{b})) + 1) + 1$. The proof is complete. ■

Proof: (Proposition) Let $x, y \in \tilde{\mathcal{O}}$ be such that $\tilde{d}(x, y) \leq M$, then we may find a continuum K containing $\{x, y\}$ such that $\tilde{d}(x, y) = \text{diam}_\varepsilon p(K)$. Note that $p(K)$ is contained in an element of \mathcal{U}_0 . Therefore, for any $n \geq 0$, $(p \circ \tilde{f}^n)(K)$ is contained in an element of \mathcal{U}_n and so $\tilde{d}(\tilde{f}^n(x), \tilde{f}^n(y)) \lesssim e^{-\varepsilon n}$. It follows easily that, for any pair of points x, y without any restriction on their distance, we have

$$\tilde{d}(\tilde{f}^n(x), \tilde{f}^n(y)) \lesssim e^{-\varepsilon n} \tilde{d}(x, y).$$

This implies that

$$\limsup_{n \rightarrow \infty} \left(\sup_{x, y \in \tilde{\mathcal{O}}} \frac{\tilde{d}(\tilde{f}^n(x), \tilde{f}^n(y))}{\tilde{d}(x, y)} \right)^{1/n} = e^{-\varepsilon} < 1.$$

Now, since $f(b) = b$ is fixed, we may define the virtual endomorphism ϕ by choosing the connecting path α (as in the definition given in §3.1) to be the constant path. Hence, $g \in \text{dom} \phi^n$ if and only if $\tilde{f}^n(g(\tilde{b}))$ belongs to the orbit $G.\tilde{b}$, and in this case, $\Phi(\phi^n(g)) = \tilde{f}^n(g(\tilde{b}))$. From the preceding Lemma 7.2, the metrics $\|\cdot\|$ and \tilde{d} are comparable, so we obtain

$$\limsup_{n \rightarrow \infty} \left(\limsup_{\|g\| \rightarrow \infty} \frac{\|\phi^n(g)\|}{\|g\|} \right)^{1/n} \leq \limsup_{n \rightarrow \infty} \left(\sup_{x, y \in \tilde{\mathcal{O}}} \frac{\tilde{d}(\tilde{f}^n(x), \tilde{f}^n(y))}{\tilde{d}(x, y)} \right)^{1/n} < 1.$$

■

8 Proof of sufficiency

From the introduction, recall that the key step is to establish

Proposition 8.1 *Let $f : S^2 \rightarrow S^2$ be a Thurston map with postcritical set P and with hyperbolic biset \mathfrak{M} . The semiconjugacy $\pi : S^2 \rightarrow \mathcal{J}$ is injective on P , and no fiber of π separates S^2 .*

The main ingredients in the proof of the claim are [HP2, Lemma 6.11, Theorem 6.15], which assert that the dynamical system $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$ is expanding, irreducible, and of *topologically finite type*. Finite type implies that there is a positive integer p such that for all $\xi \in \mathcal{J}$ and all $n > 0$, $\deg(f_\Sigma^n, \xi) \leq p$. Thus in particular f_Σ has no periodic branch points. Irreducible in this context is equivalent to the fact that given any nonempty open set W , there is an integer n with $f_\Sigma^n(W) = \mathcal{J}$.

We remark that these properties imply that $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$, as a topological dynamical system, is a so-called finite type topologically coarse expanding conformal (exc) system; cf. [HP1]. This forces strong constraints on the fibers of π . We prove the claim by a case-by-case analysis of the possibilities. We remark that there is much overlap with the flavor of the arguments in [PL].

A continuum $E \subset S^2$ we call

- *special* if $E \cap P \neq \emptyset$;
- *type III* if $S^2 - E$ has at least 3 components that intersect P
- *type II* if $S^2 - E$ has exactly 2 components that intersect P
- *type I* if $S^2 - E$ has exactly 1 component that intersects P

The map f sends fibers of π onto fibers of π . Suppose $E = \pi^{-1}(\xi)$ for some $\xi \in \mathcal{J}$.

- (i) If E is special, then $f^k(E)$ is special for all $k \geq 0$, and since P is finite, E is eventually periodic. Moreover, there are only finitely many such fibers.
- (ii) It is easily verified that if $f(E)$ is not special, then neither is E , and the type of $f(E)$ is at least the type of E .
- (iii) If E is type II and non-special then it determines a homotopy class of unoriented simple closed curve in $S^2 - P$.
- (iv) Since distinct fibers are disjoint, there can be only finitely many type III fibers, and only finitely many homotopy classes of curves arise from type II fibers.
- (v) If U is a component of $S^2 \setminus E$, then $\pi(U)$ is a component of $\mathcal{J} \setminus \{\xi\}$ and is open. This follows from the facts that π is monotone, continuous and onto and that S^2 is locally connected.

The proof of the following technical lemma is straightforward.

Lemma 8.2 *Given a continuum $E \subset S^2$, there is a “thickening” \widehat{E} of E given by an open regular neighborhood of E with the following properties.*

- *the closure of \widehat{E} is a compact surface with boundary.*
- *\widehat{E} has the same type as E .*
- *$\widehat{E} \cap P = E \cap P$.*
- *for each component \widetilde{E} of $f^{-1}(E)$, we have $\deg(f : \widetilde{E} \rightarrow \widehat{E}) = \deg(f : \widetilde{E} \rightarrow E)$, where \widetilde{E} is the component of $f^{-1}(\widehat{E})$ containing \widetilde{E} .*

In what follows, given a fiber $E = \pi^{-1}(\xi)$, the symbol \widehat{E} denotes a thickening of E satisfying the conditions of Lemma 8.2. The following paragraphs rule out certain types of fibers.

Lemma 8.3 (no periodic separating fibers) *No periodic E separates S^2 , and there are no type II fiber homotopic rel P to one of its iterates. Moreover, if E is periodic and special, then $\#E \cap P = 1$ and $E \cap P$ is periodic of the same period as E .*

Proof: Taking an iterate if necessary, we may assume that E is fixed. Since f_Σ has no periodic branch points, $\deg(f_\Gamma, \xi) = \deg(f : E \rightarrow E) = 1$, by Lemma 6.3. Then E is contained in a thickening \widehat{E} for which $\deg(\widehat{E} \rightarrow \widehat{E}) = 1$. This implies $f|_E$ is a homeomorphism which extends to a homeomorphism on a neighborhood of E .

1. **Subcase E special, $\#E \cap P \geq 2$.** Since $f : E \rightarrow E$ is a homeomorphism and P is forward invariant and finite, we may assume that $P \cap E$ is pointwise fixed. Note that $\widehat{E} \cap f(\widehat{E})$ is a neighborhood of E . Let $\gamma \subset f(\widehat{E})$ join two distinct point of $E \cap P$. Then $f^{-1} \circ \gamma$ joins the same points and is homotopic rel. P to a curve $\gamma_1 \subset f(\widehat{E})$ since $f|_E$ is 1-1. This implies the existence of an arc with arbitrarily long nontrivial iterates under the pullback relation, which is impossible by Theorem 4.1.
2. **Subcase E special and type II or type III.** Let us first establish that the boundary C of some complementary component W of \widehat{E} must be essential and nonperipheral in $S^2 - P$.
If each such W contains at most one point of P then $f^{-1}(\widehat{E})$ is connected. Thus $\deg(f : f^{-1}(\widehat{E}) \rightarrow \widehat{E}) = d > 1$ holds, and this contradicts the fact that $\deg(f : E \rightarrow E) = 1$. Hence at least one W contains two points of P . If E is type III or type II and special, then C is essential and nonperipheral.
We first note that $f^k : C \rightarrow f^k(C)$ has degree 1 up to homotopy rel. P for all $k \geq 1$. Moreover, $f^k(C)$ is also essential and nonperipheral for otherwise, C would not be. Finally, the forward orbit of C has to cycle eventually up to homotopy since there are only finitely homotopic classes of fibers of type at least II. This means that a Levy cycle can be extracted from $\{f^k(C)\}_{k \geq 0}$, contradicting Theorem 4.1.
3. **Subcase E non special and type II or type I.** Let us assume that E separates S^2 . There ia a component W of $S^2 \setminus E$ which contains all of P but at most one single point. Let \widehat{W} be the component of $S^2 \setminus \widehat{E}$ in W . Then the complement $\widehat{U} = S^2 \setminus \widehat{W}$ is a Jordan domain containing E and

at most one critical value of f . Therefore, $f^{-1}(\widehat{U})$ is a union of Jordan domains. One of them contains E , hence $U = S^2 \setminus W$. It follows that $f(U) \subset U$. But U contains a component of $S^2 \setminus E$, so $\pi(U)$ has interior by (v) and $f_\Sigma(\pi(U)) \subset \pi(U)$. which contradicts the irreducibility of f_Σ .

4. **Subcase type II E periodic up to homotopy.** We assume that E is type II and $f^k(E)$ is homotopic to E rel. P . If $\deg(f^k, E) > 1$ then $\deg(f^{nk}, E) = \deg(f_\Sigma^{nk}, \xi) \rightarrow \infty$ as $n \rightarrow \infty$, violating f_Σ being finite type. So $f^k|_E$ is injective, and after thickening, we extract a Levy cycle as above. ■

No preperiodic separating fibers nor type II and type III fibers. Let E be a preperiodic separating fiber. If all its iterates are nonspecial, then we would obtain a periodic separating fiber by (ii) and (iv), which is impossible according to Lemma 8.3. Therefore, there is some $n \geq 1$ such that $E' = f^n(E)$ is special and periodic by (i). By Lemma 8.3 again, E' does not separate S^2 and intersects P at a single point, so the inclusion $U := S^2 - (P \cup E') \hookrightarrow S^2 - P$ is a homotopy equivalence. Therefore $f^{-n}(U)$ is homotopy equivalent to the connected set $S^2 - f^{-1}(P)$. Hence U is connected, E is nonseparating, and $\#E \cap P \leq 1$.

Let us now prove that any type II and type III fibers have to cycle up to homotopy, which will establish their nonexistence.. Since fibers cannot intersect, there can be only finitely many of them up to homotopy, cf. (iv). Let E be of type at least II. If its iterates are non special, then they remain of type at least II by (ii), and hence have to cycle up to homotopy by (iv). Otherwise, one iterate becomes special, hence is preperiodic as well.

In order to complete the proof, we must prove that there are no type I fiber. We have already ruled out preperiodic ones.

No wandering separating type I fiber. All the possible remaining separating type I fibers are nonspecial and hence form a totally invariant set of S^2 . We use the same notation as in Lemma 8.3: given such a fiber E , let $\widehat{W}_E \subset W_E$ contain P , $U_E = S^2 \setminus W_E$ and $\widehat{U} = S^2 \setminus \widehat{W}$. Therefore, $f^{-1}(\widehat{U}_{f(E)})$ is a union of Jordan domains disjoint from P . One of them contains E , hence U_E . It follows that $f(U_E) \subset U_{f(E)}$. This implies that no preimage of P can enter the sets U_E : this contradicts again the irreducibility of f_Σ . Therefore, no type I fiber separates the sphere.

Proof: (Proposition 8.1) It follows from the discussion above that fibers of π are connected, do not separate the sphere, and meet P in at most one point.

Proof: (Theorem 1.3) Suppose the virtual endomorphism ϕ is contracting and f has no periodic branch points. Necessarily $\#P \geq 3$. For technical reasons, it is convenient later to assume $\#P > 3$ (see the introduction for the case $\#P = 3$). ■

According to Theorem 6.1, we may consider defining data for Σ and Γ so that both graphs are quasi-isometric. Since ϕ is contracting, it follows that Γ is hyperbolic and that the quasi-isometry defines a homeomorphism $\Phi : \mathcal{J} \rightarrow \partial\Gamma$. The map $\pi = \Phi^{-1} \circ \pi_\Gamma : S^2 \rightarrow \mathcal{J}$ semiconjugates f to f_Σ . Moreover, Lemma 5.1 and Proposition 8.1 imply that π is monotone, injective on P and that its fibers do not separate S^2 . By Moore's Theorem [Dav, Theorem 25.1], the quotient \mathcal{J} is homeomorphic to the sphere. The induced map $f_\Sigma : \mathcal{J} \rightarrow \mathcal{J}$ is thus an expanding Thurston map.

It remains to construct a homotopy from f to an expanding map g . By [Dav, Theorems 13.4, 25.1], the decomposition \mathcal{G} of S^2 by the fibers of the semiconjugacy $\pi : S^2 \rightarrow \mathcal{J}$ has the property of being *strongly shrinkable*: there is a one-parameter family of continuous maps $h_t : S^2 \rightarrow S^2, t \in [0, 1]$, called a *pseudoisotopy*, such that $h_0 = \text{id}_{S^2}$, h_t is a homeomorphism for $t \in [0, 1]$, $h_t|_{P_f} = \text{id}_{P_f}$ for all t , and the fibers of h_1 and those of π coincide. The induced homeomorphism $S^2 \rightarrow \mathcal{J}$ conjugates f_Σ to an expanding Thurston map $g : S^2 \rightarrow S^2$, and the family of maps $f_t, t \in [0, 1]$ defined as the unique continuous solution of $h_t \circ f = f_t \circ h_t$, with $f_1 = g$, gives a homotopy through Thurston maps. The proof of the sufficiency in Theorem 1.3 is complete. ■

References

- [Bar] Laurent Bartholdi. Functionally recursive groups. GAP package, <http://www.uni-math.gwdg.de/laurent/FR/>, 2011.
- [BH] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [BM] Mario Bonk and Daniel Meyer. Expanding Thurston maps. <http://arxiv.org/abs/1009.3647>, 2010.
- [CFP] James W. Cannon, William J. Floyd and Walter R. Parry. Finite subdivision rules. *Conform. Geom. Dyn.* **5**(2001), 153–196 (electronic).

- [CFPP] James W. Cannon, William J. Floyd, Walter R. Parry, and Kevin Pilgrim. Subdivision rules and virtual endomorphisms. *Geom. Dedicata* **141**(2009), 181–195.
- [Dav] Robert J. Daverman. *Decompositions of manifolds*, volume 124 of *Pure and Applied Mathematics*. Academic Press Inc., Orlando, FL, 1986.
- [DH] Adrien Douady and John Hubbard. A Proof of Thurston’s Topological Characterization of Rational Functions. *Acta. Math.* **171**(1993), 263–297.
- [HP1] Peter Haïssinsky and Kevin M. Pilgrim. Coarse expanding conformal dynamics. *Astérisque* **325** (2009), viii+139 pp. (2010).
- [HP2] Peter Haïssinsky and Kevin M. Pilgrim. Finite type coarse expanding conformal dynamics. *Groups Geom. Dyn.* **5** (2011), 603–661.
- [Nek1] Volodymyr Nekrashevych. *Self-similar groups*, volume 117 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2005.
- [Nek2] Volodymyr Nekrashevych. Combinatorial models of expanding dynamical systems. <http://arxiv.org/pdf/0810.4936.pdf>.
- [PL] Kevin Pilgrim and Tan Lei. Rational maps with disconnected Julia set. *Astérisque* (2000), xiv, 349–384. *Géométrie complexe et systèmes dynamiques* (Orsay, 1995).
- [Pil] Kevin M. Pilgrim. Julia sets as Gromov boundaries following V. Nekrashevych. *Topology Proc.* **29**(2005), 293–316. Spring Topology and Dynamical Systems Conference.
- [Rees] Mary Rees. A partial description of parameter space of rational maps of degree two. I. *Acta Math.* **168**(1992), no. 1-2, 11–87.
- [Shu] Michael Shub. Endomorphisms of compact differentiable manifolds. *Amer. J. Math.* **91**(1969), 175–199.