SUBDIVISION RULES AND VIRTUAL ENDOMORPHISMS

J. W. CANNON, W. J. FLOYD, W. R. PARRY, AND K. M. PILGRIM

Abstract. Suppose $f : S^2 \to S^2$ is a postcritically finite branched covering without periodic branch points. If f is the subdivision map of a finite subdivision rule with mesh going to zero combinatorially, then the virtual endomorphism on the orbifold fundamental group associated to f is contracting. This is a first step in a program to clarify the relationships among various notions of expansion for noninvertible dynamical systems with branching behavior.

0. Introduction

Let $T^2 = \mathbb{R}^2/\mathbb{Z}^2$ denote the real two-dimensional torus, equipped with the Euclidean Riemannian metric ds^2 inherited from the usual metric on $\mathbb{R}^2$, and suppose $f : T^2 \to T^2$ is a continuous orientation-preserving covering map. It is well-known that a necessary and sufficient condition for f to be homotopic to a covering map $g : T^2 \to T^2$ which is expanding with respect to ds^2 is that the spectrum of the induced linear map $f_* : H_1(T^2, \mathbb{R}) \to H_1(T^2, \mathbb{R})$ lies outside the closed unit disk. Thus, there is a complete homotopy-theoretic invariant for detecting those homotopy classes of coverings which contain expanding maps.

In this note, we take a first step toward a similar detection result for certain branched self-coverings of the 2-sphere to itself, called *Thurston maps*, which arise naturally in the classification of holomorphic dynamical systems in one complex variable [DH]. Our main result asserts that for certain Thurston maps, if one form of combinatorial expansion property is satisfied, then so is another. It is one part in a program to clarify the relationships between various notions of expansion for Thurston maps.

Let S^2 denote the 2-sphere equipped with an orientation. An orientation-preserving branched covering map $f : S^2 \to S^2$ of degree $d \geq 2$ has, by the Riemann-Hurwitz formula, a set B_f of $2d-2$ *branch points*, counted with multiplicity. By a branch point, we mean a point at which the local degree $\deg(f, x)$ of f at x is strictly larger than one. We denote by f^n the n-fold composition of f with itself. If the *postcritical set*

$$P_f = \bigcup_{n>0} f^n(B_f)$$

is finite, we call f a *Thurston map*. Two Thurston maps f, g are called *equivalent* provided there are orientation-preserving homeomorphisms $h_0, h_1 : (S^2, P_f) \to (S^2, P_g)$ such that $h_0 \circ f = g \circ h_1$ and h_0, h_1 are homotopic through homeomorphisms fixing P_f. The condition of being equivalent is a homotopy-theoretic one. Indeed, Nekrashevych [Nek, Theorem 6.5.2] has shown that checking equivalence can, via fundamental group

Date: August 22, 2008.

2000 *Mathematics Subject Classification.* Primary 37F10, 52C20; Secondary 20E08.

Key words and phrases. finite subdivision rule, Thurston map, virtual endomorphism.

considerations, be reduced to checking an algebraic condition. This condition can be phrased in terms of an algebraic invariant, the *virtual endomorphism* ϕ_f of the *orbifold fundamental group* associated to f. The precise definition will be given in §2.

When a Thurston map $g : S^2 \to S^2$ is expanding with respect to some complete length structure, then a very general result of Nekrashevych shows that the virtual endomorphism ϕ_g satisfies a homotopy-theoretic version of expansion which, since it is constructed by considering inverse images of f, is naturally called *contraction* [Nek, Theorem 5.5.3]. For example, any rational function $g : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ from the Riemann sphere to itself which is also a Thurston map will have the property that the associated virtual endomorphism ϕ_g is contracting. For an arbitrary Thurston map f, if ϕ_f is contracting, the limit space construction of Nekrashevych [Nek, §5] provides a synthetic construction of a topological dynamical system, living on the boundary at infinity of a negatively curved one-complex, that is naturally associated to the equivalence class of f.

Finite subdivision rules provide a wealth of concrete examples of Thurston maps. They have been extensively studied since they provide insight into how metric conformal structures arise as limits of discrete structures; see [CFP1], [CFP2], [CFP3], [CFP4], [CFKP]. The definition of a finite subdivision rule, and of related concepts, is given in §1.

Our main result is the following.

Theorem 0.1 (Mesh going to zero implies contracting)**.** *Let $\mathcal{R}$ be a finite subdivision rule whose model subdivision complex $S_{\mathcal{R}}$ is the 2-sphere and whose subdivision map f is orientation-preserving. If $\mathcal{R}$ has bounded valence, and if the mesh of $\mathcal{R}$ approaches zero combinatorially, then the virtual endomorphism ϕ_f on the orbifold fundamental group is contracting.*

The subdivision map f in the statement of the theorem is a Thurston map. The condition "bounded valence" is equivalent to the condition that no branch point of f is periodic. The condition "mesh approaching zero combinatorially" is a combinatorial expansion condition. Unfortunately, we do not know how to show that if f has mesh going to zero combinatorially, then f is homotopic to a map which is expanding with respect to some complete length structure. Therefore, we cannot appeal to Nekrashevych's result [Nek, Theorem 5.5.3] in our proof. Instead, our proof proceeds along the same outline as his, but uses a combinatorial version of length structure in place of usual lengths.

The main result is useful, since in concrete examples checking the condition of mesh going to zero is much easier than checking the condition of having contracting virtual endomorphism. As we shall show, the former involves only local calculations which are independent of the degree of the map.

There are many examples of finite subdivision rules for which the valence is unbounded and for which the virtual endomorphism is contracting. However, for such maps, the condition of mesh going to zero combinatorially does not adequately describe the relevant expansion properties.

Outline. In §1 we define and discuss finite subdivision rules. In §2, we define the orbifold, orbifold fundamental group, and virtual endomorphism ϕ_f associated to a Thurston map f. We also reduce the problem of showing that ϕ_f is contracting to showing that a certain finiteness property holds. This property, which is phrased in terms of so-called *restrictions* of elements, is interpreted geometrically in §3. The proof of Theorem 0.1 is then reduced to checking a geometric condition, Proposition 3.3. In §4, the proof of Theorem 0.1 is completed. We conclude in §5 with a brief discussion of some examples and of complexity issues.

1. Finite subdivision rules

The definition and basic theory of finite subdivision rules are given more leisurely and thoroughly in [CFP1].

Definition 1.1. *A* polygonal disk *is a cell structure on the closed 2-disk D such that there is a single 2-cell, the 1-skeleton of D is the unit circle, and there are at least three vertices. A* finite subdivision rule $\mathcal{R}$ *consists of the following: i) a finite 2-complex $S_{\mathcal{R}}$ such that $S_{\mathcal{R}}$ is the union of its closed 2-cells and each closed 2-cell $s \in S_{\mathcal{R}}$ is the image of a polygonal disk $t = t_s$ by a continuous cellular map $\psi_t\colon\ t \to S_{\mathcal{R}}$ which restricts to a homeomorphism on each open cell of t; ii) a subdivision $\mathcal{R}(S_{\mathcal{R}})$ of $S_{\mathcal{R}}$; iii) a continuous cellular map $\sigma_{\mathcal{R}}\colon\ \mathcal{R}(S_{\mathcal{R}}) \to S_{\mathcal{R}}$ whose restriction to each open cell is a homeomorphism. If $\mathcal{R}$ is a finite subdivision rule, an $\mathcal{R}$-*complex *is a CW-complex X which is the union of its closed 2-cells together with a continuous cellular map $\psi\colon\ X \to S_{\mathcal{R}}$, called the* structure map*, which restricts to a homeomorphism on each open cell.*

Suppose $\mathcal{R}$ is a finite subdivision rule. The complex $S_{\mathcal{R}}$ is called the *model subdivision complex* of $\mathcal{R}$, and the map $\sigma_{\mathcal{R}}$ is called the *subdivision map*. If s is a closed 2-cell of $S_{\mathcal{R}}$, the associated polygonal disk t_s is called a *tile type*. If X is an $\mathcal{R}$-complex, then the subdivision $\mathcal{R}(S_{\mathcal{R}})$ pulls back under the structure map $\psi\colon\ X \to S_{\mathcal{R}}$ to a subdivision $\mathcal{R}(X)$. Furthermore, $\mathcal{R}(X)$ is an $\mathcal{R}$-complex with structure map $\sigma_{\mathcal{R}} \circ f$. By iterating this construction, we can recursively subdivide $\mathcal{R}$-complexes. In particular, we can recursively subdivide the tile types and the model subdivision complex.

Definition 1.2. *A finite subdivision rule $\mathcal{R}$ has* bounded valence *if there is a uniform upper bound on the valence of vertices in the complexes $\mathcal{R}^n(S_{\mathcal{R}})$.*

Definition 1.3. *A finite subdivision rule $\mathcal{R}$ has* mesh approaching zero combinatorially *if there is a positive integer n such that:*

(1) *every edge in $S_{\mathcal{R}}$ is properly subdivided in $\mathcal{R}^n(S_{\mathcal{R}})$;*
(2) *if t is a tile type and e, e' are disjoint edges of t, then no tile of $\mathcal{R}^n(t)$ contains an edge of the subdivision of e in $\mathcal{R}^n(t)$ and an edge of the subdivision of e' in $\mathcal{R}^n(t)$.*

Suppose that $\mathcal{R}$ is a finite subdivision rule such that the model subdivision complex $S_{\mathcal{R}}$ is the 2-sphere and the subdivision map $\sigma_{\mathcal{R}}$ is orientation preserving. Then $\sigma_{\mathcal{R}}$ is a branched covering and each postcritical point is a vertex of $S_{\mathcal{R}}$, so the postcritical set is finite and $\sigma_{\mathcal{R}}$ is a Thurston map. In this case we usually denote the subdivision map by f. The condition that $\mathcal{R}$ has bounded valence is then equivalent to the condition that f has no periodic branch points.

2. Virtual endomorphisms

This section summarizes constructions of V. Nekrashevych, specialized to the case of Thurston maps. We refer to [Nek] for details. We suppress any discussion of the elegant and natural general algebraic theory and instead give only the minimal presentation needed for our proofs.

Let $f : S^2 \to S^2$ be a Thurston map. For $x, y \in S^2$ let $\nu_0(y) = \mathrm{lcm}\{\deg(f^n, x) : f^n(x) = y\}$ 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 $\mathcal{O}_i$. 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 $f : U_1 \to U_0$ 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 $\pi_1(U_i, b_i)/N_i$.

Let $f_* : \pi_1(U_1, b_1) \to \pi_1(U_0, b_0)$ be the injective homomorphism induced by the covering $f : U_1 \to U_0$. Since f sends peripheral loops to peripheral loops, it follows from the definitions of the weight functions ν_i that $f_* : N_1 \to N_0$ 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 $\overline{f}_* : \pi_1(\mathcal{O}_1, b_1) \to \pi_1(\mathcal{O}_0, b_0)$. We denote the image group $\overline{f}_*(\pi_1(\mathcal{O}_1, b_1))$ by H.

Let $\alpha : [0, 1] \to U_0$ be a path joining b_1 to b_0 and $\alpha_* : \pi_1(U_0, b_1) \to \pi_1(U_0, b_0)$ 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 $\pi_1(\mathcal{O}_0, b_1) = \pi_1(U_0, b_1)/N_1'$. Again, the map α_* descends to a well-defined isomorphism $\overline{\alpha}_* : \pi_1(\mathcal{O}_0, b_1) \to \pi_1(\mathcal{O}_0, b_0)$.

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 $\iota_* : \pi_1(U_1, b_1) \to \pi_1(U_0, b_1)$ is surjective and sends N_1 to N_0'. It easily follows that the map ι_* also descends to a surjective map $\overline{\iota}_* : \pi_1(\mathcal{O}_1, b_1) \to \pi_1(\mathcal{O}_0, b_1)$.

Definition 2.1. *The* virtual endomorphism induced by f *is the homomorphism* $\phi : H \to \pi_1(\mathcal{O}_0, b_0)$ *defined by*

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

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

The virtual endomorphism depends on the choice of basepoint b_0, preimage b_1, and homotopy class of path α. Different choices yield virtual endomorphisms which differ by pre- and/or post-composition by inner automorphisms. For $n \geq 2$, the *nth iterate* ϕ^n is the homomorphism whose domain is defined inductively by

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

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

In what follows, we denote the orbifold fundamental group $\pi_1(\mathcal{O}_0, b_0)$ by G.

Suppose S is a finite generating set for G. We denote by $||g||$ the word length of g in the generators S, and we let $\phi = \phi_f$.

Definition 2.2. *The virtual endomorphism $\phi : H \to G$ is called* contracting *if the* contraction ratio

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

The contraction ratio of the virtual endomorphism ϕ is independent of the choices used in its construction. Hence, the property of being contracting is independent of such choices, and so one may speak meaningfully about whether the virtual endomorphism of a Thurston map is contracting. Moreover, Nekrashevych has shown that the property of having a contracting virtual endomorphism is preserved under equivalence.

The contracting property is easy to state. But in practice, it is often easier to verify a certain equivalent property. In order to formulate it, we require several more definitions.

Fix a bijection $\Lambda : A \to f^{-1}(b_0)$ where A is a finite alphabet. Since b_0 is chosen to be a nonsingular point, $\#A = \deg(f)$. For $a \in A$ choose an oriented path $\lambda_a : [0,1] \to U_0$ joining b_0 to $\Lambda(a)$. For $n \in \mathbb{N}$ let A^n denote the set of words of length n in the alphabet A; set $A^0 = \{\emptyset\}$ where $\emptyset$ is the empty word. Let $A^* = \cup_n A^n$. We denote by $|w|$ the length of a word w.

The restriction $f^n : f^{-n}(U_0) \to U_0$ is an unramified covering map for all $n \geq 1$; in particular, any path β in U_0 can be lifted under any iterate f^n of f.

Given a point $x \in U_0$, a point $\tilde{x} \in f^{-n}(x)$, and a path or loop β starting at x, we denote by $f^{-n}(\beta)[\tilde{x}]$ the lift of β under f^n based at $\tilde{x}$. By induction and path-lifting, there is a map

$$\Lambda : A^* \to \bigcup_n f^{-n}(b_0)$$

given by

$$\Lambda(aw) = \text{the endpoint of the path} f^{-n}(\lambda_a)[\Lambda(w)]$$

where $a \in A$ and $w \in A^n$ are arbitrary. Notice that

$$f(\Lambda(wa)) = \Lambda(w)$$

for all $w \in A^*$ and all $a \in A$, i.e. that the dynamics acts as the right-shift.

For each $n \in \mathbb{N}$, the map $A^n \to f^{-n}(b_0) \times \{n\}$ defined by $w \mapsto \Lambda(w) \times \{n\}$, is a bijection. We obtain therefore a bijection $A^* \to \cup_n f^{-n}(b_0) \times \{n\}$. For each $n \in \mathbb{N}$, the projection onto the first coordinate gives an injection

$$f^{-n}(b_0) \times \{n\} \to f^{-n}(U_0).$$

Since the restriction $f^n : f^{-n}(U_0) \to U_0$ is an unramified covering, the fundamental group $\pi_1(U_0, b_0)$ acts by path-lifting on the fiber $f^{-n}(b_0)$ and hence, by means of the bijection constructed above, on the set A^n of words of length n. We obtain in this way an action of $\pi_1(U_0, b_0)$ on the set A^* of words of arbitrary finite length which preserves the length of a word and which acts transitively on words of a given fixed length.

We now show that this action descends to an action of G. Let γ^k be a closed loop in U_0 about some point $y \in \Sigma_0$ representing a generator g^k of the normal subgroup N_0 of $\pi_1(U_0, b_0)$ as constructed above. By the definition of the weight function ν_0, if $f^n(x) = y$ for some point $x \in \mathcal{O}_1$ and some integer $n \geq 1$, then $k = \nu_0(y)$ is a multiple of the local degree of f^n at x, and so γ^k lifts under f^n to a closed loop surrounding x. It follows that the path-lifting action of the subgroup N_0 on A^* is trivial, and hence that the action descends to a well-defined action of G on A^*. We denote by $g.w$ the image of w under the action of g.

Given $w = a_n \cdots a_2 a_1 \in A^*$ we denote by λ_w the path in U_0 starting at b_0 and given by

$$\lambda_w = \tilde{\lambda}_{a_1} * \tilde{\lambda}_{a_2} * \cdots * \tilde{\lambda}_{a_n}$$

where $\tilde{\lambda}_{a_1} = \lambda_{a_1}$ is traversed first, and where for $2 \leq i \leq n$ the path $\tilde{\lambda}_{a_i} = f^{-(i-1)}(\lambda_{a_i})[\Lambda(a_1 a_2 \cdots a_{i-1})]$. See Figure 1.

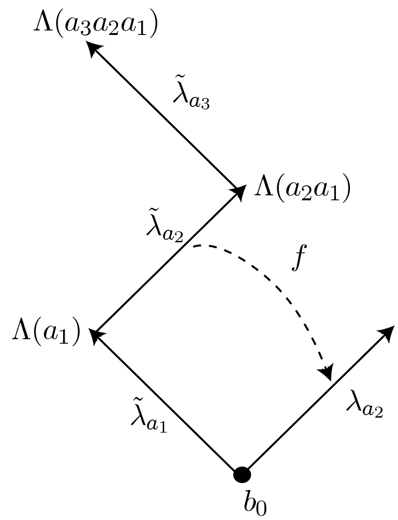

FIGURE 1. The dashed arrow shows $f(\tilde{\lambda}_{a_2}) = \lambda_{a_2}$.

Given $w \in A^n$ and $g \in G$, the *restriction* of g at w, denoted $g|_w$, is the element of G defined as follows. Represent g by a loop γ based at b_0, and let $g|_w$ denote the element of $G = \pi_1(U_0, b_0)/N_0$ represented by the path

$$\lambda_w * f^{-n}(\gamma)[\Lambda(w)] * \lambda_{g.w}^{-1}$$

where λ_w is traversed first and the exponent -1 in $\lambda_{g.w}^{-1}$ indicates that the path is traversed in the opposite direction. See Figure 2. The resulting element is well-defined independent of the choice of representative γ for g.

We are now ready to state our reformulation of the contracting property.

Proposition 2.3 ([Nek, Prop. 2.11.11])**.** *The virtual endomorphism ϕ is contracting if and only if there is a finite set $\mathcal{N} \subset G$ with the following property: for every integer $L \geq 1$, there exists an integer "magic level" $m(L)$ such that for all $g \in G$ with $||g|| \leq L$, and for all $w \in A^*$ with $|w| \geq m(L)$, the restriction $g|_w$ belongs to $\mathcal{N}$.*

3. The geometry of restriction

In this section, we reformulate the criterion for contracting spelled out in Proposition 2.3 in geometric terms.

Let f be the subdivision Thurston map arising from a finite subdivision rule $\mathcal{R}$. We assume that the basepoint b_0, bijection Λ, arcs $\lambda_a, a \in A$, and a generating set for G

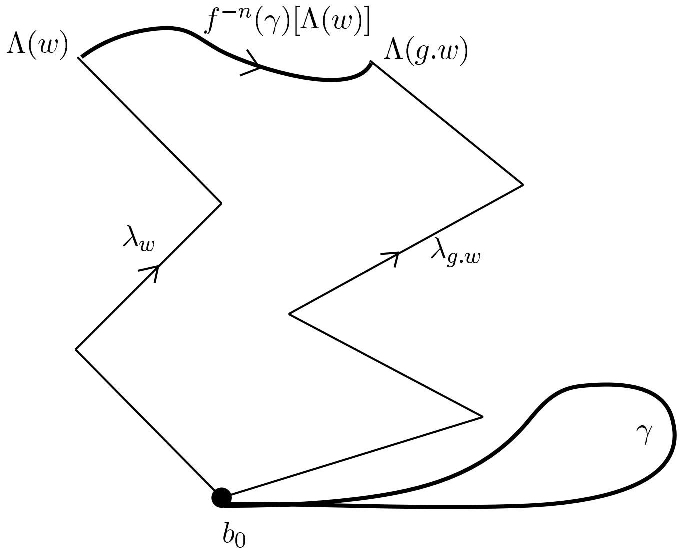

FIGURE 2. The restriction $g|_w$ is represented by $\lambda_w * f^{-n}(\gamma)[\Lambda(w)] * \lambda_{g.w}^{-1}$.

have been chosen as in the previous section. We drop the subscript 0 in the definition of the previous section and refer to the base orbifold, singular set, punctured surface, basepoint, and normal subgroup as $\mathcal{O}, \Sigma, U, b$, and N, respectively.

The *universal orbifold covering* $\pi : \widetilde{\mathcal{O}} \to \mathcal{O}$ is defined as follows. Let $\pi : \widetilde{U} \to U$ denote the covering corresponding to the subgroup N. As a topological space, $\widetilde{U}$ is homeomorphic to a plane (or disk) punctured at a countably (and possibly empty) infinite discrete set of points. A small peripheral loop about one of these punctures maps under π in a k to 1 fashion to a peripheral loop surrounding a singular point x of $\mathcal{O}$ of finite weight $k = \nu_0(x)$. By filling in the punctures, π extends to a continuous map. Doing this for all punctures yields an *orbifold universal covering map* $\pi : \widetilde{\mathcal{O}} \to \mathcal{O}$.

The group G acts freely and properly discontinuously as the group of covering transformations of the covering map $\pi : \widetilde{U} \to U$. The latter means that $g : \widetilde{U} \to \widetilde{U}$ is a homeomorphism satisfying $\pi \circ g = \pi$ for all $g \in G$. Upon filling in the punctures, each covering transformation g extends continuously to a homeomorphism $g : \widetilde{\mathcal{O}} \to \widetilde{\mathcal{O}}$ which satisfies $\pi \circ g = \pi$ where now $\pi : \widetilde{\mathcal{O}} \to \mathcal{O}$ is the orbifold universal covering. The resulting action of G on $\widetilde{\mathcal{O}}$ is again properly discontinuous. Furthermore, it acts transitively on the fibers of π. If there are singular points of finite weight, the action of G on $\widetilde{\mathcal{O}}$ is not free. For $g \in G$ and $\tilde{x} \in \widetilde{\mathcal{O}}$ we denote by $g.\tilde{x}$ the image of $\tilde{x}$ under the action of g.

Fix now a universal orbifold covering map $\pi : \widetilde{\mathcal{O}} \to \mathcal{O}$, a basepoint $b \in U$, and a preimage $\tilde{b} \in \pi^{-1}(b)$.

In the remainder of this section, we assume $\mathcal{R}$ has bounded valence. This is the case if and only if the orbifold $\mathcal{O}$ has no singular points of infinite weight, i.e. if and only if f has no periodic branch points. The universal orbifold covering $\pi : \widetilde{\mathcal{O}} \to \mathcal{O}$ then gives the underlying space $\widetilde{\mathcal{O}}$ the structure of an $\mathcal{R}$-complex with structure map

π. Therefore, we may speak meaningfully of edges, tiles, etc. at level n in $\widetilde{\mathcal{O}}$. Such tiles are compact subsets of $\widetilde{\mathcal{O}}$.

Definition 3.1 (Skinny path pseudometric)**.** *Suppose $\tilde{x}, \tilde{y} \in \widetilde{\mathcal{O}}$ and $n \in \mathbb{N}$. The* skinny path pseudodistance $d_n(\tilde{x}, \tilde{y})$ *is defined to be the minimum number m such that there exists n-tiles $\tilde{t}_0, \tilde{t}_1, \ldots, \tilde{t}_m$ in $\widetilde{\mathcal{O}}$ such that $\tilde{x} \in \tilde{t}_0, \tilde{y} \in \tilde{t}_m$, and $\tilde{t}_{i-1} \cap \tilde{t}_i \neq \emptyset$ for $i \in \{1, \ldots, m\}$.*

The skinny path pseudometrics d_n are proper in the sense that given any n, any fixed $\tilde{x} \in \widetilde{\mathcal{O}}$, and any $r > 0$, the "closed ball" $\{\tilde{y} \in \widetilde{\mathcal{O}} : d_n(\tilde{x}, \tilde{y}) \leq r\}$ is compact. The fact that the action of G on $\widetilde{\mathcal{O}}$ is properly discontinuous then immediately implies the following fact.

Proposition 3.2. *Fix an integer $r \geq 1$. Then*

$$C(r) := \#\{g \in G : d_0(\tilde{b}, g.\tilde{b}) \leq r\} < \infty.$$

Combining Propositions 3.2 and 2.3, we obtain

Proposition 3.3. *The virtual endomorphism is contracting if and only if the following condition holds. There is an integer $C \geq 1$ (depending on the choice of paths λ_a, $a \in A$) such that for every integer $L \geq 1$, there exists an integer "magic level" $m(L)$ such that for all $g \in G$ with $||g|| \leq L$, and for all $w \in A^*$ with $|w| \geq m(L)$, the restriction $g|_w$ satisfies*

$$d_0(\tilde{b}, (g|_w).\tilde{b}) \leq C.$$

4. Proof of Theorem

We are now ready to prove Theorem 0.1. We assume that we are in the setup of the previous section. Given $g \in G$ and $w \in A^*$, represent g by a loop γ, and recall from the definition that the restriction $g|_w$ is represented by the path

$$\underbrace{\lambda_w}_{1st} * \underbrace{f^{-n}(\gamma)[\Lambda(w)]}_{2nd=\tilde{\gamma}_w} * \underbrace{\lambda_{g.w}^{-1}}_{3rd}$$

traversed in the order indicated. The 1st and 3rd paths, as well as the 2nd, which we denote by $\tilde{\gamma}_w$, can be lifted to the universal cover $\widetilde{\mathcal{O}}$. Given a path $\beta \subset U$ starting at a nonsingular point x and a preimage $\tilde{x}$ of x under π, we denote by $\beta.\tilde{x}$ the endpoint of the path obtained by lifting β under π starting at $\tilde{x}$. By the triangle inequality, and using the fact that by definition $\tilde{\gamma}_w$ joins $\lambda_w.\tilde{b}$ to $\lambda_{g.w}.\tilde{b}$,

$$d_0(\tilde{b}, (g|_w).\tilde{b}) \leq d_0(\tilde{b}, \lambda_w.\tilde{b}) + d_0(\lambda_w.\tilde{b}, \tilde{\gamma}_w.(\lambda_w.\tilde{b})) + d_0(\tilde{b}, \lambda_{g.w}.\tilde{b}).$$

By Proposition 3.3, the proof is finished once we establish the following two claims.

Claim 1. *There is an integer C_1, depending on the choice of arcs λ_a, $a \in A$, such that for every $w \in A^*$,*

$$d_0(\tilde{b}, \lambda_w.\tilde{b}) \leq C_1.$$

Informally speaking, Claim 1 says that the "combinatorial 0-length" of any path λ_w is uniformly bounded independent of w.

Claim 2. *There is an integer C_2 such that for every integer $L \geq 1$, there is an integer "magic level" $m(L)$ such that for all $g \in G$ with $||g|| \leq L$, and for all $w \in A^*$ with $|w| \geq m(L)$, there exists a loop γ representing g such that*

$$d_0(\lambda_w.\tilde{b}, \tilde{\gamma}_w.(\lambda_w.\tilde{b})) \leq C_2.$$

Claim 2 says that for such a loop representing an element of g, the "combinatorial 0-length" of an iterated preimage of this loop can be made uniformly small by using a suitably high iterate of f.

Let us make this precise. Let $\beta : [0,1] \to U$ be a path in U and $n \in \mathbb{N}$. The *combinatorial n-length* $l_n(\beta)$ of β is

$$\#\{n\text{-tiles in } \widetilde{\mathcal{O}} \text{ meeting } \tilde{\beta}\} - 1$$

where $\tilde{\beta}$ is a lift of β under π to $\widetilde{\mathcal{O}}$. Since the group of covering transformations G sends n-tiles to n-tiles and acts transitively on lifts of β, the definition is independent of the chosen lift $\tilde{\beta}$. Like ordinary lengths, combinatorial n-lengths are subadditive. That is, if $\beta = \beta_1 * \beta_2$ is the concatenation of paths β_1, β_2, then

$$l_n(\beta_1 * \beta_2) \leq l_n(\beta_1) + l_n(\beta_2). \tag{1}$$

Unlike ordinary lengths, however, equality need not hold. For example, suppose x is a singular point of $\mathcal{O}$ of weight k and the 0-cell of $S_{\mathcal{R}}$ determined by x lies on the boundary of l distinct 2-cells of $S_{\mathcal{R}}$. Let γ be a small loop surrounding x. Then $l_0(\gamma^{pk}) = kl - 1$, which is independent of p; here γ^{pk} denotes the (pk)-fold concatenation of γ with itself.

The remainder of this section is devoted to the proof of these two claims. For the proof of each claim, we will need the following key facts.

Proposition 4.1. *Suppose $\tilde{x}, \tilde{y} \in \widetilde{\mathcal{O}}$. Then there is a positive integer m with the following property. Let k be a nonnegative integer, and let $\tilde{x}, \tilde{y}$ be points in $\widetilde{\mathcal{O}}$ such that $d_k(\tilde{x}, \tilde{y}) \geq 2$. Then $d_{k+m}(\tilde{x}, \tilde{y}) \geq 2 \cdot d_k(\tilde{x}, \tilde{y})$.*

Proposition 4.1 follows from Lemma 2.7 of [CFP2].

Lemma 4.2. *Suppose $n \in \mathbb{N}$, β_n is a path in U, and β_{n+1} is a lift of β_n under f. Then $l_{n+1}(\beta_{n+1}) \leq l_n(\beta_n)$.*

Proof: Consider again the pair U_0, U_1 defined in Section 2. Let $a_0 = \beta_n(0)$ denote the starting point of the path β and $a_1 = \beta_{n+1}(0)$ be the starting point of the lift β_{n+1}. Denote by $\widetilde{U}_0 = \pi^{-1}(U_0)$. The set $\widetilde{U}_0$ typically is the complement in $\widetilde{\mathcal{O}}$ of a countably infinite set of points. Then $f : (U_1, a_1) \to (U_0, a_0)$ is a covering map sending a_1 to a_0. Let $\tilde{a}_0 \in \pi^{-1}(a_0)$ and $\tilde{a}_1 \in \pi^{-1}(a_1)$. Using a standard monodromy argument, it follows easily from the definitions of the weight functions and the covering π that there exists a lift q of f^{-1} making the diagram

$$\begin{array}{ccc} (\widetilde{U}_1, \tilde{a}_1) & \stackrel{q}{\longleftarrow} & (\widetilde{U}_0, \tilde{a}_0) \\ \pi \downarrow & & \downarrow \pi \\ (U_1, a_1) & \stackrel{f}{\rightarrow} & (U_0, a_0) \end{array}$$

commute. Denote by $\tilde{\beta}_n$ and $\tilde{\beta}_{n+1}$ the lifts under π of β_n and β_{n+1} starting at $\tilde{a}_0$ and $\tilde{a}_1$, respectively. Then $q(\tilde{\beta}_n) = \tilde{\beta}_{n+1}$. By definition, f takes $(n+1)$-tiles in $\mathcal{R}^{n+1}(S_{\mathcal{R}})$

to n-tiles in $\mathcal{R}^n(S_{\mathcal{R}})$. Hence the lift q of the inverse of f takes an n-tile in $\widetilde{\mathcal{O}}$ to an $(n+1)$-tile in $\widetilde{\mathcal{O}}$. It follows that $l_{n+1}(\beta_{n+1}) \le l_n(\beta_n)$ as required. $\square$

We remark that in the conclusion of Lemma 4.2, equality need not hold. For example, suppose that x is a critical point of f at which f has degree 2. Suppose that $\nu_0(x) = 1$, that $y = f(x)$, and that $\nu_0(y) = 6$. Choose $\tilde{y} \in \widetilde{\mathcal{O}}$ such that $\pi(\tilde{y}) = y$, and set $\tilde{x} = q(\tilde{y})$, so that $\pi(\tilde{x}) = x$. Then π has degree 6 at $\tilde{y}$ and degree 1 at $\tilde{x}$. Furthermore, q has degree 3 at $\tilde{y}$. So we would expect a small simple closed curve about $\tilde{y}$ to meet 3 times as many tiles as its image under q, which winds 3 times about $\tilde{x}$.

4.1. **Proof of Claim 1.** Let $C = \max\{l_0(\lambda_a) : a \in A\}$. By Lemma 4.2 and induction, for any $n \in \mathbb{N}$, any $a \in A$, and any lift $\tilde{\lambda}_a$ of λ_a under f^{-n}, we have $l_n(\tilde{\lambda}_a) \le C$ as well. By the definition of combinatorial length, it follows that for every nonnegative integer n, for every $a \in A$, and every $w \in A^n$, the lift of $f^{-n}(\lambda_a)$ to $\widetilde{\mathcal{O}}$ based at $\lambda_w.\widetilde{b}$ meets at most C tiles at level n. So for every nonnegative integer n we have that $d_n(\lambda_w.\widetilde{b}, \lambda_{wa}.\widetilde{b}) \le C$ for every $a \in A$ and $w \in A^n$.

Now let $w \in A^*$ as in the statement of Claim 1. We may assume that $|w| > 0$. Let $w = a_n \cdots a_2 a_1$. For $k \in \{0, \ldots, n\}$ let $w_k = a_k \cdots a_2 a_1$, so that $w_0 = \emptyset$ and $w_n = w$, and let $\tilde{b}_k = \lambda_{w_k}.\tilde{b}$, so that $\tilde{b}_0 = \tilde{b}$ and $\tilde{b}_n = \lambda_w.\tilde{b}$. According to the previous paragraph,

$$d_k(\tilde{b}_k, \tilde{b}_{k+1}) \le C \tag{2}$$

for every $k \in \{0, \ldots, n-1\}$.

Let m be as in Proposition 4.1. We will prove Claim 1 for $C_1 = 2mC$. Arguing by contradiction, we suppose that $d_0(\widetilde{b}, \lambda_w.\widetilde{b}) > C_1$. We will show that this assumption implies the following:

(*) for every nonnegative integer r for which $n \ge rm$, the inequality

$$d_{rm}(\tilde{b}_{rm}, \tilde{b}_n) > C_1 \tag{3}$$

is satisfied.

This is impossible for the following reason. Suppose r is a maximal such integer, so that $0 \le n - rm < m$. If $n - rm = 0$ then $\tilde{b}_{rm} = \tilde{b}_n$ and so $d_{rm}(\tilde{b}_{rm}, \tilde{b}_n) = 0 \not> C_1$. Otherwise, we have

$$\begin{aligned}
d_{rm}(\tilde{b}_{rm}, \tilde{b}_n) &\le \sum_{j=0}^{n-rm-1} d_{rm}(\tilde{b}_{rm+j}, \tilde{b}_{rm+j+1}) && \text{by } \Delta \text{ inequality} \\
&\le \sum_{j=0}^{n-rm-1} d_{rm+j}(\tilde{b}_{rm+j}, \tilde{b}_{rm+j+1}) && \text{since } d_{rm} \le d_{rm+j} \\
&\le \sum_{j=0}^{n-rm-1} C && \text{by (2)} \\
&\le mC && \text{since } r \text{ is maximal} \\
&< 2mC = C_1.
\end{aligned}$$

We now establish condition (*) by induction on the nonnegative integer r. The case $r = 0$ follows from the assumption, to the contrary of our desired conclusion, that $d_0(\widetilde{b}, \lambda_w.\widetilde{b}) = d_0(\tilde{b}_0, \tilde{b}_n) > C_1$.

As our inductive hypothesis, we assume that the inequality (3) holds for a nonnegative integer r with $rm < n$. We may assume that $(r+1)m \leq n$. By the triangle inequality,

$$d_{rm}(\tilde{b}_{rm}, b_n) \leq d_{rm}(\tilde{b}_{rm}, \tilde{b}_{(r+1)m}) + d_{rm}(\tilde{b}_{(r+1)m}, \tilde{b}_n).$$

Rewriting this inequality, we conclude

$$\begin{aligned}
d_{rm}(\tilde{b}_{(r+1)m}, \tilde{b}_n) &\geq d_{rm}(\tilde{b}_{rm}, b_n) - d_{rm}(\tilde{b}_{rm}, \tilde{b}_{(r+1)m}) & \\
&> C_1 - d_{rm}(\tilde{b}_{rm}, \tilde{b}_{(r+1)m}) & \text{by inductive hypothesis} \\
&\geq C_1 - mC & \text{by } \Delta \text{ inequality} \\
&\geq 2mC - mC = mC & \text{by def. of } C_1 \\
&> 1. &
\end{aligned}$$

Proposition 4.1 and the second-to-the-last inequality above imply, respectively, that

$$d_{(r+1)m}(\tilde{b}_{(r+1)m}, \tilde{b}_n) \geq 2 \cdot d_{rm}(\tilde{b}_{(r+1)m}, \tilde{b}_n) > 2 \cdot mC = C_1.$$

This completes the proof of the induction step, and so the proof of Claim 1 is complete. □

4.2. **Proof of Claim 2.** For the proof of Claim 2, we will need the following Lemma.

Lemma 4.3. *There is a positive integer B such that for all $g \in G$, there exists a representative γ of g for which*

$$l_0(\gamma) \leq B \cdot ||g||.$$

Proof: Suppose $\{g_1, g_2, \ldots, g_r\}$ is the chosen set of generators used in the definition of the word length $|| \cdot ||$. For each g_i, choose a representative $\gamma_i \subset U$. Let $B = \max\{l_0(\gamma_i) : 1 \leq i \leq r\}$. Now suppose $g \in G$ is arbitrary and $g = g_{i_1} g_{i_2} \ldots g_{i_k}$ where $k = ||g||$. Then $\gamma = \gamma_{i_1} * \gamma_{i_2} * \cdots * \gamma_{i_k}$ represents g and so by (1) we have $l_0(\gamma) \leq \sum_{j=1}^{k} l_0(\gamma_{i_j}) \leq B||g||$ as required. □

Proof of Claim 2. Suppose $||g|| \leq L$. By Lemma 4.3, g is represented by a path γ for which $l_0(\gamma) \leq BL$. By Lemma 4.2, for any $n \in \mathbb{N}$, the combinatorial n-length of any lift $\tilde{\gamma}_w$ of γ under f^{-n} satisfies $l_n(\tilde{\gamma}_w) \leq BL$. For convenience, given $k \in \mathbb{N}$ denote by $d_k = d_k(\lambda_w.\tilde{b}, \tilde{\gamma}_w.(\lambda_w.\tilde{b}))$. So by definition $d_n \leq BL$ for any $n \in \mathbb{N}$. Let m be as in Proposition 4.1 and let r be a positive integer with $r \geq \log_2(L)$. Let $m(L) = rm$. Then if $n \geq m(L)$ we have that

$$BL \geq d_n \geq d_{rm}.$$

If $d_0 \geq 2$ then Proposition 4.1 implies that in turn

$$d_{rm} \geq 2^r d_0 \geq L d_0$$

and so $d_0 \leq B$. Thus we may always take $C_2 = B$. □

5. Complexity issues

The theorem below says that in order to check the condition mesh going to zero combinatorially, it suffices to check finitely many local conditions. The proof is constructive and gives, for concrete examples, an algorithm for determining when conditions (1) and (2) hold in the definition of mesh going to zero combinatorially.

Theorem 5.1. *Suppose $\mathcal{R}$ is a finite subdivision rule. Let k be the number of tiles in $S_{\mathcal{R}}$, and let l be the maximum number of edges in a tile type of $S_{\mathcal{R}}$. Then $\mathcal{R}$ has mesh going to zero combinatorially if and only if conditions (1) and (2) in the definition are satisfied whenn $= kl^2$.*

Note that the quantity n depends only on $S_{\mathcal{R}}$ and not on $\mathcal{R}(S_{\mathcal{R}})$. In particular, the bound n is independent of the degree.

Proof: Sufficiency is trivial. So suppose $\mathcal{R}$ has mesh going to zero combinatorially.

Consider first condition (1) in the definition. Form a directed graph G as follows. (In what follows, a directed graph may have a loop from a vertex to itself, but multiple loops from a vertex to itself and multiple directed edges between two given vertices are not permitted.) As vertex set, take the disjoint union, over all tiles s of $S_{\mathcal{R}}$, of the 1-cells in the tile type t_s of s. Thus, there are at most kl vertices. Join a vertex corresponding to a 1-cell e_1 of a tile of type t_1 to a vertex corresponding to a 1-cell e_2 of a tile of type t_2 if and only if the subdivision $\mathcal{R}(t_1)$ contains a tile u of type t_2 such that the 1-cell of u corresponding to e_2 coincides with e_1. Less formally: vertices e_1 and e_2 are joined if and only if $\sigma_{\mathcal{R}}$ sends e_2 homeomorphically to e_1. A vertex of G corresponding to a 1-cell e has an outgoing edge if and only if it fails to be properly subdivided. By induction, given a positive integer n, there is a directed edge-path of length n starting at a vertex corresponding to a 1-cell e if and only if e fails to be properly subdivided after n subdivisions. If condition (1) fails when $n = kl^2$, then there is a directed edge-path of length greater than or equal to the number of vertices of G. It follows that some vertex of G is visited twice in this directed edge-path, and hence G must contain a directed cycle. Therefore, there exists a 1-cell which is never properly subdivided. This contradicts the assumption that condition (1) holds for some value of n.

Now consider condition (2). We argue similarly. Form a different directed graph G as follows. As vertices, we take triples (t, e_1, e_2), where t is a tile type, and e_1, e_2 are disjoint 1-cells in t. Thus there are at most kl^2 vertices of G. Two triples $(t, e_1, e_2), (t', e'_1, e'_2)$ are joined by a directed edge in G if and only if the subdivision $\mathcal{R}(t)$ contains a tile of type t' such that the 1-cell of t' corresponding to e'_1 is contained in e_1 and the 1-cell of t' corresponding to e'_2 is contained in e_2. As before, it follows that given a positive integer n, a triple (t, e_1, e_2) fails condition (2) for this value of n if and only if there is a directed edge-path of length n starting at this triple. The proof concludes as in the previous paragraph. □

We finish with two examples.

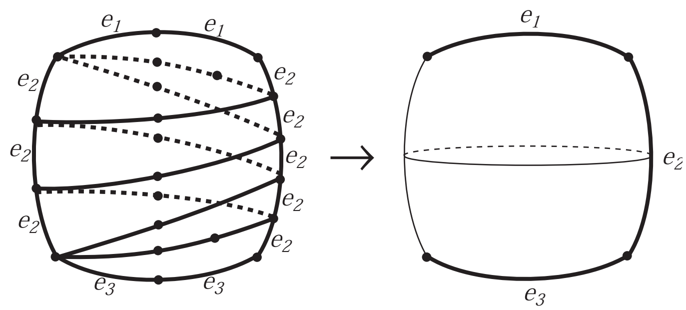

FIGURE 3. The branched map f.

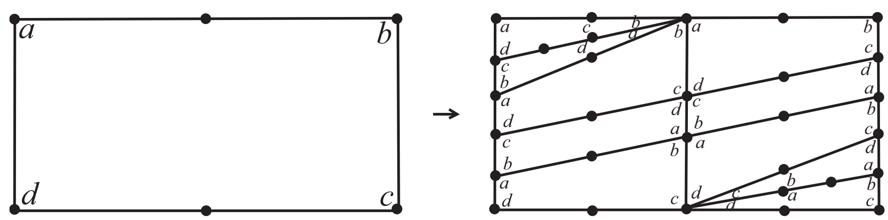

FIGURE 4. The subdivision of the tile type for $\mathcal{R}$.

Example 1. This example is analyzed in more detail in [CFPP]. Consider the branched map f shown in Figure 3. It is orientation-preserving and cellular with respect to the cell structures whose edges are drawn with bold arcs, and preserves edge labels (though for clarity not all edges are labelled). The map f is the subdivision map of a finite subdivision rule $\mathcal{R}$ with a single tile type; the tile type t and its subdvision $\mathcal{R}(t)$ are shown in Figure 4. Since none of the corners of t are properly subdivided in $\mathcal{R}(t)$, $\mathcal{R}$ has bounded valence. Since every edge of t is properly subdivided in $\mathcal{R}(t)$, the first graph constructed in the proof of Theorem 5.1 has no edges and so $\mathcal{R}$ satisfies condition (1). Similarly, since no tile of $\mathcal{R}(t)$ intersects two disjoint edges of t, the subdivision rule $\mathcal{R}$ satisfies condition (2). Hence $\mathcal{R}$ has mesh approaching zero combinatorially. By Theorem 0.1, the virtual endomorphism of f is contracting.

Example 2. Now consider the branched map g shown in Figure 5. The map g is orientation-preserving and cellular, and preserves edge labels. It is the subdivision map of a finite subdivision rule $\mathcal{Q}$ with two tile types. The subdivisions of the tile types are shown in Figure 6, where edges are labelled by the labels of their images in $S_{\mathcal{Q}}$. The two tile types t_1 and t_2 have the same subdivisions but correspond to the two possible orientations on a hexagon. Since none of the corners of t_1 and t_2 are properly subdivided in $\mathcal{Q}(t_1)$ and $\mathcal{Q}(t_2)$, $\mathcal{Q}$ has bounded valence.

The corresponding directed graph for checking condition (1) has twelve vertices and eight edges. The directed edges are

$$(t_1, b) \to (t_1, a), \ \ (t_1, c) \to (t_1, b), \ \ (t_1, e) \to (t_2, b), \ \ (t_1, f) \to (t_2, a)$$

$$(t_2, b) \to (t_2, a), \ \ (t_2, c) \to (t_2, b), \ \ (t_2, e) \to (t_1, b), \ \ (t_2, f) \to (t_1, a).$$

There are directed edge-paths of length 2 but not of length 3; that is why in Figure 7 a tile type t has edges that are not properly subdivided in $\mathcal{Q}^2(t)$ but every edge of t

is properly subdivded in $\mathcal{Q}^3(t)$. Therefore, there are no directed cycles, so condition (1) is satisfied.

It is also easy to construct the directed graph for checking condition (2). It has 36 vertices and 24 edges. Again, it turns out that there are directed edge-paths of length 2 (e.g., $((t_1, b, d), (t_1, a, c))$ followed by $((t_1, a, c), (t_1, f, b)))$, but no directed cycles.

Since $\mathcal{Q}$ satisfies conditions (1) and (2), it has mesh approaching zero combinatorially.

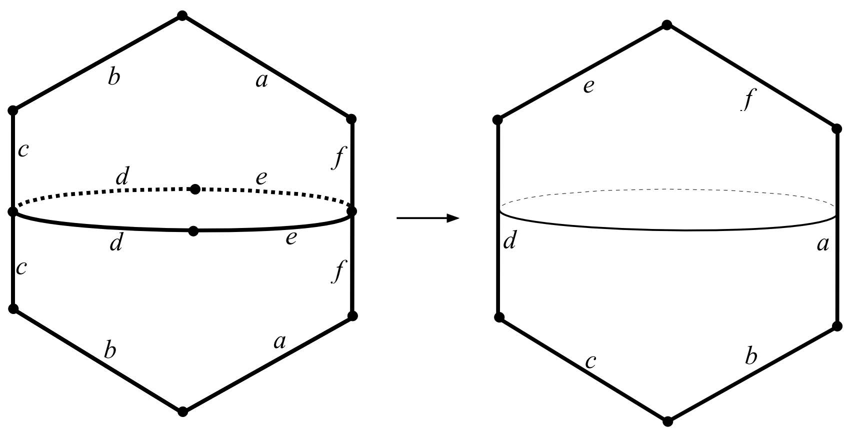

FIGURE 5. The branched map g.

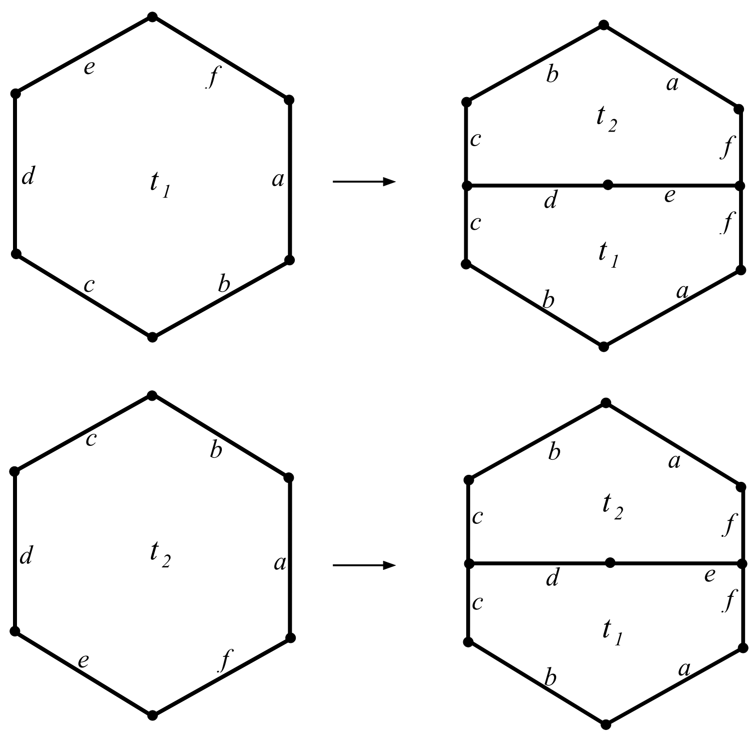

FIGURE 6. The subdivisions of the tile types for $\mathcal{Q}$.

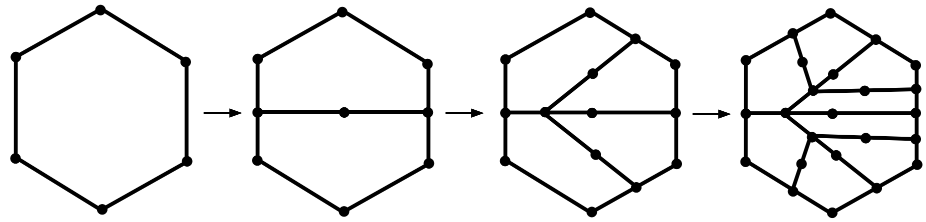

FIGURE 7. The first three subdvisions of a tile type for $\mathcal{Q}$.

References

[CFKP] 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.

[CFP1] J. W. Cannon, W. J. Floyd, and W. R. Parry. Finite subdivision rules. *Conform. Geom. Dyn.* **5**(2001), 153–196 (electronic).

[CFP2] J. W. Cannon, W. J. Floyd, and W. R. Parry. Expansion complexes for finite subdivision rules. I. *Conform. Geom. Dyn.* **10**(2006), 63–99 (electronic).

[CFP3] J. W. Cannon, W. J. Floyd, and W. R. Parry. Expansion complexes for finite subdivision rules. II. *Conform. Geom. Dyn.* **10**(2006), 326–354 (electronic).

[CFP4] J. W. Cannon, W. J. Floyd, and W. R. Parry. Constructing subdivision rules from rational maps. *Conform. Geom. Dyn.* **11**(2007), 128–136 (electronic).

[CFPP] J. W. Cannon, W. J. Floyd, W. R. Parry, and K. M. Pilgrim, A critically finite map, and the assoicated map on Teichmüller space, in preparation.

[DH] A. Douady and John Hubbard. A Proof of Thurston's Topological Characterization of Rational Functions. *Acta. Math.* **171**(1993), 263–297.

[Nek] Volodymyr Nekrashevych. *Self-similar groups*, volume 117 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2005.

Department of Mathematics, Brigham Young University, Provo, UT 84602, U.S.A.
E-mail address: cannon@math.byu.edu

Department of Mathematics, Virginia Tech, Blacksburg, VA 24061, U.S.A.
E-mail address: floyd@math.vt.edu
URL: http://www.math.vt.edu/people/floyd

Department of Mathematics, Eastern Michigan University, Ypsilanti, MI 48197, U.S.A.
E-mail address: walter.parry@emich.edu

Department of Mathematics, Indiana University, Bloomington, IN 47405, U.S.A.
E-mail address: pilgrim@indiana.edu