LATTÈS MAPS AND FINITE SUBDIVISION RULES

J. W. CANNON, W. J. FLOYD, AND W. R. PARRY

ABSTRACT. This paper is concerned with realizing Lattès maps as subdivision maps of finite subdivision rules. The main result is that the Lattès maps in all but finitely many analytic conjugacy classes can be realized as subdivision maps of finite subdivision rules with one tile type. An example is given of a Lattès map which is not the subdivision map of a finite subdivision rule with either (i) two tile types and 1-skeleton of the subdivision complex a circle or (ii) one tile type.

This paper is concerned with realizing rational maps by subdivision maps of finite subdivision rules. If $\mathcal{R}$ is an orientation-preserving finite subdivision rule such that the subdivision complex $S_\mathcal{R}$ is a 2-sphere, then the subdivision map $\sigma_\mathcal{R}$ is a postcritically finite branched map. Furthermore, $\mathcal{R}$ has bounded valence if and only if $\sigma_\mathcal{R}$ has no periodic critical points. In [1] and [4], Bonk-Meyer and Cannon-Floyd-Parry prove that if f is a postcritically finite rational map without periodic critical points, then every sufficiently large iterate of f is the subdivision map of a finite subdivision rule $\mathcal{R}$ with two tile types such that the 1-skeleton of $S_\mathcal{R}$ is a circle. Since finite subdivision rules are essentially combinatorial objects, this gives good combinatorial models for these iterates. It is especially convenient to realize a postcritically finite map by the subdivision map of a finite subdivision rule with either a single tile type or with two tile types and 1-skeleton of the subdivision complex a circle.

While passing to an iterate is not usually a serious obstacle, it would be preferable if one didn't need to do this. Suppose f is a postcritically finite rational map without periodic critical points. We consider the following questions:

(1) Is f the subdivision map of a finite subdivision rule?
(2) Is f the subdivision map of a finite subdivision rule with two tile types and 1-skeleton of the subdivision complex a circle?
(3) Is f the subdivision map of a finite subdivision rule with one tile type?

In this paper we consider these questions for Lattès maps. The main result of this paper is to answer question (3) in the affirmative for the maps in all but finitely many analytic conjugacy classes of Lattès maps. In addition, we exhibit a Lattès map of degree 2 for which the answer to questions (2) and (3) is negative, although the answer to question (1) is positive.

This paper has four sections. The first section develops the setting of Lattès maps. These results can be used to easily enumerate all Lattès maps of very small degree to test (not so easy) the above three questions for them.

Section 2 presents the pruning lemma and immediate consequences. To describe the pruning lemma, let f be a Lattès map which is the subdivision map of a finite subdivision rule with one tile type. This obtains a subdivision complex structure on the Riemann sphere whose 1-skeleton is a tree. The postcritical points of f are vertices of this tree. The pruning lemma states that it is possible to prune this tree to obtain a subtree for which every vertex with valence either 1 or 2 is a postcritical point and this subtree is the 1-skeleton of a subdivision complex for which f is the subdivision map. So if f is the subdivision map for a subdivision complex whose 1-skeleton is a tree, then f is the subdivision map for a subdivision complex whose 1-skeleton is a tree with a special form. This special form is very restrictive.

Section 3 contains the main result in Corollary 3.11, namely, that every map in almost every analytic conjugacy class of Lattès maps is the subdivision map of a finite subdivision rule with one tile type. This is essentially proven in two theorems. The first theorem, Theorem 3.1, states that every Lattès map with sufficiently large degree is the subdivision map of a finite subdivision rule with one tile type. The second theorem, Theorem 3.10, states that every nonrigid Lattès map is the subdivision map of a finite subdivision rule with one tile type. Since there are only finitely many analytic conjugacy classes of rigid Lattès maps with a given degree (see the end of Section 1), these two theorems give the main result. Theorems 3.1 and 3.10 actually say a bit more, giving additional information about the tilings.

The proof of Theorem 3.1 is much more difficult than the proof of Theorem 3.10. It roughly parallels the proof of the main theorem of [4]. The strategy is to construct a finite subdivision rule $\mathcal{R}$ with one tile type, bounded valence, and mesh approaching 0 combinatorially such that the subdivision map $\sigma_{\mathcal{R}}$ of $\mathcal{R}$ is isotopic to the Lattès map f relative to its postcritical set P_f. Once this is achieved, our results for expansion complexes [3] can be used to show that $\sigma_{\mathcal{R}}$ and f are in fact conjugate rel P_f, and so f is the subdivision map of a finite subdivision rule with one tile type. The finite subdivision rule $\mathcal{R}$ is constructed using a lift of f to the complex plane and a standard tiling of the plane by regular hexagons. Our tiling of the Riemann sphere by one tile lifts to a tiling of the plane which is combinatorially equivalent to this standard tiling by regular hexagons. The proof of Theorem 3.10 does not require an isotopy and hence does not require results for expansion complexes. In this case our tiling of the Riemann sphere by one tile lifts to a standard tiling of the plane by parallelograms.

Section 4 presents an example which shows that the main result does not hold for every Lattès map. This example is a quadratic Lattès map f for which both question (2) and question (3) are false. Question (1) is true for f, as we show that f is the subdivision map of a finite subdivision rule with two tile types, but the 1-skeleton of the subdivision complex is not a circle. The Lattès map f has a lift to the complex plane of the form $\widetilde{f}(z) = \alpha z$, where $\alpha = (1 + \sqrt{-7})/2$ and the postcritical set of f lifts to the lattice generated by 1 and α. The complex conjugate of f gives another such example. Yet another example with these properties is given by a real cubic Lattès map g. The map g has a lift to the complex plane of the form $\widetilde{g}(z) = \alpha z$, where $\alpha = \sqrt{-3}$ and the postcritical set of g lifts to the lattice

generated by 1 and $(1 + \sqrt{-3})/2$. It is possible to prove these claims for g much as we prove them for f in Section 4. The three analytic conjugacy classes of Lattès maps represented by these three maps might be the only ones for which question (3) is false.

As for question (1), as far as we know, every rational map with finite postcritical set and no periodic critical points is the subdivision map of a finite subdivision rule.

1. Definitions and basic facts for Lattès maps

Following Milnor [7, Remark 3.5], we define a Lattès map to be a rational function from the Riemann sphere $\widehat{\mathbb{C}}$ to itself such that its local degree at every critical point is 2 and there are exactly four postcritical points, none of which is also critical. Let $f\colon \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ be a Lattès map.

As in Section 3.1 of [7], it follows that there exists an analytic branched cover $\wp\colon \mathbb{C} \to \widehat{\mathbb{C}}$ such that $\wp$ is branched exactly over the postcritical points of f and the local degree of $\wp$ at every branch point is 2. (The function $\wp$ is a Weierstrass $\wp$-function up to precomposing and postcomposing with analytic automorphisms.) Let Λ be the set of branch points of $\wp$. It is furthermore true that $\wp$ is a regular branched cover, and its group of deck transformations Γ is generated by the set of all rotations of order 2 about the points of Λ. We refer to Γ as the **orbifold fundamental group** of f. Given rotations $z \mapsto 2\lambda - z$ and $z \mapsto 2\mu - z$ of order 2 about the points $\lambda, \mu \in \Lambda$, their composition, the second followed by the first, is the translation $z \mapsto z + 2(\lambda - \mu)$. We may, and do, normalize so that $0 \in \Lambda$. So Γ contains a subgroup with index 2 consisting of translations of the form $z \mapsto z + 2\lambda$ with $\lambda \in \Lambda$. It follows that Λ is a lattice in $\mathbb{C}$ and that the elements of Γ are the maps of the form $z \mapsto \pm z + 2\lambda$ for some $\lambda \in \Lambda$.

Douady and Hubbard show in [6, Proposition 9.3] that the map $f\colon \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ lifts to a map $\widetilde{f}\colon \mathbb{C} \to \mathbb{C}$ given by $\widetilde{f}(z) = \alpha z + \beta$ for some imaginary quadratic algebraic integer α (possibly an element of $\mathbb{Z}$) such that $\alpha\Lambda \subseteq \Lambda$ and some $\beta \in \Lambda$. The following lemma is devoted to determining to what extent α, β, and Λ are determined by the analytic conjugacy class of f.

Lemma 1.1. *Let f_0 be a Lattès map which is analytically conjugate to f. Let $\wp_0\colon \mathbb{C} \to \widehat{\mathbb{C}}$ be a branched cover for f_0 corresponding to $\wp$. Let Λ_0 be the set of branch points of $\wp_0$, and assume that $0 \in \Lambda_0$. Suppose that $f_0\colon \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ lifts to a map $\widetilde{f}_0\colon \mathbb{C} \to \mathbb{C}$ given by $\widetilde{f}_0(z) = \alpha_0 z + \beta_0$. Then the following hold.*

   (1) $\alpha_0 = \pm\alpha$.
   (2) $\beta_0 = \gamma\beta + \delta$ *where* $\gamma \in \mathbb{C}^\times$ *and* $\delta \in (\alpha + 1)\Lambda_0 + 2\Lambda_0$.
   (3) $\Lambda_0 = \gamma\Lambda$ *with* γ *as in statement* (2).

Conversely, if $\wp_0\colon \mathbb{C} \to \widehat{\mathbb{C}}$ is a branched cover as above, if Λ_0 is the set of branch points of $\wp_0$ with $0 \in \Lambda_0$, and if $\widetilde{f}_0(z) = \alpha_0 z + \beta_0$ such that statements (1), (2), and (3) hold, then $\widetilde{f}_0$ is the lift of a Lattès map which is analytically conjugate to f.

Proof. Let σ be a Möbius transformation such that $f_0 = \sigma \circ f \circ \sigma^{-1}$. Then σ maps the postcritical set of f to the postcritical set of f_0. So just as for f and f_0, the map $\sigma\colon \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ lifts to a map $\widetilde{\sigma}\colon \mathbb{C} \to \mathbb{C}$ given by $\widetilde{\sigma}(z) = \rho z + \nu$ for some $\rho \in \mathbb{C}^\times$ and $\nu \in \mathbb{C}$ such that $\wp_0 \circ \widetilde{\sigma} = \sigma \circ \wp$. Then $\widetilde{\sigma}$ maps Λ to Λ_0; that is, $\rho\Lambda + \nu = \Lambda_0$. In other words, the coset $\rho\Lambda + \nu$ of the group $\rho\Lambda$ equals the group Λ_0. The only way

that a coset can be a group is if it is the trivial coset, and so $\nu \in \Lambda_0$ and $\Lambda_0 = \rho\Lambda$. Since $\widetilde{f}_0$ and $\widetilde{\sigma} \circ \widetilde{f} \circ \widetilde{\sigma}^{-1}$ are both lifts of f_0, they differ by an element of Γ_0, the group of deck transformations of $\wp_0$. In other words, there exists $\lambda_0 \in \Lambda_0$ such that $\pm\widetilde{f}_0(z) + 2\lambda_0 = \widetilde{\sigma} \circ \widetilde{f} \circ \widetilde{\sigma}^{-1}(z)$ for all $z \in \mathbb{C}$. So we have the following:

$$\pm(\alpha_0 z + \beta_0) + 2\lambda_0 = \pm\widetilde{f}_0(z) + 2\lambda_0 = \widetilde{\sigma} \circ \widetilde{f} \circ \widetilde{\sigma}^{-1}(z) = \rho\widetilde{f}(\rho^{-1}(z - \nu)) + \nu$$
$$= \rho(\alpha\rho^{-1}(z - \nu) + \beta) + \nu = \alpha z + \rho\beta + (1 - \alpha)\nu.$$

Hence $\alpha_0 = \pm\alpha$, which yields statement (1). Furthermore,

$$\pm\beta_0 = \rho\beta + (1 - \alpha)\nu - 2\lambda_0.$$

We have seen that $\nu \in \Lambda_0$. So setting $\gamma = \pm\rho$ and $\delta = \pm((\alpha + 1)\nu - 2\alpha\nu - 2\lambda_0)$, we have that $\beta_0 = \gamma\beta + \delta$ with $\gamma \in \mathbb{C}^\times$ and $\delta \in (\alpha + 1)\Lambda_0 + 2\Lambda_0$. We now have verified statements (1), (2), and (3).

For the converse, it is a straightforward matter to construct $\widetilde{\sigma}$ such that $\widetilde{\sigma}$ conjugates $\widetilde{f}$ to $\widetilde{f}_0$ up to the action of Γ_0. One checks that $\widetilde{\sigma}$ descends to a rational map $\sigma \colon \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ which has local degree 1 at every point of $\widehat{\mathbb{C}}$. So σ is a Möbius transformation. It follows that $\widetilde{f}_0$ is the lift of an analytic conjugate of f.

This completes the proof of Lemma 1.1. $\qquad\square$

Lemma 1.1 implies that with an appropriate modification of β, the lattice Λ may be replaced by $\gamma\Lambda$, where γ is any nonzero complex number, without changing the analytic conjugacy class of f. In Section 7 of Chapter 2 of the number theory book [2] by Borevich and Shafarevich the lattices Λ and $\gamma\Lambda$ are said to be similar. Theorem 9 and the remark following it in Section 7 of Chapter 2 of [2] imply that every lattice in $\mathbb{C}$ is similar to a lattice with a $\mathbb{Z}$-basis consisting of 1 and τ, where τ lies in the standard fundamental domain for the action of $\mathrm{SL}(2, \mathbb{Z})$ on the upper half complex plane. More precisely, τ satisfies the following inequalities:

$$\mathrm{Im}(\tau) > 0,$$
$$(1.2) \qquad -\frac{1}{2} < \mathrm{Re}(\tau) \leq \frac{1}{2},$$
$$|\tau| \geq 1 \text{ and if } |\tau| = 1, \text{ then } \mathrm{Re}(\tau) \geq 0.$$

Moreover there is only one such τ which satisfies these inequalities. This and Lemma 1.1 imply that τ is uniquely determined by the analytic conjugacy class of f.

Corollary 1.3. *Let $\widetilde{f}_0(z) = \alpha_0 z + \beta_0$ be a linear polynomial such that $\widetilde{f}_0(\Lambda) \subseteq \Lambda$. Then $\widetilde{f}_0$ descends via the branched cover $\wp \colon \mathbb{C} \to \widehat{\mathbb{C}}$ to a rational function which is analytically conjugate to f if and only if $\alpha_0 = \pm\alpha$ and $\beta_0 = \gamma\beta + \delta$ where $\gamma\Lambda = \Lambda$, $\gamma = e^{\pm 2\pi i/n}$ with $n \in \{1, 2, 3, 4, 6\}$, and $\delta \in (\alpha + 1)\Lambda + 2\Lambda$.*

Proof. In this situation $\Lambda_0 = \Lambda$. So just as for $\widetilde{f}$, the containment $\gamma\Lambda \subseteq \Lambda$ implies that the complex number γ is in fact an imaginary quadratic algebraic integer. Because $\gamma\Lambda = \Lambda$, γ is invertible; that is, it is a unit. But all imaginary quadratic units have the form $e^{\pm 2\pi i/n}$ with $n \in \{1, 2, 3, 4, 6\}$. This discussion and Lemma 1.1 prove Corollary 1.3. $\qquad\square$

In this paragraph we consider related effects of complex conjugation. It is easy to see that the complex conjugate of a Lattès map is also a Lattès map. By applying complex conjugation to the Lattès map f, the branched cover $\wp$, and the lift $\widetilde{f}$ of f,

we see that the complex conjugate of $\widetilde{f}$ is a lift of the complex conjugate of f. With respect to finite subdivision rules, the behavior of f is the same as the behavior of $\overline{f}$, so when considering $\widetilde{f}(z) = \alpha z + \beta$, we may assume that $\text{Im}(\alpha) \geq 0$. Since $\widetilde{f}$ and $-\widetilde{f}$ both lift f, we may also assume that $\text{Re}(\alpha) \geq 0$. This shows that the restrictions put on α in the following lemma are reasonable.

Lemma 1.4. *As above, let Λ be the inverse image in $\mathbb{C}$ of the postcritical set of the Lattès map $f \colon \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$, and let $\widetilde{f}(z) = \alpha z + \beta$ be a lift of f. Suppose that 1 and τ form a $\mathbb{Z}$-basis of Λ. Multiplication by α determines an endomorphism of Λ. Let $\left[\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right]$ be the matrix of this endomorphism with respect to the ordered $\mathbb{Z}$-basis $(1, \tau)$. Suppose that $\text{Im}(\alpha) > 0$. Then $\text{Re}(\alpha) \geq 0$ and τ lies in the standard fundamental domain for the action of $SL(2, \mathbb{Z})$ on the upper half complex plane if and only if the following inequalities are satisfied:*

$$c > 0,$$
$$a \geq -\frac{c}{2},$$
$$\max\{a - c + 1, -a\} \leq d \leq a + c,$$
$$b \leq -c \ \text{ and if } b = -c, \text{ then } d \geq a.$$

Proof. From the first column of the matrix $\left[\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right]$ it follows that $\alpha = a + c\tau$. So $\tau = \frac{\alpha - a}{c}$. The matrix of the endomorphism determined by $\alpha - a$ is $\left[\begin{smallmatrix} 0 & b \\ c & d-a \end{smallmatrix}\right]$. Because the eigenvalues of this matrix are $\alpha - a$ and $\overline{\alpha - a}$, its trace is twice the real part of $\alpha - a$ and its determinant is the square of the modulus of $\alpha - a$. So $\text{Re}(\alpha - a) = \frac{d-a}{2}$ and $|\alpha - a|^2 = -bc$. Hence $\text{Re}(\tau) = \frac{d-a}{2c}$ and $|\tau|^2 = -\frac{b}{c}$. Similarly $\text{Re}(\alpha) = \frac{a+d}{2}$.

Suppose that $\text{Re}(\alpha) \geq 0$ and that the inequalities in display (1.2) hold. Since $\tau = \frac{\alpha - a}{c}$ and $\text{Im}(\alpha) > 0$, the inequality $\text{Im}(\tau) > 0$ implies that $c > 0$, giving the first inequality in the statement of the lemma. For the second inequality, we combine $\text{Re}(\alpha) \geq 0$ and $\text{Re}(\tau) \leq \frac{1}{2}$ to obtain $a + d \geq 0$ and $d - a \leq c$, hence $a - d \geq -c$. So $a \geq -\frac{c}{2}$, giving the second inequality in the statement of the lemma. Combining $-\frac{1}{2} < \text{Re}(\tau) \leq \frac{1}{2}$ and $\text{Re}(\alpha) \geq 0$ obtains $-c < d - a \leq c$ and $a + d \geq 0$, which easily gives the third inequality in the statement of the lemma. The fourth inequality follows from the fact that $|\tau| \geq 1$ with equality only if $\text{Re}(\tau) \geq 0$.

Proving the converse is straightforward.

This proves Lemma 1.4. $\square$

Lemma 1.5. (1) *In Corollary 1.3 the case $\gamma = \pm i$ occurs only when $\tau = i$, and the case $\gamma = \pm e^{\pm 2\pi i/3}$ occurs only when $\tau = e^{2\pi i/6}$.*

(2) *Let $\left[\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right]$ be as in Lemma 1.4, and let M be the reduction of $\left[\begin{smallmatrix} a+1 & b \\ c & d+1 \end{smallmatrix}\right]$ modulo 2. Then a complete list of distinct coset representatives of $(\alpha + 1)\Lambda + 2\Lambda$ in Λ is given by*

$$0 \ \text{if } rank(M) = 2,$$
$$0, 1, \tau, \tau + 1 \ \text{if } rank(M) = 0,$$
$$0, \lambda \ \text{if } rank(M) = 1,$$

where λ is any element of Λ whose image in $\Lambda/2\Lambda$ is not in the column space of M.

Proof. If $\gamma = \pm i$, then $i \in \Lambda$ because $\gamma\Lambda \subseteq \Lambda$. But then i is an integral linear combination of 1 and τ with τ in the standard fundamental domain for the action

of $\mathrm{SL}(2, \mathbb{Z})$ on the upper half complex plane. This implies that $\tau = i$. A similar argument applies when $\gamma = \pm e^{\pm 2\pi i/3}$. This proves statement (1).

Statement (2) is clear. □

Since the elements of the orbifold fundamental group are Euclidean isometries, $\widetilde{f}$ multiplies areas uniformly by the factor $\deg(f)$. Translation by β does not change areas. Multiplication by α multiplies lengths by $|\alpha|$ and areas by $|\alpha|^2 = \alpha\overline{\alpha}$. Multiplication by α also corresponds to multiplication by the matrix $\left[\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right]$, and this multiplies areas by its determinant. Therefore $\deg(f) = \alpha\overline{\alpha} = ad - bc$.

Lemma 1.6. *If a, b, c, and d satisfy the inequalities of Lemma 1.4, then $ad - bc \geq 3c^2/4$.*

Proof. The inequalities of Lemma 1.4 imply that $|a - d| \leq c$. Hence $c^2 \geq (a - d)^2 \geq (a - d)^2 - (a + d)^2 = -4ad$. Since $b \leq -c$ and $c > 0$, it follows that $ad - bc \geq -c^2/4 + c^2 = 3c^2/4$, as desired. □

Let f be a Lattès map with lift $\widetilde{f}(z) = \alpha z + \beta$ and lattice Λ as above. We say that f is **nonrigid** if $\alpha \in \mathbb{R}$, equivalently, $\alpha \in \mathbb{Z}$. We say that f is **rigid** if $\alpha \notin \mathbb{R}$. If f is nonrigid, then since multiplication by an integer stabilizes every lattice in $\mathbb{C}$, Λ can be arbitrary. There are uncountably many analytic conjugacy classes of Lattès maps for every integer $\alpha \geq 2$. On the other hand, there are only finitely many analytic conjugacy classes of rigid Lattès maps with a given degree. To see why, first note that Lemma 1.6 and the paragraph before it imply that in the rigid case, a bound on $\deg(f) = ad - bc$ puts a bound on c. We may assume that $\mathrm{Im}(\alpha) > 0$. The inequalities of Lemma 1.4 imply that a bound on c puts a lower bound on a and d. Because $b \leq -c$, $c > 0$, and $|a - d| \leq c$, a bound on $ad - bc$ puts an upper bound on both a and d. So a bound on $ad - bc$ puts a bound on c, a, d and therefore b. So if $\deg(f)$ is bounded, then there are only finitely many possibilities for a, b, c, and d. These values determine α in the upper half plane and τ. Given α and τ, there are always at most four possibilities for β up to equivalence. So if f is rigid and $\deg(f)$ is bounded, then there are only finitely many possibilities for the analytic conjugacy class of f.

2. The pruning lemma

Suppose that the Lattès map f is the subdivision map of a finite subdivision rule with one tile type. Then the 1-skeleton of our subdivision complex is a tree T in $\widehat{\mathbb{C}}$. The object of the next lemma, the pruning lemma, is to prune T in order to simplify our finite subdivision rule. Since f is a homeomorphism on open cells of the subdivision complex, all postcritical points are vertices of T. The pruning lemma implies that we may assume that all other vertices of T have valence at least 3. This lemma generalizes to more general maps and graphs, but for simplicity we content ourselves with the following statement.

Lemma 2.1 (Pruning Lemma). *Suppose that the Lattès map f is the subdivision map of a finite subdivision rule with one tile type. Then f is the subdivision map of a finite subdivision rule with one tile type such that every vertex in its subdivision complex with valence either 1 or 2 is a postcritical point.*

Proof. Let T be the 1-skeleton of the given finite subdivision rule. We temporarily view T just as a topological space with no regard to vertices or edges. Let S be

the subspace of T which is the union of all arcs joining postcritical points. Then S is connected and hence is a topological tree. We next show that $f(S) \subseteq S$. For this, let γ be an arc in T joining two postcritical points. If the restriction of f to γ is injective, then $f(\gamma)$ is an arc joining two postcritical points, and so $f(\gamma) \subseteq S$. If the restriction of f to γ is not injective, then because f maps γ into the tree T, there exists a point $p \in \gamma$ such that $f(\gamma)$ is folded at $f(p)$. But then p is a critical point and $f(p)$ is a postcritical point. We see that in general $f(\gamma)$ is a union of arcs which join two postcritical points. So $f(S) \subseteq S$. We make S into a graph by putting vertices at the postcritical points as well as the points whose complements have at least three connected components. We see that S is the 1-skeleton of a cell structure for the 2-sphere and that f is a subdivision map for this subdivision complex with one tile type. The vertices of S with valence either 1 or 2 are postcritical points.

This proves Lemma 2.1. $\qquad\square$

Lemma 2.1 shows that if the Lattès map f is the subdivision map of a finite subdivision rule with one tile type and 1-skeleton T, then we may assume that every vertex of T with valence either 1 or 2 is a postcritical point. We refer to the vertices of T which are not postcritical points as **accidental vertices**. The assumption that accidental vertices have valence at least 3 severely limits the possibilities for T. There are five of them. Figure 1 shows all five possibilities for T. The tile in $\widehat{\mathbb{C}}$ which T determines lifts to a tiling of $\mathbb{C}$, and Figure 1 indicates the tiling corresponding to every possibility for T. Dots in trees indicate postcritical points, and dots in tilings indicate inverse images of postcritical points.

We emphasize that the illustrations in Figure 1 are correct only up to isotopy. This is true in particular for the tilings. The inverse image Λ of the postcritical set of f is a lattice and the tiling is invariant under the orbifold fundamental group Γ, which imposes a certain structure, but otherwise every tiling is correct only up to an affine transformation followed by a Γ-invariant isotopy rel Λ.

3. Main results

Theorem 3.1. *Every Lattès map with sufficiently large degree is the subdivision map of a finite subdivision rule with one tile type, and the tile has the form of type* 5 *in Figure* 1.

Proof. The proof begins by fixing notation and making definitions. Then we outline the argument. Finally, we fill in the details.

Let $f \colon \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ be a Lattès map. Let Λ be a lift of the postcritical set of f, let $\widetilde{f}(z) = \alpha z + \beta$ be a lift of f, and let Γ be the orbifold fundamental group as in Section 1. As in the paragraph immediately before Lemma 1.4, we may assume that $\mathrm{Re}(\alpha) \geq 0$ and $\mathrm{Im}(\alpha) \geq 0$. As in the paragraph containing display (1.2), we may assume that Λ is a lattice with a $\mathbb{Z}$-basis consisting of 1 and τ, where τ satisfies the inequalities in display (1.2); that is, we may assume that τ lies in the standard fundamental domain for the action of $\mathrm{SL}(2, \mathbb{Z})$ on the upper half complex plane.

Since 1 and τ are linearly independent over $\mathbb{R}$, so are α^{-1} and $\alpha^{-1}\tau$. So there exists an $\mathbb{R}$-linear isomorphism $\psi \colon \mathbb{C} \to \mathbb{C}$ such that the points of the lattice $\psi(\alpha^{-1}\Lambda)$ are midpoints of edges of a standard tiling T of the plane by regular hexagons and
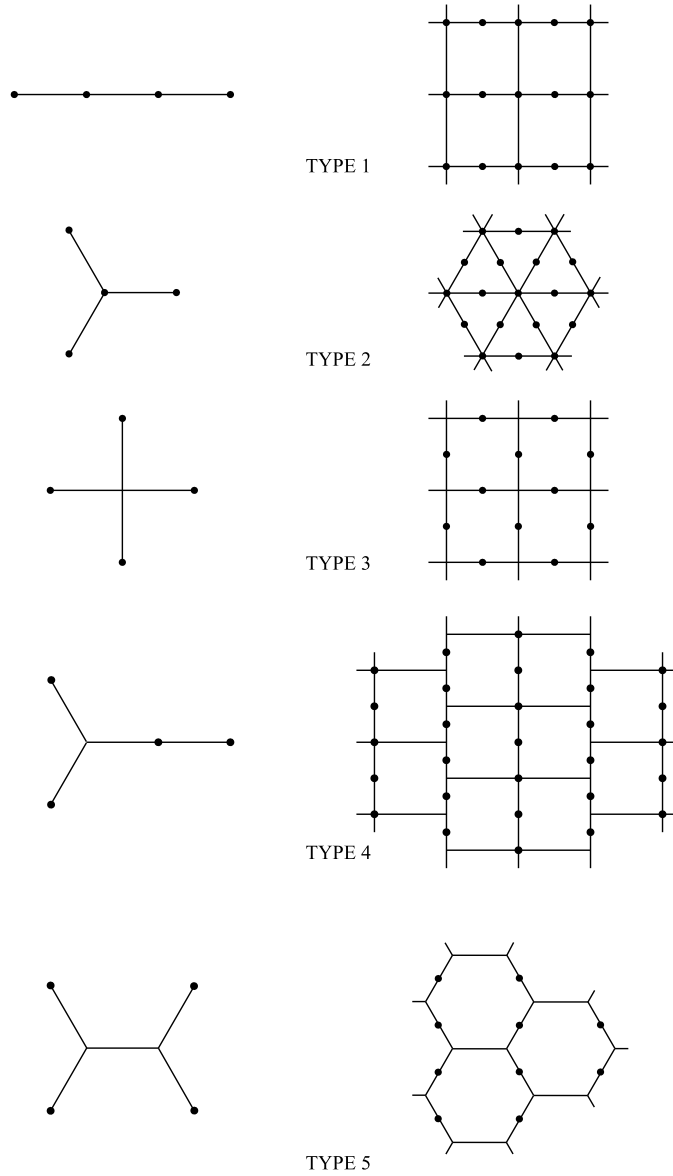
TYPE 1

TYPE 2

TYPE 3

TYPE 4

TYPE 5

FIGURE 1. The possible trees and tilings.

0, $\psi(\alpha^{-1})$, and $\psi(\alpha^{-1}\tau)$ lie in one hexagon as shown in Figure 2. Let T^* be the tiling of the plane by equilateral triangles which is dual to T. The vertices of T^* are the centers of the tiles of T. The duality between T and T^* determines a bijection between the edges of T and the edges of T^*. Let S and S^* be the tilings of the plane which are the pullbacks of T and T^* under ψ.

In this paragraph we consider approximations to lines in the plane by edge paths in S. Let L be a line in the plane, and first suppose that L does not contain a vertex of S^*. We construct a subset $\widehat{L}$ of the plane as a union of edges of S. An edge e of
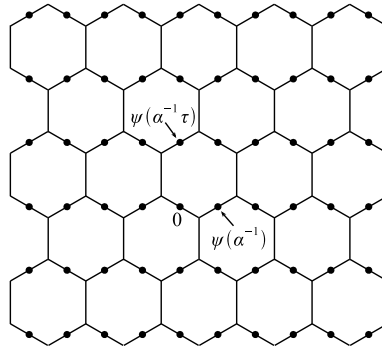
FIGURE 2. The tiling T.
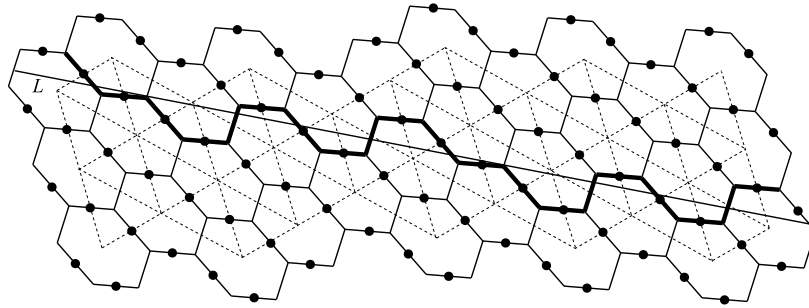
FIGURE 3. Constructing the edge path approximation to the line L.

S is contained in $\widehat{L}$ if and only if L meets the edge of S^* dual to e; see Figure 3, which shows part of S and S^*. The dots indicate points of the lattice $\alpha^{-1}\Lambda$. Part of L is shown, and part of $\widehat{L}$ is drawn with thick line segments. It is not difficult to see that $\widehat{L}$ is homeomorphic to $\mathbb{R}$. If L contains a vertex but not an edge of S^*, then we construct $\widehat{L}$ in essentially the same way by perturbing L near every vertex of S^* contained in L. If L is a union of edges of S^*, then we perturb every edge in L to a new line segment which contains no vertex of S^*. The result is again a subset of the 1-skeleton of S which is homeomorphic to $\mathbb{R}$, but in this case there are two choices for every vertex of S^* in L and $\widehat{L}$ is not unique. We call $\widehat{L}$ the **edge path approximation** to L.

We now outline the proof of Theorem 3.1; see Figure 4. Let L_1 be the line containing β and $1 + \beta$. Since $\beta \in \Lambda$, either L_1 contains the centers of both tiles of S which contain β, or L_1 does not contain the center of either tile of S which contains β. If L_1 does not contain the centers of the two tiles of S which contain β, then $\widehat{L}_1$ contains β. If L_1 contains the centers of the two tiles of S which contain β, then $\widehat{L}_1$ is not uniquely determined, but we may choose $\widehat{L}_1$ so that it contains β. So we may assume that $\widehat{L}_1$ contains β, and likewise $1 + \beta$. Let L_2 be the line containing $\tau + \beta$ and $1 + \tau + \beta$. As for $\widehat{L}_1$, we may assume that $\widehat{L}_2$ contains both $\tau + \beta$ and $1 + \tau + \beta$. We will later prove that $\widehat{L}_1$ and $\widehat{L}_2$ are usually disjoint. Let L_3 be the line containing $-\frac{1}{2} + \beta$ and $-\frac{1}{2} + \tau + \beta$. Let P_1 and P_2 be points of $\widehat{L}_1 \cap \widehat{L}_3$
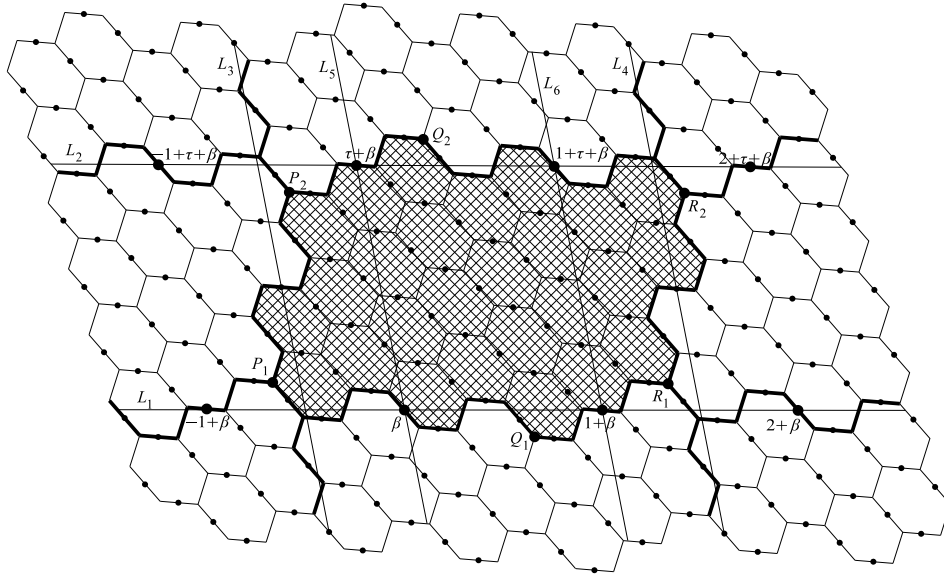
FIGURE 4. Constructing a fundamental domain when $\alpha = \frac{9+\sqrt{-23}}{2}$ and $\tau = \frac{-1+\sqrt{-23}}{4}$.

and $\widehat{L}_2 \cap \widehat{L}_3$, respectively, such that the open segment of $\widehat{L}_3$ with endpoints P_1 and P_2 is disjoint from both $\widehat{L}_1$ and $\widehat{L}_2$. We will later prove that P_1 is usually strictly between $-1+\beta$ and β in $\widehat{L}_1$ and that P_2 is usually strictly between $-1+\tau+\beta$ and $\tau + \beta$ in $\widehat{L}_2$. Let Q_1 be the image of P_1 under the rotation of order 2 about β, and let R_1 be the image of Q_1 under the rotation of order 2 about $1 + \beta$. When $\widehat{L}_1$ is uniquely determined, the rotation of order 2 about any point of $\Lambda \cap L_1$ stabilizes $\widehat{L}_1$. When $\widehat{L}_1$ is not uniquely determined, we construct $\widehat{L}_1$ by making choices between β and $1 + \beta$ and then extending so that the rotation of order 2 about any point of $\Lambda \cap L_1$ stabilizes $\widehat{L}_1$. We construct $\widehat{L}_2$ in the same way. So Q_1 is a point of $\widehat{L}_1$ usually strictly between β and $1+\beta$ and R_1 is a point of $\widehat{L}_1$ usually strictly between $1 + \beta$ and $2 + \beta$. The composition of two rotations of order 2 is a translation. We translate L_3 by the map $z \mapsto z + R_1 - P_1$ to a line L_4. The image of the closed segment of $\widehat{L}_3$ with endpoints P_1 and P_2 under this translation is the closed segment of $\widehat{L}_4$ with endpoints R_1 and R_2. We will later prove that the closed segment of $\widehat{L}_3$ with endpoints P_1 and P_2 is usually disjoint from its image under this translation. From the segment of $\widehat{L}_1$ joining P_1 and R_1, the segment of $\widehat{L}_4$ joining R_1 and R_2, the segment of $\widehat{L}_2$ joining R_2 and P_2, and the segment of $\widehat{L}_3$ joining P_2 and P_1 we obtain the hatched region F, which is a union of tiles of S. Let S' be the tiling obtained from S simply by making the points of Λ vertices. So the vertices of S are the vertices of S' with valence 3, and the points of Λ are the vertices of S' with valence 2. Whereas the tiles of S are hexagons, the tiles of S' are decagons. Let s be the tile of S' containing 0, α^{-1}, $\alpha^{-1}\tau$, and $\alpha^{-1} + \alpha^{-1}\tau$. There is a canonical pairing of the edges of s. The two edges of s containing 0 are interchanged by the rotation of order 2 about 0. The same is true at α^{-1}, $\alpha^{-1}\tau$, and $\alpha^{-1} + \alpha^{-1}\tau$. The two remaining edges of s are translates of one another. In the same way we view

F as having six vertices P_1, Q_1, R_1, P_2, Q_2, R_2 of valence 3 and four vertices β, $1 + \beta$, $\tau + \beta$, $1 + \tau + \beta$ of valence 2, which we view as decomposing the boundary of F into ten edges. The two edges of F containing β are interchanged by the rotation of order 2 about β. The same is true at $1 + \beta$, $\tau + \beta$, and $1 + \tau + \beta$. The two remaining edges of F are translates of one another. It follows that F is a fundamental domain for the orbifold fundamental group Γ. The image of every tile of S' under $\widetilde{f}$ is isotopic rel Λ to the image of F under some element of Γ, and these isotopies can be made Γ-equivariant. Because f respects the edge pairings of s and F up to isotopy, when we descend to $\widehat{\mathbb{C}}$, the tiling of F by the tiles of S' determines a finite subdivision rule $\mathcal{R}$ with one tile type. The subdivision rule $\mathcal{R}$ has bounded valence; in fact, valences of vertices are bounded by 3. We will prove that the mesh of $\mathcal{R}$ usually approaches 0 combinatorially. The result is a finite subdivision rule $\mathcal{R}$ with bounded valence and mesh approaching 0 combinatorially such that the subdivision map $\sigma_{\mathcal{R}}$ of $\mathcal{R}$ is isotopic to f rel P_f. From here we proceed as in the proof of the main theorem of [4] to conclude that $\sigma_{\mathcal{R}}$ and f are in fact conjugate rel P_f, and so f is the subdivision map of a finite subdivision rule with one tile type.

The previous paragraph reduces the proof of Theorem 3.1 to the construction of an appropriate fundamental domain F for the orbifold fundamental group Γ. This in turn reduces to verifying certain claims made in the previous paragraph. The word "usually" occurs a few times. "Usually" means that the degree of f is sufficiently large and that $\alpha \notin \{\tau, 1+\tau, 2+\tau\}$. The claims in the previous paragraph do not always hold if $\alpha \in \{\tau, 1 + \tau, 2 + \tau\}$ and sometimes when they do, the proofs below do not. The cases in which $\alpha \in \{\tau, 1+\tau, 2+\tau\}$ will be handled separately at the end of the proof. When edge path approximations are not unique, the claims mean that it is possible to choose edge path approximations so that the claims are true.

In this paragraph we make an observation concerning a type of symmetry. The outline in the next-to-last paragraph constructs a fundamental domain F for the Lattès map with lift $z \mapsto \alpha z + \beta$ and lattice generated by 1 and τ. The rotation of order 2 about $1 + \tau + \beta$ takes F to the corresponding fundamental domain for the Lattès map with lift $z \mapsto \alpha z + 1 + \tau + \beta$ and lattice generated by 1 and τ. So for example, by this symmetry if for a fixed α and τ and for every $\beta \in \Lambda$ no tile of S' meets an edge of F containing β and an edge of F containing $\tau + \beta$, then for the same fixed α and τ and for every $\beta \in \Lambda$ no tile of S' meets an edge of F containing $1 + \beta$ and an edge of F containing $1 + \tau + \beta$. We will use this symmetry to simplify arguments.

We prepare to verify the claims in the outline by turning our attention to edge path approximations.

Lemma 3.2. *Let L and L' be parallel lines in the plane.*

    (1) *If no edge of S^* meets both L and L', then $\widehat{L}$ and $\widehat{L}'$ are disjoint.*

    (2) *If there does not exist a line segment which is the union of two adjacent edges of S^* such that this line segment meets both L and L', then no tile of S' meets both $\widehat{L}$ and $\widehat{L}'$.*

Proof. Note that because the vertices of S have valence 3, if $\widehat{L}$ and $\widehat{L}'$ are not disjoint, then they have an edge of S in common. With this observation, statement (1) follows from the definition of edge path approximation. To prove statement (2), let L'' be the line parallel to both L and L' such that L'' is equidistant to L and L'.

Suppose that there does not exist a line segment which is the union of two adjacent edges of S^* such that this line segment meets both L and L'. Then no edge of S^* meets both L and L''. Likewise no edge of S^* meets both L' and L''. So statement (1) implies that $\widehat{L}$, $\widehat{L}'$, and $\widehat{L}''$ are mutually disjoint. Moreover, $\widehat{L}''$ is between $\widehat{L}$ and $\widehat{L}'$. It follows that no tile of S' meets both $\widehat{L}$ and $\widehat{L}'$.

This proves Lemma 3.2. $\hfill\square$

Lemma 3.3. *If the degree of f is sufficiently large, then no tile of S' meets both $\widehat{L}_1$ and $\widehat{L}_2$.*

Proof. We prepare to apply statement (2) of Lemma 3.2. We consider whether or not there exists a line segment which is the union of two adjacent edges of S^* such that this line segment meets both L_1 and L_2. First suppose that these two edges of S^* are translates of the line segment joining 0 and $2\alpha^{-1}$. Since $\tau + \beta \in L_2$ and L_1 is the translate of the real axis by β, for such edges we are led to consider real numbers r and s such that $\tau + \beta + 2r\alpha^{-1} = s + \beta$; that is, $\tau + 2r\alpha^{-1} = s$. If $\alpha \in \mathbb{R}$, then there are no such numbers r and s and there is no such line segment which meets both L_1 and L_2. Otherwise r and s exist, and if $|r| > 2$, then there does not exist such a line segment which meets both L_1 and L_2. The equation $\tau + 2r\alpha^{-1} = s$ implies that $\alpha = r\frac{2}{s-\tau}$. Since $\mathrm{Im}(\tau)$ is bounded from 0, the complex number $s - \tau$ is bounded from 0 uniformly in τ as s varies over $\mathbb{R}$. So $r\frac{2}{s-\tau}$ is bounded uniformly in τ for $|r| \le 2$ and $s \in \mathbb{R}$. Hence α is bounded for $|r| \le 2$ and $s \in \mathbb{R}$. But since $\deg(f) = \alpha\overline{\alpha}$, it follows that there is no such line segment meeting both L_1 and L_2 if the degree of f is sufficiently large.

In addition to having edges in the direction of $2\alpha^{-1}$, the tiling S^* has edges in the direction of $\alpha^{-1}(1 + \tau)$ and $\alpha^{-1}(-1 + \tau)$. We next perform analogous verifications for these edges. Once these verifications are complete, we may conclude that no tile of S' meets both $\widehat{L}_1$ and $\widehat{L}_2$ if the degree of f is sufficiently large.

Now we consider edges of S^* in the direction of $\alpha^{-1}(1 + \tau)$. In this case we obtain the equation $\tau + r\alpha^{-1}(1 + \tau) = s$. So $\alpha = r\frac{1+\tau}{s-\tau}$. For all real numbers s we have that

(3.4) $$|s - \tau| \ge \mathrm{Im}(\tau) = |\tau|\sin(\arg(\tau)) \ge \frac{\sqrt{3}}{2}|\tau|,$$

and so

$$\left|\frac{1 + \tau}{s - \tau}\right| \le \frac{2}{\sqrt{3}}\frac{1 + |\tau|}{|\tau|} \le \frac{4}{\sqrt{3}}.$$

So $r\frac{1+\tau}{s-\tau}$ is bounded for $|r| \le 2$ and $s \in \mathbb{R}$. Again it follows that there is no such line segment meeting both L_1 and L_2 if the degree of f is sufficiently large.

For edges of S^* in the direction of $\alpha^{-1}(-1 + \tau)$ we replace $1 + \tau$ by $-1 + \tau$ in the previous paragraph and again conclude that there is no such line segment meeting both L_1 and L_2 if the degree of f is sufficiently large. Thus if the degree of f is sufficiently large, then no tile of S' meets both $\widehat{L}_1$ and $\widehat{L}_2$.

This proves Lemma 3.3. $\hfill\square$

As in Figure 4, let L_5 be the line containing β and $\tau + \beta$ and let L_6 be the line containing $1 + \beta$ and $1 + \tau + \beta$. It would be convenient to have a lemma for L_5 and L_6 analogous to Lemma 3.3 for L_1 and L_2. Unfortunately, the situation for L_5 and L_6 is more complicated than for L_1 and L_2. This takes us to the following two lemmas.

Lemma 3.5. *Let L and L' be lines in the plane parallel to L_5 such that the distance between L and L' is at least half the distance between L_5 and L_6. If the degree of f is sufficiently large, then no line segment which is the union of two edges of S^* in the direction of α^{-1} meets both L and L'.*

Proof. We consider whether or not there exists a line segment meeting both L and L' which is the union of two edges of S^* which are translates of the line segment joining 0 and $2\alpha^{-1}$. As in the proof of Lemma 3.3, we obtain the equation $x + 2r\alpha^{-1} = s\tau$, and we wish to have no solution with $x \geq 1/2$, $|r| \leq 2$, and $s \in \mathbb{R}$. So we wish to have no solution to the equation $1 + 2r\alpha^{-1} = s\tau$ with $|r| \leq 4$ and $s \in \mathbb{R}$. Solving the last equation for α shows that $\alpha = r\frac{2}{s\tau - 1}$. The argument in display (3.4) with τ^{-1} instead of τ shows for all real numbers s that $\left|s - \tau^{-1}\right| \geq \frac{\sqrt{3}}{2}|\tau|^{-1}$, and so $|s\tau - 1| \geq \frac{\sqrt{3}}{2}$. So $r\frac{2}{s\tau - 1}$ is bounded uniformly in τ for $|r| \leq 4$ and $s \in \mathbb{R}$. Hence α is bounded for $|r| \leq 4$ and $s \in \mathbb{R}$. Since $\deg(f) = \alpha\overline{\alpha}$, this proves that there is no such line segment if the degree of f is sufficiently large.

This proves Lemma 3.5. $\qquad\square$

Lemma 3.6. *Let L and L' be lines in the plane parallel to L_5. Let a and c be the integers such that $\alpha = a + c\tau$.*

(1) *If the distance between L and L' is at least half the distance between L_5 and L_6 and if either $c \geq 11$ or $c = 0$ with $a \geq 5$, then no line segment which is the union of two edges of S^* in the direction of either $\alpha^{-1}(1 + \tau)$ or $\alpha^{-1}(-1 + \tau)$ meets both L and L'.*

(2) *If the distance between L and L' is at least twice the distance between L_5 and L_6, then usually no edge of S^* in the direction of either $\alpha^{-1}(1 + \tau)$ or $\alpha^{-1}(-1 + \tau)$ meets both L and L'.*

Proof. We prove both statements together. We first consider edges of S^* in the direction of $\alpha^{-1}(1 + \tau)$. Arguing as in the proof of Lemma 3.5, we obtain the equation $1 + r\alpha^{-1}(1 + \tau) = s\tau$. For statement (1) we want there to be no solution in real numbers r and s with $|r| \leq 4$. For statement (2) the restriction on r is $|r| \leq \frac{1}{2}$. Solving for α shows that $\alpha = -r\frac{1 + \tau}{1 - s\tau}$.

We justify every step of the next display immediately after it:

$$c\,\mathrm{Im}(\tau) = \mathrm{Im}(\alpha) \leq |\alpha| = |r|\left|\frac{1 + \tau}{1 - s\tau}\right| \leq |r|\frac{4|\tau|}{\sqrt{3}} \leq |r|\frac{8}{3}\mathrm{Im}(\tau).$$

The two equations and the first inequality are straightforward. For the second inequality, we use the fact that $1 \leq |\tau|$ to obtain $|1 + \tau| \leq 1 + |\tau| \leq 2|\tau|$. Using display (3.4) as in the proof of Lemma 3.5, we obtain $|1 - s\tau| \geq \frac{\sqrt{3}}{2}$. Combining the results of the last two sentences gives the second inequality of the display. Display (3.4) implies that $|\tau| \leq \frac{2}{\sqrt{3}}\mathrm{Im}(\tau)$. This gives the last inequality of the above display. From the display we conclude that $c \leq |r|\frac{8}{3}$. Hence if $|r| \leq 4$, then $c \leq \frac{32}{3}$ and if $|r| \leq \frac{1}{2}$, then $c \leq \frac{4}{3}$. This proves statement (1) for edges of S^* in the direction of $\alpha^{-1}(1 + \tau)$ except when $c = 0$ and it proves statement (2) for edges of S^* in the direction of $\alpha^{-1}(-1 + \tau)$ except when $c \in \{0, 1\}$.

To continue, we consider the image of the extended real line under the linear fractional transformation $z \mapsto \frac{1 + \tau}{1 - z\tau}$. One verifies that $\infty \mapsto 0$, $-1 \mapsto 1$, and $0 \mapsto 1 + \tau$. Since linear fractional transformations map circles and extended lines to circles and extended lines, it follows that the image of the extended real line is the circle C containing 0, 1, and $1 + \tau$.

Thus if $|r| \leq 4$ and $c = 0$, then $0 < \alpha \leq 4$. This completes the proof of statement (1) for edges of S^* in the direction of $\alpha^{-1}(1 + \tau)$. If $|r| \leq \frac{1}{2}$ and $c = 0$, then $0 < \alpha \leq \frac{1}{2}$. Hence statement (2) is true for edges of S^* in the direction of $\alpha^{-1}(1 + \tau)$ if $c = 0$.

It remains to prove statement (2) for edges of S^* in the direction of $\alpha^{-1}(1 + \tau)$ when $c = 1$. Recall that C contains 0, 1, and $1 + \tau$. This implies that the center of C is on the line given by $\mathrm{Re}(z) = \frac{1}{2}$. Every point of the circle C' through 0 and 1 with center $\frac{1}{2}$ has imaginary part at most $\frac{1}{2}$. Since $\mathrm{Im}(1 + \tau) \geq \frac{\sqrt{3}}{2} > \frac{1}{2}$, the imaginary part of the center of C is positive. Comparing C with C', we see that no points of the form $a - \tau$ with $a \in \mathbb{R}$ are within C. Since $\mathrm{Re}(1 + \tau) > \frac{1}{2}$, the complex number $1 + \tau$ is in the right half of C. So no complex number of the form $a + \tau$ with $a > 1$ is within C. Hence no such number is within $\frac{1}{2}C$. Since "usually" means that $a \geq 3$ if $c = 1$, this completes the proof of statement (2) for edges of S^* in the direction of $\alpha^{-1}(1 + \tau)$.

For edges of S^* in the direction of $\alpha^{-1}(-1 + \tau)$, we argue in the same way. The only modification occurs at the very end. Whereas before the circle C contains $1 + \tau$, now it contains $-1 + \tau$. This is insufficient to conclude that no complex number of the form $a + \tau$ with $a > 1$ is within C. For this we verify that the linear fractional transformation $z \mapsto \frac{-1 + \tau}{1 - z\tau}$ maps the real number $\frac{-1 + \tau + \overline{\tau}}{\tau \overline{\tau}}$ to $-\overline{\tau} = -2\mathrm{Re}(\tau) + \tau$. Now we can conclude that no complex number of the form $a + \tau$ with $a > 1$ is within C.

This proves Lemma 3.6. $\square$

We now consider the claims made in the outline of Theorem 3.1. The first claim is that $\widehat{L}_1$ and $\widehat{L}_2$ are usually disjoint. This follows from Lemma 3.3. The next claim is that P_1 is usually strictly between $-1 + \beta$ and β in $\widehat{L}_1$. This is the main content of the first statement of the next lemma. The other two statements are closely related results which will be used later.

Lemma 3.7. (1) *The point P_1 is usually strictly between $-1 + \beta$ and β in $\widehat{L}_1$ and P_1 is usually not adjacent to either $-1 + \beta$ or β in the tiling S'.*

   (2) *Usually no tile of S' meets both $\widehat{L}_3$ between P_1 and P_2 and $\widehat{L}_1$ between Q_1 and R_1.*

   (3) *Usually no tile of S' meets both $\widehat{L}_4$ between R_1 and R_2 and $\widehat{L}_1$ between P_1 and Q_1.*

Proof. We prove all three statements together. As usual, let a and c be the integers such that $\alpha = a + c\tau$. We begin by proving Lemma 3.7 for either $c = 0$ or $c \geq 11$. In this case, statement (2) of Lemma 3.2, Lemma 3.5, and statement (1) of Lemma 3.6 combine to imply that usually no tile of S' meets both $\widehat{L}_3$ and $\widehat{L}_5$. (The condition in Lemma 3.6 that $a \geq 5$ can be met by taking the degree of f to be sufficiently large.) Since $\beta \in L_5$, we may assume that $\beta \in \widehat{L}_5$. So we have that $\beta \in \widehat{L}_1 \cap \widehat{L}_5$, and it is easy to see that the points of $\widehat{L}_1$ which are on the same side of β as $1 + \beta$ either lie on $\widehat{L}_5$ or are on the same side of $\widehat{L}_5$ as $1 + \beta$. Since $\widehat{L}_3$ is usually strictly on the other side of $\widehat{L}_5$ and $P_1 \in \widehat{L}_3$, it follows that P_1 is usually strictly on the same side of β as $-1 + \beta$ in $\widehat{L}_1$ and that P_1 is usually not adjacent to β in S'. This argument also shows that usually no tile of S' meets $\widehat{L}_3$ between P_1 and P_2 and $\widehat{L}_1$ between Q_1 and R_1. We have so far proved that P_1 is usually strictly on the same side of β as $-1 + \beta$ in $\widehat{L}_1$, that P_1 is usually not adjacent to β in S', and that

statement (2) is true if either $c = 0$ or $c \geq 11$. Analogous arguments prove that P_1 is usually strictly on the same side of $-1 + \beta$ as β in $\widehat{L}_1$, that P_1 is usually not adjacent to $-1 + \beta$ in S', and that statement (3) is true if either $c = 0$ or $c \geq 11$. This proves Lemma 3.7 if either $c = 0$ or $c \geq 11$.

So suppose that $1 \leq c \leq 10$ for the rest of the proof of Lemma 3.7. We partition this case into two subcases according to whether $a < 3c$ or $a \geq 3c$. Let b and d be the integers such that $\alpha\tau = b + d\tau$. If $a < 3c$, then both a and c are bounded. With a and c bounded, Lemma 1.4 implies that d is bounded. Because $\deg(f) = ad - bc$ and this degree may be taken to be arbitrarily large, the integer b may be taken to be arbitrarily negative. So if $a < 3c$, then d is bounded and b may be taken to be arbitrarily negative. On the other hand, if $a \geq 3c$, then Lemma 1.4 implies that $d > 0$. Of course, $b < 0$.

In this paragraph suppose that $a \geq 3c$. We will prove statement (2) and half of statement (1) under this assumption. Let $X = \{x\alpha^{-1} + y\alpha^{-1}\tau + \beta : x, y \in \mathbb{R}, x \leq -\frac{3}{2}\}$, a closed half plane. Let $\overline{X}$ be the union of the tiles of S' contained in X. Both $\psi(\overline{X})$ and the image of $\psi(\overline{X})$ under the rotation of order 2 about $\psi(\beta)$ are shaded in Figure 5. There are essentially two possibilities, hence two parts to Figure 5, depending on whether the edge of T containing $\psi(\beta)$ has negative or positive slope. It is possible to choose $\widehat{L}_1$ so that the portion of $\psi(\widehat{L}_1)$ in the region shown in Figure 5 is contained in the hatched region which is pinched near $\psi(\beta)$. The lines L_1 and L_3 meet at $-\frac{1}{2} + \beta = -\frac{a}{2}\alpha^{-1} - \frac{c}{2}\alpha^{-1}\tau + \beta$, and this point is in X since $a \geq 3c \geq 3$. Moreover, this point is in the boundary of X if and only if $a = 3$ and $c = 1$. Since $b < 0$, the half of L_3 with endpoint $-\frac{1}{2} + \beta$ which contains $-\frac{1}{2} + \tau + \beta = -\frac{1}{2} + b\alpha^{-1} + d\alpha^{-1}\tau + \beta$ is in X. So the half of $\widehat{L}_3$ with endpoint P_1 which contains P_2 is contained in $\overline{X}$. Similarly, the half of $\widehat{L}_1$ with endpoint P_1 which does not contain β is contained in $\overline{X}$. So the half of $\widehat{L}_1$ with endpoint Q_1 which does not contain β is contained in the image of $\overline{X}$ under the rotation of order 2 about β. We see that P_1 is strictly on the same side of β as $-1 + \beta$ in $\widehat{L}_1$ and that P_1 is not adjacent to β in S'. Moreover no tile of S' meets both $\widehat{L}_3$ between P_1 and P_2 and $\widehat{L}_1$ between Q_1 and R_1. We have just proved statement (2) and half of statement (1) when $1 \leq c \leq 10$ and $a \geq 3c$.

Next suppose that $a < 3c$. We will prove statement (2) and the same half of statement (1) under this assumption. Recall from the definition of usually that $\alpha \notin \{\tau, 1 + \tau, 2 + \tau\}$. This together with the inequality $a < 3c$ implies that we may assume that $c \neq 1$. For every integer m we consider the closed half plane $Y_m = \{x\alpha^{-1} + y\alpha^{-1}\tau + \beta : y \geq m - \frac{1}{2}\}$. Let $\overline{Y}_m$ be the union of the tiles of S' contained in Y_m. Now let m be the ceiling of $-\frac{c}{2}$. Again L_1 and L_3 meet at $-\frac{1}{2} + \beta = -\frac{a}{2}\alpha^{-1} - \frac{c}{2}\alpha^{-1}\tau + \beta$. This point is in Y_m. Because L_3 contains $-\frac{1}{2} + \beta$ and $-\frac{1}{2} + \tau + \beta = -\frac{1}{2} + b\alpha^{-1} + d\alpha^{-1}\tau + \beta$ and because d is bounded and b can be made arbitrarily negative, the slope of $\psi(L_3)$ can be made arbitrarily close to 0. The edge path distance between $-\frac{1}{2} + \beta$ and P_1 is bounded since a and c are bounded, so it follows that usually an arbitrarily long initial portion of the directed segment of $\widehat{L}_3$ from P_1 to P_2 is in the boundary of either $\overline{Y}_m$ or $\overline{Y}_{m-1}$. In particular, P_1 is usually in the boundary of either $\overline{Y}_m$ or $\overline{Y}_{m-1}$. Moreover since $c \geq 2$, it follows that $m < 0$, and so $P_1 \notin \overline{Y}_0$. Similarly, the half of $\widehat{L}_1$ with endpoint P_1 which does not contain β is disjoint from $\overline{Y}_0$. On the other hand, $\beta \in \overline{Y}_0$, and the half of $\widehat{L}_1$ with endpoint β which contains $1 + \beta$ is even in $\overline{Y}_0$. The previous two sentences
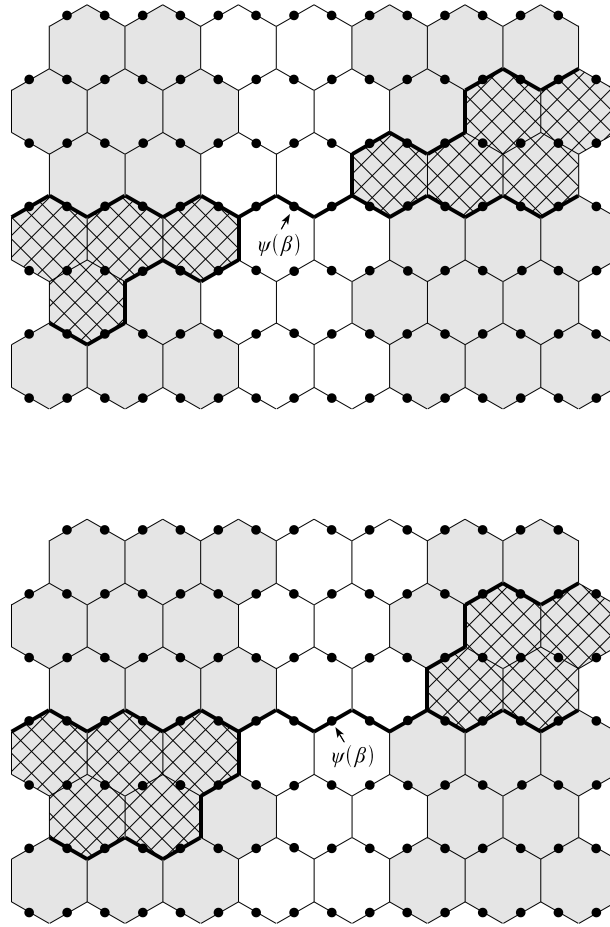
FIGURE 5. Proving Lemma 3.7.

together with the fact that Q_1 is the image of P_1 under the rotation of order 2 about β imply that the half of $\widehat{L}_1$ with endpoint Q_1 which does not contain β is in $\overline{Y}_1$. Thus P_1 is usually strictly on the same side of β as $-1 + \beta$ in $\widehat{L}_1$, it is usually not adjacent to β in S', and usually no tile of S' meets both $\widehat{L}_3$ between P_1 and P_2 and $\widehat{L}_1$ between Q_1 and R_1. This completes the proof of statement (2) and half of statement (1) when $1 \le c \le 10$.

Finally we prove statement (3) and the other half of statement (1) when $1 \le c \le 10$. As in the previous paragraph, it is still true that $-\frac{1}{2} + \beta$ is contained in Y_m, where m is the ceiling of $-\frac{c}{2}$. It is also still usually true that P_1 is in $\overline{Y}_n$ with either $n = m$ or $n = m - 1$. If $a \ge 3c$, then the slope of $\psi(L_3)$ is negative and the directed segment of $\widehat{L}_3$ from P_1 to P_2 is also in $\overline{Y}_n$. If $a < 3c$, then this slope might be negative, but usually an arbitrarily long initial portion of this directed segment is contained in $\overline{Y}_n$. Moreover, if $c \in \{1, 2\}$, then $d \ge 0$, the slope of $\psi(L_3)$ is nonpositive, and we may take $n = m$. It follows that $-c < n$ and that $-1 + \beta \notin \overline{Y}_n$. It is also true that the half of $\widehat{L}_1$ with endpoint P_1 which contains

β is contained in $\overline{Y}_n$. The last two sentences combine to prove that P_1 is usually strictly on the same side of $-1 + \beta$ as β in $\widehat{L}_1$ and that P_1 is not adjacent to $-1 + \beta$ in S'. The translation $z \mapsto z + 2 = z + 2a\alpha^{-1} + 2c\alpha^{-1}\tau$ stabilizes $\widehat{L}_1$ and it takes the directed segment of $\widehat{L}_3$ from P_1 to P_2 to the directed segment of $\widehat{L}_4$ from R_1 to R_2. So either the directed segment of $\widehat{L}_4$ from R_1 to R_2 is contained in $\overline{Y}_{n+2c}$ or at least an arbitrarily long initial portion of it is usually contained in $\overline{Y}_{n+2c}$. Since $-1 + b \notin \overline{Y}_n$, it follows that $1 + \beta \notin \overline{Y}_{n+2c}$. Since the segment of $\widehat{L}_1$ between P_1 and Q_1 is contained in the image of $\overline{Y}_{n+2c}$ under the rotation of order 2 about $1 + \beta$, it follows that usually no tile of S' meets both the segment of $\widehat{L}_4$ between R_1 and R_2 and the segment of $\widehat{L}_1$ between P_1 and Q_1.

This proves Lemma 3.7. $\qquad\square$

So P_1 is usually strictly between $-1 + \beta$ and β in $\widehat{L}_1$. The next claim in the outline is that P_2 is usually strictly between $-1 + \tau + \beta$ and $\tau + \beta$ in $\widehat{L}_2$. To see this, note that the previous claim and symmetry show that R_2 is usually strictly between $1 + \tau + \beta$ and $2 + \tau + \beta$ in $\widehat{L}_2$. This and a translation imply that P_2 is usually strictly between $-1 + \tau + \beta$ and $\tau + \beta$ in $\widehat{L}_2$.

The next claim is that the closed segment of $\widehat{L}_3$ with endpoints P_1 and P_2 is usually disjoint from its image under the translation given by $z \mapsto z + R_1 - P_1$. This follows from statement (1) of Lemma 3.2, Lemma 3.5, and statement (2) of Lemma 3.6.

In the outline of the proof of Theorem 3.1, a fundamental domain F for the orbifold fundamental group Γ is constructed. The points P_1, Q_1, R_1, P_2, Q_2, R_2, β, $1 + \beta$, $\tau + \beta$, $1 + \tau + \beta$ on the boundary of F decompose the boundary of F into ten edges, and there is a pairing of these edges. There is a corresponding decomposition of the boundary of every tile of S' into ten paired edges. This and the tiling of F by tiles of S' determine a finite subdivision rule $\mathcal{R}$. The final claim is that the mesh of $\mathcal{R}$ usually approaches 0 combinatorially. To say that the mesh of $\mathcal{R}$ approaches 0 combinatorially means that there exists a positive integer n such that the following two conditions hold, where t is the tile type of $\mathcal{R}$. Condition (1) is that every edge of t properly subdivides in $\mathcal{R}^n(t)$. Condition (2) is that no tile of $\mathcal{R}^n(t)$ has two edges contained in disjoint edges of t. The next two lemmas verify that $\mathcal{R}$ usually satisfies conditions (1) and (2), and so the mesh of $\mathcal{R}$ usually approaches 0 combinatorially. The strategy is to first show in Lemma 3.8 that if $\mathcal{R}$ satisfies condition (2), then it satisfies condition (1). Lemma 3.9 completes the verification by showing that $\mathcal{R}$ satisfies condition (2).

Lemma 3.8. *If $\mathcal{R}$ satisfies condition* (2)*, then it satisfies condition* (1)*.*

Proof. The translates of the fundamental domain F under the universal orbifold group tile the plane. Furthermore, the points P_1, Q_1, R_1, P_2, Q_2 R_2 are vertices of this tiling with valence 3. So if $v \in \{P_1, Q_1, R_1, P_2, Q_2, R_2\}$, then there exists exactly one tile s of S' such that $v \in s \subseteq F$.

If e is an edge of F which does not subdivide, then e is in fact an edge of S'. Moreover the previous paragraph implies that if s is the tile of S' with $e \subseteq s \subseteq F$, then two edges of s are contained in disjoint edges of F. Therefore if condition (1) fails, then so does condition (2).

This proves Lemma 3.8. $\qquad\square$

Lemma 3.9. *The subdivision rule $\mathcal{R}$ usually satisfies condition* (2).

Proof. To prove this, we interpret condition (2) in terms of a directed graph G which was considered previously in [5]. The vertices of G are ordered pairs (e_1, e_2), where e_1 and e_2 are disjoint edges of the tile type t of $\mathcal{R}$. There exists a directed edge from the vertex (e_1, e_2) to the vertex (e_3, e_4) if and only if $\mathcal{R}(t)$ contains a tile s such that an edge of s with edge type e_3 is contained in e_1 and an edge of s with edge type e_4 is contained in e_2. Then $\mathcal{R}$ satisfies condition (2) if and only if G contains no directed cycles. We will show that G usually has no directed cycles.

To facilitate the discussion we number the edge types of $\mathcal{R}$ so that the types of the ten edges of F are numbered in counterclockwise order beginning with the two edges of F which contain β. So the two edges containing β have types 1 and 2, the two edges containing $1 + \beta$ have types 3 and 4, the two edges containing $1 + \tau + \beta$ have types 6 and 7, and the two edges containing $\tau + \beta$ have types 8 and 9.

Let e_1 and e_2 be distinct edges of t. Let s be any tile of S', and let E_1 and E_2 be the edges of s corresponding to e_1 and e_2. Then each of E_1 and E_2 might be an edge of S or half of an edge of S. We say that e_1 and e_2 are S-**adjacent** if the edges of S containing E_1 and E_2 have a vertex in common.

Let e_1 and e_2 be disjoint edges of t which are S-adjacent. We next show that (e_1, e_2) is usually not the terminal vertex of a directed edge of G. For this we assume that s is a tile of S' contained in F such that the edges E_1 and E_2 of s corresponding to e_1 and e_2 are contained in edges of F. Statement (1) of Lemma 3.7 shows that P_1 is usually not adjacent to either $-1 + \beta$ or β in S'. Using this and symmetry, we see that usually no edge of F is a half edge of S, that is, an edge of S' which is not an edge of S. Moreover, every edge of F has an endpoint which is a vertex of S. It follows that the edges of F which contain E_1 and E_2 usually have a vertex in common. So these edges of F are usually not disjoint. Thus (e_1, e_2) is usually not the terminal vertex of a directed edge of G. Therefore if e_1 and e_2 are S-adjacent, then (e_1, e_2) is usually not contained in a directed cycle of G.

Next let e_1 and e_2 be edges of t such that one of them has edge type either 1, 2, 3, or 4 and one of them has edge type either 6, 7, 8, or 9. The corresponding edges of F are contained in $\widehat{L}_1$ and $\widehat{L}_2$. In this case Lemma 3.3 implies that (e_1, e_2) is usually not the initial vertex of a directed edge of G. Therefore (e_1, e_2) is usually not contained in a directed cycle of G.

Now we prove that G usually does not contain a directed cycle. This will be done by contradiction. So suppose that (e_1, e_2) is the initial vertex and that (e_3, e_4) is the terminal vertex of an edge in a directed cycle of G. The previous two paragraphs show that e_3 and e_4 are usually not S-adjacent and usually one of them has type either 5 or 10.

Suppose that e_1 has type 10. The above discussion of S-adjacent edges implies that the type of e_2 is usually not 1, 2, 8, or 9. Statement (2) of Lemma 3.7 implies that the type of e_2 is usually not 3 or 4. Statement (3) of Lemma 3.7 and symmetry imply that the type of e_2 is usually not 6 or 7. So the type of e_2 is usually 5. Hence if e_1 corresponds to the edge of F in $\widehat{L}_3$, then e_2 usually corresponds to the edge of F in $\widehat{L}_4$. Similarly, if e_1 corresponds to the edge of F in $\widehat{L}_4$, then e_2 usually corresponds to the edge of F in $\widehat{L}_3$. The same is true for e_3 and e_4. So we may assume that one of e_1 and e_2 corresponds to the edge of F in $\widehat{L}_3$, the other corresponds to the edge of F in $\widehat{L}_4$, and e_3 and e_4 also correspond to two edges of

S' which are dual to edges of S^* in the direction of α^{-1}. Now Lemma 3.5 shows that this is usually impossible.

Thus $\mathcal{R}$ usually satisfies condition (2). This proves Lemma 3.9.   □

This completes the proof that the mesh of $\mathcal{R}$ usually approaches 0 combinatorially. The proof of Theorem 3.1 is now complete except for the cases in which $\alpha \in \{\tau, \tau + 1, \tau + 2\}$. So suppose that $\alpha \in \{\tau, \tau + 1, \tau + 2\}$. In the usual notation, $c = 1$ and $a \in \{0, 1, 2\}$. Lemma 1.4 implies that $d \in \{a, a + 1\}$. As in the proof of Lemma 3.7, we may assume that b is arbitrarily negative.

We proceed by indicating the form of a suitable fundamental domain F. Rather than working directly with F, we find it easier to work with $\psi(F)$. By construction, the points of the lattice $\psi(\alpha^{-1}\Lambda)$ are the midpoints of the edges of T which are not vertical. So there are essentially two possibilities for $\psi(\beta)$; either $\psi(\beta)$ is the midpoint of an edge of T with positive slope or $\psi(\beta)$ is the midpoint of an edge of T with negative slope. The same holds for $\psi(\tau + \beta)$. This leads to two possibilities for the form of the edges of $\psi(F)$ with edge types 1, 2, 3, 4 and two possibilities for the form of the edges of $\psi(F)$ with edge types 6, 7, 8, 9. In the following paragraphs we indicate the form of $\psi(F)$ by giving the forms of the edges of $\psi(F)$ with these edge types. There are, in general, many ways to complete the construction of $\psi(F)$. It is a straightforward matter to do so. One then checks that this defines a fundamental domain F for the orbifold fundamental group Γ. As in the outline of the proof of Theorem 3.1, the fundamental domain F has an edge pairing which is combinatorially equivalent to the original edge pairing on the tiles of S'. As before, we obtain a finite subdivision rule $\mathcal{R}$. Finally, one checks that the mesh of $\mathcal{R}$ approaches 0 combinatorially. Lemma 3.8 still holds, and so to check that the mesh of $\mathcal{R}$ approaches 0 combinatorially, it suffices to check condition (2). We leave the details to the reader.

Suppose that $(a, c) = (2, 1)$. Figure 6 shows four edge paths. We view these as building blocks for constructing fundamental domains. The first edge path in Figure 6 gives the form of the edges of $\psi(F)$ with edge types 6, 7, 8, and 9 when the slope of the edge of T containing $\psi(\tau + \beta)$ is positive. The second edge path occurs when this slope is negative. The third edge path gives the form of the edge of $\psi(F)$ with edge types 1, 2, 3, and 4 when the slope of the edge of T containing $\psi(\beta)$ is negative. The fourth edge path occurs when this slope is positive. The points $\psi(\beta)$, $\psi(1 + \beta)$, $\psi(\tau + \beta)$, and $\psi(1 + \tau + \beta)$ are indicated by small circles. Circles indicating $\psi(\beta)$ are larger than the others. As in Figure 4, the fundamental domain F contains points P_1, Q_1, R_1, P_2, Q_2, and R_2 and their images under ψ are indicated by large dots.

For example, suppose that $(a, b, c, d) = (2, -3, 1, 2)$ and that $\beta = 0$. The slope of the edge of T containing $\psi(\beta) = 0$ is negative, so we choose the third edge path in Figure 6. Since $\psi(\tau + \beta) = \psi(\tau) = \psi(b\alpha^{-1} + d\alpha^{-1}\tau) = \psi(-3\alpha^{-1} + 2\alpha^{-1}\tau)$, the slope of the edge of T containing $\psi(\tau + \beta)$ is positive, so we also choose the first edge path in Figure 6. This case is so simple that these two choices determine $\psi(F)$, and we obtain the fundamental domain shown in the left portion of Figure 7. The subdivision of the tile type of the corresponding finite subdivision rule is shown in Figure 8. The tile type is the Gosper snowflake.

For another example we take $(a, b, c, d) = (2, -6, 1, 2)$ with $\beta = 0$. In this example, the slope of the edge of T containing $\psi(\beta) = 0$ is still negative.
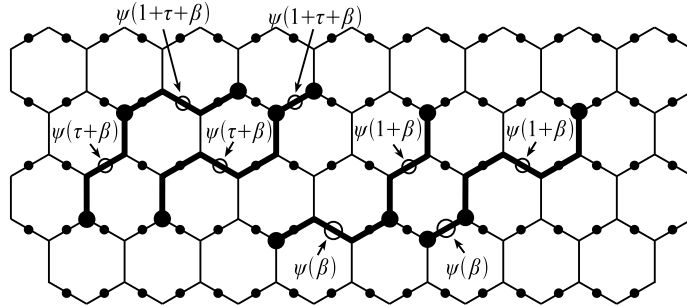
FIGURE 6. Building blocks for constructing fundamental domains for the cases in which $(a, c) = (2, 1)$.
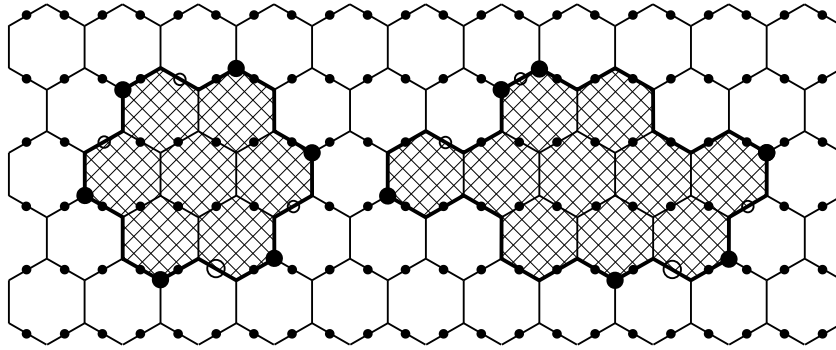

FIGURE 7. Fundamental domains for the cases in which $(a, b, c, d) = (2, -3, 1, 2)$ with $\beta = 0$ and $(a, b, c, d) = (2, -6, 1, 2)$ with $\beta = 0$.
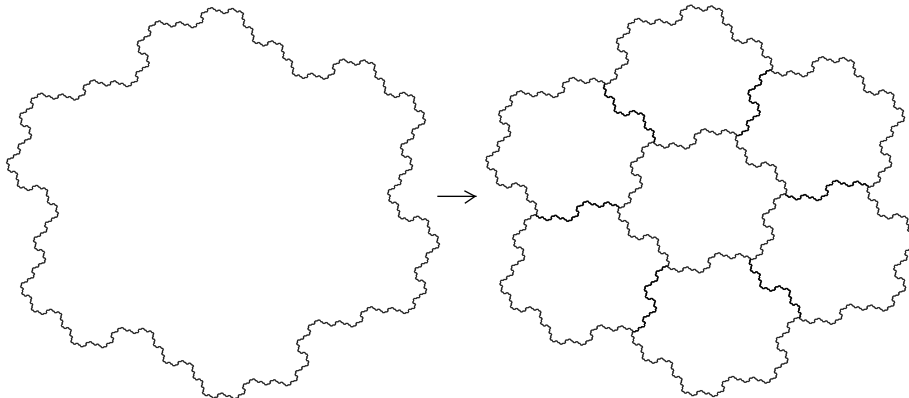

FIGURE 8. The subdivision of the tile type of the finite subdivision rule associated to the Lattès map with $(a, b, c, d) = (2, -3, 1, 2)$.
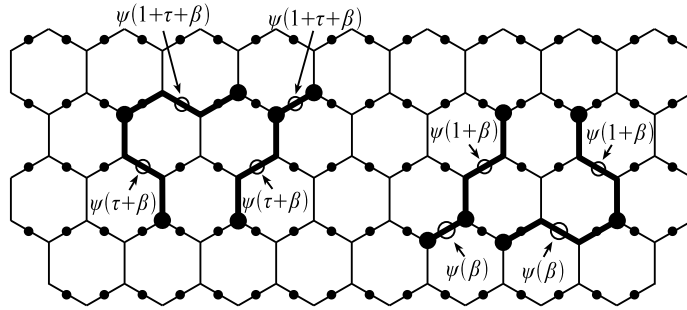
FIGURE 9. Building blocks for constructing fundamental domains
for the cases in which $(a, c) = (1, 1)$.

The slope of the edge of T containing $\psi(\tau + \beta) = \psi(\tau) = \psi(b\alpha^{-1} + d\alpha^{-1}\tau) = \psi(-6\alpha^{-1} + 2\alpha^{-1}\tau)$ is also negative, so we also choose the second edge path in Figure 6. The edge paths in Figure 6 determine 8 of the 10 edges of $\psi(F)$. The two remaining edges must be translates of one another, but they are not uniquely determined. One choice is given in the right portion of Figure 7. It is easy to check that the meshes of these finite subdivision rules approach 0 combinatorially. The situation is similar whenever $(a, c) = (2, 1)$ and the degree of f is sufficiently large.

Next suppose that $(a, c) = (1, 1)$. Figure 9 shows building blocks for this case. For example, suppose that $(a, b, c, d) = (1, -3, 1, 2)$ and $\beta = 0$. The slope of the edge of T containing $\psi(\beta) = 0$ is negative, so we choose the fourth edge path in Figure 9. Since $\psi(\tau + \beta) = \psi(\tau) = \psi(b\alpha^{-1} + d\alpha^{-1}\tau) = \psi(-3\alpha^{-1} + 2\alpha^{-1}\tau)$, the slope of the edge of T containing $\psi(\tau + \beta)$ is positive, so we also choose the second edge path in Figure 9. We complete these two edge paths in the straightforward way to the fundamental domain shown in the left portion of Figure 10. For another example we take $(a, b, c, d) = (1, -6, 1, 2)$ and $\beta = 0$. Since $\psi(\beta) = 0$, the slope of the edge of T containing $\psi(\beta)$ is negative, so we choose the fourth edge path in Figure 9. Since $\psi(\tau + \beta) = \psi(\tau) = \psi(-6\alpha^{-1} + 2\alpha^{-1}\tau)$, the slope of the edge of T containing $\psi(\tau + \beta)$ is negative, so we also choose the first edge path in Figure 9. One way to complete these two edge paths to a fundamental domain is shown in the right portion of Figure 10. Just as for the case in which $(a, c) = (2, 1)$, whenever $(a, c) = (1, 1)$ and the degree of f is sufficiently large we obtain a fundamental domain which determines a finite subdivision rule whose mesh approaches 0 combinatorially.

Finally suppose that $(a, c) = (0, 1)$. Figure 11 shows building blocks for this case. For example, suppose that $(a, b, c, d) = (0, -4, 1, 1)$ and $\beta = 0$. The slope of the edge containing $\psi(\beta) = 0$ is negative, so we choose the third edge path in Figure 11. Since $\psi(\tau + \beta) = \psi(\tau) = \psi(b\alpha^{-1} + d\alpha^{-1}\tau) = \psi(-4\alpha^{-1} + \alpha^{-1}\tau)$, the slope of the edge of T containing $\psi(\tau + \beta)$ is positive, so we also choose the second edge path in Figure 11. We complete these two edge paths as shown in the left portion of Figure 12 to obtain a fundamental domain. For another example we take $(a, b, c, d) = (0, -7, 1, 0)$ and $\beta = 1$. The slope of the edge containing $\psi(\beta) = \psi(1) = \psi(a\alpha^{-1} + c\alpha^{-1}\tau) = \psi(\alpha^{-1}\tau)$ is positive, so we choose the fourth edge path in Figure 11. Since $\psi(\tau + \beta) = \psi(\tau + 1) = \psi((a+b)\alpha^{-1} + (c+d)\alpha^{-1}\tau) = \psi(-7\alpha^{-1} + \alpha^{-1}\tau)$, the slope of the edge of T containing $\psi(\tau + \beta)$ is negative, so
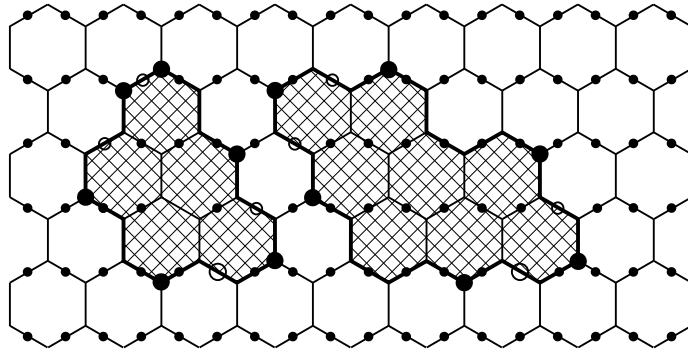
FIGURE 10. Fundamental domains for the cases in which $(a, b, c, d) = (1, -3, 1, 2)$ with $\beta = 0$ and $(a, b, c, d) = (1, -6, 1, 2)$ with $\beta = 0$.
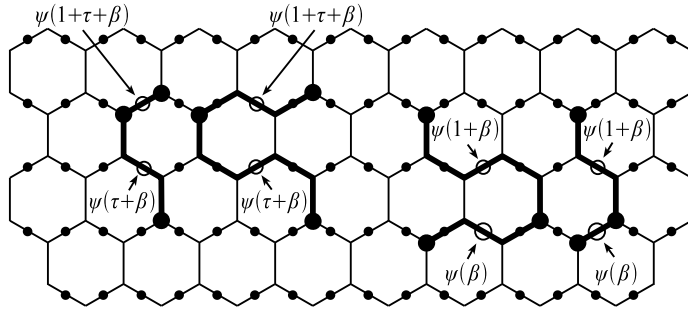

FIGURE 11. Building blocks for constructing fundamental domains for the cases in which $(a, c) = (0, 1)$.
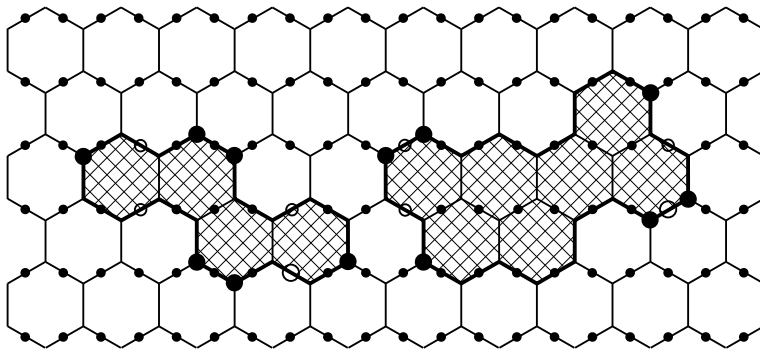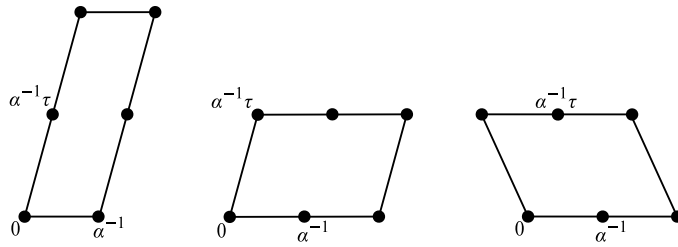

FIGURE 12. Fundamental domains for the cases in which $(a, b, c, d) = (0, -4, 1, 1)$ with $\beta = 0$ and $(a, b, c, d) = (0, -7, 1, 0)$ with $\beta = 1$.

FIGURE 13. The parallelogram P.

we choose the first edge path in Figure 11. We complete these two edge paths as shown in the right portion of Figure 12 to obtain a fundamental domain. As before, whenever $(a, c) = (0, 1)$ and the degree of f is sufficiently large, we obtain a fundamental domain which determines a finite subdivision rule whose mesh approaches 0 combinatorially.

This completes the proof of Theorem 3.1. $\qquad\square$

Theorem 3.10. *Every nonrigid Lattès map is the subdivision map of a finite subdivision rule with one tile type, and the tile has the form of type* 1 *in Figure* 1.

Proof. Let f be a nonrigid Lattès map. Let $\widetilde{f}(z) = \alpha z + \beta$ be a lift of f. Then $\alpha \in \mathbb{Z}$, and we may assume that $\alpha \geq 2$. Let 1 and τ form a basis of the usual lattice Λ with $\mathrm{Im}(\tau) > 0$. We may add any element of 2Λ to β, so we assume that $\beta \in \{0, 1, \tau, 1 + \tau\}$.

Recall that Figure 1 shows five types of planar tilings. We construct a tiling S' of the plane with type 1 so that every tile of S' is a parallelogram. Furthermore, one of the parallelograms of S' is chosen to be as in Figure 13. If $\beta = 1$, then we choose the leftmost parallelogram. If $\beta = \tau$, then we choose the middle parallelogram. If $\beta = 1 + \tau$, then we choose the rightmost parallelogram. If $\beta = 0$, then we choose any of these. Let P denote the chosen parallelogram.

The tiling S' is constructed so that β is a vertex with valence 4. From this it follows that $F = \widetilde{f}(P)$ is a union of tiles of S'. The parallelogram P is really a hexagon because it has two vertices with valence 2. The two edges which contain one of the vertices with valence 2 are interchanged by a rotation of order 2. This pairs four of the six edges of P. The remaining two edges of P are paired by a translation. In the same way, F has six edges which are paired so that $\widetilde{f}$ respects these edge pairings. As in the proof of Theorem 3.1, we obtain a finite subdivision rule $\mathcal{R}$ with one tile type. Not as in the proof of Theorem 3.1, in the present situation f is already the subdivision map of $\mathcal{R}$.

This proves Theorem 3.10. $\qquad\square$

Corollary 3.11. *Every map in almost every analytic conjugacy class of Lattès maps is the subdivision map of a finite subdivision rule with one tile type.*

Proof. This follows from Theorem 3.1, Theorem 3.10, and the last statement in Section 1; namely, that there are only finitely many analytic conjugacy classes of rigid Lattès maps with bounded degree. $\qquad\square$

4. An exceptional quadratic Lattès map

This section deals with the Lattès map f with lift $\widetilde{f}(z) = \alpha z$, where $\alpha = (1 + \sqrt{-7})/2$ and the usual lattice Λ is generated by 1 and α. We prove that f is not the subdivision map of a finite subdivision rule with one tile type, f is not the subdivision map of a finite subdivision rule with two tile types and 1-skeleton of the subdivision complex homeomorphic to a circle, and f is the subdivision map of a finite subdivision rule with two tile types.

The minimal quadratic equation satisfied by α is $\alpha^2 - \alpha + 2 = 0$. Since the constant term is $\alpha\overline{\alpha} = 2$, the map f is indeed quadratic. Moreover $\widetilde{f}(0) = 0$, $\widetilde{f}(1) = \alpha$, $\widetilde{f}(\alpha) = \alpha^2 \equiv \alpha \mod 2\Lambda$ and $\widetilde{f}(\alpha + 1) = \alpha^2 + \alpha \equiv 0 \mod 2\Lambda$. Let $A = \wp(\alpha + 1)$, $B = \wp(1)$, $X = \wp(0)$, and $Y = \wp(\alpha)$. Then $f(A) = X$, $f(B) = Y$, $f(X) = X$, and $f(Y) = Y$. Since A, B, X, and Y are the only postcritical points of f, one critical point of f maps to A and one critical point of f maps to B. We normalize so that ∞ and 0 are the critical points of f with $f(\infty) = A$ and $f(0) = B$. Figure 14 shows the mapping scheme for f. Integers next to arrows are local mapping degrees.
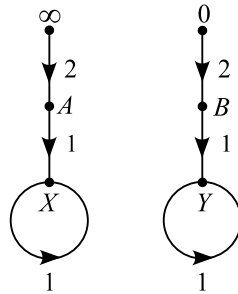

FIGURE 14. The mapping scheme of f.

We can normalize f so that $f(z) = \frac{Az^2 + B}{z^2 + 1}$. Since $X = f(X) = f(A)$, $Y = f(Y) = f(B)$, and f is even, $X = -A$ and $Y = -B$. Hence $\frac{A^3 + B}{A^2 + 1} = -A$ and $\frac{AB^2 + B}{B^2 + 1} = -B$. This gives $B = -2A^3 - A$ and $AB + 1 = -B^2 - 1$. Substituting $B = -A(2A^2 + 1)$ into the equation $AB + 1 = -B^2 - 1$ yields $0 = 2A^6 + A^4 + 1 = (A^2 + 1)(2A^4 - A^2 + 1)$. If $A = \pm i$, then $B = A$, so $2A^4 - A^2 + 1 = 0$. Hence $f(z) = \frac{Az^2 + B}{z^2 + 1}$, where $A^2 = \frac{1 \pm \sqrt{-7}}{4}$ and $B = -A(2A^2 + 1)$. That is, either $2A^2 = \alpha$ or $2A^2 = \overline{\alpha}$. (By choosing 0 and ∞ to be the fixed postcritical points, in [7, Section B.4] Milnor gives the alternate form $f(z) = z\frac{z + \alpha^2}{\alpha^2 z + 1}$ for f.)

We first prove by contradiction that f is not the subdivision map of a finite subdivision rule with one tile type. Suppose that f is the subdivision map of a finite subdivision rule $\mathcal{R}$ with one tile type. Let T be the 1-skeleton of the subdivision complex of $\mathcal{R}$. By Lemma 2.1, we may and do assume that the only vertices of T with valence 1 or 2 are postcritical points.

Lemma 4.1. *No vertex of T is a critical point of f.*

Proof. If a critical point of f is a vertex of T, then it is an accidental vertex (not a postcritical point). A tree with type 1 or type 2 has no accidental vertices. The

postcritical vertices of trees of type 3 and type 5 all have valence 1, and so a critical vertex would have valence at most 2, which is impossible. Suppose that T has type 4. Then T has just one accidental vertex, the vertex with valence 3. By symmetry we may assume that this vertex is ∞. As before, since ∞ has valence greater than 2, $f(\infty) = A$ must have valence greater than 1. So A is the vertex with valence 2. But then $f(A)$ also has valence at least 2 because the degree of f at A is 1. This is impossible. This shows that no vertex of T is a critical point of f. $\qquad\square$

Lemma 4.2. *Let p and q be distinct vertices of T such that $f(p), f(q) \in \{p, q\}$. Then the open segment (p, q) of T contains a critical point.*

Proof. If f maps both p and q to the same vertex, then because f maps the segment (p, q) into the tree T, the image of (p, q) under f must be folded somewhere. Such a fold is the image of a critical point. So (p, q) contains a critical point in this case.

Now suppose that $f(p) = p$ and $f(q) = q$ and that (p, q) contains no critical point of f. Then f does not fold the segment $[p, q]$ anywhere. It follows that f maps $[p, q]$ homeomorphically onto $[p, q]$. But then the edges in $[p, q]$ never subdivide. This implies that the mesh of our finite subdivision rule does not go to 0, contrary to the fact that f is an expanding map. So the lemma holds if $f(p) = p$ and $f(q) = q$.

Finally, the case in which $f(p) = q$ and $f(q) = p$ can be proved just as in the case in which $f(p) = p$ and $f(q) = q$.

This proves Lemma 4.2. $\qquad\square$

Lemma 4.3. *Neither A nor B is in the closed segment $[X, Y]$.*

Proof. Proceeding by contradiction, we may assume by symmetry that $A \in [X, Y]$. Since $f(A) = X$ and $f(Y) = Y$ there exists a point in (A, Y) which maps to A. This point must be ∞. Since f maps A to X with degree 1, $f([X, Y])$ contains a nontrivial segment which meets $[X, Y]$ at X. But then there exists a nontrivial segment s in T which contains a vertex of T with valence 1 and has $s \cap [X, Y] = X$. This vertex with valence 1 must be B. Since $f(X) = X$ and $f(B) = Y$, it follows that s contains a preimage of A, namely, $\infty \in s$. This is impossible. $\qquad\square$

Now we apply Lemma 4.2 to see that there exists a critical point in $[X, Y]$. By symmetry we may assume that $\infty \in [X, Y]$. Lemma 4.3 shows that $A \notin [X, Y]$. Let v be the vertex of $[X, Y]$ which is closest to A in T. If $v = Y$, then because $f(X) = X$ and $f(\infty) = A$, f maps some point in $(X, \infty) \subseteq (X, Y)$ to Y. This point must be B, which is impossible by Lemma 4.3. Hence $v \neq Y$. Likewise $v \neq X$. So v is a vertex in (X, Y) with valence at least 3.

In this paragraph we assume that v is the only vertex of T with valence at least 3. As in the previous paragraph we see that v is the vertex in $[X, Y]$ which is closest to both A and B. Lemma 4.1 implies that v is not a critical point. Hence $f(v)$ has valence at least 3, and so $f(v) = v$. Since $f(X) = X$ and $f(v) = v$, Lemma 4.2 implies that (X, v) contains a critical point p. Since the critical point maps to A or B, (X, v) contains a preimage of v. Likewise (v, Y) contains a critical point q, and (q, Y) contains a preimage of v. We now have three points mapping to v, which is impossible.

To complete this argument we assume that T contains two vertices u and v with valence 3, that one of them, v, is in (X, Y), and that either $u \in (X, Y)$ or v is the vertex in $[X, Y]$ closest to u. As in the previous paragraph, we see that

$f(u), f(v) \in \{u, v\}$. First suppose that $f(u) = u$ and $f(v) = v$. Lemma 4.2 implies that (u, v) contains a critical point; so does the segment from X to $[u, v]$ as well as the segment from Y to $[u, v]$. This is impossible. Next suppose that $f(u) = v$ and $f(v) = u$. Let $w \in (u, v)$. Then (u, v) contains a preimage of w. So does the segment from X to $[u, v]$ as well as the segment from Y to $[u, v]$. We now have three preimages of w, which is impossible. We are left with the case in which $f(u) = f(v) \in \{u, v\}$. Suppose in addition that $u \notin [X, Y]$. Recall that ∞ is a critical point in $[X, Y]$. Choose $z \in \{X, Y\}$ so that $v \notin [z, \infty]$. Then the image of $[z, \infty]$ covers $[u, v]$, and so $[z, \infty]$ contains a preimage of $f(u) = f(v)$, giving three preimages, which is impossible. So $u \in [X, Y]$. Now choose $z \in \{X, Y\}$ so that neither u nor v lies in $(z, f(v))$. Lemma 4.2 implies that $(z, f(v))$ contains a critical point p. But then (z, p) contains a preimage of $f(v)$, giving a third preimage of $f(v)$, which is once more impossible.

The proof that f is not given by a finite subdivision rule with one tile type is now complete.

We next prove by contradiction that f is not given by a finite subdivision rule with two tile types with 1-skeleton of the subdivision complex homeomorphic to a circle. Suppose that f is given by a finite subdivision rule with two tile types with 1-skeleton G of the subdivision complex homeomorphic to a circle. If some open edge of G is not in $f(G)$, then we delete that open edge to obtain a finite subdivision rule for f with one tile type. Since such a finite subdivision rule does not exist, it follows that $f(G) = G$.

In this paragraph we assume that the restriction of f to G is locally injective. Then the restriction of f to G is a covering map with degree either 1 or 2. If this degree is 2, then every point in G has two preimages, but both A and B have at most one preimage in G. If this degree is 1, then the restriction of f to G is a homeomorphism. In this case the edges of G do not subdivide and the mesh of our finite subdivision rule fails to go to 0. So the restriction of f to G is not locally injective.

Since the restriction of f to G is not locally injective, G contains a critical point. We may thus assume that f is not locally injective at $\infty \in G$. Let e_1 and e_2 be the edges of G which contain A and suppose that e_1 is doubly covered by f near ∞. But then e_2 is not in $f(G)$, for if it is, then some point other than ∞ maps to A, which is not the case. Thus f is not given by a finite subdivision rule with two tile types with 1-skeleton homeomorphic to a circle.

We finally prove that f is the subdivision map of a finite subdivision rule with two tile types. The left portion of Figure 15 shows a parallelogram which is a fundamental domain F for the orbifold fundamental group Γ of f. In addition to
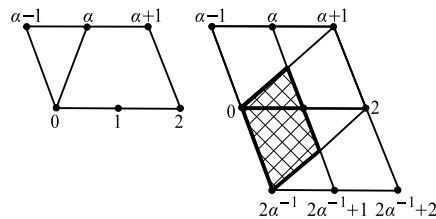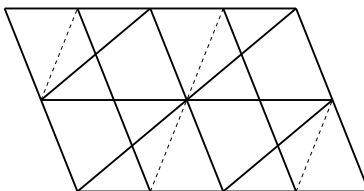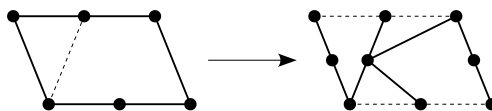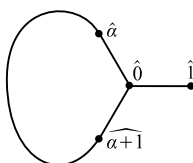

FIGURE 15. The fundamental domain F and $\widetilde{f}^{-1}(F)$.

FIGURE 16. A small part of the tiling T and $\widetilde{f}^{-1}(T)$.

FIGURE 17. The finite subdivision rule $\mathcal{R}$.

FIGURE 18. The subdivision complex for f.

the vertices at the corners of F, there are vertices at 1 and α. There is also an extra edge joining 0 and α. The hatched region in the right portion of Figure 15 shows $\widetilde{f}^{-1}(F)$. To verify this, it is useful to note that $\alpha^2 - \alpha + 2 = 0$ implies that $\alpha(1 - \alpha) = 2$ and so $1 - \alpha = 2\alpha^{-1}$. The last identity shows that $2\alpha^{-1}$ is drawn correctly. Moreover, the fact that 0, 2 and $\alpha + 1$ are vertices of the parallelogram F implies that 0, $2\alpha^{-1}$ and $\alpha^{-1} + 1$ are vertices of the parallelogram $\widetilde{f}^{-1}(F)$. Thus $\widetilde{f}^{-1}(F)$ is drawn correctly. Let T be the tiling of the plane by the images of F under Γ. Figure 16 shows a part of T drawn with dashed line segments and part of $\widetilde{f}^{-1}(T)$ drawn with solid line segments. Few dashed line segments of T are visible because most of them are obscured by the line segments of $\widetilde{f}^{-1}(T)$. Now it is clear that there exists a Γ-equivariant isotopy from T to $\widetilde{f}^{-1}(T)$ rel Λ. This determines a finite subdivision rule $\mathcal{R}$ with two tile types, one triangle and one pentagon. Figure 17 gives a schematic description of the subdivisions of the tile types of $\mathcal{R}$. The edges drawn with dashes in Figure 17 are the edges whose edge types correspond to the line segments drawn with dashes in Figure 16. One easily checks that $\mathcal{R}$ has bounded valence and mesh approaching 0 combinatorially. It now follows as in the proof of the main theorem of [4] that f is the subdivision map of a finite subdivision rule isomorphic to $\mathcal{R}$. A stereographic projection of the subdivision complex for f is shown in Figure 18 with hats indicating images of 0, 1, α, and $\alpha + 1$.

References

[1] M. Bonk and D. Meyer, *Expanding Thurston maps*, in preparation.

[2] Z. I. Borevich and I. R. Shafarevich, *Number Theory*, Academic Press, New York and London, 1966. MR0195803 (33:4001)

[3] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Expansion complexes for finite subdivision rules I*, Conform. Geom. Dyn. **10** (2006), 63–99. MR2218641 (2007c:30048)

[4] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Constructing subdivision rules from rational maps*, Conform. Geom. Dyn. **11** (2007), 128–136. MR2329140 (2008i:37093)

[5] J. W. Cannon, W. J. Floyd, W. R. Parry, and K. M. Pilgrim, *Subdivision rules and virtual endomorphisms*, Geom. Dedicata **141** (2009), 181–195. MR2520071

[6] A. Douady and J. H. Hubbard, *A proof of Thurston's topological characterization of rational functions*, Acta Math. **171** (1993), 263–297. MR1251582 (94j:58143)

[7] J. Milnor, *Pasting together Julia sets: a worked out example of mating*, Experimental Math. **13** (2004), 55–92. MR2065568 (2005c:37087)

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

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

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