

Eliminating Thurston obstructions and controlling dynamics on curves

MARIO BONK[†], MIKHAIL HLUSHCHANKA^{‡§} and ANNINA ISELI^{¶||}

[†] *Department of Mathematics, University of California, Los Angeles, CA 90095, USA*
(e-mail: mbonk@math.ucla.edu)

[‡] *Korteweg-de Vries Instituut voor Wiskunde, Universiteit van Amsterdam, 1090 GE Amsterdam, The Netherlands*

[§] *Mathematisch Instituut, Universiteit Utrecht, 3508 TA Utrecht, The Netherlands*
(e-mail: mikhail.hlushchanka@gmail.com)

[¶] *Department of Mathematics, University of Fribourg, 1700 Fribourg, Switzerland*
^{||} *Institute of Mathematics, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*
(e-mail: annina.iseli@epfl.ch)

(Received 1 February 2022 and accepted in revised form 7 November 2023)

Abstract. Every Thurston map $f: S^2 \rightarrow S^2$ on a 2-sphere S^2 induces a pull-back operation on Jordan curves $\alpha \subset S^2 \setminus P_f$, where P_f is the postcritical set of f . Here the isotopy class $[f^{-1}(\alpha)]$ (relative to P_f) only depends on the isotopy class $[\alpha]$. We study this operation for Thurston maps with four postcritical points. In this case, a Thurston obstruction for the map f can be seen as a fixed point of the pull-back operation. We show that if a Thurston map f with a hyperbolic orbifold and four postcritical points has a Thurston obstruction, then one can ‘blow up’ suitable arcs in the underlying 2-sphere and construct a new Thurston map \widehat{f} for which this obstruction is eliminated. We prove that no other obstruction arises and so \widehat{f} is realized by a rational map. In particular, this allows for the combinatorial construction of a large class of rational Thurston maps with four postcritical points. We also study the dynamics of the pull-back operation under iteration. We exhibit a subclass of our rational Thurston maps with four postcritical points for which we can give positive answer to the global curve attractor problem.

Key words: Thurston maps, obstructions, Lattès maps, intersection numbers, curve attractor

2020 Mathematics Subject Classification: 37F20 (Primary); 37F10(Secondary)

Contents

1	Introduction	2455
1.1	Blowing up Lattès maps	2457
1.2	The global curve attractor problem	2460
1.3	Organization of this paper	2463
2	Preliminaries	2463
2.1	Notation and basic concepts	2463
2.2	Branched covering maps	2465
2.3	Planar embedded graphs	2466
2.4	The Euclidean square pillow	2467
2.5	Isotopies and intersection numbers	2468
2.6	Jordan curves in spheres with four marked points	2470
3	Thurston maps	2473
3.1	The $(n \times n)$ -Lattès map	2475
3.2	Thurston's characterization of rational maps	2476
4	Blowing up arcs	2479
4.1	The general construction	2479
4.2	Blowing up the $(n \times n)$ -Lattès map	2483
5	Realizing blown-up Lattès maps	2487
5.1	The horizontal and vertical curves	2488
5.2	Ruling out other obstructions	2489
6	Essential circuit length	2494
7	Eliminating obstructions by blowing up arcs	2501
8	Global curve attractors	2506
9	Further discussion	2514
9.1	Julia sets of blown-up Lattès maps	2514
9.2	The global curve attractor problem	2515
9.3	Twisting problems	2516
9.4	Thurston maps with more than four postcritical points	2517
9.5	Other combinatorial constructions of rational maps	2517
	Acknowledgements	2517
A	Appendix. Isotopy classes of Jordan curves in spheres with four marked points	2518
	References	2531

1. Introduction

In this paper, we consider Thurston maps and the dynamics of the induced pull-back operation on Jordan curves on the underlying 2-sphere. By definition, a *Thurston map* is a branched covering map $f: S^2 \rightarrow S^2$ on a topological 2-sphere S^2 such that f is not a homeomorphism and every critical point of f (points where f is not a local homeomorphism) has a finite orbit under iteration of f . These maps are named after William Thurston who introduced them in his quest for a better understanding of the dynamics of postcritically-finite rational maps on the Riemann sphere. We refer to [BM17, Ch. 2] for general background on Thurston maps and related concepts.

For a branched covering map $f: S^2 \rightarrow S^2$, we denote by C_f the set of critical points of f and by f^n the n th iterate of f for $n \in \mathbb{N}$. Then the postcritical set of f is defined as

$$P_f = \bigcup_{n \in \mathbb{N}} \{f^n(c) : c \in C_f\}.$$

For a Thurston map f , this set has finite cardinality $2 \leq \#P_f < \infty$ (for the first inequality, see [BM17, Corollary 2.13]).

A Thurston map f often admits a description in purely combinatorial-topological terms. In this context, it is an interesting question whether f can be *realized* (in a suitable sense) by a rational map with the same combinatorics. Roughly speaking, this means that f is conjugate to a rational map ‘up to isotopy’ (see §3 for the precise definition).

It is not hard to see that each Thurston map with two or three postcritical points is realized. The situation is much more complicated for Thurston maps f with $\#P_f \geq 4$. William Thurston found a necessary and sufficient condition when a Thurston map can be realized by a rational map [DH93]. Namely, if f has an associated hyperbolic orbifold (this is always true apart from some well-understood exceptional maps), then f is realized if and only if f has no (*Thurston*) *obstruction*. Such an obstruction is given by a finite collection of disjoint Jordan curves in $S^2 \setminus P_f$ (up to isotopy) with certain invariance properties (see §3.2 for more discussion).

The ‘if’ part of this statement gives a positive criterion for f to be realized, but it is very hard to apply in practice, because, at least in principle, it involves the verification of infinitely many conditions for the map f . For this reason, in each individual case, a successful verification for a map, or a class of maps, is difficult and usually constitutes an interesting result in its own right.

We mention two results in this direction. The first one is the ‘arcs intersecting obstructions’ theorem by Pilgrim and Tan Lei [PL98, Theorem 3.2] that gives control on the position of an obstruction and has many applications in holomorphic dynamics (see, for instance, [DMRS19, PL98]). The other one is the ‘mating criterion’ by Tan Lei, Rees, and Shishikura that addresses the question when two postcritically-finite quadratic polynomials can be topologically glued together to form a rational map (see [Lei92, Ree92, Shi00]).

The investigation of obstructions of a Thurston map $f: S^2 \rightarrow S^2$ is closely related to the study of the *pull-back operation* on Jordan curves. It is easy to show that if $\alpha \subset S^2 \setminus P_f$ is a Jordan curve, then the isotopy class $[f^{-1}(\alpha)]$ (relative to P_f) only depends on the isotopy class $[\alpha]$ (see Lemma 3.4). Intuitively, the number of postcritical points of a Thurston map can be seen as a measure of its combinatorial complexity. In this paper, we focus on the simplest non-trivial case, namely Thurston maps f with $\#P_f = 4$. In this case, the pull-back operation gives rise to a well-defined map, the *slope map*, on these isotopy classes $[\alpha]$ (we will discuss this in more detail below). The search for obstructions of f amounts to understanding the fixed points of the slope map.

There exist various natural constructions that allow one to *combine* or *modify* given (rational) Thurston maps to obtain a new dynamical system. The most studied constructions are *mating* (see [Lei92, SL00]), *tuning* (see [Ree92]), and *capture* (see [Hea88, Lei97]). In this paper, we study the operation of *blowing up arcs*, originally introduced by Pilgrim and Tan Lei in [PL98]. This operation can be applied to an arbitrary Thurston

map f and results in a new Thurston map \widehat{f} that is of higher degree, but combinatorially closely related to the original map f . In particular, f and \widehat{f} have the same set of postcritical points and the same dynamics on them. Nevertheless, the dynamical behavior of Jordan curves under the pull-back operation for the original map f and the new map \widehat{f} may differ drastically.

We show that if a Thurston map $f: S^2 \rightarrow S^2$ with $\#P_f = 4$ has an obstruction α , then one can naturally modify f by blowing up certain arcs to produce a new Thurston map \widehat{f} for which this obstruction α is eliminated. The main result of this paper is the fact that then, no new obstructions arise for \widehat{f} and so it is realized by a rational map.

THEOREM 1.1. *Let $f: S^2 \rightarrow S^2$ be a Thurston map with $\#P_f = 4$ and a hyperbolic orbifold. Suppose that f has an obstruction represented by a Jordan curve $\alpha \subset S^2 \setminus P_f$, and $E \neq \emptyset$ is a finite set of arcs in $(S^2, f^{-1}(P_f))$ that satisfy the α -restricted blow-up conditions.*

Let \widehat{f} be a Thurston map obtained from f by blowing up arcs in E (with some multiplicities) so that $\lambda_{\widehat{f}}(\alpha) < 1$. Then \widehat{f} is realized by a rational map.

The technical verbiage and the notation in this formulation will be explained in subsequent sections (see in particular equation (3.3) for the definition of the ‘eigenvalue’ $\lambda_{\widehat{f}}(\alpha)$ and Definition 6.5 for α -restricted blow-up conditions).

Recently, Dylan Thurston provided a positive characterization when a Thurston map is realized, at least in the case when each critical point eventually lands in a critical cycle under iteration. He proved that such a Thurston map f is realized by a rational map if and only if there is an ‘elastic spine’ (that is, a planar embedded graph in $S^2 \setminus P_f$ with a suitable metric on it) that gets ‘looser’ under backwards iteration (see [Thu16, Thu20] for more details). In concrete cases, especially for Thurston maps that should be realized by rational maps with Julia sets homeomorphic to the Sierpiński carpet, the application of Dylan Thurston’s criterion is not so straightforward. Moreover, his criterion is only valid for Thurston maps with periodic critical points. In contrast, for some maps for which Dylan Thurston’s criterion is not applicable or hard to apply, Theorem 1.1 can be used to verify that the maps are realized. In particular, many maps obtained by blowing up Lattès maps (see below) are of this type.

1.1. Blowing up Lattès maps. We will now discuss a special case of Theorem 1.1 in detail to give the reader some intuition for the geometric ideas behind this statement and its proof.

Let \mathbb{P} be a *pillow* obtained from two copies of the unit square $[0, 1]^2 \subset \mathbb{R}^2 \cong \mathbb{C}$ glued together along their boundaries. We consider the two copies of $[0, 1]^2$ in \mathbb{P} as the front and back side of \mathbb{P} and call them the *tiles of level 0* or simply *0-tiles*. We denote by $A := (0, 0) \in \mathbb{P}$ the lower left corner of \mathbb{P} (see the right part of Figure 1). The pillow \mathbb{P} is a topological 2-sphere. Actually, if we consider \mathbb{P} as an abstract polyhedral surface, then \mathbb{P} carries a conformal structure making \mathbb{P} conformally equivalent to the Riemann sphere $\widehat{\mathbb{C}}$. See §2.4 for more discussion.

We now fix $n \in \mathbb{N}$ with $n \geq 2$. We subdivide each of the two 0-tiles of \mathbb{P} into n^2 small squares of sidelength $1/n$, called the *1-tiles*. We color these 1-tiles in a checkerboard




FIGURE 1. The (4×4) -Lattès map.




FIGURE 2. Gluing in a flap.

fashion black and white so that the 1-tile in the front 0-tile that contains the vertex A on its boundary is colored white (see the left part of Figure 1). We map this white 1-tile to the front 0-tile of the right-hand pillow by an orientation-preserving Euclidean similarity that fixes the vertex A . This similarity scales distances by the factor n . We can uniquely extend the similarity by a successive Schwarz reflection process to the whole pillow \mathbb{P} to obtain a continuous map $\mathcal{L}_n : \mathbb{P} \rightarrow \mathbb{P}$. Then on each 1-tile S , the map \mathcal{L}_n is a Euclidean similarity that sends S to the front or back 0-tile of \mathbb{P} depending on whether S is white or black. We call \mathcal{L}_n the $(n \times n)$ -Lattès map, because under a suitable conformal equivalence $\mathbb{P} \cong \widehat{\mathbb{C}}$, the map \mathcal{L}_n is conjugate to a rational map obtained from n -multiplication of a Weierstrass \wp -function. See Figure 1 for an illustration of the map \mathcal{L}_4 . Here, the marked points on the left pillow \mathbb{P} (the domain of the map) correspond to the preimage points $\mathcal{L}_4^{-1}(A)$. Note that there is exactly one preimage of A in the interior of the back side of the pillow.

It is easy to see that the $(n \times n)$ -Lattès map $\mathcal{L}_n : \mathbb{P} \rightarrow \mathbb{P}$ is a Thurston map with four postcritical points, namely, the four corners of the pillow \mathbb{P} . The map \mathcal{L}_n is realized by a rational map, because it is even conjugate to such a map.

We now modify the map \mathcal{L}_n by gluing in vertical or horizontal flaps to \mathbb{P} . This is a special case of the more general construction of blowing up arcs mentioned above. We will describe this in detail in §4, but will illustrate the procedure in Figure 2, where we show how to glue in one horizontal flap.

We cut the pillow \mathbb{P} open along a horizontal side e of one of the 1-tiles. Note that in this process, e is ‘doubled’ into two arcs e' and e'' with common endpoints. We then take two disjoint copies of the Euclidean square $[0, 1/n]^2$ and identify them along three corresponding sides to obtain a flap F . It has two ‘free’ sides on its boundary. We glue each free side to one of the arcs e' and e'' of the cut in the obvious way.




FIGURE 3. An example of a map $\widehat{\mathcal{L}}$ obtained from \mathcal{L}_4 by gluing in flaps.

In general, one can repeat this construction and glue several flaps at the location given by the arc e . We assume that this has been done simultaneously for $n_h \geq 0$ flaps along horizontal edges and $n_v \geq 0$ flaps along vertical edges. By this procedure, we obtain a ‘flapped’ pillow $\widehat{\mathbb{P}}$, which is still a topological 2-sphere (see the left part of Figure 3). By construction, it is tiled by $2n^2 + 2(n_h + n_v)$ squares of sidelength $1/n$, which we consider as the 1-tiles of $\widehat{\mathbb{P}}$. The checkerboard coloring of the base surface \mathbb{P} extends in a unique way to the new surface $\widehat{\mathbb{P}}$. The original $(n \times n)$ -Lattès map $\mathcal{L}_n : \mathbb{P} \rightarrow \mathbb{P}$ can naturally be ‘extended’ to a continuous map $\widehat{\mathcal{L}} : \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ so that each 1-tile S of $\widehat{\mathbb{P}}$ is mapped to the front or back 0-tile of \mathbb{P} (depending on the color of S) by a Euclidean similarity scaling distances by the factor n . See Figure 3 for an illustration of a map $\widehat{\mathcal{L}}$ obtained from the Lattès map \mathcal{L}_4 by gluing in flaps at a vertical and a horizontal edge. Similarly as in Figure 1, on the left, we marked the preimages of A under $\widehat{\mathcal{L}}$.

To obtain a Thurston map $f : \mathbb{P} \rightarrow \mathbb{P}$ from this construction, we need to choose a homeomorphism $\phi : \widehat{\mathbb{P}} \rightarrow \mathbb{P}$. Roughly speaking, ϕ is a homeomorphism that identifies $\widehat{\mathbb{P}}$ with \mathbb{P} and fixes each corner of the pillow. The precise choice of ϕ is somewhat technical and so we refer to §4.2 for the details. Then, one easily observes that the postcritical set P_f consists of the four corners of the pillow \mathbb{P} . The map f is uniquely determined up to Thurston equivalence (see Definition 3.2) independently of the choice of ϕ under suitable restrictions. We refer to f as a Thurston map obtained from the $(n \times n)$ -Lattès map by gluing n_h horizontal and n_v vertical flaps to \mathbb{P} .

Now the following statement is true. As we will explain, it can be seen as a special case of our main result.

THEOREM 1.2. *Let $n \in \mathbb{N}$ with $n \geq 2$ and $f : \mathbb{P} \rightarrow \mathbb{P}$ be a Thurston map obtained from the $(n \times n)$ -Lattès map \mathcal{L}_n by gluing $n_h \geq 0$ horizontal and $n_v \geq 0$ vertical flaps to \mathbb{P} , where $n_h + n_v > 0$. Then the map f has a hyperbolic orbifold. It has an obstruction if and only if $n_h = 0$ or $n_v = 0$. In particular, if $n_h > 0$ and $n_v > 0$, then f is realized by a rational map.*

If $n_h = n_v = 0$, then no flaps were glued to \mathbb{P} and the map f coincides with the original $(n \times n)$ -Lattès map \mathcal{L}_n (strictly speaking, only if we choose the homeomorphism ϕ used in the construction above to be the identity on \mathbb{P} , as we may). Then $f = \mathcal{L}_n$ has a parabolic orbifold. Therefore, Thurston’s criterion as formulated in §3.2 does not apply.

If $n_h = 0$ or $n_v = 0$, but $n_h + n_v > 0$ as in Theorem 1.2, then it is immediate to see that f has an obstruction (see §5.1) and therefore f cannot be realized by a rational map. So the interesting part of Theorem 1.2 is the claim that if $n_h > 0$ and $n_v > 0$, then f has no obstruction.

Even though Theorem 1.2 follows from our more general statement formulated in Theorem 1.1, we will give a complete proof. We will argue by contradiction and assume that a map f with $n_h > 0$ and $n_v > 0$ has an obstruction. In principle, there are infinitely many candidates represented by essential isotopy classes of Jordan curves $\alpha \subset \mathbb{P} \setminus P_f$. These isotopy classes in turn are distinguished by different slopes in $\widehat{\mathbb{Q}} = \mathbb{Q} \cup \{\infty\}$ (as will be explained in §2.5). For such an isotopy class represented by α to be an obstruction, it has to be f -invariant in the sense that $f^{-1}(\alpha)$ should contain a component $\tilde{\alpha}$ isotopic to α relative to P_f . It seems to be a very intricate problem to find all slopes in $\widehat{\mathbb{Q}}$ that give an invariant isotopy class for f . Since we have been able to decide this question only for very simple maps f , we proceed in a more indirect manner.

We assume that the Jordan curve $\alpha \subset \mathbb{P} \setminus P_f$ is f -invariant and gives an obstruction. We then investigate the mapping degrees of f on components of $f^{-1}(\alpha)$ and consider intersection numbers of some relevant curves together with a careful counting argument. We heavily use the fact that the *horizontal* and *vertical* Jordan curves (see (2.4)) are f -invariant. Ultimately, we arrive at a contradiction. See §5 for the details of this argument.

Our idea to use intersection numbers (as in Lemma 5.5) to control possible locations of obstructions and dynamics on curves is not new (see, for example, [PL98, Theorem 3.2], [CPL16, §8], and [Par18]). However, the previously available results do not provide sharp enough estimates applicable in our situation.

One can think of Theorem 1.2 in the following way. Suppose that instead of directly passing from the Lattès map \mathcal{L}_n to a map, let us now call it \widehat{f} , obtained by gluing $n_h > 0$ horizontal and $n_v > 0$ vertical flaps to \mathbb{P} , we first create an intermediate map f obtained by gluing $n_h > 0$ horizontal, but no vertical flaps. Then f has a hyperbolic orbifold and an obstruction given by a ‘horizontal’ Jordan curve α . In the passage from f to \widehat{f} , we kill this obstruction, because we glue additional vertical flaps that serve as obstacles and increase the mapping degree on some pullbacks of α . Theorem 1.1 then says that no other obstructions arise for \widehat{f} . Therefore, Theorem 1.1 generalizes Theorem 1.2 if we interpret it in the way just described. The proof of Theorem 1.1 is based on the ideas that we use to establish Theorem 1.2, but substantial refinements and extensions are required.

1.2. The global curve attractor problem. The mapping properties of Jordan curves play an important role in Thurston’s characterization of rational maps. The original proof of this statement associates with a given Thurston map $f: S^2 \rightarrow S^2$ with a hyperbolic orbifold a certain Teichmüller space \mathcal{T}_f and an analytic map $\sigma_f: \mathcal{T}_f \rightarrow \mathcal{T}_f$, called *Thurston’s pullback map*. One can show that the map f is realized by a rational map if and only if σ_f has a fixed point [DH93]. This reduction to a fixed point problem in a Teichmüller space has also been successfully applied by Thurston in other contexts such as uniformization problems and the theory of 3-manifolds (there is a rich literature on the subject; see, for example, [FLP12, Hub16, Thu88, Thu98, Ota01]).

In recent years, the pullback map σ_f and its dynamical properties have been subject to deeper investigations (see, for example, [KPS16, Lod13, Pil12, Sel12]). In particular, Selinger showed in [Sel12] that σ_f extends to the Weil–Petersson boundary of \mathcal{T}_f . The behavior of σ_f on this boundary is closely related to the behavior of Jordan curves under pull-back by f . This in turn leads to the following difficult open question in holomorphic dynamics, called the *global curve attractor problem* (see [Lod13, §9]).

Conjecture. Let $f: S^2 \rightarrow S^2$ be a Thurston map with a hyperbolic orbifold that is realized by a rational map. Then there exists a finite set $\mathcal{A}(f)$ of Jordan curves in $S^2 \setminus P_f$ such that for every Jordan curve $\gamma \subset S^2 \setminus P_f$, all pullbacks $\tilde{\gamma}$ of γ under f^n are contained in $\mathcal{A}(f)$ up to isotopy relative to P_f for all sufficiently large $n \in \mathbb{N}$.

A set of Jordan curves $\mathcal{A}(f)$, as in this conjecture, is called a *global curve attractor* of f . We will give a solution of this problem for maps as in Theorem 1.2 with $n = 2$ and $n_h, n_v \geq 1$. Unfortunately, our methods only apply for $n = 2$ and not for $n \geq 3$.

THEOREM 1.3. *Let $f: \mathbb{P} \rightarrow \mathbb{P}$ be a Thurston map obtained from the (2×2) -Lattès map by gluing $n_h \geq 1$ horizontal and $n_v \geq 1$ vertical flaps to the pillow \mathbb{P} . Then f has a global curve attractor $\mathcal{A}(f)$.*

One can show that the Julia set of a rational map, as provided by Theorem 1.2, is either a Sierpiński carpet or the whole Riemann sphere depending on whether the map has periodic critical points or not (see Proposition 9.1). Accordingly, Theorem 1.3 provides the first examples of maps with Sierpiński carpet Julia set for which an answer to the global curve attractor problem is known. In fact, we obtain such maps with arbitrarily large degrees.

Recently, Belk *et al* proved the existence of a finite global curve attractor for all postcritically finite polynomials [BLMW22]. The conjecture is also known to be true for all *critically fixed rational maps* (that is, rational maps for which each critical point is fixed) and some *nearly Euclidean Thurston maps* (that is, Thurston maps with exactly four postcritical points and only simple critical points); see [FKK⁺17, Hlu19, Lod13]. In [KL19], Kelsey and Lodge verified the conjecture for all quadratic non-Lattès maps with four postcritical points. However, for general postcritically finite rational maps, the conjecture remains wide open.

Since the maps we consider have four postcritical points, it is convenient to reformulate the global curve attractor problem by introducing the *slope map* (it is closely related to the Thurston pull-back map σ_f on the Weil–Petersson boundary of \mathcal{T}_f). To define it in the special case relevant for us, we consider the *marked pillow* (\mathbb{P}, V) , where V is the set consisting of the four corners of \mathbb{P} , and assume that $f: \mathbb{P} \rightarrow \mathbb{P}$ is a Thurston map with $P_f = V$. Up to topological conjugacy, every Thurston map with four postcritical points can be assumed to have this form.

As we already mentioned, there is a bijective correspondence between isotopy classes $[\alpha]$ of essential Jordan curves α in (\mathbb{P}, V) and slopes $r/s \in \widehat{\mathbb{Q}}$ (see Lemma 2.3). We introduce the additional symbol \odot to represent peripheral Jordan curves in (\mathbb{P}, V) . We now define the *slope map* $\mu_f: \widehat{\mathbb{Q}} \cup \{\odot\} \rightarrow \widehat{\mathbb{Q}} \cup \{\odot\}$ associated with f as follows. We set

$\mu_f(\odot) := \odot$. This corresponds to the fact that each pullback of a peripheral Jordan curve α in (\mathbb{P}, V) under f is peripheral (see Corollary 3.5(i)). If $r/s \in \widehat{\mathbb{Q}}$ is an arbitrary slope, then we choose a Jordan curve α in (\mathbb{P}, V) whose isotopy class $[\alpha]$ is represented by r/s . If all pullbacks of α under f are peripheral, we set $\mu_f(r/s) := \odot$. Otherwise, there exists an essential pullback $\tilde{\alpha}$ of α under f . Then the isotopy class $[\tilde{\alpha}]$ is independent of the choice of the essential pullback $\tilde{\alpha}$ (see Corollary 3.5(ii)) and so it is represented by a unique slope $r'/s' \in \widehat{\mathbb{Q}}$. In this case, we set $\mu_f(r/s) := r'/s'$. In this way, $\mu_f(x) \in \widehat{\mathbb{Q}} \cup \{\odot\}$ is defined for all $x \in \widehat{\mathbb{Q}} \cup \{\odot\}$. Since the map μ_f has the same source and target, we can iterate it. If $n \in \mathbb{N}_0$, then we denote by μ_f^n the n th iterate of μ_f . We will then prove the following statement.

THEOREM 1.4. *Let $f: \mathbb{P} \rightarrow \mathbb{P}$ be a Thurston map obtained from the (2×2) -Lattès map by gluing $n_h \geq 1$ horizontal and $n_v \geq 1$ vertical flaps to the pillow \mathbb{P} . Then there exists a finite set $S \subset \widehat{\mathbb{Q}} \cup \{\odot\}$ with the following property: for each $x \in \widehat{\mathbb{Q}} \cup \{\odot\}$, there exists $N \in \mathbb{N}_0$ such that $\mu_f^n(x) \in S$ for all $n \geq N$.*

Note that $P_f = V$ in this case; so our previous considerations apply and the map μ_f is defined. It is clear that the previous theorem leads to the solution of the global curve attractor problem for the maps f considered.

Proof of Theorem 1.3 based on Theorem 1.4. To obtain a finite attractor $\mathcal{A}(f)$, pick a Jordan curve in each isotopy class represented by a slope in S and add five Jordan curves that represent the isotopy classes of peripheral Jordan curves in (\mathbb{P}, V) (one for null-homotopic curves and one for each corner of \mathbb{P}). \square

For the proof of Theorem 1.4, we will establish a certain monotonicity property of the slope map μ_f for a map f as in the statement (see Proposition 8.1). Roughly speaking, this monotonicity means that up to isotopy relative to $P_f = V$, complicated essential Jordan curves in (\mathbb{P}, V) get ‘simpler’ and ‘less twisted’ if we take successive preimages under f and eventually end up in the global curve attractor.

Our methods again rely on the consideration of intersection numbers. The algebraic methods for solving the global curve attractor problem developed in [Pil12] (specifically, [Pil12, Theorem 1.4]) do not apply in general for the maps considered in Theorem 1.3 (see the discussion in §9.3).

Some of our ideas can also be used for the study of the global dynamics of the slope map for Thurston maps that are not covered by Theorem 1.4. In particular, we are able to describe the iterative behavior of μ_f for a specific obstructed Thurston map f obtained by blowing up the (2×2) -Lattès map (see §9.2 for the details). This provides an answer to a question by Pilgrim.

While it is straightforward to compute $\mu_f(x)$ for individual values $x \in \widehat{\mathbb{Q}} \cup \{\odot\}$, we have been unable to give an explicit formula for μ_f for the maps f we consider. In general, these slope maps show very complicated behavior. Currently, very few explicit computations of slope maps are known in the literature. Except for some very special situations (for example, when the slope map is constant, that is, when $\mu_f(x) = \odot$ for all $x \in \widehat{\mathbb{Q}} \cup \{\odot\}$), we are only aware of computations of slope maps for nearly Euclidian

Thurston maps in [CFPP12, §5] and [Lod13, §6]. See also [FPP18] for some general properties of the slope map μ_f .

An undergraduate student at UCLA, Darragh Glynn, performed some computer experiments to compute μ_f for maps f as in Theorem 1.2 for $n \geq 3$ (and $n_h, n_v \geq 1$ corresponding to the rational case). His results show that in these cases, the map μ_f does not have the monotonicity property as for $n = 2$, but indicate that these maps f still have a global curve attractor (see §9.2 for more discussion).

1.3. Organization of this paper. Our paper is organized as follows. In the next two sections, we review some background. In §2, we fix notation and state some basic definitions. We also discuss isotopy classes of Jordan curves in spheres with four marked points, how isotopy classes of such curves correspond to slopes in $\widehat{\mathbb{Q}}$, as well as some relevant facts about intersection numbers. Even though all of this is fairly standard, we give complete proofs in the appendix, because it is hard to track down this material in the literature with a detailed exposition.

In §3, we recall some basics about Thurston maps and the relevant concepts for a precise formulation of Thurston's characterization of rational maps for Thurston maps with four postcritical points—the only case relevant for us (see §3.2).

We explain the blow-up procedure for arcs in §4 and relate this to the procedure of gluing flaps to the pillow \mathbb{P} (see §4.2). The proof of Theorem 1.2 is then given in §5.

The proof of our main result, Theorem 1.1, requires more preparation. This is the purpose of §6. There we introduce the concept of *essential circuit length* that will allow us to formulate tight estimates for the number of essential pullbacks of a Jordan curve under a Thurston map with four postcritical points. This is formulated in the rather technical Lemma 6.2 which is of crucial importance though. The proof of Theorem 1.1 is then given in §7.

Section 8 is devoted to the proof of Theorem 1.4. In §9, we discuss some further directions related to this work. As we already mentioned, the appendix is devoted to the discussion of isotopy classes and intersection numbers of Jordan curves in spheres with four marked points.

2. Preliminaries

In this section, we discuss background relevant for the rest of the paper.

2.1. Notation and basic concepts. We denote by $\mathbb{N} = \{1, 2, \dots\}$ the set of natural numbers and by $\mathbb{N}_0 = \{0, 1, 2, \dots\}$ the set of natural numbers including 0. The sets of integers, real numbers, and complex numbers are denoted by \mathbb{Z} , \mathbb{R} , and \mathbb{C} , respectively. We write i for the imaginary unit in \mathbb{C} , and $\text{Im}(z)$ for the imaginary part of a complex number $z \in \mathbb{C}$.

As usual, $\mathbb{R}^2 := \{(x, y) : x, y \in \mathbb{R}\}$ is the Euclidean plane and $\widehat{\mathbb{C}} := \mathbb{C} \cup \{\infty\}$ is the Riemann sphere. Here and elsewhere, we write $A := B$ for emphasis when an object A is defined to be another object B . When we consider two objects A and B , and there is a natural identification between them that is clear from the context, we write $A \cong B$. For

example, $\mathbb{R}^2 \cong \mathbb{C}$ if we identify a point $(x, y) \in \mathbb{R}^2$ with $x + iy \in \mathbb{C}$. We will freely switch back and forth between these different viewpoints of $\mathbb{R}^2 \cong \mathbb{C}$.

We use the notation $\mathbb{I} := [0, 1] \subset \mathbb{R}$ for the closed unit interval, $\mathbb{D} := \{z \in \mathbb{C} : |z| < 1\}$ for the open unit disk in \mathbb{C} , and $\mathbb{Z}^2 := \{x + iy : x, y \in \mathbb{Z}\}$ for the square lattice in \mathbb{C} . If $z, w \in \mathbb{C}$, then we write $[z, w] := \{z + t(w - z) : t \in \mathbb{I}\}$ for the line segment in \mathbb{C} joining z and w . We also use the notation $[z_0, w_0] := [z_0, w_0] \setminus \{w_0\}$ and $(z_0, w_0) := [z_0, w_0] \setminus \{z_0, w_0\}$.

The cardinality of a set X is denoted by $\#X \in \mathbb{N}_0 \cup \{\infty\}$ and the identity map on X by id_X . If X is a topological space and $M \subset X$, then $\text{cl}(M)$ denotes the closure, $\text{int}(M)$ the interior, and ∂M the boundary of M in X .

Let $f : X \rightarrow Y$ be a map between sets X and Y . If $M \subset X$, then $f|_M$ stands for the restriction of f to M . If $N \subset Y$, then $f^{-1}(N) := \{x \in X : f(x) \in N\}$ is the preimage of N in X . Similarly, $f^{-1}(y) := \{x \in X : f(x) = y\}$ is the preimage of a point $y \in Y$.

Let $f : X \rightarrow X$ be a map. For $n \in \mathbb{N}$, we denote by

$$f^n := \underbrace{f \circ \dots \circ f}_{n \text{ factors}}$$

the n th iterate of f . It is convenient to define $f^0 := \text{id}_X$. For $n \in \mathbb{N}_0$, we denote by $f^{-n}(M) := \{x \in X : f^n(x) \in M\}$ and $f^{-n}(p) := \{x \in X : f^n(x) = p\}$ the preimages of a set $M \subset X$ and a point $p \in X$ under f^n , respectively.

A *surface* S is a connected and oriented topological 2-manifold. We denote its *Euler characteristic* by $\chi(S)$. Note that $\chi(S) \in \{2, 1, 0, -1, \dots\} \cup \{-\infty\}$. Throughout this paper, we use the notation S^2 for a (*topological*) 2-sphere, that is, S^2 indicates a surface homeomorphic to the Riemann sphere $\widehat{\mathbb{C}}$. An *annulus* is a surface homeomorphic to $\{z \in \mathbb{C} : 1 < |z| < 2\}$.

A *Jordan curve* α in a surface S is the image $\alpha = \eta(\partial\mathbb{D})$ of a (topological) embedding $\eta : \partial\mathbb{D} \rightarrow S$ of the unit circle $\partial\mathbb{D} = \{z \in \mathbb{C} : |z| = 1\}$ into S . An *arc* e in S is the image $e = \iota(\mathbb{I})$ of an embedding $\iota : \mathbb{I} \rightarrow S$. Then $\iota(0)$ and $\iota(1)$ are the *endpoints* of e , and we define $\partial e := \{\iota(0), \iota(1)\}$. The set $\text{int}(e) := e \setminus \partial e$ is called the *interior* of e . The notions of endpoints and interior of e only depend on e and not on the choice of the embedding ι . Note that the notation ∂e and $\text{int}(e)$ is ambiguous, because it should not be confused with the boundary and interior of e as a subset of S . For arcs e in a surface S , we will only use ∂e and $\text{int}(e)$ with the meaning just defined.

A subset U of a surface S is called an *open* or *closed Jordan region* if there exists a topological embedding $\eta : \text{cl}(\mathbb{D}) = \{z \in \mathbb{C} : |z| \leq 1\} \rightarrow S$ such that $U = \eta(\mathbb{D})$ or $U = \eta(\text{cl}(\mathbb{D}))$, respectively. In both cases, $\partial U = \eta(\partial\mathbb{D})$ is a Jordan curve in S . A *crosscut* e in an open or closed Jordan region U is an arc $e \subset \text{cl}(U)$ such that $\text{int}(e) \subset \text{int}(U)$ and $\partial e \subset \partial U$.

A *path* γ in a surface S is a continuous map $\gamma : [a, b] \rightarrow S$, where $[a, b] \subset \mathbb{R}$ is a compact (non-degenerate) interval. As is common, we will use the same notation γ for the image $\gamma([a, b])$ of the path if no confusion can arise. The path γ *joins* two sets $M, N \subset S$ if $\gamma(a) \in M$ and $\gamma(b) \in N$, or *vice versa*. A *loop* in S *based at* $p \in S$ is a path $\gamma : [a, b] \rightarrow S$ such that $\gamma(a) = \gamma(b) = p$. The loop γ is called *simple* if γ is injective on $[a, b)$. So essentially, a simple loop is a Jordan curve run through with some parameterization.

Let M, N, K be subsets of a surface S . We say that K separates M and N if every path in S joining M and N meets K . Note that here, K is not necessarily disjoint from M or N . We say that K separates a point $p \in S$ from a set $M \subset S$ if K separates $\{p\}$ and M .

Let $Z \subset S$ be a finite set of points in a surface S . Then we refer to the pair (S, Z) as a *marked surface*, and the points in Z as the *marked points* in S . The most important case for us will be when $S = S^2$ is a 2-sphere and $Z \subset S^2$ consists of four points.

A *Jordan curve* α in a marked surface (S, Z) is a Jordan curve $\alpha \subset S \setminus Z$. An *arc* e in (S, Z) is an arc $e \subset S$ with $\partial e \subset Z$ and $\text{int}(e) \subset S \setminus Z$. We say that a Jordan curve α in a marked sphere (S^2, Z) is *essential* if each of the two connected components of $S^2 \setminus \alpha$ contains at least two points of Z ; otherwise, we say that α is *peripheral*.

Let (S^2, Z) be a marked sphere with $\#Z = 4$. A *core arc* of an essential Jordan curve α in (S^2, Z) is an arc in (S^2, Z) that is contained in one of the two connected components of $S^2 \setminus \alpha$ and joins the two points in Z that lie in this component.

Let A be an annulus. Then a *core curve* of A is a Jordan curve $\beta \subset A$ such that under some homeomorphism $\varphi: A \rightarrow A'$, the curve $\beta' = \varphi(\beta)$ separates the boundary components of $A' = \{z \in \mathbb{C} : 1 < |z| < 2\}$.

2.2. Branched covering maps. Let X and Y be surfaces. Then a continuous map $f: X \rightarrow Y$ is called a *branched covering map* if for each point $q \in Y$, there exists an open set $V \subset Y$ homeomorphic to \mathbb{D} with $q \in V$ that is *evenly covered* in the following sense: for some index set $J \neq \emptyset$, we can write $f^{-1}(V)$ as a disjoint union

$$f^{-1}(V) = \bigcup_{j \in J} U_j \tag{2.1}$$

of open sets $U_j \subset X$ such that U_j contains precisely one point $p_j \in f^{-1}(q)$. Moreover, we require that for each $j \in J$, there exists $d_j \in \mathbb{N}$ and orientation-preserving homeomorphisms $\varphi_j: U_j \rightarrow \mathbb{D}$ with $\varphi_j(p_j) = 0$ and $\psi_j: V \rightarrow \mathbb{D}$ with $\psi_j(q) = 0$ such that

$$(\psi_j \circ f \circ \varphi_j^{-1})(z) = z^{d_j}$$

for all $z \in \mathbb{D}$ (see [BM17, §A.6] for more background on branched covering maps). For given f , the number d_j is uniquely determined by $p = p_j$, and called the *local degree* of f at p and denoted by $\text{deg}(f, p)$. A point $p \in X$ with $\text{deg}(f, p) \geq 2$ is called a *critical point* of f . The set of all critical points of f is a discrete set in X and denoted by C_f . If f is a branched covering map, then it is a covering map (in the usual sense) from $X \setminus f^{-1}(C_f)$ onto $Y \setminus C_f$.

In the following, suppose X and Y are compact surfaces, and $f: X \rightarrow Y$ is a branched covering map. Then $C_f \subset X$ is a finite set. Moreover, if $\text{deg}(f) \in \mathbb{N}$ denotes the topological degree of f , then

$$\sum_{p \in f^{-1}(q)} \text{deg}(f, p) = \text{deg}(f)$$

for each $q \in Y$.

If $\gamma: [a, b] \rightarrow Y$ is a path, then we call a path $\tilde{\gamma}: [a, b] \rightarrow X$ a *lift* of γ (under f) if $f \circ \tilde{\gamma} = \gamma$. Every path γ in Y has a lift $\tilde{\gamma}$ in X (see [BM17, Lemma A.18]), but in general, $\tilde{\gamma}$

is not unique. If $\gamma([a, b]) \subset Y \setminus f(C_f)$ and $x_0 \in f^{-1}(\gamma(a))$, then there exists a unique lift $\tilde{\gamma}: [a, b] \rightarrow X$ of γ under f with $\tilde{\gamma}(a) = x_0$. This easily follows from standard existence and uniqueness theorems for lifts under covering maps (see [BM17, Lemma A.6]).

If $e \subset Y$ is an arc, then an arc $\tilde{e} \subset X$ is called a *lift* of e (under f) if $f|\tilde{e}$ is a homeomorphism of \tilde{e} onto e . It easily follows from the existence and uniqueness statements for lifts of paths just discussed that if e is an arc in $(Y, f(C_f))$, $y_0 \in \text{int}(e)$, and $x_0 \in f^{-1}(y_0)$, then there exists a unique lift $\tilde{e} \subset X$ of e with $x_0 \in \tilde{e}$.

Let $V \subset Y$ be an open and connected set, and $U \subset f^{-1}(V)$ be a (connected) component of $f^{-1}(V)$. Then $f|U: U \rightarrow V$ is also a branched covering map. Each point $q \in V$ has the same number $d \in \mathbb{N}$ of preimages under $f|U$ counting local degrees. We set $\text{deg}(f|U) := d$. If the Euler characteristic $\chi(V)$ is finite, then $\chi(U)$ is also finite and we have the *Riemann–Hurwitz formula*

$$\chi(U) + \sum_{p \in U \cap C_f} (\text{deg}(f, p) - 1) = \text{deg}(f|U) \cdot \chi(V). \tag{2.2}$$

2.3. Planar embedded graphs. A *planar embedded graph* in a sphere S^2 is a pair $G = (V, E)$, where V is a finite set of points in S^2 and E is a finite set of arcs in (S^2, V) with pairwise disjoint interiors. The sets V and E are called the *vertex* and *edge sets* of G , respectively. Note that our notion of a planar embedded graph does not allow *loops*, that is, edges that connect a vertex to itself, but it does allow *multiple edges*, that is, distinct edges that join the same pair of vertices. The *degree* of a vertex v in G , denoted $\text{deg}_G(v)$, is the number of edges of G incident to v . Note that $2 \cdot \#E = \sum_{v \in V} \text{deg}_G(v)$.

The *realization* of G is the subset \mathcal{G} of S^2 given by

$$\mathcal{G} := V \cup \bigcup_{e \in E} e.$$

A *face* of G is a connected component of $S^2 \setminus \mathcal{G}$. Usually, we conflate a planar embedded graph G with its realization \mathcal{G} . Then it is understood that \mathcal{G} contains a finite set $V \subset \mathcal{G}$ of distinguished points that are the vertices of the graph. Its edges are the closures of the components of $\mathcal{G} \setminus V$.

A *subgraph* of a planar embedded graph $G = (V, E)$ is a planar embedded graph $G' = (V', E')$ with $V' \subset V$ and $E' \subset E$. A *path of length n between vertices v and v'* in G is a sequence $v_0, e_0, v_1, e_1, \dots, e_{n-1}, v_n$, where $v_0 = v, v_n = v'$, and e_k is an edge incident to the vertices v_k and v_{k+1} for $k = 0, \dots, n - 1$. A path that does not repeat vertices is called a *simple path*.

A path $v_0, e_0, v_1, e_1, \dots, e_{n-1}, v_n$ with $v_0 = v_n$ and $n \geq 2$ is called a *circuit of length n* in G and is denoted by $(e_0, e_1, \dots, e_{n-1})$. Such a circuit is called a *simple cycle* if all vertices $v_k, k = 0, \dots, n - 1$, are distinct.

A planar embedded graph G is called *connected* if any two distinct vertices of G can be joined by a path in G . Equivalently, G is connected if its realization \mathcal{G} is connected as a subset of S^2 . Note that if G is connected, then each face of G is simply connected.

As follows from [Die05, Lemma 4.2.2], the topological boundary ∂U of each face U of G may be viewed as the realization of a subgraph of G . Moreover, a walk around any




FIGURE 4. The Euclidean square pillow \mathbb{P} .

connected component of the boundary ∂U traces a circuit $(e_0, e_1, \dots, e_{n-1})$ in G such that each edge of G appears zero, one, or two times in the sequence e_0, e_1, \dots, e_{n-1} . We will say that the circuit $(e_0, e_1, \dots, e_{n-1})$ *traces* (a connected component of) the boundary ∂U . If U is simply connected, then ∂U is connected, and the length of the (essentially unique) circuit that bounds U is called the *circuit length* of U in G .

A planar embedded graph (V, E) is called *bipartite* if we can split V into two disjoint subsets V_1 and V_2 such that each edge $e \in E$ has one endpoint in V_1 and one in V_2 .

2.4. The Euclidean square pillow. As discussed in the introduction, we consider a *square pillow* \mathbb{P} obtained from gluing two identical copies of the unit square $\mathbb{I}^2 \subset \mathbb{R}^2$ along their boundaries by the identity map. Then \mathbb{P} is a topological 2-sphere. We equip \mathbb{P} with the induced path metric that agrees with the Euclidean metric on each of the two copies of the unit square. We call this metric space \mathbb{P} the *Euclidean square pillow*. The vertices and edges of the unit square \mathbb{I}^2 in \mathbb{P} are called the *vertices* and *edges* of \mathbb{P} . One copy of \mathbb{I}^2 in \mathbb{P} is called the *front* and the other copy the *back side* of \mathbb{P} . In a dynamical context, we also refer to these two copies of \mathbb{I}^2 as the *0-tiles* of \mathbb{P} . We color the front side of \mathbb{P} white, and its back side black. Finally, we equip \mathbb{P} with the orientation that agrees with the standard orientation on the front side \mathbb{I}^2 of \mathbb{P} (represented by the positively oriented standard flag $((0, 0), \mathbb{I} \times \{0\}, \mathbb{I}^2)$; see [BM17, Appendix A.4]).

We label the vertices and edges of \mathbb{P} in counterclockwise order by A, B, C, D and a, b, c, d , respectively, so that $A \in \mathbb{P}$ corresponds to the vertex $(0, 0) \in \mathbb{I}^2$ and the edge $a \subset \mathbb{P}$ corresponds to $[0, 1] \times \{0\} \subset \mathbb{I}^2$. Then a has the endpoints A and B . We can view the boundary $\partial \mathbb{I}^2$ of \mathbb{I}^2 as a planar embedded graph in \mathbb{P} with the vertex set $V := \{A, B, C, D\}$ and the edge set $E := \{a, b, c, d\}$. We call a and c the *horizontal edges*, and b and d the *vertical edges* of \mathbb{P} ; see Figure 4.

The pillow \mathbb{P} is an example of a *Euclidean polyhedral surface*, that is, a surface obtained by gluing Euclidean polygons along boundary edges by using isometries. Note that the metric on \mathbb{P} is locally flat except at its vertices, which are Euclidean conic singularities. So \mathbb{P} is an *orbifold* (see, for example, [Mil06a, Appendix E] and [BM17, Appendix A.9]).

An alternative description for the pillow \mathbb{P} can be given as follows. We consider the unit square $\mathbb{I}^2 \subset \mathbb{R}^2 \cong \mathbb{C}$ and map it to the upper half-plane in $\widehat{\mathbb{C}}$ by a conformal map, normalized so that the vertices $0, 1, 1 + i, i$ are mapped to $0, 1, \infty, -1$, respectively. By Schwarz




FIGURE 5. The map $\wp : \mathbb{C} \rightarrow \mathbb{P}$.

reflection, this map can be extended to a meromorphic function $\wp : \mathbb{C} \rightarrow \widehat{\mathbb{C}}$. Then \wp is a Weierstrass \wp -function (up to a postcomposition with a Möbius transformation) that is doubly periodic with respect to the lattice $2\mathbb{Z}^2 := \{2k + 2ni : k, n \in \mathbb{Z}\} \subset \mathbb{C}$. Actually, for $z, w \in \mathbb{C}$, we have

$$\wp(z) = \wp(w) \quad \text{if and only if } z - w \in 2\mathbb{Z}^2 \text{ or } z + w \in 2\mathbb{Z}^2. \tag{2.3}$$

We can push forward the Euclidean metric on \mathbb{C} to the Riemann sphere $\widehat{\mathbb{C}}$ by \wp . With respect to this metric, called the *canonical orbifold metric* for \wp , the sphere $\widehat{\mathbb{C}}$ is isometric to the Euclidean square pillow \mathbb{P} . In the following, we identify the pillow \mathbb{P} with $\widehat{\mathbb{C}}$ by the orientation-preserving isometry that maps the vertices A, B, C, D to $0, 1, \infty, -1$, respectively. Then we can consider $\wp : \mathbb{C} \rightarrow \widehat{\mathbb{C}} \cong \mathbb{P}$ as a map onto the pillow \mathbb{P} . Actually, \wp is the *universal orbifold covering map* for \mathbb{P} (see [BM17, §A.9] for more background). A very intuitive description of this map can be given if we color the squares $[k, k + 1] \times [n, n + 1], k, n \in \mathbb{Z}$, in checkerboard manner black and white so that $[0, 1] \times [0, 1]$ is white. Restricted to such a square S , the map \wp is an isometry that sends S to the white 0-tile of \mathbb{P} if S is white, and to the black 0-tile \mathbb{P} if S is black; see Figure 5 for an illustration. Here, the points in the complex plane \mathbb{C} marked by a black dot (on the left) are mapped to A by \wp and are elements of $\wp^{-1}(A) = 2\mathbb{Z}^2$.

2.5. Isotopies and intersection numbers. Let X and Y be topological spaces. Then a continuous map $H : X \times \mathbb{I} \rightarrow Y$ is called a *homotopy* from X to Y . For $t \in \mathbb{I}$, we denote by $H_t := H(\cdot, t) : X \rightarrow Y$ the *time- t* map of the homotopy. The homotopy H is called an *isotopy* if H_t is a homeomorphism from X onto Y for each $t \in \mathbb{I}$. If $Z \subset X$, then a homotopy $H : X \times \mathbb{I} \rightarrow Y$ is said to be a *homotopy relative to Z* if $H_t(p) = H_0(p)$ for all $p \in Z$ and $t \in \mathbb{I}$. In other words, the image of each point in Z remains fixed during the homotopy H . Isotopies relative to Z are defined in a similar way.

Two homeomorphisms $h_0, h_1 : X \rightarrow Y$ are called *isotopic (relative to $Z \subset X$)* if there exists an isotopy $H : X \times \mathbb{I} \rightarrow Y$ (relative to Z) with $H_0 = h_0$ and $H_1 = h_1$. Given $M, N, Z \subset X$, we say that M is *isotopic to N relative to Z* (or M can be isotoped into N

relative to Z), denoted by $M \sim N$ relative to Z , if there exists an isotopy $H : X \times \mathbb{I} \rightarrow X$ relative to Z with $H_0 = \text{id}_X$ and $H_1(M) = N$. Recall that id_X is the identity map on X .

Let (S, Z) be a marked surface (with a finite, possibly empty set $Z \subset S$ of marked points). If α is a Jordan curve in (S, Z) , then its isotopy class $[\alpha]$ (with (S, Z) understood) consists of all Jordan curves β in (S, Z) such that $\alpha \sim \beta$ relative to Z .

The following statement gives a sufficient condition for two Jordan curves in (S, Z) to be isotopic relative to Z .

LEMMA 2.1. *Let α and β be disjoint Jordan curves in a marked surface (S, Z) . Suppose there is an annulus $U \subset S \setminus Z$ such that $\partial U = \alpha \cup \beta$. Then α and β are isotopic relative to Z .*

Proof. This is standard and we will only give a sketch of the proof. Since Jordan curves in surfaces are tame, one can slightly enlarge the annulus U to an annulus $U' \subset S^2 \setminus Z$ that contains α and β . Then α can be isotoped into β by an isotopy on U' that is the identity near $\partial U'$. This isotopy on U' can be extended to an isotopy on S^2 relative to Z that isotopes α into β . □

Let (S, Z) be a marked surface. If α and β are arcs or Jordan curves in (S, Z) , we define their (unsigned) intersection number as

$$i(\alpha, \beta) := \inf\{\#(\alpha' \cap \beta') : \alpha \sim \alpha' \text{ relative to } Z \text{ and } \beta \sim \beta' \text{ relative to } Z\}.$$

The relevant marked surface (S, Z) here will be understood from the context, and we suppress it from our notation for intersection numbers. If we want to emphasize it, we will say that we consider intersection numbers in (S, Z) . The intersection number is always finite, because we can always reduce to the case when α and β are piecewise geodesic with respect to some Riemannian metric on S (see [Bus10, Lemma A.8]). If α and β satisfy $i(\alpha \cap \beta) = \#(\alpha \cap \beta)$, then we say that α and β are in minimal position (in their isotopy classes relative to Z).

Suppose α and β are arcs or Jordan curves in (S, Z) . Then we say that α and β meet transversely at a point $p \in \alpha \cap \beta \cap (S \setminus Z)$ (or α crosses β at p) if p is an isolated point in $\alpha \cap \beta$ and if the following condition is true for a (small) arc $\sigma \subset \alpha$ containing p as an interior point such that $\sigma \cap \beta = \{p\}$: let σ^L and σ^R be the two subarcs of σ into which σ is split by p , then with suitable orientation of β near p , the arc σ^L lies to the left and σ^R to the right of β . We say that α and β meet transversely or have transverse intersection if the set $\alpha \cap \beta$ is finite and if α and β meet transversely at each point $p \in \alpha \cap \beta \cap (S \setminus Z)$.

LEMMA 2.2. *Suppose α and β are Jordan curves or arcs in a marked surface (S, Z) . If α and β are in minimal position, then α and β meet transversely.*

Proof. This is essentially a standard fact (see, for example, [Bus10, pp. 416–417]), and we will only give an outline of the proof.

Since $\#(\alpha \cap \beta) = i(\alpha \cap \beta)$, the set $\alpha \cap \beta$ consists of finitely many isolated points. To reach a contradiction, suppose that α and β do not meet transversely at some point $p \in \alpha \cap \beta \cap (S \setminus Z)$. Then there exists an arc $\sigma \subset \alpha$ containing p as an interior point such that $\sigma \cap \beta = \{p\}$ and with the following property: if σ_1 and σ_2 denote the two subarcs of

σ into which σ is split by p , then σ_1 and σ_2 lie on the same side of β (equipped with some orientation locally near p). In other words, α touches β locally near p from one side and does not cross β at p .

We can then modify the curve α near p by an isotopy that pulls the subarc σ away from β so that the new curve α does not have the intersection point p with β while no new intersection points of α and β arise. This contradicts our assumption that for the original curve α , we have $\#(\alpha \cap \beta) = i(\alpha \cap \beta)$. □

2.6. *Jordan curves in spheres with four marked points.* If (S^2, Z) is a marked sphere where $Z \subset S^2$ consists of exactly four points, then, up to homeomorphism, we may assume that S^2 is equal to the pillow \mathbb{P} , and $Z = V = \{A, B, C, D\}$ consists of the four vertices of \mathbb{P} . We will freely switch back and forth between a general marked sphere (S^2, Z) with $\#Z = 4$ and (\mathbb{P}, V) .

We need some statements about isotopy classes of Jordan curves and arcs in (\mathbb{P}, V) and their intersections numbers. They are ‘well known’, but unfortunately we have been unable to track down a comprehensive account in the literature. Accordingly, we will provide a complete treatment. This may be of independent interest apart from the main objective of the paper. We will give the statements in this section, but will provide the details of the proofs in the appendix.

As we will see, there is a natural way to define a bijection between the set of isotopy classes $[\gamma]$ of essential Jordan curves γ in (\mathbb{P}, V) and the set of *extended rational numbers* $\widehat{\mathbb{Q}} := \mathbb{Q} \cup \{\infty\}$. Throughout this paper, whenever we write $r/s \in \widehat{\mathbb{Q}}$, we assume that $r \in \mathbb{Z}$ and $s \in \mathbb{N}_0$ are two relatively prime integers. We allow $s = 0$ here, in which case we assume $r = 1$. Then $r/s = 1/0 := \infty \in \widehat{\mathbb{Q}}$.

We say that a (straight) line $\ell \subset \mathbb{C}$ has slope $r/s \in \widehat{\mathbb{Q}}$ if it is given as

$$\ell = \{z_0 + (s + ir)t : t \in \mathbb{R}\} \subset \mathbb{C}$$

for some $z_0 \in \mathbb{C}$. We use the notation $\ell_{r/s}(z_0)$ for the unique line in \mathbb{C} with slope r/s passing through $z_0 \in \mathbb{C}$, and the notation $\ell_{r/s}$ (when the point z_0 is not important) for any line in \mathbb{C} with slope r/s .

Let $\ell_{r/s} \subset \mathbb{C}$ be any line with slope $r/s \in \widehat{\mathbb{Q}}$. If $\ell_{r/s}$ does not contain any point in the lattice $\mathbb{Z}^2 = \wp^{-1}(V)$ and so $\ell_{r/s} \subset \mathbb{C} \setminus \mathbb{Z}^2$, then $\tau_{r/s} := \wp(\ell_{r/s})$ is a Jordan curve in $\mathbb{P} \setminus V$. Actually, $\tau_{r/s}$ is a simple closed geodesic in the Euclidean square pillow \mathbb{P} (see Figure 6 for an illustration). If $\ell_{r/s}$ contains a point in \mathbb{Z}^2 , then $\xi_{r/s} := \wp(\ell_{r/s})$ is a geodesic arc in (\mathbb{P}, V) .

It is easy to see that every simple closed geodesic or geodesic arc τ in (\mathbb{P}, V) has the form $\tau = \wp(\ell_{r/s})$ for a line $\ell_{r/s} \subset \mathbb{C}$ with some slope $r/s \in \widehat{\mathbb{Q}}$. In the following, we use the notation $\tau_{r/s}$ for a simple closed geodesic and $\xi_{r/s}$ for a geodesic arc obtained in this way.

It follows from (2.3) that for fixed $r/s \in \widehat{\mathbb{Q}}$, we obtain precisely two distinct arcs $\xi_{r/s}$ and $\xi'_{r/s}$ of the form $\wp(\ell_{r/s}(z_0))$ depending on $z_0 \in \mathbb{Z}^2$. For each simple closed geodesic $\tau_{r/s}$, the arcs $\xi_{r/s}$ and $\xi'_{r/s}$ are core arcs of $\tau_{r/s}$ lying in different components of $\mathbb{P} \setminus \tau_{r/s}$ (see the appendix for more details). In particular, $\tau_{r/s}$ is always an essential Jordan curve in (\mathbb{P}, V) .




FIGURE 6. A line ℓ_2 and the corresponding Jordan curve $\tau_2 = \wp(\ell_2)$ in \mathbb{P} .

It turns out that the isotopy classes of essential Jordan curves in (\mathbb{P}, V) are closely related to the simple closed geodesics $\tau_{r/s}$.

LEMMA 2.3. *Let γ be an essential Jordan curve in (\mathbb{P}, V) . Then there exists a unique slope $r/s \in \widehat{\mathbb{Q}}$ with the following property. Let $\ell_{r/s}$ be any line in \mathbb{C} with slope r/s and $\ell_{r/s} \subset \mathbb{C} \setminus \mathbb{Z}^2$, and set $\tau_{r/s} := \wp(\ell_{r/s})$. Then $\tau_{r/s}$ is an essential Jordan curve in (\mathbb{P}, V) with $\gamma \sim \tau_{r/s}$ relative to V . Moreover, the map $[\gamma] \mapsto r/s$ gives a bijection between isotopy classes $[\gamma]$ of essential Jordan curves γ in (\mathbb{P}, V) and slopes $r/s \in \widehat{\mathbb{Q}}$.*

While this is well known (see, for example, [FM12, Proposition 2.6] or [KS94, Proposition 2.1]), we find the available proofs too sketchy. This is the reason why we provide a detailed proof in the appendix. Implicit in Lemma 2.3 is the fact that the isotopy class $[\tau_{r/s}]$ of $\tau_{r/s} = \wp(\ell_{r/s})$ only depends on r/s and not on the specific choice of the line $\ell_{r/s}$ with $\ell_{r/s} \subset \mathbb{C} \setminus \mathbb{Z}^2$ (see Lemma A.7 for an explicit statement).

Recall that a, c denote the horizontal, and b, d the vertical edges of \mathbb{P} . In the following, we denote by $\alpha^h = \tau_0$ a horizontal essential Jordan curve in (\mathbb{P}, V) (corresponding to slope 0 and separating the edges a and c of \mathbb{P}) and by $\alpha^v = \tau_\infty$ a vertical essential Jordan curve in (\mathbb{P}, V) (corresponding to slope ∞ and separating b from d). To be specific, we set

$$\alpha^h := \wp(\mathbb{R} \times \{1/2\}) \quad \text{and} \quad \alpha^v := \wp(\{1/2\} \times \mathbb{R}). \tag{2.4}$$

The following lemma summarizes the intersection properties of essential Jordan curves and arcs in (\mathbb{P}, V) .

LEMMA 2.4. *Let α and β be essential Jordan curves in (\mathbb{P}, V) and $r/s, r'/s' \in \widehat{\mathbb{Q}}$ be the unique slopes such that $\alpha \sim \tau_{r/s}$ and $\beta \sim \tau_{r'/s'}$ relative to V , where $\tau_{r/s}$ and $\tau_{r'/s'}$ are simple closed geodesics in (\mathbb{P}, V) with slopes r/s and r'/s' , respectively. Let ξ be a core arc of β , and $\xi_{r'/s'}$ be a geodesic arc in (\mathbb{P}, V) with slope r'/s' . Then the following statements are true for intersection numbers in (\mathbb{P}, V) :*

- (i) *if $r/s = r'/s'$, then $i(\alpha, \beta) = 0$, and if $r/s \neq r'/s'$, then $i(\alpha, \beta) = \#(\tau_{r/s} \cap \tau_{r'/s'}) = 2|rs' - sr'| > 0$;*
- (ii) *$i(\alpha, \xi) = \#(\tau_{r/s} \cap \xi_{r'/s'}) = \frac{1}{2}i(\alpha, \beta) = |rs' - sr'|$;*




FIGURE 7. Counting intersections of τ_2 with the horizontal curve α^h and the horizontal edges a and c .

- (iii) $i(\alpha, a) = \#(\tau_{r/s} \cap a) = |r|$, $i(\alpha, c) = \#(\tau_{r/s} \cap c) = |r|$;
- (iv) $i(\alpha, b) = \#(\tau_{r/s} \cap b) = s$, $i(\alpha, d) = \#(\tau_{r/s} \cap d) = s$;
- (v) $i(\alpha, \alpha^h) = 2|r|$ and $i(\alpha, \alpha^v) = 2s$.

We will prove this lemma in the appendix. Note that Figure 7 illustrates statements (iii) and (v) when $\alpha = \tau_2$. It follows from the lemma that $\tau_{r/s}$ for $r/s \neq 0, \infty$ is in minimal position with each of the curves $a, b, c, d, \alpha^h, \alpha^v$.

Let γ be a Jordan curve or an arc in a surface S , and $M_1, M_2 \subset S$ be two disjoint sets with $0 < \#(\gamma \cap M_j) < \infty$ for $j = 1, 2$. We say that the points in $\gamma \cap M_1 \neq \emptyset$ and $\gamma \cap M_2 \neq \emptyset$ alternate on γ if any two points in one of the sets are separated by the other, that is, any subarc $\sigma \subset \gamma$ with both endpoints in either of the sets $\gamma \cap M_1$ or $\gamma \cap M_2$ must contain a point in the other set.

More intuitively, this situation when the points in $\gamma \cap M_1$ and $\gamma \cap M_2$ alternate on an arc γ can be described as follows. Suppose we traverse γ in some (injective) parameterization starting from one of its endpoints. Then we will first meet a point in either M_1 or M_2 , say in M_1 . Then as we continue along γ , we will meet a point in M_2 , then a point in M_1 , etc. A similar remark applies when γ is a Jordan curve. Note that in this case $\#(\gamma \cap M_1) = \#(\gamma \cap M_2)$.

LEMMA 2.5. *Let $\tau = \wp(\ell_{r/s})$ be a simple closed geodesic or a geodesic arc in (\mathbb{P}, V) obtained from a line $\ell_{r/s} \subset \mathbb{C}$ with slope $r/s \in \hat{\mathbb{Q}}$. If $r/s \neq 0$, then the sets $a \cap \tau$ and $c \cap \tau$ are non-empty and finite, and the points in $a \cap \tau$ and $c \cap \tau$ alternate on τ .*

A similar statement is true if $r/s \neq \infty$ and we replace a, c with b, d , respectively. Lemma 2.5 is related to a similar statement in a more general setting that is the key to proving Lemma 2.3 (see Lemma A.3).

Proof. Define $\omega := s + ir \in \mathbb{Z}^2$. Suppose first that $\tau = \wp(\ell_{r/s})$ is a simple closed geodesic. Then $\tau = \wp([z_0, w_0])$, where $z_0 \in \ell_{r/s} \subset \mathbb{C} \setminus \mathbb{Z}^2$ and $w_0 = z_0 + 2\omega$, and the map $u \in [0, 1] \mapsto \wp(uz_0 + (1 - u)w_0)$ provides a parameterization of τ as a simple loop as follows from (2.3).

Now the set $\wp^{-1}(a \cup c)$ consists precisely of the lines $\ell_0(ni) = \{z \in \mathbb{C} : \text{Im}(z) = n\}$, $n \in \mathbb{Z}$. These lines alternate in the sense that \wp maps $\ell_0(ni)$ onto a or c depending on whether $n \in \mathbb{Z}$ is even or odd, respectively. Since $r/s \neq 0$, we have $r \in \mathbb{Z} \setminus \{0\}$, and so $\text{Im}(w_0 - z_0) = \text{Im}(2\omega) = 2r$ is a non-zero even integer. This implies that the line segment $[z_0, w_0] \subset \ell_{r/s}$ has non-empty intersections (consisting of finitely many points) with each of the sets $\wp^{-1}(a)$ and $\wp^{-1}(c)$. Moreover, the points in these intersections alternate on the segment $[z_0, w_0]$. From this, together with the fact that

$$\#(\wp^{-1}(a) \cap [z_0, w_0]) = |r| = \#(\wp^{-1}(c) \cap [z_0, w_0]),$$

the statement follows (the latter fact is needed to argue that the points in $a \cap \tau$ and $c \cap \tau$ alternate on the simple closed geodesic τ).

If τ is a geodesic arc, then there exists $z_0 \in \mathbb{Z}^2$ such that $\tau = \wp([z_0, w_0])$, where $w_0 = z_0 + \omega$. The map \wp sends $[z_0, w_0]$ homeomorphically onto τ . Again, $[z_0, w_0]$ has non-empty intersections consisting of finitely many points with each of the sets $\wp^{-1}(a)$ and $\wp^{-1}(c)$. Moreover, the points in the sets $\wp^{-1}(a) \cap [z_0, w_0] \neq \emptyset$ and $\wp^{-1}(c) \cap [z_0, w_0] \neq \emptyset$ alternate on the segment $[z_0, w_0]$. The statement also follows in this case. □

We conclude this section with a statement related to the previous considerations formulated for an arbitrary sphere with four marked points.

LEMMA 2.6. *Let (S^2, Z) be a marked sphere with $\#Z = 4$, and α, γ be essential Jordan curves in (S^2, Z) . Suppose that a_α and c_α are core arcs of α that lie in different components of $S^2 \setminus \alpha$. Then the following statements are true:*

- (i) $i(\alpha, \gamma) = 2i(a_\alpha, \gamma) = 2i(c_\alpha, \gamma)$;
- (ii) *if $i(\alpha, \gamma) > 0$, then there exists a Jordan curve γ' in (S^2, Z) with $\gamma' \sim \gamma$ relative to Z such that γ' is in minimal position with $\alpha, a_\alpha, c_\alpha$ and the points in $a_\alpha \cap \gamma' \neq \emptyset$ and $c_\alpha \cap \gamma' \neq \emptyset$ alternate on γ' .*

Proof. We may identify the marked sphere (S^2, Z) with the pillow (\mathbb{P}, V) by a homeomorphism that sends $\alpha, a_\alpha, c_\alpha$ to α^h, a, c , respectively. Then, by Lemma 2.3, the curve γ is isotopic to a simple closed geodesic $\tau_{r/s}$ with slope $r/s \in \widehat{\mathbb{Q}}$. Statement (i) then follows from Lemma 2.4(iii) and (v).

If $i(\alpha, \gamma) = i(\alpha^h, \tau_{r/s}) = 2|r| > 0$, then we can choose $\gamma' = \tau_{r/s}$ in statement (ii). Indeed, then $\gamma' = \tau_{r/s} \sim \gamma$ relative to $V = Z$, and $\gamma' = \tau_{r/s}$ is in minimal position with $\alpha^h = \tau_0, a_\alpha = a, c_\alpha = c$ as follows from Lemma 2.4 (i) and (iii). Since $r/s \neq 0$ in this case, the statement about alternation follows from Lemma 2.5. □

3. Thurston maps

Here we provide a very brief summary of some relevant definitions and facts. For more details, we refer the reader to [BM17, Ch. 2].

Let $f: S^2 \rightarrow S^2$ be a branched covering map of a topological 2-sphere S^2 . A point $p \in S^2$ is called *periodic* (for f) if $f^n(p) = p$ for some $n \in \mathbb{N}$. Recall that C_f denotes the set of all critical points of f . The union

$$P_f = \bigcup_{n \in \mathbb{N}} f^n(C_f)$$

of the orbits of critical points is called the *postcritical set* of f . Note that

$$f(P_f) \subset P_f \subset f^{-1}(P_f).$$

The map f is said to be *postcritically finite* if its postcritical set P_f is finite, in other words, if every critical point of f has a finite orbit under iteration.

Definition 3.1. A *Thurston map* is a postcritically finite branched covering map $f: S^2 \rightarrow S^2$ of topological degree $\deg(f) \geq 2$.

Natural examples of Thurston maps are given by *rational Thurston maps*, that is, postcritically finite rational maps on the Riemann sphere $\widehat{\mathbb{C}}$.

The *ramification function* of a Thurston map $f: S^2 \rightarrow S^2$ is a function $\alpha_f: S^2 \rightarrow \mathbb{N} \cup \{\infty\}$ such that $\alpha_f(p)$ for $p \in S^2$ is the lowest common multiple of all local degrees $\deg(f^n, q)$, where $q \in f^{-n}(p)$ and $n \in \mathbb{N}$ are arbitrary. In particular, $\alpha_f(p) = 1$ for $p \in S^2 \setminus P_f$ and $\alpha_f(p) \geq 2$ for $p \in P_f$.

Definition 3.2. Two Thurston maps $f: S^2 \rightarrow S^2$ and $g: \widehat{S}^2 \rightarrow \widehat{S}^2$, where \widehat{S}^2 is another topological 2-sphere, are called *Thurston equivalent* if there are homeomorphisms $h_0, h_1: S^2 \rightarrow \widehat{S}^2$ that are isotopic relative to P_f such that $h_0 \circ f = g \circ h_1$.

We say that a Thurston map is *realized* (by a rational map) if it is Thurston equivalent to a rational map. Otherwise, we say that it is *obstructed*.

The *orbifold* O_f associated with a Thurston map f is the pair (S^2, α_f) . The *Euler characteristic* of O_f is

$$\chi(O_f) := 2 - \sum_{p \in P_f} \left(1 - \frac{1}{\alpha_f(p)} \right). \tag{3.1}$$

Here we set $1/\infty := 0$.

The Euler characteristic of the orbifold O_f satisfies $\chi(O_f) \leq 0$. We call O_f *hyperbolic* if $\chi(O_f) < 0$, and *parabolic* if $\chi(O_f) = 0$.

If $f: S^2 \rightarrow S^2$ is a Thurston map, then $f(C_f \cup P_f) \subset P_f$, which implies $C_f \cup P_f \subset f^{-1}(P_f)$. The reverse inclusion is related to the parabolicity of O_f .

LEMMA 3.3. *Let $f: S^2 \rightarrow S^2$ be a Thurston map. If f has a parabolic orbifold, then $f^{-1}(P_f) = C_f \cup P_f$. Moreover, conversely, if $\#P_f \geq 4$ and $f^{-1}(P_f) \subset C_f \cup P_f$, then f has a parabolic orbifold.*

The second part follows from [DH93, Lemma 2], but for the convenience of the reader, we will provide the simple proof. Here the assumption $\#P_f \geq 4$ cannot be omitted as some examples with $\#P_f = 3$ show (such as the Thurston map arising from the ‘barycentric subdivision rule’; see [BM17, Example 12.21]).

Proof. Let α_f be the ramification function of f . Then $p \in P_f$ if and only if $\alpha_f(p) \geq 2$.

First suppose that f has a parabolic orbifold. Then $\alpha_f(q) \cdot \deg(f, q) = \alpha_f(f(q))$ for all $q \in S^2$ (see [BM17, Proposition 2.14]). So if $q \in f^{-1}(P_f)$, then $f(q) \in P_f$ which implies

$$\alpha_f(q) \cdot \deg(f, q) = \alpha_f(f(q)) \geq 2.$$

This is only possible if $\alpha_f(q) \geq 2$ in which case $q \in P_f$, or if $\deg(f, q) \geq 2$ in which case $q \in C_f$. Hence, $q \in C_f \cup P_f$, and so $f^{-1}(P_f) \subset C_f \cup P_f$. Since the reverse inclusion is true for all Thurston maps, we see that $f^{-1}(P_f) = C_f \cup P_f$ if f has a parabolic orbifold.

For the converse, suppose that $f: S^2 \rightarrow S^2$ is an arbitrary Thurston map with $\#P_f \geq 4$ and $f^{-1}(P_f) \subset C_f \cup P_f$. Let $d := \deg(f) \geq 2$. Note that $f^{-1}(P_f) \subset (C_f \setminus P_f) \cup P_f$ by our hypotheses.

Each point $p \in S^2$ has precisely d preimages counting multiplicities, that is,

$$d = \sum_{q \in f^{-1}(p)} \deg(f, q).$$

Furthermore, since $C_f \subset f^{-1}(P_f)$ and $\deg(f, q) \geq 2$ for $q \in S^2$ if and only if $q \in C_f$, the Riemann–Hurwitz formula implies

$$\begin{aligned} \#(C_f \setminus P_f) \leq \#C_f &\leq \sum_{c \in C_f} (\deg(f, c) - 1) = \sum_{q \in f^{-1}(P_f)} (\deg(f, q) - 1) \\ &= 2d - 2. \end{aligned}$$

It follows that

$$\begin{aligned} d \cdot \#P_f &= \sum_{q \in f^{-1}(P_f)} \deg(f, q) = \sum_{q \in f^{-1}(P_f)} (\deg(f, q) - 1) + \#f^{-1}(P_f) \\ &= (2d - 2) + \#f^{-1}(P_f) \leq (2d - 2) + \#(C_f \setminus P_f) + \#P_f \\ &\leq 4(d - 1) + \#P_f. \end{aligned}$$

Hence, $(d - 1) \cdot \#P_f \leq 4(d - 1)$ and so $4 \leq \#P_f \leq 4$. This implies $\#P_f = 4$ and that all the previous inequalities must be equalities. In particular, $\#(C_f \setminus P_f) = 2d - 2$, which shows that at all critical points, the local degree of f is equal to 2 and no critical points belong to P_f .

As a consequence, under iteration of f , the orbit of any point $p \in S^2$ passes through at most one critical point. It follows that we have $\alpha_f(p) = 1$ for $p \in S^2 \setminus P_f$ and $\alpha_f(p) = 2$ for $p \in P_f$. This implies that the Euler characteristic (see equation (3.1)) of the orbifold \mathcal{O}_f associated with f is equal to

$$\chi(\mathcal{O}_f) = 2 - \sum_{p \in P_f} \left(1 - \frac{1}{\alpha_f(p)}\right) = 2 - (1/2 + 1/2 + 1/2 + 1/2) = 0.$$

We conclude that f has a parabolic orbifold. □

3.1. *The $(n \times n)$ -Lattès map.* In general, a *Lattès map* is a rational Thurston map with parabolic orbifold that does not have periodic critical points. Here we provide the analytic definition for the Lattès maps that we use in this paper and interpret this from a more

geometric perspective. See [Mil06b] and [BM17, Ch. 3] for a general discussion of Lattès maps.

Let $\mathbb{P} \cong \widehat{\mathbb{C}}$ be the Euclidean square pillow and $\wp: \mathbb{C} \rightarrow \mathbb{P}$ be its universal orbifold covering map, as discussed in §2.4. Fix a natural number $n \geq 2$. It follows from (2.3) that there is a unique (and well-defined) map $\mathcal{L}_n: \mathbb{P} \rightarrow \mathbb{P}$ such that

$$\mathcal{L}_n(\wp(z)) = \wp(nz) \text{ for } z \in \mathbb{C}. \quad (3.2)$$

We call \mathcal{L}_n the $(n \times n)$ -Lattès map. In fact, \mathcal{L}_n is a rational map under the identification $\mathbb{P} \cong \widehat{\mathbb{C}}$ as discussed in §2.4.

Alternatively, we can describe the map \mathcal{L}_n in a combinatorial fashion as follows. Recall that the front side of \mathbb{P} is colored white, and the back side black. These are the two 0-tiles of \mathbb{P} , and we subdivide each of them into n^2 squares of sidelength $1/n$. We refer to these small squares as 1-tiles (with n understood), and color them in a checkerboard fashion black and white so that the 1-tile S in the white side of \mathbb{P} with the vertex A on its boundary is colored white. We map S to the white side of the pillow \mathbb{P} by an orientation-preserving Euclidean similarity (that scales by the factor n) so that the vertex A is fixed. If we extend this map by reflection to the whole pillow, we get the $(n \times n)$ -Lattès map \mathcal{L}_n (see Figure 1 for $n = 4$). The map \mathcal{L}_n sends each black or white 1-tile homeomorphically (by a similarity) onto the 0-tile in \mathbb{P} of the same color.

Based on this combinatorial description, it is easy to see that each critical point of \mathcal{L}_n has local degree 2 and that the postcritical set of \mathcal{L}_n coincides with the set of vertices of \mathbb{P} , that is, $P_{\mathcal{L}_n} = \{A, B, C, D\} = V$. One can also check that for the ramification function of \mathcal{L}_n , we have $\alpha_{\mathcal{L}_n}(p) = 2$ for each $p \in P_{\mathcal{L}_n}$. Substituting this into equation (3.1), we see that $\chi(O_{\mathcal{L}_n}) = 0$. Thus, \mathcal{L}_n has a parabolic orbifold.

3.2. Thurston's characterization of rational maps. Thurston maps can often be described from a combinatorial viewpoint as the Lattès map \mathcal{L}_n in §3.1 (see, for instance, [CFKP03] and [BM17, Ch. 12]). The question whether a given Thurston map f can be realized by a rational map is usually difficult to answer except in some special cases. William Thurston provided a sharp, purely topological criterion that answers this question. The formulation and proof of this celebrated result can be found in [DH93]. In this section, we introduce the necessary concepts and formulate the result only when $\#P_f = 4$, which is the relevant case for this paper.

In the following, let $f: S^2 \rightarrow S^2$ be a Thurston map. The map f defines a natural pullback operation on Jordan curves in (S^2, P_f) : a pullback of a Jordan curve $\gamma \subset S^2 \setminus P_f$ under f is a connected component $\tilde{\gamma}$ of $f^{-1}(\gamma)$. Since f is a covering map from $S^2 \setminus f^{-1}(P_f)$ onto $S^2 \setminus P_f$, each pullback $\tilde{\gamma}$ of γ is a Jordan curve in (S^2, P_f) . Moreover, $f|_{\tilde{\gamma}}: \tilde{\gamma} \rightarrow \gamma$ is a covering map. For some $k \in \mathbb{N}$ with $1 \leq k \leq \deg(f)$, each point $p \in \gamma$ has precisely k distinct preimages in $\tilde{\gamma}$. Here k is the (unsigned) mapping degree of $f|_{\tilde{\gamma}}$ which we denote by $\deg(f: \tilde{\gamma} \rightarrow \gamma)$.

Recall that a Jordan curve $\gamma \subset S^2 \setminus P_f$ is called *essential* if each of the two connected components of $S^2 \setminus \gamma$ contains at least two points from P_f , and is called *peripheral* otherwise.

We will need the following standard facts.

LEMMA 3.4. Let $f: S^2 \rightarrow S^2$ be a Thurston map and let γ and γ' be Jordan curves in (S^2, P_f) with $\gamma' \sim \gamma$ relative to P_f . Then there is a bijection $\tilde{\gamma} \leftrightarrow \tilde{\gamma}'$ between the pullbacks $\tilde{\gamma}$ of γ and the pullbacks $\tilde{\gamma}'$ of γ' under f such that for all pullbacks corresponding under this bijection, we have $\tilde{\gamma} \sim \tilde{\gamma}'$ relative to P_f and $\deg(f: \tilde{\gamma} \rightarrow \gamma) = \deg(f: \tilde{\gamma}' \rightarrow \gamma')$.

For the proof, see [BM17, Lemma 6.9]. A consequence of this statement is that the isotopy classes of curves in $f^{-1}(\gamma)$ relative to P_f only depend on the isotopy class $[\gamma]$ relative to P_f and not on the specific choice of γ .

COROLLARY 3.5. Let $f: S^2 \rightarrow S^2$ be a Thurston map, and γ be a Jordan curve in (S^2, P_f) .

- (i) If γ is peripheral, then every pullback of γ under f is also peripheral.
- (ii) Suppose that $\#P_f = 4$ and let $\tilde{\gamma}$ be a pullback of γ under f . If γ and $\tilde{\gamma}$ are essential, then the isotopy class $[\tilde{\gamma}]$ relative to P_f only depends on the isotopy class $[\gamma]$ relative to P_f and not on the specific choice of γ and its essential pullback $\tilde{\gamma}$.

Proof. (i) Since γ is peripheral, γ can be isotoped (relative to P_f) into a Jordan curve γ' inside a small open Jordan region $V \subset S^2$ such that $\#(V \cap P_f) \leq 1$ and V is evenly covered by the branched covering map f as in equation (2.1).

Then for each component U_j of $f^{-1}(V)$, the map $f|_{U_j}: U_j \rightarrow V$ is given by $z \in \mathbb{D} \mapsto z^{d_j} \in \mathbb{D}$ for some $d_j \in \mathbb{N}$ after orientation-preserving homeomorphic coordinate changes in the source and target. This implies that $\#(U_j \cap P_f) \leq 1$ and that each pullback of γ' in U_j is peripheral. Hence, all pullbacks of γ' under f are peripheral and the same is true for the pullbacks of γ as follows from Lemma 3.4.

(ii) Suppose $\tilde{\gamma}$ and $\tilde{\gamma}'$ are two distinct essential pullbacks of γ under f . Since these are components of $f^{-1}(\gamma)$, the Jordan curves $\tilde{\gamma}$ and $\tilde{\gamma}'$ are disjoint. Then the set $S^2 \setminus (\tilde{\gamma} \cup \tilde{\gamma}')$ is a disjoint union $S^2 \setminus (\tilde{\gamma} \cup \tilde{\gamma}') = V \cup U \cup V'$, where $V, V' \subset S^2$ are Jordan regions and $U \subset S^2$ is an annulus with $\partial U = \tilde{\gamma} \cup \tilde{\gamma}'$. Since $\tilde{\gamma}$ and $\tilde{\gamma}'$ are essential, both V and V' must contain at least two postcritical points. Now $\#P_f = 4$, and so $U \cap P_f = \emptyset$. Lemma 2.1 then implies that $\tilde{\gamma}$ and $\tilde{\gamma}'$ are isotopic relative to P_f .

It follows that the isotopy class $[\tilde{\gamma}]$ relative to P_f does not depend on the choice of the essential pullback $\tilde{\gamma}$ of γ . At the same time, Lemma 3.4 implies that $[\tilde{\gamma}]$ only depends on the isotopy class $[\gamma]$, as desired. \square

For a general Thurston map f the concept of an invariant multicurve is important to decide whether f is realized or obstructed. By definition, a *multicurve* is a non-empty finite family Γ of essential Jordan curves in $S^2 \setminus P_f$ that are pairwise disjoint and pairwise non-isotopic relative to P_f .

Suppose now that $\#P_f = 4$. Then any two essential Jordan curves in $S^2 \setminus P_f$ are either isotopic relative to P_f or have a non-empty intersection (as follows from the argument in the proof of Corollary 3.5(ii)). Thus, in this case, each multicurve Γ consists of a single essential Jordan curve γ in $S^2 \setminus P_f$. We say that an essential Jordan curve γ in $S^2 \setminus P_f$ is *f-invariant* if each essential pullback of γ under f is isotopic to γ relative to P_f .




FIGURE 8. The two pullbacks of a curve $\gamma = \tau_2$ under \mathcal{L}_2 .

Definition 3.6. Let $f : S^2 \rightarrow S^2$ be a Thurston map with $\#P_f = 4$ and let $\gamma \subset S^2 \setminus P_f$ be an essential f -invariant Jordan curve. We denote by $\gamma_1, \dots, \gamma_n$ for $n \in \mathbb{N}_0$ all the essential pullbacks of γ under f and define

$$\lambda_f(\gamma) := \sum_{j=1}^n \frac{1}{\deg(f : \gamma_j \rightarrow \gamma)}. \tag{3.3}$$

Then γ is called a (Thurston) obstruction for f if $\lambda_f(\gamma) \geq 1$.

Note that if $n = 0$, then the sum in equation (3.3) is the empty sum and so $\lambda_f(\gamma) = 0$. It immediately follows from Lemma 3.4 that whether or not γ is an obstruction for f only depends on the isotopy class $[\gamma]$ relative to P_f .

The following theorem gives a criterion when a Thurston map f with $\#P_f = 4$ is realized.

THEOREM 3.7. (Thurston’s criterion) *Let $f : S^2 \rightarrow S^2$ be a Thurston map with $\#P_f = 4$ and suppose that f has a hyperbolic orbifold. Then f is realized by a rational map if and only if f has no obstruction.*

With a suitable definition of an obstruction (as an invariant multicurve that satisfies certain mapping properties), this statement is also true for general Thurston maps with a hyperbolic orbifold; see [DH93] or [BM17, §2.6].

The example of the $(n \times n)$ -Lattès map $\mathcal{L}_n : \mathbb{P} \rightarrow \mathbb{P}$ with $n \geq 2$ shows that Theorem 3.7 is false if f has a parabolic orbifold. Indeed, let γ be any essential Jordan curve in $(\mathbb{P}, P_{\mathcal{L}_n})$, where $P_{\mathcal{L}_n} = V = \{A, B, C, D\}$ consists of the vertices of the pillow \mathbb{P} . Since only the isotopy class $[\gamma]$ relative to V matters, by Lemma 2.3, we may assume without loss of generality that $\gamma = \tau_{r/s} = \wp(\ell_{r/s}(z_0))$ with $z_0 \in \mathbb{C}, r/s \in \widehat{\mathbb{Q}}$, and $\ell_{r/s}(z_0) \subset \mathbb{C} \setminus \mathbb{Z}^2$. Here r and s are relatively prime integers and so there exist $p, q \in \mathbb{Z}$ such that $pr + qs = 1$. Let $\tilde{\omega} := -p + iq$. Using (2.3) and (3.2), one can verify that under \mathcal{L}_n , the curve γ has exactly n distinct pullbacks

$$\gamma_j = \wp(\ell_{r/s}((z_0 + 2j\tilde{\omega})/n)), \quad j = 1, \dots, n. \tag{3.4}$$

Moreover, each curve γ_j is isotopic to γ relative to $P_{\mathcal{L}_n} = V$ and $\deg(\mathcal{L}_n : \gamma_j \rightarrow \gamma) = n$ for all $j = 1, \dots, n$; see Figure 8 for an illustration. Thus, $\lambda_{\mathcal{L}_n}(\gamma) = 1$ and γ is an obstruction.

4. Blowing up arcs

Here, we describe the operation of ‘blowing up arcs’, originally introduced by Pilgrim and Tan Lei in [PL98, §2.5]. This operation allows us to define and modify various Thurston maps and plays a crucial role in this paper. We will first describe the general construction and then illustrate it for Lattès maps. As we will explain, the procedure of ‘gluing a flap’ that we introduced in §1.1 can be viewed as a special case of blowing up arcs for Lattès maps.

The construction of blowing up arcs will involve a finite collection E of arcs with pairwise disjoint interiors in a 2-sphere S^2 . We denote by V the set of endpoints of these arcs and consider (V, E) as an embedded graph in S^2 . In the construction, we will make various topological choices. The following general statement guarantees that we do not run into topological difficulties. In the formulation, we equip S^2 with a ‘nice’ metric d so that (S^2, d) is isometric to $\widehat{\mathbb{C}}$ carrying the spherical metric with length element $ds = 2|dz|/(1 + |z|^2)$.

PROPOSITION 4.1. *Let $G = (V, E)$ be a planar embedded graph in S^2 and $\mathcal{G} \subset S^2$ be its realization. Then there exists a planar embedded graph $G' = (V, E')$ in S^2 with the same vertex set such that its realization \mathcal{G}' is isotopic to \mathcal{G} relative to V and such that each edge of G' is a piecewise geodesic arc in (S^2, d) .*

An outline of the proof is given in [Bol79, Ch. I, §4]; the proposition also follows from [Bus10, Lemma A.8].

4.1. The general construction. Before we provide a formal definition, we give some rough idea of how to ‘blow up’ arcs. In the following, $f: S^2 \rightarrow S^2$ is a fixed Thurston map. Let e be an arc in S^2 such that the restriction $f|_e$ is a homeomorphism onto its image. We cut the sphere S^2 open along e and glue in a closed Jordan region D along the boundary. In this way, we obtain a new 2-sphere on which we can define a branched covering map \widehat{f} as follows: \widehat{f} maps the complement of $\text{int}(D)$ in the same way as the original map f and it maps $\text{int}(D)$ to the complement of $f(e)$ by a homeomorphism that matches the map $f|_e$. We say that \widehat{f} is obtained from f by *blowing up* the arc e with multiplicity 1.

Now we proceed to give a rigorous definition of the blow-up operation in the general case, where several arcs e are blown up simultaneously with possibly different multiplicities $m_e \geq 1$ resulting in a new Thurston map \widehat{f} . To this end, let E be a finite set of arcs in $(S^2, f^{-1}(P_f))$ with pairwise disjoint interiors such that the restriction $f|_e: e \rightarrow f(e)$ is a homeomorphism for each $e \in E$. In this case, we say that E satisfies the *blow-up conditions*.

We assume that each arc $e \in E$ has an assigned *multiplicity* $m_e \in \mathbb{N}$. Since each $e \in E$ is an arc in $(S^2, f^{-1}(P_f))$, its interior $\text{int}(e)$ is disjoint from $f^{-1}(P_f) \supset P_f$ and so $\text{int}(e)$ does not contain any critical or postcritical point of f .

For each arc $e \in E$, we choose an open Jordan region $W_e \subset S^2$ so that the following conditions hold:

- (A1) the open Jordan regions $W_e, e \in E$, are pairwise disjoint;
- (A2) for distinct arcs $e_1, e_2 \in E$, we have $\text{cl}(W_{e_1}) \cap \text{cl}(W_{e_2}) = \partial e_1 \cap \partial e_2$;

Choose regions $W_e \supset e$

Open up the arcs e
to regions $D_e \subset W_e$

Subdivide each D_e into
regions D_e^1, \dots, D_e^m




FIGURE 9. Setup for blowing up the arcs e_1 and e_2 (in the sphere on the left) with the multiplicities $m_{e_1} = 1$ and $m_{e_2} = 2$.

- (A3) $\text{int}(e) \subset W_e$ and $\partial e \subset \partial W_e$ for each $e \in E$;
- (A4) $\text{cl}(W_e) \cap f^{-1}(P_f) = e \cap f^{-1}(P_f) = \partial e$ for each $e \in E$;
- (A5) $f|_{\text{cl}(W_e)}$ is a homeomorphism onto its image for each $e \in E$.

The existence of such a choice (and also of the choices below) can easily be justified based on Proposition 4.1 and we will skip the details.

Let $e \in E$ and W_e be chosen as above. Then we choose a closed Jordan region D_e so that e is a crosscut in D_e and $D_e \setminus \partial e \subset W_e$. The two endpoints of e lie on the Jordan curve ∂D_e and partition it into two arcs, which we denote by ∂D_e^+ and ∂D_e^- . One can think of D_e as the resulting region if e has been ‘opened up’. This is illustrated in the left and middle parts of Figure 9.

To define the desired Thurston map \hat{f} , we want to collapse D_e back to e . For this, we choose a continuous map $h : S^2 \times \mathbb{I} \rightarrow S^2$ with the following properties:

- (B1) h is a *pseudo-isotopy*, that is, $h_t := h(\cdot, t)$ is a homeomorphism on S^2 for each $t \in [0, 1)$;
- (B2) h_0 is the identity map on S^2 ;
- (B3) h_t is the identity map on $S^2 \setminus \bigcup_{e \in E} W_e$ for each $t \in [0, 1]$;
- (B4) h_1 is a homeomorphism of $S^2 \setminus \bigcup_{e \in E} D_e$ onto $S^2 \setminus \bigcup_{e \in E} e$, and h_1 maps ∂D_e^+ and ∂D_e^- homeomorphically onto e for each $e \in E$.

It is easy to see that if we equip S^2 with some metric, then the set $h_t(D_e)$ Hausdorff converges to e as $t \rightarrow 1^-$. This implies that $h_1(D_e) = e$. So intuitively, the deformation process described by h collapses each closed Jordan region D_e to e at time 1 so that the points in $S^2 \setminus \bigcup_{e \in E} W_e$ remain fixed.

We now make yet another choice. For a fixed arc $e \in E$, let $m = m_e$. We choose $m - 1$ crosscuts e^1, \dots, e^{m-1} in D_e with the same endpoints as e such that these crosscuts have pairwise disjoint interiors. We set $e^0 := \partial D_e^+$ and $e^m := \partial D_e^-$. The arcs e^0, \dots, e^m subdivide the closed Jordan region D_e into m closed Jordan regions D_e^1, \dots, D_e^m , called *components* of D_e . This is illustrated in the right-hand part of Figure 9.

We may assume that the labeling is such that $\partial D_e^k = e^{k-1} \cup e^k$ for $k = 1, \dots, m$. For each $k = 1, \dots, m$, we now choose a continuous map $\varphi_k : D_e^k \rightarrow S^2$ with the following properties:

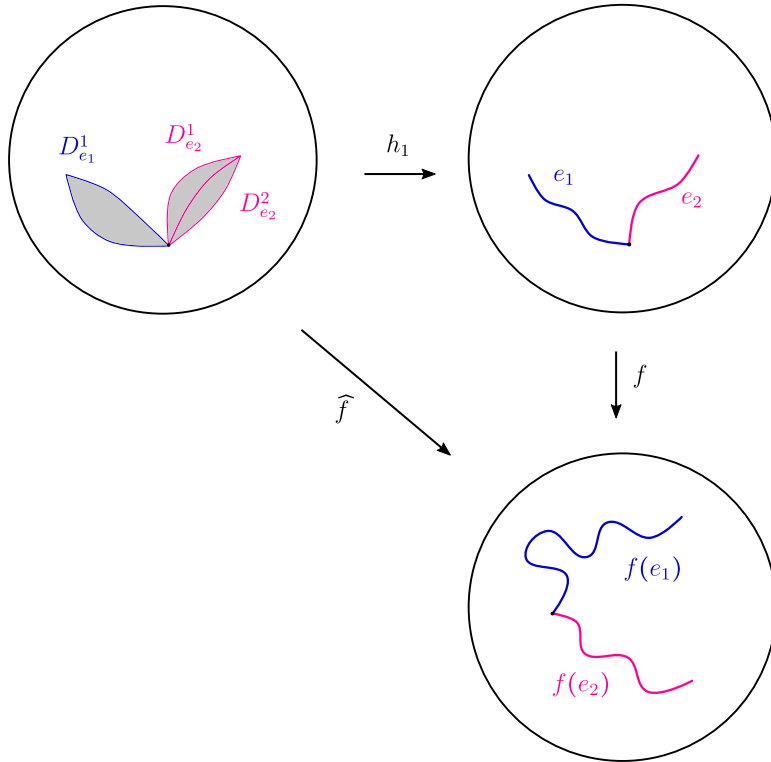


FIGURE 10. The map \widehat{f} is obtained from f by blowing up the arcs e_1 and e_2 with multiplicities $m_{e_1} = 1$ and $m_{e_2} = 2$.

- (C1) φ_k is an orientation-preserving homeomorphism of $\text{int}(D_e^k)$ onto $S^2 \setminus f(e)$ and maps e^{k-1} and e^k homeomorphically onto $f(e)$;
- (C2) $\varphi_1|e^0 = f \circ h_1|e^0$, $\varphi_m|e^m = f \circ h_1|e^m$, and $\varphi_k|e^k = \varphi_{k+1}|e^k$ for $k = 1, \dots, m - 1$.

Note that by the earlier discussion, h_1 maps $e^0 = \partial D_e^+$ and $e^m = \partial D_e^-$ homeomorphically onto e and f is a homeomorphism of e onto $f(e)$. These choices of the maps φ_k depend on e , but we suppress this in our notation for simplicity.

A map $\widehat{f}: S^2 \rightarrow S^2$ can now be defined as follows:

- (D1) if $p \in S^2 \setminus \bigcup_{e \in E} \text{int}(D_e)$, we set $\widehat{f}(p) = f(h_1(p))$;
- (D2) if $p \in D_e$ for some $e \in E$, then p lies in one of the components D_e^k of D_e and we set $\widehat{f}(p) = \varphi_k(p)$.

The matching conditions in property (C2) above immediately imply that \widehat{f} is well defined and continuous.

Definition 4.2. We say that the map $\widehat{f}: S^2 \rightarrow S^2$ as described above is obtained from the Thurston map f by *blowing up each arc $e \in E$ with multiplicity m_e* .

Figure 10 illustrates the construction of \widehat{f} . Here, we blow up the arcs e_1 and e_2 from Figure 9 with multiplicities $m_{e_1} = 1$ and $m_{e_2} = 2$. The arcs $f(e_1)$ and $f(e_2)$ share an




FIGURE 11. Setup for blowing up the arc set $E = \{e_1, e_2\}$ (on the left pillow) with $m_{e_1} = 1$ and $m_{e_2} = 2$.




FIGURE 12. The map \hat{f} obtained from $f = \mathcal{L}_2$ by blowing up the arcs in the set $E = \{e_1, e_2\}$ illustrated in Figure 11 with $m_{e_1} = 1$ and $m_{e_2} = 2$.

endpoint (since e_1 and e_2 do), but in general, they could have more points in common or even coincide. For simplicity, we chose to draw them with disjoint interiors.

By construction, \hat{f} acts in a similar way as f outside the closed Jordan regions $D_{e_1} = D_{e_1}^1$ and $D_{e_2} = D_{e_2}^1 \cup D_{e_2}^2$. More precisely, the map \hat{f} equals $f \circ h_1$ on $S^2 \setminus (\text{int}(D_{e_1}) \cup \text{int}(D_{e_2}))$, where h_1 collapses the closed Jordan regions D_{e_1} and D_{e_2} onto e_1 and e_2 , respectively. At the same time, \hat{f} maps $\text{int}(D_{e_1})$ homeomorphically onto $S^2 \setminus f(e_1)$, and each of the regions $\text{int}(D_{e_2}^1)$ and $\text{int}(D_{e_2}^2)$ homeomorphically onto $S^2 \setminus f(e_2)$.

In the next section, we want to relate ‘blowing up arcs’ with ‘gluing flaps’ as discussed in the introduction. To set this up, we consider the (2×2) -Lattès map $f = \mathcal{L}_2$. We choose two edges e_1 and e_2 of a 1-tile in \mathbb{P} as shown in the pillow on the left in Figure 11. Note that f sends e_1 and e_2 homeomorphically onto the edges c and b of \mathbb{P} , respectively. Thus, the set $E = \{e_1, e_2\}$ satisfies the blow-up conditions. Figure 11 illustrates the setup for blowing up these arcs e_1 and e_2 with the multiplicities $m_{e_1} = 1$ and $m_{e_2} = 2$. The resulting map $\hat{f}: \mathbb{P} \rightarrow \mathbb{P}$ is shown in Figure 12. The points marked by a dot on the left pillow \mathbb{P} (the domain of the map) correspond to the preimage points $\hat{f}^{-1}(V)$. The pillow on the left is subdivided into closed Jordan regions alternately colored black and white. The map \hat{f} sends each of these closed Jordan regions U homeomorphically onto the back side or front side of the pillow \mathbb{P} depending on whether U is black or white.

The following statement summarizes the main properties of maps \hat{f} as in Definition 4.2.

LEMMA 4.3. *Let $f: S^2 \rightarrow S^2$ be a Thurston map and E be a set of arcs in $(S^2, f^{-1}(P_f))$ satisfying the blow-up conditions. Suppose $\hat{f}: S^2 \rightarrow S^2$ is the map obtained by blowing up each arc $e \in E$ with multiplicity $m_e \in \mathbb{N}$.*

Then \widehat{f} is a Thurston map with $P_{\widehat{f}} = P_f$. Moreover, the map \widehat{f} is uniquely determined up to Thurston equivalence independently of the choices in the above construction. More precisely, up to Thurston equivalence \widehat{f} depends only on the original map f , the isotopy classes of the arcs in E relative to $f^{-1}(P_f)$, and the multiplicities m_e for $e \in E$.

Proof. By construction, \widehat{f} is an orientation-preserving local homeomorphism near each point $p \in S^2 \setminus f^{-1}(P_f)$. By considering the number of preimages of a generic point in S^2 , we see that the topological degree of \widehat{f} is equal to $\deg(f) + \sum_{e \in E} m_e > 0$. The fact that $\widehat{f}: S^2 \rightarrow S^2$ is a branched covering map can now be deduced from [BM17, Corollary A.14].

We have $\deg(\widehat{f}, p) = 1$ for $p \in S^2 \setminus f^{-1}(P_f)$ and

$$\deg(\widehat{f}, p) = \deg(f, p) + \sum_{\{e \in E: p \in e\}} m_e$$

for $p \in f^{-1}(P_f)$. This implies $C_f \subset C_{\widehat{f}} \subset f^{-1}(P_f)$. Since on the set $f^{-1}(P_f)$ the maps f and \widehat{f} agree, it follows that $P_{\widehat{f}} = P_f$. So \widehat{f} has a finite postcritical set, and we conclude that \widehat{f} is indeed a Thurston map.

We omit a detailed justification of the second claim that \widehat{f} is uniquely determined up to Thurston equivalence by f , the isotopy classes of the arcs E , and their multiplicities. A proof can be given along the lines of [PL98, Proposition 2]. □

Remark 4.4. If in the previous statement $E \neq \emptyset$ and $\#P_f \geq 3$, then \widehat{f} has a hyperbolic orbifold. To see this, pick an arc $e \in E$. Then $f(e)$ has at most two points in common with P_f , and so we can find a point $p \in P_f \setminus f(e) \subset P_{\widehat{f}}$. Then it follows from the construction of \widehat{f} that there exists a point q in the interior of the region D_e associated with e such that $\widehat{f}(q) = p$ and $\deg(\widehat{f}, q) = 1$. Then $q \notin C_{\widehat{f}}$, but we also have

$$q \in \text{int}(D_e) \subset S^2 \setminus f^{-1}(P_f) \subset S^2 \setminus P_f = S^2 \setminus P_{\widehat{f}}.$$

This shows that $q \in \widehat{f}^{-1}(p) \subset \widehat{f}^{-1}(P_{\widehat{f}})$, but $q \notin C_{\widehat{f}} \cup P_{\widehat{f}}$, and so \widehat{f} must have a hyperbolic orbifold by the first part of Lemma 3.3.

4.2. *Blowing up the $(n \times n)$ -Lattès map.* Let \mathbb{P} be the Euclidean square pillow and $\mathcal{L}_n: \mathbb{P} \rightarrow \mathbb{P}$ be the $(n \times n)$ -Lattès map for fixed $n \geq 2$. We denote by $C \subset \mathbb{P}$ the common boundary of the two sides of \mathbb{P} . The set C may be viewed as a planar embedded graph with the vertex set $V = \{A, B, C, D\}$ and the edge set $\{a, b, c, d\}$ in the notation from §2.4. Let $\widetilde{C} := \mathcal{L}_n^{-1}(C) \subset \mathbb{P}$ be the preimage of C under \mathcal{L}_n , viewed as a planar embedded graph with the vertex set $\mathcal{L}_n^{-1}(V)$. In the next section, we will study the question whether a Thurston map is realized by a rational map if it is obtained from \mathcal{L}_n by blowing up edges of \widetilde{C} . To facilitate this discussion, we will provide a more concrete combinatorial model for these maps.

By the definition of the map \mathcal{L}_n , the graph \widetilde{C} subdivides the pillow \mathbb{P} into $2n^2$ 1-tiles, which are squares of sidelength $1/n$. The edges of the embedded graph \widetilde{C} are precisely the sides of these squares. We call them the 1-edges of \mathbb{P} (for given n). The map \mathcal{L}_n sends each

1-edge e of \mathbb{P} homeomorphically onto one of the edges a, b, c, d of C . We call e *horizontal* if \mathcal{L}_n maps it onto a or c , and *vertical* if \mathcal{L}_n maps it onto b or d .

We take two disjoint copies of the Euclidean square $[0, 1/n]^2 \subset \mathbb{R}^2$ and identify the points on three of their sides, say the sides $\{0\} \times [0, 1/n]$, $[0, 1/n] \times \{1/n\}$, and $\{1/n\} \times [0, 1/n]$. We call the object obtained a *flap* F . Note that it is homeomorphic to the closed unit disk and has two ‘free’ sides corresponding to the two copies of $[0, 1/n] \times \{0\}$ in F .

We can cut the pillow \mathbb{P} open along one of the edges of \tilde{C} and glue in a flap F to the pillow by identifying each copy of $[0, 1/n] \times \{0\}$ in the flap with one side of the slit by an isometry (see Figure 2). In this way, we get a new polyhedral surface homeomorphic to S^2 . One can also glue multiple copies of the flap to the slit by an isometry and obtain a polyhedral surface $\hat{\mathbb{P}}$ homeomorphic to S^2 . This can be described more concretely as follows. Let e be an edge in \tilde{C} and F_1, \dots, F_m be $m \geq 1$ copies of the flap. For each $k = 1, \dots, m$, we denote the two copies of $[0, 1/n] \times \{0\}$ in the flap F_k by e'_k and e''_k . We now construct a new polyhedral surface $\hat{\mathbb{P}}$ in the following way.

- (i) First, we cut the original pillow \mathbb{P} open along the edge e .
- (ii) Then, for each $k = 1, \dots, m - 1$, we identify the edge e''_k of F_k with the edge e'_{k+1} of F_{k+1} by an isometry. We get a polyhedral surface D_e homeomorphic to a closed disk, whose boundary consists of two edges e'_1 and e''_m .
- (iii) Finally, we glue the disk D_e to the pillow \mathbb{P} cut open along e by identifying the edges e'_1 and e''_m in ∂D_e with the two sides of the slit by an isometry so that e'_1 and e''_m are identified with different sides of the slit. We obtain a polyhedral surface $\hat{\mathbb{P}}$ that is homeomorphic to a 2-sphere.

More generally, we can cut open \mathbb{P} simultaneously along several edges e of \tilde{C} and, by the method described, glue $m_e \in \mathbb{N}$ copies of the flap to the slit obtained from each edge e . If these edges e of \tilde{C} with their multiplicities m_e are given, then there is essentially only one way of gluing flaps so that the resulting object is a polyhedral surface homeomorphic to S^2 .

Let $\hat{\mathbb{P}}$ be the polyhedral surface obtained from \mathbb{P} by gluing a total number of $n_h \geq 0$ *horizontal flaps* (that is, flaps glued along horizontal edges of \tilde{C}) and a total number of $n_v \geq 0$ *vertical flaps* (that is, flaps glued along vertical edges of \tilde{C}). We call this surface a *flapped pillow*. We denote by E the set of all edges in \tilde{C} along which flaps were glued and by $m_e, e \in E$, the corresponding multiplicities. See the left part of Figure 13 for an example of a flapped pillow $\hat{\mathbb{P}}$ obtained by gluing one horizontal and two vertical flaps at the edges e_1 and e_2 from Figure 11.

The polyhedral surface $\hat{\mathbb{P}}$ is naturally subdivided into

$$2(n^2 + n_h + n_v) = 2n^2 + 2 \sum_{e \in E} m_e$$

squares of sidelength $1/n$, called the *1-tiles* of the flapped pillow $\hat{\mathbb{P}}$. The vertices and the edges of these squares are called the *1-vertices* and *1-edges* of $\hat{\mathbb{P}}$. There is a natural path metric on $\hat{\mathbb{P}}$ that agrees with the Euclidean metric on each 1-tile. The surface $\hat{\mathbb{P}}$ equipped with this metric is locally Euclidean with conic singularities at some of the 1-vertices. Such a conic singularity arises at a 1-vertex $v \in \hat{\mathbb{P}}$ if v is contained in $k_v \neq 4$ distinct 1-tiles.




FIGURE 13. A branched covering map $\widehat{\mathcal{L}}: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ induced by the flapped pillow $\widehat{\mathbb{P}}$ on the left.




FIGURE 14. The base $B(\widehat{\mathbb{P}})$ of the flapped pillow $\widehat{\mathbb{P}}$ from Figure 13 depicted in two different ways: as a subset of $\widehat{\mathbb{P}}$ and as the subset $\mathbb{P} \setminus \bigcup_{e \in E} \text{int}(e)$ of \mathbb{P} .

We will assume that $\widehat{\mathbb{P}}$ has at least one flap, that is, $n_h + n_v \geq 1$. Let F_j , $j = 1, \dots, n_h + n_v$, be the collection of flaps glued to \mathbb{P} . Each flap F_j consists of two 1-tiles in $\widehat{\mathbb{P}}$. We call the four 1-vertices that belong to F_j the *vertices* of the flap F_j . The boundary ∂F_j is a Jordan curve in $\widehat{\mathbb{P}}$ composed of two 1-edges e'_j and e''_j , which we call the *base edges* of F_j . The 1-edge in F_j that is opposite to the base edges is called the *top edge* of the flap F_j . Note that $\partial e'_j = \partial e''_j$ consists of two vertices of F_j .

We now define the *base* $B(\widehat{\mathbb{P}}) \subset \widehat{\mathbb{P}}$ of the flapped pillow as

$$B(\widehat{\mathbb{P}}) := \widehat{\mathbb{P}} \setminus \left(\bigcup_{j=1}^{n_h+n_v} (F_j \setminus \partial e'_j) \right). \tag{4.1}$$

In other words, $B(\widehat{\mathbb{P}})$ is obtained from $\widehat{\mathbb{P}}$ by removing all flaps F_j from $\widehat{\mathbb{P}}$, except that we keep the two vertices in $\partial e'_j \subset F_j$ from each flap. There is a natural identification

$$B(\widehat{\mathbb{P}}) \cong \mathbb{P} \setminus \bigcup_{e \in E} \text{int}(e) \subset \mathbb{P}. \tag{4.2}$$

This means that we can consider the base $B(\widehat{\mathbb{P}})$ both as a subset of $\widehat{\mathbb{P}}$ and of \mathbb{P} . Figure 14 illustrates these two viewpoints. This is slightly imprecise, but this point of view will be extremely convenient in the following.

We choose the orientation on $\widehat{\mathbb{P}}$ so that the induced orientation on $B(\widehat{\mathbb{P}})$ considered as a subset of $\widehat{\mathbb{P}}$ coincides with the orientation on $B(\widehat{\mathbb{P}})$ considered as a subset of the oriented sphere \mathbb{P} (if we represent orientations on surfaces by flags, as described in [BM17, §A.4], then we simply pick a positively oriented flag contained in $B(\widehat{\mathbb{P}}) \subset \mathbb{P}$ and declare it to be positively oriented in $\widehat{\mathbb{P}} \supset B(\widehat{\mathbb{P}})$).

The set $B(\widehat{\mathbb{P}}) \cong \mathbb{P} \setminus \bigcup_{e \in E} \text{int}(e)$ contains the vertex set $\mathcal{L}_n^{-1}(V) \subset \mathbb{P}$ of the graph $\widetilde{C} = \mathcal{L}_n^{-1}(C) \subset \mathbb{P}$. This means that we can naturally view each vertex of \widetilde{C} also as a 1-vertex in $\widehat{\mathbb{P}}$. Let $\widehat{A}, \widehat{B}, \widehat{C}, \widehat{D}$ be the 1-vertices of $\widehat{\mathbb{P}}$ that correspond to the vertices A, B, C, D of the original pillow, respectively. We set $\widehat{V} := \{\widehat{A}, \widehat{B}, \widehat{C}, \widehat{D}\}$ and call $\widehat{A}, \widehat{B}, \widehat{C}, \widehat{D}$ the *vertices* of $\widehat{\mathbb{P}}$.

Recall from §3.1 that the faces of the embedded graph $\widetilde{C} \subset \mathbb{P}$ are colored black and white in a checkerboard manner. This coloring induces a checkerboard coloring on the 1-tiles of the flapped pillow $\widehat{\mathbb{P}}$. The original map $\mathcal{L}_n: \mathbb{P} \rightarrow \mathbb{P}$ can now be naturally extended to a continuous map $\widehat{\mathcal{L}}: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ by reflection so that it preserves the coloring: $\widehat{\mathcal{L}}$ maps each 1-tile of $\widehat{\mathbb{P}}$ by a Euclidean similarity (scaling distances by the factor n) onto the 0-tile of \mathbb{P} with the same color; see Figure 13 for an illustration. On the base $B(\widehat{\mathbb{P}})$, the map $\widehat{\mathcal{L}}$ agrees with the original $(n \times n)$ -Lattès map \mathcal{L}_n (if we consider $B(\widehat{\mathbb{P}})$ as a subset of \mathbb{P} by the identification in (4.2)).

It is clear that $\widehat{\mathcal{L}}: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ is a branched covering map. This map is essentially the Thurston map obtained from \mathcal{L}_n by blowing up each arc $e \in E$ with multiplicity m_e . To make this more precise, we need a suitable identification of the source $\widehat{\mathbb{P}}$ with the target \mathbb{P} of $\widehat{\mathcal{L}}$ so that we obtain a self-map on \mathbb{P} . For this, we choose a natural homeomorphism $\phi: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$, which we will now define.

We view the set $\widehat{C} := \widehat{\mathcal{L}}^{-1}(C) \subset \widehat{\mathbb{P}}$ as a planar embedded graph, whose vertices and edges are precisely the 1-vertices and the 1-edges of the flapped pillow $\widehat{\mathbb{P}}$. Each 1-edge of $\widehat{\mathbb{P}}$ is homeomorphically mapped by $\widehat{\mathcal{L}}$ onto one of the edges of \mathbb{P} . Similarly as before, the 1-edges of $\widehat{\mathbb{P}}$ that are mapped by $\widehat{\mathcal{L}}$ onto a or c are called *horizontal*, while the 1-edges of $\widehat{\mathbb{P}}$ that are mapped by $\widehat{\mathcal{L}}$ onto b or d are called *vertical*.

There is a simple path of length n in the graph \widehat{C} that connects the vertices \widehat{A} and \widehat{B} . Clearly, any such path consists only of horizontal 1-edges in $\widehat{\mathbb{P}}$. We denote by \widehat{a} the realization of the chosen path in the sphere $\widehat{\mathbb{P}}$, which is an arc in $(\widehat{\mathbb{P}}, \widehat{V})$. The arc \widehat{a} may not be uniquely determined (namely, if flaps have been glued to slits obtained from edges $e \subset a$), but any two such arcs are isotopic relative to \widehat{V} . We define $\widehat{b}, \widehat{c}, \widehat{d}$ in a similar way and call \widehat{a}, \widehat{c} the *horizontal edges*, and \widehat{b}, \widehat{d} the *vertical edges* of $\widehat{\mathbb{P}}$.

We now choose an orientation-preserving homeomorphism $\phi: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ that sends $\widehat{A}, \widehat{B}, \widehat{C}, \widehat{D}$ to A, B, C, D , and $\widehat{a}, \widehat{b}, \widehat{c}, \widehat{d}$ to a, b, c, d , respectively. We define $f := \widehat{\mathcal{L}} \circ \phi^{-1}$, which is a self-map on \mathbb{P} . Clearly, f is a branched covering map on \mathbb{P} . To refer to this map, we say that $f: \mathbb{P} \rightarrow \mathbb{P}$ is *obtained from the $(n \times n)$ -Lattès map \mathcal{L}_n by gluing n_h horizontal and n_v vertical flaps to \mathbb{P}* . More informally, we call both maps f and $\widehat{\mathcal{L}}$ a *blown-up Lattès map*.

A point in $\widehat{\mathbb{P}}$ is a critical point for $\widehat{\mathcal{L}}: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ if and only if it is on the boundary of at least four 1-tiles subdividing $\widehat{\mathbb{P}}$. This implies that the set $C_{\widehat{\mathcal{L}}}$ of critical points of $\widehat{\mathcal{L}}$ consists of 1-vertices of $\widehat{\mathbb{P}}$ and that each critical point of \mathcal{L}_n is also a critical point for $\widehat{\mathcal{L}}$ (recall that

we can view each point in $\mathcal{L}_n^{-1}(V) \supset C_{\mathcal{L}_n}$ as a 1-vertex in $\widehat{\mathbb{P}}$. Moreover, if a 1-vertex of $\widehat{\mathbb{P}}$ is a critical point of $\widehat{\mathcal{L}}$, but not of \mathcal{L}_n , then it must be one of the points in \widehat{V} . For example, $\widehat{A} \in \widehat{V}$ is a critical point of $\widehat{\mathcal{L}}$ if and only if a flap was glued to an edge of \widehat{C} incident to $A \cong \widehat{A}$. In any case, since $\widehat{\mathcal{L}}$ sends the 1-vertices of $\widehat{\mathbb{P}}$ to the vertices of \mathbb{P} , the postcritical set of $f = \widehat{\mathcal{L}} \circ \phi^{-1}$ coincides with the vertex set V . Thus, f is a Thurston map.

Since we assumed that $\widehat{\mathbb{P}}$ contains at least one flap (that is, $n_h + n_v > 0$), the orbifold of the Thurston map f is hyperbolic. Indeed, each 1-vertex of $\widehat{\mathbb{P}}$ that is a critical point of \mathcal{L}_n is also a critical point of $\widehat{\mathcal{L}}$ with the same or larger local degree. Since we glued at least one flap, there is at least one 1-vertex v contained in six or more 1-tiles of $\widehat{\mathbb{P}}$. Then $\deg(\widehat{\mathcal{L}}, v) = \deg(f, v') \geq 3$, where $v' = \phi(v)$. Now

$$X := f(v') = \widehat{\mathcal{L}}(v) \in V = \{A, B, C, D\} = P_f,$$

and so for the ramification function α_f of f , we have $\alpha_f(X) \geq 3$. However, for all other points $Y \in V = P_f$, we have $\alpha_f(Y) \geq \alpha_{\mathcal{L}_n}(Y) \geq 2$. It then follows from equation (3.1) that $\chi(O_f) < 0$, and so f has indeed a hyperbolic orbifold.

The homeomorphism ϕ chosen in the definition of f is not unique, but any two such homeomorphisms are isotopic relative to \widehat{V} (this easily follows from [Bus10, Theorem A.5]). This implies that f is uniquely determined up to Thurston equivalence. This map may be viewed (up to Thurston equivalence) as the result of the blowing up operation introduced in §4.1 applied to the edges $e \in E$ with the multiplicities m_e . In particular, if we run the procedure for the map $\widehat{\mathcal{L}}$ indicated in Figure 13, then we obtain the map \widehat{f} illustrated in Figure 12 (up to Thurston equivalence).

5. Realizing blown-up Lattès maps

The goal of this section is to determine when a blown-up Lattès map is realized by a rational map. In particular, we will apply Thurston’s criterion to prove Theorem 1.2. The strategies and techniques used in the proof will highlight the main ideas needed for establishing the more general Theorem 1.1.

We fix $n \geq 2, n_h, n_v \geq 0$, and follow the notation introduced in §4.2. In particular, we denote by $\widehat{\mathbb{P}}$ a flapped pillow with n_h horizontal and n_v vertical flaps, by $\widehat{\mathcal{L}}: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ the respective blown-up $(n \times n)$ -Lattès map, and by $\phi: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ the identifying homeomorphism. Then $f: \mathbb{P} \rightarrow \mathbb{P}$ given as $f = \widehat{\mathcal{L}} \circ \phi^{-1}$ is the Thurston map under consideration. We will assume that $n_h + n_v > 0$, and so $\widehat{\mathbb{P}}$ has at least one flap. In this case, f has a hyperbolic orbifold as we have seen, and so we can apply Thurston’s criterion. For this, we consider essential Jordan curves γ in $(\mathbb{P}, P_f) = (\mathbb{P}, V)$ and study their (essential) pullbacks under f .

If γ is such a curve, then the homeomorphism ϕ sends the pullbacks of γ under $\widehat{\mathcal{L}}$ to the pullbacks of γ under f . So to understand the isotopy types and mapping properties of the pullbacks under f , we will instead look at the pullbacks of γ under $\widehat{\mathcal{L}}$. In particular, if $\widehat{\gamma}$ is a pullback of γ under $\widehat{\mathcal{L}}$, then $\deg(\widehat{\mathcal{L}}: \widehat{\gamma} \rightarrow \gamma) = \deg(f: \phi(\widehat{\gamma}) \rightarrow \gamma)$ and $\phi(\widehat{\gamma})$ is essential in $(\mathbb{P}, P_f) = (\mathbb{P}, V)$ if and only if $\widehat{\gamma}$ is essential in $(\widehat{\mathbb{P}}, \widehat{V})$, where \widehat{V} denotes the vertex set of the flapped pillow $\widehat{\mathbb{P}}$.




FIGURE 15. Pullbacks of α^h for a blown-up (4×4) -Lattès map with $n_h = n_v = 1$.

Since the mapping $\widehat{\mathcal{L}}: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ is a similarity map on each 1-tile of $\widehat{\mathbb{P}}$, the preimage $\widehat{\mathcal{L}}^{-1}(\gamma)$ of a Jordan curve γ in (\mathbb{P}, P_f) (or of any subset γ of \mathbb{P}), can be obtained in the following intuitive way: we rescale and copy the part of γ that belongs to the white 0-tile of \mathbb{P} into each white 1-tile of $\widehat{\mathbb{P}}$ and the part of γ that belongs to the black 0-tile of \mathbb{P} into each black 1-tile.

5.1. *The horizontal and vertical curves.* Recall that α^h and α^v (see (2.4)) denote the Jordan curves in $(\mathbb{P}, V) = (\mathbb{P}, P_f)$ that separate the two horizontal and the two vertical edges of \mathbb{P} , respectively. These two curves are invariant under f and will play a crucial role in the considerations of this section.

LEMMA 5.1. *Let $f = \widehat{\mathcal{L}} \circ \phi^{-1}: \mathbb{P} \rightarrow \mathbb{P}$ be a Thurston map obtained from the $(n \times n)$ -Lattès map, $n \geq 2$, by gluing $n_h \geq 0$ horizontal and $n_v \geq 0$ vertical flaps to \mathbb{P} . Then the following statements are true.*

- (i) *The Jordan curve α^h has $n + n_h$ pullbacks under f . Exactly n of these pullbacks are essential. Each of these essential pullbacks is isotopic to α^h relative to P_f .*
- (ii) *If $\tilde{\alpha}$ is one of the n essential pullbacks of α^h , then $\deg(f: \tilde{\alpha} \rightarrow \alpha^h) = n + n_{\tilde{\alpha}}$, where $n_{\tilde{\alpha}} \geq 0$ is the number of distinct vertical flaps in $\widehat{\mathbb{P}}$ that $\phi^{-1}(\tilde{\alpha})$ meets.*
- (iii) *We have*

$$\lambda_f(\alpha^h) = \sum_{\tilde{\alpha}} \frac{1}{n + n_{\tilde{\alpha}}},$$

where the sum is taken over all essential pullbacks $\tilde{\alpha}$ of α^h under f .

Analogous statements are true for the curve α^v .

Proof. Figure 15 illustrates the proof. It is obvious that the curve α^h has exactly $n + n_h$ distinct pullbacks under $\widehat{\mathcal{L}}$. Among them, there are n essential pullbacks $\widehat{\alpha}_1, \dots, \widehat{\alpha}_n$ that separate the two horizontal edges of $\widehat{\mathbb{P}}$ and thus are isotopic to each other relative to the vertex set \widehat{V} of $\widehat{\mathbb{P}}$. For each $j = 1, \dots, n$, the image $\alpha_j := \phi(\widehat{\alpha}_j)$ is isotopic to α^h . Moreover, we have $\deg(\widehat{\mathcal{L}}: \widehat{\alpha}_j \rightarrow \alpha^h) = n + n_{\alpha_j}$. The other n_h pullbacks of α^h under $\widehat{\mathcal{L}}$ are each contained in one of the horizontal flaps and thus are peripheral in $(\widehat{\mathbb{P}}, \widehat{V})$. Consequently,




FIGURE 16. Pullbacks of α^h for a blown up (4×4) -Lattès map with $n_h = 1$ and $n_v = 0$.

$$\lambda_f(\alpha^h) = \sum_{j=1}^n \frac{1}{\deg(f: \alpha_j \rightarrow \alpha^h)} = \sum_{j=1}^n \frac{1}{\deg(\widehat{\mathcal{L}}: \widehat{\alpha}_j \rightarrow \alpha^h)} = \sum_{j=1}^n \frac{1}{n + n_{\alpha_j}}.$$

This completes the proof of the lemma for the curve α^h . The proof for the curve α^v follows from similar considerations. □

The following corollary is an immediate consequence of the previous lemma.

COROLLARY 5.2. *Let $f = \widehat{\mathcal{L}} \circ \phi^{-1}: \mathbb{P} \rightarrow \mathbb{P}$ be a Thurston map obtained from the $(n \times n)$ -Lattès map, $n \geq 2$, by gluing $n_h \geq 0$ horizontal and $n_v \geq 0$ vertical flaps to \mathbb{P} . Then α^h is an obstruction (for f) if and only if $n_v = 0$, and α^v is an obstruction if and only if $n_h = 0$.*

Proof. Let us first suppose that $n_v = 0$, that is, $\widehat{\mathbb{P}}$ does not have any vertical flaps. Then by Lemma 5.1, α^h has n essential pullbacks under f , each of which is mapped onto α^h with degree n (this is illustrated in Figure 16 in a special case). Consequently, $\lambda_f(\alpha^h) = n \cdot (1/n) = 1$, which means α^h is an obstruction for f .

If $n_v > 0$, the flapped pillow $\widehat{\mathbb{P}}$ has at least one vertical flap. Then $n_{\widehat{\alpha}} > 0$ for at least one essential pullback $\widehat{\alpha}$ of α^h . Lemma 5.1 implies that

$$\lambda_f(\alpha^h) \leq (n - 1) \frac{1}{n} + \frac{1}{n + 1} < 1,$$

and so α^h is not an obstruction for f .

The proof for the vertical curve α^v is completely analogous. □

The above corollary can be read as follows: the obstruction α^h for the $(n \times n)$ -Lattès map can be eliminated by gluing a vertical flap to \mathbb{P} . Similarly, the obstruction α^v can be eliminated by gluing a horizontal flap. We will show momentarily that if both of these obstructions are eliminated (that is, if there are both horizontal and vertical flaps), then no other obstructions are present and so the map f is realized.

5.2. Ruling out other obstructions. Now we discuss what happens with the essential curves in (\mathbb{P}, P_f) that are not isotopic to the horizontal curve α^h or the vertical curve α^v .

THEOREM 5.3. *Let $f = \widehat{\mathcal{L}} \circ \phi^{-1} : \mathbb{P} \rightarrow \mathbb{P}$ be a Thurston map obtained from the $(n \times n)$ -Lattès map, $n \geq 2$, by gluing $n_h \geq 0$ horizontal and $n_v \geq 0$ vertical flaps to \mathbb{P} and assume that $n_h + n_v > 0$. If $\gamma \subset \mathbb{P} \setminus P_f$ is an essential Jordan curve that is not isotopic to either α^h or α^v , then γ is not an obstruction for f .*

Before we turn to the proof of this theorem, we first record how it implies Theorem 1.2 stated in the introduction.

Proof of Theorem 1.2. Let n, f, n_h , and n_v with $n_h + n_v > 0$ be as in the statement. We have seen in §4.2 that then f has a hyperbolic orbifold. If $n_h = 0$ or $n_v = 0$, then by Corollary 5.2, the curve α^v or the curve α^h is an obstruction, respectively.

If $n_h > 0$ and $n_v > 0$, then f has no obstruction as follows from Corollary 5.2 and Theorem 5.3. Since f has a hyperbolic orbifold, in this case, f is realized by a rational map according to Theorem 3.7. □

Corollary 5.2 and Theorem 5.3 also imply that if $n_h = 0$, then α^v is the only obstruction for f (up to isotopy relative to P_f). Similarly, α^h is the only obstruction if $n_v = 0$.

Before we go into the details, we will give an outline for the proof of Theorem 5.3. We argue by contradiction and assume that f has an obstruction given by an essential Jordan curve γ in (\mathbb{P}, P_f) that is isotopic to neither α^h nor α^v relative to P_f . Then $\lambda_f(\gamma) \geq 1$. Let $\gamma_1, \dots, \gamma_k$ for some $k \in \mathbb{N}$ be all the essential pullbacks of γ under f , which must be isotopic to γ relative to P_f .

Using facts about intersection numbers and the mapping properties of f , one can show that for the number of essential pullbacks of γ , we have $k \leq n$ and that the corresponding mapping degrees satisfy $\deg(f : \gamma_j \rightarrow \gamma) \geq n$ for all $j = 1, \dots, k$. Since $\lambda_f(\gamma) \geq 1$, it follows that there are exactly $k = n$ essential pullbacks and that $\deg(f : \gamma_j \rightarrow \gamma) = n$ for each $j = 1, \dots, n$.

This in turn implies that none of the essential pullbacks $\widehat{\gamma}_j := \phi^{-1}(\gamma_j)$ of γ under $\widehat{\mathcal{L}}$ goes over a flap in $\widehat{\mathbb{P}}$. Then all the n pullbacks $\widehat{\gamma}_1, \dots, \widehat{\gamma}_n$ belong to the base $B(\widehat{\mathbb{P}})$ of $\widehat{\mathbb{P}}$. This means that each $\widehat{\gamma}_j$ can be thought of as a pullback of γ under the original $(n \times n)$ -Lattès map \mathcal{L}_n . However, there are only n pullbacks of γ under \mathcal{L}_n , which cross all the edges of the graph $\mathcal{L}_n^{-1}(C)$, where C is the common boundary of the 0-tiles in \mathbb{P} . Consequently, the pullbacks $\widehat{\gamma}_1, \dots, \widehat{\gamma}_n$ cross all the 1-edges in the closure of the base $B(\widehat{\mathbb{P}})$. It follows that one of the pullbacks $\widehat{\gamma}_j$ must cross one of the base edges of a flap F in $\widehat{\mathbb{P}}$, which would necessarily mean that $\widehat{\gamma}_j$ goes over the flap F . This gives the desired contradiction and Theorem 5.3 follows.

In the remainder of this section, we will fill in the details for this outline. First, we establish several general facts about degrees and intersection numbers.

Let $n \in \mathbb{N}$ and $f : X \rightarrow Y$ be a map between two sets X and Y . We say that f is *at most n -to-1* if $\#f^{-1}(y) \leq n$ for each $y \in Y$. We say that f is *n -to-1* if $\#f^{-1}(y) = n$ for each $y \in Y$.

LEMMA 5.4. *Let $f : X \rightarrow Y$ be a map between two sets X and Y . Suppose $M \subset X$, $N \subset Y$, and $f|M : M \rightarrow f(M)$ is at most n -to-1 for some $n \in \mathbb{N}$. Then*

$$\#(M \cap f^{-1}(N)) \leq n \cdot \#(f(M) \cap N).$$

Proof. The map f sends each point in $M \cap f^{-1}(N)$ to a point in $f(M) \cap N$. Moreover, each point in $f(M) \cap N$ has at most n preimages in $M \cap f^{-1}(N)$ under f . The statement follows. \square

LEMMA 5.5. *Let $f: S^2 \rightarrow S^2$ be a Thurston map with $\#P_f = 4$, and γ be an essential Jordan curve in (S^2, P_f) . Suppose that $\tilde{\alpha}$ is an essential Jordan curve or an arc in (S^2, P_f) such that:*

- (i) $f(\tilde{\alpha})$ and $\tilde{\alpha}$ are isotopic relative to P_f ;
- (ii) the map $f|_{\tilde{\alpha}}: \tilde{\alpha} \rightarrow f(\tilde{\alpha})$ is at most n -to-1, where $n \in \mathbb{N}$;
- (iii) $i(f(\tilde{\alpha}), \gamma) = \#(f(\tilde{\alpha}) \cap \gamma) > 0$.

Then $k \leq n$, where $k \in \mathbb{N}_0$ denotes the number of distinct pullbacks of γ under f that are isotopic to γ relative to P_f . Moreover, if $\tilde{\alpha}$ meets a peripheral pullback of γ , then $k < n$.

In the formulation and the ensuing proof, intersection numbers are with respect to (S^2, P_f) . Note that since $f(\tilde{\alpha})$ and $\tilde{\alpha}$ are isotopic relative to P_f by assumption, $f(\tilde{\alpha})$ is of the same type as $\tilde{\alpha}$, that is, a Jordan curve or an arc in (S^2, P_f) .

Proof. Let $\gamma_1, \dots, \gamma_k$ be the distinct pullbacks of γ under f that are isotopic to γ relative to P_f . Since $f|_{\tilde{\alpha}}: \tilde{\alpha} \rightarrow f(\tilde{\alpha})$ is at most n -to-1, we can apply Lemma 5.4 and conclude that

$$\#(\tilde{\alpha} \cap f^{-1}(\gamma)) \leq n \cdot \#(f(\tilde{\alpha}) \cap \gamma) = n \cdot i(f(\tilde{\alpha}), \gamma).$$

However,

$$n \cdot i(f(\tilde{\alpha}), \gamma) \geq \#(\tilde{\alpha} \cap f^{-1}(\gamma)) \geq \sum_{j=1}^k \#(\tilde{\alpha} \cap \gamma_j) \geq \sum_{j=1}^k i(\tilde{\alpha}, \gamma_j) = k \cdot i(f(\tilde{\alpha}), \gamma). \tag{5.1}$$

Since $i(f(\tilde{\alpha}), \gamma) > 0$, we see that $k \leq n$. If $\tilde{\alpha}$ meets a peripheral pullback of γ , then the second inequality in (5.1) is strict and we actually have $k < n$. \square

The next result will lead to the strict inequality from Lemma 5.5 in the proof of Theorem 5.3.

LEMMA 5.6. *As before, let $\widehat{\mathcal{L}}: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ be the blown-up $(n \times n)$ -Lattès map, and suppose $\gamma = \tau_{r/s}$ is a simple closed geodesic in \mathbb{P} with slope $r/s \in \widehat{\mathbb{Q}} \setminus \{0, \infty\}$. Let $\widehat{\gamma}$ be a pullback of γ under $\widehat{\mathcal{L}}$. If $\widehat{\gamma}$ intersects the interior of a base edge of a flap F in $\widehat{\mathbb{P}}$, then $\widehat{\gamma}$ also intersects the top edge of F .*

Proof. We have $\gamma = \wp(\ell_{r/s})$, where $\ell_{r/s} \subset \mathbb{C} \setminus \mathbb{Z}^2$ is a straight line with slope $r/s \neq 0, \infty$. Then γ is an essential Jordan curve in (\mathbb{P}, V) . Let $\widehat{\gamma} \subset \widehat{\mathbb{P}}$ be as in the statement. As in §2.4, we denote by a, b, c, d the edges of the pillow \mathbb{P} . Let $e' \subset \widehat{\mathbb{P}}$ be a base edge of a flap F in $\widehat{\mathbb{P}}$ such that $\widehat{\gamma} \cap \text{int}(e') \neq \emptyset$. We will assume that F is a horizontal flap. Then $\widehat{\mathcal{L}}(e') = a$ or $\widehat{\mathcal{L}}(e') = c$. We will make the further assumption that $\widehat{\mathcal{L}}(e') = a$. The other cases, when $\widehat{\mathcal{L}}(e') = c$ or when F is a vertical flap, can be treated in a way that is completely analogous to the ensuing argument.




FIGURE 17. A pullback $\widehat{\gamma}$ going over a horizontal flap in $\widehat{\mathbb{P}}$.

Let $e'' \subset F$ be the base edge of F different from e' , and \tilde{e} be the top edge of F . Then $\widehat{\mathcal{L}}(e'') = a$ and $\widehat{\mathcal{L}}(\tilde{e}) = c$. Moreover,

$$F \cap \widehat{\mathcal{L}}^{-1}(a) = e' \cup e'' = \partial F \quad \text{and} \quad F \cap \widehat{\mathcal{L}}^{-1}(c) = \tilde{e}. \tag{5.2}$$

We have $\gamma = \wp(\ell_{r/s})$ with $r/s \neq 0$, and so by Lemma 2.5, the sets $a \cap \gamma$ and $c \cap \gamma$ are finite and non-empty and the points in these sets alternate on γ . Since $\widehat{\mathcal{L}}$ is a covering map from $\widehat{\gamma}$ onto γ , we conclude that the sets $\widehat{\gamma} \cap \widehat{\mathcal{L}}^{-1}(a)$ and $\widehat{\gamma} \cap \widehat{\mathcal{L}}^{-1}(c)$ are also finite and non-empty and the points in these sets alternate on $\widehat{\gamma}$.

We choose a point $p \in \widehat{\gamma} \cap \text{int}(e')$. Since $\gamma = \wp(\ell_{r/s})$, the curve γ has transverse intersections with $\text{int}(a)$. It follows that the pullback $\widehat{\gamma}$ has a transverse intersection with $\text{int}(e')$ at p , and so $\widehat{\gamma}$ crosses into the interior of the flap F at p . Therefore, if we travel along $\widehat{\gamma}$ starting at $p \in \widehat{\gamma} \cap \widehat{\mathcal{L}}^{-1}(a)$ and traverse into the interior of the flap F , we must meet $\widehat{\mathcal{L}}^{-1}(c)$ before we possibly exit F through its boundary $\partial F = e' \cup e'' = F \cap \widehat{\mathcal{L}}^{-1}(a)$ (see equation (5.2)). Now $F \cap \widehat{\mathcal{L}}^{-1}(c) = \tilde{e}$ and so this implies that the pullback $\widehat{\gamma}$ meets the top edge \tilde{e} (see Figure 17 for an illustration). The statement follows. \square

We are now ready to prove the main result of this section.

Proof of Theorem 5.3. Let $f: \mathbb{P} \rightarrow \mathbb{P}$ be a Thurston map as in the statement, obtained from the $(n \times n)$ -Lattès map \mathcal{L}_n , $n \geq 2$, by gluing $n_h \geq 0$ horizontal and $n_v \geq 0$ vertical flaps, where we assume $n_h + n_v > 0$. As described in the beginning of this section, then $f = \widehat{\mathcal{L}} \circ \phi^{-1}$, where $\widehat{\mathcal{L}}: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ is a branched covering map on the associated flapped pillow $\widehat{\mathbb{P}}$ and $\phi: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ is an identifying homeomorphism as discussed in §4.2. Note that $\widehat{\mathbb{P}}$ has at least one flap, since $n_h + n_v > 0$. Following the notation from §4.2, we denote by $\widehat{a}, \widehat{b}, \widehat{c}, \widehat{d}$ the arcs in $(\widehat{\mathbb{P}}, \widehat{V})$ that, under ϕ , correspond to the edges a, b, c, d of \mathbb{P} , respectively.

We now argue by contradiction and assume that there exists an essential Jordan curve γ in $(\mathbb{P}, P_f) = (\mathbb{P}, V)$ that is not isotopic to α^h or α^v relative to $P_f = V$, but is an obstruction for f , that is, $\lambda_f(\gamma) \geq 1$. Since we can replace γ with any curve in the same isotopy class, by Lemma 2.3, we may assume that $\gamma = \wp(\ell_{r/s})$ for a straight line $\ell_{r/s} \subset \mathbb{C} \setminus \mathbb{Z}^2$ with slope $r/s \in \mathbb{Q}$. Since γ is not isotopic to α^h or α^v relative to V , we have $r/s \neq 0, \infty$, and so $r, s \neq 0$. Then it follows from Lemma 2.4(iii) and (iv) that

$$\begin{aligned} \#(a \cap \gamma) &= i(a, \gamma) = |r| = i(c, \gamma) = \#(c \cap \gamma) > 0, \\ \#(b \cap \gamma) &= i(b, \gamma) = s = i(d, \gamma) = \#(d \cap \gamma) > 0. \end{aligned} \tag{5.3}$$

In particular, $\gamma \subset \mathbb{P} \setminus V$ intersects the interiors of all four edges of \mathbb{P} .

Let $\gamma_1, \dots, \gamma_k$ for some $k \in \mathbb{N}$ denote all the pullbacks of γ under f that are isotopic to γ relative to P_f . By construction of the blown-up map, the arc $\widehat{a} = \phi^{-1}(a)$ consists of n 1-edges in $\widehat{\mathbb{P}}$, each of which is homeomorphically mapped onto a by $\widehat{\mathcal{L}}$. This implies that the map $f|_a: a \rightarrow a$ is at most n -to-1. By (5.3), we can apply Lemma 5.5 to a and γ and conclude that $k \leq n$.

By Lemma 5.1, the horizontal curve α^h has n distinct pullbacks under f that are isotopic to α^h relative to P_f . Since

$$i(\gamma, \alpha^h) = \#(\gamma \cap \alpha^h) = 2|r| > 0 \tag{5.4}$$

by Lemma 2.4(i), we can apply Lemma 5.5 again, this time to γ_j and α^h (in the roles of $\widetilde{\alpha}$ and γ , respectively), and conclude that $n \leq \deg(f: \gamma_j \rightarrow \gamma)$ for all $j = 1, \dots, k$.

Then

$$1 \leq \lambda_f(\gamma) = \sum_{j=1}^k \frac{1}{\deg(f: \gamma_j \rightarrow \gamma)} \leq k/n \leq 1.$$

Therefore, $k = n$ and $\deg(f: \gamma_j \rightarrow \gamma) = n$ for each $j = 1, \dots, n$.

This shows that the curve γ has exactly n essential pullbacks under $\widehat{\mathcal{L}}$ given by $\widehat{\gamma}_1 := \phi^{-1}(\gamma_1), \dots, \widehat{\gamma}_n := \phi^{-1}(\gamma_n)$. Here, the isotopy classes are considered with respect to the vertex set \widehat{V} of $\widehat{\mathbb{P}}$. We will now use the second part of Lemma 5.5 to show that none of these pullbacks goes over a flap in $\widehat{\mathbb{P}}$.

Claim. For each flap F in $\widehat{\mathbb{P}}$ and each pullback $\widehat{\gamma}_j, j = 1, \dots, n$, we have $F \cap \widehat{\gamma}_j = \emptyset$.

To see that the claim is true, suppose some pullback $\widehat{\gamma}_j$ meets a flap F in $\widehat{\mathbb{P}}$. We may assume that F is a horizontal flap; the other case, when F is vertical, can be treated by similar considerations. Then F contains a peripheral pullback $\widehat{\alpha}^h$ of α^h under $\widehat{\mathcal{L}}$, which separates the union ∂F of the two base edges of F from the top edge of F . We will first show that $\widehat{\gamma}_j$ intersects $\widehat{\alpha}^h$.

Note that ∂F is a Jordan curve and that $\text{int}(F)$ does not contain any point from the vertex set \widehat{V} of $\widehat{\mathbb{P}}$. It follows that the curve $\widehat{\gamma}_j$ must intersect ∂F , because $\widehat{\gamma}_j$ is essential in $(\widehat{\mathbb{P}}, \widehat{V})$. Since the curve γ does not pass through $P_f = V$, its pullback $\widehat{\gamma}_j$ under $\widehat{\mathcal{L}}$ does not pass through any 1-vertex in $\widehat{\mathbb{P}}$. Consequently, $\widehat{\gamma}_j$ must meet the interior of one of the two base edges of F , which compose the boundary ∂F . Lemma 5.6 now implies that $\widehat{\gamma}_j$ also meets the top edge of F . Therefore, $\widehat{\gamma}_j$ meets the peripheral pullback $\widehat{\alpha}^h$ in F .

It follows that $\gamma_j = \phi(\widehat{\gamma}_j)$ meets the peripheral pullback $\phi(\widehat{\alpha}^h)$ of α^h under f . Lemma 5.5 now implies that $n < \deg(f: \gamma_j \rightarrow \gamma) = n$, which is a contradiction. This finishes the proof of the claim.

The claim implies that each essential pullback $\widehat{\gamma}_j, j = 1, \dots, n$, belongs to the base $B(\widehat{\mathbb{P}})$ of $\widehat{\mathbb{P}}$. By (4.2), we can identify $B(\widehat{\mathbb{P}})$ with the subset $\mathbb{P} \setminus \bigcup_{e \in E} \text{int}(e)$ of the original pillow \mathbb{P} , where E denotes the non-empty subset of all 1-edges of \mathbb{P} along which flaps were glued in the construction of $\widehat{\mathbb{P}}$.

Under this identification, the map $\widehat{\mathcal{L}}$ on $B(\widehat{\mathbb{P}})$ coincides with the $(n \times n)$ -Lattès map \mathcal{L}_n . Thus, we may view $\widehat{\gamma}_1, \dots, \widehat{\gamma}_n$ as pullbacks of γ under \mathcal{L}_n on the original pillow \mathbb{P} . Now γ has exactly n pullbacks under \mathcal{L}_n (see equation (3.4)). This implies that

$$\mathcal{L}_n^{-1}(\gamma) = \widehat{\gamma}_1 \cup \dots \cup \widehat{\gamma}_n.$$

Since γ meets the interior of every edge of \mathbb{P} , the set $\mathcal{L}_n^{-1}(\gamma) = \widehat{\gamma}_1 \cup \dots \cup \widehat{\gamma}_n$ meets the interior of every 1-edge of \mathbb{P} . This is impossible, because

$$\widehat{\gamma}_1 \cup \dots \cup \widehat{\gamma}_n \subset \mathcal{B}(\widehat{\mathbb{P}}) = \mathbb{P} \setminus \bigcup_{e \in E} \text{int}(e)$$

does not meet the interior of any 1-edge in $E \neq \emptyset$. This is a contradiction and the statement follows. □

6. Essential circuit length

To prove our main result, Theorem 1.1, we need some preparation, in particular, a refined version of Lemma 5.5. We will also address the question how blowing up arcs modifies the pullbacks of a curve α under natural restrictions on the blown-up arcs. First, we introduce some terminology and establish some auxiliary facts.

Let $U \subset S^2$ be an open and connected set, and $\sigma \subset S^2$ be an arc. We say that σ is an arc in U ending in ∂U if there exists an endpoint p of σ such that $\sigma \setminus \{p\} \subset U$ and $p \in \partial U$.

Let \mathcal{G} be a connected planar embedded graph in S^2 and U be one of its faces. Then U is simply connected, and so we can find a homeomorphism $\varphi: \mathbb{D} \rightarrow U$. Since we want some additional properties of φ here, it is easiest to equip S^2 with a complex structure and choose a conformal map $\varphi: \mathbb{D} \rightarrow U$.

Since ∂U is a finite union of edges of \mathcal{G} , this set is locally connected and so the conformal map φ extends to a surjective continuous map $\varphi: \text{cl}(\mathbb{D}) \rightarrow \text{cl}(U)$ [Pom92, Theorem 2.1]. This extension has the following property: if σ is an arc in U ending in ∂U , then $\varphi^{-1}(\sigma)$ is an arc in \mathbb{D} ending in $\partial\mathbb{D}$ (see [Pom92, Proposition 2.14]). For given \mathcal{G} and U , we fix, once and for all, such a map $\varphi = \varphi_{\mathcal{G},U}$ from $\text{cl}(\mathbb{D})$ onto $\text{cl}(U)$.

Let (e_1, e_2, \dots, e_n) be a circuit in \mathcal{G} that traces the boundary ∂U . Recall from §2.3 that the number n is called the circuit length of U in \mathcal{G} , and each edge $e \in \partial U$ appears exactly once or twice in the sequence e_1, e_2, \dots, e_n depending on whether the face U lies on one or both sides of e , respectively. Then there is a corresponding decomposition $\partial\mathbb{D} = \sigma_1 \cup \dots \cup \sigma_n$ of the unit circle $\partial\mathbb{D}$ into non-overlapping subarcs $\sigma_1, \dots, \sigma_n$ of $\partial\mathbb{D}$ such that $\varphi = \varphi_{\mathcal{G},U}$ is a homeomorphism of σ_m onto e_m for each $m = 1, \dots, n$.

Let $0 < \epsilon < 1$. We say that a Jordan curve $\beta \subset U$ is an ϵ -boundary of U with respect to \mathcal{G} if $\beta' := \varphi^{-1}(\beta) \subset A_\epsilon := \{z \in \mathbb{C} : 1 - \epsilon < |z| < 1\}$, and β' separates 0 from $\partial\mathbb{D}$. Then β' is a core curve of the annulus A_ϵ .

For the remainder of this section, $f: S^2 \rightarrow S^2$ is a Thurston map. All isotopies on S^2 are relative to P_f , and we consider intersection numbers in (S^2, P_f) .

Let e be an arc in (S^2, P_f) . Then we can naturally view the set $\mathcal{G} := f^{-1}(e)$ as a planar embedded graph with the vertex set $f^{-1}(\partial e)$ and the edges given by the lifts of e under f . Note that \mathcal{G} is bipartite.

LEMMA 6.1. *Let $f : S^2 \rightarrow S^2$ be a Thurston map, e be an arc and γ be a Jordan curve in (S^2, P_f) with $\#(e \cap \gamma) = i(e, \gamma)$. Suppose that \mathcal{H} is a connected subgraph of $\mathcal{G} = f^{-1}(e)$ and U is a face of \mathcal{H} . Let $2n$ with $n \in \mathbb{N}$ be the circuit length of U in \mathcal{H} . Then for each $0 < \epsilon < 1$, there exists an ϵ -boundary β of U with respect to \mathcal{H} such that $\#(\beta \cap f^{-1}(\gamma)) = 2n \cdot i(e, \gamma)$.*

Note that the circuit length of U in \mathcal{H} is even since \mathcal{H} is a bipartite graph.

Proof. Let $s := i(e, \gamma) \in \mathbb{N}_0$. Then e and γ have exactly s distinct points in common, say $p_1, \dots, p_s \in e \cap \gamma$. Each point p_j lies in $\text{int}(e)$, because $\partial e \subset P_f$ and $\gamma \subset S^2 \setminus P_f$. Since e and γ are in minimal position, they meet transversely (see Lemma 2.2), that is, if we travel along γ toward one of the points p_j (according to some orientation of γ), then near p_j , we stay on one side of e , but cross over to the other side of e if we pass p_j .

This implies that we can find disjoint subarcs $\sigma_1, \dots, \sigma_s$ of γ such that each arc σ_j contains p_j in its interior, but contains no other point in $e \cap \gamma$. Moreover, p_j splits σ_j into two non-overlapping subarcs σ_j^L and σ_j^R with the common endpoint p_j so that with some fixed orientation of e , the arc σ_j^L lies to the left and σ_j^R lies to the right of e . Note that if $\gamma' := \gamma \setminus (\sigma_1 \cup \dots \cup \sigma_s)$, then $e \cap \gamma' = \emptyset$.

For the given face U of \mathcal{H} , we fix a map $\varphi = \varphi_{\mathcal{H},U} : \text{cl}(\mathbb{D}) \rightarrow \text{cl}(U)$ as discussed in the beginning of this section. Let (e_1, \dots, e_{2n}) be a circuit in \mathcal{H} that traces the boundary ∂U . As we have already pointed out, the number of edges in the circuit is even, because \mathcal{H} is a bipartite graph. With suitable orientation of each arc e_m , the face U lies on the left of e_m . If an arc appears twice in the list e_1, \dots, e_{2n} , then it will carry opposite orientations in its two occurrences.

We want to investigate the set $f^{-1}(\gamma) \cap \text{cl}(U)$ near ∂U . Note that f maps each arc e_m homeomorphically onto e . This implies that f is a homeomorphism on a suitable Jordan region that contains e_m as a crosscut. It follows that we can pull back the local picture near points in $e \cap \gamma$ to a similar local picture for points in $e_m \cap f^{-1}(\gamma)$. So if we choose the arcs σ_j small enough, as we may assume, and pull them back by f , then it is clear that $f^{-1}(\gamma) \cap \text{cl}(U)$ can be represented in the form

$$f^{-1}(\gamma) \cap \text{cl}(U) = K \cup \bigcup_{m=1}^{2n} \bigcup_{j=1}^s \sigma_{m,j},$$

where K has positive distance to $\partial U = e_1 \cup \dots \cup e_{2n}$ (with respect to some base metric on S^2). Moreover, each $\sigma_{m,j}$ is an arc in U ending in $e_m \subset \partial U$ such that f is a homeomorphism from $\sigma_{m,j}$ onto σ_j^L or σ_j^R depending on whether $f|_{e_m} : e_m \rightarrow e$ is orientation-preserving or orientation-reversing. If we remove from each arc $\sigma_{m,j}$ its endpoint in e_m , then the half-open arcs obtained are all disjoint. Two arcs $\sigma_{m,j}$ and $\sigma_{m',j'}$ share an endpoint precisely when $j = j'$ and they arise from edges e_m and $e_{m'}$ with the same underlying set, but with opposite orientations. In this case, f sends one of these arcs to σ_j^L , and the other one to σ_j^R .

Since K has positive distance to ∂U , it is clear that if β is an ϵ -boundary of U with respect to \mathcal{H} for $\epsilon > 0$ small enough (as we may assume), then $\beta \cap K = \emptyset$. So to control $\#(\beta \cap f^{-1}(\gamma))$, we have to worry only about the intersections of β with the arcs $\sigma_{m,j}$.

Note that there are exactly $2n \cdot s = 2n \cdot i(e, \gamma)$ of these arcs. If we pull them back by the map φ , then we obtain pairwise disjoint arcs in \mathbb{D} ending in $\partial\mathbb{D}$. The statement now follows from the following fact, whose precise justification we leave to the reader: if $\alpha_1, \dots, \alpha_M$ with $M \in \mathbb{N}_0$ are pairwise disjoint arcs in \mathbb{D} ending in $\partial\mathbb{D}$, then for each $0 < \epsilon < 1$, there exists a core curve β' of the annulus $A_\epsilon = \{z \in \mathbb{C} : 1 - \epsilon < |z| < 1\}$ such that $\#(\beta' \cap (\alpha_1 \cup \dots \cup \alpha_M)) = M$. □

Now we are ready to provide a refined version of Lemma 5.5.

LEMMA 6.2. *Let $f : S^2 \rightarrow S^2$ be a Thurston map with $\#P_f = 4$, α and γ be essential Jordan curves in (S^2, P_f) , c be a core arc of α , and assume that $\#(c \cap \gamma) = i(c, \gamma) > 0$.*

Let \mathcal{H} be a connected subgraph of $\mathcal{G} := f^{-1}(c)$, and U be a face of \mathcal{H} such that for small enough $\epsilon > 0$, each ϵ -boundary β of U with respect to \mathcal{H} is isotopic to α relative to P_f . Let $2n$ with $n \in \mathbb{N}$ be the circuit length of U in \mathcal{H} .

Then $k \leq n$, where $k \in \mathbb{N}_0$ denotes the number of pullbacks of γ under f that are isotopic to γ relative to P_f . Moreover, if $\partial U \subset \mathcal{H}$ meets a peripheral pullback of γ under f , then $k < n$.

Proof. Let $\gamma_1, \dots, \gamma_k$ be the pullbacks of γ under f that are isotopic to γ relative to P_f . Then by Lemma 6.1, for sufficiently small $\epsilon > 0$, we can find an ϵ -boundary β of U with respect to \mathcal{H} such that $\beta \sim \alpha$ relative to P_f and

$$\#(\beta \cap f^{-1}(\gamma)) = 2n \cdot i(c, \gamma) = n \cdot i(\alpha, \gamma),$$

where the last equality follows from Lemma 2.6. Hence, we have

$$n \cdot i(\alpha, \gamma) = \#(\beta \cap f^{-1}(\gamma)) \geq \sum_{j=1}^k \#(\beta \cap \gamma_j) \geq \sum_{j=1}^k i(\beta, \gamma_j) = k \cdot i(\alpha, \gamma). \tag{6.1}$$

Since $i(\alpha, \gamma) = 2 \cdot i(c, \gamma) > 0$, we conclude that $k \leq n$, as desired.

To see the second statement, we have to revisit the proof of Lemma 6.1. There we identified $2n \cdot i(c, \gamma) = n \cdot i(\alpha, \gamma)$ distinct arcs σ in U ending in ∂U (they were called $\sigma_{m,j}$ in the proof). These arcs were subarcs of $f^{-1}(\gamma)$ and accounted for all possible intersections of β with $f^{-1}(\gamma)$ for sufficiently small $\epsilon > 0$; with a suitable choice of β , each of these arcs σ gave precisely one such intersection point. Now if a peripheral pullback $\tilde{\gamma} \subset f^{-1}(\gamma)$ of γ under f meets ∂U , then one of these arcs σ is a subarc of $\tilde{\gamma}$. It follows that the first inequality in (6.1) must be strict and so $k < n$. □

For the rest of this section, we fix a Thurston map $f : S^2 \rightarrow S^2$ with $\#P_f = 4$, an essential Jordan curve α in (S^2, P_f) , and core arcs a and c of α that lie in different components of $S^2 \setminus \alpha$. We can view the set $\mathcal{G} := f^{-1}(a \cup c)$ as a planar embedded graph in S^2 with the set of vertices $f^{-1}(P_f)$, and the edge set consisting of the lifts of a and c under f . Then \mathcal{G} is a bipartite graph.

Let U be the unique connected component of $S^2 \setminus (a \cup c)$. Then U is an annulus and α is its core curve. The connected components \tilde{U} of $f^{-1}(U)$ are precisely the complementary components of $\mathcal{G} = f^{-1}(a \cup c)$ in S^2 . It easily follows from the Riemann–Hurwitz




FIGURE 18. A Thurston map f . The sphere on the right shows a Jordan curve α in (S^2, P_f) and two core arcs a and c . The sphere on the left shows the pullbacks of α under f and the planar embedded graph $\mathcal{G} = f^{-1}(a \cup c)$.

formula (see equation (2.2)) that each \tilde{U} is an annulus, and that $f : \tilde{U} \rightarrow U$ is a covering map. Moreover, each such annulus contains precisely one pullback $\tilde{\alpha}$ of α under f .

This setup is illustrated in Figure 18. The points marked in black indicate the four postcritical points of f . The sphere on the right contains two core arcs a and c of a Jordan curve α in (S^2, P_f) . On the left, the lifts of a and c under f are shown in blue and magenta colors, respectively, and the pullbacks of α in green.

We call a connected component \tilde{U} of $f^{-1}(U) = S^2 \setminus \mathcal{G}$ *essential* or *peripheral*, depending on whether the unique pullback $\tilde{\alpha}$ of α contained in \tilde{U} is essential or peripheral in (S^2, P_f) , respectively. Each boundary $\partial\tilde{U}$ has exactly two connected components. One of them is mapped by f to a and the other one to c ; accordingly, we denote them by $\partial_a\tilde{U}$ and $\partial_c\tilde{U}$, respectively. Then we have

$$\partial\tilde{U} = \partial_a\tilde{U} \cup \partial_c\tilde{U}, \quad \partial_a\tilde{U} = f^{-1}(a) \cap \partial\tilde{U}, \quad \text{and} \quad \partial_c\tilde{U} = f^{-1}(c) \cap \partial\tilde{U}.$$

The sets $\partial_a\tilde{U}$ and $\partial_c\tilde{U}$ are subgraphs of \mathcal{G} . Since \tilde{U} is a connected subset of $S^2 \setminus \mathcal{G} \subset S^2 \setminus \partial_a\tilde{U}$, there exists a unique face V_a of $\partial_a\tilde{U}$ (considered as a subgraph of \mathcal{G}) such that $\tilde{U} \subset V_a$. Similarly, there exists a unique face V_c of $\partial_c\tilde{U}$ with $\tilde{U} \subset V_c$. By definition, the *circuit length* of $\partial_a\tilde{U}$ or of $\partial_c\tilde{U}$ is the circuit length of V_a in $\partial_a\tilde{U}$ or of V_c in $\partial_c\tilde{U}$, respectively.

Then the following statement is true.

LEMMA 6.3. *The circuit lengths of $\partial_a\tilde{U}$ and $\partial_c\tilde{U}$ are both equal to $2 \cdot \deg(f : \tilde{U} \rightarrow U)$.*

We call the identical circuit lengths of $\partial_a\tilde{U}$ and $\partial_c\tilde{U}$ the *circuit length of \tilde{U}* (for fixed f , α , a , and c).

Proof. It is clear that the subgraph $\partial_a\tilde{U}$ of \mathcal{G} is bipartite, and so $\partial_a\tilde{U}$ has even circuit length $2n$ with $n \in \mathbb{N}$. Let $d := \deg(f : \tilde{U} \rightarrow U)$. It is enough to show that $2n = 2d$, because the roles of $\partial_a\tilde{U}$ and $\partial_c\tilde{U}$ are symmetric, and so the same identity will then also be true for the circuit length of $\partial_c\tilde{U}$.

To see that $2n = 2d$, we use a similar idea as in the proof of Lemma 6.1. We choose a point $p \in \text{int}(a)$ and an arc $\sigma \subset S^2 \setminus c$ with $p \in \text{int}(\sigma)$ that meets a transversely in p , but has no other point with a in common. Then p splits σ into two non-overlapping subarcs

σ^L and σ^R with the common endpoint p so that with some fixed orientation of a , the arc σ^L lies to the left and σ^R lies to the right of a .

Let (e_1, \dots, e_{2n}) be a circuit in $\partial_a \tilde{U}$ that traces the boundary $\partial V_a = \partial_a \tilde{U}$, where V_a is the unique face of $\partial_a \tilde{U}$ with $\tilde{U} \subset V_a$. With suitable orientation of each arc e_m , the face \tilde{U} lies on the left of e_m . We now consider the set $f^{-1}(\sigma) \cap \text{cl}(\tilde{U})$ near $\partial_a \tilde{U}$. If we choose σ small enough (as we may), then as in the proof of Lemma 6.1, we see that

$$f^{-1}(\sigma) \cap \text{cl}(\tilde{U}) = \bigcup_{m=1}^{2n} \sigma_m, \tag{6.2}$$

where each σ_m is an arc in \tilde{U} ending in $e_m \subset \partial_a \tilde{U}$ such that f is a homeomorphism from σ_m to σ^L or σ^R depending on whether $f|_{e_m}: e_m \rightarrow a$ is orientation-preserving or orientation-reversing. If we remove from each arc σ_m its endpoint in e_m , then the half-open arcs obtained are all disjoint. However, since $f: \tilde{U} \rightarrow U$ is a d -to-1 covering map, there are precisely d distinct lifts of $\sigma^L \setminus \{p\}$ and d distinct lifts of $\sigma^R \setminus \{p\}$ under f contained in \tilde{U} . These must be precisely the half-open arcs obtained from σ_m , $m = 1, \dots, 2n$. It follows that $2n = 2d$, as desired. □

Suppose \tilde{U} is an essential component of $f^{-1}(U)$. We consider a circuit in $\partial \tilde{U} \subset \mathcal{G}$ and denote by $\mathcal{H} \subset \mathcal{G}$ the underlying graph of the circuit. In the following, we will often conflate the circuit with its underlying graph \mathcal{H} , where we think of \mathcal{H} as traversed as a circuit in some way. Since \tilde{U} is a connected set in $S^2 \setminus \mathcal{G} \subset S^2 \setminus \mathcal{H}$, there exists a unique face V of \mathcal{H} such that $\tilde{U} \subset V$. By definition, for $0 < \epsilon < 1$, an ϵ -boundary β of \tilde{U} with respect to \mathcal{H} is an ϵ -boundary of V with respect to \mathcal{H} . This is an abuse of terminology, because even for small $\epsilon > 0$, such an ϵ -boundary β may not lie in \tilde{U} , but it is convenient in the following. Note that for small enough $\epsilon > 0$, such ϵ -boundaries for fixed \mathcal{H} have the same isotopy type relative to P_f .

By definition, the *essential circuit length* of \tilde{U} is the minimal length of all circuits \mathcal{H} in $\partial \tilde{U}$ such that for all small enough $\epsilon > 0$, each ϵ -boundary of \tilde{U} with respect to \mathcal{H} is isotopic to a core curve of \tilde{U} relative to P_f . As we will see momentarily, if we run through $\partial_a \tilde{U}$ and $\partial_c \tilde{U}$ as circuits, then they have this property and so the essential circuit length of \tilde{U} is well defined. We call a circuit \mathcal{H} in $\partial \tilde{U}$ that realizes the essential circuit length an *essential circuit* for \tilde{U} .

LEMMA 6.4. *We have the inequality*

$$\text{circuit length of } \tilde{U} \geq \text{essential circuit length of } \tilde{U}.$$

For example, consider the annulus \tilde{U} containing the pullback $\tilde{\alpha}$ in Figure 18. Then the circuit length of \tilde{U} equals 6, while the essential circuit length of \tilde{U} equals 4.

Proof. Consider $\partial_a \tilde{U}$ as a circuit in $\partial \tilde{U}$. Let V_a be the unique face of $\partial_a \tilde{U}$ that contains \tilde{U} . Then, for sufficiently small $\epsilon > 0$, each ϵ -boundary β of V_a with respect to $\partial_a \tilde{U}$ necessarily separates $\partial_a \tilde{U}$ and $\partial_c \tilde{U}$, and is thus a core curve of \tilde{U} . The statement follows. □




FIGURE 19. The Thurston map f from Figure 18 and a set $E = \{e_1, e_2\}$ of arcs in $(S^2, f^{-1}(P_f))$ satisfying the α -restricted blow-up conditions.

Let \widehat{f} be a Thurston map obtained from f by blowing up some set of arcs E in $S^2 \setminus f^{-1}(P_f)$. Under certain natural assumptions on the arcs in E , we want to describe the components of $\widehat{f}^{-1}(U)$ and their properties in terms of the components of $f^{-1}(U)$. We first formulate suitable conditions that allow such a comparison.

Definition 6.5. (α -restricted blow-up conditions) Let $f: S^2 \rightarrow S^2$ be a Thurston map with $\#P_f = 4$, α be an essential Jordan curve in (S^2, P_f) , and a and c be core arcs of α that lie in different components of $S^2 \setminus \alpha$. Suppose $E \neq \emptyset$ is a finite set of arcs in $(S^2, f^{-1}(P_f))$ satisfying the blow-up conditions, that is, the interiors of the arcs in E are disjoint and $f: e \rightarrow f(e)$ is a homeomorphism for each $e \in E$.

We say that E satisfies the α -restricted blow-up conditions if

$$i(f(e), \alpha) = \#(f(e) \cap \alpha) = 1 \quad \text{and} \quad f(\text{int}(e)) \cap a = \emptyset = f(\text{int}(e)) \cap c \quad (6.3)$$

for each $e \in E$.

In other words, for each $e \in E$, the arc $f(e)$ is in minimal position with respect to α and intersects α only once, and $f(\text{int}(e)) = \text{int}(f(e))$ belongs to the annulus $U = S^2 \setminus (a \cup c)$. Note that the endpoints of $f(e)$ lie in $P_f \subset a \cup c = \partial U$; see Figure 19 for an illustration.

The condition in (6.3) is somewhat artificial, because it depends not only on α , but also on the choices of a and c . One can show that up to isotopy, it can be replaced by the more natural condition $i(f(e), \alpha) = 1$ for all $e \in E$. Since the justification of this claim involves some topological machinery that is beyond the scope of the paper, we prefer to work with (6.3).

Now the following statement is true.

LEMMA 6.6. *Let $f: S^2 \rightarrow S^2$ be a Thurston map with $\#P_f = 4$, α be an essential Jordan curve in (S^2, P_f) , a and c be core arcs of α that lie in different components of $S^2 \setminus \alpha$. Suppose a set E of arcs in $(S^2, f^{-1}(P_f))$ satisfies the α -restricted blow-up conditions and we are given multiplicities $m_e \in \mathbb{N}$ for $e \in E$.*

Then the Thurston map \widehat{f} obtained from f by blowing up each arc $e \in E$ with multiplicity m_e can be constructed so that it satisfies the following conditions.




FIGURE 20. A Thurston map \hat{f} obtained from the Thurston map f in Figures 18 and 19 by blowing up the arcs e_1 and e_2 with multiplicities $m_{e_1} = 2$ and $m_{e_2} = 1$, respectively.

- (i) $\mathcal{G} = f^{-1}(a \cup c)$ is a subgraph of $\hat{\mathcal{G}} := \hat{f}^{-1}(a \cup c)$.
- (ii) Each complementary annulus \hat{U} of $\hat{\mathcal{G}}$ is contained in a unique complementary annulus \tilde{U} of \mathcal{G} . Moreover, the assignment $\hat{U} \mapsto \tilde{U}$ is a bijection between the complementary annuli of $\hat{\mathcal{G}}$ and of \mathcal{G} .

Let \hat{U} and \tilde{U} be corresponding annuli as in condition (ii). Then the following statements are true.

- (iii) The core curves of \hat{U} and of \tilde{U} are isotopic relative to $P_f = P_{\hat{f}}$. In particular, \hat{U} is essential if and only if \tilde{U} is essential.
- (iv) If \hat{U} (and hence also \tilde{U}) is essential, then the essential circuit lengths of \hat{U} and \tilde{U} are the same. Moreover, if \mathcal{H} is an essential circuit for \hat{U} , then \mathcal{H} is an essential circuit for \tilde{U} . In particular, $\mathcal{H} \subset \partial\hat{U} \subset \partial\tilde{U}$.

In condition (i), it is understood that the planar embedded graph $\mathcal{G} = f^{-1}(a \cup c) \subset S^2$ has the vertex set $f^{-1}(P_f)$ and that $\hat{\mathcal{G}} = \hat{f}^{-1}(a \cup c)$ has the vertex set $\hat{f}^{-1}(P_{\hat{f}}) \supset f^{-1}(P_f)$.

To illustrate the lemma, we consider the Thurston map \hat{f} that is indicated in Figure 20 and obtained from the Thurston map f in Figure 18 by blowing up the arcs e_1 and e_2 in Figure 19 with multiplicities $m_{e_1} = 2$ and $m_{e_2} = 1$. Here, the lifts of a and c under the blown-up map \hat{f} are shown in blue and magenta colors on the left sphere, respectively. Comparing these figures, we immediately see that in this case, the statements of the lemma are true.

Proof. Let $e \in E$ be arbitrary. Since E satisfies the conditions in Definition 6.5, the set $f(e)$ is an arc in (S^2, P_f) , and so $f(e)$ has its endpoints in P_f . By (6.3), the arc $f(e)$ meets α precisely once and is in minimal position with respect to α . In particular, $f(e)$ meets α transversely by Lemma 2.2. This implies that the endpoints of $f(e)$ lie in different core arcs of α , and so one endpoint of $f(e)$ lies in a and the other one in c . It follows that e has one endpoint in $f^{-1}(a)$ and the other one in $f^{-1}(c)$.

The set $\text{int}(e)$ belongs to a unique annulus \tilde{U} obtained as a complementary component of $\mathcal{G} = f^{-1}(a \cup c)$. Then one endpoint of e is in $\partial_a\tilde{U} = f^{-1}(a) \cap \partial\tilde{U}$ and the other one in $\partial_c\tilde{U} = f^{-1}(c) \cap \partial\tilde{U}$. In the blow-up construction described in §4.1, we can choose

the open Jordan region W_e so that $W_e \subset \tilde{U}$ for each $e \in E$ (of course, here the annulus \tilde{U} depends on e). Now we make choices of the subsequent ingredients in the blow-up construction as discussed in §4.1. That is, for each fixed arc $e \in E$, we choose a closed Jordan region D_e inside W_e . It is subdivided into $m = m_e$ components D_e^1, \dots, D_e^m . In addition, we also choose a pseudo-isotopy $h: S^2 \times \mathbb{I} \rightarrow S^2$ satisfying conditions (B1)–(B4), as well as maps $\varphi_k: D_e^k \rightarrow S^2$, $k = 1, \dots, m = m_e$, satisfying conditions (C1) and (C2). Let \hat{f} be the Thurston map obtained by blowing up each arc $e \in E$ with multiplicity m_e according to these choices. We claim that \hat{f} satisfies all the conditions in the statement.

It immediately follows from condition (B3) and the definition of \hat{f} that $\mathcal{G} = f^{-1}(a \cup c)$ is a subgraph of $\hat{\mathcal{G}} := \hat{f}^{-1}(a \cup c)$, and so statement (i) is true.

We have $\hat{\mathcal{G}} \setminus \mathcal{G} \subset \bigcup_{e \in E} D_e$. Condition (C1) now implies that for each $e \in E$ and each $k = 1, \dots, m = m_e$, the set $\hat{\mathcal{G}} \cap D_e^k$ consists of two disjoint edges, one of which is homeomorphically mapped onto a and the other one onto c by \hat{f} . We will call these edges *a-sticks* and *c-sticks*, respectively. Each closed Jordan region D_e contains exactly m_e *a-sticks*, which have a common endpoint in $\partial e \cap f^{-1}(a)$, and exactly m_e *c-sticks* with a common endpoint in $\partial e \cap f^{-1}(c)$. The edge set of $\hat{\mathcal{G}}$ consists of all the edges of \mathcal{G} together with all the *a-sticks* and *c-sticks*.

Each complementary component \hat{U} of $\hat{\mathcal{G}}$ is equal to a unique complementary component \tilde{U} of \mathcal{G} with all the *a-* and *c-sticks* removed that are contained in $\text{cl}(\tilde{U})$. Statement (ii) follows. Furthermore, since $P_f \cap \tilde{U} = \emptyset$, statement (iii) follows as well.

To prove statement (iv), let \tilde{U} and \hat{U} be corresponding essential annuli as in statement (ii). Viewing $\partial\tilde{U}$ as a subgraph of \mathcal{G} and $\partial\hat{U}$ as a subgraph of $\hat{\mathcal{G}}$, we see that $\partial\tilde{U}$ is a subgraph of $\partial\hat{U}$. The additional edges of $\partial\hat{U}$ are exactly the *a-* and *c-sticks* contained in $\text{cl}(\tilde{U})$. It follows from the definition that the essential circuit length of \hat{U} is greater than or equal to the essential circuit length of \tilde{U} .

Let \mathcal{H} be an essential circuit for \hat{U} and suppose \mathcal{H} contains an *a-* or *c-stick* σ . Then one of the endpoints of σ has degree 1 in $\partial\hat{U}$, and so σ must appear in two consecutive positions in the circuit \mathcal{H} . Omitting these two occurrences of σ from \mathcal{H} , we get a shorter circuit \mathcal{H}' in $\partial\hat{U}$ such that for all small enough $\epsilon > 0$, each ϵ -boundary of \hat{U} with respect to \mathcal{H}' is isotopic to a core curve of \hat{U} relative to $P_{\hat{f}} = P_f$. This contradicts the choice of \mathcal{H} , and it follows that \mathcal{H} does not contain any *a-* or *c-sticks*. Consequently, $\mathcal{H} \subset \partial\tilde{U} \subset \partial\hat{U}$, and the definition of the essential circuit length together with statement (iii) imply that \mathcal{H} is an essential circuit for \tilde{U} . Statement (iv) follows. □

7. Eliminating obstructions by blowing up arcs

The goal of this section is to show that the blow-up surgery can be applied to an obstructed Thurston map f with four postcritical points in such a way that the resulting map \hat{f} is realized by a rational map. The precise formulation is given in Theorem 1.1 (see also Remark 7.2). We will prove this statement by contradiction. For this, we assume that \hat{f} has an obstruction, and will carefully analyze some related mapping degrees. This leads to a very tight situation, where in some inequalities, we actually have equality. From this, we want to conclude that f has a parabolic orbifold, in contradiction to our hypotheses in Theorem 1.1. We first formulate a related criterion for parabolicity.

LEMMA 7.1. *Let $f : S^2 \rightarrow S^2$ be a Thurston map with $\#P_f = 4$. Suppose there exists a Jordan curve α in (S^2, P_f) such that the following conditions are true:*

- (i) α is an obstruction for f ;
- (ii) α has no peripheral pullbacks under f ;
- (iii) if we choose core arcs a and c of α in different components of $S^2 \setminus \alpha$ and consider the graph $\mathcal{G} = f^{-1}(a \cup c)$, then \mathcal{G} has precisely $n \in \mathbb{N}$ essential complementary components U_1, \dots, U_n with core curves isotopic to α relative to P_f . Moreover, we assume the essential circuit length of U_j is equal to $2n$ for each $j = 1, \dots, n$.

Then f has a parabolic orbifold.

Proof. Each U_j contains precisely one pullback α_j of α under f . The curves $\alpha_1, \dots, \alpha_n$ are all the pullbacks of α under f . Then it follows from our assumptions that

$$\begin{aligned} 2 \deg(f : \alpha_j \rightarrow \alpha) &= 2 \deg(f : U_j \rightarrow U) \\ &= \text{circuit length of } U_j \quad (\text{by Lemma 6.3}) \\ &\geq \text{essential circuit length of } U_j \quad (\text{by Lemma 6.4}) \\ &= 2n, \end{aligned}$$

and so $\deg(f : \alpha_j \rightarrow \alpha) \geq n$ for $j = 1, \dots, n$. However, α is an obstruction for f , and so

$$1 \leq \lambda_f(\alpha) = \sum_{j=1}^n \frac{1}{\deg(f : \alpha_j \rightarrow \alpha)} \leq n/n = 1.$$

It follows that we have equality in all previous inequalities. In particular,

$$\text{circuit length of } U_j = \text{essential circuit length of } U_j = 2n$$

for $j = 1, \dots, n$.

We want to apply the second part of Lemma 3.3, that is, we want to show that $f^{-1}(P_f) \subset C_f \cup P_f$. To see this, let $v \in f^{-1}(P_f)$ be arbitrary. Then v is a vertex of $\mathcal{G} = f^{-1}(a \cup c)$. If v is incident with two or more edges in \mathcal{G} , then $v \in C_f$.

Otherwise, v is the endpoint of precisely one edge e in \mathcal{G} , and so $\deg_{\mathcal{G}}(v) = 1$. We claim that then $v \in P_f$; to see this, we argue by contradiction and assume that $v \notin P_f$. Since α has no peripheral pullbacks, we have

$$e \subset \mathcal{G} = \bigcup_{j=1}^n \partial U_j,$$

and so $e \subset \partial U_j$ for some U_j . Then e is contained in a circuit (e_1, \dots, e_{2n}) of length $2n$ that traces one of the components of ∂U_j . Since $\deg_{\mathcal{G}}(v) = 1$, the circuit must traverse e twice with opposite orientations, that is, the edge e appears precisely in two consecutive entries in the cycle (e_1, \dots, e_{2n}) . Erasing these two occurrences from the cycle, we obtain a new circuit in $\partial U_j \subset \mathcal{G}$ with length $2n - 2$. Let \mathcal{H} denote the underlying subgraph of \mathcal{G} corresponding to this shortened circuit and let U be the face of \mathcal{H} that contains U_j . Since the endpoint v of e does not belong to P_f , for every small enough $\epsilon > 0$, the ϵ -boundary of \mathcal{H} with respect to U is isotopic to the curve α_j relative to P_f . Then the essential circuit

length of U_j is $\leq 2n - 2$, contradicting the fact that $2n$ is the essential circuit length of U_j by our hypotheses. So we must have $v \in P_f$.

It follows that $f^{-1}(P_f) \subset C_f \cup P_f$, and so f has a parabolic orbifold by Lemma 3.3. □

We are now ready to prove our main result.

Proof of Theorem 1.1. As in the statement, suppose $f: S^2 \rightarrow S^2$ is a Thurston map with $\#P_f = 4$ and a hyperbolic orbifold. We assume that f has an obstruction given by a Jordan curve α in (S^2, P_f) . We choose core arcs a and c for α that lie in different components of $S^2 \setminus \alpha$, and assume that $E \neq \emptyset$ is a finite set of arcs in $(S^2, f^{-1}(P_f))$ that satisfies the α -restricted blow-up conditions as in Definition 6.5.

We assume that we obtained a Thurston map $\widehat{f}: S^2 \rightarrow S^2$ by blowing up arcs in E (with some multiplicities) so that $\lambda_{\widehat{f}}(\alpha) < 1$. Then $P_{\widehat{f}} = P_f$ and \widehat{f} has a hyperbolic orbifold (see Lemma 4.3 and Remark 4.4). Up to replacing \widehat{f} with a Thurston equivalent map, we may also assume that the statements in Lemma 6.6 are true for the map \widehat{f} . We now argue by contradiction and assume that \widehat{f} is not realized by a rational map. Then by Theorem 3.7, the map \widehat{f} has an obstruction given by a Jordan curve γ in (S^2, P_f) .

We set $U = S^2 \setminus (a \cup c)$. Since E satisfies the α -restricted blow-up conditions, we have $\#(f(e) \cap \alpha) = 1$ and $f(\text{int}(e)) \cap a = f(\text{int}(e)) \cap c = \emptyset$ for each $e \in E$. In other words, $f(e)$ intersects α only once and $\text{int}(f(e))$ belongs to U . Then each arc in E intersects only one pullback of α and only once.

Since γ is an obstruction for \widehat{f} , but α is not, the curves α and γ are not isotopic relative to $P_{\widehat{f}} = P_f$. So we have $i(\alpha, \gamma) > 0$ for intersection numbers in (S^2, P_f) as follows from Lemma 2.1. By Lemma 2.6(i), we have

$$i(a, \gamma) = i(c, \gamma) = \frac{1}{2}i(\alpha, \gamma) > 0. \tag{7.1}$$

As follows from Lemma 2.6(ii), by replacing γ with an isotopic curve relative to $P_{\widehat{f}} = P_f$ if necessary, we may also assume that

$$\#(\alpha \cap \gamma) = i(\alpha, \gamma), \#(a \cap \gamma) = i(a, \gamma), \#(c \cap \gamma) = i(c, \gamma) \tag{7.2}$$

and that the points in the non-empty and finite sets $a \cap \gamma$ and $c \cap \gamma$ alternate on γ .

We denote by $\alpha_1, \dots, \alpha_n$ with $n \in \mathbb{N}$ the pullbacks of α under f that are isotopic to α relative to P_f . Now we consider the graphs $\mathcal{G} = f^{-1}(a \cup c)$ and $\widehat{\mathcal{G}} = \widehat{f}^{-1}(a \cup c)$ as in §6. By Lemma 6.6, \mathcal{G} is a subgraph of $\widehat{\mathcal{G}}$. Moreover, the following facts are true for their complementary components. Each α_j is a core curve in an essential annulus U_j that is a component of $S^2 \setminus \mathcal{G}$. Each U_j contains precisely one component \widehat{U}_j of $S^2 \setminus \widehat{\mathcal{G}}$. This component is essential and contains precisely one essential pullback $\widehat{\alpha}_j$ of α under \widehat{f} . The essential circuit length of U_j is the same as the essential circuit length of \widehat{U}_j . The curves $\widehat{\alpha}_1, \dots, \widehat{\alpha}_n$ are precisely all the distinct essential pullbacks of α under \widehat{f} . They are isotopic to α relative to P_f .

Let $\gamma_1, \dots, \gamma_k$ with $k \in \mathbb{N}$ be the pullbacks of γ under \widehat{f} that are isotopic to γ relative to $P_{\widehat{f}} = P_f$. Applying Lemmas 6.3 and 6.2 (for the latter, equations (7.1) and (7.2) are important) to an essential circuit for \widehat{U}_j , we see that

$$\begin{aligned}
 \deg(f: \alpha_j \rightarrow \alpha) &= \deg(f: U_j \rightarrow U) \\
 &= \frac{1}{2} \cdot \text{circuit length of } U_j \\
 &\geq \frac{1}{2} \cdot \text{essential circuit length of } U_j \\
 &= \frac{1}{2} \cdot \text{essential circuit length of } \widehat{U}_j \\
 &\geq k
 \end{aligned} \tag{7.3}$$

for $j = 1, \dots, n$.

However, α has n distinct essential pullbacks $\widehat{\alpha}_1, \dots, \widehat{\alpha}_n$ under \widehat{f} , and so Lemma 5.5 implies that $\deg(\widehat{f}: \gamma_m \rightarrow \gamma) \geq n$ for $m = 1, \dots, k$ (again equations (7.1) and (7.2) are used here). Since α and γ are obstructions for f and \widehat{f} , respectively, we conclude that

$$1 \leq \lambda_f(\alpha) = \sum_{j=1}^n \frac{1}{\deg(f: \alpha_j \rightarrow \alpha)} \leq n/k$$

and

$$1 \leq \lambda_{\widehat{f}}(\gamma) = \sum_{m=1}^k \frac{1}{\deg(\widehat{f}: \gamma_m \rightarrow \gamma)} \leq k/n.$$

It follows that $k = n$, which forces $\deg(f: \alpha_j \rightarrow \alpha) = \deg(\widehat{f}: \gamma_j \rightarrow \gamma) = n$ for $j = 1, \dots, n$. If we combine this with (7.3), then we also see that

$$\begin{aligned}
 \text{circuit length of } U_j &= \text{essential circuit length of } U_j \\
 &= \text{essential circuit length of } \widehat{U}_j = 2n
 \end{aligned} \tag{7.4}$$

for $j = 1, \dots, n$.

We now want to apply Lemma 7.1 to our map f and the obstruction α . To verify the hypotheses of Lemma 7.1, it remains to show that α has no peripheral pullbacks under f , or equivalently, no peripheral pullbacks under \widehat{f} (see Lemma 6.6(iii)).

We argue by contradiction and assume that α has some peripheral pullbacks under \widehat{f} . Then there exists at least one peripheral annulus in the complement of $\widehat{\mathcal{G}}$. Such an annulus is disjoint from each annulus \widehat{U}_j . We can then travel from a point p of such a peripheral annulus to a point in the set $\widehat{U}_1 \cup \dots \cup \widehat{U}_n$ along an arc σ in $S^2 \setminus \widehat{f}^{-1}(P_{\widehat{f}})$ that crosses each edge in the graph $\widehat{\mathcal{G}}$ transversely. Then there is a first point q on σ where we enter the closure M of $\widehat{U}_1 \cup \dots \cup \widehat{U}_n$. The point q is necessarily an interior point of an edge e of $\widehat{\mathcal{G}}$ contained in the boundary $\partial \widehat{U}_j$ for some $j \in \{1, \dots, n\}$. Interior points of the subarc of σ between p and q that are close to q do not lie in $M \cup \widehat{\mathcal{G}}$. Hence, such points must belong to a peripheral component \widehat{U} of $\widehat{\mathcal{G}}$. Then necessarily, $e \subset \partial \widehat{U}$.

In other words, there exists an edge e in the graph $\widehat{\mathcal{G}}$ that belongs to the boundary of an essential annulus \widehat{U}_j and a peripheral annulus \widehat{U} . Clearly, $\widehat{f}(e) = a$ or $\widehat{f}(e) = c$. In the following, we will assume that $\widehat{f}(e) = c$, that is, $e \subset \partial_c \widehat{U}_j \cap \partial_c \widehat{U}$; the other case, $\widehat{f}(e) = a$, is completely analogous.

Since $i(c, \gamma) = \#(c \cap \gamma) > 0$, there exists a pullback $\widehat{\gamma}$ of γ under \widehat{f} that meets e transversely. Consequently, this pullback $\widehat{\gamma}$ meets both \widehat{U}_j and \widehat{U} . Since $\widehat{f}: \widehat{\gamma} \rightarrow \gamma$ is

a covering map and the points in $a \cap \gamma \neq \emptyset$ and $c \cap \gamma \neq \emptyset$ alternate on γ , the points in $\widehat{f}^{-1}(a) \cap \widehat{\gamma} \neq \emptyset$ and $\widehat{f}^{-1}(c) \cap \widehat{\gamma} \neq \emptyset$ alternate on $\widehat{\gamma}$. This implies that the curve $\widehat{\gamma}$ also meets the sets $\partial_a \widehat{U}_j$ and $\partial_c \widehat{U}$, and hence both components of the boundary of \widehat{U}_j and of \widehat{U} . We conclude that $\widehat{\gamma}$ meets the core curve $\widehat{\alpha}_j$ of \widehat{U}_j and the core curve $\widehat{\alpha}$ of \widehat{U} . Note that $\widehat{\alpha}$ is a peripheral pullback of α under \widehat{f} . To show that this is impossible, we consider two cases.

Case 1: $\widehat{\gamma}$ is an essential pullback of γ under \widehat{f} , say $\widehat{\gamma} = \gamma_m$ for some $m \in \{1, \dots, n\}$. Then Lemma 5.5 (for the map \widehat{f} , and $\alpha, \widehat{\gamma}$ in the roles of $\gamma, \tilde{\alpha}$, respectively) shows $n < \deg(\widehat{f}: \widehat{\gamma} \rightarrow \gamma)$, because α has n essential pullbacks and $\widehat{\gamma}$ meets a peripheral pullback of α . However, we know that $\deg(\widehat{f}: \widehat{\gamma} \rightarrow \gamma) = \deg(\widehat{f}: \gamma_m \rightarrow \gamma) = n$. This is a contradiction.

Case 2: $\widehat{\gamma}$ is a peripheral pullback of γ under \widehat{f} . Let $\mathcal{H} := \partial_c U_j$. Then it follows from Lemma 6.6 that $\mathcal{H} \subset \partial_c \widehat{U}_j$. Moreover, equation (7.4) implies that \mathcal{H} (considered as a circuit) realizes the essential circuit lengths of U_j and \widehat{U}_j , which are both equal to $2n$.

Now $e \subset \partial_c U_j = \mathcal{H}$, and so \mathcal{H} meets the peripheral pullback $\widehat{\gamma}$ of γ under \widehat{f} . The second part of Lemma 6.2 applied to \mathcal{H} implies that the number k of essential pullbacks of γ under \widehat{f} is less than n , contradicting $k = n$.

To summarize, these contradictions show that α has no peripheral pullbacks under \widehat{f} , and hence no peripheral pullbacks under f by Lemma 6.6(iii). So we can apply Lemma 7.1 and conclude that f has a parabolic orbifold. This is yet another contradiction, because f has a hyperbolic orbifold by our hypotheses. This shows that our initial assumption that \widehat{f} has an obstruction is false. Hence, \widehat{f} is realized by a rational map. This completes the proof of Theorem 1.1. □

Remark 7.2. Let $f: S^2 \rightarrow S^2$ be a Thurston map with $\#P_f = 4$ and a hyperbolic orbifold, and suppose f has an obstruction represented by a Jordan curve α in (S^2, P_f) . Then there always exist a set of arcs $E \neq \emptyset$ in $(S^2, f^{-1}(P_f))$ satisfying the α -restricted blow-up conditions and multiplicities $m_e, e \in E$, such that the corresponding blown-up map \widehat{f} satisfies $\lambda_{\widehat{f}}(\alpha) < 1$. We are then exactly in the setup of Theorem 1.1.

To see this, we first fix some core arcs a and c of α lying in different components of $S^2 \setminus \alpha$. We now choose an arc e_0 in (S^2, P_f) with $i(e_0, \alpha) = \#(e_0 \cap \alpha) = 1$ and $\text{int}(e_0) \subset S^2 \setminus (a \cup c)$. Let E be the (non-empty) set of all lifts of e_0 under f . Then E is a set of arcs in $(S^2, f^{-1}(P_f))$ and it is clear that E satisfies the α -restricted blow-up conditions. If $\tilde{\alpha}$ is any pullback of α under f , then there exists at least one arc $e \in E$ that meets $\tilde{\alpha}$ (necessarily in an interior point of e). Blowing up the arc e with some multiplicity $m_e \in \mathbb{N}$ increases the mapping degree for the corresponding pullback $\widehat{\alpha}$ under \widehat{f} by m_e and does not change the isotopy class of this pullback, that is, $[\widehat{\alpha}] = [\tilde{\alpha}]$ relative to $P_{\widehat{f}} = P_f$ (this easily follows from Lemma 6.6 and its proof). Note that each pullback $\widehat{\alpha}$ of α under \widehat{f} corresponds to a pullback $\tilde{\alpha}$ of α under f (this is essentially Lemma 6.6(ii)). It follows that if we choose the multiplicities $m_e, e \in E$, large enough, then for the Thurston map \widehat{f} , we will have $\lambda_{\widehat{f}}(\alpha) < 1$ and so α is not an obstruction for \widehat{f} . By Theorem 1.1, the map \widehat{f} is actually realized by a rational map. So by a suitable blow-up operation, an obstructed Thurston map f (with $P_f = 4$ and a hyperbolic orbifold) can be turned into a Thurston map \widehat{f} that is realized.

8. Global curve attractors

In this section, we will prove Theorem 1.4. We consider the pillow \mathbb{P} with its vertex set $V = \{A, B, C, D\}$. For the remainder of this section, $f: \mathbb{P} \rightarrow \mathbb{P}$ is a Thurston map obtained from the (2×2) -Lattès map by gluing $n_h \geq 1$ horizontal and $n_v \geq 1$ vertical flaps to \mathbb{P} . Then f is Thurston equivalent to a rational map by Theorem 1.2. In the following, all isotopies on \mathbb{P} are considered relative to $P_f = V$.

To prove Theorem 1.4, we want to show that Jordan curves in (\mathbb{P}, V) are getting ‘less twisted’ under taking preimages under f . To formalize this, we define the *complexity* $\|x\|$ of $x \in \widehat{\mathbb{Q}} \cup \{\odot\}$ as $\|x\| := 0$ for $x = \odot$ and $\|x\| := |r| + s$ for $x = r/s \in \widehat{\mathbb{Q}}$. Recall that \odot represents the isotopy classes of all peripheral curves, and that for a slope $r/s \in \widehat{\mathbb{Q}}$, we use the convention that the numbers $r \in \mathbb{Z}$ and $s \in \mathbb{N}_0$ are relatively prime and that $r = 1$ if $s = 0$. Note that $\|x\| = 0$ for $x \in \widehat{\mathbb{Q}} \cup \{\odot\}$ if and only if $x = \odot$.

The complexity admits a natural interpretation in terms of intersection numbers. To see this, recall that α^h and α^v (see (2.4)) represent simple closed geodesics in (\mathbb{P}, V) that separate the two horizontal and the two vertical edges of \mathbb{P} , respectively. Suppose the slope $r/s \in \widehat{\mathbb{Q}}$ corresponds to the isotopy class $[\gamma]$ of a (necessarily essential) Jordan curve γ in (\mathbb{P}, V) . Then by Lemma 2.4(v),

$$\|r/s\| = |r| + s = \frac{1}{2}i(\gamma, \alpha^h) + \frac{1}{2}i(\gamma, \alpha^v).$$

Moreover, if γ is peripheral, then $i(\gamma, \alpha^h) + i(\gamma, \alpha^v) = 0$, which agrees with the fact that $\|\odot\| = 0$.

As we will see, under the slope map μ_f (as defined in §1.2), complexities do not increase, and actually strictly decrease unless the slope belongs to a certain finite set. More precisely, we will show the following statement.

PROPOSITION 8.1. *Let $f: \mathbb{P} \rightarrow \mathbb{P}$ be a Thurston map obtained from the (2×2) -Lattès map by gluing $n_h \geq 1$ horizontal and $n_v \geq 1$ vertical flaps to the pillow \mathbb{P} . Then the following statements are true:*

- (i) $\|\mu_f(x)\| \leq \|x\|$ for all $x \in \widehat{\mathbb{Q}} \cup \{\odot\}$;
- (ii) $\|\mu_f(x)\| < \|x\|$ for all $x \in \widehat{\mathbb{Q}} \cup \{\odot\}$ with $\|x\| > 8$.

Since the set $\{x \in \widehat{\mathbb{Q}} \cup \{\odot\} : \|x\| \leq 8\}$ is finite, we actually have the strict inequality in statement (i) with at most finitely many exceptions. The proof of the proposition will show that $\|\mu_f(x)\| = \|x\|$ if and only if $\mu_f(x) = x$ (see Remark 8.6). As we will see below, Theorem 1.4 easily follows from Proposition 8.1.

Before we proceed with the proof of this proposition, we will establish several auxiliary results. As in §2.4, a, b, c, d are the edges of the pillow \mathbb{P} , and $\wp: \mathbb{C} \rightarrow \mathbb{P}$ denotes the Weierstrass function that is doubly periodic with respect to the lattice $2\mathbb{Z}^2$.

We are interested in simple closed geodesics and geodesic arcs τ in (\mathbb{P}, V) . Recall that every such geodesic has the form $\tau = \wp(\ell_{r/s})$ for a line $\ell_{r/s} \subset \mathbb{C}$ with slope $r/s \in \widehat{\mathbb{Q}}$. If $\ell_{r/s} \subset \mathbb{C} \setminus \mathbb{Z}^2$, then $\tau = \wp(\ell_{r/s})$ is a simple closed geodesic in (\mathbb{P}, V) , that is, $\tau \subset \mathbb{P} \setminus V$. If $\ell_{r/s}$ contains a point in \mathbb{Z}^2 , then $\tau = \wp(\ell_{r/s})$ is a geodesic arc in (\mathbb{P}, V) , that is, its interior lies in $\mathbb{P} \setminus V$ and its endpoints are in V .

LEMMA 8.2. *Let τ be a simple closed geodesic or a geodesic arc in (\mathbb{P}, V) with slope $r/s \in \widehat{\mathbb{Q}}$. We consider the 1-edges of \mathbb{P} with respect to the $(n \times n)$ -Lattès map \mathcal{L}_n , $n \geq 2$, that is, the lifts of the edges a, b, c, d of \mathbb{P} under \mathcal{L}_n . Then the following statements are true.*

- (i) *If $|r| > 2n$, then τ intersects the interior of every horizontal 1-edge of \mathbb{P} .*
- (ii) *If $s > 2n$, then τ intersects the interior of every vertical 1-edge of \mathbb{P} .*

Proof. We will only show the first part of the statement. The proof of the second part is completely analogous.

Let τ be a simple closed geodesic or a geodesic arc in (\mathbb{P}, V) with slope $r/s \in \widehat{\mathbb{Q}}$, where $|r| > 2n$. Suppose that e is a horizontal 1-edge and \tilde{e} is a lift of e under \wp . Then the arc \tilde{e} is a line segment of length $1/n$ contained in a line $\ell_0 \subset \mathbb{C}$ parallel to the real axis. To show that τ meets $\text{int}(e)$, it suffices to represent the given geodesic τ in the form $\tau = \wp(\ell_{r/s})$ for a line $\ell_{r/s} \subset \mathbb{C}$ with $\ell_{r/s} \cap \text{int}(\tilde{e}) \neq \emptyset$.

For this, we choose $p, q \in \mathbb{Z}$ such that $pr + qs = 1$ and define $\omega := 2(s + ir)$ and $\tilde{\omega} := 2(-p + iq)$. The numbers ω and $\tilde{\omega}$ form a basis of the period lattice $2\mathbb{Z}^2$ of \wp . In particular, if $\tau = \wp(\ell_{r/s}(z_0))$ for some $z_0 \in \mathbb{C}$, then $\tau = \wp(\ell_{r/s}(z_0 + j\tilde{\omega}))$ for all $j \in \mathbb{Z}$. The lines $\ell_{r/s}(z_0 + j\tilde{\omega})$, $j \in \mathbb{Z}$, are parallel and equally spaced. Actually, two consecutive lines in this family differ by a translation by $\tilde{\omega}$. Since $r \neq 0$, these lines are not parallel to the real axis and so they will cut out subsegments of equal length on the line ℓ_0 that contains \tilde{e} . To determine the length of these segments, we write $\tilde{\omega}$ in the form

$$\tilde{\omega} = u + v\omega \tag{8.1}$$

with $u, v \in \mathbb{R}$. It is easy to see that equation (8.1) implies that $u = -2/r$ (multiply equation (8.1) by the complex conjugate of ω and take imaginary parts), and so the lines in our family cut ℓ_0 into subsegments of length $|u| = 2/|r|$. Since $|u| = 2/|r| < 1/n$ by our hypotheses, one of these lines meets $\text{int}(\tilde{e})$. This implies that $\tau \cap \text{int}(e) \neq \emptyset$, as desired. □

We now want to see what happens to a geodesic arc ξ in (\mathbb{P}, V) if we take preimages under a map f as in Proposition 8.1. Unless ξ has slope in a finite exceptional set, suitable sets \mathcal{H} in the preimage $f^{-1}(\xi)$ will meet the interior of a flap glued to the pillow \mathbb{P} , and consequently a peripheral pullback of the horizontal curve $\alpha^h \subset \mathbb{P}$ or of the vertical curve $\alpha^v \subset \mathbb{P}$. We will formulate some relevant statements in a slightly more general situation. We first introduce some terminology.

Suppose $Z \subset S^2$ consists of four distinct points. We say that $K \subset S^2$ essentially separates Z if we can split Z into two disjoint subsets Z_1 and Z_2 consisting of two points each such that K separates Z_1 and Z_2 . Note that K trivially has this property if $K \cap Z$ consists of two or more points.

Now let $n \in \mathbb{N}$, $n \geq 2$, and consider the $(n \times n)$ -Lattès map $\mathcal{L}_n: \mathbb{P} \rightarrow \mathbb{P}$. If ξ is a geodesic arc in (\mathbb{P}, V) , then the preimage $\mathcal{L}_n^{-1}(\xi)$ is a disjoint union of simple closed geodesics and geodesic arcs in (\mathbb{P}, V) . Note that each connected component of $\mathcal{L}_n^{-1}(\xi)$ essentially separates V , but no proper subset of such a component does. It follows that if $K \subset \mathcal{L}_n^{-1}(\xi)$ is a connected set, then it essentially separates V if and only if K is a simple closed geodesic or a geodesic arc in (\mathbb{P}, V) .

Let $\widehat{\mathcal{L}}: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ be a branched covering map obtained from the $(n \times n)$ -Lattès map by gluing flaps to \mathbb{P} . As in §4.2, we denote by \widehat{V} the vertex set and by $B(\widehat{\mathbb{P}})$ the base of the flapped pillow $\widehat{\mathbb{P}}$. By construction, $\widehat{\mathcal{L}}$ maps each 1-tile of $\widehat{\mathbb{P}}$ by a Euclidean similarity (with scaling factor n) onto a 0-tile of \mathbb{P} . We also recall that we can naturally view the base $B(\widehat{\mathbb{P}})$ as a subset of \mathbb{P} (see (4.2)) and, with such an identification, the map $\widehat{\mathcal{L}}$ coincides with \mathcal{L}_n on $B(\widehat{\mathbb{P}})$.

Suppose that a geodesic arc ξ in (\mathbb{P}, V) joins two distinct points $X, Y \in V$. We consider $\widehat{\mathcal{G}} := \widehat{\mathcal{L}}^{-1}(\xi)$ as a planar embedded graph in $\widehat{\mathbb{P}}$ with the set of vertices $\widehat{\mathcal{L}}^{-1}(\{X, Y\})$ and the edges given by the lifts of ξ under $\widehat{\mathcal{L}}$.

LEMMA 8.3. *Let $\widehat{\mathcal{L}}: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ be a branched covering map obtained from the $(n \times n)$ -Lattès map with $n \geq 2$ by gluing $n_h \geq 1$ horizontal and $n_v \geq 1$ vertical flaps to \mathbb{P} . Suppose ξ is a geodesic arc in (\mathbb{P}, V) with slope $r/s \in \widehat{\mathbb{Q}} \setminus \{0, \infty\}$ and $\widehat{\xi}$ is a lift of ξ under $\widehat{\mathcal{L}}$.*

Let F be a flap in $\widehat{\mathbb{P}}$ with the base edges e' and e'' . If

$$\widehat{\xi} \cap (\text{int}(F) \cup \text{int}(e') \cup \text{int}(e'')) \neq \emptyset,$$

then $\widehat{\xi}$ meets a base edge and the top edge of the flap F .

Proof. The proof is similar to the proof of Lemma 5.6. Recall that a, b, c, d denote the edges of the pillow \mathbb{P} . Suppose $\xi \subset \mathbb{P}$, $\widehat{\xi} \subset \widehat{\mathbb{P}}$, and $e', e'' \subset F$ are as in the statement of the lemma. Let $\tilde{e} \subset F$ be the top edge of F .

Without loss of generality, we will assume that F is a horizontal flap. Then $\widehat{\mathcal{L}}(e') = a$ or $\widehat{\mathcal{L}}(e') = c$. We will make the further assumption that $\widehat{\mathcal{L}}(e') = a$. The other cases, when $\widehat{\mathcal{L}}(e') = c$ or when F is a vertical flap, can be treated in a way that is completely analogous to the ensuing argument. Then $\widehat{\mathcal{L}}(e'') = a$ and $\widehat{\mathcal{L}}(\tilde{e}) = c$. Moreover,

$$\widehat{\mathcal{L}}^{-1}(a \cup c) \cap F = e' \cup e'' \cup \tilde{e}. \tag{8.2}$$

Since ξ is a geodesic arc in (\mathbb{P}, V) with slope $r/s \neq 0$, by Lemma 2.5, the sets $a \cap \xi$ and $c \cap \xi$ are non-empty and finite, and the points in these sets alternate on ξ . We claim that there is a point $p \in \widehat{\xi} \cap \text{int}(F)$. By our hypotheses, this can only fail if $\widehat{\xi}$ meets either $\text{int}(e')$ or $\text{int}(e'')$ in a point q . Since the arc ξ has a transverse intersection with $\text{int}(a)$ at $\widehat{\mathcal{L}}(q)$, the arc $\widehat{\xi}$ has a transverse intersection with $\text{int}(e')$ or $\text{int}(e'')$ at q . Then $\widehat{\xi}$ meets $\text{int}(F)$ in a point p in any case.

Since the points in $a \cap \xi \neq \emptyset$ and $c \cap \xi \neq \emptyset$ alternate on ξ , the points in $\widehat{\xi} \cap \widehat{\mathcal{L}}^{-1}(a) \neq \emptyset$ and $\widehat{\xi} \cap \widehat{\mathcal{L}}^{-1}(c) \neq \emptyset$ alternate on $\widehat{\xi}$. Note that $F \cap \widehat{\mathcal{L}}^{-1}(c) = \tilde{e}$ and $F \cap \widehat{\mathcal{L}}^{-1}(a) = \partial F = e' \cup e''$. So, if we trace the arc $\widehat{\xi}$ starting from p in two different directions, we must meet a base edge of F in one direction and the top edge of F in the other direction. The statement follows. □

Now the following fact is true.

LEMMA 8.4. *Let $\widehat{\mathcal{L}}: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ be a branched covering map obtained from the $(n \times n)$ -Lattès map with $n \geq 2$ by gluing $n_h \geq 1$ horizontal and $n_v \geq 1$ vertical flaps to \mathbb{P} . Suppose that ξ is a geodesic arc in (\mathbb{P}, V) with slope $r/s \in \widehat{\mathbb{Q}}$ and \mathcal{H} is any connected subgraph of $\widehat{\mathcal{G}} = \widehat{\mathcal{L}}^{-1}(\xi)$ that essentially separates $\widehat{V} \subset \widehat{\mathbb{P}}$. If $|r| + s > 4n$, then \mathcal{H} meets a base edge and the top edge of a flap in $\widehat{\mathbb{P}}$.*

Proof. Suppose ξ is a geodesic arc in (\mathbb{P}, V) with slope $r/s \in \widehat{\mathbb{Q}}$, where $|r| + s > 4n$, and \mathcal{H} is a connected subgraph of $\widehat{\mathcal{G}} = \widehat{\mathcal{L}}^{-1}(\xi)$ that essentially separates the vertex set \widehat{V} of $\widehat{\mathbb{P}}$. Note that then $r/s \neq 0, \infty$, which will allow us to apply Lemma 8.3. Each edge of the graph $\widehat{\mathcal{G}} = \widehat{\mathcal{L}}^{-1}(\xi)$ is a lift $\widehat{\xi}$ of ξ as in this lemma.

We now argue by contradiction and suppose that there is no flap F in $\widehat{\mathbb{P}}$ such that \mathcal{H} meets both a base edge and the top edge of F . By the definition of $B(\widehat{\mathbb{P}})$ (see (4.1)) and Lemma 8.3, this means that each edge of \mathcal{H} , and thus the graph \mathcal{H} itself, is contained in $B(\widehat{\mathbb{P}})$. Consequently, we can consider \mathcal{H} as a connected subset of $\mathbb{P} \supset B(\widehat{\mathbb{P}})$. On $B(\widehat{\mathbb{P}})$, the maps $\widehat{\mathcal{L}}$ and \mathcal{L}_n are identical. Therefore, we can also regard \mathcal{H} as a connected subset of $\mathcal{L}_n^{-1}(\xi)$.

The set \mathcal{H} , now considered as a subset of \mathbb{P} , essentially separates $V \subset \mathbb{P}$. To see this, let $\widehat{V}_1, \widehat{V}_2 \subset \widehat{V} \subset \widehat{\mathbb{P}}$ be two pairs of vertices separated by \mathcal{H} in $\widehat{\mathbb{P}}$. We can identify \widehat{V}_1 with a pair V_1 and \widehat{V}_2 with a pair V_2 of vertices of \mathbb{P} . We claim that V_1 and V_2 are separated by \mathcal{H} in \mathbb{P} . Indeed, if this was not the case, then we could find a path β in \mathbb{P} that joins V_1 and V_2 without meeting \mathcal{H} . This path can be modified as follows to a path $\widehat{\beta}$ in $\widehat{\mathbb{P}}$ that joins \widehat{V}_1 and \widehat{V}_2 and does not meet $\mathcal{H} \subset \widehat{\mathbb{P}}$: if β meets some 1-edge e of \mathbb{P} to which one or several flaps are glued, then on β , there is a first point $p \in e$ and a last point $q \in e$. We now replace the part of β between p and q by a path that joins points corresponding to p and q in $\widehat{\mathbb{P}}$, travels on these flaps, and does not meet \mathcal{H} . If we make such replacements for all these edges e consecutively, then we obtain a path $\widehat{\beta}$ that joins \widehat{V}_1 and \widehat{V}_2 , but is disjoint from \mathcal{H} . However, such a path $\widehat{\beta}$ cannot exist, because \mathcal{H} separates \widehat{V}_1 and \widehat{V}_2 in $\widehat{\mathbb{P}}$.

We see that $\mathcal{H} \subset \mathcal{L}_n^{-1}(\xi)$ indeed essentially separates V . Since \mathcal{H} is connected, the discussion above (after the definition of essential separation) implies that \mathcal{H} is a simple closed geodesic or a geodesic arc in (\mathbb{P}, V) with slope r/s . Since $|r| + s > 4n$, either $|r| > 2n$ or $s > 2n$. Thus, by Lemma 8.2, the geodesic \mathcal{H} meets each horizontal 1-edge of \mathbb{P} in the first case or each vertical 1-edge of \mathbb{P} in the second case. Since $n_h \geq 1$ and $n_v \geq 1$, in either case, \mathcal{H} must meet the interior of a 1-edge along which a flap is glued and hence cannot be a subset of $B(\widehat{\mathbb{P}})$. This is a contradiction and the lemma follows. □

Remark 8.5. Suppose we are in the setup of Lemma 8.4. Then the connected set \mathcal{H} meets a base edge and the top edge of a flap, say a horizontal flap F . Then there exists a peripheral pullback $\widehat{\alpha}$ of the horizontal curve α^h under the map $\widehat{\mathcal{L}}$ that is contained in F . Let e' and e'' be the base edges of F , and \widetilde{e} be the top edge of F . Then the curve $\widehat{\alpha}$ separates $\partial F = e' \cup e''$ from $\widetilde{e} \subset F$. Since \mathcal{H} is connected and meets both \widetilde{e} and $e' \cup e''$, we conclude that $\mathcal{H} \cap \widehat{\alpha} \neq \emptyset$. If β is a connected set that traces \mathcal{H} closely, then it will also have points close to \widetilde{e} and close to $e' \cup e''$. Again this will imply that $\beta \cap \widehat{\alpha} \neq \emptyset$. This remark will become important in the proof of Proposition 8.1.

A completely analogous statement to Lemma 8.4 is true (with a very similar proof) if we assume that ξ is a simple closed geodesic in (\mathbb{P}, V) and \mathcal{H} is an essential pullback of ξ under $\widehat{\mathcal{L}}$.

We now turn to the proof of Proposition 8.1.

Proof of Proposition 8.1. Let $f: \mathbb{P} \rightarrow \mathbb{P}$ be a Thurston map obtained from the (2×2) -Lattès map by gluing $n_h \geq 1$ horizontal and $n_v \geq 1$ vertical flaps to \mathbb{P} . Then $P_f = V$, where $V = \{A, B, C, D\}$ is the set of vertices of \mathbb{P} , and A is the unique point in $P_f = V$ that is fixed by f .

To prove the first statement, let $x \in \widehat{\mathbb{Q}} \cup \{\odot\}$ be arbitrary. If $x = \odot$, then $\mu_f(\odot) = \odot$ and $\|\mu_f(\odot)\| = \|\odot\| = 0$. So in the following, we will assume that $x = r/s \in \widehat{\mathbb{Q}}$. Let $\gamma \subset \mathbb{P} \setminus V$ be a simple closed geodesic with slope $r/s \in \widehat{\mathbb{Q}}$. Then γ is an essential Jordan curve and so each of the two complementary components of γ in \mathbb{P} contains precisely two postcritical points of f . Let ξ and ξ' be core arcs of γ belonging to different components of $\mathbb{P} \setminus \gamma$. Here we may assume that ξ and ξ' are geodesic arcs in (\mathbb{P}, V) with slope r/s .

As before, we denote by α^h and α^v simple closed geodesics in (\mathbb{P}, V) that separate the two horizontal and the two vertical edges of \mathbb{P} , respectively. Then, by Lemma 2.4, we have

$$\begin{aligned} i(\gamma, \alpha^h) &= 2i(\xi, \alpha^h) = \#(\gamma \cap \alpha^h) = 2|r|, \\ i(\gamma, \alpha^v) &= 2i(\xi, \alpha^v) = \#(\gamma \cap \alpha^v) = 2s, \\ \|x\| &= |r| + s = \frac{1}{2}i(\gamma, \alpha^h) + \frac{1}{2}i(\gamma, \alpha^v). \end{aligned}$$

We call a point $p \in \mathbb{P}$ a 1-vertex if $f(p) \in P_f = \{A, B, C, D\}$. We say that a 1-vertex is of type A, B, C , or D if it is a preimage of A, B, C , or D under f , respectively.

Without loss of generality, we may assume that the core arc ξ connects the point A with a point $X \in \{B, C, D\}$. Then ξ' joins the two points in $\{B, C, D\} \setminus \{X\}$. Let $\mathcal{G} = f^{-1}(\xi \cup \xi')$, which we view as a planar embedded graph with the set of vertices $f^{-1}(V)$. Note that the degree of a vertex p in \mathcal{G} is equal to the local degree of the map f at p . In addition, the graph \mathcal{G} has the following properties.

- (P1) \mathcal{G} is a bipartite graph. In particular, 1-vertices of type A are connected only to 1-vertices of type X and *vice versa*.
- (P2) Each postcritical point of f is a 1-vertex of type A . If a 1-vertex of type A has degree ≥ 2 in \mathcal{G} , then it must be a postcritical point.

The analog of property (P1) is valid for arbitrary Thurston maps with four postcritical points. To see that property (P2) is true, note that the (2×2) -Lattès map sends each of the four vertices of \mathbb{P} to A . This remains true if we glue any number of flaps to \mathbb{P} . Moreover, gluing additional flaps can only create additional preimages of A of degree 1 in \mathcal{G} .

If every pullback of γ under f is peripheral, then $\mu_f(x) = \odot$, and so

$$\|\mu_f(x)\| = \|\odot\| = 0 < |r| + s = \|x\|. \tag{8.3}$$

Suppose γ has an essential pullback $\tilde{\gamma}$ under f . Then $\mu_f(x) \in \widehat{\mathbb{Q}}$ is the slope corresponding to the isotopy class of $\tilde{\gamma}$. By the discussion in §6, the pullback $\tilde{\gamma}$ belongs to a unique component \tilde{U} of $\mathbb{P} \setminus \mathcal{G}$. We use the notation $\partial_\xi \tilde{U} := f^{-1}(\xi) \cap \partial \tilde{U}$. Then $\partial_\xi \tilde{U}$ is a subgraph of \mathcal{G} that only contains 1-vertices of type A and X . Since $\tilde{\gamma}$ is essential, $\partial_\xi \tilde{U}$ satisfies:

- (P3) $\#(\partial_\xi \tilde{U} \cap P_f) \leq 2$.




FIGURE 21. Different combinatorial types of the graph $\partial_\xi \tilde{U}$. The subgraph in magenta corresponds to the appropriate choice of \mathcal{H} in each case. The vertices in black indicate the postcritical points.

Our goal now is to simplify the pullback $\tilde{\gamma}$ using an isotopy depending on the combinatorics of $\partial_\xi \tilde{U}$. More precisely, we will construct a curve β that is isotopic to $\tilde{\gamma}$, but has fewer intersections with α^h and α^v . To obtain a suitable curve β , we now distinguish several cases that exhaust all possibilities.

Case 1: $\partial_\xi \tilde{U}$ does not contain any simple cycle. Then $\partial_\xi \tilde{U}$ is a tree and, since $\tilde{\gamma}$ is essential, there are exactly two postcritical points in $\partial_\xi \tilde{U}$. These are 1-vertices of type A by property (P2). Let $\mathcal{H} \subset \partial_\xi \tilde{U}$ be the unique simple path that joins these two postcritical points in $\partial_\xi \tilde{U}$. By property (P1), the path \mathcal{H} must have length ≥ 2 , because the endpoints of \mathcal{H} have type A and the vertices of types A and X alternate on \mathcal{H} .

If the length of \mathcal{H} was ≥ 3 , then \mathcal{H} would contain at least one additional point p of type A apart from its endpoints. Then $\deg_{\mathcal{H}}(p) = 2$, so $\deg_{\mathcal{G}}(p) \geq 2$, which means p must be a postcritical point by property (P2). However, then $\mathcal{H} \subset \partial_\xi \tilde{U}$ contains at least three postcritical points, which contradicts property (P3). We conclude that \mathcal{H} has length 2; see Figure 21 (Case 1).

Let $\hat{U} = S^2 \setminus \mathcal{H}$. Then the annulus between $\tilde{\gamma}$ and \mathcal{H} contains no postcritical points of f , and hence for sufficiently small ϵ , each ϵ -boundary β of \hat{U} with respect to \mathcal{H} is isotopic to $\tilde{\gamma}$, as follows from Lemma 2.1.

Case 2: $\partial_\xi \tilde{U}$ contains a simple cycle. Then by property (P1), one of the vertices of such a cycle must be of type A. Since this vertex has degree equal to 2 in the cycle, and hence degree ≥ 2 in \mathcal{G} , it must be a postcritical point by property (P2). It follows that $\#(\partial_\xi \tilde{U} \cap P_f) \geq 1$. So by property (P3), either $\#(\partial_\xi \tilde{U} \cap P_f) = 1$ or $\#(\partial_\xi \tilde{U} \cap P_f) = 2$.

Case 2a: $\#(\partial_\xi \tilde{U} \cap P_f) = 1$. Since $\tilde{\gamma}$ is essential, there are exactly two postcritical points in the component of $S^2 \setminus \tilde{\gamma}$ that contains $\partial_\xi \tilde{U}$. One of them belongs to $\partial_\xi \tilde{U}$, while the other

one belongs to a face of $\partial_\xi \tilde{U}$ disjoint from \tilde{U} . This postcritical point then necessarily belongs to a face of a simple cycle \mathcal{H} in $\partial_\xi \tilde{U}$. This simple cycle \mathcal{H} then necessarily contains the unique postcritical point in $\partial_\xi \tilde{U}$ as we have seen above. Moreover, \mathcal{H} must have length 2, because otherwise, \mathcal{H} has an even length ≥ 4 by property (P1). However, then \mathcal{H} contains another 1-vertex of type A with degree ≥ 2 , which is necessarily a postcritical point by property (P2). Then $\mathcal{H} \subset \partial_\xi \tilde{U}$ contains at least two postcritical points, which contradicts our assumption for this case. So \mathcal{H} has indeed length 2; see Figure 21 (Case 2a).

Let \hat{U} denote the face of \mathcal{H} that contains \tilde{U} . Then again, the annulus between $\tilde{\gamma}$ and \mathcal{H} contains no postcritical points of f , and hence each ϵ -boundary β of \hat{U} with respect to \mathcal{H} is isotopic to $\tilde{\gamma}$ for sufficiently small ϵ .

Case 2b: $\#(\partial_\xi \tilde{U} \cap P_f) = 2$. Let \mathcal{H} be a simple path in $\partial_\xi \tilde{U}$ that joins the two postcritical points in $\partial_\xi \tilde{U}$. By the same reasoning as in Case 1, \mathcal{H} has length 2; see Figure 21 (Case 2b). Let $\hat{U} = S^2 \setminus \mathcal{H}$. Since $\tilde{\gamma}$ is essential, there are no postcritical points in the annulus between $\tilde{\gamma}$ and \mathcal{H} . Thus, each ϵ -boundary β of \hat{U} with respect to \mathcal{H} is isotopic to $\tilde{\gamma}$ for sufficiently small ϵ .

Note that in all cases, \mathcal{H} essentially separates $V = P_f$, because in all cases, \mathcal{H} separates the pairs of points in V contained in different complementary components of $\tilde{\gamma}$. Moreover, by our choice, the circuit length of \hat{U} is equal to 4 in Cases 1 and 2b, and equal to 2 in Case 2a. So in each case, it is ≤ 4 . Since ξ and α^h are in minimal position, as follows from Lemma 2.4, we can apply Lemma 6.1 to the face \hat{U} of \mathcal{H} . Hence, for each sufficiently small $\epsilon > 0$, we can always find an ϵ -boundary β of \hat{U} with respect to \mathcal{H} that is isotopic to $\tilde{\gamma}$ and satisfies $\#(\beta \cap f^{-1}(\alpha^h)) \leq 4i(\xi, \alpha^h)$.

Let $\tilde{\alpha}_1$ and $\tilde{\alpha}_2$ be the two pullbacks of α^h under f that are isotopic to α^h (there are exactly two such pullbacks by Lemma 5.1). Then in all cases, we have

$$\begin{aligned} 2i(\tilde{\gamma}, \alpha^h) &= 2i(\beta, \alpha^h) = i(\beta, \tilde{\alpha}_1) + i(\beta, \tilde{\alpha}_2) \\ &\leq \#(\beta \cap \tilde{\alpha}_1) + \#(\beta \cap \tilde{\alpha}_2) \\ &\leq \#(\beta \cap f^{-1}(\alpha^h)) \\ &\leq 4i(\xi, \alpha^h) \\ &= 2i(\gamma, \alpha^h). \end{aligned} \tag{8.4}$$

Thus, $i(\tilde{\gamma}, \alpha^h) \leq i(\gamma, \alpha^h)$. The same reasoning (with a possibly different choice of β) also shows $i(\tilde{\gamma}, \alpha^v) \leq i(\gamma, \alpha^v)$. Combining these inequalities, we conclude

$$\|\mu_f(x)\| = \frac{1}{2}i(\tilde{\gamma}, \alpha^h) + \frac{1}{2}i(\tilde{\gamma}, \alpha^v) \leq \frac{1}{2}i(\gamma, \alpha^h) + \frac{1}{2}i(\gamma, \alpha^v) = \|x\|. \tag{8.5}$$

This completes the proof of the first part of the statement.

Note that the second inequality in (8.4) is strict if β intersects a peripheral pullback of α^h . A similar statement is also true for the analogous inequality for the curve α^v . We now assume that $x = r/s \in \hat{\mathbb{Q}}$ satisfies $\|x\| > 8$. We will argue that then either

the inequality in (8.4) or the analogous inequality for α^v is strict. This will lead to $\|\mu_f(x)\| < \|x\|$.

To see this, first note that $f = \widehat{\mathcal{L}} \circ \phi^{-1}$, where $\widehat{\mathcal{L}}: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ is the associated branched covering map obtained by blowing up the (2×2) -Lattès map and $\phi: \widehat{\mathbb{P}} \rightarrow \mathbb{P}$ is a suitable homeomorphism (see §4.2 for the details). We proceed as in the first part of the proof and again represent x by a simple closed geodesic γ in (\mathbb{P}, V) with slope $x = r/s$. We may assume that γ has an essential pullback $\tilde{\gamma}$ under f , because otherwise, we have the desired strict inequality by (8.3).

We choose a geodesic core arc ξ in (\mathbb{P}, V) and a connected set $\mathcal{H} \subset f^{-1}(\xi)$ as before and define $\widehat{\mathcal{H}} := \phi^{-1}(\mathcal{H})$. Then $\widehat{\mathcal{H}}$ is a connected subset of $\widehat{\mathcal{L}}^{-1}(\xi)$ that essentially separates \widehat{V} , where $\widehat{V} = \phi^{-1}(V)$ is the set of vertices of the flapped pillow $\widehat{\mathbb{P}}$. Since $\|x\| = |r| + s > 8$, we can apply Lemma 8.4 (with $n = 2$) and conclude that the set $\widehat{\mathcal{H}}$ will meet a base edge and the top edge of some flap F in $\widehat{\mathbb{P}}$. We will assume that F is a horizontal flap, the case of a vertical flap being completely analogous.

If ϵ is small enough, then the ϵ -boundary β constructed above traces \mathcal{H} very closely in the sense that for each point in \mathcal{H} , there is a nearby point in β . The same is true for $\widehat{\beta} := \phi^{-1}(\beta)$ and $\widehat{\mathcal{H}}$. Using Remark 8.5, this implies that if ϵ is sufficiently small (as we may assume), then $\widehat{\beta}$ will meet the peripheral pullback $\widehat{\alpha}$ of α^h under $\widehat{\mathcal{L}}$ that is contained in the horizontal flap F . Consequently, β meets the peripheral pullback $\phi(\widehat{\alpha})$ of α^h under f . As we already pointed out, this leads to a strict inequality in (8.4) and thus also in (8.5). The statement follows. □

The proof of Theorem 1.4 is now easy.

Proof of Theorem 1.4. Let $f: \mathbb{P} \rightarrow \mathbb{P}$ be a Thurston map as in the statement. Then Proposition 8.1 implies that if $x \in \widehat{\mathbb{Q}} \cup \{\circ\}$ is arbitrary, then the complexities of the elements $x, \mu_f(x), \mu_f^2(x), \dots$ of the orbit of x under iteration of μ_f strictly decrease until this orbit eventually reaches the finite set $S := \{u \in \widehat{\mathbb{Q}} \cup \{\circ\} : \|u\| \leq 8\}$. From this point on, the orbit of x stays in S . The statement follows. □

Remark 8.6. The proofs of Theorems 1.3 and 1.4 show that a global curve attractor $\mathcal{A}(f)$ for f can be obtained from Jordan curves corresponding to slopes in the finite set $S = \{x \in \widehat{\mathbb{Q}} \cup \{\circ\} : \|x\| \leq 8\}$. Actually, (8.4) and (8.5) imply that $\|\mu_f(x)\| = \|x\|$ if and only if $\mu_f(x) = x$. Therefore, the *minimal* global curve attractor $\mathcal{A}(f)$ corresponds to the set $\{x \in \widehat{\mathbb{Q}} \cup \{\circ\} : \mu_f(x) = x\} \subset S$. In other words, the minimal $\mathcal{A}(f)$ consists of peripheral curves and essential curves that are invariant under f (up to isotopy).

In principle, a global curve attractor for a map f , as in Theorem 1.4, depends on the locations of the flaps. By Remark 8.6, for each concrete case, one can easily determine the exact attractor by checking if a slope $x \in \widehat{\mathbb{Q}}$ with $\|x\| \leq 8$ is invariant. For example, by using a computer program written by Darragh Glynn, we verified that for the map f corresponding to the flapped pillow in Figure 22 (with one horizontal flap and one vertical flap glued at the two 1-edges of \mathbb{P} incident to the vertex B), the invariant slopes are $0, \infty, 1, -1$.




FIGURE 22. A flapped pillow.

9. Further discussion

In this section, we briefly discuss some additional topics related to the investigations in this paper.

9.1. *Julia sets of blown-up Lattès maps.* An obvious question is what we can say about the Julia sets of the rational maps provided by Theorem 1.2 (for the definitions of the Julia and Fatou sets of rational maps and other basic notions in complex dynamics, see [Mil06a]).

PROPOSITION 9.1. *Let $g: \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$ be a rational map that is Thurston equivalent to a map $f: \mathbb{P} \rightarrow \mathbb{P}$ obtained from the $(n \times n)$ -Lattès map \mathcal{L}_n with $n \geq 2$ by gluing $n_h \geq 1$ horizontal and $n_v \geq 1$ vertical flaps to the pillow \mathbb{P} . Then the following statements are true.*

- (i) *The Julia set of g is equal to $\widehat{\mathbb{C}}$ if n is even and the vertex A is not contained in a flap, or if n is odd and none of the points in V is contained in a flap.*
- (ii) *The Julia set of g is equal to a Sierpiński carpet in $\widehat{\mathbb{C}}$ if n is even and A is contained in a flap, or if n is odd and at least one of the points in V is contained in a flap.*

Obviously, these cases cover all possibilities and so the Julia set of g is either the whole Riemann sphere $\widehat{\mathbb{C}}$ or a Sierpiński carpet, that is, a subset of $\widehat{\mathbb{C}}$ homeomorphic to the standard $1/3$ -Sierpiński carpet fractal. As we will see, in the first case, the map g has no periodic critical points, while it has periodic critical points (namely critical fixed points) in the second case.

Proof. Let g be a rational map as in the statement. To see what the Julia set of g is, we will check whether g has periodic critical points or not, and verify in the former case that g has no Levy arcs (see below for the definition). These conditions are invariant under Thurston equivalence and therefore it is enough to consider the map f . Then $P_f = V$, where $V = \{A, B, C, D\}$ is the set of vertices of \mathbb{P} . By definition of the $(n \times n)$ -Lattès map \mathcal{L}_n , for each $X \in V$, we have $\mathcal{L}_n(X) = A$ if n is even and $\mathcal{L}_n(X) = X$ if n is odd.

Since $f|_V$ agrees with $\mathcal{L}_n|_V$, this implies that for each $X \in V$, we also have $f(X) = A$ if n is even and $f(X) = X$ if n is odd.

Since the orbit of each critical point under iteration of f passes through the set $P_f = V$, this shows that any periodic critical point of f must be equal to the point A if n is even or must belong to V if n is odd. Now for $X \in V$, we have $\deg_{\mathcal{L}_n}(X) = 1$ and so

$$\deg_f(X) = n_X + \deg_{\mathcal{L}_n}(X) = n_X + 1,$$

where $n_X \in \mathbb{N}_0$ is the number of flaps that contain X . These considerations show that f , and hence also g , has no periodic critical points in case (i). Hence, the Julia set of g is the whole Riemann sphere $\widehat{\mathbb{C}}$ in this case (see [Mil06a, Corollary 19.8]).

In case (ii), the map f , and hence also g , has a critical fixed point, and so the Fatou set of g is non-empty. To show that its Julia set is a Sierpiński carpet, we use the following criterion that follows from [BD18, Lemma 4.16]: the Julia set of g is a Sierpiński carpet if and only if g , or equivalently f , has no Levy arcs. Here a *Levy arc* of f is a path α in \mathbb{P} satisfying the following conditions:

- (L1) α is an arc in (\mathbb{P}, V) , or α is a simple loop based at a point $X \in V$ such that $\alpha \setminus \{X\} \subset \mathbb{P} \setminus V$ and each component of $\mathbb{P} \setminus \alpha$ contains at least one point in V ;
- (L2) there exist $k \in \mathbb{N}$ and a lift $\tilde{\alpha}$ of α under f^k such that α and $\tilde{\alpha}$ are isotopic relative to V .

Now suppose that f has a Levy arc α with $\tilde{\alpha}$ and $k \in \mathbb{N}$ as in condition (L2). Then $f^k|_{\tilde{\alpha}}$ is a 1-to-1 map and either $i(\alpha, \alpha^h) > 0$ or $i(\alpha, \alpha^v) > 0$. Without loss of generality, we may assume that $i(\alpha, \alpha^h) = \#(\alpha \cap \alpha^h) > 0$. If α is an arc in (\mathbb{P}, V) , then we can apply Lemma 5.5 with $\gamma := \alpha^h$, $f := f^k$, and $\tilde{\alpha} := \tilde{\alpha}$ and conclude that the number of distinct pullbacks of α^h under f^k that are isotopic to α^h is at most 1. This is also true if α is a simple loop as in condition (L1) by the argument in the proof of Lemma 5.5.

We reach a contradiction, because it follows from Lemma 5.1 that α^h has $n^k > 1$ such pullbacks. Consequently, f and g do not have any Levy arcs and so the Julia set of g is a Sierpiński carpet. □

9.2. *The global curve attractor problem.* We were able to prove the existence of a finite global curve attractor only for blown-up $(n \times n)$ -Lattès maps with $n = 2$. The proof of Theorem 1.4 crucially relies on Proposition 8.1, which says that the (naturally defined) complexity of curves does not increase under the pullback operation. The latter statement is false in general for blown-up $(n \times n)$ -Lattès maps with $n \geq 3$.

Numerical computations by Darragh Glynn suggest that for some blown-up (3×3) -Lattès map f , one can have infinitely many slopes $x \in \widehat{\mathbb{Q}}$ such that $\|\mu_f(x)\| > \|x\|$. For example, consider the map f obtained from the (3×3) -Lattès map by blowing up once the horizontal and vertical edges incident to the vertex B of \mathbb{P} . Then one can prove the following general relation for the slope map μ_f :

$$\mu_f(r/s) = r'/s' \Rightarrow \mu_f(r/(s + 24r)) = r'/(s' + 22r').$$

Based on this, one can show that $\|\mu_f(x)\| > \|x\|$ for all

$$x \in \{1/(m + 24k) : m \in \{7, 8, 9, 15, 16, 17\}, k \in \mathbb{N}_0\}.$$

Actually, it seems that in this case, the slope map μ_f has orbits with arbitrarily many strict increases of complexity. For instance, we have two jumps of complexity for the orbit of slope $1/9$ under μ_f :

$$1/9 \rightarrow 3/25 \rightarrow 3/23 \rightarrow 1/7 \rightarrow 3/19 \rightarrow 3/17 \rightarrow 1/5 \rightarrow 1/5.$$

Note that this orbit stabilizes at the fixed point $1/5$ of μ_f . The numerical computations by Darragh Glynn also show that there are examples of blown-up $(n \times n)$ -Lattès maps with $n \geq 5$ for which the slope map has periodic cycles of length ≥ 2 .

It is natural to ask what one can say about the behavior of the slope map μ_f for an obstructed Thurston map f (with $\#P_f = 4$). It was already observed in [KPS16] that for a blown-up (2×2) -Lattès map f with only vertical flaps glued to the pillow \mathbb{P} , there are infinitely many (non-isotopic) invariant essential Jordan curves. Indeed, for such a map f , the curve α^v is f -invariant and satisfies $\lambda_f(\alpha^v) = 1$. One can use this to show that f commutes with T^2 (up to isotopy relative to P_f), where T is a *Dehn twist* about α^v . This implies that each curve $T^{2n}(\alpha^h)$ is f -invariant. In fact, it is easy to verify directly that each essential Jordan curve with slope $x \in \mathbb{Z} \cup \{\infty\}$ is f -invariant, or equivalently, that $\mu_f(x) = x$ for $x \in \mathbb{Z} \cup \{\infty\}$.

However, for such a blown-up (2×2) -Lattès map f with only vertical flaps glued to the pillow \mathbb{P} , the general behavior of the slope map μ_f under iteration has not been analyzed before. The considerations in the proof of the first part of Proposition 8.1 also apply in this situation. In particular, (8.4) and (8.5) are still true and show that the orbit of an arbitrary $x \in \widehat{\mathbb{Q}} \cup \{\infty\}$ under μ_f eventually lands in a fixed point of μ_f . Moreover, results in §8 provide a method to determine all fixed slopes for μ_f .

The easiest case is the map f obtained from the (2×2) -Lattès map \mathcal{L}_2 by gluing at least one vertical flap to each of the four vertical 1-edges in the ‘middle’ of the pillow \mathbb{P} . If ξ is a geodesic arc in (\mathbb{P}, V) with an endpoint in $A \in V$ and slope $x \in \mathbb{Q} \setminus \mathbb{Z}$, then each component of $\mathcal{L}_2^{-1}(\xi)$ must pass through the interior of one of the four vertical 1-edges in the middle of \mathbb{P} . Consequently, we can apply the considerations in the proof of Lemma 8.4 and in the second part of the proof of Proposition 8.1, and conclude that $\|\mu_f(x)\| < \|x\|$ for $x \in \mathbb{Q} \setminus \mathbb{Z}$. Thus, the orbit of each $x \in \widehat{\mathbb{Q}} \cup \{\circ\}$ under μ_f eventually lands in $\mathbb{Z} \cup \{\infty, \circ\}$ (that is, in a fixed point of μ_f). Since the map f is easily seen to be *expanding* (see [BM17, Definition 2.2 and Theorem 14.1]), this provides an answer to a question raised by Pilgrim of whether there is an obstructed expanding Thurston map for which one has a complete understanding of the global dynamics of the slope map.

9.3. *Twisting problems.* Many natural problems related to Thurston equivalence remain rather mysterious and are often very difficult to solve. *Twisting problems* are examples of this nature.

To explain this, suppose we are given a rational Thurston map $f: \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$. Let $\phi: \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$ be an orientation-preserving homeomorphism that fixes the postcritical set P_f pointwise. We now consider the branched covering map $g := \phi \circ f$ on $\widehat{\mathbb{C}}$, called the ϕ -*twist* of f . Then $C_g = C_f$ and g has the same dynamics on C_f as f . In particular, $P_g = P_f$; so g has a finite postcritical set and is a Thurston map.

This leads us to the natural questions: *Is g realized? And if yes, to which rational map is g equivalent depending on the isotopy type of ϕ ?* In fact, there are only finitely many rational maps g (up to Möbius conjugation) that can arise in this way from a fixed map f . A famous instance of this question, called the ‘twisted rabbit problem’, was solved by Bartholdi and Nekrashevych in [BN06] (see also [BLMW22, Lod13]).

In our context, we can ask which twists of maps, as in Theorem 1.2, are realized. We do not have an answer to this question, but it seems that this leads to non-trivial and difficult problems. For example, consider the blown-up (2×2) -Lattès map $f: \mathbb{P} \rightarrow \mathbb{P}$ corresponding to the flapped pillow $\widehat{\mathbb{P}}$ as in Figure 22. Then $\widehat{\mathbb{P}}$ has one horizontal and one vertical flap, and so f is realized by Theorem 1.2. One can check that the Jordan curve $\gamma := \wp(\ell_{3/13})$ has exactly two essential pullbacks $\gamma_1, \gamma_2 \sim \wp(\ell_{1/3})$ under f with $\deg(f: \gamma_1 \rightarrow \gamma) = 1$ and $\deg(f: \gamma_2 \rightarrow \gamma) = 2$. We now choose an orientation-preserving homeomorphism $\phi: \mathbb{P} \rightarrow \mathbb{P}$ that maps γ onto γ_1 , while fixing each point in $V = P_f$. Then the curve γ_1 is an obstruction for the twisted map $g := \phi \circ f$ with $\lambda_g(\gamma_1) = 3/2$. An analogous construction applies to some other essential Jordan curves, for instance, with slopes $3/23$ and $3/49$, and gives twists of f with an obstruction.

It follows from this discussion that the *mapping class biset* associated with the map f above is not *contracting* (see [BD17, BD18] for the definitions). Thus, the algebraic methods for solving the global curve attractor problem developed in [Pil12] (see, specifically, [Pil12, Theorem 1.4]) do not apply in general for the maps considered in Theorem 1.2.

9.4. *Thurston maps with more than four postcritical points.* While in this paper we only discuss the case of Thurston maps $f: S^2 \rightarrow S^2$ with $\#P_f = 4$, it is natural to ask if one can adapt Theorem 1.1 to the case when $\#P_f > 4$. The main difficulty is that an obstruction in this case is in general not given by a unique essential Jordan curve in (S^2, P_f) , but by a *multicurve*. Of course, this fact complicates the analysis of pullback properties of curves and their intersection numbers. However, we expect that one can naturally generalize our result for an arbitrary Thurston map: given an obstructed Thurston map f , one can eliminate all possible multicurve obstructions by successively applying the blow-up operation and obtain a Thurston map that is realized.

9.5. *Other combinatorial constructions of rational maps.* The dynamical behavior of curves under the pullback operation is an important topic in holomorphic dynamics. While in this paper we only study the realization and the global curve attractor problems, one is led to similar considerations, for example, in the study of *iterated monodromy groups* (see [HM18]). For these investigations, it is important to have explicit classes of rational maps at hand that are constructed in combinatorial fashion and against which conjectures can be tested or which lead to the discovery of general phenomena. The maps provided by Theorem 1.1 may be useful in this respect. Another interesting class of maps worthy of further investigation are Thurston maps constructed from tilings of the Euclidean or hyperbolic plane as in [BM17, Example 12.25].

Acknowledgements. The authors would like to thank Kostya Drach, Dima Dudko, Daniel Meyer, Kevin Pilgrim, and Dylan Thurston for various useful comments and remarks. We

are grateful to Darragh Glynn for allowing us to incorporate some of his numerical findings in this paper.

M.B. was partially supported by NSF grant DMS-1808856. M.H. was partially supported by the ERC advanced grant ‘HOLOGRAM’. A.I. was partially supported by the Swiss National Science Foundation (project no. 181898).

A. Appendix. Isotopy classes of Jordan curves in spheres with four marked points

In this appendix, we will provide proofs for Lemmas 2.3 and 2.4. Our presentation is rather detailed. We need some additional auxiliary facts that we will discuss first. Throughout, we will rely on the notation and terminology established in §2.

In the following, we will consider a marked sphere (S^2, Z) , where $Z \subset S^2$ consists of precisely four points. If $M \subset S^2$ and α is a Jordan curve in (S^2, Z) , then we say that α is in *minimal position with the set M* if $\#(\alpha \cap M) \leq \#(\alpha' \cap M)$ for all Jordan curves α' in (S^2, Z) with $\alpha \sim \alpha'$ relative to Z .

Let α and β be Jordan curves or arcs in (S^2, Z) . We say that subarcs $\alpha' \subset \alpha$ and $\beta' \subset \beta$ form a *bigon* U in (S^2, Z) if α' and β' have the same endpoints, but disjoint interiors, and if $U \subset S^2$ is an open Jordan region with $\partial U = \alpha' \cup \beta'$ and $U \subset S^2 \setminus Z$.

LEMMA A.1. *Let γ be a Jordan curve in a marked sphere (S^2, Z) with $\#Z = 4$, and let a and c be disjoint arcs in (S^2, Z) . Suppose γ is in minimal position with the set $a \cup c$. Then γ meets each of the arcs a and c transversely and no subarcs of γ and of a or c form a bigon U in (S^2, Z) with $U \cap (\gamma \cup a \cup c) = \emptyset$.*

Proof. These facts are fairly standard in contexts like this (see, for example, [FM12, §1.2.4]), and so we will only give an outline of the proof.

Our assumptions imply that γ meets each arc a and c transversally and has only finitely many intersections with $a \cup c$ (Lemma 2.1 and its proof apply *mutatis mutandis* to our situation). We now argue by contradiction and assume that a subarc $\gamma' \subset \gamma$ and a subarc σ of a or c form a bigon U in (S^2, Z) with $U \cap (\gamma \cup a \cup c) = \emptyset$. Note that then, $\text{cl}(U) \subset S^2 \setminus Z$. Hence, we can modify the curve γ near U by an isotopy in $S^2 \setminus Z$ that pulls the subarc γ' of γ through U and away from σ so that the new Jordan curve γ does not intersect $\sigma \subset a \cup c$ and no new intersection points with $a \cup c$ arise. This leads to a contradiction, because the original curve γ was in minimal position with $a \cup c$. \square

A topological space D is called a *closed topological disk* if there exists a homeomorphism $\eta: \text{cl}(\mathbb{D}) \rightarrow D$ of the closed unit disk $\text{cl}(\mathbb{D}) \subset \mathbb{C}$ onto D . This is an abstract version of the notion of a closed Jordan region contained in a surface. The set $\partial D := \eta(\partial \mathbb{D})$ is a Jordan curve independent of η and called the *boundary* of D . The *interior* of D is defined as $\text{int}(D) := \eta(\mathbb{D}) = D \setminus \partial D$.

Similarly as for closed Jordan regions, an arc α contained in a closed topological disk D is called a *crosscut* (in D) if $\partial \alpha \subset \partial D$ and $\text{int}(\alpha) \subset \text{int}(D)$. A crosscut α splits D into two compact and connected sets S and S' called the *sides* of α (in D) such that $D = S \cup S'$ and $S \cap S' = \alpha$. With suitable orientations of α and D , one side of α lies on the left and the other side on the right of α . Each non-empty connected set $c \subset D$ that does not meet α is contained in precisely one side of α .

If $\alpha_1, \dots, \alpha_n$ for $n \in \mathbb{N}$ are pairwise disjoint crosscuts in a closed topological disk D , then we can define an abstract graph $G = (V, E)$ in the following way: we consider each component U of $D \setminus (\alpha_1 \cup \dots \cup \alpha_n)$ as a vertex of G . We join two distinct vertices represented by components U and U' by an edge if one of the crosscuts α_j is contained in the boundary of both U and U' . Accordingly, the edges of G are in bijective correspondence with the crosscuts $\alpha_1, \dots, \alpha_n$.

LEMMA A.2. *Let $n \in \mathbb{N}$ and $\alpha_1, \dots, \alpha_n$ be pairwise disjoint crosscuts in a closed topological disk D , and let G be the graph obtained from the components of the set $D \setminus (\alpha_1 \cup \dots \cup \alpha_n)$ as described. The the following statements are true.*

- (i) *The graph G is a finite tree with at least two vertices.*
- (ii) *Let $c \subset D$ be a connected set and suppose that $c \cap \alpha_k = \emptyset$ for some $k \in \{1, \dots, n\}$. Then there exists $m \in \{1, \dots, n\}$ and a side S of α_m such that $S \setminus \alpha_m$ is disjoint from all the sets $c, \alpha_1, \dots, \alpha_n$.*

If $\alpha = \alpha_m$ and S are as in statement (ii), then there exists a subarc β of ∂D with the same endpoints as $\partial \alpha \subset D$ such that $\partial S = \alpha \cup \beta$. Then $U := \text{int}(S)$ is an open Jordan region bounded by the union of the arcs α and β whose only common points are their endpoints. This region U does not meet c nor any of the arcs $\alpha_1, \dots, \alpha_n$. In the proof of Lemma A.3, we will use such a region U to obtain a bigon in an appropriate context.

Proof. (i) This is intuitively clear, and we leave the details to the reader. By induction on the number n of crosscuts, one can show that G is a finite connected graph with at least two vertices. Since a crosscut splits D into two sides, it easily follows that the removal of any edge from G disconnects it. Hence, G cannot contain any simple cycle and must be a tree.

(ii) The graph G is a tree; so if we remove the edge corresponding to the crosscut α_k from G , then we obtain two disjoint non-empty subgraphs G_1 and G_2 of G . The connected components of $D \setminus (\alpha_1 \cup \dots \cup \alpha_n)$ corresponding to the vertices of G_1 are contained in one side S' of α_k , while the other connected components of $D \setminus (\alpha_1 \cup \dots \cup \alpha_n)$ corresponding to the vertices of G_2 lie in the other side S'' of α_k . Since c is connected and does not meet α_k , it must be contained in one of the sides of α_k , say $c \subset S''$. Then c is disjoint from S' and hence from all the sets that correspond to vertices in G_1 .

The tree G has a leaf v in $G_1 \neq \emptyset$, that is, there exists a vertex v of G_1 such that v is connected to the rest of G by precisely one edge. Then the connected component of $D \setminus (\alpha_1 \cup \dots \cup \alpha_n)$ corresponding to v has exactly one of the crosscuts, say α_m with $m \in \{1, \dots, n\}$, on its boundary. Then this component has the form $S \setminus \alpha_m$, where S is the unique side of α_m contained in S' . Then $S \setminus \alpha_m$ is disjoint from $c \subset S''$ and from all the crosscuts $\alpha_1, \dots, \alpha_n$. \square

We can now prove a statement that is the key to the understanding of isotopy classes of Jordan curves in a sphere with four marked points.

LEMMA A.3. *Let γ be a Jordan curve in a marked sphere (S^2, Z) with $\#Z = 4$, and let a and c be disjoint arcs in (S^2, Z) . Suppose γ is in minimal position with the set $a \cup c$. Then*

the sets $a \cap \gamma$ and $c \cap \gamma$ are non-empty and finite, and the points in these sets alternate on γ unless γ is peripheral or $\gamma \cap (a \cup c) = \emptyset$.

Proof. In the given setup, each of the disjoint sets ∂a and ∂c contains two points in Z . Since $\#Z = 4$, it follows that a connects two of the points in Z , while c connects the other two points.

We may assume that γ is essential and that at least one of the arcs a or c meets γ , say $a \cap \gamma \neq \emptyset$, because otherwise, we are in an exceptional situation as in the statement.

If none of the arcs a and c meets γ in more than two points, then $\#(a \cap \gamma) = 1$ and $\#(c \cap \gamma) \leq 1$. Now a and γ meet transversely by Lemma A.1. This implies that the endpoints of a lie in different components of $S^2 \setminus \gamma$. Since γ is essential, each of these components contains precisely two points of $Z = \partial a \cup \partial c$. Hence, the endpoints of c also lie in different components of $S^2 \setminus \gamma$. This implies that $c \cap \gamma \neq \emptyset$ and so $\#(c \cap \gamma) = 1$ in the case under consideration. So both a and c meet γ in exactly one point. It follows that the statement is true in this case.

We are reduced to the situation where at least one of the arcs a or c meets γ in at least two (but necessarily finitely many) points, say $n := \#(a \cap \gamma) \geq 2$. We now endow γ and a with some orientations. With the given orientation, we denote the initial point of a by x_0 and its terminal point by x_1 . Let $y_1, \dots, y_n, y_{n+1} = y_1$ denote the $n \geq 2$ intersection points of γ with a that we encounter while traversing γ once starting from some point in $\gamma \setminus a$. The same n points also appear on a . We denote them by p_1, \dots, p_n in the order they appear if we traverse a starting from x_0 . For $k = 1, \dots, n$, we denote by $\gamma[y_k, y_{k+1}]$ the subarc of γ obtained from traversing γ with the given orientation from y_k to y_{k+1} .

Claim. $c \cap \gamma[y_k, y_{k+1}] \neq \emptyset$ for each $k = 1, \dots, n$.

To see this, we argue by contradiction and assume that $c \cap \gamma[y_k, y_{k+1}] = \emptyset$ for some $k \in \{1, \dots, n\}$. Our goal now is to show that some subarcs of a and γ form a bigon U in (S^2, Z) with $U \cap (\gamma \cup a \cup c) = \emptyset$. This is a contradiction with Lemma A.1, because γ and $a \cup c$ are in minimal position.

To produce such bigon U , we want to apply Lemma A.2. To do this, we slit the sphere S^2 open along the arc a . This results in a closed topological disk D whose boundary ∂D consists of two copies a^+ and a^- of the arc a . The set $S^2 \setminus \text{int}(a)$ can be identified with $\text{int}(D)$, while each point in $\text{int}(a)$ is doubled into one corresponding point in a^+ and one in a^- .

The arcs a^+ and a^- have their endpoints x_0 and x_1 in common. We identify a^+ with the original arc a with the same orientation. Then we can think of the intersection points p_1, \dots, p_n of γ with a as lying on $a^+ = a$, while each of the points p_j has a corresponding point q_j on a^- .

Each arc $\gamma[y_j, y_{j+1}]$ corresponds to a crosscut γ_j in D for $j = 1, \dots, n$. These crosscuts have their endpoints in the set $P \cup Q$, where $P := \{p_1, \dots, p_n\}$ and $Q := \{q_1, \dots, q_n\}$. Moreover, the crosscuts $\gamma_1, \dots, \gamma_n$ are pairwise disjoint. Indeed, the only possible common intersection point of two of these arcs could be a common endpoint of two consecutive arcs γ_j and γ_{j+1} (where $\gamma_{n+1} := \gamma_1$) corresponding to $y_{j+1} \in a$; but in the process of creating D , the point y_{j+1} is doubled into the points p_ℓ and q_ℓ for some

$\ell \in \{1, \dots, n\}$. Since γ meets a transversely, one of these points will be the terminal point of γ_j , while the other one will be the initial point of γ_{j+1} , and so actually, $\gamma_j \cap \gamma_{j+1} = \emptyset$. It follows that the hypotheses of Lemma A.2 are satisfied.

It is clear that the arc c , now considered as a subset of D , does not meet the crosscut γ_k corresponding to $\gamma[y_k, y_{k+1}]$. Hence, by Lemma A.2, there exists $m \in \{1, \dots, n\}$ and a side S of γ_m in D such that $S \setminus \gamma_m$ is disjoint from c and all the arcs $\gamma_1, \dots, \gamma_n$. Then there exists an arc $\beta \subset \partial D = a^+ \cup a^-$ with the same endpoints as γ_m such that $\partial S = \gamma_m \cup \beta$. The set $S \setminus \gamma_m$ is disjoint from $\gamma_1 \cup \dots \cup \gamma_n \supset P \cup Q$, and so the arc β has its endpoints in the set $P \cup Q$, but no other points in common with $P \cup Q$.

This implies that neither x_0 nor x_1 are contained in β ; indeed, suppose $x_0 \in \beta$, for example. Then the endpoints of β and hence of γ_m are necessarily the points p_1 and q_1 . Collapsing D back to S^2 , we see that the endpoints y_m and y_{m+1} of $\gamma[y_m, y_{m+1}]$ are the same. This is a contradiction (here the assumption $n \geq 2$ is crucial). We arrive at a similar contradiction (using the points p_n and q_n) if we assume $x_1 \in \beta$. It follows that $\beta \subset \text{int}(a^+)$ or $\beta \subset \text{int}(a^-)$.

These considerations imply that if we pass back to S^2 by identifying corresponding points in a^+ and a^- , then from $\text{int}(S)$, we obtain a bigon $U \subset S^2$ bounded by the subarc $\gamma[y_m, y_{m+1}]$ of γ and a subarc $\tilde{\beta}$ of a , where U is disjoint from $\gamma \cup a \cup c$. This is impossible by Lemma A.1 since γ is in minimal position with $a \cup c$. This contradiction shows that the claim is indeed true.

The claim implies that c has at least $n \geq 2$ intersection points with γ . Hence, we can reverse the roles of a and c and get a similar statement as the claim also for the arc c . This implies that the (finitely many) points in $a \cap \gamma \neq \emptyset$ and $c \cap \gamma \neq \emptyset$ alternate on γ . \square

As in §2.4, we now consider the pillow \mathbb{P} with its set V of vertices as the marked points, and the Weierstrass function $\wp: \mathbb{C} \rightarrow \mathbb{P}$ that is doubly periodic with respect to the lattice $2\mathbb{Z}^2$. We will now revert to the notation a and c for the horizontal edges of \mathbb{P} . Recall that $\mathbb{I} = [0, 1]$.

LEMMA A.4. *Let $\alpha: \mathbb{I} \rightarrow \mathbb{C}$ be a simple loop or a homeomorphic parameterization of an arc. Suppose that the endpoints $z_0 = \alpha(0)$ and $w_0 = \alpha(1)$ of α lie in $\mathbb{C} \setminus \wp^{-1}(a \cup c)$ and that $\wp(z_0) = \wp(w_0)$. Suppose further that either $\alpha \cap \wp^{-1}(a \cup c) = \emptyset$, or all of the following conditions are true: α meets each of the lines in $\wp^{-1}(a \cup c)$ transversely, we have*

$$0 < \#(\alpha \cap \wp^{-1}(a)) = \#(\alpha \cap \wp^{-1}(c)) < \infty,$$

and the points in $\alpha \cap \wp^{-1}(a)$ and $\alpha \cap \wp^{-1}(c)$ alternate on α .

Then $w_0 - z_0 \in 2\mathbb{Z}^2$. Here $w_0 - z_0 \neq 0$ unless $\alpha \cap \wp^{-1}(a \cup c) = \emptyset$.

Proof. Note that the set $\wp^{-1}(a \cup c)$ consists precisely of the lines $L_n := \{z \in \mathbb{C} : \text{Im}(z) = n\}$, $n \in \mathbb{Z}$. Moreover, such a line L_n is mapped to a or c depending on whether n is even or odd, respectively.

We consider the second case first when $\alpha \cap \wp^{-1}(a \cup c) \neq \emptyset$. We denote by $0 < u_1 < \dots < u_k < 1$, $k \in \mathbb{N}$, all the (finitely many) u -parameter values with $\alpha(u_j) \in \wp^{-1}(a \cup c)$ for $j = 1, \dots, k$. Since α meets each line L_n transversely and the points

in $\alpha \cap \wp^{-1}(a) \neq \emptyset$ and $\alpha \cap \wp^{-1}(c) \neq \emptyset$ alternate on α , it is clear that the values $\text{Im}(\alpha(u_j))$ either strictly increase by 1 in each step $j = 1, \dots, k$, or strictly decrease by 1 in each step. Here we have strict increase if $\text{Im}(\alpha(u_1) - \alpha(0)) > 0$, and strict decrease if $\text{Im}(\alpha(u_1) - \alpha(0)) < 0$. This implies that $k > 0$ is precisely the number of lines $L_n, n \in \mathbb{Z}$, that separate z_0 and w_0 . So z_0 and w_0 lie in different components of $\mathbb{C} \setminus \wp^{-1}(a \cup c)$ which shows that $z_0 \neq w_0$.

By our hypotheses, $\#(\alpha \cap \wp^{-1}(a)) = \#(\alpha \cap \wp^{-1}(c))$, which implies that the number k of intersection points of α with $\wp^{-1}(a \cup c)$ is even. Since $\wp(z_0) = \wp(w_0)$, by (2.3), we have $w_0 = \pm z_0 + v_0$ with $v_0 \in 2\mathbb{Z}^2$. We have to rule out the minus sign here.

We argue by contradiction and assume that $w_0 = -z_0 + v_0$. Then $v_1 := \frac{1}{2}(z_0 + w_0) = \frac{1}{2}v_0 \in \mathbb{Z}^2$, and so the endpoints z_0 and w_0 of α are in symmetric position to the point $v_1 \in \mathbb{Z}^2$. This implies that the number k of lines $L_n, n \in \mathbb{Z}$, separating z_0 and w_0 is odd, contradicting what we have just seen. We conclude $w_0 = z_0 + v_0$ with $v_0 \in 2\mathbb{Z}^2$, and the statement follows in this case. Note that the exact same argument leading to $w_0 - z_0 \in 2\mathbb{Z}^2$ also applies if $\alpha \cap \wp^{-1}(a \cup c) = \emptyset$. □

We call a Jordan curve γ in (\mathbb{P}, V) *null-homotopic in (\mathbb{P}, V)* if γ can be homotoped in $\mathbb{P} \setminus V$ to a point, that is, if there exists a homotopy $H : \partial\mathbb{D} \times \mathbb{I} \rightarrow \mathbb{P} \setminus V$ such that H_0 is a homeomorphism of $\partial\mathbb{D}$ onto γ and H_1 is a constant map.

LEMMA A.5. *Let γ be an essential Jordan curve in (\mathbb{P}, V) . Then γ is not null-homotopic in (\mathbb{P}, V) .*

This statement sounds somewhat tautological, because in topology, ‘essential’ is often defined as ‘not null-homotopic’. Recall though that in our context γ is called essential if each of the two components of $\mathbb{P} \setminus \gamma$ contains precisely two of the points in V . In the proof, we will use some standard facts about winding numbers; see [Bur79, Ch. 4] for the basic definitions and background.

Proof. On an intuitive level, every homotopy contracting γ to a point must slide over all the points in one of the complementary components of γ . Hence, it cannot stay in $\mathbb{P} \setminus V$, and so γ is not null-homotopic in (\mathbb{P}, V) .

To make this more rigorous, we argue by contradiction. By the Schönflies theorem, we may identify \mathbb{P} with $\widehat{\mathbb{C}}$ and γ with $\partial\mathbb{D}$ and assume that 0 and ∞ belong to the set $Z \subset \widehat{\mathbb{C}}$ of marked points corresponding to the points in V . We now argue by contradiction and assume that there exists a homotopy $H : \partial\mathbb{D} \times \mathbb{I} \rightarrow \widehat{\mathbb{C}} \setminus Z \subset \mathbb{C} \setminus \{0\}$ such that H_0 is a homeomorphism on $\partial\mathbb{D}$ and H_1 is a constant map. Then for each $t \in \mathbb{I}$, the map $u \in \mathbb{I} \mapsto \alpha_t(u) := H_t(e^{2\pi i u})$ is a loop in $\mathbb{C} \setminus \{0\}$. Each loop $\alpha_t, t \in \mathbb{I}$, has the same winding number $\text{ind}_{\alpha_t}(0)$ around 0, because this winding number is invariant under homotopies in $\mathbb{C} \setminus \{0\}$ (see [Bur79, Theorem 4.12]). However, $\text{ind}_{\alpha_0}(0) = \pm 1$, because α_0 is a simple loop (see [Bur79, Theorem 4.42]), while $\text{ind}_{\alpha_1}(0) = 0$, because α_1 is a constant loop. This is a contradiction. □

An element x of a rank-2 lattice Γ (such as $2\mathbb{Z}^2$) in \mathbb{C} is called *primitive* if it cannot be represented in the form $x = ny$ with $y \in \Gamma$ and $n \in \mathbb{N}, n \geq 2$. Note that then, $x \neq 0$.

LEMMA A.6. Let γ be a Jordan curve in (\mathbb{P}, V) parameterized as a simple loop $\beta: \mathbb{I} \rightarrow \mathbb{P}$, and let $\alpha: \mathbb{I} \rightarrow \mathbb{C}$ be a lift of β under \wp .

- (i) If $\alpha(0) = \alpha(1)$ and α is a path in a convex subset of $\mathbb{C} \setminus \mathbb{Z}^2$, then γ is null-homotopic in (\mathbb{P}, V) .
- (ii) If γ is essential, then $\alpha(1) - \alpha(0)$ is a primitive element of $2\mathbb{Z}^2$. It is uniquely determined up to sign by the isotopy class $[\gamma]$ of γ relative to V .

Here we call α a lift of β under \wp if $\beta = \wp \circ \alpha$. In this lemma and its proof, we will carefully distinguish between a path and its image set (unlike elsewhere in the paper).

Proof. Note that since $\beta(\mathbb{I}) = \gamma \subset \mathbb{P} \setminus V$ and \wp is a covering map over $\mathbb{P} \setminus V$, a lift α of β under \wp exists. Moreover, for each choice of $z_0 \in \wp^{-1}(\beta(0))$, there exists a unique lift α of β such that $\alpha(0) = z_0$ (for these standard facts see [Hat02, §1.3, Proposition 1.30]). We will use this uniqueness property of lifts repeatedly in the following.

(i) The idea for the first part is very simple. We use a ‘straight-line homotopy’ between α and the constant path $u \in \mathbb{I} \mapsto \alpha(0)$ and push it to $\mathbb{P} \setminus V$ by applying \wp .

More precisely, we define

$$H(u, t) := \wp((1 - t)\alpha(u) + t\alpha(0))$$

for $u, t \in \mathbb{I}$. Since the path α lies in a convex set $K \subset \mathbb{C} \setminus \mathbb{Z}^2$, we have

$$\alpha_t(u) := (1 - t)\alpha(u) + t\alpha(0) \in K$$

for all $u, t \in \mathbb{I}$, and so $H(\mathbb{I} \times \mathbb{I}) \subset \wp(K) \subset \mathbb{P} \setminus V$. Hence, H is a homotopy in $\mathbb{P} \setminus V$. Moreover, $H_t(u) = (\wp \circ \alpha_t)(u)$ for all $u \in \mathbb{I}$, and so $H_t = \wp \circ \alpha_t$ for all $t \in \mathbb{I}$. In particular, $H_0 = \wp \circ \alpha_0 = \wp \circ \alpha = \beta$. Moreover, $H_1(u) = \alpha(0)$ for $u \in \mathbb{I}$, and so H_1 is a constant path.

For all $t \in \mathbb{I}$, we have

$$\alpha_t(1) - \alpha_t(0) = (1 - t)(\alpha(1) - \alpha(0)) = 0.$$

Hence, $\wp(\alpha_t(1)) = \wp(\alpha_t(0))$ for $t \in \mathbb{I}$, and so $H_t = \wp \circ \alpha_t$ is a loop in $\mathbb{P} \setminus V$ for all $t \in \mathbb{I}$. By identifying the points $(0, t)$ and $(1, t)$ for each $t \in \mathbb{I}$, we get a homotopy $\overline{H}: \partial\mathbb{D} \times \mathbb{I} \rightarrow \mathbb{P} \setminus V$ such that

$$\overline{H}(e^{2\pi i u}, t) = H(u, t)$$

for all $u, t \in \mathbb{I}$. Since $\overline{H}_0(e^{2\pi i u}) = H_0(u) = \beta(u)$ for $u \in \mathbb{I}$, we see that \overline{H}_0 is a homeomorphism of $\partial\mathbb{D}$ onto $\beta(\mathbb{I}) = \gamma$. However, \overline{H}_1 is a constant map. Hence, γ is null-homotopic in (\mathbb{P}, V) .

(ii) The proof is somewhat tedious as we have to worry about different choices of the curve in $[\gamma]$, its different parameterizations as a simple loop, and the different lifts of these parameterizations under \wp .

To prove the statement, we first consider a special case, namely we choose a Jordan curve γ_0 in (\mathbb{P}, V) that lies in the same isotopy class relative to V as γ with the additional property that γ_0 is in minimal position with the set $a \cup c$. Since γ_0 cannot be a subset of $a \cup c$, we can parameterize γ_0 as a simple loop $\beta_0: \mathbb{I} \rightarrow \mathbb{P}$ such that $\beta_0(0) = \beta_0(1) \notin$

$a \cup c$. We now consider a lift $\alpha_0: \mathbb{I} \rightarrow \mathbb{C}$ of β_0 under \wp . Note that α_0 is a simple loop or a homeomorphic parameterization on an arc in \mathbb{C} .

Since γ is essential, the Jordan curve γ_0 is also essential. Indeed, under an isotopy relative to V that deforms γ into γ_0 , the complementary components of γ in \mathbb{P} are deformed into the complementary components of γ_0 while the points in V stay fixed. Therefore, each of the two components of $\mathbb{P} \setminus \gamma_0$ contains precisely two points of V .

It follows from Lemma A.1 that γ_0 meets a and c transversely. Moreover, by Lemma A.3, either $\gamma_0 \cap (a \cup c) = \emptyset$, or the sets $a \cap \gamma_0$ and $c \cap \gamma_0$ are non-empty and finite, and the points in these sets alternate on γ_0 . We now define $\bar{\alpha}_0 := \alpha_0(\mathbb{I})$. Then $\bar{\alpha}_0$ is a Jordan curve or an arc in \mathbb{C} . Moreover, we either have $\bar{\alpha}_0 \cap \wp^{-1}(a \cup c) = \emptyset$, or $\bar{\alpha}_0$ meets each of the lines in $\wp^{-1}(a \cup c)$ transversely,

$$0 < \#(\bar{\alpha}_0 \cap \wp^{-1}(a)) = \#(\bar{\alpha}_0 \cap \wp^{-1}(c)) < \infty,$$

and the points in $\bar{\alpha}_0 \cap \wp^{-1}(a)$ and $\bar{\alpha}_0 \cap \wp^{-1}(c)$ alternate on $\bar{\alpha}_0$. Let $z_0 := \alpha_0(0)$ and $w_0 := \alpha_0(1)$. Then

$$\wp(z_0) = \wp(\alpha_0(0)) = \beta_0(0) = \beta_0(1) = \wp(\alpha_0(1)) = \wp(w_0),$$

and by the choice of β_0 , we have $z_0, w_0 \notin \wp^{-1}(a \cup c)$. Therefore, we are exactly in the situation of Lemma A.4.

It follows that $v_0 := w_0 - z_0 \in 2\mathbb{Z}^2$. Here $v_0 \neq 0$. Indeed, otherwise $z_0 = w_0$. Then the second part of Lemma A.4 implies that the arc $\bar{\alpha}_0$ does not meet $\wp^{-1}(a \cup c)$ and so it lies in a connected component of $\mathbb{C} \setminus \wp^{-1}(a \cup c)$. This component is an infinite strip, and hence a convex set, contained in $\mathbb{C} \setminus \mathbb{Z}^2$. Now part (i) implies that γ_0 is null-homotopic in (\mathbb{P}, V) . By Lemma A.5, this contradicts the fact that γ_0 is essential. We conclude that indeed $v_0 = w_0 - z_0 \neq 0$.

This shows that $\alpha_0(1) - \alpha_0(0) = w_0 - z_0$ is a non-zero element of the lattice $2\mathbb{Z}^2$. We claim that $w_0 - z_0$ is actually a primitive element of $2\mathbb{Z}^2$. To see this, we argue by contradiction and assume that $w_0 - z_0 = ny_0$ with $y_0 \in 2\mathbb{Z}^2 \setminus \{0\}$ and $n \in \mathbb{N}, n \geq 2$.

We now consider the path $\sigma: \mathbb{I} \rightarrow \mathbb{C} \setminus \{0\}$ given as $\sigma(u) = \exp(2\pi i/y_0)\alpha_0(u)$ for $u \in \mathbb{I}$. This is a loop with winding number

$$\text{ind}_\sigma(0) = \frac{1}{y_0}(\alpha_0(1) - \alpha_0(0)) = n$$

around 0. Now a simple loop in $\mathbb{C} \setminus \{0\}$ has winding number 0 or ± 1 around 0 (see [Bur79, Theorem 4.42]), and so σ cannot be simple. This implies that there are numbers $0 \leq u < u' < 1$ such that $\sigma(u) = \sigma(u')$. This in turn means that $\alpha_0(u') - \alpha_0(u) = ky_0$ for some $k \in \mathbb{Z}$. Since $y_0 \in 2\mathbb{Z}^2$, it follows that

$$\beta_0(u') = \wp(\alpha_0(u')) = \wp(\alpha_0(u)) = \beta_0(u).$$

This is impossible, since β_0 is injective on $[0, 1)$.

We have shown the first part of the statement for a particular Jordan curve γ_0 in $[\gamma]$ with a special parameterization β_0 , and a choice of a lift α_0 of β_0 under \wp . We now have to show that the number $v_0 = w_0 - z_0 = \alpha_0(1) - \alpha_0(0)$ obtained in this way only depends

on $[\gamma]$ up to sign. For this, we pick an arbitrary Jordan curve in $[\gamma]$ which we will simply call γ .

Since $\gamma_0 \sim \gamma$ relative to V , there exists an isotopy $H: \mathbb{P} \times \mathbb{I} \rightarrow \mathbb{P}$ relative to V with $H_0 = \text{id}_{\mathbb{P}}$ and $H_1(\gamma_0) = \gamma$. We now define a homotopy $\overline{H}: \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{P} \setminus V$ by setting

$$\overline{H}(u, t) := H(\beta_0(u), t)$$

for $u, t \in \mathbb{I}$. Note that \overline{H} maps into $\mathbb{P} \setminus V$ as follows from the facts that $\gamma_0 \subset \mathbb{P} \setminus V$ and H is an isotopy relative to V .

The time-0 map $\overline{H}(\cdot, 0) = \beta_0$ of the homotopy \overline{H} is the parameterization of the loop γ_0 , while the time-1 map $\beta := \overline{H}(\cdot, 1) = H_1 \circ \beta_0$ gives some parameterization of $\gamma = H_1(\gamma_0)$ as a simple loop.

By the homotopy lifting theorem (see [Hat02, Proposition 1.30]), there exists a homotopy $\tilde{H}: \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{C} \setminus \mathbb{Z}^2$ such that $\overline{H} = \wp \circ \tilde{H}$ and $\tilde{H}_0 = \tilde{H}(\cdot, 0) = \alpha_0$. Then $\alpha := \tilde{H}(\cdot, 1)$ is a lift of $\beta = \overline{H}(\cdot, 1)$ under \wp . We want to show that $\alpha(1) - \alpha(0) = v_0 = w_0 - z_0$.

To see this, we consider the paths $\sigma, \tau: \mathbb{I} \rightarrow \mathbb{C} \setminus \mathbb{Z}^2$ defined as $\sigma(t) = \tilde{H}(0, t)$ and $\tau(t) = \tilde{H}(1, t)$ for $t \in \mathbb{I}$. Then

$$\begin{aligned} \overline{\sigma}(t) &:= \wp(\sigma(t)) = \wp(\tilde{H}(0, t)) = \overline{H}(0, t) \\ &= H_t(\beta(0)) = H_t(\beta(1)) = \overline{H}(1, t) = \wp(\tilde{H}(1, t)) = \wp(\tau(t)) \end{aligned} \tag{A.1}$$

for $t \in \mathbb{I}$. Note also that

$$\tau(0) = \tilde{H}(1, 0) = \alpha_0(1) = \alpha_0(0) + v_0 = \tilde{H}(0, 0) + v_0 = \sigma(0) + v_0.$$

This implies that the paths τ and $t \in \mathbb{I} \mapsto \sigma(t) + v_0$ have the same initial points. Since $v_0 \in 2\mathbb{Z}^2$, it follows from (2.3) and (A.1) that the map \wp sends them both to $\overline{\sigma} = \wp \circ \sigma = \wp \circ \tau$, which is a path in $\mathbb{P} \setminus V$. It follows from the uniqueness of lifts under \wp that $\tau(t) = \sigma(t) + v_0$ for all $t \in \mathbb{I}$.

This implies that

$$\alpha(1) - \alpha(0) = \tilde{H}(1, 1) - \tilde{H}(0, 1) = \tau(1) - \sigma(1) = v_0,$$

as desired.

Note that for a given parameterization β of γ , the difference $\alpha(1) - \alpha(0)$ is independent up to sign of the choice of the lift α of β . Indeed, suppose α' is another lift of β under \wp . Then

$$\wp(\alpha(0)) = \beta(0) = \wp(\alpha'(0)),$$

and so by (2.3), we have

$$\alpha'(0) = \pm\alpha(0) + m_0$$

for some (fixed) choice of the sign \pm and $m_0 \in 2\mathbb{Z}^2$. Then $t \in \mathbb{I} \mapsto \pm\alpha(t) + m_0$ is a lift of β with the same initial point as α' and so we see that

$$\alpha'(t) = \pm\alpha(t) + m_0$$

for all $t \in \mathbb{I}$. This implies that

$$\alpha'(1) - \alpha'(0) = \pm(\alpha(1) - \alpha(0)), \tag{A.2}$$

as desired.

It remains to show that up to sign $\alpha(1) - \alpha(0)$ is independent of the choice of the parameterization β of γ . For this, we consider another parameterization β' of γ as a simple loop. We first assume that $\beta'(0) = \beta(0)$. Then there exists a homeomorphism $h: \mathbb{I} \rightarrow \mathbb{I}$ such that $\beta' = \beta \circ h$. Here h fixes the endpoints 0 and 1 of \mathbb{I} or interchanges them depending on whether β' parameterizes γ with the same or opposite orientation as β , respectively. In any case, $\alpha' = \alpha \circ h$ is a lift of β' under \wp . It follows that

$$\alpha'(1) - \alpha'(0) = \alpha(h(1)) - \alpha(h(0)) = \pm(\alpha(1) - \alpha(0)),$$

as desired. As we now know, this relation is independent of the specific choice of the lift α' of β' .

Finally, we have to consider the case where β' has a possibly different initial point than β , say $p_0 := \beta'(0) \in \gamma$. By what we have seen, to establish (A.2), we can choose any parameterization β' of γ as a simple loop with $\beta'(0) = p_0$ and any lift α' of β' .

We can extend our parameterization β on \mathbb{I} periodically to a continuous map $\beta: \mathbb{R} \rightarrow \gamma$ such that $\beta(u + 1) = \beta(u)$ for all $u \in \mathbb{R}$. Then β lifts under \wp to a continuous map $\alpha: \mathbb{R} \rightarrow \mathbb{C}$ which agrees with the original lift α on \mathbb{I} . Then $u \in \mathbb{R} \mapsto \alpha(u + 1) - v_0$ is also a lift of β under \wp . The initial point of this lift corresponding to $u = 0$ is equal to $\alpha(0)$. The uniqueness of lifts implies that this lift and the original lift α are the same paths and so

$$\alpha(u + 1) = \alpha(u) + v_0 \tag{A.3}$$

for all $u \in \mathbb{R}$.

We can find $u_0 \in [0, 1)$ such that $\beta(u_0) = p_0$. Then $\beta': \mathbb{I} \rightarrow \gamma$ defined as $\beta'(u) = \beta(u_0 + u)$ for $u \in \mathbb{I}$ is a parameterization of γ as a simple loop with the initial point p_0 . Under \wp , this path β' has the lift $\alpha': \mathbb{I} \rightarrow \mathbb{C}$ given by $\alpha'(u) = \alpha(u + u_0)$ for $u \in \mathbb{I}$. Then equation (A.3) implies that

$$\alpha'(1) - \alpha'(0) = \alpha(u_0 + 1) - \alpha(u_0) = v_0,$$

as desired. The proof is complete. □

We are now almost ready to prove Lemma 2.3. Before we get to this, it is useful to discuss an alternative way to view our pillow \mathbb{P} .

We consider a slope $r/s \in \widehat{\mathbb{Q}}$. Then we can choose $p, q \in \mathbb{Z}$ such that $pr + qs = 1$ and define $\omega := s + ir$ and $\tilde{\omega} := -p + iq$. The numbers ω and $\tilde{\omega}$ form a basis of $\mathbb{C} \cong \mathbb{R}^2$ over \mathbb{R} , and so every point $z \in \mathbb{C}$ can be uniquely written in the form $z = u\tilde{\omega} + v\omega$ with $u, v \in \mathbb{R}$. Accordingly, the map $z = u\tilde{\omega} + v\omega \mapsto R(z) := u\tilde{\omega} - v\omega$ for $u, v \in \mathbb{R}$ is a well-defined ‘skew-reflection’ R on \mathbb{C} . Note also that $\mathbb{Z}^2 = \{n\tilde{\omega} + k\omega : n, k \in \mathbb{Z}\}$.

We consider the parallelogram $Q := \{u\tilde{\omega} + v\omega : u \in [0, 1], v \in [-1, 1]\} \subset \mathbb{C}$. Then it follows from (2.3) that $\wp(Q) = \mathbb{P}$. Moreover, for $z, w \in Q, z \neq w$, we have $\wp(z) = \wp(w)$ if and only if $z, w \in \partial Q$ and $w = R(z)$. Intuitively, this means that the pillow \mathbb{P} can be

obtained from Q by ‘folding’ Q in its middle segment $[0, \tilde{\omega}] \subset Q$ and identifying the points on ∂Q that correspond to each other under the skew-reflection R . The map \wp sends the set $\{0, \tilde{\omega}, \tilde{\omega} + \omega, \omega\}$ bijectively onto the set $\{A, B, C, D\}$ of vertices of \mathbb{P} (but not necessarily in that order).

From this geometric picture, it is clear that for each $t \in (0, 1)$, the set

$$\tau_{r/s} := \wp([t\tilde{\omega} - \omega, t\tilde{\omega} + \omega]) = \wp(\ell_{r/s}(t\tilde{\omega}))$$

is a simple closed geodesic in (\mathbb{P}, V) . Moreover, the sets

$$\xi_{r/s} := \wp([- \omega, + \omega]) = \wp(\ell_{r/s}(0)) \text{ and } \xi'_{r/s} := \wp([\tilde{\omega} - \omega, \tilde{\omega} + \omega]) = \wp(\ell_{r/s}(\tilde{\omega}))$$

are geodesic core arcs of $\tau_{r/s}$ lying in different components of $\mathbb{P} \setminus \tau_{r/s}$. In particular, $\tau_{r/s}$ is an essential Jordan curve in (\mathbb{P}, V) .

It follows from (2.3) that $\wp(\ell_{r/s}(t\tilde{\omega})) \cap \wp(\ell_{r/s}(t'\tilde{\omega})) \neq \emptyset$ for $t, t' \in \mathbb{R}$ if and only if $t' - t \in 2\mathbb{Z}$ or $t' + t \in 2\mathbb{Z}$. In this case, we have $\wp(\ell_{r/s}(t\tilde{\omega})) = \wp(\ell_{r/s}(t'\tilde{\omega}))$. Moreover, $\tau = \wp(\ell_{r/s}(t\tilde{\omega}))$ is a simple closed geodesic $\tau_{r/s}$ in (\mathbb{P}, V) if $t \in \mathbb{R} \setminus \mathbb{Z}$, it is equal to the geodesic arc $\xi_{r/s}$ if t is an even integer, and is equal to the geodesic arc $\xi'_{r/s}$ if t is an odd integer. Note that $\ell_{r/s}(t\tilde{\omega})$ for $t \in \mathbb{R}$ contains a point in $\mathbb{Z}^2 = \{n\tilde{\omega} + k\omega : n, k \in \mathbb{Z}\}$ if and only if $t \in \mathbb{Z}$.

LEMMA A.7. *Let $\tau_{r/s}$ and $\tau'_{r/s}$ be two distinct simple closed geodesics in (\mathbb{P}, V) with slope $r/s \in \widehat{\mathbb{Q}}$. Then $\tau_{r/s}$ and $\tau'_{r/s}$ are isotopic relative to V .*

Proof. The previous considerations imply that we may assume that the geodesics are represented in the form $\tau_{r/s} = \wp(\ell_{r/s}(t\tilde{\omega}))$ and $\tau'_{r/s} = \wp(\ell_{r/s}(t'\tilde{\omega}))$ with $t, t' \in (0, 1)$, $t \neq t'$. We may assume $t < t'$. Then

$$U := \wp(\{u\tilde{\omega} + v\omega : u \in (t, t'), v \in [-1, 1]\})$$

is an annulus contained in $\mathbb{P} \setminus V$ with $\partial U = \tau_{r/s} \cup \tau'_{r/s}$. It follows from Lemma 2.1 that $\tau_{r/s}$ and $\tau'_{r/s}$ are isotopic relative to V . □

Proof of Lemma 2.3. Let γ be an essential Jordan curve in (\mathbb{P}, V) . If we parameterize γ as a simple loop $\beta: \mathbb{I} \rightarrow \mathbb{P}$ and lift β to a path $\alpha: \mathbb{I} \rightarrow \mathbb{C}$ under \wp , then by Lemma A.6, we know that $\alpha(1) - \alpha(0)$ is a primitive element of $2\mathbb{Z}^2$ uniquely determined by $[\gamma]$ up to sign. Hence, we can find relatively prime integers $r, s \in \mathbb{Z}$ such that

$$\alpha(1) - \alpha(0) = 2(s + ir). \tag{A.4}$$

By switching signs here, which corresponds to parameterizing γ with opposite orientation, we may assume that $r \in \mathbb{Z}, s \in \mathbb{N}_0$, and that $r = 1$ if $s = 0$. Note that with these restrictions on r and s , the primitive element $2(s + ir)$ of $2\mathbb{Z}^2$ corresponds to the unique slope $r/s \in \widehat{\mathbb{Q}}$, and every slope in $\widehat{\mathbb{Q}}$ arises from a unique primitive element of $2\mathbb{Z}^2$ in this form.

As before, define $\omega := s + ir$ and $\tilde{\omega} := -p + iq$, where $p, q \in \mathbb{Z}$ are chosen so that $pr + qs = 1$. We know that $\xi := \wp(\ell_{r/s}(0))$ and $\xi' := \wp(\ell_{r/s}(\tilde{\omega}))$ are disjoint geodesic arcs in (\mathbb{P}, V) . The sets $\wp^{-1}(\xi)$ and $\wp^{-1}(\xi')$ consist of parallel lines with slope r/s , and these lines alternate in the following sense: each component of $\mathbb{C} \setminus \wp^{-1}(\xi \cup \xi')$ is an infinite strip whose boundary contains one line in $\wp^{-1}(\xi)$ and one line in $\wp^{-1}(\xi')$.

We may assume that γ is in minimal position with $\xi \cup \xi'$. Moreover, we may assume that the parameterization β of γ as a simple loop was chosen so that $\beta(0) = \beta(1) \notin \xi \cup \xi'$.

Then by Lemmas A.1 and A.3, we know that either $\xi \cap \gamma = \emptyset = \xi' \cap \gamma$, or the following conditions are true: γ meets ξ and ξ' transversely, the sets $\xi \cap \gamma$ and $\xi' \cap \gamma$ are non-empty and finite, and the points in these sets alternate on γ . We claim that the latter is not possible.

Otherwise, we choose a lift α of β under \wp . Then α intersects the lines in $\wp^{-1}(\xi \cup \xi')$ transversely, we have

$$k := \#(\alpha \cap \wp^{-1}(\xi)) = \#(\alpha \cap \wp^{-1}(\xi')) \in \mathbb{N},$$

and the points in $\alpha \cap \wp^{-1}(\xi) \neq \emptyset$ and $\alpha \cap \wp^{-1}(\xi') \neq \emptyset$ alternate on α . An argument very similar to the proof of Lemma A.4 then shows that $2k > 0$ is the number of lines in the set $\wp^{-1}(\xi \cup \xi')$ that separate $\alpha(0)$ from $\alpha(1)$. However, we know that $\alpha(1) - \alpha(0) = 2(s + ir)$, and so $\alpha(1)$ and $\alpha(0)$ lie on a line with slope r/s and are not separated by any line in $\wp^{-1}(\xi \cup \xi')$. This is a contradiction.

This shows that $\xi \cap \gamma = \emptyset = \xi' \cap \gamma$. Now consider a simple closed geodesic $\tau_{r/s} = \wp(\ell_{r/s}(t\tilde{\omega}))$ with slope r/s and $0 < t < 1$. Since $\xi \cap \gamma = \emptyset$, we can choose t very close to 0 so that $\tau_{r/s} \cap \gamma = \emptyset$. Now we can apply considerations very similar to the proof of Corollary 3.5(ii). The complement of $\tau_{r/s} \cup \gamma$ in \mathbb{P} is a disjoint union $\mathbb{P} \setminus (\tau_{r/s} \cup \gamma) = W \cup U \cup W'$, where $W, W' \subset \mathbb{P}$ are open Jordan regions and $U \subset \mathbb{P}$ is an annulus with $\partial U = \tau_{r/s} \cup \gamma$. Since $\tau_{r/s}$ and γ are essential, both W and W' must contain at least two points in V . Since $\#V = 4$, we have $U \cap V = \emptyset$. Lemma 2.1 now implies that $\tau_{r/s}$ and γ are isotopic relative to V . By Lemma A.7, the curve γ is actually isotopic to each closed geodesic $\tau_{r/s}$ with slope r/s .

The map $[\gamma] \mapsto r/s$ that sends each isotopy class $[\gamma]$ to a slope $r/s \in \widehat{\mathbb{Q}}$ obtained from a primitive element in $2\mathbb{Z}^2$ associated with $[\gamma]$ according to Lemma A.6 is well defined. It is clear that it is surjective, because the isotopy class $[\tau_{r/s}]$ of a geodesic $\tau_{r/s}$ with slope $r/s \in \widehat{\mathbb{Q}}$ is sent to r/s . To see that it is injective, suppose two isotopy classes $[\gamma]$ and $[\gamma']$ are sent to the same slope $r/s \in \widehat{\mathbb{Q}}$ by this map. If $\tau_{r/s}$ is a closed geodesic with slope r/s , then by our previous discussion, we have $\gamma \sim \tau_{r/s} \sim \gamma'$ relative to V . Hence, $[\gamma] = [\gamma']$. It follows that the map $[\gamma] \mapsto r/s$ is indeed a bijection. □

Remark A.8. Suppose γ is an essential Jordan curve in (\mathbb{P}, V) and its isotopy class $[\gamma]$ relative to V corresponds to slope $r/s \in \widehat{\mathbb{Q}}$ according to Lemma 2.3. Let $\beta: \mathbb{I} \rightarrow \gamma$ be a parameterization of γ as a simple loop, and $\alpha: \mathbb{I} \rightarrow \mathbb{C}$ be a lift of β under \wp . Then equation (A.4) in the proof of Lemma A.6 shows that we always have $\alpha(1) - \alpha(0) = \pm 2(s + ir)$.

A similar statement is true for arcs in (\mathbb{P}, V) .

COROLLARY A.9. *Let ξ be an arc in (\mathbb{P}, V) . Then ξ is isotopic relative to V to a geodesic arc $\xi_{r/s}$ for some slope $r/s \in \widehat{\mathbb{Q}}$. Moreover, if $\beta: \mathbb{I} \rightarrow \xi$ is a homeomorphic parameterization of ξ and α is any lift of β under \wp , then $\alpha(1) - \alpha(0) = \pm(s + ir)$.*

Proof. The arc ξ joins two of the points in V , while the other two points in V do not lie on ξ . Hence, we may ‘surround’ ξ by an essential Jordan curve γ in (\mathbb{P}, V) such that ξ is a

core arc of γ . By Lemma 2.3, we know that γ is isotopic relative to V to a closed geodesic $\tau_{r/s}$ in \mathbb{P} with some slope $r/s \in \widehat{\mathbb{Q}}$. Hence, ξ is isotopic relative to V to a core arc ξ' of $\tau_{r/s}$.

Now any two arcs in the interior of a closed topological disk D with the same endpoints are isotopic by an isotopy that fixes the endpoints of these arcs and the points in ∂D (see [Bus10, Theorem A.6(ii)]). This implies that any two core arcs of an essential Jordan curve in (\mathbb{P}, V) are isotopic relative to V if the core arcs have the same endpoints. We know that the closed geodesic $\tau_{r/s}$ has precisely two geodesic arcs with slope r/s as core arcs in different components of $\mathbb{P} \setminus \tau_{r/s}$. Therefore, ξ' , and hence also ξ , is isotopic relative to V to a geodesic arc $\xi_{r/s}$ with slope r/s . This proves the first part of the statement.

Let $\beta: \mathbb{I} \rightarrow \xi$ be a homeomorphic parameterization of ξ and $\alpha: \mathbb{I} \rightarrow \mathbb{C}$ be a lift of β under \wp . By what we have seen, we can choose an isotopy $H: \mathbb{P} \times \mathbb{I} \rightarrow \mathbb{P}$ relative to V with $H_0 = \text{id}_{\mathbb{P}}$ and $H_1(\xi) = \xi_{r/s}$. We use this to define a homotopy $\overline{H}: \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{P}$ by setting

$$\overline{H}(u, t) := H(\beta(u), t)$$

for $u, t \in \mathbb{I}$. Note that \overline{H}_t for $t \in \mathbb{I}$ gives a homeomorphic parameterization of an arc in (\mathbb{P}, V) . These arcs have all the same endpoints. In particular, $u \in \mathbb{I} \mapsto \beta'(u) := \overline{H}_1(u) = H_1(\beta(u))$ gives a homeomorphic parameterization of $H_1(\xi) = \xi_{r/s}$.

We can lift \overline{H}_t under \wp to find a homotopy $\tilde{H}: \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{C}$ such that $\tilde{H}_0 = \alpha$ and $\wp \circ \tilde{H}_t = \overline{H}_t$ for all $t \in \mathbb{I}$. To see this, one first applies the standard homotopy lifting theorem (see [Hat02, Proposition 1.30]) to the homotopy \overline{H} restricted $(0, 1) \times \mathbb{I}$ and the covering map $\wp: \mathbb{C} \setminus \mathbb{Z}^2 \rightarrow \mathbb{P} \setminus V \supset \overline{H}((0, 1) \times \mathbb{I})$ to obtain a unique homotopy $\tilde{H}: (0, 1) \times \mathbb{I} \rightarrow \mathbb{C}$ with $\tilde{H}_0 = \alpha|_{(0, 1)}$. Now as $u_0 \rightarrow 0^+$, the set $\overline{H}((0, u_0) \times \mathbb{I})$ shrinks to the point $\beta(0) \in V$, and so the connected set $\tilde{H}((0, u_0) \times \mathbb{I})$ shrinks to a unique point in $\wp^{-1}(V) = \mathbb{Z}^2$. This point can only be $\alpha(0)$. Hence, $\overline{H}(u, t) \rightarrow \alpha(0)$ uniformly for $t \in \mathbb{I}$ as $u \rightarrow 0^+$, and similarly $\overline{H}(u, t) \rightarrow \alpha(1)$ uniformly for $t \in \mathbb{I}$ as $u \rightarrow 1^-$. This implies that we can continuously extend \tilde{H} to a homotopy on $\mathbb{I} \times \mathbb{I}$ with the desired properties by setting $\tilde{H}(0, t) = \alpha(0)$ and $\tilde{H}(1, t) = \alpha(1)$ for $t \in \mathbb{I}$.

Then $\alpha' := \tilde{H}_1$ is a lift of β' , because $\wp \circ \alpha' = \wp \circ \tilde{H}_1 = \overline{H}_1 = \beta'$. Since β' is a homeomorphic parameterization of the geodesic arc $\xi_{r/s}$, the path α' sends \mathbb{I} homeomorphically onto a subsegment of a line $\ell_{r/s} \subset \mathbb{C}$. Since α' has its endpoints in \mathbb{Z}^2 and $\alpha'((0, 1))$ is disjoint from \mathbb{Z}^2 , this implies $\alpha'(1) - \alpha'(0) = \pm(s + ir)$. Since α and α' have the same endpoints, the statement follows. \square

Proof of Lemma 2.4. In the proof all isotopies, isotopy classes, intersection numbers, etc. are for isotopies on \mathbb{P} relative to V . We will use the facts about the geodesics on (\mathbb{P}, V) discussed before Lemma A.7 without further reference.

(i) Let α and β be essential Jordan curves in (\mathbb{P}, V) as in the statement. As before, we define $\omega = s + ir$ and $\tilde{\omega} = -p + iq$, where $p, q \in \mathbb{Z}$ and $pr + qs = 1$.

First suppose that $r/s = r'/s'$. Then in the isotopy class $[\alpha] = [\beta]$, we can find simple closed geodesics with slope r/s that are disjoint, for example, the curves $\tau_{r/s} = \wp(\ell_{r/s}(\tilde{\omega}/3))$ and $\tau'_{r/s} = \wp(\ell_{r/s}(2\tilde{\omega}/3))$. It follows that $i(\alpha, \beta) = 0$ in this case.

We now assume that $r/s \neq r'/s'$. To determine $i(\alpha, \beta)$, we have to find the minimum of all numbers $\#(\alpha \cap \beta)$, where α and β range over the given isotopy classes. By applying a suitable isotopy, we can reduce to the case where α is a fixed curve in its isotopy class and we only have to take variations over β . So by Lemma 2.3, we may assume that $\alpha = \tau_{r/s} = \wp(\ell_{r/s})$ is a simple closed geodesic as in the statement. Then $\tau_{r/s} = \wp(\ell_{r/s}(t_0\tilde{\omega}))$ for some $t_0 \in (0, 1)$. The preimage $\wp^{-1}(\tau_{r/s})$ of $\tau_{r/s}$ under \wp consists of the two disjoint families

$$\mathcal{F}_1 := \{\ell_{r/s}((t_0 + 2j)\tilde{\omega}) : j \in \mathbb{Z}\} \text{ and } \mathcal{F}_2 := \{\ell_{r/s}((-t_0 + 2j)\tilde{\omega}) : j \in \mathbb{Z}\} \tag{A.5}$$

of distinct lines with slope r/s .

Now let $\tilde{\beta}: \mathbb{I} \rightarrow \mathbb{C}$ be a lift of β under \wp , where we think of β as a simple closed loop in a parameterization with suitable orientation. Then if $z_0 := \tilde{\beta}(0)$ and $w_0 := \tilde{\beta}(1)$, we have $w_0 - z_0 = 2(s' + ir')$ as follows from Remark A.8. By changing the basepoint of β if necessary, we may assume that $\tilde{\beta}(0), \tilde{\beta}(1) \notin \wp^{-1}(\tau_{r/s})$. If $\omega' := s' + ir'$, then we can write ω' uniquely in the form

$$\omega' = k\omega + n\tilde{\omega}, \tag{A.6}$$

where $k, n \in \mathbb{Z}$. Note that then, $|n| = N := |rs' - sr'| > 0$ (to see this, multiply equation (A.6) by the complex conjugate of ω and take imaginary parts).

Now each family $\mathcal{F}_j, j = 1, 2$, consists of equally spaced parallel lines with slope r/s such that consecutive lines in each family differ by a translation by $2\tilde{\omega}$. This implies that the points z_0 and $w_0 = z_0 + 2\omega'$ are separated by precisely N lines from each of the families $\mathcal{F}_j, j = 1, 2$. So $\tilde{\beta}$ must have at least $2N$ points in common with $\wp^{-1}(\tau_{r/s})$. Since $\wp \circ \tilde{\beta}$ maps $[0, 1)$ injectively onto β , we conclude that $\beta = \wp(\tilde{\beta})$ has at least $2N$ points in common with $\tau_{r/s}$. If $\beta = \tau_{r'/s'}$, then $\tilde{\beta}$ is a parameterization of the line segment $[z_0, w_0]$, and so $\tilde{\beta}$ meets $\wp^{-1}(\tau_{r/s})$ in precisely $2N$ points. This means that $\tau_{r/s}$ and $\tau_{r'/s'}$ have exactly $2N$ points in common.

It follows that for all β , we have

$$2N \leq \#(\wp^{-1}(\tau_{r/s}) \cap \tilde{\beta}) = \#(\tau_{r/s} \cap \beta),$$

and so $2N \leq i(\alpha, \beta) \leq \#(\tau_{r/s} \cap \tau_{r'/s'}) = 2N$. Thus, we have equality here and the statement follows.

(ii) This is a variant of the argument in statement (i) and we use the same notation.

Since $\beta \sim \tau_{r'/s'}$, the core arc ξ of β is isotopic to a core arc of $\tau_{r'/s'}$. Now two core arcs of a given essential Jordan curve in (\mathbb{P}, V) are isotopic relative to V if they have the same endpoints (this was pointed out in the proof of Corollary A.9). It follows that $\xi \sim \xi'_{r'/s'}$, where $\xi'_{r'/s'}$ is one of the two geodesic core arcs of $\tau_{r'/s'}$. In particular, $\xi'_{r'/s'} = \wp(\ell_{r'/s'})$ for a line $\ell_{r'/s'} \subset \mathbb{C}$ that contains a point in \mathbb{Z}^2 . Note that $\xi'_{r'/s'}$ possibly differs from the geodesic arc $\xi_{r'/s'}$ as in the statement (if $\xi'_{r'/s'}$ and $\xi_{r'/s'}$ lie in different components of $\mathbb{P} \setminus \tau_{r'/s'}$).

If $\tilde{\xi}: \mathbb{I} \rightarrow \mathbb{C}$ is a lift of ξ under \wp in suitable orientation, and $z_1 := \tilde{\xi}(0) \in \mathbb{Z}^2, w_1 := \tilde{\xi}(1) \in \mathbb{Z}^2$, then it follows from Corollary A.9 that $w_1 - z_1 = \omega' = s' + ir'$. Now equation (A.6) implies that there are exactly $N = |n| = |rs' - sr'|$ lines in $\mathcal{F}_1 \cup \mathcal{F}_2$ that separate z_1

and w_1 (essentially, this follows from the fact that the set $[0, n] \cap \{2j \pm t_0 : j \in \mathbb{Z}\}$, where $t_0 \in (0, 1)$, contains precisely $N = |n|$ points).

Let $\tilde{w}' := -p' + iq'$, where $p', q' \in \mathbb{Z}$ and $p'r' + q's' = 1$. Then, by the discussion before Lemma A.7, the map \wp sends one of the segments $[z_1, w_1]$ and $[z_1 + \tilde{w}', w_1 + \tilde{w}']$ homeomorphically onto $\xi_{r'/s'}$ depending on whether $\xi'_{r'/s'} = \xi_{r'/s'}$ or $\xi'_{r'/s'} \neq \xi_{r'/s'}$, respectively. In either case, each segment meets exactly N lines in $\mathcal{F}_1 \cup \mathcal{F}_2$.

Arguing as before, we see that

$$\#(\tau_{r/s} \cap \xi_{r'/s'}) = N \leq \#(\wp^{-1}(\tau_{r/s}) \cap \tilde{\xi}) = \#(\tau_{r/s} \cap \xi).$$

This leads to $N \leq i(\alpha, \xi) \leq \#(\tau_{r/s} \cap \xi_{r'/s'}) = N$. So we have equality here and the statement follows.

(iii)–(v) These are special cases of statements (i) and (ii). For example, $a = \wp(\mathbb{R} \times \{0\}) = \wp(\ell_0(0))$ is a core arc of α^h corresponding to slope $r'/s' = 0/1 = 0$. Hence, by statement (ii), we have

$$i(\alpha, a) = |r \cdot 1 - s \cdot 0| = |r| = \#(\tau_{r/s} \cap \xi_0) = \#(\tau_{r/s} \cap a).$$

The other statements follow from similar considerations. □

REFERENCES

[BD17] L. Bartholdi and D. Dudko. Algorithmic aspects of branched coverings. *Ann. Fac. Sci. Toulouse Math.* (6) **26**(5) (2017), 1219–1296.

[BD18] L. Bartholdi and D. Dudko. Algorithmic aspects of branched coverings IV/V. Expanding maps. *Trans. Amer. Math. Soc.* **370**(11) (2018), 7679–7714.

[BLMW22] J. Belk, J. Lanier, D. Margalit and R. R. Winarski. Recognizing topological polynomials by lifting trees. *Duke Math. J.* **171**(17) (2022), 3401–3480.

[BM17] M. Bonk and D. Meyer. *Expanding Thurston Maps (Mathematical Surveys and Monographs, 225)*. American Mathematical Society, Providence, RI, 2017.

[BN06] L. Bartholdi and V. Nekrashevych. Thurston equivalence of topological polynomials. *Acta Math.* **197**(1) (2006), 1–51.

[Bol79] B. Bollobás. *Graph Theory (Graduate Texts in Mathematics, 63)*. Springer, New York–Berlin, 1979.

[Bur79] R. B. Burckel. *An Introduction to Classical Complex Analysis*. Vol. 1. Birkhäuser, Basel–Stuttgart, 1979.

[Bus10] P. Buser. *Geometry and Spectra of Compact Riemann Surfaces*. Birkhäuser, Boston, MA, 2010.

[CFKP03] J. W. Cannon, W. J. Floyd, R. Kenyon and W. R. Parry. Constructing rational maps from subdivision rules. *Conform. Geom. Dyn.* **7** (2003), 76–102.

[CFPP12] J. W. Cannon, W. J. Floyd, W. R. Parry and K. M. Pilgrim. Nearly Euclidean Thurston maps. *Conform. Geom. Dyn.* **16** (2012), 209–255.

[CPL16] G. Cui, W. Peng and Tan Lei. Renormalizations and wandering Jordan curves of rational maps. *Comm. Math. Phys.* **344**(1) (2016), 67–115.

[DH93] A. Douady and J. H. Hubbard. A proof of Thurston’s topological characterization of rational functions. *Acta Math.* **171**(2) (1993), 263–297.

[Die05] R. Diestel. *Graph Theory (Graduate Texts in Mathematics, 173)*, 3rd edn. Springer, Berlin, 2005.

[DMRS19] K. Drach, Y. Mikulich, J. Rückert and D. Schleicher. A combinatorial classification of postcritically fixed Newton maps. *Ergod. Th. & Dynam. Sys.* **39**(11) (2019), 2983–3014.

[FKK⁺17] W. Floyd, G. Kelsey, S. Koch, R. Lodge, W. Parry, K. M. Pilgrim and E. Saenz. Origami, affine maps, and complex dynamics. *Arnold Math. J.* **3**(3) (2017), 365–395.

[FLP12] A. Fathi, F. Laudenbach and V. Poénaru. *Thurston’s Work on Surfaces (Mathematical Notes, 48)*. Princeton University Press, Princeton, NJ, 2012.

[FM12] B. Farb and D. Margalit. *A Primer on Mapping Class Groups (Princeton Mathematical Series 49)*. Princeton University Press, Princeton, NJ, 2012.

- [FPP18] W. Floyd, W. Parry and K. M. Pilgrim. Rationality is decidable for nearly Euclidean Thurston maps. *Geom. Dedicata* **213** (2021), 487–512.
- [Hat02] A. Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2002.
- [Hea88] J. E. Head. The combinatorics of Newton’s method for cubic polynomials. *PhD Thesis*, Cornell University, 1988.
- [HM18] M. Hlushchanka and D. Meyer. Exponential growth of some iterated monodromy groups. *Proc. Lond. Math. Soc. (3)* **116**(6) (2018), 1489–1518.
- [Hlu19] M. Hlushchanka. Tischler graphs of critically fixed rational maps and their applications. *Preprint*, 2019, [arXiv:1904.04759](https://arxiv.org/abs/1904.04759).
- [Hub16] J. H. Hubbard. *Teichmüller Theory and Applications to Geometry, Topology, and Dynamics. Vol. 2*. Matrix Editions, Ithaca, NY, 2016.
- [KS94] L. Keen and C. Series. The Riley slice of Schottky space. *Proc. Lond. Math. Soc. (3)* **69**(1) (1994), 72–90.
- [KL19] G. Kelsey and R. Lodge. Quadratic Thurston maps with few postcritical points. *Geom. Dedicata* **201** (2019), 33–55.
- [KPS16] S. Koch, K. M. Pilgrim and N. Selinger. Pullback invariants of Thurston maps. *Trans. Amer. Math. Soc.* **368**(7) (2016), 4621–4655.
- [Lei92] Tan Lei. Matings of quadratic polynomials. *Ergod. Th. & Dynam. Sys.* **12**(3) (1992), 589–620.
- [Lei97] Tan Lei. Branched coverings and cubic Newton maps. *Fund. Math.* **192**(3) (1997), 207–260.
- [Lod13] R. Lodge. Boundary values of the Thurston pullback map. *Conform. Geom. Dyn.* **17** (2013), 77–118.
- [Mil06a] J. Milnor. *Dynamics in One Complex Variable (Annals of Mathematics Studies, 160)*, 3rd edn. Princeton University Press, Princeton, NJ, 2006.
- [Mil06b] J. Milnor. On Lattès maps. *Dynamics on the Riemann Sphere*. Eds. P. G. Hjorth and C. L. Petersen. European Mathematical Society, Zürich, 2006, pp. 9–43.
- [Ota01] J.-P. Otal. *The Hyperbolization Theorem for Fibered 3-Manifolds (SMF/AMS Texts and Monographs, 7)*. American Mathematical Society, Providence, RI, 2001.
- [Par18] W. Parry. NET map slope functions. *Preprint*, 2018, [arXiv:1811.01274](https://arxiv.org/abs/1811.01274).
- [Pil12] K. M. Pilgrim. An algebraic formulation of Thurston’s characterization of rational functions. *Ann. Fac. Sci. Toulouse Math. (6)* **21**(5) (2012), 1033–1068.
- [PL98] K. M. Pilgrim and Tan Lei. Combining rational maps and controlling obstructions. *Ergod. Th. & Dynam. Sys.* **18**(1) (1998), 221–245.
- [Pom92] C. Pommerenke. *Boundary Behaviour of Conformal Maps (Grundlehren der mathematischen Wissenschaften, 299)*. Springer, Berlin, 1992.
- [Ree92] M. Rees. A partial description of parameter space of rational maps of degree two. I. *Acta Math.* **168**(1–2) (1992), 11–87.
- [Sel12] N. Selinger. Thurston’s pullback map on the augmented Teichmüller space and applications. *Invent. Math.* **189**(1) (2012), 111–142.
- [Shi00] M. Shishikura. On a theorem of M. Rees for matings of polynomials. *The Mandelbrot Set, Theme and Variations (London Mathematical Society Lecture Note Series, 274)*. Ed. Tan Lei. Cambridge University Press, Cambridge, 2000, pp. 289–305.
- [SL00] M. Shishikura and Tan Lei. A family of cubic rational maps and matings of cubic polynomials. *Exp. Math.* **9**(1) (2000), 29–53.
- [Thu88] W. P. Thurston. On the geometry and dynamics of diffeomorphisms of surfaces. *Bull. Amer. Math. Soc. (N.S.)* **19**(2) (1988), 417–431.
- [Thu98] W. P. Thurston. Hyperbolic structures on 3-manifolds, II: Surface groups and 3-manifolds which fiber over the circle. *Preprint*, 1998, [arXiv:math/9801045](https://arxiv.org/abs/math/9801045).
- [Thu16] D. P. Thurston. From rubber bands to rational maps: a research report. *Res. Math. Sci.* **3** (2016), Paper no. 15.
- [Thu20] D. P. Thurston. A positive characterization of rational maps. *Ann. of Math. (2)* **192**(1) (2020), 1–46.