

Standardly embedded train tracks and pseudo-Anosov maps with minimum expansion factor

Eriko Hironaka and Chi Cheuk Tsang

Abstract. We show that given a fully punctured pseudo-Anosov map $f : S \rightarrow S$ whose punctures lie in at least two orbits under the action of f , the expansion factor $\lambda(f)$ satisfies the inequality

$$\lambda(f)^{|\chi(S)|} \geq \mu^4 \approx 6.85408,$$

where $\mu = \frac{1+\sqrt{5}}{2} \approx 1.61803$ is the golden ratio. The proof involves a study of standardly embedded train tracks, and the Thurston symplectic form defined on their weight space.

Contents

1. Introduction	1
2. Background	9
3. Standardly embedded train tracks	14
4. The Thurston symplectic form	24
5. Computing the radical	28
6. Proof of the main theorem	38
7. Sharpness of the main theorem	40
8. Discussion and further questions	50
References	53

1. Introduction

Let $S = S_{g,s}$ be an oriented surface of finite type with genus g and s punctures, such that $\chi(S) = 2 - 2g - s$ is negative, and let $f : S \rightarrow S$ be a homeomorphism of S . By the Nielsen–Thurston classification of mapping classes, up to isotopy, f either preserves an essential multicurve (and is *periodic* or *reducible*), or it is *pseudo-Anosov*. In the latter case, there exists a transverse pair of measured, singular *stable* and *unstable* foliations ℓ^s, ℓ^u such that f stretches the measure of ℓ^u by $\lambda(f)$ and contracts the measure of ℓ^s by $\frac{1}{\lambda(f)}$. The number $\lambda(f) > 1$ here is known as the *expansion factor* of f .

Mathematics Subject Classification 2020: 37D40 (primary); 37E30, 57M50 (secondary).

Keywords: pseudo-Anosov maps, expansion factor, fully punctured, train tracks, standardly embedded, Thurston symplectic form.

Thurston’s train track theory links the dynamics of pseudo-Anosov maps with that of Perron–Frobenius matrices and implies that $\lambda(f)$ is a bi-Perron algebraic unit with degree bounded in terms of $\chi(S)$ [14]. In particular, for fixed (g, s) , $\lambda(f)$ attains a minimum $\lambda_{g,s} > 1$. Furthermore, each pseudo-Anosov map determines a closed geodesic on the moduli space $\mathcal{M}(S)$ with respect to the Teichmüller metric, whose length equals $\log(\lambda(f))$. Thus, $\log(\lambda_{g,s})$ is the length of the shortest geodesics on $\mathcal{M}(S_{g,s})$ [2]. However, so far $\lambda_{g,s}$ is known only for small g and s (see [10, 19, 26]).

In this paper, we focus on the normalized expansion factor

$$L(S, f) = \lambda(f)^{|\chi(S)|}$$

and consider pseudo-Anosov maps that are *fully punctured*, meaning the singularities of ℓ^s and ℓ^u lie on the punctures of S . Let μ be the golden ratio $\frac{1+\sqrt{5}}{2} \approx 1.61803$. Our main result is the following.

Theorem 1.1. *Let $f : S \rightarrow S$ be a fully punctured pseudo-Anosov map with at least two puncture orbits, then $L(S, f)$ satisfies the sharp inequality*

$$L(S, f) \geq \mu^4 \approx 6.85408.$$

Here by a puncture orbit, we mean an orbit of the action of f on the punctures of S . If we replace the punctures of S with boundary components, then the number of puncture orbits equals the number of boundary components of the mapping torus of f . We remark that the assumption of f having at least two puncture orbits is necessary. See Section 7.4 for explicit counterexamples.

For the rest of this introduction, we will first give some background on what is known about the pattern of minimum expansion factors and its relations to the topology and geometry of the mapping torus. Then we will state a more precise version of Theorem 1.1, discuss some special cases, and finally explain the tools that are used in the proof.

1.1. Background on small expansion factors and mapping tori

The expansion factor is a measure of the dynamical complexity of a pseudo-Anosov map. More concretely, it is equal to the exponential of the entropy, which measures how much the map mixes points on the surface. Thus, the minimum expansion factor is related to the complexity of the underlying surface and of the mapping torus, with the latter being a fibered hyperbolic 3-manifold [38]. In the following, we describe what is known about this relation.

1.1.1. Minimum expansion factor as a function on the (g, s) plane. Applying properties of Perron–Frobenius matrices, Penner [32] observed that $\log \lambda_{g,s}$ is bounded below by the reciprocal of a linear function in g and s

$$\log \lambda_{g,s} \geq \frac{\log 2}{12g - 12 + 4s}.$$

Furthermore, for $s = 0$, Penner constructed a sequence of examples giving an upper bound, yielding

$$\log \lambda_{g,0} \leq \frac{C}{g}$$

for some $C > 0$, which together with the lower bound implies

$$\log \lambda_{g,0} \asymp \frac{1}{|\chi(S_{g,0})|}.$$

This asymptotic behavior has been shown to persist along other lines in the (g, s) -plane, namely, $s = mg$, where $m > 0$ [40]; $g = 0$ [21]; $g = 1$ [39]; and $s = s_0$ for fixed $s_0 > 0$ [41]. For lines $g = g_0$ with fixed $g_0 \geq 2$, however, [39] shows that

$$\log \lambda_{g,s} \asymp \frac{\log |\chi(S_{g,s})|}{|\chi(S_{g,s})|}.$$

We also mention a result of Agol, Leininger, and Margalit [4] that concerns the minimum expansion factor among pseudo-Anosov maps on an oriented closed surface with fixed genus g whose action on the first homology preserves a k -dimensional subspace. We denote this minimum expansion factor as $\lambda_{g,0,k}$. Their result is that

$$\log \lambda_{g,0,k} \asymp \frac{k+1}{g}.$$

See [6] for a partial generalization of this result to punctured surfaces.

Let λ_K be the minimum expansion factor for pseudo-Anosov maps $f : S \rightarrow S$ satisfying $|\chi(S)| = K$.

Question 1.2 (McMullen [29]). Does the sequence $(\lambda_K)^K$ converge, and if so does it converge to μ^4 ?

As we will see in this paper, it is possible to use the theory of digraphs and standardly embedded train tracks to get information about minimum expansion factors in the fully punctured case. To translate results from the fully punctured case to the general case, the following question is of interest.

Question 1.3. For fixed K , where, on the level sets in the (g, s) -plane where $2g + s - 2 = K$, is the minimum λ_K realized? Is there a bound s_0 , such that λ_K is always achieved by $\lambda_{g,s}$ for $s \leq s_0$?

A positive answer to Question 1.3 would suggest that the minimum expansion factors λ_K° for fully punctured pseudo-Anosov maps $f : S \rightarrow S$ satisfying $|\chi(S)| = K$ should have the same asymptotic behavior as λ_K .

We note that Theorem 1.9 answers Question 1.3 in the affirmative if we restrict to fully punctured pseudo-Anosov mapping classes with even Euler characteristic and at least two puncture orbits.

1.1.2. Mapping tori. Pseudo-Anosov mapping classes (S, f) can be partitioned into flow-equivalence classes, that is, collections of pseudo-Anosov mapping classes whose mapping tori and suspension flows are equivalent. Thurston's fibered face theory gives a way to parameterize elements of a flow-equivalence class by primitive integral points on a cone over a top-dimensional face F , called a *fibered face*, of the Thurston norm ball in $H^1(M; \mathbb{R})$. Here, the Thurston norm of a primitive integral element $a \in \text{cone}(F)$ associated with (S_a, f_a) equals $|\chi(S_a)|$, and hence its projection to F is given by $\bar{a} = \frac{1}{|\chi(S_a)|}a$ (see [37]).

It follows that the rational elements on F are in one-to-one correspondence with elements of the flow-equivalence class. By a result of Fried, the normalized expansion factor $L(S_a, f_a)$, considered as a function of integral classes a , extends to a continuous concave function defined for all $a \in F$ that goes to infinity towards the boundary of F [15]. Thus, in particular, for any compact subset of F , the corresponding pseudo-Anosov maps have bounded normalized expansion factor.

For a hyperbolic 3-manifold M with fibered face F , let $C_{M,F}$ be the infimum of $L(S, f)$, where (S, f) is associated with an element of $\text{cone}(F)$. One consequence of Fried and Thurston's fibered face theory is the following.

Theorem 1.4 (Fried–Thurston [15, 16, 37]). *If $b_1(M) \geq 2$, then for any $\varepsilon > 0$ there are infinitely many fibrations $M \rightarrow S^1$ whose monodromy (S, f) satisfies*

$$L(S, f) < C_{M,F} + \varepsilon.$$

Given $C > 1$, the pseudo-Anosov mapping classes (S, f) that satisfy $L(S, f) < C$ must have mapping tori M and associated fibered face F with $C_{M,F} \leq C$. The following result is often referred to as the universal finiteness property and states that after removing the singular orbits of the suspension pseudo-Anosov flow, the number of such pairs (M, F) is finite.

Theorem 1.5 (Farb–Leininger–Margalit [12]). *Given any C , there is a finite set of fibered, hyperbolic 3-manifolds Ω_C such that if (S, f) is a fully punctured pseudo-Anosov map, and $L(S, f) < C$, then the mapping torus of (S, f) lies in Ω_C (see also [3] and [5]).*

The above theorem also implies finiteness for families of defining polynomials that can arise for small expansion factors. Given a polynomial $p(t)$, let $|p|$ be the complex norm of the largest root of $p(t)$.

Theorem 1.6 (McMullen [29]). *Let M be a fibered hyperbolic 3-manifold, with $n = b_1(M) \geq 2$ and let $F \subset H^1(M; \mathbb{R})$ be a fibered face. Then there is a polynomial $\Theta \in \mathbb{Z}[t_1, \dots, t_n]$ with the property that for any primitive integral $a = (a_1, \dots, a_n)$ in the fibered cone over F , $|\Theta(t^{a_1}, \dots, t^{a_n})|$ is the expansion factor of the monodromy of (S, f) corresponding to a .*

The polynomial Θ is known as the *Teichmüller polynomial* of the fibered face.

Define C_M to be the infimum of $L(S, f)$, where (S, f) is the monodromy of a fibration $M \rightarrow S^1$. The topological invariant C_M of a fibered hyperbolic 3-manifold M is related to the geometry of M as shown by the following theorem.

Theorem 1.7 (Kojima–McShane [25]). *For any hyperbolic, fibered 3-manifold M , C_M satisfies the following inequality*

$$C_M \geq \exp\left(\frac{\text{vol}(M)}{3\pi}\right).$$

Remark 1.8. Theorem 1.7 implies, for example, that if (S, f) satisfies $L(S, f) = \mu^4$, the volume of the mapping torus M must satisfy

$$\text{vol}(M) \leq 12\pi \log(\mu) \approx 18.14123.$$

For comparison, the magic manifold, which realizes many of the smallest known expansion factors, has volume ≈ 5.33349 according to SnapPy [11]. Further work comparing known pseudo-Anosov maps with small expansion factor and the volume of their mapping tori can be found in [1, 23].

1.2. Refinement of the main theorem and some special cases

For positive integers a, b with $a < b$, define

$$LT_{a,b}(t) = t^{2b} - t^{b+a} - t^b - t^{b-a} + 1.$$

This family of polynomials was first noticed by Lanneau and Thiffeault [27] to play a role in the study of minimum expansion factors.

The more precise version of our main theorem is as follows.

Theorem 1.9. *Let $f : S \rightarrow S$ be a fully punctured pseudo-Anosov map with at least two puncture orbits where $|\chi(S)| = K \geq 3$. Then $\lambda(f)$ satisfies the inequality*

$$\lambda(f) \geq |LT_{1, \frac{K}{2}}|,$$

for K even, and

$$\lambda(f)^K \geq 8,$$

for K odd.

Moreover, for each even K , equality is achieved by a fully punctured pseudo-Anosov map with $|\chi(S)| = K$ and $s \leq 4$.

Sharpness in Theorem 1.9 follows from computations in [1, 20, 23].

1.2.1. Orientable pseudo-Anosov mapping classes. A pseudo-Anosov mapping class (S, f) is called *orientable* if its stable and unstable foliations are orientable. These have the property that the expansion factor of f is equal to the largest eigenvalue of its action on the integral homology of S .

Lanneau and Thiffeault [27] asked whether the minimum expansion factor λ_g^+ for even genus g orientable pseudo-Anosov mapping classes is given by

$$\lambda_g^+ = |LT_{1,g}|$$

and showed that $|LT_{1,g}|$ is a lower bound for λ_g^+ for $g = 2, 4, 8, 10$. For the case when $g \geq 2$ is even and not divisible by 3, $|LT_{1,g}|$ is an upper bound for λ_g^+ [20].

Theorem 1.10. *Let (S, f) be an orientable pseudo-Anosov mapping class on a genus $g \geq 2$ closed surface with exactly two singularities, each of which is fixed by f . Then $\lambda(f)$ satisfies the inequality*

$$\lambda(f) \geq |LT_{1,g}|$$

and this inequality is sharp when g is even and not divisible by 3.

1.2.2. Braid monodromies. Consider (S, f) where $S = S_{0,n+1}$, and one puncture p_∞ is fixed by f . A mapping class of this type is called a *braid monodromy on n strands*. Given a standard braid β on n strands, there is an associated braid monodromy, and vice versa (see, e.g., [8]). For convenience, we use the standard Artin braid group generators to specify a braid and its associated monodromy.

Theorem 1.11. *Let β be a fully punctured pseudo-Anosov braid monodromy on n strands. Then $\lambda(\beta)$ satisfies the inequality*

$$\lambda(\beta)^{n-1} \geq \mu^4.$$

More precisely, for $n \geq 4$, $\lambda(\beta)$ satisfies the inequality

$$\lambda(\beta) \geq |LT_{1, \frac{n-1}{2}}|$$

for n odd, and

$$\lambda(\beta)^{n-1} \geq 8$$

for n even.

The 3-strand braid $\beta = \sigma_1 \sigma_2^{-1}$ satisfies $\lambda(\beta) = \mu^2$ (see [20]), while the 5-strand braid $\beta = \sigma_1 \sigma_2 \sigma_3 \sigma_4 \sigma_1 \sigma_2$ satisfies $\lambda(\beta) = |LT_{1,2}|$ (see [19]). Since both of these braid monodromies are fully punctured, this shows that Theorem 1.11 is sharp. For larger number of strands, however, there are no known examples achieving the lower bound. Instead, the smallest known expansion factors of braid monodromies follow the pattern found in [21] (see also [23]). This invites the following question.

Question 1.12. Is it true that aside from $\sigma_1\sigma_2^{-1}$ and $\sigma_1\sigma_2\sigma_3\sigma_4\sigma_1\sigma_2$, the expansion factor of a fully punctured pseudo-Anosov braid monodromy on n strands is strictly larger than $|LT_{1, \frac{n-1}{2}}|$?

A positive answer to Question 1.3 would suggest a positive answer to Question 1.12 as well. This is because the former would imply that the smallest values of $\lambda_{g,s}$ are concentrated in a vertical band $s \leq s_0$ on the (g, s) -plane, whereas the braid monodromies concern the horizontal line $g = 0$ which moves away from this vertical band.

We remark that the answer to Question 1.12 is “no” if the word “fully punctured” is dropped: It has been pointed out to us by Eiko Kin that the 5-strand braid $\beta = (\sigma_1\sigma_2\sigma_3\sigma_4)^2\sigma_4\sigma_3$ satisfies $\lambda(\beta) = |LT_{1,2}|$ as well (but it is not fully punctured).

1.3. Techniques used in the proof

Our proof of Theorem 1.1 involves a study of properties of Perron–Frobenius-directed graphs and of standardly embedded train tracks.

1.3.1. Perron–Frobenius digraphs. To any non-negative square integer matrix $M \geq 0$, one can associate a directed graph, or *digraph*, Γ where each vertex corresponds to a row (or column) of M , and the number of directed edges counted with multiplicity from one vertex to another equals the corresponding (row, column) entry of M . The matrix M is *Perron–Frobenius* if it has the property that $M^n > 0$ for large enough n . This translates to the condition that Γ is *strongly connected*, that is, any ordered pair of vertices can be connected by a directed edge path, and *aperiodic*, that is, the greatest common divisor of the lengths of directed cycles is 1.

The characteristic polynomial θ_Γ of Γ is defined to be the characteristic polynomial of M . We say that a polynomial is *reciprocal* if the set of roots (counted with multiplicities) is closed under inverses. This is equivalent to the polynomial being palindromic, that is, the list of its coefficients forms a palindrome, up to a sign.

Theorem 1.13 (McMullen [30]). *Let Γ be a Perron–Frobenius digraph of rank $K \geq 3$ with reciprocal characteristic polynomial θ_Γ . Then*

$$|\theta_\Gamma| \geq |LT_{1, \frac{K}{2}}|$$

for K even, and

$$|\theta_\Gamma|^K \geq 8$$

for K odd.

1.3.2. Expansion factors and train tracks. Recall that to compute expansion factors of pseudo-Anosov maps $f : S \rightarrow S$ one uses train tracks to translate the geometric dynamics of f to combinatorial information. In their most general form train tracks are embedded graphs on surfaces with smoothings at vertices indicating directions which paths are

allowed to take. A foliation on a surface is said to be *carried* by a train track τ if any leaf can be isotoped to a smooth path on τ . Thurston showed that if $f : S \rightarrow S$ is pseudo-Anosov, then there is a train track τ embedded on S which carries the unstable foliation of f , and for which the map f determines a graph map on τ which sends vertices to vertices and edges to edge paths. Thus, f naturally induces a linear map

$$f_* : \mathbb{R}^{\mathcal{E}} \rightarrow \mathbb{R}^{\mathcal{E}},$$

where $\mathcal{E}$ is the set of edges of τ and $\mathbb{R}^{\mathcal{E}}$ can be thought of as choices of weights defined on the edges. Under an appropriate choice of τ , f_* is Perron–Frobenius and the spectral radius of f_* is the expansion factor of f (see, e.g., [14]).

One might be tempted to apply Theorem 1.13 to the digraph associated with this f_* . However, there are some problems with this. Firstly, notice that one also needs to show that f_* has reciprocal characteristic polynomial. A natural idea for doing so is to make use of the *Thurston symplectic form* ω . However, ω is actually degenerate on $\mathbb{R}^{\mathcal{E}}$, so reciprocity is not immediate here. Even if one can show reciprocity, the dimension of $\mathbb{R}^{\mathcal{E}}$ is rather large, so the bound obtained from Theorem 1.13 would not be very good.

One can instead restrict f_* to the subspace of *allowable weights* $\mathcal{W} \subset \mathbb{R}^{\mathcal{E}}$, consisting of weights that satisfy what are called the branching conditions. This is the minimal condition required so that any simple closed loop which is carried by τ will define a positive vector in $\mathcal{W}$. This would reduce the dimension and remove a large part of the degeneracy of ω , but it is now unclear whether the action of f_* on $\mathcal{W}$ is Perron–Frobenius or not. Indeed, there is not even a natural basis of $\mathcal{W}$ that we can write the action of f_* as a matrix in, making it difficult to apply digraph techniques directly.

These problems are avoided by considering only pseudo-Anosov maps that are carried by a special kind of train track based on the Bestvina–Handel algorithm [7, 9], which we define next.

1.3.3. Standardly embedded train tracks.

Definition 1.14. Let $f : S \rightarrow S$ be a fully punctured pseudo-Anosov map. A train track τ that carries f is *standardly embedded* if its edges can be partitioned into two types: $\mathcal{E}_{\text{real}}$, the set of *real edges*, and $\mathcal{E}_{\text{inf}}$, the set of *infinitesimal edges*, such that:

- (1) Each connected component of the complement of τ is homeomorphic to a once-punctured disc.
- (2) The action of f permutes the set of infinitesimal edges.
- (3) The linear map on $\mathbb{R}^{\mathcal{E}_{\text{real}}}$ induced by f

$$f_*^{\text{real}} : \mathbb{R}^{\mathcal{E}_{\text{real}}} \rightarrow \mathbb{R}^{\mathcal{E}_{\text{real}}}$$

is Perron–Frobenius with respect to the standard basis.

For computations and properties of standardly embedded train tracks, see, for example, [13, 21, 28, 34].

It follows that given a pseudo-Anosov map f with a standardly embedded train track τ that carries it, the expansion factor $\lambda(f)$ is the spectral radius of the restriction f_*^{real} of f_* to $\mathbb{R}^{\mathcal{E}_{\text{real}}}$. Furthermore, the degree of the characteristic polynomial of f_*^{real} equals $|\chi(S)|$.

Our main technical result is the following.

Theorem 1.15. *If $f : S \rightarrow S$ is a fully punctured pseudo-Anosov map with at least two puncture orbits, then there exists a standardly embedded train track τ such that*

- (i) τ carries f , and
- (ii) the characteristic polynomial of f_*^{real} is reciprocal.

The proof of (ii) uses an explicit description of the radical of the Thurston symplectic form on the weight space of τ . We refer to Proposition 5.5 for the technical statement used in our proof. Later in Proposition 8.1 we give an extension that applies to general train tracks carrying a pseudo-Anosov map.

Our main theorem is then obtained by applying Theorem 1.13 to the digraph associated with f_*^{real} for such a standardly embedded train track.

1.4. Organization of paper

We set up some basic definitions in Section 2. In Section 3, we define standardly embedded train tracks and prove Theorem 1.15. We note that unlike in the introduction, we will define standardly embedded train tracks as a class of objects separate from surfaces and pseudo-Anosov maps. This will make the discussion easier in Sections 4 and 5, where we prove that the characteristic polynomial is reciprocal by studying the Thurston symplectic form.

We prove our main theorem in Section 6. In Section 7, we explore the sharpness of the main theorem, in particular demonstrating examples that show the last statement of Theorem 1.9. Finally, in Section 8, we discuss some questions and future directions.

2. Background

2.1. Pseudo-Anosov maps

We recall the definition of a pseudo-Anosov map. More details can be found in [14].

Definition 2.1. A *finite-type surface* is an oriented closed surface with finitely many points, which we call the *punctures*, removed.

A homeomorphism f on a finite-type surface S is said to be *pseudo-Anosov* if there exists a pair of singular measured foliations (ℓ^s, μ^s) and (ℓ^u, μ^u) such that:

- (1) Away from a finite collection of *singular points*, which includes the punctures, ℓ^s and ℓ^u are locally conjugate to the foliations of $\mathbb{R}^2$ by vertical and horizontal lines, respectively.

- (2) Near a singular point, ℓ^s and ℓ^u are locally conjugate to either
- the pull back of the foliations of $\mathbb{R}^2$ by vertical and horizontal lines by the map $z \mapsto z^{\frac{n}{2}}$, respectively, for some $n \geq 3$, or
 - the pull back of the foliations of $\mathbb{R}^2 \setminus \{(0, 0)\}$ by vertical and horizontal lines by the map $z \mapsto z^{\frac{n}{2}}$, respectively, for some $n \geq 1$.

In this case, we say that the singular point is *n-pronged*.

- (3) $f_*(\ell^s, \mu^s) = (\ell^s, \lambda^{-1}\mu^s)$ and $f_*(\ell^u, \mu^u) = (\ell^u, \lambda\mu^u)$ for some $\lambda = \lambda(f) > 1$.

We call (ℓ^s, μ^s) and (ℓ^u, μ^u) the *stable* and *unstable* measured foliations, respectively. We call $\lambda(f)$ the *expansion factor* of f and call $\lambda(f)^{|\chi(S)|}$ the *normalized expansion factor* of f .

Definition 2.2. f is said to be *fully punctured* if the set of singular points equals the set of punctures. In this case, we will denote the set of punctures by $\mathcal{X}$. Note that f acts on $\mathcal{X}$ by some permutation, hence it makes sense to talk about the orbits of punctures under the action of f , or *puncture orbits* for short.

We recall two standard facts about pseudo-Anosov maps which will be used in Section 3.

Definition 2.3. A *half-leaf* of ℓ^s is a properly embedded copy of $[0, \infty)$ in a leaf of ℓ^s . A *half-leaf* of ℓ^u is similarly defined.

Proposition 2.4. *We have the following properties of f , ℓ^s , and ℓ^u .*

- *The set of periodic points of f is dense.*
- *Every half-leaf of ℓ^s and of ℓ^u is dense.*

Proof. This is Proposition 9.20 and Proposition 9.6 of [14], respectively. ■

2.2. Train tracks

We define train tracks and set up some terminology.

Definition 2.5. A (*ribbon*) *train track* τ is a finite graph endowed with two additional pieces of data:

- (1) A cyclic ordering on the set $\mathcal{E}_v$ of half-edges incident to each vertex v .
- (2) A partition $\mathcal{E}_v = \mathcal{E}_v^1 \sqcup \mathcal{E}_v^2$ into two nonempty cyclically consecutive subsets, which we refer to as the *smoothing* at v .

We denote the set of vertices of τ by $\mathcal{V}(\tau)$ and denote the set of edges of τ by $\mathcal{E}(\tau)$. We will sometimes refer to the vertices of τ as the *switches*.

We will think of the two pieces of additional data on each $\mathcal{E}_v$ as represented by having a small oriented disc at v with the half-edges being arranged around v as determined by

the cyclic order, and with a tangent line at v such that the half-edges in $\mathcal{E}_v^1$ are tangent to one side of the tangent line and those in $\mathcal{E}_v^2$ tangent to the other side (see Figure 1, left).

In fact, one can form an oriented surface N by taking the union of these small oriented discs and a rectangular strip connecting the discs, respecting their orientations, along each edge. N can be foliated by properly embedded intervals transverse to the edges of τ . We call N the *tie neighborhood* of τ and call the intervals foliating N the *ties* (see Figure 1, right).

We refer to a pair of adjacent elements in some $\mathcal{E}_v^\beta$ as a *cusps* at v . Equivalently, consider the complementary regions of τ in its tie neighborhood. The cusps are in one-to-one correspondence with the nonsmooth points on the boundary that meet τ .

In general, let e and e' be two elements in $\mathcal{E}_v^\beta$ for some switch v and some $\beta \in \mathbb{Z}/2$. e is said to *lie on the left of* e' if for some (hence any) $f \in \mathcal{E}_v^{\beta+1}$, (e, e', f) is oriented counterclockwise in $\mathcal{E}_v$. We show one example of this in Figure 1 (left).

We make two remarks regarding this definition.

Firstly, train tracks are frequently defined to be subsets of surfaces in the literature. However, it will be to our advantage to define them independently of a fixed surface. To connect with the usual definition, we have to consider embeddings of the tie neighborhoods into surfaces. The following definition illustrates an example of this.

Definition 2.6. Let l be a lamination on finite-type surface S . Let τ be a train track with tie neighborhood N , and let $\iota : \tau \hookrightarrow S$ be an embedding. $\iota(\tau)$ is said to *carry* l if ι can be extended into an orientation preserving embedding $\iota : N \hookrightarrow S$ such that the leaves of l are contained in $\iota(N)$ and are transverse to the ties. $\iota(\tau)$ is said to *fully carry* l if in addition every tie intersects some leaf of l .

Secondly, the underlying graph of a train track is sometimes required to be trivalent. Here, we allow our train tracks to have arbitrary valence since this is the natural setting for discussing standardly embedded train tracks.

We then define maps of train tracks.

Definition 2.7. An edge path or cycle c of τ is said to be *carried by* τ if at each vertex v on c , the incoming and outgoing edges do not lie in the same $\mathcal{E}_v^\beta$.




Figure 1. A train track and its tie neighborhood near a switch.

A map $f : \tau \rightarrow \tau'$ is said to be a *train track map* if f sends switches of τ to switches of τ' and edges of τ to edge paths carried by τ' , such that the induced map $D_v f : \mathcal{E}_v \rightarrow \mathcal{E}_{f(v)}$ of half-edges at each switch preserves the cyclic ordering and sends elements in distinct $\mathcal{E}_v^\beta$ to distinct $\mathcal{E}_{f(v)}^{\beta'}$.

Suppose $f : \tau \rightarrow \tau'$ is a train track map. Let N and N' be the tie neighborhoods of τ and τ' , respectively. Then, f induces an embedding $N \hookrightarrow N'$ which sends ties in N into ties in N' (see Figure 2). For convenience of notation, we will denote this induced embedding by f as well.

Definition 2.8. Let $f : \tau \rightarrow \tau'$ be a train track map. The *transition matrix* of f is the matrix $f_* \in M_{\mathcal{E}(\tau') \times \mathcal{E}(\tau)}(\mathbb{Z})$ whose entries are defined by

$$(f_*)_{e',e} = \# \text{ times } f(e) \text{ passes through } e'.$$

2.3. Reciprocal Perron–Frobenius matrices

We recall some terminology regarding matrices.

Definition 2.9. A matrix B is said to be *positive* if all of its entries are positive.

Let $A \in M_{n \times n}(\mathbb{Z}^{\geq 0})$ be a matrix with non-negative integer entries. A is said to be *Perron–Frobenius* if there exists $k \geq 1$ such that A^k is positive.

Perron–Frobenius matrices arise naturally in the study of pseudo-Anosov dynamics and are frequently applied with the following classical theorem.

Theorem 2.10 (Perron–Frobenius theorem). *Let A be a Perron–Frobenius matrix. Let λ be the largest real eigenvalue of A . Then λ is equal to the spectral radius of A and there exists a positive λ -eigenvector v of A . Moreover, up to scalar multiplication by a positive number, v is the unique positive eigenvector of A .*

Another property of matrices which we will use in this paper is reciprocity.




Figure 2. A train track map $\tau \rightarrow \tau'$ induces an embedding of tie neighborhoods $N \hookrightarrow N'$.

Definition 2.11. A linear map $T : V \rightarrow V$ is said to be *reciprocal* if its eigenvalues (taken with multiplicity) are invariant under $\mu \mapsto \mu^{-1}$. Equivalently, T is reciprocal if and only if its characteristic polynomial $p(t)$ is reciprocal, that is, it satisfies $p(t) = \pm t^n p(t^{-1})$ where $n = \dim V$.

We remark that even though we will freely regard matrices as linear transformations on Euclidean spaces using the standard bases, the reader should observe that reciprocity is a property of *linear transformations*, whereas being Perron–Frobenius is strictly a property of *matrices*.

For matrices that are both Perron–Frobenius and reciprocal, McMullen proved the following sharp inequality regarding their spectral radius.

Theorem 2.12 ([30]). *Let $A \in M_{n \times n}(\mathbb{Z}^{\geq 0})$, $n \geq 2$, be a reciprocal Perron–Frobenius matrix with spectral radius $\rho(A)$. Then*

$$\rho(A)^n \geq \mu^4.$$

More precisely, for $n \geq 3$, we have

$$\rho(A) \geq \left| LT_{1, \frac{n}{2}} \right|$$

if n is even, and

$$\rho(A)^n \geq 8$$

if n is odd.

Proof. In [30], the first statement is stated as Theorem 7.1. The second statement is stated as Theorem 1.1 for n even and stated in P.33 under the proof of Theorem 7.1 for n odd. ■

To conclude this subsection, we state some elementary facts that will help us establish reciprocity.

Proposition 2.13. *We have the following examples and properties of reciprocal matrices.*

- (1) *A symplectic matrix is reciprocal.*
- (2) *Suppose $P \in M_{n \times n}(\mathbb{Z})$ is a signed permutation matrix and $V \subset \mathbb{R}^n$ is a subspace invariant under P , then $P|_V : V \rightarrow V$ is reciprocal.*
- (3) *Suppose matrix A admits a block decomposition $\begin{bmatrix} B & * \\ 0 & C \end{bmatrix}$. Then two of A , B , or C being reciprocal implies that the remaining one is reciprocal as well.*

Proof. In (1), a *symplectic matrix* is a matrix $A \in M_{n \times n}(\mathbb{R})$ satisfying $\omega(Av, Av') = \omega(v, v')$ for all $v, v' \in \mathbb{R}^n$, for some symplectic form ω . It is an elementary fact that these are reciprocal (see, e.g., the proof of [18, Theorem 2.1]).

In (2), by a *signed permutation matrix*, we mean that the only nonzero entries of P are $P_{i, \sigma(i)} = \pm 1$ for some permutation $\sigma \in S_n$. Such matrices are orthogonal (with respect to

the standard inner product on $\mathbb{R}^n$), hence the restriction $P|_V$ is also orthogonal, and it is another elementary fact that orthogonal maps are reciprocal (again see [18, Theorem 2.1]).

The statement in (3) follows from the observation that the set of eigenvalues of A is the union of that of B and C (taken with multiplicity). ■

3. Standardly embedded train tracks

For the rest of this paper, we fix the following setting: Let $f : S \rightarrow S$ be a fully punctured pseudo-Anosov map. Let (ℓ^s, μ^s) and (ℓ^u, μ^u) be the stable and unstable measured foliations of f , respectively. Let λ be the expansion factor of f .

In this section, we will define standardly embedded train tracks and explain how they are related to certain Markov partitions that encode the dynamics of f .

3.1. Definition of standardly embedded train tracks

We first define the notion of standardly embedded train tracks. The definition we use is adapted from [10].

Let τ be a train track with tie neighborhood N . Each complementary region of τ in N is homeomorphic to an annulus, with one boundary component along ∂N and the other boundary component along τ . We call the boundary component along τ a *boundary component of τ* . Under this terminology, the boundary components of τ are in one-to-one correspondence with the boundary components of N . We denote the collection of boundary components of τ by $\partial\tau$.

We say that a boundary component c of τ is n -pronged if it contains n cusps. We say that c is odd/even-pronged if it is n -pronged for odd/even n , respectively.

Definition 3.1. Let $\partial\tau = \partial_I\tau \sqcup \partial_O\tau$ be a partition of the boundary components of τ into a nonempty set of *inner boundary components* and a nonempty set of *outer boundary components*, respectively.

A train track τ is said to be *standardly embedded with respect to $(\partial_I\tau, \partial_O\tau)$* if its set of edges $\mathcal{E}$ can be partitioned into a set of *infinitesimal edges* $\mathcal{E}_{\text{inf}}$ and a set of *real edges* $\mathcal{E}_{\text{real}}$, such that:

- The smoothing at each vertex v is defined by separating the infinitesimal edges and the real edges.
- The union of infinitesimal edges is a disjoint union of cycles, which we call the *infinitesimal polygons*.
- The infinitesimal polygons are exactly the inner boundary components of τ .

We show one example of a standardly embedded train track in Figure 3, where the infinitesimal polygons are drawn in gray.




Figure 3. An example of a standardly embedded train track. The infinitesimal polygons are drawn in gray.

Proposition 3.2. *Let τ be a standardly embedded train track. The number of real edges of τ is equal to $-\chi(\tau)$.*

Proof. It follows from the definition that $|\mathcal{E}_{\text{inf}}| = |\mathcal{V}|$, so τ has $|\mathcal{V}| + |\mathcal{E}_{\text{real}}|$ edges in total. Hence, $\chi(\tau) = |\mathcal{V}| - (|\mathcal{V}| + |\mathcal{E}_{\text{real}}|) = -|\mathcal{E}_{\text{real}}|$. ■

3.2. Train track partitions

We now define the notion of train track partitions. These will be used to construct standardly embedded train tracks in the next subsection.

Definition 3.3. A *rectangle* is defined to be a subset R of S such that $(R, \ell^s|_R, \ell^u|_R)$ is homeomorphic to

$$([0, 1] \times [0, 1]) \setminus \{x_1, \dots, x_n\}, \text{Vert, Hor}),$$

where $\{x_1, \dots, x_n\}$ is a (possibly empty) collection of points on the boundary of $[0, 1] \times [0, 1]$, Vert is the foliation by vertical lines, and Hor is the foliation by horizontal lines. See Figure 4, where the empty dots denote omitted points. In this paper, we will draw stable leaves in red and unstable leaves in blue.




Figure 4. A rectangle in S .

We call the two sides of ∂R that lie along leaves of ℓ^s the *stable sides* of R and write $\partial_s R$ for the union of the two stable sides. Similarly, we call the two sides of ∂R that lie along leaves of ℓ^u the *unstable sides* of R and write $\partial_u R$ for the union of the two unstable sides. In Figure 4, the stable sides are the vertical sides while the unstable sides are the horizontal sides.

A *partition* is defined to be a finite collection of rectangles $\{R_i\}$ that have disjoint interiors and cover S . A partition $\{R_i\}$ is *Markov* if it satisfies

- $f\left(\bigcup_i \partial_s R_i\right) \subset \bigcup_i \partial_s R_i$, and
- $f^{-1}\left(\bigcup_i \partial_u R_i\right) \subset \bigcup_i \partial_u R_i$.

Definition 3.4. Let x be a puncture of S . A *stable prong* at x is a connected subset of a stable leaf that limits to x . A *stable star* at x is a maximal disjoint union of prongs at x . In particular, a stable star at x has n connected components when x is an n -pronged puncture. A *side* of a stable star is the union of two adjacent prongs.

An *unstable prong* and an *unstable star* at x are similarly defined.

See Figure 5 for an example. Here, x is a 5-pronged puncture, so both the stable and unstable stars at x in the figure have five prongs. We have also indicated a side of the unstable star in dark blue.

Definition 3.5. Let $\mathcal{X} = \mathcal{X}_I \sqcup \mathcal{X}_O$ be a partition of the set of punctures of S into a nonempty set of *inner punctures* and a nonempty set of *outer punctures*, respectively.

A *train track partition with respect to* $(\mathcal{X}_I, \mathcal{X}_O)$ consists of

- a partition $\{R_i\}$, along with
- a stable star σ_x^s at every inner puncture $x \in \mathcal{X}_I$, and
- an unstable star σ_x^u at every outer puncture $x \in \mathcal{X}_O$

such that:

- Each stable star σ_x^s is disjoint from each unstable star σ_x^u .




Figure 5. The local picture of ℓ^s and ℓ^u at a 5-pronged puncture x . We have indicated a side of the unstable star at x in dark blue.

- Each stable side of each rectangle R_i lies along the closure of some side of some stable star σ_x^s , and each point on each σ_x^s meets the stable side of some R_i .
- Each unstable side of each rectangle R_i lies along the closure of some side of some unstable star σ_x^u , and each point on each σ_x^u meets the unstable side of some R_i .

A train track partition $(\{R_i\}, \{\sigma_x^s\}, \{\sigma_x^u\})$ is *Markov* if the partition $\{R_i\}$ is Markov.

Notice that the collection of stable and unstable stars actually determines the partition: The rectangles are the complementary regions of their union.

Also notice that if the train track partition is Markov, then f must preserve the set of inner and outer punctures, respectively. In particular, f has at least two puncture orbits. We will show below that this is in fact a sufficient condition. The following construction will be used in the proof.

Construction 3.6. Let $\mathcal{X} = \mathcal{X}_I \sqcup \mathcal{X}_O$ be some partition of $\mathcal{X}$ into two nonempty subsets.

Let $\widehat{\sigma}_x^u$ be the unstable star at each $x \in \mathcal{X}_O$ for which each of its prongs has μ^s -length 1. Let σ_x^s be the stable star at each $x \in \mathcal{X}_I$ that is maximal with respect to the property that it is disjoint from $\bigcup_{x \in \mathcal{X}_O} \widehat{\sigma}_x^u$. That is, each σ_x^s is defined by extending the stable prongs at x until it bumps into some $\widehat{\sigma}_x^u$. Then define σ_x^u to be the unstable star at each $x \in \mathcal{X}_O$ that is maximal with respect to the property that it contains $\widehat{\sigma}_x^u$ and is disjoint from $\bigcup_{x \in \mathcal{X}_I} \sigma_x^s$. That is, each σ_x^u is defined by extending the prongs of $\widehat{\sigma}_x^u$ until it bumps into some σ_x^s .

We claim that each complementary region of $\bigcup_{x \in \mathcal{X}_I} \sigma_x^s \cup \bigcup_{x \in \mathcal{X}_O} \sigma_x^u$ is a rectangle. By construction, the punctures never lie in the interior of a complementary region, hence each complementary region is foliated by properly embedded intervals that are the restriction of the stable leaves. Meanwhile, a stable star σ_x^s and an unstable star σ_x^u can only meet in the interior of the σ_x^s and at the endpoint of the σ_x^u or vice versa. That is, they meet in a $\perp$ form. Hence, each complementary region has convex corners along its boundary, which implies the claim.

By taking the partition to be the collection of complementary regions, we obtain a train track partition, which we denote by $\mathcal{M}_{(\mathcal{X}_I, \mathcal{X}_O)}$.

We remark that this construction of $\mathcal{M}_{(\mathcal{X}_I, \mathcal{X}_O)}$ is natural. More precisely, if $f_1 : S_1 \rightarrow S_1$ and $f_2 : S_2 \rightarrow S_2$ are fully punctured pseudo-Anosov maps, and $g : S_1 \rightarrow S_2$ is a homeomorphism that sends the stable and unstable measured foliations $(\ell_1^{s/u}, \mu_1^{s/u})$ of f_1 to $(\ell_2^{s/u}, \mu_2^{s/u})$ of f_2 , then for any partition $\mathcal{X}_1 = \mathcal{X}_{1,I} \sqcup \mathcal{X}_{1,O}$ of the set of punctures of S_1 , we have

$$g_* \mathcal{M}_{(\mathcal{X}_{1,I}, \mathcal{X}_{1,O})} = \mathcal{M}_{(g_* \mathcal{X}_{1,I}, g_* \mathcal{X}_{1,O})}.$$

More generally, if $f : S \rightarrow S$ is a fully punctured pseudo-Anosov map, $\pi : \tilde{S} \rightarrow S$ is a finite cover, and $\tilde{f} : \tilde{S} \rightarrow \tilde{S}$ is a lift of f , with its stable and unstable measured foliations lifted from that of f , then for any partition $\mathcal{X} = \mathcal{X}_I \sqcup \mathcal{X}_O$, we have

$$\pi^*(\mathcal{M}_{(\mathcal{X}_I, \mathcal{X}_O)}) = \mathcal{M}_{(\pi^* \mathcal{X}_I, \pi^* \mathcal{X}_O)}.$$

We will be applying this fact implicitly in the sequel.

Proposition 3.7. *If $\mathcal{X}_I$ and $\mathcal{X}_O$ are f -invariant, then $\mathcal{M}(\mathcal{X}_I, \mathcal{X}_O)$ is Markov.*

Proof. In Construction 3.6, notice that $f^{-1}(\bigcup_{x \in \mathcal{X}_O} \widehat{\sigma}_x^u) \subset \bigcup_{x \in \mathcal{X}_O} \widehat{\sigma}_x^u$ since f expands μ^s -lengths. This implies that $f^{-1}(\bigcup_{x \in \mathcal{X}_I} \sigma_x^s) \supset \bigcup_{x \in \mathcal{X}_I} \sigma_x^s$, since each $f^{-1}(\sigma_x^s)$ is obtained by extending the stable prongs at x until it bumps into some $f^{-1}(\widehat{\sigma}_x^u)$. Similarly, this in turn implies that $f^{-1}(\bigcup_{x \in \mathcal{X}_O} \sigma_x^u) \subset \bigcup_{x \in \mathcal{X}_O} \sigma_x^u$. Now the proposition follows from the observation that $\bigcup_i \partial_s R_i = \bigcup_{x \in \mathcal{X}_I} \sigma_x^s$ and $\bigcup_i \partial_u R_i = \bigcup_{x \in \mathcal{X}_O} \sigma_x^u$. ■

3.3. From train track partitions to standardly embedded train tracks

Train track partitions give rise to standardly embedded train tracks via the following construction.

Construction 3.8. Let $\mathcal{M} = (\{R_i\}, \{\sigma_x^s\}, \{\sigma_x^u\})$ be a train track partition with respect to $(\mathcal{X}_I, \mathcal{X}_O)$.

Define a graph $\tau_{\mathcal{M}}$ by taking a set of vertices in one-to-one correspondence with the sides of the stable stars σ_x^s . The edges of $\tau_{\mathcal{M}}$ will come in two types: infinitesimal and real. The infinitesimal edges are in one-to-one correspondence with the prongs of σ_x^s , with their endpoints at the two vertices corresponding to the two sides the prong lies in. The real edges are in one-to-one correspondence with the rectangles R_i , with their endpoints at the two vertices corresponding to the two sides the stable sides of R_i lie along.

The smoothing at each vertex is defined by separating the infinitesimal edges and the real edges. Each vertex is incident to exactly two infinitesimal half-edges, so it suffices to order the real half-edges: This order is determined by the position of rectangles at the corresponding side. It is straightforward to check that this makes $\tau_{\mathcal{M}}$ into a standardly embedded train track.

We illustrate a local picture of this construction in Figure 6.

When $\mathcal{M} = \mathcal{M}(\mathcal{X}_I, \mathcal{X}_O)$, we write $\tau(\mathcal{X}_I, \mathcal{X}_O) = \tau_{\mathcal{M}(\mathcal{X}_I, \mathcal{X}_O)}$.

Construction 3.9. When $\mathcal{M}$ is a Markov train track partition, we can in addition define a train track map $f_{\mathcal{M}} : \tau_{\mathcal{M}} \rightarrow \tau_{\mathcal{M}}$ as follows:

$f_{\mathcal{M}}$ maps the vertex corresponding to a side s of a stable star to the vertex corresponding to the side containing $f(s)$, maps the infinitesimal edge corresponding to a stable prong p to the infinitesimal edge corresponding to the prong containing $f(p)$, and maps the real edge corresponding to a rectangle R_i to the edge path corresponding to the sequence of rectangles and prongs passed through by $f(R_i)$.

When $\mathcal{M} = \mathcal{M}(\mathcal{X}_I, \mathcal{X}_O)$, we write $f(\mathcal{X}_I, \mathcal{X}_O) = f_{\mathcal{M}(\mathcal{X}_I, \mathcal{X}_O)}$.

The train track $\tau_{\mathcal{M}}$ and the train track map $f_{\mathcal{M}}$ capture the dynamics of f in the following sense.




Figure 6. A local example of Construction 3.8.

Proposition 3.10. *For every train track partition $\mathcal{M}$, there is an embedding ι of the tie neighborhood N of $\tau_{\mathcal{M}}$ into S that is a homotopy equivalence, such that $\iota(\tau_{\mathcal{M}})$ fully carries the unstable lamination of f (obtained by blowing air into the leaves of the unstable foliation ℓ^u that contain the punctures).*

Moreover, if $\mathcal{M}$ is Markov, then $f\iota$ and $\iota f_{\mathcal{M}}$ are homotopic embeddings of N in S .

Proof. Cut S along the unstable stars σ_x^u , which does not change the homeomorphism type of S . The resulting space $\widehat{S}$ can be obtained by gluing the rectangles R_i along their stable sides to the stable stars σ_x^s . Since the pattern of gluing corresponds exactly to that used to glue real edges to infinitesimal polygons of $\tau_{\mathcal{M}}$, we can define an embedding $\iota : \tau_{\mathcal{M}} \hookrightarrow S$ by sending the infinitesimal polygons into a neighborhood of the corresponding stable stars, then sending the real edges into their corresponding rectangles.

After blowing air into the leaves of the unstable foliation ℓ^u that contain the punctures, we can assume that the leaves of the unstable lamination are contained in $\widehat{S}$. In this case, we can extend ι into an embedding of N such that the ties are transverse to the leaves.

With this ι , $f\iota$ will send each the union of ties in N intersecting an infinitesimal polygon into a neighborhood of the image of the corresponding stable star under f , which by definition is the stable star corresponding to the image of the infinitesimal polygon under $f_{\mathcal{M}}$. Similarly, $f\iota$ will send the union of ties that intersect a real edge into the sequence of rectangles passed through by the image of the corresponding rectangle under f , which by definition is the sequence of rectangles corresponding to the image of the real edge under $f_{\mathcal{M}}$. From this one can construct a homotopy $f\iota \simeq \iota f_{\mathcal{M}}$. ■

In the sequel, we will always use ι in Proposition 3.10 to embed $\tau_{\mathcal{M}}$ in S if needed.

Remark 3.11. The proof of Proposition 3.10 shows that the boundary components of $\tau_{\mathcal{M}}$ are in one-to-one correspondence with the punctures of S . Namely, a boundary component c corresponds to the puncture p which $\iota(c)$ is homotopic into. Under this correspondence, c is an element of $\partial_{I/O}\tau_{\mathcal{M}}$ if and only if p is an element of $\mathcal{X}_{I/O}$, respectively. Also, c is n -pronged if and only if p is n -pronged.

The fact that $f_{\mathcal{M}}$ captures the dynamics of f implies that we can compute the expansion factor of f using the transition matrix of $f_{\mathcal{M}}$. More specifically, notice that by definition, $f_{\mathcal{M}}$ maps each infinitesimal edge to a single infinitesimal edge, hence if we list the infinitesimal edges in front of the real edges, the transition matrix of $f_{\mathcal{M}}$ will be of the form

$$f_{\mathcal{M}*} = \begin{bmatrix} P & * \\ 0 & f_{\mathcal{M}*}^{\text{real}} \end{bmatrix},$$

where P is a permutation matrix. We call $f_{\mathcal{M}*}^{\text{real}}$ the *real transition matrix* of $f_{\mathcal{M}}$.

Proposition 3.12. *For any Markov train track partition $\mathcal{M}$, $f_{\mathcal{M}*}^{\text{real}}$ is Perron–Frobenius.*

Proof. The rows and columns of $f_{\mathcal{M}*}^{\text{real}}$ are indexed by the real edges of $\tau_{\mathcal{M}}$, which in turn correspond to the rectangles R_i in the train track partition $\mathcal{M}$. Under this correspondence, the (R_j, R_i) -entry of $(f_{\mathcal{M}*}^{\text{real}})^k$ is the number of times $f^k(R_i)$ crosses R_j . We claim that for each i there exists k_i such that the (R_j, R_i) -entry of $(f_{\mathcal{M}*}^{\text{real}})^{k_i}$ is positive for each j . This would imply that $(f_{\mathcal{M}*}^{\text{real}})^{\prod k_i}$ is positive, hence $f_{\mathcal{M}*}^{\text{real}}$ is Perron–Frobenius.

To show the claim, we fix an i . By Proposition 2.4, there exists a periodic point z in the interior of R_i , say of period p . Consider a short interval I lying along the unstable leaf passing through z , with one endpoint on z and contained within R_i . Up to doubling p , we have $I \subset f^p(I)$ and $\mu^s(f^p(I)) = \lambda^p \mu^s(I)$.

Now $\bigcup_{s=0}^{\infty} f^{sp}(I)$ is a half-leaf, hence by Proposition 2.4, is dense in S . So there exists $s > 0$ such that $f^{sp}(I)$ meets the interior of each rectangle, which implies that $f^{sp}(R_i)$ crosses every rectangle and proves the claim. ■

Proposition 3.13. *For any Markov train track partition $\mathcal{M}$, the spectral radius of $f_{\mathcal{M}*}^{\text{real}}$ is the expansion factor of f .*

Proof. We will directly define an eigenvector u of $f_{\mathcal{M}*}^{\text{real}}$. The entries of u can be indexed by the rectangles R_i in $\mathcal{M}$ as above. We define the R_i -entry of u to be the μ^s -measure of an unstable side of R_i . By definition, the R_i -entry of $f_{\mathcal{M}*}^{\text{real}}u$ is the μ^s measure of the image of an unstable side of $f(R_i)$, which is λ times the R_i -entry of u . Hence, u is a λ -eigenvector of $f_{\mathcal{M}*}^{\text{real}}$.

We have shown that $f_{\mathcal{M}*}^{\text{real}}$ is Perron–Frobenius in Proposition 3.12. Here, u is positive, hence by Theorem 2.10, λ is the spectral radius of $f_{\mathcal{M}*}^{\text{real}}$. ■

Let $\mathcal{X} = \mathcal{X}_I \sqcup \mathcal{X}_O$ be a partition of the set of punctures of S into nonempty f -invariant sets. Propositions 3.2 and 3.10 imply that $f_{(\mathcal{X}_I, \mathcal{X}_O)^*}^{\text{real}}$ is a $|\chi(S)|$ -by- $|\chi(S)|$ matrix, while Proposition 3.12 implies that it is Perron–Frobenius. Our goal in Sections 4 and 5 is to show that $f_{(\mathcal{X}_I, \mathcal{X}_O)^*}^{\text{real}}$ is reciprocal, for then we can apply Theorem 2.12 and prove our main theorem.

3.4. Invariant standardly embedded train tracks

In this subsection, we will explain how for every Markov train track partition $\mathcal{M}$, $\tau_{\mathcal{M}}$ is an f -invariant train track, which implies Theorem 1.15. To do so, we have to discuss elementary moves on train tracks.

Definition 3.14. Let τ be a train track. Suppose e is an edge of τ with endpoints at switches v_1, v_2 . We can define a new train track τ' by declaring an interior point v of e as a new vertex and replacing e with two edges e_1, e_2 connecting v_1 and v_2 , respectively, with v . The ordering and smoothing at v_i is unchanged, with the half-edge determined by e_i replacing that determined by e . For v , we take the unique choice of ordering and smoothing.

Note that τ' is homeomorphic to τ as a topological space but nonisomorphic as a graph. Nevertheless, there is a natural train track map $\tau \rightarrow \tau'$. We refer to this map as *the subdivision move on e* (see Figure 7, top).

Definition 3.15. Let τ be a train track. Suppose e_1, e_2 are two edges of τ which determine a cusp at switch v , and having their other endpoint on switches v'_1 and v'_2 , respectively. Define a new train track τ' by combining v'_1 and v'_2 into one switch v' and replacing e_1 and e_2 by a single edge e connecting v and v' . Without loss of generality suppose that e_1 lies to the left of e_2 at v and the half-edge of e_i lies in $\mathcal{E}_{v_i}^2$ for both $i = 1, 2$. Then the ordering on $\mathcal{E}_{v'}^1 = \mathcal{E}_{v'_1}^1 \sqcup \mathcal{E}_{v'_2}^1$ is determined by placing all the half-edges in $\mathcal{E}_{v'_1}^1$,




Figure 7. Elementary moves. Top: The subdivision move on e . Bottom: The elementary folding move on (e_1, e_2) .

in their original order, to the left of those in $\mathcal{E}_{v'_2}^1$, in their original order. The ordering on $\mathcal{E}_{v'}^2 = (\mathcal{E}_{v'_1}^2 \setminus \{e_1\}) \sqcup \{e\} \sqcup (\mathcal{E}_{v'_2}^2 \setminus \{e_2\})$ is determined by placing all the half-edges in $\mathcal{E}_{v'_2}^2 \setminus \{e_2\}$, in their original order, to the left of e , which is in turn placed to the left of all the half-edges in $\mathcal{E}_{v'_1}^1 \setminus \{e_1\}$, in their original order.

There is a train track map $\tau \rightarrow \tau'$ defined by sending v_1 and v_2 to v , sending e_1 and e_2 to e , and sending the remaining vertices and edges to themselves. We refer to this map as *the elementary folding move on (e_1, e_2)* (see Figure 7, bottom).

We refer to a subdivision move or an elementary folding move as an *elementary move* in general.

Definition 3.16. Let τ be a train track and with tie neighborhood N , and let $\iota : N \hookrightarrow S$ be an embedding. Suppose there exists train track maps $f_1, \dots, f_n, \sigma$ where each f_i is an elementary move and σ is an isomorphism of train tracks, such that $f\iota$ and $\iota\sigma f_n \cdots f_1$ are homotopic embeddings of N in S . Then $\iota(\tau)$ is said to be *f -invariant*.

When ι is understood, we will just say that τ is *f -invariant*.

Proposition 3.18 will do most of the heavy lifting in showing that $\tau_{\mathcal{M}}$ is *f -invariant*. We will also be applying Proposition 3.18 and its corollary, Proposition 3.19, in Section 4 and Section 5.

Definition 3.17. Let $\mathcal{M} = (\{R_i\}, \{\sigma_x^s\}, \{\sigma_x^u\})$ and $\mathcal{M}' = (\{R'_i\}, \{\sigma'_x{}^s\}, \{\sigma'_x{}^u\})$ be two train track partitions with respect to $(\mathcal{X}_I, \mathcal{X}_O)$ and $(\mathcal{X}'_I, \mathcal{X}'_O)$, respectively. We say that $\mathcal{M}$ is *wider* than $\mathcal{M}'$ if $\bigcup_{x \in \mathcal{X}_I} \sigma_x^s \subset \bigcup_{x \in \mathcal{X}'_I} \sigma'_x{}^s$ and $\bigcup_{x \in \mathcal{X}_O} \sigma_x^u \supset \bigcup_{x \in \mathcal{X}'_O} \sigma'_x{}^u$. Notice this implies that $\mathcal{X}_I \subset \mathcal{X}'_I$ and $\mathcal{X}_O \supset \mathcal{X}'_O$.

Proposition 3.18. *If $\mathcal{M}$ is wider than $\mathcal{M}'$ then there exists elementary moves $f_1, \dots, f_n$ such that $f_n \cdots f_1$ maps $\tau_{\mathcal{M}}$ to $\tau_{\mathcal{M}'}$.*

Proof. For notational convenience we set $\sigma_x^s = \emptyset$ for $x \in \mathcal{X}_O$ and $\sigma_x^u = \emptyset$ for $x \in \mathcal{X}_I$, and similarly $\sigma'_x{}^s = \emptyset$ for $x \in \mathcal{X}'_O$ and $\sigma'_x{}^u = \emptyset$ for $x \in \mathcal{X}'_I$. Then we will have $\sigma_x^s \subset \sigma'_x{}^s$ and $\sigma_x^u \supset \sigma'_x{}^u$ for every x .

We label the prongs of all the σ_x^u as $p_1, \dots, p_N$ and label the prong of the σ_x^u that is contained in p_i as p'_i . The idea of the proof is that in contracting each p_i into p'_i , we combine some rectangles, and this determines corresponding elementary moves (see Figure 8). The precise description of the proof is rather technical. The reader may wish to skip it on the first reading.

Define subsets p_i^j of p_i by setting $p_i^0 = p_i$ and inductively defining p_i^{j+1} to be the subset maximal with respect to the property of containing p'_i and being properly contained in p_i^j , and having an endpoint lying on $\bigcup_x \sigma_x^s$. Suppose $p'_i = p_i^{n_i}$. Write $U_{\sum_{i=1}^{k-1} n_i + j} = \bigcup_{i=1}^{k-1} p'_i \cup p_k^j \cup \bigcup_{i=k}^N p_i$.

Notice that for each j , $(\bigcup_x \sigma_x^s) \setminus U_j$ is a union of intervals. We call each such interval a *prong* of $(\bigcup_x \sigma_x^s) \setminus U_j$. Each interval has two sides in the ℓ^u direction. We consider




Figure 8. Performing elementary folding moves on τ_j to obtain τ_{j+1} .

two sides to such intervals to be equivalent if there is a rectangle in the complement of $\bigcup_x \sigma_x'^s \cup U_j$ with a stable side incident to the two sides. We call each equivalence class of sides a *side* of $(\bigcup_x \sigma_x'^s) \setminus U_j$. Thus, each prong lies on two sides, and for each j , $(\bigcup_x \sigma_x'^s) \setminus U_{j+1}$ has one or two less sides than $(\bigcup_x \sigma_x'^s) \setminus U_j$. Intuitively, as j increases, U_j shrinks and the sides of $(\bigcup_x \sigma_x'^s) \setminus U_j$ get combined, until we hit $j = \sum_{i=1}^N n_i$ and $(\bigcup_x \sigma_x'^s) \setminus U_{\sum_{i=1}^N n_i}$ becomes exactly the union of stable stars that are the $\sigma_x'^s$ and the definitions of prongs and sides agree with our previous usage.

We will define train tracks τ_j for each $j = 0, \dots, \sum_{i=1}^N n_i$ such that

- the vertices of τ_j are in one-to-one correspondence with the sides of $(\bigcup_x \sigma_x'^s) \setminus U_j$, and
- the edges of τ_j are in one-to-one correspondence with the rectangles in the complement of $\bigcup_x \sigma_x'^s \cup U_j$ and the prongs of $(\bigcup_x \sigma_x'^s) \setminus U_j$,

and such that each τ_{j+1} is obtained from τ_j via elementary folding moves.

First, to define τ_0 , consider the components of $(\bigcup_x \sigma_x'^s) \setminus U_0 = (\bigcup_x \sigma_x'^s) \setminus (\bigcup_x \sigma_x^u)$ that do not lie in $\bigcup_x \sigma_x^s$. Each of these is an interval lying in a rectangle R_i . For each of these intervals, we subdivide twice the real edge of $\tau_{\mathcal{M}}$ corresponding to the rectangle the interval lies in. That is, we add two new vertices, each corresponding to a side of $(\bigcup_x \sigma_x'^s) \setminus U_0$.

Inductively, suppose τ_j is defined. When going from U_j to U_{j+1} , two rectangles in the complement of $\bigcup_x \sigma_x'^s \cup U_j$ are combined into one in $\bigcup_x \sigma_x'^s \cup U_{j+1}$. We fold the two edges of τ_j corresponding to these two rectangles. If $(\bigcup_x \sigma_x'^s) \setminus U_{j+1}$ has one less component than $(\bigcup_x \sigma_x'^s) \setminus U_j$, then two prongs of $(\bigcup_x \sigma_x'^s) \setminus U_j$ are combined into one. In this case, we also fold the two edges of τ_j corresponding to these two prongs. After these one or two elementary folding moves, we obtain τ_{j+1} (see Figure 8).

Continuing inductively, by the time we reach $\tau_{\sum_{i=1}^N n_i}$, this will be a train track with vertices in one-to-one correspondence with the sides of $\bigcup_x \sigma_x'^s$ and edges in one-to-one correspondence with the rectangles R'_i and the prongs of $\bigcup_x \sigma_x'^s$, thus $\tau_{\sum_{i=1}^N n_i} = \tau_{\mathcal{M}'}$. ■

Proposition 3.19. *If $\mathcal{X}_I \subset \mathcal{X}'_I$, then there are elementary folding moves $f_1, \dots, f_n$ such that $f_n \cdots f_1$ maps $\tau(\mathcal{X}_I, \mathcal{X}_O)$ to $\tau(\mathcal{X}'_I, \mathcal{X}'_O)$.*

Proof. If $\mathcal{X}_I \subset \mathcal{X}'_I$, then $\mathcal{X}_O \supset \mathcal{X}'_O$, so $\bigcup \widehat{\sigma}_x^u \supset \bigcup \widehat{\sigma}_x'^u$ in Construction 3.6. This implies that $\bigcup \sigma_x^s \subset \bigcup \sigma_x'^s$, which in turn implies that $\bigcup \sigma_x^u \supset \bigcup \sigma_x'^u$. That is, $\mathcal{M}(\mathcal{X}_I, \mathcal{X}_O)$ is wider than $\mathcal{M}(\mathcal{X}'_I, \mathcal{X}'_O)$. So this proposition follows from Proposition 3.18. ■

Lemma 3.20. *For any Markov train track partition $\mathcal{M}$, the train track map $f_{\mathcal{M}} : \tau_{\mathcal{M}} \rightarrow \tau_{\mathcal{M}}$ can be written as a composition of train track maps $\sigma f_n \cdots f_1$ where each f_i is an elementary move and σ is an isomorphism of train tracks.*

Proof. Note that $\mathcal{M}$ being Markov implies that $\mathcal{M}$ is wider than $f^{-1}(\mathcal{M})$. So we can apply Proposition 3.18 to get elementary moves $f_1, \dots, f_n$ such that $f_n \cdots f_1$ maps $\tau_{\mathcal{M}}$ to $\tau_{f^{-1}(\mathcal{M})}$. Meanwhile, since f sends $f^{-1}(\mathcal{M})$ to $\mathcal{M}$, there is an induced isomorphism of train tracks $\sigma : \tau_{f^{-1}(\mathcal{M})} \rightarrow \tau_{\mathcal{M}}$.

It remains to show that $f_{\mathcal{M}} = \sigma f_n \cdots f_1$. This will follow from checking that these maps send each edge e to the same edge path.

Suppose e is infinitesimal and corresponds to a prong p . Then $f_n \cdots f_1$ sends e to the infinitesimal edge in $\tau_{f^{-1}(\mathcal{M})}$ corresponding to the prong containing p , which σ sends to the infinitesimal edge in $\tau_{\mathcal{M}}$ corresponding to the prong containing $f(p)$. This is equal to the image of e under $f_{\mathcal{M}}$.

Suppose e is real and corresponds to a rectangle R . Then $f_n \cdots f_1$ sends e to the edge path in $\tau_{f^{-1}(\mathcal{M})}$ corresponding to the sequence of rectangles and prongs passed by an unstable side of R , which σ sends to the edge path in $\tau_{\mathcal{M}}$ corresponding to the sequence of rectangles and prongs passed by an unstable side of $f(R)$. This is equal to the image of e under $f_{\mathcal{M}}$. ■

Combining Proposition 3.7, Proposition 3.10, and Lemma 3.20, we have the following.

Proposition 3.21. *Let $f : S \rightarrow S$ be a fully punctured pseudo-Anosov map with at least two puncture orbits. Let $\mathcal{X} = \mathcal{X}_I \sqcup \mathcal{X}_O$ be some partition of the set of punctures into two nonempty f -invariant subsets. Then there exists an f -invariant train track τ that is standardly embedded with respect to $(\partial_I \tau, \partial_O \tau)$, where the boundary components in $\partial_{I/O} \tau$ are homotopic into the punctures in $\mathcal{X}_{I/O}$, respectively.*

Theorem 1.15 follows from Proposition 3.21.

4. The Thurston symplectic form

In this section, we make the first steps towards showing that $f_{(\mathcal{X}_I, \mathcal{X}_O)^*}^{\text{real}}$ is reciprocal. The idea is to consider the space of weights on the train track, on which the Thurston symplectic form can be defined. There is a slight misnomer here: The Thurston symplectic form

is in general degenerate hence not a symplectic form. However, we will only set up the theory in this section, leaving the task of addressing this to Section 5.

4.1. Weight space

Definition 4.1. Let τ be a train track. A system of *weights* on τ is an assignment of a real number $w(e)$ to each edge e of τ such that at each switch v ,

$$\sum_{e \in \mathcal{E}_v^1} w(e) = \sum_{e \in \mathcal{E}_v^2} w(e).$$

It is convenient to think of the weight of a branch as the width of the branch. From this point of view, the equation above simply states that the total width on the two sides of a switch match up.

The *weight space* of τ is the linear space of all systems of weights on τ , which we denote by $\mathcal{W}(\tau)$.

Consider the space $\mathbb{R}^{\mathcal{E}(\tau)}$ of all assignments of real numbers to the edges of τ . Define the linear map $T_v : \mathbb{R}^{\mathcal{E}(\tau)} \rightarrow \mathbb{R}$ by $T_v(w) = \sum_{e \in \mathcal{E}_v^1} w(e) - \sum_{e \in \mathcal{E}_v^2} w(e)$ for each switch v , and define $T_{\mathcal{V}} : \mathbb{R}^{\mathcal{E}(\tau)} \rightarrow \mathbb{R}^{\mathcal{V}(\tau)}$ by $T_{\mathcal{V}}(w) = (T_v(w))_{v \in \mathcal{V}(\tau)}$. Then the weight space $\mathcal{W}(\tau)$ is the kernel of $T_{\mathcal{V}}$. In other words, we have the exact sequence

$$0 \longrightarrow \mathcal{W}(\tau) \longrightarrow \mathbb{R}^{\mathcal{E}(\tau)} \xrightarrow{T_{\mathcal{V}}} \mathbb{R}^{\mathcal{V}(\tau)}.$$

We remark that there is no canonical way to label $\mathcal{E}_v^1$ and $\mathcal{E}_v^2$ for each v , so each T_v can only be canonically defined up to a sign. This is not very significant for our purposes, but for concreteness let us simply fix some labeling of $\mathcal{E}_v^1$ and $\mathcal{E}_v^2$ for each v .

Proposition 4.2. Let $f : \tau \rightarrow \tau'$ be a train track map. There exists a signed permutation matrix $P \in M_{\mathcal{V}(\tau') \times \mathcal{V}(\tau)}(\mathbb{R})$ which fits into the commutative diagram

$$\begin{array}{ccccccc} 0 & \longrightarrow & \mathcal{W}(\tau) & \longrightarrow & \mathbb{R}^{\mathcal{E}(\tau)} & \xrightarrow{T_{\mathcal{V}}} & \mathbb{R}^{\mathcal{V}(\tau)} \\ & & \downarrow \mathcal{W}(f) & & \downarrow f_* & & \downarrow P \\ 0 & \longrightarrow & \mathcal{W}(\tau') & \longrightarrow & \mathbb{R}^{\mathcal{E}(\tau')} & \xrightarrow{T_{\mathcal{V}}} & \mathbb{R}^{\mathcal{V}(\tau')} \end{array}.$$

Here, f_* is the transition matrix of f and $\mathcal{W}(f)$ is the restriction of f_* to $\mathcal{W}(\tau)$.

Proof. Recall that f sends switches to switches. Define

$$P_{v',v} = \begin{cases} 1, & \text{if } v' = f(v) \text{ and } D_v f \text{ sends } \mathcal{E}_v^1 \text{ into } \mathcal{E}_{v'}^1, \\ -1, & \text{if } v' = f(v) \text{ and } D_v f \text{ sends } \mathcal{E}_v^1 \text{ into } \mathcal{E}_{v'}^2, \\ 0, & \text{otherwise.} \end{cases}$$

We have to check that $PT_{\mathcal{V}} = T_{\mathcal{V}}f_*$. For convenience, let us denote the unit vector in $\mathbb{R}^{\mathcal{E}(\tau)}$ corresponding to $e \in \mathcal{E}(\tau)$ by e as well. Then $PT_{\mathcal{V}}(e)$ is the vector having two ± 1 entries at the images of the endpoints of e under f , whereas $T_{\mathcal{V}}f_*(e)$ is the vector having two ± 1 entries at the endpoints of the image of e under f . The position of the only two nonzero entries hence coincide. It can also be checked that their signs coincide respectively. ■

Note that by Proposition 4.2 and the discussion above Proposition 3.12 (along with Proposition 2.13), in order to show that $f_{(\mathcal{X}_I, \mathcal{X}_O)^*}^{\text{real}}$ is reciprocal, it suffices to show that $\mathcal{W}(f_{(\mathcal{X}_I, \mathcal{X}_O)})$ is reciprocal.

4.2. The Thurston symplectic form

In this subsection we define the Thurston symplectic form on the weight space of a train track. This is done for trivalent train tracks in [33, Section 3.2]; we simply generalize the discussion to general train tracks.

Definition 4.3. Let $w_1, w_2 \in \mathcal{W}(\tau)$ be two systems of weights on τ . We define

$$\omega(w_1, w_2) = \sum_{v \in \mathcal{V}(\tau)} \sum_{e_1 \text{ left of } e_2} (w_1(e_1)w_2(e_2) - w_1(e_2)w_2(e_1)),$$

where the second summation is taken over all pairs e_1, e_2 in $\mathcal{E}_v$ for which e_1 is on the left of e_2 .

Then ω is clearly a skew-symmetric bilinear form on $\mathbb{R}^{\mathcal{E}(\tau)}$ hence on $\mathcal{W}(\tau)$. We call ω the *Thurston symplectic form*.

The definition of the Thurston symplectic form ω is motivated from the algebraic intersection number. We refer to [33, Lemma 3.2.2] for an explanation of this. The property of ω that matters to us is that it is preserved by the type of train track maps that we study.

Lemma 4.4. *If $f : \tau \rightarrow \tau'$ is a subdivision move, then $\mathcal{W}(f)$ is an isomorphism that preserves the Thurston symplectic form ω .*

Proof. We use the notation as in Definition 3.14, and write $w' = (\mathcal{W}(f))(w)$.

The map $\mathcal{W}(f)$ is an isomorphism since $w(e)$ can be recovered as $w'(e_1) = w'(e_2)$.

When computing ω in τ' , v does not make a contribution, since it only meets two half-edges. For the other vertices, since $w'(e_1) = w'(e_2) = w(e)$, their contributions remain the same. ■

Lemma 4.5. *If $f : \tau \rightarrow \tau'$ is an elementary folding move, then $\mathcal{W}(f)$ is an isomorphism that preserves the Thurston symplectic form ω .*

Proof. We use the notation as in Definition 3.15, and write $w'_i = (\mathcal{W}(f))(w_i)$.

The map $\mathcal{W}(f)$ is an isomorphism since $w_i(e_j)$ can be recovered as

$$\sum_{e \in \mathcal{E}_{v_j}^1} w'_i(e) - \sum_{e \in \mathcal{E}_{v_j}^2 \setminus \{e_j\}} w'_i(e).$$

When computing ω in τ' , the contributions from v and v' are as follows (we omit the summands when they are $w_1(e')w_2(e'') - w_1(e'')w_2(e')$ and omit writing “ e' ” to the left of e'' ” under the summations):

$$\begin{aligned} & \sum_{e', e'' \in \mathcal{E}_v^2} + \sum_{e', e'' \in \mathcal{E}_v^1} + \sum_{e', e'' \in \mathcal{E}_{v'}^2} + \sum_{e', e'' \in \mathcal{E}_{v'}^1} \\ &= \sum_{e', e'' \in \mathcal{E}_v^2} + \left(\sum_{e', e'' \in \mathcal{E}_v^1} -(w_1(e_1)w_2(e_2) - w_1(e_2)w_2(e_1)) \right) \\ &+ \left(\sum_{e', e'' \in \mathcal{E}_{v'_1}^2} + \sum_{e', e'' \in \mathcal{E}_{v'_2}^2} + \sum_{e' \in \mathcal{E}_{v'_2}^2, e'' \in \mathcal{E}_{v'_1}^2} -(w_1(e_2)w_2(e_1) - w_1(e_1)w_2(e_2)) \right) \\ &+ \left(\sum_{e', e'' \in \mathcal{E}_{v'_1}^1} + \sum_{e', e'' \in \mathcal{E}_{v'_2}^1} + \sum_{e' \in \mathcal{E}_{v'_1}^1, e'' \in \mathcal{E}_{v'_2}^1} \right) \\ &= \sum_{e', e'' \in \mathcal{E}_v^2} + \sum_{e', e'' \in \mathcal{E}_v^1} + \sum_{e', e'' \in \mathcal{E}_{v'_1}^2} + \sum_{e', e'' \in \mathcal{E}_{v'_2}^2} \\ &+ \left(\sum_{e' \in \mathcal{E}_{v'_2}^2} w_1(e') \right) \left(\sum_{e'' \in \mathcal{E}_{v'_1}^2} w_2(e'') \right) - \left(\sum_{e'' \in \mathcal{E}_{v'_1}^2} w_1(e'') \right) \left(\sum_{e' \in \mathcal{E}_{v'_2}^2} w_2(e') \right) \\ &+ \sum_{e', e'' \in \mathcal{E}_{v'_1}^1} + \sum_{e', e'' \in \mathcal{E}_{v'_2}^1} + \left(\sum_{e' \in \mathcal{E}_{v'_1}^1} w_1(e') \right) \left(\sum_{e'' \in \mathcal{E}_{v'_2}^1} w_2(e'') \right) \\ &- \left(\sum_{e'' \in \mathcal{E}_{v'_2}^1} w_1(e'') \right) \left(\sum_{e' \in \mathcal{E}_{v'_1}^1} w_2(e') \right) \\ &= \sum_{e', e'' \in \mathcal{E}_v^2} + \sum_{e', e'' \in \mathcal{E}_v^1} + \sum_{e', e'' \in \mathcal{E}_{v'_1}^2} + \sum_{e', e'' \in \mathcal{E}_{v'_2}^2} + \sum_{e', e'' \in \mathcal{E}_{v'_1}^1} + \sum_{e', e'' \in \mathcal{E}_{v'_2}^1}, \end{aligned}$$

which is the contribution from v, v'_1, v'_2 when computing ω in τ .

The contributions from the rest of the vertices stay the same, so $\mathcal{W}(f)$ preserves ω . ■

Proposition 4.6. *For any Markov train track partition $\mathcal{M}$, $\mathcal{W}(f_{\mathcal{M}}) : \mathcal{W}(\tau_{\mathcal{M}}) \rightarrow \mathcal{W}(\tau_{\mathcal{M}})$ is an isomorphism that preserves ω .*

Proof. This follows from Lemma 3.20, Lemma 4.4, Lemma 4.5, and the fact that an isomorphism of train tracks induces an isomorphism which preserves the Thurston symplectic form on the weight space. ■

In general, Lemmas 4.4 and 4.5 essentially say that we can modify a train track as much as we like using elementary moves when studying its weight space. Hence, we make the following definition.

Definition 4.7. Consider the equivalence relation on the set of all train tracks that is generated by there being an elementary move between two train tracks. We say that two train tracks in the same equivalence class are *equivalent*.

In particular, we have the following lemma, which will be very useful when making computations in Section 5.

Lemma 4.8. Let $\mathcal{X} = \mathcal{X}_I \sqcup \mathcal{X}_O$ and $\mathcal{X} = \mathcal{X}'_I \sqcup \mathcal{X}'_O$ be two partitions of the set of punctures of S . Then $\tau_{(\mathcal{X}_I, \mathcal{X}_O)}$ and $\tau_{(\mathcal{X}'_I, \mathcal{X}'_O)}$ are equivalent.

Proof. This follows immediately from Proposition 3.19. ■

5. Computing the radical

As mentioned in the last section, the Thurston symplectic form ω on the weight space $\mathcal{W}(\tau)$ is not actually a symplectic form in general. Its failure of being one is measured by its *radical* $\text{rad}(\omega) = \{w_0 \in \mathcal{W}(\tau) : \omega(w, w_0) = 0 \text{ for every } w \in \mathcal{W}(\tau)\}$.

Now, if ω were symplectic, we would be able to show that $\mathcal{W}(f_{(\mathcal{X}_I, \mathcal{X}_O)})$ is reciprocal by just applying Proposition 2.13 (1). In general, to apply this approach, we need to understand what $\text{rad}(\omega)$ is. In this section, we make the computations to determine this.

5.1. Radical elements

Definition 5.1. Suppose c is an even-pronged boundary component of τ . Label the intervals in the complement of the cusps by $I_1, \dots, I_n$ in a cyclic order. We define the *radical element* of c , denoted by r_c , to be the element of $\mathcal{W}(\tau)$ where we assign to each edge on I_k a weight of $(-1)^k$ (see Figure 9, left). Here, we assign weights with multiplicity, meaning if an edge appears multiple times on possibly multiple I_k , then the weight we assign to it is the sum of the weights that is assigned to it each time it appears in some I_k .

In the degenerate case when c is 0-pronged, r_c is the element of $\mathcal{W}(\tau)$ that assigns each edge on c a weight of 1.

We remark that there is no canonical way to label the I_k , so r_c can only be canonically defined up to a sign. This is not very significant for our purposes, we can simply fix a labeling where appropriate.




Figure 9. Left: A radical element r_c . Right: Showing that the terms cancel each other out when computing $\omega(w, r_c)$.

Proposition 5.2. *Each r_c lies in the radical of ω .*

Proof. Let w be an element of $\mathcal{W}(\tau)$, we have to check that $\omega(w, r_c) = 0$. Let us first assume that c is embedded in τ for simplicity. Let v be a vertex on c that lies at a cusp. Suppose half-edges e_1 and e_2 determine the cusp, where e_1 is to the left of e_2 , and suppose that r_c assigns the weight $(-1)^i$ to e_i . Then the contribution of v in $\omega(w, r_c)$ is

$$\begin{aligned}
 & \sum_{e \text{ left of } e_1} (w(e) \times (-1) - w(e_1) \times 0) + (w(e) \times 1 - w(e_2) \times 0) \\
 & + \sum_{e_2 \text{ left of } e} (w(e_2) \times 0 - w(e) \times 1) + (w(e) \times 1 - w(e_2) \times 1) \\
 & + w(e_1) \times 1 - w(e_2) \times (-1) \\
 & = w(e_1) + w(e_2).
 \end{aligned}$$

If instead r_c assigns the weight $(-1)^{i+1}$ to e_i , then the contribution of v is $-w(e_1) - w(e_2)$.

Now let v be a vertex on c that does not lie at a cusp. Let $e_\beta \in \mathcal{E}_v^\beta$, $\beta = 1, 2$, be the two edges that lie on c , such that e_1 is the rightmost half-edge in $\mathcal{E}_v^1$ and e_2 is the leftmost half-edge in $\mathcal{E}_v^2$. Suppose that r_c assigns the weight 1 to e_1 and e_2 . Then the contribution of v in $\omega(w, r_c)$ is

$$\begin{aligned}
 & \sum_{e \text{ left of } e_1} (w(e) \times 1 - w(e_1) \times 0) + \sum_{e_2 \text{ left of } e} (w(e_2) \times 0 - w(e) \times 1) \\
 & = \left(\sum_{e \in \mathcal{E}_v^1} w(e) - w(e_1) \right) - \left(\sum_{e \in \mathcal{E}_v^2} w(e) - w(e_2) \right) \\
 & = w(e_2) - w(e_1).
 \end{aligned}$$

If instead r_c assigns the weight -1 to e_1 and e_2 , then the contribution of v is $w(e_1) - w(e_2)$.

When adding together the contributions from all vertices, the terms cancel out in pairs, giving us 0. We schematically illustrate how the canceling occurs in Figure 9 (right).

If c is not embedded, that is, some switch of τ meets c more than once, then the above computation still holds with some more careful bookkeeping. For example, if a vertex v lies on two cusps of c , say determined by pairs of half-edges (e_1^1, e_2^1) and (e_1^2, e_2^2) , respectively, then one can check that the contribution of v is $\pm(w(e_1^1) + w(e_2^1)) \pm (w(e_1^2) + w(e_2^2))$, depending on whether r_c assigns $\pm(-1)^i$ to e_i^1 and e_i^2 , respectively. We let the reader fill in the details. Alternatively, one can pass to a finite cover of τ where c is embedded and perform the computation there. ■

The radical elements behave nicely with respect to elementary moves. To state this precisely, first note that if τ' is obtained from τ by an elementary move, then the boundary components of τ' are in natural one-to-one correspondence with those of τ .

Proposition 5.3. *Suppose c' is a boundary component of τ' that corresponds to a boundary component c of τ . Then r_c maps to $r_{c'}$ under the isomorphism $\mathcal{W}(\tau) \rightarrow \mathcal{W}(\tau')$ induced by the elementary move.*

Proof. The proposition is clear for subdivision moves, and clear for elementary folding moves if, in the notation of Definition 3.15, the cusp determined by (e_1, e_2) is not in c . If the cusp is in c , then, up to a sign, r_c assigns the weight $(-1)^i$ to e_i , hence upon folding these add up to 0, which is the weight assigned by $r_{c'}$. ■

Proposition 5.4. *Let $\mathcal{X} = \mathcal{X}_I \sqcup \mathcal{X}_O$ be a partition of the set of punctures into two nonempty f -invariant sets. Then the train track map $f(\mathcal{X}_I, \mathcal{X}_O) : \tau(\mathcal{X}_I, \mathcal{X}_O) \rightarrow \tau(\mathcal{X}_I, \mathcal{X}_O)$ preserves $\text{span}\{r_c\}$ and acts on it via a reciprocal map.*

Proof. Lemma 3.20 and Proposition 5.3 imply that $\mathcal{W}(f(\mathcal{X}_I, \mathcal{X}_O))$ preserves $\text{span}\{r_c\}$.

For the second part of the statement, we claim that $\{r_c\}$ are linearly independent unless all punctures are even-pronged, in which case there is only at most one relation of the form $\sum_c \pm r_c = 0$.

Notice that by Proposition 5.3, the validity of the claim is invariant under equivalence of train tracks. Hence by Lemma 4.8, we can assume that $\mathcal{X}_O$ consists of a single element corresponding to $c_O \in \partial\tau$.

If there is an odd-pronged puncture, we can further assume that c_O is odd-pronged. In this case, all even-pronged boundary components of the train track are disjoint, hence $\{r_c\}$ are linearly independent.

If all punctures are even-pronged, then all even-pronged boundary components except for c_O are disjoint. For an edge e on an infinitesimal polygon c , r_c and r_{c_O} , and only r_c and r_{c_O} , are nonzero on it. In fact, they assign weights ± 1 to e , so the ratio of their coefficients is uniquely determined to be one of ± 1 . This proves the claim.

Returning to the proof of the second part of the proposition, if $\{r_c\}$ are linearly independent, then by Lemma 3.20 and Proposition 5.3, $f(\mathcal{X}_I, \mathcal{X}_O)$ acts on $\text{span}\{r_c\}$ by a signed permutation matrix, so this follows from Proposition 2.13 (2).

If $\{r_c\}$ are not linearly independent, then consider the commutative diagram

$$\begin{array}{ccccccc}
 0 & \longrightarrow & \langle \sum \pm r_c \rangle & \longrightarrow & \mathbb{R}^{\mathcal{X}} & \longrightarrow & \text{span}\{r_c\} \longrightarrow 0 \\
 & & \downarrow \pm 1 & & \downarrow P & & \downarrow \mathcal{W}(f(\mathcal{X}_I, \mathcal{X}_O)) \\
 0 & \longrightarrow & \langle \sum \pm r_c \rangle & \longrightarrow & \mathbb{R}^{\mathcal{X}} & \longrightarrow & \text{span}\{r_c\} \longrightarrow 0
 \end{array} ,$$

where P is a signed permutation matrix, hence is reciprocal, as in the last case. By Proposition 2.13 (3), $\mathcal{W}(f(\mathcal{X}_I, \mathcal{X}_O))$ acts on $\text{span}\{r_c\}$ via a reciprocal map. ■

The significance of Proposition 5.4 comes from the following result, whose proof will occupy the rest of this section.

Proposition 5.5. *Let $\mathcal{X} = \mathcal{X}_I \sqcup \mathcal{X}_O$ be a partition of the set of punctures into two nonempty sets. For the train track $\tau(\mathcal{X}_I, \mathcal{X}_O)$, we have*

$$\text{rad}(\omega) = \text{span}\{r_c\}, \quad (5.1)$$

where c ranges over all even-pronged boundary components of τ .

The strategy to proving Proposition 5.5 is to modify the train track into a convenient form before making concrete computations. This strategy was already used in the proof of Proposition 5.4, where we modified the train track up to equivalence. This applies equally well in the setting of Proposition 5.5, as we have the following observation.

Lemma 5.6. *Suppose τ and τ' are equivalent, then (5.1) holds for τ if and only if it holds for τ' .*

Proof. This follows from Lemma 4.5 and Proposition 5.3. ■

Let us call a standardly embedded train track *floral* if it has only one infinitesimal polygon. Visually, the single infinitesimal polygon forms the pistil while the real edges form the petals of a flower. See Figure 10 for an example. Hence using Lemmas 4.8 and 5.6, we can assume that $\mathcal{X}_I$ consists of a single element, that is, $\tau(\mathcal{X}_I, \mathcal{X}_O)$ is floral.

Here, the proof divides into two cases. Case 1 is if there is an odd-pronged puncture. In this case, we can take the single infinitesimal polygon of $\tau(\mathcal{X}_I, \mathcal{X}_O)$ to be odd-pronged. We show that *admissible deletion* of real edges preserves (5.1) (Lemmas 5.9 and 5.10). This allows us to modify our train track into a simple form where we can explicitly verify (5.1). This case is tackled in Section 5.2.

Case 2 is if all punctures are even-pronged. In this case, we would like to repeat the reasoning in case 1, but here we must first pass to the orientable cover (Lemma 5.12) before the arguments work. This case is tackled in Section 5.3.

5.2. Case 1: There is an odd-pronged puncture

As explained above, in this case we can assume that $\tau(x_I, x_O)$ is floral with an odd-pronged infinitesimal polygon, or *odd floral* for short. See Figure 10 for an example.

For odd floral train tracks, we can construct some convenient elements in the weight space for which we can test elements of $\text{rad}(\omega)$ against.

Construction 5.7. Let τ be an odd floral train track. Let e be a real edge in τ . Let c_I be the infinitesimal polygon of τ . Label the vertices of c_I by $v_1, \dots, v_n$ in a cyclic order such that the endpoints of e lie on v_1 and v_k . Up to flipping the ordering, we can assume that k is even. Also label the edges of c_I by $e_1, \dots, e_n$ such that e_i connects v_i to v_{i+1} .

We define $w_e \in \mathcal{W}(\tau)$ to be the element that assigns $(-1)^{i+1}$ to e_i for $i = 1, \dots, k-1$, assigns 1 to e , and assigns 0 to all other edges.

We now introduce the operation of admissible deletions. This is a general way of modifying train tracks, but is particularly useful for proving Proposition 5.5 when applied within the realm of floral train tracks, as we will see.

Construction 5.8. Let τ be a train track and e be an edge of τ . If e is not the only half-edge at both of its end points, then we can delete e from τ to get a new train track τ' (see Figure 11). We call this operation an *admissible deletion* (of the edge e).

Suppose e meets boundary components c_1 and c_2 of τ . If $c_1 \neq c_2$, then deleting e combines c_1 and c_2 into one boundary component c' of τ' . If c_i is n_i -pronged, then c' is $(n_1 + n_2 - 2)$ -pronged. On the other hand, if $c_1 = c_2 =: c$, then deleting e splits c into two boundary components c'_1 and c'_2 . In this case, if c'_i is n'_i -pronged, then c is $(n'_1 + n'_2 + 2)$ -pronged.




Figure 10. An odd floral train track.




Figure 11. Admissible deletion of an edge from a train track.

If w' is a system of weights on τ' , then we can define a system of weights w on τ by $w(e') = w'(e')$ for $e' \neq e$ and $w(e) = 0$. Conversely, if w is a system of weights on τ with $w(e) = 0$, then by restricting w to the remaining edges, we get a system of weights w' on τ' . This allows us to identify $\mathcal{W}(\tau')$ as a subspace of $\mathcal{W}(\tau)$. Moreover, this inclusion preserves the Thurston symplectic form ω .

Our next task is to show that admissible deletion of a real edge from an odd floral train track preserves (5.1). We split into two cases according to whether the admissible deletion combines two boundary components into one (Lemma 5.9) or splits one boundary component into two (Lemma 5.10).

Lemma 5.9. *Suppose τ and τ' are odd floral train tracks where τ' is obtained from admissible deletion of a real edge e from τ . Suppose deleting e combines two boundary components into one. Then (5.1) holds for τ if and only if it holds for τ' .*

Proof. There are two cases here. Case 1 is if at least one of c_i , say c_1 , is even-pronged. In this case, we claim that $\mathcal{W}(\tau) = \mathcal{W}(\tau') \oplus \langle r_{c_1} \rangle$. Indeed, since $c_1 \neq c_2$, $r_{c_1}(e) \neq 0$, so $\mathcal{W}(\tau') \cap \langle r_{c_1} \rangle = 0$. Meanwhile, $\dim(\mathbb{R}^{\mathcal{E}(\tau)}/\mathbb{R}^{\mathcal{E}(\tau')}) = 1$ and $\mathcal{W}(\tau)/\mathcal{W}(\tau')$ can be identified with a subspace of $\mathbb{R}^{\mathcal{E}(\tau)}/\mathbb{R}^{\mathcal{E}(\tau')}$, so $\dim(\mathcal{W}(\tau)/\mathcal{W}(\tau')) \leq 1$, which proves the claim.

To prove the lemma in this case, first suppose that (5.1) holds for τ . Then it follows from the claim for any $w'_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau')$, $w'_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau)$ as well. Hence, $w'_0 = \sum_c a_c r_c$. But $w'_0(e) = 0$, so we have $a_{c_1} = 0$ if c_2 is odd-pronged, and $a_{c_1} = a_{c_2}$ if c_2 is even-pronged (under an appropriate choice of signs for r_{c_i}). Together with the fact that $\pm r_{c'} = r_{c_1} + r_{c_2}$ when c_2 is even-pronged, this shows that $w'_0 = \sum_{c'} a_{c'} r_{c'}$ in $\mathcal{W}(\tau')$.

Conversely, suppose that (5.1) holds for τ' . Then for any $w_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau)$, we can consider $w_0 - w_0(e)r_{c_1}$. This lies in $\mathcal{W}(\tau')$ hence lies in $\text{rad}(\omega)$ in $\mathcal{W}(\tau')$. By hypothesis, we then have $w_0 - w_0(e)r_{c_1} \in \text{span}\{r_{c'}\}$, which together with the fact again that $\pm r_{c'} = r_{c_1} + r_{c_2}$ when c_2 is even-pronged, we have $w_0 \in \text{span}\{r_c\}$ in $\mathcal{W}(\tau)$.

Case 2 is if both c_i are odd-pronged. In this case c' is even-pronged. It can be shown by the same reasoning as in the last case that $\mathcal{W}(\tau) = \mathcal{W}(\tau') \oplus \langle w_e \rangle$, where w_e is defined in Construction 5.7.

We compute $\omega(r_{c'}, w)$ for $w \in \mathcal{W}(\tau)$. Suppose first for simplicity that e meets c' in two of its cusps p and p' . For $i = 1, 2$, let e_i be the half-edge at p which is adjacent to e on c_i , and let e'_i be the half-edge at p' which is adjacent to e on c_i . Without loss of generality suppose that e_1 lies to the left of e and $r_{c'}$ assigns $(-1)^i$ to e_i , then since c_i are odd-pronged, $r_{c'}$ assigns $(-1)^{i+1}$ to e'_i (see Figure 12, left).

By the computation made in Proposition 5.2, the total contribution to $\omega(r_{c'}, w)$ from the vertices on c' aside from p and p' is $w(e_1) + w(e_2) + w(e'_1) + w(e'_2)$. The contribution from p is $-w(e) - w(e_2) - w(e) - w(e_1)$, and the contribution from p' is $-w(e) - w(e'_2) - w(e) - w(e'_1)$. Adding these together, we see that $\omega(r_{c'}, w) = -4w(e)$.




Figure 12. The situation in case 2 of Lemma 5.9. Left: if e meets c' in two of its cusps. Right: if e meets c' away from its cusps.

If one of the endpoints of e lies in the complement of the cusps on c' instead, we define e_i as above, but with p or p' being the endpoint of e . Suppose p is such an endpoint, say e_1 does not determine a cusp with e on c_1 , and suppose that $r_{c'}$ assigns 1 to e_1 and e_2 (see Figure 12, right). Then the total contribution to $\omega(r_{c'}, w)$ from the vertices on c' aside from p and p' is $-w(e_1) + w(e_2) + w(e'_1) + w(e'_2)$. The contribution from p is $(-\sum_{e \in \mathcal{E}_p^\beta} w(e) + w(e_1)) + (\sum_{e \in \mathcal{E}_p^\beta} w(e) - w(e_2) - 2w(e))$, and the contribution from p' is $-w(e) - w(e'_2) - w(e) - w(e'_1)$. So we still have $\omega(r_{c'}, w) = -4w(e)$.

We are now ready to prove the lemma in this case. Suppose (5.1) holds for τ . Let $w'_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau')$. Let $a = \omega(w'_0, w_e)$, then $w'_0 + \frac{a}{4}r_{c'} \in \text{rad}(\omega)$ in $\mathcal{W}(\tau')$ and $\omega(w'_0 + \frac{a}{4}r_{c'}, w_e) = a + \frac{a}{4}(-4w_e(e)) = 0$, so $w'_0 + \frac{a}{4}r_{c'} \in \text{rad}(\omega)$ in $\mathcal{W}(\tau)$, implying that $w'_0 \in \text{span}\{r_c\} - \frac{a}{4}r_{c'} \subset \text{span}\{r_{c'}\}$ in $\mathcal{W}(\tau')$.

Conversely, suppose (5.1) holds for τ' . Let $w_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau)$. Then $0 = \omega(r_{c'}, w_0) = -4w_0(e)$. Hence, we can treat w_0 as an element of $\mathcal{W}(\tau')$, hence an element of $\text{rad}(\omega)$ in $\mathcal{W}(\tau')$. Thus, $w_0 \in \text{span}\{r_{c'}\}$ in $\mathcal{W}(\tau')$. Say $w_0 = \sum_{c'} a_{c'}r_{c'}$. To establish that $w_0 \in \text{span}\{r_c\}$ in $\mathcal{W}(\tau)$, we need to show that $a_{c'} = 0$. This follows since $0 = \omega(w_0, w_e) = \omega(a_{c'}r_{c'}, w_e) = -4a_{c'}$. ■

Lemma 5.10. *Suppose τ and τ' are odd floral train tracks where τ' is obtained from an admissible deletion of a real edge e from τ . Suppose deleting e splits a boundary component into two. Then (5.1) holds for τ if and only if it holds for τ' .*

Proof. There are three cases here. The proof of each case is similar to one of the cases in Lemma 5.9.

Case 1 is if c is odd-pronged. In this case one of c'_i , say c'_1 , is even-pronged while the other is odd-pronged.

We follow the strategy of case 2 in Lemma 5.9. It can be shown as before that $\mathcal{W}(\tau) = \mathcal{W}(\tau') \oplus \langle w_e \rangle$. Also, one can compute that $\omega(r_{c'_1}, w) = -2w(e)$ for $w \in \mathcal{W}(\tau)$ (for an appropriate choice of sign for $r_{c'_1}$).

Suppose (5.1) holds for τ . Let $w'_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau')$. Let $a = \omega(w'_0, w_e)$, then $w'_0 + \frac{a}{2}r_{c'_1} \in \text{rad}(\omega)$ in $\mathcal{W}(\tau')$ and $\omega(w'_0 + \frac{a}{2}r_{c'_1}, w_e) = a + \frac{a}{2}(-2w_e(e)) = 0$, so $w'_0 + \frac{a}{2}r_{c'_1} \in \text{rad}(\omega)$ in $\mathcal{W}(\tau)$, implying that $w'_0 \in \text{span}\{r_c\} - \frac{a}{2}r_{c'_1} \subset \text{span}\{r_{c'}\}$ in $\mathcal{W}(\tau')$.

Conversely, suppose (5.1) holds for τ' . Let $w_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau)$. Then $0 = \omega(r_{c'_1}, w_0) = -2w_0(e)$. Hence, we can treat w_0 as an element of $\mathcal{W}(\tau')$, hence an element of $\text{rad}(\omega)$ in $\mathcal{W}(\tau')$. Thus, $w_0 \in \text{span}\{r_{c'_1}\}$ in $\mathcal{W}(\tau')$. Say $w_0 = \sum_{c'} a_{c'} r_{c'}$. To establish that $w_0 \in \text{span}\{r_c\}$ in $\mathcal{W}(\tau)$, we need to show that $a_{c'_1} = 0$. This follows since $0 = \omega(w_0, w_e) = \omega(a_{c'_1} r_{c'_1}, w_e) = -2a_{c'_1}$.

Case 2 is if c is even-pronged while c'_1 and c'_2 are odd-pronged.

We follow the strategy of case 1 in Lemma 5.9. It can be shown as before that $\mathcal{W}(\tau) = \mathcal{W}(\tau') \oplus \langle r_c \rangle$, since c'_i being odd-pronged implies that $r_c(e) = 2$ (under an appropriate choice of sign for r_c).

Suppose that (5.1) holds for τ . For any $w'_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau')$, $w'_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau)$ as well. Hence, $w'_0 = \sum_c a_c r_c$. But $w'_0(e) = 0$ so we have $a_c = 0$. This shows that $w'_0 = \sum_{c'} a_{c'} r_{c'}$ in $\mathcal{W}(\tau')$.

Conversely, suppose that (5.1) holds for τ' . Then for any $w_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau)$, we can consider $w_0 - \frac{w_0(e)}{2}r_c$. This lies in $\mathcal{W}(\tau')$ hence lies in $\text{rad}(\omega)$ in $\mathcal{W}(\tau')$. By hypothesis, we then have $w_0 - \frac{w_0(e)}{2}r_c \in \text{span}\{r_{c'}\}$, which gives $w_0 \in \text{span}\{r_c\}$ in $\mathcal{W}(\tau)$.

Finally, case 3 is if c is even-pronged while c'_1 and c'_2 are even-pronged.

We follow the strategy of case 2 in Lemma 5.9. It can be shown as before that $\mathcal{W}(\tau) = \mathcal{W}(\tau') \oplus \langle w_e \rangle$. Also, one can compute that $\omega(r_{c'_1}, w) = -2w(e)$ for $w \in \mathcal{W}(\tau)$. Using the fact that $r_c = r_{c'_1} + r_{c'_2}$, we have $\omega(r_{c'_2}, w) = 2w(e)$ (for appropriate signs for $r_c, r_{c'_1}, r_{c'_2}$).

Suppose (5.1) holds for τ . Let $w'_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau')$. Let $a = \omega(w'_0, w_e)$, then $w'_0 + \frac{a}{2}r_{c'_1} \in \text{rad}(\omega)$ in $\mathcal{W}(\tau')$ and $\omega(w'_0 + \frac{a}{2}r_{c'_1}, w_e) = a + \frac{a}{2}(-2w_e(e)) = 0$, so $w'_0 + \frac{a}{2}r_{c'_1} \in \text{rad}(\omega)$ in $\mathcal{W}(\tau)$, implying that $w'_0 \in \text{span}\{r_c\} - \frac{a}{2}r_{c'_1} \subset \text{span}\{r_{c'}\}$ in $\mathcal{W}(\tau')$.

Conversely, suppose (5.1) holds for τ' . Let $w_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau)$. Then $0 = \omega(r_{c'_1}, w_0) = -2w_0(e)$. Hence, we can treat w_0 as an element of $\mathcal{W}(\tau')$, hence an element of $\text{rad}(\omega)$ in $\mathcal{W}(\tau')$. Thus, $w_0 \in \text{span}\{r_{c'_1}\}$ in $\mathcal{W}(\tau')$. Say $w_0 = \sum_{c'} a_{c'} r_{c'}$. To establish that $w_0 \in \text{span}\{r_c\}$ in $\mathcal{W}(\tau)$, we need to show that $a_{c'_1} = a_{c'_2}$. This follows since $0 = \omega(w_0, w_e) = \omega(a_{c'_1} r_{c'_1} + a_{c'_2} r_{c'_2}, w_e) = -2a_{c'_1} + 2a_{c'_2}$. ■

Now notice that if τ_1 and τ_2 are two floral train tracks whose unique infinitesimal polygons have the same number of prongs, then they can be related by a sequence of admissible deletion of real edges. This is because upon fixing an identification of their infinitesimal polygons, one can first add all the real edges of τ_2 to τ_1 (which is the reverse of deleting those edges), with the cyclic ordering of half-edges determined by, say, placing all the half-edges of τ_2 to the left of those of τ_1 , then deleting the real edges of τ_1 . See Figure 13 for an example of this procedure.

Hence, Lemmas 5.9 and 5.10 implies that in order to prove Proposition 5.5 in this case, we can simply establish (5.1) for one floral train track with an n -pronged infinitesimal polygon for every odd n .




Figure 13. Any two floral train tracks whose unique infinitesimal polygons have the same number of prongs are related by a sequence of removals of real edges.

To this end, we choose the floral train track τ_n illustrated in Figure 14 left. That is, τ_n has an n -pronged infinitesimal polygon, say with vertices labeled $v_1, \dots, v_n$ in a cyclic way, and n real edges $e_1, \dots, e_n$, each e_i having endpoints on v_i and v_{i+1} , with e_{i-1} lying to the left of e_i at v_i for every i . Then τ_n has $n + 2$ boundary components, two of them n -pronged and n of them 0-pronged. It is straightforward to check that $\mathcal{W}(\tau_n)$ is generated by $\{r_c\}$ for all the 0-pronged boundary components c . Hence, $\text{rad}(\omega) = \text{span}\{r_c\}$ must hold. This concludes the proof of Proposition 5.5 in this case.

5.3. Case 2: All punctures are even-pronged

As explained under Proposition 5.5, the proof in this case is very similar to the last case. The only difference is the preparatory step of taking a twofold cover.

Definition 5.11. A train track is said to be *orientable* if its edges can be oriented in a way such that at each switch v , all the edges in one $\mathcal{E}_v^\beta$ are oriented into v while all the edges in $\mathcal{E}_v^{\beta+1}$ are oriented out of v .

For example, the train tracks τ_{x_I, x_O} we have been considering are orientable if and only if the unstable foliation ℓ^u is orientable in the usual sense.

Now, up to passing to a twofold cover, ℓ^u can always be made orientable. More specifically, one can define a 1-cocycle $\alpha \in H^1(S, \mathbb{Z}/2)$ by $\alpha(\gamma) = 0$ if and only if ℓ^u is




Figure 14. The train tracks we use to demonstrate (5.1) at the end of our modifications.

orientable in a neighborhood of a curve γ . Then the *orientable twofold cover* $\tilde{S}$ is determined by α . By embedding τ in S as usual, α also determines the *orientable twofold cover* $\tilde{\tau} \rightarrow \tau$.

The special property when there are no odd-pronged punctures is that $\alpha(c) = 0$ for every boundary component c of τ . So if $\tilde{\tau}$ is the orientable twofold cover, then every boundary component of τ lifts homomorphically to $\tilde{\tau}$. Correspondingly, we have the 2-to-1 map $\tilde{\mathcal{X}} \rightarrow \mathcal{X}$. In particular, we can lift a partition $\mathcal{X} = \mathcal{X}_I \sqcup \mathcal{X}_O$ to $\tilde{\mathcal{X}} = \tilde{\mathcal{X}}_I \sqcup \tilde{\mathcal{X}}_O$.

By naturality of Construction 3.6, $\tau_{(\tilde{\mathcal{X}}_I, \tilde{\mathcal{X}}_O)}$ is the orientable twofold cover of $\tau_{(\mathcal{X}_I, \mathcal{X}_O)}$. Hence by the following lemma, we can assume that the $\tau_{(\mathcal{X}_I, \mathcal{X}_O)}$ we are dealing with is orientable.

Lemma 5.12. *Let $\pi : \tilde{\tau} \rightarrow \tau$ be a finite normal covering of train tracks such that each boundary component of τ lifts homeomorphically to $\tilde{\tau}$. Then (5.1) holds for τ if (5.1) holds for $\tilde{\tau}$.*

Proof. Let d be the degree of the covering and let G be the group of deck transformations. We can define operators $\pi_* : \mathcal{W}(\tilde{\tau}) \rightarrow \mathcal{W}(\tau)$, $\pi^* : \mathcal{W}(\tau) \rightarrow \mathcal{W}(\tilde{\tau})$, and $s : \mathcal{W}(\tilde{\tau}) \rightarrow \mathcal{W}(\tilde{\tau})$ by

$$\begin{aligned} (\pi_*(\tilde{w}))(e) &= \sum_{\pi(\tilde{e})=e} \tilde{w}(e) \\ (\pi^*(w))(\tilde{e}) &= w(\pi(\tilde{e})) \\ (s(\tilde{w}))(\tilde{e}) &= \sum_{g \in G} \tilde{w}(g\tilde{e}). \end{aligned}$$

We have the following properties:

- $\pi^*(r_c) = \sum_{\pi(\tilde{c})=c} r_{\tilde{c}}$
- $\pi_*\pi^*(w) = dw$
- $\omega(w_1, \pi_*(\tilde{w}_2)) = \omega(\pi^*(w_1), \tilde{w}_2)$
- $s\pi^*(w) = d\pi^*(w)$

which imply that π^* is injective and $\pi^*(\text{rad}(\omega)) \subset \text{rad}(\omega)$.

Now suppose that (5.1) holds for $\tilde{\tau}$. Let $w_0 \in \text{rad}(\omega)$ in $\mathcal{W}(\tau)$. $\pi^*(w_0) \in \text{rad}(\omega)$ so $\pi^*(w_0) = \sum_{\tilde{c}} a_{\tilde{c}} r_{\tilde{c}}$. Hence

$$d\pi^*(w_0) = s\pi^*(w_0) = \sum_{\tilde{c}} \left(\sum_{g \in G} a_{g\tilde{c}} \right) r_{\tilde{c}} = \pi^* \left(\sum_c \left(\sum_{\pi(\tilde{c})=c} a_{\tilde{c}} \right) r_c \right),$$

which implies that $w_0 = \frac{1}{d} \sum_c (\sum_{\pi(\tilde{c})=c} a_{\tilde{c}}) r_c \in \text{span}\{r_c\}$ by injectivity of π^* . ■

Together with Lemma 4.8, we can assume that $\tau_{(\mathcal{X}_I, \mathcal{X}_O)}$ is orientable and floral (which implies that the infinitesimal polygon c_I is even-pronged), or *orientable floral* for short. See Figure 15 for an example of such a train track. In this case, if we label the vertices of




Figure 15. An orientable floral train track.

c_I by $v_1, \dots, v_n$ in a counterclockwise order, then by orientability, the endpoints on each real edge must lie on v_i and v_j for i, j of different parity.

Construction 5.7 can be repeated for orientable floral train tracks word-by-word. Notice here that we require τ to be orientable in order for k to be even, in the notation of Construction 5.7. Also notice that in this case both cyclic orderings give an even k . The w_e defined under the two choices will differ by r_{c_I} . This is not very significant for our purposes, but for concreteness we can simply fix some choice for each real edge e .

The proofs of Lemma 5.9 and Lemma 5.10 then carry over word for word to show the following lemma.

Lemma 5.13. *Suppose τ and τ' are orientable floral train tracks where τ' is obtained from an admissible deletion of a real edge e from τ . Then (5.1) holds for τ if and only if it holds for τ' .*

Hence by the reasoning at the end of the last subsection, we can simply establish (5.1) for one orientable floral train track with an n -pronged infinitesimal polygon for every even n .

To this end, we choose the floral train track τ_n illustrated in Figure 14 right. That is, τ_n has an n -pronged infinitesimal polygon, say with vertices labeled $v_1, \dots, v_n$ in a cyclic way, and n real edges $e_1, \dots, e_n$, each e_i having endpoints on v_i and v_{i+1} , with e_{i-1} lying to the left of e_i at v_i for every i . Then τ_n has $n + 2$ boundary components, two of them n -pronged and n of them 0-pronged. It is straightforward to check that $W(\tau_n)$ is generated by $\{r_c\}$ for all these boundary components c . Hence, $\text{rad}(\omega) = \text{span}\{r_c\}$ must hold. This concludes the proof of Proposition 5.5.

6. Proof of the main theorem

We gather all the ingredients to prove our main theorem, which we restate below for the reader's convenience.

Theorem 6.1. *Let $f : S \rightarrow S$ be a fully punctured pseudo-Anosov map with at least two puncture orbits. Then the normalized expansion factor $L(S, f) = \lambda(f)^{|\chi(S)|}$ satisfies the inequality*

$$L(S, f) \geq \mu^4.$$

More precisely, for $|\chi(S)| \geq 3$, we have

$$\lambda(f) \geq |LT_{1, \frac{K}{2}}|$$

if $|\chi(S)| = K$ is even, and

$$\lambda(f)^K \geq 8$$

if $|\chi(S)| = K$ is odd.

Proof. Take some partition $\mathcal{X} = \mathcal{X}_I \sqcup \mathcal{X}_O$ of the set of punctures of S into two nonempty f -invariant sets. Consider the standardly embedded train track $\tau = \tau(\mathcal{X}_I, \mathcal{X}_O)$, the train track map $f = f(\mathcal{X}_I, \mathcal{X}_O)$, and the matrix $f_*^{\text{real}} = f_{(\mathcal{X}_I, \mathcal{X}_O)*}^{\text{real}}$. Proposition 3.12 shows that f_*^{real} is Perron–Frobenius.

Meanwhile, consider the weight space $\mathcal{W}(\tau)$. We have the commutative diagram

$$\begin{array}{ccccccc} 0 & \longrightarrow & \text{rad}(\omega) & \longrightarrow & \mathcal{W}(\tau) & \longrightarrow & \mathcal{W}(\tau)/\text{rad}(\omega) \longrightarrow 0 \\ & & \downarrow \mathcal{W}(f) & & \downarrow \mathcal{W}(f) & & \downarrow \mathcal{W}(f) \\ 0 & \longrightarrow & \text{rad}(\omega) & \longrightarrow & \mathcal{W}(\tau) & \longrightarrow & \mathcal{W}(\tau)/\text{rad}(\omega) \longrightarrow 0 \end{array}.$$

Propositions 5.4 and 5.5 imply that the restriction of $\mathcal{W}(f)$ to $\text{rad}(\omega)$ is reciprocal. On the other hand, $\mathcal{W}(\tau)/\text{rad}(\omega)$ inherits the form ω , which is now symplectic. The induced map of $\mathcal{W}(f)$ on $\mathcal{W}(\tau)/\text{rad}(\omega)$ preserves ω hence is symplectic, thus reciprocal by Proposition 2.13 (1). By Proposition 2.13 (3), $\mathcal{W}(f) : \mathcal{W}(\tau) \rightarrow \mathcal{W}(\tau)$ is reciprocal.

We also have the following commutative diagram from Proposition 4.2:

$$\begin{array}{ccccccc} 0 & \longrightarrow & \mathcal{W}(\tau) & \longrightarrow & \mathbb{R}^\mathcal{E} & \xrightarrow{T_V} & T_V(\mathbb{R}^\mathcal{E}) \longrightarrow 0 \\ & & \downarrow \mathcal{W}(f) & & \downarrow f_* & & \downarrow P|_{T_V(\mathbb{R}^\mathcal{E})} \\ 0 & \longrightarrow & \mathcal{W}(\tau) & \longrightarrow & \mathbb{R}^\mathcal{E} & \xrightarrow{T_V} & T_V(\mathbb{R}^\mathcal{E}) \longrightarrow 0 \end{array}.$$

We have deduced that $\mathcal{W}(f) : \mathcal{W}(\tau) \rightarrow \mathcal{W}(\tau)$ is reciprocal above. By Proposition 2.13 (2), $P|_{T_V(\mathbb{R}^\mathcal{E})}$ is reciprocal. So by Proposition 2.13 (3), f_* is reciprocal.

Finally, by the discussion above Proposition 3.12, f_* being reciprocal implies that f_*^{real} is reciprocal. Hence, f_*^{real} is a $|\chi(S)|$ -by- $|\chi(S)|$ reciprocal Perron–Frobenius matrix.

By Proposition 3.13, $\lambda(f)$ is the spectral radius of f_*^{real} . Hence, the theorem follows from Theorem 2.12. $\blacksquare$

7. Sharpness of the main theorem

In this section, we discuss the sharpness of the main theorem. We give two families of pseudo-Anosov maps realizing the lower bounds in Theorem 1.9 for even $\chi(S)$. We do this in two ways. In Section 7.1, we describe folding sequences of train tracks which determine the maps. In Section 7.2, we describe classes in certain fibered faces which determine the same maps. In Section 7.3, we give examples showing that Theorems 1.11 and 1.9 for odd $\chi(S)$ are not sharp in general. In Section 7.4, we give examples showing that the assumption of f having at least two puncture orbits is necessary in Theorem 1.1.

7.1. Examples for even $\chi(S)$: Train tracks

In this subsection, we will show that Theorem 1.9 is sharp in the cases when $\chi(S)$ is even, by demonstrating *folding sequences* of standardly embedded train tracks, that is, sequences of the form $\tau_0 \xrightarrow{f_1} \cdots \xrightarrow{f_n} \tau_n \xrightarrow{\sigma} \tau_0$ where each f_i is the composition of a subdivision move and an elementary folding move involving one of the subdivided edges, and σ is an isomorphism of train tracks.

Given such a sequence, the induced map $\sigma f_n \cdots f_1$ on the tie neighborhood N of τ_1 will determine a mapping class on the punctured surface S that is the interior of N . If the real transition matrix f_*^{real} of $\sigma f_n \cdots f_1$ is Perron–Frobenius, then this mapping class contains a (unique) fully punctured pseudo-Anosov map $f : S \rightarrow S$, for which τ_0 is an invariant train track. For a more detailed explanation of recovering pseudo-Anosov maps from folding sequences of train tracks, see for example [7].

As described in Section 3, one can then compute the expansion factor $\lambda(f)$ of f as the spectral radius of f_*^{real} . Here, we will make this computation using the method described in [30]. Namely, we write down the directed graph associated with f_*^{real} , compute its curve complex G , then compute the clique polynomial $Q_G(t)$ of G . Theorems 1.2 and 1.4 of [30] state that the smallest positive root of $Q_G(t)$ is equal to $\frac{1}{\lambda(f)}$. Here, as shown in the proof of Theorem 1.1, f_*^{real} will be reciprocal, so we can more directly compute $\lambda(f)$ as the largest positive root of $Q_G(t)$.

Our first family of examples is shown in Figure 16. Here, each train track has a single infinitesimal polygon with $3k$ cusps, which we have used as the center of reference, and $2k$ real edges. We indicated each fold f_i by highlighting the relevant edges in bright red (before) and dark red (after). The train track isomorphism σ is induced by a rotation of the center infinitesimal polygon.

When $k = 2$, one computes the real transition matrix to be

$$f_*^{\text{real}} = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 2 \\ 0 & 1 & 0 & 0 \end{bmatrix}$$




Figure 16. Folding sequence of train tracks in the first family of examples.

and when $k \geq 3$, one computes it to be

$$f_*^{\text{real}} = \left[\begin{array}{cc|c|cccc} 0 & 0 & I_{2k-6} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ \hline 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \end{array} \right],$$

where I_n denotes the n by n identity matrix.

The corresponding directed graphs and curve complexes for $k = 2$ and $k \geq 3$ are shown in Figure 17, top and bottom, respectively. Here a number n besides a directed edge $\rightarrow$ is shorthand for n consecutive edges $\rightarrow \cdots \rightarrow$. Meanwhile, a number besides a vertex in the curve complex denotes its weight.

For each $k \geq 2$, the clique polynomial of the corresponding curve complex is $LT_{1,k} = t^{2k} - t^{k+1} - t^k - t^{k-1} + 1$. Hence, the induced fully punctured pseudo-Anosov maps in this family of examples have expansion factors $|LT_{1,k}|$.




Figure 17. The associated directed graphs to the real transition matrices of Figure 16 and their curve complexes.

Notice that this first family of examples already shows the sharpness statement in Theorem 1.9. However, we are actually able to find a second family of examples that attain equality in Theorem 1.9 for even $|\chi(S)|$. We will describe this second family next, following the same format.

Unfortunately, for number theoretical reasons, it is difficult to present a single picture that illustrates all the members of this second family; we need to split into five subcases depending on the value of $k \pmod{5}$, where $|\chi(S)| = 2k$.

For $k \equiv 3 \pmod{5}$, the folding sequence is shown in Figure 18.

For consistency, we have drawn Figure 18 in the style as Figure 16. Namely, we used the single infinitesimal polygon as the center of reference, we highlighted each fold in red, and the train track isomorphism is induced by a rotation of the infinitesimal polygon.

However, we caution that there are a few differences: This time the infinitesimal polygon has $2k + 1$ cusps (but there are still $2k$ real edges). Real edges that connect a shown cusp to a cusp in the $\cdots$ range are truncated. Also, the number of cusps in each $\cdots$ differs. The fact that the last map is a train track isomorphism will determine how these real edges should be connected and how many cusps each $\cdots$ should contain.

For $k \equiv 4, 0$, and $1 \pmod{5}$, the folding sequences are shown in Figure 19. These are drawn in the same style as Figure 18.




Figure 18. Folding sequence of train tracks in the second family of examples.

Finally, for $k \equiv 2 \pmod{5}$, the folding sequence is shown in Figure 20.

Figure 20 is drawn in a slightly different style from the previous pictures by necessity. There are now five infinitesimal polygons, each having $\frac{2k+1}{5}$ cusps. The train track isomorphism now permutes the five infinitesimal polygons; we have arranged these so that this permutation is induced by a rotation. However, we caution that the isomorphism is not just a rotation of the picture; one also needs to rotate one of the infinitesimal polygons.

The real transition matrices of these subcases admit a more consistent description. When $k = 2$, we have

$$f_*^{\text{real}} = \begin{bmatrix} 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}$$




Figure 19. Folding sequence of train tracks in the second family of examples (part I).




Figure 19. Folding sequence of train tracks in the second family of examples (part II).

and when $k \geq 3$, we have

$$f_{*}^{\text{real}} = \left[\begin{array}{ccc|c|ccc} 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & I_{2k-6} & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \end{array} \right]$$

The corresponding directed graphs and curve complexes for $k = 2$ and $k \geq 3$ are shown in Figure 21, top and bottom, respectively.

For each $k \geq 2$, the clique polynomial of the corresponding curve complex is $LT_{1,k} = t^{2k} - t^{k+1} - t^k - t^{k-1} + 1$. Hence, the induced fully punctured pseudo-Anosov maps in this family of examples have expansion factors $|LT_{1,k}|$.

We point out that similar train track maps were obtained in [22]. In particular, there is some overlap between our first family and [22, Example 4.2] and between our second family and [22, Example 4.4]. See also the discussion in the next subsection.




Figure 20. Folding sequence of train tracks in the second family of examples.

7.2. Examples for even $\chi(S)$: Fibered face theory

In this subsection, we will show the sharpness statement in Theorem 1.9 again, but this time using the tool of fibered face theory. This will also explain the source of our train track maps in Section 7.1.

We first provide a brief review of fibered face theory, referring to [29] for details. Let M be a fibered hyperbolic 3-manifold, with $n = b_1(M) \geq 2$ and let $F \subset H^1(M; \mathbb{R})$ be a fibered face. To every primitive integral point $a \in \text{cone}(F)$, there is a fibration of M over S^1 whose monodromy is a pseudo-Anosov mapping class (S_a, f_a) .

Let $\Delta_M \in \mathbb{Z} H_1(M)$ be the Alexander polynomial of M . The *Newton polytope* of Δ_M is the convex hull of points in $H_1(M)$ that have nonzero coefficient in Δ_M . The *Alexander norm* of an element $a \in H^1(M)$ is defined to be

$$\|a\| = \sup_{g, h} \langle a, g - h \rangle,$$

where g, h range over the Newton polytope of Δ_M .

It is a well-known fact that the Euler characteristic of S_a can be calculated by

$$|\chi(S_a)| = \|a\|.$$




Figure 21. The associated directed graphs to the real transition matrices of Figures 18–20, and their curve complexes.

In other words, the Thurston norm agrees with the Alexander norm on $\text{cone}(F)$. See, for example, [29, Theorem 7.1].

Meanwhile, there is another polynomial $\Theta_F \in \mathbb{Z}H_1(M)$, called the *Teichmüller polynomial* associated with the fibered face F , with the property that the expansion factor of the pseudo-Anosov monodromy (S_a, f_a) corresponding to $a = (a_1, \dots, a_n)$ can be calculated by

$$\lambda(f_a) = |\Theta_F(t^{a_1}, \dots, t^{a_n})|.$$

See, for example, [29, Theorem 5.1].

Now consider the L6a2 link complement, which we denote by M_1 . It can be computed (e.g., by using the veering triangulation eLMkbcdddddde_2100, see [17, 31] for the details of such a computation) that the Alexander polynomial of M_1 is

$$\Delta_1(a, b) = b^2 + b(a^2 - a + 1) + a^2$$

under some choice of basis (a, b) for $H_1(M_1)$, and the Teichmüller polynomial associated with the fibered cone $C_1 = \text{cone}\{a^* + 2b^*, -a^*\}$ is

$$\Theta_1(a, b) = b^2 - b(a^2 + a + 1) + a^2.$$

Hence, the Thurston norm on C_1 is given by $-2a + 2b$.

Consider the class $x_1 = b^*$. The corresponding pseudo-Anosov monodromy $f_{1,1}$ is defined on a surface S_1 with $|\chi(S_1)| = \|x_1\| = 2$. The expansion factor of $f_{1,1}$ is the largest root of $t^2 - 3t + 1$, which is μ^2 . Hence, $f_{1,1}$ attains equality in Theorem 1.1.

Now consider the class $x_k = a^* + (k+1)b^* \in C_1$, for $k \geq 2$. The corresponding pseudo-Anosov monodromy $f_{1,k}$ is defined on a surface S_k with $|\chi(S_k)| = \|x_k\| = 2k$ and its expansion factor is the largest root of $t^{2k+2} - t^{k+1}(t^2 + t + 1) + t^2 = t^2 L_{T_{1,k}}$, which is $|L_{T_{1,k}}|$. Hence, $f_{1,k}$ attains equality in Theorem 1.9.

Indeed, the first family of train track maps we described in Section 7.1 is computed from $f_{1,k}$. Here, we will not demonstrate this computation since the details are rather tedious. It suffices to say that we essentially followed the methodology in [22]; see, in particular, [22, Example 4.2]. We also remark that the maps $f_{1,k}$ are considered in [20] as well, even though invariant train tracks were not provided there.

Similarly, consider the L13n5885 link (= the Whitehead sister link = the $(-2, 3, 8)$ -pretzel link) complement, which we denote by M_2 . It can be computed (e.g., by using the veering triangulation fLLQcbdddeehhbghh_01110) that the Alexander polynomial of M_2 is

$$\Delta_2(a, b) = b^2 + b(a^2 + a + 1) + a^2$$

under some choice of basis (a, b) for $H_1(M_2)$, and the Teichmüller polynomial associated with the fibered cone $C_2 = \text{cone}\{a^* + 2b^*, -a^*\}$ is

$$\Theta_2(a, b) = b^2 - b(a^2 + a + 1) + a^2.$$

Hence, the Thurston norm on C_2 is given by $-2a + 2b$.

Consider the class $x_1 = b^*$. The corresponding pseudo-Anosov monodromy $f_{2,1}$ is defined on a surface S_1 with $|\chi(S_1)| = \|x_1\| = 2$. The expansion factor of $f_{2,1}$ is the largest root of $t^2 - 3t + 1$, which is μ^2 . Hence, $f_{2,1}$ attains equality in Theorem 1.1.

Now consider the class $x_k = a^* + (k+1)b^* \in C_2$, for $k \geq 2$. The corresponding pseudo-Anosov monodromy $f_{2,k}$ is defined on a surface S_k with $|\chi(S_k)| = \|x_k\| = 2k$ and its expansion factor is the largest root of $t^{2k+2} - t^{k+1}(t^2 + t + 1) + t^2 = t^2 L_{T_{1,k}}$, which is $|L_{T_{1,k}}|$. Hence, $f_{2,k}$ attains equality in Theorem 1.9.

The second family of train track maps we described in Section 7.1 is computed from $f_{2,k}$. We refer to [22, Example 4.4] for the methodology of our computation. We also remark that some of the maps $f_{2,k}$ were considered in [1, 24], even though invariant train tracks were not provided in those works.

We summarize all the examples we have discussed in Table 1. In the table, we list a puncture as having singularity type p if it has p prongs, and punctures with the same color belong to the same orbit.

Finally, we remark that both M_1 and M_2 can be obtained by Dehn filling a single fibered 3-manifold M , commonly known as the *magic manifold*. Indeed, all the known examples of small expansion factor maps are realized as monodromies of M and its Dehn fillings along a component [23].

g	s	Range of k	Description of f	Singularity type
0	4	–	Fiberings of L6a2	$(1, 1, 1, 1)$
k	2	$k \geq 2, k \equiv 1, 2 \pmod{3}$		$(3k, k)$
$k-1$	4	$k \geq 3, k \equiv 0 \pmod{3}$		$(3k, \frac{k}{3}, \frac{k}{3}, \frac{k}{3})$
1	2	–	Fiberings of L13n5885	$(2, 2)$
k	2	$k \geq 4, k \equiv 0, 1, 4 \pmod{5}$		$(2k+1, 2k-1)$
$k-2$	6	$k \geq 2, k \equiv 2 \pmod{5}$		$(\frac{2k+1}{5}, \frac{2k+1}{5}, \frac{2k+1}{5}, \frac{2k+1}{5}, \frac{2k+1}{5}, 2k-1)$
$k-2$	6	$k \geq 3, k \equiv 3 \pmod{5}$		$(2k+1, \frac{2k-1}{5}, \frac{2k-1}{5}, \frac{2k-1}{5}, \frac{2k-1}{5}, \frac{2k-1}{5})$

Table 1. Examples of maps that attain the lower bound in Theorem 6.1.

7.3. Braids and odd $\chi(S)$

Note that the only braid monodromies that appear in Table 1 are $f_{1,1}$ and $f_{2,2}$. It can be checked that these are the simplest hyperbolic 3-braid $\sigma_1\sigma_2^{-1}$ (see [20]) and the 5-braid of minimal expansion factor $\sigma_1\sigma_2\sigma_3\sigma_4\sigma_1\sigma_2$ (see [19]), respectively. These imply that Theorem 1.11 is sharp for $n = 3, 5$.

We do not currently know whether Theorem 1.11 is sharp for odd $n \geq 7$ (see Question 1.12). Also note that the results of [26] on braids of minimum expansion factor do not provide an answer here, since many of those braids are not fully punctured.

So far we have only discussed sharpness in the cases when $|\chi(S)|$ is even. This is because the cases when $|\chi(S)|$ is odd are likely not sharp. For example, it is shown in [26] that the 6-braid of minimum expansion factor $\sigma_2\sigma_1\sigma_2\sigma_1(\sigma_1\sigma_2\sigma_3\sigma_4\sigma_5)^2$ is fully punctured and has normalized expansion factor $\lambda_{0,7}^5 \approx 15.14 > 8$. See Section 8 for some discussion on how one might try to sharpen the bound for odd $|\chi(S)|$.

7.4. Single orbit of punctures

In this subsection, we describe some examples that show that the assumption of f having at least two punctures orbits in Theorem 1.1 is necessary.

Most of the examples are defined on the once-punctured torus $S_{1,1}$. To describe them, write $S_{1,1}$ as $(\mathbb{R}^2 \setminus \mathbb{Z}^2)/\mathbb{Z}^2$ and notice that an element $A \in SL(2, \mathbb{Z})$ induces a map f_A on $S_{1,1}$. It is a classical fact that if $|\text{tr} A| > 2$, then f_A is pseudo-Anosov with expansion factor given by the spectral radius of A . Meanwhile, since $|\chi(S_{1,1})| = 1$, the normalized expansion factor is equal to the expansion factor, and since $S_{1,1}$ only has one puncture, any map defined on it must only have one puncture orbit.

g	s	Description of f	Singularity type	$L(S, f)$
1	1	Induced by $\pm \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}$	(2)	$\frac{3+\sqrt{5}}{2} \approx 2.62$
1	1	Induced by $\pm \begin{bmatrix} 3 & 2 \\ 1 & 1 \end{bmatrix}$	(2)	$\frac{4+\sqrt{12}}{2} \approx 3.73$
5	1	Fibering of K12n242	(18)	(Lehmer's number) ⁹ ≈ 4.31
1	1	Induced by $\pm \begin{bmatrix} 4 & 3 \\ 1 & 1 \end{bmatrix}$	(2)	$\frac{5+\sqrt{21}}{2} \approx 4.79$
2	1	Fill in 2-pronged puncture of $f_{1,2}$	(6)	$ LT_{1,2} ^3 \approx 5.10$
1	1	Induced by $\pm \begin{bmatrix} 5 & 4 \\ 1 & 1 \end{bmatrix}$	(2)	$\frac{6+\sqrt{32}}{2} \approx 5.83$

Table 2. Examples of fully punctured pseudo-Anosov maps f with only one puncture orbit for which the bound in Theorem 1.1 fails.

With this understanding, notice that the maps induced by $\pm \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}$, $\pm \begin{bmatrix} 3 & 2 \\ 1 & 1 \end{bmatrix}$, $\pm \begin{bmatrix} 4 & 3 \\ 1 & 1 \end{bmatrix}$, and $\pm \begin{bmatrix} 5 & 4 \\ 1 & 1 \end{bmatrix}$ have expansion factors $\frac{3+\sqrt{5}}{2} \approx 2.62$, $\frac{4+\sqrt{12}}{2} \approx 3.73$, $\frac{5+\sqrt{21}}{2} \approx 4.79$, and $\frac{6+\sqrt{32}}{2} \approx 5.83$, respectively, each of them being strictly less than $\mu^4 \approx 6.85$. This shows that the inequality in Theorem 1.1 fails if the assumption “with at least two punctures orbits” is removed.

There are (at least) two other examples that can be used to show this. Consider the K12n242 knot (= the $(-2,3,7)$ -pretzel knot) complement. This 3-manifold has a unique fibering, with monodromy $f : S_{5,1} \rightarrow S_{5,1}$. The expansion factor of f is given by the largest real root of $t^{10} + t^9 - t^7 - t^6 - t^5 - t^4 - t^3 + t + 1$, also known as *Lehmer's number*, which is ≈ 1.18 . Hence, the normalized expansion factor of f is ≈ 4.31 , which is strictly less than $\mu^4 \approx 6.85$.

Meanwhile consider the map $f_{1,2}$ defined in Section 7.2. According to Table 1, it is defined on $S_{2,2}$, where one of the punctures is 6-pronged while the other is 2-pronged. If we fill in the 2-pronged puncture, we would still get a fully punctured pseudo-Anosov map $\overline{f}_{1,2}$ with the same expansion factor but now defined on $S_{2,1}$. Hence, the normalized expansion factor of $\overline{f}_{1,2}$ is $|LT_{1,2}|^3 \approx 5.10$, which is strictly less than $\mu^4 \approx 6.85$.

We summarize these examples in Table 2.

8. Discussion and further questions

We first note the following generalization of Proposition 5.5. (Recall Definition 5.1 for the definition of the radical elements r_c .)

Proposition 8.1. *Suppose τ is a train track that fully carries the unstable lamination of a pseudo-Anosov map. Then the radical of the Thurston symplectic form ω on τ is given by*

$$\text{rad}(\omega) = \text{span}\{r_c\},$$

where c ranges over all even-pronged boundary components of τ .

In particular,

$$\dim \text{rad}(\omega) = \# \text{ even-pronged boundary components} - \varepsilon$$

for $\varepsilon = 0$ or 1 .

Proof. Up to puncturing the pseudo-Anosov map f at an orbit of a (nonsingular) periodic point, we can arrange for f to have at least two puncture orbits. Correspondingly, we can modify τ by slitting a small segment of an edge at every punctured point (see Figure 22).

It is straightforward to check that this slitting operation preserves $\text{rad}(\omega) = \text{span}\{r_c\}$, hence we can assume that f has at least two puncture orbits.

Then by Proposition 5.5, $\text{rad}(\omega) = \text{span}\{r_c\}$ is true for at least one train track τ' fully carrying the unstable lamination of f . Now by a theorem of Stallings [35], τ and τ' are related by a sequence of elementary moves, hence by Lemma 5.6, $\text{rad}(\omega) = \text{span}\{r_c\}$ is true for τ as well.

For the statement about the dimension of $\text{rad}(\omega)$, recall that the proof of Proposition 5.4 shows that the radical elements r_c have at most one relation. ■

Question 8.2. Is the equation $\text{rad}(\omega) = \text{span}\{r_c\}$, where c ranges over all even-pronged boundary components of τ , true for any train track τ ?

Given the invariance of the equation under a wide array of operations on train tracks, it seems reasonable to expect a positive answer. However, if the answer is negative, then Proposition 8.1 would be an obstruction for train tracks to carry the unstable lamination of a pseudo-Anosov map.

Next, we propose some questions about small expansion factors which one might hope to tackle using the ideas in this paper.

Firstly, recall that Table 1 records some maps that attain the lower bound in Theorem 6.1. These being the only examples we are aware of, it is a natural question to ask if they are actually the only possible examples.

Question 8.3. Let $f : S \rightarrow S$ be a fully punctured pseudo-Anosov mapping class with at least two puncture orbits. Suppose $|\chi(S)| = 2k \geq 4$ and $L(f) = |LT_{1,k}|^{2k}$. Must f be one of the maps listed in Table 1?

Recall that for the examples in Table 1, we calculated that $L(S, f) = |LT_{1,k}|^{2k}$ using the techniques described in [30]. More specifically, starting with the matrix f_*^{real} we get from the train track map, we consider its associated digraph, compute its curve complex G , compute the clique polynomial of G , and find its largest real root.



Figure 22. Puncturing at a nonsingular point corresponds to slitting a small segment of an edge.

Now, McMullen has actually classified the weighted graphs G that could occur in this computation, provided that $L(S, f) = |LT_{1,k}|^{2k}$. Hence, one can hope to work backwards from that information and classify the matrices f_*^{real} that could occur in this computation, then classify the train track maps themselves.

If the answer to Question 8.3 is “yes,” then this means that it is possible to sharpen Theorem 1.9 for the other values of (g, s) . A natural candidate here would be to replace $|LT_{1,k}|$ by $|LT_{3,k}|$, since in McMullen’s analysis in [30], this is the second smallest value for the largest real root of the clique polynomial that could occur in the computation (at least for large enough values of k). We remark that the first author has found maps that attain this expansion factor in [20].

If this can be done, then one could repeat Question 8.3 for this second-best bound to ask for the exact range of (g, s) it applies to. Then repeating this scheme as far as possible, one could hope to paint a nice picture to the minimum expansion factor problem when restricted to fully punctured maps.

With our current understanding of Perron–Frobenius matrices (and pseudo-Anosov maps), one can hope to carry out this program for even $|\chi(S)|$. However, the same question for odd $|\chi(S)|$ would likely require new ideas.

Question 8.4. What is the sharpest lower bound for odd $|\chi(S)|$ in Theorem 1.9?

Indeed, this is down to the fact that in Theorem 2.12, the case when n is odd is likely not sharp. This is in turn due to the approach taken in [30], where McMullen simply shows that $\rho(A)^n < 8$ is impossible by showing that there would not be an appropriate curve complex G in the computation, via analyzing some small graphs. To understand the sharpest lower bound possible when n is odd, one would likely have to extend the analysis to some slightly larger graphs.

As another potential direction for the ideas in this paper, we turn our attention to the hypothesis that f has at least two puncture orbits in Theorem 1.1. Notice that all the examples in Table 1 have exactly two puncture orbits. In particular, if the answer to Question 8.3 is “yes,” then it must be possible to sharpen Theorem 1.1 if we replace the hypothesis by, say, “at least three puncture orbits.” In general, we ask the following question.

Question 8.5. What is the sharpest lower bound in Theorem 1.1 if one replaces “at least two puncture orbits” by “at least q puncture orbits” for some number $q \geq 3$?

As before, the approach to answering Question 8.5 that we have in mind is to dive into the analysis of [30]. The condition on the number of puncture orbits should impose some conditions on the Perron–Frobenius digraph, which might translate to some conditions on the corresponding curve complex G . By restricting to graphs satisfying this condition (potentially broadening the analysis if necessary), one can hope to get lower bounds in this scenario.

Another approach would be to generalize the definition of standardly embedded train tracks. For standardly embedded train tracks, the set of punctures is naturally divided

into the sets of inner and outer punctures. If there is a type of train track that naturally treats, say, three classes of punctures distinctly, then the resulting dynamics could also give interesting lower bounds.

We remark that whether the sharpest such lower bound is attained is also an interesting question. Sun [36] showed that there exists 3-manifolds M with fibered faces F such that the minimum normalized expansion factor on F is attained at an irrational point, that is, not attained by a pseudo-Anosov map corresponding to a rational point on F . If this is the case for the 3-manifolds that produce the sharpest lower bound in Question 8.5 for some particular value of q , then such a lower bound will not be attained.

Acknowledgments. We would like to thank Ian Agol and Curtis McMullen for their support and encouragement throughout this project; Luya Wang, Anna Parlak, Saul Schleimer, and Henry Segerman for helpful conversations; and Eiko Kin and Livio Liechti for their comments on an earlier version of this paper. This work was done while the first author was a participant of the Complex Dynamics program at the MSRI and later while she served at the NSF.

Funding. The first author was partially supported by a grant from the Simons Foundation #426722. The second author was partially supported by a grant from the Simons Foundation #376200.

References

- [1] J. W. Aaber and N. Dunfield, [Closed surface bundles of least volume](#). *Algebr. Geom. Topol.* **10** (2010), no. 4, 2315–2342 Zbl [1205.57018](#) MR [2745673](#)
- [2] W. Abikoff, *The real analytic theory of Teichmüller space*. Lecture Notes in Math. 820, Springer, Berlin, 1980 vii+144 pp. Zbl [0452.32015](#) MR [0590044](#)
- [3] I. Agol, [Ideal triangulations of pseudo-Anosov mapping tori](#). In *Topology and geometry in dimension three*, pp. 1–17, Contemp. Math. 560, American Mathematical Society, Providence, RI, 2011 Zbl [1335.57026](#) MR [2866919](#)
- [4] I. Agol, C. J. Leininger, and D. Margalit, [Pseudo-Anosov stretch factors and homology of mapping tori](#). *J. Lond. Math. Soc. (2)* **93** (2016), no. 3, 664–682 Zbl [1388.37033](#) MR [3509958](#)
- [5] I. Agol and C. C. Tsang, Dynamics of veering triangulations: infinitesimal components of their flow graphs and applications. [v1] 2022, [v2] 2024, arXiv:[2201.02706v2](#)
- [6] H. Baik, J. Baik, C. Kim, and P. Tranchida, [Topological entropy of pseudo-Anosov maps on punctured surfaces vs. homology of mapping tori](#). *Glas. Mat. Ser. III* **57** (2022), no. 2, 291–312 Zbl [1512.37042](#) MR [4541300](#)
- [7] M. Bestvina and M. Handel, [Train-tracks for surface homeomorphisms](#). *Topology* **34** (1995), no. 1, 109–140 Zbl [0837.57010](#) MR [1308491](#)
- [8] J. S. Birman, [Braids, links, and mapping class groups: based on lecture notes by James Cannon](#). Annals of Mathematics Studies, No. 82, Princeton University Press, Princeton, NJ; University of Tokyo Press, Tokyo, 1975, IX, 229 p. Zbl [0305.57013](#) MR [0375281](#)

- [9] P. Brinkmann, [An implementation of the Bestvina-Handel algorithm for surface homeomorphisms](#). *Exp. Math.* **9** (2000), no. 2, 235–240 Zbl [0982.57005](#) MR [1780208](#)
- [10] J.-H. Cho and J.-Y. Ham, [The minimal dilatation of a genus-two surface](#). *Exp. Math.* **17** (2008), no. 3, 257–267 Zbl [1153.37375](#) MR [2455699](#)
- [11] M. Culler, N. M. Dunfield, M. Goerner, and J. R. Weeks, SnapPy, a computer program for studying the geometry and topology of 3-manifolds. <http://snappy.computop.org> visited on 7 May 2024
- [12] B. Farb, C. J. Leininger, and D. Margalit, [Small dilatation pseudo-Anosov homeomorphisms and 3-manifolds](#). *Adv. Math.* **228** (2011), no. 3, 1466–1502 Zbl [1234.37022](#) MR [2824561](#)
- [13] E. Farber, B. Reinoso, and L. Wang, Fixed-point-free pseudo-Anosov homeomorphisms, knot Floer homology and the cinquefoil. [v1] 2022, [v2] 2023, arXiv:2203.01402v2
- [14] A. Fathi, F. Laudenbach, and V. Poénaru, *Thurston's work on surfaces*. Mathematical Notes (Princeton) 48, Princeton University Press, Princeton, NJ, 2012, xiii, 255 p. Zbl [1244.57005](#) MR [3053012](#)
- [15] D. Fried, [The geometry of cross sections to flows](#). *Topology* **21** (1982), no. 4, 353–371 Zbl [0594.58041](#) MR [670741](#)
- [16] D. Fried, [Growth rate of surface homeomorphisms and flow equivalence](#). *Ergodic Theory Dynam. Systems* **5** (1985), no. 4, 539–563 Zbl [0603.58020](#) MR [829857](#)
- [17] A. Giannopolous, S. Schleimer, and H. Segerman, A census of veering structures. <https://math.okstate.edu/people/segerman/veering.html> visited on 7 May 2024
- [18] B. H. Gross and C. T. McMullen, [Automorphisms of even unimodular lattices and unramified Salem numbers](#). *J. Algebra* **257** (2002), no. 2, 265–290 Zbl [1022.11016](#) MR [1947324](#)
- [19] J.-Y. Ham and W. T. Song, [The minimum dilatation of pseudo-Anosov 5-braids](#). *Exp. Math.* **16** (2007), no. 2, 167–179 Zbl [1151.37037](#) MR [2339273](#)
- [20] E. Hironaka, [Small dilatation mapping classes coming from the simplest hyperbolic braid](#). *Algebr. Geom. Topol.* **10** (2010), no. 4, 2041–2060 Zbl [1221.57028](#) MR [2728483](#)
- [21] E. Hironaka and E. Kin, [A family of pseudo-Anosov braids with small dilatation](#). *Algebr. Geom. Topol.* **6** (2006), 699–738 Zbl [1126.37014](#) MR [2240913](#)
- [22] E. Kin, Dynamics of the monodromies of the fibrations on the magic 3-manifold. *New York J. Math.* **21** (2015), 547–599 Zbl [1341.57008](#) MR [3386537](#)
- [23] E. Kin, S. Kojima, and M. Takasawa, [Minimal dilatations of pseudo-Anosovs generated by the magic 3-manifold and their asymptotic behavior](#). *Algebr. Geom. Topol.* **13** (2013), no. 6, 3537–3602 Zbl [1306.37042](#) MR [3248741](#)
- [24] E. Kin and M. Takasawa, [Pseudo-Anosovs on closed surfaces having small entropy and the Whitehead sister link exterior](#). *J. Math. Soc. Jpn.* **65** (2013), no. 2, 411–446 Zbl [1270.57044](#) MR [3055592](#)
- [25] S. Kojima and G. McShane, [Normalized entropy versus volume for pseudo-Anosovs](#). *Geom. Topol.* **22** (2018), 2403–2426 Zbl [1398.57024](#) MR [3784525](#)
- [26] E. Lanneau and J.-L. Thiffeault, [On the minimum dilatation of braids on punctured discs](#). *Geom. Dedicata* **152** (2011), 165–182 Zbl [1230.37047](#) MR [2795241](#)
- [27] E. Lanneau and J.-L. Thiffeault, [On the minimum dilatation of pseudo-Anosov homeomorphisms on surfaces of small genus](#). *Ann. Inst. Fourier (Grenoble)* **61** (2011), no. 1, 105–144 Zbl [1237.37027](#) MR [2828128](#)
- [28] E. Lanneau and F. Valdez, [Computing the Teichmüller polynomial](#). *J. Eur. Math. Soc. (JEMS)* **19** (2017), no. 12, 3867–3910 Zbl [1380.37070](#) MR [3730517](#)

- [29] C. T. McMullen, [Polynomial invariants for fibered 3-manifolds and Teichmüller geodesics for foliations](#). *Ann. Sci. Ecole Norm. Sup. (4)* **33** (2000), no. 4, 519–560 Zbl [1013.57010](#) MR [1832823](#)
- [30] C. T. McMullen, [Entropy and the clique polynomial](#). *J. Topol.* **8** (2015), no. 1, 184–212 Zbl [1353.37033](#) MR [3335252](#)
- [31] A. Parlak, [Computation of the taut, the veering and the teichmüller polynomials](#). *Exp. Math.* **33** (2024), no. 1, 1–26 Zbl [07815479](#) MR [4713861](#)
- [32] R. C. Penner, [Bounds on least dilatations](#). *Proc. Amer. Math. Soc.* **113** (1991), no. 2, 443–450 Zbl [0726.57013](#) MR [1068128](#)
- [33] R. C. Penner and J. L. Harer, [Combinatorics of train tracks](#). *Annals of Mathematics Studies* 125, Princeton University Press, Princeton, NJ, 1992, xii+216 pp. Zbl [0765.57001](#) MR [1144770](#)
- [34] W. T. Song, K. H. Ko, and J. E. Los, [Entropies of braids](#). *J. Knot Theory Ramifications* **11** (2002), no. 4, 647–666 Zbl [1010.57004](#) MR [1915500](#)
- [35] J. R. Stallings, [Topology of finite graphs](#). *Invent. Math.* **71** (1983), no. 3, 551–565 Zbl [0521.20013](#) MR [695906](#)
- [36] H. Sun, [A transcendental invariant of pseudo-Anosov maps](#). *J. Topol.* **8** (2015), no. 3, 711–743 Zbl [1334.57012](#) MR [3394315](#)
- [37] W. P. Thurston, A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.* **59** (1986), no. 339, i–vi and 99–130 Zbl [0585.57006](#) MR [0823443](#)
- [38] W. P. Thurston, [On the geometry and dynamics of diffeomorphisms of surfaces](#). *Bull. Amer. Math. Soc.* **19** (1988), no. 2, 417–431 Zbl [0674.57008](#) MR [0956596](#)
- [39] C.-Y. Tsai, [The asymptotic behavior of least pseudo-Anosov dilatations](#). *Geom. Topol.* **13** (2009), no. 4, 2253–2278 Zbl [1204.37043](#) MR [2507119](#)
- [40] A. D. Valdivia, Sequences of pseudo-Anosov mapping classes and their asymptotic behavior. *New York J. Math.* **18** (2012), 609–620 Zbl [1350.57021](#) MR [2967106](#)
- [41] M. Yazdi, [Pseudo-Anosov maps with small stretch factors on punctured surfaces](#). *Algebr. Geom. Topol.* **20** (2020), no. 4, 2095–2128 Zbl [1450.37032](#) MR [4127091](#)

Received 28 December 2022.

Eriko Hironaka

Department of Mathematics, Florida State University, 208 Love Building, 1017 Academic Way, Tallahassee, FL 32301-4510, USA; hironaka@math.fsu.edu

Chi Cheuk Tsang

Department of Mathematics, University of California, Berkeley, 970 Evans Hall #3840, Berkeley, CA 94720-3840, USA; chicheuk@math.berkeley.edu