

POLYNOMIAL INVARIANTS OF PSEUDO-ANOSOV MAPS

JOAN BIRMAN, PETER BRINKMANN, AND KEIKO KAWAMURO

ABSTRACT. We investigate the structure of the characteristic polynomial $\det(xI - T)$ of a transition matrix T that is associated to a train track representative of a pseudo-Anosov map $[F]$ acting on a surface. As a result we obtain three new polynomial invariants of $[F]$, one of them being the product of the other two, and all three being divisors of $\det(xI - T)$. The degrees of the new polynomials are invariants of $[F]$ and we give simple formulas for computing them by a counting argument from an invariant train-track. We give examples of genus 2 pseudo-Anosov maps having the same dilatation, and use our invariants to distinguish them.

1. INTRODUCTION

Let S be an orientable 2-manifold of genus g , closed or with finitely many punctures, where the genus and the number of punctures are chosen so that S admits a hyperbolic structure. The *modular group* $\text{Mod}(S)$ is the group $\pi_0(\text{Diff}(S))$, where admissible homeomorphisms preserve orientation. If a mapping class $[F] \in \text{Mod}(S)$ is *pseudo-Anosov* or *pA*, then there exists a representative $F : S \rightarrow S$, a pair of invariant transverse measured foliations (\mathcal{F}^u, μ_u) , (\mathcal{F}^s, μ_s) , and a real number λ , the *dilatation* of $[F]$, such that F multiplies the transverse measure μ_u (resp. μ_s) by λ (resp. $\frac{1}{\lambda}$). The real number $\lambda(F)$ is an invariant of the conjugacy class of $[F]$ in $\text{Mod}(S)$.

In this paper we introduce a new approach to the study of invariants of $[F]$, when $[F]$ is pA. Our work was in part motivated by recent efforts (see [7] and a host of papers that were inspired by it) to understand precisely which real numbers λ occur in this setting. It is known that if one fixes S , then as F is varied its dilatation $\lambda(F)$ takes on a minimum value $\lambda_{S,min}$, where by this we mean that any pA map on S has dilatation $\geq \lambda_{S,min}$, also $\lambda_{S,min}$ is realized by some pA map F . The number $\lambda_{S,min}$ is of great interest. Many attempts have been made to use the approach in [7] to find $\lambda_{S,min}$ for closed surfaces of arbitrary genus g , however it appeared to us (after attending a workshop in April 2010 addressed to the study of $\lambda_{S,min}$), that a new approach was needed. The main result in this paper is, for an arbitrary pA map F , the determination of two integer polynomials, both containing $\lambda(F)$ as their largest real root, and the proof that both are invariants of the given pA mapping class. Both will be seen to have a known topological meaning.

Date: January 30, 2012.

2000 Mathematics Subject Classification. Primary 57M25, 57M27; Secondary 57M50.

Key words and phrases. Pseudo-Anosov map; Train track; transition matrix; dilatation.

The third author was partially supported by NSF grants DMS-0806492 and DMS-0635607.

The study of these two polynomials, rather than of λ itself, is the new approach that we have in mind.

Measured train tracks are a partially combinatorial device that Thurston introduced to encode the essential properties of $(\mathcal{F}^u, \mu_u), (\mathcal{F}^s, \mu_s)$. A train track τ is a branched 1-manifold that is embedded in the surface S . It is made up of vertices (called *switches*) and smooth edges (called *branches*), disjointly embedded in S . See §1.3 of [8]. Given a pA map $[F]$, there exists a train track $\tau \subset S$ that *fills* the surface, i.e., the complement of τ consists of (possibly punctured) discs, and τ is left invariant by $[F]$. Moreover, τ is equipped with a transverse measure (resp. tangential measure) that is related to the transverse measure μ_u on \mathcal{F}_u (resp. μ_s on \mathcal{F}_s).

In [1] Bestvina and Handel gave an algorithmic proof of Thurston's classification theorem for mapping classes. Their proof shows that, if $[F]$ is a pA map of S , then one may construct, algorithmically, a graph G , homotopic to S when S is punctured, and an induced map $f : G \rightarrow G$, that we call a *train track map*. For every $r \geq 1$ the restriction of f^r to the interior of every edge is an immersion. It takes a vertex of G to a vertex, and takes an edge to an edge-path which has no backtracking. Let e_1, \dots, e_n be the unoriented edges of G . Knowing G and $f : G \rightarrow G$, they construct a somewhat special measured train track τ , and we will always assume that our τ comes from their construction. The *transition matrix* T is an $n \times n$ matrix whose entry $T_{i,j}$ is the number of times the edge path $f(e_j)$ passes over e_i in either direction, so that all entries of T are non-negative integers. If $[F]$ is pA, then T is irreducible and it has a dominant real eigenvalue λ , the *Perron-Frobenius eigenvalue* [6]. The eigenvalue λ is the dilatation of $[F]$. The left (resp. right) eigenvectors of T determine tangential (resp. transversal) measures on τ , and eventually determine μ_s (resp. μ_u).

In this paper we study the structure of the characteristic polynomial $\det(xI - T)$ of the transition matrix T . Our work depends crucially on the Bestvina-Handel algorithm, however we look for the structure needed to find a measured train track in T , and not in the matrix $T' = \begin{bmatrix} N & A \\ 0 & T \end{bmatrix}$ that appears in §3.3 and §3.4 of [1]. Bestvina and Handel use T' to build the invariant foliations associated to f . As is well-known, all of the information needed for that construction is already present in T , and we shall build on that fact. With that in mind, let $V(G)$ be the vector space of real weights on the edges of the Bestvina-Handel graph G . Let $f_* : V(G) \rightarrow V(G)$ be the map induced from the train track map $f : G \rightarrow G$. Let $\chi(f_*) = \det(xI - T)$. It is well-known that $\chi(f_*)$ depends on the choice of $f : G \rightarrow G$ within its conjugacy class $[F]$.

The first new result in this paper is the discovery that, after dividing $\chi(f_*)$ by a polynomial that is determined by the way that a train-track map acts on certain vertices of G , one obtains a quotient polynomial which is a topological invariant of $[F]$. This polynomial arises via an f_* -invariant direct sum decomposition of the \mathbb{R} -vector space of transverse measures on τ . It is the characteristic polynomial of the action of f_* on one of the summands. We call it the *homology polynomial* of $[F]$ for reasons that will become clear in a moment. We will construct examples of pA maps on a surface of genus 2 which have the same dilatation, but are distinguished by their homology polynomials.

Like $\chi(f_*)$, our homology polynomial is the characteristic polynomial of an integer matrix, although (unlike T) that matrix is not in general non-negative. We now describe how we found it. We define and study an f_* -invariant subspace $W(G, f) \subset V(G)$. The subspace $W(G, f)$ is chosen so that weights on edges determine a transverse measure on the train track τ associated to G and $f : G \rightarrow G$. We study $f_*|_{W(G, f)}$. See page 427, lines 20- to 13- of [11], where the mathematics that underlies $f_*|_{W(G, f)}$ is described by Thurston. Our first contribution in this paper is to make the structure that Thurston described there concrete and computable, via an enhanced form of the Bestvina-Handel algorithm. This allows us to prove that the characteristic polynomial $\chi(f_*|_{W(G, f)})$ is an invariant of the mapping class $[F]$ in $\text{Mod}(S)$. This polynomial $\chi(f_*|_{W(G, f)})$ is our homology polynomial and we denote it by $h(x)$. It contains the dilatation of $[F]$ as its largest real root, and so is divisible by the minimum polynomial of λ . Its degree depends upon a careful analysis of the action of f_* on the vertices of G .

Investigating the action of f_* on $W(G, f)$, we show that $W(G, f)$ supports a skew-symmetric form that is f_* -invariant. The existence of the symplectic structure was known to Thurston and also was studied by Penner-Harer in [8], however it is unclear to us whether it was known to earlier workers that it could have degeneracies. See Remark 3.7. We discovered via examples that degeneracies do occur. In §3 we investigate the radical Z of the skew-symmetric form, i.e. the totally degenerate subspace of the skew-symmetric form, and arrive at an f_* -invariant decomposition of $W(G, f)$ as $Z \oplus (W(G, f)/Z)$. This decomposition leads to a product decomposition of the homology polynomial as a product of two additional new polynomials, with both factors being invariants of $[F]$. We call the first of these new polynomials, $p(x) = \chi(f_*|_Z)$, the *puncture* polynomial because it is a cyclotomic polynomial that relates to the way in which the pA map F permutes certain punctures in S . As for $s(x) = \chi(f_*|_{W(G, f)/Z})$, our *symplectic* polynomial, we know that it arises from the action of f_* on the symplectic space $W(G, f)/Z$, but we do not fully understand it at this writing. Sometimes the symplectic polynomial is irreducible, in which case it is the minimum polynomial of λ . However we will give examples to show that it can be reducible, and even an example where it is symplectically reducible. Thus its relationship to the minimum polynomial of λ is not completely clear at this writing.

Summarizing, we will prove:

Theorem 1.1. *Let $[F]$ be a pA mapping class in $\text{Mod}(S)$, with Bestvina-Handel train track map $f : G \rightarrow G$ and transition matrix T .*

- (1) *The characteristic polynomial $\chi(f_*)$ of T has a divisor, the homology polynomial $h(x)$ which is an invariant of $[F]$. It contains λ as its largest real root, and is associated to an induced action of F_* on $H_1(X, \mathbb{R})$, where X is the surface S when τ is orientable and its orientation cover \tilde{S} when τ is non-orientable.*
- (2) *The homology polynomial $h(x)$ decomposes as a product $p(x) \cdot s(x)$ of two polynomials, each a topological invariant of $[F]$.*
 - (a) *The first factor, the puncture polynomial $p(x)$, records the action of f_* on the radical of a skew-symmetric form on $W(G, f)$. It has topological meaning related to the way in which F permutes certain punctures in the surface S .*

- It is a palindromic or anti-palindromic polynomial, and all of its roots are on the unit circle.*
- (b) *The second factor, the symplectic polynomial $s(x)$, records the action of f_* on the non-degenerate symplectic space $W(G, f)/Z$. It contains λ as its largest real root. It is palindromic. If irreducible, it is the minimum polynomial of λ , but it is not always irreducible.*
 - (3) *The homology polynomial $h(x)$, being a product of the puncture and symplectic polynomials, is palindromic or anti-palindromic.*

The proof of Theorem 1.1 can be found in §2 and §3 below. In §4 we give several applications, and prove that our three invariants behave nicely when the defining map F is replaced by a power F^k . The paper ends, in §5 with a set of examples which give concrete meaning to our ideas. The first such example, Example 5.1, defines three distinct maps F_1, F_2, F_3 on a surface of genus 2, chosen so that all three have the same dilatation. Two of the three pairs are distinguished by any one of our three invariants. The third map was chosen so that it probably is not conjugate to the other two, however our invariants could not prove that.

2. PROOF OF PART (1) OF THEOREM 1.1

We begin our work in §2.1 by recalling some well-known facts from [1] that relate to the construction of the train track τ by adding infinitesimal edges to the graph G . After that, in §2.2, we introduce the space $W(G, f)$ of transverse measures, which plays a fundamental role throughout this paper. Rather easily, we will be able to prove our first decomposition and factorization theorem. Thus at the end of §2.2 we have our homology polynomial in hand, but we do not know its meaning, have not proved it is an invariant, and don't know how to compute it. In §2.3 we prepare for the work ahead by constructing a basis for $W(G, f)$. We also learn how to find the matrix for the action of f_* on the basis. With that in hand, in §2.4 we identify the vector space $W(G, f)$ with a homology space. We will be able to prove Corollary 2.25, which asserts that the homology polynomial $h(x)$ is a topological invariant of the conjugacy class of our pA map $[F]$ in $\text{Mod}(S)$. (Later in §5, we will use it to distinguish examples of pA maps acting on the same surface and having the same dilatation.)

2.1. Preliminaries. It will be assumed that the reader is familiar with the basic ideas of the algorithm of Bestvina-Handel [1]. The mapping class $[F]$ will always be pA. Further, assume that we are given the graph $G \subset S$, homotopic to S , and a train track map $f : G \rightarrow G$. We note that if S is closed, the action of $[F]$ always has periodic points with finite order, and the removal of a periodic orbit will not affect our results, therefore without loss of generality we may assume that S is finitely punctured.

Following ideas in [1] we construct a train track τ from $f : G \rightarrow G$ by equipping the vertices of G with additional structure: Let $e_1, e_2 \subset G$ be two (non-oriented) edges originating at the same vertex v . Edges e_1 and e_2 belong to the same *gate* at v if for some $r > 0$, the edge-paths $f^r(e_1)$ and $f^r(e_2)$ have a nontrivial common initial segment.

If e_1 and e_2 belong to different gates at v and there exists some exponent $r > 0$ and an edge e so that $f^r(e)$ contains e_2e_1 or e_1e_2 as a subpath, then we connect the gates associated to e_1, e_2 with an *infinitesimal edge*. In this way, a vertex $v \in G$ with k gates becomes an infinitesimal k -gon in the train track τ . While this k -gon may be missing one-side, the infinitesimal edges must connect all the gates at each vertex, see §3.3 of [1]. In addition to the infinitesimal edges, τ also has *real edges* corresponding to the edges of G . Hence, a *branch* of τ , in the sense of Penner-Harer [8], is either an infinitesimal edge or a real edge.

It is natural to single out the following properties of the vertices of G :

Definition 2.1 (Vertex types). See Figure 1. A vertex of G is *odd* (resp. *even*) if its corresponding infinitesimal complete polygon in τ has an odd (resp. *even*) number of sides, and it is *partial* if its infinitesimal edges form a polygon in τ with one side missing. Partial vertices include the special case where v has only two gates connected by one infinitesimal edge; we call such vertices *evanescent*. The symbol w_i (resp. x_i) will denote the weight of i -th gate (resp. infinitesimal edge).




FIGURE 1. Shaded disks enclose infinitesimal (partial) polygons in τ that correspond to the vertices of G .

Remark 2.2. In §5 the reader can find several examples illustrating the graph G with infinitesimal polygons associated to particular pA mapping classes. In those illustrations the vertices of the graphs have been expanded to shaded discs which show the structure at the vertices. In the sketch of a train track that XTrain generates, the branches at each gate do not appear to be tangent to each other. This was done for ease in drawing the required figures. The reader should keep in mind that all the branches at each gate are tangent to each other.

We recall properties of non-evanescent vertices that are preserved under a train track map:

Lemma 2.3. (Proposition 3.3.3 of [1]) *For $k \geq 3$, let O_k be the set of odd vertices with k gates, E_k the set of even vertices with k gates, and P_k the set of partial vertices with k gates. Then the restriction of $f : G \rightarrow G$ to each of these sets is a permutation of the set.*

Moreover, for each (non evanescent) vertex v with at least three gates, f induces a bijection between the gates at v and the gates at $f(v)$ that preserves the cyclic order.

Remark 2.4. The number of evanescent vertices of a train track representative is *not* an invariant of the underlying mapping class. Examples exist where a train track map has a representative with evanescent vertices, and another without.

Definition 2.5 (Orientable and non-orientable train tracks). Choose an orientation on each branch of a train track τ . A train track is *orientable* if we can orient all the branches so that, at every switch, the angle between each incoming branch and each outgoing branch is π . For example, see the train tracks τ_1 and τ_2 that are given in Figure 7 of §5. After adding the infinitesimal edges, one sees that τ_1 is orientable, but τ_2 is not.

Here is an easy observation.

Lemma 2.6. *If G has an odd vertex, then the corresponding train track τ is non-orientable. This condition is sufficient, but not necessary.*

Proof. If v is an odd vertex, then there exists no consistent orientation for the corresponding infinitesimal polygon in τ . The example in Figure 7 shows that the condition is not necessary. \square

Remark 2.7. We do not know any immediate visual criterion beyond the one in Lemma 2.6 for detecting non-orientability. The two train tracks τ_1, τ_2 in Figure 7 of this paper both have 2 vertices, one even and one partial, and τ_1 is orientable whereas τ_2 is non-orientable. If all the vertices are partial, then the train track may be either orientable (see Example 4.2 of [2]) or non-orientable (see Example 5.3); also, a non-orientable train track may have odd, even, and partial vertices at the same time (see Example 5.4).

2.2. The space $W(G, f)$ and the first decomposition. Given a graph G of n edges, one always has an \mathbb{R} -vector space $V(G) \simeq \mathbb{R}^n$ of weights on G . Our goal in this section is to define a subspace $W(G, f) \subset V(G)$ of ‘transverse measures on G ’. This space is the natural projection of the measured train track τ to a space of measures on G . It will play a fundamental role in our work.

In our setting, all train tracks are *bi-recurrent*, that is, recurrent and transversely recurrent, cf. p.20 of [8]. To define our space $W(G, f)$, we apply Penner-Harer’s work described in §3.2 of [8], where bi-recurrence is assumed.

Let $V(\tau) \cong \mathbb{R}^{n+n'}$, where n (resp. n') is the number of the real (resp. infinitesimal) edges of train track τ . Penner-Harer defined a subspace $W(\tau) \subset V(\tau)$ of assignments of (possibly negative) real numbers, one to each branch of τ , which satisfy the *switch conditions*. That is, if $\eta \in W(\tau)$ then at each switch of τ , the sum of the weights on the incoming branches equals to the sum on the outgoing branches. For example, in Figure 2-(A), $\eta(a) = \eta(b_1) + \eta(b_2)$.

Definition 2.8. There is a natural surjection $\pi : \tau \rightarrow G$ which is defined by collapsing all the infinitesimal (partial) polygons to their associated vertices in G and taking each real edge in τ to the corresponding edge in G . Let $W(G, f) = \pi_*(W(\tau))$. That is, $W(G, f) \subset V(G)$ is the subspace whose elements admit an extension to a (possibly



FIGURE 2. (A) A switch of valence 3. (B) A switch of valence 5.

negative) transverse measure on τ . The name $W(G, f)$ has been chosen to reflect the fact that our subspace depends not only on G , but also on the action $f : G \rightarrow G$.

Here is a useful criterion for an element of $V(G)$ to be in $W(G, f)$:

Lemma 2.9. *An element $\eta \in V(G)$ belongs to $W(G, f)$ if and only if for each non-odd vertex the alternating sum of the weights at the incident gates is zero.*

Proof. Assume that $\eta \in V(G)$ belongs to $W(G, f)$. Let $v \in G$ be a vertex with k gates, i.e., v corresponds to an infinitesimal k -gon, possibly partial, in the train track τ . Let $w_0, \dots, w_{k-1} \in \mathbb{R}$ be the weights of η at the incident gates of v , and let $x_0, \dots, x_{k-1} \in \mathbb{R}$ (or x_0, \dots, x_{k-2} if v is a partial vertex) be the weights of the infinitesimal edges. See Figure 1. The weights on the infinitesimal edges may turn out to be negative real numbers. We determine when an assignment of weights to the real edges admits an extension to the infinitesimal edges that satisfies the switch conditions.

If v is odd or even with k gates, then the switch condition imposes:

$$\begin{aligned} x_{k-1} + x_0 &= w_0, \\ x_0 + x_1 &= w_1, \\ x_1 + x_2 &= w_2, \\ &\vdots \\ x_{k-2} + x_{k-1} &= w_{k-1}. \end{aligned}$$

If k is odd, this system of equations has a unique solution, regardless of the weights w_i . If k is even, the system is consistent if and only if $\sum_{i=0}^{k-1} (-1)^i w_i = 0$.

If v is partial with k gates, then the switch condition imposes:

$$\begin{aligned} x_0 &= w_0, \\ x_0 + x_1 &= w_1, \\ x_1 + x_2 &= w_2, \\ &\vdots \\ x_{k-3} + x_{k-2} &= w_{k-2}, \\ x_{k-2} &= w_{k-1}. \end{aligned}$$

This system has a unique solution if and only if $\sum_{i=0}^{k-1} (-1)^i w_i = 0$. □

Lemma 2.10. *$W(G, f)$ is an invariant subspace of $V(G)$ under f_* , i.e., $f_*(W(G, f)) \subseteq W(G, f)$.*

Proof. Suppose v is a non-odd vertex and mapped to a non-odd vertex $f(v)$. Let $\eta \in W(G, f)$. By Lemma 2.9, the alternating weight sum of η at the incident gates of v is 0. Lemma 2.3 implies that all the weights of η at the infinitesimal edges for v is inherited to the weights of $f_*\eta$ at the infinitesimal edges for $f(v)$.

In addition, we account for an edge $e \subset G$ whose image $f(e) = e_0e_1 \cdots e_k$ passes through the vertex $f(v)$. Assume that η has weight $w = \eta(e)$ at the edge e . If a sub-edge-path e_ie_{i+1} passes through $f(v)$ then edges e_i and e_{i+1} belong to adjacent gates at $f(v)$ and the contribution of e_ie_{i+1} to the alternating sum of weights of the gates at $f(v)$ is $w - w = 0$.

Therefore, the alternating weight sum for $f_*\eta$ at the incident gates for $f(v)$ is 0. By Lemma 2.9, $f_*\eta \in W(G, f)$. \square

The dimension of the vector space $W(G, f)$ can be computed combinatorially by inspecting a train track associated to the pair (G, f) :

Lemma 2.11. (1) *If τ is orientable, then*

$$\begin{aligned} \dim W(G, f) &= \#(\text{edges of } G) - \#(\text{vertices of } G) + 1, \\ W(G, f) &\cong Z_1(G; \mathbb{R}) \cong H_1(G; \mathbb{R}) \cong H_1(S; \mathbb{R}). \end{aligned}$$

In particular, the switch conditions are precisely the cycle conditions.

(2) *If τ is non-orientable, then*

$$\dim W(G, f) = \#(\text{edges of } G) - \#(\text{non-odd vertices of } G).$$

Proof. Assume that τ is orientable. By Lemma 2.6, G has no odd vertices. For $\eta \in W(G, f)$ and a non-odd vertex v , let

$$w_i^v := w_i^v(\eta) = \text{the weight of } \eta \text{ at the } i^{\text{th}} \text{ gate of the vertex } v.$$

By Lemma 2.9, $\sum_i (-1)^i w_i^v = 0$. We number the gates at v so that the orientation of the real edges at $2i$ -th (resp. $(2i+1)$ -th) gate is inward (resp. outward). If G has m vertices, v_1, \dots, v_m , then we have a system of m equations:

$$\begin{aligned} w_0^{v_1} + w_2^{v_1} + w_4^{v_1} + \cdots &= w_1^{v_1} + w_3^{v_1} + w_5^{v_1} + \cdots, \\ w_0^{v_2} + w_2^{v_2} + w_4^{v_2} + \cdots &= w_1^{v_2} + w_3^{v_2} + w_5^{v_2} + \cdots, \\ &\vdots \\ w_0^{v_m} + w_2^{v_m} + w_4^{v_m} + \cdots &= w_1^{v_m} + w_3^{v_m} + w_5^{v_m} + \cdots. \end{aligned}$$

The sum of the left hand sides is equal to the sum of the right hand sides. Since τ is oriented, the last equation follows from the other $m-1$ equations, i.e., the switch conditions are *not* independent. Therefore,

$$\begin{aligned} \dim W(G, f) &= \dim V(G) - (m-1) \\ &= \#(\text{edges of } G) - \#(\text{vertices of } G) + 1. \end{aligned}$$

With respect to the orientations of the edges of G , let $\partial : C_1(G; \mathbb{R}) \rightarrow C_0(G; \mathbb{R})$ be the boundary map of the chain complex. There is a natural isomorphism $V(G) \cong C_1(G; \mathbb{R})$.

If $\gamma \in Z_1(G; \mathbb{R})$ is a cycle, then the cycle condition $\partial\gamma = 0$ is equivalent to the alternating sum condition $\sum_i (-1)^i w_i^v(\gamma) = 0$ at each vertex $v \in G$. By Lemma 2.9 we obtain $W(G, f) \cong Z_1(G; \mathbb{R})$. In fact, the Euler characteristic of s -punctured genus g surface S is; $\chi(S) = 2 - 2g - s = \#(\text{vertices of } G) - \#(\text{edges of } G)$. Thus $\dim W(G, f) = 2g + s - 1 = \dim H_1(S; \mathbb{R})$.

In the non-orientable case, the switch conditions are satisfied if and only if the alternating sum of the weights of gates around each even or partial vertex is zero (Lemma 2.9). Moreover, all these conditions are independent of each other (see Lemma 2.1.1 of [8]), so that the number of independent constraints is the number of non-odd vertices. Since $\dim V(G)$ is the number of edges, statement (2) follows. \square

Now we know that $V(G)$ can be identified with the 1-chains $C_1(G)$ in the orientable case, we can extend the boundary map $\partial : C_1(G) \rightarrow C_0(G)$ to the following map δ on $V(G)$:

Definition 2.12. (The map δ) Assume that the graph G has m non-odd vertices, v_1, \dots, v_m . Define a linear map $\delta : V(G) \rightarrow \mathbb{R}^m$, $\eta \mapsto \delta(\eta)$ such that

$$(k^{\text{th}} \text{ entry of the vector } \delta(\eta)) = \sum_i (-1)^i w_i^{v_k}(\eta),$$

the alternating sum of the weights of η at the gates incident to the vertex v_k . These weights satisfy the following conditions:

- If τ is oriented, then we determine the sign of each gate to be compatible with the orientation of the real edges of τ . The alternating sum is defined without ambiguity. For an example, in Figure 7, left sketch, a plus (resp. minus) sign may be assigned at each gate, according as the real edges are oriented toward (resp. away from) the gate.
- If τ is non-orientable, we assign alternating signs to the incident gates for each non-odd vertex. Since τ is non-orientable, the assignments will be local and not global. Clearly, the alternating sum depends on the choice of the sign assignment.

Lemma 2.13. *In both the orientable and non-orientable cases $W(G, f) \cong \ker \delta$. Moreover, if m is the number of non-odd vertices of G , then*

$$\dim(\text{im } \delta) = \begin{cases} m - 1, & \text{if } \tau \text{ is orientable,} \\ m, & \text{if } \tau \text{ is non-orientable.} \end{cases}$$

Proof. Lemma 2.9, immediately implies that $W(G, f) \cong \ker \delta$. The dimension count follows from Lemma 2.11. \square

Finally we prove:

Theorem 2.14 (First Decomposition). *Let $h(x) = \chi(f_*|_{W(G, f)})$, then*

$$(2.1) \quad V(G) \cong W(G, f) \oplus \text{im } \delta.$$

$$(2.2) \quad \chi(f_*) = h(x) \chi(f_*|_{\text{im } \delta}).$$

The degree of $h(x)$ (resp. $\chi(f_*|_{\text{im } \delta})$) is the dimension of $W(G, f)$ (resp. $\text{im } \delta$), as given in Lemma 2.11 (resp. Lemma 2.13).

Proof. From Lemma 2.13, identifying $\text{im } \delta$ with the quotient $V(G)/W(G, f)$ we obtain (2.1). By the same argument as in the proof of Lemma 2.10, we obtain $f_*(\text{im } \delta) \subset \text{im } \delta$. This along with $f_*(W(G, f)) \subset W(G, f)$ yields (2.2). \square

2.3. A basis for $W(G, f)$ and the matrix for the action of f_* on $W(G, f)$. The goal of this section is to find a basis for $W(G, f) \subset V(G)$ and learn how to determine the action of f_* on the basis. With regard to the basis, it will be convenient to consider three cases separately: the cases when τ is orientable, when τ is non-orientable and has odd vertices, and when τ is non-orientable but has no odd vertices. That is accomplished in §2.3.1, 2.3.2 and 2.3.3. Having the basis in hand, in §2.3.4 we learn how to compute the action of f_* on the basis elements.

2.3.1. Basis for $W(G, f)$, orientable case. If the train track τ associated to G is orientable, we choose an orientation for τ , thereby inducing an orientation on the edges of G , and choose a maximal spanning tree $Y \subset G$. Every vertex of G will be in Y . We consider all edges e of G that are not in Y , and construct a set of vectors $\{\eta_e \in V(G)\}$ and prove that the constructed set is a basis for $W(G, f) \subset V(G)$.

For each $e \in G \setminus Y$, find the unique shortest path in Y joining the endpoints of e . The union of this path and e forms an oriented loop L_e in which the edge e appears exactly once, the orientation being determined by that on e . If the orientation of edge $e' \subset G$ agrees (resp. disagrees) with the orientation of $e' \subset L_e$, then we assign a weight of 1 (resp. -1) to e' . In particular, $e \subset L_e$ has weight 1. The edges not in L_e are assigned weight of 0. In this way we obtain a vector $\eta_e \in V(G)$ whose entries are the assigned weights. By construction, η_e satisfies the criterion for a transverse measure, described in Lemma 2.9, therefore η_e is an element of $W(G, f)$.

We now show that $\{\eta_{e_1}, \dots, \eta_{e_l}\}$ is a basis for $W(G, f)$ where e_1, \dots, e_l are the edges of $G \setminus Y$ and $l = \#(\text{edges of } G) - \#(\text{vertices of } G) + 1$. Note that if $e_i, e_j \in (G \setminus Y)$ with $i \neq j$, then $e_j \notin L_{e_i}$. Therefore $\eta_{e_i}(e_j) = 1$ if and only if $i = j$, because all edges not in L_{e_i} have weight 0, i.e., the vectors $\eta_{e_1}, \dots, \eta_{e_l}$ are linearly independent. Consulting Lemma 2.11 we see that we have the right number of linearly independent elements, so we have found a basis for $W(G, f)$.

Example 2.15. Go to §5 below and see Example 5.1 and its accompanying Figure 7-(1). The train track for this example is orientable. The space $V(G)$ has dimension 5. Order the edges of G as a, b, c, d, e . The edge a and the vertices v_0, v_1 form a maximal tree $Y \subset G$, with edges $b, c, d, e \notin Y$, so that $W(G, f)$ has dimension 4. We have the loops $L_b = ab$; $L_c = \bar{a}c$; $L_d = \bar{a}d$; $L_e = ae$, so that $W(G, f)$ has basis $\eta_b = (1, 1, 0, 0, 0)'$; $\eta_c = (-1, 0, 1, 0, 0)'$; $\eta_d = (-1, 0, 0, 1, 0)'$; $\eta_e = (1, 0, 0, 0, 1)'$, where ‘prime’ means transpose.

2.3.2. Basis for $W(G, f)$, non-orientable case with odd vertices. If G has an odd vertex v_0 then choose a maximal spanning tree $Y \subset G$. Let V be the set of vertices of G . Define

a height function $h : V \rightarrow \mathbb{N} \cup \{0\}$ by

$$h(v) = (\text{the distance between } v \text{ and } v_0 \text{ in } Y).$$

We obtain a forest $Y' \subset Y$ by removing from Y all the edges each of which connects an odd vertex and the adjacent vertex of smaller height. See Figure 3. The forest Y'




FIGURE 3. A tree Y (left) and a forest Y' (right). Hollow dots v_0, \dots, v_4 , are odd vertices. Black dots are non-odd vertices.

contains all the vertices of G with exactly one odd vertex in each connected component.

Now, let e be an edge that is not in Y' . We can find two (possibly empty) shortest paths in Y' each of which connects an endpoint of e to an odd vertex. The union of e and the two paths forms an arc, which we denote by L_e . (If both of the endpoints of e belong to the same tree component of Y' , then L_e becomes a loop containing one odd vertex.) Assign a weight of 1 to e and weight of 0 to the edges that are not in L_e . To the other edges in L_e , we assign weights of ± 1 so that at each non-odd vertex the criterion of transverse measure (Lemma 2.9) is satisfied. This defines an element η_e of $W(G, f)$.

Let e_1, \dots, e_l be the edges of $G \setminus Y'$. By the construction, we have $\eta_{e_i}(e_j) = 1$ if and only if $i = j$, so the vectors $\eta_{e_1}, \dots, \eta_{e_l}$ are linearly independent. Since $l = \#(\text{edges of } G) - \#(\text{non-odd vertices of } G)$, Lemma 2.11 tells us that $l = \dim W(G, f)$, hence $\{\eta_{e_1}, \dots, \eta_{e_l}\}$ is a basis of $W(G, f)$.

Example 2.16. See Example 5.4. The graph G has two odd vertices v_0 and v_4 . We choose a maximal tree Y whose edges are a, c, d, j . This gives us a forest Y' with two components. One consists of the single vertex v_4 , and the other consists of vertices v_0, v_1, v_2, v_3 , and edges a, c, j . The edge h is not in Y' , and its endpoints are v_1 and v_4 . The associated arc $L_h = ch$. Since c and h share the same gate at v_1 , η_h satisfies $\eta_h(c) = -1$, $\eta_h(h) = 1$, and 0 for rest of the edges. The graph G has 10 edges a, b, \dots, j , with a, c, j in Y' . The vector space $W(G, f)$ has dimension 7. The edge h is not in Y' , and its endpoints are v_1 and v_4 . Then $L_h = ch$. Since c and h share the same gate at v_1 , η_h satisfies $\eta_h(c) = -1$, $\eta_h(h) = 1$, and 0 for rest of the edges. The edge b is also not in Y' , and its endpoints are v_2 and v_3 , so $L_b = cjbac$ is an arc whose endpoints coincide at v_0 and η_b satisfies $\eta_b(a) = \eta_b(b) = \eta_b(j) = 1$, and 0 for rest of the edges. The other five basis elements are constructed in a similar way.

2.3.3. Basis for $W(G, f)$, non-orientable case with no odd vertices. In this case, we can find a simple loop $\mathcal{L}_0 \subset G$ that does not admit an orientation consistent with the train track τ . If \mathcal{L}_0 misses any vertices of G , then we define \mathcal{L}_1 by adding an edge with exactly one vertex in \mathcal{L}_0 . If \mathcal{L}_1 misses any vertices, we define \mathcal{L}_2 by adding an edge with exactly one vertex in \mathcal{L}_1 , etc. Ultimately we obtain a connected subgraph \mathcal{L} that is homotopy equivalent to a circle and contains all vertices of G .

If e is an edge outside \mathcal{L} , we can find paths in \mathcal{L} from each endpoint of e to the loop \mathcal{L}_0 , resulting in a path L_\bullet that contains e and with endpoints in \mathcal{L}_0 . Now we can find paths L_1, L_2 in \mathcal{L}_0 so that $L_1 \cup L_2 = \mathcal{L}_0$ and the endpoints of L_1, L_2 , called v_a and v_b , agree with those of L_\bullet . It is possible that v_a and v_b are the same vertex, in which case we set $L_2 = \emptyset$. (This happens in Example 2.17, where our construction is applied to Example 5.3.)

Let $\eta^0 \in V(G)$ be a vector which assigns 1 to edge e , ± 1 to the other edges in L_\bullet , and 0 to the edges not in L_\bullet , so that the alternating weight sum is 0 at all the vertices but v_a, v_b .

Next, let $\eta^1 \in V(G)$ (resp. $\eta^2 \in V(G)$) be a vector which assigns ± 1 to the edges of L_1 (resp. L_2) and 0 to the other edges not in L_1 (resp. L_2) so that $\eta^0 + \eta^1$ (resp. $\eta^0 + \eta^2$) has the alternating sum of weights equal to 0 at all the vertices of G but v_a . In particular, at vertex v_b , for $i = 1, 2$, (the alternating sum of weights of $\eta^0 + \eta^i$) = 0, hence

$$(\text{the alternating sum of weights of } \eta^1 - \eta^2 \text{ at } v_b) = 0.$$

While, at vertex v_a

$$(\text{the alternating sum of weights of } \eta^1 - \eta^2 \text{ at } v_a) = \pm 2.$$

For, if it were 0 then the loop \mathcal{L}_0 can admit an orientation consistent with τ , which is a contradiction. Therefore, at v_a ,

$$(\text{the signed weight of } \eta^1) = -(\text{the signed weight of } \eta^2).$$

This means, at v_a , we have (the alternating sum of weights of $\eta^0 + \eta^1$) = 0 if and only if (the alternating sum of weights of $\eta^0 + \eta^2$) $\neq 0$.

If (the alternating sum of weights of $\eta^0 + \eta^1$) = 0 then define $\eta_e := \eta^0 + \eta^1$ and $L_e := L_\bullet \cup L_1$. Otherwise define $\eta_e := \eta^0 + \eta^2$ and $L_e := L_\bullet \cup L_2$. By construction, η_e satisfies the alternating sum condition of Lemma 2.9, hence $\eta_e \in W(G, f)$.

Note that the number of edges in \mathcal{L} is equal to the number of vertices of G , so by Lemma 2.11 the number l of edges outside \mathcal{L} is equal to $\dim W(G, f)$. Suppose e_1, \dots, e_l are the edges of $G \setminus \mathcal{L}$, then $\eta_{e_i}(e_j) = 1$ if and only if $i = j$, i.e., vectors $\eta_{e_1}, \dots, \eta_{e_l}$ are linearly independent. This proves that $\{\eta_{e_1}, \dots, \eta_{e_l}\}$ is a basis of $W(G, f)$.

Example 2.17. See Example 5.3. The partial vertex v_0 has ten gates. We assign signs alternatively to the gates, which imposes orientations on the edges b, c, e . However, the gates of each edge a, d or f have the same sign, hence a, d, f do not admit consistent orientations. We may choose the loop $\mathcal{L} = \mathcal{L}_0$ to be the union of v_0 and the edge f . For

the edge b , the loop L_b consists of a single edge b and v_0 . The element η_b has

$$\eta_b(b) = 1, \text{ and } \eta_b(a) = \eta_b(c) = \eta_b(d) = \eta_b(e) = \eta_b(f) = 0.$$

For the edge a , the loop $L_a = a \cup f$ and η_a has

$$\eta_a(a) = \eta_a(f) = 1, \text{ and } \eta_a(b) = \eta_a(c) = \eta_a(d) = \eta_a(e) = 0.$$

The alternating sum of the weights of the ten gates is zero for both η_a and η_b . Lemma 2.9 guarantees that $\eta_a, \eta_b \in W(G, f)$

2.3.4. The matrix for the action of f_* on $W(G, f)$. With respect to the basis of $W(G, f)$ described in §2.3.1, 2.3.2 and 2.3.3, let A denote the matrix representing the map $f_*|_{W(G, f)}$. With this A we can compute the homology polynomial $h(x)$.

We compute A explicitly as follows: Let e_1, \dots, e_n be the edges of G . Let ζ_1, \dots, ζ_n be the standard basis of $V(G) \cong \mathbb{R}^n$, where $\zeta_i(e_j) = \delta_{i,j}$ the Kronecker delta. Let $l = \dim W(G, f)$. Suppose that $\{\eta_1, \dots, \eta_l\}$ is a basis constructed as in §2.3.1, 2.3.2 and 2.3.3. Reordering the labels, if necessary, we may assume $\eta_j = \eta_{e_j}$ for $j = 1, \dots, l$. Now let Q be an $n \times l$ matrix whose entries $q_{i,j} \in \mathbb{Z}$ satisfy $\eta_j = \sum_{i=1}^n q_{i,j} \zeta_i$. Let T be the transition matrix for the train track map $f : G \rightarrow G$, and let $P : \mathbb{R}^n \rightarrow \mathbb{R}^l$ be the projection onto the first l coordinates. Then A is an $l \times l$ integer matrix with $A = PTQ$. See Example 5.1 for the calculation of the matrix A . We note that Corollary 2.24, in the next section, implies that $A \in GL(l, \mathbb{Z})$.

Remark 2.18. A related question is the computation of the ‘vertex polynomial’ $f_*|_{\text{im } \delta}$, even though that polynomial is not a topological invariant and so is only of passing interest. We mention it because in special cases it may be easiest to compute the homology polynomial from the characteristic polynomial of T by dividing by the vertex polynomial. See examples 5.2 and 5.4 below, where the computation of the vertex polynomial is carried out in two cases, using data that is supplied by XTrain.

2.4. The orientation cover and the homology polynomial. While we have the first decomposition theorem in hand, and have learned how to compute the homology polynomial, that is the characteristic polynomial of the action of f_* on $W(G, f)$, we have not proved that it is an invariant of $[F]$ and we do not understand its topological meaning. All that will be remedied in this section. Our work begins by recalling the definition of the orientation cover \tilde{S} of S , introduced by Thurston (p.427 of [11]), see also p.184 of [8]. After that we will establish several of its properties. See Proposition 2.21. In Theorem 2.23 and Corollary 2.24 we study the homology space $H_1(\tilde{S}; \mathbb{R})$ and its relationship to our vector space $W(G, f)$ in the case when τ is non-orientable. At the end of the section, in Corollary 2.25, we establish the important and fundamental result that the homology polynomial is an invariant of the mapping class $[F]$ in $\text{Mod}(S)$.

Definition 2.19. (Angle between two branches at a switch) For a switch in τ , fix a very small neighborhood that only contains the switch and the branches meeting at the switch. Within this neighborhood, we orient the branches in the direction outward from the switch. This allows us to define the *angle* between two branches that meet at the switch. Since they always meet tangentially, this angle is either 0 or π . If the angle

is 0, then we say that the branches form a *corner*. For example, the angle between the branches a and b_1 in Figure 2 (B) is π , whereas the angle between the branches b_1 and b_2 is 0 and b_1, b_2 form a corner.

Definition 2.20. (The orientation cover) Let $F : S \rightarrow S$ be a pA homeomorphism with non-orientable train track τ . Add a puncture to S for each 2-cell of $S \setminus \tau$ corresponding to an odd or even vertex of G . The resulting surface S' deformation retracts to τ , i.e., $\pi_1(S') = \pi_1(\tau)$. Each (not necessarily smooth) loop $\gamma \subset \tau$ consists of branches of τ . We define a homomorphism $\theta : \pi_1(\tau) \rightarrow \mathbb{Z}/2\mathbb{Z}$ which maps a loop in τ to 0 if and only if it has an even number of corners (see Definition 2.19). The *orientation cover* $p : \tilde{S} \rightarrow S$ associated to τ is obtained from the double cover \tilde{S}' of S corresponding to $\ker \theta$ by filling in the punctures in \tilde{S}' that do not belong to the original punctures of S .

At the same time, the non-orientable train track τ lifts to an orientable train track $\tilde{\tau} \subset \tilde{S}$. Collapsing the infinitesimal (partial) polygons in $\tilde{\tau}$ to vertices, we obtain a graph \tilde{G} , that is a double branched cover of G .

Note that the branch points of $p : \tilde{S} \rightarrow S$ are precisely the odd vertices of G . Intuitively, the effect of passing to the orientation cover is a partial unrolling of loops in τ that do not admit a consistent orientation. For example, Figure 4 illustrates what happens near the branch point for a vertex of valence three.




FIGURE 4. Unrolling an odd vertex.

Proposition 2.21. *Assume that τ is non-orientable.*

- (1) *The orientation cover $\tilde{\tau}$ is an orientable train track. The natural involution $\iota : \tilde{S} \rightarrow \tilde{S}$, or the deck transformation, reverses the orientation of $\tilde{\tau}$.*
- (2) *A puncture of S corresponds to two punctures of \tilde{S} if and only if a loop around the puncture is homotopic to a loop in τ with an even number of corners. Otherwise, the puncture lifts to one puncture in \tilde{S} .*

Proof. (1) By Definition 2.20 we have $p_*(\pi_1(\tilde{\tau})) = \ker \theta$, hence every loop in $\tilde{\tau}$ has an even number of corners, and so $\tilde{\tau}$ can be consistently oriented. If the involution ι did not reverse the orientation of $\tilde{\tau}$, then the orientation of $\tilde{\tau}$ would induce a consistent orientation of τ , but τ is not orientable.

(2) A loop in τ lifts to two loops in $\tilde{\tau}$ if and only if it has an even number of corners. If it has an odd number of corners, then its concatenation with itself has a unique lift. \square

Proposition 2.22. *A train track map $f : G \rightarrow G$ has two lifts $\tilde{f}_{\text{op}} : \tilde{G} \rightarrow \tilde{G}$, orientation preserving, and $\tilde{f}_{\text{or}} : \tilde{G} \rightarrow \tilde{G}$, orientation reversing. They are related to each other by $\tilde{f}_{\text{or}} = \iota \cdot \tilde{f}_{\text{op}}$, where $\iota : \tilde{G} \rightarrow \tilde{G}$ is the deck translation. Let n be the number of edges in G . Then there exist $n \times n$ non-negative matrices A, B such that $A + B = T$, the transition matrix of f , and \tilde{f}_{op} and \tilde{f}_{or} are represented as: $\tilde{f}_{\text{op}} = \begin{bmatrix} A & B \\ B & A \end{bmatrix}$ and $\tilde{f}_{\text{or}} = \begin{bmatrix} B & A \\ A & B \end{bmatrix}$. Their characteristic polynomials are $\chi((\tilde{f}_{\text{op}})_*) = \chi(f_*) \det(A - B)$ and $\chi((\tilde{f}_{\text{or}})_*) = \chi(f_*) \det(B - A)$.*

The above proposition confirms that pA maps, $F : S \rightarrow S$; $\tilde{F}_{\text{op}} : \tilde{S} \rightarrow \tilde{S}$; and $\tilde{F}_{\text{or}} : \tilde{S} \rightarrow \tilde{S}$, all have the same dilatation.

Proof. Even though the train track τ is non-orientable, we assign an orientation to each edge of G . Let e_1, \dots, e_n be the oriented edges of G . We denote the lifts of $e_k \subset G$ by $\tilde{e}_k, \tilde{e}'_k \subset \tilde{G}$. Since $\tilde{\tau}$ is orientable, we choose an orientation, which induces orientations of $\tilde{e}_k, \tilde{e}'_k$. Denote the orientation cover by $p : \tilde{G} \rightarrow G$. Proposition 2.21-(1) implies that there are two choices: $p(\tilde{e}_k) = e_k$ or $p(\tilde{e}_k) = \bar{e}_k$, where \bar{e}_k is the edge $e_k \subset G$ with reversed orientation. We choose to assume that $p(\tilde{e}_k) = e_k$, which implies that $p(\tilde{e}'_k) = \bar{e}_k$.

We define an orientation preserving lift $\tilde{f}_{\text{op}} : \tilde{G} \rightarrow \tilde{G}$ in the following way. For an edge $e \subset G$, let $f(e)_{\text{head}}$ (resp. $f(e)_{\text{tail}}$) denote the first (resp. last) letter of the word $f(e)$. For each twin edges $\tilde{e}, \tilde{e}' \subset \tilde{G}$, we choose

$$(2.3) \quad \tilde{f}_{\text{op}}(\tilde{e})_{\text{head}} := \widetilde{f(e)_{\text{head}}}, \quad \tilde{f}_{\text{op}}(\tilde{e}')_{\text{head}} := \widetilde{f(e)_{\text{tail}}}.$$

Next, we define the word $\tilde{f}_{\text{op}}(\tilde{e})$ to be the word $f(e)$ with each letter e_i in $f(e)$ replaced by \tilde{e}_i or \tilde{e}'_i so that the resulting word corresponds to a connected edge-path in \tilde{G} . Due to the choice (2.3), the choice between \tilde{e}_i and \tilde{e}'_i is uniquely determined. The word $\tilde{f}_{\text{op}}(\tilde{e}')$ is given by the word $\tilde{f}_{\text{op}}(\tilde{e})$ read from the right to left, then replace \tilde{e}_i by \tilde{e}'_i and \tilde{e}'_i by \tilde{e}_i .

We define an orientation reversing train track map by $\tilde{f}_{\text{or}} := \iota \cdot \tilde{f}_{\text{op}}$.

Let $\{\zeta_1, \dots, \zeta_n, \zeta'_1, \dots, \zeta'_n\}$ be the standard basis of $V(\tilde{G}) \simeq \mathbb{R}^{2n}$, where ζ_k, ζ'_k correspond to $\tilde{e}_k, \tilde{e}'_k \subset \tilde{G}$ respectively. From the constructions of \tilde{f}_{op} and \tilde{f}_{or} , with respect to this basis, their transition matrices are of the form $\begin{bmatrix} A & B \\ B & A \end{bmatrix}$ and $\begin{bmatrix} B & A \\ A & B \end{bmatrix}$ respectively, for some non-negative $n \times n$ matrices A and B satisfying $A + B = T$. The formulae on characteristic polynomials follow from basic row and column reductions. \square

In Example 5.3 there is a sketch of the orientation cover of a non-orientable train track which has no odd vertices. Also one can see explicit computations of \tilde{f}_{op} and \tilde{f}_{or} and matrices A, B .

We are finally in a position to understand the topological meaning of $W(G, f)$:

Theorem 2.23. *Assume that τ is non-orientable. Let $\iota : \tilde{S} \rightarrow \tilde{S}$ be the involution of the orientation cover. Let E^+ and E^- be the eigenspaces of $\iota_* : H_1(\tilde{S}; \mathbb{R}) \rightarrow H_1(\tilde{S}; \mathbb{R})$ corresponding to the eigenvalues 1 and -1 , so that $H_1(\tilde{S}; \mathbb{R}) \cong E^+ \oplus E^-$. Then $E^+ \cong H_1(S; \mathbb{R})$ and $E^- \cong W(G, f)$.*

Proof. Fix an orientation of $\tilde{\tau}$ once and for all. This determines an orientation of \tilde{G} . Since ι is an involution, the only possible eigenvalues are ± 1 , and ι_* is diagonalizable.

For a homology class $\xi \in H_1(S; \mathbb{R})$, let $\tilde{\xi} \in H_1(\tilde{S}; \mathbb{R})$ denote its lift to the orientation cover. Since $p \cdot \iota = p$, we have $\iota_* \tilde{\xi} = \tilde{\xi}$. Thus $H_1(S; \mathbb{R}) \subseteq E^+$.

Each edge $e \subset G$ has two lifts \tilde{e} and $\tilde{e}' \subset \tilde{G}$. Let ζ_e be the basis element of $V(G)$ corresponding to the (unoriented) edge e , and let $\zeta_{\tilde{e}}$ be the basis element of $C_1(\tilde{G}; \mathbb{R}) \cong C_1(\tilde{S}; \mathbb{R})$ corresponding to the (oriented) edge \tilde{e} . Define a homomorphism: $\phi : V(G) \rightarrow C_1(\tilde{S}; \mathbb{R})$ by $\zeta_e \mapsto \zeta_{\tilde{e}} + \zeta_{\tilde{e}'}$. Recall that a basis element η_e of $W(G, f)$, introduced in § 2.3.2 and § 2.3.3, has a corresponding arc or loop $L_e \subset G$ where η_e assigns weight of ± 1 satisfying the alternating sum condition. Different edges e_1, e_2 correspond to distinct L_{e_1} and L_{e_2} . Moreover, when L_e is an arc, its end points are odd vertices which are branch points of the orientation cover, so the lift of L_e is a closed curve in \tilde{S} . Hence, the restriction of ϕ to $W(G, f)$ is injective. Assume $\eta_e = \sum_i \eta_e(e_i) \zeta_{e_i}$. By Proposition 2.21-(1), the involution takes $\iota : \tilde{e} \mapsto -\tilde{e}'$. We have:

$$\iota_* \phi(\eta_e) = \iota_* \left(\sum_i \eta_e(e_i) (\zeta_{\tilde{e}_i} + \zeta_{\tilde{e}'_i}) \right) = \sum_i \eta_e(e_i) (-\zeta_{\tilde{e}'_i} - \zeta_{\tilde{e}_i}) = -\phi(\eta_e),$$

i.e., $\phi(W(G, f)) \subseteq E^-$.

Comparison of Euler characteristics along with Lemma 2.11 shows that

$$\dim H_1(\tilde{S}; \mathbb{R}) = \dim H_1(S; \mathbb{R}) + \dim W(G, f),$$

which implies $H_1(S; \mathbb{R}) \cong E^+$ and $\phi(W(G, f)) \cong E^-$. \square

In Lemma 2.10 we proved that $f_*(W(G, f)) \subseteq W(G, f)$. In fact, a stronger statement holds.

Corollary 2.24. *The restriction map $f_*|_{W(G, f)} : W(G, f) \rightarrow W(G, f)$ is an isomorphism.*

Proof. Regardless of the orientability of τ , the fact that $F : S \rightarrow S$ is a homeomorphism implies that the induced map $F_* : H_1(S; \mathbb{R}) \rightarrow H_1(S; \mathbb{R})$ is an isomorphism.

Suppose that τ is orientable. The isomorphism $W(G, f) \cong H_1(S; \mathbb{R})$ in Lemma 2.11 allows us to identify $f_*|_{W(G, f)}$ with F_* , which is an isomorphism.

Suppose that τ is non-orientable. Let $\{\eta_e\}_{e \in E}$ be a basis of $W(G, f)$ constructed as in § 2.3.2 and § 2.3.3. Since the map $\phi : W(G, f) \rightarrow E^-$ in the proof of Theorem 2.23 is an isomorphism, the set $\{\phi(\eta_e)\}_{e \in E}$ is a basis of E^- . Let $\tilde{F} : \tilde{S} \rightarrow \tilde{S}$ be a lift of $F : S \rightarrow S$. It induces an isomorphism $\tilde{F}_* : H_1(\tilde{S}; \mathbb{R}) \rightarrow H_1(\tilde{S}; \mathbb{R})$ and a train track map $\tilde{f} : \tilde{G} \rightarrow \tilde{G}$.

Since \tilde{S} deformation retracts to \tilde{G} , we can identify \tilde{F}_* with $\tilde{f}_* : H_1(\tilde{G}; \mathbb{R}) \rightarrow H_1(\tilde{G}; \mathbb{R})$. Using the same notation as in the proof of Theorem 2.23, we have

$$\begin{aligned} \iota_* \tilde{f}_*(\phi(\eta_e)) &= \iota_* \tilde{f}_* \left(\sum_i \eta_e(e_i) (\zeta_{\tilde{e}_i} + \zeta_{\tilde{e}_i'}) \right) \\ &= \iota_* \sum_i \eta_e(e_i) (\zeta_{\tilde{f}_*(\tilde{e}_i)} + \zeta_{\tilde{f}_*(\tilde{e}_i')}) \\ &= \sum_i \eta_e(e_i) (-\zeta_{\tilde{f}_*(\tilde{e}_i')} - \zeta_{\tilde{f}_*(\tilde{e}_i)}) = -\tilde{f}_*(\phi(\eta_e)). \end{aligned}$$

Hence $\tilde{F}_*(E^-) = \tilde{f}_*(E^-) \subset E^-$. Since $H_1(\tilde{S}; \mathbb{R})$ is finite dimensional and \tilde{F}_* is an isomorphism, we obtain that $\tilde{F}_*|_{E^-} = \tilde{f}_*|_{W(G,f)}$ is an isomorphism. \square

The topological invariance of $\chi(f_*|_{W(G,f)})$ was stated as a conjecture in an earlier draft. Reading that draft, Jeffrey Carlson pointed the authors to a connection they had missed, making our conjecture an immediate consequence of Theorem 2.23. We are grateful for his help.

Corollary 2.25. *Let $h(x)$ be the characteristic polynomial for $f_*|_{W(G,f)}$. It is the homology polynomial. Then $h(x)$ is an invariant of the pA mapping class $[F]$.*

Proof. Let $\tilde{F} : \tilde{S} \rightarrow \tilde{S}$ be a lift of the pA map $F : S \rightarrow S$. By Theorem 2.23 we have $h(x) = \chi(f_*|_{W(G,f)}) = \chi(\tilde{F}_*|_{E^-})$. Since the eigenspace E^- is an invariant of $[F]$, so is the polynomial $h(x)$. \square

This concludes the proof of Part (1) of Theorem 1.1.

3. PROOF OF PARTS (2) AND (3) OF THEOREM 1.1

Having established the meaning of $W(G, f)$ and the invariance of the homology polynomial, our next goal is to understand whether it is irreducible, and if not to understand its factors. At the same time, we will investigate its symmetries.

With those goals in mind, we show that there is a well-defined and f_* -invariant skew-symmetric form on the space $W(G, f)$. See Proposition 3.3. We define the subspace $Z \subset W(G, f)$ to be the space of degeneracies of this skew-symmetric form. We are able to interpret the action of f_* on Z geometrically, as being a permutation of certain punctures on S . In Theorem 3.8 we will prove that the space $W(G, f)$ has a decomposition into summands that are invariant under the action of f_* , and that as a consequence the homology polynomial decomposes as a product of two polynomials, $p(x)$ and $s(x)$. We call them the *puncture* and *symplectic* polynomials. Like the homology polynomial, both are invariants of $[F]$ in $\text{Mod}(S)$. We also establish their symmetries in Theorem 3.8, and understand the precise meaning of the puncture polynomial. The symplectic polynomial contains λ as its largest real root. When irreducible, it coincides with the minimum polynomial of λ , but in general it is not irreducible. At this writing we do not understand when it is or is not reducible.

3.1. Lifting the basis elements for $W(G, f)$ to $W(\tau)$. Our work begins with a brief diversion, to establish a technical result that will be needed in the sections that follow. We have shown how to construct basis elements $\eta_{e_1}, \dots, \eta_{e_l}$ for $W(G, f)$. We now build on this construction to give an explicit way to lift each $\eta_e \in W(G, f)$ to an element $\eta'_e \in W(\tau)$ in such a way that $\pi_*(\eta'_e) = \eta_e$, where $\pi_* : W(\tau) \rightarrow W(G, f)$ is the natural surjection. Although there are infinitely many lifts of any given basis element η_e , our construction of a specific η'_e will be useful later. The issues to be faced in lifting η_e to η'_e are the assignment of weights to the infinitesimal edges.

Definition 3.1. Suppose that $v \in G$ is a vertex with k gates numbered $0, \dots, k-1$, counterclockwise. For $i = 0, \dots, k-1$, we define a *transitional* element $\sigma_i \in V(\tau)$ that assigns 1 to the i -th infinitesimal edge for the vertex v , and 0 to the remaining branches of τ . In other words, in Figure 1, $x_i = 1$ and $x_j = 0$ for $j \neq i$.

Suppose that $v \in G$ is an odd vertex with k gates. For $i = 0, \dots, k-1$, we define a *terminal* element $\omega_i \in V(\tau)$ which assigns $\pm \frac{1}{2}$ to the incident infinitesimal edges for v so that the i -th gate has weight $w_i = 1$ and j -th ($j \neq i$) gate has weight $w_j = 0$, cf. Figure 5; and assigns 0 for rest of the branches of τ .




FIGURE 5. Transitional and terminal elements.

For both the orientable and the non-orientable case, our basis element $\eta_e \in W(G, f)$ is a vector whose entries are ± 1 or 0. Recall that the edges whose weights are ± 1 form a loop or an arc, denoted by L_e in § 2.3.1, 2.3.2 and 2.3.3.

At an even or a partial vertex of L_e , suppose L_e goes through the i -th and j -th gates ($i \leq j$). To η_e we add consecutive transitional elements $\sigma_i, \sigma_{i+1}, \dots, \sigma_{j-1}$ with alternating signs so that the switch condition is satisfied (cf. the upper middle circle in Figure 5). Repeat this procedure for all the non-odd vertices of L_e . If L_e is a loop, it yields an element η'_e of $W(\tau)$.

When L_e is an arc (i.e., τ is non-orientable with odd vertices), the two endpoints of L_e are odd. Suppose L_e enters the i -th gate of an odd vertex. After adding transitional

elements as above, we further add terminal element ω_i or $-\omega_i$ so that the switch condition is satisfied at all the incident gates of the odd vertex. Proceed in this way for the other odd vertex as well, and we obtain an element η'_e of $W(\tau)$.

3.2. A skew-symmetric form on $W(G, f)$. In this section we define a skew-symmetric form $\langle \cdot, \cdot \rangle$ on $W(G, f)$. To get started, we slightly modify the skew-symmetric form on $W(\tau)$ introduced by Penner-Harer (p.182 of [8]): For a branch $b \subset \tau$ and $\eta \in W(\tau)$, let $\eta(b)$ denote the weight that η assigns to b . At a switch of valence k ($k \geq 3$), label the branches a, b_1, \dots, b_{k-1} as in Figure 2. The cyclic order of a, b_1, \dots, b_{k-1} is determined by the orientation of the surface and the embedding of τ in S . We define a skew-symmetric form:

$$\langle \eta, \zeta \rangle_{W(\tau)} := \frac{1}{2} \sum_{\substack{\text{switches} \\ \text{in } \tau}} \sum_{i < j} \begin{vmatrix} \eta(b_i) & \eta(b_j) \\ \zeta(b_i) & \zeta(b_j) \end{vmatrix}, \quad \text{for } \eta, \zeta \in W(\tau).$$

Recall the surjective map $\pi : \tau \rightarrow G$ collapsing the infinitesimal (partial) polygons to vertices.

Definition 3.2. For $\eta, \zeta \in W(G, f)$ there exist $\eta', \zeta' \in W(\tau)$ so that $\pi_*(\eta') = \eta$ and $\pi_*(\zeta') = \zeta$. We define a skew-symmetric form on $W(G, f)$ by:

$$\langle \eta, \zeta \rangle_{W(G, f)} := \langle \eta', \zeta' \rangle_{W(\tau)}$$

Proposition 3.3. *The skew-symmetric form $\langle \cdot, \cdot \rangle_{W(G, f)}$ has the following properties:*

- (1) *It is well-defined.*
- (2) *When τ is orientable, $\langle \eta, \zeta \rangle_{W(G, f)}$ is the homology intersection number of 1-cycles associated to η and ζ .*
- (3) *When τ is non-orientable, recall that E^\pm are the eigenspaces of the deck transformation $\iota : \tilde{S} \rightarrow \tilde{S}$ for the orientation cover studied in Theorem 2.23. Since $p : \tilde{S} \rightarrow S$ is a double branched cover, we have the following results, to be compared with p.187 of [8]:*
 - (a) *The restriction of the intersection form on $H_1(\tilde{S}; \mathbb{R})$ to E^+ is twice the intersection form on $H_1(S; \mathbb{R})$.*
 - (b) *The restriction of the intersection form on $H_1(\tilde{S}; \mathbb{R})$ to E^- is twice the skew-symmetric form $\langle \cdot, \cdot \rangle_{W(G, f)}$.*
- (4) *For all $\eta, \zeta \in W(G, f)$, we have $\langle f_*\eta, f_*\zeta \rangle_{W(G, f)} = \langle \eta, \zeta \rangle_{W(G, f)}$.*

Proof. (1) It suffices to show that for any $\eta' \in \ker \pi_* \subset W(\tau)$ and $\zeta' \in W(\tau)$, the product $\langle \eta', \zeta' \rangle_{W(\tau)} = 0$. Since $\pi_*(\eta') = \vec{0}$, η' assigns 0 to any real edge $b \subset \tau$, it follows that the weight of η' at any gate is 0. If a vertex $v \in G$ is odd or partial, η' assigns 0 to any infinitesimal edges associated to the vertex v . Therefore, v does not contribute to $\langle \eta', \zeta' \rangle_{W(\tau)}$. If a vertex $v \in G$ is even with k gates, the weights x_0, \dots, x_{k-1} that η' assigns to the infinitesimal edges for v form an alternating sequence: $x_i = (-1)^i x_0$. The

contribution of the even vertex v to $\langle \eta', \zeta' \rangle_{W(\tau)}$ is:

$$\frac{1}{2} \sum_{i=0}^{k-1} \begin{vmatrix} x_i & x_{i-1} \\ y_i & y_{i-1} \end{vmatrix} = \frac{x_0}{2} \sum_{i=0}^{k-1} \begin{vmatrix} (-1)^i & (-1)^{i-1} \\ y_i & y_{i-1} \end{vmatrix} = \frac{x_0}{2} \sum_{i=0}^{k-1} (-1)^i (y_i + y_{i-1}) = 0,$$

where y_i are the weights assigned by ζ' and indices are modulo k .

(2) Assertion (2) is established in Lemma 3.2.2 of [8].

(3) Assertion (3) follows directly from Theorem 2.23.

(4) If τ is orientable, then we identify $W(G, f) \cong H_1(S; \mathbb{R})$. Since $F : S \rightarrow S$ is homeomorphism, the homology intersection number is preserved under $f_* : H_1(S; \mathbb{R}) \rightarrow H_1(S; \mathbb{R})$ and the assertion follows.

If τ is non-orientable, then, passing to the orientation cover, $\tilde{F}_* : H_1(\tilde{S}; \mathbb{R}) \rightarrow H_1(\tilde{S}; \mathbb{R})$ preserves the homology intersection number. The assertion then follows from (i) $\tilde{F}_*|_{E^-} = f_*|_{W(G, f)}$ and (ii) assertion (3) of this proposition. \square

Knowing that $\langle \cdot, \cdot \rangle_{W(G, f)}$ is well-defined, we can compute $\langle \eta_1, \eta_2 \rangle_{W(G, f)} = \langle \eta'_1, \eta'_2 \rangle_{W(\tau)}$ by using the basis elements $\eta_1, \eta_2 \in W(G, f)$ discussed in § 2.3.1, 2.3.2 and 2.3.3, and their particular extensions $\eta'_1, \eta'_2 \in W(\tau)$ introduced in § 3.1. For this, it is convenient to study how transitional and terminal elements contribute to the skew-symmetric form. Straight forward calculation of determinants at incident gates yields the following:

Proposition 3.4. *Let v be a vertex with k incident gates, numbered $0, \dots, k-1$, counter-clockwise. We have $\langle \sigma_i, \sigma_j \rangle = \langle \sigma_0, \sigma_{j-i} \rangle$, $\langle \omega_i, \sigma_j \rangle = \langle \omega_0, \sigma_{j-i} \rangle$ and $\langle \omega_i, \omega_j \rangle = \langle \omega_0, \omega_{j-i} \rangle$ for $0 \leq i \leq j \leq k-1$. Moreover,*

$$\langle \sigma_0, \sigma_i \rangle = \begin{cases} -\frac{1}{2} & \text{if } i = 1 \\ \frac{1}{2} & \text{if } i = k-1 \\ 0 & \text{otherwise} \end{cases} \quad \langle \omega_0, \sigma_i \rangle = \begin{cases} -\frac{1}{2} & \text{if } i = 0 \\ \frac{1}{2} & \text{if } i = k-1 \\ 0 & \text{otherwise} \end{cases}$$

$$\text{and} \quad \langle \omega_0, \omega_i \rangle = \begin{cases} \frac{(-1)^i}{2} & \text{if } i \neq 0 \\ 0 & \text{if } i = 0. \end{cases}$$

3.3. Degeneracies of the skew-symmetric form and the second decomposition.

In this section, we investigate the totally degenerate subspace of $W(G, f)$,

$$Z := \{ \eta \in W(G, f) \mid \langle \zeta, \eta \rangle = 0 \text{ for all } \zeta \in W(G, f) \},$$

the *radical* of the skew-symmetric form. It will lead us, almost immediately, to the second decomposition theorem and another new invariant of pA maps. We begin by showing how Z has already appeared in our work, in a natural way.

Proposition 3.5. *Let s be the number of punctures of S .*

- (1) *If τ is orientable, $\dim Z = s - 1$.*

(2) If τ is non-orientable, then

$$\begin{aligned} \dim Z &= \#(\text{punctures of } S \text{ that correspond to two punctures in } \tilde{S}) \\ &= \#(\text{punctures of } S \text{ represented by loops in } \tau \\ &\quad \text{with even numbers of corners}). \end{aligned}$$

Proof. In the orientable case, recall Lemma 2.11 which states $W(G, f) \cong H_1(S; \mathbb{R})$, and Lemma 3.2.2 of [8] which shows that our skew-symmetric form agrees with the homology intersection form. Hence the space Z is generated by the homology classes of s loops around the punctures. Because their sum is null-homologous, they are linearly dependent and $\dim Z = s - 1$.

In the non-orientable case, let \tilde{S} be the orientation cover of S (Definition 2.20). Recall the eigen spaces E^\pm for the deck transformation $\iota : \tilde{S} \rightarrow \tilde{S}$, cf. Theorem 2.23. Let s (resp. r) be the number of punctures of S that lift to two (resp. single) punctures in \tilde{S} , and let $\alpha_1, \beta_1, \dots, \alpha_s, \beta_s, \gamma_1, \dots, \gamma_r$ be the homology classes of loops around the punctures of \tilde{S} , chosen so that $\iota_* \alpha_i = \beta_i$ for all $i = 1, \dots, s$ and $\iota_* \gamma_j = \gamma_j$ for all $j = 1, \dots, r$ and oriented so that their sum is zero. The radical of the homology intersection form on $H_1(\tilde{S}; \mathbb{R})$ is spanned by

$$\begin{aligned} &\text{span}\{\alpha_1, \beta_1, \dots, \alpha_s, \beta_s, \gamma_1, \dots, \gamma_r\} \\ &= \text{span}\{\alpha_1 - \beta_1, \dots, \alpha_s - \beta_s, \alpha_1 + \beta_1, \dots, \alpha_s + \beta_s, \gamma_1, \dots, \gamma_r\}. \end{aligned}$$

Note that $\alpha_i - \beta_i \in E^-$ and $\alpha_i + \beta_i, \gamma_j \in E^+$ for all $i = 1, \dots, s$ and $j = 1, \dots, r$. By Theorem 2.23 and assertion (3) of Proposition 3.3, we obtain that $Z = \text{span}\{\alpha_1 - \beta_1, \dots, \alpha_s - \beta_s\}$. Clearly, $\alpha_1 - \beta_1, \dots, \alpha_s - \beta_s$ are linearly independent, i.e., $\dim Z = s$. \square

Corollary 3.6. *Assume that S is once punctured and that τ is non-orientable. Then $\dim Z = 1$ if and only if $\dim W(G, f)$ is odd.*

Proof. The induced skew-symmetric form on $W(G, f)/Z$ is non-degenerate, and so the dimension of $W(G, f)/Z$ is even. Thus, $\dim Z$ is odd if and only if $\dim W(G, f)$ is odd. Since S has exactly one puncture, Proposition 3.5 yields that $\dim Z \leq 1$, and the corollary follows. \square

Remark 3.7. Straightforward modifications of our arguments show that if τ is a non-orientable train track (not necessarily induced by a train track map), then the dimension of $\text{rad } W(\tau)$ is the number of complementary regions of τ with even numbers of corners. In particular, if τ is complete, then the complement of τ consists of triangles and monogons (Theorem 1.3.6 of [8]), and so the skew-symmetric form on $W(\tau)$ is non-degenerate in this case. Although not explicitly stated, this is the case covered by Theorem 3.2.4 of [8].

Theorem 3.8 (Second Decomposition). *Let $p(x)$ (resp. $s(x)$) be the characteristic polynomial of $f_*|_Z$ (resp. $f_*|_{W(G, f)/Z}$). The map f_* preserves the direct sum decomposition $W(G, f) \cong Z \oplus (W(G, f)/Z)$ so that $h(x) = p(x)s(x)$. Moreover, we have:*

- (1) *The polynomial $p(x)$ is an invariant of the pA mapping class $[F] \in \text{Mod}(S)$. The restriction $f_*|_Z$ encodes how F permutes the punctures whose projections to τ have even numbers of corners. In particular, $f_*|_Z$ is a periodic map, so that all the roots of $p(x)$ are roots of unity and the polynomial $p(x)$ is palindromic or anti-palindromic.*
- (2) *The polynomial $s(x)$ is an invariant of $[F]$. The skew-symmetric form $\langle \cdot, \cdot \rangle_{W(G,f)}$ naturally induces a symplectic form on $W(G, f)/Z$. The map f_* induces a symplectomorphism of $W(G, f)/Z$. Hence $s(x)$ is palindromic.*
- (3) *The homology polynomial $h(x)$ is either palindromic or anti-palindromic.*

Proof. Suppose $\eta \in Z$. By assertion (4) of Proposition 3.3, we have $0 = \langle \eta, \zeta \rangle = \langle f_*(\eta), f_*(\zeta) \rangle$ for all $\zeta \in W(G, f)$. By Corollary 2.24, $f_*|_{W(G,f)}$ is surjective, and so $f_*(\eta) \in Z$. Thus f_* preserves the decomposition $(W(G, f)/Z) \oplus Z$.

(1) The restriction $f_*|_Z$ is periodic because of Proposition 3.5 and the fact that $[F]$ permutes the punctures of S . Hence all the roots of $p(x)$ are roots of unity. Moreover, if μ is a root of $p(x)$, then $\frac{1}{\mu} = \bar{\mu}$ is also a root of $p(x)$ because $p(x) \in \mathbb{R}[x]$. This implies that $p(x)$ is palindromic or anti-palindromic.

(2) By the definition of Z , the skew-symmetric form induces a non-degenerate form on $W(G, f)/Z$. This together with Proposition 3.2-(4) implies that the polynomial $s(x)$ is palindromic. It is an invariant of $[F]$ because it is the quotient of two polynomials, both of which have been proved to be invariants.

(3) The homology polynomial $h(x)$ is either palindromic or anti-palindromic because it is a product of two polynomials, one of which is palindromic and the other of which is either palindromic or anti-palindromic. \square

This concludes the proof of parts (2) and (3) of Theorem 1.1. Since the proof of part (1) was completed in §2, it follows that Theorem 1.1 has been proved.

4. APPLICATIONS

In this section we give several applications of Theorem 1.1. Corollary 4.1 summarizes the numerical class invariants of $[F]$ that, as a consequence of Theorem 1.1, can be computed from the train track τ by simple counting arguments. Corollary 4.3 is an application to fibered hyperbolic knots in 3-manifolds. Corollary 4.5 shows that our three polynomials behave very nicely under the passage $[F] \rightarrow [F^k]$.

Corollary 4.1. *Under the same notation as in Theorem 1.1, let n, v, v_o be the number of edges, vertices, odd vertices respectively in the graph G . Let s (resp. r) be the number of punctures of S which are represented by loops in τ with even (resp. odd) numbers of corners. Let g (resp. \tilde{g}) be the genus of S (resp. its orientation cover \tilde{S}).*

- (1) *The orientability (or non-orientability) of τ is a class invariant of $[F]$.*

(2) *The degree of the homology polynomial is a class invariant. It is:*

$$\deg h(x) = \begin{cases} n - v + 1 & \text{if } \tau \text{ is orientable,} \\ n - v + v_o & \text{if } \tau \text{ is not orientable.} \end{cases}$$

(3) *The degree of the puncture polynomial is a class invariant. It is:*

$$\deg p(x) = \begin{cases} s - 1 & \text{if } \tau \text{ is orientable,} \\ s & \text{if } \tau \text{ is not orientable.} \end{cases}$$

(4) *The degree of the symplectic polynomial is a class invariant. It is:*

$$\deg s(x) = \begin{cases} 2g & \text{if } \tau \text{ is orientable,} \\ 2(\tilde{g} - g) & \text{if } \tau \text{ is not orientable.} \end{cases}$$

Remark 4.2. Assertion (4) implies that the dilatation of $[F]$ is the largest real root of a polynomial of degree $2d$, where $2d \leq 2g$ (resp. $2(\tilde{g} - g)$). However, this bound is not sharp because, as will be seen in Example 5.2, the symplectic polynomial is not necessarily irreducible.

Proof. Assertion (1) is clear. See Lemma 2.11 for (2). In the orientable case each puncture is represented by a loop in τ with an even number of corners. This, together with Proposition 3.5, implies (3).

To prove (4), let v_e, v_p be the number of even vertices, partial vertices in G , respectively. If τ is orientable, the assertion is clear. If τ is non-orientable, the Euler characteristics of S and \tilde{S} are

$$\begin{aligned} \chi(S) &= v_o + v_e + v_p - n = 2 - 2g - (r + s), \\ \chi(\tilde{S}) &= v_o + 2v_e + 2v_p - 2n = 2 - 2\tilde{g} - (r + 2s). \end{aligned}$$

From Lemma 2.11 and Proposition 3.5, we have:

$$\dim(W(G, f)/Z) = (n - v_e - v_p) - s = 2(\tilde{g} - g).$$

Note that the sum of the degrees of the symplectic and puncture polynomials is the degree of the homology polynomial. \square

Corollary 4.3. (1) *The symplectic polynomial $s(x)$ is an invariant of fibered hyperbolic links in 3-manifolds.*

(2) *Assume that M is a homology 3-sphere and $K \subset M$ is a fibered hyperbolic knot whose monodromy admits an orientable train track. Let $\Delta_K(x)$ denote the Alexander polynomial of K . Then*

$$s(x) = \begin{cases} \Delta_K(x) & \text{if } f \text{ is orientation preserving,} \\ \Delta_K(-x) & \text{if } f \text{ is orientation reversing.} \end{cases}$$

Proof. (1) Let $L \subset M$ be a link. Thurston proved that a 3-manifold $M \setminus L$ is fibered over S^1 with a pA monodromy $[F] \in \text{Mod}(S)$ if and only if $M \setminus L$ is hyperbolic. Combining his result with assertion (2) of Theorem 3.8, we obtain the first claim.

(2) Let $F_* : H_1(S; \mathbb{R}) \rightarrow H_1(S; \mathbb{R})$ be the induced map. By assertion (1) of Proposition 3.5, the space Z is trivial. Lemma 2.11 tells us that $W(G, f)/Z \cong W(G, f) \cong H_1(S; \mathbb{R})$. Since (\pm) -gate is mapped to a (\mp) -gate if and only if $f : G \rightarrow G$ is orientation reversing, we have

$$f_*|_{W(G, f)/Z} = \begin{cases} F_* & \text{if } f \text{ is orientation preserving,} \\ -F_* & \text{if } f \text{ is orientation reversing.} \end{cases}$$

The fact that $\Delta_K(\pm x) = \chi(\pm F_*)$ yields the statement. \square

Remark 4.4. Corollary 4.3-(2) can be seen as a refinement of Rykken's Theorem 3.3 in [10].

Our final application is to prove that our three polynomials, $h(x)$, $p(x)$ and $s(x)$ behave in a very nice way under the passage $[F] \rightarrow [F^n]$:

Corollary 4.5. *Let $n > 0$. If $f : G \rightarrow G$ represents a pA mapping class $[F]$, then f^n represents $[F^n]$. Suppose $s([F], x) = \prod_i (x - z_i)$ and $p([F], x) = \prod_j (x - w_j)$ where $z_i, w_j \in \mathbb{C}$, then*

$$\begin{aligned} s([F^n], x) &= \prod_i (x - z_i^n), \\ p([F^n], x) &= \prod_j (x - w_j^n), \\ h([F^n], x) &= \prod_i (x - z_i^n) \prod_j (x - w_j^n). \end{aligned}$$

Proof. Note that the pA maps $[F]$ and $[F^n]$ act on the same surface and share the same graph G and the associated train track τ . The direct sum decomposition in Theorem 3.8 tells us that $f_*^n|_{W(G, f^n)/Z} = (f_*|_{W(G, f)/Z})^n$ and $f_*^n|_Z = (f_*|_Z)^n$. Since z_i and w_j are eigenvalues of $f_*|_{W(G, f)/Z}$ and $f_*|_Z$ respectively, the desired equations follow. The product decomposition for the homology polynomial follows from the fact that it is a product of the other two polynomials. \square

5. EXAMPLES

All of our examples were analyzed with the software package XTrain [2], with some help from Octave [3]. This package is an adaptation of the Bestvina-Handel algorithm to once-punctured surfaces. Our illustrations show a train track τ (in the sense of [1]) embedded in a once-punctured surface. Regardless of whether τ admits an orientation, we equip individual edges of the graph G with a direction for the purpose of specifying the map $f : G \rightarrow G$, although they coincide when τ is orientable. We remark that the limitations of the available software, at this time, to once-punctured surfaces means that the puncture polynomial in our examples is always either 1 or $x - 1$.

A surface is shown as a fundamental domain in the Poincaré model for \mathbb{H}^2 , with the identification pattern on the boundary given by the labels on edges. For example, a side crossed by an edge labeled a will be identified with the side crossed by the edge labeled

\bar{a} . These labels also indicate the direction of an edge; a is the first half, \bar{a} the second one. The shaded regions in the pictures contain all the infinitesimal edges associated with a vertex, as in Figure 1. To recover the graph G , collapse each shaded region to a point. In each example, we give the associated train track map on edges of G . That map determines the maps on the vertices.

Example 5.1. We illustrate Corollary 2.25 with a triplet of examples. Sketches (1), (2) and (3) of Figure 6 show three copies of a once-punctured genus 2 surface $S = S_{2,1}$, each containing two simple closed curves u_i and v_i . In all three cases $u_i \cup v_i$ fills S , that is the complement of the union of the two curves is a family of discs. In all three cases the geometric intersection $i(u_i \cap v_i) = 6$. We use our curves to define three diffeomorphisms




FIGURE 6. Curves on a surface of genus 2.

of S by the formula $F_i = T_{v_i}^{-1} T_{u_i}$ where T_c denotes a Dehn twist about a simple closed curve c . By a theorem of Thurston (Theorem 14.1 of [4]) there is a representation of the free subgroup of $\text{Mod}(S)$ generated by T_{u_i} and T_{v_i} in $\text{PSL}(2, \mathbb{R})$ which sends the product F_i to the matrix $\begin{pmatrix} 1 & 0 \\ -6 & 1 \end{pmatrix}^{-1} \begin{pmatrix} 1 & 6 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 6 \\ 6 & 37 \end{pmatrix}$. By the Thurston's theorem, F_i is pA and its dilation is the largest real root of the characteristic polynomial $x^2 - 38x + 1$ of this matrix, that is $37.9737 \dots$ in all three cases. We ask two questions: Are F_1, F_2, F_3 conjugate in $\text{Mod}(S)$, and if they are not conjugate can our invariants distinguish them?

See [2] for a choice of *standard curves* a_0, d_0, c_0, d_1, c_1 on S . To obtain the needed input data for the computer software XTrain [2] we must express our maps as products of Dehn twists about these curves. For simplicity, we denote the Dehn twist T_{a_0} by the same symbol a_0 . Using A_i, C_i, D_i for the inverse Dehn twists of a_i, c_i, d_i , we find, after a small calculation, that:

$$\begin{aligned}
 F_1 &= c_0 d_0 d_1 A_0 C_1 c_1 d_1 c_0 d_0 A_0 D_0 C_0 D_1 C_1 c_1 a_0 D_1 D_0 C_0 \\
 &\quad c_1 b_0 (C_1 D_1)^6 c_1 d_1 c_0 d_0 a_0 D_0 C_0 D_1 C_1 (d_1 c_1)^6 B_0 C_1 \\
 F_2 &= c_1 (d_1 c_1)^6 a_0 c_1 d_1 c_0 d_0 A_0 D_0 C_0 D_1 C_1 A_0 (C_1 D_1)^6 C_1 c_1 d_1 c_0 d_0 a_0 D_0 C_0 D_1 C_1 \\
 F_3 &= A_0 D_0 C_0 d_0 a_0 B_1 D_1 c_0 d_1 b_1 A_0 D_0 B_1 D_1 C_0 d_1 b_1 d_0 a_0 B_1 D_1 C_0 d_1 b_1 A_0 D_0 C_0 d_0 a_0 (d_1 c_1)^6
 \end{aligned}$$

Focussing on F_1 and F_2 first, XTrain tells us that the associated transition matrices are:

$$T_1 = \begin{bmatrix} 15 & 7 & 14 & 23 & 16 \\ 10 & 6 & 10 & 16 & 11 \\ 4 & 2 & 5 & 7 & 5 \\ 2 & 1 & 2 & 3 & 1 \\ 10 & 5 & 10 & 16 & 12 \end{bmatrix} \quad \text{and} \quad T_2 = \begin{bmatrix} 6 & 10 & 5 & 6 & 10 \\ 5 & 11 & 5 & 7 & 10 \\ 5 & 10 & 6 & 6 & 10 \\ 6 & 12 & 6 & 7 & 12 \\ 5 & 10 & 5 & 5 & 11 \end{bmatrix}$$

In both cases $\chi(f_*) = \det(xI - T_i) = x^5 - 41x^4 + 118x^3 - 118x^2 + 41x - 1$, with largest real root $37.9737\dots$, as expected.

XTrain tells that the train tracks τ_1, τ_2 for our two examples are the ones that are illustrated in Figure 7. With the train tracks τ_1, τ_2 in hand we can see, immediately, that




FIGURE 7. Train tracks for the maps F_1, F_2 of Example 5.1.

F_1 and F_2 are inequivalent, because τ_1 is orientable and τ_2 is not. Also, by Lemma 2.11 the dimension of $W(G_i, f_i)$, which is the degree of the homology polynomial, is 4 (resp. 3) when $i = 1$ (resp. 2).

We compute the homology and symplectic polynomials of F_1 and F_2 explicitly.

For F_1 , a basis of $W(G_1, f_1)$, $\{\eta_b, \eta_c, \eta_d, \eta_e\}$, were computed in Example 2.15. Set $e_1 = b, e_2 = c, e_3 = d, e_4 = e, e_5 = a$ and follow the instructions in §2.3.4 for finding the

matrix A_1 representing $(f_1)_*|_{W(G_1, f_1)}$. We obtain:

$$A_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 6 & 10 & 16 & 11 & 10 \\ 2 & 5 & 7 & 5 & 4 \\ 1 & 2 & 3 & 1 & 2 \\ 5 & 10 & 16 & 12 & 10 \\ 7 & 14 & 23 & 16 & 15 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & -1 & -1 & 1 \end{bmatrix} = \begin{bmatrix} 16 & 0 & 6 & 21 \\ 6 & 1 & 3 & 9 \\ 3 & 0 & 1 & 3 \\ 15 & 0 & 6 & 22 \end{bmatrix}$$

Its characteristic polynomial is $1 - 40x + 78x^2 - 40x^3 + x^4 = (-1 + x)^2(1 - 38x + x^2)$, which is the homology polynomial. Corollary 4.1-(4) tells that the symplectic polynomial has degree 4, hence it coincides with the homology polynomial.

For F_2 , we apply §2.3.3 and set the non-orientable loop $\mathcal{L}_0 = a \cup v_0$ and subgraph $\mathcal{L} = \mathcal{L}_1 = a \cup b \cup v_0 \cup v_1$. Edges c, d, e are not in \mathcal{L} . We reorder the edges and call $e_1 = c, e_2 = d, e_3 = e, e_4 = a, e_5 = b$. With this order, basis vectors of $W(G_2, f_2)$ are $\eta_a = (1, 0, 0, 1, 2)'$; $\eta_c = (0, 1, 0, 0, 0)'$; $\eta_e = (0, 0, 1, 0, -1)'$; where “prime” means the transpose. Following the instructions in §2.3.4, we obtain the matrix A_2 representing $(f_2)_*|_{W(G_2, f_2)}$.

$$A_2 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} 6 & 6 & 10 & 5 & 10 \\ 6 & 7 & 12 & 6 & 12 \\ 5 & 5 & 11 & 5 & 10 \\ 5 & 6 & 10 & 6 & 10 \\ 5 & 7 & 10 & 5 & 11 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \\ 2 & 0 & -1 \end{bmatrix} = \begin{bmatrix} 31 & 6 & 0 \\ 36 & 7 & 0 \\ 30 & 5 & 1 \end{bmatrix}$$

The homology polynomial is $\det(xI - A_2) = -1 + 39x - 39x^2 + x^3 = (-1 + x)(1 - 38x + x^2)$, which means $\dim W(G_2, f_2) = 3$. By Corollary 3.6, the symplectic polynomial has degree 2, so it is $1 - 38x + x^2$.

We turn to F_3 . From XTrain, we learn that the homology polynomials for F_2 and F_3 are the same. However, observe that the curves u_1, v_1, u_2, v_2, v_3 are all non-separating on S , but u_3 is separating. From this it follows that there cannot be an element $F' \in \text{Mod}(S)$ that maps u_3 to either u_i or v_i , $i = 1, 2$. Thus $[F_3]$ is very likely not conjugate to either $[F_1]$ or $[F_2]$ in $\text{Mod}(S)$.

Example 5.2 (Figure 8). This example shows that the symplectic polynomial need not be irreducible over the rationals. The monodromy of the hyperbolic knot 8_9 [9] is represented by the following train track map:


$$\begin{array}{lll} a : (v_3, v_2) \mapsto e & b : (v_2, v_0) \mapsto g & c : (v_1, v_3) \mapsto b \\ d : (v_1, v_2) \mapsto bi & e : (v_0, v_1) \mapsto hed & f : (v_2, v_1) \mapsto d \\ g : (v_1, v_3) \mapsto fgh & h : (v_3, v_0) \mapsto ic & i : (v_0, v_1) \mapsto a \end{array}$$

We have that $\chi(f_*) = x^9 - 2x^8 + x^7 - 4x^5 + 4x^4 - x^2 + 2x - 1$. Fix a basis of $\text{im } \delta$, $\{v_0 - v_1, v_2 - v_0, v_3 - v_0\}$. With respect to this basis, $f|_{\text{im } \delta}$ is represented as:

$$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 0 & 0 \\ 1 & -1 & -1 \end{bmatrix}$$

whose characteristic polynomial is $x^3 + x^2 - x - 1$. Therefore $h(x)$ is:

$$(x^9 - 2x^8 + x^7 - 4x^5 + 4x^4 - x^2 + 2x - 1) / (x^3 + x^2 - x - 1) = x^6 - 3x^5 + 5x^4 - 7x^3 + 5x^2 - 3x + 1.$$

FIGURE 8. Example 5.2: The knot 8_9 . Orientable train track.

Proposition 3.5-(1) tells us that $p(x) = 1$, so that $s(x) = x^6 - 3x^5 + 5x^4 - 7x^3 + 5x^2 - 3x + 1$. Since $f : G \rightarrow G$ is orientation preserving, by Corollary 4.3-(2), we have $\Delta_{8_9} = x^6 - 3x^5 + 5x^4 - 7x^3 + 5x^2 - 3x + 1$. It further factors as $(x^3 - 2x^2 + x - 1)(x^3 - x^2 + 2x - 1)$. It is interesting that these factors are no longer palindromic, and one contains the dilatation λ as a root and the other $1/\lambda$.

A similar analysis based on the hyperbolic knot 8_{10} shows that $s(x)$ need not even be symplectically irreducible. It has a non-orientable train track and

$$s(x) = (x+1)^2(x^{10} - 3x^9 + 3x^8 - 4x^7 + 5x^6 - 5x^5 + 5x^4 - 4x^3 + 3x^2 - 3x + 1).$$

The train tracks in the remaining examples are all non-orientable.

Example 5.3 (Figure 9). This example was suggested to us by Robert Penner. Figure 9a shows five curves c_1, \dots, c_5 on the genus 3 surface. The example is the product of positive Dehn twists about c_1, c_2, c_3 and inverse twists about c_4, c_5 . The map f acts as follows:

$$\begin{array}{ll} a \mapsto ada & b \mapsto da\bar{c}\bar{d}\bar{a}b \\ c \mapsto c\bar{e}\bar{b}adc & d \mapsto da\bar{c}\bar{d}\bar{a}b e f c\bar{e}\bar{b}ad \\ e \mapsto e f c\bar{e}\bar{b}ad a d a d a \bar{c}\bar{d}\bar{a} b e & f \mapsto e f c\bar{e}\bar{b}ad a d a d a \bar{c}\bar{d}\bar{a} b e f \end{array}$$

There is exactly one vertex (which is partial) and the train track is non-orientable. The only vertex is fixed by the map and $\chi(f_*|_{\text{im } \delta}) = x - 1$. We have $\dim W(G, f) = 5$, hence the skew-symmetric product is degenerate and $\dim Z = 1$, which means $p(x) = x - 1$. The characteristic polynomial factors as $\chi(f_*) = (x^4 - 11x^3 + 22x^2 - 11x + 1)(x - 1)(x - 1)$. The symplectic polynomial $s(x) = x^4 - 11x^3 + 22x^2 - 11x + 1$ is irreducible, hence in

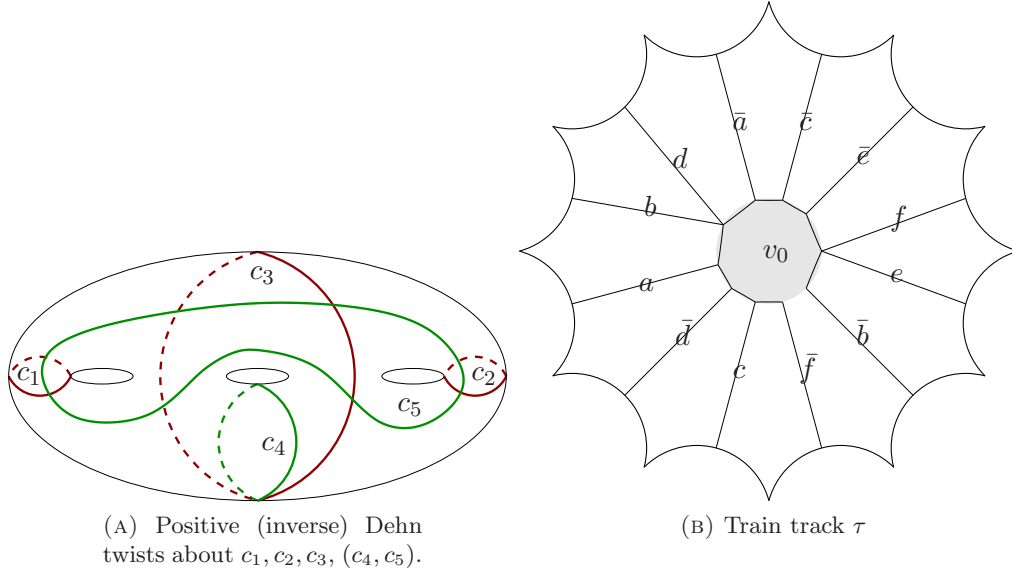



FIGURE 9. Example 5.3: Penner's pA map, non-orientable train track.

this example it is necessarily the minimum polynomial of its dilatation. (Note that this was not the case in Example 5.2 above.)

If one orients the real and infinitesimal edges of the train track τ in Figure 9b locally so that all orientations are consistent around the single partial vertex v_0 , one sees that the loops a, d and f do not have globally consistent orientations, whereas the loops b, c and e do. Since G has no odd vertices, the orientation cover is just an ordinary double cover as illustrated in Figure 10. Each edge, say $a \subset G$, lifts to two copies $a, a' \subset \tilde{G}$ and the vertex $v_0 \in G$ lifts to $v_0, v_1 \in \tilde{G}$. We choose an orientation for $\tilde{\tau}$. As claimed in Proposition 2.21, twin edges have opposite orientations. Proposition 2.22 implies that there are two covering maps; orientation preserving and orientation reversing. The following is the orientation preserving train track map $\tilde{f}_{\text{op}} : \tilde{G} \rightarrow \tilde{G}$:

$$\begin{array}{ll}
 a : (v_1, v_0) \mapsto ada & a' : (v_1, v_0) \mapsto a'd'a' \\
 b : (v_0, v_0) \mapsto dac'd'a'b & b' : (v_1, v_1) \mapsto b'adca'd' \\
 c : (v_1, v_1) \mapsto ce'b'adc & c' : (v_0, v_0) \mapsto c'd'a'bec' \\
 d : (v_0, v_1) \mapsto dac'd'a'befce'b'ad & d' : (v_0, v_1) \mapsto d'a'bec'f'e'b'adca'd' \\
 e : (v_0, v_0) \mapsto efce'b'adadac'd'a'be & e' : (v_1, v_1) \mapsto e'b'adca'd'a'd'a'bec'f'e' \\
 f : (v_0, v_1) \mapsto efce'b'adadac'd'a'bef & f' : (v_0, v_1) \mapsto f'e'b'adca'd'a'd'a'bec'f'e'
 \end{array}$$

FIGURE 10. The orientation cover $\tilde{\tau}$ of τ in Figure 9.

We order the edges $a, b, \dots, f, a', b', \dots, f'$. Then the transition matrix of \tilde{f}_{op} has the form $\begin{bmatrix} A & B \\ B & A \end{bmatrix}$ where

$$A = \begin{bmatrix} 2 & 1 & 1 & 2 & 3 & 3 \\ 0 & 1 & 0 & 1 & 1 & 1 \\ 0 & 0 & 2 & 1 & 1 & 1 \\ 1 & 1 & 1 & 2 & 2 & 2 \\ 0 & 0 & 0 & 1 & 2 & 2 \\ 0 & 0 & 0 & 1 & 1 & 2 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 0 & 1 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Note that $A + B$ is the transition matrix for $f : G \rightarrow G$. Its characteristic polynomial is

$$(-1 + x)^4(1 - 4x + x^2)(1 - 3x + x^2)(1 - 11x + 22x^2 - 11x^3 + x^4),$$

and the symplectic polynomial is

$$\chi(\tilde{f}_{\text{op}*}|_{W(\tilde{G}, \tilde{f})/\tilde{Z}}) = (-1 + x)^2(1 - 4x + x^2)(1 - 3x + x^2)(1 - 11x + 22x^2 - 11x^3 + x^4).$$

The following is the orientation reversing train track map $\tilde{f}_{\text{or}} : \tilde{G} \rightarrow \tilde{G}$.

$$\begin{array}{ll} a : (v_1, v_0) & \mapsto \bar{a}'\bar{d}'\bar{a}' \\ b : (v_0, v_0) & \mapsto \bar{d}'\bar{a}'\bar{c}\bar{d}\bar{a}\bar{b}' \\ c : (v_1, v_1) & \mapsto \bar{c}'\bar{e}\bar{b}\bar{a}'\bar{d}'\bar{c}' \\ d : (v_0, v_1) & \mapsto \bar{d}'\bar{a}'\bar{c}\bar{d}\bar{a}\bar{b}'\bar{e}'\bar{f}'\bar{c}'\bar{e}\bar{b}\bar{a}'\bar{d}' \\ e : (v_0, v_0) & \mapsto \bar{e}'\bar{f}'\bar{c}'\bar{e}\bar{b}\bar{a}'\bar{d}'\bar{a}'\bar{d}'\bar{a}'\bar{c}\bar{d}\bar{a}\bar{b}'\bar{e}' \\ f : (v_0, v_1) & \mapsto \bar{e}'\bar{f}'\bar{c}'\bar{e}\bar{b}\bar{a}'\bar{d}'\bar{a}'\bar{d}'\bar{a}'\bar{c}\bar{d}\bar{a}\bar{b}'\bar{e}'\bar{f}' \end{array} \quad \begin{array}{ll} a' : (v_1, v_0) & \mapsto \bar{a}\bar{d}\bar{a} \\ b' : (v_1, v_1) & \mapsto \bar{b}\bar{a}'\bar{d}'\bar{c}'\bar{a}\bar{d} \\ c' : (v_0, v_0) & \mapsto \bar{c}\bar{d}\bar{a}\bar{b}'\bar{e}'\bar{c} \\ d' : (v_0, v_1) & \mapsto \bar{d}\bar{a}\bar{b}'\bar{e}'\bar{c}'\bar{f}\bar{e}\bar{b}\bar{a}'\bar{d}'\bar{c}'\bar{a}\bar{d} \\ e' : (v_1, v_1) & \mapsto \bar{e}\bar{b}\bar{a}'\bar{d}'\bar{c}'\bar{a}\bar{d}\bar{a}\bar{d}\bar{a}\bar{b}'\bar{e}'\bar{c}'\bar{f}\bar{e} \\ f' : (v_0, v_1) & \mapsto \bar{f}\bar{e}\bar{b}\bar{a}'\bar{d}'\bar{c}'\bar{a}\bar{d}\bar{a}\bar{d}\bar{a}\bar{b}'\bar{e}'\bar{c}'\bar{f}\bar{e} \end{array}$$

We observe that the transition matrix for \tilde{f}_{or} is $\begin{bmatrix} B & A \\ A & B \end{bmatrix}$. Its characteristic polynomial is

$$(-1 + x)^2(1 + x)^2(1 + 3x + x^2)(1 + 4x + x^2)(1 - 11x + 22x^2 - 11x^3 + x^4),$$

and the symplectic polynomial is

$$\chi(\tilde{f}_{\text{or}*}|_{W(\tilde{G}, \tilde{f})/\tilde{Z}}) = (1+x)^2(1+3x+x^2)(1+4x+x^2)(1-11x+22x^2-11x^3+x^4).$$

The dilatation cannot distinguish the pA maps $F : S \rightarrow S$, $\tilde{F}_{\text{op}} : \tilde{S} \rightarrow \tilde{S}$ and $\tilde{F}_{\text{or}} : \tilde{S} \rightarrow \tilde{S}$, but our symplectic polynomial can distinguish the three.

It seems to be an open question to describe all the ways to construct all pA maps having a fixed dilatation.

Example 5.4 (Figure 11). The following map shows that even, odd, and partial vertices can coexist in the same (non-orientable) train track:




FIGURE 11. Example 5.4: A train track with even, odd, and partial vertices.

$$\begin{array}{lll} a \mapsto a\bar{g} & b \mapsto ja & c \mapsto fc \\ d \mapsto ca\bar{g}ibja\bar{g}e & e \mapsto \bar{g}eha\bar{g}ibja\bar{g}ef & f \mapsto \bar{d} \\ g \mapsto \bar{i} & h \mapsto c & i \mapsto b \\ j \mapsto \bar{g}eha\bar{g}ibj \end{array}$$

The characteristic polynomial factors as

$$s(x)p(x)\chi(f_*|_{\text{im}\delta}) = (x^6 - 3x^5 + x^4 - 5x^3 + x^2 - 3x + 1)(x-1)(x^3 - x^2 - x + 1).$$

The factor $\chi(f_*|_{\text{im}\delta}) = (x^3 - x^2 - x + 1)$ is the characteristic polynomial of the matrix

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix},$$

which describes how non-odd vertices v_1, v_2, v_3 are permuted by $[F]$.

ACKNOWLEDGMENT

The authors would like to thank Jeffrey Carlson for his help with the proof of Corollary 2.25, which was stated as a conjecture in an earlier draft. They thank Dan Margalit for suggesting a way to find examples of non-conjugate maps with the same dilatation. They would also like to thank Mladen Bestvina, Nathan Dunfield, Jordan Ellenberg, Ji-young Ham, Eriko Hironaka, Eiko Kin, Chris Leininger, Robert Penner, and Jean-Luc Thiffeault for their generosity in sharing their expertise and their patience in responding to questions.

REFERENCES

- [1] M. Bestvina and M. Handel, *Train tracks for surface homeomorphisms*, Topology **34**, No. 1 (1995), 109-140.
- [2] P. Brinkmann, *An implementation of the Bestvina-Handel algorithm for surface homeomorphisms*, Experimental Mathematics **9**, No. 2 (2000), 235-240.
- [3] J. W. Eaton, GNU Octave Manual, Network Theory Limited, 2002.
- [4] B. Farb and D. Margalit, *A PRIMER ON MAPPING CLASS GROUPS*
- [5] A. Fathi, F. Laudenbach and V. Poénaru, *TRAVAUX DE THURSTON SUR LES SURFACES*, volume 66 of Astérisque. Société Mathématique de France, Paris, 1979. Séminaire Orsay, With an English summary.
- [6] F. R. Gantmacher, *THE THEORY OF MATRICES: Vol II*, Chelsea Publishing Co., New York, NY, 1959. Translated by K. A. Hirsch.
- [7] C. McMullen, *Polynomial invariants for fibered 3-manifolds and Teichmüller geodesic for foliations*, Annales Scientifique de l'École Normal Supérieure, Quatrième Série **33** (2000), 519-560.
- [8] R. C. Penner with J. L. Harer, *COMBINATORICS OF TRAIN TRACKS*, Princeton Univ. Press (1992), Princeton, N.J.
- [9] D. Rolfsen, *KNOTS AND LINKS*, Publish or Perish Inc., Houston Texas, 1990. Corrected reprint of 1976 original.
- [10] E. Rykken, *Expanding factors for pseudo-Anosov homeomorphisms*, Michigan Math. Journal, **46**, No. 2 (1999), 281-296.
- [11] W. P. Thurston, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. **19**, No. 2 (1988), 417-431.

DEPARTMENT OF MATHEMATICS, COLUMBIA UNIVERSITY

E-mail address: jbm@math.columbia.edu

GOOGLE INC.

E-mail address: peter.brinkmann@gmail.com

DEPARTMENT OF MATHEMATICS, THE UNIVERSITY OF IOWA

E-mail address: kawamuro@iowa.uiowa.edu