ORIENTABLE MAPS AND POLYNOMIAL INVARIANTS OF FREE-BY-CYCLIC GROUPS

SPENCER DOWDALL, RADHIKA GUPTA, AND SAMUEL J. TAYLOR

ABSTRACT. We relate the McMullen polynomial of a free-by-cyclic group to its Alexander polynomial. To do so, we introduce the notion of an *orientable* fully irreducible outer automorphism φ and use it to characterize when the homological stretch factor of φ is equal to its geometric stretch factor.

CONTENTS

1. INTRODUCTION

This paper introduces the property of *orientability* for fully irreducible free group automorphisms and characterizes how it is reflected in both the stretch factors of the automorphism and in the Alexander and McMullen polynomials of the associated free-by-cyclic group.

To motivate our work, let us recall the theory from the mapping class group setting. A pseudo-Anosov homeomorphism φ of an orientable hyperbolic surface comes equipped with two notions of stretch factor: The *geometric* stretch factor or dilatation λ_φ is the value by which φ stretches the leaves of its expanding foliation; λ_φ is also characterized as the exponential growth rate of the geodesic length of any curve under iteration by φ, and its logarithm $\log \lambda_\varphi$ is both the topological entropy of φ and the translation length for its action on Teichmüller space. The *homological* stretch factor ρ_φ is instead the spectral radius for the action of φ on the first homology of the surface. It is well-known that these two numbers are equal if and only if the invariant foliations of φ are transversely orientable. Indeed, that transverse orientability implies equality of the stretch factors was observed by Thurston [Thu88], and the reverse implication was proven by Band and Boyland [BB07, Lemma 4.3].

In fact, a much stronger symmetry persists when φ preserves these transverse orientations. Associated to the mapping torus of a pseudo-Anosov homeomorphism there are two important polynomial invariants that arise: the classical Alexander polynomial [Ale28], which captures homological information about monodromies, and the Teichmüller polynomial [McM00] of the fibered cone, which encodes the associated geometric stretch factors. In the orientable case, McMullen proved these invariants are closely related in that the Alexander polynomial

divides the Teichmüller polynomial [McM00, Theorem 7.1]. More recently, Parlak [Par21] extended McMullen's work to explicitly determine the quotient polynomial, giving a precise relation between the two polynomials and the associated stretch factors.

1.1. Orientable automorphisms.

In this article, we extend the above results to the free group setting, replacing the pseudo-Anosov homeomorphism with a fully irreducible automorphism of a finite rank free group $\mathbb{F}$ and replacing the mapping torus M_φ by the associated free-by-cyclic group. In fact, we define what it means for a fully irreducible endomorphism to be *orientable* and prove this occurs exactly when the geometric and homological stretch factors agree.

The condition is best understood in the context of graph maps. For a graph map $f\colon G \to G$, we define the geometric stretch factor λ_f to be the spectral radius of the transition matrix of f, and the homological stretch factor ρ_f to be the spectral radius of the induced map f_* on $H_1(G;\mathbb{R})$. An *orientation* on G is a choice of positive orientation for each edge of G. We say f is *positively orientable* (resp. *negatively orientable*) if there exists an orientation on G such that every positive edge maps to a positive (resp. negative) edge path. We often abbreviate this terminology to pos/neg-orientable or ± 1–orientable. Combining these notions, the map is called *orientable* if it respects some orientation on G, in that it is either positively orientable or negatively orientable. See §2 for additional details.

Theorem A. *Suppose $f\colon G \to G$ is a graph map whose transition matrix has a positive power. Then $\lambda_f = \rho_f$ if and only if f is orientable. Moreover, λ_f (resp. $-\lambda_f$) is an eigenvalue of f_* if and only if f is pos-orientable (resp. neg-orientable).*

By combining this with well-known properties of fully irreducible endomorphisms and their expanding laminations (see §3), we obtain the following corollary, which is a simplified statement of Theorem 3.6:

Theorem B. *For a fully irreducible free group endomorphism φ, the following are equivalent:*

- *the homological and geometric stretch factors of φ are equal;*
- *some (every) irreducible train track representative of φ is orientable;*
- *the expanding lamination $\mathcal{L}_\varphi^+$ of φ is orientable.*

1.2. Splittings and the BNS-invariant.

To frame our results relating Alexander and McMullen polynomials, we start with a corollary characterizing when geometric and homological stretch factors agree for monodromies of free-by-cyclic groups.

Any group endomorphism $\phi\colon B \to B$ determines a *generalized HNN extension*

$$B*_\phi := \langle B, t \mid t^{-1}bt = \phi(b) \text{ for all } b \in B \rangle$$

that comes equipped with a projection $B*_\phi \to \mathbb{Z}$ sending B to 0 and t to 1; we call this element of $\mathrm{Hom}(B*_\phi, \mathbb{Z}) = H^1(B*_\phi; \mathbb{Z})$ the *dual cohomology class*.

In the case that Γ is finitely generated, there is an open, $\mathbb{R}_+$–invariant subset $\mathrm{BNS}(\Gamma)$ of $H^1(\Gamma;\mathbb{R})\backslash\{0\}$ that governs finite generation of kernels of such maps $\Gamma \to \mathbb{Z}$. This *Bieri–Neumann–Strebel invariant* [BNS87] is defined to contain precisely those rational classes $u\colon \Gamma \to \mathbb{Q}$ for which $\ker(u)$ is finitely generated over the u–positive monoid, meaning there is an element $t \in \Gamma$ with $u(t) > 0$ so that $\ker(u)$ is generated by the positive conjugates $\{t^n X t^{-n} \mid n \geqslant 0\}$ of some finite subset $X \subset \ker(u)$. Unpacking this, one finds that $u\colon \Gamma \to \mathbb{Z}$ lies in $\mathrm{BNS}(\Gamma)$ if and only if it is the dual class of an HNN splitting $\Gamma \cong B*_\phi$ over a *finitely generated* base group B [BNS87, Proposition 4.3]. We will say that any such splitting $\Gamma \cong B*_\phi$ (with B finitely generated) is *dual* to u and call ϕ a *monodromy* associated to u.

Now suppose that Γ is free-by-cyclic. For any splitting $\Gamma \cong B*_\phi$ dual to a class $u \in \text{BNS}(\Gamma)$, the group B is free. Although the monodromy ϕ is not uniquely determined by u, its geometric and homological stretch factors are determined by u and these are denoted by $\lambda(u)$ and $\rho(u)$, respectively. See §6 for details.

Theorem C. *Consider the free-by-cyclic group $\Gamma = \mathbb{F}*_\varphi$, where φ is a fully irreducible automorphism, and let $\mathcal{C}$ be the component of $\text{BNS}(\Gamma)$ containing the dual class of $\mathbb{F}*_\varphi$. Then the set of all primitive integral $u \in \mathcal{C}$ such that $\lambda(u) = \rho(u)$ is either (1) the entire cone, (2) exactly one residue class mod 2, or (3) empty.*

When φ is additionally atoroidal, it is known that *every* monodromy associated to any primitive integral class $u \in \mathcal{C}$ is also fully irreducible [DKL17a, Mut21]; hence in light of Theorem A, the classes with $\lambda(u) = \rho(u)$ described in Theorem C are precisely those whose monodromy endomorphisms are orientable; see Remark 6.3.

1.3. **Relating the Alexander and McMullen polynomials.** The trichotomy in Theorem C is obtained by explicitly relating two polynomial invariants attached to $\Gamma = \mathbb{F}*_\varphi$. The first of these, the Alexander polynomial, was originally defined by Alexander in 1928 [Ale28] as a knot invariant, and its definition was later extended by McMullen to all finitely generated groups [McM02]. It is an invariant of the group Γ and is formally defined, up to a unit, as an element Δ_Γ of the ring $\mathbb{Z}[H]$, where $H = H_1(\Gamma; \mathbb{Z})/\text{torsion}$; see §4 or the next subsection §1.4 below for details.

The second of these, the McMullen polynomial, was defined by Dowdall–Kapovich–Leininger in [DKL17b] as an analog of McMullen's Teichmüller polynomial from [McM00]. It is defined in terms of an irreducible train track representative $f : G \to G$ of a fully irreducible automorphism φ of $\pi_1(G) \cong \mathbb{F}$ and an associated *folded mapping torus $X^\dagger$* that comes equipped with a natural semiflow $\psi^\dagger$ and is homotopy equivalent to the usual mapping torus of f. The *McMullen polynomial* $\mathfrak{m} \in \mathbb{Z}[H]$ is an algebraic invariant of this dynamical system $(X^\dagger, \psi^\dagger)$. The key features of $\mathfrak{m}$ are that it picks out the component $\mathcal{C}$ of the BNS-invariant containing the dual class of the splitting $\Gamma = \mathbb{F}*_\varphi$ and calculates the geometric stretch factors $\lambda(u)$ of the classes $u \in \mathcal{C}$, as in Theorem C. We refer the reader to [DKL17b] and §5 for more details.

To connect these, in §5.4 we introduce an additional *vertex polynomial*

$$\mathfrak{p} = \prod_i (1 - z_i) \in \mathbb{Z}[H],$$

of the dynamical system $(X^\dagger, \psi^\dagger)$, where $\{z_i\}$ are the elements in $H = H_1(X^\dagger; \mathbb{Z})/\text{torsion}$ represented by the finitely many closed orbits of $\psi^\dagger$ through the vertices of $X^\dagger$. Our next theorem shows that these three polynomials are related in a precise way that moreover reflects the positive, negative, or non-orientability of f:

Theorem D. *Let $f : G \to G$ be an irreducible train track map representing a fully irreducible automorphism φ of $\pi_1(G) \cong \mathbb{F}$, with associated free-by-cyclic group Γ and folded mapping torus $(X^\dagger, \psi^\dagger)$. If $\text{rank}(H_1(X, \mathbb{R})) \geqslant 2$, then up to a unit:*

    (1) *if f is pos-orientable, then $\mathfrak{m} = \Delta_\Gamma \cdot \mathfrak{p}$, and*

    (2) *if f is neg-orientable, then there is an orientation involution $\iota : \mathbb{Z}[H] \to \mathbb{Z}[H]$ for which $\iota(\mathfrak{m}) = \Delta_\Gamma \cdot \mathfrak{p}$.*

    (3) *in general, $\mathfrak{m} = \Delta_\Gamma \cdot \mathfrak{p}$ mod 2.*

If instead $\text{rank}(H_1(\Gamma; \mathbb{R})) = 1$, then each equations holds after multiplying $\mathfrak{m}$ by $(z-1)$, where z generates $H_1(\Gamma, \mathbb{Z})$.

Moreover, Theorem D applies not only to the original monodromy, but also to any monodromy dual to a primitive integral class in the corresponding component $\mathcal{C}$ of the BNS-invariant, so that its 3 conclusions pair with the trichotomy of Theorem C. Indeed, the involution ι in the neg-orientable case comes from an *orientation class* that depends only on $\mathcal{C}$ (Proposition 6.7). This more general perspective is taken in Theorems 6.4, 6.8 and 6.9, which, respectively, give more detailed treatments of the conclusions (1), (2) and (3) and, collectively, prove Theorem D.

1.4. **Computing the Alexander polynomial.** Although the Alexander polynomial can always be efficiently computed from a group presentation via Fox Calculus (see [But07] or [Hir11] for relevant examples), our proof of Theorem D requires a computation of Δ_Γ that is compatible with the definition of $\mathfrak{m}$. Our final result provides this by giving an explicit 'determinant formula' for the Alexander polynomial that is inspired by McMullen's formula for the Teichmüller polynomial.

To explain, recall that any graph map $f\colon G \to G$ determines a *mapping torus*

$$X = X_f := \frac{G \times [0,1]}{(x,1) \sim (f(x),0)},$$

whose fundamental group Γ splits as a generalized HNN extension, which we write as $\Gamma \cong \pi_1(G) *_f$. Here, the associated dual class of the splitting is obtained by pulling back the fundamental class of the circle under the associated fibration $X \to \mathbb{S}^1$.

Our determinant formula computes the Alexander polynomial Δ_Γ of $\Gamma = \pi_1(X)$ directly in terms of the graph map f. Let $\widetilde{X} \to X$ be the universal free abelian cover of X with deck group H. Fix a connected component $\widetilde{G}_0$ of the full preimage of G in $\widetilde{X}$ and fix a lift $\widetilde{f}\colon \widetilde{G}_0 \to \widetilde{G}_0$ of f. The deck group of $\widetilde{G}_0 \to G$ is denoted by K and is naturally identified with the image of the inclusion induced homomorphism $\pi_1(G) \to \pi_1(X) \to H$. Moreover, the choice of lift $\widetilde{f}\colon \widetilde{G}_0 \to \widetilde{G}_0$ determines (see §4.4):

- a splitting $H = K \oplus \langle z \rangle$, where z translates against the lifted semiflow,
- $\mathbb{Z}[K]$–valued matrices $\widetilde{M}$ and $\widetilde{P}$ representing the action of $\widetilde{f}$ on the free $\mathbb{Z}[K]$–modules of 1–chains and 0–chains of $\widetilde{G}_0$, respectively.

Theorem E (Determinant formula for Δ_Γ). *Let $f\colon G \to G$ be a graph map whose mapping torus X has fundamental group Γ. Then up to a unit in $\mathbb{Z}[H]$,*

$$\Delta_\Gamma = \frac{\det(zI - \widetilde{M})}{\det(zI - \widetilde{P})} \cdot r,$$

where $r = 1 \in \mathbb{Z}[H]$ if $\mathrm{rank}(H_1(X,\mathbb{R})) \geqslant 2$ and $r = (z-1)$ if $\mathrm{rank}(H_1(X,\mathbb{R})) = 1$.

2. Orientable graph maps and stretch factors

Let G be a finite graph with vertex set $V = V(G)$ and edge set $E = E(G)$. A *graph map* $f\colon G \to G$ is a continuous map sending vertices to vertices and edges to nondegenerate edge paths. For careful definitions of standard notions like 'graph' or 'edge path,' we refer the reader to [DKL17b, §2].

An *orientation* on G is a choice of a positive orientation for each edge of G. With such a choice, we say an edge path γ is *positive* if it crosses each edge with its positive orientation, and is *negative* if it crosses each edge with its negative orientation. In general, if $\gamma\colon [0,1] \to G$ is an edge path, then $\overline{\gamma}$ will denote the same path with the reverse orientation.

Definition 2.1 (Orientability). A graph map is *positively orientable* (resp. negatively orientable) if there exists an orientation on G such that every positive edge is mapped to a positive (resp. negative) edge path. We abbreviate these as pos/neg or ± 1–orientable. The map is *orientable* if it is either pos-orientable or neg-orientable.

Note that if f is neg-orientable, then f^2 is pos-orientable. Since a consistently oriented edge path in G cannot backtrack, we observe that an orientable map is necessarily a train track map (see §3.3). See §7 for examples.

Choosing an orientation on each edge gives a basis for the vector space $C_1(G;\mathbb{R})$ of simplicial 1–chains. Under this identification $C_1(G;\mathbb{R}) \cong \mathbb{R}^E$, the induced map on simplicial 1–chains is expressed as a matrix $M\colon \mathbb{R}^E \to \mathbb{R}^E$ with integer coefficients. The real first homology of G is the subspace $H_1(G;\mathbb{R}) \leqslant \mathbb{R}^E$ of 1–cycles, and we denote the induced map by $f_*\colon H_1(G;\mathbb{R}) \to H_1(G;\mathbb{R})$. The *homological stretch factor* of f is defined to be the spectral radius ρ_f of f_*. Note that this also equals the spectral radius of M; indeed the eigenvalues of M that are not eigenvalues of f_* are eigenvalues of the action of f on $C_0(G;\mathbb{R})$, which are roots of unity.

For an arbitrary graph map $f\colon G \to G$, we will denote the transition matrix of f by A, whose (i,j) entry is the number of times the edge path $f(e_j)$ crosses the edge e_i with either orientation. The *(geometric) stretch factor* λ_f of f is defined to be the spectral radius of A. Equivalently, since A is non-negative, λ_f is the eigenvalue of A of largest modulus. Since the absolute values of the entries of M are bounded by the entries of A, we have $\rho_f \leqslant \lambda_f$ by Gelfand's formula for the spectral radius: for any square matrix A, the spectral radius of A is equal to $\lim_{r\to\infty} \|A^r\|^{1/r}$. The following fact relating M and A follows from the definitions.

Fact 2.2. *If $f\colon G \to G$ is ± 1–orientable, then $M = \pm A$ after a choice of orientation.*

Recall that a non-negative square matrix A is *irreducible* if for each pair of indices (i,j) there exists $k \geqslant 1$ such that the (i,j) entry of A^k is nonzero. The matrix is moreover *primitive* if there exists $k \geqslant 1$ such that A^k is a positive matrix. Note that A is primitive if and only if A^k is irreducible for all $k \geqslant 1$. We extend this terminology by saying a graph map $f\colon G \to G$ is *irreducible/primitive* if its transition matrix is irreducible/primitive. Relatedly, f is *expanding* if for each edge e of G the combinatorial lengths of the edge paths $f^k(e)$ tend to infinity as $k \to \infty$. By the Perron-Frobenius theorem (see [Gan59]), if f is a primitive graph map, then $\lambda_f > 1$.

The main goal of this section is to prove the following theorem from the introduction:

Theorem A. *Suppose that $f\colon G \to G$ is a primitive graph map. Then $\lambda_f = \rho_f$ if and only if f is orientable. Moreover, λ_f (resp. $-\lambda_f$) is an eigenvalue of f_* if and only if f is pos-orientable (resp. neg-orientable).*

2.1. **The oriented edge-double.** Motivated by the use of double branched covers in [BB07, Lemma 4.3], we define the *oriented edge double* $\widehat{G}$ of G as follows: Let $p\colon \widehat{G} \to G$ be a bijection on vertices and 2-to-1 over edges such that for each oriented edge e of G, there are oriented edges e_+ and e_- of $\widehat{G}$ labeled so that $p(e_+) = e$ and $p(e_-) = \overline{e}$. There is a natural continuous involution $\sigma\colon \widehat{G} \to \widehat{G}$ defined by $e_+ \mapsto \overline{e}_-$ and $e_- \mapsto \overline{e}_+$. Notice that by construction σ is neg-orientable and satisfies $p = p \circ \sigma$.

The key feature of $\widehat{G}$ is that *every* edge path in G has a unique *positive lift* to $\widehat{G}$. That is, for any edge path $\gamma\colon [0,1] \to G$ there is a unique positive edge path $\widehat{\gamma}\colon [0,1] \to \widehat{G}$ such that $\gamma = p \circ \widehat{\gamma}$. Similarly γ has a unique *negative lift*, which is simply the reverse of the unique positive lift of $\overline{\gamma}$ and is given by $\sigma \circ \widehat{\gamma}$. In particular, e_+ and e_- are, respectively, the positive lifts of e and $\overline{e}$.

Since σ is an involution, the eigenvalues of $\sigma_*\colon H_1(\widehat{G};\mathbb{R}) \to H_1(\widehat{G};\mathbb{R})$ are ± 1 and we write E_1, E_{-1} for the corresponding eigenspaces.

Lemma 2.3. *We have $H_1(\widehat{G};\mathbb{R}) = E_1 \oplus E_{-1}$. Furthermore, $p_*\colon H_1(\widehat{G};\mathbb{R}) \to H_1(G;\mathbb{R})$ restricts to an isomorphism $E_1 \xrightarrow{\cong} H_1(G)$ and has kernel $\ker(p_*) = E_{-1} \cong C_1(G;\mathbb{R})$.*

Proof. For any class $c \in H_1(\widehat{G})$, we clearly have $c + \sigma_*(c) \in E_1$ and $c - \sigma_*(c) \in E_{-1}$. Since

$$c = \frac{c + \sigma_*(c)}{2} + \frac{c - \sigma_*(c)}{2},$$

the direct sum splitting $H_1(\widehat{G}) = E_1 \oplus E_{-1}$ follows. The containment $E_{-1} \leqslant \ker(p_*)$ follows immediately from $p = p \circ \sigma$, since if $c \in E_{-1}$ then $p_*(c) = p_*(\sigma_*(c)) = p_*(-c) = -p_*(c)$.

There are natural injective maps $s_\pm\colon C_1(G;\mathbb{R}) \to C_1(\widehat{G};\mathbb{R})$ on 1–chains defined on each basis edge $e \in C_1(G;\mathbb{R})$ by $s_\pm(e) = \frac{1}{2}(e_+ \mp e_-) = \frac{1}{2}(e_+ \pm \sigma_*(e_+))$. The image of s_+ is contained in the fixed set of σ_*, that is $\sigma_* \circ s_+ = s_+$. Since s sends cycles to cycles, it therefore restricts to a map $s_+\colon H_1(G;\mathbb{R}) \to E_1 \leqslant H_1(\widehat{G};\mathbb{R})$. Further, we see that for any $c \in C_1(G;\mathbb{R})$ we have $p_*(s_+(c)) = c$. Similarly the image of s_- is clearly contained in $E_{-1} \leqslant H_1(\widehat{G})$. Now any $\widehat{c}_\pm \in E_{\pm 1}$ necessarily has the form $\widehat{c}_\pm = \sum_{e\in G} a_e(e_+ \mp e_-) = \sum_{e \in G} 2a_e s_\pm(e)$. In particular s_- is surjective onto E_{-1}, giving an isomorphism $C_1(G;\mathbb{R}) \to E_{-1}$. We also see that $p_*(\widehat{c}_+) = \sum_{e\in G} 2a_e e$ and hence $s_+(p_*(\widehat{c}_+)) = \sum_{e\in G} 2a_e s_+(e) = \widehat{c}_+$. It follows that p_* restricts to an isomorphism on $E_1 \to H_1(G;\mathbb{R})$ and that additionally $\ker(p_*) \leqslant E_{-1}$. $\quad\square$

Any graph map $f\colon G \to G$ has a unique pos-orientable lift $\widehat{f}\colon \widehat{G} \to \widehat{G}$, preserving the specified orientation of $\widehat{G}$, satisfying $p \circ \widehat{f} = f \circ p$. This sends e_+ to the unique positive lift of $f(e)$, and sends e_- to the unique positive lift of $f(\overline{e})$. More generally, for any positive edge path γ in $\widehat{G}$ we have

$$\widehat{f} \circ \gamma = \widehat{f \circ p \circ \gamma}.$$

We denote the transition matrix of $\widehat{f}$ by $\widehat{A}$ and note that $\widehat{A}$ is also the induced map on simplicial 1-chains (see Fact 2.2). Notice that for any positive edge e of $\widehat{G}$, the paths $\sigma(\widehat{f}(e))$ and $\widehat{f}(\sigma(e))$ are both negative lifts of $f(p(e))$. As such lifts are unique, we conclude that $\widehat{f}$ and σ commute:

$$\widehat{f} \circ \sigma = \sigma \circ \widehat{f}.$$

Consequently, $\widehat{f}_*$ necessarily preserves the splitting $H_1(\widehat{G}) = E_1 \oplus E_{-1}$ from Lemma 2.3; in fact $\widehat{f}_* = f_* \oplus A$ with the transition matrix A describing the action on $C_1(G;\mathbb{R}) \cong E_{-1}$.

Lemma 2.4. *Let $f\colon G \to G$ be a primitive graph map and let $\widehat{f}\colon \widehat{G} \to \widehat{G}$ be its orientation lift. Then $\lambda_{\widehat{f}} = \lambda_f$ and $\widehat{f}$ is primitive if and only if f is not orientable.*

Proof. As noted previously, for any square matrix B with spectral radius μ, we have that $\mu = \lim_{r\to\infty} \|B^r\|^{1/r}$ where $\|\cdot\|$ is any matrix norm. Let us take $\|\cdot\|$ to be the maximum absolute column sum (i.e. the ℓ^1 operator norm). Now observe that the corresponding column sums of (all powers of) $\widehat{A}$ and A are equal. Indeed, this follows from the definitions of $\widehat{A}, A$

and the fact that for each edge e of G, $p \circ \widehat{f}^n(e_+) = f^n(e)$ and $p \circ \widehat{f}^n(e_-) = f^n(\overline{e})$. Therefore $\|\widehat{A}^r\| = \|A^r\|$ for all $r \geqslant 1$ and we conclude $\lambda_{\widehat{f}} = \lambda_f$.

Let $D(A)$ and $D(\widehat{A})$ be the directed graphs such that their adjacency matrices are A and $\widehat{A}$, respectively. The vertices of $D(A)$ and $D(\widehat{A})$ are labeled by the edges of G and $\widehat{G}$, respectively. There is a natural simplicial map $D(\widehat{A}) \to D(A)$ that is 2-to-1 over vertices (induced by $\widehat{G} \to G$) and also 2-to-1 over edges. It is an isomorphism on the link of each vertex, so it is a (possibly disconnected) 2-fold covering. It follows that f is orientable if and only if $D(\widehat{A})$ is disconnected, and each component of $D(\widehat{A})$ gives an f-invariant orientation of G.

Since A is irreducible, $D(A)$ is strongly connected, i.e. for any two vertices, there is a directed path from one to the other. We now understand the connectivity of $D(\widehat{A})$. By path lifting, any directed path from e to e' in $D(A)$, lifts to a directed path in $D(\widehat{A})$ from e_+ to either e'_+ or e'_- (and similarly for e_-). Moreover, if there is a directed path from e_+ to $e'_{\pm}$, then there is a corresponding directed path from e_- to $e'_{\mp}$. This can be obtained by applying the deck transformation of $D(\widehat{A}) \to D(A)$. Hence, there are directed loops starting at each vertex of $D(\widehat{A})$, and if there is a directed path from v to w then there is a directed path form w to v. For example, suppose that there is a directed path from e_+ to e'_+, but no directed path from e'_+ to e_+. Since $D(A)$ is strongly connected, there is a directed path from e' to e and hence there must be a directed path from e'_+ to e_-. This implies there is a directed path from e'_- to e_+. Also there is a directed path from e_- to e'_-. Thus, we have constructed a directed path from e'_+ to e_- to e'_- to e_+, giving a contradiction.

The above discussion implies that $D(\widehat{A})$ is either strongly connected or it has two components each of which is isomorphic to $D(A)$. In the later case we are done since f is orientable. So we henceforth assume that $D(\widehat{A})$ is strongly connected, or equivalently that $\widehat{A}$ is irreducible.

Since A is primitive the gcd of the lengths of directed cycles of $D(A)$ is 1 [BR91, Theorem 3.4.4]. Since $D(\widehat{A}) \to D(A)$ is a 2-fold covering, the gcd of lengths of directed cycles of $D(\widehat{A})$ is at most 2. If it is 1, then $\widehat{A}$ is also primitive and we are done. Otherwise the gcd is 2 so that $\widehat{A}^2$ is reducible and $D(\widehat{A}^2)$ has two connected components. Therefore, we must have that f^2 is pos-orientable by the argument above.

We will now show that f is in fact neg-orientable. Notice that if $\widehat{A}$ is not primitive, then every odd length cycle of $D(A)$ fails to lift to $D(\widehat{A})$. This means the edge path $f^k(e)$ can map over e preserving orientation only when k is even. Therefore $D(\widehat{A})$ is bipartite with each directed edge going from one recurrent component of $D(\widehat{A}^2)$ to the other. Each of these recurrent components orients the graph map f^2 and with this orientation of G, f maps each edge over a negatively oriented edge path. Hence, f is neg-orientable. $\square$

2.2. **Proof of Theorem A.** We now give the proof of Theorem A using the oriented edge-double construction.

Proof of Theorem A. Let $f \colon G \to G$ be a primitive graph map and first suppose f is orientable. Recall A is the primitive transition matrix of f with spectral radius $\lambda := \lambda_f > 1$ and so there is an eigenvector $v \in \mathbb{R}^E$ such that $Av = \lambda v$. Fact 2.2 then implies that v represents a 1-chain (using our identification $C_1(G; \mathbb{R}) \cong \mathbb{R}^E$) such that $Mv = \pm \lambda v$, where the negative sign occurs exactly in the neg-orientable case. In fact, $v \in H_1(G; \mathbb{R}) \leqslant C_1(G; \mathbb{R})$

since otherwise $\partial\colon C_1(G;\mathbb{R}) \to C_0(G;\mathbb{R})$ maps v to an eigenvector whose eigenvalue has absolute value $\lambda > 1$, contradicting that f acts on $C_0(G;\mathbb{R})$ as a permutation. This shows that $f_*(v) = \pm\lambda v$. Since it is always the case that $\rho_f \leqslant \lambda_f$, we conclude that $\lambda_f = \rho_f$ when f is orientable.

Next suppose that f is not orientable. Then by Lemma 2.4, $\widehat{f}\colon \widehat{G} \to \widehat{G}$ is orientable and has primitive transition matrix $\widehat{A}$. Hence by the first paragraph above (and Lemma 2.4) we have $\lambda := \lambda_f = \lambda_{\widehat{f}} = \rho_{\widehat{f}}$. By the Perron-Frobenius theorem, λ is a simple eigenvalue of $\widehat{A}$, i.e. it has algebraic multiplicity 1, and every other eigenvalue μ of $\widehat{A}\colon H_1(\widehat{G}) \to H_1(\widehat{G})$ has $|\mu| < \lambda$. Moreover, there is a λ–eigenvector $\widehat{v} \in H_1(\widehat{G})$ that lies in the positive cone of $C_1(\widehat{G})$ spanned by the 1-chains e_+ and e_- for all $e \in EG$. We claim that $\widehat{v} \in E_{-1}$, where recall $H_1(\widehat{G}) = E_1 \oplus E_{-1}$. Indeed, since σ commutes with $\widehat{f}$, $\sigma_*(\widehat{v})$ is also a λ–eigenvector. Since this eigenspace is one dimensional $\sigma_*(\widehat{v}) = \pm\widehat{v}$. Now σ_* sends the positive cone to the negative cone, therefore it must be that $\widehat{v} \in E_{-1} = \ker(p_*)$ as claimed. From this, we see that any eigenvector of $\widehat{A}$ in E_1 corresponding to an eigenvalue μ has $|\mu| < \lambda$. By Lemma 2.3, $p_*\colon E_1 \to H_1(G)$ is a linear isomorphism that conjugates the action of $\widehat{f}_*$ on E_1 to the action of f_* on $H_1(G)$. We conclude that each eigenvalue μ of f_* has $|\mu| < \lambda$. Hence, $\rho_f < \lambda_f$ and the proof is complete.

The moreover statement now follows from Fact 2.2, since if f is ± 1–orientable then $M = \pm A$ and so $\pm\lambda_f$ is evidently a zero of the polynomial $\det(uI - f_*)$ in the variable u. $\qquad\square$

3. Orientability for free group endomorphisms

In this section we introduce a notion of orientability for fully irreducible free group endomorphisms and characterize this property in terms of stretch factors, train track representatives, and expanding laminations.

3.1. Stretch factors for endomorphisms.
Let $\mathbb{F}$ be a finite-rank free group and $\varphi\colon \mathbb{F} \to \mathbb{F}$ any endomorphism. The *(geometric) stretch factor* of φ is defined as

$$\lambda_\varphi = \sup_{w \in \mathbb{F}} \liminf_{n \to \infty} \sqrt[n]{\|\varphi^n(w)\|},$$

where the supremum is over all elements of $\mathbb{F}$ and $\|\cdot\|$ is the length of the conjugacy class with respect to some (or any) fixed generating set of $\mathbb{F}$. The *homological stretch factor* of φ is defined as the spectral radius for the action of φ_* on $H_1(\mathbb{F};\mathbb{R})$ and is calculated as

$$\rho_\varphi = \sup_{v \in H_1(\mathbb{F},\mathbb{R})} \liminf_{n \to \infty} \sqrt[n]{\|\varphi_*^n(v)\|},$$

where $\|\cdot\|$ denotes some (or any) fixed norm on the vector space $H_1(\mathbb{F};\mathbb{R})$.

Since the conjugacy length on $\mathbb{F}$ and norm on $H_1(\mathbb{F};\mathbb{R})$ can be chosen such that $\|w^{\mathrm{ab}}\| \leqslant \|w\|$ for any element $w \in \mathbb{F}$ (where w^{ab} denotes the image in the abelianization $H_1(\mathbb{F};\mathbb{Z})$ of $\mathbb{F}$), it is immediate that $\rho_\varphi \leqslant \lambda_\varphi$ for any endomorphism.

3.2. Dual splittings.
Recall from §1.2 that any endomorphism $\phi\colon B \to B$ of a group determines a generalized HNN extension $B*_\phi$ along with a dual class given by the projection $B*_\phi \to \mathbb{Z}$. We say that two such HNN extensions are *equivalent* if there is an isomorphism $B_1*_{\phi_1} \cong B_2*_{\phi_2}$ that respects the projections to $\mathbb{Z}$. Note that endomorphisms of distinct groups B_1, B_2 may determine equivalent HNN extensions and that, in principle, such endomorphisms ϕ_1, ϕ_2 may have distinct stretch factors. However, in the case of finite-rank free groups, we have the following relationships between equivalent splittings:

Lemma 3.1. *Let $\varphi \colon \mathbb{F} \to \mathbb{F}$ be an endomorphism of a finite-rank free group.*

- *There exists an injective endomorphism $\phi \colon B \to B$ of a free group with $\mathrm{rank}(B) \leqslant \mathrm{rank}(\mathbb{F})$ such that $\mathbb{F} \ast_\varphi$ and $B \ast_\phi$ are equivalent.*
- *If $\mathbb{F} \ast_\varphi$ is equivalent to $B \ast_\phi$ for some endomorphism $\phi \colon B \to B$ of a finite-rank free group, then $\lambda_\varphi = \lambda_\phi$ and $\rho_\varphi = \rho_\phi$. In fact, the characteristic polynomials of φ and ϕ acting on homology agree up to multiplication by t^k.*

Proof. The first facts are explained in §2.4 of [DKL17b]. Specifically, Corollary 2.8 of that paper provides an injective endomorphism $\phi \colon B \to B$ as in the first item and Proposition 2.9 implies the invariance of geometric stretch factors in the second. The fact that $\rho_\varphi = \rho_\phi$ with equal characteristic polynomials up to a factor is a consequence of our determinant formula for computing the Alexander polynomial, which is an invariant of the group $\mathbb{F} \ast_\varphi$; see Corollary 4.9 and Remark 4.10 below. $\qquad \square$

3.3. **Train track representatives.** A *topological representative* of an endomorphism $\varphi \colon \mathbb{F} \to \mathbb{F}$ is a graph map $f \colon G \to G$ for which there is a choice of basepoint $v \in G$, path β from $f(v)$ to v, and isomorphism $\mathbb{F} \cong \pi_1(G, v)$ that conjugates φ to the endomorphism $\gamma \mapsto \bar{\beta} f(\gamma) \beta$ of $\pi_1(G, v)$. It follows easily from the definitions that if $f \colon G \to G$ represents φ, then $\lambda_\varphi \leqslant \lambda_f$ and $\rho_\varphi = \rho_f$.

In part to obtain more precise control on geometric stretch factors, Bestvina and Handel [BH92] introduced the notion of a *train track map*, which is a graph map $f \colon G \to G$ such that for every edge e and power $k \geqslant 1$ the edge path $f^k(e)$ does not backtrack. The following connection to geometric stretch factors is well-known and is originally due to Bestvina–Handel in the case of automorphisms.

Lemma 3.2 ([DKL17b, Proposition 2.2]). *If f is an irreducible train track representative of a free group endomorphism φ, then $\lambda_\varphi = \lambda_f$.*

3.4. **Full irreducibility.** Recall that an endomorphism $\varphi \colon \mathbb{F} \to \mathbb{F}$ of a free group is *fully irreducible* if for any proper nontrivial free factor $A \leqslant \mathbb{F}$, $\varphi^n(A)$ is not conjugate into A for any $n \geqslant 1$. Note that fully irreducible endomorphisms are necessarily injective (see, e.g., [Mut20, Lemma 2.2]).

In light of Theorem A, it will be important to know that the train track maps we consider are automatically primitive. The following lemma is essentially well-known, at least in the automorphism case, but we provide a proof for completeness and because it is not always stated correctly in the literature.

Lemma 3.3. *Suppose an endomorphism $\varphi \colon \mathbb{F} \to \mathbb{F}$ is either fully irreducible or represented by a pseudo-Anosov homeomorphism. Then φ admits a train track representative $f \colon G \to G$. Moreover, the following are equivalent for any such representative:*

- (1) *f has no invariant forests,*
- (2) *f is primitive,*
- (3) *f is expanding, and*
- (4) *f is irreducible.*

Note that any train track representative satisfies these conditions after collapsing the invariant forest. In the sequel, we usually refer to such train tracks as irreducible.

Proof. It is a seminal result of Bestvina and Handel that automorphisms which are fully irreducible [BH92] or represented by a pseudo-Anosov [BH95, Theorem 4.1.3] admit a train track representative. The existence for non-surjective endomorphisms was proven by Reynolds [Rey11] and independently Mutanguha [Mut20]. It is clear that $(2) \implies (3)$ and that $(2) \implies (4) \implies (1)$. To see $(1) \implies (2)$, suppose f does not have an invariant forest and let A

be its transition matrix. To show that A is primitive, it suffices to show that A^n is irreducible for each $n \geqslant 1$. Otherwise, there is $n \geqslant 1$ and edges e, e' of G such that $(f^n)^j(e)$ does not cross e' for any $j \geqslant 1$. Since there is no invariant forest, the lengths of the nonbacktracking paths $(f^n)^j(e)$ go to infinity in j. These iterates then fill a subgraph G' of G which is not a forest and for which $f^n(G') \subset G'$. Since G' does not contain e', it represents a proper free factor which is an immediate contradiction when φ is fully irreducible. If instead φ is induced by a pseudo-Anosov homeomorphism $h \colon S \to S$ of a punctured surface with $\pi_1(S) \cong \mathbb{F}$, then G' must be a rank 1 subgraph whose embedded core α represents a peripheral curve. Indeed, these peripheral curves are the only free factors that are invariant under a pseudo-Anosov. By construction, α is legal and hence its length grows under iteration. This is a contradiction to the fact that h permutes the peripheral curves of S.

The final implication (3) $\implies$ (4) is similar: For any edge e of G the union of images $f^j(e)$ forms an f–invariant subgraph $G' \subset G$. Since f is an expanding train track map, the nonbacktracking edge paths $f^j(e)$ become arbitrarily long. Hence G' cannot be a forest, since that would force these long immersed paths $f^j(e)$ to embed in the finite graph G'. If G' is a proper subgraph of G, then we obtain a contradiction by the argument above. This shows that f is irreducible. $\qquad\square$

3.5. Orientable laminations. Here we briefly review the theory of expanding laminations and attracting trees for free group endomorphisms. This allows for an intrinsically motivated definition of orientability and completes the analogy with the pseudo-Anosov theory.

In [BFH97, Section 1], Bestvina–Feighn–Handel define a canonical expanding lamination $\mathcal{L}_\varphi^+$ for every fully irreducible automorphism φ of $\mathbb{F}$. One may check that the *auto*morphism hypothesis is only used in their argument to ensure the existence of an irreducible train track representative $f \colon G \to G$. Since endomorphisms also admit such representatives (Lemma 3.3), the canonical construction of $\mathcal{L}_\varphi^+$ from [BFH97] in fact applies to any fully irreducible endomorphism $\varphi \colon \mathbb{F} \to \mathbb{F}$. The lamination is realized as a $\mathbb{F}$–invariant collection of properly embedded bi-infinite lines (called *leaves*) in the universal cover $\widetilde{G}$ that is invariant under the lifted train track map $\widetilde{f} \colon \widetilde{G} \to \widetilde{G}$.

The fully irreducible endomorphism φ also comes equipped with a canonical *attracting* $\mathbb{R}$–*tree* T_φ^+ equipped with an isometric $\mathbb{F}$–action, which is the unique attracting fixed point for the action φ on the closure of Culler–Vogtmann outer space; see [BFH97, LL03, Rey11, Mut20]. The tree may be obtained from the train track representative f via an explicit construction that provides an $\mathbb{F}$–equivariant, 1–Lipschitz quotient map $F \colon \widetilde{G} \to T_\varphi^+$ that restricts to an isometry on each leaf of $\mathcal{L}_\varphi^+$. Further, φ is represented by λ_φ–homothety $f^+ \colon T_\varphi^+ \to T_\varphi^+$ of this tree such that $F \circ \widetilde{f} = f^+ \circ F$; see [HM11, §2.5].

Definition 3.4. The expanding lamination $\mathcal{L}_\varphi^+$ of a fully irreducible endomorphism φ is *orientable* if its leaves may each be $\mathbb{F}$–equivariantly oriented so that their images in the attracting tree T_φ^+ have compatible orientations whenever they overlap.

Remark 3.5. This generalizes the notion of transverse orientability for the contracting measured foliation $\mathcal{F}_-$ of a pseudo-Anosov surface homeomorphism ψ. Indeed, one may lift to the universal cover and collapse the leaves of $\widetilde{\mathcal{F}}_-$ to obtain an $\mathbb{R}$–tree T_ψ^+ in which each leaf of $\widetilde{\mathcal{F}}_+$ maps to a properly embedded bi-infinite line. This collection of lines in T_ψ^+ is orientable in the sense of Definition 3.4 if and only if the contracting foliation $\mathcal{F}_-$ of ψ is transversely orientable on the surface.

3.6. **Characterizing orientability.** We are now ready to establish Theorem B by proving that the various manifestations of orientability all agree for fully irreducible endomorphisms.

Theorem 3.6. *For a fully irreducible free group endomorphism $\varphi \colon \mathbb{F} \to \mathbb{F}$, the following conditions are equivalent:*

(1) *$\lambda_\varphi = \rho_\varphi$ with $+\lambda_\varphi$ (resp. $-\lambda_\varphi$) an eigenvalue of φ_* acting on $H_1(\mathbb{F})$,*
(2) *φ can be represented by a positively (resp. negatively) orientable graph map,*
(3) *every irreducible train track representative of φ is positively (resp. negatively) orientable,*
(4) *the expanding lamination $\mathcal{L}_\varphi^+$ of φ is orientable with the action of φ on the attracting tree T_φ^+ preserving (resp. reversing) the orientations on the leaves of $\mathcal{L}_\varphi^+$.*

Accordingly, we say a fully irreducible endomorphism is *(positively/negatively) orientable* if it satisfies the equivalent conditions of Theorem 3.6. To prove the characterization in terms of laminations, we will use the following lemma:

Lemma 3.7. *For any arc J in T_φ^+, there is a finite chain of leaves ℓ_i of the expanding lamination $\mathcal{L}_\varphi^+$ such that $J \subset \ell_1 \cup \cdots \cup \ell_n$ with each $\ell_i \cap \ell_{i+1}$ a nondegenerate arc (i.e it contains at least two points).*

Proof. In the case that φ is an automorphism, it is well-known that the action $\mathbb{F} \curvearrowright T_\varphi^+$ is *indecomposable* [CH12, Theorem 2.1]. Following Guirardel [Gui08], this means that for any other nondegenerate arc $I \subset T_\varphi^+$, there are $g_1, \ldots g_k \in F_n$ such that $J \subset g_1 I \cup \ldots \cup g_k I$ and $g_i I \cap g_{i+1} I$ is nondegenerate. Hence choosing I to be any arc in a leaf of $\mathcal{L}_\varphi^+$ completes the proof.

When φ is non-surjective, there is an irreducible train track $f \colon G \to G$ that is an immersion with connected Whitehead graphs [Mut20]. Thus by construction the attracting tree T_φ^+ is simply the universal cover $\widetilde{G}$ with an appropriate metric. Recall that the Whitehead graph of a vertex of $\widetilde{G}$ records which turns are taken by leaves of the lamination. Since these graphs are connected, for any vertex of $\widetilde{G}$ we may choose a chain of leaves $\ell_1, \ldots, \ell_k$ with nondegenerate overlaps $\ell_i \cap \ell_{i+1}$ that connect any pair of directions at that vertex. Concatenating such chains for the finitely many vertices along the arc J produces a chain covering J as in the lemma. $\qed$

This has two important consequences when $\mathcal{L}_\varphi^+$ is orientable: First, the orientation on any leaf determines the orientation on all other leaves; thus the orientation is unique up to a global reversal, and the action of f^+ on T_φ^+ either preserves or reverses the orientation. Second, any oriented arc α in T_φ^+ can be decomposed as a concatenation of subarcs of leaves of $\mathcal{L}_\varphi^+$. Since leaves are consistently oriented, each subarc inherits a signed length from the metric on T_φ^+, and we may add to get a signed length $\int_\alpha d\mathcal{L}_\varphi^+$ for α. This assignment $\alpha \mapsto \int_\alpha d\mathcal{L}_\varphi^+$ is $\mathbb{F}$–equivariant and satisfies $\int_{f^+(\alpha)} d\mathcal{L}_\varphi^+ = \pm\lambda_\varphi \cdot \int_\alpha d\mathcal{L}_\varphi^+$, where the sign is positive or negative depending on whether the λ_φ–homothety f^+ preserves or reverses the orientation on $\mathcal{L}_\varphi^+$.

Lemma 3.8. *If $\mathcal{L} = \mathcal{L}_\varphi^+$ is orientable and f^+ preserves (resp. reverses) its orientation and fixes the point $x \in T_\varphi^+$, then the homomorphism $\xi \colon \mathbb{F} \to \mathbb{R}$ determined by*

$$g \mapsto \int_{[x,gx]} d\mathcal{L} \qquad for \ g \in \mathbb{F}$$

defines a λ_φ-eigenvector (resp. $-\lambda_\varphi$-eigenvector) of φ^* in $H^1(\mathbb{F}; \mathbb{R})$.

Proof. First note that by considering the corresponding tripod in T_φ^+, we see that for any $w, y, z \in T_\varphi^+$,

$$\int_{[w,z]} d\mathcal{L} = \int_{[w,y]} d\mathcal{L} + \int_{[y,z]} d\mathcal{L}.$$

Using $\mathbb{F}$–equivariance, this translates into the fact that for any $g, h \in \mathbb{F}$, $\xi(gh) = \xi(g) + \xi(h)$ by setting $w = x$, $y = gx$, $z = ghx$. To complete the proof, observe that

$$\xi(\varphi(g)) = \int_{[x,f^+(gx)]} d\mathcal{L} = \int_{f^+([x,gx])} d\mathcal{L} = \pm\lambda_\varphi \cdot \xi(g),$$

as required. $\square$

Proof of Theorem 3.6. By Lemma 3.3 every irreducible train track representative of φ is primitive. Hence, the equivalence of (1), (2), and (3) follows immediately from Theorem A. Since every subarc of a leaf of $\mathcal{L}_\varphi^+$ in $\widetilde{G}$ is contained in $\widetilde{f}^n(e)$ for some edge of $\widetilde{G}$, it is clear that f being orientable allows one to consistently and equivariantly orient the leaves of $\mathcal{L}_\varphi^+$ in T_φ^+; thus (3) implies (4). Finally, Lemma 3.8 shows that (4) implies (1). $\square$

4. Computing the Alexander polynomial

In this section, we prove Theorem E which gives a determinant formula for the Alexander polynomial of the mapping torus of a graph map. See Example 7.3 for an explicit computation.

4.1. **Fitting ideals and invariants.** Let R be a unique factorization domain (UFD). Throughout, we will use the notation $p \doteq q$ to mean that two elements p, q of R are equal up to multiplication by a unit. Let N be a finitely generated R-module. If N has a presentation,

$$R^m \xrightarrow{M} R^n \to N \to 0,$$

then the *ith Fitting ideal* (also called the *ith determinantal ideal*) $\mathrm{Fit}_i(N) \subset R$ is the ideal generated by the $(n-i) \times (n-i)$ minors of M. It is a fact that these ideals do not depend on the choice of presentation.

Recall that since R is a UFD the greatest common divisor (gcd) of elements of R is well-defined up to a unit in R. For an R-module N as above, the gcd of the elements of $\mathrm{Fit}_i(N)$, $\gcd(\mathrm{Fit}_i(N))$, is called the *ith Fitting invariant* of N.

Let H be a finitely generated free abelian group. Then the group ring $\mathbb{Z}[H]$ is a UFD and its *augmentation ideal* $\mathcal{A}$ is defined to be the kernel of the augmentation homomorphism $\mathbb{Z}[H] \to \mathbb{Z}$ defined by extending the assignment $h \mapsto 1$ for $h \in H$.

Finally, note that any $P \in \mathbb{Z}[H]$ can be written in the form $P = \sum_{h \in H} a_h h$, where each $a_h \in \mathbb{Z}$ and $a_h = 0$ for all but finitely many h. Such a $P \in \mathbb{Z}[H]$ is said to be *symmetric* if, up to a unit in $\mathbb{Z}[H]$, it is equal to $\sum_{h \in H} a_h h^{-1}$. In general the set $\mathrm{supp}(P) = \{h \in H : a_h \neq 0\}$ is called the *support* of P. The *specialization* of P at a class $u \in \mathrm{Hom}(H, \mathbb{Z})$ is the the single-variable Laurent polynomial

$$P^u(t) = \sum_{h \in H} a_h t^{u(h)}.$$

4.2. **Alexander polynomial generalities.** Let X be a connected cell complex with $\Gamma = \pi_1(X)$ and let $H = H_1(X)/\text{torsion}$.

Let $\widetilde{X}^{ab}$ be the universal free abelian cover of X, i.e. the cover corresponding to the kernel of $\Gamma \to H$ whose deck group is H. Fix $x \in X^{(0)}$ and let $\widetilde{x}$ be its full preimage in $\widetilde{X}^{ab}$. The homology group $H_1(\widetilde{X}^{ab}, \widetilde{x})$ naturally has the structure of a $\mathbb{Z}[H]$–module and is called the *Alexander module*. Its first Fitting invariant is called the *Alexander polynomial* and is denoted by Δ_X. In other words,

$$\Delta_X \doteq \gcd \text{Fit}_1(H_1(\widetilde{X}^{ab}, \widetilde{x})),$$

which is an element of $\mathbb{Z}[H]$, defined up to a unit $\pm h \in \mathbb{Z}[H]$. We note that if T is a maximal tree in $X^{(1)}$ and $\widetilde{T}$ is its full preimage in $\widetilde{X}^{ab}$, then we see that the $\mathbb{Z}[H]$–modules $H_1(\widetilde{X}^{ab}, \widetilde{x})$ and $H_1(\widetilde{X}^{ab}, \widetilde{T})$ are isomorphic by considering the long exact sequence associated to the triple $(\widetilde{X}^{ab}, \widetilde{T}, \widetilde{x})$.

In fact, the Alexander polynomial of X depends only on its fundamental group (see [McM02, Section 2]), and so in what follows we also write Δ_Γ to denote the Alexander polynomial Δ_X, where $\Gamma = \pi_1(X)$.

4.3. **Mapping tori and homology.** Let $f\colon G \to G$ be a graph map on a finite graph G. Fix an arbitrary orientation on each edge of G and let $V = VG$ and $E = EG$ respectively denote the sets of vertices and (oriented) edges of G. Let $X := X_f$ be the mapping torus with the usual cell structure, so that $\pi_1(X) = \pi_1(G){*}_f$ is the associated generalized HNN-extension.

In a bit more detail: Give $G \times [0,1]$ the cell structure induced from the product cell structure so that the oriented 2-cells induce the given orientation on the edges of $G \times \{0\}$ and the opposite orientation on the edges of $G \times \{1\}$. Then subdivide $G \times \{1\}$ so that $f\colon G \times \{1\} \to G \times \{0\}$ maps open cells homeomorphically onto their images. The mapping torus X is then the quotient

$$X = G \times [0,1]/(x,1) \sim (f(x), 0),$$
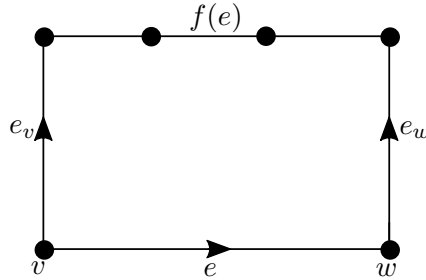
with its induced cellular structure. See Figure 1.

FIGURE 1. A 2-cell of X_f with indicated orientations on edges.

The mapping torus comes with a natural map $X \to S^1$. In the cell structure we view edges of G as "horizontal" and the other edges as "vertical," oriented so that each maps to a positively oriented based loop in S^1 representing the generator $1 \in \mathbb{Z}$. Let $\mathcal{E}$ denote the set of (oriented) edges in X. Each vertex $v \in V$ is the tail (i.e., initial vertex) of a unique vertical edge e_v in X. This gives a bijection $v \mapsto e_v$ which we use to identify V with the set of vertical edges of X. When we wish to be explicit, we denote this bijection by ν. Similarly each edge $e \in E$ gives rise to an oriented 2–cell σ_e of X whose boundary crosses e positively and the image path $f(e)$ negatively. In this way we see that the 1–cells of X are

in bijective correspondence with $E \sqcup V$, and the 2–cells with the edges E of G. With these identifications, the $\mathbb{Z}$–modules of 0–chains in X are $C_0(X) = C_0(G) \cong \mathbb{Z}[V]$, the 1–chains are $C_1(X) \cong \mathbb{Z}[\mathcal{E}] \cong \mathbb{Z}[E \sqcup V]$, and the 2–chains are $C_2(X) \cong \mathbb{Z}[E]$.

The map $f \colon G \to G$ restricts to a map $V \to V$ which we express as an integer $V \times V$ matrix P whose (v, w) entry equals 1 if $f(w) = v$ and is zero otherwise. That is, the map that f induces on 0–chains $C_0(G) \cong \mathbb{Z}[V] = \mathbb{Z}^V$ is simply the linear map $\mathbb{Z}^V \to \mathbb{Z}^V$ given by $x \mapsto Px$. Similarly, the map that f induces on $C_1(G) \cong \mathbb{Z}^E$ is expressed as an integer $E \times E$ matrix M whose (e, e') entry is the number of times the edge path $f(e')$ crosses e, counted with signs. The boundary maps on chains may then be written as

$$
\begin{array}{ccccc}
C_2(X) & \xrightarrow{\quad \partial_2 \quad} & C_1(X) & \xrightarrow{\quad \partial_1 \quad} & C_0(X) \\
\Big\downarrow{\cong} & \begin{pmatrix} \pi_E \partial_2 \\ \pi_V \partial_2 \end{pmatrix} & \Big\downarrow{\cong} & \begin{pmatrix} \partial_1|_E & \partial_1|_V \end{pmatrix} & \Big\downarrow{\cong} \\
\mathbb{Z}^E & \xrightarrow{\hspace{2.5cm}} & \mathbb{Z}^E \oplus \mathbb{Z}^V & \xrightarrow{\hspace{2cm}} & \mathbb{Z}^V
\end{array}
$$

That is, $\partial_1|_E \colon \mathbb{Z}^E \to \mathbb{Z}^V$ is the restriction of ∂_1, viewed as an $E \times V$ matrix, and similarly for $\partial_1|_V$. Likewise $\pi_E \partial_2$ is the $E \times E$ matrix describing the projection $\mathbb{Z}^E \xrightarrow{\partial_2} \mathbb{Z}^E \oplus \mathbb{Z}^V \xrightarrow{\pi_E} \mathbb{Z}^E$, and similarly for $\pi_V \partial_2$. Note that the boundary of a vertical edge e_v corresponding to its tail vertex $v \in V$, is $f(v) - v$. Hence $\partial_1|_V = P - I$. Similarly, if an edge $e \in E \subset C_1(X)$ has image $f(e) = a_1 e_1 + \cdots + a_k e_k \in C_1(X)$, then the corresponding 2–cell has boundary

$$\partial_2(\sigma_e) = e - a_1 e_1 - \cdots - a_k e_k + \nu \circ \partial_1(e).$$

That is to say we have $\pi_E \partial_2 = I - M$, and $\pi_V \partial_2 = \partial_1|_E$.

4.4. The universal free abelian cover of X.

To calculate the Alexander polynomial for $\pi_1(X)$, let $H = H_1(X)/\text{torsion}$ and as usual let $\widetilde{X}^{ab} \to X$ be the universal free abelian cover with deck group H. In what follows, we abbreviate $\widetilde{X}^{ab}$ by $\widetilde{X}$. Again, we view lifts of edges of G as being horizontal in $\widetilde{X}$, and lifts of the oriented vertical edges in X as being oriented upward.

Let $\widetilde{G} \subset \widetilde{X}$ be the full preimage of G, and choose some connected component $\widetilde{G}_0$. Then $\widetilde{G}_0 \to G$ is the free abelian cover with deck group K, where K is the image of the homomorphism $\pi_1(G) \to \pi_1(X) \to H$. Alternatively, K can be intrinsically characterized as the quotient of $H_1(G)$ isomorphic to $\text{Hom}(H^1(G)^f, \mathbb{Z})$ by duality. Here $H^1(G)^f$ is the cohomology of G that is fixed by f (see [DKL17b, Section 8]).

The graph map $f \colon G \to G$ lifts to $\widetilde{G}_0$ and we fix such a lift

$$\widetilde{f} \colon \widetilde{G}_0 \to \widetilde{G}_0$$

once and for all. This lift determines a splitting $H = K \oplus \langle z \rangle$ where z is the deck transformation chosen so that for any $\widetilde{x} \in \widetilde{G}_0$, $\widetilde{f}(\widetilde{x})$ is the first point in $\widetilde{G}_0$ encountered by the lift of the arc $x \times [0, 1]$ starting at $z\widetilde{x}$. Here x is the image of $\widetilde{x}$ under the covering $\widetilde{X} \to X$. Note that the splitting $H = K \oplus \langle z \rangle$ depends on our choice of $\widetilde{f}$ and z has been chosen to map to -1 under $\pi_1(X) \to \pi_1(S^1)$. Hence, z translates downward as a deck transformation of $\widetilde{X}$. For notational convenience, we set $\bar{z} = z^{-1} \in H$. For each $i \in \mathbb{Z}$ let $\widetilde{G}_i = \bar{z}^i \widetilde{G}_0$; this decomposes $\widetilde{G}$ into its connected components $\widetilde{G} = \bigsqcup_i \widetilde{G}_i$, with $\widetilde{G}_{i+1}$ situated above $\widetilde{G}_i$.

For each cell in G, we fix arbitrary lifts in $\widetilde{X}$ as follows. For each $v \in V$ we fix a lift $\widetilde{v} \in \widetilde{G}_0$ and we let $\widetilde{e}_v$ denote the lift of e_v with tail $\widetilde{v}$. For each edge $e \in E$ we fix, once and for all, a lift $\widetilde{e} \in \widetilde{G}_0$ which may or may not be based at the chosen lifts of vertices of G. We then let $\widetilde{\sigma}_e$ be the lift of the 2–cell σ_e whose boundary crosses $\widetilde{e}$ positively and $\bar{z}\widetilde{f}(\widetilde{e}) \subset \widetilde{G}_1$ negatively.

Note that this chooses a preferred lift to $\widetilde{X}$ for each edge of $\mathcal{E}$ which we call the *base lifts*. Since H acts as deck transformations on $\widetilde{X}$, the $\mathbb{Z}$–modules of chains in $\widetilde{X}$ are naturally free $\mathbb{Z}[H]$–modules. Using our chosen base lifts as basis elements thus gives $\mathbb{Z}[H]$–module isomorphisms

$$C_2(\widetilde{X}) \cong \mathbb{Z}[H]^E, \quad C_1(\widetilde{X}) \cong \mathbb{Z}[H]^E \oplus \mathbb{Z}[H]^V, \quad \text{and} \quad C_0(\widetilde{X}) \cong \mathbb{Z}[H]^V.$$

The action of K on $\widetilde{G}_0$ by deck transformations similarly gives $C_1(\widetilde{G}_0) \cong \mathbb{Z}[K]^E$ and $C_0(\widetilde{G}_0) \cong \mathbb{Z}[K]^V$. With these identifications, the maps that $\widetilde{f}$ induces on $C_1(\widetilde{G}_0)$ and $C_0(\widetilde{G}_0)$ are expressed by $\mathbb{Z}[K]$–matrices $\widetilde{M}$ and $\widetilde{P}$ that respectively specialize to M and P under the ring morphism $\mathbb{Z}[K] \to \mathbb{Z}$ induced by the trivial group homomorphism $K \to \{1\}$.

4.5. The chain complex of $\widetilde{X}$.

We now examine the boundary maps $\widetilde{\partial}_i$ for the cellular homology of $\widetilde{X}$. Firstly, let us write $\widetilde{B} = \widetilde{\partial}_1|_{\mathbb{Z}[H]^E}$ for the matrix representing the restriction of $\widetilde{\partial}_1$ to the horizontal edges of $\widetilde{X}$. Since the basis horizontal edges $\widetilde{e}$ and their boundary vertices lie in $\widetilde{G}_0$, we see that $\widetilde{B}$ is in fact a $\mathbb{Z}[K]$–valued $V \times E$ matrix. With this notation, each basis 2–cell $\widetilde{\sigma}_e$ and each basis vertical 1–cell $\widetilde{e}_v$ has boundary

$$\widetilde{\partial}_2(\widetilde{\sigma}_e) = \widetilde{\nu} \circ \widetilde{\partial}_1(\widetilde{e}) + \widetilde{e} - \bar{z}\widetilde{f}(\widetilde{e}) = \widetilde{B}\widetilde{e} + \widetilde{e} - \bar{z}\widetilde{M}\widetilde{e} \qquad \text{and} \qquad \widetilde{\partial}_1(\widetilde{e}_v) = \bar{z}\widetilde{P}\widetilde{v} - \widetilde{v}.$$

Therefore the chain complex that computes the cellular homology of $\widetilde{X}$ takes the form

$$
\begin{array}{ccccc}
C_2(\widetilde{X}) & \xrightarrow{\quad\widetilde{\partial}_2\quad} & C_1(\widetilde{X}) & \xrightarrow{\quad\widetilde{\partial}_1\quad} & C_0(\widetilde{X}) \\
\Big\downarrow{\scriptstyle\cong} & & \Big\downarrow{\scriptstyle\cong} & & \Big\downarrow{\scriptstyle\cong} \\
\mathbb{Z}[H]^E & \xrightarrow{\left(\begin{smallmatrix} I-\bar{z}\widetilde{M} \\ \widetilde{B} \end{smallmatrix}\right)} & \mathbb{Z}[H]^E \oplus \mathbb{Z}[H]^V & \xrightarrow{\left(\begin{smallmatrix} \widetilde{B} & \bar{z}\widetilde{P}-I \end{smallmatrix}\right)} & \mathbb{Z}[H]^V
\end{array}
$$

The chain condition $\widetilde{\partial}_1\widetilde{\partial}_2 = 0$ has a useful interpretation in terms of the coboundary maps $\widetilde{\delta^1}\colon C^0(\widetilde{X}) \to C^1(\widetilde{X})$ and $\widetilde{\delta^2}\colon C^1(\widetilde{X}) \to C^2(\widetilde{X})$ on cochains, where our bases induce isomorphisms $C^i(\widetilde{X}) \cong C_i(\widetilde{X}) \cong \mathbb{Z}[H]^{d_i}$ with respect to which $\widetilde{\delta^i} = \widetilde{\partial}_i^{\,t}$ is simply the transpose of the boundary matrix. For any vertex w in $\widetilde{X}$, the vector $\widetilde{\delta^1}(w) \in C^1(\widetilde{X})$ is the sum of those edges of $\widetilde{X}$ that terminate at w minus the sum of edges that start at w. Similarly, for each edge η of $\widetilde{X}$, the vector $\widetilde{\delta^2}(\eta) \in C^2(\widetilde{X})$ is the sum of 2–cells whose boundary crosses η positively minus those whose boundary crosses η negatively. For short hand, let us set $R_\eta = \widetilde{\delta^2}(\eta) \in C^2(\widetilde{X}) \cong \mathbb{Z}[H]^E$. Thus the row of the matrix $\widetilde{\partial}_2$ corresponding to an index $x \in \mathcal{E}$ is precisely $R_{\widetilde{x}}$. Now for any cochain $w \in C^0(\widetilde{X})$, we may express $\widetilde{\delta^1}(w) \in C^1(\widetilde{X})$ as a vector in $\mathbb{Z}[H]^{\mathcal{E}}$ and apply $\widetilde{\delta^2}$ to obtain a $\mathbb{Z}[H]$–linear dependence relation between the rows of $\widetilde{\partial}_2$:

$$(4.1) \qquad \widetilde{\delta^1}(w) = \sum_{x \in \mathcal{E}} a_x \widetilde{x} \quad \implies \quad 0 = \widetilde{\delta^2}(\widetilde{\delta^1}(w)) = \sum_{x \in \mathcal{E}} a_x R_{\widetilde{x}}.$$

4.6. Relative homology and vertex cycles.

We will use relative homology to calculate the Alexander polynomial, and for this it will be convenient to use a well-chosen maximal tree in $X^{(1)}$. Let $\mathcal{V} \subset X$ be the union of all vertical edges, decomposed into its connected components $\mathcal{V} = \mathcal{V}_1 \sqcup \cdots \sqcup \mathcal{V}_m$. Note that each $\mathcal{V}_k$ is a directed embedded *vertex cycle*, denoted $\mathcal{V}_k'$, with trees hanging off. Here the loop corresponds to a periodic orbit $v, f(v), \ldots, f^{\ell_k}(v) = v$ of length $\ell_k \geqslant 1$ for the action of f on $V = VG$ and the hanging trees correspond to other vertices $w \in V$ that map into the orbit $\{f(v), \ldots, f^{\ell_k}(v)\}$. Let $c_k = [\mathcal{V}_k] = [\mathcal{V}_k'] \in H$ denote the homology class defined by the cycle in $\mathcal{V}_k$.

Lemma 4.2. *We have* $\det(I - \bar{z}\widetilde{P}) \doteq \prod_{i=1}^{m}(1 - c_i)$.

In §5.4, this formula will be used to define the vertex polynomial.

Proof. The components $\mathcal{V}_1 \sqcup \cdots \sqcup \mathcal{V}_m$ determine a block diagonal decomposition of $I - \bar{z}\widetilde{P}$ into blocks $I - \bar{z}\widetilde{P}_i$ for $1 \leq i \leq m$ and so it suffices to show that $\det(I - \bar{z}\widetilde{P}_i) = 1 - c_i$. As mentioned above, $\mathcal{V}_i$ is an embedded loop $\mathcal{V}_i'$ along with trees hanging off. If we index the vertices so that they are ordered in a way that is compatible with their distance to $\mathcal{V}_i'$, then it follows that $\det(I - \bar{z}\widetilde{P}_i) = \det(I - \bar{z}\widetilde{P}_i')$, where $\widetilde{P}_i'$ is the submatrix of $\widetilde{P}_i$ associated to the vertices of $\mathcal{V}_i'$.

Now let v be a vertex in $\mathcal{V}_i'$ and suppose that the vertical loop $\mathcal{V}_i'$ has the vertex sequence $v = v_0, v_1, \ldots, v_{\ell_k} = v$. If we lift the vertical loop $\mathcal{V}_i'$ to $\widetilde{X}$ starting at $\widetilde{v}$, we obtain a path from $\widetilde{v}$ to $c_i\widetilde{v}$. The vertex sequence of this lift is

$$\widetilde{v} = \widetilde{v}_0, g_1\widetilde{v}_1, \ldots, g_{\ell_k}\widetilde{v}_{\ell_k} = c_i\widetilde{v},$$

where $\widetilde{v}_i$ is the base lift of v_i and $g_i \in H$. By our choice of lift $\widetilde{f}\colon \widetilde{G}_0 \to \widetilde{G}_0$ and splitting $H = K \oplus \langle z \rangle$, it follows that $g_{j+1}\widetilde{v}_{j+1} = \bar{z}\widetilde{P}(g_j\widetilde{v}_j)$. So if we set $x_j = g_j^{-1}g_{j+1}$, then $x_j\widetilde{v}_{j+1} = \bar{z}\widetilde{P}(\widetilde{v}_j)$. With this, and the fact that the matrices $\widetilde{P}$ and $\widetilde{P}_i'$ agree on the base lifts of vertices in $\mathcal{V}_i'$, an easy computation now gives the needed equality

$$\det(I - \bar{z}\widetilde{P}_i') = 1 - \prod_{j=0}^{\ell_k - 1} x_j = 1 - c_i. \qquad \square$$

Choose a maximal tree T_k inside of each component $\mathcal{V}_k$; since $\mathcal{V}_k$ is a loop with trees hanging off, this amounts to deleting exactly one edge η_k so that $\mathcal{V}_k = T_k \cup \eta_k$. The union $T_1 \sqcup \cdots \sqcup T_m$ is then a forest in X, and we extend this to a maximal tree $T \subset X$ by adding horizontal edges that connect distinct components. We then use $\mathcal{E}\backslash T$ to denote the set of edges of X that are not contained in T and note that $|\mathcal{E}\backslash T| = |E| + 1$.

Let $\widetilde{T}$ be the full preimage of T in $\widetilde{X}$. The $\mathbb{Z}[H]$–module of relative 1–chains for the pair $(\widetilde{X}, \widetilde{T})$ is identified with $C_1(\widetilde{X}, \widetilde{T}) \cong \mathbb{Z}[H]^{\mathcal{E}\backslash T}$ and the chain complex for relative homology becomes

$$C_2(\widetilde{X}, \widetilde{T}) \cong \mathbb{Z}[H]^E \xrightarrow{A} C_1(\widetilde{X}, \widetilde{T}) \cong \mathbb{Z}[H]^{\mathcal{E}\backslash T} \longrightarrow C_0(\widetilde{X}, \widetilde{T}) = 0$$

where A is simply the $(\mathcal{E}\backslash T) \times E$ matrix obtained from $\widetilde{\partial}_2$ by deleting the rows corresponding to edges in T. It follows that the $\mathbb{Z}[H]$–module $H_1(\widetilde{X}, \widetilde{T})$ may be presented as

$$(4.3) \qquad \mathbb{Z}[H]^E \xrightarrow{A} \mathbb{Z}[H]^{\mathcal{E}\backslash T} \longrightarrow H_1(\widetilde{X}, \widetilde{T}) \to 0.$$

The Alexander polynomial $\Delta_\Gamma \in \mathbb{Z}[H]$ is thus the gcd of the $E \times E$ minors of A. We emphasize that here and throughout all identifications are made using the base lifts fixed in §4.4.

Since H acts freely on the connected components of $\widetilde{T}$, fixing a component $\widetilde{T}'$ gives an isomorphism $H_0(\widetilde{T}) \cong \mathbb{Z}[H]$ as $\mathbb{Z}[H]$–modules. Let $\alpha\colon \mathbb{Z}[H]^{\mathcal{E}\backslash T} \to \mathbb{Z}[H]$ be the composition

$$(4.4) \qquad \mathbb{Z}[H]^{\mathcal{E}\backslash T} \cong C_1(\widetilde{X}, \widetilde{T}) \longrightarrow H_1(\widetilde{X}, \widetilde{T}) \xrightarrow{\sigma} H_0(\widetilde{T}) \cong \mathbb{Z}[H],$$

where σ is the connecting map for the long exact sequence of the pair $(\widetilde{X}, \widetilde{T})$. Note that if e is any edge of $\mathcal{E}\backslash T$ and $\widehat{e}$ is its (nonbase) lift to $\widetilde{X}$ whose tail is contained in $\widetilde{T}'$, then $\alpha(\widehat{e}) = g - 1$, where g is represented by the loop in $e \cup T$ crossing e positively. From this, we see that the image of α is the usual augmentation ideal $\mathcal{A} \subset \mathbb{Z}[H]$.

4.7. **The fundamental relations and the main lemma.** For each $i = 1, \ldots, m$, let $w_i \in C^0(\widetilde{X})$ be the sum of the vertices of $\widetilde{T}_i'$, which is the lift of T_i contained in the fixed component $\widetilde{T}'$. The coboundary of w_i then has the form $\widetilde{\delta}^1(w_i) = \sum_{x \in \mathcal{E}} b_x^i \widetilde{x}$ for some ring elements $b_x^i \in \mathbb{Z}[H]$. Let us calculate these coefficients b_x^i for $x \in \mathcal{E}$.

Vertical edges: Recall the set of vertical edges $\mathcal{V}$ is a disjoint union $\mathcal{V}_1 \sqcup \cdots \sqcup \mathcal{V}_m$. If $x \in \mathcal{V}_k$ for some $k \neq i$, then each lift $\hat{x}$ of x is disjoint from $\widetilde{T}_i'$ and hence has $\widetilde{\delta}^1(w_i)(\hat{x}) = 0$; thus $b_x^i = 0$ in this case. If $x \in \mathcal{V}_i$, then it is either in T_i or is equal to η_i. In the first case, there is exactly one lift $\hat{x}$ that intersects $\widetilde{T}_i'$, but this lift has *both* its tip and tail in w_i meaning $\widetilde{\delta}^1(w_i)(\hat{x}) = w_i(\widetilde{\partial}_1(\hat{x})) = 0$. Hence $b_x^i = 0$ for each vertical edge x in T_i.

For $x = \eta_i$, there are precisely two lifts incident on $\widetilde{T}_i'$. Writing $\widetilde{\eta}_i$ for the base lift of η_i, there is one translate, say $g\widetilde{\eta}_i$, whose tip lies in $\widetilde{T}_i'$ and another whose tail lies in $\widetilde{T}_i'$. Recall that $c_i \in H$ is the homology class of the unique nontrivial loop in $T \cup \eta_i$ that crosses η_i once positively. Lifting this loop to $\widetilde{X}$, we see that the tip of $g\widetilde{\eta}_i$ lying in $\widetilde{T}_i'$ implies the tail of $g\widetilde{\eta}_i$ lies in $c_i^{-1}\widetilde{T}'$. Thus $c_i g\widetilde{\eta}_i$ is the unique lift of η_i whose tail lies in $\widetilde{T}_i'$. Since $\widetilde{\delta}^1(w_i)$ evaluates to 1 on $g\widetilde{\eta}_i$ and to -1 on $c_i g\widetilde{\eta}_i$, we conclude that the coefficient of $\widetilde{\eta}_i$ in $\widetilde{\delta}^1(w_i)$ is $b_{\eta_i}^i = g(1 - c_i)$. Notice that by definition of $\alpha \colon C_1(\widetilde{X}, \widetilde{T}) \to H_0(\widetilde{T}) \cong \mathbb{Z}[H]$ we have $\alpha(g\widetilde{\eta}_i) = 1 - c_i^{-1}$ since the tip and tail of $g\widetilde{\eta}_i$ lie in $\widetilde{T}'$ and $c_i^{-1}\widetilde{T}'$. Therefore we have $b_{\eta_i}^i = -g^2 c_i^{-1} \alpha(\widetilde{\eta}_i) = -g_i \alpha(\widetilde{\eta}_i)$ for some unit $g_i \in \mathbb{Z}[H]$.

Horizontal edges: First let x be a horizontal edge in the maximal tree T. If both end points both of x lie outside T_i, then $b_x^i = 0$ since each lift $\hat{x}$ of x is disjoint from $\widetilde{T}_i'$ and hence has $\widetilde{\delta}^1(w_i)(\hat{x}) = 0$. If instead x has exactly one end point in T_i, then x has exactly one lift $\hat{x}$ that intersects $\widetilde{T}_i'$, and this lift has only one endpoint in $\widetilde{T}_i'$; hence $b_x^i \in \mathbb{Z}[H]$ is a unit. Let D_i denote the set of horizontal edges in T that have exactly one end point in T_i. Finally, the nature of the coefficient b_x^i for $x \in E \backslash T$ will not be relevant for us.

Combining these observations, we see that

$$\widetilde{\delta}^1(w_i) = -g_i \alpha(\widetilde{\eta}_i)\widetilde{\eta}_i + \sum_{x \in D_i \cup (E \backslash T)} b_x^i \widetilde{x}$$

where b_x^i is a unit in $\mathbb{Z}[H]$ for each x in D_i and g_i is also a unit. The Observation (4.1) therefore gives us a relation:

$$(4.5) \qquad g_i \alpha(\widetilde{\eta}_i) R_{\eta_i} = \sum_{x \in D_i \cup E \backslash T} b_x^i R_x \qquad \text{for each } i = 1, \ldots, m.$$

That is: For each vertical edge η_i in $\mathcal{E} \backslash T$, multiplying the corresponding row R_{η_i} of A by $g_i \alpha(\widetilde{\eta}_i)$ produces a linear combination of the rows R_x for horizontal x in $\mathcal{E} \backslash T$ plus a linear combination *with unit coefficients* of the rows R_x of $\widetilde{\partial}_2$ corresponding to x in D_i.

Lemma 4.6. *Let A_k denote the matrix obtained from A (from equation (4.3)) by removing the row for the vertical edge $\eta_k \in \mathcal{E} \backslash T$. Then $\det(A_k) \det(I - \bar{z}\widetilde{P}) \doteq \alpha(\widetilde{\eta}_k) \det(I - \bar{z}\widetilde{M})$.*

Proof. We need to relate the determinant of A_k to that of $I - \bar{z}\widetilde{M}$, which is the submatrix of $\widetilde{\partial}_2$ corresponding to the horizontal edges E. The matrix A_k already contains all rows R_x of $\widetilde{\partial}_2$ corresponding to horizontal edges x in $E \backslash T$. However, A_k is missing the rows corresponding to horizontal edges in T and in their place has extra rows corresponding to vertical edges of $\mathcal{E} \backslash T$ of the form η_i except η_k. We will recursively use the fundamental relation (4.5) to replace each row for a vertical edge of A_k with the row for a missing horizontal edge.

Consider $\det(I-\bar{z}\widetilde{P})A_k$. By Lemma 4.2 we have $\det(I-\bar{z}\widetilde{P}) \doteq \prod_{i=1}^{m}(1-c_i) = \prod_{i=1}^{m}(g_i\alpha(\widetilde{\eta}_i))$. For each $i \neq k$, let us multiply the η_i–row of A_k by the $g_i\alpha(\widetilde{\eta}_i)$ factor from $\det(I-\bar{z}\widetilde{P})$. This produces a matrix B' whose η_i–row is the linear combination $\sum_{x\in D_i\cup(E\backslash T)} b_x^i R_x$. Since B' contains the row R_x for each x in $E\backslash T$, we may perform row operations to kill the terms $b_x^i R_x$ with $x \in E\backslash T$ in the η_i–row of B'. These operations yield a new matrix B whose row for $x \in E\backslash T$ is R_x and whose row for η_i with $i \neq k$ is $\sum_{x\in D_i} b_x^i R_x$. Since these row operations do not change determinants, we further see that

$$\alpha(\widetilde{\eta}_k)\det(B) \doteq \det(I-\bar{z}\widetilde{P})\det(A_k).$$

It thus remains to show that $\det(B) \doteq \det(I-\bar{z}\widetilde{M})$. For this, let Λ denote the set of horizontal edges of T, that is, $\Lambda = D_1 \cup \cdots \cup D_m$. Impose a partial order on Λ defined by $e \leqslant e'$ if the path in the tree T from e' to T_k passes through e. For each $i \neq k$, there is a unique tree path γ from T_i to T_k. This path γ crosses exactly one edge of D_i, which we denote $y_i \in D_i$. Since the path from any other x in D_i to T_k must follow γ and thus pass through y_i, we see that $y_i \leqslant x$ for all $x \in D_i$; hence y_i is the minimal element of D_i. Conversely, each horizontal edge in T is equal to the minimal edge $y_i \in D_i$ for a unique index $i \neq k$. Thus we have a bijection between Λ and $\{1,\ldots,m\}\backslash\{k\}$, so that $\Lambda = \{y_i \mid i \neq k\}$.

Say that a partition $\Lambda = \Lambda_1 \sqcup \Lambda_2$ is $\leqslant$–compatible if there do not exist $\lambda_j \in \Lambda_j$ such that $\lambda_2 \leqslant \lambda_1$; that is, if for every pair $(\lambda_1,\lambda_2) \in \Lambda_1 \times \Lambda_2$ either $\lambda_1 \leqslant \lambda_2$ or else the pair is not ordered. Given such a partition let $B(\Lambda_1 \sqcup \Lambda_2)$ be the matrix obtained from B as follows: for each $i \neq k$ with $y_i \in \Lambda_2$, replace the row $\sum_{x\in D_i} b_x^i R_x$ in B corresponding to η_i with simply the row R_{y_i} of $I-\bar{z}\widetilde{M}$ corresponding to $y_i \in D_i$. Clearly we have $B(\Lambda \sqcup \varnothing) = B$ and $B(\varnothing \sqcup \Lambda) = I-\bar{z}\widetilde{M}$.

We claim by induction that $\det(B(\Lambda_1 \sqcup \Lambda_2)) \doteq \det(B)$ for every $\leqslant$–compatible partition $\Lambda = \Lambda_1 \sqcup \Lambda_2$. This is immediate for the trivial partition $\Lambda \sqcup \varnothing$. Now let $\Lambda_1 \sqcup \Lambda_2$ be any $\leqslant$–compatible partition. If $\Lambda_2 \neq \varnothing$, we may choose a $\leqslant$–minimal element $\lambda \in \Lambda_2$. The adjusted partition $\Lambda_1' \sqcup \Lambda_2' = (\Lambda_1 \cup \{\lambda\}) \sqcup (\Lambda_2\backslash\{\lambda\})$ is then $\leqslant$–compatible, and by induction we may assume $\det(B(\Lambda_1' \sqcup \Lambda_2')) \doteq \det(B)$. Let $i \neq k$ be the index so that $\lambda = y_i$. The matrices $B(\Lambda_1 \sqcup \Lambda_2)$ and $B(\Lambda_1' \sqcup \Lambda_2')$ then only differ in the row corresponding to η_i, which in the latter is $\sum_{x\in D_i} b_x^i R_x$. Consider any edge $x \in D_i$ with $x \neq y_i$. Then by definition of y_i we have $y_i \leqslant x$. Therefore, the fact that $\Lambda_1 \sqcup \Lambda_2$ is $\leqslant$–compatible and $y_i \in \Lambda_2$ implies we cannot have $x \in \Lambda_1$. Hence $x \in \Lambda_2$ and, since $x \neq y_i = \lambda$, furthermore $x \in \Lambda_2'$. This means $B(\Lambda_1' \sqcup \Lambda_2')$ contains the row R_x for each $x \in D_i$ with $x \neq y_i$. Thus we may perform a row operation on $B(\Lambda_1' \sqcup \Lambda_2')$ that uses the row R_x to kill the $b_x^i R_x$ term in the η_i–row of $B(\Lambda_1' \sqcup \Lambda_2')$. Applying these operations for each $x \in D_i\backslash\{y_i\}$ produces a new matrix whose η_i–row is $b_{y_i}^i R_{y_i}$. Multiplying this row by the unit $(b_{y_i}^i)^{-1}$ then yields $B(\Lambda_1 \sqcup \Lambda_2)$. Thus we have transformed $B(\Lambda_1' \sqcup \Lambda_2')$ into $B(\Lambda_1 \sqcup \Lambda_2)$ via row operations that only effect determinants up to a unit. This proves the claim and, in particular, that $\det(I-\bar{z}\widetilde{M}) = \det(B)$ up to a unit in $\mathbb{Z}[H]$, as needed. $\square$

4.8. **The determinant formula.** We will make use of the following lemma which can be found in work of Hironaka [Hir11, Lemma 8]. The argument also appears in the proof of [McM02, Theorem 5.1].

Lemma 4.7. *Suppose*

$$\mathbb{Z}[H]^n \xrightarrow{A} \mathbb{Z}[H]^{n+1} \xrightarrow{\alpha} \mathbb{Z}[H],$$

is a sequence of $\mathbb{Z}[H]$–modules such that $\alpha \circ A = 0$. Let Δ_i be the determinant of the matrix obtained from A after removing the ith row. Then

(4.8) $$\alpha(e_i)\Delta_j = \pm\alpha(e_j)\Delta_i,$$

for all $1 \leqslant i, j \leqslant n + 1$. Here $\{e_i : 1 \leqslant i \leqslant n + 1\}$ is the standard free basis of $\mathbb{Z}[H]^{n+1}$.

We are now ready to prove the determinant formula for the Alexander polynomial:

Theorem E (Determinant formula for Δ_Γ). *Let $f \colon G \to G$ be a graph map whose mapping torus X has fundamental group Γ. Then*

$$\Delta_\Gamma \doteq \frac{\det(zI - \widetilde{M})}{\det(zI - \widetilde{P})} \cdot r,$$

where $r = 1 \in \mathbb{Z}[H]$ if $\mathrm{rank}(H_1(X, \mathbb{R})) \geqslant 2$ and $r = (z - 1)$ if $\mathrm{rank}(H_1(X, \mathbb{R})) = 1$.

Proof of Theorem E. Recall that $|\mathcal{E}\backslash T| = |E| + 1$ and, for $1 \leqslant i \leqslant |E| + 1$, let Δ_i be the minor of A (from equation (4.3)) obtained by removing the ith row. Then the Alexander polynomial is by definition

$$\Delta_\Gamma = \gcd\{\Delta_i : 1 \leqslant i \leqslant |E| + 1\}.$$

With $\alpha \colon \mathbb{Z}[H]^{\mathcal{E}\backslash T} \to \mathbb{Z}[H]$ from equation (4.4), we note that $\alpha \circ A = 0$ and so we apply Lemma 4.7 to conclude that

$$\alpha(\widetilde{e}_i)\Delta_j = \pm\alpha(\widetilde{e}_j)\Delta_i,$$

where $\{\widetilde{e}_i\}$ are the base lifts of the edges of $\mathcal{E}\backslash T$. Fix some index, say k, representing a vertical edge η_k of $\mathcal{E}\backslash T$ so that $\widetilde{e}_k = \widetilde{\eta}_k$ as in Lemma 4.6. Then

$$\begin{aligned}
\alpha(\widetilde{\eta}_k)\Delta_\Gamma &= \gcd\{\alpha(\widetilde{\eta}_k)\Delta_i : 1 \leqslant i \leqslant |E| + 1\} \\
&= \gcd\{\alpha(\widetilde{e}_i) : 1 \leqslant i \leqslant |E| + 1\} \cdot \Delta_k.
\end{aligned}$$

But according to Lemma 4.6, we know that,

$$\Delta_k \det(I - \bar{z}\widetilde{P}) \doteq \alpha(\widetilde{\eta}_k) \det(I - \bar{z}\widetilde{M}).$$

After rearranging, we conclude that

$$\Delta_\Gamma \doteq \frac{\det(zI - \widetilde{M})}{\det(zI - \widetilde{P})} \cdot \gcd\{\alpha(\widetilde{e}_i) : 1 \leqslant i \leqslant |E| + 1\},$$

and so it suffices to show that $\gcd\{\alpha(\widetilde{e}_i) : 1 \leqslant i \leqslant |E| + 1\} = 1$ if $\mathrm{rank}(H) \geqslant 2$ and $\gcd\{\alpha(\widetilde{e}_i) : 1 \leqslant i \leqslant |E| + 1\} = z - 1$ if $H = \langle z \rangle$.

For this we recall the following: if e_i is any edge of $\mathcal{E}\backslash T$ and $\widehat{e}_i$ is its (nonbase) lift to $\widetilde{X}$ whose tail is contained in $\widetilde{T}'$, then $\alpha(\widehat{e}_i) = g_i - 1$, where g_i is represented by the loop in $e \cup T$ crossing e positively. Since $\widehat{e}_i$ is a translate of the base lift $\widetilde{e}_i$ we have that, up to a unit, $\alpha(\widetilde{e}_i) = g_i - 1$. Moreover, since the $\{g_i\}$ correspond to edges of $X^{(1)}$ outside of the maximal tree T, we note that $H = \langle g_1, \ldots, g_{|\mathcal{E}\backslash T|} \rangle$.

If $\mathrm{rank}(H) \geqslant 2$, then there are g_i and g_j that generate a rank 2 subgroup of H and hence $(g_i - 1)$ and $(g_j - 1)$ are relatively prime. In particular, $\gcd\{\alpha(\widetilde{e}_i)\} = \gcd\{(g_i - 1)\} = 1$, as required. If $\mathrm{rank}(H) = 1$, then $g_i = \bar{z}^{k_i}$ and since the g_i generate H, in fact $\gcd\{k_i\} = 1$. Hence, we also have that $\gcd\{(\bar{z}^{k_i} - 1)\} = (\bar{z} - 1)$ and so $\gcd\{\alpha(\widetilde{e}_i)\} = (z - 1)$. This completes the proof. $\qquad\square$

4.9. **Characteristic polynomials of monodromy.** The following corollary shows that specialization of the Alexander polynomial gives the characteristic polynomial of the monodromy's action on homology, up to replacing the indeterminate by its inverse. Similar statements appear in the literature where the characteristic polynomial considered is that of the deck transformation on the homology of the associated cyclic cover. See e.g. Milnor [Mil68, Assertion 4].

Recall that $f_*\colon H_1(G;\mathbb{Z}) \to H_1(G;\mathbb{Z})$ is the induced map on first homology of G.

Corollary 4.9. *Let X be the mapping torus of the graph map $f\colon G \to G$ and let u be its dual class. Then*
$$\det(t^{-1}I - f_*) \doteq (1 - t)^p \cdot \Delta_\Gamma^u(t),$$
where $p = 1$ if $\mathrm{rank}(H_1(X)) \geqslant 2$ and $p = 0$ if $\mathrm{rank}(H_1(X)) = 1$.

Proof. As before, let M represent the action of f on 1–chains $\mathbb{Z}^E$ and $P\colon \mathbb{Z}^V \to \mathbb{Z}^V$ denote the action of f on 0–chains.

Let $\partial\colon \mathbb{Z}^E \to \mathbb{Z}^V$ be the usual boundary map on the oriented edges, and let $\alpha\colon \mathbb{Z}^V \to \mathbb{Z}$ be the augmentation map defined by $\sum a_v v \mapsto \sum a_v$. Then there is a diagram with exact rows:

$$
\begin{array}{ccccccccccc}
0 & \longrightarrow & H_1(G) & \longrightarrow & \mathbb{Z}^E & \xrightarrow{\ \partial\ } & \mathbb{Z}^V & \xrightarrow{\ \alpha\ } & \mathbb{Z} & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle f_*} & & \downarrow{\scriptstyle M} & & \downarrow{\scriptstyle P} & & \downarrow{\scriptstyle 1} & & \\
0 & \longrightarrow & H_1(G) & \longrightarrow & \mathbb{Z}^E & \xrightarrow{\ \partial\ } & \mathbb{Z}^V & \xrightarrow{\ \epsilon\ } & \mathbb{Z} & \longrightarrow & 0
\end{array}
$$

From this, we have the equation of characteristic polynomials (in the variable t^{-1}),
$$(t^{-1} - 1) \cdot \det(t^{-1}I - M) = \det(t^{-1}I - f_*) \cdot \det(t^{-1}I - P).$$

The proof is completed by applying Theorem E and noting that the u–specializations of $\det(zI - \widetilde{M})$, $\det(zI - \widetilde{P})$, and Δ_Γ are $\det(t^{-1}I - M)$, $\det(t^{-1}I - P)$, and $\Delta_\Gamma^u(z)$, respectively (since $u(z) = -1$). $\qquad\square$

Remark 4.10. Corollary 4.9 has the following immediate consequence: If graph maps $f\colon G \to G$ and $f'\colon G' \to G'$ produce equivalent mapping tori, in the sense that there is an isomorphism $\phi\colon \pi_1(X_f) \to \pi_1(X_{f'})$ for which the dual classes satisfy $u' \circ \phi = u$, then the characteristic polynomials of f_* and f'_* are the same, up to a unit in $\mathbb{Z}[t^{\pm1}]$. In particular, the homological stretch factors of f and f' agree.

5. The folded mapping torus, cross sections, and the McMullen polynomial

Throughout this section, we fix a fully irreducible automorphism φ of $\mathbb{F}$ and an irreducible train track map $f\colon G \to G$ representing φ. As in §4.3, this determines mapping torus $X = X_f$ equipped with a semiflow $\psi_t\colon X \to X$, coming from the local upward flow on $G \times [0, 1]$, whose fundamental group Γ is the free-by-cyclic group determined by φ. As before, we also set $H = H_1(X)/\text{torsion}$. In this section, we briefly recall the main constructions and required results from [DKL15, DKL17b].

5.1. **The folded mapping torus and its cross sections.** To construct cross sections representing different splittings of Γ, Dowdall–Kapovich–Leininger [DKL15] constructed a modified mapping torus equipped with its own semiflow. Given a fixed factorization of the graph map $f\colon G \to G$ into a sequence of Stallings folds, the authors produce a 2-complex $X^\dagger = X_f^\dagger$ which they call the *folded mapping torus*. Just as for the mapping torus, $X^\dagger$ comes

equipped with a semiflow $\psi_t^\dagger \colon X^\dagger \to X^\dagger$ and the two semiflows are related by a natural flow-equivariant quotient map $q \colon X \to X^\dagger$ with the property that the set of vertex leaves of X is mapped bijectively to the set of vertex leaves of $X^\dagger$. Here, a *vertex leaf* is a connected component of the set of points whose forward orbit under the semiflow meets a vertex. The map q is a homotopy equivalence and we use it to once and for all identify the fundamental groups and homology groups of the 2-complexes.

Following [DKL17b, §5.1], a *cross section* of $X^\dagger$ is a finite embedded graph Θ that is transverse to the flow such that every flowline hits Θ infinitely often in the sense that $\{s \in \mathbb{R}_{\geqslant 0} \mid \psi_s^\dagger(x) \in \Theta\}$ is unbounded for every $x \in X^\dagger$. In this case Θ has continuous *first return map* $f_\Theta \colon \Theta \to \Theta$ that sends $x \in \Theta$ to the next point at which the flowline $\psi_s^\dagger(x)$ intersects Θ.

Remark 5.1. We caution that in general a cross section Θ need not be π_1–injective and the first return map f_Θ need not be a homotopy equivalence.

For our purposes, we will only consider cross sections that are *compatible* with $\psi^\dagger$, which means the intersection of Θ with the 1–skeleton of $X^\dagger$ is a finite set contained in the vertex leaves. In this case, Θ may be equipped with a finite *standard graph structure* [DKL17b, Definition 7.3] in which every edge lies in a 2–cell of $X^\dagger$ and has endpoints in vertex leaves. Then, the first return map f_Θ is a train track map with irreducible transition matrix A_Θ and spectral radius $\lambda(f_\Theta) > 1$ [DKL17b, Proposition 7.7]. According to [DKL15, Corollary 7.9], the train track map f_Θ satisfies the additional property that the Whitehead graphs at each of its vertices are connected (note that the proof of [DKL15, Corollary 7.9] does not use the standing assumption of [DKL15, Convention 7.6] that f is atoroidal).

For any cross section Θ of $X^\dagger$, the semiflow $\psi^\dagger$ can be reparameterized so that the return time to Θ is constant and equal to 1 (see [DKL17b, Definition 5.1]). Then, just as above, there is a natural flow preserving quotient map q_Θ from the mapping torus X_{f_Θ} of the first return map f_Θ onto $X^\dagger$ determined by

$$(5.2) \qquad \Theta \times [0,1] \to X^\dagger$$
$$(x,t) \mapsto \psi_t^\dagger(x).$$

Since the induced flow preserving map $q_\Theta \colon X_{f_\Theta} \to X^\dagger$ is a π_1–isomorphism, it is a homotopy equivalence. Compatibility of the cross section Θ implies that the set of vertex leaves of X_{f_Θ} is mapped bijectively to the set of vertex leaves of $X^\dagger$, just as for the original map q defined above. In what follows, all cross sections are assumed to be compatible.

For any cross section Θ, the homotopy equivalence $q_\Theta \colon X_{f_\Theta} \to X^\dagger$ identifies the fundamental groups and homology groups of the 2-complexes. Hence, any cross section determines a *dual cohomology class*

$$[\Theta] \in H^1(X_{f_\Theta}; \mathbb{Z}) \cong H^1(X^\dagger; \mathbb{Z}) = \mathrm{Hom}(\Gamma, \mathbb{Z}) = \mathrm{Hom}(H, \mathbb{Z}).$$

In fact, the classes dual to cross sections are precisely the primitive integral points of an open, rational cone $\mathcal{C}_\mathfrak{m}$ called the *cone of cross sections* (see Theorem 5.8). To explain, we must first recall the McMullen polynomial.

5.2. **The McMullen polynomial and the cone of cross sections.** We next describe a polynomial invariant of the dynamical system $(X^\dagger, \psi^\dagger)$ that was constructed in [DKL15, §4.1]. Continue to let $H = H_1(X^\dagger)/\text{torsion}$ and let $\widetilde{X}^\dagger \to X^\dagger$ be the free abelian cover with deck group H and with lifted semiflow $\widetilde{\psi}_t^\dagger \colon \widetilde{X}^\dagger \to \widetilde{X}^\dagger$.

Let us define a *transversal* of $\widetilde{\psi}^{\dagger}$ to be an arc τ contained in a 2–cell of $\widetilde{X}^{\dagger}$ that is transverse to the flowlines and has its endpoints in vertex leaves. The action of H on $\widetilde{X}^{\dagger}$ by deck transformations induces an action on transversals and makes the free $\mathbb{Z}$–module $\mathbb{F}(\widetilde{\psi}^{\dagger})$ on the set of transversals into a $\mathbb{Z}[H]$–module.

The *module of transversals* $T(\widetilde{\psi}^{\dagger})$ is the quotient of $\mathbb{F}(\psi^{\dagger})$ by the submodule generated by subdivision relations $\tau - \tau_1 - \tau_2$ for all transversals where $\tau = \tau_1 \cup \tau_2$ and $\tau_1 \cap \tau_2$ is a point in a vertex leaf, and flow relations $\tau - \tau'$ whenever a transversal τ flows homeomorphically on to τ'. Proposition 4.3 of [DKL17b] proves that $T(\widetilde{\psi}^{\dagger})$ is a finitely presented $\mathbb{Z}[H]$–module.

Definition 5.3. The *McMullen polynomial* $\mathfrak{m} \in \mathbb{Z}[H]$ of the system $(X^{\dagger}, \psi^{\dagger})$ is the 0^{th} Fitting invariant of $T(\widetilde{\psi}^{\dagger})$; that is, gcd of the 0^{th} Fitting ideal of the $\mathbb{Z}[H]$–module $T(\widetilde{\psi}^{\dagger})$.

Remark 5.4. We use the natural action of H on transversals given by $h \cdot \tau = h(\tau)$ for $h \colon \widetilde{X}^{\dagger} \to \widetilde{X}^{\dagger}$ a deck transformation. This differs from the convention used in [DKL17b], where H was declared to act on transversals by taking preimages. As a result, our McMullen polynomial $\mathfrak{m}$ differs from the polynomial $\mathfrak{m}_{\text{DKL}}$ introduced in [DKL17b] by the ring automorphism $\text{inv} \colon \mathbb{Z}[H] \to \mathbb{Z}[H]$ induced by the automorphism $h \mapsto h^{-1}$ of the abelian group H; that is, $\mathfrak{m} = \text{inv}(\mathfrak{m}_{\text{DKL}})$. The convention here is chosen to more closely parallel the construction of the Alexander polynomial, whereas the convention in [DKL17b] was used to parallel features of the Teichmüller polynomial in [McM00].

While Definition 5.3 is rather opaque, [DKL15, Theorem D] proves that $\mathfrak{m}$ may be explicitly calculated as a determinant with respect to any compatible cross section $\Theta \subset X^{\dagger}$. To explain this, let $\widetilde{\Theta} \subset \widetilde{X}^{\dagger}$ be the full preimage and let $\widetilde{\Theta}_0$ be a designated connected component. Choose also a lift $\widetilde{f}_{\Theta} \colon \widetilde{\Theta}_0 \to \widetilde{\Theta}_0$ of f_{Θ} to $\widetilde{\Theta}_0$. By using the flow preserving homotopy equivalence $q_{\Theta} \colon X_{f_{\Theta}} \to X^{\dagger}$ from Equation (5.2), we see that exactly as in §4.4, these choices determine a splitting $H \cong K_{\Theta} \oplus \langle z_{\Theta} \rangle$ of the deck group H as well as an identification $C_1(\widetilde{\Theta}_0) \cong \mathbb{Z}[K_{\Theta}]^{E\Theta}$ with respect to which the transition matrix of $\widetilde{f}_{\Theta}$ is expressed as an $E\Theta \times E\Theta$ matrix $\widetilde{A}_{\Theta}$ with entries in $\mathbb{Z}[K_{\Theta}]$ that specializes to A_{Θ} under the augmentation map $\mathbb{Z}[K] \to \mathbb{Z}$. With this notation, we have the following calculation of $\mathfrak{m}$:

Theorem 5.5 (Determinant formula [DKL17b, Theorem D]). *Let Θ be any connected, compatible cross section of $(X^{\dagger}, \psi^{\dagger})$ with associated splitting $H = K_{\Theta} \oplus \langle z_{\Theta} \rangle$ as above. Then the McMullen polynomial of $(X^{\dagger}, \psi^{\dagger})$ is*

$$\mathfrak{m} \doteq \det(z_{\Theta} I - \widetilde{A}_{\Theta}).$$

Remark 5.6. In the coordinates $\mathbb{Z}[H] \cong \mathbb{Z}[\mathbf{t}^{\pm}, x^{\pm}]$ used in [DKL17b], the determinant formula is written $\mathfrak{m}_{\text{DKL}} \doteq \det(xI - \widetilde{A}(\mathbf{t}))$, where x is the inverse of our element z_{Θ} and $\widetilde{A}(\mathbf{t})$ describes $\widetilde{f}_{\Theta}$ acting on $C_1(\widetilde{\Theta}_0)$ with respect to the module structure in which deck transformations act by taking preimages. Hence $\text{inv} \colon \mathbb{Z}[H] \to \mathbb{Z}[H]$ sends x to z_{Θ} and $\widetilde{A}(\mathbf{t})$ to $\widetilde{A}_{\Theta}$, so that we indeed have $\mathfrak{m} = \text{inv}(\mathfrak{m}_{\text{DKL}}) \doteq \text{inv}(\det(xI - A(\mathbf{t}))) = \det(z_{\Theta} I - \widetilde{A}_{\Theta})$.

Notice that the dual cohomology class $[\Theta] \colon H \to \mathbb{Z}$ sends z_{Θ} to -1 and K_{Θ} to 0; hence the induced ring map $\mathbb{Z}[H] \to \mathbb{Z}[\mathbb{Z}] \cong \mathbb{Z}[t^{\pm}]$ sends z_{Θ} to t^{-1} and $\widetilde{A}_{\Theta}$ to the $E\Theta \times E\Theta$ integer matrix A_{Θ}. Applying Theorem 5.5, we conclude that the specialization of $\mathfrak{m}$ at $[\Theta]$ (see §4.1) is simply the reciprocal characteristic polynomial of the transition matrix A_{Θ} of f_{Θ}:

$$(5.7) \qquad\qquad \mathfrak{m}^{[\Theta]}(t) \doteq \det(t^{-1} I - A_{\Theta}).$$

To summarize: the first return map f_Θ is a train track map with irreducible transition matrix A_Θ whose characteristic polynomial is $\mathfrak{m}^{[\Theta]}(t^{-1})$. Hence the geometric stretch factor $\lambda(A_\Theta) > 1$ of f_Θ is the reciprocal of the smallest root of $\mathfrak{m}^{[\Theta]}(t)$.

The construction picks out a distinguished term of the McMullen polynomial; namely the $z^{|E|}$ term of $\det(zI - \widetilde{A}_{G^\dagger})$ in the coordinates $H = K \oplus \langle z \rangle$ adapted to the base cross section $G^\dagger$. Writing $\mathfrak{m} = a_0 h_0 + \cdots + a_n h_n \in \mathbb{Z}[H]$ with $a_0 h_0$ this distinguished term, the polynomial determines a corresponding *cone of cross sections*

$$\mathcal{C}_\mathfrak{m} = \{u \in H^1(X^\dagger; \mathbb{R}) \mid u(h_0) < u(h_i) \text{ for each } i = 1, \ldots, n\}.$$

Notice that this is an open polyhedral cone with finitely many rationally defined sides. The significance of $\mathcal{C}_\mathfrak{m}$ (and a justification of its name) is captured by the following main result from [DKL17b]:

Theorem 5.8 ([DKL17b, Meta-Theorem I]). *An integral cohomology class $u \in H^1(X^\dagger; \mathbb{Z})$ is dual to a cross section if and only if $u \in \mathcal{C}_\mathfrak{m}$; this cross section may be chosen to be compatible with $\psi^\dagger$ and is connected if and only if u is primitive.*

Further, $\mathcal{C}_\mathfrak{m} \subset H^1(X^\dagger; \mathbb{R})$ is equal to the set of cohomology classes that are positive on every closed orbit of the semiflow $\psi^\dagger$, and also to the component of the BNS-invariant $\mathrm{BNS}(\Gamma)$ containing the dual class $\pi_1(G)_{f_*} \to \mathbb{Z}$ of the original HNN splitting of Γ.

5.3. McMullen polynomials via Perron polynomials.

In this section, we present an alternative characterization of the McMullen polynomial (Proposition 5.11) that will be essential in understanding the relation between the McMullen polynomial and the Alexander polynomial in the negatively orientable setting.

Let D be the directed graph with a vertex for each edge of G and a directed edge from e to e' for each time $f(e)$ crosses e' with either orientation. That is, D is the directed graph whose adjacency matrix is A, and so in particular D is strongly connected. Recalling the cell structure of X_f discussed in §4.3, there is an embedding $i \colon D \to X_f$ obtained by mapping each vertex of D to the midpoint of the corresponding horizontal edge of X_f and mapping the edges out of the e-vertex into the 2–cell σ_e in the obvious way. See Figure 2. The map $i \colon D \to X_f$ is an embedding and we often identify D with its image.
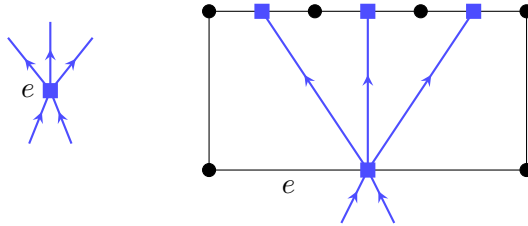

FIGURE 2. The image of D inside the 2-cell σ_e whose bottom edge is e. The original vertices of X_f are black dots and the midpoint vertices (i.e. the images of vertices of D) are blue squares.

Lemma 5.9. *The embedding $i \colon D \to X_f$ is π_1–surjective.*

Proof. Identify D with its image under i. Begin by subdividing the horizontal edges of X_f at the vertices of D and call these *midpoint vertices*. It suffices to show that if p is an edge path in $X_f^{(1)}$ that meets the midpoint vertices exactly at its endpoints then it is homotopic rel vertices into D.

First assume that p is horizontal, i.e. contained in $G \subset X_f$. Then p contains exactly one vertex v of G and two half edges with midpoints m_a and m_b; we denote these (directed) half edges by $m_a v$ and $v m_b$, respectively, so that $p = m_a v * v m_b$. Note that if m is the midpoint vertex of an edge e with both its endpoints at v, then the notation mv is ambiguous, but in what follows the proper meaning will be clear from context.

If the path p lies in the top of a 2-cell, then the claim is clear since p is homotopic rel endpoints to a (nondirected) path in D of length 2. See Figure 2. Otherwise, we consider the vertices m_a and m_b in the Whitehead graph $\mathrm{Wh}_G(v)$. Since f represents a fully irreducible automorphism, $\mathrm{Wh}_G(v)$ is connected and so there is a path $m_a, m_1, \ldots, m_k, m_b$ in $\mathrm{Wh}_G(v)$. By definition of the Whitehead graph, for each i, the corresponding path $m_i v m_{i+1}$ in G is contained in the image of an edge under f^n for some $n \geqslant 1$. We claim that these paths are each homotopic into D. If $n = 1$, then $m_i v m_{i+1}$ lies at the top of a 2-cell, so by the previous case, it is homotopic into D. Otherwise, the turn of G at v associated to $m_i v m_{i+1}$ is the image of the turn at a vertex w associated to a path $m_i' w m_{i+1}'$ in G, and $m_i' w m_{i+1}'$ is traversed by the image of an edge under f^{n-1}. By induction, $m_i' w m_{i+1}'$ is homotopic into D and if we denote by c_i and c_{i+1} the 2–cells whose bottom edges contain the midpoint vertices m_i' and m_{i+1}', respectively, we see that c_i and c_{i+1} share the vertical edge from w to v. Moreover, c_j contains a directed edge d_j of D from m_j' to m_j for each $j \in \{i, i+1\}$. Hence, $m_i v m_{i+1}$ is homotopic in $c_i \bigcup c_{i+1}$ to the path $\overline{d}_i * m_i' w m_{i+1}' * d_{i+1}$ and hence homotopic into D. As $p = m_a v m_b$ is homotopic to the concatenation $m_a v m_1 * m_1 v m_2 * \ldots * m_k v m_b$, this shows that every horizontal p is homotopic into D.

Next, suppose that p contains a single vertical edge e. If this path lies in the boundary of a single 2-cell, then it is again clear from Figure 2 that p is homotopic into D. Otherwise, p has the form $m_a v * e * w m_b$, where e is a vertical edge from v to w. Let σ be the 2-cell containing the vertical edge e and m_a along its bottom edge. Let m_c be the midpoint at the top of σ nearest e. Then p is homotopic to $m_a v * e * w m_c * m_c w m_b$, which we can handle by the previous cases. Finally, if p contains multiple vertical edges, then it homotopic to a concatenation of paths of this form $m_a v * e * w m_b$ and we finish as above. $\qquad\square$

For the embedding $i \colon D \to X_f$, we denote by i_* both the induced group homomorphism $H_1(D) \to H$ and ring homomorphism $\mathbb{Z}[H_1(D)] \to \mathbb{Z}[H]$.

Now let P_D be the Perron polynomial of D, defined as $P_D := \det(I - A)$ where A is the symbolic transition matrix of D given by

$$A_{xy} = \sum_{\partial e = y - x} e,$$

with coefficients in $\mathbb{Z}[C_1(D)]$. We refer the reader to [McM15] for additional details and recall that according to [McM15, Section 3] or [AKHR15, Theorem 2.14],

(5.10) $$P_D = 1 + \sum_{\sigma} (-1)^{|\sigma|} \sigma \in \mathbb{Z}[H_1(D)],$$

where the sum varies over nonempty oriented multi-cycles σ in D and $|\sigma|$ is the number of components of the multi-cycle.

The next proposition provides an explicit connection between polynomial invariants defined in [DKL17b] and [AKHR15]. The first claim is essentially proven by [DKL17b, Theorem 12.10], but we give a direct proof using a result from [LMT20].

Proposition 5.11. *With $i\colon D \to X_f$ as above, $\mathfrak{m} \doteq i_*(P_D)$. Moreover, $\mathfrak{m}' = i_*(P_D)$ is the unique normalization of $\mathfrak{m}$ such that $\mathrm{supp}(\mathfrak{m}')$ contains 1 and for some (any) $u \in \mathcal{C}_\mathfrak{m}$, u is positive on $\mathrm{supp}(\mathfrak{m}')\backslash\{1\}$.*

Proof. First, Lemma 5.9 implies that the homomorphism $i_*\colon H_1(D) \to H$ is surjective. By [LMT20, Proposition 4.2], $i_*(P_D)$ is equal to the polynomial P_{D,i_*} defined as follows: consider the preimage $\widetilde{D}$ of D in the universal free abelian cover $\widetilde{X}_f$ with deck group H. The polynomial $P_{D,i_*} \in \mathbb{Z}[H]$ is defined as $\det L$, where L is the endomorphism of the free $\mathbb{Z}[H]$-module generated by H–orbits of vertices of $\widetilde{D}$ given by $L(v) = v - (v_1 + \ldots + v_l)$. Here, the vertices $v_1, \ldots, v_l$ are exactly the endpoints (with multiplicity) of the directed edges of $\widetilde{D}$ out of v. Since D is the graph associated to the adjacency matrix A of f, we observe that $L = I - \overline{z}\widetilde{A}$. From Theorem 5.5, we conclude the required equality:

$$i_*(P_D) = P_{D,i_*} = \det(I - \overline{z}\widetilde{A}) = \overline{z}^k\mathfrak{m}.$$

For the moreover statement, we first claim that $\mathfrak{m}' = i_*(P_D)$ has the required properties. This follows from the fact that P_D has the properties that $1 \in \mathrm{supp}(P_D)$ by eq. (5.10) and each $a \in \mathrm{supp}(P_D)\backslash\{1\}$ has $u(i(a)) > 0$ for any $u \in \mathcal{C}_\mathfrak{m}$ by [DKL17b, Proposition 6.3]. Therefore, $1 \in \mathrm{supp}(i_*(P_D))$ and the claim is established.

Next suppose that $\mathfrak{m}''$ is another such normalization. Since $\mathfrak{m}'$ and $\mathfrak{m}''$ both contain 1 in their support and differ by a unit, $\mathfrak{m}' = h\,\mathfrak{m}''$ implies that $h \in \mathrm{supp}(\mathfrak{m}')$ and $h^{-1} \in \mathrm{supp}(\mathfrak{m}'')$. This must mean that $h = 1$ since any $u \in \mathcal{C}_\mathfrak{m}$ has to be positive on h and h^{-1}. $\qquad\square$

Proposition 5.12. *Let $\mathfrak{m}'$ be the normalization of $\mathfrak{m}$ from Proposition 5.11. Then support of $\mathfrak{m}'$ generates H as an abelian group.*

Proof. From Equation (5.10), we observe that the support of P_D contains (the homology classes of) all of the simple directed cycles of D and these generate $H_1(D)$ as an abelian group (see e.g. [LMT20, Lemma 5.8]) . It follows from Lemma 5.9 that the i–images of these directed cycles generate H and so $i_*(\mathrm{supp}(P_D))$ generates H.

By Proposition 5.11, $\mathrm{supp}(\mathfrak{m}')$ is equal to $\mathrm{supp}(i_*(P_D))$, which may be properly contained in $i_*(\mathrm{supp}(P_D))$. Nevertheless, we now show that the span of $\mathrm{supp}(i_*(P_D))$ is equal to the the span of $i_*(\mathrm{supp}(P_D))$. We will consider the class $u \in H^1(X_f)$ dual to the projection $X_f \to S^1$ as a 'norm'; it has the property that for each directed cycle y of D, $u(i(y)) > 0$.

Say $x \in \mathrm{supp}(P_D)$ *survives* if $i_*(x) \in \mathrm{supp}(\mathfrak{m}')$. Then we are left to show that the image of every simple cycle is a combination of the images of surviving simple cycles.

Suppose to the contrary that there is a simple directed cycle $x \in \mathrm{supp}(P_D)$ such that $i(x)$ is not in the span of the images of surviving simple directed cycles. Among all such cycles, let x be one such that $u(x)$ is minimal. Evidently x does not survive. Since x is simple and not surviving, there is a $y \in \mathrm{supp}(P_D)$ represented by a multi-cycle σ with an even number of components $|\sigma|$ such that $i_*(x) = i_*(y)$. Hence, y can be written as a disjoint union $d_1 \cup \ldots \cup d_k$ of simple directed cycles in D with $k \geqslant 2$. Then $i_*(x) = \prod i_*(d_j)$ and so $u(i_*(d_j)) < u(i_*(x))$ for each j (recalling that $u(i_*(d_j)) > 0$). But since x is u–minimal, we see that each $i_*(d_j)$ must be the span of surviving simple cycles. As $i_*(x)$ is a combination of the $i_*(d_j)$, we have a contradiction and the proof is complete. $\qquad\square$

5.4. **The vertex polynomial.** For the folded mapping torus $X^\dagger$, we define the *vertex cycles* of $X^\dagger$ to be the closed orbits of $\psi^\dagger$ that pass through vertices of $X^\dagger$. The vertex cycles form a finite collection of disjoint, embedded closed orbits.

Let $c_1, \ldots, c_m$ be the vertex cycles of X. The *vertex polynomial* $\mathfrak{p}$ of $X^\dagger$ is defined to be the element

$$\mathfrak{p} = \prod_{i=1}^{m} (1 - c_i) \in \mathbb{Z}[H],$$

where we have identified c_i with its class in H.

The vertex polynomial can also be explicitly computed as a determinant similar to the McMullen polynomial. For this, we recall the setup before Theorem 5.5: Fix a (compatible) cross section Θ of the semiflow $\psi^\dagger$ on $X^\dagger$. Let $\widetilde{\Theta} \subset \widetilde{X}^\dagger$ be the full preimage and let $\widetilde{\Theta}_0$ be a designated connected component. Choose also a lift $\widetilde{f}_\Theta \colon \widetilde{\Theta}_0 \to \widetilde{\Theta}_0$ of f_Θ to $\widetilde{\Theta}_0$. As before this lift determines a splitting $H \cong K_\Theta \oplus \langle z_\Theta \rangle$ of the deck group H. Letting $V\Theta$ denote the set of vertices of Θ, choosing base lifts of these to $\widetilde{\Theta}_0$ again gives an identification $C_0(\widetilde{\Theta}_0) \cong \mathbb{Z}[K_\Theta]^{V\Theta}$ with respect to which $\widetilde{f}_\Theta$'s action on vertices is given as an $V\Theta \times V\Theta$ matrix $\widetilde{P}_\Theta$ with entries in $\mathbb{Z}[K_\Theta]$.

Lemma 5.13. *With notation as above, $\mathfrak{p} \doteq \det(z_\Theta I - \widetilde{P}_\Theta)$.*

Proof. Let X_{f_Θ} be the folded mapping torus of the return map f_Θ and let $q_\Theta \colon X_{f_\Theta} \to X^\dagger$ be flow preserving homotopy equivalence from (5.2). Since Θ is compatible, q_Θ induces a bijection between the vertex cycles of the flows. Hence $\mathfrak{p} = \prod_{i=1}^m (1 - [\mathcal{V}_i'])$, where $\{\mathcal{V}_1', \ldots, \mathcal{V}_m'\}$ is the (disjoint) collection of closed orbits of X_{f_Θ} through vertices. From Lemma 4.2, we then see that $\mathfrak{p} \doteq \det(zI - \widetilde{P})$. Since q_Θ identifies the splitting $H = K \oplus \langle z \rangle$ (from §4.4) with the splitting $H = K_\Theta \oplus \langle z_\Theta \rangle$, this completes the proof. $\square$

6. Relating the polynomials and applications

In this section we prove our main theorems relating the McMullen and Alexander polynomials. This is done first in the positively orientable case in §6.1, then in the negatively orientable case in §6.2, and finally in the general case in §6.3.

For the entirety of this section we fix a fully irreducible automorphism $\varphi \colon \mathbb{F} \to \mathbb{F}$ represented by an irreducible train track $f \colon G \to G$. Let $X^\dagger$ be an associated folding mapping torus with fundamental group $\Gamma = \pi_1(X^\dagger)$, let $\mathfrak{m}, \Delta_\Gamma \in \mathbb{Z}[H]$ be its McMullen and Alexander polynomials, and $\mathcal{C}_\mathfrak{m} \subset H^1(X^\dagger; \mathbb{R})$ its cone of sections.

Since $\mathcal{C}_\mathfrak{m}$ is a component of $\mathrm{BNS}(\Gamma)$ (Theorem 5.8), we know that every integral class $u \in \mathcal{C}_\mathfrak{m}$ is dual to a splitting of Γ as a generalized HNN extension $B *_\phi$ over a finitely generated free group B (see §1.2). In this case, we say that ϕ is a (not necessarily unique) *monodromy* associated to u. Lemma 3.1 guarantees that the monodromy may be chosen to be injective and that the stretch factors ρ_ϕ, λ_ϕ and characteristic polynomial of ϕ_* acting on $H_1(B; \mathbb{Z})$ (up to a monomial factor) depend only on the class u and not on the chosen monodromy:

Definition 6.1. Given a primitive integral class $u \colon \Gamma \to \mathbb{Z}$ in $\mathcal{C}_\mathfrak{m}$, we define its homological and geometric *stretch factors*, respectively denoted $\rho(u)$ and $\lambda(u)$, to be the stretch factors ρ_ϕ and λ_ϕ of any monodromy ϕ associated to u. We call u *orientable* if $\lambda(u) = \rho(u)$; such classes are moreover either *positively* or *negatively orientable*, respectively, if $\lambda(u)$ or $-\lambda(u)$ is a root of the characteristic polynomial of an associated monodromy ϕ.

The next proposition justifies this terminology and demonstrates that the various ways a cohomology class can be orientable all agree.

Proposition 6.2. *The following are equivalent for primitive integral classes $u \in \mathcal{C}_\mathfrak{m}$:*

   (1) *u is orientable, i.e. $\lambda(u) = \rho(u)$,*

(2) *any injective monodromy φ_u associated to u has an orientable graph map representative,*

(3) *for any compatible cross section Θ_u dual to u, the associated first return map $f_u \colon \Theta_u \to \Theta_u$ is orientable.*

Moreover, positive (negative) orientability in one case implies the same in all cases.

Proof. If φ is atoroidal, then Γ is hyperbolic and φ_u is also fully irreducible by [DKL17a, Theorem 1.2] or [Mut21, Theorem 4.5]. Otherwise, φ is toroidal and represented by a pseudo-Anosov homeomorphism on a once-punctured surface [BH92, Theorem 4.1], in which case $\mathcal{C}_{\mathfrak{m}}$ is the cone over a fibered face of the Thurston norm ball and all associated monodromies, including φ_u, represent pseudo-Anosov homeomorphisms [Thu86] (see also [Mut21, Theorem 3.4]). In either case, Lemma 3.3 implies that φ_u admits a primitive train track representative F_u.

The first return map $f_u \colon \Theta_u \to \Theta_u$ is an irreducible train track map with connected Whitehead graphs. Hence, it is also primitive by [Mut20, Proposition 2.6]. By definition, the homomorphism on π_1 induced by f_u is a monodromy associated to u, although it may not be injective. Thus Lemma 3.1 implies $\lambda_{f_u} = \lambda(u) = \lambda_{F_u}$ and $\rho_{f_u} = \rho(u) = \rho_{F_u}$. The proposition now follows immediately from Theorem A. $\qquad\square$

Remark 6.3. As indicated in the proof above, if φ is atoroidal, then any injective monodromy φ_u associated to a class $u \in \mathcal{C}_{\mathfrak{m}}$ is also a fully irreducible free group endomorphism. Hence in this case Theorem B implies the conditions in Proposition 6.2 are also equivalent to the monodromy φ_u itself being (pos/neg) orientable.

6.1. **Positively orientable case.** In this subsection, we relate the two polynomial invariants of a positively orientable fully irreducible automorphism.

Theorem 6.4. *Suppose the fully irreducible automorphism φ is positively orientable. If $\operatorname{rank}(H_1(\Gamma)) \geqslant 2$, then the McMullen and Alexander polynomials are related by*

$$\mathfrak{m} \doteq \Delta_\Gamma \cdot \mathfrak{p},$$

Otherwise, the equation is $\mathfrak{m} \cdot (1 - z) \doteq \Delta_\Gamma \cdot \mathfrak{p}$, where z generates $H_1(\Gamma)/\text{torsion}$.

Proof. By Fact 2.2, for a positively orientable graph map the induced map on cellular 1-chains and the transition matrix are equal. Lifting to a map $\widetilde{f} \colon \widetilde{G}_0 \to \widetilde{G}_0$, the matrices used in the definition of $\mathfrak{m}$ and in Theorem E are equal, i.e. $\widetilde{A} = \widetilde{M}$. This, together with the characterizations of $\mathfrak{m}$ in Theorem 5.5 and $\mathfrak{p}$ in Lemma 5.13, completes the proof. $\qquad\square$

Theorem 6.5. *If some primitive integral class in $\mathcal{C}_{\mathfrak{m}}$ is positively orientable then so is every primitive integral class in $\mathcal{C}_{\mathfrak{m}}$.*

Proof. Let u be a primitive integral class in $\mathcal{C}_{\mathfrak{m}}$ represented by a cross section $\Theta = \Theta_u$ compatible with the induced semiflow (Theorem 5.8) and let $f_\Theta \colon \Theta \to \Theta$ denote its first return map. As in Equation (5.2), there is a flow preserving homotopy equivalence

$$q_\Theta \colon X_{f_\Theta} \to X^\dagger$$

sending vertex leaves to vertex leaves, which we use to identify $H_1(X_{f_\Theta})$ and $H_1(X^\dagger)$. Also, let $\widetilde{A}_\Theta, \widetilde{M}_\Theta, \widetilde{P}_\Theta$ be the $\mathbb{Z}[H]$-valued matrices appearing in the characterization of $\mathfrak{m}$ from Theorem 5.5 and the determinant formula for Δ_Γ (Theorem E), associated to f_Θ.

Suppose that the class u is positively orientable so that by Proposition 6.2 the map f_Θ is also positively orientable. This implies that $\widetilde{A}_\Theta = \widetilde{M}_\Theta$ just as in Theorem 6.4 and using

Theorem E and Lemma 5.13 we conclude that

$$\mathfrak{m} \doteq \Delta_\Gamma \cdot \mathfrak{p}.$$

Now let w be any other primitive integral class in $\mathcal{C}_\mathfrak{m}$. The above formula gives the equality of specializations: $\mathfrak{m}^w(t) \doteq \Delta_\Gamma^w(t) \cdot \mathfrak{p}^w(t)$. Moreover, from the definition of $\mathfrak{p}$ we see that all the zeroes of $\mathfrak{p}^w(t)$ are roots of unity. Hence, if $\lambda > 1$ is equal to $\lambda(w)$, then λ^{-1} is the smallest root of $\mathfrak{m}^w(t)$ (5.7) and hence of $\Delta_\Gamma^w(t)$. By Corollary 4.9, this gives that λ is the largest root of the characteristic polynomial of any monodromy associated to w. This implies that w is positively orientable and completes the proof. $\square$

6.2. **Negatively orientable case.** We now relate the polynomial invariants of a negatively orientable fully irreducible outer automorphism. Given that $\mathcal{C}_\mathfrak{m}$ contains a negatively orientable class, we first define an involution $\iota\colon \mathbb{Z}[H] \to \mathbb{Z}[H]$.

Let $u \in \mathcal{C}_\mathfrak{m}$ be a negatively orientable primitive integral class. If $\Theta = \Theta_u$ is a cross section dual to u, then its return map f_Θ is negatively orientable by Proposition 6.2. In this case, with notation as in the proof of Theorem 6.5, $\widetilde{A}_\Theta = -\widetilde{M}_\Theta$.

Define $\epsilon_u\colon H \to \{1, -1\}$ to be the homomorphism given by setting $\epsilon_u(h) = 1$ if $u(h)$ is even and $\epsilon_u(h) = -1$ if $u(h)$ is odd. In other words, ϵ_u is the reduction of u mod 2 (where $\mathbb{Z}/2 = \{1, -1\}$) and we often write $\epsilon_u = (-1)^u$. Define a ring automorphism $\iota_u\colon \mathbb{Z}[H] \to \mathbb{Z}[H]$ extending $h \mapsto \epsilon_u(h)h$. Note that $\iota_u\colon \mathbb{Z}[H] \to \mathbb{Z}[H]$ is an involution. We observe that

$$(6.6) \qquad \iota_u(\mathfrak{m}) \doteq \iota_u(\det(z_\Theta I - \widetilde{A}_\Theta))$$

$$= \det(-z_\Theta I - \widetilde{A}_\Theta))$$

$$\doteq \det(z_\Theta I - \widetilde{M}_\Theta)$$

$$\doteq \Delta_\Gamma \cdot \mathfrak{p}$$

A priori, however, the class ϵ_u depends on u. The following lemma states that each negatively orientable class in $\mathcal{C}_\mathfrak{m}$ determines the same class, which we call the *orientation class* $\epsilon\colon \Gamma \to \{-1, 1\}$. We define the associated ring involution $\iota\colon \mathbb{Z}[H] \to \mathbb{Z}[H]$ accordingly.

Proposition 6.7. *If u and ζ are negatively orientable classes in $\mathcal{C}_\mathfrak{m}$, then $\epsilon_u = \epsilon_\zeta$. Hence, if $\mathcal{C}_\mathfrak{m}$ contains a single negatively orientable class, then $\iota(\mathfrak{m}) = \Delta_\Gamma \cdot \mathfrak{p}$.*

Proof. By Equation (6.6), $\iota_u(\mathfrak{m}) \doteq \iota_\zeta(\mathfrak{m})$. We claim that if $\mathfrak{m}'$ is the normalization defined in Proposition 5.11, then $\iota_u(\mathfrak{m}') = \iota_\zeta(\mathfrak{m}')$ as elements of $\mathbb{Z}[H]$. Indeed, we know that $\mathfrak{m}' \doteq (\iota_u \circ \iota_\zeta)(\mathfrak{m}')$. But both sides of this equation satisfy the normalization in Proposition 5.11 (since ι_u and ι_ζ do not change the support) and so must be equal by uniqueness. This proves the claim.

Finally, if $\mathfrak{m}' = \sum a_g g$, then equating coefficients gives that $\epsilon_u(g) = \epsilon_\zeta(g)$ for each $g \in \mathrm{supp}(\mathfrak{m}')$. Since $\mathrm{supp}(\mathfrak{m}')$ generates H by Proposition 5.12, we see that $\epsilon_u = \epsilon_\zeta$. The second statement now follows from Equation (6.6). $\square$

Theorem 6.8. *Suppose that $\mathcal{C}_\mathfrak{m}$ contains a negatively orientable class u. Then for any other primitive integral class $\zeta \in \mathcal{C}_\mathfrak{m}$, the following are equivalent:*

(1) $\lambda(\zeta) = \rho(\zeta)$,
(2) ζ *is negatively orientable,*
(3) $\zeta = u \bmod 2$.

Note that if $\mathcal{C}_\mathfrak{m}$ is a cone of cross sections that contains a class u such that $\lambda(u) = \rho(u)$, then $\mathcal{C}_\mathfrak{m}$ is covered by either Theorem 6.5 or Theorem 6.8.

Proof. By Theorem 6.5, there are no positively orientable first return maps in the cone $\mathcal{C}_{\mathfrak{m}}$, so (1) and (2) are equivalent by definition.

Now suppose that $\zeta = u \bmod 2$. Then $\epsilon = (-1)^\zeta$. By Proposition 6.7, we know $\iota(\mathfrak{m}) \doteq \Delta_\Gamma \cdot \mathfrak{p}$, and so specializing to ζ gives

$$\iota(\mathfrak{m})^\zeta(t) \doteq \Delta_\Gamma^\zeta(t) \cdot \mathfrak{p}^\zeta(t),$$

where $p^\zeta(t)$ is a product of cyclotomic polynomials and $\Delta_\Gamma^\zeta(t)$ is, up to a factor of $t^k(t-1)$, the characteristic polynomial of any monodromy associated to ζ (see Corollary 4.9).

We claim that $\iota(\mathfrak{m})^\zeta(t) = \mathfrak{m}^\zeta(-t)$. Indeed, if $\mathfrak{m} = \sum a_g g \in \mathbb{Z}[H]$, then

$$\begin{aligned}
\iota(\mathfrak{m})^\zeta(t) &= \sum \epsilon(g) a_g t^{\zeta(g)} \\
&= \sum (-1)^{\zeta(g)} a_g t^{\zeta(g)} \\
&= \sum a_g (-t)^{\zeta(g)} \\
&= \mathfrak{m}^\zeta(-t).
\end{aligned}$$

Using equation (5.7) and Corollary 4.9, we conclude as in the proof of Theorem 6.5 that ζ is negatively orientable and hence (3) implies (2). Since (2) $\implies$ (3) follows from Proposition 6.7, this completes the proof. $\qquad\square$

With these facts in hand, we can prove Theorem C from the introduction.

Proof of Theorem C. By Theorem 5.8, the component $\mathcal{C}$ of BNS(Γ) containing the dual class of $\mathbb{F} *_\varphi$ is equal to the cone of cross sections $\mathcal{C}_{\mathfrak{m}}$. If some primitive integral $u \in \mathcal{C}$ is positively orientable, then so is every such class in $\mathcal{C}$ (Theorem 6.5) and so $\lambda(u) = \rho(u)$ for the entire cone. This is case (1) from the theorem statement.

If some primitive integral $u \in \mathcal{C}$ is negatively orientable, then Theorem 6.8 gives that the primitive integral $\zeta \in \mathcal{C}$ with $\lambda(\zeta) = \rho(\zeta)$ are exactly those that equal $u \bmod 2$, thus giving case (2). The final alternative is that no classes in $\mathcal{C}$ are orientable and this is exactly case (3). $\qquad\square$

6.3. **An equation mod** 2.

Here we show that the conclusion of Theorem 6.4 holds in total generality after reducing mod 2.

Theorem 6.9. *Let φ be a fully irreducible automorphism, with associated free-by-cyclic group Γ. If* $\mathrm{rank}(H_1(\Gamma)) \geqslant 2$, *then*

$$\mathfrak{m} \doteq \Delta_\Gamma \cdot \mathfrak{p} \pmod 2$$

Otherwise, the equation holds after multiplying $\mathfrak{m}$ by $(z-1)$, where z generates $H_1(\Gamma)/\mathrm{torsion}$.

We begin by noting that the proof of Theorem E holds with coefficients in $\mathbb{Z}/2$; that is, using the group ring $\mathbb{Z}/2[H]$. In short, if we denote the first Fitting invariant of $H_1(\widetilde{X}, \widetilde{T}; \mathbb{Z}/2)$ by Δ_2, then

$$\Delta_2 \doteq (z-1)^l \cdot \frac{\det(zI - \widetilde{M_2})}{\det(zI - \widetilde{P_2})},$$

where $l = 0$ if $\mathrm{rank}(H_1(X)) \geqslant 2$ and $l = 1$ otherwise. Here $\widetilde{M_2}$ and $\widetilde{P_2}$ are the mod 2 reductions of $\widetilde{M}$ and $\widetilde{P}$, respectively, from §4.4. More formally, if we denote by $r \colon \mathbb{Z}[H] \to \mathbb{Z}/2[H]$ the mod 2 reduction homomorphism and use the same notation for the corresponding ring homomorphism between matrix rings, then $\widetilde{M_2} = r(\widetilde{M})$ and $\widetilde{P_2} = r(\widetilde{P})$.

If we denote the adjacency matrix for $\tilde{f}$ (from §5.2) by $\tilde{A}$ and similarly set $\tilde{A}_2 = r(A)$, then $\tilde{A}_2 = \widetilde{M}_2$ as matrices with coefficients in $\mathbb{Z}/2[H]$. Then note that

$$r(\mathfrak{m}) \doteq r(\det(zI - \tilde{A})) = \det(zI - \tilde{A}_2) = \det(zI - \widetilde{M}_2),$$

and similarly $r(\mathfrak{p}) = \det(zI - \tilde{P}_2)$. We conclude that

$$(z-1)^l \cdot r(\mathfrak{m}) \doteq \Delta_2 \cdot r(\mathfrak{p}).$$

It remains to prove that $r(\Delta) \doteq \Delta_2$ in $\mathbb{Z}/2[H]$, where $\Delta = \Delta_\Gamma$. First, note that by the universal coefficients theorem and the fact that $H_0(\tilde{X}, \tilde{T}) = 0$,

$$H_1(\tilde{X}, \tilde{T}; \mathbb{Z}/2) = H_1(\tilde{X}, \tilde{T}) \otimes \mathbb{Z}/2$$
$$= H_1(\tilde{X}, \tilde{T}) \otimes_{\mathbb{Z}[H]} \mathbb{Z}/2[H].$$

Then by the general theory of Fitting ideals (e.g. [Eis95, Corollary 20.5]),

$$\mathrm{Fit}_1(H_1(\tilde{X}, \tilde{T}) \otimes_{\mathbb{Z}[H]} \mathbb{Z}/2[H]) = r(\mathrm{Fit}_1(H_1(\tilde{X}, \tilde{T}))).$$

To conclude we need the following corollary to the proof of Theorem E. It is similar to [McM02, Theorem 5.1] in the 3-manifold setting.

Corollary 6.10. *Let $\mathcal{A} \subset \mathbb{Z}[H]$ be the augmentation ideal. Then the Alexander ideal equals*

$$\mathrm{Fit}_1(H_1(\tilde{X}, \tilde{T})) = (\Delta) \cdot \mathcal{A}^p,$$

where (Δ) is the ideal generated by Δ, and where $p = 1$ if $\mathrm{rank}(H_1(X)) \geqslant 2$ and $p = 0$ if $\mathrm{rank}(H_1(X)) = 1$.

Proof. Using notation from Theorem E, we have $\mathrm{Fit}_1(H_1(\tilde{X}, \tilde{T})) = (\Delta_1, \ldots, \Delta_{|E|+1})$. Fix some index, say k, representing a vertical edge η_k of $\mathcal{E} \backslash T$ so that $\tilde{e}_k = \tilde{\eta}_k$. Then from the proof of Theorem E we have

$$(z-1)^{1-p}\Delta_k = \alpha(\tilde{\eta}_k)\Delta_\Gamma.$$

Now using Lemma 4.7 to relate Δ_i and Δ_k, for every index $i \in \{1, \ldots, |E|+1\}$ we get

$$(z-1)^{1-p}\Delta_i = \alpha(\tilde{e}_i)\Delta_\Gamma.$$

Since the image of α is the augmentation ideal, for $\mathrm{rank}(H_1(X)) \geqslant 2$ and $1 - p = 0$ we get

$$(\Delta_1, \ldots, \Delta_{|E|+1}) = (\Delta_\Gamma) \cdot \mathcal{A}.$$

For $\mathrm{rank}(H_1(X)) = 1$ and $1 - p = 1$, we showed in the proof of Theorem E that the gcd of $\alpha(\tilde{e}_i)$ is equal to $(z-1)$. Since $\mathbb{Z}[z^{\pm 1}]$ is a PID, we have $\mathcal{A} = ((z-1))$ and thus

$$(\Delta_1, \ldots, \Delta_{|E|+1}) = (\Delta_\Gamma). \qquad \square$$

Hence, we see that

$$\mathrm{Fit}_1(H_1(\tilde{X}, \tilde{T}; \mathbb{Z}/2)) = r(\Delta) \cdot r(\mathcal{A}^p),$$

and so $\Delta_2 = r(\Delta) \cdot \gcd r(\mathcal{A}^p)$.

If $p = 0$, then clearly $\Delta_2 = r(\Delta)$. Otherwise, $p = 1$ and $r(\mathcal{A})$ is the ideal in $\mathbb{Z}/2[H]$ generated by $1 - g_i$, where recall the elements g_i from the proof of Theorem E. In this case, however, there are g_j and g_k that generate a rank 2 subgroup of H and so $(1 - g_j)$ and $(1 - g_k)$ are still relatively prime in $\mathbb{Z}/2[H]$. We conclude that $\gcd r(\mathcal{A}) = 1$. This shows that $\Delta_2 = r(\Delta)$ and completes the proof that

$$(z-1)^l \cdot r(\mathfrak{m}) \doteq r(\Delta) \cdot r(\mathfrak{p}).$$

6.4. **Newton polytopes and the cone of sections.** We conclude with an observation, which will be useful in §7 below, that in the orientable case the cone of sections can be computed directly from the Alexander polynomial. Recall that the *Newton polytope* of an element $\mathfrak{q} \in \mathbb{Z}[H]$ is the convex hull $N(\mathfrak{q}) \subset H_1(X; \mathbb{R})$ of the elements $h \in H$ appearing with nonzero coefficient in $\mathfrak{q}$ (see [McM00, Appendix]). The *dual cone* of a vertex $v \in N(\mathfrak{q})$ is by definition the set of cohomology classes $u \in H^1(X; \mathbb{R})$ that achieve a maximum value on $N(\mathfrak{q})$ precisely at the vertex v.

Letting $\mathrm{inv} \colon \mathbb{Z}[H] \to \mathbb{Z}[H]$ be the homomorphism sending h to h^{-1}, Theorem 5.8 says that the cone of cross sections $\mathcal{C}_{\mathfrak{m}}$ of the folded mapping torus $X^{\dagger}$ (which is also a component of $\mathrm{BNS}(\Gamma)$) is the dual cone of a vertex of $N(\mathrm{inv}(\mathfrak{m}))$. When cone of sections contains an orientable class, the same holds for the Alexander polynomial:

Lemma 6.11. *Suppose the cone $\mathcal{C}_{\mathfrak{m}}$ contains an orientable primitive class, that is a class with $\lambda(u) = \rho(u)$. Then $\mathcal{C}_{\mathfrak{m}}$ is equal to the dual cone of a vertex of $N(\mathrm{inv}(\Delta_{\Gamma}))$ (namely the unique vertex whose dual cone contains u).*

Proof. Let $\iota \colon \mathbb{Z}[H] \to \mathbb{Z}[H]$ be the involution defined in §6.2 in case that u is negatively orientable, and let ι denote the identity if u is positively orientable. By Theorem 6.4 and Proposition 6.7, we thus have $\mathfrak{m} \doteq \iota(\Delta_{\Gamma}) \cdot \iota(\mathfrak{p})$. By Theorem 6.5 and Theorem 6.8 we also have $\lambda(\zeta) = \rho(\zeta)$ for every primitive integral $\zeta \in \mathcal{C}_{\mathfrak{m}}$ that agrees with u mod 2. Thus Corollary 4.9 says $\lambda(\zeta)$ is the reciprocal of the smallest root of the specialization $\iota(\Delta_{\Gamma})^{\zeta}(t)$ or, equivalently, the largest root of the specialization of $\mathrm{inv}(\iota(\Delta_{\Gamma}))$ at ζ (see Remark 5.6). By [DKL17b], $\log(\lambda(\zeta))$ tends to infinity as ζ tends to the boundary of $\mathcal{C}_{\mathfrak{m}}$. Now by using [McM00, Theorem A.1], we conclude that $\mathcal{C}_{\mathfrak{m}}$ is equal to the dual cone of a vertex of $N(\mathrm{inv}(\iota(\Delta_{\Gamma}))) = N(\mathrm{inv}(\Delta_{\Gamma}))$. $\square$

7. Examples

In this section we give several fully irreducible automorphisms φ that illustrate aspects of the theory developed in the paper. These also serve to contrast the situation for surface homeomorphisms and to highlight that the four stretch factors λ_{φ}, ρ_{φ}, $\lambda_{\varphi^{-1}}$, $\rho_{\varphi^{-1}}$ of an automorphism and its inverse are in general unrelated to each other. To this end, we say φ has *stretch factor symmetry* if $\lambda_{\varphi} = \lambda_{\varphi^{-1}}$.

Recall that a fully irreducible automorphism φ of a free group is *geometric* if it may be represented by a pseudo-Anosov homeomorphism on a (punctured) surface. In this case, the pseudo-Anosov property automatically implies $\lambda_{\varphi} = \lambda_{\varphi^{-1}}$, and one additionally has $\rho_{\varphi} = \rho_{\varphi^{-1}}$ provided the surface itself is orientable. These basic equalities can be deduced from the stronger facts that the Teichmüller [McM00] and Alexander [Bla57, Tur75] polynomials of the fibered 3-manifold determined by φ are symmetric. In addition, if φ is orientable then $\lambda_{\varphi} = \rho_{\varphi}$ and $\lambda_{\varphi^{-1}} = \rho_{\varphi^{-1}}$.

The fully irreducible automorphism φ is known to be geometric if and only if both its attracting and repelling trees are geometric [Gui05, HM07]; see also [CH12]. The automorphism is *parageometric* when the attracting tree is geometric but the repelling tree is not, and in this case Handel–Mosher have shown $\lambda_{\varphi} \neq \lambda_{\varphi^{-1}}$ [HM07].

The examples below show that orientable fully irreducible automorphisms need not exhibit the stretch factor symmetry enjoyed by pseudo-Anosovs, and that orientability of φ and φ^{-1} are in general independent. Throughout, we use Thierry Coulbois's train track Sage package [Cou] to check that the example graph maps represent fully irreducible automorphisms and to compute their stretch factors. The package can also be used to find periodic Nielsen paths in a stable train track representative of a fully irreducible automorphism φ. We then use

[BF94, Theorem 3.2] to determine if the attracting tree T_φ^+ is geometric or nongeometric: T_φ^+ is geometric if and only if a stable train track representative contains an indivisible orbit of periodic Nielsen paths.

Example 7.1. Here we give an example of a parageometric fully irreducible automorphism that is (positively) orientable but whose inverse is not orientable.

Let φ be the fully irreducible automorphism of $F_3 = \langle a, b, c \rangle$ given by

$$a \overset{\varphi}{\to} abbc, \qquad b \overset{\varphi}{\to} bcabbc, \qquad c \overset{\varphi}{\to} CBcabbc,$$

where throughout we use capital letters to denote inverse elements or reversed edges. This is represented by a train track graph map $g \colon G \to G$, where G is the graph with two vertices $0, 1$ and four edges $0 \overset{a}{\to} 0$, $0 \overset{b}{\to} 1$, $1 \overset{c}{\to} 0$, and $1 \overset{d}{\to} 0$, where the marking is $a \leftrightarrow a$, $b \leftrightarrow bd$, $c \leftrightarrow Dc$, and where the map g is defined on edges by

$$g(a) = abdbc, \qquad g(b) = bcab, \qquad g(c) = cabdbc, \qquad g(d) = dbc.$$

This graph map is clearly positively orientable, and thus $\lambda_\varphi = \rho_\varphi \sim 4.61$ both equal the largest root of the characteristic polynomial $1 - 3t + 7t^2 - 6t^3 + t^4$. The inverse automorphism φ^{-1} is represented by a train track map g' on a graph G' with two vertices $0, 1$ and four edges $0 \overset{a}{\to} 1$, $0 \overset{b}{\to} 1$, $1 \overset{c}{\to} 0$, $0 \overset{f}{\to} 1$. Here the marking is $a \leftrightarrow aF$, $b \leftrightarrow bF$, $c \leftrightarrow fc$ and the map g' is given by

$$g'(a) = acf, \qquad g'(b) = bFCFa, \qquad g'(c) = AfcfA, \qquad g'(f) = b.$$

From this one easily calculates that $\lambda_{\varphi^{-1}} \sim 3.08$ but that $\rho_{\varphi^{-1}} \sim 2.15$. According to Theorem B, neither the graph map g' nor automorphism φ^{-1} is orientable.

Example 7.2. Our next example shows that even if a fully irreducible outer automorphism and its inverse are both (negatively) orientable, they need not be geometric. In fact, here φ^{-1} is parageometric with $\lambda_\varphi \sim 2.17$ and $\lambda_{\varphi^{-1}} \sim 3.72$. Note that squaring provides a positively orientable pair φ^2, φ^{-2} that also fails to be geometric.

For $F_3 = \langle a, b, c \rangle$, let φ be given by $a \to BA$, $b \to CAA$, $c \to B$. This map is in fact a negatively orientable train track map on the 3-petal rose labeled a, b, c. The automorphism φ^{-1} is given by $a \to Ac$, $B \to c$, $C \to bAcAc$, which is also a negatively orientable train track representative. The characteristic polynomials of φ and φ^{-1} on homology are $-1 + 3t + t^2 - t^3$ and $-1 + 3t + 3t^2 - t^3$, respectively.

Example 7.3. Our final example provides fully irreducible outer automorphisms φ and φ^{-1} that are negatively orientable and satisfy stretch factor symmetry. That is, $\lambda_\varphi = \lambda_{\varphi^{-1}} = \rho_\varphi = \rho_{\varphi^{-1}} \sim 3.73$ and yet φ is not a geometric automorphism. Indeed, the attracting and repelling trees associated to φ are nongeometric and the McMullen polynomial of the free-by-cyclic group Γ associated to φ is not symmetric.

We also use this example to detail the computation of the Alexander and McMullen polynomials using Theorem E and Proposition 6.7, and to illustrate how orientability varies in the cone of sections in keeping with Theorem 6.8.

Let G be the graph in Figure 3 with two vertices v, w, and with four oriented edges a, c, d, f from w to v and two oriented edges b, e from v to w. Let $g \colon G \to G$ be the, clearly negatively orientable, graph map defined on edges as

$$a \mapsto DEABF, \ \ b \mapsto EDE, \ \ c \mapsto ABF, \ \ d \mapsto CEDBA, \ \ e \mapsto BAB, \ \ f \mapsto CED.$$
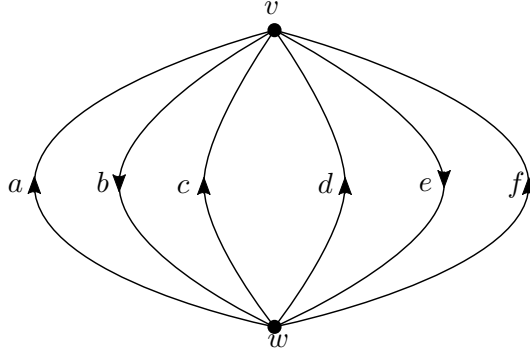
FIGURE 3. Graph G

Choosing f as the maximal subtree gives a basis aF, fb, cF, dF, fe of $F_5 = \pi_1(G, w)$; denoting these as a, b, c, d, e for brevity, the corresponding automorphism is

$$a \overset{\varphi}{\mapsto} DEABdec, \; b \overset{\varphi}{\mapsto} CEDEDE, \; c \overset{\varphi}{\mapsto} ABdec, \; d \overset{\varphi}{\mapsto} CEDBAdec, \; e \overset{\varphi}{\mapsto} CEDBAB$$

Let X be the mapping torus of g and H its homology. Recall from §4.4 that K is the image of $\pi_1(G) \to \pi_1(X) \to H$. A direct computation using the standard group presentation of $\Gamma = F_5 \rtimes_\varphi \mathbb{Z}$ gives $[aF] = [bf] = [cF] \neq 0$ and $[dF] = [ef] = 0$ in H and hence in K. Denoting the non-trivial homology class by α, the free abelian cover $\widetilde{G}_0 \to G$ with deck group K is therefore as depicted in Figure 4. Fix base lifts $\tilde{v}, \tilde{w}, \tilde{a}, \tilde{b}, \ldots$ of the cells of G as indicated in the figure, and let $\tilde{g} \colon \widetilde{G}_0 \to \widetilde{G}_0$ to be the unique lift of g such that $\tilde{g}(\tilde{v}) = \tilde{w}$. This determines a splitting $H = \langle \alpha \rangle \oplus \langle z \rangle$, as described in §4.4, where z maps to -1 and α to 0 under the map $\pi_1(X) \to \mathbb{Z}$.
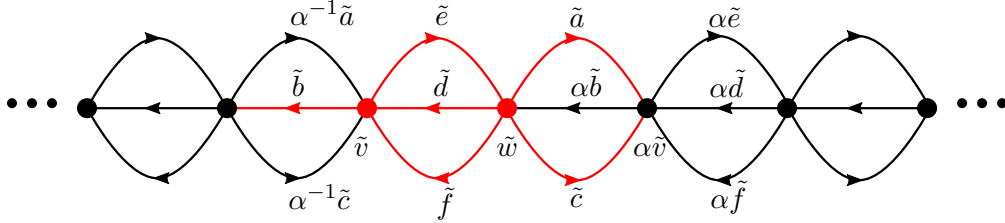
FIGURE 4. The free abelian cover $\widetilde{G}_0$

We now compute the $\mathbb{Z}[K]$–matrices $\widetilde{P}$ and $\widetilde{M}$ describing the action of $\tilde{g}$ on $C_1(\widetilde{G}_0)$. We already know $\tilde{g}(\tilde{v}) = \tilde{w}$. To find $\tilde{g}(\tilde{w})$ we look at the image of the edge $\tilde{e}$ with end points $\tilde{v}$ and $\tilde{w}$. Since $\tilde{g}(\tilde{v}) = \tilde{w}$, $\tilde{g}(\tilde{e})$ starts at $\tilde{w}$, traverses the edges $\alpha\tilde{B}, \tilde{A}, \alpha\tilde{B}$ in that order and ends at the vertex $\alpha\tilde{v}$. Therefore, $\tilde{g}(\tilde{w}) = \alpha\tilde{v}$. The images $\tilde{g}(\tilde{a}), \tilde{g}(\tilde{b}), \ldots$ of the other base edges are found in the same manner. We thus compute:

$$\widetilde{P} = \begin{pmatrix} 0 & \alpha \\ 1 & 0 \end{pmatrix}, \quad \text{and} \quad \widetilde{M} = \begin{pmatrix} -1 & 0 & -1 & -1 & -1 & 0 \\ -\alpha & 0 & -\alpha & -\alpha & -2\alpha & 0 \\ 0 & 0 & 0 & -1 & 0 & -1 \\ -\alpha & -1 & 0 & -1 & 0 & -1 \\ -\alpha & -2 & 0 & -1 & 0 & -1 \\ -\alpha & 0 & -\alpha & 0 & 0 & 0 \end{pmatrix}.$$
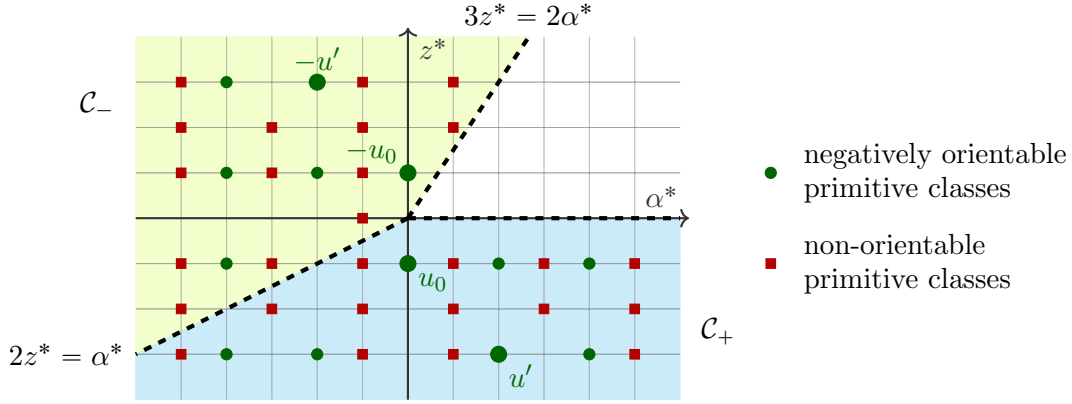
FIGURE 5. Cones of cross sections $\mathcal{C}_+$ and $\mathcal{C}_-$ for φ and φ^{-1}.

From this we can easily compute the polynomial invariants associated to φ. By the determinant formula in Theorem E, the Alexander polynomial is

$$\Delta_\Gamma \doteq \frac{\det(zI - \widetilde{M})}{\det(zI - \widetilde{P})} = \frac{z^6 + 2z^5 + z^4(1 - 8\alpha) + z^2(8\alpha^2 - \alpha) - 2\alpha^2 z - \alpha^3}{z^2 - \alpha}$$

$$= z^4 + 2z^3 + z^2(1 - 7\alpha) + 2\alpha z + \alpha^2.$$

The McMullen polynomial satisfies $i(\mathfrak{m}) \doteq \Delta_\Gamma \cdot \mathfrak{p}$, where $\epsilon \colon H = \langle \alpha \rangle \oplus \langle z \rangle \to \mathbb{Z}/2\mathbb{Z}$ is the map $\epsilon(\alpha) = 1$ and $\epsilon(z) = -1$ and $i \colon \mathbb{Z}[H] \to \mathbb{Z}[H]$ is the involution given by $i(h) = \epsilon(h)h$. Thus by Proposition 6.7 we compute that

$$\mathfrak{m} \doteq i(\Delta_\Gamma \cdot \mathfrak{p}) \doteq z^6 - 2z^5 + z^4(1 - 8\alpha) + z^2(8\alpha^2 - \alpha) + 2\alpha^2 z - \alpha^3.$$

One may check that the inverse φ^{-1} is represented by the train track map

$$a \mapsto DIFEHAB, \quad b \mapsto EDIFE, \quad c \mapsto AB, \quad d \mapsto GBIA, \quad e \mapsto BIABI,$$
$$f \mapsto CED, \qquad g \mapsto H, \qquad h \mapsto IFG, \quad i \mapsto IF$$

on the graph with 4 vertices and 9 edges $4 \xrightarrow{a} 1$, $2 \xrightarrow{b} 4$, $0 \xrightarrow{c} 1$, $2 \xrightarrow{d} 3$, $3 \xrightarrow{e} 0$, $0 \xrightarrow{f} 1$, $4 \xrightarrow{g} 0$, $1 \xrightarrow{h} 3$, $1 \xrightarrow{i} 2$. This map is clearly negatively orientable and may be used to confirm the stretch factor symmetry $\lambda_{\varphi^{-1}} = \lambda_\varphi$.

In cohomology $H^1(\Gamma) = \langle \alpha^* \rangle \oplus \langle z^* \rangle$, the splittings $F_5 *_\varphi \cong \Gamma \cong F_5 *_{\varphi^{-1}}$ correspond to the classes $u_0 = (0, -1)$ and $-u_0 = (0, 1)$. According to Lemma 6.11 the two components of $\mathrm{BNS}(\Gamma)$ containing $\pm u_0$ are both dual cones of vertices of the Newton polytope of $\mathrm{inv}(\Delta_\Gamma)$, which is the convex hull of the homology classes $(0, -4)$, $(0, -3)$, $(0, -2)$, $(-1, -2)$, $(-1, -1)$, $(-2, 0)$. These cones $\mathcal{C}_\pm$ for $\varphi^\pm$ are thus as depicted in Figure 5. According to Theorem 6.8 every primitive class u in $\mathcal{C}_+$ or $\mathcal{C}_-$ is either negatively orientable or non-orientable depending on whether u agrees with u_0 mod 2. Notice that while the pair $\pm u_0$ has stretch factor symmetry, this need not hold for other classes in the cones. For example, the classes $u' = (2, -3)$ and $-u' = (-2, 3)$ are negatively orientable so, by Corollary 4.9, their stretch factors are the reciprocals of the smallest roots of their respective specializations of Δ_Γ. Hence one calculates that $\lambda(u') = \rho(u') \sim 1.43092$ but that $\lambda(-u') = \rho(-u') \sim 1.36225$.

References

[AKHR15] Yael Algom-Kfir, Eriko Hironaka, and Kasra Rafi. Digraphs and cycle polynomials for free-by-cyclic groups. *Geom. Topol.*, 19(2):1111–1154, 2015.

[Ale28] James W Alexander. Topological invariants of knots and links. *Transactions of the American Mathematical Society*, 30(2):275–306, 1928.

[BB07] Gavin Band and Philip Boyland. The Burau estimate for the entropy of a braid. *Algebr. Geom. Topol.*, 7:1345–1378, 2007.

[BF94] Mladen Bestvina and Mark Feighn. Outer limits. 1994. preprint, available at http://andromeda.rutgers.edu/~feighn/papers/outer.pdf.

[BFH97] Mladen Bestvina, Mark Feighn, and Michael Handel. Laminations, trees, and irreducible automorphisms of free groups. *Geom. Funct. Anal.*, 7(2):215–244, 1997.

[BH92] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Ann. of Math. (2)*, 135(1):1–51, 1992.

[BH95] M. Bestvina and M. Handel. Train-tracks for surface homeomorphisms. *Topology*, 34(1):109–140, 1995.

[Bla57] Richard C Blanchfield. Intersection theory of manifolds with operators with applications to knot theory. *Annals of Mathematics*, pages 340–356, 1957.

[BNS87] Robert Bieri, Walter D. Neumann, and Ralph Strebel. A geometric invariant of discrete groups. *Invent. Math.*, 90(3):451–477, 1987.

[BR91] Richard A. Brualdi and Herbert J. Ryser. *Combinatorial matrix theory*, volume 39 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1991.

[But07] Jack O. Button. Mapping tori with first Betti number at least two. *J. Math. Soc. Japan*, 59(2):351–370, 2007.

[CH12] Thierry Coulbois and Arnaud Hilion. Botany of irreducible automorphisms of free groups. *Pacific J. Math*, 256(2):291–307, 2012.

[Cou] Thierry Coulbois. Train track package. https://www.i2m.univ-amu.fr/perso/thierry.coulbois/train-track/.

[DKL15] Spencer Dowdall, Ilya Kapovich, and Christopher J. Leininger. Dynamics on free-by-cyclic groups. *Geom. Topol.*, 19(5):2801–2899, 2015.

[DKL17a] Spencer Dowdall, Ilya Kapovich, and Christopher J. Leininger. Endomorphisms, train track maps, and fully irreducible monodromies. *Groups Geom. Dyn.*, 11(4):1179–1200, 2017.

[DKL17b] Spencer Dowdall, Ilya Kapovich, and Christopher J. Leininger. McMullen polynomials and Lipschitz flows for free-by-cyclic groups. *J. Eur. Math. Soc. (JEMS)*, 19(11):3253–3353, 2017.

[Eis95] David Eisenbud. *Commutative algebra*, volume 150 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995. With a view toward algebraic geometry.

[Gan59] F. R. Gantmacher. *The theory of matrices. Vols. 1, 2.* Chelsea Publishing Co., New York, 1959. Translated by K. A. Hirsch.

[Gui05] Vincent Guirardel. Cœur et nombre d'intersection pour les actions de groupes sur les arbres. *Ann. Sci. École Norm. Sup. (4)*, 38(6):847–888, 2005.

[Gui08] Vincent Guirardel. Actions of finitely generated groups on R-trees. *Annales de l'institut Fourier*, 58(1):159–211, 2008.

[Hir11] Eriko Hironaka. Computing alexander polynomials using monodromy, 2011.

[HM07] Michael Handel and Lee Mosher. The expansion factors of an outer automorphism and its inverse. *Transactions of the American Mathematical Society*, 359(7):3185–3208, 2007.

[HM11] Michael Handel and Lee Mosher. Axes in outer space. *Mem. Amer. Math. Soc.*, 213(1004):vi+104, 2011.

[LL03] Gilbert Levitt and Martin Lustig. Irreducible automorphisms of have north–south dynamics on compactified outer space. *Journal of the Institute of Mathematics of Jussieu*, 2(1):59–72, 2003.

[LMT20] Michael Landry, Yair N Minsky, and Samuel J Taylor. A polynomial invariant for veering triangulations. Preprint arXiv:2008.04836, 2020.

[McM00] Curtis T. McMullen. Polynomial invariants for fibered 3-manifolds and Teichmüller geodesics for foliations. *Ann. Sci. École Norm. Sup. (4)*, 33(4):519–560, 2000.

[McM02] Curtis T. McMullen. The Alexander polynomial of a 3-manifold and the Thurston norm on cohomology. *Ann. Sci. École Norm. Sup. (4)*, 35(2):153–171, 2002.

[McM15] Curtis T McMullen. Entropy and the clique polynomial. *Journal of Topology*, 8(1):184–212, 2015.

[Mil68]   John Milnor. Infinite cyclic coverings. In *Conference on the Topology of Manifolds*, volume 13, pages 115–133, 1968.

[Mut20]   Jean Pierre Mutanguha. Irreducible nonsurjective endomorphisms of F_n are hyperbolic. *Bulletin of the London Mathematical Society*, 52(5):960–976, 2020.

[Mut21]   Jean Pierre Mutanguha. Irreducibility of a free group endomorphism is a mapping torus invariant. *Comment. Math. Helv.*, 96(1):47–63, 2021.

[Par21]   Anna Parlak. The taut polynomial and the Alexander polynomial. Preprint arXiv:2101.12162, 2021.

[Rey11]   Patrick Reese Reynolds. *Dynamics of irreducible endomorphisms of F_n*. ProQuest LLC, Ann Arbor, MI, 2011. Thesis (Ph.D.)–University of Illinois at Urbana-Champaign.

[Thu86]   William P Thurston. A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.*, 59(339):99–130, 1986.

[Thu88]   William P Thurston. On the geometry and dynamics of diffeomorphisms of surfaces. *Bulletin of the American Aathematical Aociety*, 19(2):417–431, 1988.

[Tur75]   Vladimir Georgievich Turaev. The alexander polynomial of a three-dimensional manifold. *Mathematics of the USSR-Sbornik*, 26(3):313, 1975.

DEPARTMENT OF MATHEMATICS, VANDERBILT UNIVERSITY, NASHVILLE, TN
Email address: `spencer.dowdall@vanderbilt.edu`

SCHOOL OF MATHEMATICS, TATA INSTITUTE OF FUNDAMENTAL RESEARCH, MUMBAI
Email address: `rgupta@tifr.res.in`

DEPARTMENT OF MATHEMATICS, TEMPLE UNIVERSITY, PHILADELPHIA, PA
Email address: `samuel.taylor@temple.edu`