Geometry & Topology 19 (2015) 1111–1154

Digraphs and cycle polynomials for free-by-cyclic groups

YAEL ALGOM-KFIR
ERIKO HIRONAKA
KASRA RAFI

Let $\phi \in \mathrm{Out}(F_n)$ be a free group outer automorphism that can be represented by an expanding, irreducible train-track map. The automorphism ϕ determines a free-by-cyclic group $\Gamma = F_n \rtimes_\phi \mathbb{Z}$ and a homomorphism $\alpha \in H^1(\Gamma; \mathbb{Z})$. By work of Neumann, Bieri, Neumann and Strebel, and Dowdall, Kapovich and Leininger, α has an open cone neighborhood $\mathcal{A}$ in $H^1(\Gamma; \mathbb{R})$ whose integral points correspond to other fibrations of Γ whose associated outer automorphisms are themselves representable by expanding irreducible train-track maps. In this paper, we define an analog of McMullen's Teichmüller polynomial that computes the dilatations of all outer automorphisms in $\mathcal{A}$.

57M20

1 Introduction

There is continually growing evidence of a powerful analogy between the mapping class group $\mathrm{Mod}(S)$ of a closed oriented surface S of finite type and the group of outer automorphisms $\mathrm{Out}(F_n)$ of free groups F_n. A recent advance in this direction can be found in work of Dowdall, Kapovich and Leininger [6] who developed an analog of the fibered face theory of surface homeomorphisms due to Thurston [16] and Fried [8]. In this paper we develop the analogy further by defining a version of McMullen's Teichmüller polynomial for surface automorphisms defined in [12] in the setting of outer automorphisms.

Fibered face theory for free-by-cyclic groups

A free-by-cyclic group

$$\Gamma = F_n \rtimes_\phi \mathbb{Z}$$

is a semidirect product defined by an element $\phi \in \mathrm{Out}(F_n)$. If $x_1, \dots, x_n$ are generators of F_n and $\phi_\circ \in \mathrm{Aut}(F_n)$ is a representative automorphism in the class ϕ, then Γ has a finite presentation

$$\langle x_1, \dots, x_n, s \mid s x_i s^{-1} = \phi_\circ(x_i),\ i = 1, \dots, n \rangle.$$

Published: 10 April 2015 DOI: 10.2140/gt.2015.19.1111

There is a distinguished homomorphism $\alpha_\phi\colon \Gamma \to \mathbb{Z}$ induced by projection to the second coordinate. That is, α_ϕ is an element of $H^1(\Gamma;\mathbb{Z})$ and F_n is the kernel of α_ϕ.

The deformation theory of free-by-cyclic groups started with the work of Neumann [14] and Bieri, Neumann and Strebel [3], where they showed there is an open cone U in $H^1(\Gamma;\mathbb{R})$ so that for all rational $u \in H^1(X,\mathbb{R})$, $u \in U$ if and only if $\ker(u)$ is finitely generated. Dowdall, Kapovich and Leininger [6] showed that the deformation can be understood geometrically in a possibly smaller cone.

More precisely, assume $\phi \in \mathrm{Out}(F_n)$ is representable by an expanding irreducible train-track map (see Kapovich [10], Dowdall, Kapovich and Leininger [6] and Section 4.1 for definitions). The outer automorphism $\phi \in \mathrm{Out}(F_n)$ may admit many train-track representatives f and every train-track representative can be decomposed into a sequence of folds f (see Stallings [15]) which is also nonunique. Dowdall, Kapovich and Leininger showed the following (see [6, Theorem A]).

Theorem 1.1 *For $\phi \in \mathrm{Out}(F_n)$ that is representable by an expanding irreducible train-track map and an associated folding sequence f, there is an open cone neighborhood $\mathcal{A}_\mathsf{f}$ of α_ϕ in $\mathrm{Hom}(\Gamma;\mathbb{R})$, such that, all primitive integral elements $\alpha \in \mathcal{A}$, are associated to a free-by-cyclic decomposition*

$$\Gamma = F_{n_\alpha} \rtimes_{\phi_\alpha} \mathbb{Z},$$

where $\alpha = \alpha_{\phi_\alpha}$ and $\phi_\alpha \in \mathrm{Out}(F_{n_\alpha})$ is also representable by an expanding irreducible train-track map.

We call $\mathcal{A}_\mathsf{f}$ a DKL–*cone* associated to ϕ.

Main result

Our main theorem is an analog of results in McMullen [12] in the setting of the outer automorphism groups (see below for more on the motivation behind the result). For a given ϕ, there are many DKL–cones associated to ϕ since A_f depends on the choice of the train-track representative f and folding sequence f. We show that there is a more unified picture. Namely, there is a cone $\mathcal{T}_\phi$ depending only on ϕ that contains every cone $\mathcal{A}_\mathsf{f}$. The cone $\mathcal{T}_\phi$ is the support of a convex, real analytic, homogenous function L of degree -1 whose restriction to every cone $\mathcal{A}_\mathsf{f}$ is the logarithm of dilatation function. Moreover, this function can be computed via specialization of a single polynomial Θ that also depends only on ϕ.

Our approach is combinatorial. We associate a labeled digraph to the folding sequence f. This gives a combinatorial description of f and in turn defines a cycle polynomial θ

and the cone $\mathcal{T}_\phi$. We analyze the effect of certain elementary moves on digraphs and show that their associated cycle polynomial and cone remain unchanged under these elementary moves. We show that as we pass to different fibrations of Γ corresponding to other integral points of $\mathcal{A}_{\mathsf{f}}$, the digraph changes by elementary moves, as do the digraphs associated to different folding sequences f. This establishes the independence of θ and $\mathcal{T}_\phi$ from the choice of folding sequences f. The polynomial Θ is a factor of the cycle polynomial θ determined by the log dilatation and does not depend on the choice of train track map.

We establish some terminology before stating the main theorem more precisely. For $\phi \in \mathrm{Out}(F_n)$ that is representable by an expanding irreducible train-track map and a nontrivial $\gamma \in F_n$, the growth rate of cyclically reduced word-lengths of $\phi^k(\gamma)$ is exponential, with a base $\lambda(\phi) > 1$ that is independent of γ and f. The constant $\lambda(\phi)$ is called the *dilatation* (or *expansion factor*) of ϕ.

Let G be a finitely generated free abelian group of rank k and let

$$\theta = \sum_{g \in G} a_g g, \quad a_g \in \mathbb{Z}$$

be an element of the group ring $\mathbb{Z}G$. For $\alpha \in \mathrm{Hom}(G;\mathbb{Z})$, the *specialization* of θ at α is the single-variable integer polynomial

$$\theta^{(\alpha)}(x) = \sum_{g \in G} a_g x^{\alpha(g)} \in \mathbb{Z}[x].$$

The *house* of a polynomial $p(x) \in \mathbb{Z}[x]$ is defined by

$$|p| = \max\{|\mu| \mid \mu \in \mathbb{C}, p(\mu) = 0\}.$$

Recall that, for a polynomial, there is an associated *Newton polyhedron* defined by a finite system of inequalities as described in Remark 2.7.

Theorem A *Let $\phi \in \mathrm{Out}(F_n)$ be an outer automorphism that is representable by an expanding irreducible train-track map, $\Gamma = F_n \rtimes_\phi \mathbb{Z}$ and let $G = \Gamma^{\mathrm{ab}}/\mathrm{torsion}$. Then there exists an element $\Theta \in \mathbb{Z}G$ (well-defined up to an automorphism of $\mathbb{Z}G$) with the following properties.*

(1) *There is an open cone $\mathcal{T}_\phi \subset \mathrm{Hom}(G;\mathbb{R})$ dual to a vertex of the Newton polyhedron of Θ so that for any expanding irreducible train-track representative $f\colon \tau \to \tau$ and any folding decomposition f of f, we have*

$$\mathcal{A}_{\mathsf{f}} \subset \mathcal{T}_\phi.$$

(2) *For any integral $\alpha \in \mathcal{A}_{\mathsf{f}}$, we have*
$$|\Theta^{(\alpha)}| = \lambda(\phi_\alpha).$$

(3) *The function*
$$L(\alpha) = \log |\Theta^{(\alpha)}|,$$
which is defined on the primitive integral points of $\mathcal{A}_{\mathsf{f}}$, extends to a real analytic, convex function on $\mathcal{T}_\phi$ that is homogeneous of degree -1 and goes to infinity toward the boundary of any affine planar section of $\mathcal{T}_\phi$.

(4) *The element Θ is minimal with respect to property (2), that is, if $\theta \in \mathbb{Z}G$ satisfies*
$$|\theta^{(\alpha)}| = \lambda(\phi_\alpha)$$
on the integral elements of some open subcone of $\mathcal{T}_\phi$, then Θ divides θ.

Remark B In their original paper [6], Dowdall, Kapovich and Leininger also show that $\log(\lambda(\phi_\alpha))$ is convex and has degree -1 and in the subsequent paper [7], using a different approach from ours, they give an independent definition of an element $\Theta_{\mathrm{DKL}} \in \mathbb{Z}G$ such that $\lambda(\phi_\alpha) = |\Theta_{\mathrm{DKL}}^{(\alpha)}|$ for $\alpha \in A_{\mathsf{f}}$. Property (4) of Theorem A implies that Θ divides Θ_{DKL}.

Remark C Thinking of G as an abelian group generated by $t_1, \dots, t_k$, we can identify the elements of G with monomials in the symbols $t_1, \dots, t_k$, and hence $\mathbb{Z}G$ with Laurent polynomials in $\mathbb{Z}(t_1, \dots, t_k)$. Thus we can associate to $\theta \in \mathbb{Z}G$ a polynomial $\theta(t_1, \dots, t_k) \in \mathbb{Z}(t_1, \dots, t_k)$. Identifying $\mathrm{Hom}(G; \mathbb{Z})$ with $\mathbb{Z}^k$, each element $\alpha = (a_1, \dots, a_k)$ defines a specialization of $\theta = \theta(t_1, \dots, t_k)$ by
$$\theta^{(\alpha)}(x) = \theta(x^{a_1}, \dots, x^{a_k}).$$
For ease of notation, we mainly use the group ring notation through most of this paper.

Motivation from pseudo-Anosov mapping classes on surfaces

Let S be a closed oriented surface of negative finite Euler characteristic. A *mapping class* $\phi = [\phi_\circ]$ is an isotopy class of homeomorphisms
$$\phi_\circ \colon S \to S.$$
The mapping torus $X_{(S,\phi)}$ of the pair (S, ϕ) is the quotient space
$$X_{(S,\phi)} = S \times [0,1]/(x,1) \sim (\phi_\circ(x), 0).$$
Its homeomorphism type is independent of the choice of representative $\phi_\circ$ for ϕ. The mapping torus $X_{(S,\phi)}$ has a distinguished fibration $\rho_\phi \colon X_{(S,\phi)} \to S^1$ defined by

projecting $S\times[0,1]$ to its second component and identifying endpoints. Conversely, any fibration $\rho\colon X\to S^1$ of a 3–manifold X over a circle can be written as the mapping torus of a unique mapping class (S,ϕ), with $\rho=\rho_\phi$. The mapping class (S,ϕ) is called the *monodromy* of ρ.

Thurston's fibered face theory [16] gives a parameterization of the fibrations of a 3–manifold X over the circle with connected fibers by the primitive integer points on a finite union of disjoint convex cones in $H^1(X;\mathbb{R})$, called *fibered cones*. Thurston showed that the mapping torus of any pseudo-Anosov mapping class is hyperbolic, and the monodromy of any fibered hyperbolic 3–manifold is pseudo-Anosov. It follows that the set of all pseudo-Anosov mapping classes partitions into subsets corresponding to integral points on fibered cones of hyperbolic 3–manifolds.

By results of Fried [8] (cf Matsumoto [11] and McMullen [12]) the function $\log\lambda(\phi)$ defined on integral points of a fibered cone $\mathcal{T}$ extends to a continuous convex function

$$\mathcal{Y}\colon \mathcal{T}\to\mathbb{R}$$

that is a homogeneous of degree -1, and goes to infinity toward the boundary of any affine planar section of $\mathcal{T}$. McMullen's *Teichmüller polynomial* [12] is an element Θ_{Teich} in the group ring $\mathbb{Z}G$, defined up to units, where $G=H_1(X;\mathbb{Z})/\mathrm{torsion}$. The group ring $\mathbb{Z}G$ can be thought of as a ring of Laurent polynomials in the generators of G considered as a multiplicative group. Thus we can also think of Θ_{Teich} as a polynomial defined up to multiplication by monomials. The Teichmüller polynomial Θ_{Teich} has the property that the dilatation $\lambda(\phi_\alpha)$ of each mapping class ϕ_α, for $\alpha\in\mathcal{T}$, is the house of a specialization of Θ_{Teich}. Furthermore, the cone $\mathcal{T}$ and the function $\mathcal{Y}$ are determined by Θ_{Teich}. Our work is a step towards reproducing this picture in the setting of $\mathrm{Out}(F_n)$.

Organization of paper

In Section 2 we establish some preliminaries about Perron–Frobenius digraphs D with edges labeled by a free abelian group G. Each digraph D determines a cycle complex C_D and cycle polynomial θ_D in the group ring $\mathbb{Z}G$. Under certain extra conditions, we define a cone $\mathcal{T}$, which we call the McMullen cone, and show that

$$L(\alpha)=\log|\theta_D^{(\alpha)}|,$$

which is defined for integral elements of $\mathcal{T}$, extends to a homogeneous function of degree -1 that is real analytic and convex on $\mathcal{T}$ and goes to infinity toward the boundary of affine planar sections of $\mathcal{T}$. Furthermore, we show the existence of a distinguished factor Θ_D of θ_D with the property that

$$|\Theta_D^{(\alpha)}| = |\theta_D^{(\alpha)}|,$$

and Θ_D is minimal with this property. Our proof uses a key result of McMullen (see [12, Appendix A]).

In Section 3 we define branched surfaces $(X, \mathfrak{C}, \psi)$, where X is a 2–complex with a semiflow ψ, and cellular structure $\mathfrak{C}$ satisfying compatibility conditions with respect to ψ. To a branched surface we associate a dual digraph D and a G–labeled cycle complex C_D, where $G = H_1(X; \mathbb{Z})/\text{torsion}$, and a cycle function $\theta_D \in \mathbb{Z}G$. We show that θ_D is invariant under certain allowable cellular subdivisions and homotopic modifications of $(X, \mathfrak{C}, \psi)$.

In Sections 4 and 5 we study the branched surfaces associated to the train-track map f and folding sequence f defined in [6], called respectively the mapping torus and folded mapping torus. We use the invariance under allowable cellular subdivisions and modifications established in Sections 2 and 3 to show that the cycle functions for these branched surfaces are equal. The results of Section 2 applied to the mapping torus for f imply the existence of Θ_ϕ and $\mathcal{T}_\phi$ in Theorem A. An argument in [6] implies that further subdivisions of the folded mapping torus give rise to mapping tori for train-track maps corresponding to ϕ_α, and we use this to show that $\lambda(\phi_\alpha) = |\Theta_\phi^{(\alpha)}|$ for $\alpha \in \mathcal{A}_{\mathsf{f}}$. We further compare the definition of the DKL–cone $\mathcal{A}_{\mathsf{f}}$ and $\mathcal{T}_\phi$ to show inclusion $\mathcal{A}_{\mathsf{f}} \subset \mathcal{T}_\phi$, and thus complete the proof of Theorem A.

We conclude in Section 6 with an example where A_{f} is a proper subcone of $\mathcal{T}_\phi$.

Acknowledgements The authors would like to thank J Birman, S Dowdall, N Dunfield, A Hadari, C Leininger, C McMullen and K Vogtman for helpful discussions.

The second author was partially supported by collaboration grant number 209171 from the Simons Foundation, the third author was partially supported by an NSERC Discovery grant number RGPIN 435885.

2 Digraphs, cycle complexes and eigenvalues of G–matrices

This section contains definitions and properties of digraphs, and a key result of McMullen that will be useful in our proof of Theorem A.

2.1 Digraphs, cycle complexes and their cycle polynomials

We recall basic results concerning digraphs (see Gantmacher [9] and Cvetković and Rowlinson [5] for more details).

Definition 2.1 A *digraph* D is a finite directed graph with at least two vertices. Given an ordering $v_1, \dots, v_m$ of the vertices of D, the *adjacency matrix* of D is the matrix

$$M_D = [a_{i,j}],$$

where $a_{i,j} = m$ if there are m directed edges from v_i to v_j. The *characteristic polynomial* P_D is the characteristic polynomial of M_D and the *dilatation* $\lambda(D)$ of D is the spectral radius of M_D:

$$\lambda(D) = \max\{|e| \mid e \text{ is an eigenvalue of } M_D\}.$$

Conversely, any square $m \times m$ matrix $M = [a_{i,j}]$ with nonnegative integer entries determines a digraph D with $M_D = M$. The digraph D has m vertices and $a_{i,j}$ vertices from the i^{th} to the j^{th} vertex.

Definition 2.2 For a matrix M, let a_{ij}^m be the ij^{th} entry of M^m. A nonnegative matrix M with real entries is called *expanding* if

$$\limsup_{m\to\infty} a_{ij}^m = \infty.$$

A digraph D is *expanding* if its directed adjacency matrix M_D is expanding.

An eigenvalue of M is *simple* if its algebraic multiplicity is 1. Note that several simple eigenvalues may have the same norm. The following theorem is well known (see, for example, [9]).

Theorem 2.3 *Let M be a matrix and $\lambda(M)$ the spectral radius of M. If M is expanding, then it has a simple eigenvalue with norm equal to $\lambda(M)$ and it has an associated eigenvector that is strictly positive. In addition, for every i and j, we have*

$$\limsup_{m\to\infty} (a_{ij}^m)^{1/m} = \lambda(M).$$

Definition 2.4 A *simple cycle* α on a digraph D is an isotopy class of embeddings of the circle S^1 to D oriented compatibly with the directed edges of D. A *cycle* is a disjoint union of simple cycles. The *cycle complex* C_D of a digraph D is the collection of cycles on D thought of as a simplicial complex, whose vertices are the simple cycles.

The cycle complex C_D has a measure which assigns to each cycle its length in D, that is, if γ is a cycle on C_D, then its *length* $\ell(\gamma)$ is the number of vertices (or equivalently the number of edges) of D on γ, and, if $\sigma = \{\gamma_1, \dots, \gamma_s\}$, then

$$\ell(\sigma) = \sum_{i=1}^{s} \ell(\gamma_i).$$

Let $|\sigma| = s$ be the *size* of σ. The *cycle polynomial* of a digraph D is given by
$$\theta_D(x) = 1 + \sum_{\sigma \in C_D} (-1)^{|\sigma|} x^{-\ell(\sigma)}.$$

Theorem 2.5 (Coefficient theorem for digraphs [5]) *Let D be a digraph with m vertices, and P_D the characteristic polynomial of the directed adjacency matrix M_D for D. Then*
$$P_D(x) \;=\; x^m \theta_D(x).$$

Proof Let $M_D = [a_{i,j}]$ be the directed adjacency matrix for D. Then
$$P_D(x) = \det(xI - M_D).$$

Let S_{V} be the group of permutations of the vertices V of D. For $\pi \in S_V$, let $\mathrm{fix}(\pi) \subseteq \mathsf{V}$ be the set of vertices fixed by π, and let $\mathrm{sign}(\pi)$ be -1 if π is an odd permutation and 1 if π is even. Then
$$P_D(x) = \sum_{\pi \in S_{\mathsf{V}}} \mathrm{sign}(\pi) A_\pi,$$
where
$$A_\pi = \prod_{v \notin \mathrm{fix}(\pi)} (-a_{v,\pi(v)}) \prod_{v \in \mathrm{fix}(\pi)} (x - a_{v,v}). \tag{1}$$

There is a natural map $\Sigma\colon C_D \to S_{\mathsf{V}}$ from the cycle complex C_D to the permutation group S_{V} on the set V defined as follows. For each simple cycle γ in D passing through the vertices $\mathsf{V}_\gamma \subset \mathsf{V}$, there is a corresponding cyclic permutation $\Sigma(\gamma)$ of V_γ. That is, if $\mathsf{V}_\gamma = \{v_1, \dots, v_\ell\}$ contains more than one vertex and is ordered according to their appearance in the cycle, then $\Sigma(\gamma)(v_i) = v_{i+1 \pmod{\ell}}$. If V_γ contains one vertex, we say γ is a *self-edge*. For self-edges γ, $\Sigma(\gamma)$ is the identity permutation. Let $\sigma = \{\gamma_1, \dots, \gamma_s\}$ be a cycle on D. Then we define $\Sigma(\sigma)$ to be the product of disjoint cycles
$$\Sigma(\sigma) = \Sigma(\gamma_1) \circ \cdots \circ \Sigma(\gamma_\ell).$$

The polynomial A_π in (1) can be rewritten in terms of the cycles σ of C_D with $\Sigma(\sigma) = \pi$. First we rewrite A_π as
$$A_\pi = \sum_{\nu \subset \mathrm{fix}(\pi)} x^{|\mathrm{fix}(\pi) - \nu|} \prod_{v \notin \mathrm{fix}(\pi)} (-a_{v,\pi(v)}) \prod_{v \in \nu} (-a_{v,v}). \tag{2}$$

Let $\pi \in S_{\mathsf{V}}$ be in the image of Σ. For a cycle $\sigma \in C_D$, let $\nu(\sigma) \subset \mathsf{V}$ be the subset vertices at which σ has a self-edge.

For $\nu \subset \mathrm{fix}(\pi)$, let
$$P_{\pi,\nu} = \{\sigma \in C_D \mid \Sigma(\sigma) = \pi \text{ and } \nu(\sigma) = \nu\}.$$

Then we claim that the number of elements in P_{π_ν} is

$$\prod_{v \notin \mathrm{fix}(\pi)} a_{v,\pi(v)} \prod_{v \in \nu} a_{v,v}. \tag{3}$$

Let $\sigma \in C_D$ be such that $\Sigma(\sigma) = \pi$. Then for each $v \in \mathsf{V} \smallsetminus \mathrm{fix}(\pi)$, there is a choice of $a_{v,\pi(v)}$ edges from v to $\pi(v)$, and for each $v \in \mathrm{fix}(\pi)$ σ either contains no self-edge, or one of $a_{v,v}$ possible self-edges at v. This proves (3).

For each $\sigma \in C_D$, we have

$$\ell(\sigma) = m - |\mathrm{fix}(\Sigma(\sigma))| + |\nu(\sigma)|.$$

Thus the summand in (2) associated to $\pi \in S_{\mathsf{V}} \smallsetminus \mathrm{id}$ and $\nu \subset \mathrm{fix}(\pi)$ is given by

$$\begin{aligned} x^{|\mathrm{fix}(\pi)|-|\nu|} \prod_{v \notin \mathrm{fix}(\pi)} (-a_{v,\pi(v)}) \prod_{v \in \nu} (-a_{v,v}) &= (-1)^{m-|\mathrm{fix}(\pi)|+|\nu|} \sum_{\sigma \in P_{\pi,\nu}} x^{m-\ell(\sigma)} \\ &= \sum_{\sigma \in P_{\pi,\nu}} (-1)^{\ell(\sigma)} x^{m-\ell(\sigma)}, \end{aligned}$$

and similarly for $\pi = \mathrm{id}$ we have

$$A_\pi = \prod_{v \in \mathsf{V}} (x - a_{v,v}) = x^m + \sum_{\sigma \in P_{\pi,\nu}} (-1)^{\ell(\sigma)} x^{m-\ell(\sigma)}.$$

For each $\sigma \in C_D$, $\mathrm{sign}(\Sigma(\sigma)) = (-1)^{\ell(\sigma)-|\sigma|}$. Putting this together, we have

$$\begin{aligned} P_D(x) = \sum_{\pi \in S_{\mathsf{V}}} \mathrm{sign}(\pi) A_\pi &= x^m + \sum_{\pi \in S_{\mathsf{V}}} \sum_{\sigma \in C_D \mid \Sigma(\sigma) = \pi} (-1)^{\ell(\sigma)-|\sigma|} (-1)^{\ell(\sigma)} x^{m-\ell(\sigma)} \\ &= x^m + \sum_{\sigma \in C_D} (-1)^{|\sigma|} x^{m-\ell(\sigma)}. \end{aligned}$$

This completes the proof. □

2.2 McMullen cones

Each group ring element partitions $\mathrm{Hom}(G;\mathbb{R})$ into a union of cones defined below.

Definition 2.6 (cf McMullen [13]) Let G be a finitely generated free abelian group. Given an element $\theta = \sum_{g \in G} a_g g \in \mathbb{Z}G$, the *support* of θ is the set

$$\mathrm{Supp}(\theta) = \{g \in G \mid a_g \neq 0\}.$$

Let $\theta \in \mathbb{Z}G$ and $g_0 \in \mathrm{Supp}(\theta)$ the *McMullen cone* of θ for g_0 is the set

$$\mathcal{T}_\theta(g_0) = \{\alpha \in \mathrm{Hom}(G;\mathbb{R}) \mid \alpha(g_0) > \alpha(g) \text{ for all } g \in \mathrm{Supp}(\theta) \smallsetminus \{g_0\}\}.$$

Remark 2.7 The elements of G can be identified with a subset of the dual space

$$\widehat{\mathrm{Hom}(G;\mathbb{R})} = \mathrm{Hom}(\mathrm{Hom}(G;\mathbb{R}),\mathbb{R})$$

to $\mathrm{Hom}(G;\mathbb{R})$. Let $\theta \in \mathbb{Z}G$ be any element. The convex hull of $\mathrm{Supp}(\theta)$ in $\widehat{\mathrm{Hom}(G;\mathbb{R})}$ is called the *Newton polyhedron* $\mathcal{N}$ of θ. Let $\widehat{\mathcal{N}}$ be the dual of $\mathcal{N}$ in $\mathrm{Hom}(G;\mathbb{R})$. That is, each top-dimensional face of $\widehat{\mathcal{N}}$ corresponds to a vertex $g \in \mathcal{N}$, and each α in the cone over this face has the property that $\alpha(g) > \alpha(g')$ where g' is any vertex of $\mathcal{N}$ with $g \neq g'$. Thus the McMullen cones $\mathcal{T}_\theta(g_0)$ for $g_0 \in \mathrm{Supp}(\theta)$ are the cones over the top-dimensional faces of the dual to the Newton polyhedron of θ.

2.3 A coefficient theorem for H–labeled digraphs

Throughout this section let H be the free abelian group with k generators and let $\mathbb{Z}H$ be its group ring. Let $G = H \times \langle s \rangle$, where s is an extra free variable. Then the Laurent polynomial ring $\mathbb{Z}H(u)$ is canonically isomorphic to $\mathbb{Z}G$, by an isomorphism that sends s to u.

We generalize the results of Section 2.1 to the setting of H–labeled digraphs.

Definition 2.8 Let C be a simplicial complex. An *H–labeling* of C is a map

$$h\colon C \to H$$

compatible with the simplicial complex structure of H, ie

$$h(\sigma) = \sum_{i=1}^{\ell} h(v_i),$$

for $\sigma = \{v_1, \dots, v_\ell\}$. An *$H$–complex* $\mathcal{C}^H$ is an abstract simplicial complex together with a H–labeling.

Definition 2.9 The *cycle function* of an H–labeled complex $\mathcal{C}^H$ is the element of $\mathbb{Z}H$ defined by

$$\theta_{\mathcal{C}^H} = 1 + \sum_{\sigma \in \mathcal{C}^H} (-1)^{|\sigma|} h(\sigma)^{-1}.$$

Definition 2.10 An *H–digraph* $\mathcal{D}^H$ is a digraph D along with a map

$$h\colon \mathcal{E}_D \to H,$$

where $\mathcal{E}_D$ is the set of edge of D. The digraph D is the *underlying digraph* of $\mathcal{D}^H$.

An H–labeling on a digraph induces an H–labeling on its cycle complex. Let γ be a simple cycle on D. Then up to isotopy, γ can be written as

$$\gamma = e_0 \cdots e_{k-1}$$

for some collection of edges $e_0, \ldots, e_{k-1}$ cyclically joined end to end on D. Let

$$h(\gamma) = h(e_0) + \cdots + h(e_{k-1}),$$

and for $\sigma = \{\gamma_1, \ldots, \gamma_\ell\}$, let

$$h(\sigma) = \sum_{i=1}^{\ell} h(\gamma_i).$$

Denote the labeled cycle complex by $\mathcal{C}_{\mathcal{D}}^H$. The *cycle polynomial* $\theta_{\mathcal{D}^H}$ of $\mathcal{D}^H$ is given by

$$\theta_{\mathcal{D}^H}(u) = 1 + \sum_{\sigma \in \mathcal{C}\mathcal{D}^H} (-1)^{|\sigma|} h(\sigma)^{-1} u^{-\ell(\sigma)} \in \mathbb{Z}H[u] = \mathbb{Z}G.$$

The cycle polynomial of $\theta_{\mathcal{D}^H}(u)$ contains both the information about the associated labeled complex $\mathcal{C}_{\mathcal{D}}^H$ and the length functions on cycles on D. One observes the following by comparing Definitions 2.9 and 2.10.

Lemma 2.11 *The cycle polynomial of the H–labeled digraph $\mathcal{D}^H$, and the cycle function of the labeled cycle complex $\mathcal{C}_{\mathcal{D}}^H$ are related by*

$$\theta_{\mathcal{C}_{\mathcal{D}}^H} = \theta_{\mathcal{D}^H}(1).$$

Definition 2.12 An element $\theta \in \mathbb{Z}H$ is *positive*, denoted $\theta > 0$, if

$$\theta = \sum_{h \in H} a_h h,$$

where $a_h \geq 0$ for all $h \in H$, and $a_h > 0$ for at least one $h \in H$. If θ is positive or 0 we say that it is nonnegative and write $\theta \geq 0$.

A matrix M^H with entries in $\mathbb{Z}H$ is called an *H–matrix.* If all entries are nonnegative, we write $M^H \geq 0$ and if all entries are positive we write $M^H > 0$.

Lemma 2.13 *There is a bijective correspondence between H–digraphs $\mathcal{D}^H$ and nonnegative H–matrices $M_{\mathcal{D}}^H$, so that $M_{\mathcal{D}}^H$ is the directed incidence matrix for $\mathcal{D}^H$.*

Proof Given a labeled digraph $\mathcal{D}^H$, let E_{ij} be the set of edges from the i^{th} vertex to the j^{th} vertex. We form a matrix $M_{\mathcal{D}}^H$ with entries in $\mathbb{Z}H$ by setting

$$a_{ij} = \sum_{e \in E_{ij}} h(e),$$

where $h(e)$ is the H–label of the edge e.

Conversely, given an $n \times n$ matrix M^H with entries in $\mathbb{Z}H$, let $\mathcal{D}^H$ be the H–digraph with n vertices $v_1, \dots, v_n$ and, for each i, j with $m_{i,j} = \sum_{h \in H} a_g g \geq 0$, it has a_h directed edges from v_i to v_j labeled by h. The directed incidence matrix $M_{\mathcal{D}}^H$ equals M as desired. □

The proof of the next theorem is similar to that of the Theorem 2.5 and is left to the reader.

Theorem 2.14 (Coefficients theorem for H–labeled digraphs) *Let $\mathcal{D}^H$ be an H–labeled digraph with m vertices, and let $P_{\mathcal{D}}(u) \in \mathbb{Z}H[u]$ be the characteristic polynomial of its incidence matrix. Then*

$$P_{\mathcal{D}}(u) = u^m \theta_{\mathcal{D}^H}(u).$$

2.4 Expanding H–matrices

In this section we recall a key theorem of McMullen on leading eigenvalues of specializations of expanding H–matrices (see [12, Appendix A]). McMullen's theorem is stated for Perron–Frobenius matrices, but the proof extends to expanding matrices.

Definition 2.15 A labeled digraph $\mathcal{D}^H$ is called *expanding* if the underlying digraph $\mathcal{D}$ is expanding. The H–matrix $M_{\mathcal{D}}^H$ is defined to be *expanding* if the associated labeled digraph $\mathcal{D}^H$ is expanding.

For the rest of this section, we fix an expanding H–labeled digraph $\mathcal{D}^H$. Consider an element $\mathrm{t} \in \mathrm{Hom}(H, \mathbb{R}_+)$. Define $M_{\mathcal{D}}^H(\mathrm{t})$ to be the real valued matrix obtained by applying t to the entries of $M_{\mathcal{D}}^H$ (where t is extended linearly to $\mathbb{Z}H$). Equivalently, identify H with the space of monomials in k variables $t_1, \dots, t_k$. This gives a natural identification of $\mathrm{Hom}(H, \mathbb{R}_+)$ with $\mathbb{R}_+^k$, where the i^{th} coordinate in $\mathbb{R}_+^k$ is associated to the variable t_i. Then $M_{\mathcal{D}}^H(\mathrm{t})$ is the matrix obtained by replacing t_i with i^{th} coordinate of $\mathrm{t} \in \mathbb{R}_+^k = \mathrm{Hom}(H, \mathbb{R}_+)$.

Note that, since $\mathcal{D}^H$ is expanding, for every $\mathrm{t} \in \mathbb{R}_+^k$, the real valued matrix $M_{\mathcal{D}}^H(\mathrm{t})$ is also expanding. Define a function

$$E\colon \mathbb{R}_+^k \to \mathbb{R}_+, \quad E(\mathrm{t}) = \lambda(M_{\mathcal{D}}^H(\mathrm{t})).$$

Identifying the ring $\mathrm{Hom}(H,\mathbb{R})$ with $\mathbb{R}^k$, there is a natural map

$$\exp\colon \mathrm{Hom}(H,\mathbb{R}) \to \mathrm{Hom}(H,\mathbb{R}_+),$$

where, for $\mathsf{w} = (w_1,\dots,w_k) \in \mathbb{R}^k$,

$$\exp(\mathsf{w}) = (e^{w_1},\dots,e_k^{w}).$$

Define

$$\delta\colon \mathbb{R}^k \to \mathbb{R} \quad \text{by} \quad \delta(\mathsf{w}) = \log E(\exp(\mathsf{w})).$$

Note that the graph of the function δ lives in $\mathbb{R}^k \times \mathbb{R}$ which can be naturally identified with $\mathrm{Hom}(G;\mathbb{R})$, where we recall that $G = H \times \langle s \rangle$.

Theorem 2.16 [12, Theorem A.1] *For an expanding H–labeled digraph $\mathcal{D}^H$, we have the following.*

1. *The function δ is real analytic and convex.*
2. *The graph of δ meets every ray through the origin of $\mathbb{R}^k \times \mathbb{R}$ at most once.*
3. *For $Q(u)$ any factor of $P_{\mathcal{D}}(u)$, where $Q(E(\mathsf{t})) = 0$ for all $\mathsf{t} \in \mathbb{R}^k_+$, and for $d = \deg(Q)$, the set of rays passing through the graph of δ in $\mathbb{R}^k \times \mathbb{R}$ coincides with the McMullen cone $\mathcal{T}_Q(u^d)$.*

Definition 2.17 For any expanding H–labeled digraph $\mathcal{D}^H$, let $d = \deg(P_{\mathcal{D}})$. We refer to the cone $\mathcal{T} = \mathcal{T}_{P_{\mathcal{D}}}(u^d)$ as the *McMullen cone* for the element $P_{\mathcal{D}} \in \mathbb{Z}G$. Alternatively we refer to it as the McMullen cone for the H–matrix $M_{\mathcal{D}}^H$.

Theorem 2.18 (McMullen [12]) *For any expanding H–labeled digraph $\mathcal{D}^H$ the map*

$$L\colon \mathrm{Hom}(G;\mathbb{Z}) \to \mathbb{R} \quad \textit{defined by} \quad L(\alpha) = \log |P_{\mathcal{D}}^{(\alpha)}|,$$

extends to a homogeneous of degree -1, real analytic, convex function on the McMullen cone $\mathcal{T}$ for the element $P_{\mathcal{D}}$. It goes to infinity toward the boundary of affine planar sections of $\mathcal{T}$.

Theorem 2.18 summarizes results taken from [12] given in the context of mapping classes on surfaces. For the convenience of the reader, we give a proof here.

Proof The function L is real analytic since the house of a polynomial is an algebraic function in its coefficients. Homogeneity of $L(z)$ follows from the following observation: ρ is a root of $Q(x^{\mathsf{w}}, x^s)$ if and only if $\rho^{1/c}$ is a root of $Q(x^{c\mathsf{w}}, x^{cs})$. Thus

$$L(cz) = \log |Q(x^{c\mathsf{w}}, x^{cs})| = c^{-1} \log |Q(x^{\mathsf{w}}, x^s)| = c^{-1} L(z).$$

By homogeneity of L, the values of L are determined by the values at any level set, one of which is the graph of $\delta(\mathsf{w})$. To prove convexity of L, we show that level sets of L are convex, ie the line connecting two points on a level set lies above the level set. Let $\Gamma = \{z = (\mathsf{w}, s) \mid L(z) = 1\}$ and $\Gamma' = \{z = (\mathsf{w}, s) \mid s = \delta(\mathsf{w})\}$. We show that $\Gamma = \Gamma'$. It then follows that, since Γ' is a graph of a convex function by Theorem 2.16, Γ is convex.

We begin by showing that $\Gamma' \subset \Gamma$ (cf [12, proof of Theorem 5.3]). If $\beta = (\mathsf{a}, b) \in \Gamma'$ then $\delta(\mathsf{a}) = b$, hence $Q(e^{\mathsf{a}}, e^{b}) = 0$ and $|Q(e^{\mathsf{a}}, e^{b})| \geq e$. Let

$$r = L(\beta) = \log |Q(e^{\mathsf{a}}, e^{b})|.$$

Since $b = \delta(\mathsf{a})$, by the convexity of the function δ, we have $rb \geq \delta(r\mathsf{a})$. On the other hand, $Q(e^{r\mathsf{a}}, e^{rb}) = 0$ hence e^{rb} is an eigenvalue of $M(e^{r\mathsf{a}})$ so

$$rb \leq \log E(e^{r\mathsf{a}}) = \delta(r\mathsf{a}).$$

We get that $rb = \delta(r\mathsf{a})$. The points $(\mathsf{a}, b), (r\mathsf{a}, rb)$ both lie on the same line through the origin so by Theorem 2.16(2), they are equal. Thus $r = 1 = L(\beta)$, and hence $\beta \in \Gamma$.

To show that $\Gamma \subset \Gamma'$ in $\mathcal{T}$, note that every ray in $\mathcal{T}$ initiating from the origin intersects Γ because it intersects Γ' by part (3) of Theorem 2.16. Because L is homogeneous, level sets of L intersect every ray from the origin at most once. Therefore, in $\mathcal{T}$, $\Gamma = \Gamma'$ and is the graph of a convex function.

We now show that if L is a homogeneous function of degree -1, and has convex level sets then L is convex (cf [12, Corollary 5.4]). This is equivalent to showing that $1/L(z)$ is concave on $\mathcal{T}$. Let $z_1, z_2 \in T$ lie on distinct rays through the origin, and let

$$z_3 = sz_1 + (1-s)z_2.$$

Let c_i, $i = 1, 2, 3$, be constants so that $z_i' = c_i^{-1} z_i$ is in the level set $L(c_i^{-1} z_i) = 1$. Let p lie on the line $[z_1', z_2']$ and on the ray through z_3. Then p has the form

$$p = rz_1' + (1-r)z_2'$$

for $0 < r < 1$. If

$$r = \frac{sc_1}{sc_1 + (1-s)c_2},$$

then we have

$$p = \frac{z_3}{sc_1' + (1-s)c_2'}.$$

Since the level set for $L(z) = 1$ is convex, p is equal to or above z_3/c_3, and we have

$$1/(sc_1 + (1-s)c_2) \geq 1/c_3. \tag{4}$$

Thus

$$1/L(z_3) = c_3 \geq sc_1 + (1-s)c_2 = s/L(z_2) + (1-s)/L(z_3). \tag{5}$$

Thus $1/L(z)$ is concave, and hence $L(z)$ is convex.

Let z_n be a sequence of points on an affine planar section of $\mathcal{T}$ approaching the boundary of $\mathcal{T}$. Let c_n be such that $c_n^{-1} z_n$ is in the level set $L(z) = 1$. Then $L(z_n) = c_n^{-1}$ for all n. But z_n is bounded, while the level set $L(z) = 1$ is asymptotic to the boundary of $\mathcal{T}$. Therefore, $1/L(z_n)$ goes to 0 as n goes to infinity. □

Remark 2.19 If the level set $L(z) = 1$ is strictly convex, then $L(z)$ is strictly convex. Indeed, if $L(z) = 1$ is strictly convex, then the inequality in (4) is strict, and hence the same holds for (5).

2.5 Distinguished factor of the characteristic polynomial

We define a distinguished factor of the characteristic polynomial of a Perron–Frobenius H–matrix.

Proposition 2.20 *Let P be the characteristic polynomial of a Perron–Frobenius H–matrix. Then P has a factor Q with the following properties.*

(1) *For all integral elements α in the McMullen cone $\mathcal{T}$,*
$$|P^{(\alpha)}| = |Q^{(\alpha)}|.$$

(2) *The polynomial Q is minimal, ie if $Q_1 \in \mathbb{Z}H[u]$ satisfies $|Q^{(\alpha)}| = |Q_1^{(\alpha)}|$ for all α ranging among the integer points of an open subcone of $\mathcal{T}$, then Q divides Q_1.*

(3) *The cones $\mathcal{T}_P(u^d)$ and $\mathcal{T}_Q(u^r)$ are equal, where d is the degree of P and r is the degree of Q as elements of $\mathbb{Z}H[u]$.*

Definition 2.21 Given a Perron–Frobenius H–matrix M^H, the polynomial Q is called the *distinguished factor* of the characteristic polynomial of M^H.

Lemma 2.22 *Let $F(\mathrm{t})\colon \mathbb{R}^k \to \mathbb{R}$ be a function. Then*
$$I_F = \{\theta \in \mathbb{Z}(\mathrm{t})[u] \mid \theta(\mathrm{t}, F(\mathrm{t})) = 0 \textit{ for all } \mathrm{t} \in \mathbb{R}^k\}$$
is a principal ideal.

Proof Let $\mathbb{Q}(\mathrm{t})[u]$ be the ring of polynomials in the variable u over the quotient field $\mathbb{Q}(\mathrm{t})$ of $\mathbb{Z}(\mathrm{t})$. Since $\mathbb{Q}(\mathrm{t})[u]$ is a principal ideal domain, I_F generates a principal ideal $\overline{I}_F$ in $\mathbb{Q}(\mathrm{t})[u]$.

Let $\overline{\theta}_1$ be a generator of $\overline{I}_F$; then $\overline{\theta}_1 = \theta_1(\mathrm{t}, u)/\sigma(\mathrm{t})$ with $\theta_1 \in I_F$. Thus $\overline{\theta}_1(\mathrm{t}, F(\mathrm{t})) = 0$ for all t. If I_F is the zero ideal then there is nothing to prove, therefore we suppose it is not. Let $\overline{\theta}_1(\mathrm{t}, u) = \nu(\mathrm{t}, u)/\delta(\mathrm{t})$, where ν and δ are relatively prime in $\mathbb{Q}(\mathrm{t})[u]$, a unique factorization domain. Since $\theta_1(\mathrm{t}, F(\mathrm{t})) = 0$ for all t, $\nu(\mathrm{t}, F(\mathrm{t})) = 0$ for all t, and hence $\nu \in I_F$.

Since I_F is not the zero ideal then $\overline{I}_F$ is not the zero ideal, hence $\overline{\theta}_1 \neq 0$ which implies that $\nu \neq 0$. Let $\theta \in I_F$ be any polynomial. Since $\overline{\theta}_1$ divides θ, then ν divides $\theta\delta$. Since ν and δ are relatively prime, ν divides θ. We've shown that ν divides all elements of I_F. Thus ν is a principal generator. □

Proof of Proposition 2.20 The proposition follows from Lemma 2.22 by declaring Q to be the generator of I_L for $L\colon \mathcal{T} \to \mathbb{R}$ defined in Theorem 2.18. □

3 Branched surfaces with semiflows

In this section we associate a digraph and an element $\theta_{X,\mathfrak{C},\psi} \in \mathbb{Z}G$ to a branched surface $(X, \mathfrak{C}, \psi)$. We show that this element is invariant under certain kinds of subdivisions of $\mathfrak{C}$.

3.1 The cycle polynomial of a branched surface with a semiflow

Definition 3.1 Given a 2–dimensional CW–complex X, a *semiflow* on X is a continuous map $\psi\colon X \times \mathbb{R}_+ \to X$ satisfying:

(i) $\psi(\,\cdot\,, 0)\colon X \to X$ is the identity.
(ii) $\psi(\,\cdot\,, t)\colon X \to X$ is a homotopy equivalence for every $t \geq 0$.
(iii) $\psi(\psi(x, t_0), t_1) = \psi(x, t_0 + t_1)$ for all $t_0, t_1 \geq 0$.

A cell-decomposition $\mathfrak{C}$ of X is ψ*–compatible* if the following hold.

(1) Each 1–cell is either contained in a flow line (*vertical*), or transversal to the semiflow at every point (*transversal*).
(2) For every vertex $p \in \mathfrak{C}^{(0)}$, the image of the *forward flow* of p,
$$\{\psi(p, t) \mid t \in \mathbb{R}_{>0}\},$$
is contained in $\mathfrak{C}^{(1)}$.

A *branched surface* is a triple $(X, \mathfrak{C}, \psi)$, where X is a 2–complex with semiflow ψ and a ψ–compatible cellular structure $\mathfrak{C}$.

Remark 3.2 We think of branched surfaces as flowing downwards. From this point of view, property (2) implies that every 2–cell $c \in \mathfrak{C}^{(2)}$ has a unique *top* 1*–cell*, that is, a 1–cell e such that each point in c can be realized as the forward orbit of a point on e.

Definition 3.3 Let e be a 1–cell on a branched surface $(X, \mathfrak{C}, \psi)$ that is transverse to the flow at every point. A *hinge* containing e is an equivalence class of homeomorphisms $\kappa\colon [0,1] \times [-1,1] \hookrightarrow X$ so that:

(1) The half segment $\Delta = \{(x, 0) \mid x \in I\}$ is mapped onto e.
(2) The image of the interior of the Δ intersects $\mathfrak{C}^{(1)}$ only in e.
(3) The vertical line segments $\{x\} \times [-1,1]$ are mapped into flow lines on X.

Two hinges κ_1, κ_2 are *equivalent* if there is an isotopy rel Δ between them. The 2–cell on $(X, \mathfrak{C}, \psi)$ containing $\kappa([0,1] \times [0,1])$ is called the *initial cell* of κ and the 2–cell containing the point $\kappa([0,1] \times [-1,0])$ is called the *terminal cell* of κ.

An example of a hinge is illustrated in Figure 1.

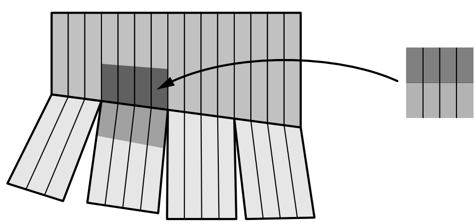

Figure 1: A hinge on a branched surface

Definition 3.4 Let $(X, \mathfrak{C}, \psi)$ be a branched surface. The *dual digraph* D of $(X, \mathfrak{C}, \psi)$ is the digraph with a vertex for every 2–cell and an edge for every hinge κ from the vertex corresponding to its initial 2–cell to the vertex corresponding to its terminal 2–cell. The dual digraph D for $(X, \mathfrak{C}, \psi)$ embeds into X

$$D \hookrightarrow X$$

so that each vertex is mapped into the interior of the corresponding 2–cell, and each directed edge is mapped into the union of the two-cells corresponding to its initial and end vertices, and intersects the common boundary of the 2–cells at a single point. The embedding is well-defined up to homotopies of X to itself.

An example of an embedded dual digraph is shown in Figure 2. In this example, there are three edges emanating from v with endpoints at w_1, w_2 and w_3. It is possible that $w_i = w_j$ for some $i \neq j$, of that $w_i = v$ for some i. These cases can be visualized using Figure 2, by identifying the corresponding 2–cells.

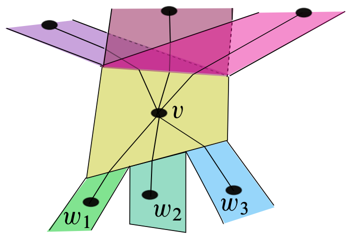

Figure 2: A section of an embedded dual digraph

Let $G = H_1(X;\mathbb{Z})/\text{torsion}$, thought of as the integer lattice in $H_1(X,\mathbb{R})$. The embedding of D in X determines a G–labeled cycle complex $\mathcal{C}_D^G$ where for each $\sigma \in \mathcal{C}_D^G$ and $g(\sigma)$ is the homology class of the cycle σ considered as a 1–cycle on X.

Definition 3.5 Given a branched surface $(X, \mathfrak{C}, \psi)$, the *cycle function* of $(X, \mathfrak{C}, \psi)$ is the group ring element

$$\theta_{X,\mathfrak{C},\psi} = 1 + \sum_{\sigma \in \mathcal{C}_D^G} (-1)^{|\sigma|} g(\sigma)^{-1} \in \mathbb{Z}G.$$

Then we have

$$\theta_{X,\mathfrak{C},\psi} = \theta_{\mathcal{C}_D^G}(1),$$

where $\theta_{\mathcal{C}_D^G}(u)$ is the cycle polynomial of $\mathcal{C}_D^G$.

3.2 Subdivision

We show that the cycle function of $(X, \mathfrak{C}, \psi)$ is not invariant under certain kinds of cellular subdivisions.

Definition 3.6 Let $p \in \mathfrak{C}^{(1)}$ be a point in the interior of a transversal edge in $\mathfrak{C}^{(1)}$. Let $x_0 = p$ and inductively define $x_i = \psi(x_{i-1}, s_i)$, for $i = 1, \dots, r$, so that

$$s_i = \min\{s \mid \psi(x_{i-1}, s) \text{ has endpoint in } \mathfrak{C}^{(1)}\}.$$

The *vertical subdivision of X along the forward orbit of p* is the cellular subdivision $\mathfrak{C}'$ of $\mathfrak{C}$ obtained by adding the edges $\psi(x_{i-1}, [0, s_i])$, for $i = 1, \dots, r$, and subdividing the corresponding 2–cells. If x_r is a vertex in the original skeleton $\mathfrak{C}^{(0)}$ of X, then we say the vertical subdivision is *allowable*.

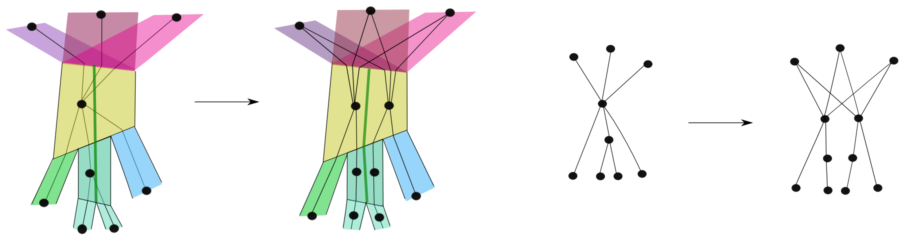

Figure 3: An allowable vertical subdivision, and effect on the directed dual digraph

Proposition 3.7 *Let $(X, \mathfrak{C}', \psi)$ be obtained from $(X, \mathfrak{C}, \psi)$ by allowable vertical subdivision. Then the cycle function $\theta_{X,\mathfrak{C},\psi}$ and $\theta_{X,\mathfrak{C}',\psi}$ are equal.*

We establish a few lemmas before proving Proposition 3.7.

Lemma 3.8 *Let $(X, \mathfrak{C}', \psi)$ be obtained from $(X, \mathfrak{C}, \psi)$ by allowable vertical subdivision. Let D' and D be the dual digraphs for $(X, \mathfrak{C}', \psi)$ and $(X, \mathfrak{C}, \psi)$. There is a quotient map $q\colon D' \to D$ that is induced by a continuous map from X to itself that is homotopic to the identity, and in particular the diagram*

$$\begin{tikzcd} H_1(D';\mathbb{Z}) \arrow[r, "q_*"] \arrow[dr] & H_1(D;\mathbb{Z}) \arrow[d] \\ & H_1(X;\mathbb{Z}) \end{tikzcd}$$

commutes.

Proof Working backwards from the last vertically subdivided cell to the first, each allowable vertical subdivision decomposes into a sequence of allowable vertical subdivisions that involve only one 2–cell. An illustration is shown in Figure 4.

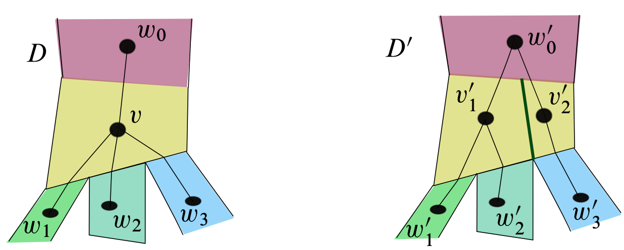

Figure 4: Vertical subdivision of one cell

Let v be the vertex of D corresponding to the cell c of X that contains the new edge. The digraph D' is constructed from D by the following steps:

(1) Each vertex $u \neq v$ in D lifts to a well-defined vertex u' in D'. The vertex $v \in D$ lifts to two vertices v_1', v_2' in D'.

(2) For each edge ε of D neither of whose endpoints u and w equal v, the quotient map is one-to-one over ε, and hence there is only one possible lift ϵ' from u' to w'.

(3) For each edge ε from $w \neq v$ to v there are two edges $\varepsilon_1', \varepsilon_2'$ where ε_i' begins at w' and ends at v_i'.

(4) For each outgoing edge ε from v to w (where v and w are possibly equal), there is a representative κ of the hinge corresponding to ε that is contained in the union of two 2–cells in the $\mathfrak{C}'$. This determines a unique edge ε' on D' that lifts ε.

There is a continuous map homotopic to the identity from X to itself that restricts to the identity on every cell other than c or c_w, where c_w corresponds to a vertex w with an edge from w to v in D. On $c \cup c_w$ the map merges the edges $\varepsilon_1', \varepsilon_2'$ so that their endpoints v_i' merge to the one vertex v. □

Lemma 3.9 *The quotient map $q : D \to D'$ induces an inclusion*

$$q^*\colon C_D \hookrightarrow C_{D'}$$

which preserves lengths, sizes, and labels, so that for $\sigma \in C_D$, $q(q^*(\sigma)) = \sigma$.

Proof Again we may assume that the subdivision involves a vertical subdivision of one 2–cell c corresponding to the vertex $v \in D$ and then use induction. It is enough to define lifts of simple cycles on D to a simple cycle in D'. All edges in D from u to w with $w \neq v$ have a unique lift in D'. Thus, if γ does not contain v then there is a unique γ' in D' such that $q(\gamma') = \gamma$. Suppose that γ contains v. If γ consists of a single edge ε, then ε is a self-edge from v to itself, and ε has two lifts: a self-edge from v_1' to v_1' and an edge from v_1' to v_2', where v_1' is the vertex corresponding to the initial cell of the hinge containing ε. Thus, there is a well-defined self-edge γ' lifting γ (see Figure 5).

Now suppose γ is not a self-edge and contains v. Let $w_1, \dots, w_{\ell-1}$ be the vertices in γ other than v in their induced sequential order. Let ε_i be the edge from w_{i-1} to w_i for $i = 2, \dots, \ell-1$. Then since none of the ε_i have initial or endpoint v, they have unique lifts ε_i' in D'. Since the vertical subdivision is allowable, there is one vertex, say v_1', above v with an edge ε_1' from v_1' to w_2'. Let ε_ℓ' be the edge from $w_{\ell-1}'$ to v_1' (cf Figure 4). Let γ' be the simple cycle with edges $\varepsilon_1', \dots, \varepsilon_\ell'$.

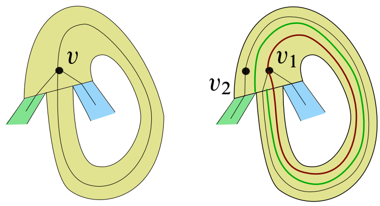

Figure 5: Vertical subdivision when digraph has a self edge

Since the lift of a simple cycle is simple, the lifting map determines a well-defined map $q^*\colon C_D \to C_{D'}$ that satisfies $q \circ q^* = \mathrm{id}$ and preserves size. The commutative diagram in Lemma 3.8 implies that the images of σ and $q^*(\sigma)$ in G are the same, and hence their labels are the same. $\square$

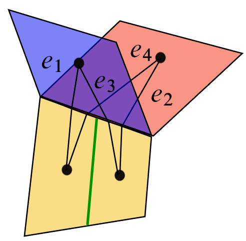

Figure 6: A switching locus

Lemma 3.10 *Let D' be obtained from D by an allowable vertical subdivision on a single 2–cell. The set of edges of each $\sigma \in C_{D'} \smallsetminus q^*(C_D)$ contains exactly one matched pair.*

Proof Since $\sigma' \notin q^*(C_D)$, the quotient map q is not injective on σ'. Thus $q(\sigma')$ must contain two distinct edges $\varepsilon_1, \varepsilon_2$ with endpoint v, and these have lifts ε_1' and ε_2' on σ'. Since σ' is a cycle, ε_1' and ε_2' must have distinct endpoints, hence one is v_1' and one is v_2'. There cannot be more than one matched pair on σ', since σ' can pass through each v_i' only once. $\square$

Definition 3.11 Let D' be obtained from D by an allowable vertical subdivision on a single 2–cell. Let v be the vertex corresponding to the subdivided cell, and let v_1' and v_2' be its lifts to D'.

For any pair of edges $\varepsilon_1', \varepsilon_2'$ with endpoints at v_1' and v_2' and distinct initial points w_1' and w_2', there is a corresponding pair of edges η_1', η_2' from w_1' to v_2' and from w_2'

to v_1'. Write

$$\mathrm{op}\{\varepsilon_1', \varepsilon_2'\} = \{\eta_1', \eta_2'\}.$$

We call the pair $\{\varepsilon_1', \varepsilon_2'\}$ a *matched pair*, and $\{\eta_1', \eta_2'\}$ its *opposite*. (See Figure 6).

Lemma 3.12 *If $\sigma' \in C_{D'}$ contains a matched pair, the edge-path obtained from σ' by exchanging the matched pair with its opposite is a cycle.*

Proof It is enough to observe that the set of endpoints and initial points of a matched pair and its opposite are the same. □

Define a map $\mathfrak{r}\colon C_{D'} \to C_{D'}$ to be the map that sends each $\sigma \in C_{D'}$ to the cycle obtained by exchanging each appearance of a matched pair on $\sigma' \in C_{D'}$ with its opposite.

Lemma 3.13 *The map $\mathfrak{r}$ is a simplicial map of order two that preserves length and labels. It also fixes the elements of $q^*(C_D)$, and changes the parity of the size of elements in $C_{D'} \setminus q^*(C_D)$.*

Proof The map $\mathfrak{r}$ sends cycles to cycles, and hence simplices to simplices. Since op has order 2, it follows that $\mathfrak{r}$ has order 2. The total number of vertices does not change under the operation op. It remains to check that the homology class of σ' and $\mathfrak{r}(\sigma')$ as embedded cycles in X are the same, and that the size switches parity.

There are two cases. Either the matched edges lie on a single simple cycle γ' or on different simple cycles γ_1', γ_2' on σ'.

In the first case, $\mathfrak{r}(\{\gamma'\})$ is a cycle with 2 components $\{\gamma_1', \gamma_2'\}$. As 1–chains we have

$$\beta = \mathfrak{r}(\sigma') - \sigma' = \gamma_1' + \gamma_2' - \gamma' = \eta_1' + \eta_2' - \varepsilon_1' - \varepsilon_2'. \tag{6}$$

In X, β bounds a disc (see Figure 6), thus $g(\gamma') = g(\gamma_1') + g(\gamma_2')$, and hence

$$g(\sigma') = g(\mathfrak{r}(\sigma')). \tag{7}$$

The simple cycle γ' is replaced by two simple cycles γ_1' and γ_2', and hence the size of σ' and $\mathfrak{r}(\sigma')$ differ by one.

Now suppose σ' contains two cycles γ_1' and γ_2', one passing through v_1' and the other passing through v_2'. Then $\mathfrak{r}(\sigma')$ contains a simple cycle γ' in place of $\gamma_1' + \gamma_2'$, so the size decreases by one. By (6) we have (7) for σ' of this type. □

Proof of Proposition 3.7 By Lemma 3.9, the quotient map $q\colon D' \to D$ induces an injection of $q^*\colon C_D \hookrightarrow C_{D'}$ defined by the lifting map, and this map preserves labels. We thus have

$$\theta_{X,\mathfrak{C},\psi} = 1 + \sum_{\sigma \in C_D} (-1)^{|\sigma|} g(\sigma)^{-1} = 1 + \sum_{\sigma' \in q^*(C_D)} (-1)^{|\sigma'|} g(\sigma')^{-1}.$$

The cycles in $C_{D'} \smallsetminus q^*(C_D)$ partition into $\sigma', \mathfrak{r}(\sigma')$, and by Lemma 3.13 the contributions of these pairs in $\theta_{X,\mathfrak{C}',\psi}$ cancel with each other. Thus, we have

$$\theta_{X,\mathfrak{C}',\psi} = 1 + \sum_{\sigma' \in C_{D'}} (-1)^{|\sigma'|} g(\sigma')^{-1} = 1 + \sum_{\sigma' \in q^*(C_D)} (-1)^{|\sigma'|} g(\sigma')^{-1} = \theta_{X,\mathfrak{C},\psi}. \qquad \square$$

Definition 3.14 Let $(X, \mathfrak{C}, \psi)$ be a branched surface and c a 2–cell. Let p, q be two points on the boundary 1–chain ∂c of c that do not lie on the same 1–cell of $\mathfrak{C}$. Assume that p and q each have the property that

(i) it lies on a vertical edge, or

(ii) its forward flow under ψ eventually lies on a vertical 1–cell of $(X, \mathfrak{C})$.

The *transversal subdivision* of $(X, \mathfrak{C}, \psi)$ at $(c; p, q)$ is the new branched surface $(X, \mathfrak{C}', \psi)$ obtained from $\mathfrak{C}$ by doing the (allowable) vertical subdivisions of $\mathfrak{C}$ defined by p and q, and doing the additional subdivision induced by adding a 1–cell from p to q.

Lemma 3.15 *Let $(X, \mathfrak{C}, \psi)$ be a branched surface, and let $(X, \mathfrak{C}', \psi)$ be a transversal subdivision. Then the corresponding cycle functions are the same.*

Proof By first vertically subdividing $\mathfrak{C}$ along the forward orbits of p and q if necessary, we may assume that p and q lie on different vertical 1–cells on the boundary of c. Let v be the vertex of D corresponding to c. Then D' is obtained from D by substituting the vertex v by a pair v'_1, v'_2 that are connected by a single edge. Each edge ε from $w \neq v$ to v is replaced by an edge ε' from w' to v'_1 and edge ε from v to $u \neq v$ is replaced by an edge from v'_2 to u'. Each edge from v to itself is substituted by an edge from v_2 to v_1. The cycle complexes of D and D' are the same, and their labelings are identical. Thus the cycle function is preserved. $\square$

3.3 Folding

Let $(X, \mathfrak{C}, \psi)$ be a branched surface with a flow. Let c_1 and c_2 be two cells with the property that their boundaries ∂c_1 and ∂c_2 both contain the segment $e_1 e_2$, where e_1

is a vertical 1–cell and e_2 is a transversal 1–cell of $\mathfrak{C}$. Let p be the initial point of e_1 and q the end point of e_2. Then p and q both lie on vertical 1–cells, and hence $(c_1; p, q)$ and $(c_2; p, q)$ define a composition of transversal subdivisions $\mathfrak{C}_1$ of $\mathfrak{C}$. For $i = 1, 2$, let e_3^i be the new 1–cell on c_i, and let $\Delta(e_1, e_2, e_3^i)$ be the triangle c_i bounded by the 1–cells e_1, e_2 and e_3^i.

Definition 3.16 The quotient map $F\colon X \to X'$ that identifies $\Delta(e_1, e_2, e_3^1)$ and $\Delta(e_1, e_2, e_3^2)$ (see Figure 7) is called the *folding map* of X. The quotient X' is endowed with the structure of a branched surface $(X', \mathfrak{C}', \psi')$ induced by $(X, \mathfrak{C}_1, \psi)$.

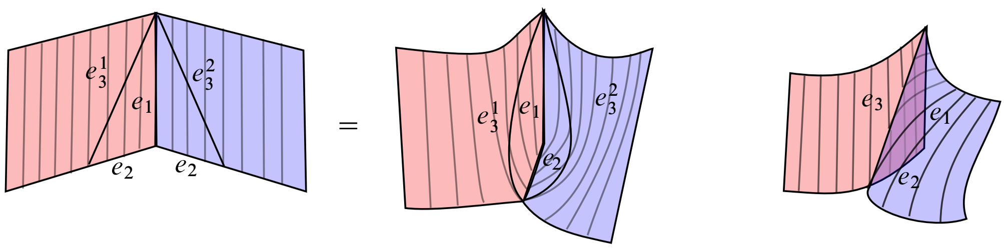

Figure 7: The left and middle diagrams depict the two 2–cells sharing the edges e_1 and e_2; the right diagram is the result of folding.

The following proposition is easily verified (see Figure 7).

Proposition 3.17 *The quotient map F associated to a folding is a homotopy equivalence, and the semiflow $\psi\colon X \times \mathbb{R}_+ \to X$ induces a semiflow $\psi'\colon X \times \mathbb{R}_+ \to X$.*

Definition 3.18 Given a folding map $F\colon X \to X'$, there is an induced branched surface structure $(X', \mathfrak{C}', \psi')$ on X given by taking the minimal cellular structure on X' for which the map F is a cellular map and deleting the image of e_2 if there are only two hinges containing e_2 on X.

Remark 3.19 In the case that c_1, c_2 are the only cells above e_2, folding preserves the dual digraph D.

Lemma 3.20 *Let $F\colon X \to X'$ be a folding map, and let $(X', \mathfrak{C}', \psi')$ be the induced branch surface structure of the quotient. Then*

$$\theta_{X,\mathfrak{C},\psi} = \theta_{X',\mathfrak{C}',\psi'}.$$

Proof Let D be the dual digraph of $(X, \mathfrak{C}, \psi)$ and D' the dual digraph of $(X', \mathfrak{C}', \psi')$. Assume that there are at least three hinges containing e_2. Then D' is obtained from D by gluing two adjacent half edges (see Figure 8), a homotopy equivalence. Thus, $\mathcal{CD}^G = \mathcal{CD}'^G$, and the cycle polynomials are equal. □

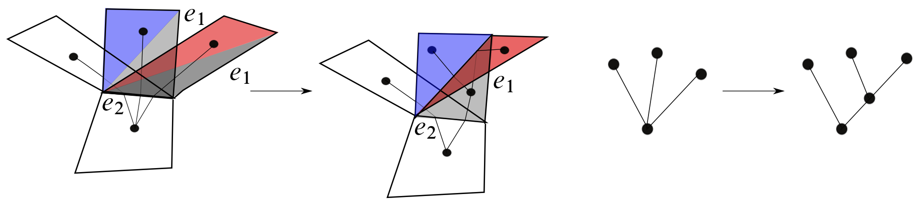

Figure 8: Effect of folding on the digraph

4 The branched surface of an automorphism

Throughout this section, let $\phi \in \mathrm{Out}(F_n)$ be an element that can be represented by an expanding irreducible train-track map $f\colon \tau \to \tau$. Let $\Gamma = F_n \rtimes_\phi \mathbb{Z}$, and $G = \Gamma^{\mathrm{ab}}/\text{torsion}$. We shall define the mapping torus $(Y_f, \mathfrak{C}, \psi)$ associated f, and prove that its cycle polynomial $\theta_{Y_f,\mathfrak{C},\psi}$ has a distinguished factor Θ with a distinguished McMullen cone $\mathcal{T}$. We show that the logarithm of the house of Θ specialized at integral elements in $\mathcal{T}$ extends to a homogeneous of degree -1, real analytic concave function L on an open cone in $\mathrm{Hom}(G,\mathbb{R})$, and satisfies a universality property.

4.1 Free group automorphisms and train-track maps

In this section we give some background definitions for free group automorphisms, and their associated train-track representatives following [6]. We also recall some sufficient conditions for a free group automorphism to have an expanding irreducible train-track map due to work of Bestvina and Handel [2].

Definition 4.1 A *topological graph* is a finite 1–dimensional cellular complex. For each edge e, an orientation on e determines an *initial* and *terminal* point of e. Given an oriented edge e, we denote by $\overline{e}$, the edge e with opposite orientation. An *edge path* on a graph is an ordered sequence of edges $e_1 \cdots e_\ell$, where the endpoint of e_i is the initial point of e_{i+1}, for $i = 1, \dots, \ell - 1$. The edge path has *back-tracking* if $e_i = \overline{e}_{i+1}$ for some i. The length of an edge path $e_1 \cdots e_\ell$ is ℓ.

Definition 4.2 A *graph map* $f\colon \tau \to \tau$ is a continuous map from a graph τ to itself that sends vertices to vertices, and is a local embedding on edges. A graph map assigns to each edge $e \in \tau$ an edge path $f(e) = e_1 \cdots e_\ell$ with no backtracking. Identify the fundamental group $\pi_1(\tau)$ with a free group F_n. A graph map f *represents* an element $\phi \in \mathrm{Out}(F_n)$ if ϕ is conjugate to f_* as an element of $\mathrm{Out}(F_n)$.

Remark 4.3 In many definitions of a graph map one is also allowed to collapse an edge, but for this exposition, graph maps send edges to nonconstant edge-paths.

Definition 4.4 A graph map $f\colon \tau \to \tau$ is a *train-track map* if:

(i) $f^k(e)$ has no back-tracking for all edges e of τ and $k \geq 1$.

(ii) f is a homotopy equivalence.

Definition 4.5 Given a train-track map $f\colon \tau \to \tau$, let $\{e_1, \dots, e_k\}$ be an ordering of the oriented edges of τ, and let D_f be the digraph whose vertices v_e correspond to the undirected edges e of τ, and whose edges from e_i to e_j correspond to each appearance of e_j and $\overline{e}_j$ in the edgepath $f(e_i)$. The *transition matrix* M_f of D_f is the directed adjacency matrix

$$M_f = [a_{i,j}],$$

where $a_{i,j}$ is equal to the number of edges from v_{e_i} to v_{e_j}.

Definition 4.6 If $f\colon \tau \to \tau$ be a train-track map, the *dilatation* of f is given by the spectral radius of M_f

$$\lambda(f) = \max\{|\mu| \mid \mu \text{ is an eigenvalue of } M_f\}.$$

Definition 4.7 A train-track map $f\colon \tau \to \tau$ is *irreducible* if its transition matrix M_f is irreducible, it is *expanding* if the lengths of edges of τ under iterations of f are unbounded.

Remark 4.8 A Perron–Frobenius matrix is irreducible and expanding, but the converse is not necessarily true.

Example 4.9 Let τ be the rose with four petals a, b, c and d. Let $f\colon \tau \to \tau$ be the train-track map associated to the free group automorphism

$$\begin{aligned} a &\mapsto cdc, & b &\mapsto cd, \\ c &\mapsto aba, & d &\mapsto ab. \end{aligned} \tag{8}$$

The train-track map f has transition matrix

$$M_f = \begin{bmatrix} 0 & 0 & 2 & 1 \\ 0 & 0 & 1 & 1 \\ 2 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix},$$

which is an irreducible matrix, and hence f is irreducible. The train-track map is expanding, since its square is block diagonal, where each block is a 2×2 Perron–Frobenius matrix. On the other hand, f is clearly not PF, since no power of M_f is positive.

Definition 4.10 Fix a generating set $\Omega = \{\omega_1, \dots, \omega_n\}$ of F_n. Then each $\gamma \in F_n$ can be written as a *word* in Ω,

$$\gamma = \omega_{i_1}^{r_1} \cdots \omega_{i_\ell}^{r_\ell}, \tag{9}$$

where $\omega_{i_1}, \dots, \omega_{i_\ell} \in \Omega$ and $r_j \in \{1, -1\}$. The word is *reduced* if there are no cancelations, that is $\omega_{i_j}^{r_j} \neq \omega_{i_{j+1}}^{-r_{j+1}}$ for $j = 1, \dots, \ell-1$. The word length $\ell_\Omega(\gamma)$ is the length ℓ of a reduced word representing γ in F_n. The *cyclically reduced word length* $\ell_{\Omega,\mathrm{cyc}}(\gamma)$ of γ represented by the word in (9) is the minimum word length of the elements

$$\gamma_j = \omega_{i_j}^{r_j} \omega_{i_{j+1}}^{r_{j+1}} \cdots \omega_{i_\ell}^{r_\ell} \omega_{i_1}^{r_1} \cdots \omega_{i_{j-1}}^{r_{j-1}},$$

for $j = 1, \dots, \ell-1$.

Proposition 4.11 *Let $\phi \in \mathrm{Out}(F_n)$ be represented by an expanding irreducible train-track map f, and let $\gamma \in F_n$ be a nontrivial element. Then either ϕ acts periodically on the conjugacy class of γ in F_n, or the growth rate satisfies*

$$\lambda_{\Omega,\mathrm{cyc}}(\gamma) = \lim_k \ell_{\Omega,\mathrm{cyc}}(\phi^k(\gamma))^{1/k} = \lambda(f),$$

and in particular, it is independent of the choice of generators, and of γ.

Proof See, for example [2, Remark 1.8]. □

In light of Proposition 4.11, we make the following definition.

Definition 4.12 Let $\phi \in \mathrm{Out}(F_n)$ be an element that is represented by an expanding irreducible train-track map f. Then we define the *dilatation* of ϕ to be

$$\lambda(\phi) = \lambda(f).$$

Remark 4.13 An element $\phi \in \mathrm{Out}(F_n)$ is *hyperbolic* if $F_n \rtimes_\phi \mathbb{Z}$ is word-hyperbolic. It is *atoroidal* if there are no periodic conjugacy classes of elements of F_n under iterations of ϕ. By a result of Brinkmann [4], ϕ is hyperbolic if and only if ϕ is atoroidal.

Definition 4.14 An automorphism $\phi \in \mathrm{Out}(F_n)$ is *reducible* if ϕ leaves the conjugacy class of a proper free factor in F_n fixed. If ϕ is not reducible it is called *irreducible*. If ϕ^k is irreducible for all $k \geq 1$, then ϕ is *fully irreducible*.

Theorem 4.15 (Bestvina and Handel [2]) *If $\phi \in \mathrm{Out}(F_n)$ is irreducible, then ϕ can be represented by an irreducible train track map, and if ϕ is fully irreducible, then it can be represented by a* PF *train track map.*

Remark 4.16 Theorem A deals with an automorphism ϕ that can be represented by an irreducible and expanding train-track map. It does not follow that for such an automorphism every train-track representative is expanding and irreducible. For example, consider the automorphism ϕ from Example 4.9. Let τ' be a graph constructed from an edge e with two distinct endpoints v and w by attaching at v two loops labeled a and b and attaching at w two loops c and d. The map $f'\colon \tau' \to \tau'$ defined by (8) and $e \mapsto \overline{e}$ represents the same automorphism ϕ as in Example 4.9. However, since e is invariant, the map is not irreducible and not expanding.

If we assume that ϕ is fully irreducible, then all train-track representatives are expanding. Indeed, let $f'\colon \tau' \to \tau'$ be a train-track representative of ϕ. Then f' is irreducible because an invariant subgraph will produce a ϕ–invariant free factor. It is now enough to show that some edge is expanding. Let α be an embedded loop in τ'. We can think of α as a conjugacy class in F_n. Then by Proposition 4.11 either α is periodic or α grows exponentially. However, α cannot be periodic since α represents a free factor of F_n. Therefore, α grows exponentially, hence some edge grows exponentially and because f' is irreducible, all edges grow exponentially.

4.2 The mapping torus of a train-track map

In this section we define the branched surface $(X_f, \mathfrak{C}_f, \psi_f)$ associated to an irreducible expanding train-track map f.

Definition 4.17 The *mapping torus* (Y_f, ψ_f) associated to $f\colon \tau \to \tau$ is the branched surface where Y_f is the quotient of $\tau \times [0,1]$ by the identification $(t,1) \sim (f(t),0)$, and ψ_f is the semiflow induced by the product structure of $\tau \times [0,1]$. Write

$$q\colon \tau \times [0,1] \to Y_f$$

for the quotient map. The map to the circle induced by projecting $\tau \times [0,1]$ to the second coordinate induces a map $\rho\colon Y_f \to S^1$.

Definition 4.18 The ψ_f–compatible cellular decomposition $\mathfrak{C}_f$ for Y_f is defined as follows. For each edge e, let v_e be the initial vertex of e (the edges e are oriented by the orientation on τ). The 0–cells of $\mathfrak{C}_f$ are $q(v_e \times \{0\})$, the 1–cells are of the form $s_e = q(v_e \times [0,1])$ or $t_e = q(e \times \{0\})$, and the 2–cells are $c_e = q(e \times [0,1])$, where e ranges over the oriented edges of τ. For this cellular decomposition of Y_f, the collection $\mathcal{V}$ of s_e is the set of *vertical* 1–cells and the collection $\mathcal{E}$ of 1–cells t_e is the set of *horizontal* 1–cells.

By this definition $(Y_f, \mathfrak{C}_f, \psi_f)$ is a branched surface. Let $\theta_{Y_f,\mathfrak{C}_f,\psi_f}$ be the associated cycle function (Definition 3.5).

Proposition 4.19 *The digraph D_f for the train-track map f and the dual digraph of $(Y_f, \mathfrak{C}_f, \psi_f)$ are the same, and we have*

$$\lambda(\phi) = \left|\theta^{(\alpha)}_{Y_f,\mathfrak{C}_f,\psi_f}\right|,$$

where $\alpha\colon \Gamma \to \mathbb{Z}$ is the projection associated to ϕ.

Proof Each 2–cell c of $(Y_f, \mathfrak{C}_f, \psi_f)$ is the quotient of one that can be drawn as in Figure 9, and hence there is a one-to-one correspondence between 2–cells and edges of τ. One can check that for each time $f(e)$ passes over the edge e_i, there is a corresponding hinge between the cell $q(e \times [0,1])$ and the cell $q(e_i \times [0,1])$. This gives a one-to-one correspondence between the directed edges of D_f and the edges of the dual digraph.

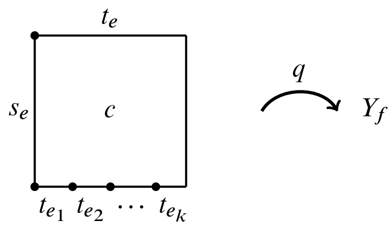

Figure 9: A cell of the mapping torus of a train-track map

Recall that $\lambda(\phi) = \lambda(f)$ is the spectral radius of M_f (Definition 4.12). By Theorem 2.5, the characteristic polynomial of D_f satisfies

$$P_{D_f}(x) = x^m \theta_{D_f}(x).$$

Each edge of D_f has length one with respect to the map α, and hence for each cycle $\sigma \in C_{D_f}$, the number of edges in σ equals $\ell_\alpha(\sigma)$. It follows that $\theta_{D_f}(x)$ is the specialization by α of the cycle function $\theta_{Y_f,\mathfrak{C}_f,\psi_f}$, and we have

$$\lambda_{\mathrm{PF}}(D_f) = |P_{D_f}| = |\theta_{D_f}| = \left|\theta^{(\alpha)}_{Y_f,\mathfrak{C}_f,\psi_f}\right|. \qquad \square$$

In the following sections, we study the behavior of $|\theta^{(\alpha)}_{Y_f,\mathfrak{C}_f,\psi_f}|$ as we let α vary in $\mathrm{Hom}(\Gamma;\mathbb{R})$.

4.3 Application of McMullen's theorem to cycle polynomials

Fix a train-track map $f\colon \tau \to \tau$. Recall: $\theta_f = \theta_{Y_f,\mathfrak{C}_f,\pi_f} = 1 + \sum_{\sigma \in \mathcal{C}^G_{D_f}} (-1)^{|\sigma|} g(\sigma)^{-1}$. Thus the McMullen cone $\mathcal{T}_{\theta_f}(1)$ is given by

$$\begin{aligned}\mathcal{T}_{\theta_f}(1) &= \{\alpha \in \mathrm{Hom}(G;\mathbb{R}) \mid \alpha(g) > 0 \text{ for all } g \in \mathrm{Supp}(\theta)\}\\ &= \{\alpha \in \mathrm{Hom}(G;\mathbb{R}) \mid \alpha(g) > 0 \text{ for all } g \in G \text{ such that } a_g \neq 0\}\end{aligned}$$

(see Definition 2.6). We write $\mathcal{T}_f = \mathcal{T}_{\theta_f}(1)$ for simplicity when the choice of cone associated to θ_f is understood.

Proposition 4.20 *Let $\mathcal{T}_f$ be the McMullen cone for θ_f. The map*

$$\delta\colon \mathrm{Hom}(G;\mathbb{R}) \to \mathbb{R}$$

defined by

$$\delta(\alpha) = \log|\Theta^{(\alpha)}|,$$

extends to a homogeneous of degree -1, real analytic, convex function on $\mathcal{T}_f$ that goes to infinity toward the boundary of affine planar sections of $\mathcal{T}_f$. Furthermore, θ_f has a factor Θ with the properties:

(1) *For all $\alpha \in \mathcal{T}_f$,*
$$\left|\theta_f^{(\alpha)}\right| = |\Theta^{(\alpha)}|.$$

(2) *The polynomial Q is minimal, ie if $\theta \in \mathbb{Z}G$ satisfies $|\theta^{(\alpha)}| = |\theta_f^{(\alpha)}|$ for all α ranging among the integer points of an open subcone of $\mathcal{T}_f$, then Θ divides θ.*

To prove Proposition 4.20 we write $G = H \times \langle s\rangle$ and identify θ_f with the characteristic polynomial P_f of an expanding H–matrix M_f. Then Proposition 4.20 follows from Theorem 2.18.

Let

$$G = H_1(Y_f;\mathbb{Z})/\mathrm{torsion} = \Gamma^{\mathrm{ab}}/\mathrm{torsion},$$

and let H be the image of $\pi_1(\tau)$ in G induced by the composition

$$\tau \to \tau \times \{0\} \hookrightarrow \tau \times [0,1] \xrightarrow{q} Y_f.$$

Let $\rho_*\colon G \to \mathbb{Z}$ be the map corresponding to $\rho\colon Y_f \to S^1$.

Lemma 4.21 *The group G has decomposition as $G = H \times \langle s\rangle$, where $\rho_*(s) = 1$.*

Proof The map ρ_* is onto $\mathbb{Z}$ and its kernel equals H. Take any $s \in \rho_*^{-1}(1)$. Then since $s \notin H$, and G/H is torsion free, we have $G = H \times \langle s\rangle$. $\square$

We call s a *vertical generator* of G with respect to ρ, and identify $\mathbb{Z}G$ with the ring of Laurent polynomials $\mathbb{Z}H(u)$ in the variable u with coefficients in $\mathbb{Z}H$, by the map $\mathbb{Z}G \to \mathbb{Z}H(u)$ determined by sending $s \in \mathbb{Z}G$ to $u \in \mathbb{Z}H(u)$.

Definition 4.22 Given $\theta \in \mathbb{Z}G$, the *associated polynomial* $P_\theta(u)$ of θ is the image of θ in $\mathbb{Z}H(u)$ defined by the identification $\mathbb{Z}G = \mathbb{Z}H(u)$.

The definition of support for an associated polynomial P_θ is analogous to the one for θ.

Definition 4.23 The *support* of an element $P_\theta \in \mathbb{Z}H(u)$ is given by

$$\mathrm{Supp}(P_\theta) = \{hu^r \mid h, r \text{ are such that } (h, s^r) \in \mathrm{Supp}(\theta)\}.$$

Let $P_{\theta_f} \in \mathbb{Z}H(u)$ be the Laurent polynomial associated to θ_f. Instead of realizing P_{θ_f} directly as a characteristic polynomial of an H–labeled digraph, we start with a more natural labeling of the digraph D_f.

Let $C_1 = \mathbb{Z}^{\mathcal{V}\cup\mathcal{E}}$ be the free abelian group generated by the positively oriented edges of Y_f, which we can also think of as 1–chains in $\mathfrak{C}^{(1)}$ (see Definition 4.18). Let $Z_1 \subseteq C_1$ be the subgroup corresponding to closed 1–chains. The map ρ induces a homomorphism $\rho_*\colon C_1 \to \mathbb{Z}$.

Let $\nu\colon Z_1 \to G$ be the quotient map. The map ν determines a ring homomorphism

$$\begin{aligned}\nu_*\colon \mathbb{Z}Z_1 &\to \mathbb{Z}G,\\ \sum_{g\in Z_1} a_g g &\mapsto \sum_{g\in J} a_g \nu(g).\end{aligned}$$

This extends to a map from $\mathbb{Z}Z_1(u)$ to $\mathbb{Z}G(u)$.

Let $K_1 \subseteq Z_1$ be the kernel of $\rho_*|_{Z_1}\colon Z_1 \to \mathbb{Z}$. Then H is the subgroup of G generated by $\nu(K_1)$. Let ν^H be the restriction of ν to K_1. Then ν^H similarly defines

$$\nu_*^H\colon \mathbb{Z}K_1 \to \mathbb{Z}H,$$

the restriction of ν_* to $\mathbb{Z}K_1$, and this extends to

$$\nu_*^H\colon \mathbb{Z}K_1(u) \to \mathbb{Z}H(u).$$

Proposition 4.24 *There is a Perron–Frobenius K_1–matrix $M_f^{K_1}$, whose characteristic polynomial $P_f^{K_1}(u) \in \mathbb{Z}K_1[u]$ satisfies*

$$P_{\theta_f}(u) = u^{-m}\nu_*^H P_f^{K_1}(u).$$

To construct $M_f^{K_1}$, we define a K_1–labeled digraph with underlying digraph D_f. Let s be a vertical generator relative to ρ_*. Choose any element $s' \in Z_1$ mapping to the vertical generator $s \in G$. Write each $s_e \in \mathcal{V}$ as $s_e = s'k_e$, where $k_e \in K_1$. Label edges of the digraph D_f by elements of C_1 as follows. Let $f(e) = e_1 \cdots e_r$. Then for each $i = 1, \dots, r$, there is a corresponding hinge κ_i whose initial cell corresponds to e and whose terminal cell corresponds to e_i. Take any edge η on D_f emanating from v_e. Then η corresponds to one of the hinges κ_i, and has initial vertex v_e and terminal vertex v_{e_i}. For such an η, define

$$g(\eta) = s_e t_{e_1} \cdots t_{e_{i-1}} = s' k_e t_{e_1} \cdots t_{e_{i-1}} = s' k(\eta),$$

where $k(\eta) \in K_1$. This defines a map from the edges D_f to C_1 giving a labeling $\mathcal{D}_f^{C_1}$. It also defines a map from edges of D_f to K_1 by $\eta \mapsto k(\eta)$. Denote this labeling of D_f by $D_f^{K_1}$.

Definition 4.25 Given a labeled digraph $\mathcal{D}^G$, with edge labels $g(\varepsilon)$ for each edge ε of the underlying digraph D, the *conjugate digraph* $\widehat{\mathcal{D}}^G$ of $\mathcal{D}^G$ is the digraph with same underlying graph D, and edge labels $g(\varepsilon)^{-1}$ for each edge ε of D.

Let $\widehat{\mathcal{D}}_f^{K_1}$ be the conjugate digraph of $\mathcal{D}_f^{K_1}$, and let $\widehat{M}_f^{K_1}$ be the directed adjacency matrix for $\widehat{\mathcal{D}}_f^{K_1}$.

Lemma 4.26 *The cycle function $\theta_f \in \mathbb{Z}G$ and the characteristic polynomial $\widehat{P}_f(u) \in \mathbb{Z}K_1[u]$ of $\widehat{M}^{K_1}$ satisfy*

$$\nu_*^H(\widehat{P}_f(u)) = u^m P_{\theta_f}(u).$$

Proof By the coefficient theorem for labeled digraphs (Theorem 2.14) we have

$$\widehat{P}_f(u) = u^m \theta_{\mathcal{D}_f^{K_1}} = u^m \biggl(1 + \sum_{\sigma \in C_{D_f}} (-1)^{|\sigma|} k(\sigma)^{-1} u^{-\ell(\sigma)}\biggr).$$

Since $g(\sigma) = k(\sigma) s^{\ell(\sigma)}$, a comparison of $\widehat{P}_f$ with θ_f gives the desired result. □

Proof of Proposition 4.20 Let M_f be the matrix with entries in $\mathbb{Z}H$ given by taking $\widehat{M}^{K_1}$ and applying ν^H to its entries. Then the characteristic polynomial P_f of M_f is related to the characteristic polynomial $\widehat{P}_f$ of $\widehat{M}^{K_1}$ by

$$P_f(u) = \nu_*^H(\widehat{P}_f(u)).$$

Thus, Lemma 4.26 implies

$$P_f(u) = u^m P_{\theta_f}(u),$$

and hence the properties of Theorem 2.16 applied to $\widehat{P}_f$ also hold for θ_f. □

5 The folded mapping torus and its DKL-cone

We start this section by defining a folded mapping torus and stating some results of Dowdall, Kapovich and Leininger on deformations of free group automorphisms. We then proceed to finish the proof of the main theorem.

5.1 Folding maps

In [15] Stallings introduced the notion of a folding decomposition of a train-track map.

Definition 5.1 (Stallings [15]) Let τ be a topological graph, and v a vertex on τ. Let e_1, e_2 be two distinct edges of τ meeting at v, and let q_1 and q_2 be their other endpoints. Assume that q_1 and q_2 are distinct vertices of τ. The *fold* of τ at v, is the image τ_1 of a quotient map $\mathsf{f}_{(e_1,e_2:v)}\colon \tau \to \tau_1$ where q_1 and q_2 are identified as a single vertex in τ_1 and the two edges e_1 and e_2 are identified as a single edge in τ_1. The map $\mathsf{f}_{(e_1,e_2:v)}$ is called a *folding map*.

It is not hard to check the following.

Lemma 5.2 *Folding maps on graphs are homotopy equivalences.*

Definition 5.3 A *folding decomposition* of a graph map $f\colon \tau \to \tau$ is a decomposition

$$f = hf_k \cdots f_1,$$

where $f_i\colon \tau_{i-1} \to \tau_i$ for $i = 1, \dots, k$ are folding maps on a sequence of graphs $\tau_0, \dots, \tau_k$, where τ_0 is obtained by edge subdivision from τ, and $h\colon \tau_k \to \tau_k$ is a homeomorphism. We denote the folding decomposition by $(f_1, \dots, f_k; h)$.

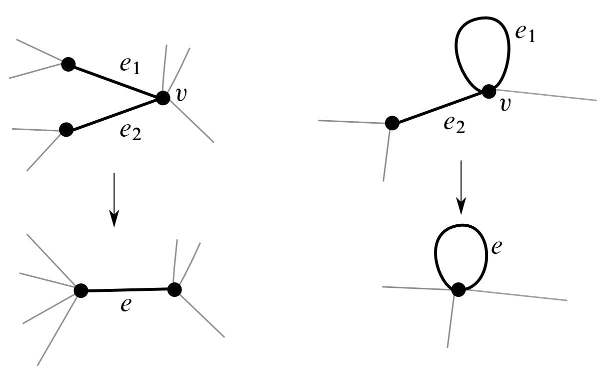

Figure 10: Two examples of folding maps

Lemma 5.4 (Stallings [15]) *Every homotopy equivalence of a graph to itself has a (nonunique) folding decomposition. Moreover, the homeomorphism at the end of the decomposition is uniquely determined.*

Decompositions of a train-track map into a composition of folding maps gives rise to a branched surface that is homotopy equivalent to Y_f.

Let $f\colon \tau \to \tau$ be a train-track map with a folding decomposition $\mathsf{f} = (f_1, \dots, f_k; h)$, where $f_i\colon \tau_{i-1} \to \tau_i$ is a folding map, for $i = 1, \dots, k$, $\tau = \tau_0 = \tau_k$, and $h\colon \tau \to \tau$ is a homeomorphism.

For each $i = 0, \dots, k$, define a 2–complex X_i and semiflow ψ_i as follows. Say f_i is the folding map on τ folding e_1 onto e_2 at their common endpoint v. Let q be the initial vertex of both e_1 and e_2, and q_i the terminal vertex of e_i. Let X_i be the quotient of $\tau_{i-1} \times [0,1]$ obtained by identifying the triangles

$$[(q,0),(q,1),(q_1,1)] \quad \text{on} \quad e_1 \times [0,1]$$

with

$$[(q,0),(q,1),(q_2,1)] \quad \text{on} \quad e_2 \times [0,1].$$

The semiflow ψ_i is defined by the second coordinate of $\tau_{i-1} \times [0,1]$. By the definitions, the image of $\tau_{i-1} \times \{1\}$ in X_i under the quotient map is τ_i.

Let X_{f} be the union of pieces $X_0 \cup \dots \cup X_k$ so that the image of $\tau_{i-1} \times \{1\}$ in X_{i-1} is attached to the image of $\tau_i \times \{0\}$ in X_i by their identifications with τ_i, and the image of $\tau_k \times \{1\}$ in X_k is attached to the image of $\tau_0 \times \{0\}$ in X_0 by h.

Each X_i has a semiflow induced by its structure as the quotient of $\tau_i \times [0,1]$. This induces a semiflow ψ_{f} on X_{f}. The cellular structure on X_{f} is defined so that the 0–cells correspond to the images in X_i of $(q,0), (q,1), (q_1,1)$ and $(q_2,1)$. The transversal 1–cells of $\mathfrak{C}_{\mathsf{f}}$ correspond to the images in X_i of edges $[(q,0),(q_i,1)]$, for $i = 1, 2$. The vertical 1–cells of $\mathfrak{C}_{\mathsf{f}}$ are the forward flows of all the vertices of X_{f}. The vertical and transversal 1–cells form the boundaries of the 2–cells of $\mathfrak{C}_{\mathsf{f}}$.

Definition 5.5 (cf [6]) A *folded mapping torus* associated to a folding decomposition f of a train-track map is the branched surface $(X_{\mathsf{f}}, \mathfrak{C}_{\mathsf{f}}, \psi_{\mathsf{f}})$ defined above.

Lemma 5.6 *If $(X_{\mathsf{f}}, \mathfrak{C}_{\mathsf{f}}, \psi_{\mathsf{f}})$ is a folded mapping torus, then there is a cellular decomposition of X_{f} so that the following holds:*

- (i) *The 1–skeleton $\mathfrak{C}_{\mathsf{f}}^{(1)}$ is a union of oriented 1–cells meeting only at their endpoints.*
- (ii) *Each 1–cell has a distinguished orientation so that the corresponding tangent directions are either tangent to the flow (*vertical case*) or positive but skew to the flow (*diagonal case*).*
- (iii) *The endpoint of any vertical 1–cell is the starting point of another vertical 1–cell.*

Proof The cellular decomposition of $X_{\mathfrak{f}}$ has transversal 1–cells corresponding to the folds, and vertical 1–cells corresponding to the flow suspensions of the endpoints of the diagonal 1–cells. □

5.2 Simple example

We give a simple example of a train-track map, a folding decomposition and their associated branched surfaces.

Consider the train-track in Figure 11, and the train-track map corresponding to the free group automorphism $\phi \in \mathrm{Out}(F_2)$ defined by

$$a \mapsto ba,$$
$$b \mapsto bab.$$

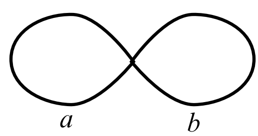

Figure 11: Two petal rose

Then the corresponding train-track map $f\colon \tau \to \tau$ sends the edge a over b and a, and the edge b over b then a then b. The corresponding mapping torus is shown on the left of Figure 12.

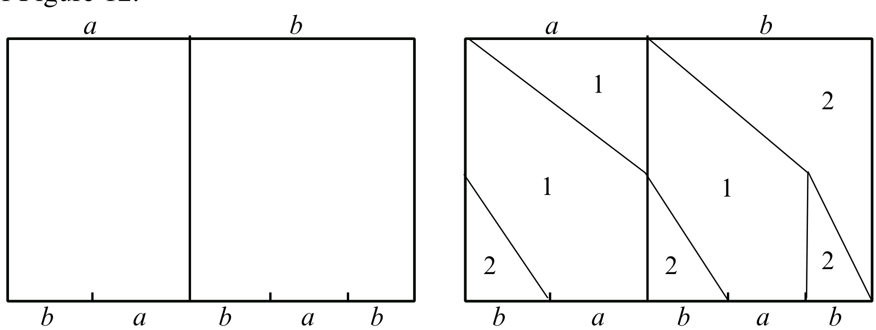

Figure 12: Mapping torus and folded mapping torus

A folding decomposition is obtained from f by subdividing the edge a twice and the edge b three times. The first fold identifies the entire edge a with two segments of the edge b. This yields a train-track that is homeomorphic to the original. The second fold identifies the edge b to one segment of the edge a. The resulting folded mapping torus is shown on the right of Figure 12. Here cells labeled with the same number are identified.

5.3 Dowdall–Kapovich–Leininger's theorem

We first recall that the elements $\alpha \in H^1(X_{\mathsf{f}};\mathbb{R})$ can be represented by cocycle classes $z\colon H_1(X_{\mathsf{f}};\mathbb{R}) \to \mathbb{R}$.

Definition 5.7 Given a branched surface $\mathsf{X} = (X_{\mathsf{f}}, \mathfrak{C}_{\mathsf{f}}, \psi_{\mathsf{f}})$, orient the edges of $\mathfrak{C}_{\mathsf{f}}$ positively with respect to the semiflow ψ_{f}. The associated *positive cone* for X in $H^1(X;\mathbb{R})$, denoted $\mathcal{A}_{\mathsf{f}}$, is given by

$$\mathcal{A}_{\mathsf{f}} = \{\alpha \in H^1(X_{\mathsf{f}};\mathbb{R}) \mid \text{there is a } z \in \alpha \text{ so that } z(e) > 0 \text{ for all } e \in \mathfrak{C}_{\mathsf{f}}^{(1)}\}.$$

Theorem 5.8 (Dowdall, Kapovich and Leininger [6]) *Let f be an expanding irreducible train-track map, f a folding decomposition of f and $(X_{\mathsf{f}}, \mathfrak{C}_{\mathsf{f}}, \psi_{\mathsf{f}})$ the folded mapping torus associated to f. For every integral $\alpha \in A_{\mathsf{f}}$ there is a continuous map $\eta_\alpha\colon X_{\mathsf{f}} \to S^1$ with the following properties.*

1. *Identifying $\pi_1(X_{\mathsf{f}})$ with Γ and $\pi_1(S^1)$ with $\mathbb{Z}$, $(\eta_\alpha)_* = \alpha$.*
2. *The restriction of η_α to a semiflow line is a local diffeomorphism. The restriction of η_α to a flow line in a 2–cell is a nonconstant affine map.*
3. *For all simple cycles c in X_{f} oriented positively with respect to the flow, $\ell(\eta_\alpha(c)) = \alpha[c])$ where $[c]$ is the image of c in G.*
4. *Suppose $x_0 \in S^1$ is not the image of any vertex, denote $\tau_\alpha := \eta_\alpha{-}1(x_0)$. If α is primitive τ_α is connected, and $\pi_1(\tau_\alpha) \cong \ker(\alpha)$.*
5. *For every $p \in \tau_\alpha \cap (\mathfrak{C}_{\mathsf{f}})^{(1)}$, there is an $s \geq 0$ so that $\psi(p,s) \in (\mathfrak{C}_{\mathsf{f}})^{(0)}$.*
6. *The flow induces a map of first return $f_\alpha\colon \tau_\alpha \to \tau_\alpha$, which is an expanding irreducible train-track map.*
7. *The assignment that associates to a primitive integral $\alpha \in A_{\mathsf{f}}$ the logarithm of the dilatation of f_α can be extended to a continuous and convex function on A_{f}.*

Proof This is a compilation of results of [6]. □

5.4 The proof of main theorem

In this section, we prove Theorem A. A crucial step to our proof is that the mapping torus $\mathsf{Y} = (Y_f, \mathfrak{C}_f, \psi_f)$ and the folded mapping torus $\mathsf{X} = (X_{\mathsf{f}}, \mathfrak{C}_{\mathsf{f}}, \psi_{\mathsf{f}})$ both have the same cycle polynomial.

Proposition 5.9 *The cycle functions θ_{Y} of $(Y_f, \mathfrak{C}_f, \psi_f)$ and θ_{X} of $(X_{\mathsf{f}}, \mathfrak{C}_{\mathsf{f}}, \psi_{\mathsf{f}})$ coincide.*

Proof We observe that $(X_{\mathfrak{f}}, \mathfrak{C}_{\mathfrak{f}}, \psi_{\mathfrak{f}})$ can be obtained from the mapping torus of the train-track map $(Y_f, \mathfrak{C}_f, \psi_f)$ by a sequence of folds, vertical subdivisions and transversal subdivision, as defined in Sections 3.2 and 3.3. The reverse of these folds is shown in Figure 13.

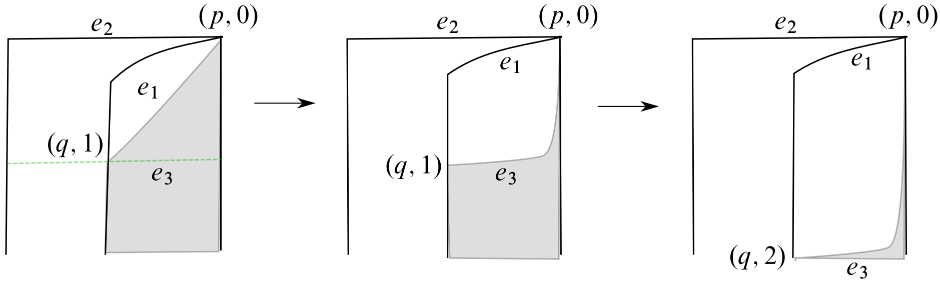

Figure 13: Vertical unfolding

The proposition now follows from Proposition 3.7 and Lemmas 3.15 and 3.20. $\square$

Proposition 5.10 *Let $\theta_{\mathfrak{f}}$ be the cycle polynomial of the* DKL *mapping torus. Then*

$$\mathcal{A}_{\mathfrak{f}} \subseteq \mathcal{T}_{\theta_{\mathfrak{f}}}(1).$$

Proof We need to show that, for every $\sigma \in \mathcal{C}X_{\mathfrak{f}}$ with $|\sigma| = 1$, we have $\alpha(g(\sigma)) > 0$. Then for all nontrivial $g \in \mathrm{Supp}(\theta_{\mathfrak{f}})$, we have $\alpha(g) > 0$, and hence $\alpha \in \mathcal{T}_{\theta_{\mathfrak{f}}}(1) = \mathcal{T}$. Let c be a closed loop in D. The embedding of D in $X_{\mathfrak{f}}$ described in Definition 3.4 induces an orientation on the edges of D that is compatible with the flow ψ. For each edge μ of c, item (2) in Theorem 5.8 implies $\ell(\eta_\alpha \mu)) > 0$ and item (3) in Theorem 5.8 implies

$$\alpha([c]) = \ell(\eta_\alpha(c)) = \sum_{\mu \in c} \ell(\eta_\alpha(\mu)) > 0. \qquad \square$$

Proposition 5.11 *Let $(X_{\mathfrak{f}}, \mathfrak{C}_{\mathfrak{f}}, \psi_{\mathfrak{f}})$ be the folded mapping torus, $\theta_{\mathfrak{f}}$ its cycle polynomial and $\mathcal{A}_{\mathfrak{f}}$ the* DKL*–cone. For all primitive integral $\alpha \in A_{\mathfrak{f}}$, we have*

$$\lambda(\phi_\alpha) = \left|\theta_{\mathfrak{f}}^{(\alpha)}\right|.$$

Proof Embed τ_α in $X_{\mathfrak{f}}$ transversally as in Theorem 5.8(4), and perform a vertical subdivision so that the intersections of τ_α with $(X_{\mathfrak{f}})^{(1)}$ are contained in the 0–skeleton (we can do this by Theorem 5.8(5)). Perform transversal subdivisions to add the edges of τ_α to the 1–skeleton. Then perform a sequence of foldings and unfoldings to move the branching of the complex into τ_α, and remove the extra edges. Denote the new branched surface by

$$X_{\mathfrak{f}}^{(\alpha)} = (X_{\mathfrak{f}}, \mathfrak{C}_{\mathfrak{f}}^{(\alpha)}, \psi_{\mathfrak{f}}).$$

These operations preserve the cycle polynomials of the respective 2–complexes, therefore we denote all of these polynomials by θ (in particular $\theta_{\mathsf{f}} = \theta$).

Let $f_\alpha\colon \tau_\alpha \to \tau_\alpha$ be the map induced by the first return map, and D_α its digraph. Then f_α defines a train-track map representing ϕ_α, and $\lambda(\phi_\alpha) = \lambda(D_\alpha)$.

The (unlabeled) digraph $D_{\mathsf{f}}^{(\alpha)}$ of the new branched surface $(X_{\mathsf{f}}, \mathfrak{C}_{\mathsf{f}}^{(\alpha)}, \psi_{\mathsf{f}})$ is identical to D_α. For a cycle c in D_α, let $\ell(c)$ be the number of edges in c. Then $\ell(c)$ equals the number of 1–cells in $\tau_\alpha \cap (X_{\mathsf{f}}^{(\alpha)})^{(1)}$, and by Theorem 5.8, items (4) and (3),

$$\ell(c) = \ell(\eta_\alpha(c)) = \alpha([c]).$$

Thus $\ell(\sigma) = \alpha(g(\sigma))$ for every $\sigma \in C_{D_\alpha}$. Let $P_\alpha(x)$ be the characteristic polynomial of the directed incidence matrix associated to D_α. By the coefficients theorem for digraphs (Theorem 2.5) we have

$$P_\alpha(x) = x^m + \sum_{\sigma \in C_D} (-1)^{|\sigma|} x^{m-\ell(\sigma)} = x^m \Bigl(1 + \sum_{\sigma \in C_{D_\alpha}} (-1)^{|\sigma|} x^{\alpha(g(\sigma))}\Bigr) = x^m \theta^{(\alpha)}.$$

Therefore

$$\lambda(\phi_\alpha) = |P_\alpha| = |\theta^{(\alpha)}|. \qquad \square$$

We are now ready to prove our main result.

Proof of Theorem A Choose an expanding train-track representative f of ϕ, and a folding decomposition f of f. As before, let $\mathsf{Y} = (Y_f, \mathfrak{C}_f, \psi_f)$ be the mapping torus of f, and $\mathsf{X} = (X_{\mathsf{f}}, \mathfrak{C}_{\mathsf{f}}, \psi_{\mathsf{f}})$ the folded mapping torus. By Proposition 5.9 their cycle function $\theta_{\mathsf{Y}}, \theta_{\mathsf{X}}$ are equal, and we will call them θ.

Let Θ be the minimal factor of θ defined in Proposition 4.20, and let $\mathcal{T} = \mathcal{T}_\Theta(1)$ be the McMullen cone. By Proposition 5.10, $A_{\mathsf{f}} \subseteq \mathcal{T}$, and by Proposition 5.11, $\lambda(\phi_\alpha) = |\Theta^{(\alpha)}|$. By Proposition 4.20, $|\Theta_\phi^{(\alpha)}| = |\Theta^{(\alpha)}|$ in $\mathcal{T}$ so we have $\lambda(\alpha) = |\Theta_\phi^{(\alpha)}|$ for all $\alpha \in A_{\mathsf{f}}$. Item (2) of Proposition 4.20 implies part (2) of Theorem A. If f' is another folding decomposition of another expanding irreducible train-track representative f' of ϕ, we get another distinguished factor $\Theta_{\mathsf{f}'}$. Since the cones $\mathcal{T}_{\mathsf{f}}$ and $\mathcal{T}_{\mathsf{f}'}$ must intersect, it follows by the minimality properties of Θ_{f} and $\Theta_{\mathsf{f}'}$ in Proposition 4.20 that they are equal. Item (3) of Proposition 4.20 completes the proof. $\square$

6 Example

In this section, we compute the cycle polynomial for an explicit example, and compare the DKL– and McMullen cones.

Consider the rose with four directed edges a, b, c, d and the map

$$f = \begin{cases} a \to B \to adb, \\ c \to D \to cbd. \end{cases}$$

Capital letters indicate the relevant edge in the opposite orientation to the chosen one. It is well known (see eg the first and third authors [1, Proposition 2.6]) that if $f\colon \tau \to \tau$ is a graph map, and τ is a graph with $2m$ directed edges, and for every edge e of τ, the path $f^{2m}(e)$ does not have back-tracking (see Definition 4.1), then f is a train-track map. One can verify that f is a train-track map.

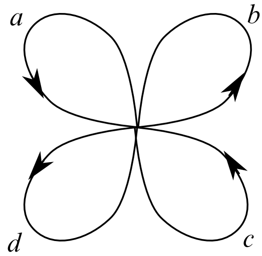

Figure 14: Four petal rose with directed edges

The train-track transition matrix is given by

$$M_f = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 1 \end{bmatrix}.$$

The associated digraph is shown in Figure 15.

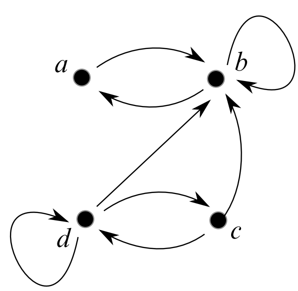

Figure 15: Digraph associated to the train-track map f

The matrix M_f is nonnegative and M_f^3 is positive. Thus M_f is a Perron–Frobenius matrix and f is a PF train-track map. By Theorem 1.1, α_ϕ has an open cone neighborhood, the DKL–cone $\mathcal{A}_{\mathfrak{f}} \subset \mathrm{Hom}(\Gamma; \mathbb{R})$, whose primitive integral elements of

$\mathcal{A}_{\mathfrak{f}}$ correspond to free group automorphisms that can be represented by expanding irreducible train-track maps.

Remark 6.1 The outer automorphism ϕ represented by f is reducible. Consider the free factor $\langle bA, ad, ac\rangle$. Then

$$f(bA) = BDAb, \quad f(ad) = BDBC = aBDAaBCA, \quad f(ac) = BD = bAad.$$

Therefore this factor is invariant up to conjugacy. Although ϕ is reducible, f is expanding and irreducible, and we can apply both Theorem 5.8 and Theorem A.

Identifying the fundamental group of the rose with F_4 we choose the basis a, b, c, d of F_4. The free-by-cyclic group corresponding to $[f_*]$ has the presentation

$$\Gamma = \langle a, b, c, d, s' \mid a^{s'} = B, b^{s'} = BDA, c^{s'} = D, d^{s'} = DBC\rangle.$$

Let $G = \Gamma^{\mathrm{ab}}$ and for $w \in \Gamma$ we denote by $[w]$ its image in G. Then

$$[a] = -[b] = [d] = -[c].$$

Thus $G = \mathbb{Z}^2 = \langle t, s\rangle$ where $t = [a]$ and $s = [s']$. We decompose f into four folds

$$\tau = \tau_0 \xrightarrow{f_1} \tau_1 \xrightarrow{f_2} \tau_2 \xrightarrow{f_3} \tau_3 \xrightarrow{f_4} \tau_4 \cong \tau,$$

where all the graphs τ_i are roses with four petals. f_1 folds all of a with the first third of b, to the edge a_1 of τ_1, the other edges will be denoted b_1, c_1, d_1. f_1 folds the edge c_1 with the first third of the edge d_1. With the same notation scheme, f_2 folds the edge c_2 with half of the edge b_2 and f_3 folds the edge a_3 with half of the edge d_3. Figure 16 shows the folded mapping torus $X_{\mathfrak{f}}$ for this folding sequence.

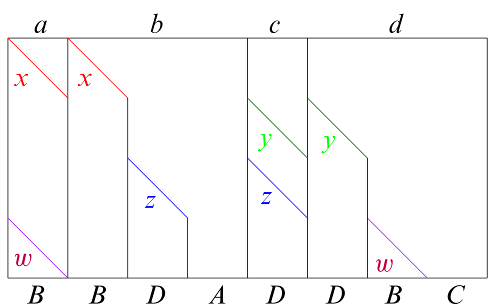

Figure 16: The complex X

The cell structure $\mathfrak{C}_{\mathfrak{f}}$ has 4 vertices, 8 edges: $s_1, s_2, s_3, s_4, x, y, z, w$, and four 2–cells: c_x, c_y, c_z, c_w. The 2–cells are sketched in Figure 17.

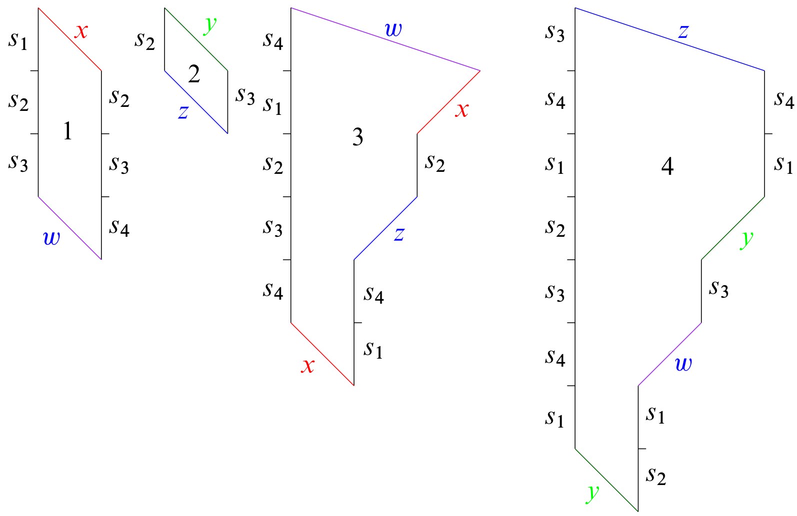

Figure 17: The discs in X

Let C_1 be the free abelian group generated by the edges of $X_{\mathfrak{f}}$, and let F be the maximal tree consisting of the edges s_1, s_2, s_3, then $Z_1 \subset C_1$ is generated by x, y, z, w and $s_1 + s_2 + s_3 + s_4$. The quotient homomorphism $\nu\colon Z_1 \to G$ is given by collapsing the maximal tree and considering the relations given by the two cells. The map is given by $\nu(s_1 + s_2 + s_3 + s_4) = s$ and

$$\nu(x) = t, \quad \nu(y) = \nu(z) = -t, \quad \nu(w) = t + s.$$

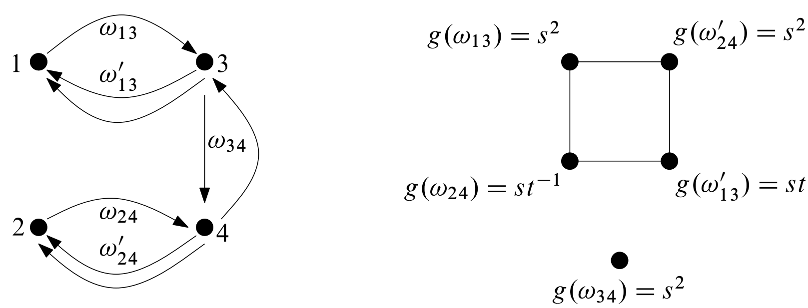

Figure 18: The dual digraph is on the left and the labeled cycle complex is on the right.

The dual digraph D to X is shown on the left of Figure 18. There are five cycles: ω_{13} and ω'_{13} the two distinct cycles containing 1 and 3, ω_{24} and ω'_{24} the two distinct cycles containing 2 and 4, and ω_{34} is the cycle containing 3 and 4. The cycle complex is shown on the right of Figure 18:

$$\begin{aligned}\theta_f &= 1-(s^{-2}+s^{-1}t^{-1}+s^{-2}+s^{-1}t+s^{-2})+(s^{-3}t+s^{-2}+s^{-3}t^{-1}+s^{-4})\\ &= 1+s^{-4}-2s^{-2}-s^{-1}t^{-1}-s^{-1}t+s^{-3}t+s^{-3}t^{-1}.\end{aligned}$$

Note that Θ_ϕ might be a proper factor of this polynomial. However, for the sake of computing the support cone (and the dilatations of ϕ_α for different $\alpha \in \mathcal{A}_\mathsf{f}$) we may use θ_f.

Computing the McMullen cone In order to simplify notation, for $\alpha \in \mathrm{Hom}(G,\mathbb{R})$ and $g \in G$ we denote $g^\alpha = \alpha(g)$. The cone $\mathcal{T}_\phi$ in $H^1(G,\mathbb{R})$ is given by

$$\begin{aligned}\mathcal{T}_\phi &= \{\alpha \in \mathrm{Hom}(G,\mathbb{R}) \mid g^\alpha < 0^\alpha \text{ for all } g \in \mathrm{Supp}(\theta_f)\}\\ &= \{\alpha \in \mathrm{Hom}(G,\mathbb{R}) \mid (-4s)^\alpha(-2s)^\alpha(-s-t)^\alpha < 0, (-s+t)^\alpha(-3s+t)^\alpha(-3s-t)^\alpha < 0\}.\end{aligned}$$

Therefore, the McMullen cone is

$$\mathcal{T}_\phi = \{\alpha \in \mathrm{Hom}(G,\mathbb{R}) \mid s^\alpha > 0 \text{ and } |t^\alpha| < s^\alpha\}. \tag{10}$$

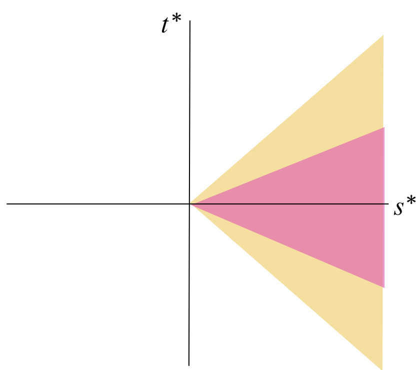

Figure 19: The McMullen cone $\mathcal{T}$ (outer) and DKL–cone $\mathcal{A}_\mathsf{f}$ (inner)

Computing the DKL–cone We now compute the DKL–cone $\mathcal{A}_\mathsf{f}$. A cocycle a represents an element in $\alpha \in \mathcal{A}_\mathsf{f}$ if it evaluates positively on all edges in X_f. We use the notation $\mathsf{a}(e) = e^\mathsf{a}$. Thus for a a positive cocycle we have

$$s_1^\mathsf{a}, s_2^\mathsf{a}, s_3^\mathsf{a} > 0,$$
$$s_4^\mathsf{a} > 0 \Longrightarrow s^\mathsf{a} - s_1^\mathsf{a} + s_2^\mathsf{a} + s_3^\mathsf{a} > 0 \Longrightarrow s^\mathsf{a} > s_1^\mathsf{a} + s_2^\mathsf{a} + s_3^\mathsf{a} > 0.$$

Now by considering the cell structure given by all edges in Figure 19 and recalling that $[a]=[d]=t$ and $[b]=[c]=-t$ we have

$$x = t + s_1, \quad w = t + s_4, \quad y = s_2 - t, \quad z = s_3 - t.$$

The diagonal edges x, w give us

$$0 < x^{\mathsf{a}} = t^{\mathsf{a}} + s_1^{\mathsf{a}} \quad \text{and} \quad 0 < w^{\mathsf{a}} = t^{\mathsf{a}} + s_4^{\mathsf{a}},$$

so

$$t^\alpha - \frac{s_1^{\mathsf{a}} + s_4^{\mathsf{a}}}{2} > -\frac{s^\alpha}{2}.$$

The other diagonal edges give us

$$0 < z^{\mathsf{a}} = s_3^{\mathsf{a}} - t^{\mathsf{a}} \quad \text{and} \quad 0 < y^{\mathsf{a}} = s_2^{\mathsf{a}} - t^{\mathsf{a}},$$

hence

$$t^\alpha < \frac{s_2^{\mathsf{a}} + s_3^{\mathsf{a}}}{2} < \frac{s^\alpha}{2}.$$

We obtain the cone

$$\left\{ s^\alpha > 0 \ \text{ and } \ |t^\alpha| < \frac{s^\alpha}{2} \right\}. \tag{11}$$

If α is in this cone there is a positive cocycle representing α. Therefore $\mathcal{A}_{\mathsf{f}}$ is equal to the cone in (11) and is strictly contained in the cone $\mathcal{T}_\phi$ (see (10) and Figure 19).

References

[1] **Y Algom-Kfir**, **K Rafi**, *Mapping tori of small dilatation irreducible train-track maps* `arXiv:1209.5635`

[2] **M Bestvina**, **M Handel**, *Train tracks and automorphisms of free groups*, Ann. of Math. 135 (1992) 1–51 MR1147956

[3] **R Bieri**, **W D Neumann**, **R Strebel**, *A geometric invariant of discrete groups*, Invent. Math. 90 (1987) 451–477 MR914846

[4] **P Brinkmann**, *Hyperbolic automorphisms of free groups*, Geom. Funct. Anal. 10 (2000) 1071–1089 MR1800064

[5] **D Cvetković**, **P Rowlinson**, *The largest eigenvalue of a graph: A survey*, Linear and Multilinear Alg. 28 (1990) 3–33 MR1077731

[6] **S Dowdall**, **I Kapovich**, **C J Leininger**, *Dynamics on free-by-cyclic groups* `arXiv:1301.7739`

[7] **S Dowdall**, **I Kapovich**, **C J Leininger**, *McMullen polynomials and Lipschitz flows for free-by-cyclic groups* `arXiv:1310.7481`

[8] **D Fried**, *Flow equivalence, hyperbolic systems and a new zeta function for flows*, Comment. Math. Helv. 57 (1982) 237–259 MR684116

[9] **F R Gantmacher**, *The theory of matrices, Vol.* 1, 2, Chelsea Publ., Providence, RI (1959) MR0107649

[10] **I Kapovich**, *Detecting quasiconvexity: Algorithmic aspects*, from: "Geometric and computational perspectives on infinite groups", (G Baumslag, D Epstein, R Gilman, H Short, C Sims, editors), DIMACS Ser. Discrete Math. Theoret. Comput. Sci. 25, Amer. Math. Soc. (1996) 91–99 MR1364182

[11] **S Matsumoto**, *Topological entropy and Thurston's norm of atoroidal surface bundles over the circle*, J. Fac. Sci. Univ. Tokyo Sect. IA Math. 34 (1987) 763–778 MR927609

[12] **C T McMullen**, *Polynomial invariants for fibered* 3*–manifolds and Teichmüller geodesics for foliations*, Ann. Sci. École Norm. Sup. 33 (2000) 519–560 MR1832823

[13] **C T McMullen**, *The Alexander polynomial of a* 3*–manifold and the Thurston norm on cohomology*, Ann. Sci. École Norm. Sup. 35 (2002) 153–171 MR1914929

[14] **W D Neumann**, *Normal subgroups with infinite cyclic quotient*, Math. Sci. 4 (1979) 143–148 MR562003

[15] **J R Stallings**, *Topology of finite graphs*, Invent. Math. 71 (1983) 551–565 MR695906

[16] **W P Thurston**, *A norm for the homology of* 3*–manifolds*, Mem. Amer. Math. Soc. 339, Amer. Math. Soc. (1986) 99–130 MR823443

Mathematics Department, University of Haifa
Mount Carmel, Haifa 31905, Israel

Department of Mathematics, Florida State University
208 Love Building, 1017 Academic Way, Tallahassee, FL 32301-4510, USA

Department of Mathematics, University of Toronto
Room 6290, 40 George Street, Toronto ON M5S 2E4, Canada

algomy@gmail.com, hironaka@math.fsu.edu, rafi@math.toronto.edu

Proposed: Ian Agol Received: 2 February 2014
Seconded: Walter Neumann, Benson Farb Revised: 2 February 2014