



Spencer Dowdall · Ilya Kapovich · Christopher J. Leininger

McMullen polynomials and Lipschitz flows for free-by-cyclic groups

Received December 4, 2014

Abstract. Consider a group G and an epimorphism $u_0: G \rightarrow \mathbb{Z}$ inducing a splitting of G as a semidirect product $\ker(u_0) \rtimes_{\varphi} \mathbb{Z}$ with $\ker(u_0)$ a finitely generated free group and $\varphi \in \text{Out}(\ker(u_0))$ representable by an expanding irreducible train track map. Building on our earlier work [DKL], in which we realized G as $\pi_1(X)$ for an Eilenberg–MacLane 2-complex X equipped with a semi-flow ψ , and inspired by McMullen’s Teichmüller polynomial for fibered hyperbolic 3-manifolds, we construct a polynomial invariant $m \in \mathbb{Z}[H_1(G; \mathbb{Z})/\text{torsion}]$ for (X, ψ) and investigate its properties.

Specifically, m determines a convex polyhedral cone $\mathcal{C}_X \subset H^1(G; \mathbb{R})$, a convex, real-analytic function $\mathfrak{H}: \mathcal{C}_X \rightarrow \mathbb{R}$, and *specializes* to give an integral Laurent polynomial $m_u(\zeta)$ for each integral $u \in \mathcal{C}_X$. We show that \mathcal{C}_X is equal to the “cone of sections” of (X, ψ) (the convex hull of all cohomology classes dual to sections of ψ), and that for each (compatible) cross section $\Theta_u \subset X$ with first return map $f_u: \Theta_u \rightarrow \Theta_u$, the specialization $m_u(\zeta)$ encodes the characteristic polynomial of the transition matrix of f_u . More generally, for *every* class $u \in \mathcal{C}_X$ there exists a geodesic metric d_u and a codimension-1 foliation Ω_u of X defined by a “closed 1-form” representing u transverse to ψ so that after reparametrizing the flow ψ_s^u maps leaves of Ω_u to leaves via a local $e^{s\mathfrak{H}(u)}$ -homothety.

Among other things, we additionally prove that \mathcal{C}_X is equal to (the cone over) the component of the BNS-invariant $\Sigma(G)$ containing u_0 , and consequently each primitive integral $u \in \mathcal{C}_X$ induces a splitting of G as an ascending HNN-extension $G = Q_u *_{\phi_u}$ with Q_u a finite-rank free group and $\phi_u: Q_u \rightarrow Q_u$ injective. For any such splitting, we show that the stretch factor of ϕ_u is exactly given by $e^{\mathfrak{H}(u)}$. In particular, \mathcal{C}_X and \mathfrak{H} depend only on the group G and epimorphism u_0 .

Keywords. Free-by-cyclic groups, BNS-invariant, train track maps, stretch factors

S. Dowdall: Department of Mathematics, Vanderbilt University,
1326 Stevenson Center, Nashville, TN 37240, U.S.A.;
e-mail: spencer.dowdall@vanderbilt.edu, <http://www.math.vanderbilt.edu/~dowdals/>

I. Kapovich: Department of Mathematics, University of Illinois at Urbana-Champaign,
1409 West Green Street, Urbana, IL 61801, USA;
e-mail: kapovich@math.uiuc.edu, <http://www.math.uiuc.edu/~kapovich/>

C. J. Leininger: Department of Mathematics, University of Illinois at Urbana-Champaign,
1409 West Green Street, Urbana, IL 61801, USA;
e-mail: clein@math.uiuc.edu, <http://www.math.uiuc.edu/~clein/>

Mathematics Subject Classification (2010): Primary 20F65; Secondary 57M07, 57M20, 37B99, 37E99

Contents

1. Introduction	3254
2. Background	3266
3. Setup	3273
4. The module of transversals and the McMullen polynomial	3277
5. Cross sections, closed 1-forms, and cohomology	3279
6. A homological characterization of \mathcal{S}	3287
7. The algebra and dynamics of cross sections to ψ	3290
8. Cross sections and homology	3296
9. Graph modules	3300
10. The isomorphism of modules	3302
11. Specialization, characteristic polynomials, and stretch factors	3305
12. Real-analyticity, convexity, and divergence	3310
13. Closed 1-forms and foliations	3319
14. Lipschitz flows	3324
Appendix A. Characterizing sections	3335
Appendix B. Local boundedness of the stretch function	3346
References	3351

1. Introduction

Given an outer automorphism φ of the finite-rank free group F_N , consider the free-by-cyclic group

$$G = G_\varphi = F_N \rtimes_\varphi \mathbb{Z} = \langle F_N, r \mid r^{-1}wr = \Phi(w) \text{ for all } w \in F_N \rangle$$

(here $\Phi \in \text{Aut}(F_N)$ is any representative of φ). The epimorphisms $G \rightarrow \mathbb{Z}$ are exactly the primitive integral points of the vector space $H^1(G; \mathbb{R}) = \text{Hom}(G, \mathbb{R})$, and every such $u: G \rightarrow \mathbb{Z}$ gives rise to a *splitting*

$$1 \rightarrow \ker(u) \rightarrow G \xrightarrow{u} \mathbb{Z} \rightarrow 1$$

of G and an associated *monodromy* $\varphi_u \in \text{Out}(\ker(u))$ generating the outer action of \mathbb{Z} on $\ker(u)$. We are mostly interested in the case where $\ker(u)$ is finitely generated, or more generally, finitely generated over the monoid generated by φ_u . In this case there is a canonically defined stretch factor $\Lambda(u)$ measuring the exponential growth rate of φ_u acting on $\ker(u)$; see §3 for the precise definitions of φ_u and $\Lambda(u)$.

We aim to illuminate the dynamical and topological properties of splittings of G_φ in the case where φ may be represented by an expanding irreducible train track map f . For example, one broad goal is to understand how φ_u and $\Lambda(u)$ vary with u . We began this undertaking in [DKL] where, inspired by the work of Thurston and Fried on hyperbolic 3-manifolds and their fibrations over the circle, we defined a $K(G, 1)$ space called the *folded mapping torus* $X = X_f$ that comes equipped with a semiflow ψ . This determines an open convex cone $\mathcal{A} = \mathcal{A}_f \subset H^1(X; \mathbb{R}) = \text{Hom}(G, \mathbb{R})$ containing the natural projection $u_0: G \rightarrow G/F_N \cong \mathbb{Z}$ associated with the given splitting $G = F_N \rtimes_\varphi \mathbb{Z}$. The cone \mathcal{A} has the property that every primitive integral $u \in \mathcal{A}$ has $\ker(u)$ a finitely generated free group and monodromy φ_u which is again represented by an expanding irreducible train track

map f_u . Moreover, if the original automorphism $\varphi \in \text{Out}(F_N)$ is hyperbolic and fully irreducible, then so is φ_u for every primitive integral $u \in \mathcal{A}$. We also proved in [DKL] that there exists a convex, continuous, degree -1 homogeneous function $\mathfrak{h}: \mathcal{A} \rightarrow \mathbb{R}_+$ such that $\mathfrak{h}(u)$ is equal to the topological entropy of f_u and also to $\log(\Lambda(u))$ for all primitive integral u . We refer the reader to [DKL] for a detailed discussion of this material and other related results such as those in [AKR, BS, Gau1, Gau2, Gau3, Wan].

Comparing our results on G and (X, ψ) from [DKL] with those on a fibered hyperbolic 3-manifold leads to several additional questions. The capstone of the 3-manifold picture is McMullen's *Teichmüller polynomial* [McM1] which encapsulates nearly all the relevant information regarding fibrations of the manifold. In particular, it detects a canonically defined cone—which has topological, dynamical and algebraic significance—and gives specific information about every class in this cone. It is thus tantalizing to ask: Is there an analogous polynomial in the free-by-cyclic setting and, more importantly, what does it tell you about G and (X, ψ) ? Any such analog would naturally determine a cone in $H^1(X; \mathbb{R})$ and a convex, real-analytic function on that cone. What is the relationship of these to the cone \mathcal{A} and function \mathfrak{h} found in [DKL] or to the canonically defined stretch function Λ ? Is there a geometric, algebraic or dynamical characterization of the cone (e.g., in terms of foliations of X , Fried's cone of homology directions [Fri2], or the BNS-invariant)? What, if any, topological, geometric, or dynamical structure is there associated to the *irrational* points of the cone?

The centerpiece of the present paper is the construction of exactly such a polynomial \mathfrak{m} , which we call the *McMullen polynomial*, together with a detailed analysis of \mathfrak{m} that answers all of the above questions. Our multivariate, integral Laurent polynomial \mathfrak{m} is constructed as an element of the integral group ring $\mathbb{Z}[H]$, where $H = H_1(G; \mathbb{R})/\text{torsion} \cong \mathbb{Z}^b$, for $b = b_1(G)$. As such, \mathfrak{m} has a natural *specialization* \mathfrak{m}_u for every $u \in H^1(X; \mathbb{R})$; this is a single variable Laurent polynomial when u is an integral class and a “power sum” in general. We now briefly summarize the main results of the paper and explain the information packaged in \mathfrak{m} . Over the following several pages, we then provide detailed and expanded statements of our results together with discussion of related and other recent results in the literature and a more in depth comparison with the motivating 3-manifold setting.

Convention 1.1. For the entirety of this paper (except §2), $G = G_\varphi$ will denote the free-by-cyclic group associated to an outer automorphism $\varphi \in \text{Out}(F_N)$ represented by an expanding irreducible train track map $f: \Gamma \rightarrow \Gamma$ as above, and $X = X_f$ will denote the folded mapping torus built from f .

Meta-Theorem

I. Canonical cone. *There is an open, convex, rationally defined, polyhedral cone $\mathcal{C}_X \subset H^1(G; \mathbb{R}) = H^1(X; \mathbb{R})$ containing \mathcal{A} , which can be characterized by any of the following:*

- (1) *the dual cone of a distinguished vertex of the Newton polytope of \mathfrak{m} (Theorem E);*
- (2) *the open convex hull of rays through integral $u \in H^1(X; \mathbb{R})$ dual to sections of ψ (Proposition 5.12);*

- (3) $\{u \in H^1(X; \mathbb{R}) \mid u \text{ is represented by a "flow-regular closed 1-form"}\}$ (Definition 5.9);
- (4) $\{u \in H^1(X; \mathbb{R}) \mid u \text{ is positive on all closed orbits of } \psi\}$ (Definition 6.1);
- (5) $\{u \in H^1(X; \mathbb{R}) \mid u \text{ is positive on a specific finite set of closed orbits of } \psi\}$ (Theorem A);
- (6) $\{u \in H^1(X; \mathbb{R}) \mid u \text{ is represented by a cellular 1-cocycle which is positive on a specific set of 1-cells of the "circuitry cell structure" of } X\}$ (Proposition A.2);
- (7) *the cone on the component of the BNS-invariant containing u_0* (Theorem I).

Item (1) shows that the cone \mathcal{C}_X is determined by the McMullen polynomial \mathfrak{m} . Given $u \in \mathcal{C}_X$ primitive integral, the section $\Theta_u \subset X$ of ψ dual to u provided by (2) has first return map we denote $f_u: \Theta_u \rightarrow \Theta_u$. By (7) we can realize G (not necessarily uniquely) as an ascending HNN-extension $G = Q_u *_{\phi_u}$ dual to u over a finitely generated free group Q_u . For any class $u \in \mathcal{C}_X$ (not necessarily integral), the flow-regular closed 1-form provided by (3) determines a "foliation" Ω_u of X .

II. Specialization: entropy and stretch factors. *There is a convex, real-analytic function $\mathfrak{H}: \mathcal{C}_X \rightarrow \mathbb{R}$ that is homogeneous of degree -1 , extends the previously defined function on \mathcal{A} , and tends to ∞ at $\partial \mathcal{C}_X$ such that for every primitive integral $u \in \mathcal{C}_X$:*

- (1) $e^{\mathfrak{H}(u)} > 1$ is equal to the largest root of $\mathfrak{m}_u(\zeta)$.
- (2) $f_u: \Theta_u \rightarrow \Theta_u$ is an expanding irreducible train track map "weakly representing" ϕ_u .
- (3) The characteristic polynomial of the transition matrix for f_u is $\pm \zeta^k \mathfrak{m}_u(\zeta)$ for some $k \in \mathbb{Z}$.
- (4) The entropy of f_u is given by $h(f_u) = \mathfrak{H}(u)$.
- (5) The stretch factor of ϕ_u is $\Lambda(u) = \lambda(\phi_u) = e^{\mathfrak{H}(u)}$.

III. Determinant formula. *For any primitive integral $u \in H^1(X; \mathbb{R})$, the first return map $f_u: \Theta_u \rightarrow \Theta_u$ determines a matrix $A_u(\mathbf{t}) = A_u(t_1, \dots, t_{b-1})$ with entries in the ring $\mathbb{Z}[t_1^{\pm 1}, \dots, t_{b-1}^{\pm 1}]$ of Laurent polynomials, for which $A_u(1, \dots, 1)$ is the transition matrix for f_u and whose characteristic polynomial is \mathfrak{m} :*

$$\mathfrak{m}(x, \mathbf{t}) = \det(xI - A_u(\mathbf{t})) \in \mathbb{Z}[x, \mathbf{t}] \cong \mathbb{Z}[H].$$

IV. Foliations and flows. *From the construction of \mathfrak{m} , for any class $u \in \mathcal{C}_X$ (not necessarily integral), there is a geodesic metric d_u on X and a reparameterization ψ^u of ψ such that:*

- (1) The flowlines $s \mapsto \psi_s^u(\xi)$ are d_u -geodesics for all $\xi \in X$.
- (2) The "leaves" of Ω_u are (possibly infinite) graphs.
- (2) ψ^u maps leaves to leaves.
- (3) The restriction of ψ_s^u to any leaf is a local $e^{s\mathfrak{H}(u)}$ -homothety with respect to the induced path metric.

Some remarks

- We are indebted to Nathan Dunfield who suggested the scheme to prove Meta-Theorem I(7).

- We consider left actions instead of right actions, and so the BNS-invariant discussed here agrees with $\Sigma^1(G; \mathbb{Z})$ from [BR], but is $-\Sigma_{G'}(G)$ from [BNS] (see [BR, Section 1.3]).
- Despite the ostensible dependence on the dynamical system (X, ψ) , Meta-Theorem I–II shows that \mathcal{C}_X and \mathfrak{H} in fact depend only on the group G and the initial cohomology class $u_0: G \rightarrow \mathbb{Z}$.
- We record that Theorem A, Proposition 5.10, and Proposition A.2 give the precise meaning to, and prove the equivalence of, (2)–(6) in Meta-Theorem I. Meanwhile Theorems E and I prove the equivalence of these with (1) and (7). Meta-Theorem II follows from Theorems B, C, and F. Meta-Theorem III is just Theorem D, and Meta-Theorem IV is an abbreviated version of Theorem H.

Remark 1.2. Contemporaneously with the release of this paper, Algom-Kfir, Hironaka, and Rafi [AKHR] independently introduced a related polynomial $\Xi \in \mathbb{Z}[H]$ that provides information about G and X . We stress that the results of the present paper (outlined above) have very little overlap with those of [AKHR]. Specifically, the polynomial Ξ similarly determines a cone $\mathcal{T} \subset H^1(G; \mathbb{R})$ and a convex, real-analytic function $L: \mathcal{T} \rightarrow \mathbb{R}$. Algom-Kfir, Hironaka, and Rafi prove that \mathcal{A} (the cone defined in [DKL]) is contained in \mathcal{T} and that $L(u) = \log(\Lambda(u))$ for every integral class $u \in \mathcal{A}$. By real-analyticity, this implies that $L = \mathfrak{H}$ and $\mathcal{T} = \mathcal{C}_X$. Rather than study classes $u \in \mathcal{T} \setminus \mathcal{A}$ directly, the analysis in [AKHR] focuses on the construction of alternative “ \mathcal{A} -cones” inside \mathcal{T} , and also produces an interesting interpretation of Ξ in terms of a “cycle polynomial” (which we learned from the authors of [AKHR], and employ in the proof of Theorem E below). However, [AKHR] provides no geometric, algebraic, or dynamical interpretation of the significance of the cone \mathcal{T} (cf. Meta-Theorem I), nor for the value $L(u)$ when u is not in an \mathcal{A} -cone (cf. Meta-Theorem II); note that Example 8.3 shows that \mathcal{C}_X is not a union of \mathcal{A} -cones, in general. Furthermore, [AKHR] does not provide any discussion of, nor structure for, the irrational points of \mathcal{C}_X (cf. Meta-Theorem IV).

We additionally remark that in general the polynomial Ξ from [AKHR] is not equal to m . In fact, Ξ has the advantage of depending only on the pair (G, u_0) and not on the folded mapping torus (X, ψ) . The construction of [AKHR] furthermore ensures that Ξ is a factor of m . However, Ξ consequently cannot in general compute the characteristic polynomials for first return maps of cross sections to ψ (cf. Meta-Theorem II(3)).

We now turn to a more detailed discussion of the main results and ideas in this paper.

1.1. Cross sections

Some of the central objects of our current investigation are the *cross sections* of ψ ; these are embedded graphs in X which are transverse to the flow and intersect every flowline of ψ ; see §5.1 for details. Every section $\Theta \subset X$ determines a dual cohomology class $u = [\Theta]$, and we construct an open convex cone $\mathcal{S} \subset H^1(X; \mathbb{R})$ containing \mathcal{A} with the property that an integral element $u \in H^1(X; \mathbb{R})$ is dual to a section if and only if $u \in \mathcal{S}$ (Proposition 5.10). This *cone of sections* \mathcal{S} has a convenient definition in terms of closed 1-forms (in the sense of [FGS]) satisfying an additional property we call *flow-*

regularity (§5.3). To aid in our analysis of \mathcal{S} , in §6 we introduce another cone, the *Fried cone* $\mathcal{D} \subset H^1(X; \mathbb{R})$, consisting of those cohomology classes which are positive on all closed orbits of ψ ; in fact, \mathcal{D} is determined by finitely many closed orbits (Proposition 6.3). The definition of \mathcal{S} in terms of flow-regular closed 1-forms easily implies that $\mathcal{S} \subseteq \mathcal{D}$. In Appendix A we give a combinatorial characterization of \mathcal{D} (Proposition A.2) that is similar in spirit to the definition of \mathcal{A} and which allows us to show the reverse containment $\mathcal{D} \subseteq \mathcal{S}$. Our first theorem summarizes these properties of \mathcal{S} and \mathcal{D} :

Theorem A (Cone of sections). *There is an open convex cone $\mathcal{S} \subset H^1(X; \mathbb{R}) = H^1(G; \mathbb{R})$ containing \mathcal{A} (and thus containing u_0) such that a primitive integral class $u \in H^1(X_f; \mathbb{R})$ is dual to a section of ψ if and only if $u \in \mathcal{S}$. Moreover, \mathcal{S} is equal to the Fried cone \mathcal{D} , and there exist finitely many closed orbits $\mathcal{O}_1, \dots, \mathcal{O}_k$ of ψ such that $u \in H^1(X; \mathbb{R})$ lies in \mathcal{S} if and only if $u(\mathcal{O}_i) > 0$ for each $1 \leq i \leq k$. In particular, \mathcal{S} is an open, convex, polyhedral cone with finitely many rationally defined sides.*

Remark 1.3. We emphasize that the closed orbits \mathcal{O}_i appearing in Theorem A are explicitly defined in terms of the data used to construct X (see §6).

This theorem mirrors Thurston’s result in the 3-manifold setting, where the cone of sections is the cone over a face of the (polyhedral) Thurston norm ball [Thu] and hence is defined by finitely many rational inequalities, and also Fried’s characterization of this cone in terms of “homology directions” [Fri2]. In his unpublished thesis [Wan], Wang uses Fried’s notion of homology directions to provide a similar characterization of the cone of sections for the mapping torus of a free group automorphism. Since there are semiflow equivariant maps between the mapping torus and X , Theorem A implies Wang’s result in his setting.

For every primitive integral class $u \in \mathcal{S}$ we henceforth use Θ_u to denote a cross section dual to u (see Convention 7.8). For technical reasons we impose an additional assumption on the section Θ_u we call “ \mathcal{F} -compatibility”, but as we show in §A.6, every primitive integral $u \in \mathcal{S}$ is dual to such a section. Every cross section $\Theta \subset X$ has an associated *first return map* $\Theta \rightarrow \Theta$ induced by the semiflow, and we will use $f_u: \Theta_u \rightarrow \Theta_u$ to denote the first return map of Θ_u .

1.2. Splittings

We will see that, in general, the sections Θ_u dual to primitive integral $u \in \mathcal{S}$ are not as nicely behaved as those dual to elements of \mathcal{A} . In particular, the inclusion $\Theta_u \hookrightarrow X$ need not be π_1 -injective and f_u need not be a homotopy equivalence. Despite these shortcomings, our next theorem shows that for each section Θ_u , the induced endomorphism $(f_u)_*$ of $\pi_1(\Theta_u)$ nevertheless provides a great deal of information about the corresponding splitting of G : namely, it gives a realization of this splitting as an ascending HNN-extension over a finitely generated free group and provides a way to calculate $\Lambda(u)$.

For the statement, we need the following terminology regarding arbitrary endomorphisms ϕ of a finitely generated free group F : As we explain in §2.4, each such ϕ naturally descends to an injective endomorphism ϕ of the *stable quotient* of F by the *stable kernel*

consisting of all elements that become trivial under some power of ϕ . Additionally, each endomorphism ϕ has a *growth rate* or *stretch factor* defined as

$$\lambda(\phi) = \sup_{g \in F} \liminf_{n \rightarrow \infty} \sqrt[n]{\|\phi^n(g)\|_A},$$

where A is any free basis of F and $\|\phi^n(g)\|_A$ denotes the cyclically reduced word length of $\phi^n(g)$ with respect to A . One can show that this definition is independent of the free basis A ; see §2.3 for details. With this terminology, we prove that the first return maps $f_u : \Theta_u \rightarrow \Theta_u$ provide the following information about the splittings of G determined by primitive integral classes $u \in \mathcal{S}$.

Theorem B (Splittings and ascending HNN-extensions). *Let $u \in \mathcal{S}$ be a primitive integral class with \mathcal{F} -compatible dual section $\Theta_u \subset X$ and first return map $f_u : \Theta_u \rightarrow \Theta_u$. Let Q_u be the stable quotient of $(f_u)_*$ and let $\phi_u = (f_u)_*$ be the induced endomorphism of Q_u . Then:*

- (1) f_u is an expanding irreducible train track map.
- (2) Q_u is a finitely generated free group and $\phi_u : Q_u \rightarrow Q_u$ is injective.
- (3) G may be written as an HNN-extension

$$G \cong Q_u *_{\phi_u} = \langle Q_u, r \mid r^{-1}qr = \phi_u(q) \text{ for all } q \in Q_u \rangle$$

such that $u : G \rightarrow \mathbb{Z}$ is given by the assignment $r \mapsto 1$ and $Q_u \mapsto 0$.

- (4) If $J_u \leq \ker(u)$ denotes the image of $\pi_1(\Theta_u)$ induced by the inclusion $\Theta_u \hookrightarrow X$, then there is an isomorphism $Q_u \rightarrow J_u$ conjugating ϕ_u to $\Phi_u|_{J_u}$ for some automorphism $\Phi_u \in \text{Aut}(\ker(u))$ representing the monodromy φ_u .
- (5) The topological entropy of f_u is equal to $\log(\lambda(\phi_u)) = \log(\lambda(\Phi_u|_{J_u}))$ and also to $\log(\Lambda(u))$.
- (6) $\ker(u)$ is finitely generated if and only if ϕ_u is an automorphism, in which case we have $\ker(u) \cong Q_u$ and $\varphi_u = [\phi_u] \in \text{Out}(Q_u)$.

Remark 1.4. Throughout this paper—for example in Theorem B(1)—we use “train track map” to mean a graph map that satisfies the usual dynamical properties of train track maps in $\text{Out}(F_N)$ theory but which may not be a homotopy equivalence; see §2.1 for the precise definition.

As noted above, the inclusion $\Theta_u \subset X$ is not necessarily π_1 -injective and f_u may fail to be a homotopy equivalence in two different ways. First, it can happen that $(f_u)_* : \pi_1(\Theta_u) \rightarrow \pi_1(\Theta_u)$ may be injective but nonsurjective, in which case $Q_u = \pi_1(\Theta_u)$, $\phi_u = (f_u)_*$ and $\ker(u)$ is not finitely generated even though $u \in \mathcal{S}$. Second, it can happen that $(f_u)_*$ is noninjective, in which case the Hopficity of the free group $\pi_1(\Theta_u)$ necessitates that $(f_u)_*$ is nonsurjective as well. We will see (in Remark 8.4) that both of these possibilities can in fact occur when $u \in \mathcal{S} \setminus \mathcal{A}$. In the second case it can moreover happen that the injective endomorphism $\phi_u : Q_u \rightarrow Q_u$ is surjective and consequently an automorphism of $Q_u \cong \ker(u)$. These kinds of phenomena are not present in the 3-manifold setting; see §1.7 for a more detailed discussion.

1.3. The McMullen polynomial

Having shown that a section Θ_u dual to $u \in \mathcal{S}$ leads to an algebraic description of the corresponding splitting of G , we now turn our attention to the dynamical properties of these splittings. In this regard, our main result is the introduction of a polynomial invariant \mathfrak{m} , termed the *McMullen polynomial*, that simultaneously encodes information about all of the first return maps f_u and the stretch factors of their associated endomorphisms $(f_u)_*$. This result is analogous to McMullen's construction of the Teichmüller polynomial in the setting of fibered hyperbolic 3-manifolds [McM1]. Let

$$H = H_1(G; \mathbb{Z})/\text{torsion} \cong \mathbb{Z}^b,$$

where $b = b_1(G)$ is the first Betti number of G . The McMullen polynomial of (X, ψ) is a certain element \mathfrak{m} in the group ring $\mathbb{Z}[H]$, that is, it is a finite sum

$$\mathfrak{m} = \sum_{h \in H} a_h h \in \mathbb{Z}[H]$$

of elements of H with integral coefficients. Note that if we express the elements $h \in H$ in terms of a multiplicative basis t_1, \dots, t_b , then $\mathfrak{m} = \mathfrak{m}(t_1, \dots, t_b)$ becomes an integral Laurent polynomial in t_1, \dots, t_b .

For every integral class $u \in H^1(G; \mathbb{Z})$, one may then form the *specialization* of \mathfrak{m} at u , which is the 1-variable integral Laurent polynomial

$$\mathfrak{m}_u(\zeta) = \sum_{h \in H} a_h \zeta^{u(h)} \in \mathbb{Z}[\zeta^{\pm 1}].$$

Our next theorem establishes an intimate relationship between these specializations \mathfrak{m}_u and the first return maps f_u of sections Θ_u dual to $u \in \mathcal{S}$.

Theorem C (The McMullen polynomial and its specializations). *There exists an element $\mathfrak{m} \in \mathbb{Z}[H]$ as above with the following properties. Let $u \in \mathcal{S} \subset H^1(X; \mathbb{R})$ be a primitive integral class with dual compatible section Θ_u , first return map f_u , and injective endomorphism $\phi_u = \overline{(f_u)_*}$ as in Theorem B. Then:*

- (1) *The specialization $\mathfrak{m}_u(\zeta) \in \mathbb{Z}[\zeta^{\pm 1}]$ is, up to multiplication by $\pm \zeta^k$ for some $k \in \mathbb{Z}$, equal to the characteristic polynomial of the transition matrix of f_u .*
- (2) *The largest positive root λ_u of \mathfrak{m}_u is equal to the stretch factor $\lambda(\phi_u)$ of ϕ_u , to the spectral radius of the transition matrix of f_u , and to the stretch factor of $(f_u)_*$. In particular, $\lambda_u = \Lambda(u)$. Additionally, $\log(\lambda_u)$ is the topological entropy f_u .*

Remark 1.5. Using different methods, Algom-Kfir, Hironaka and Rafi [AKHR] have independently produced a polynomial and have obtained a similar result to Theorem C, but only for primitive integral elements $u \in \mathcal{A}$. Their polynomial has an additional minimality property which guarantees that it is a factor of \mathfrak{m} and that it depends only on the outer automorphism f_* represented by $f: \Gamma \rightarrow \Gamma$. Because of this last property, their polynomial need not specialize to the characteristic polynomial of the train track map f_u for primitive integral $u \in \mathcal{A}$ (even up to units; cf. Theorem 11.3 and Example 11.5).

The McMullen polynomial m in Theorem C is constructed in the following manner, which is in some sense “dual” to McMullen’s construction of the Teichmüller polynomial [McM1]. Let \tilde{X} denote the universal torsion-free abelian cover of X (which has deck group H) and consider the foliation \mathcal{F} of \tilde{X} by (lifted) flowlines of ψ . In §4 we then construct a *module of transversals to \mathcal{F}* ; this is a module over the group ring $\mathbb{Z}[H]$ and is denoted by $T(\mathcal{F})$. A key fact regarding this module, which we prove in §10, is that $T(\mathcal{F})$ is finitely presented as a $\mathbb{Z}[H]$ -module. The McMullen polynomial of (X, ψ) is then defined to be the g.c.d. of the Fitting ideal of $T(\mathcal{F})$; see §4.1 for details.

This abstract definition of m is, however, rather opaque and thus somewhat unsatisfying. In particular, it gives no indication as to why m should enjoy the properties described in Theorem C. Ultimately, these properties follow from the fact that $T(\mathcal{F})$ encodes a great deal of information about the semiflow ψ , but it is not readily apparent how this information is imparted to the fitting ideal and consequently to m .

To remedy this situation we give an alternative description of the McMullen polynomial, showing that m is a very concrete and explicitly computable object. This description is in terms of *graph modules* for sections of ψ . More precisely, for each primitive integral class $u \in \mathcal{S}$ with dual compatible section Θ_u , we consider the preimage $\hat{\Theta}_u \subset \tilde{X}$ in the universal torsion-free abelian cover and define a corresponding *graph module* $T(\hat{\Theta}_u)$; this module is a certain quotient of the free module on the edges of $\hat{\Theta}_u$ that encodes information about the first return map of $\hat{\Theta}_u$ to itself.

The module $T(\hat{\Theta}_u)$ may be explicitly described as follows: Choosing a component $\tilde{\Theta}_u \subset \hat{\Theta}_u$, the stabilizer of $\tilde{\Theta}_u$ in H is a rank $b - 1$ subgroup $H_u \subset H$, and H splits (noncanonically) as $H \cong H_u \oplus \mathbb{Z}$. There is a canonical submodule of the free module of the edges of $\hat{\Theta}_u$ consisting of finite sums of edges of $\tilde{\Theta}_u$, and if we choose a finite set E of H_u -orbit representatives of the edges of $\tilde{\Theta}_u$, then this submodule is naturally isomorphic to the *finitely generated* free $\mathbb{Z}[H_u]$ -module $\mathbb{Z}[H_u]^E$. If we choose a multiplicative basis t_1, \dots, t_{b-1} for H_u , this latter module may be regarded as the finitely generated free module $\mathbb{Z}[t_1^{\pm 1}, \dots, t_{b-1}^{\pm 1}]^E$ over the ring of integer Laurent polynomials in $b - 1$ variables. The first return map $f_u: \Theta_u \rightarrow \Theta_u$ then lifts to a map $\tilde{\Theta}_u \rightarrow \tilde{\Theta}_u$ which has a well-defined $|E| \times |E|$ “transition matrix” $A_u(t_1, \dots, t_{b-1})$ with entries in $\mathbb{Z}[t_1^{\pm 1}, \dots, t_{b-1}^{\pm 1}]$ (see §9.1).

Choosing a multiplicative generator x for the complement of $H_u < H$ produces an isomorphism between $\mathbb{Z}[H]$ and the ring $\mathbb{Z}[t_1^{\pm 1}, \dots, t_{b-1}^{\pm 1}, x^{\pm 1}]$ of integral Laurent polynomials. With respect to this isomorphism, the McMullen polynomial m is given by the following “determinant formula”, which is analogous to McMullen’s formula for the Teichmüller polynomial in [McM1] and which implies Meta-Theorem III.

Theorem D (Determinant formula). *For any primitive integral $u \in \mathcal{S}$ we have*

$$m(t_1, \dots, t_{b-1}, x) = \det(xI - A_u(t_1, \dots, t_{b-1}))$$

up to units in $\mathbb{Z}[t_1^{\pm 1}, \dots, t_{b-1}^{\pm 1}, x^{\pm 1}]$.

While the compatibility condition imposed on the cross sections Θ_u can be computationally expensive, as it can lead to the addition of a large number of valence-2 vertices (see Definition 7.3), it is essential for the above discussion because the module of transver-

sals $T(\mathcal{F})$, and consequently \mathfrak{m} , is sensitive to the choice of distinguished *vertex leaves* of \mathcal{F} . These leaves are determined by the original graph structure on Γ , and subdividing Γ can indeed affect the McMullen polynomial by introducing extra factors. To illuminate this dependence and increase computational flexibility, we derive a secondary determinant formula (Theorem 11.3) that explains exactly how the McMullen polynomial \mathfrak{m} changes under the addition of new vertex leaves resulting from any subdivision of Γ or of any cross section Θ_u .

1.4. The McMullen cone \mathcal{C}_X

The McMullen polynomial $\mathfrak{m} \in \mathbb{Z}[H]$ naturally determines yet another convex cone $\mathcal{C}_X \subset H^1(X; \mathbb{R})$, which we term the *McMullen cone*. Specifically, \mathcal{C}_X is the dual cone associated to a certain vertex of the “Newton polytope” of $\mathfrak{m} \in \mathbb{Z}[H]$ (see Definition 12.6). The next result states that \mathcal{C}_X is in fact equal to the cone of sections \mathcal{S} . The proof of this result appeals to Theorems A and C and a formula for \mathfrak{m} as a “cycle polynomial” introduced by Algom-Kfir, Hironaka and Rafi [AKHR]. The proof of this formula involves the use of directed graphs labeled by homology classes, which also appears in the work of Hadari [Had1], and is related to McMullen’s recent work on the “clique polynomial” of a weighted directed graph [McM2]. Thus in addition to supplying dynamical information about all cross sections of ψ , our polynomial invariant also detects exactly which integral cohomology classes are dual to sections.

Theorem E (McMullen polynomial detects \mathcal{S}). *The McMullen cone \mathcal{C}_X is equal to the cone of sections \mathcal{S} .*

Theorem C shows that for primitive integral points u of \mathcal{S} , the canonically defined stretch factor $\Lambda(u)$ can be calculated, without any mention of splittings of G or cross sections of ψ , in terms of the specializations \mathfrak{m}_u of the McMullen polynomial. Therefore it is perhaps unsurprising that the algebraic object \mathfrak{m} might impose strong regularity on the function Λ . Indeed, following McMullen [McM1], we use \mathfrak{m} and properties of Perron–Frobenius matrices with entries in the ring of integer Laurent polynomials to prove that the assignment $u \mapsto \log(\lambda(\phi_u)) = \log(\Lambda(u))$ extends to a real-analytic, convex and homogeneous function on the entire cone \mathcal{S} . Thus we obtain a new proof of [DKL, Theorem D] and also extend that result to \mathcal{S} . Together with Theorem E, our argument additionally shows that this function blows up on the boundary of \mathcal{S} .

Theorem F (Convexity of stretch factors). *There exists a real-analytic, degree -1 homogeneous function $\mathfrak{H}: \mathcal{S} \rightarrow \mathbb{R}$ such that:*

- (1) $1/\mathfrak{H}$ is positive and concave, hence \mathfrak{H} is convex.
- (2) For every primitive integral $u \in \mathcal{S} \subset H^1(X; \mathbb{R})$ with dual compatible section Θ_u , first return map f_u , and injective endomorphism $\phi_u = \overline{(f_u)_*}$ as in Theorem B we have

$$\mathfrak{H}(u) = \log(\Lambda(u)) = \log(\lambda(\phi_u)) = \log(\lambda(f_u)) = h(f_u).$$

- (3) $\mathfrak{H}(u)$ tends to infinity as $u \rightarrow \partial\mathcal{S}$.

Remark 1.6. In [AKHR] the authors have also constructed a function defined on a cone containing \mathcal{A} , and have proved a result similar to Theorem F (with (2) holding for $u \in \mathcal{A}$). Real-analyticity of \mathfrak{H} and of the analogous function from [AKHR], combined with the fact that these functions necessarily agree on $\mathcal{A} \subset \mathcal{C}_X$, ensures that the function and cone constructed in [AKHR] agree with \mathfrak{H} and \mathcal{C}_X , respectively. McMullen has obtained similar results in a purely graph-theoretic setting [McM2, Theorem 5.2]. In [Had2], Hadari generalizes and further analyzes the kinds of polynomials produced here (and in [AKHR], [McM1], and [McM2]).

1.5. Foliations and Lipschitz flows

Up to this point our focus has been on properties of sections, or equivalently on the integral points of \mathcal{S} . However, there is a rich geometric, topological, and dynamical structure associated to the nonintegral points as well. To describe this structure, we recall that any element $u \in \mathcal{S}$ is represented by a closed 1-form ω^u that is flow-regular. Associated to ω^u is a “foliation of X transverse to ψ ” denoted by Ω_u with leaves that are (typically infinite) graphs. Concretely, the 1-form determines a (flow-regular, u -equivariant) function on the universal torsion-free abelian cover $\tilde{X} \rightarrow \mathbb{R}$, and each “leaf” of Ω_u is the image in X of a fiber of this map (see §13). The leaves $\Omega_{y,u}$ are thus in one-to-one correspondence with points $y \in \mathbb{R}/u(\pi_1(G))$. As with sections, we are able to draw the strongest conclusions about the closed 1-forms representing points in \mathcal{A} .

Theorem G (π_1 -injective foliations). *Given $u \in \mathcal{A}$, there exists a flow-regular closed 1-form ω^u representing u with associated foliation Ω_u of X having the following property. There is a reparameterization of ψ , denoted ψ^u , such that for each $y \in \mathbb{R}$ the inclusion of the fiber $\Omega_{y,u} \rightarrow X$ is π_1 -injective and induces an isomorphism $\pi_1(\Omega_{y,u}) \cong \ker(u)$. Furthermore, for every $s \geq 0$ the restriction*

$$\psi_s^u: \Omega_{y,u} \rightarrow \Omega_{y+s,u}$$

of ψ_s^u to any leaf $\Omega_{y,u}$ is a homotopy equivalence.

This result should be compared to the situation for a fibered hyperbolic 3-manifold M , where elements in the cone on a fibered face of the Thurston norm ball are represented by closed, nowhere vanishing 1-forms. The kernel of such a 1-form is then tangent to a *taut foliation* of M , which has π_1 -injective leaves by the Novikov–Rosenberg Theorem (see e.g. [Cal]). Furthermore, appropriately reparameterizing the suspension flow on M will map leaves to leaves by homeomorphisms.

Another application of the McMullen polynomial m mirrors McMullen’s construction of a “Teichmüller flow” for each cohomology class in the cone on a fibered face of the Thurston norm ball [McM1, Theorems 1.1 and 9.1]. In the setting of a fibered 3-manifold M , there is a metric on M so that the reparameterized flow mapping leaves to leaves is actually a Teichmüller mapping on each leaf.

In our setting, associated to any $u \in \mathcal{S}$, we construct a metric on X so that the reparameterized flow, which maps leaves of Ω_u to leaves, has “constant stretch factor”. See §5.3 and §14 for precise definitions.

Theorem H (Lipschitz flows). *For every $u \in \mathcal{S}$, let $\mathfrak{H}(u)$ be as in Theorem F, ω^u a tame flow-regular closed 1-form representing u , ψ^u the associated reparameterization of ψ , and Ω_u the foliation defined by ω^u . Then there is a geodesic metric d_u on X such that:*

- (1) *The semiflow-lines $s \mapsto \psi_s^u(\xi)$ are local geodesics for all $\xi \in X$.*
- (2) *The metric d_u induces a path metric on each (component of a) leaf $\Omega_{y,u}$ of the foliation Ω_u defined by ω^u , making it into a (not necessarily finite) simplicial metric graph.*
- (3) *The restrictions of the reparameterized semiflow to any leaf*

$$\{\psi_s^u : \Omega_{y,u} \rightarrow \Omega_{y+s,u}\}_{s \geq 0}$$

are λ^s -homotheties on the interior of every edge, where $\lambda = e^{\mathfrak{H}(u)}$.

- (4) *The restriction of d_u to the interior of any 2-cell of X is locally isometric to a constant negative curvature Riemannian metric.*

The proof of this theorem resembles McMullen’s construction of Teichmüller flows in some ways. In particular, the metric d_u in this theorem relies on the construction of a kind of “twisted transverse measure” on the foliation \mathcal{F} of X associated to the ray in \mathcal{S} through u (see §14.1). These measures are similar in spirit to the “affine laminations” of Oertel [Oer] and the “twisted measured laminations” of McMullen [McM1] in the context of hyperbolic surfaces and 3-manifolds.

1.6. Relation to the BNS-invariant

Recall that a rational point of $H^1(G; \mathbb{R})$ projects into the BNS-invariant $\Sigma(G)$ if and only if $\ker(u)$ is finitely generated as a $\langle \varphi_u \rangle_+$ -module [BNS] (see the remarks following the Meta-Theorem). We note that the stretch function Λ is naturally defined on the rational BNS-cone $\mathbb{Q}\Sigma(G)$ consisting of classes $u \in H^1(G; \mathbb{R})$ with $[u] \in \Sigma(G)$ and $u(G)$ discrete (see Definition 3.6). Theorem B implies that the rational points of \mathcal{S} project into $\Sigma(G)$, and Theorem 5.11 shows that the rational points of \mathcal{A} in fact land in the symmetrized BNS-invariant $\Sigma_s(G) = \Sigma(G) \cap -\Sigma(G)$. Our above investigations into the foliations Ω_u of X dual to nonrational points $u \in \mathcal{S}$ can be used to show that these inclusions hold for irrational points as well; namely $\mathcal{S} = \mathcal{C}_X$ projects into $\Sigma(G)$ (Proposition 13.2) and \mathcal{A} projects into $\Sigma_s(G)$ (Corollary 13.7). Combining the properties of the stretch function Λ provided by Theorem F and Proposition 3.7, we moreover find that \mathcal{C}_X projects onto a full connected component of $\Sigma(G)$:

Theorem I (McMullen polynomial detects a component of $\Sigma(G)$). *The McMullen cone \mathcal{C}_X projects onto a full component of the BNS-invariant $\Sigma(G)$. That is, $\{[u] \mid u \in \mathcal{C}_X\}$ is a connected component of $\Sigma(G)$.*

This shows that the component of $\Sigma(G)$ containing u_0 is polyhedral and that Λ extends to a real-analytic function on the cone over that component. It is an interesting question whether every component of $\Sigma(G)$ has this structure:

Question 1.7. Does every component of $\Sigma(G)$ contain a class $u: G \rightarrow \mathbb{Z}$ such that $\ker(u)$ is finitely generated and φ_u can be represented by an expanding irreducible train track map?

If the answer to Question 1.7 is “yes,” then the theory developed in this paper would imply that $\Sigma(G)$ is a union of rationally defined polyhedra and that $\Lambda: \widehat{\mathbb{Q}\Sigma}(G) \rightarrow \mathbb{R}_+$ admits a real-analytic, convex extension.

In a recent paper [CL, Theorem 5.2] Cashen and Levitt computed $\Sigma(G)$ for the case where G is the mapping torus of a polynomially growing automorphism ϕ of F_N . They showed that in that situation $\Sigma(G)$ is centrally symmetric and consists of the complement of finitely many rationally defined hyperplanes in $H^1(G; \mathbb{R})$. For a polynomially growing ϕ we have $\lambda(\phi) = 1$, and such an ϕ does not admit an expanding train track representative. Thus the results of [CL] are disjoint from the results in the present paper.

1.7. Contrasting the 3-manifold setting

Our results illustrate some qualitatively different behavior for free-by-cyclic groups as compared with similar results for 3-manifold groups. For a fibered hyperbolic 3-manifold M , there are several natural cones that one may consider in $H^1(M; \mathbb{R})$, such as the cone determined by McMullen’s Teichmüller polynomial, the analog of the cone of sections \mathcal{S} (which can be viewed as the dual of Fried’s [Fri2] cone of “homology directions”), and the components of both the BNS-invariant $\Sigma(\pi_1(M))$ and its symmetrization $\Sigma_s(\pi_1(M))$ containing u_0 . These cones all turn out to be the same and are in fact equal to the cone on a “fibered face” of the Thurston norm ball [Thu, McM1, Fri2, BNS].

Above we have seen that in the free-by-cyclic setting the positive cone \mathcal{A} is contained in a component of $\Sigma_s(G)$, and the McMullen cone \mathcal{C}_X is equal to a component of $\Sigma(G)$. However, our computations show that in general this component of $\Sigma_s(G)$ can be strictly smaller than $\mathcal{C}_X = \mathcal{S}$. While it was already known that, unlike the 3-manifold case, the BNS-invariant of a free-by-cyclic group need not be symmetric (precisely because a free-by-cyclic group can also split as a strictly ascending HNN-extension of a finite-rank free group, as an example in Brown’s 1987 paper [Bro] illustrates), one still might have hoped that the geometric property of being dual to a section was sufficient to ensure containment in $\Sigma_s(G)$. Evidently this is not the case.

In particular, in the “running example” group $G = G_\varphi$ that we analyze in detail throughout this paper, the positive cone $\mathcal{A} \subseteq H^1(G; \mathbb{R})$ is equal to a component of $\Sigma_s(G)$ but is a *proper* subcone of $\mathcal{S} = \mathcal{C}_X$. We moreover exhibit a specific primitive integral element $u_1 \in \mathcal{S} \cap \partial\mathcal{A}$ such that $\ker(u_1)$ is not finitely generated but which does, in accordance with Theorem B, induce a splitting of G as a strictly ascending HNN-extension over a finitely generated free group. Thus u_1 belongs to $\mathcal{S} = \mathcal{C}_X$ but not to $\Sigma_s(G)$; see Examples 5.6, 7.11, 8.3, and 11.2 for the relevant computations regarding u_1 . For the running example we also exhibit a primitive integral class $u_2 \in \mathcal{S} \setminus \overline{\mathcal{A}}$ with dual section Θ_2 such that the first return map $f_2: \Theta_2 \rightarrow \Theta_2$ fails to be a homotopy equivalence but nevertheless the induced endomorphism descends to an automorphism of the finitely generated free group $\ker(u_2)$; see Examples 5.7 and 7.15 for the relevant computations.

This shows that \mathcal{S} can contain multiple distinct components of the symmetrized BNS-invariant $\Sigma_S(G)$.

Remark 1.8. In Example 5.6 and later in Example 7.11 and Remark 7.14, we find a section $\Theta_1 \subset X$ dual to a class $u_1 \in \mathcal{S}$ (mentioned above) with $\ker(u_1)$ infinitely generated and which gives rise to a splitting of G as a strictly ascending HNN-extension of a finite rank free group over an injective but nonsurjective endomorphism. This example illustrates that some of the results similar to Theorem B that are claimed in Wang’s thesis [Wan] are incorrect. Specifically, [Wan, Lemma 1.3] produces, for any section of the semiflow on the mapping torus, a kind of “fibration” of a related complex. According to [Wan, Lemma 4.1.3], this would imply u_1 has finitely generated kernel (equal to the fundamental group of a fiber of the associated “fibration”), and thus u_1 would define a splitting of G as a (f.g. free)-by-cyclic group. However, our computations show that this is not the case.

2. Background

2.1. Graphs and graph maps

We briefly review the relevant terminology regarding graph maps, referring the reader to [DKL, §2] for a more detailed discussion. A continuous map $f: \Gamma \rightarrow \Gamma$ of a finite graph Γ is said to be a *topological graph map* if it sends vertices to vertices and edges to edge paths. We always assume that our graphs do not have valence-1 vertices. If Γ is equipped with either a metric structure or a (weaker) linear structure, then f is furthermore said to be a (*linear*) *graph map* if it enjoys a certain piecewise linearity on edges. For simplicity, we also typically assume that f is a *combinatorial graph map*, which means that Γ is given a metric structure in which every edge has length one, and the restriction of f to each edge e is a local d_e -homothety, where d_e is the number of edges in the edge path $f(e)$. The map is called *expanding* if for every edge e of Γ , the combinatorial lengths of the edge paths $f^n(e)$ tend to infinity as $n \rightarrow \infty$.

Every graph map $f: \Gamma \rightarrow \Gamma$ has an associated *transition matrix* $A(f)$, which is the $|E\Gamma| \times |E\Gamma|$ matrix whose (e, e') -entry records the number of times the edge path $f(e')$ crosses the edge e (in either direction). We caution that, while this definition of $A(f)$ agrees with the definition given in [BH], it gives the transpose of what was used to mean the transition matrix in [DKL]. The graph map f is then said to be *irreducible* if its transition matrix $A(f)$ is irreducible. Moreover, we use $\lambda(f)$ to denote the spectral radius of $A(f)$.

A *train track map* is a linear graph map $f: \Gamma \rightarrow \Gamma$ such that f^k is locally injective at each valence-2 vertex and on each edge of Γ for all $k \geq 1$. Note that this definition is slightly more general than the standard notion of a train track map in $\text{Out}(F_N)$ theory (in particular, more general than the one used in [DKL]) since here we do not require f to be a homotopy equivalence. A train track map that is irreducible always has $\lambda(f) \geq 1$; furthermore this inequality is strict if and only if f is expanding.

If $f: \Gamma \rightarrow \Gamma$ is an irreducible train track map with $\lambda(f) > 1$, then by Corollary A.7 of [DKL] there exists a volume-1 metric structure \mathcal{L} on Γ with respect to which the map f is a local $\lambda(f)$ -homothety on every edge of Γ . As in [DKL], we call \mathcal{L} the *canonical metric structure* on Γ , and the corresponding metric $d_{\mathcal{L}}$ on Γ is the *canonical eigenmetric*.

2.2. Markings and representatives

When Γ is a finite connected graph and $f: \Gamma \rightarrow \Gamma$ is topological graph map which is a homotopy equivalence, it induces an automorphism f_* of the free group $\pi_1(\Gamma)$ that is well-defined up to conjugacy; accordingly we say that f is a *topological representative* of the outer automorphism $[f_*] \in \text{Out}(\pi_1(\Gamma))$. Even when f is not a homotopy equivalence, there is still a sense in which f represents an endomorphism of a free group, as we now explain.

A *marking* on the free group F_N consists of a finite connected graph Γ without valence-1 vertices and an isomorphism $\alpha: F_N \rightarrow \pi_1(\Gamma)$. If \mathcal{L} is a metric structure on Γ , then the pair (α, \mathcal{L}) is called a *marked metric structure* on F_N . Lifting \mathcal{L} to the universal cover $\tilde{\Gamma}$ defines a metric $d_{\tilde{\mathcal{L}}}$ on $\tilde{\Gamma}$ such that $T = (\tilde{\Gamma}, d_{\tilde{\mathcal{L}}})$ is an \mathbb{R} -tree equipped with a free and discrete isometric action (via α) of F_N by covering transformations. For an element $g \in F_N$, let $\|g\|_T := \inf_{x \in T} d_{\tilde{\mathcal{L}}}(x, gx)$ denote the *translation length* of g on T . Note that $\|g\|_T$ is equal to the \mathcal{L} -length of the unique immersed loop in Γ which is freely homotopic to the loop $\alpha(g) \in \pi_1(\Gamma)$. If $A = \{a_1, \dots, a_N\}$ is a free basis of F_N , Γ is a rose with N petals corresponding to the elements of A , and \mathcal{L} is the metric structure assigning every edge of Γ length 1, then T is exactly the Cayley tree of F_N corresponding to A . In this case for $g \in F_N$ we have $\|g\|_T = \|g\|_A$, where $\|g\|_A$ is the cyclically reduced word length of g with respect to A . Note that if $g, g' \in F_N$ are conjugate elements then $\|g\|_T = \|g'\|_T$.

A *topological representative* of a free group endomorphism $\phi: F_N \rightarrow F_N$ consists of a marking $\alpha: F_N \rightarrow \pi_1(\Gamma, v)$ (where $v \in V\Gamma$) and a topological graph map $f: \Gamma \rightarrow \Gamma$ such that for some inner automorphism τ of F_N we have $\tau \circ \phi = \alpha^{-1} \circ f_* \circ \alpha$. Here $f_*: \pi_1(\Gamma, v) \rightarrow \pi_1(\Gamma, v)$ is the endomorphism of $\pi_1(\Gamma, v)$ defined as $f_*(\gamma) = \beta f(\gamma) \beta^{-1}$ for some fixed edge path β in Γ from v to $f(v)$. Thus the only difference with the standard definition of a topological representative is that here the graph map $f: \Gamma \rightarrow \Gamma$ need not be a homotopy equivalence. Indeed, we note that a topological representative f as above will be a homotopy equivalence if and only if ϕ is an actual automorphism of F_N . As usual, we often suppress the marking and, by abuse of notation, simply talk about $f: \Gamma \rightarrow \Gamma$ being a topological representative of ϕ . Most of the free group train track theory is developed for automorphisms of F_N , but topological and train track representatives of endomorphisms of F_N appear, for example, in [AKR, DV, Rey].

2.3. Growth

Let $\alpha: F_N \rightarrow \pi_1(\Gamma)$ be a marking on F_N , let \mathcal{L} be a metric structure on Γ , and let $T = (\tilde{\Gamma}, d_{\tilde{\mathcal{L}}})$ be the corresponding \mathbb{R} -tree as above. For any endomorphism $\phi: F_N \rightarrow F_N$

and any element $g \in F_N$ set

$$\lambda(\phi; g, T) := \liminf_{n \rightarrow \infty} \sqrt[n]{\|\phi^n(g)\|_T}.$$

It is not hard to check that the \liminf in the above formula is in fact always a limit. Moreover, if T' is an \mathbb{R} -tree corresponding to another marked metric structure on F_N , then the trees T and T' are F_N -equivariantly quasi-isometric. This fact implies that $\lambda(\phi; g, T) = \lambda(\phi; g, T')$ for all $g \in F_N$. Thus we may unambiguously define $\lambda(\phi; g) := \lambda(\phi; g, T)$, where T is the \mathbb{R} -tree corresponding to any marked metric structure on F_N .

Definition 2.1 (Growth of an endomorphism). Let $\phi: F_N \rightarrow F_N$ be an arbitrary endomorphism. The *growth rate* or *stretch factor* of ϕ is defined to be

$$\lambda(\phi) := \sup_{g \in F_N} \lambda(\phi; g).$$

We say that ϕ is *exponentially growing* if $\lambda(\phi) > 1$.

It is not hard to check that if ϕ is not exponentially growing then $\lambda(\phi; g) = 1$ or $\lambda(\phi; g) = 0$ for every $g \in F_N$. The case $\lambda(\phi; g) = 0$ is possible since ϕ need not be injective; indeed, $\lambda(\phi; g) = 0$ if and only if $\phi^k(g) = 1 \in F_N$ for some $k \geq 1$. Furthermore, since the translation length of an element for an isometric group action is invariant under conjugation in the group, Definition 2.1 implies that $\lambda(\phi) = \lambda(\tau\phi) = \lambda(\phi\tau)$ for any inner automorphism $\tau \in \text{Inn}(F_N)$ of F_N .

Any irreducible train track representative of ϕ (assuming such a representative exists) can be used to compute the growth rate $\lambda(\phi)$:

Proposition 2.2. *Let $\phi: F_N \rightarrow F_N$ be an endomorphism and let $f: \Gamma \rightarrow \Gamma$ be an irreducible train track representative of ϕ . Then $\lambda(\phi) = \lambda(f)$ and $\log(\lambda(f)) = h(f)$, where $h(f)$ is the topological entropy of f .*

Proof. If $\lambda(f) = 1$, then f is a simplicial automorphism of Γ and hence ϕ is an automorphism of F_N which has finite order in $\text{Out}(F_N)$. Therefore $\lambda(\phi) = 1$ and $\lambda(\phi) = \lambda(f)$, as required. Also, in this case it is easy to see that $h(f) = 0$, and so $\log(\lambda(f)) = h(f)$.

Suppose now that $\lambda(f) > 1$, so that f is expanding. Proposition A.1 of [DKL] then implies $\log(\lambda(f)) = h(f)$. Let \mathcal{L} be the canonical metric structure on Γ . Then for every edge e of Γ we have $\mathcal{L}(f(e)) = \lambda(f)\mathcal{L}(e)$, and hence for every edge path γ in Γ ,

$$\mathcal{L}(f_{\#}(\gamma)) \leq \mathcal{L}(f(\gamma)) = \lambda(f)\mathcal{L}(\gamma),$$

where $f_{\#}(\gamma)$ is the tightened form of $f(\gamma)$. This implies that if T is the \mathbb{R} -tree corresponding to (Γ, \mathcal{L}) , then for every $g \in F_N$ we have $\|\phi(g)\|_T \leq \lambda(f)\|g\|_T$. Therefore for $n \geq 1$ we have $\|\phi^n(g)\|_T \leq \lambda(f)^n\|g\|_T$ and hence $\lambda(\phi) \leq \lambda(f)$.

Since f is an expanding graph map, for some edge e there exists $m \geq 1$ such that $f^m(e)$ contains at least two occurrences of the same oriented edge. Therefore there exists a nontrivial immersed loop γ in Γ such that γ is a subpath of $f^m(e)$. Since f is a train track map, for every $n \geq 1$ the loop $f^n(\gamma) = (f^n)_{\#}(\gamma)$ is immersed in Γ , and so

$$\mathcal{L}((f^n)_{\#}(\gamma)) = \mathcal{L}(f^n(\gamma)) = \lambda(f)^n\mathcal{L}(\gamma).$$

Taking $g_0 \in F_N$ to be any element whose conjugacy class is represented by γ , we see that $\|\phi^n(g_0)\|_T = \lambda(f)^n \|g_0\|_T$ for all $n \geq 1$. Hence $\lambda(\phi; g_0, T) = \lambda(f)$. Together with $\lambda(\phi) \leq \lambda(f)$, this implies that $\lambda(\phi) = \lambda(f)$. \square

Proposition 2.3. *Let $\phi: F_N \rightarrow F_N$ be an arbitrary free group endomorphism. Given any $i \geq 0$, set $J = \phi^i(F_N)$ and let $\xi = \phi|_J: J \rightarrow J$. Then $\lambda(\phi) = \lambda(\xi)$.*

Proof. Pick a free basis A of F_N and let $T = T_A$ be the Cayley graph of F_N with respect to A . Let $i \geq 0$ be arbitrary and let $J = \phi^i(F_N)$. Thus J is a finitely generated free group of rank $\leq N$. Choose a free basis B for J and let T' be the Cayley graph of J with respect to B . Let $\xi = \phi|_J: J \rightarrow J$. For any $g \in F_N$ the definitions imply that $\lambda(\phi; g, T) = \lambda(\phi; \phi^i(g), T)$. The subgroup $J \leq F_N$ is quasi-isometrically embedded in F_N . Hence there exists $C \geq 1$ such that

$$\frac{1}{C} \|w\|_T \leq \|w\|_{T'} \leq C \|w\|_T$$

for all $w \in J$. Therefore for all $n \geq 1$ and $g \in F_N$ we have $\phi^{i+n}(g) = \xi^n(\phi^i(g))$ and

$$\frac{1}{C} \|\phi^{n+m}(g)\|_T \leq \|\xi^n(\phi^m(g))\|_{T'} \leq C \|\phi^{n+m}(g)\|_T.$$

Hence

$$\lambda(\phi; g, T) = \lambda(\phi; \phi^m(g), T) = \lambda(\xi; \phi^m(g), T').$$

The definition of the growth rate now implies that $\lambda(\phi) = \lambda(\xi)$. \square

2.4. Endomorphisms and HNN-like presentations

In this subsection we elaborate on the observation of Kapovich [Kap] about the algebraic structure of HNN-like presentations based on arbitrary (and possibly noninjective) endomorphisms.

Let $\phi: G \rightarrow G$ be an arbitrary endomorphism of any group G . We then use the notation $G*_\phi$ to denote the group given by the ‘‘HNN-like’’ presentation

$$G*_\phi := \langle G, r \mid r^{-1}gr = \phi(g) \text{ for all } g \in G \rangle. \tag{2.4}$$

The generator r here is called the *stable letter* of $G*_\phi$, and the group has a (natural) projection $G*_\phi \rightarrow \mathbb{Z}$ defined by sending $r \mapsto 1$ and $G \mapsto 0$.

Presentations as above with ϕ noninjective are difficult to work with, since in that case G does not embed into $G*_\phi$ and Britton’s Lemma (the normal form theorem for HNN-extensions) does not hold. However, as we will see below, presentation (2.4) does define a group which is a genuine HNN-extension along an injective endomorphism of a quotient group of G .

To this end, define the *stable kernel* of $\phi: G \rightarrow G$ to be the normal subgroup

$$K_\phi := \bigcup_{i=1}^\infty \ker(\phi^i) \triangleleft G$$

obtained as the union of the increasing chain $\ker(\phi) \leq \ker(\phi^2) \leq \dots$ of subgroups $\ker(\phi^i) \triangleleft G$. We also denote $\bar{G}^\phi = G/K_\phi$ and call \bar{G}^ϕ the *stable quotient* of ϕ .

Proposition 2.5. *Let G be a group and $\phi: G \rightarrow G$ be an arbitrary endomorphism. Let $G*_\phi$ be as in presentation (2.4) above and let $\bar{G} = \bar{G}^\phi$. Then:*

- (1) $\phi: G \rightarrow G$ descends to an injective endomorphism $\bar{\phi}: \bar{G} \rightarrow \bar{G}$.
- (2) The natural homomorphism $G \rightarrow G*_\phi$ has kernel equal to K_ϕ , so that the image of G in $G*_\phi$ is canonically identified with \bar{G} .
- (3) The quotient map $G \rightarrow \bar{G}$ (which we denote $g \mapsto \bar{g}$) induces an isomorphism

$$G*_\phi \rightarrow \bar{G}*_\phi = \langle \bar{G}, \bar{r} \mid \bar{r}^{-1}\bar{g}\bar{r} = \bar{\phi}(\bar{g}) \text{ for all } \bar{g} \in \bar{G} \rangle,$$

whose composition with the projection $\bar{G}*_\phi \rightarrow \mathbb{Z}$ yields the natural projection of $G*_\phi$. Thus $G*_\phi$ is canonically isomorphic to the (genuine) HNN-extension $\bar{G}*_\phi$ over the injective endomorphism $\bar{\phi}$.

Proof. From the definition of K_ϕ we see that $g \in K_\phi$ if and only if $\phi(g) \in K_\phi$. Therefore ϕ does indeed descend to an endomorphism $\bar{\phi}: \bar{G} \rightarrow \bar{G}$, and moreover $\bar{\phi}$ is injective. Thus (1) is established.

If $g \in K_\phi$ then $g \in \ker(\phi^n)$ for some $n \geq 1$. Thus $\phi^n(g) = 1$ in G and hence in the group $G*_\phi$ we have $g = r^n\phi^n(g)r^{-n} = 1$. Therefore, by applying a Tietze transformation, we can rewrite presentation (2.4) of $G*_\phi$ as

$$\begin{aligned} G*_\phi &= \langle G, r \mid r^{-1}gr = \phi(g) \text{ for all } g \in G \rangle \\ &= \langle G, r \mid r^{-1}gr = \phi(g) \text{ for all } g \in G; \text{ and } g = 1 \text{ for all } g \in K_\phi \rangle \\ &= \langle G/K_\phi, \bar{r} \mid \bar{r}^{-1}\bar{g}\bar{r} = \bar{\phi}(\bar{g}) \text{ for all } \bar{g} \in G/K_\phi \rangle = \bar{G}*_\phi. \end{aligned}$$

This implies parts (2) and (3) of the proposition. □

The following proposition implies that in certain situations the isomorphism $G*_\phi \cong \bar{G}*_\phi$ from Proposition 2.5 can actually be seen from *inside* G .

Proposition 2.6. *Let $\phi: G \rightarrow G$ be an arbitrary endomorphism, and let \bar{G} and $\bar{\phi}$ be as in Proposition 2.5. Suppose that $i \geq 1$ is such that for $J = \phi^i(G)$ the endomorphism $\xi = \phi|_J: J \rightarrow J$ is injective. Then $K_\phi = \ker(\phi^i)$ and there is a canonical isomorphism $\tau: \bar{G} \rightarrow J$ conjugating $\bar{\phi}$ to ξ .*

Proof. Since $\xi = \phi|_J$ is injective by assumption, it follows that $\phi^n|_J$ is injective for all $n \geq 0$. Therefore, if $g \in G$ is such that $\phi^i(g) \in J$ is nontrivial, then $\phi^{n+i}(g) = \phi^n(\phi^i(g))$ is also nontrivial for all $n \geq 0$. This implies that $\phi^i(k) = 1 \in J$ for all $k \in K_\phi$, and thus $K_\phi \subset \ker(\phi^i)$. As $\ker(\phi^i) \subset K_\phi$ by definition, it follows that $\ker(\phi^i) = K_\phi$.

Thus $\bar{G} = G/\ker(\phi^i)$ and so the First Isomorphism Theorem provides a canonical isomorphism $\tau: \bar{G} = G/K_\phi \rightarrow J$ given by $\tau(gK_\phi) = \phi^i(g)$ for all $g \in G$. Then for any $g \in G$ we have

$$\xi(\tau(gK_\phi)) = \xi(\phi^i(g)) = \phi^{i+1}(g) = \tau(\phi(g)K_\phi) = \tau(\bar{\phi}(gK_\phi)).$$

Thus τ indeed conjugates $\bar{\phi}$ to ξ , and the statement of the proposition follows. □

It was observed by Kapovich [Kap] that the assumption of Proposition 2.6 is always satisfied when G is a finite-rank free group F_N .

Proposition 2.7. *Let $\phi: F_N \rightarrow F_N$ be an arbitrary endomorphism of a finite-rank free group F_N . Then there exists $i \geq 0$ such that $b_1(\phi^i(F_N)) = b_1(\phi^{i+1}(F_N))$. Furthermore, if $J = \phi^i(F_N)$ for any such i , then $\xi = \phi|_J: J \rightarrow J$ is injective.*

Proof. For every $m \geq 1$ the group $\phi^m(F_N)$ is a subgroup of F_N and it is also a homomorphic image of $\phi^{m-1}(F_N)$. Therefore each image $\phi^m(F_N)$ is a finitely generated free group and the integral sequence $\{b_1(\phi^m(F_N))\}$ is nonincreasing. Therefore there exists $i \geq 0$ such that $b_1(\phi^i(F_N)) = b_1(\phi^{i+1}(F_N))$ as claimed.

For any such $i \geq 0$, the free groups $\phi^i(F_N)$ and $\phi^{i+1}(F_N)$ have the same rank and are thus isomorphic. Moreover the homomorphism ϕ maps $\phi^i(F_N)$ onto $\phi^{i+1}(F_N)$. Since finitely generated free groups are Hopfian, it follows that ϕ maps $\phi^i(F_N)$ isomorphically onto $\phi^{i+1}(F_N)$ and thus the restriction of $\phi|_{\phi^i(F_N)}$ is injective. \square

Combining the above with the results from §2.3, we obtain the following corollary, which, in the case when G is free, gives two (essentially equivalent) useful ways of realizing the group $G*_{\phi}$ from (2.4) as an ascending HNN-extension of a finitely generated free group along an injective endomorphism.

Corollary 2.8. *Let $G = F_N$, let $\phi: G \rightarrow G$ be an arbitrary endomorphism, and let $\bar{G} = \bar{G}^{\phi}$ and $\bar{\phi}$ be as in Proposition 2.5. Then \bar{G} is a finite-rank free group and there exists $J = \phi^i(G)$ for some $i \geq 0$ such that $\xi = \phi|_J: J \rightarrow J$ is injective. Furthermore, the respective stretch factors satisfy $\lambda(\phi) = \lambda(\bar{\phi}) = \lambda(\xi)$, and there are canonical isomorphisms*

$$G*_{\phi} \cong \bar{G}*_{\bar{\phi}} \cong J*\xi$$

respecting the natural projections of these groups to \mathbb{Z} .

Proof. The existence of such an $i \geq 0$ follows from Proposition 2.7. Proposition 2.6 then provides an isomorphism $\tau: \bar{G} \rightarrow J$ showing that \bar{G} is a finite-rank free group. Furthermore, since τ conjugates $\bar{\phi}$ to ξ , it also induces an isomorphism $\bar{G}*_{\bar{\phi}} \rightarrow J*\xi$ and shows that $\lambda(\xi) = \lambda(\bar{\phi})$. Lastly, the equality $\lambda(\phi) = \lambda(\xi)$ follows from Proposition 2.3, and the isomorphism $G*_{\phi} \cong \bar{G}*_{\bar{\phi}}$ from Proposition 2.5. \square

2.5. Invariance of HNN stretch factor

Consider a finitely presented group D that splits as an HNN-extension

$$D = A*_{\phi} = \langle A, r \mid r^{-1}ar = \phi(a) \text{ for all } a \in A \rangle,$$

where A is a finite-rank free group and $\phi: A \rightarrow A$ is an injective endomorphism. The natural projection $u: D \rightarrow \mathbb{Z}$ defined by $r \mapsto 1$ and $A \mapsto 0$ is then said to be dual to the splitting $D = A*_{\phi}$.

Note that if $\ker(u) = ncl_D(A)$ is not finitely generated, then there are infinitely many splittings of D that are all dual to the same u . Indeed, in this case $\phi: A \rightarrow A$ is nonsurjective and so we may choose a finitely generated subgroup $C \leq A$ such that $\phi(A) \leq C \leq A$. Set $B = rCr^{-1}$, so that $C \leq A \leq B$ and $r^{-1}Br = C \leq B$. Defining $\psi: B \rightarrow B$ by $\psi(b) = r^{-1}br$, it is not hard to see that the HNN-extension $B*_{\psi} = \langle B, r \mid r^{-1}br = \psi(b) \rangle$

for all $b \in B$) gives another splitting $D = B*\psi$, which is again dual to u . Notice that this construction can be used to produce splittings $D = B*\phi$ dual to u where the rank of B is any integer greater than the rank of A . Other variations and iterations of this construction are also possible.

This shows that if a homomorphism $u: D \rightarrow \mathbb{Z}$ is dual to an ascending HNN-extension splitting $D = A*\phi$ of D , then the injective endomorphism $\phi: A \rightarrow A$ defining the splitting is in no way canonical. Nevertheless, its stretch factor $\lambda(\phi)$ is uniquely determined by the homomorphism u :

Proposition 2.9. *Suppose that $u \in \text{Hom}(D, \mathbb{Z})$ is dual to two splittings $D = A*\phi$ and $D = B*\psi$, where $\phi: A \rightarrow A$ and $\psi: B \rightarrow B$ are injective endomorphisms of finite-rank free groups A and B . Then $\lambda(\phi) = \lambda(\psi)$.*

Proof. By assumption we may write

$$\langle A, r \mid r^{-1}ar = \phi(a) \text{ for all } a \in A \rangle = D = \langle B, s \mid s^{-1}bs = \psi(b) \text{ for all } b \in B \rangle,$$

where $u(r) = u(s) = 1$ and $u(A) = u(B) = 0$. In particular, $\ker(u) = ncl_D(A) = ncl_D(B)$.

If ϕ is an automorphism of A , then $A = \ker(u) = B$, so that ψ defines an automorphism of A in the same outer automorphism class as ϕ . In this case we obviously have $\lambda(\phi) = \lambda(\psi)$. Thus we may assume neither ϕ nor ψ is surjective.

The fact that $u(s) = u(r)$ implies $s = ra'$ for some $a' \in \ker(u) = \bigcup_{i=0}^{\infty} r^i Ar^{-i}$. Thus there exists $k \geq 1$ such that $a_0 := r^{-k}a'r^k \in A$. Conjugating the HNN presentation $D = B*\psi$ by the element $r^{-k} \in D$, which does not change the dual homomorphism $u: D \rightarrow \mathbb{Z}$ and preserves the stretch factor of the defining endomorphism, we may henceforth assume $s = ra_0$ for some $a_0 \in A$. Note that we then have $s^{-1}As \leq A$.

We claim that $s^m As^{-m} = r^m Ar^{-m}$ for all $m \geq 0$. This is obviously true for $m = 0$, so by induction assume it holds for some $m \geq 0$. Then

$$s(s^m As^{-m})s^{-1} = (ra_0)(r^m Ar^{-m})(a_0^{-1}r^{-1}) = r^{m+1}Ar^{-m-1},$$

where we have used the fact that $a_0(r^m Ar^{-m})a_0^{-1} = r^m Ar^{-m}$ since $a_0 \in r^m Ar^{-m}$. The claim follows.

We therefore have $\ker(u) = \bigcup_{i=0}^{\infty} s^i As^{-i}$. Since $B \leq \ker(u)$ is finitely generated and $s^{-1}As \leq A$, it follows that there exists $n_0 \geq 0$ for which $s^{-n_0}Bs^{n_0} \leq A$. Thus, conjugating the splitting $D = B*\psi$ by the element $s^{-n_0} \in D$, which does not change the stable letter s of the presentation, we may assume that $B \leq A$.

Consider the injective endomorphism $\phi': A \rightarrow A$ defined by $\phi'(a) = s^{-1}as \in A$. Notice that for any $a \in A$, the elements $\phi'(a) = s^{-1}as$ and $\phi(a) = r^{-1}ar$ are conjugate in A (since $s = ra_0$) and therefore have the same cyclically reduced word length with respect to any free basis of A . Thus $\|\phi'(a)\|_A = \|\phi(a)\|_A$ for all $a \in A$, and so we obviously have $\lambda(\phi) = \lambda(\phi')$. On the other hand, since ψ is defined as $\psi(b) = s^{-1}bs$ and $B \leq A$, we see that $\psi = \phi'|_B$. Therefore, letting K denote the maximum value of

$\| \cdot \|_A$ over all elements in a free basis of B , we find that

$$\begin{aligned} \lambda(\psi) &= \sup_{b \in B} \liminf_{n \rightarrow \infty} \sqrt[n]{\|\psi^n(b)\|_B} \leq \sup_{b \in B} \liminf_{n \rightarrow \infty} \sqrt[n]{K \|\psi^n(b)\|_A} \\ &\leq \sup_{a \in A} \liminf_{n \rightarrow \infty} \sqrt[n]{\|\phi^n(a)\|_A} = \lambda(\phi'). \end{aligned}$$

Thus $\lambda(\psi) \leq \lambda(\phi') = \lambda(\phi)$. By symmetry we have $\lambda(\phi) \leq \lambda(\psi)$ as well, and so the proposition follows. \square

We remark that, in view of Corollary 2.8, the conclusion of Proposition 2.9 holds even if one omits the requirement that the endomorphisms ϕ and ψ be injective.

3. Setup

Let Γ be a finite graph with no valence-1 vertices, and let $f: \Gamma \rightarrow \Gamma$ be an expanding irreducible train track map representing an outer automorphism $\varphi \in \text{Out}(F_N)$ of the rank- N free group $F_N \cong \pi_1(\Gamma)$. Let $G = F_N \rtimes_{\varphi} \mathbb{Z}$ be the free-by-cyclic group determined by the outer automorphism φ . Explicitly, G is defined up to isomorphism by choosing a representative $\Phi \in \text{Aut}(F_N)$ of φ and setting

$$G = F_N \rtimes_{\varphi} \mathbb{Z} := \langle w, r \mid r^{-1}wr = \Phi(w) \text{ for } w \in F_N \rangle. \tag{3.1}$$

The first cohomology $H^1(G; \mathbb{R})$ of G is simply the set of homomorphisms $\text{Hom}(G, \mathbb{R})$, and an element $u \in H^1(G; \mathbb{R})$ is said to be *primitive integral* if $u(G) = \mathbb{Z}$. In this situation, the element u determines a splitting (i.e., a split extension)

$$1 \rightarrow \ker(u) \rightarrow G \xrightarrow{u} \mathbb{Z} \rightarrow 1 \tag{3.2}$$

of G and a corresponding *monodromy* $\varphi_u \in \text{Out}(\ker(u))$. Namely, if $t_u \in G$ is any element such that $u(t_u) = 1$, then the conjugation $g \mapsto t_u^{-1}gt_u$ defines an automorphism of $\ker(u) \triangleleft G$ whose outer automorphism class is the monodromy φ_u . It is easy to see that $\varphi_u \in \text{Out}(\ker(u))$ depends only on u and not on the choice of t_u . The homomorphism associated to the original splitting $G = F_N \rtimes_{\varphi} \mathbb{Z}$ (that is, the homomorphism to \mathbb{Z} with kernel $F_N \triangleleft G$ and sending the stable letter r to $1 \in \mathbb{Z}$) will be denoted by u_0 ; its monodromy is the given outer automorphism φ .

We are particularly interested in the case where $\ker(u)$ is finitely generated, which automatically implies that $\ker(u)$ is free [GMSW] and thus φ_u is a free group automorphism. However, even when $\ker(u)$ is not finitely generated, we will see that in many cases φ_u has a naturally associated injective endomorphism—of a finitely generated free subgroup of $\ker(u)$ —with respect to which the splitting (3.2) is realized as an ascending HNN-extension (see §7.2).

3.1. The folded mapping torus

Given a train track map f as above, in [DKL, §4] we constructed a $K(G, 1)$ polyhedral 2-complex $X = X_f$, called the *folded mapping torus* of f . This 2-complex comes equipped with a semiflow ψ and a natural map $\eta: X \rightarrow \mathbb{S}^1$ that is a local isome-

try on flowlines and whose induced map on fundamental groups is the homomorphism $u_0 = \eta_*: G \rightarrow \mathbb{Z}$. The 1-skeleton of X consists of *vertical 1-cells* which are arcs of flowlines, and *skew 1-cells* which are transverse to the flowlines. Each 2-cell of X is a trapezoid whose top and bottom edges consist of skew 1-cells and whose sides consist of vertical 1-cells (it may be that one side is degenerate).

As in [DKL], we will assume that f is a combinatorial graph map (see §2.1), though this is only done to simplify the discussion and exposition. We briefly recall the construction of X : Let $Z_f = \Gamma \times [0, 1]/(x, 1) \sim (f(x), 0)$ be the usual mapping torus of f . There is a natural suspension semiflow on Z_f given by flowing in the $[0, 1]$ direction. The folded mapping torus X is constructed as an explicit flow-respecting quotient of Z_f . In particular, the original graph Γ may be identified with the image of $\Gamma \times \{0\} \subset \Gamma \times [0, 1]$ in X . In this way, Γ is realized as a cross section of the flow on X , and the first return map of ψ to this cross section is exactly f . Furthermore, the natural map $Z_f \rightarrow \mathbb{S}^1$ descends to a map $\eta: X \rightarrow \mathbb{S}^1$ whose induced map on the fundamental groups is just the homomorphism $u_0: G \rightarrow \mathbb{Z}$.

We will work in the universal torsion-free abelian cover $p: \tilde{X} \rightarrow X$, which is the cover corresponding to the subgroup

$$\ker(G \rightarrow H_1(G; \mathbb{Z})/\text{torsion}) < G.$$

Denote the covering group of $\tilde{X} \rightarrow X$ by $H \cong H_1(G; \mathbb{Z})/\text{torsion} \cong \mathbb{Z}^b$, where $b = b_1(G) = \text{rk}(H)$. The semiflow ψ lifts to a semiflow $\tilde{\psi}$ on \tilde{X} . The cell structure on X determines one on \tilde{X} so that the covering map is cellular, and we will use the same terminology (skew, vertical, etc.) to describe objects in \tilde{X} with the same meaning as in X .

Example 3.3 (Running example). We recall here the “running example” that was introduced in [DKL, Example 2.2] and developed throughout that paper. In this paper we will continue our analysis of that example as we use it to illustrate key ideas. Let Γ denote the graph in Figure 1, which has four edges oriented as shown and labeled $\{a, b, c, d\} = E_+\Gamma$. We consider a graph map $f: \Gamma \rightarrow \Gamma$ under which the edges of Γ map to the combinatorial edge paths $f(a) = d$, $f(b) = a$, $f(c) = b^{-1}a$, and $f(d) = ba^{-1}db^{-1}ac$. It is straightforward to check that f is a train track map.

The graph map f induces an automorphism $\varphi = f_*$ of the free group $F_3 = F(\gamma_1, \gamma_2, \gamma_3) \cong \pi_1(\Gamma)$ generated by the loops $\gamma_1 = b^{-1}a$, $\gamma_2 = a^{-1}d$ and $\gamma_3 = c$.




Fig. 1. The example train track map. Left: Original graph Γ . Right: Subdivided graph Δ with labels. Our map $f: \Gamma \rightarrow \Gamma$ is the composition of the “identity” $\iota: \Gamma \rightarrow \Delta$ with the map $f': \Delta \rightarrow \Gamma$ that sends edges to edges preserving labels and orientations.




Fig. 2. The folded mapping torus X_f , with its trapezoid cell structure, for the example $f : \Gamma \rightarrow \Gamma$. The top is glued to the bottom as described by the labeling, and shaded cells with the same shading and shape are identified.

Explicitly, the automorphism φ is given by $\varphi(\gamma_1) = \gamma_2$, $\varphi(\gamma_2) = \gamma_2^{-1}\gamma_1^{-1}\gamma_2\gamma_1\gamma_3$, and $\varphi(\gamma_3) = \gamma_1$. Moreover, one may verify that φ is hyperbolic and fully irreducible.

Consider now the group extension $G = G_\varphi$ defined by the presentation (3.1). As this presentation contains the relations $\gamma_2 = r^{-1}\gamma_1r$ and $\gamma_3 = r\gamma_1r^{-1}$, we see that G is in fact generated by γ_1 and r . Writing $\gamma = \gamma_1$ and performing Tietze transformations to eliminate the generators γ_2 and γ_3 , one may obtain the following two-generator one-relator presentation for G :

$$G = \langle \gamma, r \mid \gamma^{-1}r\gamma^{-1}r\gamma^{-1}r^{-1}\gamma r\gamma r\gamma r^{-3} = 1 \rangle. \tag{3.4}$$

Using the specified train track representative f of φ , we construct the corresponding folded mapping torus $X = X_f$. This folded mapping torus X , equipped with its trapezoid cell structure, is illustrated in Figure 2. See the examples in [DKL] for more details. The covering group $H \cong H_1(G; \mathbb{Z})$ of the universal torsion-free abelian cover $\tilde{X} \rightarrow X$ is then just \mathbb{Z}^2 . Moreover, H is freely generated by the image (under $G \rightarrow G^{\text{ab}}$) of the stable letter $r \in G$ and the image of γ_1 (which is also the image of γ_2 and of γ_3).

For the purposes of analyzing examples, we also note that the trapezoidal cell structure on X is, in general, a subdivision of a cell structure with fewer cells. The 1-cells of this *unsubdivided* cell structure are again either vertical or skew, and so may be oriented so that restriction of the orientation to a 1-cell of the subdivision agrees with its original orientation. In particular, a positive (or nonnegative) 1-cocycle for the canonical cell structure will give rise to one for the unsubdivided cell structure. We do not bother with a formal definition of this cell structure as it will only be used to simplify our discussion of examples, where it will be described.

Example 3.5. Let us employ the observations of the preceding paragraph to build a simplified cell structure on the folded mapping torus from Example 3.3. This cell structure is shown in Figure 3 where we have removed the duplicate polygons to further simplify the picture.




Fig. 3. A simpler picture of the folded mapping torus with the unsubdivided cell structure. The six vertices v_1, \dots, v_6 are those at heights $1, \dots, 6$ as indicated.

There are just six vertices in the unsubdivided cell structure, and we label them v_1, \dots, v_6 according to the heights as illustrated in Figure 3. There are twelve 1-cells, with at most one 1-cell between every pair of vertices. Consequently, we can label the oriented 1-cells by their endpoints, and they are

$$\left\{ \begin{array}{l} [v_1, v_3], [v_3, v_5], [v_5, v_1], [v_2, v_4], [v_4, v_6], [v_6, v_1], \\ [v_1, v_2], [v_2, v_3], [v_3, v_4], [v_4, v_5], [v_5, v_6], [v_6, v_1] \end{array} \right\}.$$

We let $\{v_1^*, \dots, v_6^*\}$ denote the dual basis of 0-cochains. Similarly, if $[v_i, v_j]$ is an oriented 1-cell, we let $[v_i^*, v_j^*]$ denote the dual 1-cochain, thus determining a basis of the 1-cochains.

There are also six 2-cells obtained by gluing the polygons in the figure. Orienting the 2-cells, we can read off the vertices along the boundary. Since there is at most one 1-cell between any two vertices, this ordered list of vertices uniquely determines the loop in the 1-skeleton which is the boundary of the 2-cell. The boundaries of the 2-cells are:

$$\left\{ \begin{array}{l} (v_1, v_3, v_5, v_4, v_2, v_6), (v_1, v_2, v_6, v_5), (v_2, v_4, v_6, v_5, v_3), (v_1, v_3, v_4, v_2), \\ (v_3, v_5, v_1, v_6, v_4), (v_1, v_3, v_5, v_1, v_2, v_3, v_1, v_6, v_5, v_4, v_3, v_2, v_6, v_4, v_5) \end{array} \right\}.$$

3.2. The BNS-invariant and the stretch function

Recall from [BNS] that the *BNS-invariant* of our finitely generated group G is defined to be the subset $\Sigma(G)$ of the sphere $S(G) = (H^1(G; \mathbb{R}) \setminus \{0\})/\mathbb{R}_+$ consisting of those directions $[u]$ for which $G' = [G, G]$ is finitely generated over a finitely generated submonoid of $\{g \in G \mid u(G) \geq 0\}$ (where G acts on G' by conjugation—see the remarks following the Meta-Theorem). It follows from [BNS, Proposition 4.3] that for $u \in H^1(G; \mathbb{R})$ primitive integral, the ray $[u]$ is in $\Sigma(G)$ if and only if there exists $t \in G$ with $u(t) = 1$ and a

finitely generated subgroup $B \leq \ker(u)$ such that $t^{-1}Bt \leq B$ and $G = \langle B, t \rangle$. Defining $\phi: B \rightarrow B$ by $\phi(b) = t^{-1}bt$, we then see that G splits as an ascending HNN-extension

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

that is dual to u (i.e., $u: G \rightarrow \mathbb{Z}$ is given by $B \mapsto 0$ and $t \mapsto 1$). Theorem 2.6 (with Remark 2.7) of [GMSW] shows that in this case B is necessarily a free group. Therefore ϕ is a free group endomorphism and we may consider its stretch factor $\lambda(\phi)$. By Proposition 2.9, $\lambda(\phi)$ is independent of the particular choices of $t \in G$ and $B \leq \ker(u)$ and in fact only depends on u ; thus we may unambiguously define $\Lambda(u) := \lambda(\phi)$.

Definition 3.6 (Stretch function). Let us define the *rational BNS-cone* of G to be the set

$$\widehat{\mathbb{Q}\Sigma}(G) = \{u \in H^1(G; \mathbb{R}) \mid [u] \in \Sigma(G) \text{ and } u(G) \text{ is discrete}\}.$$

Given $u \in \widehat{\mathbb{Q}\Sigma}(G)$, there is a unique $k > 0$ such that $u' = ku$ is primitive integral, and we define $\Lambda(u) := \Lambda(u')^k$ (with $\Lambda(u')$ as defined above). This gives a well-defined function

$$\Lambda: \widehat{\mathbb{Q}\Sigma}(G) \rightarrow \mathbb{R}_+,$$

which we call the *stretch function* of G . It is a canonical invariant of the free-by-cyclic group G and satisfies the homogeneity property $\Lambda(ku) = \Lambda(u)^{1/k}$ for all $k > 0$.

One of the main goals of this paper is to understand and explicitly compute Λ on a large part of $\widehat{\mathbb{Q}\Sigma}(G)$. Along the way, we will need the following crucial feature of the stretch function:

Proposition 3.7. *The stretch function Λ is locally bounded. That is, for all $u \in \widehat{\mathbb{Q}\Sigma}(G)$, there exists a neighborhood $U \subset \widehat{\mathbb{Q}\Sigma}(G)$ of u and a finite number $R > 0$ such that $\Lambda(u') \leq R$ for all $u' \in U$.*

This result will be obtained by expanding on some ideas in [DKL]. As these considerations are somewhat far from our current discussion, the proof of Proposition 3.7 is relegated to Appendix B.

4. The module of transversals and the McMullen polynomial

The flowlines of $\tilde{\psi}$ intersected with each trapezoidal 2-cell determine a 1-dimensional foliation of \tilde{X} . The two skew 1-cells of each trapezoid are transverse to the flowlines, and the vertical 1-cells are arcs of flowlines. The arcs of flowlines in a 2-cell will be called *plaque arcs*. A maximal, path connected, countable union of plaque arcs will be called a *leaf*. We will refer to this decomposition of \tilde{X} into leaves as a *foliation* of \tilde{X} and denote it \mathcal{F} (we also use \mathcal{F} to denote the actual foliation of any 2-cell). This foliation descends to a foliation on X for which the leaves are the images of the leaves in \tilde{X} , and by abuse of notation we also refer to this foliation as \mathcal{F} whenever convenient.

Recall that the union of the vertical 1-cells is preserved by ψ . Moreover, since the vertices of $\Gamma \subset X$ all lie on vertical 1-cells, and since the preimages of these vertices

under all powers of f form a dense subset of Γ , it follows that the set of points that eventually flow into vertical 1-cells is a dense subset of \tilde{X} . This set is, by definition, a union of leaves, and we refer to these as the *vertex leaves* of \mathcal{F} .

Definition 4.1 (Transversal). A *transversal* τ to \mathcal{F} is an arc contained in a 2-cell of \tilde{X} which is transverse to the foliation \mathcal{F} and has both endpoints contained in vertex leaves. We do not view a single point as an arc.

Let $\mathbb{F}(\mathcal{F})$ denote the free \mathbb{Z} -module generated by all transversals τ to \mathcal{F} . Following McMullen, we define the module of transversals to \mathcal{F} , denoted $T(\mathcal{F})$, to be the largest quotient $\mathbb{F}(\mathcal{F}) \rightarrow T(\mathcal{F})$ in which the images $[\tau]$ of transversals τ satisfy the following *basic relations*:

- (1) $[\tau] = [\tau_1] + [\tau_2]$ if $\tau = \tau_1 \cup \tau_2$ and $\tau_1 \cap \tau_2$ is a single point in a vertex leaf,
- (2) $[\tau] = [\tau']$ if τ flows homeomorphically onto τ' in the sense that there exists a continuous function $\nu: \tau \rightarrow [0, \infty)$ such that $w \mapsto \tilde{\psi}_{\nu(w)}(w)$ defines a homeomorphism from τ onto τ' .

More precisely, we make the following definition:

Definition 4.2 (Module of transversals). Let $R \leq \mathbb{F}(\mathcal{F})$ be the submodule generated by all elements $\tau - \tau_1 - \tau_2$ and $\tau - \tau'$, with $\tau, \tau', \tau_1, \tau_2$ as in the basic relations (1)–(2) above. The *module of \mathcal{F}* is defined to be the quotient $T(\mathcal{F}) = \mathbb{F}(\mathcal{F})/R$.

The covering group H acts on the set of transversals by taking preimages: given an element $h \in H$ and a transversal τ , we have $h \cdot \tau = h^{-1}(\tau)$. This is naturally a right action, and it is sometimes convenient to write

$$\tau \cdot h = h \cdot \tau = h^{-1}(\tau).$$

Since H is abelian, the distinction between left versus right is not important, but taking preimages (as opposed to images) will be important. This makes both $\mathbb{F}(\mathcal{F})$ and $T(\mathcal{F})$ into $\mathbb{Z}[H]$ -modules, where $\mathbb{Z}[H]$ is the integral group ring of H . We also note that the quotient $\mathbb{F}(\mathcal{F}) \rightarrow T(\mathcal{F})$ is a $\mathbb{Z}[H]$ -module homomorphism (not just a \mathbb{Z} -module homomorphism).

4.1. The McMullen polynomial

We now define a multivariable polynomial which is analogous to the Teichmüller polynomial defined by McMullen in the 3-manifold setting [McM1]. In later sections we will see that this polynomial invariant encodes much information about cross sections to ψ and various splittings of G . The definition relies on the following proposition, which will follow from the results in §10.

Proposition 4.3. *The $\mathbb{Z}[H]$ -module $T(\mathcal{F})$ is finitely presented.*

Choose any finite presentation of $T(\mathcal{F})$ as a $\mathbb{Z}[H]$ -module, say with m generators and r relations. This gives an exact sequence

$$\mathbb{Z}[H]^r \xrightarrow{D} \mathbb{Z}[H]^m \rightarrow T(\mathcal{F}) \rightarrow 0$$

where D is an $m \times r$ matrix with entries in $\mathbb{Z}[H]$. Recall that the *fitting ideal* of $T(\mathcal{F})$ is the ideal $\mathcal{I} \leq \mathbb{Z}[H]$ generated by all $m \times m$ minors of D , and that this ideal is independent of the chosen finite presentation of $T(\mathcal{F})$ [Lan, Ch. XIII, §10], [Nor].

Definition 4.4 (The McMullen polynomial). Define $\mathfrak{m} \in \mathbb{Z}[H]$ to be the g.c.d. of the fitting ideal \mathcal{I} of $T(\mathcal{F})$, which is well-defined up to multiplication by units in $\mathbb{Z}[H]$ (note that $\mathbb{Z}[H]$ is a unique factorization domain). Explicitly, if p_1, \dots, p_k denote the minors generating \mathcal{I} , then

$$\mathfrak{m} = \gcd\{p \in \mathcal{I}\} = \gcd\{p_1, \dots, p_k\}.$$

Viewing $\mathbb{Z}[H]$ as the ring of integral Laurent polynomials in b variables, we can think of \mathfrak{m} as a Laurent polynomial which we call the *McMullen polynomial* of (X, ψ) . Note that this definition depends only on \tilde{X} , \mathcal{F} , and the action of H .

In the process of proving Proposition 4.3 we will see that \mathfrak{m} enjoys many of the properties that McMullen's Teichmüller polynomial does for 3-manifolds.

5. Cross sections, closed 1-forms, and cohomology

We will see that the McMullen polynomial is intimately related to the cross sections of the semiflow ψ . In fact, these cross sections will play a crucial role in our analysis and understanding of \mathfrak{m} . To this end, we discuss here the definitions and general properties of cross sections and the duality between cross sections and cohomology. Along the way we recall the notion of a closed 1-form on a topological space, and describe a class of closed 1-forms which interact well with ψ . We then introduce the cone of sections \mathcal{S} and give two descriptions of it—one in terms of cross sections and the other in terms of closed 1-forms. In Appendix A we will describe a procedure for explicitly building cross sections dual to certain cohomology classes. This provides a characterization of the classes in \mathcal{S} in combinatorial terms and results in a concrete description of \mathcal{S} allowing us to prove Theorem A.

5.1. Cross sections and flow-regular maps

Given an open subset $W \subset X$ (or \tilde{X}), we say that a continuous map $\eta': W \rightarrow Y$ with $Y = \mathbb{S}^1$ or $Y = \mathbb{R}$ is *flow-regular* if for any $\xi \in X$ the map $\{s \in \mathbb{R}_{\geq 0} \mid \psi_s(\xi) \in W\} \rightarrow Y$ defined by $s \mapsto \eta'(\psi_s(\xi))$ is an orientation preserving local diffeomorphism. With this terminology, we say that a finite embedded graph $\Theta \subset X$ is *transverse* to ψ if there is a neighborhood W of Θ and a flow-regular map $\eta': W \rightarrow \mathbb{S}^1$ such that $\Theta = (\eta')^{-1}(x_0)$ for some $x_0 \in \mathbb{S}^1$.

Definition 5.1 (Cross section). A finite embedded graph $\Theta \subset X$ which is transverse to ψ is called a *cross section* (or simply *section*) of ψ if every flowline intersects Θ infinitely often, that is, if $\{s \in \mathbb{R}_{\geq 0} \mid \psi_s(\xi) \in \Theta\}$ is unbounded for every $\xi \in X$. The cross section then has an associated *first return map* $f_\Theta: \Theta \rightarrow \Theta$, which is the continuous map defined by sending $\xi \in \Theta$ to $\psi_{T(\xi)}(\xi)$, where $T(\xi)$ is the minimum of $\{s > 0 \mid \psi_s(\xi) \in \Theta\}$.

If Θ is a cross section, there exists a reparameterization ψ^Θ of the semiflow ψ such that this first return map is exactly the restriction of the time-1 map ψ_1^Θ to Θ . Explicitly, if K is larger than $T(\xi)$ for all $\xi \in \Theta$, then we may reparameterize ψ inside the flow-regular neighborhood W to obtain a semiflow ψ' for which the first return map of Θ to itself is exactly ψ'_K . The desired reparameterization is then given by $\psi_t^\Theta = \psi'_{t/K}$.

Proposition-Definition 5.2 (Flow-regular maps representing cross sections). *For any cross section $\Theta \subset X$ there exists a flow-regular map $\eta_\Theta: X \rightarrow \mathbb{S}^1$ such that $\Theta = \eta_\Theta^{-1}(0)$. (Thus the flow-regular map $W \rightarrow \mathbb{S}^1$ witnessing the fact that Θ is transverse to ψ may in fact be taken to have domain all of X .) We say that any such η_Θ represents Θ .*

Proof. Note that the function

$$\tau(\xi) = \min\{s > 0 \mid \psi_s^\Theta(\xi) \in \Theta\} \quad \text{for } \xi \in X$$

is bounded above. Indeed, $\tau \equiv 1$ on Θ by construction, and the fact that Θ is transverse to ψ implies that τ is also bounded on an open neighborhood U of Θ . As τ is continuous and therefore bounded on the compact set $X \setminus U$, the global boundedness of τ follows. This implies that every biinfinite flowline (that is, a map $\gamma: \mathbb{R} \rightarrow X$ satisfying $\psi_s^\Theta(\gamma(t)) = \gamma(s + t)$) must intersect Θ infinitely often in the backward direction. Since every point of X lies on a biinfinite flowline, it follows that

$$X = \bigcup_{s \geq 0} \psi_s^\Theta(\Theta) = \bigcup_{0 \leq s \leq 1} \psi_s^\Theta(\Theta). \tag{5.3}$$

In particular, we see that $\tau(\xi) \leq 1$ for all $\xi \in X$, and $\tau(\xi) = 1$ if and only if $\xi \in \Theta$. The assignment $\xi \mapsto 1 - \tau(\xi)$, which is continuous on $X \setminus \Theta$, thus descends to a continuous map $\eta_\Theta: X \rightarrow \mathbb{R}/\mathbb{Z} = \mathbb{S}^1$ that satisfies $\eta_\Theta(\psi_s^\Theta(\xi)) = s + \eta_\Theta(\xi)$ for all $\xi \in X$ and $s \geq 0$. This map η_Θ is a local isometry on each ψ^Θ -flowline, and we also have $\Theta = \eta_\Theta^{-1}(0)$ by construction. Since ψ^Θ is just a reparameterization of ψ , it follows that η_Θ is also flow-regular with respect to the original flow ψ . \square

5.2. Cross sections and cohomology

Every cross section Θ of ψ determines a homomorphism $[\Theta]: G \rightarrow \mathbb{Z}$ as follows: Let $\eta': W \rightarrow \mathbb{S}^1$ be a flow-regular map for which $\Theta = (\eta')^{-1}(x_0)$ for some $x_0 \in \mathbb{S}^1$. Taking a sufficiently small neighborhood $W' \subset W$ of Θ , any closed loop $\gamma: \mathbb{S}^1 \rightarrow X$ may then be homotoped so that $t \mapsto \eta'(\gamma(t))$ is a local homeomorphism on $\gamma^{-1}(W')$. (In fact, one may perform the homotopy by applying ψ judiciously inside W to arrange, for example, that each component of $\gamma \cap W'$ is an arc of a flowline). The value of $[\Theta]$ on γ is then equal to the number of components of $\gamma^{-1}(W')$ on which $\eta' \circ \gamma$ is orientation preserving, minus the number on which it is orientation reversing. Alternatively, if $\eta_\Theta: X \rightarrow \mathbb{S}^1$ represents Θ as in Proposition-Definition 5.2, then $[\Theta] = (\eta_\Theta)_*$.

Definition 5.4 (Duality). The cross section Θ and the corresponding integral cohomology class $[\Theta] \in H^1(X; \mathbb{Z})$ are said to be *dual* to each other.

For any cell structure Y on X , there is a natural way to represent the class $[\Theta]$ by a cellular 1-cocycle $z \in Z^1(Y; \mathbb{Z})$: Adjusting Θ by a homotopy if necessary, first choose a small neighborhood $W' \subset W$ of Θ , as above, such that W' is disjoint from the 0-skeleton $Y^{(0)}$. For each 1-cell σ of Y , one may then find an arc $\gamma: [0, 1] \rightarrow Y$ that is homotopic to σ rel $\partial\sigma$ and for which the assignment $t \mapsto \eta'(\gamma(t))$ is a local homeomorphism on each component of $\gamma^{-1}(W')$. The value of the cocycle z on σ is then defined, as above, to be the number of components of $\gamma^{-1}(W')$ on which $\eta' \circ \gamma$ is orientation preserving minus the number on which it is orientation reversing. By definition, z then agrees with $[\Theta]$ on any 1-cycle representing a closed loop in X . Therefore we see that z is in fact a cocycle (since the boundary ∂T of any 2-cell T is null-homotopic and thus satisfies $z(\partial T) = [\Theta](\partial T) = 0$) and that the cohomology class of z is equal to $[\Theta]$.

Proposition 5.5. *A cross section Θ of ψ is connected if and only if its dual cohomology class $[\Theta]$ is primitive.*

Proof. Let $f_\Theta: \Theta \rightarrow \Theta$ denote the first return map. First suppose that Θ is a disjoint union of $k > 1$ connected components $\Theta_1, \dots, \Theta_k$. By continuity, f_Θ must map each Θ_i into some other Θ_j , and thus determines a self-map ζ of $\{1, \dots, k\}$.

We claim that ζ must be surjective. To see this, suppose for contradiction that 1 was not in the image of ζ . It then follows that every flowline $\{\psi_s(\xi) \mid s \geq 0, \xi \in X\}$ intersects the graph Θ_1 at most once, for otherwise there would be some point of Θ_1 that was mapped back into Θ_1 by some iterate of f_Θ . This fact shows that the subgraph $\Theta' = \Theta_2 \cup \dots \cup \Theta_k$, which is automatically transverse to ψ , is also a section of ψ . Equation (5.3) then shows that $X = \bigcup_{s \geq 0} \psi_s(\Theta')$; in particular, there exist a point $\xi \in \Theta'$ and a time $s \geq 0$ for which $\psi_s(\xi) \in \Theta_1$. This contradicts the assumption that 1 is not in the image of ζ .

Since ζ is surjective, it is also injective and therefore a permutation of the set $\{1, \dots, k\}$. In fact, it must be a cyclic permutation: For each i , the set X_i of points $\xi \in X$ that eventually flow into Θ_i is connected. If the action of ζ partitioned $\{1, \dots, k\}$ into multiple orbits, then the corresponding sets X_i would give a nontrivial decomposition of X into disjoint closed subsets, contradicting the connectedness of X . It now follows that each component Θ_i of Θ is itself a cross section and so determines a dual cohomology class $[\Theta_i]$. However, it is easy to see that these classes are all equal (for example, by using the fact that $\psi_1^\Theta(\Theta_i) = \Theta_{\zeta(i)}$), and thus their sum $[\Theta]_1 + \dots + [\Theta_k]$, which necessarily equals $[\Theta]$, is not primitive.

Conversely, suppose that Θ is connected. Choose a point $\xi \in \Theta$ and consider the loop $\gamma \subset X$ obtained by concatenating the flowline from ξ to $f_\Theta(\xi)$ with a path in Θ back to ξ . Taking a homotopic loop that is transverse to Θ (for example, $\psi_\epsilon(\gamma)$ for some small $\epsilon > 0$) we see that γ intersects Θ once with positive orientation. Therefore $[\Theta](\gamma) = 1$, showing that $[\Theta]$ is primitive. \square

Example 5.6 (The cocycle z_1). Let us construct a cross section and the dual cohomology class for the running example. Using the unsubdivided cell structure introduced in Example 3.5 and illustrated in Figure 3 consider the cocycle

$$z_1 = [v_2^*, v_3^*] + [v_3^*, v_5^*] + 2[v_4^*, v_6^*] + [v_5^*, v_1^*] + [v_6^*, v_2^*].$$

Evaluating this expression on the boundary of each 2-cell shows that it is indeed a 1-cocycle, and we may construct a cross section dual to the class $[z_1]$ as follows: Place $z_1(\sigma)$ vertices along each 1-cell σ of X , so in our case the cells $[v_2, v_3]$, $[v_3, v_5]$, $[v_5, v_1]$, and $[v_6, v_2]$ each get one vertex and $[v_4, v_6]$ gets two vertices. Since z_1 satisfies the cocycle condition, it is possible to “connect the dots” in each 2-cell, which yields a graph Θ_1 intersecting $X^{(1)}$ in exactly these points. In this case it is moreover possible, as illustrated in Figure 4, to construct Θ_1 so that it is transverse to the flow. Every flowline of ψ is seen to hit Θ_1 infinitely often by inspection, so Θ_1 is in fact a cross section. Applying the recipe following Definition 5.4 to Θ_1 yields exactly the cocycle z_1 , so Θ_1 is indeed dual to $[z_1]$ as desired.




Fig. 4. A cross section Θ_1 dual to the class $[z_1]$. The vertices of X are small and the vertices of Θ_1 are larger.

The abstract graph Θ_1 is illustrated on the right of Figure 4. We note that here we have used the standard graph structure of Definition 7.3 below to ensure that the first return map sends vertices to vertices. Following flowlines, we can then calculate the first return map $f_1 = f_{\Theta_1} : \Theta_1 \rightarrow \Theta_1$ as indicated in Figure 5. We thus find that the characteristic polynomial of the transition matrix $A(f_1)$ is

$$\zeta^{10}(\zeta^9 - \zeta^5 - \zeta^4 - \zeta^3 - \zeta^2 - \zeta - 2).$$

One may also verify directly, or by using Peter Brinkmann’s software package `xtrain`,¹ that f_1 is a train track map (cf. Proposition 7.7 below)

By applying Stallings folds to the graph in the middle we produce the graph on the right. This folding operation κ is a homotopy equivalence, and f_1 factors as the composition of κ with the locally injective immersion from the graph on the right into Θ_1 . Since this immersion is locally injective but not injective, we see that $(f_1)_*$ is an injective endomorphism that is *not* an automorphism.

¹ Available at <http://gitorious.org/xtrain>.




Fig. 5. The first return map $f_1 : \Theta_1 \rightarrow \Theta_1$ for the section Θ_1 dual to $[z_1]$ and the result of Stallings folds κ on the right.




Fig. 6. A cross section Θ_2 dual to the class $[z_2]$. Again the vertices of X are small and the vertices of Θ_2 are larger.

Example 5.7 (The cocycle z_2). Here we exhibit another cross section for the running example. Consider the graph Θ_2 illustrated in Figure 6. As in Example 5.6 above, one may verify by inspection that Θ_2 is indeed a cross section and that its dual cohomology




Fig. 7. The first return map $f_2: \Theta_2 \rightarrow \Theta_2$.

class $[\Theta_2]$ is represented by the cocycle

$$z_2 = [v_1^*, v_2^*] + [v_2^*, v_3^*] - [v_4^*, v_5^*] - [v_5^*, v_6^*] + [v_1^*, v_3^*] + 2[v_6^*, v_2^*].$$

Notice that the cells $[v_4, v_5]$ and $[v_5, v_6]$ cross Θ_2 “against the flow”, that is, in the opposite direction as the flowlines do. The abstract graph Θ_2 along with its first return map f_2 are illustrated in Figure 7. Here we have again used the standard graph structure from Definition 7.3 to ensure that f_2 sends vertices to vertices. Applying Stallings folds to the graph on the right produces a graph with rank 4, showing that $(f_2)_*$ is neither injective nor surjective. We have also verified, using xtrain, that f_2 is an expanding irreducible train track map and calculated that the characteristic polynomial of $A(f_2)$ is $\zeta^{13}(\zeta^6 - 3\zeta^3 - 3\zeta - 1)$.

5.3. Closed 1-forms and cohomology

Cross sections provide a useful tool for analyzing certain cohomology classes, but working with cross sections confines us to studying only the integral classes. To consider more general cohomology classes we must recall some additional machinery borrowed from differential geometry.

Following [FGS], a *closed 1-form* ω on a topological space Y is a collection of functions

$$\omega = \{\omega_U : U \rightarrow \mathbb{R}\}_{U \in \mathcal{J}}$$

defined on some open cover \mathcal{J} of Y with the compatibility property that if $U \cap V \neq \emptyset$ then there exist a locally constant function $\Delta_{U,V}$ such that on $U \cap V$, $\omega_V(\xi) = \omega_U(\xi) + \Delta_{U,V}(\xi)$. We assume that the collection of functions and open sets is maximal in the sense that any other function on an open set $W \rightarrow \mathbb{R}$ that also satisfies this compatibility property is already in ω . Any collection of functions satisfying the compatibility condition

is contained in a maximal collection which is unique up to addition of locally constant functions, so we typically ignore this technicality when no confusion arises.

Given a continuous path $\delta: [a, b] \rightarrow Y$, take a partition $a = a_0 < a_1 < \dots < a_k = b$ such that for all $j = 1, \dots, k$, $\delta([a_{j-1}, a_j]) \subset U_j$ for some $U_j \in \mathcal{J}$. Then we can integrate ω over δ :

$$\int_{\delta} \omega = \sum_{j=1}^k [\omega_{U_j}(a_j) - \omega_{U_j}(a_{j-1})].$$

This integral is independent of the covering used and does not change if we homotope δ relative to its endpoints. In particular, the integral over a null-homologous loop is zero, and in this way a closed 1-form defines a 1-dimensional cohomology class. If ω represents the cohomology class u , then we will sometimes write $\omega = \omega^u$ or $u = [\omega]$. Another consequence of these observations is that if U is a simply connected open subset of Y then by the maximality condition on the functions defining ω , U is a domain for a function $\omega_U \in \omega$.

Closed 1-forms can be added and multiplied by scalars (by applying these operations to the defining functions on simply connected domains, for example), and hence the closed 1-forms on Y form a real vector space. If Y is a CW-complex then the map from closed 1-forms to $H^1(Y; \mathbb{R})$ is an epimorphism.

Given a continuous function $\beta: Y' \rightarrow Y$ between topological spaces and a closed 1-form ω on Y , we can pull ω back by β to a closed 1-form $\beta^*(\omega)$. This is the collection of functions obtained from the compositions

$$\omega_U \circ \beta|_{\beta^{-1}(U)}: \beta^{-1}(U) \rightarrow \mathbb{R}$$

and extended to a maximal collection.

For the universal torsion-free abelian cover $p: \tilde{X} \rightarrow X$ and any closed 1-form ω^u on X representing $u \in H^1(X; \mathbb{R})$ the pull-back $\tilde{\omega}^u = p^*(\omega^u)$ contains a globally defined function $\tilde{\eta}_u: \tilde{X} \rightarrow \mathbb{R} \in \tilde{\omega}^u$. Composing with any covering transformation $h: \tilde{X} \rightarrow \tilde{X}$, $h \in H$, gives rise to another such function $\tilde{\eta}_u \circ h$, and these differ by translation by $u(h)$,

$$\tilde{\eta}_u \circ h(\xi) = \tilde{\eta}_u(\xi) + u(h).$$

That $\tilde{\eta}_u$ and $\tilde{\eta}_u \circ h$ differ by a translation follows from the definition of a closed 1-form. To see that the translation is by $u(h)$, take a loop $\gamma: [0, 1] \rightarrow X$ in $G = \pi_1(X)$ representing the homology class h . The covering transformation h sends the initial point $\xi_0 = \tilde{\gamma}(0)$ of a lift of γ to the terminal point $\tilde{\gamma}(1) = h(\xi_0)$ and so

$$\tilde{\eta}_u \circ h(\xi_0) - \tilde{\eta}_u(\xi_0) = \tilde{\eta}_u(\tilde{\gamma}(1)) - \tilde{\eta}_u(\tilde{\gamma}(0)) = \int_{\tilde{\gamma}} \tilde{\omega}^u = \int_{\gamma} \omega^u = u(\gamma) = u(h).$$

Conversely, given a function $\tilde{\eta}_u: \tilde{X} \rightarrow \mathbb{R}$ equivariant with respect to the action of H on \mathbb{R} by translation determined by the homomorphism $u: H \rightarrow \mathbb{R}$, it is easy to see that this function belongs to the pull-back $p^*(\omega^u)$ for a unique closed 1-form ω^u on X representing u .

5.4. Flow-regular closed 1-forms and sections

We will be interested in closed 1-forms that interact with ψ in some meaningful way. To that end, we make the following

Definition 5.8 (Flow-regular closed 1-form). Say that a closed 1-form ω is *flow-regular* provided all the functions $\omega_U \in \omega$ are flow-regular. If $\gamma_\xi(s) = \psi_s(\xi)$ is a flowline, then the ω -height of the arc $\gamma_\xi([0, s_0])$ of the flowline is defined to be

$$\int_{\gamma_\xi([0, s_0])} \omega.$$

We say that a flow-regular closed 1-form ω is *tame* if the restriction to every 2-cell is smooth and the restriction to each skew 1-cell has finitely many critical points.

Given a flow-regular closed 1-form ω , it is possible to construct another one which is tame and which represents the same cohomology class. The point is to first adjust ω along the skew 1-cells so that it is smooth with finitely many critical points, but without changing the integral of ω over the 1-cell. Then we further adjust the 1-form over every 2-cell. When convenient we will pass from an arbitrary flow-regular closed 1-form to a tame one.

For every section Θ , the integral class $[\Theta]$ may be represented by a closed flow-regular 1-form $\omega^{[\Theta]}$. Indeed, consider a lift of the flow-regular map $\eta_\Theta: X \rightarrow \mathbb{S}^1$ to a map

$$\tilde{\eta}_\Theta: \tilde{X} \rightarrow \mathbb{R}.$$

Then for every $h \in H$ we have $\tilde{\eta}_\Theta \circ h = \tilde{\eta}_\Theta + [\Theta](h)$, and hence $\tilde{\eta}_\Theta$ determines a closed 1-form $\omega^{[\Theta]}$ on X . Since η_Θ is flow regular, so is $\tilde{\eta}_\Theta$ and hence also $\omega^{[\Theta]}$.

Conversely, a flow-regular closed 1-form $\omega^{[\Theta]}$ representing an integral class determines a flow-regular map $\eta_\Theta: X \rightarrow \mathbb{S}^1$ by integrating along paths from a fixed basepoint, and any fiber of this map is a section.

Given this correspondence between sections and integral classes represented by closed flow-regular 1-forms, the following definition is quite natural.

Definition 5.9 (Cone of sections). Let $\mathcal{S} \subset H^1(X; \mathbb{R})$ denote the set of all cohomology classes represented by flow-regular closed 1-forms. We refer to \mathcal{S} as the *cone of sections*.

The discussion above proves the following.

Proposition 5.10. *The integral classes in \mathcal{S} are precisely the set of classes in $H^1(X; \mathbb{R})$ that are dual to sections of ψ .* \square

5.5. Convexity and the \mathcal{A} cone

In [DKL] we introduced the nonempty, open, convex cone $\mathcal{A} \subset H^1(X; \mathbb{R})$ containing the epimorphism u_0 associated to the original splitting $G = F_N \rtimes_\varphi \mathbb{Z}$ and consisting of cohomology classes that may be represented by cellular 1-cocycles that are positive on all 1-cells of X . We then proved that every primitive integral class $u \in \mathcal{A}$ is dual to a cross section of ψ with some nice additional properties. More precisely, Theorem B of [DKL] gives:

Theorem 5.11. *Given a positive cellular 1-cocycle $z \in Z^1(X; \mathbb{R})$ representing a primitive integral element $u \in \mathcal{A} \subset H^1(X; \mathbb{R}) = H^1(G; \mathbb{R})$, there exists a flow-regular map $\eta_u: X \rightarrow \mathbb{S}^1$ with $(\eta_u)_* = u: G \rightarrow \pi_1(\mathbb{S}^1) = \mathbb{Z}$ and a fiber $\Theta_u = \eta_u^{-1}(y_0)$ for some $y_0 \in \mathbb{S}^1$ such that*

- (1) Θ_u is a finite, connected topological graph such that the inclusion $\Theta_u \subset X$ is π_1 -injective and $\pi_1(\Theta_u) = \ker(u) \leq \pi_1(X) = G$, and
- (2) Θ_u is a section of ψ that is dual to u .

From this theorem and basic properties of flow-regular closed 1-forms we deduce the following.

Proposition 5.12. *The cone of sections \mathcal{S} is an open convex cone in $H^1(X; \mathbb{R})$ containing \mathcal{A} . Moreover, \mathcal{S} is the convex hull of the rays through integral classes dual to sections.*

Proof. If ω_1 and ω_2 are closed flow-regular 1-forms, then so is $t_1\omega_1 + t_2\omega_2$ for any $t_1, t_2 \in \mathbb{R}_+$, thus \mathcal{S} is a convex cone. According to Theorem 5.11, every primitive integral element of \mathcal{A} is dual to a section, and hence the rays through integral classes of \mathcal{A} are contained in \mathcal{S} . Since \mathcal{A} is an open convex cone, these rays are dense and in fact \mathcal{A} is the convex hull of this set of rays. By convexity of \mathcal{S} , it follows that $\mathcal{A} \subset \mathcal{S}$.

We now prove that \mathcal{S} is open. Given $u \in \mathcal{S}$, choose a tame, flow-regular closed 1-form ω representing u . The derivative of ω along flowlines is well-defined (independent of the choice of function in ω since any two differ by a constant) and we denote it $D_\psi\omega$. Since ω is smooth on each 2-cell, $D_\psi\omega$ is a continuous function on X . By compactness, $D_\psi\omega$ is bounded above and below by positive constants, and we let $\epsilon > 0$ be a lower bound.

Next, since \mathcal{A} is open, we can choose tame, flow-regular closed 1-forms $\omega_1, \dots, \omega_b$ representing elements of \mathcal{A} which form a basis for $H^1(X; \mathbb{R})$. Let $K > 0$ be such that $D_\psi\omega_j \leq K$ for all $j = 1, \dots, b$. Now the set

$$U = \left\{ \left[\omega + \sum_{j=1}^b \epsilon_j \omega_j \right] \mid |\epsilon_j| < \epsilon / (bK) \text{ for all } j = 1, \dots, b \right\}$$

is an open neighborhood of $[\omega]$ in $H^1(X; \mathbb{R})$. Furthermore, $D_\psi(\omega + \sum \epsilon_j \omega_j)$ with $|\epsilon_j| < \epsilon / (bK)$ is bounded below by $\epsilon - \sum |\epsilon_j K| > \epsilon - bK\epsilon / (bK) = 0$, and hence this closed 1-form is flow-regular. It follows that U is contained in \mathcal{S} and hence \mathcal{S} is open. The last statement of the proposition is now immediate. □

As we will see in Example 8.3, \mathcal{A} can be a proper subcone of \mathcal{S} .

6. A homological characterization of \mathcal{S}

In this section we introduce a new cone $\mathcal{D} \subset H^1(X; \mathbb{R})$ and investigate its properties. Studying \mathcal{D} ultimately leads to a homological characterization of \mathcal{S} and the proof of Theorem A. Let $\mathfrak{D}(\psi)$ denote the set of closed orbits \mathcal{O} of ψ . Each $\mathcal{O} \in \mathfrak{D}(\psi)$ may be thought of as an integral homology class on X ; for emphasis, we will sometimes denote this class by $[\mathcal{O}] \in H_1(X; \mathbb{Z})$. Since this set of homology classes mimics Fried’s cone of homology directions [Fri2], we call its dual the *Fried cone*:

Definition 6.1. The *Fried cone* of (X, ψ) is defined to be the set

$$\mathcal{D} = \{u \in H^1(X; \mathbb{R}) \mid u(\mathcal{O}) > 0 \text{ for all } \mathcal{O} \in \mathfrak{D}(\psi)\}.$$

To further analyze \mathcal{D} , we now describe a way to organize the set $\mathfrak{D}(\psi)$ that reveals some additional structure. Recall that $f: \Gamma \rightarrow \Gamma$ is an expanding irreducible train track map and that the (e, e') -entry of its transition matrix $A = A(f)$ records the number of times the edge path $f(e')$ crosses the edge e in either direction. The matrix $A(f)$ has an associated transition graph $\mathcal{G}(f)$ whose vertex set is $E\Gamma$ and which has $A(f)_{(e, e')}$ directed edges from e' to e for each pair $(e, e') \in E\Gamma^2 = (V\mathcal{G}(f))^2$.

Consider a combinatorial edge path $t_1 \cdots t_k$ in $\mathcal{G}(f)$ consisting of directed edges t_i from e_{i-1} to e_i . Such an edge path is said to be a *circuit* in $\mathcal{G}(f)$ if $e_k = e_0$ and all of the vertices e_0, \dots, e_{k-1} are distinct. Notice that in this case the edges t_1, \dots, t_k are also distinct. Let \mathcal{Y} denote the set of circuits in $\mathcal{G}(f)$. Circuits that differ by a cyclic permutation of its edges are not considered distinct, thus $t_1 \cdots t_k$ and $t_2 \cdots t_k t_1$ define the same element of \mathcal{Y} . Since $\mathcal{G}(f)$ is a finite graph, the set \mathcal{Y} of circuits is finite.

Suppose that $y = t_1 \cdots t_k \in \mathcal{Y}$ is a circuit in $\mathcal{G}(f)$ with t_i directed from e_{i-1} to e_i . Then by definition of $\mathcal{G}(f)$, the directed edge t_i corresponds to a particular occurrence of e_i in the combinatorial edge path $f(e_{i-1}) \subset \Gamma$. Similarly, the entire circuit determines a particular occurrence of $e_k = e_0$ in the edge path $f^k(e_0) \subset \Gamma$. By the linearity of f , there is then a subinterval $\alpha \subset e_0$ such that the restriction $f^k|_\alpha$ maps α homeomorphically and affinely onto this particular occurrence of e_0 in the edge path $f^k(e_0)$. The map $f^k|_\alpha: \alpha \rightarrow e_0$ then necessarily has a unique fixed point $p \in \alpha$, and we let

$$\mathcal{O}_y = \{\psi_s(p) \mid s \geq 0\} = \{\psi_s(p) \mid 0 \leq s \leq k\}$$

denote the closed orbit of ψ through $p \in \Gamma \subset X$. Notice that this orbit is independent of the cyclic ordering on the directed edges t_1, \dots, t_k of the circuit y . For each $y \in \mathcal{Y}$ we thus have a well-defined closed orbit $\mathcal{O}_y \subset X$.

More generally, each closed combinatorial edge path $w = t_1 \cdots t_k$ in $\mathcal{G}(f)$ (not necessarily embedded) determines a closed orbit \mathcal{O}_w in the same way. Conversely, any closed orbit of ψ crosses a sequence e_1, \dots, e_k of edges of $\Gamma \subset X$ and so determines a closed combinatorial edge path in $\mathcal{G}(f)$. Thus we have a surjective function from the set of closed edge paths in $\mathcal{G}(f)$ (up to cyclic permutation) to the set of closed orbits $\mathfrak{D}(\psi)$; this function is a bijection off the preimage of closed orbits through vertices of Γ .

Lemma 6.2. *Given any closed edge path w in $\mathcal{G}(f)$, we may write w as a union of circuits $y_1 \cup \cdots \cup y_n$ with $y_1, \dots, y_n \in \mathcal{Y}$ and correspondingly write*

$$[\mathcal{O}_w] = [\mathcal{O}_{y_1}] + \cdots + [\mathcal{O}_{y_n}] \in H_1(X; \mathbb{Z}).$$

Note that since $[w] = [y_1] + \cdots + [y_n] \in H_1(\mathcal{G}(f); \mathbb{Z})$, it follows that the function from closed edge paths in $\mathcal{G}(f)$ to $\mathfrak{D}(\psi)$ descends to an additive map at the level of homology.

Proof of Lemma 6.2. If $w = t_1 \cdots t_k$ is already a circuit, then there is nothing to do. Otherwise, we decompose w as a union of circuits in the following recursive way. Let $w_1 = w$. Since w_1 is closed, there must exist a subpath $t_i \cdots t_j$, for some $1 \leq i \leq j \leq k$,

that is a circuit. In this case, we may “factor out” the subpath $y_1 = t_i \cdots t_j \in \mathcal{Y}$ to obtain a shorter closed path $w_2 = t_1 \cdots t_{i-1} t_{j+1} \cdots t_k$ in $\mathcal{G}(f)$. Since w_2 is again closed, we may factor out another circuit $y_2 \in \mathcal{Y}$, giving yet a shorter closed path w_3 . This process continues until we reach a closed path w_n that is itself a circuit $y_n = w_n \in \mathcal{Y}$.

At the level of homology, we may perform an analogous factoring procedure for the orbit $\mathcal{O}_w \subset X$. Specifically, we note that the closed orbits \mathcal{O}_w and \mathcal{O}_{y_1} start on some edge e of Γ , cross the same set of edges (and at the same occurrences inside iterates of f on e) before returning back to e when \mathcal{O}_{y_1} closes up. At this time \mathcal{O}_w switches and starts crossing the same edges as \mathcal{O}_{w_2} , finally returning to e as both \mathcal{O}_w and \mathcal{O}_{w_1} close up. Thus, the closed orbit \mathcal{O}_w is homotopic to a closed curve constructed out of the closed orbits \mathcal{O}_{y_1} and \mathcal{O}_{w_2} together with a collection of subarcs of the edge e . In particular, $[\mathcal{O}_w] - [\mathcal{O}_{y_1}] - [\mathcal{O}_{w_2}] \in H_1(X; \mathbb{Z})$ is represented by a cocycle contained in e which one may check is null-homologous, and hence $[\mathcal{O}_w] - [\mathcal{O}_{y_1}] - [\mathcal{O}_{w_2}] = 0 \in H_1(X; \mathbb{Z})$. Repeating this procedure for $[\mathcal{O}_{w_2}]$ and continuing recursively, we obtain the desired formula

$$[\mathcal{O}_w] = [\mathcal{O}_{y_1}] + \cdots + [\mathcal{O}_{y_n}]. \quad \square$$

As a consequence, we see that \mathcal{D} can be defined by requiring positivity on only finitely many orbits (cf. [Fri2, Theorem I] and [Wan, Theorem 1.6]).

Proposition 6.3. *The Fried cone \mathcal{D} is given by*

$$\mathcal{D} = \{u \in H^1(X; \mathbb{R}) \mid u(\mathcal{O}_y) > 0 \text{ for all } y \in \mathcal{Y}\}.$$

In particular, \mathcal{D} is an open, convex, polyhedral cone with finitely many rationally defined sides.

Proof. Since $\{\mathcal{O}_y\}_{y \in \mathcal{Y}} \subset \mathfrak{D}(\psi)$, the set in question clearly contains \mathcal{D} . For the other containment, we suppose $u(\mathcal{O}_y) > 0$ for all $y \in \mathcal{Y}$ and prove $u \in \mathcal{D}$. Every closed orbit $\mathcal{O} \in \mathfrak{D}(\psi)$ has the form \mathcal{O}_w for some closed path w in $\mathcal{G}(f)$. By Lemma 6.2 we may then write $[\mathcal{O}_w] = [\mathcal{O}_{y_1}] + \cdots + [\mathcal{O}_{y_n}]$, from which we obtain

$$u(\mathcal{O}_w) = u(\mathcal{O}_{y_1}) + \cdots + u(\mathcal{O}_{y_n}) > 0.$$

This proves the first claim of the proposition. The second claim is an immediate consequence of the first. □

Integrating a closed, flow-regular 1-form over a closed orbit yields a positive number, and hence we have

Proposition 6.4. *The Fried cone contains the cone of sections: $\mathcal{S} \subseteq \mathcal{D}$.*

Proof. Given any class $u \in \mathcal{S}$, we can represent it by a closed flow-regular 1-form $\omega = \omega^u$. For any closed curve γ we then have $u(\gamma) = \int_\gamma \omega^u$. In particular, for each closed orbit $\mathcal{O} \in \mathfrak{D}(\psi)$, this gives

$$u(\mathcal{O}) = \int_{\mathcal{O}} \omega^u > 0$$

by flow-regularity. Therefore $u \in \mathcal{D}$. □

We will see that the containment in this proposition is actually an equality. Proving this fact requires that we construct a section dual to any integral class $u \in \mathcal{D}$. This is a fairly technical construction which we carry out in Appendix A (Proposition A.12). With this result in hand, we may now prove Theorem A from the introduction.

Theorem A (Cone of sections). *There is an open convex cone $\mathcal{S} \subset H^1(X; \mathbb{R}) = H^1(G; \mathbb{R})$ containing \mathcal{A} (and thus containing u_0) such that a primitive integral class $u \in H^1(X_f; \mathbb{R})$ is dual to a section of ψ if and only if $u \in \mathcal{S}$. Moreover, \mathcal{S} is equal to the Fried cone \mathcal{D} , and there exist finitely many closed orbits $\mathcal{O}_1, \dots, \mathcal{O}_k$ of ψ such that $u \in H^1(X; \mathbb{R})$ lies in \mathcal{S} if and only if $u(\mathcal{O}_i) > 0$ for each $1 \leq i \leq k$. In particular, \mathcal{S} is an open, convex, polyhedral cone with finitely many rationally defined sides.*

Proof. The first assertion follows from Propositions 5.10 and 5.12. According to Proposition 6.4, we have $\mathcal{S} \subseteq \mathcal{D}$. On the other hand, Proposition A.12 shows that all rational points of \mathcal{D} are contained in \mathcal{S} . Since \mathcal{D} is the convex hull of its rational points, we also have $\mathcal{D} \subseteq \mathcal{S}$, and thus $\mathcal{S} = \mathcal{D}$ as claimed. If we take the finite set $\{\mathcal{O}_1, \dots, \mathcal{O}_k\}$ of closed orbits to be $\{\mathcal{O}_y\}_{y \in \mathcal{Y}}$, the remaining claims follow from Proposition 6.3. \square

7. The algebra and dynamics of cross sections to ψ

In this section we study the first return maps to cross sections of ψ . While technical considerations require that we restrict to certain nice cross sections, this is not a serious restriction as every integral class in \mathcal{S} is seen to be dual to such a section. In §7.2 we go on to describe how each of these first return maps, which need not be homotopy equivalences, gives a realization of G as an HNN-extension. This discussion culminates in the proof of Theorem B.

7.1. First return maps

Topologically, a cross section is merely a subset $\Theta = (\eta')^{-1}(x_0)$, for some flow-regular map $\eta': X \rightarrow \mathbb{S}^1$ and $x_0 \in \mathbb{S}^1$, which abstractly has the structure of an embedded topological graph. To aid in our analysis of the first return map f_Θ , here we describe a topological graph structure on Θ that is tailored to the complex X and the flow ψ . We then show that Θ moreover supports a linear graph structure with respect to which f_Θ is an expanding irreducible train track map.

Suppose that $\eta': X \rightarrow \mathbb{S}^1$ is a flow-regular map and $\Theta = (\eta')^{-1}(x_0)$ is a cross section of ψ for some $x_0 \in \mathbb{S}^1$. It follows that for each $x \in \mathbb{S}^1$ the subset $\Theta_x := (\eta')^{-1}(x)$ is again a cross section and all of these cross sections $\{\Theta_x\}_{x \in \mathbb{S}^1}$ are homotopic to each other.

Definition 7.1 (Compatible cross section). A cross section $\Theta \subset X$ is said to be *compatible with the foliation \mathcal{F}* (or more succinctly *\mathcal{F} -compatible*) if $\Theta \cap X^{(1)}$ consists of a finite set of points that all lie on vertex leaves of \mathcal{F} .

Remark 7.2. Since the union of vertex leaves of \mathcal{F} is dense in X , it is always possible to adjust any given cross section by a homotopy so that it is \mathcal{F} -compatible. Nevertheless, when we construct cross sections in §A.6 below, we will explicitly arrange for the constructed cross sections to be \mathcal{F} -compatible.

Every \mathcal{F} -compatible cross section may be equipped with a convenient topological graph structure:

Definition 7.3 (Standard graph structure). Let Θ be an \mathcal{F} -compatible cross section. The *standard (topological) graph structure* on Θ is then defined as follows. Firstly, Θ naturally inherits a topological graph structure in which the (finite) vertex set is $\Theta \cap X^{(1)}$ and every edge is an embedded arc in the interior of a 2-cell of X . We subdivide this initial graph structure by declaring the vertex set of Θ to be

$$V\Theta = \left(\bigcup_{s \geq 0} \psi_s(\Theta \cap X^{(1)}) \right) \cap \Theta.$$

Since \mathcal{F} -compatibility ensures that each point $\xi \in \Theta \cap X^{(1)}$ lies on a vertex leaf and thus that the ψ -orbit of ξ eventually becomes periodic, we see that this is indeed a finite subdivision making Θ into a finite topological graph.

Remark 7.4. We emphasize that if an \mathcal{F} -compatible cross section Θ is equipped with its standard graph structure, then each edge of Θ is an embedded arc contained in a 2-cell of X with both endpoints lying in vertex leaves of \mathcal{F} . In particular each edge of Θ is a transversal to \mathcal{F} ; this fact will play an important role in §10 below.

With these definitions, the arguments from [DKL, §7] establishing [DKL, Theorem C] essentially go through verbatim to prove that the first return map to an \mathcal{F} -compatible cross section is an expanding irreducible train track map. Rather than repeat the entire proof, here we briefly go through the argument, recalling the relevant results from [DKL].

Lemma 7.5. *Let Θ be an \mathcal{F} -compatible cross section equipped with its standard graph structure and first return map $f_\Theta: \Theta \rightarrow \Theta$. Then for all $n \geq 1$ the map f_Θ^n is locally injective on the interior of each edge of Θ . Moreover, $f_\Theta(V\Theta) \subset V\Theta$, and f_Θ is a regular topological graph map.*

Proof. The fact that $f_\Theta(V\Theta) \subset V\Theta$ is ensured by the definition of the standard graph structure (Definition 7.3). Lemma 7.2 of [DKL] shows that for every arc $\alpha: (-\delta, \delta) \rightarrow \sigma$ in the interior of a 2-cell σ of X and transverse to the flow ψ , the assignment $(t, s) \mapsto \psi_s(\alpha(t))$ defines a locally injective map $(-\delta, \delta) \times \mathbb{R}_{\geq 0} \rightarrow X$. Since each edge of Θ is by construction an embedded arc contained in the interior of 2-cell and is transverse to ψ , the proof of [DKL, Lemma 7.3] goes through verbatim to establish the remaining claims of the lemma. □

Lemma 7.6. *Let Θ be an \mathcal{F} -compatible cross section equipped with its standard graph structure. Then there exists a linear structure on Θ with respect to which $f_\Theta: \Theta \rightarrow \Theta$ is a linear graph map with irreducible transition matrix $A(f_\Theta)$ and spectral radius $\lambda(f_\Theta) > 1$.*

Proof. By assumption, the original map $f : \Gamma \rightarrow \Gamma$ defining the folded mapping torus X is an expanding irreducible train track map. The proof of [DKL, Lemma 7.4] thus shows that the topological graph map $f_\Theta : \Theta \rightarrow \Theta$ is “expanding on all scales” (see Appendix A.2 of [DKL]). It then follows from [DKL, Theorem A.3] that there exists a linear structure Λ on Θ with respect to which f_Θ is a linear graph map, the transition matrix $A(f_\Theta)$ is irreducible, and the spectral radius $\lambda(f_\Theta)$ is larger than 1. \square

Proposition 7.7. *Let Θ be an \mathcal{F} -compatible cross section equipped with the linear graph structure provided by Lemma 7.6. Then the first return map $f_\Theta : \Theta \rightarrow \Theta$ is an expanding irreducible train track map.*

Proof. By Lemmas 7.5–7.6 we know that f_Θ is a regular linear graph map for which $A(f_\Theta)$ is irreducible, $\lambda(f_\Theta) > 1$, and each iterate f_Θ^n is locally injective on each edge of Θ . To prove that all iterates of f_Θ are locally injective near the degree-2 vertices of Θ , we can follow the proof of [DKL, Lemma 2.12] verbatim with one difference: the assumption that the map was a homotopy equivalence was only used to guarantee that it was surjective. Here surjectivity of f_Θ is immediate from the fact that $A(f_\Theta)$ is irreducible. Therefore the linear graph map $f_\Theta : \Theta \rightarrow \Theta$ satisfies all the conditions of being an expanding irreducible train track map, and so the claim holds. \square

Convention 7.8. Henceforth, for any primitive integral element $u \in \mathcal{S}$, we let $\eta_u : X \rightarrow \mathbb{S}^1$ denote any flow-regular map whose induced map on the fundamental groups is $(\eta_u)_* = u : G \rightarrow \mathbb{Z}$ and whose fiber $\Theta_u \subset X$ is an \mathcal{F} -compatible cross section. Note that such a map η_u always exists by Proposition A.12. We equip Θ_u with the standard graph structure; the first return map $f_u = f_{\Theta_u} : \Theta_u \rightarrow \Theta_u$ is then an expanding irreducible train track map by Proposition 7.7. Lastly, we use ψ^u to denote the reparameterized semiflow for which the restriction of the time-1 map ψ_1^u to Θ_u is exactly f_u .

7.2. Sections and endomorphisms

Here we explain how every section Θ of ψ gives rise to splitting of G as a (possibly ascending) HNN-extension

$$G = Q_{[\Theta]} *_{\phi_{[\Theta]}} \tag{7.9}$$

of a finitely generated free group $Q_{[\Theta]}$ along an injective endomorphism $\phi_{[\Theta]}$ with respect to which $[\Theta] : G \rightarrow \mathbb{Z}$ is exactly the natural projection $Q_{[\Theta]} *_{\phi_{[\Theta]}} \rightarrow \mathbb{Z}$ sending the stable letter to 1 and $Q_{[\Theta]}$ to 0. This general situation should be compared with the special case $[\Theta] \in \mathcal{A}$, where Theorem 5.11 shows that we can find Θ such that in the splitting (7.9) we can take $Q_{[\Theta]} = \pi_1(\Theta)$ and let $\phi_{[\Theta]} = (f_\Theta)_*$ be the *isomorphism* induced by the first return map.

Let $u \in \mathcal{S}$ be any primitive integral class with dual cross section Θ_u as in Convention 7.8. After choosing a basepoint $v \in \Theta_u$ and a path $\beta \subset \Theta_u$ from $f_u(v)$ to v , we let $(f_u)_* : \pi_1(\Theta_u, v) \rightarrow \pi_1(\Theta_u, v)$ be the endomorphism defined by $(f_u)_*(\gamma) = \beta f(\gamma) \beta^{-1}$. If we let $\tau_u \in G = \pi_1(X, v)$ be the loop obtained by concatenating the flowline from v to $f_u(v)$ with β , then $u(\tau_u) = 1$ and therefore the conjugation $g \mapsto \tau_u^{-1} g \tau_u$ defines an automorphism $\Phi_u \in \text{Aut}(\ker(u))$ representing the monodromy φ_u . Furthermore, the splitting

(3.2) gives a realization of G as the HNN-extension $G = \ker(u) *_{\Phi_u}$ for which the natural projection is $u: G \rightarrow \mathbb{Z}$

The inclusion $\iota: \Theta_u \hookrightarrow X$ now induces a homomorphism

$$\iota_*: \pi_1(\Theta_u, v) \rightarrow \ker(u) \leq \pi_1(X, v) = G$$

that semiconjugates $(f_u)_*$ to Φ_u . Despite the fact that Θ_u need not π_1 -inject into X and that $(f_u)_*$ need not be an automorphism, van Kampen’s Theorem still applies to yield:

Lemma 7.10. *Let Θ_u be a section of ψ . Then, with the above notation, the homomorphism $\iota_*: \pi_1(\Theta_u, v) \rightarrow \ker(u)$ induces an isomorphism*

$$\langle \pi_1(\Theta_u, v), r \mid r^{-1}\gamma r = (f_u)_*(\gamma) \text{ for all } \gamma \in \pi_1(\Theta_u, v) \rangle \rightarrow \ker(u) *_{\Phi_u} = G$$

with respect to which $u: G \rightarrow \mathbb{Z}$ is given by sending $\pi_1(\Theta_u, v)$ to 0 and r to 1.

We can now prove Theorem B. For the statement, recall that the stable kernel K_ϕ of an arbitrary endomorphism $\phi: W \rightarrow W$ was defined in §2.4 and that ϕ naturally descends to a homomorphism $\bar{\phi}$ of the corresponding stable quotient W/K_ϕ .

Theorem B (Splittings and ascending HNN-extensions). *Let $u \in \mathcal{S}$ be a primitive integral class with \mathcal{F} -compatible dual section $\Theta_u \subset X$ and first return map $f_u: \Theta_u \rightarrow \Theta_u$. Let \mathcal{Q}_u be the stable quotient of $(f_u)_*$ and let $\phi_u = \overline{(f_u)_*}$ be the induced endomorphism of \mathcal{Q}_u . Then:*

- (1) f_u is an expanding irreducible train track map.
- (2) \mathcal{Q}_u is a finitely generated free group and $\phi_u: \mathcal{Q}_u \rightarrow \mathcal{Q}_u$ is injective.
- (3) G may be written as an HNN-extension

$$G \cong \mathcal{Q}_u *_{\phi_u} = \langle \mathcal{Q}_u, r \mid r^{-1}qr = \phi_u(q) \text{ for all } q \in \mathcal{Q}_u \rangle$$

such that $u: G \rightarrow \mathbb{Z}$ is given by the assignment $r \mapsto 1$ and $\mathcal{Q}_u \mapsto 0$.

- (4) If $J_u \leq \ker(u)$ denotes the image of $\pi_1(\Theta_u)$ induced by the inclusion $\Theta_u \hookrightarrow X$, then there is an isomorphism $\mathcal{Q}_u \rightarrow J_u$ conjugating ϕ_u to $\Phi_u|_{J_u}$ for some automorphism $\Phi_u \in \text{Aut}(\ker(u))$ representing the monodromy φ_u .
- (5) The topological entropy of f_u is equal to $\log(\lambda(\phi_u)) = \log(\lambda(\Phi_u|_{J_u}))$ and also to $\log(\Lambda(u))$.
- (6) $\ker(u)$ is finitely generated if and only if ϕ_u is an automorphism, in which case we have $\ker(u) \cong \mathcal{Q}_u$ and $\varphi_u = [\phi_u] \in \text{Out}(\mathcal{Q}_u)$.

Proof. We use the notation introduced before Lemma 7.10; in particular, we implicitly use $v \in \Theta_u$ as a basepoint and take $\Phi_u \in \text{Aut}(\ker(u))$ to be conjugation by $\tau_u \in \pi_1(X) = G$. Then Φ_u represents the monodromy $\varphi_u \in \text{Out}(\ker(u))$. Item (1) was established by Proposition 7.7, and by Proposition 2.2 this further implies that the topological entropy of f_u is equal to $\log(\lambda((f_u)_*))$. Corollary 2.8 then implies item (2) and additionally proves $\lambda(\phi_u) = \lambda((f_u)_*)$, thus establishing the first part of (5). The remaining assertions $\lambda(\phi_u) = \lambda(\Phi_u|_{J_u})$ and $\Lambda(u) = \lambda(\phi_u)$ of (5) will follow from (4) and (3), respectively.

Since $\iota_*: \pi_1(\Theta_u) \rightarrow \ker(u)$ semiconjugates $(f_u)_*$ to Φ_u , which is injective, it follows that the stable kernel $K_{(f_u)_*}$ is contained in $\ker(\iota_*)$. Therefore ι_* factors through the stable quotient Q_u , thus yielding a pair of homomorphisms

$$\pi_1(\Theta_u) \rightarrow Q_u \rightarrow \ker(u).$$

Since these homomorphisms respectively semiconjugate $(f_u)_*$ to $\phi_u = \overline{(f_u)_*}$ and ϕ_u to Φ_u , they induce canonical homomorphisms

$$\pi_1(\Theta_u)_{*(f_u)_*} \rightarrow Q_u *_{\phi_u} \rightarrow \ker(u) *_{\Phi_u} = G$$

that respect the natural projections of these groups to \mathbb{Z} . The above composition is an isomorphism by Lemma 7.10, and thus each map is itself an isomorphism. This proves (3). Since Q_u embeds into $Q_u *_{\phi_u}$, this furthermore implies that $Q_u \rightarrow \ker(u)$ is injective. Therefore $Q_u \rightarrow \ker(u)$ is an isomorphism onto its image, which, by definition, is equal to J_u . This isomorphism establishes (4).

It remains to prove (6). The fact that $u: G \rightarrow \mathbb{Z}$ is the natural projection $Q_u *_{\phi_u} \rightarrow \mathbb{Z}$ defined by $r \mapsto 1$ and $Q_u \mapsto 0$ implies that $\ker(u)$ is equal to the normal closure $ncl(Q_u)$ of Q_u in G . The HNN-extension presentation $G = Q_u *_{\phi_u}$ further implies that $ncl(Q_u) = \bigcup_{n=0}^{\infty} r^n Q_u r^{-n}$. Thus if ϕ_u is an automorphism, then $\ker(u)$ is equal to Q_u and is thus finitely generated by (2). On the other hand, if ϕ_u is not an automorphism, then $\phi_u: Q_u \rightarrow Q_u$ is necessarily nonsurjective (by (2)), and thus $ncl(Q_u) = \bigcup_{n=0}^{\infty} r^n Q_u r^{-n}$ is the union of an infinite strictly increasing chain of subgroups of $\ker(u)$. This shows that $\ker(u)$ is not finitely generated and completes the proof of (6). \square

Example 7.11 (The ascending HNN-extension for $u_1 = [z_1]$). Let us calculate the splitting described by Theorem B(2) for the homomorphism $u_1 = [z_1]$, where

$$z_1 = [v_2^*, v_3^*] + [v_3^*, v_5^*] + 2[v_4^*, v_6^*] + [v_5^*, v_1^*] + [v_6^*, v_2^*]$$

is the cocycle introduced in Example 5.6. Recall from (3.4) that in our running example, G has the two-generator one-relator presentation

$$G = \langle \gamma, r \mid \gamma^{-1} r \gamma^{-1} r \gamma^{-1} r^{-1} \gamma r \gamma r \gamma r^{-3} = 1 \rangle. \tag{7.12}$$

Thus we may apply Brown’s method [Bro] to compute the kernel $\ker(u_1)$ and corresponding monodromy $\varphi_1 = \varphi_{u_1}$. Working in the unsubdivided cell structure on X , the loops $\gamma, r \subset X$ generating G may be realized as 1-chains

$$\gamma = [[v_1, v_2] + [v_2, v_4] + [v_4, v_5] - [v_3, v_5] - [v_1, v_3]], \quad r = [[v_1, v_3] + [v_3, v_5] + [v_5, v_1]].$$

Pairing these with z_1 we find that $u_1(\gamma) = -1$ and $u_1(r) = 2$. Setting $\beta = r\gamma^2$ and $\tau = \gamma^{-1}$, we therefore have $u_1(\tau) = 1$ and $u_1(\beta) = 0$. Conjugating $k \mapsto \tau^{-1}k\tau$ by τ thus yields an automorphism $\Phi_1 \in \text{Aut}(\ker(u_1))$ representing the monodromy φ_1 . Noting that $\gamma = \tau^{-1}$ and $r = \beta\tau^2$, we may rewrite the defining relation of (7.12) in terms of the generators β and τ to obtain the following presentation for G :

$$G = \langle \beta, \tau \mid \tau\beta\tau^3\beta\tau\beta^{-1}\tau^{-1}\beta\tau\beta\tau^{-1}\beta^{-1}\tau^{-2}\beta^{-1}\tau^{-2}\beta^{-1} = 1 \rangle.$$

We now represent G as an HNN-extension of another one-relator group, following the method introduced by McCool and Schupp [MS]. For each $i \in \mathbb{Z}$ denote $\beta_i = \tau^{-i} \beta \tau^i$. Then the defining relation above may be rewritten as

$$\beta_{-1} \beta_{-4} \beta_{-5}^{-1} \beta_{-4} \beta_{-5} \beta_{-4}^{-1} \beta_{-2}^{-1} \beta_0^{-1} = 1.$$

Hence we can rewrite the above presentation of G as

$$G = \left\langle \beta_{-5}, \dots, \beta_0, \tau \mid \begin{array}{l} \beta_{-1} \beta_{-4} \beta_{-5}^{-1} \beta_{-4} \beta_{-5} \beta_{-4}^{-1} \beta_{-2}^{-1} \beta_0^{-1} = 1, \\ \tau^{-1} \beta_i \tau = \beta_{i+1} \text{ for } i = -5, \dots, -1 \end{array} \right\rangle. \tag{7.13}$$

Consider now the one-relator group

$$B = \langle \beta_{-5}, \dots, \beta_0 \mid \beta_{-1} \beta_{-4} \beta_{-5}^{-1} \beta_{-4} \beta_{-5} \beta_{-4}^{-1} \beta_{-2}^{-1} \beta_0^{-1} = 1 \rangle.$$

By Magnus' Freiheitssatz theorem for the one-relator group B , the subgroups $L = \langle \beta_{-5}, \dots, \beta_{-1} \rangle \leq B$ and $J = \langle \beta_{-4}, \dots, \beta_0 \rangle \leq B$ are free groups of rank 5 and the indicated generating sets of L and J are their free bases. Hence the map $\phi_1: L \rightarrow J$, sending $\beta_i \mapsto \beta_{i+1}$ for $i = -5, \dots, -1$, is an isomorphism. Moreover, since the defining relation of B implies that $\beta_0 = \beta_{-1} \beta_{-4} \beta_{-5}^{-1} \beta_{-4} \beta_{-5} \beta_{-4}^{-1} \beta_{-2}^{-1}$, we see that the generator β_0 may be eliminated from the presentation of B and thus $L = \langle \beta_{-5}, \dots, \beta_{-1} \rangle$ is in fact equal to B . Therefore $\phi_1: L \rightarrow J$ is actually an injective endomorphism $\phi_1: B \rightarrow B$ of the rank-5 free group B . This shows that the presentation (7.13) gives a splitting of G as an HNN-extension

$$G = B *_{\phi_1} = \langle B, \tau \mid \tau^{-1} b \tau = \phi_1(b) \text{ for all } b \in B \rangle.$$

Note that with respect to the free basis $B = \langle \beta_{-5}, \dots, \beta_{-1} \rangle$, the injective endomorphism ϕ_1 is given by $\phi_1(\beta_{-1}) = \beta_{-1} \beta_{-4} \beta_{-5}^{-1} \beta_{-4} \beta_{-5} \beta_{-4}^{-1} \beta_{-2}^{-1}$ and $\phi_1(\beta_i) = \beta_{i+1}$ for $i = -5, \dots, -2$. Using the Stallings folding method [KM, Sta] we find that ϕ_1 is not surjective: The fact that β_{-5} occurs in $\phi_1(\beta_{-1})$ twice leads to the conclusion that the rose with petals labeled by the words $\phi_1(\beta_{-5}), \dots, \phi_1(\beta_{-1})$ does not fold onto the rose with petals labeled $\beta_{-5}, \dots, \beta_{-1}$. Therefore (7.13) is a strictly ascending HNN-extension presentation of G with base $B = F(\beta_{-5}, \dots, \beta_{-1})$ and stable letter τ along the injective but nonsurjective endomorphism ϕ_1 of B .

It follows that $\ker(u_1)$ is not finitely generated because it is equal to the infinite union $ncl_G(B) = \bigcup_{i=0}^{\infty} \tau^i B \tau^{-i}$ of a strictly increasing chain of subgroups. We also note that the free group endomorphism ϕ_1 is just the restriction $\Phi_1|_B$ of Φ_1 to $B \leq G$. Using the software package xtrain, one may verify that ϕ_1 does admit an expanding irreducible train track representative, and the spectral radius of the transition matrix for this representative is ≈ 1.35827 . By Proposition 2.2 we may thus conclude that ϕ_1 has a stretch factor of $\lambda(\phi_1) \approx 1.35827$.

Remark 7.14. In light of the proof of Theorem B, one may also see that $\ker(u_1)$ is infinitely generated by noting, as we have done in Example 5.6, that the first return map f_1 induces an injective but not surjective endomorphism $(f_1)_*$ of $\pi_1(\Theta_1)$: The fact that

$(f_1)_*$ is injective implies that its stable kernel is trivial and Θ_1 π_1 -injects into G . Thus the HNN-extension provided by Theorem B(2) is simply

$$G \cong \pi_1(\Theta_1) *_{(f_1)_*}.$$

As $(f_1)_*$ is not surjective, Theorem B(6) now implies that $\ker(u_1)$ is not finitely generated.

Example 7.15 (The splitting for $u_2 = [z_2]$). Let us also calculate the splitting of G for the homomorphism $u_2 = [z_2]$ determined by the cocycle

$$z_2 = [v_1^*, v_2^*] + [v_2^*, v_3^*] - [v_4^*, v_5^*] - [v_5^*, v_6^*] + [v_1^*, v_3^*] + 2[v_6^*, v_2^*]$$

introduced in Example 5.7. Again, by realizing the generators $\gamma, r \in G$ as 1-cycles and pairing with z_2 , we find that $u_2(\gamma) = -1$ and $u_2(r) = 1$. Setting $\beta = r\gamma$ and $\tau = \gamma^{-1}$, we may rewrite the presentation (7.12) for G in terms of the generators β and τ as we did in (7.13) above to find that

$$G = \left\langle \beta_{-4}, \dots, \beta_0, \tau \mid \begin{array}{l} \beta_{-1}\beta_{-3}\beta_{-4}^{-1}\beta_{-3}^2\beta_{-2}^{-1}\beta_{-1}^{-1}\beta_0^{-1} = 1, \\ \tau^{-1}\beta_i\tau = \beta_{i+1} \text{ for } i = -4, \dots, -1 \end{array} \right\rangle, \quad (7.16)$$

where we again denote $\beta_i = \tau^{-i}\beta\tau^i$ for $i \in \mathbb{Z}$. Setting

$$B' = \langle \beta_{-4}, \dots, \beta_0 \mid \beta_{-1}\beta_{-3}\beta_{-4}^{-1}\beta_{-3}^2\beta_{-2}^{-1}\beta_{-1}^{-1}\beta_0^{-1} = 1 \rangle,$$

Magnus' Freiheitssatz now tells us that the subgroups $L' = \langle \beta_{-4}, \dots, \beta_{-1} \rangle \leq B'$ and $J' = \langle \beta_{-3}, \dots, \beta_0 \rangle \leq B'$ are freely generated by the indicated generating sets, and the assignment $\beta_i \mapsto \beta_{i+1}$ for $i = -4, \dots, -1$ gives an isomorphism $\phi_2: L' \rightarrow J'$. However, as the generator β_0 of B' is evidently superfluous, we in fact have $L' = B'$. Thus ϕ_2 is actually an endomorphism $\phi_2: B' \rightarrow B'$, and with respect to the free basis $B' = \langle \beta_{-4}, \dots, \beta_{-1} \rangle$ it takes the form $\phi_2(\beta_{-1}) = \beta_{-1}\beta_{-3}\beta_{-4}^{-1}\beta_{-3}^2\beta_{-2}^{-1}\beta_{-1}^{-1}$ and $\phi_2(\beta_i) = \beta_{i+1}$ for $i = -4, -3, -2$. Here the Stallings folding method shows that ϕ_2 is both injective and surjective, and therefore (7.13) presents G as an HNN-extension $G \cong B' *_{\phi_2}$ over the free base group B' along the automorphism ϕ_2 . It now follows that $\ker(u_2) = B'$ is finitely generated and in fact has rank 4. Theorem B then implies that ϕ_2 represents the monodromy $\varphi_2 = \varphi_{u_2} \in \text{Out}(\ker(u_2))$, and the program xtrain calculates that the stretch factor here is $\lambda(\varphi_2) = \lambda(\phi_2) \approx 1.632992$.

8. Cross sections and homology

In this section we return our attention to the abelian group $H = H_1(G; \mathbb{Z})/\text{torsion}$ and describe bases for H and associated coordinate systems on $H^1(G; \mathbb{R}) = H^1(X; \mathbb{R})$ that are tailored to each integral class $u \in \mathcal{S}$. Viewing a given primitive integral class $u \in \mathcal{S}$ as a surjective homomorphism $G \rightarrow \mathbb{Z}$, we let $\tilde{X}_u \rightarrow X$ denote the infinite cyclic cover corresponding to the kernel of u . The inclusion $\Theta_u \rightarrow X$ lifts to an inclusion $\Theta_u \rightarrow \tilde{X}_u$ since $\pi_1(\Theta_u) < \ker(u)$. Because H is the maximal, abelian, torsion-free quotient of G ,

the homomorphism u factors as a composition $G \rightarrow H \rightarrow \mathbb{Z}$, and we let $H_u < H$ denote the kernel of the homomorphism $H \rightarrow \mathbb{Z}$.

Observe that H_u is the image in H of $\pi_1(\tilde{X}_u) = \ker(u) < G$ under the homomorphism $G \rightarrow H$. In fact, we claim that H_u is the image of $\pi_1(\Theta_u)$ (if $u \in \mathcal{A}$, then $\pi_1(\Theta_u) = \pi_1(\tilde{X}_u)$ and this is obvious). Before we prove the claim, we note that $\pi_1(\Theta_u)$ is only well-defined up to conjugacy in general, but any two conjugates have the same image in H . Now, given an arbitrary element $[\alpha]$ of H_u , let α be a loop in \tilde{X}_u sent to $[\alpha]$. Lifting the semiflow ψ to \tilde{X}_u we flow α into some lift of Θ_u in \tilde{X}_u , producing a loop α' in some lift of Θ_u , freely homotopic to α . Pushing down to X , we find $[\alpha] = [\alpha'] \in H_u$, and hence $[\alpha]$ is in the image of $\pi_1(\Theta_u)$, as required.

Write $p: \tilde{X} \rightarrow X$, let $\tilde{\Theta}_u = p^{-1}(\Theta_u) \subset \tilde{X}$, and let $\tilde{\Theta}_u$ denote some component. Since $\tilde{X} \rightarrow \tilde{X}_u$ is the cover corresponding to the kernel of $\pi_1(\tilde{X}_u) \rightarrow H_u$, and $\pi_1(\Theta_u)$ maps onto H_u , H_u is precisely the stabilizer in H of $\tilde{\Theta}_u$. Furthermore, H/H_u is the covering group of $\tilde{X}_u \rightarrow X$, and so $H/H_u \cong \mathbb{Z}$. It follows that there is a (noncanonical) splitting $H = H_u \oplus \mathbb{Z}$. Since H has rank b , H_u has rank $b - 1$.

A useful alternative way to describe this splitting is the following. Consider the subgroup of cohomology fixed by f_u , denoted

$$H^1(\Theta_u; \mathbb{Z})^{f_u} \subset H^1(\Theta_u; \mathbb{Z}).$$

Given a homotopy class α of loops in Θ_u , evaluating cohomology classes on α defines an element of $\text{Hom}(H^1(\Theta_u; \mathbb{Z})^{f_u}, \mathbb{Z})$. This is a quotient of $H_1(\Theta_u; \mathbb{Z})$, and in fact it is the maximal quotient on which f_u acts trivially. On the other hand, the inclusion $\Theta_u \subset X$ induces a map $H_1(\Theta_u; \mathbb{Z}) \rightarrow H$, and it is straightforward to check that this map factors as a composition $H_1(\Theta_u; \mathbb{Z}) \rightarrow \text{Hom}(H^1(\Theta_u; \mathbb{Z})^{f_u}, \mathbb{Z}) \rightarrow H$. Moreover, this second map is actually an injection, so since $H_1(\Theta_u; \mathbb{Z})$ is the abelianization of $\pi_1(\Theta_u)$, from the discussion above, we have a natural isomorphism

$$H_u = \text{Hom}(H^1(\Theta_u; \mathbb{Z})^{f_u}, \mathbb{Z}).$$

We choose a basis s_1, \dots, s_{b-1} for H_u and $w \in H$ a generator for the complement. Here the group structure is additive, but it will also be convenient to have a multiplicative basis. Specifically, we let $t_1, \dots, t_{b-1}, x \in H_1(X; \mathbb{R}_+)$ be given by $t_i = e^{s_i}$ and $x = e^w$. Formally, this just means that t_i and x are the images of s_i and w , respectively, under the isomorphism $H_1(X; \mathbb{R}) \rightarrow H_1(X; \mathbb{R}_+)$ determined by the isomorphism $\exp: \mathbb{R} \rightarrow \mathbb{R}_+$ of the coefficient groups. When convenient we write $\mathbf{s} = (s_1, \dots, s_{b-1})$ and $\mathbf{t} = e^{\mathbf{s}} = (t_1, \dots, t_{b-1})$. In this way, the integral group ring of H is naturally identified with the ring of Laurent polynomials

$$\mathbb{Z}[H] \cong \mathbb{Z}[t_1^{\pm 1}, \dots, t_{b-1}^{\pm 1}, x^{\pm 1}] = \mathbb{Z}[\mathbf{t}^{\pm 1}, x^{\pm 1}].$$

It is also convenient to think of t_1, \dots, t_{b-1}, x as generators for the covering group H of $\tilde{X} \rightarrow X$, and we will do so. Finally, we assume that the generator x is chosen so that it translates \tilde{X} positively with respect to $\tilde{\psi}$.

It will be convenient to have a concise means for referring to all of these choices, so we make the following

Definition 8.1. Given a primitive integral class $u \in \mathcal{S}$ with dual cross section Θ_u , we say that the splitting $H = H_u \oplus \mathbb{Z}$ and the associated bases \mathbf{s}, w and $\mathbf{t} = e^{\mathbf{s}}, x = e^w$ described above are *adapted* to u .

Convention 8.2. The pairing between homology and cohomology allows us to view the homology classes $s_1, \dots, s_{b-1}, w, t_1, \dots, t_{b-1}, x$ as real-valued functions on $H^1(X; \mathbb{R})$. In fact, $(\mathbf{s}, w): H^1(X; \mathbb{R}) \rightarrow \mathbb{R}^b$ are linear coordinate functions, and in these coordinates

$$u = (\mathbf{s}(u), w(u)) = (\mathbf{0}, 1).$$

Furthermore, the formal expressions $\mathbf{t} = e^{\mathbf{s}}$ and $x = e^w$ can be interpreted as a functional relation

$$(\mathbf{t}, x) = (e^{\mathbf{s}}, e^w): H^1(X; \mathbb{R}) \rightarrow \mathbb{R}_+^b.$$

The bases \mathbf{t} and \mathbf{t}, x for H_u and H , respectively, determine isomorphisms of integral group rings with rings of integral Laurent polynomials:

$$\begin{aligned} \mathbb{Z}[H_u] &\cong \mathbb{Z}[\mathbf{t}^{\pm 1}] = \mathbb{Z}[t_1^{\pm 1}, \dots, t_{b-1}^{\pm 1}], \\ \mathbb{Z}[H] &\cong \mathbb{Z}[\mathbf{t}^{\pm 1}, x^{\pm 1}] = \mathbb{Z}[t_1^{\pm 1}, \dots, t_{b-1}^{\pm 1}, x^{\pm 1}]. \end{aligned}$$

Whenever we have chosen bases adapted to u , we will freely use these identifications.

Example 8.3 (The cones $\mathcal{A} \subset \mathcal{S}$). Let us return to the running example introduced in Example 3.3. If we take u_0 to be the homomorphism associated to the original splitting $G = F_N \rtimes_{\varphi} \mathbb{Z}$, then Θ_{u_0} is the original graph Γ shown in Figure 1. Let $H_0 = H_{u_0}$ as above giving rise to the splitting $H \cong H_0 \oplus \mathbb{Z}$. Here $H_0 \leq H$ is the cyclic group $\langle t \rangle$ generated by the image of either γ_1, γ_2 , or γ_3 in G under $G \rightarrow H = G^{\text{ab}}$, and its complement $\mathbb{Z} = \langle x \rangle$ is generated by the image of the stable letter $r \in G$.

Let s, w be the corresponding additive elements with $e^s = t$ and $e^w = x$, which we represent by 1-cycles

$$s = [v_1, v_3] - [v_2, v_3] - [v_1, v_2] \quad \text{and} \quad w = [v_1, v_3] + [v_3, v_5] + [v_5, v_1].$$

Using (s, w) as coordinates on $H^1(X; \mathbb{R})$ we can express the classes u_0, u_1, u_2 as

$$u_0 = (0, 1), \quad u_1 = (-1, 2), \quad u_2 = (-1, 1).$$

All of these classes are dual to cross sections (see Examples 5.6–5.7) and are therefore contained in \mathcal{S} . By construction we also know that \mathcal{A} contains u_0 . On the other hand, since $\ker(u_1)$ is infinitely generated (see Example 7.11) we necessarily have $u_1 \notin \mathcal{A}$. This gives us some information about \mathcal{A} and \mathcal{S} , and we may in fact calculate these cones explicitly as follows.

First recall that every class in \mathcal{A} may be represented by a positive cocycle on the unsubdivided cell structure. We note also that the coboundary operator $\delta: C^1(X; \mathbb{R}) \rightarrow C^2(X; \mathbb{R})$ is linear, and hence finding 1-cocycles that represent points on the boundary

of \mathcal{A} amounts to finding $c \in C^1(X; \mathbb{R})$ that maximizes/minimizes the linear function $c \mapsto c([v_1, v_3] - [v_2, v_3] - [v_1, v_2])$, subject to the system of linear equations and inequalities

$$\begin{cases} \delta(c) = 0, \\ c([v_1, v_3] + [v_3, v_5] + [v_5, v_1]) = 1, \\ c([v_i, v_j]) \geq 0 \quad \text{for all } i, j \text{ with } [v_i, v_j] \text{ a 1-cell.} \end{cases}$$

Here the first equation corresponds to the requirement that c be a cocycle. The third inequality is the requirement that the cocycle is nonnegative. The second condition is a normalization that the cocycle evaluates to 1 on w . The function we are trying to maximize/minimize is the s -coordinate, and so if $\bar{\mathcal{A}}$ contained the ray through $(1, 0)$ or $(-1, 0)$, then there would be no maximum or minimum, respectively, subject to the constraints. This is a linear programming problem that can easily be solved in Mathematica or by hand. Doing so, we find that the minimum and maximum values of the s -coordinate subject to these constraints are $-1/2$ and $1/2$, respectively. The \mathcal{A} cone is therefore as illustrated in Figure 8, and we note that u_1 in fact lies on its boundary $\partial\mathcal{A}$.

We now calculate \mathcal{S} using the homological characterization given by Theorem A. Following the explicit construction of §6, we find that in the case of our running example the finitely many closed orbits \mathcal{O}_i appearing in the statement of Theorem A consist of seven closed loops which are freely homotopic in X to the following seven elements of $\pi_1(X) = \langle \gamma_1, \gamma_2, \gamma_3, r \rangle$:

$$\left\{ \begin{array}{l} r\gamma_3^{-1}\gamma_1^{-1}, \quad r^2\gamma_3^{-1}\gamma_1^{-1}\gamma_2^{-1}, \quad r^2\gamma_3^{-1}, \quad r^3, \\ r^2\gamma_3^{-1}r^2, \quad r^2\gamma_3^{-1}\gamma_1^{-1}\gamma_2^{-1}r, \quad r^2\gamma_3^{-1}\gamma_1^{-1}r \end{array} \right\}.$$

In terms of the additive basis $\{s, w\}$, the images of these elements in H are

$$(-2, 1), \quad (-3, 2), \quad (-1, 2), \quad (0, 3), \quad (-1, 4), \quad (-3, 3), \quad (-2, 3).$$

The cohomology classes u which are positive on all of these vectors are exactly those whose (s, w) -coordinates satisfy both $w(u) > 0$ and $w(u) > 2s(u)$; such classes comprise the cone of sections \mathcal{S} and are illustrated in Figure 8.




Fig. 8. The cones $\mathcal{A} \subset \mathcal{S} \subset H^1(X; \mathbb{R})$ for our running example.

Remark 8.4. When $u \in \mathcal{A}$, Theorem 5.11 ensures that $\ker(u)$ is finitely generated. However, Theorem B makes no claims about the finite or nonfinite generation of $\ker(u)$ when $u \in \mathcal{S} \setminus \mathcal{A}$. Indeed, since $\ker(u_1)$ is infinitely generated and $\ker(u_2)$ is finitely generated (see Examples 7.11 and 7.15), the above Example 8.3 illustrates that both outcomes are possible for classes in $\mathcal{S} \setminus \mathcal{A}$.

9. Graph modules

Here we describe another, more combinatorial module that will connect the McMullen polynomial to the geometry of splittings of G . Let $\tilde{\psi}^u$ be the lift of the reparameterized flow $\psi^u = \psi^{\Theta_u}$ so that $\tilde{\psi}_1^u$ maps $\tilde{\Theta}_u$ to $x(\tilde{\Theta}_u)$ (recall that $x \in H$ is a generator for the complement of H_u , which translates positively with respect to $\tilde{\psi}^u$). Thus $x^{-1} \circ \tilde{\psi}_1^u$ restricts to a graph map on $\tilde{\Theta}_u$; in fact, $x^{-1} \circ \tilde{\psi}_1^u$ is just a lift of the first return map $f_u : \Theta_u \rightarrow \Theta_u$ of the flow ψ^u to $\Theta_u \subset X$.

The covering group H acts on (unoriented) edges of $\hat{\Theta}_u$, as with transversals, by taking preimages, so that $h \in H$ acts on an edge e by

$$h \cdot e = e \cdot h = h^{-1}(e).$$

We can organize the edges of $\hat{\Theta}_u$ nicely as follows. Let $\bar{E} = \{\bar{\sigma}_1, \dots, \bar{\sigma}_m\}$ denote the set of (unoriented) edges of Θ_u . For each $\bar{\sigma}_j$ we choose an edge $\sigma_j \in p^{-1}(\bar{\sigma}_j)$ contained in $\tilde{\Theta}_u$, and we denote this set by $E = \{\sigma_1, \dots, \sigma_m\}$. Because the action of H is free, the action on every orbit is also free. Thus the free \mathbb{Z} -module on the set of edges of $\hat{\Theta}_u$, which is naturally a $\mathbb{Z}[H]$ -module, can be identified with the module

$$\mathbb{Z}[H]^E = \left\{ \sum_{j=1}^m \sum_{h \in H} a_{j,h} h \cdot \sigma_j \mid \begin{array}{l} a_{j,h} \in \mathbb{Z} \text{ with } a_{j,h} = 0 \\ \text{for all but finitely many } h \end{array} \right\}.$$

Furthermore, since the elements of E were chosen to lie in $\tilde{\Theta}_u$, we see that the set of edges in $\hat{\Theta}_u$ is naturally identified with the H_u -orbit of E , and hence the free \mathbb{Z} -module on the edges of $\hat{\Theta}_u$ is precisely $\mathbb{Z}[H_u]^E$.

We now define the module of $\hat{\Theta}_u$, denoted $T(\hat{\Theta}_u)$, to be the largest quotient $\mathbb{Z}[H]^E \rightarrow T(\hat{\Theta}_u)$ in which the image $[e]$ of any edge e satisfies the relations

$$[e] = [e_1] + \dots + [e_k],$$

where the edge path $\tilde{\psi}_1^u(e) = e_1 \dots e_k$ is the image of e under $\tilde{\psi}_1^u$. We refer to the above relations as the *basic relations* of $T(\hat{\Theta}_u)$. More precisely:

Definition 9.1 (Module of $\hat{\Theta}_u$). Let $M \leq \mathbb{Z}[H]^E$ denote the submodule generated by all elements of the form $e - e_1 - \dots - e_k$, with e, e_1, \dots, e_k as above. Note that the set of such elements is preserved by H . The *module of $\hat{\Theta}_u$* is defined to be the quotient $\mathbb{Z}[H]$ -module $T(\hat{\Theta}_u) = \mathbb{Z}[H]^E / M$.

9.1. Connection to transition matrices

Here we develop an alternative description of the module $T(\widehat{\Theta}_u)$ closely related to the transition matrix of f_u , ultimately culminating in a finite presentation. Let $\mathbf{t} = t_1, \dots, t_{b-1}, x$ be a multiplicative basis for H adapted to u as in Definition 8.1. For each edge $e \in \widetilde{\Theta}_u$, translating the edge path $\widetilde{\psi}_1^u(e) = e_1 \cdots e_k$ by x^{-1} gives a new edge path $x^{-1} \circ \widetilde{\psi}_1^u(e) = e'_1 \cdots e'_k = x^{-1}(e_1 \cdots e_k)$, which now lies in $\widetilde{\Theta}_u$. Since $x^{-1} \circ \widetilde{\psi}_1^u(e) = \widetilde{\psi}_1^u(x^{-1}(e)) = \widetilde{\psi}_1^u(x \cdot e)$, the corresponding basic relation becomes

$$x \cdot [e] = [e'_1] + \cdots + [e'_k].$$

Now, for $e = \sigma_j \in E$, if we rewrite the sum on the right as a $\mathbb{Z}[H_u]$ -linear combination of the edges in E , this relation becomes

$$x \cdot [\sigma_j] = \sum_{i=1}^s A_{ij} \cdot [\sigma_i].$$

This gives a matrix $A = (A_{ij})$ with entries in $\mathbb{Z}[H_u]$; to clarify the dependence on u , we may sometimes write this matrix as $A_u = (A_{u,ij})$. Given the isomorphism with the ring $\mathbb{Z}[H_u] = \mathbb{Z}[\mathbf{t}^{\pm 1}]$ of integral Laurent polynomials, A_u becomes a matrix $A_u(\mathbf{t})$ with coefficients in $\mathbb{Z}[\mathbf{t}^{\pm 1}]$.

Observe that $x^{-1} \circ \widetilde{\psi}_1^u$ is a lift of the first return map $f_u: \Theta_u \rightarrow \Theta_u$, and so the following is not so surprising.

Proposition 9.2. *For every primitive integral $u \in \mathcal{S}$, the integral matrix $A_u(\mathbf{1})$, obtained by evaluating $A_u(\mathbf{t})$ at $\mathbf{t} = \mathbf{1} = (1, \dots, 1)$, is exactly the transition matrix of the train track map f_u .*

Proof. The (i, j) -entry of A_u is a \mathbb{Z} -linear combination of edges in the H_u -orbit of σ_i . Moreover, for each edge $e \in H_u \cdot \sigma_i$, the coefficient of e in $A_{u,ij}$ is the number of times the path $x^{-1} \circ \widetilde{\psi}_1^u(\sigma_j)$ crosses e . As these edges are exactly those which project to the edge $\bar{\sigma}_i$ in Θ_u , the sum of these coefficients (namely $A_{z,ij}(\mathbf{1})$) is just the number of times the projected path $p(x^{-1} \circ \widetilde{\psi}_1^u(\sigma_j)) \subset \Theta_u$ crosses $\bar{\sigma}_i$. But this projected path is exactly $f_u(\bar{\sigma}_j)$, and so the number of times it crosses $\bar{\sigma}_i$ is the (i, j) -entry of the transition matrix for f_u . \square

Furthermore, as the submodule $M \leq \mathbb{Z}[H]^E$ of relations defining $T(\widehat{\Theta}_u)$ is generated, as a $\mathbb{Z}[H]$ -module, by the elements

$$x \cdot \sigma_j - \sum_{i=1}^m A_{ij}(\mathbf{t}) \cdot \sigma_i$$

for $j = 1, \dots, m$, we have the following useful presentation for $T(\widehat{\Theta}_u)$.

Proposition 9.3. *The module $T(\widehat{\Theta}_u)$ is finitely presented as a $\mathbb{Z}[H]$ -module by*

$$\mathbb{Z}[H]^E \xrightarrow{D_u} \mathbb{Z}[H]^E \rightarrow T(\widehat{\Theta}_u) \rightarrow 0,$$

where D_u is the square $m \times m$ matrix $xI - A_u(\mathbf{t})$ with entries in $\mathbb{Z}[H]$.

We thus have the following immediate corollary:

Corollary 9.4. *The g.c.d. of the fitting ideal of $T(\widehat{\Theta}_u)$ is given by*

$$\det(xI - A_u(\mathbf{t})) \in \mathbb{Z}[H].$$

Example 9.5. Here we return again to our running example and calculate A_u and $\det(xI - A_u)$ in this case. We take $u = u_0$ so that Θ_u is the original fiber Γ ; the infinite cyclic cover $\widetilde{\Theta}_u \rightarrow \Gamma$ is then as shown in Figure 9. Here the generator t of H_u is pictured as translating to the right. The edges \bar{E} of Γ are labeled as in Figure 1, and we have used the same labels for the representative edges E of $\widetilde{\Theta}_u$; the rest of the edges of $\widetilde{\Theta}_u$ are obtained (and labeled) as translates of these representative edges under H_u .




Fig. 9. The infinite cyclic cover $\widetilde{\Theta}_u$ of Θ_u for our running example $f: \Gamma \rightarrow \Gamma$ when $u = u_0$ is the homomorphism associated to the original splitting $G = F_3 \rtimes_{\varphi} \mathbb{Z}$.

The module $T(\widehat{\Theta}_u)$ is then generated by the edges $\{\sigma_1, \dots, \sigma_4\} = \{a, b, c, d\}$. For each $\sigma_j \in \{a, b, c, d\}$, we calculate $x^{-1} \circ \psi_1(\sigma_j)$ by lifting the edge path $f(\bar{\sigma}_j) \subset \Gamma$ to $\widetilde{\Theta}_u$ and expressing it in the form $\sum_{i=1}^4 A_{ij}(t) \cdot [\sigma_j]$ with $A_{ij}(t) \in \mathbb{Z}[t^{\pm}]$. This gives the matrix

$$A(t) = \begin{pmatrix} 0 & 1 & t^{-1} & t + t^3 \\ 0 & 0 & 1 & t^2 + t^3 \\ 0 & 0 & 0 & t \\ 1 & 0 & 0 & t^2 \end{pmatrix}.$$

Notice that $A(1)$ is the transpose of the transition matrix for f found in [DKL, Example 2.9]. Calculating $\det(xI - A(t))$, we see that the fitting ideal of $T(\widehat{\Theta}_u)$ is generated by the polynomial

$$m(t, x) = x^4 - t^2x^3 - t^3x^2 - tx^2 - t^3x - t^2x - x - t \in \mathbb{Z}[H].$$

10. The isomorphism of modules

In this section we prove that there is a $\mathbb{Z}[H]$ -module isomorphism between $T(\widehat{\Theta}_u)$ and $T(\mathcal{F})$ for every primitive integral class $u \in \mathcal{S}$. We first note that by construction of Θ_u , each edge of $\widehat{\Theta}_u$ is in fact a transversal to \mathcal{F} (see Remark 7.4 and Convention 7.8). This

gives a natural (injective) $\mathbb{Z}[H]$ -module map $S: \mathbb{Z}[H]^E \rightarrow \mathbb{F}(\mathcal{F})$. We also recall from Definitions 4.2 and 9.1 that the $\mathbb{Z}[H]$ -modules $T(\widehat{\Theta}_u)$ and $T(\mathcal{F})$ are the quotients of $\mathbb{Z}[H]^E$ and $\mathbb{F}(\mathcal{F})$ by the submodules M and R , respectively. We can now state the main result of this section.

Proposition 10.1. *The $\mathbb{Z}[H]$ -module map $S: \mathbb{Z}[H]^E \rightarrow \mathbb{F}(\mathcal{F})$ descends to a $\mathbb{Z}[H]$ -module isomorphism $\bar{S}: T(\widehat{\Theta}_u) \rightarrow T(\mathcal{F})$.*

Remark 10.2. We emphasize that this isomorphism holds for every primitive integral class $u \in \mathcal{S}$. In particular, the graph modules $T(\widehat{\Theta}_u)$ are all isomorphic.

Proof of Proposition 10.1. First note that since H preserves the set of basic relations generating M , we see that elements of M are precisely \mathbb{Z} -linear combinations of elements of the form $e - e_1 - \dots - e_k$, where e is an edge of $\widehat{\Theta}_u$ and $\tilde{\psi}_1^u(e)$ is the edge path $e_1 \dots e_k$ in $\widehat{\Theta}_u$. Similarly, the elements of R are exactly the \mathbb{Z} -linear combinations of elements of the form $\tau - \tau_1 - \dots - \tau_k$ or $\tau - \tau'$, where $\tau, \tau', \tau_1, \dots, \tau_k$ are transversals to \mathcal{F} as in the basic relations defining R (more precisely, $\tau - \tau_1 - \dots - \tau_k$ is obtained from the first relation by repeated subdivision).

To see that S descends to the quotient, we need only check that $S(M)$ is contained in R . Consider a basic relation $m = e - e_1 - \dots - e_k$ in M for edges e, e_1, \dots, e_k in $\widehat{\Theta}_u$. By definition, the semiflow $\tilde{\psi}_1^u$ then maps e to the edge path $e_1 \dots e_k$. Let $\tau = S(e)$ and $\tau_i = S(e_i)$. Then τ may be subdivided as the concatenation $\tau'_1 \dots \tau'_k$, where each τ'_i flows homeomorphically onto τ_i via $\tilde{\psi}_1^u$. Thus we have

$$S(m) = \tau - \tau_1 - \dots - \tau_k = (\tau - \tau'_1 - \dots - \tau'_k) + \sum_{i=1}^k (\tau'_i - \tau_i) \in R.$$

To see that \bar{S} is surjective, let τ be any transversal, and let $[\tau]$ denote its equivalence class in $T(\mathcal{F})$. The transversal τ may be flowed forward locally homeomorphically onto an edge path $e_1 \dots e_k$ in $\widehat{\Theta}_u$ (necessarily this edge path lies in one component of $\widehat{\Theta}_u$). This identification of τ with $e_1 \dots e_k$ gives a subdivision of τ into transversals τ_1, \dots, τ_k , where each τ_i flows forward homeomorphically onto the edge e_i . Therefore

$$[\tau] = [\tau_1] + \dots + [\tau_k] = [e_1] + \dots + [e_k] = \bar{S}([e_1] + \dots + [e_k]).$$

Remark 10.3. Note that this step uses the fact that transversals have endpoints on vertex leaves of \mathcal{F} .

The only remaining difficulty is to show that \bar{S} is injective. Suppose that $w \in \mathbb{Z}[H]^E$ is such that $S(w) \in R$. It suffices to show that $w \in M$. By assumption, we may write $S(w)$ as a finite sum of basic relations in R :

$$S(w) = \sum_{\alpha \in A} r^\alpha$$

where A is a finite index set and each r^α is either a basic subdivision relation $\pm(\tau^\alpha - \tau_1^\alpha - \dots - \tau_{k_\alpha}^\alpha)$, or a basic flow relation $\pm(\tau^\alpha - \tau'^\alpha)$ in which τ^α flows forward homeomorphically onto τ'^α .

Since the sum $\sum_{\alpha} r^{\alpha}$ involves only finitely many transversals, we may find some connected component C of $\widehat{\Theta}_u$ (the image of $\widetilde{\Theta}_u$ by a sufficiently high power of $x \in H$) such that each of these transversals flows forward into C . After translating by an even higher power of x , we may furthermore assume that C is disjoint from each of these transversals (so that C is *strictly* higher than the transversals in question) and the endpoints of these transversals all flow into vertices of C .

Now, using the relations in M we can push w up into the graph C . More precisely, the assumption on $S(w)$ implies that each edge e occurring with nonzero coefficient in w lives in a component of $\widehat{\Theta}_u$ that is *below* C . Adding a relation in M , we may effectively replace e with $e_1 + \dots + e_k$, where $\widetilde{\psi}_1^u(e)$ is the edge path $e_1 \dots e_k$. Doing this repeatedly for each edge in w (and then again until all edges have been pushed up into C), we obtain an element $m \in M$ (a sum of basic relations) such that the element $w + m \in \mathbb{Z}[H]^E$ is a \mathbb{Z} -linear combination of only edges contained in C . More precisely, every edge $e \in \widehat{\Theta}_u$ appearing with a nonzero coefficient in $w + m$ lives in the component C .

We claim that $w + m = 0 \in \mathbb{Z}[H]^E$ (so that $w \in M$ as desired). We have explicitly constructed m as a sum of particular basic relations that only involve edges in the level C or below. Therefore, the above explicit proof that $S(M) \subset R$ shows us how to express $S(m)$ as a finite sum $\sum_{\alpha \in B} r^{\alpha}$ of basic relations $r^{\alpha} \in R$ that only involve transversals that are either edges of C or lie strictly below C . We may now write

$$S(w + m) = \sum_{\alpha \in A \sqcup B} r^{\alpha},$$

which is an equality in $\mathbb{F}(\mathcal{F})$. In particular, for every transversal τ to \mathcal{F} that lies strictly below C , we know that the net coefficient of τ on the right hand side must be zero.

Let τ be any transversal occurring in any of the basic relations r^{α} , $\alpha \in A \sqcup B$. By construction, τ flows homeomorphically onto some edge path $e_1^{\tau} \dots e_{k_{\tau}}^{\tau}$ in C . (It may be, if τ is in one of the new relations coming from m , that τ is already a single edge in C , but every other transversal occurring in an r^{α} lies strictly below C .) Flowing the vertices of this edge path backwards onto τ we obtain a canonical subdivision $\rho_1^{\tau} \dots \rho_{k_{\tau}}^{\tau}$ of the transversal τ . Now, for each occurrence $\pm\tau$ of τ in a basic relation r^{α} , $\alpha \in A \sqcup B$, let us add the basic relation $\delta_{\pm\tau} = \pm(\rho_1^{\tau} + \dots + \rho_{k_{\tau}}^{\tau} - \tau) \in R$. Notice that the sum of all added relations is 0 (in $\mathbb{F}(\mathcal{F})$) since each τ either lies strictly below C and thus occurs with net coefficient 0 in $S(w + m)$, or is contained in C in which case $\delta_{\tau} = \rho_1^{\tau} - \tau = 0$ by construction (i.e., such a transversal does not get subdivided). Letting Ω_{α} denote the set of (signed) transversals τ occurring in the relation r^{α} and grouping the new added relations δ_{τ} along with the old, we may now write

$$S(w + m) = \sum_{\alpha \in A \sqcup B} \left(r^{\alpha} + \sum_{\tau \in \Omega_{\alpha}} \delta_{\tau} \right).$$

Let us regroup terms and see what we have. First consider a basic subdivision relation $r^{\alpha} = \tau^{\alpha} - \tau_1^{\alpha} - \dots - \tau_{k_{\alpha}}^{\alpha}$. Let us denote the canonical subdivision of τ_i^{α} as $\beta_{i,1} \dots \beta_{i,j_i}$. Then since $\tau_1^{\alpha} \dots \tau_{k_{\alpha}}^{\alpha}$ was already a subdivision of τ^{α} , we see that the canonical subdivision of τ^{α} must be

$$\tau^{\alpha} = (\beta_{1,1} \dots \beta_{1,j_1})(\beta_{2,1} \dots \beta_{2,j_2}) \dots (\beta_{k_{\alpha},1} \dots \beta_{k_{\alpha},j_{k_{\alpha}}}).$$

Therefore we have complete cancellation in the term $r^\alpha + \sum_{\tau \in \Omega_\alpha} \delta_\tau$. That is, for each such subdivision relation r^α we have

$$r^\alpha + \sum_{\tau \in \Omega_\alpha} \delta_\tau = \tau^\alpha - \tau_1^\alpha - \dots - \tau_{k_\alpha}^\alpha - \sum_{i=1}^{k_\alpha} (\beta_{i,1} + \dots + \beta_{i,j_i} - \tau_i^\alpha) + (\beta_{1,1} + \dots + \beta_{1,j_1} + \dots + \beta_{\alpha_k,1} + \dots + \beta_{\alpha_k,j_{\alpha_k}} - \tau^\alpha) = 0.$$

Now consider a flow relation $r^\alpha = \pm(\tau - \tau')$ in which τ flows forward homeomorphically onto τ' . In this case τ and τ' both flow homeomorphically onto the same edge path in C , and so in their canonical subdivisions $\beta_1 \cdots \beta_{k_\tau}$ and $\beta'_1 \cdots \beta'_{k_{\tau'}}$ we must have $k_\tau = k_{\tau'}$ and β_i flows homeomorphically onto β'_i for each i (after possibly reindexing). Therefore, the term $r^\alpha + \sum_{\Omega_\alpha} \delta_\tau$ may be reorganized

$$r^\alpha + \sum_{\tau \in \Omega_\alpha} \delta_\tau = \pm(\tau - \tau') \pm (\beta_1 + \dots + \beta_{k_\tau} - \tau) \mp (\beta'_1 + \dots + \beta'_{k_{\tau'}} - \tau') = \pm((\beta_1 - \beta'_1) + \dots + (\beta_{k_\tau} - \beta'_{k_{\tau'}}))$$

into a sum of basic flow relations. Furthermore, each transversal appearing in this sum flows homeomorphically onto a single edge of C .

The above two paragraphs show that, in $\mathbb{F}(\mathcal{F})$, we may write $S(w + m)$ as a finite sum $\sum_{i=1}^N r_i$ of basic flow relations $r_i = \tau_i - \tau'_i$ between two transversals that both flow homeomorphically onto a single edge of C . Let Ω denote the set of all transversals appearing in $\sum_i r_i$, and partition Ω into equivalence classes in which two transversals are equivalent if they both flow homeomorphically onto the same edge of C . Notice that since the two transversals appearing in a relation $r_i = \tau_i - \tau'_i$ are in the same equivalence class, the sum of the coefficients of all transversals in a given equivalence class must be zero. On the other hand, for each transversal $\tau \in \Omega$ that is not an edge of C , the net coefficient of τ in $\sum_i r_i$ must also be zero (since $\sum_i r_i = S(w + m)$ in $\mathbb{F}(\mathcal{F})$, and the only edges of $\widehat{\Theta}_u$ with nonzero coefficient in $w + m$ lie in C). As each equivalence class contains at most one transversal $\tau \in \Omega$ which is an edge of C , these two facts imply that the net coefficient of such a τ (i.e., of a transversal which is an edge of C) must also be zero. Therefore, for each $\tau \in \Omega$, the net coefficient of τ in $\sum_i r_i$ is zero. This proves that $S(w + m) = \sum_i r_i = 0$ in $\mathbb{F}(\mathcal{F})$. On the other hand, S is injective and hence $w + m = 0$ in $\mathbb{Z}[H]^E$ as required. \square

11. Specialization, characteristic polynomials, and stretch factors

Propositions 10.1 and 9.3 together prove Proposition 4.3 and thus verify that the McMullen polynomial m is well-defined. Furthermore, Corollary 9.4 now provides an explicit way to calculate m : Given a primitive integral class $u \in \mathcal{S}$, let \mathbf{t}, x be a basis for $H = H_u \oplus \mathbb{Z}$ adapted to u as in Definition 8.1. Viewing m as an integral Laurent polynomial $m(\mathbf{t}, x) \in \mathbb{Z}[\mathbf{t}^\pm, x^{\pm 1}]$, we can now prove the “determinant formula” from the introduction.

Theorem D (Determinant formula). *For any primitive integral $u \in \mathcal{S}$ we have*

$$m(\mathbf{t}, x) = \det(xI - A_u(\mathbf{t}))$$

up to units in $\mathbb{Z}[\mathbf{t}^{\pm 1}, x^{\pm 1}]$.

Proof. Since the g.c.d. of the fitting ideal of a module depends only on the isomorphism type of the module up to units, the claim follows from Corollary 9.4 and Proposition 10.1. □

The determinant formula also tells us about characteristic polynomials of transition matrices for the expanding irreducible train track maps $f_u : \Theta_u \rightarrow \Theta_u$. The relationship can be stated in terms of the specializations of m . To describe these, first note that any cohomology class $u \in H^1(G; \mathbb{R})$ (not necessarily integral) determines a homomorphism $u : H \rightarrow \mathbb{R}$. If we write the McMullen polynomial as a finite sum

$$m = \sum_{h \in H} a_h h \in \mathbb{Z}[H],$$

then the *specialization* of m at u is the single variable power sum given by

$$m_u(\zeta) = \sum_{h \in H} a_h \zeta^{u(h)}.$$

The exponents here are real numbers, but if u is integral, this is in fact an integral Laurent polynomial in ζ .

Corollary 11.1 (Specialization). *For any primitive integral $u \in \mathcal{S}$, the specialization $m_u(\zeta)$ is equal (up to a factor $\pm \zeta^j$ for some $j \in \mathbb{Z}$) to the characteristic polynomial of the transition matrix for the train track map $f_u : \Theta_u \rightarrow \Theta_u$.*

Proof. We first recall that the kernel of the homomorphism $u : H \rightarrow \mathbb{Z}$ is exactly $H_u < H$. By definition of the basis t_1, \dots, t_{b-1}, x of $H = H_u \oplus \mathbb{Z}$, we see that $u(t_i) = 0$ for each $i = 1, \dots, b - 1$ and $u(x) = 1$. Thus by Theorem D we have

$$m_u(\zeta) = m(\zeta^0, \dots, \zeta^0, \zeta^1) = \det(\zeta^1 I - A_u(\zeta^0, \dots, \zeta^0)) = \det(\zeta I - A_u(1, \dots, 1)).$$

But $A_u(1, \dots, 1)$ is the transition matrix for f_u by Proposition 9.2. □

We may now prove Theorem C from the Introduction:

Proof of Theorem C. The existence of m follows from Proposition 4.3, and Corollary 11.1 implies that the specialization $m_u(\zeta)$ at a primitive integral class $u \in \mathcal{S}$ is, up to a factor, equal to the characteristic polynomial of f_u . Thus the largest root of $m_u(\zeta)$ is the spectral radius $\lambda(f_u)$ of the transition matrix of f_u . Since f_u is an irreducible train track map, Proposition 2.2 gives $h(f_u) = \log(\lambda(f_u))$, so that the equalities $\lambda(f_u) = \lambda((f_u)_*) = \lambda(\phi_u) = \Lambda(u)$ follow from Theorem B(5). □

Example 11.2 (Specializing at u_1 and u_2). We now use Theorem C to calculate the stretch factors associated to the classes u_1 and u_2 . Using the splitting $H = \mathbb{Z}[s] \oplus \mathbb{Z}[w]$ of H adapted to the class $u = u_0$, we find that the McMullen polynomial, as calculated in Example 9.5, is

$$m(t, x) = x^4 - t^2x^3 - t^3x^2 - tx^2 - t^3x - t^2x - x - t \in \mathbb{Z}[H].$$

Since $(s(u_1), w(u_1)) = (-1, 2)$, the specialization of m at u_1 becomes

$$\begin{aligned} m_{u_1}(\zeta) &= \zeta^8 - \zeta^{-2}\zeta^6 - \zeta^{-3}\zeta^4 - \zeta^{-1}\zeta^4 - \zeta^{-3}\zeta^2 - \zeta^{-2}\zeta^2 - \zeta^2 - \zeta^{-1} \\ &= \zeta^{-1}(\zeta^9 - \zeta^5 - \zeta^4 - \zeta^3 - \zeta^2 - \zeta - 2). \end{aligned}$$

Up to a power of ζ , this is exactly the characteristic polynomial of $A(f_1)$ for the first return map f_1 calculated in Example 5.6. Furthermore, the largest positive root of m_{u_1} is $\lambda_{u_1} \approx 1.35827$, which we note agrees with the stretch factor of the injective endomorphism ϕ_1 computed in Example 7.11.

On the other hand, as $(s(u_2), w(u_2)) = (-1, 1)$, the specialization of m at u_1 is

$$\begin{aligned} m_{u_2}(\zeta) &= \zeta^4 - \zeta^{-2}\zeta^3 - \zeta^{-3}\zeta^2 - \zeta^{-1}\zeta^2 - \zeta^{-3}\zeta - \zeta^{-2}\zeta - \zeta - \zeta^{-1} \\ &= \zeta^{-2}(\zeta^6 - 3\zeta^3 - 3\zeta - 1), \end{aligned}$$

which, again up to a power of ζ , is exactly the characteristic polynomial of $A(f_2)$ for the first return map f_2 found in Example 5.7. We also find that the largest positive root of m_{u_2} is $\lambda_{u_2} \approx 1.63299$, in agreement with the stretch factor of the monodromy $\varphi_2 = \varphi_{u_2}$ calculated in Example 7.15. Thus the data provided by m for the classes u_1 and u_2 agrees with the conclusion of Corollary 11.1.

11.1. Subdivision

Recall from §4 that the foliation \mathcal{F} of X has distinguished vertex leaves. These play an important role in the definition of the module $T(\mathcal{F})$ (see Definitions 4.1–4.2), and this is reflected in the McMullen polynomial m .

For example, suppose that $\Theta \subset X$ is a cross section equipped with any topological graph structure for which f_Θ sends vertices to vertices, and let $T(\widehat{\Theta})$ be the corresponding graph module as constructed in §9. If all vertices of Θ lie on vertex leaves of \mathcal{F} , then the proof of Proposition 10.1 still goes through verbatim and gives an isomorphism $T(\widehat{\Theta}) \rightarrow T(\mathcal{F})$. Thus any finite presentation of $T(\widehat{\Theta})$ may be used to calculate the McMullen polynomial m .

However, if $V\Theta$ contains points that do not lie on vertex leaves, then Proposition 10.1 fails (see Remarks 7.4 and 10.3) and so the fitting ideal of $T(\widehat{\Theta})$ is conceivably different from that of $T(\mathcal{F})$. Nevertheless, if we *declare* the leaves of \mathcal{F} through $V\Theta$ to be vertex leaves—effectively creating a new foliation \mathcal{F}' with a new module $T(\mathcal{F}')$ and corresponding McMullen polynomial m' —then Proposition 10.1 again gives an isomorphism $T(\widehat{\Theta}') \rightarrow T(\mathcal{F}')$ that may be used to calculate m' .

Theorem 11.3 below explains exactly how the McMullen polynomial changes in this situation. To set up notation, let Θ_u be an \mathcal{F} -compatible cross section dual to $u \in \mathcal{S}$ and equipped with the standard graph structure, and let Θ'_u be the graph structure obtained by subdividing Θ_u along a finite set $\mathcal{V} \subset \Theta_u$ that is preserved by f_u . Any such set \mathcal{V} is necessarily disjoint from the vertex leaves of \mathcal{F} , because any point contained in a vertex leaf eventually flows into a vertical 1-cell of X and thus onto a vertex of Θ_u . Let \mathcal{F}' be the foliation obtained by declaring the leaves through \mathcal{V} to be vertex leaves, and let m' be the corresponding McMullen polynomial (i.e., the g.c.d. of the fitting ideal of $T(\mathcal{F}')$).

Let $\{t, x\}$ be a basis of H adapted to u . If $\tilde{\mathcal{V}} \subset \tilde{\Theta}_u$ denotes the preimage of \mathcal{V} , then the group $H_u \leq H$ acts freely on $\tilde{\mathcal{V}}$. Thus if $\mathcal{V}_0 = \{v_1, \dots, v_k\} \subset \tilde{\mathcal{V}}$ is any set of orbit representatives, then the free \mathbb{Z} -module on $\tilde{\mathcal{V}}$ is isomorphic to the finitely generated free $\mathbb{Z}[H_u]$ -module on \mathcal{V}_0 , namely $\mathbb{Z}[H_u]^{\mathcal{V}_0} \cong \mathbb{Z}[H_u]^k$. Choosing a lift $\tilde{f}_u: \tilde{\Theta}'_u \rightarrow \tilde{\Theta}_u$ of f_u then induces a $\mathbb{Z}[H_u]$ -linear map $\mathbb{Z}[H_u]^{\mathcal{V}_0} \rightarrow \mathbb{Z}[H_u]^{\mathcal{V}_0}$, which we may represent by a matrix $B_u(\mathbf{t})$ with entries in $\mathbb{Z}[H_u]$ as follows: Choose an orientation on each edge of Θ_u and lift it to an orientation on the edges of $\tilde{\Theta}_u$. For each $v_j \in \mathcal{V}_0$, we have $\tilde{f}_u(v_j) = h \cdot v_i$ for unique $v_i \in \mathcal{V}_0$ and $h \in H_u$. Accordingly we set $B_{u,ij} = \pm h$, where the sign \pm records whether or not \tilde{f}_u preserves the orientation at v_j . All other entries in the j -column of B_u are set to 0.

We may now state the relationship between m, m' , and $B_u(\mathbf{t})$.

Theorem 11.3. *Suppose Θ'_u is obtained by subdividing Θ_u along the finite f_u -invariant set \mathcal{V} as described above. Let \mathcal{F}' denote the foliation in which the leaves through \mathcal{V} are declared to be vertex leaves, and let $m, m' \in \mathbb{Z}[H]$ be the McMullen polynomials corresponding to the modules $T(\mathcal{F})$ and $T(\mathcal{F}')$. Then with $B_u(\mathbf{t})$ as above, up to units in $\mathbb{Z}[H]$,*

$$m'(\mathbf{t}, x) = m(\mathbf{t}, x) \cdot \det(xI - B_u(\mathbf{t})).$$

Proof. By Proposition 10.1 we have $T(\mathcal{F}') \cong T(\widehat{\Theta}'_u)$, thus the McMullen polynomial m' may be calculated using any finite presentation of $T(\widehat{\Theta}'_u)$. By choosing this presentation carefully, the above formula will become obvious.

As in §9, let $\bar{E} = \{\bar{\sigma}_1, \dots, \bar{\sigma}_m\}$ be the edges of Θ_u , and let $E = \{\sigma_1, \dots, \sigma_m\}$ denote chosen lifts in $\tilde{\Theta}_u$. These choices determine the transition matrix $A_u(\mathbf{t})$ as in §9.1. Recall also that, as described above, \mathcal{V}_0 is any set of orbit representatives for the H_u action on the preimage $\tilde{\mathcal{V}} \subset \tilde{\Theta}'_u$ of \mathcal{V} , and $B_u(\mathbf{t})$ describes the $\mathbb{Z}[H_u]$ -linear action of the lift $\tilde{f}_u: \tilde{\Theta}_u \rightarrow \tilde{\Theta}_u$ on $\tilde{\mathcal{V}}$.

We now choose convenient generators for $T(\widehat{\Theta}'_u)$. As described above, fix an orientation on each edge of Θ_u and lift it to an orientation on the edges of $\tilde{\Theta}_u$. Each vertex $v_j \in \mathcal{V}_0$ is contained in a unique edge σ of $\tilde{\Theta}_u$; let α_j^- denote the initial subarc of σ that terminates at v_j . Similarly, let α_j^+ denote the terminal subarc. The arcs $\sigma_i, \alpha_j^\pm \subset \tilde{\Theta}_u$ all consist of edges of $\tilde{\Theta}'_u$ and as such may be thought of as elements of the free \mathbb{Z} -module on the edges of $\tilde{\Theta}'_u$. Moreover, since we can add and subtract H -translates of these to obtain representatives of every H -orbit, the $m + k$ elements $\{\sigma_1, \dots, \sigma_m, \alpha_1^-, \dots, \alpha_k^-\}$ freely generate this module. There is a bijection between this set and $E \cup \mathcal{V}_0$ which we use as a notational convenience.

By Definition 9.1, we now see that $T(\widehat{\Theta}'_u)$ is the quotient of $\mathbb{Z}[H]^{E \cup \mathcal{V}_0}$ by a certain submodule M of relations. Namely each arc $\beta \in E \cup \mathcal{V}_0$ must satisfy the relation

$$[\beta] = [\widetilde{\psi}_1^u(\beta)],$$

where β and $\widetilde{\psi}_1^u(\beta)$ are shorthand for the sum of the edges of $\widehat{\Theta}'_u$ comprising those edge paths. Let us express these relations in terms of our generating set. For $\sigma_j \in E$ the corresponding relator may, as in §9.1, be rewritten as $x \cdot \sigma_j - \sum_i A_{ij}(\mathbf{t}) \cdot \sigma_i$; notice that this does not involve any of the generators α_i^- .

By definition of B_u , for each $v_j \in \mathcal{V}_0$ we have $\widetilde{f}_u(v_j) = \sum_i B_{u,ij} \cdot v_i = h \cdot v$ for some $h \in H_u$ and $v \in \mathcal{V}_0$. Notice that $\widetilde{\psi}_1^u(x \cdot \alpha_j^-)$ is an edge path in $\widehat{\Theta}'_u$ that starts at a vertex of $\widetilde{\Theta}_u$, crosses several complete edges of Θ_u , and then terminates at $\widetilde{f}_u(v_j) = h \cdot v$. Therefore in $\mathbb{Z}[H]^{E \cup \mathcal{V}_0}$ this edge path may be expressed as

$$\widetilde{\psi}_1^u(x \cdot \alpha_j^-) = h_1 \cdot \sigma_1 + \dots + h_m \cdot \sigma_m + \sum_i B_{u,ij}(\mathbf{t}) \cdot \alpha_i^-$$

for some coefficients $h_i \in H$ (here we have used the fact that $\alpha_i^- + \alpha_i^+$ is a complete edge of $\widehat{\Theta}_u$). Since the above $m + k$ relators (one for each element of $E \cup \mathcal{V}_0$) generate the submodule M of all relations, it follows that the $\mathbb{Z}[H]$ -module $T(\widehat{\Theta}'_u)$ is presented by a block matrix of the form

$$\begin{pmatrix} xI - A_u(\mathbf{t}) & * \\ 0 & xI - B_u(\mathbf{t}) \end{pmatrix}.$$

Since $m(\mathbf{t}, x) = \det(xI - A_u(\mathbf{t}))$ by Theorem D, the result follows. □

Corollary 11.4. *If m and m' are McMullen polynomials as in Theorem 11.3, then for any integral class $u \in \mathcal{S}$ the specialization $m_u(\zeta)$ is a factor of $m'_u(\zeta)$.*

Proof. For any $u \in H^1(G; \mathbb{Z})$, the assignment

$$\sum_{h \in H} c_h h \mapsto \sum_{h \in H} c_h \zeta^{u(h)}$$

defines a ring homomorphism $\mathbb{Z}[H] \rightarrow \mathbb{Z}[\zeta^{\pm 1}]$. By definition of specialization, this homomorphism sends m to m_u and m' to m'_u . Therefore, since Theorem 11.3 implies that m' is a product of m with another element of $\mathbb{Z}[H]$, it follows that m'_u is the product of m_u and another element of $\mathbb{Z}[\zeta^{\pm 1}]$. □

Example 11.5. There is a fixed point in the d -edge of our running example, and we can subdivide Γ at this point to produce Γ' . Looking at Figure 9, if we choose our representative point in the edge labeled d , then the image is contained in the edge labeled t^2d . Thus $B(t) = t^2$, and we see that

$$m'(t, x) = m(t, x)(x - t^2).$$

12. Real-analyticity, convexity, and divergence

Theorem D of [DKL] showed that the logarithm of the stretch factor of φ_u , for primitive integral $u \in \mathcal{A}$, extends to a continuous, convex, and degree -1 homogeneous function $\mathfrak{h}: \mathcal{A} \rightarrow \mathbb{R}$. This theorem was an analogue of Fried's result in the setting of fibered hyperbolic 3-manifolds [Fri1]. Following McMullen's more recent approach to this 3-manifold result [McM1], we now use the McMullen polynomial m to give an alternative proof of [DKL, Theorem D] and to moreover extend this result to the entire cone of sections \mathcal{S} .

This discussion involves yet another cone $\mathcal{C}_X \subset H^1(X; \mathbb{R})$, termed the *McMullen cone*, that is naturally determined by the McMullen polynomial. We use the properties of m to show that the above function $\mathfrak{h}: \mathcal{A} \rightarrow \mathbb{R}$ naturally extends, with the same properties, to the entire cone \mathcal{C}_X and that this extension diverges at $\partial\mathcal{C}_X$. On the other hand, the relationship between m and the primitive classes $u \in \mathcal{S}$ (as illuminated by Theorem C) shows that this extension $\mathfrak{h}: \mathcal{C}_X \rightarrow \mathbb{R}$ is finite on \mathcal{S} and therefore $\mathcal{S} \subseteq \mathcal{C}_X$. Finally, we prove that \mathfrak{h} in fact diverges as we approach the boundary of \mathcal{S} by establishing the equality $\mathcal{S} = \mathcal{C}_X$, which is Theorem E. This proof combines the characterization of \mathcal{S} provided by Theorem A with another formula for m that is due to [AKHR].

12.1. McMullen's Perron–Frobenius theory

Given a primitive integral element $u \in \mathcal{S}$, let $H = H_u \oplus \mathbb{Z}$ and $\mathbf{s}, \mathbf{t} = e^s, w$, and $x = e^w$ be adapted to u as in Definition 8.1. We then express elements $u \in H^1(X; \mathbb{R})$ in terms of their (\mathbf{s}, w) -coordinates following Convention 8.2:

$$u = (\mathbf{s}(u), w(u)) = (\mathbf{s}, w).$$

Given a nonzero polynomial $\mathfrak{p}(\mathbf{t}, x) \in \mathbb{Z}[\mathbf{t}^{\pm 1}, x^{\pm 1}]$, we express this polynomial as

$$\mathfrak{p}(\mathbf{t}, x) = \sum_{\mathbf{j} \in J} a_{\mathbf{j}}(\mathbf{t}x)^{\mathbf{j}}$$

where $\mathbf{j} = (j_1, \dots, j_b)$ runs over some finite index set J , $(\mathbf{t}x)^{\mathbf{j}} = t_1^{j_1} \cdots t_{b-1}^{j_{b-1}} x^{j_b}$, and $a_{\mathbf{j}} \neq 0$ for all $\mathbf{j} \in J$. Then given $\mathbf{j} \in J$, the *dual cone of $a_{\mathbf{j}}(\mathbf{t}x)^{\mathbf{j}}$ for \mathfrak{p}* is defined by

$$\mathcal{C}(\mathfrak{p}, a_{\mathbf{j}}(\mathbf{t}x)^{\mathbf{j}}) = \{(\mathbf{s}, w) \in \mathbb{R}^b \mid \mathbf{j} \cdot (\mathbf{s}, w) > \mathbf{j}' \cdot (\mathbf{s}, w) \text{ for all } \mathbf{j}' \in J \setminus \{\mathbf{j}\}\}$$

Here, “ \cdot ” denotes the standard dot product on \mathbb{R}^b . The cone $\mathcal{C}(\mathfrak{p}, a_{\mathbf{j}}(\mathbf{t}x)^{\mathbf{j}})$ is usually referred to as the dual cone of the Newton polytope of \mathfrak{p} associated to $a_{\mathbf{j}}(\mathbf{t}x)^{\mathbf{j}}$. The dual cones for terms $a_{\mathbf{j}}(\mathbf{t}x)^{\mathbf{j}}$ and $a_{\mathbf{j}'}(\mathbf{t}x)^{\mathbf{j}'}$ of the same polynomial \mathfrak{p} are either disjoint or equal.

As (\mathbf{s}, w) are coordinates on $H^1(X; \mathbb{R})$, we can view $\mathcal{C}(\mathfrak{p}, a_{\mathbf{j}}(\mathbf{t}x)^{\mathbf{j}}) \subset H^1(X; \mathbb{R})$. This cone depends only on $\mathfrak{p} \in \mathbb{Z}[H]$ and $(\mathbf{t}x)^{\mathbf{j}} \in H$. Furthermore, if $\mathfrak{q} = \pm(\mathbf{t}x)^{\mathbf{i}}\mathfrak{p}$, then

$$\mathcal{C}(\mathfrak{q}, \pm a_{\mathbf{j}}(\mathbf{t}x)^{\mathbf{j}+\mathbf{i}}) = \mathcal{C}(\mathfrak{p}, a_{\mathbf{j}}(\mathbf{t}x)^{\mathbf{j}}) \subset H^1(X; \mathbb{R}).$$

A polynomial in $\mathbb{Z}[\mathbf{t}^{\pm 1}]$ is *nonnegative* if all coefficients are nonnegative, and *positive* if it is nonnegative and nonzero. Following [McM1], a matrix $P(\mathbf{t})$ with entries in $\mathbb{Z}[\mathbf{t}^{\pm 1}]$ is said to be *Perron–Frobenius* if each entry is nonnegative, for every i, j there exists $k > 0$ such that $(P(\mathbf{t})^k)_{ij}$ is positive, and $P(\mathbf{1})$, the matrix obtained by evaluating \mathbf{t} at $(1, \dots, 1)$, is not a permutation matrix. In [McM1, Appendix A.1], McMullen proves Theorem 12.2.

Remark 12.1. The above definition of Perron–Frobenius matrix is slightly weaker than that given in [McM1], but the proof in that paper goes through in this setting with only minor modifications.

Theorem 12.2 (McMullen [McM1]). *Let $E(\mathbf{t})$ be the leading eigenvalue of a $k \times k$ Perron–Frobenius matrix $P(\mathbf{t})$ with entries in $\mathbb{Z}[\mathbf{t}^{\pm 1}]$. Then:*

- (1) *The function $\xi(\mathbf{s}) = \log(E(e^{\mathbf{s}}))$ is a convex function of $\mathbf{s} \in \mathbb{R}^{b-1}$.*
- (2) *The graph $\mathfrak{G} = \{(\mathbf{s}, w) \mid w = \xi(\mathbf{s})\}$ meets each ray from the origin in $\mathbb{R}^{b-1} \times \mathbb{R}$ at most once.*
- (3) *The rays passing through \mathfrak{G} coincide with the dual cone $\mathcal{C}(\mathbf{p}, x^k) \subset \mathbb{R}^{b-1} \times \mathbb{R}$ for $\mathbf{p} = \det(xI - P(\mathbf{t})) \in \mathbb{Z}[\mathbf{t}^{\pm 1}, x^{\pm 1}]$.*

Note that the terms of $\det(xI - P(\mathbf{t}))$ other than x^k have a smaller exponent on x , and hence $(\mathbf{0}, 1) \in \mathbb{R}^{b-1} \times \mathbb{R}$ is contained in $\mathcal{C}(\mathbf{p}, x^k)$.

Remark 12.3. In fact the statement in [McM1] is sharper, but the above theorem will suffice for our purposes.

Lemma 12.4. *For any primitive integral class $u \in \mathcal{S}$, the transition matrix*

$$A_u(\mathbf{t}) \in M_{m \times m}(\mathbb{Z}[\mathbf{t}^{\pm 1}])$$

is a Perron–Frobenius matrix with entries in $\mathbb{Z}[\mathbf{t}^{\pm 1}] = \mathbb{Z}[H_u]$.

Proof. The entries of $A_u(\mathbf{t})$ are nonnegative Laurent polynomials in $\mathbb{Z}[\mathbf{t}^{\pm 1}]$. By Proposition 9.2, $A_u(\mathbf{1})$ is the transition matrix for f_u which is an expanding irreducible train track map by Theorem 5.11. Therefore, $A_u(\mathbf{1})$ is not a permutation matrix. Moreover, for every i, j there exists k such that $(A_u(\mathbf{1})^k)_{ij} > 0$, which implies $(A_u(\mathbf{t})^k)_{ij}$ is a positive Laurent polynomial, as required. \square

12.2. The McMullen cone

From Theorem D, we have $\mathfrak{m} = \det(xI - A_u(\mathbf{t}))$, and we let

$$\mathcal{C}_X(u) = \mathcal{C}(\det(xI - A_u(\mathbf{t})), x^m) = \mathcal{C}(\mathfrak{m}, x^m) \subset H^1(X; \mathbb{R})$$

be the cone from Theorem 12.2 which contains u . Since any two primitive integral elements $u, u' \in \mathcal{S}$ define the same polynomial $\mathfrak{m} \in \mathbb{Z}[H]$ up to units in $\mathbb{Z}[H]$, it follows that either $\mathcal{C}_X(u) = \mathcal{C}_X(u')$ or $\mathcal{C}_X(u) \cap \mathcal{C}_X(u') = \emptyset$.

Lemma 12.5. *For any two primitive integral elements $u, u' \in \mathcal{S}$, we have $\mathcal{C}_X(u) = \mathcal{C}_X(u')$. Furthermore, writing $\mathcal{C}_X = \mathcal{C}_X(u)$ for any (hence every) primitive integral $u \in \mathcal{S}$, we have $\mathcal{S} \subseteq \mathcal{C}_X$.*

Proof. Suppose $\mathcal{C}_X(u) \neq \mathcal{C}_X(u')$ for two primitive integral classes $u, u' \in \mathcal{S}$, so that as mentioned above $\mathcal{C}_X(u) \cap \mathcal{C}_X(u') = \emptyset$. Because the cones are defined by linear inequalities with integer coefficients, it follows that the line segment in $H^1(X; \mathbb{R})$ between u and u' meets the boundary of $\overline{\mathcal{C}_X(u)}$ at some rational point. This rational point is a rational multiple of some primitive integral point $u'' \in \mathcal{S}$, and hence u'' lies in the boundary of $\overline{\mathcal{C}_X(u)}$. Since the cones are open, it follows that $\mathcal{C}_X(u'') \cap \mathcal{C}_X(u) \neq \emptyset$, and so $\mathcal{C}_X(u) = \mathcal{C}_X(u'') \ni u''$. This contradicts the fact that u'' is in the boundary of $\overline{\mathcal{C}_X(u)}$. Therefore, $\mathcal{C}_X(u) = \mathcal{C}_X(u')$ for all primitive integral $u, u' \in \mathcal{S}$, and we let this common cone be denoted \mathcal{C}_X .

It follows that every primitive integral point and hence every rational line of \mathcal{S} is contained in \mathcal{C}_X . By convexity, \mathcal{C}_X then contains the convex hull of such rays, which, by Proposition 5.12, is equal to \mathcal{S} . \square

Definition 12.6. We call the cone $\mathcal{C}_X \subseteq H^1(G; \mathbb{R})$ constructed in Lemma 12.5 the *McMullen cone* associated to the folded mapping torus X .

In fact the containment $\mathcal{S} \subseteq \mathcal{C}_X$ can be promoted to an equality:

Theorem E (McMullen polynomial detects \mathcal{S}). *The McMullen cone \mathcal{C}_X is equal to the cone of sections \mathcal{S} .*

The proof of Theorem E is somewhat involved and so is postponed to the end of this section. For now we use it to prove that \mathfrak{H} extends to all of \mathcal{S} with many nice properties.

Example 12.7 (Computing the McMullen cone \mathcal{C}_X). The McMullen polynomial for our running example is, as calculated in Example 9.5, given by

$$m(t, x) = x^4 - t^2x^3 - t^3x^2 - tx^2 - t^3x - t^2x - x - t \in \mathbb{Z}[H].$$

Accordingly, we calculate that the dual cone $\mathcal{C}(m(t, x), x^4)$ is the interior of the convex hull of the rays through $(1, 2)$ and $(-1, 0)$; this is exactly equal to the cone of sections \mathcal{S} as calculated in Example 8.3, and so \mathcal{C}_X has already been illustrated in Figure 8. We note that the equality $\mathcal{C}_X = \mathcal{S}$ here is in agreement with Theorem E.

12.3. Stretch factors

In the notation we have established, the specialization of $m(\mathbf{t}, x)$ at any point $(\mathbf{s}, w) \in H^1(X; \mathbb{R})$ is given by

$$m_{(\mathbf{s}, w)}(\zeta) = m(\zeta^{\mathbf{s}}, \zeta^w).$$

We now define a function $k: \mathcal{S} \rightarrow \mathbb{R}_+ \cup \{\infty\}$ by setting

$$k(\mathbf{s}, w) = \sup\{\zeta \in \mathbb{R}_+ \mid m_{(\mathbf{s}, w)}(\zeta) = 0\}$$

for $(s, w) \in \mathcal{S}$. If there does not exist $\zeta \in \mathbb{R}_+$ such that $m_{(s,w)}(\zeta) = 0$, we interpret the above definition as $k(s, w) = \infty$. From the definition we have $k(qu) = k(u)^{1/q}$ for any $u \in \mathcal{S}$ and $q \in \mathbb{R}_+$ (with the convention that $\infty^{1/q} = \infty$).

Note that for any primitive integral element $(s, w) = u \in \mathcal{S}$, Theorem C shows that $k(s, w) = \lambda(f_u)$, the Perron–Frobenius eigenvalue of the transition matrix $A(f_u)$ of the expanding irreducible train track map $f_u: \Theta_u \rightarrow \Theta_u$. Thus for such $(s, w) \in \mathcal{S}$ we have $k(s, w) > 1$, and consequently $k > 1$ on rays through integral points in \mathcal{S} . A priori, k could be equal to infinity on rays through nonintegral points. The next proposition shows that this is in fact not the case, and $1 < k(s, w) < \infty$ for all $(s, w) \in \mathcal{S}$.

Proposition 12.8. *Let $A_u(\mathbf{t})$ be the Perron–Frobenius matrix from Lemma 12.4, and let \mathfrak{G} be the corresponding graph as in Theorem 12.2. Then:*

- (1) *For any $(s, w) \in \mathfrak{G}$ we have $m_{(s,w)}(e) = 0$.*
- (2) *\mathfrak{G} is the level set $k(s, w) = e$ in \mathcal{S} .*
- (3) *For all $(s, w) \in \mathcal{S}$, $1 < k(s, w) < \infty$ and $k: \mathcal{S} \rightarrow (1, \infty)$ is real-analytic.*

Proof. Fix $(s, w) \in \mathfrak{G}$, so that $w = \log(E(e^s))$. According to Theorem 12.2 the ray through (s, w) is contained in \mathcal{C}_X , and by Theorem E it follows that $(s, w) \in \mathcal{S}$. We claim that $k(s, w) = e$. To see this, first observe that

$$m_{(s,w)}(e) = m(e^s, e^w) = m(e^s, e^{\log(E(e^s))}) = m(e^s, E(e^s)) = 0.$$

This last equality follows from the fact that $E(e^s)$ is an eigenvalue of $A_u(e^s)$, and $m(e^s, x)$ is its characteristic polynomial. Thus (1) is verified.

By definition, we now have $e \leq k(s, w)$. Suppose that $k(s, w) = k_0 > e$ for some (s, w) . Set $\lambda_0 = \log(k_0) > 1$. We then have

$$0 = m_{(s,w)}(k_0) = m(k_0^s, k_0^w) = m(e^{\log(k_0^s)}, e^{\log(k_0^w)}) = m(e^{\lambda_0 s}, e^{\lambda_0 w}) = m_{(\lambda_0 s, \lambda_0 w)}(e).$$

But this says that $x = e^{\lambda_0 w}$ is a zero of $m(e^{\lambda_0 s}, x)$ and hence is an eigenvalue of $A_u(e^{\lambda_0 s})$. By definition $E(e^{\lambda_0 s})$ is the largest such eigenvalue, and thus

$$E(e^{\lambda_0 s}) \geq e^{\lambda_0 w}.$$

Therefore

$$\log(E(e^{\lambda_0 s})) \geq \lambda_0 w.$$

We claim that this is a contradiction. To see this, first observe that by Theorem 12.2 the ray through (s, w) intersects \mathfrak{G} in exactly one point, namely (s, w) . For all $\lambda > 1$ we must have $\lambda(s, w)$ above the graph \mathfrak{G} —that is, $\lambda w > \log(E(e^{\lambda s}))$. Since $\lambda_0 > 1$, we arrive at the desired contradiction. Thus (2) is verified.

To see that (3) holds, let $(s, w) \in \mathcal{S}$ be arbitrary. Then there exists a unique $q > 0$ such that $(s_*, w_*) = (s/q, w/q) \in \mathfrak{G}$. Then $k(s_*, w_*) = e$ and

$$k(s, w) = k(qs_*, qw_*) = e^{1/q}, \quad \text{so that } 1 < k(s, w) < \infty.$$

Since \mathfrak{G} is the graph of a real-analytic function (see the proof of Theorem 12.2 in [McM1, Appendix A]), the above formula also shows that the function $k: \mathcal{S} \rightarrow (1, \infty)$ is real-analytic. Compare this with [McM1, proof of Corollary 5.4]. □

We now prove Theorem F from the introduction:

Theorem F (Convexity of stretch factors). *There exists a real-analytic, degree -1 homogeneous function $\mathfrak{H}: \mathcal{S} \rightarrow \mathbb{R}$ such that:*

- (1) $1/\mathfrak{H}$ is positive and concave, hence \mathfrak{H} is convex.
- (2) For every primitive integral $u \in \mathcal{S} \subset H^1(X; \mathbb{R})$ with dual compatible section Θ_u , first return map f_u , and injective endomorphism $\phi_u = \overline{(f_u)_*}$ as in Theorem B, we have

$$\mathfrak{H}(u) = \log(\Lambda(u)) = \log(\lambda(\phi_u)) = \log(\lambda(f_u)) = h(f_u).$$

- (3) $\mathfrak{H}(u)$ tends to infinity as $u \rightarrow \partial\mathcal{S}$.

Proof. In view of Proposition 12.8 and Lemma 12.5, there is a real-analytic function $k: \mathcal{S} \rightarrow (1, \infty)$ such that for every $u = (\mathbf{s}, w) \in \mathcal{S}$ we have

$$k(\mathbf{s}, w) = \sup\{\zeta \in \mathbb{R}_+ \mid m_{(\mathbf{s}, w)}(\zeta) = 0\}.$$

Define $\mathfrak{H}: \mathcal{S} \rightarrow \mathbb{R}$ as $\mathfrak{H}(u) := \log k(u)$ for $u \in \mathcal{S}$. Then for any primitive integral $u = (\mathbf{s}, w) \in \mathcal{S}$, Theorem C implies that $\mathfrak{H}(u) > 0$ is equal to the topological entropy $h(f_u)$ of f_u , and $e^{\mathfrak{H}(u)}$ is equal to the spectral radius of the transition matrix of f_u and also to the stretch factor $\Lambda(u) = \lambda(\phi_u)$, proving (2). Since the function $k(u)$ satisfies $k(qu) = k(u)^{1/q}$ for any $u \in \mathcal{S}$ and $q > 0$, it follows that $\mathfrak{H}(u) = \log k(u)$ is homogeneous of degree -1 , that is, $\mathfrak{H}(qu) = \mathfrak{H}(u)/q$ for any $u \in \mathcal{S}$ and $q > 0$. By Proposition 12.8(2), we know that \mathfrak{G} is the level set of $\mathfrak{H}(u) = 1$. Therefore the concavity of $1/\mathfrak{H}$ follows from the convexity of the hypersurface \mathfrak{G} , as proved in [McM1, proof of Corollary 5.4], thus (1) follows.

All that remains is to prove (3). For this, we suppose $\{u_n\} \subset \mathcal{S}$ is a sequence in \mathcal{S} converging to $u \in \partial\mathcal{S}$ for which $\mathfrak{H}(u_n)$ is bounded. By homogeneity of \mathfrak{H} there exists such a sequence with $u_n \in \mathfrak{G}$ for all n . But since \mathfrak{G} is a closed subset of $H^1(X; \mathbb{R})$, it follows that $u \in \mathfrak{G} \cap \partial\mathcal{S}$. This is impossible since $\mathfrak{G} \subset \mathcal{C}_X$ by Theorem 12.2, $\mathcal{C}_X = \mathcal{S}$ by Theorem E, and $\mathcal{S} \cap \partial\mathcal{S} = \emptyset$ since \mathcal{S} is an open cone. \square

12.4. The McMullen cone is the cone of sections

We now return to the proof of Theorem E and show that $\mathcal{C}_X = \mathcal{S}$.

We fix attention on the original monodromy $f: \Gamma \rightarrow \Gamma$ defining X . Let $u_0 \in H^1(X; \mathbb{R})$ be the associated primitive integral class, $H = H_0 \oplus \mathbb{Z}$ with $H_0 = H_{u_0}$ the corresponding splitting, and \mathbf{s}, w and $\mathbf{t} = e^{\mathbf{s}}, x = e^w$ the bases adapted to u_0 as in Definition 8.1. Let $\bar{E} = \{\bar{\sigma}_1, \dots, \bar{\sigma}_m\}$ be representative edges in $\tilde{\Gamma} \subset \widehat{\Gamma}$ of the H -orbit of edges. Recall from §9 the construction of the matrix $A(\mathbf{t}) = A_{u_0}(\mathbf{t})$ with entries in $\mathbb{Z}[H_0] = \mathbb{Z}[\mathbf{t}^{\pm 1}]$ in terms of E so that $A(\mathbf{1}) = A(1, \dots, 1)$ is the transition matrix of $f: \Gamma \rightarrow \Gamma$ by Proposition 9.2.

Next we let \mathcal{G} denote the directed graph associated to $A(\mathbf{1})$, and \mathcal{Y} the set of all closed oriented circuits in \mathcal{G} , as in §6. Following the construction in [AKHR] (and also used by Hadari [Had1]), we now associate to the matrix $A(\mathbf{t})$ a labeling on the edges of \mathcal{G} by elements of H_0 (this labeling differs slightly from [AKHR] due to our convention on the action of H on $T(\mathcal{F}) \cong T(\widehat{\Gamma})$). First note that $A_{ij}(\mathbf{t}) \in \mathbb{Z}[\mathbf{t}^{\pm 1}]$ is the coefficient of $[\sigma_i]$ in the expression of $x \cdot [\sigma_j]$ as a $\mathbb{Z}[\mathbf{t}^{\pm 1}]$ -linear combination of the elements of E .

Convention 12.9. For the remainder of this subsection we assume that $A_{ij}(\mathbf{t})$ is written as a positive sum of elements of H , that is, a sum of monomials in $t_1^{\pm 1}, \dots, t_{b-1}^{\pm 1}$ with coefficient 1. We refer to these simply as the *terms* or *monomials* of $A_{ij}(\mathbf{t})$. We also consider the monomial 1, so $A_{ij}(\mathbf{t})$ may also have several terms which are just 1.

The terms in $A_{ij}(\mathbf{t})$ are in 1-1 correspondence with the edges of \mathcal{G} from the j^{th} vertex to the i^{th} vertex (again by appealing to Proposition 9.2), and we use any such correspondence to label those edges by the terms of $A_{ij}(\mathbf{t})$. Denote the labeled graph as $\mathcal{G}(\mathbf{t})$.

Given a circuit $y \in \mathcal{Y}$, let $|y|$ denote the length of the circuit; that is, $|y|$ is the number of vertices in y . We also observe that y gives rise to a monomial $p_y(\mathbf{t})$ which is a product of the monomials in the circuit $y \subset \mathcal{G}(\mathbf{t})$.

Next let \mathcal{Y}' denote the set of subsets of pairwise disjoint circuits in \mathcal{G} ,

$$\mathcal{Y}' = \{y = \{y_1, \dots, y_k\} \subset \mathcal{Y} \mid y_i \cap y_j = \emptyset \text{ for all } i \neq j\}.$$

Every $y = \{y_1, \dots, y_k\} \in \mathcal{Y}'$ also has a length, $|y| = |y_1| + \dots + |y_k|$, and an associated monomial $p_y(\mathbf{t}) = p_{y_1}(\mathbf{t}) \cdots p_{y_k}(\mathbf{t})$. We also record the number of circuits in y as $\#(y) = k$. For a set with one element $\{y\} \in \mathcal{Y}'$, we write $y \in \mathcal{Y}'$.

Finally, for every $n \geq 1$ we let $\mathcal{Y}'^{(n)}$ denote the subset of \mathcal{Y}' consisting of y with $|y| = n$. Thus \mathcal{Y}' is the disjoint union of $\mathcal{Y}'^{(n)}$ over all n from 1 up to m , the number of vertices of \mathcal{G} .

We now have the following very useful formula which is essentially the ‘‘cycle polynomial’’ of [AKHR], but differs slightly, again due to our conventions on the actions of H on $T(\widehat{\Gamma})$. This formula is also closely related to the ‘‘clique polynomial’’ recently studied by McMullen [McM2] in the general setting of directed graphs.

Theorem 12.10 ([AKHR]). *With the notation above we have*

$$m(\mathbf{t}, x) = \det(xI - A(\mathbf{t})) = x^m + \sum_{n=1}^m \left(\sum_{y \in \mathcal{Y}'^{(n)}} (-1)^{\#(y)} p_y(\mathbf{t}) \right) x^{m-n}. \tag{12.11}$$

For completeness, and since our definition differs slightly from that in [AKHR], we include the proof of this theorem.

Proof of Theorem 12.10. We use the definition of the determinant as a product

$$\prod_{\rho \in \mathcal{S}_m} \left(\text{sgn}(\rho) \prod_{i=1}^m (xI - A(\mathbf{t}))_{\rho(i)i} \right) \tag{12.12}$$

over all permutations ρ of the set $\{1, \dots, m\}$, where $\text{sgn}(\rho) \in \{\pm 1\}$ is the sign of the permutation. Expanding this expression out (without cancellation), we analyze each nonzero monomial $q_0(\mathbf{t}, x)$ of this expression and produce exactly one monomial in the expansion of (12.11) (also without cancellation).

The monomial $q_0(\mathbf{t}, x)$ is obtained from some permutation ρ by taking the product for $i = 1, \dots, m$ of some choice of a single term of the polynomial $(xI - A(\mathbf{t}))_{\rho(i)i}$,

for each i . To better understand $q_0(\mathbf{t}, x)$, first write ρ in its disjoint cycle notation, and to simplify the notation, assume that this representation of ρ takes the following form:

$$(m\ m - 1 \dots k_k + 1)(k_k \dots k_{k-1} + 1) \dots (k_2 \dots k_1 + 1)(k_1)(k_1 - 1) \dots (2)(1).$$

Here we are composing *right to left* instead of left to right as is classical because this better matches our other conventions. In particular $(3\ 2\ 1)$ sends 1 to 2, 2 to 3 and 3 to 1.

Next note that the only cycles in this decomposition that can contribute a power of x to $q_0(\mathbf{t}, x)$ are the 1-cycles, since these correspond precisely to entries on the diagonal. We further assume that by reindexing if necessary (and without loss of generality) these 1-cycles are the initial k_0 appearing, for some $0 \leq k_0 \leq k_1$. Thus $q_0(\mathbf{t}, x) = x^{k_0} q_1(\mathbf{t})$. To see what $q_1(\mathbf{t})$ is, we write the product of monomials defining it as a product, over each of the remaining cycles in the disjoint cycle representation of ρ , of some monomials of the associated entries of the matrix. So, for the remaining 1-cycles $(k_1), (k_1 - 1), \dots, (k_0 + 1)$ we obtain monomials

$$m_{k_1}(\mathbf{t}), \dots, m_{k_0+1}(\mathbf{t})$$

where for each $i = k_0 + 1, \dots, k_1$, the monomial $m_i(\mathbf{t})$ is a term of $-A_{ii}(\mathbf{t})$.

Every other cycle $(k_{i+1} \dots k_i + 1)$ gives us a monomial which is a product of monomials

$$\mathbf{m}_{k_i+1}(\mathbf{t}) = m_{k_i+1, k_{i+1}}(\mathbf{t})m_{k_{i+1}, k_{i+1}-1}(\mathbf{t}) \dots m_{k_i+3, k_i+2}(\mathbf{t})m_{k_i+2, k_i+1}(\mathbf{t}),$$

and each $m_{i,j}(\mathbf{t})$ is a term of $-A_{ij}(\mathbf{t})$. Then

$$q_1(\mathbf{t}) = \text{sgn}(\rho)\mathbf{m}_{k_k+1}(\mathbf{t}) \dots \mathbf{m}_{k_1+1}(\mathbf{t})m_{k_1}(\mathbf{t}) \dots m_{k_0+1}(\mathbf{t}).$$

Now we observe that, because $q_1(\mathbf{t}) \neq 0$, each of our cycles (except the first k_0) determines a circuit in $\mathcal{G}(\mathbf{t})$: specifically, the 1-cycle (j) determines a circuit y of length 1 around the loop edge of $\mathcal{G}(\mathbf{t})$ labeled by $-m_j(\mathbf{t})$. That is, the monomial associated to this y is precisely $p_y(\mathbf{t}) = -m_j(\mathbf{t})$.

On the other hand, the cycle $(k_{i+1} \dots k_i + 1)$ determines a circuit y of length $|y| = k_{i+1} - k_i$ (equal to the length of the cycle) labeled by the monomials (which as we traverse the circuit are read right to left in the following list)

$$(-m_{k_i+1, k_{i+1}}(\mathbf{t})), (-m_{k_{i+1}, k_{i+1}-1}(\mathbf{t})), \dots, (-m_{k_i+3, k_i+2}(\mathbf{t})), (-m_{k_i+2, k_i+1}(\mathbf{t})).$$

The monomial associated to this circuit is precisely

$$p_y(\mathbf{t}) = (-1)^{|y|}\mathbf{m}_{k_i+1}(\mathbf{t}).$$

Finally, note that because these are disjoint cycles, the vertices of the circuits are all distinct, hence the circuits are pairwise disjoint, and therefore their union determines an element $\mathbf{y} \in \mathcal{Y}'$. Write $\mathbf{y} = \{y_1, \dots, y_r\}$ where $r = \#(\mathbf{y})$ is the number of circuits in \mathbf{y} . The previous calculations prove that

$$q_1(\mathbf{t}) = \text{sgn}(\rho)(-1)^{|\mathbf{y}|}p_{\mathbf{y}}(\mathbf{t}).$$

Since the sign of a permutation is the product of the signs of its disjoint cycles, and the sign of a j -cycle is $(-1)^{j+1}$, we see that

$$\text{sgn}(\rho) = (-1)^{|y_1|+\dots+|y_r|+r} = (-1)^{|y|+\#\langle y \rangle}.$$

Combining this with the previous equation, and the fact that $q_0(\mathbf{t}, x) = q_1(\mathbf{t})x^{k_0}$ and $k_0 = m - |y|$, we find that our original monomial has the form

$$q_0(\mathbf{t}, x) = (-1)^{\#\langle y \rangle} p_y(\mathbf{t})x^{m-|y|}.$$

This is a term in the sum (12.11). Reversing this discussion, we can find a corresponding term in the expansion of the determinant, proving (12.11). \square

We need one more ingredient before proving Theorem E. Recall from §6 that every circuit $y \in \mathcal{Y}$ corresponds to a closed orbit $\mathcal{O}_y \subset X$ of ψ . This closed orbit defines an element of H which we denote $[\mathcal{O}_y]$.

Lemma 12.13. *For every $y \in \mathcal{Y}$ we have $[\mathcal{O}_y] = (p_y(\mathbf{t}))^{-1}x^{|y|}$.*

Here $p_y(\mathbf{t})$ is a monomial with coefficient 1 so that it and its inverse can be viewed as elements of H .

Proof of Lemma 12.13. To simplify the proof, we suppose that the closed circuit y runs around the vertices $1, 2, \dots, k - 1, k, 1$ in order, where $k = |y|$. Recall that $\bar{E} = \{\bar{\sigma}_1, \dots, \bar{\sigma}_k\}$ are the edges of Γ with σ_i an edge in the preimage of $\bar{\sigma}_i$. The closed orbit \mathcal{O}_y starts at a point $\xi \in \bar{\sigma}_1$, then runs through $\bar{\sigma}_2, \dots, \bar{\sigma}_k$ before returning to $\bar{\sigma}_1$ at the point $\psi_k(\xi) = \xi$. To see what $[\mathcal{O}_y]$ is as an element of H , we lift the path $s \mapsto \psi_s(\xi)$ for $s \in [0, k]$ to a path in \tilde{X} , and find the covering transformation taking the initial point to the terminal point. The lift is obtained by simply picking a point $\tilde{\xi} \in \tilde{X}$ in the preimage of ξ and considering the lifted flow $s \mapsto \tilde{\psi}_s(\tilde{\xi})$ for $s \in [0, k]$. Since the covering group is abelian, it does not matter which point $\tilde{\xi}$ we pick, so we assume $\tilde{\xi} \in \sigma_1$.

Now consider the monomials on the edges of the circuit y , and denote these $p_1(\mathbf{t}), \dots, p_k(\mathbf{t})$. By definition, $p_i(\mathbf{t}) \in H_0$ has the property that the edge path $x^{-1}(\tilde{\psi}_1(\sigma_i))$ contains the edge $p_i(\mathbf{t}) \cdot \sigma_{i+1} = (p_i(\mathbf{t}))^{-1}(\sigma_{i+1})$ for all i , with $i + 1$ taken modulo k (recall the convention for the action of H on the edges of $\tilde{\Gamma}$ in §9). Therefore, we see that $\tilde{\psi}_1(\sigma_i)$ contains the edge $(x^{-1}p_i(\mathbf{t})) \cdot \sigma_{i+1}$ for all i . From this, we have

$$\begin{aligned} \tilde{\psi}_k(\sigma_1) \supset x^{-1}p_1(\mathbf{t}) \cdot \tilde{\psi}_{k-1}(\sigma_2) \supset x^{-2}p_1(\mathbf{t})p_2(\mathbf{t}) \cdot \tilde{\psi}_{k-2}(\sigma_3) \supset \dots \\ \dots \supset x^{-k+1}p_1(\mathbf{t})p_2(\mathbf{t}) \cdot \dots \cdot p_{k-1}(\mathbf{t}) \cdot \tilde{\psi}_1(\sigma_k) \supset x^{-k}p_1(\mathbf{t})p_2(\mathbf{t}) \cdot \dots \cdot p_k(\mathbf{t}) \cdot \sigma_1. \end{aligned}$$

It follows that $\tilde{\psi}_k(\tilde{\xi})$ lies on this edge, and hence the covering transformation h taking $\tilde{\xi}$ to $\tilde{\psi}_k(\tilde{\xi})$ is precisely

$$h = x^k(p_y(\mathbf{t}))^{-1} = x^{|y|}(p_y(\mathbf{t}))^{-1}$$

as required. \square

Proof of Theorem E. We have already seen in Lemma 12.5 that $\mathcal{S} \subseteq \mathcal{C}_X$. We let $u \in \partial\mathcal{S}$ and prove that $u \in \partial\mathcal{C}_X$. By convexity of \mathcal{C}_X this will prove the theorem. We note that the containment of cones implies $u \in \mathcal{C}_X$, so it suffices to prove that $u \notin \mathcal{C}_X$. Let $y \in \mathcal{Y}$ be the shortest circuit (that is, with minimal $|y|$) for which $u([\mathcal{O}_y]) = 0$. Since $\mathcal{S} = \mathcal{D}$ by Theorem A, we are guaranteed that some such \mathcal{O}_y exists.

Claim 12.14. *In $m(x, \mathbf{t}) = \det(xI - A(\mathbf{t}))$, the term $p_y(\mathbf{t})x^{m-|y|}$ has a nonzero coefficient c_0 (after expanding and cancellation).*

Proof. We note that there is a term of this form after expanding (12.11) and before cancelling, so we just need to make sure it survives the cancellation. Note that after expanding, the coefficient of $p_y(\mathbf{t})x^{m-|y|}$ is just $(-1)^{\#(y)} = -1$. Thus, for there to be cancellation, there must be some $\mathbf{y} = \{y_1, \dots, y_k\} \in \mathcal{Y}'$ such that $p_y(\mathbf{t})x^{m-|y|} = p_{\mathbf{y}}(\mathbf{t})x^{m-|y|}$, but for which the coefficient after expanding is $+1$ (all coefficients are equal to ± 1 , and if they were all -1 , there could be no cancellation). But since the coefficient is $(-1)^{\#(y)} = 1$, we must have $\#(\mathbf{y}) = k > 1$.

Since $p_y(\mathbf{t}) = p_{y_1}(\mathbf{t}) \cdots p_{y_k}(\mathbf{t})$, applying Lemma 12.13 we have

$$[\mathcal{O}_{y_1}] \cdots [\mathcal{O}_{y_k}] = (p_{y_1}(\mathbf{t}))^{-1}x^{|y_1|} \cdots (p_{y_k}(\mathbf{t}))^{-1}x^{|y_k|} = (p_y(\mathbf{t}))^{-1}x^{|y|} = (p_{\mathbf{y}}(\mathbf{t}))^{-1}x^{|y|}.$$

On the other hand, $u \in \partial\mathcal{S} = \partial\mathcal{D}$ and so $u([\mathcal{O}_{y_i}]) \geq 0$ for all $i = 1, \dots, k$. The previous equation, together with the fact that u is a homomorphism, implies

$$0 = u((p_y(\mathbf{t}))^{-1}x^{|y|}) = u((p_{\mathbf{y}}(\mathbf{t}))^{-1}x^{|y|}) = u([\mathcal{O}_{y_1}]) + \cdots + u([\mathcal{O}_{y_k}]).$$

Since the terms on the right are all greater than or equal to 0, they must all be zero. This contradicts the minimality assumption on the length of y (since each y_i has length less than that of y), which completes the proof of the claim. \square

Write $p_y(\mathbf{t}) = t_1^{j_1} \cdots t_{b-1}^{j_{b-1}}x^{|y|}$. Lemma 12.13 implies that in the additive basis for H we have

$$[\mathcal{O}_y] = -j_1s_1 - j_2s_2 - \cdots - j_{b-1}s_{b-1} + |y|w.$$

Thus the condition $u([\mathcal{O}_y]) = 0$ implies

$$|y|w(u) = j_1s_1(u) + \cdots + j_{b-1}s_{b-1}(u). \tag{12.15}$$

We set $\mathbf{j} = (j_1, \dots, j_{b-1}, m - |y|)$ so that $c_0(\mathbf{t}x)^{\mathbf{j}}$ is a nonzero term of $m(\mathbf{t}, x)$ according to Claim 12.14. The ‘‘leading term’’ x^m is the one that defines the McMullen cone

$$\mathcal{C}_X = \mathcal{C}(m, x^m)$$

as described in §12.2. Writing $\mathbf{j}' = (0, \dots, 0, m)$, we have $x^m = (\mathbf{t}x)^{\mathbf{j}'}$. However, for the class u , (12.15) implies that if we write u in the coordinates $(\mathbf{s}(u), w(u))$, we have

$$\begin{aligned} \mathbf{j} \cdot (\mathbf{s}(u), w(u)) &= (j_1, \dots, j_{b-1}, m - |y|) \cdot (\mathbf{s}(u), w(u)) \\ &= j_1s_1(u) + \cdots + j_{b-1}s_{b-1}(u) + (m - |y|)w(u) \\ &= |y|w(u) + mw(u) - |y|w(u) = mw(u) = \mathbf{j}' \cdot (\mathbf{s}(u), w(u)). \end{aligned}$$

But then u cannot be in $\mathcal{C}(m, x^m)$ since this requires $\mathbf{j}' \cdot (\mathbf{s}(u), w(u)) > \mathbf{j} \cdot (\mathbf{s}(u), w(u))$. This completes the proof of Theorem E. \square

13. Closed 1-forms and foliations

Given $u \in \mathcal{S}$, let $\omega = \omega^u$ be a tame, flow-regular closed 1-form representing u . In this section, we describe the foliation Ω_u associated to ω^u . This is analogous to the foliation tangent to the kernel of a closed 1-form in the classical smooth manifold setting. This discussion quickly leads to a proof of Theorem I. After this, we specialize to the case $u \in \mathcal{A}$, where ω^u can be chosen so that it and Ω_u satisfy stronger geometric and topological conditions analogous to those satisfied in the 3-manifold setting.

For each $\xi \in X$, setting $\gamma_\xi(s) = \psi_s(\xi)$ we let

$$r_\xi(s) = \int_{\gamma_\xi([0,s])} \omega^u.$$

This can be used to define a reparameterization of ψ by $\psi_s^u(\xi) = \psi_{r_\xi^{-1}(s)}(\xi)$ so that arcs of flowlines map isometrically to \mathbb{R} by functions in ω^u . We lift this reparameterization ψ_s^u to $\tilde{\psi}_s^u$ in \tilde{X} .

Let $(\tilde{\eta}_u : \tilde{X} \rightarrow \mathbb{R}) \in p^*(\omega^u)$ be defined on all of \tilde{X} . Then as in §5.3 we have

$$\tilde{\eta}_u \circ h(\xi) = \tilde{\eta}_u(\xi) + u(h)$$

for all $\xi \in \tilde{X}$ and $h \in H$. Let $\tilde{\Omega}_{y,u}$ denote the fiber $\tilde{\eta}_u^{-1}(y)$ for $y \in \mathbb{R}$. This is naturally a graph with vertex set $\tilde{\Omega}_{y,u}^{(0)}$ equal to the intersection with the 1-skeleton $X^{(1)}$. Hence the valence of $\tilde{\Omega}_{y,u}$ is uniformly bounded. Each of the fibers $\tilde{\Omega}_{y,u}$ covers a graph $\Omega_{y,u}$ that includes into X .

We define the foliation Ω_u defined by ω^u to be the decomposition of X into the sets $\Omega_{y,u}$, which we call the leaves. We note that $\Omega_{y,u} = \Omega_{y',u}$ if and only if y and y' differ by an element of $u(H)$. As such, the leaves of Ω_u are in a one-to-one correspondence with elements of $\mathbb{R}/u(H)$. Thus, when we write $\Omega_{y,u}$, it should be understood that y is representing the coset $y + u(H)$ in $\mathbb{R}/u(H)$.

Lemma 13.1. *The reparameterized semiflow ψ^u maps leaves to leaves, $\psi_s^u(\Omega_{y,u}) = \Omega_{y+s,u}$.*

Proof. We verify that the corresponding statement holds in \tilde{X} , and then the lemma follows. By construction of the reparameterization, we know that $s \mapsto \tilde{\eta}_u(\tilde{\psi}_s^u(\xi))$ is an isometry. From this fact we deduce that if $\xi \in \tilde{\Omega}_{y,u}$ then

$$\tilde{\eta}_u(\tilde{\psi}_s^u(\xi)) = \tilde{\eta}_u(\xi) + s = y + s.$$

So, $\tilde{\psi}_s^u(\xi) \in \tilde{\Omega}_{y+s,u}$, as required. □

We note that in the special case when u is primitive integral, the leaves are compact (in fact they are sections). Indeed, in this case ω^u defines a map $\eta_u : X \rightarrow \mathbb{R}/u(H) = \mathbb{R}/\mathbb{Z} = \mathbb{S}^1$, the descent of $\tilde{\eta}_u : \tilde{X} \rightarrow \mathbb{R}$, and $\Omega_{y,u} = \Theta_{y,u}$.

Recall that $\Sigma(G) \subset (H^1(G; \mathbb{R}) \setminus \{0\})/\mathbb{R}_+$ denotes the BNS-invariant of G , which consists precisely of those $u \in H^1(G; \mathbb{R})$ for which H is finitely generated over a finitely generated submonoid of $u^{-1}([0, \infty)) < G$ (see the remarks following the Meta-Theorem). There is an alternative, more geometric description of $\Sigma(G)$ due to Bieri–Neumann–Strebel [BNS, Theorem G] which allows us to deduce the following.

Proposition 13.2. *The cone $\mathcal{S} \subset H^1(X; \mathbb{R}) = H^1(G; \mathbb{R})$ projects into $\Sigma(G)$.*

Proof. The construction of [BNS] associates to any $u \in H^1(G; \mathbb{R}) \setminus \{0\}$ a u -equivariant continuous map $\chi': \tilde{X} \rightarrow \mathbb{R}$. The authors prove that $u \in \Sigma(G)$ if and only if $\pi_1(\chi'^{-1}([0, \infty)))$ surjects onto $\pi_1(\tilde{X})$ (or if $\chi'^{-1}([0, \infty))$ is disconnected, one considers the fundamental group of the unbounded component). They further note that modifying χ' to χ'' by an equivariant, bounded homotopy does not change the property of π_1 -surjectivity. Therefore, any continuous u -equivariant map $\tilde{X} \rightarrow \mathbb{R}$ will suffice in applying this characterization.

Now suppose $u \in \mathcal{S}$ and let $(\tilde{\eta}_u: \tilde{X} \rightarrow \mathbb{R}) \in \tilde{\omega}^u$, where ω^u is a closed flow-regular 1-form representing u . According to [BNS, Theorem G], we need only prove that $\pi_1(\tilde{\eta}_u^{-1}([0, \infty)))$ surjects onto $\pi_1(\tilde{X})$. For this we take any loop $\gamma \in \tilde{X}$, and let $-s \in \mathbb{R}$ be the minimum value of $\tilde{\eta}_u$ on γ . Then $\psi_s^u(\gamma) \subset \tilde{\eta}_u^{-1}([0, \infty))$, and $\psi_s^u(\gamma)$ is clearly homotopic to γ . □

In light of Proposition 3.7 and Theorem F, it follows that \mathcal{C}_X in fact detects a full component of $\Sigma(G)$:

Theorem I (McMullen polynomial detects the component of $\Sigma(G)$). *The McMullen cone \mathcal{C}_X projects onto a full component of the BNS-invariant $\Sigma(G)$. That is, $\{[u] \mid u \in \mathcal{C}_X\}$ is a connected component of $\Sigma(G)$.*

Proof. Let $\widehat{\Sigma}(G) = \{u \in H^1(G; \mathbb{R}) \mid [u] \in \Sigma(G)\}$, and let $\widehat{\Sigma}_0(G)$ denote the component containing u_0 . Proposition 13.2 shows that $\mathcal{C}_X \subset \widehat{\Sigma}_0(G)$. Supposing the reverse inclusion fails, $\widehat{\Sigma}_0(G)$ necessarily contains a point of $\partial\mathcal{C}_X$. Since $\widehat{\Sigma}_0(G)$ is open and the faces of \mathcal{C}_X are rationally defined, this implies that $\widehat{\Sigma}_0(G)$ contains a rational point $u \in \partial\mathcal{C}_X \cap \widehat{\mathbb{Q}\Sigma}(G)$. By Theorem F we know that $\log(\Lambda)$ agrees with \mathfrak{h} and is consequently unbounded on $\mathcal{C}_X \cap \widehat{\mathbb{Q}\Sigma}(G) \cap U$ for any open neighborhood U of u . As this contradicts Proposition 3.7, we conclude $\widehat{\Sigma}_0(G) \subset \mathcal{C}_X$ as claimed. □

13.1. In the cone \mathcal{A}

For $u \in \mathcal{A}$, the closed 1-form ω^u can be chosen to enjoy better properties than for an arbitrary $u \in \mathcal{S}$ (as is true for the flow-regular maps to \mathbb{S}^1 in the integral case). Before we describe the properties we will want for these closed 1-forms, we recall the following from [DKL, Definition 5.2].

Definition 13.3 (Local model). A *local model* for ψ is a subset of X of the form

$$\mathcal{M}(K, s_0) = \bigcup_{0 \leq s \leq s_0} \psi_s(K)$$

where $s_0 > 0$ and K is a closed contractible neighborhood of a point of Γ_t , for some t . We call K the *bottom* of \mathcal{M} and $\psi_{s_0}(K)$ the *top*, and we call \mathcal{M} minus the top and bottom the *flow-interior* of \mathcal{M} . For any $0 \leq s_1 \leq s_2 \leq s_0$,

$$\bigcup_{s_1 \leq s \leq s_2} \psi_s(K)$$

is also a local model, and we call it a *local submodel* of $\mathcal{M}(K, s_0)$.

The local models were used in [DKL] to ensure that for integral classes $u \in \mathcal{A}$, the reparameterized flows ψ_s^u mapped each fiber $\eta_u^{-1}(y) = \Theta_{y,u}$ to the fiber $\Theta_{y+s,u}$, for $y, y + s \in \mathbb{S}^1$, by a homotopy equivalence [DKL, Corollary 6.20]. The key point was that we could cover $X^{(1)}$ by finitely many flow-interiors of local models so that each component of the fiber $\Theta_{y,u}$ intersected with any of the local models \mathcal{M} in this finite set is contained in a fiber of the original fibration $\eta: X \rightarrow \mathbb{S}^1$.

Definition 13.4 (Nonsingular closed 1-form). We will say that a closed 1-form ω on X is *nonsingular* if it is flow-regular and:

- (1) For any point ξ in the 1-skeleton, there is a simply connected neighborhood U of ξ containing a local model $\mathcal{M}(K, s_0)$ such that the fibers of ω_U intersected with $\mathcal{M}(K, s_0)$ have the form $\psi_s(K)$ for some $0 \leq s \leq s_0$ (and are hence contained in fibers of the original fibration $\eta: X \rightarrow \mathbb{S}^1$).
- (2) The integral over any oriented 1-cell of X is positive.

In particular, if ω is a nonsingular closed 1-form, then by the second condition the 1-cocycle z determined by integrating ω over 1-cells is positive, and hence $u = [z] \in \mathcal{A}$. We write $\omega = \omega^z = \omega^u$ in this case.

If $u \in \mathcal{A}$ is primitive integral and ω^u is constructed as in §5.4 from the flow-regular map η_u from Theorem 5.11, then [DKL, Lemma 6.16] guarantees that ω^u is nonsingular. This fact can be used to show that every element of \mathcal{A} is represented by a nonsingular closed 1-form.

Proposition 13.5. *For every class $u \in \mathcal{A}$, there is a nonsingular closed 1-form ω^u representing u .*

Proof. Suppose $u_1, \dots, u_n \in \mathcal{A}$ are integral classes such that u can be written as a positive linear combination of these elements,

$$u = \sum_{i=1}^n \zeta_i u_i.$$

Let ω^{u_i} be a nonsingular closed 1-form as above. Then u is represented by the closed 1-form

$$\omega^u = \sum_{i=1}^n \zeta_i \omega^{u_i}.$$

Since positive combinations of flow-regular maps are flow-regular, ω^u is flow-regular.

If z_i is the positive cocycle representing u_i obtained from ω_i , then the cocycle z representing u obtained from ω^u is precisely $\sum \zeta_i z_i$. Therefore, the second condition of nonsingularity is satisfied.

By choosing the covering of $X^{(1)}$ in [DKL, Lemma 6.16] by sufficiently small local models (see also [DKL, Proposition 5.6]), we can assume that the same covering works for all of u_1, \dots, u_n , and they are each contained in simply connected open sets in X . The first property of nonsingularity then also follows. □

Now for any $u \in \mathcal{A}$, we assume that u is represented by a nonsingular closed 1-form and Ω_u is the foliation defined by ω^u . We also write $\tilde{\eta}_u : \tilde{X} \rightarrow \mathbb{R}$ with fibers $\tilde{\Omega}_{y,u} \subset \tilde{X}$ covering the leaves $\Omega_{y,u}$.

Theorem 13.6. *Suppose ω^u is a nonsingular closed 1-form on X defining $u \in \mathcal{A}$. Then for every $y \in \mathbb{R}$, the inclusion of the fibers $\tilde{\Omega}_{y,u} = \tilde{\eta}_u^{-1}(y) \rightarrow \tilde{X}$ is a homotopy equivalence (in particular, the fibers are connected). Furthermore, for every $s \geq 0$, the restriction of the reparameterized flow*

$$\tilde{\psi}_s^u : \tilde{\Omega}_{y,u} \rightarrow \tilde{\Omega}_{y+s,u}$$

is a homotopy equivalence.

Before we give the proof, we observe the following corollary. Here $\Sigma_s(G) = \Sigma(G) \cap -\Sigma(G)$ is the *symmetrized* BNS-invariant.

Corollary 13.7. *The cone \mathcal{A} projects into $\Sigma_s(G)$.*

Proof. For $u \in \mathcal{A}$, the inclusion $\tilde{\Omega}_{0,u} = \tilde{\eta}_u^{-1}(\{0\}) \rightarrow \tilde{X}$ is a homotopy equivalence by Theorem 13.6. Therefore, $\tilde{\eta}_u^{-1}([0, \infty))$ and $\tilde{\eta}_u^{-1}((-\infty, 0])$ are π_1 -surjective. As in the proof of Proposition 13.2, we can appeal to [BNS, Theorem G] to conclude that both u and $-u$ are in $\Sigma(G)$ and hence $u \in \Sigma_s(G)$. \square

Remark 13.8. For a fibered hyperbolic 3-manifold, every cohomology class in the cone on the fibered face of a hyperbolic 3-manifold is represented by a nowhere vanishing closed 1-form that evaluates positively on the vector field defining the suspension flow (of any fiber in the cone on the face). In particular, the kernel of the 1-form is tangent to a taut foliation, and so by the Novikov–Rosenberg Theorem the leaves π_1 -inject into the 3-manifold (see e.g. [Cal]). Furthermore, after reparameterizing the flow (so that the vector field evaluates to 1 on the 1-form) the leaves map homeomorphically from one onto another.

Proof. We first note that when u is a scalar multiple of a primitive integral $u' \in \mathcal{A}$, then $\tilde{\Omega}_{y,u}$ is a component of the preimage in \tilde{X} of the graph $\Theta_{y',u'}$ (though the vertex set for $\tilde{\Omega}_{y,u}$ might be a proper subset of those coming from the standard graph structure of $\Theta_{y',u'}$ if $\Theta_{y',u'}$ is \mathcal{F} -compatible). In this case, the conclusion follows from [DKL, Corollaries 6.20 and 6.21]. We therefore assume that u is not a scalar multiple of an integral element. We will prove the last statement, and then deduce the rest from this.

We assume that all local models considered in what follows satisfy the condition (1) of nonsingularity for ω^u . We choose a finite cover \mathbb{M} of $X^{(1)}$ by flow-interiors of such local models satisfying the following conditions for some $\ell > 0$:

- (i) Any $\mathcal{M} \in \mathbb{M}$ intersects $X^{(1)}$ in a contractible subset containing at most one vertex.
- (ii) Given $\mathcal{M}(K, s_0), \mathcal{M}(K', s'_0) \in \mathbb{M}$, if $\psi_s(K) \cap \psi_{s'}(K') \neq \emptyset$ for some $s \in [0, s_0]$ and $s' \in [0, s'_0]$, then $\psi_s(K) \cap \psi_{s'}(K) \cap X^{(1)} \neq \emptyset$ is a single point.
- (iii) Any arc of a flowline $\gamma_\xi(s) = \psi_s(\xi)$ that nontrivially intersects $X^{(1)}$, and for which the integral of ω^u is at most ℓ , is contained in a local model.

Appealing to [DKL, Proposition 5.6] we can find a finite cover of $X^{(1)}$ by flow-interiors of local models so that the first and second conditions are satisfied. This is done by first choosing sufficiently small, pairwise disjoint flow-interiors of local models, $\mathcal{M}_1, \dots, \mathcal{M}_k$, each containing one 0-cell of X , and intersecting $X^{(1)}$ in a contractible set. Then we appropriately choose very tiny flow-interiors of local models about each point $\xi \in X^{(1)} \setminus (\mathcal{M}_1 \cup \dots \cup \mathcal{M}_k)$ which are also all disjoint from the 0-cells. These are chosen so that they intersect $X^{(1)}$ in an arc (necessarily contained in a single 1-cell), and are so small that any two which intersect different 1-cells of $X^{(1)}$ are disjoint. By compactness, we can take a finite subcover \mathbb{M} , which we do, noting that the flow-interiors covering the 0-cells chosen first are necessarily in this set. The existence of $\ell > 0$ such that the third condition is satisfied is then also a consequence of compactness of $X^{(1)}$.

Now we lift each local model in \mathbb{M} to a local model in \tilde{X} via the covering map $p: \tilde{X} \rightarrow X$ (which is possible since each $\mathcal{M} \in \mathbb{M}$ is contained in a simply connected open subset of X), and denote the collection by $\tilde{\mathbb{M}}$. Observe that conditions (i)–(iii) are also satisfied by $\tilde{\mathbb{M}}$ with the same $\ell > 0$ (though this is no longer a finite cover).

Fix any $y \in \mathbb{R}$, look at the fiber $\tilde{\Omega}_{y,u} = \tilde{\eta}_u^{-1}(y)$, and consider the set

$$\tilde{\Omega}_{[y,y+\ell],u} = \bigcup_{0 \leq s \leq \ell} \psi_s^u(\tilde{\Omega}_{y,u}) = \tilde{\eta}_u^{-1}([y, y + \ell]).$$

Suppose that the arc $\gamma_\xi([0, \ell])$ of the flowline through $\xi \in \tilde{\Omega}_{y,u}$ nontrivially intersects $\tilde{X}^{(1)}$. Then this arc is contained in a local model $\tilde{\mathcal{M}} \in \tilde{\mathbb{M}}$ by (iii). The intersection $\tilde{\Omega}_{[y,y+\ell],u} \cap \tilde{\mathcal{M}}$ is a local submodel $\tilde{\mathcal{M}}' \subset \tilde{\mathcal{M}}$ (most likely not a local model in $\tilde{\mathbb{M}}$).

Note that $\tilde{\mathcal{M}}' \cap \tilde{X}^{(1)}$ is a component of $\tilde{\Omega}_{[y,y+\ell],u} \cap \tilde{X}^{(1)}$ by (i). We enumerate the components of $\tilde{\Omega}_{[y,y+\ell],u} \cap \tilde{X}^{(1)}$ and choose $\tilde{\mathcal{M}}_i \in \tilde{\mathbb{M}}$ containing the i^{th} component of $\tilde{\Omega}_{[y,y+\ell],u} \cap \tilde{X}^{(1)}$ and let

$$\tilde{\mathcal{M}}'_i = \tilde{\mathcal{M}}_i \cap \tilde{\Omega}_{[y,y+\ell],u}.$$

By (ii), $\tilde{\mathcal{M}}'_i \cap \tilde{\mathcal{M}}'_j \neq \emptyset$ if and only if $\tilde{\mathcal{M}}_i \cap \tilde{\mathcal{M}}_j \cap \tilde{X}^{(1)} \neq \emptyset$. Since $\tilde{\mathcal{M}}'_i \cap \tilde{X}^{(1)}$ is the i^{th} component, it follows that $\tilde{\mathcal{M}}'_i \cap \tilde{X}^{(1)} = \tilde{\mathcal{M}}'_j \cap \tilde{X}^{(1)}$ and hence $i = j$.

For any $0 \leq s \leq \ell$, we can subdivide $\tilde{\Omega}_{y,u}$ and $\tilde{\Omega}_{y+s,u}$ so that the intersections of these with any $\tilde{\mathcal{M}}'_i$ are subgraphs $\tilde{\Omega}_{y,u,i} \subset \tilde{\Omega}_{y,u}$ and $\tilde{\Omega}_{y+s,u,i} \subset \tilde{\Omega}_{y+s,u}$, respectively. The restriction of the reparameterized flow $\tilde{\psi}_s^u: \tilde{\Omega}_{y,u,i} \rightarrow \tilde{\Omega}_{y+s,u,i}$ is controlled rel boundary in the sense of [DKL, Definition 5.3]. In any complementary component, $\tilde{\psi}_s^u$ restricts to a homeomorphism (since this is contained in the interior of a 2-cell where the semiflow restricts to an honest local flow). We can now apply [DKL, Proposition 5.5] to guarantee that $\tilde{\psi}_s^u$ restricts to a homotopy equivalence from $\tilde{\Omega}_{y,u}$ to $\tilde{\Omega}_{y+s,u}$.

Remark 13.9. Proposition 5.5 of [DKL] was only stated for finite graphs because that was the only case of interest. The proof is valid in the current setting as well.

Thus we have proved the last part of the theorem for any $y \in \mathbb{R}$ and any $0 \leq s \leq \ell$. Since $\tilde{\psi}_s^u$ is a semiflow, it also holds for any $s \geq 0$.

Next we claim that every fiber is connected. To see this, observe that since the semiflow restricts to a homotopy equivalence of fibers, it maps distinct components to dis-

tinct components. Thus flowing forward and backward (i.e. taking the preimage under the semiflow) for any two components of a fiber defines disjoint open sets. Thus we have a bijection between the components of a fiber and the components of \tilde{X} ; but \tilde{X} is connected, and hence so is every fiber.

Finally, to see that the inclusion $\tilde{\Omega}_{y,u} \rightarrow \tilde{X}$ is a homotopy equivalence, we note that both spaces are Eilenberg–MacLane spaces for (possibly infinitely generated) free groups, so we need only prove that the map is an isomorphism on fundamental groups. Any loop in \tilde{X} based at a point ξ_0 of $\tilde{\Omega}_{0,u}$ is contained in the region between two fibers $\tilde{\Omega}_{y,u}$ and $\tilde{\Omega}_{y',u}$ for some $y < 0 < y'$. This loop flows forward onto a loop in $\tilde{\Omega}_{y',u}$, but since the flow restricts to homotopy equivalences between fibers, the loop is in the image of $\pi_1(\Omega_{0,u}, \xi_0)$. Similarly, for any loop in $0, u$ based at ξ_0 , suppose there is a null-homotopy of this loop to a point. This null-homotopy can be flowed forward to some fiber $\tilde{\Omega}_{y,u}$, and since the flow restricts to a homotopy equivalence $\tilde{\Omega}_{0,u} \rightarrow \tilde{\Omega}_{y,u}$, the loop must have already been null-homotopic in $\tilde{\Omega}_{0,u}$. This completes the proof. \square

Theorem 13.6 quickly implies our theorem from the introduction.

Theorem G (π_1 -injective foliations). *Given $u \in \mathcal{A}$, there exists a flow-regular closed 1-form ω^u representing u with associated foliation Ω_u of X having the following property. There is a reparameterization of ψ , denoted ψ^u , so that for each $y \in \mathbb{R}$ the inclusion of the fiber $\Omega_{y,u} \rightarrow X$ is π_1 -injective and induces an isomorphism $\pi_1(\Omega_{y,u}) \cong \ker(u)$. Furthermore, for every $s \geq 0$ the restriction*

$$\psi_s^u : \Omega_{y,u} \rightarrow \Omega_{y+s,u}$$

of ψ_s^u to any leaf $\Omega_{y,u}$ is a homotopy equivalence.

Proof. Let ω^u be the nonsingular closed 1-form from Theorem 13.6.

We observe that the image of $\ker(u)$ in H , which we denote $\ker_H(u)$, preserves the fiber $\tilde{\Omega}_{y,u} = \tilde{\eta}_u^{-1}(y)$ in \tilde{X} via the action of H on \tilde{X} . Indeed, $\ker_H(u)$ is precisely the stabilizer in H of $\tilde{\Omega}_{y,u}$. Since the inclusion of $\tilde{\Omega}_{y,u} \rightarrow \tilde{X}$ is a homotopy equivalence by Theorem 13.6, so is $\tilde{\Omega}_{y,u}/\ker_H(u) \rightarrow \tilde{X}/\ker_H(u)$. The further quotient $\tilde{X}/\ker_H(u) \rightarrow X$ maps $\tilde{\Omega}_{y,u}/\ker_H(u)$ homeomorphically onto $\Omega_{y,u}$ since it is the quotient by the action of $H/\ker_H(u) = H/\text{stab}_H(\tilde{\Omega}_{y,u})$. Since $\tilde{X}/\ker_H(u) \rightarrow X$ is the cover corresponding to $\ker(u)$, the inclusion $\Omega_{y,u}$ induces an isomorphism $\pi_1(\Omega_{y,u}) \cong \pi_1(\tilde{X}/\ker_H(u)) = \ker(u) < G$, as required.

By Theorem 13.6 $\tilde{\psi}_s^u : \tilde{\Omega}_{y,u} \rightarrow \tilde{\Omega}_{y+s,u}$ is a homotopy equivalence for all $y, s \in \mathbb{R}$, $s \geq 0$. Since $\tilde{X} \rightarrow X$ is a covering map (hence restricting to a covering map on all fibers of $\tilde{\eta}$) and since this covering map is semiflow equivariant, it follows that $\psi_s^u : \Omega_{y,u} \rightarrow \Omega_{y+s,u}$ is also a homotopy equivalence. \square

14. Lipschitz flows

Associated to every primitive integral class $u \in \mathcal{S}$, there is a rich geometric structure on X described by the following.

Proposition 14.1. *For every primitive integral $u \in \mathcal{S}$, there exists a geodesic metric d_u on X such that for the reparameterized semiflow ψ^u we have:*

- (1) *The semiflow-line $s \mapsto \psi_s^u(\xi)$ is a local geodesic for all $\xi \in X$.*
- (2) *The metric d_u restricts to a path metric on each fiber $\Theta_{y,u}$ and the maps*

$$\{\psi_s^u : \Theta_{y,u} \rightarrow \Theta_{y+s,u}\}_{s \geq 0}$$

are λ^s -homotheties on the interior of every edge, where $\lambda = \lambda(f_u)$ is the stretch factor of the first return map f_u .

The metric can be built explicitly by taking the canonical eigenmetric on Θ_u , defining a metric on the fibers such that (2) holds, and then extending to a metric on X so that the semiflow-lines from (1) have unit speed and are “perpendicular” to the fibers (in the interior of each 2-cell, this makes sense, and then the 2-cells are glued together by isometries).

The main result of this section is that one can carry out an analogous construction for any cohomology class $u \in \mathcal{S}$. This mirrors the *Teichmüller flow* from Theorem 1.1 of [McM1] for fibered hyperbolic 3-manifolds. There the flow mapped leaves of the foliation tangent to the kernel of a 1-form to leaves and was a Teichmüller mapping. In our construction, all maps are Lipschitz with constant stretch factor, and so we call these *Lipschitz flows*.

Theorem H (Lipschitz flows). *For every $u \in \mathcal{S}$, let $\mathfrak{F}(u)$ be as in Theorem F, ω^u a tame flow-regular closed 1-form representing u , ψ^u the associated reparameterization of ψ , and Ω_u the foliation defined by ω^u . Then there is a geodesic metric d_u on X such that:*

- (1) *The semiflow-lines $s \mapsto \psi_s^u(\xi)$ are local geodesics for all $\xi \in X$.*
- (2) *The metric d_u determines a path metric on each (component of a) leaf of $\Omega_{y,u}$ of the foliation Ω_u defined by ω^u making it into a (not necessarily finite) simplicial metric graph.*
- (3) *The restrictions of the reparameterized semiflow*

$$\{\psi_s^u : \Omega_{y,u} \rightarrow \Omega_{y+s,u}\}_{s \geq 0}$$

are λ^s -homotheties on the interior of every edge, where $\lambda = e^{\mathfrak{F}(u)}$.

- (4) *The induced path metric from d_u on the interior of any 2-cell U is locally isometric to a constant negative curvature Riemannian metric.*

In the special case where $u \in \mathcal{S}$ is primitive integral, the 1-form ω^u is given by restricting the fibration $\eta_u : X \rightarrow \mathbb{S}^1$ to simply connected open sets and lifting the resulting maps to the universal covering $\mathbb{R} \rightarrow \mathbb{S}^1$. In particular, the leaves $\Omega_{y,u}$ are precisely the fibers $\Theta_{y,u}$ and the metric d_u satisfies the conclusion of Proposition 14.1. Thus, Proposition 14.1 will follow from Theorem H. We recall that the label y for the leaves $\Omega_{y,u}$ makes sense as an element of $\mathbb{R}/u(H)$.

14.1. Twisted transverse measures on \mathcal{F} .

For the remainder of this section we fix a primitive integral class $u_0 \in \mathcal{A}$, which we may as well assume is the class associated to our original fibration $\eta: X \rightarrow \mathbb{S}^1$ with $\Theta_{u_0} = \Gamma$ and the first return map $f_{u_0} = f: \Gamma \rightarrow \Gamma$. We write $H_0 = H_{u_0} < H$, and let $\mathbf{s} = (s_1, \dots, s_{b-1}), (t_1, \dots, t_{b-1}) = \mathbf{t} = e^{\mathbf{s}} = (e^{s_1}, \dots, e^{s_{b-1}}), w, x = e^w$ be adapted to u_0 as in Definition 8.1.

Any class $u \in H^1(X; \mathbb{R}) = \text{Hom}(G, \mathbb{R})$ can be used to turn the additive group \mathbb{R} into a module over $\mathbb{Z}[H]$ as follows. Since $u: G \rightarrow \mathbb{R}$ factors through H , we can view the exponential of u as a homomorphism $\rho = \rho_u = e^u: H \rightarrow \mathbb{R}_+$ to the multiplicative group of positive real numbers. This defines an action of H on \mathbb{R} by $h \cdot y = \rho(h)y$ for $h \in H$ and $y \in \mathbb{R}$. Let \mathbb{R}_u denote \mathbb{R} with this module structure. Given $\mathbf{j} = (j_1, \dots, j_{b-1})$, if we write $\mathbf{t}^{\mathbf{j}} = t_1^{j_1} \dots t_{b-1}^{j_{b-1}} \in H$, then we can think of this as a function $H^1(X; \mathbb{R}) \rightarrow \mathbb{R}_+$, and

$$\rho(\mathbf{t}^{\mathbf{j}}) = \mathbf{t}(u)^{\mathbf{j}} = t_1(u)^{j_1} \dots t_{b-1}(u)^{j_{b-1}}.$$

Similarly, $\rho(x^j) = x(u)^j$.

Remark 14.2. To clarify the notation we note that $\mathbf{t}(u) \in \mathbb{R}_+^n$, while for $\mathbf{j} = (j_1, \dots, j_{b-1})$ we have $\mathbf{t}(u)^{\mathbf{j}} \in \mathbb{R}_+$.

The construction of the metric d_u uses the following proposition. Recall that $\mathfrak{G} \subset \mathbb{C}_X \subset H^1(X; \mathbb{R}) \cong \mathbb{R}^b$ is the graph of $\mathbf{s} \mapsto \log(E(e^{\mathbf{s}}))$ as in Theorem 12.2.

Proposition 14.3. For any $u \in \mathfrak{G}$, there exists a homomorphism of $\mathbb{Z}[H]$ -modules

$$\mu = \mu_u: T(\mathcal{F}) \rightarrow \mathbb{R}_u$$

such that for every transversal τ we have $\mu([\tau]) > 0$. Moreover, μ is unique up to scaling.

The condition that μ is a $\mathbb{Z}[H]$ -module homomorphism means that it is a group homomorphism, and for every $h \in H$ and transversal τ we have

$$\mu(h \cdot [\tau]) = \rho(h)\mu([\tau]).$$

Proof of Proposition 14.3. Recall that $u_0 \in \mathcal{A}$ is our initial primitive integral class coming from the graph $\Gamma = \Theta_{u_0}$. As in Section 9, we let $E = \{\sigma_1, \dots, \sigma_m\}$ be the representative edges in $\tilde{\Gamma} \subset \widehat{\Gamma}$ and $A(\mathbf{t}) = A_{u_0}(\mathbf{t})$ be the associated matrix for the action $\mathbb{Z}[H_0]^E \rightarrow \mathbb{Z}[H_0]^E$ for the lift $x^{-1} \circ \tilde{\psi}_1$ of $f: \Gamma \rightarrow \Gamma$.

We will construct a homomorphism of $\mathbb{Z}[H]$ -modules

$$\tilde{\mu}: \mathbb{Z}[H]^E \rightarrow \mathbb{R}_u$$

so that for every edge e of $\widehat{\Gamma}$ we have $\tilde{\mu}(e) > 0$ and if $\tilde{\psi}_1(e) = e_1 \dots e_j$, then

$$\tilde{\mu}(e) = \tilde{\mu}(e_1) + \dots + \tilde{\mu}(e_j). \tag{14.4}$$

By definition of $T(\widehat{\Gamma})$, such a homomorphism will descend to a homomorphism of $\mathbb{Z}[H]$ -modules

$$\mu: T(\widehat{\Gamma}) \cong T(\mathcal{F}) \rightarrow \mathbb{R}_u.$$

Since every transversal τ flows forward into an edge path of $\widehat{\Gamma}$, hence is equivalent in $T(\mathcal{F})$ to a positive linear combination of edges of $\widehat{\Gamma}$, we will have $\mu([\tau]) > 0$ as required.

Now, since $u \in \mathfrak{G}$, we have $w(u) = \log(E(e^{s(u)}))$ or equivalently

$$x(u) = E(\mathbf{t}(u)).$$

Let $\mathbf{U}(\mathbf{t})$ be the (left) Perron–Frobenius eigenvector associated to $E(\mathbf{t})$. Thus, we have

$$\mathbf{U}(\mathbf{t})A(\mathbf{t}) = E(\mathbf{t})\mathbf{U}(\mathbf{t}).$$

Given $\ell \in \{1, \dots, m\}$, define

$$\bar{\mu}(\sigma_\ell) = \mathbf{U}_\ell(\mathbf{t}(u)).$$

For an arbitrary element $\mathbf{t}^j x^j$ with $\mathbf{t}^j = t_1^{j_1} \cdots t_{b-1}^{j_{b-1}} x^j$ and $\ell \in \{1, \dots, m\}$ we define

$$\bar{\mu}(\mathbf{t}^j x^j \cdot \sigma_\ell) = \rho(\mathbf{t}^j x^j) \bar{\mu}(\sigma_\ell) = \mathbf{t}(u)^j x(u)^j \mathbf{U}_\ell(\mathbf{t}(u)).$$

Since $\mathbb{Z}[H]^E$ is a free abelian group, there is a unique extension to a group homomorphism. By construction, this is a $\mathbb{Z}[H]$ -module homomorphism:

$$\begin{aligned} \bar{\mu}(\mathbf{t}^{j'} x^{j'} \cdot (\mathbf{t}^j x^j \cdot \sigma_\ell)) &= \bar{\mu}(\mathbf{t}^{j'+j} x^{j'+j} \cdot \sigma_\ell) = \mathbf{t}(u)^{j'+j} x(u)^{j'+j} \bar{\mu}(\sigma_\ell) \\ &= \mathbf{t}(u)^{j'} x(u)^{j'} (\mathbf{t}(u)^j x(u)^j \bar{\mu}(\sigma_\ell)) = \rho(\mathbf{t}^{j'} x^{j'}) \bar{\mu}(\mathbf{t}^j x^j \cdot \sigma_\ell). \end{aligned}$$

Finally, to prove (14.4), we note that an arbitrary edge e has the form $\mathbf{t}^j x^j \cdot \sigma_\ell$ and so by definition of $A(\mathbf{t})$ and the fact that $\bar{\mu}$ is a homomorphism, we are left to prove that for every $\ell = 1, \dots, m$ we have

$$\bar{\mu}(x \cdot \sigma_\ell) = \bar{\mu}\left(\sum_{i=1}^m A_{i,\ell}(\mathbf{t}) \cdot \sigma_i\right).$$

For this, we note that since $x(u) = E(\mathbf{t}(u))$, we have

$$\begin{aligned} \bar{\mu}(x \cdot \sigma_\ell) &= x(u) \mathbf{U}_\ell(\mathbf{t}(u)) = E(\mathbf{t}(u)) \mathbf{U}_\ell(\mathbf{t}(u)) = \sum_{i=1}^m A_{i,\ell}(\mathbf{t}(u)) \mathbf{U}_i(\mathbf{t}(u)) \\ &= \sum_{i=1}^m A_{i,\ell}(\rho(t_1), \dots, \rho(t_{b-1})) \bar{\mu}(\sigma_i) = \bar{\mu}\left(\sum_{i=1}^m A_{i,\ell}(\mathbf{t}) \cdot \sigma_i\right) \end{aligned}$$

as required. Thus, $\bar{\mu}$ descends to $\mu: T(\widehat{\Gamma}) = T(\mathcal{F}) \rightarrow \mathbb{R}_u$.

To prove the uniqueness, suppose μ' is some other homomorphism with $\mu'([\tau]) > 0$ for all transversals τ . We note that any homomorphism is determined by its values on $[\sigma_1], \dots, [\sigma_m]$, and so we need only verify that there is some $\lambda > 0$ such that

$$\mu'([\sigma_\ell]) = \lambda \mu([\sigma_\ell])$$

for all $\ell = 1, \dots, m$. For this, we write $\sigma = ([\sigma_1], \dots, [\sigma_m])$ as a row vector, and write $\mu'(\sigma) = (\mu'([\sigma_1]), \dots, \mu'([\sigma_m]))$ and $x \cdot \sigma = (x \cdot [\sigma_1], \dots, x \cdot [\sigma_m])$. With this convention, the relations can be expressed in matrix form as

$$x \cdot \sigma = \sigma \cdot A(\mathbf{t}).$$

Then we have

$$\begin{aligned} \mu'(\sigma) &= \mu'(x^{-1}x \cdot \sigma) = x(u)^{-1}\mu'(x \cdot \sigma) \\ &= E(\mathbf{t}(u))^{-1}\mu'(\sigma \cdot A(\mathbf{t})) = E(\mathbf{t}(u))^{-1}\mu'(\sigma)A(\mathbf{t}(u)). \end{aligned}$$

This shows that $\mu'(\sigma)$ is an eigenvector for the Perron–Frobenius eigenvalue $E(\mathbf{t}(u))$. Since $\mu(\sigma) = \mathbf{U}(\mathbf{t}(u))$ is the unique eigenvector for $E(\mathbf{t}(u))$ up to scalar multiples, it follows that $\mu'(\sigma)$ is a scalar multiple of $\mu(\sigma)$, and we are done. \square

We think of μ as defining a “transverse measure” on \mathcal{F} which is twisted by the action of H and the homomorphism ρ_u . It is only defined on transversals, and not arbitrary arcs in 2-cells transverse to \mathcal{F} . This transverse measure can be extended to any arc, as we explain in Remark 14.9.

Before we proceed, we note that μ satisfies a monotonicity condition.

Lemma 14.5. *For any pair of transversals τ, τ' , if $\tau \subseteq \tau'$ then*

$$\mu([\tau]) \leq \mu([\tau']),$$

with equality if and only if $\tau = \tau'$.

If $\tau_1, \dots, \tau_n \subset \tau$ are subtransversals of the transversal τ , pairwise intersecting at most in their endpoints, then

$$\mu([\tau_1]) + \dots + \mu([\tau_n]) \leq \mu([\tau]).$$

Equality holds if and only if $\tau = \tau_1 \cup \dots \cup \tau_n$.

Proof. We can subdivide τ' as $\tau' = \tau_1\tau_2\tau_3$, where $\tau_2 = \tau$, and where either or both of τ_1, τ_3 may be empty. Then

$$\mu(\tau') = \mu(\tau_1) + \mu(\tau_2) + \mu(\tau_3) \geq \mu(\tau_2) = \mu(\tau).$$

Equality holds if and only if $\mu(\tau_1) = \mu(\tau_3) = 0$ and hence if and only if τ_1 and τ_3 are empty, which is to say $\tau = \tau_1 = \tau'$.

The last statement holds by a similar reasoning via subdividing τ into a concatenation of subtransversals consisting of τ_1, \dots, τ_n , together with a (possibly empty) collection of other subtransversals. \square

14.2. Constructing the metric

Recall that the foliation Ω_u is the descent to X of the foliation by fibers $\tilde{\Omega}_u$ of \tilde{X} . We want to define a metric on fibers $\tilde{\Omega}_{y,u}$ using μ and then push it down to a metric on the leaves of Ω_u , but there is a technical difficulty we must address first. To begin, we just define a notion of length for certain paths.

Given a transversal τ contained in some $\tilde{\Omega}_{y,u}$, we define the μ -length of τ to be

$$\ell_\mu(\tau) = e^y \mu([\tau]).$$

We can similarly define the μ -length of any path $\tau' = \tilde{\psi}_s^u(\tau) \subset \tilde{\Omega}_{s+y,u}$, where τ is a transversal in $\tilde{\Omega}_{y,u}$, by

$$\ell_\mu(\tau') = e^{y+s} \mu([\tau]).$$

We note that while τ is really a subset of \tilde{X} , we need to think of τ' as a parameterized path (though the parameterization is unimportant) since the restriction of $\tilde{\psi}_s^u$ to τ need not be injective. It is straightforward to check that this length is independent of the transversal τ that flows into τ' . We will call such a path τ' a *flowed transversal*.

Lemma 14.6. *For any $y, s \in \mathbb{R}$ and $s \geq 0$ and any flowed transversal τ' we have*

$$\ell_\mu(\tilde{\psi}_s^u(\tau')) = e^s \ell_\mu(\tau').$$

Proof. Take a transversal $\tau \in \tilde{\Omega}_{y,u}$ and s_0 such that $\tilde{\psi}_{s_0}^u(\tau) = \tau'$. Then

$$\ell_\mu(\tilde{\psi}_s^u(\tau')) = e^{y+s_0+s} \mu([\tau]) = e^s (e^{y+s_0} \mu([\tau])) = e^s \ell_\mu(\tau'). \quad \square$$

Lemma 14.7. *The H -action preserves the μ -lengths of flowed transversals:*

$$\ell_\mu(\tau) = \ell_\mu(h(\tau)).$$

Proof. Since $\tilde{\psi}^u$ commutes with the action H , appealing to Lemma 14.6 we may assume that τ is a transversal contained in $\tilde{\Omega}_{y,u}$. Suppose $h \in H$, and then $h(\tau) = h^{-1} \cdot \tau$ is contained in the fiber $\tilde{\Omega}_{y+u(h),u}$. Because μ is a homomorphism of $\mathbb{Z}[H]$ -modules we have

$$\ell_\mu(h(\tau)) = e^{y+u(h)} \mu(h^{-1} \cdot [\tau]) = e^y e^{u(h)} e^{-u(h)} \mu([\tau]) = e^y \mu([\tau]) = \ell_\mu(\tau). \quad \square$$

By a transversal τ in a leaf $\Omega_{y,u} \subset X$ we mean the image of a transversal in $\tilde{\Omega}_{y,u}$ under the covering projection $p: \tilde{X} \rightarrow X$. We similarly define a flowed transversal τ in $\Omega_{y,u} \subset X$ and note that by Lemma 14.7, the μ -length $\ell_\mu(\tau)$ of any such τ is well-defined as the μ -length back in \tilde{X} . Moreover, from Lemma 14.6 we have

$$\ell_\mu(\psi_s^u(\tau)) = e^s \ell_\mu(\tau)$$

for all $s \in \mathbb{R}$ and any transversal τ in any leaf $\Omega_{y,u}$.

Lemma 14.8. *Let $y, \epsilon \in \mathbb{R}$ with $\epsilon > 0$ and $\xi \in \Omega_{y,u}$ be given. Then there is a finite set of paths τ_1, \dots, τ_j in $\Omega_{y,u}$ whose union forms a neighborhood of ξ , which are each flowed transversals, and for which the sum of the μ -lengths is less than ϵ . If ξ is not a vertex of $\Omega_{y,u}$ then we can take $j = 1$ and $\tau = \tau_1$ to be a transversal.*

Proof. We give the proof for ξ not a vertex.

We suppose that the conclusion is false. Thus there is a sequence of flowed transversals $\{\tau_n\}$ in $\Omega_{y,u}$ containing ξ in their interiors for which

$$\ell_\mu(\tau_n) \geq \epsilon$$

and for which the intersection is $\{\xi\}$. Since ξ is not a vertex of $\Omega_{y,u}$, it lies in the interior of a 2-cell, and hence these flowed transversals can be taken to be transversals. Such a sequence exists since vertex leaves of \mathcal{F} are dense.

Suppose first that ξ lies on a periodic orbit of ψ^u . That is, for some $r > 0$ we have $\psi_{rj}^u(\xi) = \xi$ for all integers $j \geq 0$. Then for any j , there is some n such that τ_n is so small that

$$\psi_{rj}^u(\tau_n) \subset \tau_1.$$

On the other hand, together with Lemma 14.5 this implies

$$\ell_\mu(\tau_1) \geq \ell_\mu(\psi_{rj}^u(\tau_n)) = e^{rj} \ell_\mu(\tau_n) > e^{rj} \epsilon.$$

Since the left- and right-hand sides do not depend on n , and the inequality between these holds for all j , this contradicts the fact that $\ell_\mu(\tau_1)$ is a finite number.

A similar argument works if ξ flows into a point which is periodic, so we assume that this is not the case. Therefore, we can assume that $\{\psi_s^u(\xi)\}_{s \geq 0}$ accumulates somewhere. It follows that there is a sequence of times $s_j \rightarrow \infty$ such that $\psi_{s_j}^u(\xi) \in \Omega_{y',u}$ for all j and some y' , and that these points are contained in a single transversal τ in $\Omega_{y',u}$. Then, for any $j \geq 0$, we can find n such that τ_n is so small that

$$\psi_{s_1}^u(\tau_n), \psi_{s_2}^u(\tau_n), \dots, \psi_{s_j}^u(\tau_n)$$

are pairwise disjoint transversals inside τ . Now we note that $\ell_\mu(\tau)$ is some finite number, whereas by Lemma 14.5 we have

$$\ell_\mu(\tau) \geq \ell_\mu(\psi_{s_1}^u(\tau_n)) + \dots + \ell_\mu(\psi_{s_j}^u(\tau_n)) \geq j\epsilon.$$

Since the right-hand side tends to infinity, this is a contradiction.

This completes the proof for the case where ξ is not a vertex. The case where ξ is a vertex is similar and we leave it to the reader. \square

Remark 14.9. Lemma 14.8 can also be used to extend μ to any arc transverse to \mathcal{F} (say, contained in a 2-cell, so that transversality makes sense). The reason is that we can approximate any such arc by transversals, and the lemma tells us that the limit of the values of μ will be independent of the sequence of transversals used.

For every y and every edge e of $\Omega_{y,u}$, we next define a path metric called the μ -path metric on e inducing the given topology, so that the length of any transversal is the μ -length. For any two points ξ, ξ' contained in e , we consider a sequence of transversals whose endpoints limit to ξ and ξ' . By Lemma 14.5, the limit of the μ -lengths of these edges is independent of the sequence, and we call it the μ -path distance between ξ and ξ' . The

transversals and flowed transversals define a basis for the topology on e , and so Lemmas 14.5 and 14.8, together with the fact that every transversal has positive μ -length (Proposition 14.3), further implies that the μ -path distance is a path metric inducing the given topology.

We now define a metric on X for which the induced path metric on every edge e of $\Omega_{y,u}$ is precisely the μ -path metric. This metric on X is obtained as a quotient path metric of a collection of “nice” subsets of the hyperbolic plane, and will be constructed in a couple steps.

First, for every $y \in \mathbb{R}$ and every transversal $\tau \subset \Omega_{y,u} \subset X$ let $\beta_\tau : [0, r_0] \rightarrow \tau$ denote a unit speed parameterization with respect to the μ -path metric (note that τ is contained in some edge). Next, let $s_0 > 0$ and define

$$B_{\tau,s_0} : [0, r_0] \times [0, s_0] \rightarrow X$$

by

$$B_{\tau,s_0}(r, s) = \psi_s^\mu(\beta_\tau(r)).$$

We give $[0, r_0] \times [0, s_0]$ the Riemannian metric $e^{2s}dr^2 + ds^2$, and note that this is isometric to the region in the hyperbolic plane bounded by two asymptotic geodesics and two tangent horocycles with its induced path metric. We call this the *hyperbolic metric* on $[0, r_0] \times [0, s_0]$. Also observe that B_{τ,s_0} maps $[0, r_0] \times \{s\}$ into $\Omega_{y+s,u}$ and for any arc of $[0, r_0] \times \{s\}$ that maps into an edge e of $\Omega_{y+s,u}$, the map on this arc is a path isometry to the μ -path metric e .

We will also want to consider this same construction in a slightly more general setting, namely when τ is an injective flowed transversal. In this case, τ will be a concatenation of finitely many transversals and the μ -path metric makes sense on each of these and extends to a metric on all of τ . We then proceed by choosing a unit speed parameterization $B_\tau : [0, r_0] \rightarrow \tau$ and constructing the map $B_{\tau,s_0} : [0, r_0] \times [0, s_0] \rightarrow X$ as above.

We call any B_{τ,s_0} which is a homeomorphism onto its image a *hyperbolic model map*. Suppose $\delta : [0, 1] \rightarrow X$ is a path contained in the image of some hyperbolic model map B_{τ,s_0} . Define the d_u -length of δ to be

$$\ell_u(\delta) = \ell(B_{\tau,s_0}^{-1}(\delta))$$

where on the right, ℓ denotes length in the hyperbolic metric on $[0, r_0] \times [0, s_0]$. Given a path $\delta : [0, 1] \rightarrow X$ which can be subdivided into subpaths, each of which is contained in the image of some hyperbolic model map, we define the d_u -length to be the sum of the d_u -lengths of the subpaths. We call such a path a *tame path*. The d_u -length of any tame path is independent of the set of hyperbolic model maps used to calculate its length—this follows from the fact that on the overlap of any two, the composition $B_{\tau,s_0} \circ B_{\tau',s'_0}^{-1}$ (where defined) is a path isometry.

Now we define a pseudo-metric d_u on X by

$$d_u(\xi, \xi') = \inf\{\ell_u(\delta) \mid \delta \text{ is a tame path with } \delta(0) = \xi, \delta(1) = \xi'\}.$$

Proposition 14.10. *The pseudo-metric d_u is a geodesic metric on X . Furthermore, for every point ξ in the interior of a 2-cell, there is a hyperbolic model map B_{τ,s_0} and a disk $\Delta \in [0, r_0] \times [0, s_0]$ isometric to a disk in the hyperbolic plane such that $B_{\tau,s_0}|_{\Delta}$ maps isometrically onto a neighborhood of ξ .*

For any other point ξ , there exists a finite set $\{B_{\tau_i,s_{0,i}}\}_{i=1}^n$ of hyperbolic model maps and disks $\Delta_i \subset [0, r_{0,i}] \times [0, s_{0,i}]$ isometric to disks in the hyperbolic plane such that $B_{\tau_i,s_{0,i}}$ restricts to an isometric embedding and the union of the images defines a neighborhood of ξ .

Proof. For a point ξ in the interior of a 2-cell, we can find a hyperbolic model map that is a homeomorphism onto a neighborhood of ξ . Restricting to a sufficiently small disk provides the required neighborhood.

For a point ξ which is not in the interior of a 2-cell, we choose a finite number of hyperbolic model maps $\{B_{\tau_i,s_{0,i}}\}_{i=1}^n$ for which the images, denoted $\{V_i\}_{i=1}^n$, contain ξ , and the union $V = V_1 \cup \dots \cup V_n$ is a (closed) neighborhood of ξ . By shrinking the model maps if necessary, we can assume that (1) the domains are all the same, so $s_{0,i} = s_0$ and $r_{0,i} = r_0$ for some $r_0, s_0 > 0$ and all i , (2) the defining transversals (or flowed transversals) τ_1, \dots, τ_n are contained in a single leaf $\Omega_{y,u}$, and (3) $\xi = B_{\tau_i,s_0}(r_0/2, s_0/2)$ for all i . These conditions are easy to arrange and primarily serve to simplify the notation.

We note that the ‘‘top’’ $V_i^+ = B_{\tau_i,s_0}([0, r_0], \times \{s_0\}) \subset V_i$ is contained in the leaf $\Omega_{y+s_0,u}$ for all i , and the union $V^+ = V_1^+ \cup \dots \cup V_n^+$ forms a neighborhood of the point $\psi_{s_0/2}^u(\xi)$ (which is also the point $B_{\tau_i,s_0}(r_0/2, s_0)$, for all i) inside $\Omega_{y+s_0,u}$. By shrinking the models further if necessary, we can assume that V^+ is a contractible subset of the graph $\Omega_{y+s_0,u}$, hence V^+ is a tree.

Claim 14.11. *For each $i = 1, \dots, n$, there is a map $\Pi_i : V \rightarrow V_i$ such that*

$$\Pi_i \circ B_{\tau_j,s_0} : [0, r_0] \times [0, s_0] \rightarrow [0, r_0] \times [0, s_0]$$

satisfies the following conditions with respect to the hyperbolic metrics on domain and range:

- (i) *for every i, j , $\Pi_i \circ B_{\tau_j,s_0}$ is 1-Lipschitz, and*
- (ii) *for every i , $\Pi_i \circ B_{\tau_i,s_0}$ is the identity.*

See Figure 10.

Proof of Claim. For each i , the top V_i^+ is an arc, and we first define the restriction of Π_i to V^+ to be the ‘‘closest point projection’’ $V^+ \rightarrow V_i^+$; since V^+ is a tree, this projection is independent of the choice of a path metric used to define it. For each $i, j = 1, \dots, n$, we compose Π_i with the model maps and their inverses to obtain a map $\pi_{i,j} : [0, r_0] \rightarrow [0, r_0]$ given by

$$\pi_{i,j}(r) = B_{\tau_i,s_0}^{-1}(\Pi_i(B_{\tau_j,s_0}(r, s_0))).$$

Observe that $\pi_{i,i}$ is the identity. Furthermore, when $i \neq j$, $\pi_{i,j}$ is 1-Lipschitz: in fact, because of the way closest point projections work in a tree, we can subdivide $[0, r_0]$ into at most three segments so that on each segment $\pi_{i,j}$ is either an isometry, or constant.




Fig. 10. The maps $B_{\tau_1, s_0}, B_{\tau_2, s_0}$ from $[0, r_0] \times [0, s_0]$ (left) to the neighborhood V (middle), and the projections Π_1 and Π_2 back onto $[0, r_0] \times [0, s_0]$. In this example, the map $\Pi_2 \circ B_{\tau_1, s_0}$ maps the upper left rectangle to the lower right rectangle, and is given by $(r, s) \mapsto (r', s)$ where $r' = r$ if $r \geq r_0/2$ and $r' = r_0/2$ if $r \leq r_0/2$.

Now, given any point $\xi_0 \in V$ we define $\Pi_i(\xi_0)$ to be the point obtained by first flowing up for some time s to V^+ , projecting by closest point projection to V_i^+ , then *within* V_i flowing backward a time s . This last map in the composition is possible because inside V_i , the semiflow restricts to a local flow: since B_{τ_i, s_0} is a homeomorphism onto its image, this local flow is just the conjugate of the local flow ν on $[0, r_0] \times [0, s_0]$ defined by $\nu_t(r, s) = (r, s + t)$. With this notation we can write down the map Π_i :

$$\Pi_i(\xi_0) = \nu_{-s}(B_{\tau_i, s_0}^{-1}(\Pi_i(\psi_s^u(\xi_0))))$$

where $s \geq 0$ is the smallest number so that $\psi_s(\xi_0) \in V^+$.

For any $i, j = 1, \dots, n$, the composition $\Pi_i \circ B_{\tau_j, s_0}$ takes the simple form

$$\Pi_i \circ B_{\tau_j, s_0}(r, s) = (\pi_{i, j}(r), s).$$

It now follows from the properties of $\pi_{i, j}$ mentioned above that Π_i satisfies (i) and (ii) from the claim, for all i and j . □

Now suppose that $(r, s), (r', s') \in [0, r_0] \times [0, s_0]$ are any two points. If δ is a tame path between $B_{\tau_i, s_0}(r, s)$ and $B_{\tau_i, s_0}(r', s')$ contained in V , then $\Pi_i(\delta)$ is a path between (r, s) and (r', s') with hyperbolic length no more than the d_u -length of δ (since Π is 1-Lipschitz on each hyperbolic model where we compute lengths). Therefore, with respect to the path pseudo-metric on V induced by d_u -lengths, the map B_{τ_i, s_0} preserves distances. In particular, this pseudo-metric on V is a metric.

Finally, let $\epsilon > 0$ be half the hyperbolic distance from $(r_0/2, s_0/2)$ to the boundary $[0, r_0] \times \{0, s_0\} \cup \{0, r_0\} \times [0, s_0]$, and let Δ be a hyperbolic ball of radius ϵ centered at $(r_0/2, s_0/2)$.

Claim 14.12. *The restriction of B_{τ_i, s_0} to Δ is distance preserving with respect to the d_u pseudo-metric.*

Proof of Claim. Observe that from the discussion above, in the path metric on V induced by d_u -lengths, the union of the images of Δ ,

$$B_{\tau_1, s_0}(\Delta) \cup \dots \cup B_{\tau_n, s_0}(\Delta),$$

is precisely the ϵ -ball about ξ . Furthermore, for any $(r, s), (r', s') \in \Delta$, a path δ between $B_{\tau_i, s_0}(r, s)$ and $B_{\tau_i, s_0}(r', s')$ which leaves V must contain two subpaths, each connecting the frontier of $B_{\tau_i, s_0}(\Delta) \cup \dots \cup B_{\tau_n, s_0}(\Delta)$ to the frontier of V , completely contained in V . Considering appropriate projections Π_j and Π_k , we see that these subpaths must each have length at least ϵ , and so the path δ has length at least 2ϵ . Since the hyperbolic distance between (r, s) and (r', s') is less than 2ϵ , the hyperbolic geodesic in Δ between these points projects to a path whose d_u -length is strictly shorter than any path not entirely contained in V . But then distances between points in $B_{\tau_i, s_0}(\Delta)$ are computed in terms of d_u -lengths of paths in V . We have already seen that B_{τ_i, s_0} is distance preserving in the d_u -path metric on V , and so B_{τ_i, s_0} is distance preserving as a map from Δ into X . \square

All that remains is to see that d_u is actually a metric, and not a pseudo-metric—as a path metric on a compact space, it will necessarily be geodesic. For that, we just need to know that no two points have distance 0. This follows from the arguments just given: a path from some point ξ to any other point ξ' is either entirely contained in the neighborhood of ξ built from hyperbolic models, and hence is bounded away from zero, or else must exit the neighborhood and hence has length bounded away from 0 (by the radius of the hyperbolic disk Δ). This completes the proof. \square

Corollary 14.13. *For any $y \in \mathbb{R}$ the path metric on $\Omega_{y,u}$ induced by d_u makes $\Omega_{y,u}$ into a simplicial metric graph. Moreover, the path metric on any edge e induced by d_u is precisely the μ -path metric on e .*

Proof. Because ω^u is tame, the restriction to any skew 1-cell has only finitely many critical points. Therefore the components of intersection of the neighborhoods from Proposition 14.10 with a leaf $\Omega_{y,u}$ define neighborhoods of points in $\Omega_{y,u}$ which are obtained by gluing finitely many flowed transversals of the form $B_{\tau_i, s_0}([0, r_0] \times \{s\})$ along finitely many arcs. Furthermore, the metric d_u induces the μ -path metric on these flowed transversals. Since every point in any $\Omega_{y,u}$ has such a neighborhood, the corollary follows. \square

Proof of Theorem H. Suppose first that $u \in \mathcal{S}$ lies in \mathfrak{G} so that $\lambda = e^{\mathfrak{S}(u)} = e$. Let d_u be the metric constructed above. Proposition 14.10 guarantees that d_u is a metric, and provides enough local information to guarantee that d_u satisfies all the properties of the theorem as we now explain.

First, the flowlines of the local flow $v_t(r, s) = (r, t + s)$ on the $[0, r_0] \times [0, s_0]$ are unit speed geodesics in the hyperbolic metric on $[0, r_0] \times [0, s_0]$, and the hyperbolic model maps B_{τ_i, s_0} conjugate the local flow to the restriction of the semiflow on the image. Since B_{τ_i, s_0} is an isometric embedding, (1) follows. Part (2) follows from Corollary 14.13. In the hyperbolic metric on $[0, r_0] \times [0, s_0]$, v_s maps $[0, r_0] \times \{s'\}$ to $[0, r_0] \times \{s' + s\}$ and is

an e^s -homothety. Since $\lambda = e$ and B_{τ_t, s_0} conjugates v to the restriction of ψ^u , (3) is also true. Finally, (4) is an immediate consequence of the first part of Proposition 14.10 since it says that the metric on the interior of every 2-cell is locally isometric to the hyperbolic plane, hence has constant curvature -1 .

For an arbitrary class $u \in \mathcal{S}$, we use the fact that \mathcal{S} is contained in the McMullen cone \mathcal{C}_X , hence the ray through u intersects \mathfrak{G} in a unique point u' . Let $t > 0$ be such that $u' = tu$. Then $\lambda = e^{\mathfrak{H}(u)} = e^{\mathfrak{H}(\frac{1}{t}u')} = e^{t\mathfrak{H}(u')} = e^t$. Let $\omega^{u'}$ be the closed 1-form used above since $u' \in \mathfrak{G}$. Then $\omega^u = t\omega^{u'}$ represents u , and the foliation Ω_u is exactly the same as the foliation $\Omega_{u'}$ except that the leaves are given by $\Omega_{y,u} = \Omega_{ty,u'}$ (since $\tilde{\eta}_{u'} = t\tilde{\eta}_u$). Similarly, the reparameterization ψ^u is related to $\psi^{u'}$ by $\psi_s^u = \psi_{ts}^{u'}$.

Now let $d_u = (1/t)d_{u'}$. Since $\psi_s^u = \psi_{ts}^{u'}$, the ψ_s^u -flowlines are local geodesics, and on any edge of any leaf of $\Omega_u (= \Omega_{u'})$, $\psi_s^u = \psi_{ts}^{u'}$ is still an e^{ts} -homothety (this is unaffected by scaling the metric). Since $\lambda = e^t = e^{\mathfrak{H}(u)}$, this completes the proof. \square

Appendix A. Characterizing sections

The goal of this appendix is to prove Proposition A.12, which explicitly constructs dual cross sections for all integral classes in the Fried cone \mathcal{D} . To facilitate the construction, we give a combinatorial characterization of \mathcal{D} in terms of a positivity condition for cellular cocycle representatives with respect to a certain trapezoidal subdivision of X . This characterization is similar in spirit to the definition of the cone \mathcal{A} from [DKL]. This allows us to borrow many of the techniques from [DKL] for constructing sections and apply them to classes in the Fried cone \mathcal{D} . The constructions require a number of modifications and new ideas, and while the resulting sections still provide a wealth of information about the associated cohomology class, they do not (and cannot) behave as nicely as those from Theorem 5.11 in general.

We begin by recalling some of the relevant notions about trapezoidal subdivisions from [DKL]. Then we introduce the combinatorial characterization of \mathcal{D} and analyze some of its consequences. Next we describe a procedure for subdividing the cell structure and refining the cocycle so that it more closely resembles a cocycle representing an element of \mathcal{A} . From this we finally construct a section dual to any primitive integral class $u \in \mathcal{D}$ following the construction in [DKL] with appropriate modifications.

A.1. Trapezoidal subdivisions

Our construction of cross sections will involve working with various trapezoidal subdivisions of the cell structure on X . To this end we recall from [DKL, §6.3] the relevant structure of these subdivisions as well as some terminology and notation for working with them. A *trapezoidal cell structure* on X is one in which every 1-cell is either vertical or skew, and every 2-cell is a trapezoid. Each 1-cell e is equipped with a globally defined positive orientation so that the restriction $\eta|_e$ is orientation preserving. The boundary of a trapezoidal 2-cell T consists of four arcs $\ell_-(T)$, $\ell_+(T)$, $e_-(T)$, and $e_+(T)$ which we refer to as the left, right, bottom and top arcs of T , respectively. Each of these arcs is a

union of 1-cells and may therefore be regarded as a cellular 1-chain in X . The trapezoid induces a T -orientation on each of these arcs in which $\ell_{\pm}(T)$ and $e_{-}(T)$ are given the positive orientation; the sides $\ell_{\pm}(T)$ are then distinguished by the convention that $e_{-}(T)$ is oriented from $\ell_{-}(T)$ to $\ell_{+}(T)$. The T -orientation on the top arc $e_{+}(T)$ is defined so that it is also oriented from $\ell_{-}(T)$ to $\ell_{+}(T)$. This T -orientation on $e_{+}(T)$ may not agree with the globally defined positive orientation, and we define the *sign* $\zeta(T) \in \{\pm 1\}$ of T so that the 1-cells comprising the 1-chain $\zeta(T)e_{+}(T)$ appear with positive sign.

In any trapezoidal cell structure we require that the sides $\ell_{\pm}(T)$ of each trapezoid consist of vertical 1-cells (however $\ell_{+}(T)$ may degenerate to a vertex), and that the bottom arc $e_{-}(T)$ consists of a *single* skew 1-cell. Conversely, each skew 1-cell is equal to $e_{-}(T)$ for a unique trapezoid T . On the other hand, the top arc $e_{+}(T)$ of a trapezoid may consist of several skew 1-cells. Finally, we note that each 0-cell is the initial endpoint of a unique (positively oriented) vertical 1-cell.

A.2. A combinatorial characterization of the Fried cone \mathcal{D}

Recall that $f: \Gamma \rightarrow \Gamma$ is an expanding irreducible train track map with transition matrix $A(f)$ and associated transition graph $\mathcal{G}(f)$. The circuits \mathcal{Y} of $\mathcal{G}(f)$ determine a finite set of closed orbits $\{\mathcal{O}_y\}_{y \in \mathcal{Y}}$ as explained in §6.

We subdivide X along the closed orbits \mathcal{O}_y for all circuits $y \in \mathcal{Y}$. More precisely, for a given circuit $y \in \mathcal{Y}$, each skew 1-cell σ of X that intersects \mathcal{O}_y is subdivided by adding a 0-cell at $\sigma \cap \mathcal{O}_y$, and each trapezoidal 2-cell T of X that intersects \mathcal{O}_y is subdivided into two trapezoids by adding the vertical 1-cell $T \cap \mathcal{O}_y$. The cell structure obtained by performing these subdivisions for each $y \in \mathcal{Y}$ is called the *circuitry cell structure* on X and will be denoted X^c . Since \mathcal{Y} is finite, we note that X^c is indeed a finite trapezoidal subdivision of X .

Definition A.1 (Vertical positivity). A cellular 1-cocycle $z \in Z^1(Y; \mathbb{R})$ of a trapezoidal subdivision Y of X is said to be *vertically positive* if $z(\sigma) > 0$ for every vertical 1-cell σ of Y .

Proposition A.2. *The Fried cone \mathcal{D} is equal to the set of cohomology classes that can be represented by vertically positive 1-cocycles of the circuitry cell structure X^c .*

Proof. Recall from Proposition 6.3 that $\mathcal{D} \subset H^1(X; \mathbb{R})$ consists exactly of those classes which are positive on all orbits \mathcal{O}_y for $y \in \mathcal{Y}$. Since each \mathcal{O}_y consists of vertical 1-cells of X^c , it follows immediately that every vertically positive cocycle of the circuitry cell structure X^c represents a class in \mathcal{D} .

It remains to prove that each class in \mathcal{D} can be represented by a vertically positive cocycle of X^c . Let $u \in \mathcal{D}$ so that $u(\mathcal{O}) > 0$ for every closed orbit $\mathcal{O} \in \mathfrak{D}(\psi)$ of ψ . Let us say that a vertical 1-cell of X^c is *periodic* if it lies on a closed orbit of ψ , and *nonperiodic* otherwise. As closed orbits of ψ are necessarily disjoint, we see that the periodic 1-cells form a disjoint union of finitely many closed orbits $\gamma_1, \dots, \gamma_k$ of ψ (some of which are of the form \mathcal{O}_y for $y \in \mathcal{Y}$). The (oriented) nonperiodic 1-cells, on the other hand, form

a disjoint union of oriented rooted trees T_1, \dots, T_n (oriented towards the root) with the root of each tree lying on one of the closed orbits γ_i .

Choose any cocycle $z_0 \in Z^1(X^c; \mathbb{R})$ representing u . Since $u \in \mathcal{D}$, we then have $z_0(\gamma_i) > 0$ for each of the closed orbits $\gamma_1, \dots, \gamma_k$. Therefore, by adding coboundaries associated to the vertices of γ_i to distribute the value $z_0(\gamma_i)$ equally among all periodic 1-cells comprising γ_i , we obtain a cohomologous cocycle z'_0 that is positive on all periodic 1-cells.

Having accomplished this, we may now freely add coboundaries associated to any vertex in any tree T_j , except for the root, without affecting the value of z'_0 on the periodic 1-cells. Since each T_j is a tree, it is straightforward to add such coboundaries—working from the root down—so that the resulting cocycle z is positive on all nonperiodic 1-cells as well. Thus z is a vertically positive representative of u , as desired. \square

Recall that our objective is to prove $\mathcal{D} \subseteq \mathcal{S}$ by constructing cross sections dual to all integral classes $u \in \mathcal{D}$. The first step towards this goal is the following lemma, which is a quantified strengthening of Proposition 6.3.

Lemma A.3 (Positive on flowlines). *Given $u \in \mathcal{D}$, there exists $\kappa > 0$ with the following property. Let X' be any trapezoidal subdivision of X and suppose that $\xi \in X$ and $t > 0$ are such that ξ and $\psi_t(\xi)$ both lie on the same skew 1-cell σ of X' . If $\gamma \subset X$ denotes the closed loop obtained by concatenating the flowline from ξ to $\psi_t(\xi)$ with the arc of σ connecting $\psi_t(\xi)$ to ξ , then*

$$u(\gamma) \geq \kappa(t - 1).$$

Proof. Let $\delta > 0$ be the minimum of $u(\mathcal{O}_y)$ for $y \in \mathcal{Y}$ and let $L \geq 1$ denote the longest (combinatorial) length of any circuit in $\mathcal{G}(f)$. Set $\kappa = \delta/L$.

Choose the smallest $s \geq 0$ such that $q = \psi_s(\xi) \in \Gamma \subset X$, and consider the homotopic loop $\psi_s(\gamma)$. Let $e \in E\Gamma$ be the edge of Γ containing q . By pushing the “slanted” portion of the loop $\psi_s(\gamma)$ (running along $\psi_s(\sigma)$) onto e , we see that there is some integer $k \geq 1$ (near to t) such that $\psi_s(\gamma)$ is homotopic to the closed loop γ' obtained by concatenating the flowline from q to $\psi_k(q)$ with the arc along e from $\psi_k(q)$ to q . In fact, $|t - k|$ here is equal to the length of the image of the arc along σ from $\psi_t(\xi)$ to ξ under the map $\eta: X \rightarrow \mathbb{S}^1$. Since this is a subarc of some skew 1-cell of X , and by construction (see [DKL, §4.4]) every skew 1-cell of X projects under $\eta: X \rightarrow \mathbb{S}^1$ to an arc of length less than 1, we have $k \in (t - 1, t + 1)$.

There is a subinterval $\alpha \subset e$ containing q on which f^k is an affine homeomorphism onto e , and we let $q_0 \in \alpha$ denote the fixed point of f^k . There is a closed orbit \mathcal{O}_w through q_0 , which is homotopic to γ' and which also crosses Γ k times. It follows that the associated closed path w in $\mathcal{G}(f)$ has combinatorial length k . By Lemma 6.2 we can now write

$$[\gamma] = [\gamma'] = [\mathcal{O}_w] = [\mathcal{O}_{y_1}] + \dots + [\mathcal{O}_{y_n}]$$

for some circuits $y_i \in \mathcal{Y}$ whose combinatorial lengths sum to k . Since each y_i has length at most L , we have

$$u(\gamma) = u(\mathcal{O}_{y_1}) + \dots + u(\mathcal{O}_{y_n}) \geq n\delta \geq k\delta/L \geq (t - 1)\delta/L = \kappa(t - 1). \quad \square$$

A.3. Height

Recall that $p: \tilde{X} \rightarrow X$ is the universal torsion-free abelian cover of X ; the induced homomorphism $p_*: H_1(\tilde{X}; \mathbb{R}) \rightarrow H_1(X; \mathbb{R})$ is then trivial by definition of \tilde{X} . Given a trapezoidal subdivision Y of X , let \tilde{Y} denote the lifted trapezoidal cell structure on the universal torsion-free abelian cover and choose a vertex $v_0 \in \tilde{Y}^{(0)}$ to serve as a basepoint. Every cocycle $z \in Z^1(Y; \mathbb{R})$ then determines a height function $h_z: \tilde{Y}^{(0)} \rightarrow \mathbb{R}$ whose value at a vertex $v \in \tilde{Y}^{(0)}$ is defined to be

$$h_z(v) = \tilde{z}(\Sigma),$$

where $\tilde{z} = p^*(z)$ is the pull-back of z and Σ is any cellular 1-chain on \tilde{Y} with boundary $\partial\Sigma = v - v_0$. If Σ' is any other 1-chain with $\partial\Sigma' = v - v_0$, then $\Sigma' - \Sigma$ is a cycle and so $p_*(\Sigma' - \Sigma)$ is null-homologous in $H_1(X; \mathbb{R})$. Therefore $\tilde{z}(\Sigma' - \Sigma) = z(p_*(\Sigma' - \Sigma)) = 0$, showing that $h_z(v)$ is well-defined. We note that for any two vertices $v, v' \in \tilde{Y}^{(0)}$, the difference $h_z(v') - h_z(v)$ is equal to $\tilde{z}(\Sigma)$ for any 1-chain in \tilde{Y} with $\partial\Sigma = v' - v$. For notational convenience, we similarly define the height of a 1-cell σ of \tilde{Y} as $h_z(\sigma) = \min\{h_z(t(\sigma)), h_z(o(\sigma))\}$.

Lemma A.4. *Let Y be a trapezoidal subdivision of X with a cellular cocycle $z \in Z^1(Y; \mathbb{R})$. Then there exist $L \geq 0$ such that if \tilde{T} is a trapezoid of \tilde{Y} and σ_{\pm} are skew 1-cells along the arcs $e_{\pm}(\tilde{T})$, then $|h_z(\sigma_+) - h_z(\sigma_-)| \leq L$.*

Proof. For any trapezoid T of Y , let $M(T)$ denote the quantity

$$M(T) = \sum_{\sigma \in \partial T} |z(\sigma)|,$$

where the sum is over all 1-cells that occur in the 1-chain ∂T . Take L to be the maximum value of $M(T)$ among all (finitely many) trapezoids T of Y . For skew 1-cells σ_{\pm} as in the statement of the lemma, let v_{\pm} denote the vertex of σ_{\pm} achieving $h_z(\sigma_{\pm})$. Then there exists a 1-chain Σ consisting of cells in $\partial\tilde{T}$ for which $\partial\Sigma = v_+ - v_-$. By the above definition of L , it then follows that

$$|h_z(\sigma_+) - h_z(\sigma_-)| = |h_z(v_+) - h_z(v_-)| = |\tilde{z}(\Sigma)| \leq L. \quad \square$$

Combining this with Lemma A.3, we can now prove the following.

Lemma A.5 (Linear ascent). *Let Y be a trapezoidal subdivision of X and let $z \in Z^1(Y; \mathbb{R})$ be a cocycle representing a class $[z] \in \mathcal{D}$. Then there exist $\kappa > 0$ and $N \geq 0$ satisfying the following property. Suppose that $\xi \in \tilde{Y}$ and $t > 0$ are such that ξ and $\tilde{\psi}_t(\xi)$ lie on skew 1-cells σ and σ' of \tilde{Y} , respectively. Then*

$$h_z(\sigma') - h_z(\sigma) \geq \kappa t - N.$$

Proof. Let us declare two skew 1-cells of \tilde{Y} to be equivalent if they cover the same 1-cell of Y , and let \mathcal{W} denote the corresponding set of equivalence classes. Note that \mathcal{W} is finite. Let $\kappa > 0$ be the constant provided by Lemma A.3, and let $L \geq 0$ be as in Lemma A.4.

Write $\alpha_1 = \sigma$ and let $W_1 \in \mathcal{W}$ be its equivalence class. Set $s_1 = 0$ and define

$$t_1 = \max\{s \in [0, t] \mid \tilde{\psi}_s(\xi) \text{ lies in a skew 1-cell in } W_1\}$$

to be the last time that the flowline from ξ to $\tilde{\psi}_{t_1}(\xi)$ hits a skew 1-cell in W_1 . Denote the skew 1-cell containing $\tilde{\psi}_{t_1}(\xi)$ by $\beta_1 \in W_1$. If $t_1 = t$ then we stop, otherwise let

$$s_2 = \min\{s \in (t_1, t] \mid \tilde{\psi}_s(\xi) \text{ lies in a skew 1-cell}\}$$

be the next time that the flowline hits a skew 1-cell. We again let α_2 denote the skew 1-cell containing $\tilde{\psi}_{s_2}(\xi)$ and let $W_2 \in \mathcal{W}$ be its equivalence class. Notice that $W_2 \neq W_1$ by definition of t_1 and s_2 . We now similarly define

$$t_2 = \max\{s \in [0, t] \mid \tilde{\psi}_s(\xi) \text{ lies in a skew 1-cell in } W_2\}$$

to be the last time the flowline hits a skew 1-cell in W_2 and let $\beta_2 \in W_2$ denote that particular 1-cell. Note that $t_2 \geq s_2$ by definition.

Continuing in this manner, we find a sequence of times

$$0 = s_1 \leq t_1 < s_2 \leq t_2 < \dots < s_n \leq t_n = t$$

such that for each $1 \leq i \leq n$ the points $\tilde{\psi}_{s_i}(\xi)$ and $\tilde{\psi}_{t_i}(\xi)$ lie on skew 1-cells α_i and β_i , respectively, that are both members of the same equivalence class W_i , and moreover these equivalence classes W_1, \dots, W_n are all distinct. In particular we note that $n \leq |\mathcal{W}|$.

Now, the trapezoidal cell structure on X was explicitly constructed so that every flowline must hit a skew 1-cell of X (and thus also of Y) within time 2 [DKL, §4.4]. Therefore for each $1 \leq i < n$ we have $s_{i+1} - t_i \leq 2$ by definition of s_{i+1} . Furthermore, since the flowline from $\tilde{\psi}_{t_i}(\xi)$ to $\tilde{\psi}_{s_{i+1}}(\xi)$ is disjoint from all other skew 1-cells, it must be that β_i and α_{i+1} lie along the bottom and top arcs $e_-(\tilde{T})$ and $e_+(\tilde{T})$, respectively, for some trapezoidal 2-cell \tilde{T} of \tilde{Y} . Lemma A.4 therefore implies $|h_z(\alpha_{i+1}) - h_z(\beta_i)| \leq L$ for all $1 \leq i < n$.

On the other hand, since α_i and β_i are in the same equivalence class, Lemma A.3 implies that $h_z(\beta_i) - h_z(\alpha_i) \geq \kappa(t_i - s_i - 1)$. Therefore we may conclude

$$\begin{aligned} h_z(\sigma') - h_z(\sigma) &= \sum_{i=1}^n (h_z(\beta_i) - h_z(\alpha_i)) + \sum_{i=1}^{n-1} (h_z(\alpha_{i+1}) - h_z(\beta_i)) \\ &\geq \sum_{i=1}^n \kappa(t_i - s_i - 1) - L(n - 1) \\ &= \kappa \left(t - \sum_{i=1}^{n-1} s_{i+1} - t_i \right) - \kappa n - L(n - 1) \geq \kappa t - (3\kappa + L)|\mathcal{W}|. \quad \square \end{aligned}$$

This ‘‘linear ascent’’ is a crucial feature of classes in the Fried cone. For an integral class $u \in \mathcal{D}$ with cover $X_u \rightarrow X$ corresponding to $\ker(u)$, this shows that lifting any biinfinite flowline to X_u yields a path that exits both ends of X_u (cf. [Fri2, Theorem A] and [Wan, Theorem 1.1]).

A.4. Depth

In order to construct cross sections dual to certain cohomology classes, we will need to represent those classes by cocycles that enjoy certain properties. Here we introduce the first of these properties and show that classes in the cone \mathcal{D} have such representatives.

A biinfinite flowline in \tilde{X} is a path $\tilde{\gamma}: \mathbb{R} \rightarrow \tilde{X}$ such that for all $t \in \mathbb{R}$ and $s \geq 0$ we have $\tilde{\psi}_s(\tilde{\gamma}(t)) = \tilde{\gamma}(s + t)$. In this case we say that the set $\tilde{\gamma}((-\infty, 0])$ is a *backward flowline* of $\tilde{\gamma}(0)$. Notice that every point of \tilde{X} has infinitely many backward flowlines.

For a trapezoidal subdivision Y of X and a vertex $v \in \tilde{Y}^{(0)}$, a skew 1-cell σ of \tilde{Y} is said to be *behind* v if there is a backward flowline of v that intersects σ .

Definition A.6 (Depth). Let Y be a trapezoidal subdivision of X and let $z \in Z^1(Y; \mathbb{R})$ be a cellular cocycle. The z -depth of a vertex $v \in \tilde{Y}^{(0)}$ is defined to be the number of 1-cells σ of \tilde{Y} for which there exists some biinfinite flowline $\tilde{\gamma}$ through $v = \tilde{\gamma}(0)$ and a time $t < 0$ such that

- $\tilde{\gamma}(t) \in \sigma$, and
- $h_z(v) \leq h_z(\sigma')$ for every skew 1-cell σ' of \tilde{Y} that intersects $\tilde{\gamma}([t, 0])$.

The equivariance of $\tilde{\psi}$ and \tilde{z} under deck transformations of \tilde{Y} implies that vertices of \tilde{Y} that project to the same point of Y necessarily have the same z -depth. Thus we also define the z -depth of a vertex of Y to be the z -depth of any of its lifts.

While the z -depth of a vertex could in general be infinite, this is never the case when the cocycle z represents a class in \mathcal{D} :

Lemma A.7. *Let Y be a trapezoidal subdivision of X and let $z \in Z^1(Y; \mathbb{R})$ be a cocycle for which $[z] \in \mathcal{D}$. Then for any $M \geq 0$ and any vertex v of \tilde{Y} , there are only finitely many skew 1-cells σ of \tilde{Y} behind v satisfying*

$$h_z(\sigma) \geq h_z(v) - M.$$

In particular, every vertex of \tilde{Y} has finite z -depth.

Proof. Let $\kappa > 0$ and $N \geq 0$ be the constants provided by Lemma A.5. Since every flowline of X eventually hits a skew 1-cell, we can find $s \geq 0$ such that $\tilde{\psi}_s(v)$ lies in a skew 1-cell σ' of \tilde{Y} . Set $L = \kappa^{-1}(h_z(\sigma') - h_z(v) + M + N)$.

Suppose that $\tilde{\gamma}$ is any biinfinite flowline through $v = \tilde{\gamma}(0)$, and $t < 0$ is such that $\tilde{\gamma}(t)$ lies on a skew 1-cell σ of \tilde{Y} satisfying $h_z(\sigma) \geq h_z(v) - M$. Then by Lemma A.5 we have

$$h_z(\sigma') - h_z(\sigma) \geq \kappa(s - t) - N,$$

which together with the assumption on σ implies

$$\kappa(s - t) - N \leq h_z(\sigma') - h_z(v) + M.$$

In particular we see that

$$-t \leq \kappa^{-1}(h_z(\sigma') - h_z(v) + M + N) - s \leq L.$$

The above paragraph shows that any skew 1-cell σ of \tilde{Y} behind v that satisfies $h_z(\sigma) \geq h_z(v) - M$ must intersect $\tilde{\gamma}([-L, 0])$ for some backward flowline of v . But if we declare two biinfinite flowlines to be equivalent if they agree on $[-L, 0]$, then there are only finitely many equivalence classes of biinfinite flowlines $\tilde{\gamma}$ satisfying $\tilde{\gamma}(0) = v$. Thus, taking the union of $\tilde{\gamma}([-L, 0])$ over all flowlines $\tilde{\gamma}$ satisfying $\tilde{\gamma}(0) = v$, we see that there are only finitely many skew 1-cells intersecting this set. \square

Once one knows that every vertex has finite depth, it is straightforward to build a refined cocycle for which every vertex in fact has zero depth. We recall here that a *refinement* of a cellular cocycle $z \in Z^1(Y; \mathbb{R})$ to a subdivision Y' of Y is a cocycle $z' \in Z^1(Y'; \mathbb{R})$ such that $z'(\sigma) = z(\sigma)$ for every 1-cell σ of Y , where in the expression $z'(\sigma)$ we regard σ as a 1-chain in Y' .

Lemma A.8 (A shallow refinement). *Let Y be a trapezoidal subdivision of X and let $z \in Z^1(Y; \mathbb{R})$ be a vertically positive cocycle for which $[z] \in \mathcal{D}$. Then there exists a trapezoidal subdivision Y' of Y and a vertically positive refinement $z' \in Z^1(Y'; \mathbb{R})$ of z such that every vertex of Y' has z' -depth zero.*

Proof. Let K denote the largest depth of any vertex in Y , and choose a vertex \tilde{v} of \tilde{Y} with z -depth K . We may assume $K \geq 1$. Lemma A.7 implies that the set

$$\{h_z(\tilde{v}) - h_z(\tilde{\sigma}) \mid \tilde{\sigma} \text{ is a skew 1-cell behind } \tilde{v}\}$$

is discrete. Therefore we may choose $\delta > 0$ such that every skew 1-cell $\tilde{\sigma}$ behind \tilde{v} satisfying $h_z(\tilde{\sigma}) < h_z(\tilde{v})$ also satisfies $h_z(\tilde{\sigma}) < h_z(\tilde{v}) - \delta$.

Since \tilde{v} has z -depth equal to $K \geq 1$, Definition A.6 implies that there exists a biinfinite flowline $\tilde{\gamma}$ with $\tilde{\gamma}(0) = \tilde{v}$, a time $t < 0$, and a skew 1-cell $\tilde{\sigma}$ of \tilde{Y} satisfying $h_z(\tilde{v}) \leq h_z(\tilde{\sigma})$ such that $\tilde{\gamma}(t) \in \tilde{\sigma}$ and $\tilde{\gamma}([t, 0])$ is disjoint from all skew 1-cells of \tilde{Y} . Letting \tilde{T} denote the trapezoid of \tilde{Y} for which $e_-(\tilde{T}) = \tilde{\sigma}$, it follows that $\tilde{v} \in e_+(\tilde{T})$. Now, if the arc $\tilde{\gamma}([t, 0])$ were equal to $\ell_\pm(\tilde{T})$ and thus consisting of vertical 1-cells, then the vertical positivity of z would imply $h_z(\tilde{\sigma}) < h_z(\tilde{v})$. As this is not the case, the arc $\tilde{\gamma}([t, 0])$ must therefore traverse the interior of \tilde{T} .

Recalling that $p: \tilde{Y} \rightarrow Y$ is the covering map, let $T = p(\tilde{T})$ and $\sigma = p(\tilde{\sigma})$. It follows that $p(\tilde{v})$ lies in the interior of $e_+(T)$ and that $w = p(\tilde{\gamma}(t))$ is an interior point of $e_-(T) = \sigma$. We now subdivide σ into two skew 1-cells α_1, α_2 by adding a vertex at w and subdivide T into two trapezoids T_1 and T_2 by adding the vertical 1-cell $\beta = p(\tilde{\gamma}([t, 0]))$. Call the resulting subdivision Y' and let \tilde{w} and $\tilde{\alpha}_i$ respectively denote the unique lifts of w and α_i that are contained in \tilde{T} . We refine z to a cocycle $z' \in Z^1(Y'; \mathbb{R})$ by setting $z'(\beta) = \delta$ and defining $z'(\alpha_i)$ subject to the cocycle condition for T_i . Notice that z' is again vertically positive.

Let us now consider the z' -depth of vertices in Y' . First recall that $h_z(\tilde{v}) \leq h_z(\tilde{\sigma}) = \min\{h_z(t(\tilde{\sigma})), h_z(o(\tilde{\sigma}))\}$. Since $h_{z'}(\tilde{w}) = h_{z'}(\tilde{v}) - \delta$ by construction, it follows that $h_{z'}(\tilde{w}) < h_z(t(\tilde{\sigma})), h_z(o(\tilde{\sigma}))$. Therefore $h_{z'}(\tilde{\alpha}_i) < h_z(\tilde{\sigma})$ for $i = 1, 2$, and the same holds for all equivariant translates of $\tilde{\alpha}_1, \tilde{\alpha}_2, \tilde{\sigma}$. In particular, whenever $\tilde{\alpha}_i$ is behind some vertex \tilde{v}_0 of \tilde{Y} satisfying $h_z(\tilde{\sigma}) < h_z(\tilde{v}_0)$, then also $h_{z'}(\tilde{\alpha}_i) < h_{z'}(\tilde{v}_0)$ for $i = 1, 2$. This shows that the z' -depth of every vertex of \tilde{Y} is bounded above by its z -depth. Moreover,

since $h_{z'}(\tilde{\alpha}_i) = h_{z'}(\tilde{v}) - \delta$, neither of the new skew 1-cells $\tilde{\alpha}_i$ contribute to the z' -depth of \tilde{v} , whereas $\tilde{\sigma}$ did contribute to the z -depth of \tilde{v} . Thus the z' -depth of \tilde{v} is strictly less than its z -depth.

It remains to analyze the z' -depth of the new vertex w . Notice that the set of skew 1-cells of \tilde{Y}' behind \tilde{w} is a subset of the ones behind \tilde{v} . Furthermore any skew 1-cell $\tilde{\sigma}_0$ of \tilde{Y}' behind \tilde{v} satisfying $h_{z'}(\tilde{\sigma}_0) < h_{z'}(\tilde{v})$ also satisfies $h_{z'}(\tilde{\sigma}_0) < h_{z'}(\tilde{w}) = h_{z'}(\tilde{v}) - \delta$ by the choice of δ . Thus the set of skew 1-cells of \tilde{Y}' counting towards the z' -depth of \tilde{w} is a subset of those counting towards the z' -depth of \tilde{v} . This shows that the z' -depth of \tilde{w} is bounded by the z' -depth of \tilde{v} , which we have seen is at most $K - 1$.

Recall that K was chosen to be the maximal z -depth of any vertex in Y . We have constructed a subdivision Y' and cocycle refinement $z' \in Z^1(Y'; \mathbb{R})$ such that no vertex has z' -depth more than K and the number of vertices in Y' with z' -depth equal to K is strictly less than the corresponding number of vertices in Y . By repeating this recursively, we eventually arrive at the desired subdivision in which every vertex has depth zero. \square

A.5. Consistent signage and the refinement procedure

We now introduce the second property we require of our cocycles in order to construct dual cross sections. Recall that if T is a trapezoidal 2-cell in a trapezoidal subdivision Y of X , then we may write its top arc as a cellular 1-chain $e_+(T) = \zeta(T)(\sigma_1 + \dots + \sigma_k)$ where each σ_i is a (positively oriented) skew 1-cell in Y .

Definition A.9 (Consistently signed). Let Y be a trapezoidal subdivision of X and let $z \in Z^1(Y; \mathbb{R})$ be a cocycle. We say that z is *consistently signed* if it is vertically positive and for every trapezoid T of Y with top arc $e_+(T) = \zeta(T)(\sigma_1 + \dots + \sigma_k)$, all of the numbers $z(\sigma_i)$ are nonzero and have the same sign.

Before showing how to arrange for a cocycle to be consistently signed, we first recall the subdivision procedure introduced in [DKL, §6.4]. Briefly, if Y is a trapezoidal subdivision of X , then following [DKL, Definition 6.8] the *standard subdivision* \hat{Y} of Y is obtained as follows: For each trapezoid T of Y whose top arc $e_+(T) = \zeta(T)(\sigma_1 + \dots + \sigma_k)$ contains at least $k \geq 2$ skew 1-cells, we subdivide the bottom arc $e_-(T)$ (which is a single skew 1-cell) into k skew 1-cells that flow homeomorphically onto the skew 1-cells $\sigma_1, \dots, \sigma_k$ comprising $e_+(T)$, and we add $k - 1$ vertical 1-cells through T connecting the newly added vertices on $e_-(T)$ to the already existing vertices along $e_+(T)$. (Note that the first step of the standard subdivision procedure, in which all invariant bands are subdivided, has no effect here since f is an expanding irreducible train track map; see [DKL, Remark 6.7].)

Lemma A.10. *Let Y be a trapezoidal subdivision of X and let $z \in Z^1(Y; \mathbb{R})$ be a vertically positive cocycle such that every vertex of v has z -depth equal to zero. Then there exists a trapezoidal subdivision Y' of Y and a consistently signed cocycle $z' \in Z^1(Y'; \mathbb{R})$ that refines z .*

Proof. In fact we will take Y' to be the standard subdivision \hat{Y} . Let T' be any trapezoidal 2-cell of Y' . Then, by the nature of the standard subdivision procedure, the top arc

$\zeta(T')e_+(T')$ of T' , when considered as a 1-chain in Y , is equal to a single skew 1-cell σ of Y . This skew 1-cell σ is in turn equal to $e_-(T)$ for a unique trapezoid T of Y . Writing the top arc of T as a concatenation of skew 1-cells of Y , we let v_0, v_1, \dots, v_k be the vertices of Y along $e_+(T)$ (traversed in that order with the T -orientation), so that v_0 is the top endpoint of $\ell_-(T)$ and v_k is the top endpoint of $\ell_+(T)$. Then in the subdivision procedure, new vertices u_1, \dots, u_{k-1} are placed along σ (indexed so that u_i flows onto v_i), and for each $1 \leq i < k$ a vertical 1-cell β_i is added connecting u_i to v_i . We also let $u_0 = o(\sigma)$ and $u_k = t(\sigma)$ denote the bottom endpoints of $\ell_-(T)$ and $\ell_+(T)$, respectively. Thus in Y' , the skew 1-cell σ is subdivided into the k skew 1-cells $\alpha_i = (u_{i-1}, u_i)$ for $1 \leq i \leq k$.

With this notation, writing the top arc of T' as a 1-chain in Y' we now have $e_+(T') = \zeta(T')(\alpha_1 + \dots + \alpha_k)$. Thus we must define z' on the new 1-cells α_i and β_i so that z' refines z and all of the numbers $z(\alpha_1), \dots, z(\alpha_k)$ have the same sign. Notice that σ is behind each of the vertices v_0, \dots, v_k . Since every vertex of Y has zero z -depth by assumption, we therefore have $h_z(\sigma) < h_z(v_i)$ for each $0 \leq i \leq k$. Given any partition of unity $\rho_1 + \dots + \rho_k = 1$ with each $\rho_i > 0$, if we define $z'(\alpha_i) = \rho_i z(\sigma)$ for $1 \leq i \leq k$, then the numbers $h_{z'}(u_1), \dots, h_{z'}(u_{k-1})$ will interpolate between $h_{z'}(u_0) = h_z(o(\sigma))$ and $h_{z'}(u_k) = h_z(t(\sigma))$. Therefore, since each v_i satisfies $h_z(\sigma) < h_z(v_i)$, we may choose the partition $\rho_1 + \dots + \rho_k = 1$ so that $h_{z'}(u_i) < h_z(v_i)$ for all $1 \leq i < k$ (in fact, we could even arrange to have $h_{z'}(u_i) < \min\{h_z(v_0), \dots, h_z(v_k)\}$ for each $1 \leq i < k$). Defining each $z'(\alpha_i)$ in this way, we then set $z'(\beta) = h_z(v_i) - h_{z'}(u_i) > 0$ for $1 \leq i < k$ and leave the value of z unchanged on the vertical 1-cells of $\ell_{\pm}(T)$. This ensures that z' satisfies the cocycle condition and that it is moreover a vertically positive refinement of z . \square

In [DKL, §6.5] we also introduced an accompanying *standard refinement* procedure for positive cocycles. However, the reader may check that Definitions 6.10 (skew ratios) and 6.11 (standard refinement) of [DKL] make sense verbatim if each occurrence of the word “positive” is replaced with “consistently signed”. The proofs of Lemmas 6.12–6.13 of [DKL] also go through verbatim in this more general context to give the following.

Lemma A.11. *Let Y_0 be a trapezoidal subdivision of X and $z_0 \in Z^1(Y_0; \mathbb{R})$ a consistently signed cocycle. Consider the sequence of standard subdivisions Y_0, Y_1, \dots with corresponding consistently signed cocycle refinements defined recursively by $Y_{n+1} = \widehat{Y}_n$ and $z_{n+1} = \widehat{z}_n$. Then the numbers*

$$S_n = \max\{|z_n(\sigma)| \mid \sigma \text{ is a skew 1-cell of } Y_n\}$$

tend to zero as $n \rightarrow \infty$.

A.6. Constructing cross sections

Finally, we construct cross sections dual to all integral classes $u \in \mathcal{D}$. For the proof, we first recall from [DKL, Definition 6.14] that a trapezoidal 2-cell T is said to be *unconstrained* by a cocycle z if

$$\max\{0, z(e_-(T))\} \leq \min\{z(\ell_-(T)), z(\ell_-(T)) + z(e_+(T))\}.$$

By the cocycle condition, this is equivalent to $|z(e_-(T))| \leq z(\ell_\pm(T))$ when z is vertically positive. We also recall that a trapezoidal 2-cell is said to be *degenerate* if its right side $\ell_+(T)$ consists of a single vertex (as in [DKL, Figure 8]).

Proposition A.12. *Given a primitive integral class $u \in \mathcal{D}$, there exists a flow-regular map $\eta_u: X \rightarrow \mathbb{S}^1$ with $(\eta_u)_* = u: G \rightarrow \pi_1(\mathbb{S}^1) = \mathbb{Z}$ and a fiber $\Theta_u = \eta_u^{-1}(y_0)$ for some $y_0 \in \mathbb{S}^1$ such that*

- (1) Θ_u is a finite, connected topological graph such that $\iota_*(\pi_1(\Theta_u)) \leq \ker(u)$, where ι_* is the homomorphism induced by the inclusion $\iota: \Theta_u \hookrightarrow X$,
- (2) Θ_u is a section of ψ dual to u , and
- (3) Θ_u is \mathcal{F} -compatible (and so the first return map $f_u: \Theta_u \rightarrow \Theta_u$ is an expanding irreducible train track map by Proposition 7.7).

Remark A.13. We note that the sections produced by Proposition A.12 will generally not have all the properties ensured by Theorem 5.11. Namely, for $u \in \mathcal{A}$ the fiber Θ_u , as provided by Theorem 5.11, will π_1 -inject into X and the first return map will be a homotopy equivalence. However, both of these properties will fail whenever $u \in \mathcal{D}$ lies outside of $\widehat{\mathcal{A}}$.

Proof. By Proposition A.2, we can represent $u \in \mathcal{D}$ by a vertically positive cocycle $z \in Z^1(X^c, \mathbb{R})$. Using Lemma A.8, we may then pass to a trapezoidal subdivision of X^c and a corresponding vertically positive refinement of z with respect to which every vertex has zero depth. By Lemma A.10, we may pass to a further subdivision Y_0 on which there is a refinement $z_0 \in Z^1(Y_0; \mathbb{R})$ of z that is consistently signed. We then consider the sequence Y_n of standard subdivisions with corresponding refinements z_n defined recursively by $Y_{n+1} = \widehat{Y}_n$ and $z_{n+1} = \widehat{z}_n$.

Let D_n denote the set of degenerate trapezoids in Y_n . Notice that the cardinality of D_n is constant independent of n . Indeed, when a degenerate trapezoid T of Y_n is subdivided, exactly one of the resulting trapezoids $T' \subset T$ is degenerate, and its degenerate side $\ell_+(T')$ is equal to $\ell_+(T)$. Thus if we define

$$U_N = \bigcup_{T \in D_n} \overline{T} \quad \text{and} \quad V_n = \{\ell_+(T) \mid T \in D_n\}$$

to be the union of degenerate trapezoids and the set of “degenerate sides” of trapezoids in Y_n , respectively, then all of the sets V_0, V_1, \dots are equal. Furthermore, the infinite intersection $\bigcap_n U_n$ is exactly V_0 . This follows from the fact that the union $\bigcup_n Y_n^{(0)}$ is dense in every skew 1-cell of Y_0 , which is a consequence f being an expanding irreducible train track map. Therefore, we may choose an index N such that all trapezoids in D_N have disjoint closures. Set $U = U_N$.

Let $M > 0$ denote the minimum of $z_N(\sigma)$ over all vertical 1-cells of Y_N . By Lemma A.11, we may choose $n \gg N$ so that every skew 1-cell σ of Y_n has $|z_n(\sigma)| < M$. Furthermore, as in [DKL, proof of Proposition 6.15], we additionally have $z_n(\beta) \geq M$ for every vertical 1-cell β of Y_n that is not contained in U . Therefore, every trapezoidal 2-cell T of Y_n that is not contained in U satisfies $|z_n(e_-(T))| \leq z_n(\ell_\pm(T))$ and is consequently unconstrained by z_n .

Set $Y = Y_n$ and by abuse of notation let $z = z_n$. We now use this trapezoidal subdivision to construct the desired map $\eta_u : Y \rightarrow \mathbb{S}^1$. The construction proceeds largely as in [DKL, proof of Lemma 6.16]: Choose a maximal tree $Q \subset Y^{(1)}$ and fix a vertex $v_0 \in Q$ to serve as a basepoint. Define a map $\hat{\eta}_u : Q \rightarrow \mathbb{R}$ so that $\hat{\eta}_u(v_0) = 0$ and for every 1-cell e in Q the restriction to e is a diffeomorphism onto its image and

$$\hat{\eta}_u(t(e)) - \hat{\eta}_u(o(e)) = z(e).$$

While this equation may fail for 1-cells not contained in Q , since z represents an integral cohomology class we will still have

$$\hat{\eta}_u(t(e)) - \hat{\eta}_u(o(e)) - z(e) \in \mathbb{Z}$$

for such 1-cells. Therefore, the composition of $\hat{\eta}_u$ with $\pi : \mathbb{R} \rightarrow \mathbb{S}^1$ may be extended to a map $\eta'_u : Y^{(1)} \rightarrow \mathbb{S}^1$ with the property that for any 1-cell e , any lift $\tilde{\eta}'_u|_e : e \rightarrow \mathbb{R}$ is an injective local diffeomorphism for which

$$\tilde{\eta}'_u|_e(t(e)) - \tilde{\eta}'_u|_e(o(e)) = z(e).$$

Since the trapezoids contained in U may not be unconstrained by z , we must take care to ensure that η'_u can be extended to a flow-regular map. Let T be a trapezoidal 2-cell contained in U . Flowing the bottom arc of T onto the top using ψ gives a homeomorphism $h_T : e_-(T) \rightarrow e_+(T)$ which is orientation preserving with respect to the T -orientation on $e_{\pm}(T)$. Consider any lift $\tilde{\eta}'_u|_{\partial T}$ of the restriction of η'_u to ∂T . Since z is consistently signed, this maps the arcs $e_{\pm}(T)$ homeomorphically onto intervals $I_{\pm} = \tilde{\eta}'_u|_{\partial T}(e_{\pm}(T)) \subset \mathbb{R}$ which we equip with orientations so that these identifications $e_{\pm}(T) \rightarrow I_{\pm}$ are orientation preserving. Let $A : I_+ \rightarrow I_-$ be the unique orientation preserving (with respect to the above orientations) affine homeomorphism from I_+ to I_- . We now redefine η'_u on the 1-cell $e_-(T)$ by declaring

$$\eta'_u|_{e_-(T)} = \pi \circ A \circ \tilde{\eta}'_u|_{\partial T} \circ h_T.$$

Notice that this agrees with the original value of η'_u at the endpoints of $e_-(T)$. Since the interior of $e_-(T)$ is disjoint from the closure of every other trapezoid contained in U (this follows from the choice of N), we may freely make these adjustments on the bottom arcs of all trapezoids in U without conflict. Let us denote the resulting (adjusted) map by $\eta_u : Y^{(1)} \rightarrow \mathbb{S}^1$. Then for any trapezoid T contained in U and any lift $\tilde{\eta}_u|_{\partial T} : \partial T \rightarrow \mathbb{R}$, we have

$$\tilde{\eta}_u|_{\partial T}(\xi) < \tilde{\eta}_u|_{\partial T}(h_T(\xi)) \tag{A.14}$$

for every interior point $\xi \in e_-(T)$. In fact, since all other trapezoids are unconstrained by z , this property holds for every trapezoid of Y .

Choose a point $y_0 \in \mathbb{S}^1 \setminus \eta_u(Y^{(0)})$. Since the vertex leaves of \mathcal{F} are dense in X , after adjusting η_u by a (small) homotopy rel $Y^{(0)}$, maintaining all of the properties above, we may assume that $(\eta_u)^{-1}(y_0) \subset Y^{(1)}$ is a finite set of points lying on vertex leaves of \mathcal{F} .

We now extend η_u over the 2-cells of Y . Consider any trapezoid T and a lift $\tilde{\eta}_u|_{\partial T} : \partial T \rightarrow \mathbb{R}$ of $\eta_u|_{\partial T}$ to the universal cover. As in [DKL, proof of Lemma 6.16],

we may identify \overline{T} with the Euclidean trapezoid provided by [DKL, Proposition 4.20] and then extend $\tilde{\eta}_u|_{\partial T}$ over \overline{T} in such a way that the fibers are straight (Euclidean) lines. Equation (A.14) then ensures that each fiber has nonvertical slope and is thus transverse to the foliation of \overline{T} by vertical flowlines. We then extend η_u to the interior of T by declaring it to equal $\pi \circ \tilde{\eta}_u|_{\overline{T}}$ on \overline{T} . Extending in this way over every 2-cell of Y , we obtain a map $\eta_u : Y \rightarrow \mathbb{S}^1$.

Since the fibers of η_u are transverse to the foliation of X by flowlines, we see that η_u is flow-regular on the interior of each 2-cell. However, since η_u was extended to the various 2-cells independently, it need not be flow-regular on a neighborhood of $Y^{(1)}$. Nevertheless, adjusting by a homotopy rel $Y^{(1)}$, we may smooth out and “straighten” the fibers of η_u in a neighborhood of $Y^{(1)}$ so that $\eta_u : Y \rightarrow \mathbb{S}^1$ is flow-regular on all of Y .

The map $\eta_u : Y \rightarrow \mathbb{S}^1$ now satisfies all of the conclusions of the proposition. By construction the fiber $\Theta_u = \eta_u^{-1}(y_0) \subset X$ is a connected topological graph which, since η_u is flow-regular, is necessarily a cross section dual to $(\eta_u)_*$. Indeed, Lemma A.5 explicitly shows that every flowline hits Θ_u infinitely often, and Proposition 5.5 ensures that Θ_u is connected. Moreover, our choice of $y_0 \in \mathbb{S}^1$ guarantees that $\Theta_u \cap Y^{(1)}$ is contained in vertex leaves of \mathcal{F} and thus Θ_u is \mathcal{F} -compatible. Finally, the desired relationship $(\eta_u)_* = u$ follows from the fact that for each 1-cell e of Y , any lift of the restriction $\eta_u|_e$ maps e homeomorphically onto an interval whose endpoints differ by $z(e)$. This in turn implies that $\iota_*(\pi_1(\Theta_u)) \leq \ker(u)$. □

Appendix B. Local boundedness of the stretch function

Our goal in this appendix is to prove Proposition 3.7 and thereby show that the stretch function $\Lambda : \widehat{\mathbb{Q}\Sigma}(G) \rightarrow \mathbb{R}_+$ defined in §3.2 is locally bounded. To this end, let $u_1 \in H^1(G; \mathbb{R})$ be a primitive integral point that lies in $\Sigma(G)$. Letting $G = B *_{\phi_1}$ be an HNN-extension compatible with u_1 (that is, with $B \leq \ker(u_1)$ a finite rank free group and $\phi_1 = \varphi_{u_1}|_B : B \rightarrow B$ an injective endomorphism), we then have $\Lambda(u_1) = \lambda(\phi_1)$ by definition.

B.1. An expanding representative

Let $f_1 : \mathcal{R} \rightarrow \mathcal{R}$ be a regular, irreducible expanding graph map on a rose \mathcal{R} with an identification $B \cong \pi_1(\mathcal{R})$ so that $(f_1)_* = \phi_1$. To simplify the discussion, we assume (as we may) that the transition matrix is not only irreducible, but is in fact positive. Let v_0 be the wedge point of \mathcal{R} and for $K \geq 0$ set

$$\mathcal{V}_k = \bigcup_{i=1}^K f_1^{-k}(v_0) \quad \text{and} \quad \mathcal{V}_\infty = \bigcup_{k \geq 0} \mathcal{V}_k.$$

Since f_1 has a positive transition matrix, \mathcal{V}_K nontrivially intersects in the interior of every 1-cell of \mathcal{R} for every $K \geq 1$.

Lemma B.1. *For any $K \geq 0$, $0 \leq k \leq K + 1$, and any arc e in the complement of \mathcal{V}_K , the restriction of f_1^k to e is locally injective.*

Proof. The map f_1 is locally injective on the interior of every edge of \mathcal{R} since f_1 is regular. It follows that local injectivity of f_1^2 can fail only at points $f_1^{-1}(v_0) = \mathcal{V}_1$; hence f_1^2 is locally injective on any arc in the complement of \mathcal{V}_1 . As compositions of locally injective maps are locally injective, it follows by induction that failure of local injectivity of f_1^k can only occur at points of \mathcal{V}_{k-1} , and there are no such points in e . \square

For technical reasons which will be clear later, we give each 1-cell e of \mathcal{R} an auxiliary linear structure determined by a characteristic map $[0, 1] \rightarrow e$ which maps the rational points bijectively onto $\mathcal{V}_\infty \cap e$. This is possible because both are countable dense sets (for \mathcal{V}_∞ this is because f_1 is expanding and irreducible). Furthermore, we assume, as we may, that the new linear structure and the original linear structure define the same smooth structure on the edges. It follows that for any arc α of an edge e with endpoints in \mathcal{V}_∞ , a linear map $[0, 1] \rightarrow \alpha$ also sends $\mathbb{Q} \cap [0, 1]$ onto $\mathcal{V}_\infty \cap \alpha$ (since such a linear map differs from the restriction of $[0, 1] \rightarrow e$ by a rationally defined affine map of $\mathbb{R} \rightarrow \mathbb{R}$).

B.2. The mapping torus and its cell structure

Let Y be the mapping torus of f_1 , and let ψ denote the suspension semiflow. By van Kampen’s Theorem, $\pi_1(Y) = G$. For a map $\eta: Y \rightarrow \mathbb{S}^1 = \mathbb{R}/\mathbb{Z}$ or to \mathbb{R} , if the composition of any flowline $x \mapsto \psi_s(x)$ with η is smooth, we will say that η is ψ -smooth, and the map $\psi(\eta): Y \rightarrow \mathbb{R}$, defined as the derivative of η with respect to the flow parameter, will be called the ψ -derivative.

For any $K \geq 1$, we define a cell structure on Y , denoted Y_K , as follows. We embed \mathcal{R} into Y as $\mathcal{R} \times \{0\}$ in $\mathcal{R} \times [0, 1]$. The 0-cells are the image of \mathcal{V}_K by this embedding. The 1-cells are the edges of the resulting subdivision of \mathcal{R} (viewed in Y) together with the following arcs of flowlines:

$$\left\{ \bigcup_{0 \leq s \leq 1} \psi_s(v) \right\}_{v \in \mathcal{V}_K}.$$

We refer to these two types of 1-cells as *skew* and *vertical*, respectively. We give the skew 1-cells the linear structure induced from the auxiliary linear structure on \mathcal{R} , and the vertical 1-cells the linear structure for which the flow parameter is linear. The complement of the 1-skeleton $Y_K^{(1)}$ is a union of topological disks which we define to be the 2-cells of Y_K . Whenever we write Y we will assume it is given the cell structure Y_1 , unless otherwise stated.

Lemma B.2. *For any $K \geq 1$, let $\alpha: (0, 1) \rightarrow Y_K$ be an arc contained in (the interior of) a 2-cell transverse to ψ . Let $H_\alpha: (0, 1) \times [0, K) \rightarrow Y$ be given by $H_\alpha(t, s) = \psi_s(\alpha(t))$. Then H_α is locally injective.*

Proof. The bottom of the 2-cell is a skew 1-cell of Y_K , and we let $\sigma: (0, 1) \rightarrow \mathcal{R} \subset Y_K$ be a characteristic map for this cell. We can similarly define $H_\sigma: (0, 1) \times [0, K + 1) \rightarrow Y$,

which is locally injective by Lemma B.1. Note that H_α “factors through” H_σ in the sense that there is an embedding $h: (0, 1) \times [0, K) \rightarrow (0, 1) \rightarrow [0, K + 1)$ such that $H_\sigma \circ h = H_\alpha$. Since H_σ is locally injective, so is H_α . \square

For any $K \geq 1$, the boundary of each 2-cell of Y_K is a union of two vertical 1-cells, a single skew 1-cell on the “bottom”, and a finite union of skew 1-cells on the “top”. Note that the interior of the bottom and top are each contained in a single original edge of \mathcal{R} , and with respect to the auxiliary linear structure, the flow identifies the bottom and top by a diffeomorphism. We use this, together with the flow parameter, to define a smooth structure on each 2-cell of Y_K , viewed as a product.

B.3. Maps to \mathbb{S}^1

Let $\eta_{1,0}: Y \rightarrow \mathbb{S}^1 = \mathbb{R}/\mathbb{Z}$ be the projection onto the second coordinate, plus $1/2$. Then the preimage of 0 is identified with $\mathcal{R} \times \{1/2\} \cong \mathcal{R}$. We perturb $\eta_{1,0}$ outside $\mathcal{R} \times \{1/2\}$ to a map η_1 so that

- the restriction to the 1-cells of \mathcal{R} is affine and nonconstant and sends vertices to rational points, i.e. points of \mathbb{Q}/\mathbb{Z} ,
- η_1 is a local diffeomorphism on each flowline with ψ -derivative in the interval $(.9, 1.1)$, and
- η_1 is smooth on each 2-cell of Y_2 .

We use η_1 to prescribe a preferred orientation on every edge by requiring η_1 to be orientation preserving. The map η_1 determines a rational cellular 1-cocycle $z_1 \in Z^1(Y; \mathbb{Q})$ that is positive on each 1-cell of Y .

Next, let z_2, \dots, z_b be rational cellular cocycles such that $[z_1], \dots, [z_b]$ is a basis for the integral lattice in $H^1(G; \mathbb{R}) = H^1(Y; \mathbb{R})$. For each $z_i, i = 2, \dots, b$, define a map $\eta_i: Y \rightarrow \mathbb{S}^1$ representing z_i so that

- the restriction to each 1-cell is affine, and sends vertices to rational points,
- η_i is ψ -smooth with ψ -derivative bounded above and below, and
- η_i is smooth on each 2-cell of Y_2 .

Now for any integer b -tuple $\mathbf{a} = (a_1, \dots, a_b)$ we have a map

$$\eta_{\mathbf{a}} = \sum_{i=1}^b a_i \eta_i: Y \rightarrow \mathbb{S}^1, \tag{B.3}$$

and this determines a 1-cocycle $z_{\mathbf{a}}$. Let $r > 0$ be such that if $a_1 > r \sum_{i=2}^b |a_i|$, then $z_{\mathbf{a}}$ is positive and the ψ -derivative of $\psi(\eta_{\mathbf{a}}) = \sum_i a_i \psi(\eta_i)$ is bounded below by $a_1/2$. Consequently, $\eta_{\mathbf{a}}$ is flow regular and so

$$\Upsilon_r = \left\{ \left[\sum_{i=1}^b a_i z_i \right] \in H^1(Y; \mathbb{R}) \mid a_i \in \mathbb{R}, a_1 > r \sum_{i=2}^b |a_i| \right\}$$

is a cone in $H^1(Y; \mathbb{R})$ for which every integral element is represented by the flow regular map $\eta_{\mathbf{a}}$ as in (B.3) that is linear on the 1-cells of Y_2 , smooth on the 2-cells of Y_2 , and defines a positive cocycle $z_{\mathbf{a}}$. Furthermore, we take r even larger if necessary so that any integral point in $[z_{\mathbf{a}}] \in \Upsilon_r$ with $\mathbf{a} = (a_1, \dots, a_b) \neq (1, 0, \dots, 0)$ has $a_1 \geq 2$.

Lemma B.4. *For any integral b -tuple \mathbf{a} with $[z_{\mathbf{a}}] \in \Upsilon_r$ and any skew 1-cell e , $\eta_{\mathbf{a}}^{-1}(0) \cap e \subset \mathcal{V}_{\infty}$. (Here we use the identification of \mathcal{R} with $\mathcal{R} \times \{0\}$ to view \mathcal{V}_{∞} inside Y .)*

Proof. The 1-cell e is the image of a 1-cell in the subdivision of \mathcal{R} along \mathcal{V}_1 . Because of our choice of linear structure, the characteristic map $\sigma : [0, 1] \rightarrow e$ has $\sigma(\mathbb{Q} \cap [0, 1]) \subset \mathcal{V}_{\infty}$. Compose this with $\eta_{\mathbf{a}}$ to produce a map $\sigma_{\mathbf{a}} = \eta_{\mathbf{a}} \circ \sigma$, and choose a lift to \mathbb{R} ,

$$\tilde{\sigma}_{\mathbf{a}} : [0, 1] \rightarrow \mathbb{R}.$$

Because each η_i was assumed to be affine with vertices sent to rational points, it follows that

$$\tilde{\sigma}_{\mathbf{a}}(x) = cx + d,$$

where $c, d \in \mathbb{Q}$. Suppose $\eta_{\mathbf{a}}(\sigma(x)) = 0$. Then $cx + d \in \mathbb{Z}$, and hence $x \in \mathbb{Q}$ so that $\sigma(x) \in \mathcal{V}_{\infty}$ as required. \square

B.4. Reparameterized semiflow and first return maps

Fix an integral b -tuple \mathbf{a} with $[z_{\mathbf{a}}] \in \Upsilon_r$. We reparameterize the semiflow ψ as $\psi^{\mathbf{a}}$ so that the composition of a flowline with $\eta_{\mathbf{a}}$ is a local isometry to \mathbb{S}^1 . The first return map of ψ (or equivalently of $\psi^{\mathbf{a}}$) to $\Theta_{\mathbf{a}} := \eta_{\mathbf{a}}^{-1}(0)$ is then the restriction of the time-1 map, $\psi_1^{\mathbf{a}}$, and we denote it by $f_{\mathbf{a}} : \Theta_{\mathbf{a}} \rightarrow \Theta_{\mathbf{a}}$.

Next, let $Z = \Theta_{\mathbf{a}} \cap Y_2^{(1)}$ be the intersection of $\Theta_{\mathbf{a}}$ with the 1-skeleton of Y_2 , which is a finite set of points. From Lemma B.4, every point of Z flows forward into v_0 , and we let $Z_{\mathbf{a}}$ be the intersection of $\Theta_{\mathbf{a}}$ with the flow lines starting at Z . This is also a finite subset of $\Theta_{\mathbf{a}}$, and $f_{\mathbf{a}}(Z_{\mathbf{a}}) \subset Z_{\mathbf{a}}$ by construction.

Proposition B.5. *Given an integral b -tuple \mathbf{a} with $\Theta_{\mathbf{a}}$, $Z_{\mathbf{a}}$, $\psi^{\mathbf{a}}$ and $f_{\mathbf{a}}$ be as above. Then $\Theta_{\mathbf{a}}$ is a graph with vertex set $Z_{\mathbf{a}}$ and $f_{\mathbf{a}} : \Theta_{\mathbf{a}} \rightarrow \Theta_{\mathbf{a}}$ is a regular graph map. Furthermore, $f_{\mathbf{a}}$ is expanding and irreducible.*

Proof. The complement of $Z_{\mathbf{a}}$ in $\Theta_{\mathbf{a}}$ is a collection of arcs, each contained in a 2-cell, transverse to ψ since $f_{\mathbf{a}}$ is flow-regular, thus $\Theta_{\mathbf{a}}$ is a graph. We have already observed that $f_{\mathbf{a}}(Z_{\mathbf{a}}) \subset Z_{\mathbf{a}}$. Next we prove that the restriction of $f_{\mathbf{a}}$ to any edge e of $\Theta_{\mathbf{a}}$ is locally injective. If $\mathbf{a} = (1, 0, \dots, 0)$, then $f_{\mathbf{a}} = f : \mathcal{R} \rightarrow \mathcal{R}$, with \mathcal{R} subdivided at \mathcal{V}_2 , and we are done. So, we assume that $\mathbf{a} \neq (1, 0, \dots, 0)$, and hence $a_2 \geq 2$.

Choose a parameterization $\alpha : (0, 1) \rightarrow e$, and consider the maps

$$H_{\alpha}^{\mathbf{a}} : (0, 1) \times [0, 1] \rightarrow Y \quad \text{and} \quad H_{\alpha} : (0, 1) \times [0, 2) \rightarrow Y$$

given by $H_{\alpha}^{\mathbf{a}}(t, s) = \psi_s^{\mathbf{a}}(\alpha(t))$ and $H_{\alpha}(t, s) = \psi_s(\alpha(t))$. Because the ψ -derivative of $\eta_{\mathbf{a}}$ is at least $a_2/2 \geq 2/2 = 1$, it follows that the reparameterization at any point x is

defined by a monotone function $h_x : [0, \infty) \rightarrow [0, \infty)$, so that $\psi_s^{\mathbf{a}}(x) = \psi_{h_x(s)}(x)$ with $h'_x \leq 1$. In particular, there is an embedding $h_\alpha : (0, 1) \times [0, 1] \rightarrow (0, 1) \times [0, 2)$ such that $H_\alpha^{\mathbf{a}} = H_\alpha \circ h_\alpha$. According to Lemma B.2, H_α is locally injective, and therefore so is $H_\alpha^{\mathbf{a}}$. Consequently, $f_{\mathbf{a}} = \psi_1^{\mathbf{a}}|_{\Theta_{\mathbf{a}}}$ is locally injective on e , and since e was arbitrary, $f_{\mathbf{a}}$ is regular.

The fact that $f_{\mathbf{a}}$ is expanding and irreducible follows, just as in §7.1, from the assumption that $f_1 : \mathcal{R} \rightarrow \mathcal{R}$ is expanding and irreducible: any arc α_0 contained in an edge e_0 of $\Theta_{\mathbf{a}}$ flows forward by ψ onto a nondegenerate arc α_1 contained in an edge e_1 of \mathcal{R} . This in turn eventually flows forward onto all of \mathcal{R} (since f_1 is expanding and irreducible), and consequently can be flowed further forward to hit any point of $\Theta_{\mathbf{a}}$. \square

B.5. Bounding the stretch factor

Note that $\Upsilon_r \subset H^1(Y; \mathbb{R}) = H^1(G; \mathbb{R})$ is an open cone containing u_1 . For any primitive integral b -tuple \mathbf{a} with $[z_{\mathbf{a}}] \in \Upsilon_r$, consider the graph $\Theta_{\mathbf{a}}$ and first return map $f_{\mathbf{a}}$. By construction $\Theta_{\mathbf{a}}$ is a cross section of ψ and is dual to the class $[z_{\mathbf{a}}]$. Since $[z_{\mathbf{a}}]$ is primitive integral, Proposition 5.5 ensures that $\Theta_{\mathbf{a}}$ is connected. Van Kampen’s theorem therefore gives $G = \pi_1(\Theta_{\mathbf{a}})_{*(f_{\mathbf{a}})_*}$ and so just as in the train track case (§7.2) we may write $G = Q_{\mathbf{a}} * \phi_{\mathbf{a}}$, where $\phi_{\mathbf{a}}$ is the descent to the stable quotient $Q_{\mathbf{a}}$ of $(f_{\mathbf{a}})_* : \pi_1(\Theta_{\mathbf{a}}) \rightarrow \pi_1(\Theta_{\mathbf{a}})$. In particular, this proves that $u_{\mathbf{a}} = [z_{\mathbf{a}}] \in \widehat{\mathbb{Q}\Sigma}(G)$ and that $\Lambda(u_{\mathbf{a}}) = \lambda(\phi_{\mathbf{a}})$. Furthermore $\log(\Lambda(u_{\mathbf{a}}))$ is bounded from above by the log of the Perron–Frobenius eigenvalue of the transition matrix of $f_{\mathbf{a}}$, which is the entropy $h(f_{\mathbf{a}})$:

$$\log(\Lambda(u_{\mathbf{a}})) = \log(\lambda((f_{\mathbf{a}})_*)) \leq h(f_{\mathbf{a}}). \tag{B.6}$$

For any integer $k > 0$, $h(f_{k\mathbf{a}}) = (1/k)h(f_{\mathbf{a}})$, so that $\mathbf{a} \mapsto h(f_{\mathbf{a}})$ naturally extends by degree -1 homogeneity to the rational rays, and hence (B.6) holds for all rational points in Υ_r .

Proposition B.7. *For any rational $u = [z_{\mathbf{a}}] \in \Upsilon_r \subset H^1(G; \mathbb{R})$, there is a neighborhood of $U \subset \Upsilon_r$ of u and a constant $R > 0$ so that $\Lambda \leq R$ on all rational points of U .*

Since every point of $\widehat{\mathbb{Q}\Sigma}(G)$ is contained in such an open cone neighborhood Υ_r , this proves Proposition 3.7.

Proof of Proposition B.7. From (B.6), it suffices to prove that the assignment $[z_{\mathbf{a}}] \mapsto h(f_{\mathbf{a}})$ extends continuously (and hence homogeneously with degree -1) to all of Υ_r . For this we can repeat the arguments given in [DKL, Section 8], essentially verbatim. We sketch this here, but refer the reader to [DKL] for the details.

Continuity follows by first passing to the natural extension of (Y, ψ) , which we denote (Y', Ψ) . Here Ψ is an honest flow, and there is a projection $\pi : Y' \rightarrow Y$ which is equivariant with respect to \mathbb{R}_+ , that is, $\pi(\Psi_s(x')) = \psi_s(\pi(x'))$ for all $x' \in Y'$ and $s \geq 0$. The preimage of a section $\Theta_{\mathbf{a}}$ is a section $\Theta'_{\mathbf{a}} = \pi^{-1}(\Theta_{\mathbf{a}})$ of Ψ , and π conjugates the associated first return map $f'_{\mathbf{a}} : \Theta'_{\mathbf{a}} \rightarrow \Theta'_{\mathbf{a}}$ to $f_{\mathbf{a}}$. Moreover, the entropies are equal, $h(f_{\mathbf{a}}) = h(f'_{\mathbf{a}})$, and so it suffices to prove $[z_{\mathbf{a}}] \mapsto h(f'_{\mathbf{a}})$ extends continuously.

For this, we proceed as in [DKL]: Applying the Variational Principal and Abramov's Theorem, we conclude that for any integral class $[z_{\mathbf{a}}] \in \Upsilon_r$,

$$\frac{1}{h(f'_{\mathbf{a}})} = \inf_{\mu} \frac{\mu_{\mathbf{a}}(Y')}{h_{\mu}(\Psi)},$$

where μ ranges over all Ψ -invariant probability measures on Y' , $h_{\mu}(\Psi)$ is the measure-theoretic entropy of Ψ with respect to μ , and $\mu_{\mathbf{a}}$ is an auxiliary measure associated to μ and \mathbf{a} via a measure $\bar{\mu}_0$ on the section $\Theta'_{\mathbf{a}}$; see [DKL] for details. The key property here is that given integral classes $[z_{\mathbf{a}}], [z_{\mathbf{b}}] \in \Upsilon_r$,

$$\mu_{\mathbf{a}} + \mu_{\mathbf{b}} = \mu_{\mathbf{a}+\mathbf{b}}.$$

This follows from the fact that $\eta_{\mathbf{a}} + \eta_{\mathbf{b}} = \eta_{\mathbf{a}+\mathbf{b}}$, as in [DKL].

For each μ , this shows that $[z_{\mathbf{a}}] \mapsto \mu_{\mathbf{a}}(Y')$ extends from the integral points of Υ_r , to a linear function $\mu_*: H^1(G; \mathbb{R}) \rightarrow \mathbb{R}$. Consequently, we obtain a continuous, concave, positive function on all of Υ_r given by

$$W([z_{\mathbf{a}}]) = \inf_{\mu} \frac{\mu_*([z_{\mathbf{a}}])}{h_{\mu}(\Psi)},$$

where the infimum is over all Ψ -invariant probability measures on Y' . Thus $[z_{\mathbf{a}}] \mapsto 1/W([z_{\mathbf{a}}])$ is the desired continuous extension of $[z_{\mathbf{a}}] \mapsto h(f'_{\mathbf{a}})$ to all of Υ_r . \square

Acknowledgments. We are grateful to Nathan Dunfield for many useful discussions regarding the Alexander norm and the BNS-invariant, and in particular for suggesting the approach to proving Theorem I. We would also like to thank Curtis McMullen for his help with clarifying some results from Perron–Frobenius theory related to his earlier work and to the proof of Theorem F, Asaf Hadari for an interesting conversation regarding labeled graphs, and Robert Bieri for his reference to [BR] clarifying the sign convention for the BNS-invariant.

The first author was partially supported by the NSF postdoctoral fellowship, NSF MSPRF no. 1204814. The second author was partially supported by the NSF grant DMS-0904200, DMS-1405146 and by the Simons Foundation Collaboration grant no. 279836. The third author was partially supported by the NSF grant DMS-1207183 and acknowledges support from NSF grants DMS 1107452, 1107263, 1107367 (the “GEAR Network”).

References

- [AKHR] Algom-Kfir, Y., Hironaka, E., Rafi, K.: Digraphs and cycle polynomials for free-by-cyclic groups. *Geom. Topol.* **19**, 1111–1154 (2015) [Zbl 1354.57007](#) [MR 3336279](#)
- [AKR] Algom-Kfir, Y., Rafi, K.: Mapping tori of small dilatation expanding train-track maps. *Topology Appl.* **180**, 44–63 (2015) [Zbl 1309.57017](#) [MR 3293265](#)
- [BH] Bestvina, M., Handel, M.: Train tracks and automorphisms of free groups. *Ann. of Math.* (2) **135**, 1–51 (1992) [Zbl 0757.57004](#) [MR 1147956](#)
- [BNS] Bieri, R., Neumann, W. D., Strebel, R.: A geometric invariant of discrete groups. *Invent. Math.* **90**, 451–477 (1987) [Zbl 0642.57002](#) [Zbl 0642.57002](#)
- [BR] Bieri, R., Renz, B.: Valuations on free resolutions and higher geometric invariants of groups. *Comment. Math. Helv.* **63**, 464–497 (1988) [Zbl 0654.20029](#) [MR 0960770](#)

- [BS] Brinkmann, P., Schleimer, S.: Computing triangulations of mapping tori of surface homeomorphisms. *Experiment. Math.* **10**, 571–581 (2001) [Zbl 1004.57014](#) [MR 1881758](#)
- [Bro] Brown, K. S.: Trees, valuations, and the Bieri–Neumann–Strebel invariant. *Invent. Math.* **90**, 479–504 (1987) [Zbl 0663.20033](#) [MR 0914847](#)
- [Cal] Calegari, D.: *Foliations and the Geometry of 3-manifolds*. Oxford Math. Monogr., Oxford Univ. Press, Oxford (2007) [Zbl 1118.57002](#) [MR 2327361](#)
- [CL] Cashen, C. H., Levitt, G.: Mapping tori of free group automorphisms, and the Bieri–Neumann–Strebel invariant of graphs of groups. *J. Group Theory* **19**, 191–216 (2016) [Zbl 06551804](#) [MR 3466593](#)
- [DV] Dicks, W., Ventura, E.: *The Group Fixed by a Family of Injective Endomorphisms of a Free Group*. *Contemp. Math.* 195, Amer. Math. Soc., Providence, RI (1996) [Zbl 0845.20018](#) [MR 1385923](#)
- [DKL] Dowdall, S., Kapovich, I., Leininger, C. J.: Dynamics on free-by-cyclic-groups. *Geom. Topol.* **19**, 2801–2899 (2015) [Zbl 1364.20026](#) [MR 3416115](#)
- [FGS] Farber, M., Geoghegan, R., Schütz, D.: Closed 1-forms in topology and geometric group theory. *Uspekhi Mat. Nauk* **65**, no. 1, 145–176 (2010) (in Russian) [Zbl 1245.57005](#) [MR 2655245](#)
- [Fri1] Fried, D.: Flow equivalence, hyperbolic systems and a new zeta function for flows. *Comment. Math. Helv.* **57**, 237–259 (1982) [Zbl 0503.58026](#) [MR 0684116](#)
- [Fri2] Fried, D.: The geometry of cross sections to flows. *Topology* **21**, 353–371 (1982) [MR 0670741](#) [Zbl 0594.58041](#)
- [Gau1] Gautero, F.: Dynamical 2-complexes. *Geom. Dedicata* **88**, 283–319 (2001) [Zbl 1001.57006](#) [MR 1877221](#)
- [Gau2] Gautero, F.: Cross sections to semi-flows on 2-complexes. *Ergodic Theory Dynam. Systems* **23**, 143–174 (2003) [Zbl 1140.37311](#) [MR 1971200](#)
- [Gau3] Gautero, F.: Combinatorial mapping-torus, branched surfaces and free group automorphisms. *Ann. Scuola Norm. Sup. Pisa Cl. Sci. (5)* **6**, 405–440 (2007) [Zbl 1173.20017](#) [MR 2370267](#)
- [GMSW] Geoghegan, R., Mihalik, M. L., Sapir, M., Wise, D. T.: Ascending HNN extensions of finitely generated free groups are Hopfian. *Bull. London Math. Soc.* **33**, 292–298 (2001) [Zbl 1028.20020](#) [MR 1817768](#)
- [Had1] Hadari, A.: Personal communication
- [Had2] Hadari, A.: The spectra of polynomial equations with varying exponents. *Topology Appl.* **212**, 105–121 (2016) [Zbl 1355.57020](#) [MR 3553753](#)
- [Kap] Kapovich, I.: A remark on mapping tori of free group endomorphisms. [arXiv:0208189](#) (2002)
- [KM] Kapovich, I., Myasnikov, A.: Stallings foldings and subgroups of free groups. *J. Algebra* **248**, 608–668 (2002) [Zbl 1001.20015](#) [MR 1882114](#)
- [Lan] Lang, S.: *Algebra*. 2nd ed., Addison-Wesley, Reading, MA (1984) [Zbl 0712.00001](#) [MR 0783636](#)
- [MS] McCool, J., Schupp, P. E.: On one relator groups and HNN extensions. *J. Austral. Math. Soc.* **16**, 249–256 (1973) [Zbl 0288.20046](#) [MR 0338186](#)
- [McM1] McMullen, C. T.: Polynomial invariants for fibered 3-manifolds and Teichmüller geodesics for foliations. *Ann. Sci. École Norm. Sup. (4)* **33**, 519–560 (2000) [Zbl 1013.57010](#) [MR 1832823](#)
- [McM2] McMullen, C. T.: Entropy and the clique polynomial. *J. Topol.* **8**, 184–212 (2015) [Zbl 1353.37033](#) [MR 3335252](#)

- [Nor] Northcott, D. G.: *Finite Free Resolutions*. Cambridge Tracts in Math. 71, Cambridge Univ. Press, Cambridge (1976) [Zbl 0328.13010](#) [MR 0460383](#)
- [Oer] Oertel, U.: Affine laminations and their stretch factors. *Pacific J. Math.* **182**, 303–328 (1998) [Zbl 0909.57008](#) [MR 1609595](#)
- [Rey] Reynolds, P.: Dynamics of irreducible endomorphisms of F_n . [arXiv:1008.3659](#) (2010)
- [Sta] Stallings, J. R.: Topology of finite graphs. *Invent. Math.* **71**, 551–565 (1983) [Zbl 0521.20013](#) [MR 0695906](#)
- [Thu] Thurston, W. P.: A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.* **59**, no. 339, i–vi and 99–130 (1986) [Zbl 0585.57006](#) [MR 0823443](#)
- [Wan] Wang, Z.: Mapping tori of outer automorphisms of free groups. PhD thesis, Rutgers Univ. (2002) [MR 2703150](#)