Geometry & Topology 24 (2020) 1717–1750

Homological eigenvalues of lifts of pseudo-Anosov mapping classes to finite covers

ASAF HADARI

Let Σ be a compact orientable surface of finite type with at least one boundary component. Let $f \in \mathrm{Mod}(\Sigma)$ be a pseudo-Anosov mapping class. We prove a conjecture of McMullen by showing that there exists a finite cover $\widetilde{\Sigma} \to \Sigma$ and a lift $\widetilde{f}$ of f such that $\widetilde{f}_*\colon H_1(\widetilde{\Sigma}, \mathbb{Z}) \to H_1(\widetilde{\Sigma}, \mathbb{Z})$ has an eigenvalue off the unit circle.

20C12, 57M05, 57M60

1 Introduction

Let Σ be a compact orientable surface and let $\mathrm{Mod}(\Sigma)$ be its mapping class group — the group of isotopy classes of orientation-preserving diffeomorphisms from Σ to itself that fix the boundary pointwise. The finite-dimensional representation theory of $\mathrm{Mod}(\Sigma)$ is a nascent field of study. These groups have extensive collections of finite-dimensional representations, but many basic questions remain mysterious. For instance, for most Σ it is not known whether or not $\mathrm{Mod}(\Sigma)$ is a linear group.

The largest known collection of representations are the *homological representations* which are associated to finite covers $\pi\colon \Sigma' \to \Sigma$. The first of these, associated to the trivial cover, is the standard homological representation $\mathrm{Mod}(\Sigma) \to \mathrm{GL}(H_1(\Sigma, \mathbb{Q}))$ given by the induced action on first homology. The kernel of this representation is called the *Torelli group*.

More generally, fix a point $* \in \Sigma$. Let $\mathrm{Mod}(\Sigma, *)$ be the group of orientation-preserving self-diffeomorphisms of the pair $(\Sigma, *)$ that fix the boundary pointwise, mod isotopies that fix $*$ and the boundary pointwise. The group $\mathrm{Mod}(\Sigma, *)$ acts on $\pi_1(\Sigma, *)$ by automorphisms.

Let $K < \pi_1(\Sigma, *)$ be a finite-index subgroup, and let $\pi\colon \Sigma' \to \Sigma$ be the associated finite cover. Let $G_K = \{f \in \mathrm{Mod}(\Sigma, *) : f(K) = K\}$. The group G_K is a finite-index subgroup of $\mathrm{Mod}(\Sigma,*)$. We have a natural map $\rho_K\colon G_K \to \mathrm{GL}(H_1(K;\mathbb{Z}))$. Topologically, every element $f \in G_K$ can be lifted to a diffeomorphism $f'\colon \Sigma' \to \Sigma'$. The diffeomorphism f' induces a map $f'_*\colon H_1(\Sigma';\mathbb{Q}) \to H_1(\Sigma';\mathbb{Q})$. The transformation f'_* is $\rho_K(f')$.

The representations ρ_K are called homological representations. They have been studied extensively by many authors. For example, Grunewald, Larsen, Lubotzky and Malestein [6] use these representation to construct several different infinite families of arithmetic quotients of mapping class groups. In [16], Putman and Wieland exhibit a connection between properties of homological representations and the virtual first Betti number of $\mathrm{Mod}(\Sigma)$. In work of Lubotzky and Meiri [12; 13], and separately in work of Malestein and Souto [14] these representation were used to describe generic properties of random elements of $\mathrm{Mod}(\Sigma)$.

In addition to providing information about the group $\mathrm{Mod}(\Sigma)$ as a whole, these representations also provide in formation about individual elements. For example, Koberda [9] and later Koberda and Mangahas [11] showed that the family of homological representations can detect the Nielsen–Thurston classification of a mapping class.

It is natural to try to understand whether or not the topological invariants associated to a mapping class f can be recovered from its homological representations. McMullen studied this question for pseudo-Anosov mapping classes by considering the following invariant. Fix $f \in \mathrm{Mod}(\Sigma,*)$, a pseudo-Anosov mapping class. Given a finite-index subgroup $K < \pi_1(\Sigma,*)$, let σ_K be the spectral radius of the operator $\rho_K(f)$ (that is, the modulus of its largest eigenvalue). It is a simple exercise to show that σ_K is at least 1 and at most λ, the dilatation of f. In [15], McMullen shows that if the invariant foliations of f have a singularity with an odd number of prongs, then $\sup \sigma_K < \lambda$, where the supremum is taken over all finite-index subgroups K. McMullen asked the following question, whose positive resolution has become a well-known conjecture:

Question 1.1 In the notation above, is $\sup \sigma_K > 1$?

In a previous paper [7], we provided evidence for this conjecture by proving the following:

Theorem 1.2 (Hadari) *Suppose that Σ has at least one boundary component. Then, for any infinite-order element $f \in \mathrm{Mod}(\Sigma)$, there is a finite cover $\pi\colon \Sigma' \to \Sigma$ to which f lifts to a map f' such that $f'_*\colon H_1(\Sigma';\mathbb{Q}) \to H_1(\Sigma';\mathbb{Q})$ has infinite order.*

In this paper, we use a strategy inspired by the proof in [7] to provide the following answer to McMullen's question:

Theorem 1.3 *Suppose that Σ has at least one boundary component. Then, for any $f \in \mathrm{Mod}(\Sigma)$ with positive topological entropy, there exists a regular finite cover $\pi\colon \Sigma' \to \Sigma$ to which f lifts to a map f' such that $f'_*\colon H_1(\Sigma';\mathbb{Q}) \to H_1(\Sigma';\mathbb{Q})$ has eigenvalues off of the unit circle. Furthermore, if f is pseudo-Anosov, this cover can be taken to have a solvable deck group.*

We also provide an analogous result for automorphisms of free groups.

Theorem 1.4 *Let $n \geq 2$ and let $\overline{f} \in \mathrm{Out}(F_n)$ be a fully irreducible automorphism. Then there exists a representative f of $\overline{f}$ and finite-index subgroup $K \lhd F_n$ such that $f(K) = K$, and $f_*\colon H_1(K;\mathbb{Q}) \to H_1(K;\mathbb{Q})$ has eigenvalues off of the unit circle. The subgroup K can be taken such that F_n/K is solvable.*

Note that for any surface Σ, there is a natural map $\mathrm{Mod}(\Sigma) \to \mathrm{Out}(\pi_1(\Sigma))$. In our statement of Theorem 1.3, we restrict ourselves to surfaces with boundary. These surfaces have free fundamental groups. Given a pseudo-Anosov mapping class on a surface with boundary, the surface Σ is homotopy equivalent to an invariant train track graph $\Gamma \subset \Sigma$, and the map $\overline{f}$ is induced by a continuous function $\varphi\colon \Gamma \to \Gamma$. The map φ fixes some point $* \in \Gamma$. Taking this point to be our basepoint, we get a representative of $\overline{f}$, called a *train track representative*. A similar notion exists for fully irreducible automorphisms (see Fathi, Laudenback and Poénaru [4] for surfaces, and Bestvina and Handel [2] for $\mathrm{Out}(F_n)$). We will deduce Theorems 1.3 and 1.4 from the following theorem, whose proof will take up the majority of this paper:

Theorem 1.5 *Let $n \geq 2$ and let $\overline{f} \in \mathrm{Out}(F_n)$ be either a fully irreducible automorphism or the image of a pseudo-Anosov mapping class. Let f be a train track representative of $\overline{f}$. Then there exists a finite-index subgroup $K \lhd F_n$ such that $f(K) = K$ and $f_*\colon H_1(K;\mathbb{Z}) \to H_1(K;\mathbb{Z})$ has eigenvalues off of the unit circle. Furthermore, we can choose K such that F_n/K is solvable.*

Remark 1.6 As we were concluding the writing of this paper, Yi Liu also published an independent proof of McMullen's conjecture. Like the proof in this paper, his proof is to some extent inspired by our proof in [7] but aside from this initial inspiration the two proofs are very different. The end results are also somewhat different. Liu's proof covers the case of closed surfaces, which the proof in this paper does not (our proof fails for closed surfaces in exactly one spot, Lemma 4.5). The proof in this paper covers the $\mathrm{Out}(F_n)$ case, which Liu's does not, and provides the extra information that the finite cover can be taken to be solvable.

1.1 Strategy and organization of the proof

Theorem 1.5 is nontrivial only when all of the eigenvalues of f_* are roots of unity. By replacing f with a power of itself, we can assume that all of its eigenvalues are 1, and in particular it has a 1–eigenspace.

To such an automorphism we introduce a matrix A_f, which we call the *equivariant Magnus matrix*. It is related to the Magnus representation of f (see Sakasai [17] and Suzuki [18] for definitions). The entries of this matrix are polynomials, which we view as elements of the group ring of some quotient H_f of $H_1(F_n, \mathbb{Z})$.

Given a matrix whose entries are polynomials in the variables $X_1^{\pm 1}, \dots, X_m^{\pm 1}$, we can substitute numbers $\xi_1, \dots, \xi_m$ for $X_1, \dots, X_m$ to get a matrix with entries in $\mathbb{C}$. This is called the *specialization of the matrix at* $\xi_1, \dots, \xi_m$. The equivariant Magnus matrix has the property that its specialization at roots of unity contain information about $\rho_K(f)$ for a certain collection of abelian covers K. One particular connection is that if we specialize A_f at roots of unity and get a matrix that has eigenvalues off of the unit circle, then $\rho_K(f)$ has eigenvalues off of the unit circle for some abelian cover K.

We now face the question of having to tell when a matrix with polynomial coefficients has a specialization at roots of unity with eigenvalues off of the unit circle. One possible approach is to look at the trace of such a matrix. If the trace of the matrix is in some sense large (say if the L^2 norm of its coefficients is greater than the dimension of the matrix) then it is possible to use the Fourier transform on abelian groups to find a specialization as required.

In Section 2 we introduce the equivariant Magnus matrix, discuss the connection between its specializations and the homological representations of f, and provide two criteria for finding a specialization that has eigenvalues off of the unit circle. The remainder of the proof shows how to find a sequence of covers where the trace of the equivariant Magnus matrices becomes larger and larger in the sense of Section 2.

In Section 3 we introduce a combinatorial object, called the *transition graph*, which encodes a great deal of information about f. In particular, when f is pseudo-Anosov we can associate to this transition graph a convex polygon $\mathcal{S}^e\varphi \subset H_f \otimes \mathbb{R}$, which we call the *equivariant shadow*, and which we use extensively in our proof. This polygon is related to the norm ball of the Thurston norm (this is explained in Lemma 3.19). When f is pseudo-Anosov, we can calculate the dimension of this polygon, and hence get a lower bound on the number of its vertices.

Morally speaking, we expect the convex hull of the support of $\operatorname{Tr} A_f$ to be some homothetic image of the polygon $\mathcal{S}^e\varphi$, and, when this is the case, $\operatorname{Tr} A_f$ is large in the sense of Section 2 due to our estimate of the number of vertices. When this is not the case, it is due to some cancellation occurring at the vertices of this polygon. In Section 4 we discuss this cancellation, and show that for every vertex it is possible to find a nilpotent cover where it does not cancel. Finally, in Section 5 we collect several important technical lemmas, and complete the proof of our Theorems 1.3, 1.4 and 1.5.

2 The Magnus matrix and its specializations

In this section we introduce a key concept, the equivariant Magnus matrix; and two lemmas, the anchoring lemma and the L^2–trace lemma, which will be our central tools in proving Theorem 1.5.

Throughout the section, let $f \in \operatorname{Aut}(F_n)$ be a train track representative of a pseudo-Ansov mapping class or a fully irreducible automorphism. Suppose that f is induced by the continuous map $\varphi\colon \Gamma \to \Gamma$, fixing the basepoint $*$.

2.1 The f–equivariant torsion-free universal abelian cover

Let $G = F_n \rtimes_f \mathbb{Z}$. Consider the endomorphism $i\colon F_n \to G$ given by $i(w) = (w, 0)$.

Definition 2.1 Let $h\colon G \to H_1(G;\mathbb{Q})$ be the natural map, and let $U_f = h \circ i$. Let $\widetilde{\Gamma}_f$ be the cover corresponding to U_f. We call $\widetilde{\Gamma}_f$ the *f–equivariant torsion-free universal abelian cover of* Γ.

We begin by making several simple observations. Since $[F_n, F_n] < U_f$, we have that $H_f = F_n/U_f$ is a finitely generated abelian group. Since $[G, G] \lhd G$, the definition of semidirect products gives that $f(U_f) = U_f$. Thus f acts on the group H_f.

There is an f–equivariant isomorphism between H_f and the image of F_n in the group $H_1(G;\mathbb{Q}) \cong G_{\mathrm{ab}} \otimes \mathbb{Q}$, where the action of f on G is given by conjugation. Since conjugation acts trivially on G_{ab}, we get that f acts trivially on H_f. Let f_* denote the automorphism induced by f on H $= H_1(F_n,\mathbb{Z})$. The group H_f is an abelian quotient of F_n and is thus a quotient of H. Indeed, we have the natural identification $H_f \cong \operatorname{coker}(I_n - f_*)$. Note that since $G_{\mathrm{ab}}\otimes\mathbb{Q}$ is torsion-free, so is H_f.

2.2 The equivariant Magnus matrix of f

Let $V_f = C_1(\widetilde{\Gamma}_f,\mathbb{C})$ be the space of simplicial 1–chains with coefficients in $\mathbb{C}$ in the cover $\widetilde{\Gamma}_f$.

The group H_f acts on $\widetilde{\Gamma}_f$ by deck transformations and thus permutes the edges of $\widetilde{\Gamma}_f$. This gives V_f the structure of an H_f–module.

Pick a spanning tree T of Γ and let $\widetilde{T}$ be a lift of T to a tree in $\widetilde{\Gamma}_f$. For any vertex v of Γ, let $\tilde{v}$ be the lift of v incident at $\widetilde{T}$. The action of the group H_f on $\widetilde{\Gamma}_f$ by deck transformations gives a transitive permutation on the set of all preimages of v. By identifying $\tilde{v}$ with the element $0\in H_f$, we can identify every preimage of v with an element of H_f.

Given any oriented edge η of Γ, the choice of the lift $\widetilde{T}$ gives a bijection ι_η between the collection of lifts of η in $\widetilde{\Gamma}_f$ and H_f given by reading off the label of the origin vertex of a lift. This identification induces an H_f–module isomorphism

$$(\mathbb{C}[H_f])^{E(\Gamma)} \cong V_f.$$

If we set $m = \#E(\Gamma)$ then the above isomorphism is given by

$$\Bigl(\sum a_{1,i}h_{1,i},\dots,\sum a_{m,i}h_{m,i}\Bigr) \to \sum_{j=1}^{m}\sum a_{j,i}\iota_{\eta_i}(h_{j,i}),$$

where $a_{j,i}\in\mathbb{C}$ and $h_{j,i}\in H_f$.

Let φ_f be the lift of φ to $\widetilde{\Gamma}_f$ that fixes $\tilde{*}$. Since $\tilde{\varphi}$ maps edges to edge-paths in $\widetilde{\Gamma}_f$, it induces a map $\tilde{\varphi}_*\colon V_f\to V_f$. We call this map the *equivariant Magnus representation of f on Γ*.

Because f acts trivially on H_f, we get that $\tilde{\varphi}_*$ commutes with the action of H_f on V_f, and thus induces an H_f–module homomorphism $V_f\to V_f$.

Under the identification $V_f\cong(\mathbb{C}[H_f])^{E(\Gamma)}$, this homomorphism is given by multiplication by an $m\times m$ matrix $A_f\in M_m(\mathbb{C}[H_f])$, where $m=\#E(\Gamma)$. We call the matrix A_f the *equivariant Magnus matrix of f on Γ*.

2.3 Specializations of A_f and abelian covers

Write $H_f \cong \mathbb{Z}^d$. Viewing $\mathbb{Z}^d$ as a multiplicative group, we can write $\mathbb{C}[\mathbb{Z}^d] \cong \mathbb{C}[X_1^{\pm 1}, \dots, X_d^{\pm 1}]$.

Definition 2.2 Let $\xi\colon \mathbb{Z}^d \to \mathbb{C}^\times$ be a homomorphism. Write $\xi_i = \xi(X_i)$. Let $t \in \mathbb{C}[\mathbb{Z}^d]$. Using the identification $\mathbb{C}[\mathbb{Z}^d] \cong \mathbb{C}[X_1^{\pm 1}, \dots, X_d^{\pm 1}]$, we can view t as a rational function in the variables $X_1, \dots, X_d$. By plugging in the number ξ_i for the variable X_i, we get a number $t(\xi) \in \mathbb{C}$, which we call the *specialization of t at ξ*.

Definition 2.3 If $A \in M_k(\mathbb{C}[\mathbb{Z}^d])$ for some k and $\xi\colon \mathbb{Z}^d \to \mathbb{C}$ is a homomorphism, then we can define the *specialization of A at ξ* to be the $k \times k$ matrix whose (i,j)–coordinate is $A_{i,j}(\xi)$.

The space V_f is the space of formal linear combinations of edges in $\widetilde{\Gamma}_f$ with finite support and coefficients in $\mathbb{C}$. Let $W_f = \mathbb{Z}^{E(\widetilde{\Gamma}_f)}$, the space of (not necessarily finitely supported) formal linear combinations of edges in $\widetilde{\Gamma}_f$. The deck group H_f acts on the edges of $\widetilde{\Gamma}_f$ by permutations. This action gives W_f an H_f–module structure.

Definition 2.4 Let $\xi\colon H_f \to \mathbb{C}^\times$ be a homomorphism. Define

$$W_{f,\xi} = \{t \in W_f : h \cdot t = \xi(h)t \text{ for all } h \in H_f\}.$$

Notice that for every ξ, the space $W_{f,\xi}$ is an $m = \#E(\Gamma)$–dimensional vector space. Indeed, let $\{\eta_1, \dots, \eta_m\}$ be the collection of edges of Γ. Fix an arbitrary collection $\{\widetilde{\eta}, \dots, \widetilde{\eta}_m\}$ of lifts of the edges of Γ to $\widetilde{\Gamma}_f$. Any edge in $\widetilde{\Gamma}_f$ is the image of one of the $\widetilde{\eta}_i$ under a deck transformation. Given any $t \in W_{f,\xi}$ and an edge ζ of $\widetilde{\Gamma}_f$ such that $\zeta = h \cdot \widetilde{\eta}_i$, the coefficient of ζ in t is $\xi(h)$ times the coefficient of $\widetilde{\eta}_i$. Thus, we have an obvious identification $W_{f,\xi} = \mathbb{C}^{\{\widetilde{\eta}_1, \dots, \widetilde{\eta}_m\}}$.

Since the homomorphism $\widetilde{\varphi}_*$ is an H_f–module homomorphism, it acts on every space $W_{f,\xi}$ as an H_f–module homomorphism, which we denote by $A_{f,\xi}$.

If $\{\widetilde{\eta}_1, \dots, \widetilde{\eta}_m\}$ is the set described above, then every element of $W_{f,\xi}$ can be written as

$$t = \sum_{i=1}^{m} \sum_{h \in H_f} a_i \xi(h)(h \cdot \widetilde{\eta}_i).$$

We can then define

$$A_{f,\xi}(t) = \sum_{i=1}^{m} \sum_{h \in H_f} a_i \xi(h)(h \cdot A_f \widetilde{\eta}_i)$$

Lemma 2.5 *Under the identification* $W_{f,\xi} = \mathbb{C}^{\{\tilde{\eta}_1,\dots,\tilde{\eta}_m\}}$*, the matrix corresponding to the linear transformation* $A_{f,\xi}$ *is the specialization* $A_f(\xi)$*.*

Proof Write $A_f \tilde{\eta}_i = \sum_j w_{i,j} \tilde{\eta}_j$, with $w_{i,j} \in \mathbb{C}[H_f]$. Then

$$A_f \cdot \sum_{h \in H_f} \xi(h)(h \cdot \tilde{\eta}_i) = \sum_{h \in H_f} \xi(h) \sum_j (h \cdot w_{i,jj}) \tilde{\eta}_j.$$

Switching the order of the summands and using the fact that $h \cdot t = \xi(h) t$ for any $t \in W_{f,\xi}$ and $h \in H_f$ now gives the result. □

Now suppose $\xi \colon H_f \to \mathbb{C}^\times$ has finite image. Let $k = |\xi(H_f)|$ be the size of the image group. Let $\Gamma_k \to \Gamma$ be the cover corresponding to the kernel of the homomorphism $F_n \to H_f / k H_f$ given by reduction mod k.

Let $t \in W_{f,\xi}$, and let η be an edge in Γ_k. Given any two lifts $\tilde{\eta}_1$ and $\tilde{\eta}_2$ of the edge η to $\tilde{\Gamma}_f$, the coefficients of $\tilde{\eta}_1$ and $\tilde{\eta}_2$ in t are the same. Call this number a_η. Write $\overline{t} = \sum_\eta a_\eta \eta \in C_1(\Gamma_k, \mathbb{C})$.

The action of $H_f / k H_f$ on Γ_k by deck transformations induces an H_f–module structure on $C_1(\Gamma_k, \mathbb{C})$. The map $t \to \overline{t}$ is an H_f–module isomorphism. Call its image $\overline{W}_{f,\xi}$.

The map φ lifts to a map φ_k of Γ_k. Since this map is H_f–equivariant, it fixes the space $\overline{W}_{f,\xi}$. Because the map $t \to \overline{t}$ is an isomorphism, the matrix giving the induced action $(\varphi_k)_*$ on this space is $A_f(\xi)$.

2.4 The anchoring lemma and the L^2–trace lemma

Definition 2.6 Let L be a lattice in $\mathbb{Z}^d$. Let $t = \sum a_i h_i \in \mathbb{C}[\mathbb{Z}^d]$, where $a_i \in \mathbb{C}$ and $h_i \in \mathbb{Z}^d$. Define

$$t(L) = \sum_{h_i \in L} a_i.$$

Definition 2.7 Led $B \in M_d(\mathbb{C}[\mathbb{Z}^d])$. Let $t_k = \operatorname{Tr}[B^k]$. Write $t_k = \sum_{h \in \mathbb{Z}^d} a(k,h) h$, where $a(k,h) \in \mathbb{C}$ is the coefficient of h. We say that B is *anchored* if there is some lattice L and some integer k such that

$$t_k(L) = \left| \sum_{h \in L} a(k,h) \right| > d.$$

Definition 2.8 We say that f is *anchored in* Γ if A_f is anchored.

The following lemma relates specializations to lattices:

Lemma 2.9 *Let $t \in \mathbb{C}[\mathbb{Z}^d]$, $t = \sum_{h \in \mathbb{Z}^d} a_h h$, and let L be a lattice. As in Definition 2.6, let $t(L) = \sum_{h \in L} a_h$. Let N_L be the set of all $\xi\colon \mathbb{Z}^d \to \mathbb{C}^\times$ such that $\xi|_L = 1$. Then*

$$t(L) = \frac{1}{|N_L|} \sum_{\xi \in N_L} t(\xi).$$

Proof Since the functions $t \to t(L)$ and $t \to t(\xi)$ are linear in t, it's enough to prove the lemma for the case where t is a monomial. Suppose $t = ah$, with $a \in \mathbb{C}$ and $h \in \mathbb{Z}^d$. Let $\bar{h}$ be the image of h in the finite abelian group $G = \mathbb{Z}^d / L$.

By definition, $\sum_{\xi \in N_L} t(\xi) = a \sum_{\chi \in G^\times} \chi(\bar{h})$, where $G^\times$ is the group of characters of $G^\times$. Denote the trace of the regular representation of G by ρ_G. Since G is abelian, $\rho_G = \sum_{\chi \in G^\times} \chi$. Thus,

$$\sum_{\xi \in N_L} t(\xi) = a \rho_G(\bar{h}).$$

The left-hand side is equal to 0 if $\bar{h} \neq e$ and $a \cdot |G| = a \cdot |G^\times| = a \cdot |N_L|$ if $\bar{h} = e$. Since $\bar{h} = e$ if and only if $h \in L$, this concludes the proof. □

Lemma 2.10 *Let $\varphi_k^\#\colon C_1(\Gamma_k, \mathbb{C}) \to C_1(\Gamma_k, \mathbb{C})$ be the map induced by φ_k. Suppose this map has an eigenvalue with absolute value > 1. Then the same is true for the induced map $\varphi_k^*\colon H_1(\Gamma_k, \mathbb{C}) \to H_1(\Gamma_k, \mathbb{C})$.*

Proof let $U = C_1(\Gamma_k, \mathbb{C})$, and let $W \subset U$ be the subspace spanned by all closed paths in Γ_k. We have a natural identification $W \cong H_1(\Gamma_k, \mathbb{C})$. The space U is spanned by elements of the form e, where e ranges over all edges of Γ_k. Pick any norm $\|\cdot\|$ on U, and let $\lambda > 1$ be the spectral radius of the action of φ_k on U. Then there exists an edge e such that

$$\limsup_{j \to \infty} \frac{1}{j} \log \|\varphi^j(e)\| = \log \lambda.$$

Let L_λ be the direct sum of all generalized eigenspaces corresponding to eigenvalues with absolute value λ. Setting $v_j = \varphi^j(e)/\|\varphi^j(e)\|$, we have that the distance from v_j to L_λ goes to 0 as $j \to \infty$. Note that $\varphi^j(e)$ corresponds to a path in Γ, and any such path can be closed to a loop using a bounded number of edges. Thus, the distance of v_j from W goes to 0 as $j \to \infty$. Therefore $L_\lambda \cap W \neq \{0\}$. Since this space is φ_k–invariant, it contains an eigenvector with eigenvalue of absolute value $\lambda, \dots$. □

Lemma 2.11 (the anchoring lemma) *If f is anchored in Γ then there exists an abelian cover $\Gamma_k \to \Gamma$ to which φ lifts to a map φ_k such that $(\varphi_k)_*\colon H_1(\Gamma_k, \mathbb{C}) \to H_1(\Gamma_k, \mathbb{C})$ has eigenvalues off the unit circle.*

Proof Let i be an integer and L a lattice such that $\mathrm{Tr}(A_f^i)(L) > m$. Write $t_i = \mathrm{Tr}(A_f^i)$. By Lemma 2.9, $t_i(L) = (1/|N_L|)\sum_{\xi\in N_L} t(\xi) > m$. Since $(1/|N_L|)\sum_{\xi\in N_L} t(\xi)$ is an average, there exists a $\xi \in N_L$ such that $|t_i(\xi)| > m$.

By definition of N_L, $\xi(H_f)$ is finite (since $L < \ker\xi$). Let $\Gamma_k \to \Gamma$ be the cover constructed above. The space $\overline{W}_{f,\xi}$ is an m–dimensional, φ_k–invariant subspace of $C_1(\Gamma_k, \mathbb{C})$. Furthermore, we have that $|t_i(\xi)| = |\mathrm{Tr}(A_{f,\xi})| > m$. Thus, the map induced by φ_k on $C_1(\Gamma_k, \mathbb{C})$ has eigenvalues off the unit circle. The proof now follows from Lemma 2.10. □

Given $t = \sum a_i h_i \in \mathbb{C}[\mathbb{Z}^d]$, write $\|t\|_2 = \sqrt{\sum_i a_i^2}$.

Lemma 2.12 (the L^2–trace lemma) *If there exists an i such that $\|\mathrm{Tr}(A_f^i)\|_2 > m$ then there exists an abelian cover $\Gamma_k \to \Gamma$ to which φ lifts to a map φ_k such that $(\varphi_k)_*\colon H_1(\Gamma_k, \mathbb{C}) \to H_1(\Gamma_k, \mathbb{C})$ has eigenvalues off the unit circle.*

Proof Let i be the number given in the statement of the theorem, and let $t = \mathrm{Tr}(A_f^i)$. For any homomorphism $\xi\colon H_f \to \mathbb{C}^\times$, we defined the specialization of t at ξ, $t(\xi)$. Let H_f^* be the set of all ξ such that $|\xi(h)| = 1$ for all $h \in H_f$. The function $\hat{t}\colon H_f^* \to \mathbb{C}$ given by $\xi \to t(\xi)$ is called the *Fourier transform of t*.

Setting $H_f \cong \mathbb{Z}^d$, and considering $\mathbb{Z}^d$ as a multiplicative group in $X_1^{\pm 1}, \dots, X_d^{\pm 1}$, we think of t as a rational function in $X_1, \dots, X_d$ and ξ as a d–tuple $(\xi_1, \dots, \xi_d)$ all of whose coordinates have modulus 1. The Fourier transform $\hat{t}$ is the function that plugs in the d–tuple $(\xi_1, \dots, \xi_d)$ into the rational function t. Thus, $\hat{t}$ is a continuous function from the torus $\mathbb{T}^d \cong H_f^*$ to $\mathbb{C}$.

By Plancherel's theorem, $\|t\|_2 = \|\hat{t}\|_2$ where the right-hand norm is the norm in $L^2(\mathbb{T}^d)$ measured using the Haar measure of the torus. By our assumption, we have that $\|\hat{t}\|_2 > m$. Hence, there is a point $\xi \in H_f^*$ such that $|t(\xi)| > m$. Since $\hat{t}$ is a continuous function, this point can be taken to have coordinates that are all roots of unity.

Since all the coordinates of ξ are roots of unity, we have that $\xi(H_f)$ is a finite set. Let $\Gamma_k \to \Gamma$ be the cover associated to ξ, as defined in the end of Section 2.3. We now proceed in an identical manner to the proof of Lemma 2.11.

The space $\overline{W}_{f,\xi}$ is an m–dimensional, φ_k–invariant subspace of $C_1(\Gamma_k, \mathbb{C})$. Furthermore, we have that $|t_i(\xi)| = |\mathrm{Tr}(A_{f,\xi})| > m$. Thus, the map induced by φ_k on $C_1(\Gamma_k, \mathbb{C})$ has eigenvalues off the unit circle. The proof now follows once again from Lemma 2.10. □

3 The transition graph of φ

3.1 The transition graph and associated objects

The transition graph is a technical gadget that we use to encode information about the map φ.

Definition 3.1 Let $E(\Gamma) = \{e_1, \dots, e_m\}$ be the edge set of Γ. Pick, once and for all, an orientation on each edge of Γ. Construct a directed graph $\mathcal{T} = \mathcal{T}[\Gamma, \varphi]$, called the *transition graph of* φ, in the following way. The vertex set of $\mathcal{T}$ is $\{e_1, \dots, e_m\}$. Connect the vertex e_i to the vertex e_j with $e(i, j)$ directed edges, where $e(i, j)$ is the number of times $\varphi(e_i)$ traverses the edge e_j (in either direction).

We will associate several useful objects to the graph $\mathcal{T}$ which we discuss in this section.

Definition 3.2 Pick a *decorating function* $\boldsymbol{d}\colon E(\mathcal{T}) \to \mathbb{N}$ such that whenever $E_{i,j}$ is the set of edges in $\mathcal{T}$ from e_i to e_j, then $\boldsymbol{d}|_{E_{i,j}}$ is a bijection onto the set $\{1, \dots, e(i, j)\}$. We think of each of the edges η connecting e_i to e_j as corresponding to the $\boldsymbol{d}(\eta)^{\text{th}}$ time that $\varphi(e_i)$ traverses e_j.

Definition 3.3 Extend the decorating function to a function $\boldsymbol{d}\colon \mathcal{P}(\mathcal{T}) \to \mathbb{N}$, where $\mathcal{P}(\mathcal{T})$ is the set of edge paths in $\mathcal{T}$, by requiring that $\boldsymbol{d}$ restricted to the set of all paths connecting e_i to e_j of length k is a bijection onto $\{1, \dots, e(i, j; k)\}$, where $e(i, j; k)$ is the number of times $\varphi^k(e_i)$ traverses e_j, in either direction. We think of a path $\eta_1 \dots \eta_k$ connecting e_i to e_j as corresponding to the $\boldsymbol{d}(\eta_1 \dots \eta_k)^{\text{th}}$ time that $\varphi^k(e_i)$ traverses e_j.

Definition 3.4 Let $\mathcal{P}(\Gamma)$ be the set of paths in Γ. Define a *path function* $\boldsymbol{p}\colon \mathcal{P}(\mathcal{T}) \to \mathcal{P}(\Gamma)$ in the following way. Let $\eta_1 \dots \eta_k$ be a path in $\mathcal{T}$ connecting e_1 to e_j. We write $\varphi^k(e_i) = abc$, where $a, b, c \in \mathcal{P}(\Gamma)$ and b the path of length 1 traversing e_j that corresponds to $\boldsymbol{d}(\eta_1 \dots \eta_k)$. We define $\boldsymbol{p}(\eta_1 \dots \eta_k) = a$ if b traverses e_j in the positive direction and $\boldsymbol{p}(\eta_1 \dots \eta_k) = ab$ if b traverses e_j in the negative direction. Note that this convention assures that the endpoint of $\boldsymbol{p}(\eta_1 \dots \eta_k)$ is the initial point of e_j.

Definition 3.5 Let $\pi\colon \Gamma_0 \to \Gamma$ be a regular cover to which we can lift φ to a map φ_0. Denote the deck group of this cover by $\mathcal{D}$. Let $V = V(\Gamma)$. Choose a lift V_0 of the set V to Γ_0. Every vertex $w \in V(\Gamma_0)$ satisfies $w = \sigma(v)$ for some $v \in V_0$ and $\sigma \in \mathcal{D}$. We say that σ is the *address* of w, and write $\sigma = \boldsymbol{a}(w)$.

Definition 3.6 Given a regular cover $\pi\colon \Gamma_0 \to \Gamma$ as above and lift V_0 of V, define a function $\boldsymbol{t}_\pi\colon \mathcal{P}(\mathcal{T}) \to \mathcal{D}$, called a *translation function*, by setting $\boldsymbol{t}_\pi(\eta_1 \dots \eta_k) = \boldsymbol{a}(w_2)\boldsymbol{a}(w_1)^{-1}$, where w_1 and w_2 are respectively the initial and terminal points of $\boldsymbol{p}(\eta_1 \dots \eta_k)$. We will most often be concerned with the translation function for the cover $\pi\colon \widetilde{\Gamma}_f \to \Gamma$. In this case, we will omit the subscript π.

Definition 3.7 Define a *sign function* $\boldsymbol{s}\colon E(\Gamma) \to \{\pm 1\}$ in the following way. Let η be an edge connecting e_i to e_j. Set $\boldsymbol{s}(\eta) = 1$ if and only if the $\boldsymbol{d}(\eta)^{\text{th}}$ time $\varphi(e_i)$ traverses e_j is in the positive direction. We extend the definition of $\boldsymbol{s}$ to edge paths in $\mathcal{T}$ by setting $\boldsymbol{s}(\eta_1 \dots \eta_k) = \boldsymbol{s}(\eta_1) \dots \boldsymbol{s}(\eta_k)$.

Example 3.8 Suppose $\Gamma = S^1 \vee S^1$ is the rose on two petals. The fundamental group of this graph is $F_2 = \langle a, b \rangle$, where a and b are loops about the two petals based at their intersection. Consider the inner automorphism $a \to bab^{-1}$, $b \to b$. This automorphism is induced by a function $\varphi\colon \Gamma \to \Gamma$ which is the identity on the loop b and sends a to bab^{-1}. The transition graph $\mathcal{T}$ has two vertices: one labeled a and one labeled b. Identify $H_1(\Gamma; \mathbb{Z}) \cong \mathbb{Z}^2$. The vertex b has one outgoing edge η_1 which connects it to itself. We have that $\boldsymbol{s}(\eta_b) = 1$, $\boldsymbol{p}(\eta_1)$ is the trivial path and $\boldsymbol{t}(\eta_1) = 0$. The vertex a has three outgoing edges. One, η_2, connects it to itself and two, η_3 and η_4, connect it to b. We have that $\boldsymbol{s}(\eta_2) = 1$, $\boldsymbol{p}(\eta_2)$ is the path b and $\boldsymbol{t}(\eta_2) = (0, 1)$. Of the edges connecting a to b we have that $\boldsymbol{s}(\eta_3) = 1$, $\boldsymbol{s}(\eta_4) = -1$, $\boldsymbol{p}(\eta_1)$ is the trivial path, $\boldsymbol{p}(\eta_4) = bab^{-1}$, $\boldsymbol{t}(\eta_3) = 0$ and $\boldsymbol{t}(\eta_4) = (1, 0)$.

The above definitions allow us to give a different description of the matrix A_f defined in Section 2.2.

Observation 3.9 Let $1 \le i, j, \le m$. Let $E_{i,j}$ be the set of edges in $\mathcal{T}$ connecting e_i to e_j. Let $\boldsymbol{t}$ be the translation corresponding to the cover $\widetilde{\Gamma}_f \to \Gamma$. Then

$$(A_f)_{i,j} = \sum_{\eta \in E_{i,j}} \boldsymbol{s}(\eta)\boldsymbol{t}(\eta),$$

where $\boldsymbol{t}(\eta)$ is understood as an element of the group ring of H_f supported at one point.

This observation can be seen by using the definition of A_f to calculate $A_f\widetilde{e}$, where $\widetilde{e}$ is a lift of the edge e in Γ.

3.2 Vertex subgraphs and extremal subgraphs of $\mathcal{T}$

The graph $\mathcal{T}$ has important subgraphs, which we call *extremal subgraphs* and *vertex subgraphs*, that play a major role in our proof. Before we define them, we require an observation, which follows from the fact that f acts trivially on H_f.

Observation 3.10 The map $\boldsymbol{t}\colon \mathcal{P}(\mathcal{T}) \to H_f$, where $\mathcal{P}(\mathcal{T})$ is viewed as a groupoid under concatenation, is a homomorphism of groupoids.

Let $\boldsymbol{t}$ be the translation function corresponding to the cover $\widetilde{\Gamma}_f \to \Gamma$.

Definition 3.11 For any path $\overline{\eta} = \eta_1 \dots \eta_k$, define $\boldsymbol{t}_n(\overline{\eta})$, the *normalized translation of* $\overline{\eta}$, to be $\boldsymbol{t}_n(\overline{\eta}) = \frac{1}{k}\boldsymbol{t}(\eta) \in H_f \otimes \mathbb{R}$.

Definition 3.12 A *based cycle* in $\mathcal{T}$ is closed path. A *cycle* is the equivalence class of a based cycle, under the relation identifying two based cycles that differ by a cyclic permutation of their edges. One corollary of Observation 3.10 is that the function $\boldsymbol{t}$ is well defined on cycles.

Let $\mathcal{C}$ be the set of cycles in $\mathcal{T}$ and let $\mathcal{C}_s$ be the set of simple cycles in $\mathcal{T}$ (a cycle is simple if it gives an embedding of S^1 into $\mathcal{T}$). As a corollary of Observation 3.10, we get that $\boldsymbol{t}_n(\mathcal{C})$ is contained in the convex hull of $\boldsymbol{t}_n(\mathcal{C}_s)$. Since $\mathcal{C}_s$ is a finite set, this convex hull is a polytope. We call this polytope the *equivariant shadow of* φ and denote it by $\mathcal{S}^e\varphi$.

Every vertex u is in $H_f \otimes \mathbb{Q}$. Since $\mathcal{S}^e\varphi^k = k\mathcal{S}^e\varphi$, by replacing f with some power of itself we can assume that every vertex of $\mathcal{S}^e\varphi$ has integer vertices. In our proof of Theorem 1.5 we will show that it suffices to prove the theorem for f^k for some integer k. Therefore, we can and will assume in the sequel that every vertex of $\mathcal{S}^e\varphi$ is integral.

Definition 3.13 Let $\omega\colon H_f \otimes \mathbb{R} \to \mathbb{R}$ be a linear transformation. Let M_ω be the maximal value ω takes on $\mathcal{S}^e\varphi$. Let $\mathcal{T}_\omega$ be the union of all $\gamma \in \mathcal{C}$ such that $\omega(\boldsymbol{t}_n(\gamma)) = M_\omega$. We call the graph $\mathcal{T}_\omega$ *the extremal subgraph of* $\mathcal{T}$ *corresponding to* ω.

Since every vertex of a convex polytope is the maximal set of some linear function, we have a special kind of extremal subgraph, called a *vertex subgraph*.

Definition 3.14 Let $u \in H_f \otimes \mathbb{R}$ be a vertex $\mathcal{S}^e\varphi$. Let $\mathcal{T}_u$ be the union of all $\gamma \in \mathcal{C}$ such that $\boldsymbol{t}_n(\gamma) = u$. We call $\mathcal{T}_u$ the *vertex subgraph corresponding to* u.

Lemma 3.15 *Let $\omega\colon H_f \otimes \mathbb{R} \to \mathbb{R}$ be as above. Let $\gamma \in \mathcal{C}$. Then $\omega(\boldsymbol{t}_n(\gamma)) = M_\omega$ if and only if γ is a cycle in $\mathcal{T}_\omega$.*

Proof The "only if" direction is just the definition of the subgraph $\mathcal{T}_\omega$. We will prove the "if" direction. For any path δ in $\mathcal{T}_\omega$, define $g(\delta) = \omega \boldsymbol{t}_n(\delta)$. Suppose that γ is a cycle in $\mathcal{T}_\omega$ with $g(\gamma) < M_\omega$.

Choose a based cycle $\eta_1 \dots \eta_k$ in the equivalence class of γ. For every $1 \le i \le k$, the edge η_i is part of the graph $\mathcal{T}_\omega$, and hence is contained in some cycle whose normalized translation is M_ω. If η_i connects vertices e to e', we can thus find a path ζ_i connecting e' to e such that $g(\eta_i \zeta_i) = M_\omega$. Let l_i be the length of ζ_i.

Let ζ be the path $\zeta_k \dots \zeta_1$. Since $g(\eta_i \zeta_i) = M_\omega$, we have that

$$\omega(\boldsymbol{t}(\zeta_i)) = (l_k + 1) M_\omega - \omega(\boldsymbol{t}(\eta_i)).$$

Let $l = l_1 + \dots + l_k$. Our assumption that $g(\gamma) < M_\omega$ gives that

$$g(\zeta) = \frac{1}{l} \sum_{i=1}^{k} \bigl((l_k + 1) M_\omega - \omega(\boldsymbol{t}(\eta_i))\bigr) = \Bigl(1 + \frac{k}{l}\Bigr) M_\omega - \frac{k}{l} g(\gamma) > M_\omega.$$

This is a contradiction of our definition of M_ω. □

Observation 3.16 The notion of vertex subgraphs is central to our proof, and we will need to use it in a more general context than the one outlined above. Note that the proof of Lemma 3.15 did not use any properties of $\boldsymbol{t}$, aside from the fact that it is homomorphism from the groupoid of paths to an abelian group. Thus we can define extremal and vertex graphs with respect to any such homomorphism. We can take this a step further. Any function from cycles to an abelian group that is additive on based cycles which are based at the same point can be extended to a homomorphism from the groupoid of paths. Thus, even in this more general situation, we can still define extremal and vertex subgraphs.

Definition 3.17 Let u be a vertex of $\mathcal{S}^e \varphi$. Let e_i and e_j be vertices of $\mathcal{T}$. Let $E_{i,k}(u)$ be the set of edges in $\mathcal{T}_u$ connecting e_i to e_j. Define the *vertex matrix of* u, or $A_{f,u}$, by setting

$$(A_{f,u})_{i,j} = \sum_{\eta \in E_{i,j}(u)} \boldsymbol{s}(\eta) \boldsymbol{t}(\eta).$$

Similarly, we can define a matrix for any subgraph of $\mathcal{T}$.

Definition 3.18 The vertex u is said to be *stable* if $A_{f,u}$ is not nilpotent.

3.3 Subgraphs and covers

Let $\pi\colon \Gamma' \to \Gamma$ be a cover to which φ can be lifted to a map φ'. The transition graph of φ' is a cover of the graph $\mathcal{T}$. We will denote it by $\mathcal{T}[\pi]$. If $\mathcal{T}_0$ is a subgraph of $\mathcal{T}$, we will denote it by $\mathcal{T}_0[\pi]$. Given an vertex subgraph $\mathcal{T}_u \subset \mathcal{T}$, we denote its lift to $\mathcal{T}[\pi]$ by $\mathcal{T}_u[\pi]$, and let $A_{f,u}[\pi]$ be the associated matrix. We say that u is *stable in the cover* π if $A_{f,u}[\pi]$ is not nilpotent.

3.4 The dimension of $\mathcal{S}^e\varphi$

The group $G = F_n \rtimes_f \mathbb{Z}$ is the fundamental of a mapping torus M_f. If $f\colon \Sigma \to \Sigma$ is a surface diffeomorphism, then this mapping torus is a 3–manifold. If f is a free group automorphism then we form the mapping torus $M_f = \Gamma \times I/{\sim}$, where $(x,0) \sim (\varphi(x),1)$.

We can write $H_1(M_f;\mathbb{Z}) \cong H_1^e(\Gamma;\mathbb{Z}) \oplus \mathbb{Z}$, where $H_1^e(\Gamma;\mathbb{Z})$ is the image of $H_1(\Gamma;\mathbb{Z})$ in $H_1(M_f;\mathbb{Z})$. Let γ be a cycle in $\mathcal{T}$. Following Fried (who used an equivalent definition), we call $(\boldsymbol{t}_n(\gamma),1) \in H_1(M_f;\mathbb{R})$ a *homological direction*. For f pseudo-Anosov and Σ compact, Fried studied the cone on all homological directions and related it to the Thurston norm.

Given a 3–manifold M_f that fibers over the circle, Thurston [19] defines a seminorm τ on $H_2(M_f,\partial M_f;\mathbb{R})$.

The corresponding norm on $H_2(M_f,\partial M_f;\mathbb{R})/\ker\tau$ is a convex polytope. One of the top-dimensional faces of this polytope is called the *fibered face*. If f is pseudo-Anosov and Σ has $b \geq 1$ boundary components, then

$$\dim H_2(M_f,\partial M_f;\mathbb{R})/\ker\tau = \dim H_2(M_f,\partial M_f;\mathbb{R}) - (b-1).$$

Let $\mathcal{C}$ be the cone on the set of homological directions. In [5], Fried proves that this cone has the same dimension as a cone on the fibered face (in fact, he proves a stronger claim: the two cones are dual). By Poincaré duality and the universal coefficient theorem, $\dim H_2(M_f,\partial M_f;\mathbb{R}) = \dim H_1(M_f;\mathbb{R})$. So, for a surface diffeomorphism we get that $\dim \mathcal{S}^e\varphi = \dim H_f \otimes \mathbb{R} + 1 - b$. A simpler statement holds for the case where $f \in \mathrm{Out}(F_n)$. In this case, in [3], Dowdall, Leininger and Kapovich prove that the cone $\mathcal{C}$ is $\dim H_1(M_f;\mathbb{R})$–dimensional. This is also proved separately by Algom-Kfir, Hironaka and Rafi in [1]. This means that $\dim \mathcal{S}^e\varphi$ is $\dim H_f \otimes \mathbb{R}$. We summarize this discussion in the following lemma:

Lemma 3.19 (dimension of $\mathcal{S}^e\varphi$) *If $f\colon \Sigma \to \Sigma$ is a pseudo-Anosov mapping class and Σ has $b \geq 1$ boundary components, then $\dim \mathcal{S}^e\varphi = \dim H_f \otimes \mathbb{R} + (1-b)$. If $f \in \mathrm{Out}(F_n)$ is fully irreducible then $\dim \mathcal{S}^e\varphi = \dim H_f \otimes \mathbb{R}$.*

4 Stabilizing vertex subgraphs

Our goal in this section is to describe a process we call vertex stabilization, in which we start with a vertex of $\mathcal{S}^e\varphi$ and find a cover in which it is stable. Our method uses properties of nilpotent groups.

4.1 Nilpotent groups

Let G be a finitely generated group. Define $G_1 = G$, and for every i set $G_{i+1} = [G, G_i]$. The group G is said to be *nilpotent* if G_n is trivial for some value of n. The sequence of subgroups G_i is called the *lower central series of G.*

In [10], Koberda introduces a modified form of the lower central series, called the *torsion-free lower central series*. This is a series of the form

$$\cdots \lhd G_3^{\mathrm{TF}} \lhd G_2^{\mathrm{TF}} \lhd G_1^{\mathrm{TF}} = G$$

such that

(a) the groups G_i^{TF} are characteristic in G;

(b) the groups $N_i = G/G_i^{\mathrm{TF}}$ are nilpotent;

(c) the groups $L_i = G_i^{\mathrm{TF}}/G_{i+1}^{\mathrm{TF}}$ are finitely generated torsion-free abelian groups that are central in N_{i+1}.

Koberda shows that if G is of the form $G = F \rtimes \mathbb{Z}$ where F is a surface group or a free group, then $\bigcap G_i^{\mathrm{TF}} = \{e\}$.

4.2 Nilpotent stabilization

Let $G = F_n \rtimes_f \mathbb{Z}$, and let $i\colon F_n \to G$ be given by $i(w) = (w, 0)$. Let $G_1^{\mathrm{TF}} \geq G_2^{\mathrm{TF}} \geq \cdots$ be the torsion-free lower central series of G. For every $j \geq 1$, set $K_j = i^{-1}(G_j^{\mathrm{TF}}) \lhd F_n$, $N_j = F_n/K_j$ and $L_j = K_j/K_{j+1}$.

Let π_j be the cover of Γ corresponding to N_j. Denote the corresponding translation function by $\boldsymbol{t}_j$. The subgroups K_j are all f–invariant and thus f acts on the groups N_j and L_j. Since f acts trivially on H_f, it is a standard fact that it acts trivially on each L_j.

Definition 4.1 A subgraph $\mathcal{T}' \subseteq \mathcal{T}$ is called *j–stable* if it has nontrivial cycles and for infinitely many k there exists a p such that

$$\sum_{\gamma} \boldsymbol{s}(\gamma)\boldsymbol{t}_j(\gamma) \neq 0 \in \mathbb{C}[N_j],$$

where the sum is taken over all based cycles of length k in $\mathcal{T}'$ based at p. A vertex u of $\mathcal{S}^e\varphi$ is said to be j–stable if its vertex subgraph is j–stable.

Since $N_1 = H_f$, saying that the vertex u is stable is equivalent to saying that it is 1–stable.

Definition 4.2 A subgraph $\mathcal{T}' \subseteq \mathcal{T}$ is called *j–consistent* if for every vertex p of $\mathcal{T}'$ there exists $x_p \in N_j$ and an integer d dividing the lengths of all cycles in $\mathcal{T}'$ such that, for any cycle γ of length k based at p,

$$\boldsymbol{t}_j(\gamma) = f^k(x_p)\cdots f^{3d}(x_p)f^{2d}(x_p)f^d(x_p).$$

Note that a vertex subgraph is an example of a 1–consistent subgraph.

Lemma 4.3 *Let $\mathcal{T}' \subseteq \mathcal{T}$ be a j–consistent subgraph. Then there exists a $(j+1)$–consistent subgraph $\mathcal{T}'' \subseteq \mathcal{T}'$. Furthermore, the subgraph $\mathcal{T}''$ has the property that if $\mathcal{T}''$ is $(j+1)$–stable then $\mathcal{T}'$ is also $(j+1)$–stable.*

Proof Let $x_p \in N_j$ be the elements provided by the definition of j–consistency. For each p, pick $y_p \in N_{j+1}$ whose image in N_j is x_p. For any l divisible by d, let $P_l(p) = f^l(y_p)\cdots f^d(y_p)$.

Let γ_p be a based cycle of length k in $\mathcal{T}'$ that is based at the vertex p. The *deviation of γ_p*, or $\Delta(\gamma_p)$, is given by the equation

$$\Delta(\gamma_p) = P_k^{-1}(p)\boldsymbol{t}_{j+1}(\gamma_p).$$

For any path $\delta = \delta_1\delta_2$ in the graph Γ we have that $\varphi(\delta) = \varphi(\delta_1)\varphi(\delta_2)$. Thus, for any path δ in the graph $\mathcal{T}$, and any edge η whose initial point is the endpoint of δ we have that $\boldsymbol{p}(\delta\eta) = \varphi(\delta)\boldsymbol{p}(\eta)$. More generally, if δ' is a path of length l whose initial point is the endpoint of δ then $\boldsymbol{p}(\delta\delta') = \varphi^l(\delta)\boldsymbol{p}(\delta')$.

Now let β_p be a based cycle of length l in $\mathcal{T}'$ that is also based at p. By the previous paragraph,

$$\boldsymbol{t}_{j+1}(\gamma_p\beta_p) = f^l(\boldsymbol{t}_{j+1}(\gamma_p))\boldsymbol{t}_{j+1}(\beta_p) = f^l(P_k(p)\Delta(\gamma_p))P_l(p)\Delta(\beta_p).$$

Since f acts trivially on L_j, and L_j is central in N_{j+1}, we get $\boldsymbol{t}_{j+1}(\gamma_p)\beta_p = P_{k+l}(p)\Delta(\gamma_p)\Delta(\beta_p)$. Thus, we have that

$$\Delta(\gamma_p\beta_p) = \Delta(\gamma_p)\Delta(\beta_p).$$

Since L_j is an abelian group, it follows from the above calculation that if β_p is obtained from γ_p by cyclic reordering such that both are cycles based at p, then $\Delta(\gamma_p) = \Delta(\beta_p)$.

Let M be the vector space $L_j^{V(\mathcal{T})} \otimes \mathbb{R}$. Given an unbased cycle γ in $\mathcal{T}'$, define the *basepoint-free deviation of* γ or $\overline{\Delta}(\gamma)$ to be the following element of M. Set $\overline{\Delta}(\gamma)[p] = 0$ if γ doesn't pass through p. Otherwise, set $\overline{\Delta}(\gamma)[p]$ to be the image of $\Delta(\gamma_p)$ in $L_j \otimes \mathbb{R}$, where γ_p is a basing of the loop γ at p. Note that this function depends on j. We say that $\overline{\Delta}$ is the j–level basepoint-free deviation function.

The function $\overline{\Delta}$ is additive on cycles in $\mathcal{T}'$. Furthermore, since G is finitely generated, the vector space $L_j \otimes \mathbb{R}$ is finite-dimensional, and hence M is finite-dimensional. By Observation 3.16, we can use the map $\overline{\Delta}$ to choose a vertex subgraph $\mathcal{T}'' \subseteq \mathcal{T}'$ corresponding to the vertex $v \in M$.

Since L_j is torsion-free, we have an inclusion $V(\mathcal{T}')^{L_j} \subset M$. Fix a vertex p of $\mathcal{T}$ and let v_p be the p–coordinate of v. For any cycle γ_p based at p of length k we have that $\boldsymbol{t}_j(\gamma_p) = P_k(p)\Delta_p(\gamma_p)$. We must therefore have that $kv \in L_j$. If we write $v_p = \frac{1}{q}w_p$ with $w_p \in L_j$ and $q \in \mathbb{N}$, we get that $q|k$.

Set d' to be the greatest common divisor of the lengths of all loops in $\mathcal{T}''$. We can take $w_p = d'v$. We have that $d' = ad$ for some integer a. Define $z_p = f^{ad}(y_p)\cdots f^d(y_p)\cdot w_p$. This choice of the z_p makes the graph $\mathcal{T}''$ consistent. Indeed, for any cycle γ_p of length $k = bd'$ we have that

$$\boldsymbol{t}(\gamma_p) = P_k(p)\Delta_p(\gamma_p) = P_k(p)kv_p = P_k(p)bw_p = f^k(z_p)\cdots f^{d'}(z_p).$$

Now suppose that $\mathcal{T}''$ is $(j+1)$–stable. Since it is a vertex subgraph of $\mathcal{T}'$, by Observation 3.16 and Lemma 3.15 we have that $\mathcal{T}'$ is $(j+1)$–stable. □

Definition 4.4 A subgraph $\mathcal{T}' \subseteq \mathcal{T}$ is called a *j–vertex subgraph* if there is a sequence of subgraphs $\mathcal{T}' = \mathcal{T}_j \subseteq \cdots \subseteq \mathcal{T}_1$ such that for all i, $\mathcal{T}_{i+1}$ is $(i+1)$–consistent and a vertex subgraph in $\mathcal{T}_i$ with respect to the i–level map $\overline{\Delta}$. Note that if $\mathcal{T}_{i+1}$ is an i–vertex subgraph then so is $\mathcal{T}_l$ for any $l < i$.

Lemma 4.5 (nilpotent stability) *Let u be a vertex of $\mathcal{S}^e\varphi$. There exists a $j \geq 1$ such that u is j–stable.*

Proof By repeated application of Lemma 4.3, we can find for every j a j–vertex subgraph $\mathcal{T}_{u,j} \subseteq \mathcal{T}_u$.

Since $\mathcal{T}_u$ is a finite graph, any sequence of subgraphs must stabilize, say at $\mathcal{T}_{u,N}$. That is, $\mathcal{T}_{u,k} = \mathcal{T}_{u,N}$ for every $k > N$. Given two based loops β and γ in $\mathcal{T}_N$ based at the same point and of the same length we must have that $\boldsymbol{t}_k(\gamma) = \boldsymbol{t}_k(\beta)$ for every $k \geq N$ (otherwise we could choose a further vertex subgraph).

Let $\boldsymbol{t}_\infty$ be the translation function corresponding to the universal cover of Γ. Since φ is a train track representative, for any edge e of Γ the path $\varphi^l(e)$ is immersed in Γ. Since $\pi(\Gamma)$ is a free group, this means that, given two different based cycles γ_1 and γ_2 in $\mathcal{T}$ of the same length and based at the same point, we must have that $\boldsymbol{t}_\infty(\gamma_1) \neq \boldsymbol{t}_\infty(\gamma_2)$. Since the sequence $\{K_j\}_j$ satisfies that $\bigcap K_j = \{1\}$, the graph $\mathcal{T}_{u,N}$ is a disjoint collection of cycles. Any such collection is obviously N–stable. The result now follows. □

4.3 Upgrading nilpotent stabilization

4.3.1 The trace of powers lemma Fix $r, l \in \mathbb{N}$ and $A \in M_l(\mathbb{C}[\mathbb{Z}^r])$. For any integer k, let $t_k = \mathrm{Tr}[A^k]$.

Lemma 4.6 (the trace of powers lemma) *Given any lattice $L < \mathbb{Z}^r$ there exists a collection $\alpha_1, \ldots, \alpha_s \in \mathbb{C}$ (which depends on L) and a number $C > 0$ such that*

$$t_k(L) = C \sum_i \alpha_i^k.$$

Furthermore, if A is not nilpotent then there exists a number $N > 0$ such that, for all $j > N$, $t_k(j\mathbb{Z}^r) \neq 0$ for infinitely many values of k.

Proof Fix a lattice L. Recall from the proof of Lemma 4.6 that there exists a finite set $N_L \subset (\mathbb{C}^\times)^r$ such that, for any k,

$$t_k(L) = \frac{1}{|N_L|} \sum_{\xi \in N_L} t_k(\xi).$$

Set $p(x) = \det(xI_r - A) \in \mathbb{C}[\mathbb{Z}^r][x]$. We think of p as a polynomial in the variable x with coefficients in $\mathbb{C}[\mathbb{Z}^r]$. Let $\xi\colon \mathbb{Z}^r \to \mathbb{C}^\times$. The map ξ extends linearly to a ring homomorphism $\mathbb{C}[\mathbb{Z}^r][x] \to \mathbb{C}[x]$. The image of p under this homomorphism is called the *specialization of p at ξ* and is denoted by p_ξ. Note that p_ξ is a degree r polynomial.

Pick r (not necessarily continuous) functions $\rho_1, \ldots, \rho_r\colon (\mathbb{C}^\times)^r \to \mathbb{C}$ such that, for any ξ, $\rho_1(\xi), \ldots, \rho_r(\xi)$ is the collection of all roots of the polynomial p_ξ, counted with multiplicity.

For any ξ, the numbers $\rho_1(\xi), \dots, \rho_r(\xi)$ are the roots of the characteristic polynomial of the matrix $A(\xi)$. Thus, $\rho_1^k(\xi), \dots, \rho_r^k(\xi)$ are the roots of the characteristic polynomial of the matrix $A^k(\xi)$. Therefore, $t_k(\xi) = \sum_i \rho_i^k(\xi)$. It follows that

$$t_k(L) = \frac{1}{|N_L|} \sum_{\xi \in N_L} \sum_i \rho_i^k(\xi).$$

This shows the first claim of the lemma. We now show the second claim. Since A is not nilpotent, $t_k \neq 0$ for some value of k. Recall that the Fourier transform $\hat{t}_k$ is the restriction of the function $\xi \to t_k(\xi)$ to $(S^1)^r$, where $S^1 = \{z \in \mathbb{C} : |z| = 1\}$. This is a continuous function since t_k has finite support. Since $t_k \neq 0$, this function is not the zero function.

Write $N_j = N_{j\mathbb{Z}^r}$. Using the definition of N_j, as it appears in Lemma 4.6 we have

$$N_j = \{\xi : \xi|_{j\mathbb{Z}^r} = 1\}.$$

Thus, N_j is the set of all points of order dividing j in the torus $(S^1)^r$. Thus, for any $\epsilon > 0$, the set N_j forms an ϵ–net in $(S^1)^r$ for all sufficiently large j. In particular, by continuity of the trace of A, for all sufficiently large j there exists $\xi \in N_j$ and $1 \leq i \leq r$ such that $\rho_i(\xi) \neq 0$. The second claim now follows from the elementary fact that if $\alpha_1, \dots, \alpha_s$ are not all 0, then $\sum \alpha_i^k \neq 0$ for infinitely many values of k. □

Note that an identical proof holds if we replace the lattice L with a translate of itself. We get the following:

Lemma 4.7 *Given any lattice $L < \mathbb{Z}^r$ and a vector $\overline{w} \in \mathbb{Z}^r$, there exists a collection $\alpha_1, \dots, \alpha_s \in \mathbb{C}$ (which depends on L) and a number $C > 0$ such that*

$$t_k(L + \overline{w}) = C \sum_i \alpha_i^k.$$

Furthermore, if A is not nilpotent then there exists a number $N > 0$ such that, for all $j > N$ and for all $\overline{w}$, $t_k(j\mathbb{Z}^r + \overline{w}) \neq 0$ for infinitely many values of k.

4.3.2 k–covers and nilpotent quotients Let $\mathcal{G}$ be a residually torsion-free, finitely generated group. Let $\{\mathcal{G}^{\mathrm{TF}}\}_i$ be its torsion-free lower central series. For $x_1, \dots, x_i \in \mathcal{G}$, let

$$[x_1, \dots, x_i] = [\dots[[x_1, x_2], x_3], \dots, x_i].$$

Let S be a finite generating set for Γ. It is a standard fact that, for any j, $L_j = Z(\mathcal{G}/\mathcal{G}_j^{\mathrm{TF}})$ is generated by elements of the form $[a_1, \ldots, a_j]$, where $a_i \in S$. We require the following simple lemma:

Lemma 4.8 *Let $j, k \in \mathbb{N}$. For any $1 \le i \le j$ and, for any $a_1, \ldots, a_j \in \mathcal{G}$,*

$$[a_1, \ldots, a_j]^k \equiv_j [a_1, \ldots, a_i^k, \ldots, a_j],$$

where $\equiv_j$ is understood as having equal images in L_j.

Proof We prove the claim inductively on j. The claim is obvious for $j = 1$. Assume we've proved the claim for all numbers up to j. We now prove it for $j+1$. Repeated application of the basic commutator identities $[x, zy] = [x, y] \cdot [x, z]^y$ together with the fact that conjugation acts trivially on L_i for all i gives that

$$[a_1, \ldots, a_j^k] = [a_1, \ldots, a_j] \cdot [a_1, \ldots, a_j]^y \cdots [a_1, \ldots, a_j]^{y^{k-1}} \equiv_j [a_1, \ldots, a_j]^k.$$

For $i < j$ the inductive claim gives us a $w \in \mathcal{G}_j^{\mathrm{TF}}$ such that

$$\begin{aligned} [a_1 \ldots, a_i^k, \ldots, a_j] &= [[a_1, \ldots, a_i^j, \ldots, a_{j-1}], a_j] \\ &= [[a_1, \ldots, a_{j-1}]^k w, a_j] \equiv_j [[a_1, \ldots, a_{j-1}]^k, a_j]. \end{aligned}$$

Using the identity $[y, x] = [x, y]^{-1}$ and the fact that $[a_1, \ldots, a_j^k] \equiv_j [a_1, \ldots, a_j]^k$ now yields the result. □

For any k, let $\mathcal{G}[k]$ be the kernel of the natural map $\mathcal{G} \to H_1(\mathcal{G}, \mathbb{Z}/k\mathbb{Z})$. Let $\{\mathcal{G}[k]_j^{\mathrm{TF}}\}_j$ be the torsion-free lower central series of $\mathcal{G}[k]$, and let

$$L_j[k] = Z(\mathcal{G}[k]/\mathcal{G}[k]_j^{\mathrm{TF}}).$$

The inclusion $\mathcal{G}[k] \to \mathcal{G}$ induces a natural map $L_j[k] \to L_j$. As a corollary to Lemma 4.8, we have the following:

Corollary 4.9 *For any $j, k \in \mathbb{N}$, the image of the natural map $L_j[k] \to L_j$ is $k^j L_j$.*

4.3.3 The upgrade lemma Fix a vertex u of $\mathcal{S}^e\varphi$. For any integer k, let $G[k]$ be the kernel of the map $G \to H_1(G, \mathbb{Z}/k\mathbb{Z})$. Let $F[k] = i^{-1}(G_k)$. Let $\pi_k \colon \Gamma_k \to \Gamma$ be the cover of Γ corresponding to this map. The map φ^k lifts to Γ_k. Call this lift φ_k. Let $\mathcal{T}_u^k$ be the vertex subgraph corresponding to u in the transition graph of φ^k.

Lemma 4.10 (the upgrade lemma) *Let $j \in \mathbb{N}$, and suppose u is a $(j+1)$–stable vertex for φ. Then, for all but finitely many $k \in \mathbb{N}$, the graph $\mathcal{T}_u^k[\pi_k]$ is a j–stable extremal subgraph.*

Proof Let $\mathcal{T}' \subseteq \mathcal{T}_u$ be a j–vertex subgraph. As in Lemma 4.3, let $\mathcal{M} = V(\mathcal{T})^{L_{j+1}} \otimes \mathbb{R}$, and let $\overline{\Delta}$ be the level $j+1$ basepoint-free deviation function. Since the function $\overline{\Delta}$ is additive on cycles, we can extend it to in $\mathcal{T}_u$, and produce a corresponding matrix $B \in M_{|V(\mathcal{T})|}(\mathbb{C}[\mathcal{M}])$ such that, for any k,

$$\mathrm{Tr}(B^k) = \sum_{\gamma} \boldsymbol{s}(\gamma)\overline{\Delta}(\gamma),$$

where the sum is taken over all cycles of length k in $\mathcal{T}'$ and $\overline{\Delta}(\gamma)$ is understood as an element of $\mathbb{C}[M]$. Since V_{j+1} is torsion-free, we have a natural inclusion map ${L_{j+1}}^V(\mathcal{T}) \subset M$. Let $m = |V(\mathcal{T})|$. Write $\mathcal{L} = {L_{j+1}}^V(\mathcal{T}) \cong (\mathbb{Z}^r)^m$, where $r = \mathrm{rank}(L_{j+1})$. Note that, by construction, $\mathrm{Tr}(B^k) \in \mathcal{L}$ for every k.

Since $\mathcal{T}'$ is stable, the matrix B is not nilpotent. Let N be the number provided by Lemma 4.6. Fix $k > N$. For any s, let $\boldsymbol{t}_s$ be the translation function corresponding to the map $G \to G/G_s^{\mathrm{TF}}$, and $\boldsymbol{t}[k]_s$ be the translation corresponding to the map $G[k] \to G[k]/G^{\mathrm{TF}}[k]_s$.

For any i, let $T_i = \sum_{\gamma} \boldsymbol{s}(\gamma)\boldsymbol{t}_{j+1}(\gamma)$, where the sum is taken over all cycles of length i in $\mathcal{T}'$. For any x in the support of T_i, let a_x^i be its coefficient. Since $\mathcal{T}'$ is a j–vertex subgraph, for any x and y in the support of T_i we have that $xy^{-1} \in L_{j+1}$. Let $T_i[x,k] = \sum_y a_y^i$, where the sum is taken over all y such that $xy^{-1} \in k^{j+1}L_{j+1}$.

Consider the sum $\sum_{\tilde{\gamma}} \boldsymbol{s}(\tilde{\gamma})\boldsymbol{t}[k]_{j+1}(\tilde{\gamma})$, where the sum is taken over the lifts to $\mathcal{T}'[\pi_k]$ of all cycles of length i in $\mathcal{T}'$. Pick y in the support of this sum such that the image of y in G/G_{j+1}^{TF} is x. Let $\bar{y}$ be the image of y in $G[k]/G^{\mathrm{TF}}[k]_j$ under the natural inclusion map.

By Corollary 4.9, the coefficient of $\bar{y}$ in $\sum_{\tilde{\gamma}} \boldsymbol{s}(\tilde{\gamma})\boldsymbol{t}[k]_j(\tilde{\gamma})$ is equal to $T_i[x,k]$. Since $\mathcal{T}'$ is a j–vertex, this in turn is equal to $\mathrm{Tr}(B^i)(\mathcal{L}+\overline{w})$ for some $\overline{w} \in \mathcal{L}$. By Lemma 4.6, this is not equal to 0 for infinitely many values of i. This concludes the proof. □

4.3.4 Weak vertex subgraphs and weakly consistent subgraphs

Observation 4.11 We would like to use Lemma 4.5 together with repeated applications of Lemma 4.10 to produce a cover where a given vertex u of $\mathcal{S}^e\varphi$ is stable. However, Lemma 4.10 inputs a cover where a given *vertex* is $(j+1)$–stable and outputs a further cover where a given *face* is j–stable. It is quite possible that the vertices in this face are not themselves j–stable, which prevents the use of an inductive step. To circumvent this issue, we introduce a slight technical generalization called *weakly vertex subgraphs and weakly consistent subgraphs*.

Definition 4.12 A subgraph $\mathcal{T}' \subset \mathcal{T}$ is said to be *of vertex type* if there exists some $v \in H_f \otimes \mathbb{R}$ such that $\mathcal{T}'$ consists all loops γ with $\boldsymbol{t}_n(\gamma) = v$.

Definition 4.13 A subgraph $\mathcal{T}' \subset \mathcal{T}$ of vertex type is said to be a *weak j–level vertex* if it is j–stable and there exists some linear transformation $T\colon H_f \otimes \mathbb{R} \to \mathbb{R}$ such that for any $v' \in H_f \otimes \mathbb{R}$ with $T(v') > T(v)$, the subgraph of vertex type corresponding to v' is not j–stable.

Observation 4.14 Let v and v' be as in the above definition. By replacing f with a power of itself, we may assume that for all such v', for any vertex p and, for any k, $\sum \boldsymbol{s}(\gamma_p)\boldsymbol{t}_j(\gamma_p) = 0$, where the sum is taken over all based loops of length k in the vertex type subgraph corresponding to v' that are based at p and $\boldsymbol{t}_j$ is the translation function corresponding to N_j.

Definition 4.15 A subgraph $\mathcal{T}' \subseteq \mathcal{T}$ is called *weakly j–consistent* if for every vertex p of $\mathcal{T}'$ there exists $x_p \in N_j$ and an integer d dividing the lengths of all cycles in $\mathcal{T}'$ such that, for all sufficiently large k,

$$\sum_{\gamma} \boldsymbol{s}(\gamma)\boldsymbol{t}_j(\gamma) = C_k f^k(x_p) \cdots f^{3d}(x_p) f^{2d}(x_p) f^d(x_p),$$

where the above sum is taken in the group ring of N_j, the sum is taken over all based cycles of length k in $\mathcal{T}'$ that are based at p, C_k is some number and $\boldsymbol{t}_j$ is the translation function corresponding to N_j.

Suppose $\mathcal{T}''$ is a weakly j–consistent subgraph of the weak j–level vertex subgraph $\mathcal{T}'$ corresponding to $v \in H_f \otimes \mathbb{R}$. For any vertex p and any based loop γ_p satisfying $\boldsymbol{t}_n(\gamma_p) = v$, define the deviation of γ_p or $\Delta(\gamma_p)$ exactly as in Lemma 4.3. The same calculation as the one done in Lemma 4.3 shows that Δ remains constant under cyclic reordering of based cycles based at p, and that it is additive on such cycles. Exactly as in Lemma 4.3, define the basepoint-free deviation of a cycle γ satisfying $\boldsymbol{t}_n(\gamma) = v$. We think of $\bar{\Delta}(\gamma)$ as an element of the group ring $\mathbb{C}[L_j^{V(\mathcal{T})}]$. Extend the definition of $\bar{\Delta}$ to cycles γ in $\mathcal{T}''$ not satisfying $\boldsymbol{t}_n(\gamma) = v$ by setting $\bar{\Delta}(\gamma) = 0$.

Observation 4.16 Since v is a j–level weak vertex subgraph, we can find a matrix $B \in M_{|V(\mathcal{T})|}(\mathbb{C}[L_j^{V(\mathcal{T})}])$ such that, for any k,

$$\mathrm{Tr}(B^k) = \sum_{\gamma} \boldsymbol{s}(\gamma)\bar{\Delta}(\gamma),$$

where the sum is taken over all cycles of length k in $\mathcal{T}''$ satisfying $\boldsymbol{t}_n(\gamma) = p$. The same proof as Lemma 4.10 now gives us that, for all but finitely many k's, the graph $CT''[\pi_k]$ is $(j-1)$–stable.

4.3.5 The vertex stabilization lemma

Lemma 4.17 (the vertex stabilization lemma) *For any vertex u of $\mathcal{S}^e\varphi$, there exists a solvable cover $\Gamma' \to \Gamma$ to which φ lifts such that ku is a stable vertex in Γ' for some k.*

Proof By Lemma 4.5, the vertex u is j–stable for some j. By Lemma 4.10, we can find some k_1–cover π_{k_1} such that the extremal subgraph $\mathcal{T}_u[\pi_{k_1}]$ is $(j-1)$–stable. Since this extremal subgraph is $(j-1)$–stable, it has a weak $(j-1)$–level vertex u_1. Let $\mathcal{T}_{u_1}$ be the corresponding graph. By applying Observation 4.16 we can find a k_2 and a k_2–cover π_{k_2} of π_{k_1} such that $\mathcal{T}_{u_1}[\pi_{k_2}]$ is $(j-2)$–stable. This means that $\mathcal{T}_u[\pi_{k_2}]$ is $(j-2)$–stable.

Proceeding inductively in this manner, we can find a cover π in which $\mathcal{T}_u$ is 1–stable, and hence stable. Note that at each step, the abelian covers we take are characteristic since they are simply the covers obtained by reducing homology modulo some integer. Since π is obtained by iterating characteristic covers, it is a regular cover. Since it is obtained by iterating abelian covers, it is solvable. □

5 Proof of theorems

5.1 Lemmas

In this section we collect several technical lemmas that we require for our proof.

5.1.1 The cyclic deformation lemma

Lemma 5.1 (cyclic deformation lemma) *Let $\mathcal{T}' \subseteq \mathcal{T}$ be a weak $(j+1)$–stable vertex subgraph that is weakly j–consistent. Let $\psi\colon G \to \mathbb{Z}/q\mathbb{Z}$ be a homomorphism into a cyclic group of prime order q such that $i(F_n) \not\leq \ker\psi$ and such that $\mathcal{T}'$ is j–stable in the cover corresponding to $\ker\psi$. Then, for all sufficiently large primes p, there are homomorphisms $\psi_p\colon G \to \mathbb{Z}/p\mathbb{Z}$ such that $i(F_n) \not\leq \ker\psi_p$ and $\mathcal{T}'$ is j–stable in the cover corresponding to $\ker\psi_p$.*

Proof Pick a basis $a_1,\dots,a_l$ for H_f, and extend it to a minimal generating set B for $H_1(G;\mathbb{Z})$ such that $\psi\colon H_1(G,\mathbb{Z})\to\mathbb{Z}/q\mathbb{Z}$ sends a_1 to 1 and all other generators to 0. For any prime p, let ψ_p be the homomorphism $\psi_p\colon H_1(G,\mathbb{Z})\to\mathbb{Z}/p\mathbb{Z}$ sending a_1 to 1 and all other generators to 0. So $\psi=\psi_q$.

Let $L_{j+1}=G^{\mathrm{TF}}_{j+1}/G^{\mathrm{TF}}_{j+2}$. Denote by $L_{j+1}[\psi_p]$ the image of the $(j+1)^{\mathrm{st}}$ term of the lower central series of $\ker\psi_p$ in L_{j+1}. Denote by L^*_{j+1} the dual group of L_{j+1} (that is, the group of all characters on L_{j+1} with image in the unit circle). Let $\chi\in L^*_{j+1}$ be such that $L_{j+1}[\psi_q]\le\ker\chi$.

We begin by showing that for all $\epsilon>0$, there exist infinitely many primes p, and elements $\chi_p\in L^*_{j+1}$, such that $L_{j+1}[\psi_p]\le\ker\chi_p$, and the distance from χ to χ_p is less than ϵ (here we're using the distance on L^*_{j+1} induced by an embedding into $\mathbb{C}^{\mathrm{rank}(L_{j+1})}$).

The group L_{j+1} is generated by the images of elements of the form $[b_1,\dots,b_{j+1}]$, where $b_i\in B$. By Lemma 4.8, the lattice $L_{j+1}[\psi_p]$ is generated by elements of the form $p^s[b_1,\dots,b_j]$, where s is the number of times that a_1 appears in $b_1,\dots,b_j$.

Let $R=\mathrm{span}_{\mathbb{Z}}([b_1,\dots,b_{j+1}])$ be the lattice of formal linear combinations of generators as above. For any $i\ge 0$ let Ri be the sublattice generated by all elements where a_1 appears i times. Let U_i be the image of R_i in L_{j+1}. Then $L_{j+1}[\psi_p]=\sum p^iU_i$.

By definition, we have that $U_0\le\ker\xi$, and must have that $U_0\le\ker\xi_p$ for every p. Given a lattice $L\le\mathbb{Z}^r$, we can find a direct sum decomposition $\mathbb{Z}^r\cong\bigoplus M_i$, with $L\cong\bigoplus(L\cap M_i)$ and $L\cap M_i$ a finite-index subgroup of M_i. Since every element of $L_{j+1}/L_{j+1}[\psi_p]$ has order that is a power of p and $U_0\le L_{j+1}[\psi_p]$ for all p, we must have that U_0 is a direct summand of L_{j+1}. Write $L_{j+1}=U_0\oplus V$.

For every p, let $N_p\subset L^*_{j+1}$ be the set of all characters ζ such that $U_0\le\ker\zeta$ and $\zeta(V)$ is contained in the set of p^{th} roots of unity. Let N_∞ be the set of all characters ζ such that $U_0\le\ker\zeta$. We have that $\xi\in N_\infty$, and any $\zeta\in N_p$ satisfies $L_{j+1}[\psi_p]\le\ker\zeta$. We now conclude by noting that for every $\epsilon>0$, the set N_p is ϵ–dense in N_∞ for all sufficiently large p.

We now proceed similarly to the proof of Lemma 4.10. Let $m=|V(\mathcal{T})|$ and $\mathcal{L}=\mathbb{C}[L^V_{j+1}(\mathcal{T})]$. As in Lemma 4.10, there is a $B\in M_m(\mathcal{L})$ such that, for any k,

$$\operatorname{Tr}B^k=\sum_\gamma s(\gamma)\bar{\Delta}(\gamma),$$

where the sum is taken over all loops γ of length k in $\mathcal{T}'$. Since $\mathcal{T}'$ is $(j+1)$–stable, the matrix B is not nilpotent.

As in the proof of Lemma 4.6, let $p(x)$ be the characteristic polynomial of B, and let $\rho_1,\dots,\rho_m$ be a collection of m roots of $p(x)$. Since $\mathcal{T}'$ is j–stable in the cover corresponding to $\ker\psi$, then, as in Lemma 4.10, there exists $\chi \in L^*_{j+1}$ with $L_{j+1}[\psi] \le \ker\chi$, and a root ρ_i such that $\rho_i(\chi) \neq 0$.

For every k, $\operatorname{Tr} B^k[\chi] = \sum_i \rho_i(\chi)^k$. Thus, we can find a k where $\operatorname{Tr} B^k[\chi] \neq 0$. The polynomial $\operatorname{Tr} B^k$ is continuous, and thus there exists an open set U containing χ where $\operatorname{Tr} B^k$ is nonzero at every point. For every $\chi' \in U$ there exists an i such that $\rho_i(\chi') \neq 0$.

Let $N_{\psi_p} = \{\xi \in L^*_{j+1} : L_{j+1}[\psi_p] \le \ker\xi\}$. By the above claim, there are infinitely many values of p such that $N_{\psi_p} \cap U \neq \varnothing$. For each such p, there are infinitely many values of l such that

$$\sum_i \sum_{\chi' \in N_{\psi_p}} \rho_i(\chi')^l \neq 0.$$

Thus, as in the proof of Lemma 4.10, the graph $\mathcal{T}'$ is j–stable in the cover corresponding to $\ker\psi_p$. □

5.1.2 The cyclic cover multiplicity lemma

Lemma 5.2 (cyclic cover multiplicity lemma) *Let $\pi\colon \Gamma_p \to \Gamma$ be a cyclic cover of degree p, for a prime number p. Let u be a vertex of φ such that the image of u in $\mathbb{Z}/p\mathbb{Z}$ is not 0. Let $\mathcal{T}_u$ be the vertex subgraph of u and $A_u[\pi]$ be the matrix corresponding to $\mathcal{T}_u[\pi]$. For any k, let $t^k[\pi,u]$ be the trace of the matrix $A_u[\pi]^k$. Then, as an element of the additive group of $\mathbb{Z}[H_1(\Gamma_p,\mathbb{Z})]$, $t^k(\pi,u)$ is divisible by p.*

Proof For any integer k,

$$\operatorname{Tr}(A_u[\pi]^k) = \sum_\gamma \boldsymbol{t}(\gamma),$$

where the above sum is taken over all *based* cycles γ of length k in $\mathcal{T}_u$. Since $\boldsymbol{t}(\gamma)$ does not depend on the choice of basepoint of a cycle, we can rewrite the above as

$$\operatorname{Tr}(A_u[\pi]^k) = \sum_\delta n_\delta \boldsymbol{s}(\delta)\boldsymbol{t}(\delta),$$

where the sum is taken over all unbased cycles of length k in $\mathcal{T}_u$ and n_δ is an integer.

Let δ be an unbased cycle. The group $\mathbb{Z}$ acts transitively on the set of based cycles corresponding to δ by cyclic rotations. Let s_δ be the size of the image of $\mathbb{Z}$ under this action. We call s_δ the *cyclic stabilizer of* δ. Then $n_\delta = k/s_\delta$.

An alternative characterization of n_δ is the following. Let γ be a based representative of δ. Let γ' be the minimal based subcycle of γ (sharing the same basepoint) such that $\gamma = (\gamma')^l$. Then

$$n_\delta = \frac{k}{l} = \frac{k}{s_\delta}\operatorname{length}(\gamma').$$

The based cycle α' projects to a based cycle α in $\mathcal{T}_u$. By the definition of $\mathcal{T}_u$, $\boldsymbol{t}(\alpha) = \operatorname{length}(\alpha)u$. Since u projects to a generator of $\mathbb{Z}/p\mathbb{Z}$, and γ' is a cycle in the cover π, we must have that $\operatorname{length}(\alpha) = \operatorname{length}(\gamma')$ is divisible by p. By the above, this means that n_δ is divisible by p, as required. □

5.1.3 Invariance of trace Let $\pi\colon \widetilde{\Gamma} \to \Gamma$ be a finite regular cover, to which φ lifts to a map $\widetilde{\varphi}\colon \widetilde{\Gamma} \to \widetilde{\Gamma}$. Let $\mathcal{D}$ be the deck group of the cover π. Since $\widetilde{\varphi}$ is a lift of φ, we have that $\widetilde{\varphi}\mathcal{D} = \mathcal{D}\widetilde{\varphi}$. Thus, some power of $\widetilde{\varphi}$ commutes with every element of $\mathcal{D}$. Note that the group $\mathcal{D}$ also acts on $H_{\widetilde{f}}$, the homology of the corresponding mapping torus.

Given an edge η of the train track graph of f, and a lift $\widetilde{\eta}$ of η to $\widetilde{\Sigma}$, for any $\sigma \in \mathcal{D}$ we have that $\widetilde{\varphi}\sigma(\widetilde{\eta}) = \sigma\widetilde{\varphi}(\widetilde{\eta})$. Thus, by the definition of $A_f[\pi]$, we have the following:

Lemma 5.3 (invariance of trace) *Let $\pi\colon \widetilde{\Gamma} \to \Gamma$ be a finite regular cover with deck group $\mathcal{D}$ to which φ lifts. Then there is some integer k such that $\operatorname{Tr}(A_f[\pi]^k)$ is $\mathcal{D}$–invariant. Furthermore, if $h \in \operatorname{support}(\operatorname{Tr}(A_f[\pi]^k))$ then the multiplicity of h in $\operatorname{Tr}(A_f[\pi]^k)$ is divisible by $n_h = \{\delta \in \mathcal{D} : \delta_* h = h\}$.*

5.1.4 The polytope/lattice lemmas

Lemma 5.4 (the polytope/lattice lemma) *Let $P \subset \mathbb{R}^d$ be a d–dimensional convex polytope. There exists a lattice $L' \subset \mathbb{R}^n$ and a translate L of L' such that $P \cap L$ consists of at least $d+1$ points, all of which are vertices of P, and the convex hull of the points in the intersection is a d–dimensional polytope. Furthermore, for any vertex v of P, L can be chosen so that $v \in P \cap L$.*

Proof We will prove the claim by induction on d. The claim is obvious for $d = 1$. Assume inductively that we've proved it for all dimensions up to $d - 1$.

Fix a vertex v of P. Without loss of generality, we can take v to be the origin and $v \in Q$. Let F_1 be a $(d-1)$–dimensional face of P incident at v. Let F_2 be the opposite face of F_1 (by this we mean that there exists a linear map $\omega\colon \mathbb{R}^n \to \mathbb{R}$ whose minimal value on P is achieved precisely on F_1 and whose maximal value is achieved precisely on F_2).

The vertex v is also a vertex of F_1. Let $W = \operatorname{span}(F_1)$, and let X be the set of all $\eta \in W^*$ whose minimum on F_1 is achieved exactly at v. The set X is an open set. Pick a point $u \in F_2$, and let T' be the linear operator $T'(x) = x - \omega(x)u$. The map T' sends F_2 to W. For any $\eta \in X$, there is a face of $T'(F_2)$ on which $\eta \circ T'$ achieves its minimum. Since X is open, we can choose $\eta \in X$ so this minimum is achieved at exactly one point. Call this point v'.

Define a new projection operator $T(x) = x - \omega(x)v'$. Let Q be the convex hull of $F_1 \cup T(F_2)$. This polytope is $(d-1)$–dimensional. Note that, by construction, v is a vertex of Q. Apply the induction assumption to Q to get a lattice $L_0 \subset W$. Let $C_0 = L_0 \cap Q$. Let $L \subset \mathbb{R}^d$ be the lattice generated by all the vertices in $T^{-1}(C_0)$. Note that since C_0 has at least d points, and, since $T(v') = v$ and $v \in C_0$, we get that L intersects P in at least $d+1$ points. By construction, L is a lattice as required. □

We require a slightly more specialized version of Lemma 5.4 that allows some more control over the vertices of P that belong to L.

Corollary 5.5 *Let $P \subset \mathbb{R}^d$ be a convex polytope, one of whose vertices is the origin. Suppose we have two maps $T\colon \mathbb{R}^d \to \mathbb{R}^m$ and $S\colon \mathbb{R}^d \to \ker T$ such that for every vertex v of $T(P)$, the set $T^{-1}(v)$ consists of a single vertex and such that the only vertices in $\ker S$ are vertices of the above form. Then there exists a lattice L as in Lemma 5.4 such that $L \cap V(P)$ contains at least $m+1$ vertices that project to vertices of $T(P)$.*

Proof Apply Lemma 5.4 to the polytopes $T(P)$ and $S(P)$. Let C_S and C_T be the corresponding sets of vertices. Let L be the lattice generated by the following set: for every $0 \neq v \in C_S$, pick a single vertex in $S^{-1}(v)$, and then add all the vertices of the form $T^{-1}(C_T)$. This is a lattice as required. □

5.1.5 The positive vertices lemma

Lemma 5.6 (positive vertices lemma) *Suppose that every vertex of $\mathcal{S}^e\varphi$ is stable. Then there exists a $k > 0$ such that, for every vertex v of $\mathcal{S}^e\varphi$ with vertex matrix A_v, $\operatorname{Tr} A_v^k = a_v(kv)$, where $a_v > 0$.*

Proof By Lemma 4.6, for every v there exists a collection of numbers $X_v = \{\alpha_1, \dots, \alpha_r\}$ such that $\operatorname{Tr} A_v^k = \sum \alpha_i^k$. Let $X = \bigcup_v X_v$. Let

$$Y = \left\{ \frac{\alpha}{|\alpha|} \;\middle|\; 0 \neq \alpha \in X \right\}.$$

Enumerate the elements of Y as $Y = \{\beta_1, \dots, \beta_t\}$, and let

$$\bar{y} = (\beta_1, \dots, \beta_t) \in \mathbb{T}^t = (S^1)^t.$$

There exists a number M such that the i^{th} component of $\bar{y}^M$ is 1 for every i where β_i is a root of unity. Say $\beta_i = 1$ for $0 \le i \le r$. For every other i, the coordinate β_i is not a root of unity. Thus, the set $\{\bar{y}^{Mj}\}_{j=1}^{\infty}$ is dense in the subtorus $\mathbb{T}^{t-r}$ obtained from $\mathbb{T}^t$ by setting the first r coordinates to 1.

The set of points in $\mathbb{T}^{t-r}$ where all components have a positive real component is open. Thus, there exists a j such that $\mathfrak{Re}(\beta_i^{Mj}) > 0$ for every i. Since each β_i is of the form $\beta_i = \alpha_i/|\alpha_i|$, we get that $\mathfrak{Re}(\alpha_i^{Mj}) > 0$ for every $0 \neq \alpha \in X$. Set $k = Mj$.

Since every vertex is stable, there is some $0 \neq \alpha \in X_v$ for every v. Since $\operatorname{Tr} A_v^k = a_v(kv)$ for some integer a_v, we must have that $a_v > 0$ for every v, as required. □

5.1.6 The unbounded vertices in cyclic covers lemmas

Lemma 5.7 (unbounded vertices for surface diffeomorphisms) *Let $f \colon \Sigma \to \Sigma$ be a pseudo-Anosov mapping class. Let b be the number of boundary components of Σ. For ever prime p there exists a finite cyclic cover $\pi \colon \widetilde{\Sigma} \to \Sigma$ of degree p with b boundary components and a lift g of some power of f to $\widetilde{\Sigma}$ such that $g_* \colon H_1(\widetilde{\Sigma}, \mathbb{R}) \to H_1(\widetilde{\Sigma}, \mathbb{R})$ has eigenvalues off of the unit circle, or*

$$\# V(\mathcal{S}^e g) \ge p - b - 1.$$

Proof Choose a prime $p > 0$, and a homomorphism $\psi \colon \pi_1(\Sigma) \to \mathbb{Z}/p\mathbb{Z}$ such that every boundary component of Σ is sent to $1 \in \mathbb{Z}/p\mathbb{Z}$. Let $\widetilde{\Sigma}$ be the cover corresponding to $\ker \psi$. Some power of f lifts to the cover $\widetilde{\Sigma}$. By replacing with a further cover, we may assume that it commutes with the deck group.

By a theorem of Kronecker, a matrix $A \in M_n(\mathbb{Z})$ that does not have eigenvalues off of the unit circle has only roots of unity as eigenvalues. If the map induced by the lift of f on $H_1(\widetilde{\Sigma}; \mathbb{R})$ does not have eigenvalues off of the unit circle, we can replace it with a further power such that all of the eigenvalues are 1. By replacing f with the relevant power, we may assume that all of the eigenvalues of $f_*\colon H_1(\widetilde{\Sigma};\mathbb{R}) \to H_1(\widetilde{\Sigma};\mathbb{R})$ are also 1. Call this lift g. By construction, $\widetilde{\Sigma}$ has the same number of boundary components as Σ.

Form the mapping torus M_g. The operator g_* commutes with the action of the deck group on $H_1(\widetilde{\Sigma},\mathbb{R})$. Call this deck group $\mathcal{D}$. For every j, the space $V_j = \ker(g_* - I)^j$ is $\mathcal{D}$–invariant. The number $\dim V_{j+1} - \dim V_j$ is the number of Jordan blocks in the Jordan normal form of g_* of size greater than j.

Let $W = \ker(H_1(\widetilde{\Sigma},\mathbb{R}) \to H_1(\Sigma,\mathbb{R}))$. The spaces $V_j \cap W$ are all $\mathcal{D}$–invariant, and have dimension divisible by $p-1$. Some of the spaces V_j must intersect W nontrivially by assumption. If $\dim(V_j \cap W) = \dim(V_{j+1} \cap W)$, then W must contain 1–eigenvectors. Thus, there must be at least $p-1$ such eigenvectors. If $\dim(V_j \cap W) < \dim(V_{j+1} \cap W)$ for some j, then $\dim(V_{j+1} \cap W) - \dim(V_j \cap W)$ is divisible by p, and there must be at least p Jordan blocks. Thus, we have that the multiplicity of 1 as an eigenvalue of g_* must be at least $p-1$. This proves the result by Lemma 3.19. □

Lemma 5.8 (unbounded vertices for fully irreducible automorphisms) *Let $f \in \mathrm{Aut}(F_n)$ be a train track representative of a fully irreducible automorphism, and let Γ be the corresponding train track graph. For every prime p there exists a finite cyclic cover $\pi\colon \widetilde{\Gamma} \to \Gamma$ of degree p and a lift g of some power of f to $\widetilde{\Gamma}$ such that $g_*\colon H_1(\widetilde{\Gamma},\mathbb{R}) \to H_1(\widetilde{\Gamma},\mathbb{R})$ has eigenvalues off of the unit circle, or*

$$\#V(\mathcal{S}^e g) \geq p-1.$$

Proof Choose a prime $p > 0$, and a nontrivial homomorphism $\psi\colon \pi_1(\Gamma) \to \mathbb{Z}/p\mathbb{Z}$. Let $\widetilde{\Gamma}$ be the cover corresponding to $\ker\psi$. Some power of f lifts to the cover $\widetilde{\Gamma}$. Since this power of f lifts to $\widetilde{\Gamma}$, it normalizes the deck group of $\widetilde{\Gamma} \to \Gamma$. By passing to a further power, we may assume that it commutes with the deck group. If $f^*\colon H_1(\widetilde{\Gamma};\mathbb{R}) \to H_1(\widetilde{\Gamma};\mathbb{R})$ does not have eigenvalues off of the unit circle, we can replace it with a further power such that all of the eigenvalues are 1. Call this lift g. We now proceed in a nearly identical manner to the previous lemma.

Let $\mathcal{D}$ be the deck group of the cover $\widetilde{\Gamma} \to \Gamma$. Form the mapping torus M_g. The operator g_* commutes with the action of $\mathcal{D}$ group on $H_1(\widetilde{\Sigma},\mathbb{R})$. For every j, the space $V_j = \ker(g_* - I)^j$ is $\mathcal{D}$–invariant. The number $\dim V_{j+1} - \dim V_j$ is the number of Jordan blocks in the Jordan normal form of g_* of size greater than j.

Let $W = \ker(H_1(\widetilde{\Sigma},\mathbb{R}) \to H_1(\Sigma,\mathbb{R}))$. The spaces $V_j \cap W$ are all $\mathcal{D}$–invariant, and have dimension divisible by $p-1$. Some of the spaces V_j must intersect W

nontrivially by assumption. If $\dim(V_j \cap W) = \dim(V_{j+1} \cap W)$, then W must contain 1–eigenvectors. Thus, there must be at least $p-1$ such eigenvectors. If $\dim(V_j \cap W) < \dim(V_{j+1} \cap W)$ for some j, then $\dim(V_{j+1} \cap W) - \dim(V_j \cap W)$ is divisible by p, and there must be at least p Jordan blocks. Thus, we have that the multiplicity of 1 as an eigenvalue of g_* must be at least $p-1$. This proves the result by the conclusion of Lemma 3.19 for fully irreducible automorphisms. □

5.2 Completing the proofs

We begin by noting that in the proof of Theorem 1.5, it is enough to prove the theorem for some power of f. Indeed, suppose $K \lhd F_n$ is an f^k–invariant subgroup for some k such that $f_*^k\colon H_1(K;\mathbb{C}) \to H_1(K;\mathbb{Z})$ has eigenvalues off of the unit circle. Let $K' = \bigcap \alpha(K)$, where the intersection is taken over all automorphisms $\alpha \in \mathrm{Aut}(F_n)$. The group K' is characteristic in F_n, and if F_n/K is solvable then so is F_n/K'. Since K' is characteristic, $f(K') = K'$.

The transfer map transfer map $T\colon H_1(K;\mathbb{C}) \to H_1(k';\mathbb{C})$ is f^k–equivariant, and thus $f_*^k\colon H_1(K';\mathbb{C}) \to H_1(K';\mathbb{C})$ has eigenvalues off of the unit circle. The same must then hold for f.

We now prove Theorem 1.5 by dividing it into two cases.

Proposition 5.9 (the pseudo-Anosov case) *Let $n \geq 2$ and let $\bar{f} \in \mathrm{Out}(F_n)$ be the image of a pseudo-Anosov mapping class. Let f be a train track representative of $\bar{f}$. Then there exists a finite-index subgroup $K \lhd F_n$ such that $f(K) = K$, and $f_*\colon H_1(K;\mathbb{Z}) \to H_1(F_n;\mathbb{Z})$ has eigenvalues off of the unit circle. Furthermore, we can choose K such that F_n/K is solvable.*

Proof If $f_*\colon H_1(\Sigma) \to H_1(\Sigma)$ has eigenvalues off of the unit circle, we are done. If not, replace f with a power of itself such that all of the eigenvalues of f_* are 1.

If $\mathcal{S}^e\varphi$ has a nonstable vertex v, then by the stabilization lemma, Lemma 4.17, there exists a solvable cover $\pi\colon \Sigma' \to \Sigma$ to which f lifts such that v is stable in π. Since every solvable cover is the intersection of cyclic covers of prime order, we can find an intermediate cover $\Sigma' \to \Sigma'' \to \Sigma$ to which f lifts such that v is not stable in the cover Σ'' and Σ' is a cyclic cover of Σ'' of prime order. By the cyclic deformation lemma, Lemma 5.1, there exist infinitely many primes p and cyclic p–covers of Σ'' where v is stable. Let $\pi_p\colon \Sigma_p \to \Sigma$ be such a cover corresponding to the prime p.

Note v is not stable in Σ'', and, by Lemma 5.3, the support of $\operatorname{Tr} A_v[\pi_p]^k$ contains an orbit of size p for some k. By the cyclic multiplicity lemma, Lemma 5.2, the coefficient of every point in the support is divisible by p. Thus, the L^2 norm of $\operatorname{Tr} A[\pi_p]^k$ is at least $\sqrt{p \cdot p^2} = p^{3/2}$. Suppose Σ'' is a cover of degree M. Then $\dim H_1(\Sigma_p;\mathbb{R}) \le Mp \dim H_1(\Sigma;\mathbb{R})$. In particular, if we take $p \gg 0$, we get that the L^2 norm of $\operatorname{Tr} A[\pi_p]^k$ is greater than the first Betti number of the corresponding mapping torus. This concludes the proof by the L^2–trace lemma, Lemma 2.12.

The same reasoning holds if v is stable in the cover $\pi'\colon \Sigma' \to \Sigma$, $\pi_p\colon \Sigma_p \to \Sigma'$ is a cyclic p–cover for $p \gg 0$ and the lift of $\mathcal{T}_v$ to π_p is a face graph that is not a vertex graph (because, once again, the trace will contain an orbit of size p).

For any cover $\pi\colon \Sigma' \to \Sigma$ to which f lifts, let d_π be the first Betti number of the mapping torus, and let

$$\gamma(\pi) = \frac{\sup_{k,L} \operatorname{Tr} A[\pi]^k(L)}{d_\pi},$$

where the supremum is taken over all integers k and lattices L.

Let γ_1 be γ of the trivial cover. If $\mathcal{S}^e\varphi$ has any nonstable vertices, we are done. If we can find a vertex v such that for some sufficiently large prime p there is a cyclic cover π_p such that $\mathcal{T}_v[\pi_p]$ is not a vertex subgraph, we are also done. Otherwise, by the unbounded vertices for surface diffeomorphisms lemma, Lemma 5.7, we can find a prime $p \gg 0$ and a cyclic p–cover $\pi_p\colon \Sigma_p \to \Sigma$ to which f lifts to a map f_p such that $\mathcal{S}^e f_p$ is at least $(p-1-b)$–dimensional, where b is the number of boundary components of Σ. Furthermore, Σ_p has the same number of boundary components as Σ.

Every vertex of $\mathcal{S}^e f$ is stable and lifts to a vertex of $\mathcal{S}^e$. As before, if any of the vertices of $\mathcal{S}^e f_p$ is not stable, we are done. Otherwise, by the positive vertices lemma, Lemma 5.6, we can replace f with a power f^k such that the coefficient of each vertex in $A[\pi_p]^k$ is positive. By Lemma 5.3, the coefficient of each vertex that is a lift from a vertex in Σ is divisible by p. The coefficient of every other vertex is at least 1.

By the polytope/lattice lemma, Lemma 5.4, and Corollary 5.5, we can find a lattice L such that $\operatorname{Tr} A[\pi_p]^k \cap L$ consists of a set of at least $p-b$ vertices, $\dim \mathcal{S}^e f + 1$ of which are lifts of vertices of $\mathcal{S}^e f$. Thus,

$$\gamma(\pi_p) \ge \frac{1}{p} \operatorname{Tr} A[\pi_p]^k(L) \ge \gamma_1 + \frac{p-b-\dim(\mathcal{S}^e f)}{p}$$

By taking $p \gg 0$, this can be made arbitrarily close to $\gamma_1 + 1$. Repeat the same argument for π_p. If it has unstable vertices, or vertices that lift to faces in cyclic covers, then we are done. Otherwise, for any $\epsilon > 0$ we can find $p' \gg p$ and a cover $\pi_{p'}$ of Σ_p such that $\gamma(\pi_{p''}) \geq \gamma_1 + 2 - \epsilon$. Iterating this process, we see that the set $\gamma(\pi)$ is unbounded over all solvable covers π. This concludes the proof by the anchoring lemma, Lemma 2.11. □

Proposition 5.10 (the fully irreducible case) *Let $n \geq 2$ and let $\bar{f} \in \mathrm{Out}(F_n)$ be fully irreducible. Let f be a train track representative of $\bar{f}$. Then there exists a finite-index subgroup $K \lhd F_n$ such that $f(K) = K$, and $f_*\colon H_1(K;\mathbb{Z}) \to H_1(F_n;\mathbb{Z})$ has eigenvalues off of the unit circle. Furthermore, we can choose K such that F_n/K is solvable.*

Proof The proof here is nearly identical to the proof of Proposition 5.9 except that instead of using the unbounded vertices lemma for surface diffeomorphisms, we use the unbounded vertices for fully irreducible automorphisms lemma, Lemma 5.8. □

We now need only to deduce Theorem 1.3 from Theorem 1.5.

Proof To conclude the proof of Theorem 1.3, we now need only address the case that f has positive topological entropy but is not a pseudo-Anosov mapping class. By the Nielsen–Thurston classification, we can replace f with a power of itself such that there exists a subsurface $\Sigma' \subset \Sigma$ such that Σ' is f–invariant and f restricted to Σ' is a pseudo-Anosov mapping class. Pick a basepoint $* \in \Sigma'$.

By a theorem of Marshall Hall [8], there exists a finite-index subgroup $K < \pi_1(\Sigma, *)$ such that $\pi_1(\Sigma', *) \leq K$, and $\pi_1(\Sigma', *)$ is a free subfactor of K. Replace f with a power that fixes the subgroup K. The theorem now follows from applying Theorem 1.5 to $f|_{K'}$, and noting that $H_1(\Sigma';\mathbb{C})$ injects into $H_1(K;\mathbb{C})$. □

References

[1] **Y Algom-Kfir**, **E Hironaka**, **K Rafi**, *Digraphs and cycle polynomials for free-by-cyclic groups*, Geom. Topol. 19 (2015) 1111–1154 MR

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

[3] **S Dowdall**, **I Kapovich**, **C J Leininger**, *McMullen polynomials and Lipschitz flows for free-by-cyclic groups*, J. Eur. Math. Soc. 19 (2017) 3253–3353 MR

[4] **A Fathi**, **F Laudenbach**, **V Poénaru** (editors), *Travaux de Thurston sur les surfaces*, Astérisque 66–67, Soc. Math. France, Paris (1979) MR

[5] **D Fried**, *The geometry of cross sections to flows*, Topology 21 (1982) 353–371 MR

[6] **F Grunewald**, **M Larsen**, **A Lubotzky**, **J Malestein**, *Arithmetic quotients of the mapping class group*, Geom. Funct. Anal. 25 (2015) 1493–1542 MR

[7] **A Hadari**, *Every infinite order mapping class has an infinite order action on the homology of some finite cover*, preprint (2015) arXiv

[8] **M Hall, Jr**, *Coset representations in free groups*, Trans. Amer. Math. Soc. 67 (1949) 421–432 MR

[9] **T Koberda**, *Asymptotic linearity of the mapping class group and a homological version of the Nielsen–Thurston classification*, Geom. Dedicata 156 (2012) 13–30 MR

[10] **T Koberda**, *Residual properties of fibered and hyperbolic 3–manifolds*, Topology Appl. 160 (2013) 875–886 MR

[11] **T Koberda**, **J Mangahas**, *An effective algebraic detection of the Nielsen–Thurston classification of mapping classes*, J. Topol. Anal. 7 (2015) 1–21 MR

[12] **A Lubotzky**, **C Meiri**, *Sieve methods in group theory, II: The mapping class group*, Geom. Dedicata 159 (2012) 327–336 MR

[13] **A Lubotzky**, **C Meiri**, *Sieve methods in group theory, III:* $\mathrm{Aut}(F_n)$, Int. J. Algebra Comput. 22 (2012) art. id. 1250062 MR

[14] **J Malestein**, **J Souto**, *On genericity of pseudo-Anosovs in the Torelli group*, Int. Math. Res. Not. 2013 (2013) 1434–1449 MR

[15] **C T McMullen**, *Entropy on Riemann surfaces and the Jacobians of finite covers*, Comment. Math. Helv. 88 (2013) 953–964 MR

[16] **A Putman**, **B Wieland**, *Abelian quotients of subgroups of the mappings class group and higher Prym representations*, J. Lond. Math. Soc. 88 (2013) 79–96 MR

[17] **T Sakasai**, *A survey of Magnus representations for mapping class groups and homology cobordisms of surfaces*, from "Handbook of Teichmüller theory, III" (A Papadopoulos, editor), IRMA Lect. Math. Theor. Phys. 17, Eur. Math. Soc., Zürich (2012) 531–594 MR

[18] **M Suzuki**, *Geometric interpretation of the Magnus representation of the mapping class group*, Kobe J. Math. 22 (2005) 39–47 MR

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

Department of Mathematics, University of Hawaii
Honolulu, HI, United States

hadari.asaf@gmail.com

Proposed: Étienne Ghys
Seconded: Martin Bridson, Benson Farb

Received: 11 December 2017
Revised: 25 September 2019