


Geometry and Arithmetic of Pseudo-Anosov Stretch Factors




Thèse d'habilitation présentée à la Faculté des sciences et de médecine de
l'Université de Fribourg (Suisse) pour l'obtention de la *venia legendi* par le

Dr. Livio Liechti
de Landiswil (Suisse)

Fribourg, 2025

Acceptée par la Faculté des sciences et de médecine de l'Université de Fribourg (Suisse)



Prof. Michael Walch
Doyen
Faculté des sciences et de médecine
Université de Fribourg – Suisse

Fribourg, le 22 septembre 2025

Contents

| | |
|--|-----|
| Preface | iii |
| Acknowledgements | v |
| Teaser | vii |
| Chapter 1. Introduction | 1 |
| 1. Algebraic degrees of stretch factors | 3 |
| 2. Algebraic characterisation of stretch factors | 9 |
| 3. Spectral radii of integer matrices | 11 |
| 4. Minimal stretch factors | 14 |
| Organisation | 25 |
| Chapter 2. Stretch factor degrees | 27 |
| 1. Irreducibility criteria | 27 |
| 2. Nonsplitting criteria | 31 |
| 3. Warm-up: proof of Thurston's claim | 34 |
| 4. Strata of Abelian differentials | 39 |
| 5. Connected components of strata | 46 |
| 6. Proof of Theorem 1.3 | 59 |
| 7. The case $g = 2$ of Theorems 1.1 and 1.2 | 69 |
| 8. Explicit pseudo-Anosov maps | 70 |
| Chapter 3. Algebraic characterisation of stretch factors | 77 |
| 1. Perron and bi-Perron numbers | 77 |
| 2. Coxeter transformations and the Thurston–Veech construction | 83 |
| 3. Galois conjugates of bi-Perron numbers | 87 |
| Chapter 4. Spectral radii of integer matrices | 89 |
| 1. The clique polynomial | 89 |
| 2. Skew-reciprocity and minimal spectral radii | 91 |
| 3. A reformulation of the question of Schinzel and Zassenhaus | 103 |
| 4. Symplectic vs. antisymplectic matrices | 105 |

| | |
|--|-----|
| Chapter 5. Minimal stretch factors on nonoriented surfaces | 107 |
| 1. Construction of pseudo-Anosov maps on nonorientable surfaces | 107 |
| 2. Orientation-reversing pseudo-Anosov mapping classes on odd genus surfaces | 114 |
| 3. Restrictions on polynomials | 118 |
| 4. Elimination of polynomials | 123 |
| Chapter 6. Minimal stretch factors in Penner's construction | 133 |
| 1. Nonorientable surfaces | 133 |
| 2. Stretch factor theory of the cycle | 135 |
| 3. Even genus minimal stretch factors | 141 |
| 4. Stretch factor theory of the enriched cycle | 145 |
| 5. Odd genus minimal stretch factors | 161 |
| Chapter 7. Minimal stretch factors in the fully-punctured case | 165 |
| 1. Background | 165 |
| 2. Skew-reciprocity up to cyclotomic factors | 169 |
| 3. Analysis of small curve graphs | 174 |
| 4. Sharpness of Theorem 1.29 | 188 |
| 5. Proof of Theorem 1.31 | 191 |
| 6. Discussion and further questions | 191 |
| Bibliography | 195 |

Preface

This thesis comprises some of the results I have obtained since completing my PhD studies in 2017. For the first sixteen months after my PhD, I have been supported by an SNSF Early Postdoc.Mobility grant, with grant number 175260 and project title *Geometric and homological dilatation of surface homeomorphisms*. Afterwards, I have been employed at the University of Fribourg, first as a postdoc and then as a Maître-Assistant.

This habilitation thesis contains the content of eight of my articles which are available in published form or accepted for publication. The content of Chapter 2 is taken from the joint published article [30] with Erwan Lanneau, as well as the joint preprint [31] with Erwan Lanneau. The content of Chapter 3 is taken from the joint published article [38] with Joshua Pankau. The content of Chapter 4 is taken from my published articles [36, 37]. The content of Chapter 5 and Chapter 6 is taken from the joint published articles [39, 41] with Balázs Strenner. Finally, the content of Chapter 7 is taken from the joint published article [32] with Erwan Lanneau and Chi Cheuk Tsang.

Acknowledgements

This habilitation thesis would not have been possible without a great deal of inspiration and advice throughout the years. I would like to thank all the people who influenced me, knowingly or not, by their mathematics, their advice or by their handling of the particularities of pursuing an academic path. I particularly thank Sebastian Baader, Eriko Hironaka, Ruth Kellerhals and Julien Marché.

This habilitation thesis would not have been possible either without my collaborators. I would like to thank them all, and especially those whose joint work with me found its way into this habilitation thesis: Erwan Lanneau, Joshua Pankau, Balázs Strenner and Chi Cheuk Tsang.

I would like to thank all people who offered feedback and comments to my work. Particularly, I thank Dan Margalit and Curt McMullen as well as the anonymous referees for their many comments on the articles whose content is contained in this habilitation thesis.

Even though knot theory features only tangentially in this habilitation thesis, the Swiss knots community is the place where I am most at home mathematically. Thank you everyone!

I have worked in Fribourg the majority of time after finishing my PhD. I would like to thank everyone at the Math department for the good time as well as the trust and the support throughout the years.

Finally, I thank my family, and in particular Matteo and Enea for teaching me more than I could have ever imagined.

Teaser

Let $A \in \text{GL}_2(\mathbb{Z})$ be a matrix, acting on the real plane \mathbb{R}^2 by a linear isomorphism. Since this action preserves the integer lattice $\mathbb{Z}^2 \subset \mathbb{R}^2$, it descends to a diffeomorphism $\psi_A : \mathbb{R}^2/\mathbb{Z}^2 \rightarrow \mathbb{R}^2/\mathbb{Z}^2$ of the torus $\mathbb{R}^2/\mathbb{Z}^2$, defined by

$$\psi_A(v + \mathbb{Z}^2) = Av + \mathbb{Z}^2.$$

The dynamics of such a map ψ_A are largely dictated by the eigenvalues of the matrix A , that is, the zeroes of its characteristic polynomial

$$\chi_A(t) = t^2 - \text{tr}(A)t + \det(A),$$

which are algebraic integers of degree one or two. There are three possibilities for the eigenvalues of A :

- (1) if A has two distinct eigenvalues λ_1, λ_2 on the unit circle, then λ_1 and λ_2 are roots of unity. In particular, the diffeomorphism ψ_A is a periodic map. Since the only roots of unity that are algebraic integers of degree at most two are the first, second, third, fourth and sixth roots of unity, the map ψ_A has order one, two, three, four or six;
- (2) if A has a double eigenvalue λ on the unit circle, then $\lambda = +1$ or $\lambda = -1$. In this case, A is not necessarily of finite order, for example if A is a Jordan block of size two. However, it follows from the eigenvector equation $Av = \lambda v$ that A possesses an eigenvector with rational slope r , which in turn defines a simple closed curve on the torus $\mathbb{R}^2/\mathbb{Z}^2$, by projecting the line $L \subset \mathbb{R}^2$ of slope r to the torus $\mathbb{R}^2/\mathbb{Z}^2$. The map ψ_A fixes the simple closed curve on the torus and is therefore called *reducible*;
- (3) the third possibility is for A to have a pair of irrational real eigenvalues of inverse moduli. In this case, the map ψ_A is called *Anosov*. Such maps have intriguing dynamical properties:
 - (a) the set of periodic points of $\psi_A : \mathbb{R}^2/\mathbb{Z}^2 \rightarrow \mathbb{R}^2/\mathbb{Z}^2$ is a dense subset of the torus $\mathbb{R}^2/\mathbb{Z}^2$. Indeed, the set of all rational points

of $\mathbb{R}^2/\mathbb{Z}^2$ with a fixed denominator q is preserved by ψ_A , and hence every such point is periodic of period at most q^2 . Taking the union of these sets for all denominators provides a dense set of periodic points;

- (b) there exists a dense orbit for the map ψ_A . Indeed, since the eigenvalues λ_1 and λ_2 of A are irrational, so are the slopes of the associated eigendirections L_1 and L_2 in \mathbb{R}^2 . Intuitively, the denseness of the projected eigendirection of A means that iterations of the map ψ_A smear any open set of the torus $\mathbb{R}^2/\mathbb{Z}^2$ more and more densely across the torus. More precisely, using compactness of the torus we can show that for any two open subsets $U, V \subset \mathbb{R}^2/\mathbb{Z}^2$, there exists a positive integer $N \in \mathbb{N}$ such that for all powers $n > N$, the intersection of $\psi_A^n(U)$ with V is nontrivial. This property is called *topological mixing* and a standard application of the Baire category theorem then provides a dense forward orbit of $\psi_A : \mathbb{R}^2/\mathbb{Z}^2 \rightarrow \mathbb{R}^2/\mathbb{Z}^2$.

Properties (a) and (b) together are an indicator of *chaos* of a dynamical system: an ever so tiny difference in the starting condition of the dynamical system described by such a map ψ_A can make the difference between lying on a periodic or a dense orbit—two completely different dynamical situations. In particular, if the map ψ_A described an actual physical dynamical system, it would be effectively impossible to determine by measurement whether the system evolves on a periodic or a dense orbit, simply because every measurement is made with an error.

For the three types of maps ψ_A induced by matrices $A \in \text{GL}_2(\mathbb{Z})$, the Anosov case turns out to be the most dynamically rich. The modulus of the eigenvalues which is greater than one is called the *stretch factor* of the Anosov map and is crucial for the behaviour of the map.

Since stretch factors of Anosov maps are very specific quadratic algebraic integers, we are able to readily answer all questions about them. For example, we clearly see that each stretch factor of an Anosov map is a real quadratic algebraic unit > 1 . By describing explicit matrices in $\text{GL}_2(\mathbb{Z})$ of all possible traces and determinant ± 1 , this becomes an exact characterisation of the Anosov stretch factors: they are exactly the real quadratic algebraic units > 1 .

The stretch factor of an Anosov map ψ_A corresponds to the amount of stretching along the expanding eigendirection. In particular, it is natural to think that the simplest Anosov maps are the ones with smallest stretch factors. Again, since Anosov stretch factors are real quadratic algebraic units > 1 , it is straightforward to determine the minimal ones:

- among orientation-preserving Anosov maps, that is, Anosov maps obtained from matrices A with determinant $+1$, the stretch factor is an increasing function of the modulus of the trace, and the minimal such value is three. This implies that the minimal stretch factor is the square of the golden ratio

$$\frac{3 + \sqrt{5}}{2} = \varphi^2;$$

- among orientation-reversing Anosov maps, that is, Anosov maps obtained from matrices A with determinant -1 , the stretch factor is again an increasing function of the modulus of the trace, only this time the minimal such value is one. In particular, the minimal stretch factor in this case is the golden ratio

$$\frac{1 + \sqrt{5}}{2} = \varphi.$$

We have completely characterised the algebraic integers that are Anosov stretch factors, we have determined the minimal stretch factors both in the orientation-preserving and in the orientation-reversing case and we have seen that the minimal stretch factor in the orientation-reversing case is smaller than in the orientation-preserving case.

There are two ways to generalise the setting we have just discussed.

The first and fairly obvious generalisation is to consider matrices of larger size. For example, one could consider the spectral radii of matrices in $\mathrm{GL}_n(\mathbb{Z})$ instead of $\mathrm{GL}_2(\mathbb{Z})$. The setting becomes much more difficult than in dimension two, and even the speed of convergence to 1 of the minimal spectral radius $\rho_n > 1$, as $n \rightarrow \infty$, was unknown for a long time. Indeed, in 1965, Schinzel and Zassenhaus asked [56] whether there exists a constant c such that $\rho_n > 1 + \frac{c}{n}$, a question that remained open until Dimitrov's short and beautiful number-theoretic positive answer from the very last days of the year 2019 [12], see also the description in Chapter 4 of the book by McKee and Smyth [43].

By previous work of Breusch [7] from 1951, the conjecture of Schinzel and Zassenhaus reduces to the study of spectral radii of symplectic matrices in $\mathrm{Sp}_{2g}(\mathbb{Z})$. This class of matrices presents another generalisation of the situation above. Indeed, in the two-dimensional case, we have $\mathrm{SL}_2(\mathbb{Z}) = \mathrm{Sp}_2(\mathbb{Z})$. In the context of symplectic matrices, the determinant is always positive. To have an analogue of the full picture of $\mathrm{GL}_2(\mathbb{Z})$, we can contrast symplectic matrices with *antisymplectic matrices*, that is, matrices which reverse the standard symplectic form. We will see in this habilitation thesis that the comparison of these two types of matrices provides a reformulation of the question of Schinzel and Zassenhaus that is close to the comparison of orientation-preserving and orientation-reversing diffeomorphisms of surfaces. Furthermore, we will see that for the restricted classes of primitive matrices or irreducible matrices, such a comparison holds and can be proven without referring to Dimitrov’s theorem. Maybe one day new methods might allow this strategy to provide a more geometrically inspired proof of Dimitrov’s theorem.

The second generalisation is to consider an analogue of Anosov diffeomorphisms of more complicated surfaces than the torus. While the torus is the only closed orientable surface on which Anosov diffeomorphisms exist, there exists a notion of *pseudo-Anosov maps* for closed surfaces of higher genus, with very similar dynamical properties as Anosov maps.

We now give a definition of pseudo-Anosov maps. This definition is technical, and it can safely be skipped for the sake of this teaser, keeping the Anosov diffeomorphisms ψ_A of the torus in mind as an example that we would like to generalise. For a much more in depth treatment of pseudo-Anosov maps, we refer to the original sources [65, 16] as well as the ‘Primer’ [14].

A homeomorphism f of a surface S of finite type is *pseudo-Anosov* if there exists a pair of singular transverse measured foliations (\mathcal{F}^s, μ^s) and (\mathcal{F}^u, μ^u) of S such that:

- (1) away from a finite collection of *singular points*, which includes the potential punctures of S , \mathcal{F}^s and \mathcal{F}^u are locally conjugate to the foliations of \mathbb{R}^2 by vertical and horizontal lines, respectively;
- (2) near a singular point, \mathcal{F}^s and \mathcal{F}^u are locally conjugate to either
 - the pullback of the foliations of \mathbb{R}^2 by vertical and horizontal lines by the map $z \mapsto z^{\frac{n}{2}}$, respectively, for some $n \geq 3$, or

- the pullback of the foliations of $\mathbb{R}^2 \setminus \{(0,0)\}$ by vertical and horizontal lines by the map $z \mapsto z^{\frac{n}{2}}$, respectively, for some positive integer $n \geq 1$;
- (3) the map f preserves the leaves of the foliations \mathcal{F}^s and \mathcal{F}^u , and induces the following action on the transverse measures:

$$f_*(\mathcal{F}^s, \mu^s) = (\mathcal{F}^s, \lambda^{-1}\mu^s) \text{ and } f_*(\mathcal{F}^u, \mu^u) = (\mathcal{F}^u, \lambda\mu^u)$$

for some positive real number $\lambda = \lambda(f) > 1$.

The measured foliations (\mathcal{F}^s, μ^s) and (\mathcal{F}^u, μ^u) are the *stable* and *unstable* measured foliations, respectively. The number $\lambda(f)$ is the *stretch factor* of f , sometimes also called *dilatation* or *expansion factor*.

As in properties (a) and (b) for Anosov maps ψ_A , the set of periodic points of a pseudo-Anosov map is dense and there exist dense orbits. Pseudo-Anosov maps also come together with a stretch factor that determines many dynamical aspects of the maps. Is it possible to characterise the numbers which arise as stretch factors in this more general case? Or to determine the minimal stretch factors for every fixed surface? Or to compare the minimal stretch factors in the orientation-preserving and the orientation-reversing setting? In general, these questions turn out to be very intricate. In one direction, however, we arrive at a conclusive result in the last chapter of this habilitation thesis: if we fix an even Euler characteristic for orientable surfaces, then in a special case that we call *fully punctured with at least two puncture orbits*, orientation-reversing pseudo-Anosov maps realise smaller stretch factors than orientation-preserving ones. This can be seen as a geometric analogue of our reformulation of Dimitrov's theorem.

CHAPTER 1

Introduction

Surface homeomorphisms are the natural symmetries of surfaces. For a fixed surface, these symmetries up to isotopy form a group, where the group law is the composition of maps. This group is called the *mapping class group* and its elements are called *mapping classes*. Its study was initiated by Dehn and Nielsen in the early 20th century, and revolutionised in the 70s by Thurston's classification of surface mapping classes: every mapping class that is not periodic or reducible is pseudo-Anosov [65].

By Thurston's work, a mapping class is pseudo-Anosov if it asymptotically (under iteration) stretches every isotopy class of essential simple closed curves by its stretch factor λ , which turns out to be an algebraic integer [65]. By work of Hubbard and Masur, to each pseudo-Anosov map is associated a quadratic differential on a Riemann surface, describing invariant stable and unstable measured foliations for the map [24].

Orientation-preserving pseudo-Anosov mapping classes play a crucial role in the arithmetic and geometry of moduli spaces. The moduli space \mathcal{M}_g of genus g Riemann surfaces is a central object in modern mathematics. As a set, it is defined to be the set of Riemann surfaces of genus g , up to biholomorphism. The importance of \mathcal{M}_g is evident already for $g = 1$: by the uniformisation theorem, \mathcal{M}_1 equals the set of complex tori \mathbb{C}/Λ , up to biholomorphism. In turn, by the classical embedding based on the Weierstrass \wp -function, this equals the set of nonsingular elliptic curves in the complex projective plane, up to projective transformation.

There are various ways to topologise moduli spaces—however, one can metrize them right away using Teichmüller's theorems. These state that in every isotopy class of homeomorphisms between two Riemann surfaces, there exists a unique quasiconformal representative with least quasiconformal dilatation [64]. Teichmüller's theorems can be used to define a metric on Teichmüller space \mathcal{T}_g , and then in turn to define a metric on \mathcal{M}_g , which is the quotient of \mathcal{T}_g by the action of the mapping class group of the closed genus g surface. A geodesic in this metric is described via a quadratic

differential by exponentially stretching one of its associated foliations and shrinking the other. The crucial point is that if one comes back to the same point in Moduli space along such a geodesic, then by covering theory there is an element in the mapping class group that realises the corresponding deck transformation in Teichmüller space. This element is an orientation-preserving pseudo-Anosov map with stretch factor λ , and the length of the closed geodesic is $\log(\lambda)$.

The upshot is that in the case of moduli space \mathcal{M}_g equipped with the Teichmüller metric, one of the very classic geometric endeavors, namely the comprehension of the length spectrum, is equivalent to understanding stretch factors of orientation-preserving pseudo-Anosov maps.

The foundational result for the arithmetic and geometric aspects of pseudo-Anosov stretch factors that we consider in this habilitation thesis is due to Thurston from the 1980s: the stretch factor of a pseudo-Anosov map on an orientable surface of finite type is an algebraic integer of degree bounded from above by the dimension of the Teichmüller space for the surface [65].

It is a natural question to inquire about the exact set of algebraic degrees realised by pseudo-Anosov stretch factors on any given surface—possibly under additional constraints or using only certain types of pseudo-Anosov maps. Here, the algebraic degree of the stretch factor λ is simply the degree of the field extension $\mathbb{Q}(\lambda) : \mathbb{Q}$. This degree is called the *stretch factor degree*.

The second natural question is to seek a characterisation of the set of all pseudo-Anosov stretch factors. Fried showed in 1985 that every stretch factor is a bi-Perron number, where a *bi-Perron number* λ is a real algebraic unit > 1 all of whose Galois conjugates have modulus in the open interval (λ^{-1}, λ) , except for λ itself and possibly one of $\pm\lambda^{-1}$. Furthermore, Fried asked whether pseudo-Anosov stretch factors are, up to powers, exactly the bi-Perron numbers [19].

The fact that for a given surface the set of pseudo-Anosov stretch factors consists of bi-Perron numbers of bounded degree also implies that there exists a minimal pseudo-Anosov stretch factor on every fixed surface. Indeed, the set of bi-Perron numbers of bounded degree is a discrete subset of the real numbers \mathbb{R} . In particular, this implies that for every surface, there exists a smallest pseudo-Anosov stretch factor. In the orientation-preserving

case, this stretch factor has a very geometric interpretation: as we have discussed above, it corresponds to the length of the shortest closed geodesic in a moduli space of Riemann surfaces equipped with the Teichmüller metric.

Finally, it is possible to study a more algebraic version of the question of finding minimal stretch factors. Indeed, stretch factors are Perron–Frobenius eigenvalues of nonnegative primitive matrices, and it is natural to study the minimisation problem in this purely algebraic context; compare with the question of Schinzel and Zassenhaus discussed in the teaser.

In the following, we describe in more detail the four facets of stretch factors we study this habilitation thesis: their algebraic degrees, their algebraic characterisation, the minimisation problem on a given surface and algebraic analogues of minimisation.

1. Algebraic degrees of stretch factors

From the arithmetic perspective, perhaps the most fundamental question about stretch factors concerns their algebraic degree. For simplicity, we choose to deal with closed orientable surfaces in this context. Recall that it is known that the stretch factor of a pseudo-Anosov mapping class on the closed orientable surface of genus g is an algebraic integer of degree at most $6g - 6$. This result is from the 1980s and is due to Thurston [65].

It was not until 2017 that Strenner precisely determined the set of numbers appearing as the degrees of pseudo-Anosov stretch factors on every given closed orientable surface [62]: it consists of all even numbers up to $6g - 6$ and all odd numbers > 1 and up to $3g - 3$.

Apart from the field $\mathbb{Q}(\lambda)$, there is another field which plays a central role in this context: $\mathbb{Q}(\lambda + \lambda^{-1})$. This field is called the *trace field* and is uniquely determined by the quadratic differential associated with the pseudo-Anosov map, or, equivalently, the pair of invariant foliations $(\mathcal{F}^u, \mathcal{F}^s)$, see [26, 20]. The degree of the field extension $\mathbb{Q}(\lambda + \lambda^{-1}) : \mathbb{Q}$ is called the *trace field degree*, and is bounded from above by $3g - 3$.

Stretch factor degrees and trace field degrees are closely related. Since λ is a root of the polynomial $t^2 - (\lambda + \lambda^{-1})t + 1$, the degree of λ over $\mathbb{Q}(\lambda + \lambda^{-1})$ is either one or two. The degree one case corresponds to pseudo-Anosov maps with vanishing SAF invariant by a result of Calta and Schmidt [10]. Finding such examples is difficult, but a large class is given by the lifts of pseudo-Anosov maps on nonorientable surfaces to the orientable double cover; this is a result due to Strenner [63].

Formulated in terms of the trace field, Strenner [62] showed the optimal result that all integers $1 \leq d \leq 3g - 3$ are realised as the trace field degree for a pseudo-Anosov map on a closed orientable surface of genus g .

We inquire about two specifications of Strenner's result: the first is to study trace field degrees and stretch factor degrees in the Torelli groups, the second is to study trace field degrees and stretch factor degrees when restricting the pseudo-Anosov maps to those with associated quadratic differential of specific types. These results are obtained in joint work with Lanneau [30, 31].

1.1. Trace field degrees in Torelli groups. The *Torelli group* $\mathcal{I}(S_g)$ is the kernel of the symplectic (or homological) representation of the mapping class group $\text{Mod}(S_g)$. In [42, Problem 10.6], Margalit asks which stretch factor degrees arise for pseudo-Anosov mapping classes in the Torelli group. Our first main result completely answers the question of trace field degrees arising in Torelli groups.

THEOREM 1.1 ([31], Theorem 1). *Every integer $1 \leq d \leq 3g - 3$ arises as the trace field degree of a pseudo-Anosov mapping class in the Torelli group $\mathcal{I}(S_g)$, for every genus $g \geq 2$.*

We obtain a result on stretch factor degrees in the Torelli groups by showing that for all trace field degrees there exist instances where the field extension $\mathbb{Q}(\lambda) : \mathbb{Q}(\lambda + \lambda^{-1})$ has degree two.

THEOREM 1.2 ([31], Theorem 2). *Every even integer $2 \leq 2d \leq 6g - 6$ arises as the stretch factor degree of a pseudo-Anosov mapping class in the Torelli group $\mathcal{I}(S_g)$, for every genus $g \geq 2$.*

1.2. The Thurston–Veech construction. We prove Theorem 1.1 and Theorem 1.2 using the Thurston–Veech construction¹. In the generality we state it, this construction of pseudo-Anosov maps appeared independently at roughly the same time in two papers by Thurston and Veech [65, 68].

A *multicurve* is a disjoint union of simple closed curves, and a pair of multicurves $\alpha, \beta \subset S_g$ *fills* the closed surface S_g if α and β intersect transversally and if the complement $S_g \setminus (\alpha \cup \beta)$ is a union of topological discs none of which is a bigon. This in particular implies that each pair α_i

¹In the literature, this construction is often attributed to Thurston alone and therefore often called *Thurston's construction*.

and β_j of components realises the minimal number of intersection points within their respective isotopy classes.

Given a pair of filling multicurves $\alpha, \beta \subset S_g$, the Thurston–Veech construction provides pseudo-Anosov mapping classes in the subgroup $\langle T_\alpha, T_\beta \rangle$ of $\text{Mod}(S_g)$ generated by multitwists along the multicurves α and β . Here, the *multitwist* T_α is a composition of positive Dehn twists along all the components of α . For the precise statement of the Thurston–Veech construction, let $X = (|\alpha_i \cap \beta_j|)_{ij}$ be the matrix encoding the number of intersections of the components of α and β . Furthermore, let r^2 be the Perron–Frobenius eigenvalue of the primitive matrix XX^\top . The Thurston–Veech construction [65, 68, 14] states that there exists a representation

$$\rho : \langle T_\alpha, T_\beta \rangle \rightarrow \text{PSL}_2(\mathbb{R})$$

mapping T_α to the matrix $\begin{pmatrix} 1 & r \\ 0 & 1 \end{pmatrix}$ and mapping T_β to the matrix $\begin{pmatrix} 1 & 0 \\ -r & 1 \end{pmatrix}$. Furthermore, $f \in \langle T_\alpha, T_\beta \rangle$ is pseudo-Anosov if and only if $\rho(f)$ has spectral radius > 1 , and in this case the stretch factor $\lambda(f)$ of f equals exactly the spectral radius of $\rho(f)$.

In his seminal 1988 Bulletin paper [65], Thurston provides the upper bound of $6g - 6$ on the algebraic degree of a pseudo-Anosov stretch factor $\lambda(f)$ and claims, without proof, that “*the examples of [65, Theorem 7] show that this bound is sharp*”. The referenced examples are exactly the pseudo-Anosov maps in $\langle T_\alpha, T_\beta \rangle$ obtained via the Thurston–Veech construction.

Margalit remarked in 2011 what Strenner wrote down in his article on stretch factor degrees [62], namely that no proof of Thurston’s claim has ever been published. In recent joint work with Lanneau [31], we are finally able to substantiate Thurston’s claim. Furthermore, we can even do so for pseudo-Anosov maps in the Torelli group. Our precise statement for the Thurston–Veech construction is the following.

THEOREM 1.3 ([31], Theorem 3). *Let $g \geq 2$ and $1 \leq d \leq 3g - 3$ be integers. Then there exists a pseudo-Anosov map on S_g arising from the Thurston–Veech construction with trace field degree d and stretch factor degree $2d$. For $g \geq 3$, the pseudo-Anosov maps can be chosen in the Torelli group $\mathcal{I}(S_g)$.*

Clearly, Theorem 1.3 implies Theorem 1.1 and Theorem 1.2 for $g \geq 3$. Our proof of the Torelli case of Theorem 1.3 does not work for $g = 2$, and

for this situation we directly prove Theorem 1.1 and Theorem 1.2 by using ad hoc examples, see Section 7 of Chapter 2.

Proof strategy for Theorem 1.3. For a pair $\alpha, \beta \subset S_g$ of filling multicurves, let $X = (|\alpha_i \cap \beta_j|)_{ij}$ be the matrix encoding the number of intersections of the components of α and β .

The matrix XX^\top is primitive, hence by Perron–Frobenius theory its spectral radius equals its largest eigenvalue and is therefore an algebraic integer. Let d be its algebraic degree. We call the number d the *multicurve intersection degree* of α and β .

Our proof is based on the following existence result.

THEOREM 1.4 ([31], Theorem 4). *Let $\alpha, \beta \subset S$ be a pair of filling multicurves having multicurve intersection degree d . For $\varepsilon \in \mathbb{Z} \setminus \{0\}$, there exists $n \in \mathbb{Z}_{>0}$ such that the mapping class $T_\alpha^n \circ T_\beta^{n\varepsilon}$ is pseudo-Anosov with stretch factor λ of degree $2d$.*

Assuming Theorem 1.4, what remains to be done in order to prove the first part of Theorem 1.3 is to construct all multicurve intersection degrees $1 \leq d \leq 3g - 3$ on S_g for $g \geq 2$. By the Thurston–Veech construction, the trace field degree of the resulting pseudo-Anosov maps equals exactly the multicurve intersection degree of the multicurves α and β used in the construction [65, 68]. We are done by setting $\varepsilon = 1$ in Theorem 1.4.

In order to prove the Torelli part of Theorem 1.3, we construct multicurves α and β realising the multicurve intersection degrees $1 \leq d \leq 3g - 3$ in such a way that $T_\alpha \circ T_\beta^{-1}$ is an element of $\mathcal{I}(S_g)$. We will ensure this by choosing multicurves α and β which consist of components that are separating or that come in bounding pairs, where for each bounding pair one of the curves is a component of α and the other is a component of β . We can then finish the proof of Theorem 1.3 by setting $\varepsilon = -1$ in Theorem 1.4. To see this, note that if $T_\alpha \circ T_\beta^{-1}$ is an element of $\mathcal{I}(S_g)$, then so is $T_\alpha^n \circ T_\beta^{-n}$.

We note that the types of examples of multicurves α and β we use in order to construct examples in the Torelli group $\mathcal{I}(S_g)$ cannot yield multicurve intersection degrees greater than one on the surface S_2 . Indeed, there exist no bounding pairs on S_2 and a multicurve can have at most one separating component.

Finding the right multicurves α and β is the main technical challenge remaining to prove Theorem 1.3. This is done in Chapter 2.

1.3. A natural field extension from the curve perspective. The arguments we need in order to construct certain trace field degrees or stretch factor degrees crucially depend on finding suitable pairs of multicurves. Actually, there is an even more natural field extension that we can associate with the geometric situation.

A pair of filling multicurves $\alpha, \beta \subset S_g$ naturally determines a bipartite graph whose vertices correspond to curve components and the number of edges between each pair of vertices equals the number of intersection points of the respective curve components. The adjacency matrix of this bipartite graph is $\Omega = \begin{pmatrix} 0 & X \\ X^\top & 0 \end{pmatrix}$. Clearly, the square root $\sqrt{\mu}$ of the spectral radius μ of the matrix XX^\top equals the spectral radius of Ω . We call the algebraic degree of $\sqrt{\mu}$ the *multicurve bipartite degree* of α and β . Similarly to the field extension $\mathbb{Q}(\lambda) : \mathbb{Q}(\lambda + \lambda^{-1})$, also the field extension

$$\mathbb{Q}(\sqrt{\mu}) : \mathbb{Q}(\mu) = \mathbb{Q}(\lambda + \lambda^{-1})$$

has degree one or two.

In this context, we prove the following result.

THEOREM 1.5 ([**31**], Theorem 5). *Every even integer $2 \leq 2d \leq 6g - 6$ is realised as a multicurve bipartite degree on S_g for $g \geq 2$.*

1.4. Trace field degrees of Abelian differentials. We now restrict our study to pseudo-Anosov maps whose associated quadratic differential is the square of an Abelian differential. This corresponds to transversely orientable invariant foliations for the pseudo-Anosov map. The set of nonzero Abelian differentials admits a stratification, where a stratum is the set of Abelian differentials having prescribed singularity multiplicities (k_1, \dots, k_n) , where the sum of multiplicities satisfies $\sum_{i=1}^n k_i = 2g - 2$. Moreover, the Abelian differentials form a natural topological space, and it turns out that certain strata have multiple connected components, which in 2003 were classified by Kontsevich and Zorich [**27**]. In this context, we obtain a stronger result justifying Thurston's claim in every connected component of every stratum of Abelian differentials, not just for a given genus.

THEOREM 1.6 ([**30**], Theorem 1). *Every even integer $2 \leq 2d \leq 2g$ is realised as the stretch factor degree of a product of two affine multitwists on a surface in every connected component of every stratum of Abelian differentials on Riemann surfaces of genus g .*

We will deduce Theorem 1.6 from the following result asserting that choosing a connected component of a stratum of Abelian differentials poses no restriction on the degree of trace fields.

THEOREM 1.7 ([30], Theorem 2). *Every integer $1 \leq d \leq g$ is realised as the trace field degree of a product of two affine multitwists on a surface in every connected component of every stratum of Abelian differentials on Riemann surfaces of genus g .*

We note that Theorem 1.7 is optimal. Indeed, the trace field degree of an Abelian differential is at most g , see [26, 20].

Square-tiled surfaces. The case where the degree of the field extension $\mathbb{Q}(\lambda + \lambda^{-1}) : \mathbb{Q}$ equals one plays a special role in Teichmüller theory, and our theorems are well-known in this context. The translation surfaces admitting such pseudo-Anosov maps are also called *arithmetic surfaces*, or *square-tiled surfaces*, since they are torus coverings [20]. In particular, this implies that the field extension $\mathbb{Q}(\lambda) : \mathbb{Q}(\lambda + \lambda^{-1})$ has degree two.

Idea of the proof of Theorems 1.6 and 1.7. Let $\mathcal{H}(k_1, k_2, \dots, k_m)$ be a given stratum of Abelian differentials in genus g . Fix some $2 \leq d \leq g$. This is the degree of a trace field we want to construct. In the Thurston–Veech construction, the stretch factor λ of $T_\alpha \circ T_\beta$ is related to the geometric intersection matrix of α and β as follows: the number $\lambda + \lambda^{-1} + 2$ equals the Perron–Frobenius eigenvalue of XX^\top . In order to control the degree of $\lambda + \lambda^{-1}$, we therefore need to control the degree of the Perron–Frobenius eigenvalue of XX^\top . Roughly, our strategy consists of the following four steps.

Step 1) construct examples. For positive integers $y, y_i, i = 1, \dots, g-1$, we start by constructing a square-tiled surface $(X, \omega) \in \mathcal{H}(k_1, k_2, \dots, k_m)$ depending on y, y_i . We think of the numbers y, y_i as variables that we specify later on. Applying the Thurston–Veech construction using the core curves of the horizontal and vertical annuli of (X, ω) gives us a matrix XX^\top of size $g \times g$.

Step 2) specify the y_i . The characteristic polynomial $p_g(t, y) \in \mathbb{Z}[t, y]$ of the matrix XX^\top satisfies $p_g(t, y) = (t-2)^{g-d} p_d(t, y)$ if we set $g-d+1$ of the $g-1$ parameters y_i equal to 2. Furthermore, if all the other y_i are pairwise different, then $p_d(t, y)$ is shown to be irreducible in $\mathbb{Z}[t, y]$ in Section 4.2 of Chapter 2.

Step 3) specify y . Hilbert’s irreducibility theorem [28] furnishes infinitely many integer specifications of y such that $p_d(t, y) \in \mathbb{Z}[t]$ is irreducible. By our construction, all these choices of parameters correspond to surfaces in $\mathcal{H}(k_1, k_2, \dots, k_m)$. Furthermore, the trace field is generated by the Perron–Frobenius eigenvalue of XX^\top , which has degree d as desired.

Step 4) Apply the nonsplitting criterion. Finally, we apply Theorem 2.8 to deduce that the stretch factor λ of $T_\alpha \circ T_\beta$ is of degree $2d$ for all the specifications of y_i and y as above.

This description of the strategy does not yet take into account the connected components we want to reach, but basically the same idea can be applied in order to deal with all connected components. However, we need to take variations of the families of examples we consider in order to find surfaces belonging to all of them. This is dealt with in Section 5 of Chapter 2.

2. Algebraic characterisation of stretch factors

An algebraic characterisation of the numbers that occur as stretch factors is still unknown. Recall that Fried showed in 1985 that every stretch factor is a bi-Perron number, and asked whether every bi-Perron number has a power that is a pseudo-Anosov stretch factor [19]. Fried’s question remains completely open².

In 2020, Pankau positively answered Fried’s question for the class of *Salem numbers*, which are the bi-Perron numbers with all other Galois conjugates of modulus at most one, and with at least one Galois conjugate of modulus exactly one [48].

In joint work with Pankau, we generalised Pankau’s positive answer to Fried’s question to all bi-Perron numbers with real or unimodular Galois conjugates [38]. To date, this remains the strongest known result concerning Fried’s question. In fact, our theorem is more precise in that it offers a precise geometric characterisation of the bi-Perron numbers with real or unimodular Galois conjugates. Indeed, this condition on the Galois conjugates turns out to precisely characterise the pseudo-Anosov stretch factors that arise from the Thurston–Veech construction, and the spectral radii of bipartite Coxeter transformations.

²It is not even known whether the formulation with “has a power” is actually needed, and Fried’s problem is often cited without allowing for powers. We note that already the version allowing for powers suffices to ensure that every bi-Perron number arises as the growth rate of an isomorphism of $\pi_1(S)$ induced by a surface homeomorphism

THEOREM 1.8 ([38], Theorem 1). *For a bi-Perron number λ , the following are equivalent:*

- (a) *all Galois conjugates of λ are contained in $\mathbb{S}^1 \cup \mathbb{R}$;*
- (b) *for some positive integer k , λ^k is the stretch factor of a pseudo-Anosov homeomorphism arising from the Thurston–Veech construction;*
- (c) *for some positive integer k , λ^k is the spectral radius of a bipartite Coxeter transformation of a bipartite Coxeter diagram with simple edges.*

Theorem 1.8 is of optimal quality. Indeed, the smallest stretch factor of a pseudo-Anosov homeomorphism arising from the Thurston–Veech construction as well as the smallest spectral radius > 1 of a Coxeter transformation are both equal to Lehmer’s number $\lambda_L \approx 1.17628$ by work of Leininger [33] and McMullen [44], respectively. On the other hand, no such lower bound exists for bi-Perron numbers all of whose Galois conjugates are contained in $\mathbb{S}^1 \cup \mathbb{R}$. This is the content of the following proposition.

PROPOSITION 1.9 ([38], Proposition 2). *There exist bi-Perron numbers arbitrarily close to 1 all of whose Galois conjugates are contained in $\mathbb{S}^1 \cup \mathbb{R}_{>0}$.*

PROOF. Choose any $\varepsilon > 0$. By Robinson’s work on Chebyshev polynomials [53], there exist infinitely many algebraic integers that lie, together with all their Galois conjugates, in the interval $[-2 + \varepsilon, 2 + 2\varepsilon]$. On the other hand, by a result due to Pólya described in Schur [57], only finitely many algebraic integers lie, together with all their Galois conjugates, in the interval $[-2 + \varepsilon, 2]$. It follows that in the interval $(2, 2 + 2\varepsilon]$, there exist infinitely many Perron numbers all of whose Galois conjugates are contained in the interval $[-2 + \varepsilon, 2 + 2\varepsilon]$. Let $p(t)$ be the minimal polynomial of such a Perron number and define the polynomial $f(t) = t^{\deg(p)}p(t + t^{-1})$. Then every root x of $f(t)$ is related to some root y of $p(t)$ by $x + x^{-1} = y$ and vice versa. In particular, all the roots of $f(t)$ are contained in $\mathbb{S}^1 \cup \mathbb{R}_{>0}$. Furthermore, if $2 < y < 2 + 2\varepsilon$, then $1 < x < 1 + \varepsilon + \sqrt{2\varepsilon + \varepsilon^2}$, assuming without loss of generality that $x > x^{-1}$. Now, let x_0 be the maximal real root of $f(t)$. By construction, no other root of $f(t)$ is as small as x_0^{-1} in modulus, so x_0 is a bi-Perron number all of whose Galois conjugates are contained in $\mathbb{S}^1 \cup \mathbb{R}_{>0}$. Choosing ε arbitrarily small yields the desired result. \square

Proposition 1.9 shows that we have to allow the topology of the underlying surface to change in order to have the flexibility needed to prove

Theorem 1.8. In the general context of Fried's question, we cannot hope for a positive answer if we fix the topology of the underlying surface, even if we fix the degrees of stretch factors that may appear: for example, there are cubic bi-Perron numbers that are not realised as the stretch factor of any mapping class on a genus three surface, see the work of Yazdi [69].

3. Spectral radii of integer matrices

As we have seen in the teaser, orientation-reversing integer linear dynamical systems can be simpler than orientation-preserving ones in the following sense: among all matrices $A \in \mathrm{GL}_2(\mathbb{Z})$ with $\det(A) = -1$, the smallest spectral radius > 1 is the golden ratio φ , while among matrices A with $\det(A) = 1$, the smallest spectral radius > 1 is φ^2 . We generalise this comparison to reciprocal and skew-reciprocal matrices of any even dimension, under the assumption of either primitivity or nonnegativity and irreducibility.

A matrix is *nonnegative* if all its coefficients are nonnegative. Such a matrix is *primitive* if some power has strictly positive coefficients. A matrix is *irreducible* if it is not conjugate via a permutation matrix to an upper triangular block matrix. We call a matrix *reciprocal* if the set of its eigenvalues (counted with multiplicity) is invariant under the transformation $t \mapsto t^{-1}$. Finally, we call a matrix *skew-reciprocal* if the set of its eigenvalues (counted with multiplicity) is invariant under the transformation $t \mapsto -t^{-1}$. Important examples of reciprocal or skew-reciprocal matrices are symplectic or antisymplectic matrices, respectively.

We find out that with the exception of dimension six in the primitive case, the skew-reciprocal matrices always realise smaller spectral radii > 1 than the reciprocal ones.

THEOREM 1.10 ([37], Theorem 1.1). *Let $g \geq 1$ and $g \neq 3$. Among primitive matrices $A \in \mathrm{GL}_{2g}(\mathbb{Z})$, the skew-reciprocal ones realise a smaller spectral radius > 1 than the reciprocal ones. For $g = 3$, the reciprocal matrices realise a smaller spectral radius than the skew-reciprocal ones.*

THEOREM 1.11 ([37], Theorem 1.2). *Let $g \geq 1$. Among nonnegative irreducible $A \in \mathrm{GL}_{2g}(\mathbb{Z})$, the skew-reciprocal ones realise a smaller spectral radius > 1 than the reciprocal ones.*

Naturally, the following question arises by dropping the hypotheses of primitivity or irreducibility.

QUESTION 1.12 ([37], Question 1.3). *Let $g \geq 1$. Do the skew-reciprocal matrices $A \in \text{GL}_{2g}(\mathbb{Z})$ realise a smaller spectral radius > 1 than the reciprocal ones?*

The proofs of Theorems 1.10 and 1.11 are based on McMullen's calculation of the minimal possible spectral radii > 1 for primitive and nonnegative irreducible reciprocal matrices [45]. In fact, we carry out the same calculation for skew-reciprocal matrices in order to determine the minimal spectral radii > 1 that arise, and compare the values with McMullen's result. The following two theorems summarise our results on these minimal spectral radii.

THEOREM 1.13 ([37], Theorem 1.4). *Let $g \geq 1$. The minimal spectral radius > 1 among all the skew-reciprocal nonnegative irreducible matrices $A \in \text{GL}_{2g}(\mathbb{Z})$ is realised by the by the largest root λ_{2g} of the polynomial*

$$t^{2g} - t^g - 1$$

in case g is odd, and of the polynomial

$$t^{2g} - t^{g+1} - t^{g-1} - 1$$

in case g is even.

THEOREM 1.14 ([37], Theorem 1.5). *Let $g \geq 2$. The minimal spectral radius > 1 among skew-reciprocal primitive matrices $A \in \text{GL}_{2g}(\mathbb{Z})$ is realised by the by the largest root μ_{2g} of the polynomial*

$$t^{2g} - t^{g+2} - t^{g-2} - 1$$

in case g is odd, and of the polynomial

$$t^{2g} - t^{g+1} - t^{g-1} - 1$$

in case g is even.

Given Theorems 1.14 and 1.13, Theorems 1.10 and 1.11 follow readily.

PROOF OF THEOREMS 1.10 AND 1.11. The normalised sequence $(\mu_{2g})^{2g}$ converges to the square of the silver ratio, $\sigma^2 = 3+2\sqrt{2}$, while the normalised sequence $(\lambda_{2g})^{2g}$ has two accumulation points: again the square of the silver ratio, for even g , but also the square of the golden ratio, $\varphi^2 = \frac{3+\sqrt{5}}{2}$, for odd g . To see this, note that $(\mu_{2g})^g$ is the largest real zero of the function

$$f(t) = t^2 - t^{1+\frac{2}{g}} - t^{1-\frac{2}{g}} - 1$$

in case g is odd, and of the function

$$f(t) = t^2 - t^{1+\frac{1}{g}} - t^{1-\frac{1}{g}} - 1$$

in case g is even. Clearly for $g \rightarrow \infty$ any real zero > 1 converges to the larger root of the polynomial $t^2 - 2t - 1$, which is σ . Hence $(\mu_{2g})^{2g}$ converges to $\sigma^2 = 3 + 2\sqrt{2}$. The argument for the sequence λ_g is analogous.

In either case, these accumulation points are all smaller than the analogous smallest possible accumulation point in the case of reciprocal matrices. Indeed, McMullen proves that the minimal accumulation point for the normalised sequence of spectral radii for reciprocal matrices is φ^4 , where φ is the golden ratio [45]. This number is strictly larger than the square of the silver ratio, so asymptotically the result is given. We finish the proof Theorems 1.10 and 1.11 by using the monotonicity of the sequences of normalised spectral radii, and compare these sequences for small g . It turns out that the only case where a normalised sequence of the skew-reciprocal matrices is larger than the accumulation point φ^4 of the normalised sequence of the reciprocal matrices is in the case $g = 3$ of primitive matrices. This finishes the proof for $g \neq 3$. For $g = 3$ in the primitive case, we simply check that $(\mu_6)^6 > 8.18$, whereas the smallest normalised spectral radius in the reciprocal case is ≈ 7.57 by McMullen's result [45]. This finishes the proof also in the case $g = 3$. \square

Application to Dimitrov's theorem. Let $f \in \mathbb{Z}[t]$ be a monic polynomial. The *house* $\overline{|f|}$ of f is the largest modulus among the zeroes of f ,

$$\overline{|f|} = \max_{f(\alpha)=0} |\alpha|.$$

QUESTION 1.15 (Schinzel–Zassenhaus [56]). *Does there exist a universal constant $c > 0$ such that any house > 1 of an irreducible monic integer polynomial is at least $1 + \frac{c}{d}$, where d is the degree of the polynomial?*

The question of Schinzel and Zassenhaus has been answered positively by Dimitrov [12] in 2019.

Our aim is to present a reformulation of the question of Schinzel and Zassenhaus in terms of a comparison of reciprocal and skew-reciprocal polynomials that is similar to the geometric idea of comparing stretch factors of orientation-preserving and orientation-reversing pseudo-Anosov maps.

A polynomial $f \in \mathbb{Z}[t]$ of degree $2d$ is *reciprocal*³ if $f(t) = \pm t^{2d} f(t^{-1})$. This condition is equivalent to the roots of f being invariant under the transformation $t \mapsto -t$. Similarly, a polynomial $f \in \mathbb{Z}[t]$ of even degree $2d$ is *skew-reciprocal* if $f(t) = \pm t^{2d} f(-t^{-1})$. Again, this condition is equivalent to the roots of f being invariant under the transformation $t \mapsto -t^{-1}$.

Let λ_i and $\tilde{\lambda}_i$ be the smallest houses larger than 1 among all monic integer reciprocal and skew-reciprocal polynomials of degree 2^i , respectively. Let $r_i = 2^i \log(\lambda_i)$ and $s_i = 2^i \log(\tilde{\lambda}_i)$.

THEOREM 1.16. *There exists a universal constant $c > 0$ such that any house larger than 1 of an irreducible monic integer polynomial is bounded from below by $1 + \frac{c}{d}$, where d is the degree of the polynomial, exactly if the subset $\left\{ \frac{q_N}{q_n} \in \mathbb{R} : n, N \in \mathbb{N}, n < N \right\} \subset \mathbb{R}$ is bounded away from zero by a constant, where $q_m = \prod_{i=1}^m \frac{r_i}{s_i}$.*

The easiest way to fulfil the second statement in Theorem 1.16 is if for all but finitely many i , we have $r_i \geq s_i$, compare with Question 1.12. Our Theorems 1.10 and 1.11 prove this stronger condition in the case of primitive and nonnegative irreducible matrices, respectively, without using Dimitrov's theorem. Expanding this result to the class of all reciprocal and skew-reciprocal matrices in $\mathrm{GL}_{2^i}(\mathbb{Z})$ would provide a new proof Dimitrov's theorem. In a more geometric direction, we prove a similar comparison theorem for normalised pseudo-Anosov stretch factors in the fully-punctured case, comparing orientation-preserving and orientation-reversing maps, see Section 4.3 of this introduction.

4. Minimal stretch factors

One of the most basic problems on stretch factors is to find the minimal ones on every surface. This problem is notoriously difficult: among closed orientable surfaces, it is only solved for genus one and two in the case of orientation-preserving maps [11], and only for genus one in the case of orientation-reversing maps.

QUESTION 1.17 (Minimal stretch factor question). *What is the minimal stretch factor among all pseudo-Anosov maps on a fixed surface?*

³We note that in the literature, the definition of reciprocity of a polynomial often supposes a positive sign in the defining equation for the coefficients. We allow for the possibility of a negative sign as well since it seems more natural to us in the context of spectral radii of integer matrices.

As we have argued at the beginning of this introduction, answering Question 1.17 in the orientation-preserving context for the closed orientable surface of genus g is equivalent to finding the shortest closed geodesic in the moduli space \mathcal{M}_g equipped with the Teichmüller metric. Furthermore, this task bears resemblance with finding 3-manifolds of small hyperbolic volume or finding integer polynomials with small Mahler measure. It is even more closely related to the problem of finding the minimal spectral radius among symplectic matrices of a given size, compare with Section 3 of this introduction.

Lanneau and Thiffeault found the minimal stretch factors on the orientable closed surfaces of genus 1, 2, 3, 4, 5, 7 and 8, under the additional hypothesis that the associated unstable foliations are orientable [29].

There are also certain constructions of pseudo-Anosov maps restricted to which Question 1.17 has been answered: as we have already mentioned, Leininger found the minimal stretch factor among all stretch factors arising from the Thurston–Veech construction [33]. For Penner’s construction, Question 1.17 is answered by work of the author [34] which was part of the PhD thesis [35].

Recently, Hironaka and Tsang significantly advanced our understanding of Question 1.17 in the fully-punctured case, that is, when the singularities of the invariant foliations of the pseudo-Anosov map lie on the punctures of the surface, under the additional hypothesis that there are at least two puncture orbits. More precisely, for all fixed even Euler characteristics, they answered the analogue Question 1.17, and showed that for all Euler characteristics, the normalised stretch factor is always bounded from below by $\varphi^4 \approx 6.854$, where φ is the golden ratio [22].

The focus in this habilitation thesis is to extend the three types of results, namely orientable foliations, Penner’s construction and the fully-punctured case, to the setting of orientation-reversing pseudo-Anosov maps or pseudo-Anosov maps on nonorientable surfaces. This is based on joint work with Strenner [39, 41] as well as joint work with Lanneau and Tsang [32].

4.1. Orientable foliations. Let N_g be the closed nonorientable surface of genus g , that is, the connected sum of g real projective planes $P^2(\mathbb{R})$, and let $\delta^+(N_g)$ be the minimal stretch factor among pseudo-Anosov homeomorphisms of N_g with an orientable invariant foliation. Note that at most

one of the invariant foliations can be orientable, otherwise the surface itself would have to be orientable as well.

THEOREM 1.18 ([39], Theorem 1.1). *The values and minimal polynomials of $\delta^+(N_g)$ for $g = 4, 5, 6, 7, 8, 10, 12, 14, 16, 18$ and 20 are as follows:*

| g | $\delta^+(N_g) \approx$ | Minimal polynomial of $\delta^+(N_g)$ | singularity type |
|-----|-------------------------|---------------------------------------|------------------|
| 4 | 1.83929 | $t^3 - t^2 - t - 1$ | (6) |
| 5 | 1.51288 | $t^4 - t^3 - t^2 + t - 1$ | (4,4,4) |
| 6 | 1.42911 | $t^5 - t^3 - t^2 - 1$ | (10) |
| 7 | 1.42198 | $t^6 - t^5 - t^3 + t - 1$ | (4,4,4,4,4) |
| 8 | 1.28845 | $t^7 - t^4 - t^3 - 1$ | (14) |
| 10 | 1.21728 | $t^9 - t^5 - t^4 - 1$ | (18) |
| 12 | 1.17429 | $t^{11} - t^6 - t^5 - 1$ | (22) |
| 14 | 1.14551 | $t^{13} - t^7 - t^6 - 1$ | (26) |
| 16 | 1.12488 | $t^{15} - t^8 - t^7 - 1$ | (30) |
| 18 | 1.10938 | $t^{17} - t^9 - t^8 - 1$ | (34) |
| 20 | 1.09730 | $t^{19} - t^{10} - t^9 - 1$ | (38) |

The table also contains the singularity type of the minimising pseudo-Anosov map. For example, (4,4,4) means that the pseudo-Anosov map has three 4-pronged singularities.

Based on this result, we conjecture the following.

CONJECTURE 1.19 ([39], Conjecture 1.2). *For all $k \geq 2$, $\delta^+(N_{2k})$ is the largest root of*

$$t^{2k-1} - t^k - t^{k-1} - 1.$$

We believe that the minimal stretch factors in the case of genus 9 and 11 are the ones shown in the following table. We will discuss supporting evidence in Section 4.4 of Chapter 5.

CONJECTURE 1.20 ([39], Conjecture 1.3). *The approximate values and minimal polynomials of $\delta^+(N_g)$ for $g = 9, 11$ are as follows:*

| g | $\delta^+(N_g) \approx$ | Minimal polynomial of $\delta^+(N_g)$ | singularity type |
|-----|-------------------------|--|------------------|
| 9 | 1.35680 | $t^8 - t^5 - t^4 - t^3 - 1$ | (16) |
| 11 | 1.22262 | $\frac{t^{12} - t^7 - t^6 - t^5 - 1}{t^2 + t + 1}$ | (8,8,8) |

We now turn to the case of orientation-reversing pseudo-Anosov maps on closed orientable surfaces. For our main result in this context, denote

by $\delta_{rev}^+(S_g)$ the minimal stretch factor among orientation-reversing pseudo-Anosov homeomorphisms of the closed orientable surface S_g of genus g that have orientable invariant foliations.

THEOREM 1.21 ([39], Theorem 1.4). *The values and minimal polynomials of $\delta_{rev}^+(S_g)$ for $g = 1, 3, 5, 7, 9$ and 11 are as follows:*

| g | $\delta_{rev}^+(S_g) \approx$ | Minimal polynomial of $\delta_{rev}^+(S_g)$ | singularity type |
|-----|-------------------------------|--|------------------|
| 1 | 1.61803 | $t^2 - t - 1$ | no singularities |
| 3 | 1.25207 | $\frac{t^8 - t^5 - t^3 - 1}{t^2 + 1}$ | (4,4,4,4) |
| 5 | 1.15973 | $\frac{t^{12} - t^7 - t^5 - 1}{t^2 + 1}$ | (6,6,6,6) |
| 7 | 1.11707 | $\frac{t^{16} - t^9 - t^7 - 1}{t^2 + 1}$ | (8,8,8,8) |
| 9 | 1.09244 | $\frac{t^{20} - t^{11} - t^9 - 1}{t^2 + 1}$ | (10,10,10,10) |
| 11 | 1.07638 | $\frac{t^{24} - t^{13} - t^{11} - 1}{t^2 + 1}$ | (12,12,12,12) |

Moreover, we have

$$\delta_{rev}^+(S_g) \geq \delta_{rev}^+(S_{g-1})$$

for $g = 2, 4, 6, 8$ and 10.

Based on these results, it is natural to conjecture the following.

CONJECTURE 1.22 ([39], Conjecture 1.5). *For all $k \geq 2$, $\delta_{rev}^+(S_{2k-1})$ is the largest root of*

$$t^{4k} - t^{2k+1} - t^{2k-1} - 1.$$

Theorem 1.21 shows that $\delta_{rev}^+(S_g)$ fails to be strictly decreasing at every other step for small values of g . We conjecture that in fact the value of $\delta_{rev}^+(S_g)$ strictly increases at every other step. We will discuss evidence for this after Proposition 5.23.

CONJECTURE 1.23 (Conjecture 1.6 from [39]). *For all $k \geq 1$, we have*

$$\delta_{rev}^+(S_{2k}) > \delta_{rev}^+(S_{2k-1}).$$

All of our minimising examples have a power that arises from Penner's construction of pseudo-Anosov mapping classes, see Section 4.2 below for a description of this construction. As an illustrative example, for $g = 3$ our minimising examples are obtained by taking a positive Dehn twist along the core curve of one of the glued annuli showed in Figure 1.1, composed with a rotation of the surface by three 'clicks' in the clockwise direction.

In sharp contrast, none of the classical minimal stretch factor examples in the orientation-preserving context have a power that arises from Penner's construction. This is because these stretch factors have Galois conjugates




FIGURE 1.1. A particular view of the surface $S_{k,4}$.

on the unit circle. However, Shin and Strenner showed in [59] that examples with this property do not have a power arising from Penner's construction.

One may wonder what the reason for this discrepancy is. A heuristic reason for why Galois conjugates of small stretch factors *should* lie on the unit circle is that every pseudo-Anosov stretch factor λ is a bi-Perron algebraic unit. Recall that these are real numbers larger than 1 whose Galois conjugates lie in the annulus $\lambda^{-1} \leq |z| \leq \lambda$. If λ is close to 1, this annulus is a thin neighborhood of the unit circle, so it seems natural for some Galois conjugates to lie on the unit circle.

However, in Section 3 of Chapter 5 we will prove the following theorem that explains why the nonorientable cases are different.

THEOREM 1.24 ([39], Theorem 1.10). *If f is a pseudo-Anosov map on a nonorientable surface or an orientation-reversing pseudo-Anosov map on an orientable surface, then the stretch factor of f does not have Galois conjugates on the unit circle.*

4.2. Penner's construction. We study the minimal stretch factors among pseudo-Anosov mapping classes arising from a construction by products of Dehn twists along suitable simple closed curves, due to Penner [50], which we recall now for the case of closed surfaces which need not be orientable.

We assume all mentioned two-sided simple closed curves c to be equipped with a homeomorphism φ_c from a regular neighbourhood U_c of c to the standard annulus $A = \mathbf{S}^1 \times [0, 1]$. The *Dehn twist* T_c along c is then defined to be the identity outside U_c and $\varphi_c^{-1} \circ T_A \circ \varphi_c$ inside U_c , where T_A is the standard right Dehn twist of the annulus A , sending an arc that crosses the core curve of A to an arc that crosses the core curve but also winds around it once in the positive direction. In the case of an oriented surface, the Dehn twist T_c along a curve c is positive or negative if the homeomorphism φ_c is orientation-preserving or orientation-reversing, respectively.

In Penner's construction for orientable surfaces, we ask that if two curves intersect, one should be twisted along positively and the other should be twisted along negatively. The notion of a positive or a negative Dehn twist does not make sense on a nonorientable surface, but one can still ask that locally at any intersection point, the twisting should go in different directions: we say two curves c_1 and c_2 intersect *inconsistently* if for every point $p \in c_1 \cap c_2$ the pullbacks of the orientation of A by φ_{c_1} and φ_{c_2} disagree.

THEOREM 1.25 (Penner's construction). *Let $\{c_i\}$ be a collection of at least two two-sided curves which intersect inconsistently and without bigons, and whose union fills a closed surface S . Let \mathcal{P} be the monoid generated by the Dehn twists T_{c_i} . Define $\rho : \mathcal{P} \rightarrow \mathrm{SL}(n, \mathbb{Z})$ by*

$$\rho(T_{c_i}) = I + R_{c_i},$$

and extend linearly, where the matrices R_{c_i} are obtained from the geometric intersection matrix Ω of the curves $\{c_1, \dots, c_n\}$ by setting all entries to zero which are not in the row corresponding to c_i . Then each $\phi \in \mathcal{P}$ such that every c_i gets twisted along at least once is pseudo-Anosov and its stretch factor equals the Perron–Frobenius eigenvalue of $\rho(\phi)$.

For more details and proofs, see Penner's original article [50] or Fathi's alternative approach [15]. Penner's construction for nonorientable surface is also explained in more detail by Strenner [62].

Let N_g be the nonorientable closed surface of genus g , and let $\sigma = 1 + \sqrt{2}$ be the silver ratio.

THEOREM 1.26 ([41], Theorem 1.1). *For the sequence of minimal stretch factors $\delta_P(N_g)$ among pseudo-Anosov homeomorphisms arising from Penner's construction on the nonorientable closed surface of genus g , the two limits $\lim_{k \rightarrow \infty} \delta_P(N_{2k})$ and $\lim_{k \rightarrow \infty} \delta_P(N_{2k+1})$ exist, and*

- (a) $\lim_{k \rightarrow \infty} \delta_P(N_{2k}) = 3 + 2\sqrt{2} = \sigma^2$,
- (b) $\lim_{k \rightarrow \infty} \delta_P(N_{2k+1}) > 3 + 2\sqrt{2}$.

We prove Theorem 1.26 by finding, for each nonorientable closed surface, a pseudo-Anosov mapping class which has minimal stretch factor among all pseudo-Anosov mapping classes arising from Penner's construction on this surface. When the genus is even, we give a concrete description of the minimal stretch factors $\delta_P(N_{2k})$.

THEOREM 1.27 ([41], Theorem 1.2). *For $k \geq 2$, the minimal stretch factor $\delta_P(N_{2k})$ equals the largest real solution of the equation*

$$t - t^{\frac{k}{2k-1}} - t^{\frac{k-1}{2k-1}} - 1 = 0.$$

Alternatively, $\delta_P(N_{2k})$ equals the $2k - 1$ st power of the largest real root of the integral polynomial $t^{2k-1} - t^k - t^{k-1} - 1$.

Such a description is possible due to a rotational symmetry of the mapping classes realising $\delta_P(N_{2k})$. Unfortunately, the mapping classes realising $\delta_P(N_{2k+1})$ do not have such a rotational symmetry, so we do not obtain such a concrete description of the minimal stretch factors $\delta_P(N_{2k+1})$. However, we do have a description of $\delta_P(N_{2k+1})$ as the largest eigenvalue of a certain product of matrices, so using a computer we can compute $\delta_P(N_{2k+1})$ for at least up to $k = 100$. These computations strongly suggest the following conjecture.

CONJECTURE 1.28 ([41], Conjecture 1.3). *The limit $\lim_{k \rightarrow \infty} \delta_P(N_{2k+1})$ is the largest real root of the polynomial $t^4 - 8t^3 + 13t^2 - 8t + 1$, approximately 6.071360241468951.*

For example, $\delta_P(N_{101})$ approximates the conjectured limit until 9 decimal places and $\delta_P(N_{151})$ approximates it until at least 15 decimal places. It is striking that the coefficients of this polynomial are Fibonacci numbers, considering that the golden ratio makes a frequent appearance in the literature on minimal stretch factors.

The minimising examples. Figure 1.2 depicts the genus six and seven nonorientable closed surfaces (with one open disc removed) as a surface

obtained by glueing together five or six twisted annuli, respectively. The minimising examples are obtained by applying a Dehn twist along the core curve of each of those annuli. The order of the twisting should be, in a sense we make precise later, as bipartite as possible. We find that for each nonorientable closed surface, the mapping classes of this kind minimise the stretch factor among pseudo-Anosov mapping classes arising from Penner's construction, where the size of the cycle of annuli glued together is determined by the genus of the surface (Theorems 6.16 and 6.28).




FIGURE 1.2. A connected sum of six (on the left) and seven (on the right) copies of $P^2(\mathbb{R})$ minus a disc.

Orientable versus nonorientable surfaces. A remarkable difference between the case of orientable and nonorientable surfaces is that in the case of orientable closed surfaces S_g , the limit of the sequence of minimal stretch factors $\delta_P(S_g)$ arising from Penner's construction exists, whereas in the nonorientable case, the limit of the sequence $\delta_P(N_g)$ does not. An orientable double cover argument implies that any accumulation point of $\delta_P(N_g)$ must be at least $3 + 2\sqrt{2} = \sigma^2$, which is the limit in the orientable case by work of the author [34] which is part of the PhD thesis [35]. In fact, the limit in the orientable case is the same as for the even genus subsequence in the nonorientable case.

Fibred link techniques for odd genus. A special rotational symmetry of the even genus examples helps us to determine their stretch factors. In order to deal with nonorientable surfaces of odd genus, where such a strong symmetry is lacking, we use the theory of fibred links in \mathbb{S}^3 . More precisely, we describe the stretch factor of these examples as the largest real root of the Alexander polynomial of some fibred link in \mathbb{S}^3 , obtained by plumbing

positive and negative Hopf bands to a disc. In order to distinguish between different product orders of Dehn twists, we calculate the difference of the associated Alexander polynomials, by using the skein relation for the Alexander polynomial. We find that this difference is of a very particular form, which allows us to deduce the monotonicity properties needed to single out the examples with minimal stretch factors.

4.3. The fully-punctured case. For a pseudo-Anosov map $f : S \rightarrow S$ of a surface S of finite type, we call $\lambda(f)^{|\chi(S)|}$ the *normalised stretch factor* of f . The pseudo-Anosov map f is said to be *fully-punctured* if the set of singular points equals the set of punctures. In this case, we will denote the set of punctures by \mathcal{X} . Note that f acts on \mathcal{X} by some permutation, hence it makes sense to talk about the orbits of punctures under the action of f , or *puncture orbits* for short.

Recently, Hironaka and Tsang significantly advanced our understanding of the analogous problem in the fully-punctured case, under the additional hypothesis that there are at least two puncture orbits. More precisely, for all even Euler characteristics, they determined the precise minimal stretch factor, and showed that for all Euler characteristics, the normalised stretch factor is always bounded from below by $\varphi^4 \approx 6.854$, where φ is the golden ratio [22].

Furthermore, Tsang determined the shape of the set of normalised stretch factors of fully-punctured orientation-preserving pseudo-Anosov maps [67]: it is the union of six isolated points and a dense subset of $[\varphi^4, \infty)$.

We explore analogous questions for orientation-reversing pseudo-Anosov maps. Recall our notation $\sigma = 1 + \sqrt{2}$ for the silver ratio. Our first result in the fully-punctured case is the following.

THEOREM 1.29 ([32], Theorem 1.1). *Let $f : S \rightarrow S$ be an orientation-reversing fully-punctured pseudo-Anosov map on a finite-type orientable surface. Suppose $-\chi(S) \geq 4$ and f has at least two puncture orbits, then the normalised stretch factor of f satisfies*

$$\lambda(f)^{-\chi(S)} \geq \sigma^2 \approx 5.828.$$

Moreover, this inequality is asymptotically sharp, in the sense that for every integer $k \geq 2$, there exists an orientation-reversing fully-punctured pseudo-Anosov map $f_k : \Sigma_k \rightarrow \Sigma_k$, where $\chi(\Sigma_k) = -2k$, such that we have the limit $\lim_{k \rightarrow \infty} \lambda(f_k)^{-\chi(\Sigma_k)} = \sigma^2$.

Sharpness in Theorem 1.29 follows from explicit examples, see Proposition 7.12. For all even Euler characteristics $-2k$, $k \geq 2$, we construct pseudo-Anosov maps that have as their stretch factor the minimal spectral radius among all skew-reciprocal primitive matrices $A \in \mathrm{GL}_{2k}(\mathbb{Z})$. For $k \neq 3$, this stretch factor is smaller than the lower bound for the orientation-preserving case determined by Hironaka and Tsang [22, Theorem 1.9]. We obtain the following corollary, comparing the stretch factors of orientation-preserving and orientation-reversing fully-punctured pseudo-Anosov maps on surfaces with even Euler characteristic.

COROLLARY 1.30 ([32], Corollary 1.2). *Fix a positive integer $k \neq 3$. Then, among all fully-punctured pseudo-Anosov maps with at least two puncture orbits on surfaces with even Euler characteristic $-2k$, the orientation-reversing ones realise a smaller minimal stretch factor than the orientation-preserving ones.*

The only case of Corollary 1.30 not covered by Proposition 7.12 is the case $k = 1$. However, in this case the result follows from a direct comparison: as per [22], the smallest normalised stretch factor in the orientation-preserving case is φ^4 , while there exists an orientation-reversing map on the four-punctured sphere with normalised stretch factor φ^2 , see Remark 7.10.

In the case $k = 3$, the comparison might very well go the other way around. Indeed, this would be the case if our Conjecture 7.13 about the minimal normalised stretch factors of orientation-reversing fully-punctured pseudo-Anosov maps holds.

The comparison in Corollary 1.30 agrees with the comparison of the minimal spectral radii of reciprocal and skew-reciprocal primitive matrices $A \in \mathrm{GL}_{2k}(\mathbb{Z})$ presented in Theorems 1.10 and 1.11. By Theorem 1.16, extending this comparison to the class of all reciprocal and skew-reciprocal matrices $A \in \mathrm{GL}_{2i}(\mathbb{Z})$ would yield a new proof of a theorem of Dimitrov [12], so Corollary 1.30 can be seen as a geometric version of this proof idea.

Concerning the set of normalised stretch factors of orientation-reversing fully-punctured pseudo-Anosov maps, we have the following result.

THEOREM 1.31 ([32], Theorem 1.3). *The smallest three elements of the set of normalised stretch factors of orientation-reversing fully-punctured pseudo-Anosov maps are φ, σ and φ^2 . Furthermore, the set of normalised stretch factors of orientation-reversing fully-punctured pseudo-Anosov maps contains a dense subset of $[\sigma^2, \infty)$.*

Outline of the proof of Theorem 1.29 and Theorem 1.31

For Theorem 1.29, we follow the approach by Hironaka and Tsang [22]. They develop the theory of standardly embedded train tracks, and study the action on the real edges of the train track induced by the pseudo-Anosov map. In the orientation-preserving case, this allows them to reduce the problem to the study of minimal spectral radii of reciprocal primitive matrices of given sizes, which are well-known by work of McMullen [45], at least in even dimensions.

Initially, our hope was that the adaptation of McMullen's result to skew-reciprocal matrices in Chapter 4 directly allows us to conclude an analogous result for orientation-reversing pseudo-Anosov maps. However, it turns out that in this case the action induced on the real edges need not be skew-reciprocal, but is instead *skew-reciprocal up to cyclotomic factors*: the eigenvalues are invariant under the transformation $t \mapsto -t^{-1}$, except for possibly some roots of unity. See Remark 7.7 for an explicit example.

The main technical contribution we need is to extend certain cases of the analysis of minimal spectral radii of skew-reciprocal primitive matrices of Chapter 4 to the more general class of matrices that are skew-reciprocal up to cyclotomic factors.

For the proof of Theorem 1.31, we study which elements in the set of normalised stretch factors of orientation-preserving fully-punctured pseudo-Anosov maps, determined by Tsang [67], admit an orientation-reversing square root. To this end, we use Theorem 1.24, stating that the Galois conjugates of the stretch factor of an orientation-reversing pseudo-Anosov map cannot lie on the unit circle.

However, our approach is not practicable for analysing the dense subset of normalised stretch factors of orientation-preserving maps in the interval $[\varphi^4, \sigma^4)$. This is the reason why we are currently unable to describe the normalised stretch factors of orientation-reversing maps in (φ^2, σ^2) . In particular, it remains an open problem to determine whether σ^2 is the minimal accumulation point of the set of normalised stretch factors of fully-punctured orientation-reversing pseudo-Anosov maps.

Organisation

In Chapter 2, we give the proofs of Theorems 1.1, 1.2, 1.3, 1.4, 1.5, 1.6, and 1.7 concerning trace field degrees and stretch factor degrees.

In Chapter 3, we give the proof of Theorem 1.8, characterising up to powers the algebraic integers that arise as stretch factors in the Thurston–Veech construction.

In Chapter 4, we prove Theorems 1.13 and 1.14 on spectral radii of integer matrices, as well as Theorem 1.16 on the comparison of the minimal spectral radii of reciprocal and skew-reciprocal matrices.

In Chapter 5, we discuss our results on minimal pseudo-Anosov stretch factors under the additional hypothesis of an orientable invariant foliation for nonorientable surfaces (Theorem 1.18) and for orientation-reversing maps of orientable surfaces (Theorem 1.21). We also prove Theorem 1.24 on the Galois conjugates of stretch factors in these two settings.

In Chapter 6, we give proofs of Theorems 1.26 and 1.27 on minimal stretch factors in Penner’s construction on nonorientable closed surfaces.

In Chapter 7, we prove Theorems 1.29 and 1.31 on the minimal normalised stretch factors of orientation-reversing pseudo-Anosov maps in the fully-punctured setting.

CHAPTER 2

Stretch factor degrees

In this chapter, we give the proofs of Theorems 1.1, 1.2, 1.3, 1.4, 1.5, 1.6 as well as 1.7 concerning trace field degrees and stretch factor degrees. The material is copied, adapted and consolidated from our two joint articles with Lanneau [30, 31].

We recall that our proof strategy necessitates us to control the degrees of two field extensions: first, we need to control the degree of the field extension

$$\mathbb{Q}(\lambda + \lambda^{-1}) : \mathbb{Q},$$

which we call the trace field degree. In a second step, we need to control whether the degree of the field extension

$$\mathbb{Q}(\lambda) : \mathbb{Q}(\lambda + \lambda^{-1})$$

equals one or two. We call the latter case *nonsplitting* because the stretch factor and its inverse do not split into two distinct minimal polynomials.

1. Irreducibility criteria

The goal of this section is to present an algebraic criterion that allows us to deduce that certain characteristic polynomials of matrices of the form XX^\top are irreducible.

PROPOSITION 2.1. *Let M be a square integer matrix, and let N be the principal submatrix of M obtained by deleting the first row and the first column. If M and N have no common eigenvalue, and if M has a simple eigenvalue ρ , then the characteristic polynomial of $\widetilde{M} = M + ay^p E_{11}$ is an irreducible element of $\mathbb{Z}[t, y]$, for all $p \geq 1$ and for all $0 \neq a \in \mathbb{Z}$.*

PROOF. Our goal is to use Eisenstein's criterion on the characteristic polynomial $\chi_{\widetilde{M}} \in \mathbb{Z}[t, y] \cong (\mathbb{Z}[t])[y]$, viewing it as a polynomial in the variable y and coefficients in $\mathbb{Z}[t]$. We calculate

$$\chi_{\widetilde{M}}(t, y) = \det(t \cdot \text{Id} - \widetilde{M}) = -y^p a \chi_N(t) + \chi_M(t)$$

and notice that $a\chi_N$ and χ_M are relatively prime in $\mathbb{Z}[t]$. Indeed, χ_M has leading coefficient $+1$ and no root in common with χ_N by our assumption that M and N have no eigenvalue in common. In particular, they have no common factor, which shows that $\chi_{\widetilde{M}} \in (\mathbb{Z}[t])[y]$ is primitive. In order to apply Eisenstein's criterion, let $\mu_\rho \in \mathbb{Z}[t]$ be the minimal polynomial of the simple eigenvalue ρ of M . By assumption, μ_ρ divides χ_M exactly once, but it does not divide χ_N since χ_M and χ_N have no common root. In particular, Eisenstein's criterion applies to show that the characteristic polynomial $\chi_{\widetilde{M}} \in (\mathbb{Z}[t])[y] \cong \mathbb{Z}[t, y]$ is irreducible. \square

REMARK 2.2. In the previous statement, one can easily replace $\chi_{\widetilde{M}}(t)$ by $\chi_{\widetilde{M}}(t^n)$ for any integer $n > 0$. Indeed $\chi_M(t^n)$ and $\chi_N(t^n)$ are still coprime and $\mu_\rho(t^n)$ divides $\chi_M(t^n)$ exactly once, so Eisenstein's criterion applies.

REMARK 2.3. Oscillatory matrices satisfy a stronger version of Perron–Frobenius theory, namely *all* the eigenvalues are positive real, simple, and they strictly interlace when taking a principal submatrix [2]. Hence, Proposition 2.1 applies very cleanly to this class of matrices.

We use Proposition 2.1 on the following two cases, described in Lemma 2.4 and Lemma 2.5.

LEMMA 2.4. For $n \geq 1$, let

$$N = \left(\begin{array}{c|ccc} a_1 & a_2 & \dots & a_n \\ \hline a_2 & & & \\ \vdots & & * & \\ a_n & & & \end{array} \right), \quad M = \left(\begin{array}{c|ccc} 0 & \alpha a_1 & \dots & \alpha a_n \\ \hline \alpha a_1 & & & \\ \vdots & & N & \\ \alpha a_n & & & \end{array} \right)$$

be square integer matrices with $a_1 \geq 1$ and $0 \neq \alpha \in \mathbb{Q}$. If M is nonnegative and irreducible, and if $\chi_N \in \mathbb{Z}[t]$ is irreducible, then the characteristic polynomial of the matrix $\widetilde{M} = M + ay^2E_{11}$ is irreducible in $\mathbb{Z}[t, y]$ for all $0 \neq a \in \mathbb{Z}$.

PROOF. In order to use Proposition 2.1, we need to show that M has a simple eigenvalue and that M and N share no eigenvalue. The former holds since M is nonnegative and irreducible, and in particular has a Perron–Frobenius eigenvalue which is simple. For the latter, we compute

$$\chi_M(t) = t\chi_N(t) + q(t),$$

where $q(t) \in \mathbb{Z}[t]$ is of degree at most $n - 1$. We claim that it is not the zero polynomial either. Indeed, we directly verify

$$\begin{aligned} q(0) &= \det \left(\begin{array}{c|ccc} 0 & -\alpha a_1 & \dots & -\alpha a_n \\ \hline -\alpha a_1 & & & \\ \vdots & & & \\ -\alpha a_n & & & -N \end{array} \right) \\ &= \det \left(\begin{array}{c|ccc} \alpha^2 a_1 & 0 & \dots & 0 \\ \hline 0 & & & \\ \vdots & & & \\ 0 & & & -N \end{array} \right) = \pm \alpha^2 a_1 \det N \neq 0. \end{aligned}$$

Now if there existed a common root $\lambda \in \mathbb{C}$ of χ_M and χ_N , then λ would also be a root of $q(t)$. But since χ_N is irreducible of degree n and $q(t)$ is a nonzero polynomial of degree at most $n - 1$, this is impossible. \square

LEMMA 2.5. For $n, m \geq 1$, let

$$A = \left(\begin{array}{c|ccc} a_1 & a_2 & \dots & a_n \\ \hline a_2 & & & \\ \vdots & & & * \\ a_n & & & \end{array} \right), \quad B = \left(\begin{array}{c|ccc} b_1 & b_2 & \dots & b_m \\ \hline b_2 & & & \\ \vdots & & & * \\ b_m & & & \end{array} \right)$$

be square integer matrices of dimension n and m , respectively, $a_1, b_1 \geq 1$. Furthermore, let $\alpha, \beta \neq 0$ such that

$$M = \left(\begin{array}{c|ccc|ccc} 0 & \alpha a_1 & \dots & \alpha a_n & \beta b_1 & \dots & \beta b_m \\ \hline \alpha a_1 & & & & & & \\ \vdots & & & A & & & \\ \alpha a_n & & & & & & \\ \hline \beta b_1 & & & & & & \\ \vdots & & & & & & B \\ \beta b_m & & & & & & \end{array} \right)$$

is a matrix with integer coefficients. If M is nonnegative and irreducible, and if $\chi_A, \chi_B \in \mathbb{Z}[t]$ are irreducible and distinct, then the characteristic polynomial of $\widetilde{M} = M + ay^2 E_{11}$ is irreducible in $\mathbb{Z}[t, y]$ for all $0 \neq a \in \mathbb{Z}$.

PROOF. The proof is similar to the proof of Lemma 2.4: the only thing to verify is that no eigenvalue of A or of B is also an eigenvalue of M . Again,

we compute

$$\chi_M(t) = t\chi_A(t)\chi_B(t) \pm q_1(t)\chi_B(t) \pm q_2(t)\chi_A(t),$$

where $q_1(t) \in \mathbb{Z}[t]$ is of degree at most $n - 1$ and $q_2(t) \in \mathbb{Z}[t]$ is of degree at most $m - 1$. This is seen by developing the first column of the matrix $tI - M$. The first coefficient is responsible for the summand $t\chi_A(t)\chi_B(t)$, the next n coefficients are responsible for the summand $\pm q_1(t)\chi_B(t)$ and the final m coefficients are responsible for the summand $\pm q_2(t)\chi_A(t)$. We claim that neither among $q_1(t)$ and $q_2(t)$ is the zero polynomial. Indeed, by developing the first column of the matrix $tI - M$, and evaluating at $t = 0$, we directly calculate

$$\begin{aligned} q_1(0) &= \det \left(\begin{array}{c|ccc} 0 & -\alpha a_1 & \dots & -\alpha a_n \\ -\alpha a_1 & & & \\ \vdots & & & \\ -\alpha a_n & & & \end{array} \right) \\ &= \det \left(\begin{array}{c|ccc} \alpha^2 a_1 & 0 & \dots & 0 \\ 0 & & & \\ \vdots & & & \\ 0 & & & \end{array} \right) = \pm \alpha^2 a_1 \det A, \end{aligned}$$

which is not zero since χ_A is irreducible. Similarly, $q_2(0) \neq 0$. Now if there existed a common root $\lambda \in \mathbb{C}$ of χ_M and χ_A , then λ would also be a root of either χ_B or q_1 . Since χ_A and χ_B are irreducible and distinct, we must have $q_1(\lambda) = 0$. But since χ_A is irreducible of degree n , and $q_1(t)$ is a nonzero polynomial of degree at most $n - 1$, this is impossible. Similarly, no root of χ_B can be a root of χ_M , which concludes the proof. \square

REMARK 2.6. One can formulate Lemma 2.5 in more generality. Namely, instead of two blocks A and B , one can assume any number $k \geq 2$ of blocks A_1, \dots, A_k of blocks of respective sizes n_1, \dots, n_k . In this case, all the k characteristic polynomials χ_{A_i} of the blocks need to be assumed irreducible and pairwise distinct. The argument for irreducibility is then identical by considering

$$\chi_M(t) = t \prod_{i=1}^k \chi_{A_i} + \sum_{i=1}^k \pm q_i(t) \prod_{j \neq i} \chi_{A_j},$$

where $q_i(t) \in \mathbb{Z}[t]$ is of degree at most $n_i - 1$ and nonzero.

2. Nonsplitting criteria

In this section, we give two different nonsplitting criteria providing that under certain controllable conditions, the degree of the field extension

$$\mathbb{Q}(\lambda) : \mathbb{Q}(\lambda + \lambda^{-1})$$

equals two.

2.1. The first nonsplitting criterion. The goal of this section is to present an algebraic criterion that allows us to deduce that the degree of the field extension $\mathbb{Q}(\lambda) : \mathbb{Q}(\lambda + \lambda^{-1})$ equals two for certain completely explicit products of multitwists. Let $\alpha = \alpha_1 \cup \dots \cup \alpha_n$ and $\beta = \beta_1 \cup \dots \cup \beta_m$ be two multicurves with n and m components, respectively, that fill a surface S and intersect minimally. Let X be their geometric intersection matrix, that is, the $n \times m$ matrix whose ij -th coefficient equals the geometric intersection number of α_i and β_j . We assume that the Perron–Frobenius eigenvalue μ^2 of XX^\top is of degree d . Furthermore, we let $\Omega = \begin{pmatrix} 0 & X \\ X^\top & 0 \end{pmatrix}$. For a symmetric matrix A , we denote by $\sigma(A)$ its signature, that is, the number of positive eigenvalues minus the number of negative eigenvalues. We will also denote by $\text{null}(A)$ its nullity, that is, the dimension of its kernel.

LEMMA 2.7. *The following properties hold.*

- (1) *The number $\sigma(\Omega + 2I) + \text{null}(\Omega + 2I)$ equals the number of eigenvalues of Ω in the interval $[-2, 2]$.*
- (2) *The eigenvalues λ_i of $M = \begin{pmatrix} I & X \\ 0 & I \end{pmatrix} \begin{pmatrix} I & 0 \\ -X^\top & I \end{pmatrix}$ are related to the eigenvalues μ_i of Ω by the equation $\mu_i^2 = 2 - \lambda_i - \lambda_i^{-1}$.*

PROOF. The first property is exactly Lemma 3.7 in the authors' PhD thesis [35]. The second property is Proposition 3.3(b) in [35]; as the proof in [35] does not explicitly deal with the case where M is not diagonalisable, we present a complete argument here. We first calculate

$$M = \begin{pmatrix} I & X \\ 0 & I \end{pmatrix} \begin{pmatrix} I & 0 \\ -X^\top & I \end{pmatrix} = \begin{pmatrix} I - XX^\top & X \\ -X^\top & I \end{pmatrix}$$

and note that its inverse is given by

$$M^{-1} = \begin{pmatrix} I & -X \\ X^\top & I - X^\top X \end{pmatrix}.$$

One directly verifies the equation $\Omega^2 = 2I - M - M^{-1}$. In order to obtain the same equation for all the eigenvalues (counting multiplicity), we change

basis such that M is in Jordan normal form. Note that in the new basis, also the matrix M^{-1} becomes a block diagonal matrix, where all the blocks are of upper triangular form and correspond to the Jordan blocks of M . In particular, also the matrix Ω^2 becomes upper triangular in the new basis, and the equation for the eigenvalues, $\mu_i^2 = 2 - \lambda_i + \lambda_i^{-1}$ (counting multiplicity), can be read off from the diagonal entries of the matrix equation. \square

Our criterion for the construction of pseudo-Anosov maps with stretch factors of controlled degree is the following.

THEOREM 2.8. *Let $\alpha, \beta \subset S$ be a pair of filling multicurves. Let X be their geometric intersection matrix, let d be their multicurve intersection degree and let $\Omega = \begin{pmatrix} 0 & X \\ X^\top & 0 \end{pmatrix}$. If*

$$\dim(\Omega) > \sigma(\Omega + 2I) + \text{null}(\Omega + 2I) > \dim(\Omega) - 2d,$$

then the mapping class $T_\alpha \circ T_\beta$ is pseudo-Anosov with stretch factor λ of degree $2d$.

REMARK 2.9. This criterion is particularly strong in case $n = m = d$, that is, when α and β have the same number of components and if the characteristic polynomial of the matrix XX^\top is irreducible. In this case,

$$2n > \sigma(\Omega + 2I) + \text{null}(\Omega + 2I) > 0$$

is sufficient to ensure that the mapping class $T_\alpha \circ T_\beta$ is pseudo-Anosov with stretch factor λ of degree $2d$.

PROOF. We first ensure that the mapping class $T_\alpha \circ T_\beta$ is pseudo-Anosov. If $n+m > \sigma(\Omega+2I)+\text{null}(\Omega+2I)$, then Ω has an eigenvalue outside the interval $[-2, 2]$ by (1) of Lemma 2.7. In particular, the dominating eigenvalue μ of Ω is larger than 2 and the matrix product $\begin{pmatrix} 1 & \mu \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -\mu & 1 \end{pmatrix}$ is hyperbolic, as its trace $2 - \mu^2$ is larger than 2 in modulus. Hence, the mapping class $T_\alpha \circ T_\beta$ is pseudo-Anosov by the Thurston–Veech construction [65, 68].

Now, let λ be the stretch factor of the mapping class $T_\alpha \circ T_\beta$. By the Thurston–Veech construction, we have $\lambda + \lambda^{-1} = \mu^2 - 2$. In particular, we directly observe $\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\mu^2)$. Furthermore, the degree of the field extension $\mathbb{Q}(\lambda) : \mathbb{Q}(\lambda + \lambda^{-1})$ is either 1 or 2. It equals 2, which is what we want to show, exactly if λ and λ^{-1} are Galois conjugates.

We now finish the proof by arguing that λ and λ^{-1} are indeed Galois conjugates. By (2) of Lemma 2.7, the dilatation λ is also the leading eigenvalue of $-M$, where M is the matrix product given in the (2) of Lemma 2.7.

In particular, the Galois conjugates of λ are among the eigenvalues $-\lambda_i$ of the matrix $-M$. These eigenvalues are in turn related to the eigenvalues μ_i of Ω by the equation $\mu_i^2 = 2 + \lambda_i + \lambda_i^{-1}$, again by Lemma 2.7. Since we have $\sigma(\Omega + 2I) + \text{null}(\Omega + 2I) > n + m - 2d$, the matrix Ω has at most $2d - 1$ eigenvalues outside the interval $[-2, 2]$. Via the correspondence in Lemma 2.7, the matrix $-M$ has at most $2d - 1$ eigenvalues that do not lie on the unit circle. In particular, one of the $2d$ Galois conjugates of λ or λ^{-1} (including λ and λ^{-1} themselves) must be on the unit circle by the pigeonhole principle. Thus the minimal polynomial of λ or λ^{-1} (and hence of both) is reciprocal and it follows that λ and λ^{-1} are Galois conjugates. \square

2.2. The second nonsplitting criterion. In this section we prove Theorem 1.4, which is an algebraic criterion that allows us to deduce that the degree of the field extension $\mathbb{Q}(\lambda(f)) : \mathbb{Q}(\lambda(f) + \lambda(f)^{-1})$ equals two for certain (nonexplicit) pseudo-Anosov maps f which are a product of multitwists.

PROOF OF THEOREM 1.4. By the Thurston–Veech construction, there exists a representation $\rho : \langle T_\alpha, T_\beta \rangle \rightarrow \text{PSL}_2(\mathbb{R})$ mapping T_α to the matrix $\begin{pmatrix} 1 & r \\ 0 & 1 \end{pmatrix}$ and T_β to the matrix $\begin{pmatrix} 1 & 0 \\ -r & 1 \end{pmatrix}$, where $r^2 = \mu$ is the spectral radius of the matrix XX^\top for the multicurves α and β . Furthermore, the stretch factor $\lambda(f)$ of $f \in \langle T_\alpha, T_\beta \rangle$ equals the spectral radius of $\rho(f)$. Now, let us consider the product of multitwists $f = T_\alpha^{2n} \circ T_\beta^{2n\varepsilon}$. A direct computation provides that the trace of $\rho(f)$ equals $\text{tr}(\rho(f)) = 2 - \varepsilon(2nr)^2$. Thus, $\lambda(f) + \lambda(f)^{-1} = |2 - \varepsilon(2nr)^2|$ and hence

$$\mathbb{Q}(\lambda(f) + \lambda(f)^{-1}) = \mathbb{Q}(\mu) = K.$$

Note that by assumption, the degree of the field extension $K : \mathbb{Q}$ is d , the multicurve intersection degree of α and β .

Since $\lambda = \lambda(f)$ solves the quadratic equation $t^2 - (\lambda + \lambda^{-1})t + 1 = 0$, λ has degree 1 or 2 over K . All what we need to do is find $n \in \mathbb{Z}_{>0}$ such that $\lambda \notin K$, or equivalently such that the discriminant

$$D = (2 - \varepsilon(2nr)^2)^2 - 4 = 16 \cdot n^2 \cdot ((n\varepsilon\mu)^2 - \varepsilon\mu)$$

of the quadratic equation is not a square in K . We will proceed by contradiction. Let $\mu' = \varepsilon\mu$ and let us assume that $(n\mu')^2 - \mu'$ is a square in $K = \mathbb{Q}(\mu')$ for every $n > 0$. Since the expression is invariant under the transformation $n \mapsto -n$, we can assume the expression is a square for every $n \in \mathbb{Z} \setminus \{0\}$.

Let $P = a_d t^d + a_{d-1} t^{d-1} + \cdots + a_1 t + a_0 \in \mathbb{Q}[t]$ be the minimal polynomial of μ' over \mathbb{Q} with $a_d = 1$. The Thurston–Veech construction implies that μ is an eigenvalue of a square matrix, so μ' is also an eigenvalue of a square matrix with integer coefficients, implying that $P \in \mathbb{Z}[t]$. Thus, μ' and $n^2\mu' - 1$ are algebraic integers, and we can consider their norms $N(\cdot)$. Let us compute the norm of $n^2\mu' - 1$. The monic polynomial $n^{2d}P\left(\frac{t+1}{n^2}\right) \in \mathbb{Z}[t]$ has degree d and admits $n^2\mu' - 1$ as a root. Moreover the fields $\mathbb{Q}(\mu')$ and $\mathbb{Q}(n^2\mu' - 1)$ coincide, so they have the same degree over \mathbb{Q} , namely d . Hence $n^{2d}P\left(\frac{t+1}{n^2}\right)$ is the minimal polynomial of $n^2\mu' - 1$. Inspecting its constant term, we deduce

$$N(n^2\mu' - 1) = (-1)^d \sum_{k=0}^d a_k n^{2d-2k}.$$

Since the norm is multiplicative, we have

$$N((n\mu')^2 - \mu') = N(n^2\mu' - 1) \cdot N(\mu') = a_0 \sum_{k=0}^d a_k n^{2d-2k} = Q(n^2),$$

where

$$Q(t) = a_0 \sum_{k=0}^d a_k t^{d-k}.$$

Since by assumption $(n\mu')^2 - \mu'$ is a square for every $n \in \mathbb{Z} \setminus \{0\}$, we deduce that $Q(n^2)$ is a square for every $n \in \mathbb{Z} \setminus \{0\}$. We now show that also $Q(0) = N(-\mu')$ is a square. Indeed, for any prime integer p , the reduction modulo $n = p$ of $N((n\mu')^2 - \mu')$ gives that $Q(0) = N(-\mu')$ is a quadratic residue. Thus it is also a square in \mathbb{Z} . Hence $Q(t) \in \mathbb{Z}[t]$ is a polynomial taking integral square value at every integer specialisation. By a result of Murty [47, Theorem 4], $Q(t^2)$ is the square of a polynomial.

Moreover, we observe that $Q(t) = a_0 \cdot t^d P\left(\frac{1}{t}\right) \in \mathbb{Z}[t]$. In particular $Q\left(\frac{1}{\mu'}\right) = 0$ and Q has degree d over \mathbb{Q} . Since μ' and $\frac{1}{\mu'}$ generate the field extension $K : \mathbb{Q}$, the minimal polynomial q of $\frac{1}{\mu'}$ over \mathbb{Q} is also of degree d . This implies that Q must be a rational multiple of q , and it is in particular separable. Now each root $0 \neq a \in \mathbb{C}$ of Q gives rise to two distinct roots $\pm\sqrt{a}$ of $Q(t^2)$, and conversely. Thus $Q(t^2)$ is also separable, and cannot be a square. This concludes the proof of the theorem. \square

3. Warm-up: proof of Thurston's claim

As a first illustration of our method, it is the goal of this section to provide a pair of filling multicurves α and β on S_g with multicurve intersection degree $3g - 3$. By Theorem 1.4, this validates Thurston's claim that

the product of two multitwists can realise the maximal possible algebraic degrees of stretch factors: $6g - 6$.

Recall that the matrix encoding the number of intersections of the components of α and β is denoted by $X = (|\alpha_i \cap \beta_j|)_{ij}$ (see Section 1.2). In order to read off the matrix XX^\top from our figures, we use the following formula for its coefficients, which is a direct consequence of the definition of matrix multiplication:

$$(XX^\top)_{ij} = \sum_k |\alpha_i \cap \beta_k| \cdot |\beta_k \cap \alpha_j|.$$

We start by realising, on the surface of genus $g \geq 1$ with 2 boundary components, a pair of filling multicurves α and β such that $\chi_{XX^\top} \in \mathbb{Z}[t]$ is irreducible and of degree $3g - 1$. On a surface with boundary, we say that a pair of multicurves is filling if they intersect transversally and if their complement consists of boundary-parallel annuli and discs none of which is a bigon. We proceed by induction on g .

For $g = 1$ with two boundary components. We consider the filling pair of multicurves α and β shown in Figure 2.1, where one of the components of β has $y - 1$ parallel copies. Here, we think of y as a variable that we specify later on.




FIGURE 2.1. Two multicurves α (in red) and β (in blue) on the surface of genus one with two boundary components. The multicurve β contains $y - 1$ parallel copies of one of its components.

We directly calculate

$$XX^\top = \begin{pmatrix} 4 & 2 \\ 2 & y \end{pmatrix}.$$

Observe that X is a matrix of size $2 \times y$ (the multicurve β has y components). We have $\chi_{XX^\top}(t) = t^2 - (4+y)t + 4(y-1)$ with discriminant $y^2 - 8y + 32$, which is not a square if $y \geq 12$. Indeed, in this case we have

$$(y-3)^2 = y^2 - 6y + 9 > y^2 - 8y + 32 > y^2 - 8y + 16 = (y-4)^2.$$

In particular, for $y \geq 12$ the polynomial χ_{XX^\top} is irreducible.

For $g > 1$ and two boundary components. For the inductive step, assume we have constructed on the surface of genus $g \geq 1$ with 2 boundary components a filling pair of multicurves α', β' such that the characteristic polynomial $\chi' = \chi_{XX^\top} \in \mathbb{Z}[t]$ is irreducible and of degree $3g - 1$. Furthermore, assume that α'_1 is a simple closed curve that encircles all the handles of the surface, as illustrated in Figure 2.2. Take a surface of genus 1 and two boundary components, as in the case of genus $g = 1$, see Figure 2.1, and denote its filling pair of multicurves by α'' and β'' . Now glue its right boundary component to the left boundary component of the surface of genus g , and add two new curves α_0 and β_0 to the multicurves. The curve α_0 encircles all the handles of the newly formed surface, and the curve β_0 twice intersects α_0 but no other multicurve component. Again, see Figure 2.2 for an illustration. By construction, the pair of multicurves α, β is again filling.




FIGURE 2.2. Two surfaces of genus g and 1, respectively, and two boundary components, glued together along one of their boundary components. The curves α'_1 and α''_1 are shown, each encircling all the handles of their respective surface. The new curve α_0 encircles all the handles of the newly formed surface, and the new curve β_0 runs along the glued boundary component.

Let A be the matrix XX^\top for the pair of multicurves α', β' , and let B be the matrix XX^\top for the pair of multicurves α'', β'' . We define the multicurves

$$\begin{aligned}\alpha &= \alpha_0 \cup \alpha' \cup \alpha'' \\ \beta &= \beta_0 \cup \beta' \cup \beta''\end{aligned}$$

A quick computation gives

$$A = \left(\begin{array}{c|ccc} a_1 & a_2 & \dots & a_n \\ \hline a_2 & & & \\ \vdots & & & \\ a_n & & * & \end{array} \right), \quad B = \begin{pmatrix} 4 & 2 \\ 2 & b \end{pmatrix}.$$

Let us choose b such that χ_B is irreducible and distinct from χ_A . We may assume inductively that $a_1 = 4a$. In the multicurve β , we take $y^2 - a - 1 \geq 1$ parallel copies of β_0 , for $y > 0$ large enough. The matrix XX^\top for the multicurves α and β takes the form

$$XX^\top = \left(\begin{array}{c|ccc|cc} 4y^2 & a_1 & \dots & a_n & 4 & 2 \\ \hline a_1 & & & & & \\ \vdots & & & & & \\ a_n & & & & & \\ \hline 4 & & & & 4 & 2 \\ 2 & & & & 2 & b \end{array} \right).$$

By Lemma 2.5, $\chi_{XX^\top} \in \mathbb{Z}[t, y]$ is irreducible (recall that χ_A is irreducible). Hence, by Hilbert's irreducibility theorem [28], there exist infinitely many specifications of y (and in particular infinitely many specifications of y such that $y^2 - a - 1 > 0$) with $\chi_{XX^\top} \in \mathbb{Z}[t]$ irreducible. This polynomial is of degree

$$3g - 1 + 3 = 3(g + 1) - 1,$$

which is exactly what we wanted to show. Finally, to justify our inductive assumption on the top-left coefficient of the matrix A , note that the top-left coefficient of the matrix XX^\top is again a multiple of 4. Furthermore, we note that a_0 is again a curve that encircles all handles of the newly built surface of genus $g + 1$ with two boundary components.

The closed case for $g \geq 2$. Take any example of a filling pair of multicurves α' and β' we constructed on the surface of genus $g - 1 \geq 1$ with two boundary components in Section 3. In particular, α'_1 encircles all the $g - 1$

handles of the surface, and if

$$A = \left(\begin{array}{c|ccc} a_1 & a_2 & \dots & a_n \\ \hline a_2 & & & \\ \vdots & & * & \\ a_n & & & \end{array} \right)$$

is the matrix XX^\top for the multicurves α' and β' , we have $a_1 = 4a$ and $\chi_A(t)$ is irreducible. We identify the two boundary components of the surface to increase the genus by one. Let α_0 be a longitude of the created handle, and let β_0 run along the glued boundary, see Figure 2.3. Define the two new multicurves

$$\begin{aligned} \alpha &= \alpha_0 \cup \alpha' \\ \beta &= \beta_0 \cup \beta', \end{aligned}$$

where we take $y^2 - a$ copies of β_0 . By construction, the pair of multicurves α, β is again filling. The matrix XX^\top for the multicurves α and β




FIGURE 2.3. The closed surface of genus g is obtained by identifying the left and the right boundary component of the depicted closed surface of genus $g - 1$ with two boundary components. The curve α_0 is a longitude of the created handle, and the curve β_0 runs along the glued boundary.

takes the form

$$XX^\top = \left(\begin{array}{c|ccc} y^2 & \frac{a_1}{2} & \dots & \frac{a_n}{2} \\ \hline \frac{a_1}{2} & & & \\ \vdots & & A & \\ \frac{a_n}{2} & & & \end{array} \right),$$

and $\chi_{XX^\top} \in \mathbb{Z}[t, y]$ is irreducible by Lemma 2.4. By Hilbert's irreducibility theorem [28], there exist infinitely many specifications of y (and in particular infinitely many specifications of y such that $y^2 - a > 0$) with $\chi_{XX^\top} \in \mathbb{Z}[t]$ irreducible. This polynomial is of degree $3(g - 1) - 1 + 1 = 3g - 3$.

4. Strata of Abelian differentials

In this section, we present a proof of Theorem 1.6 and Theorem 1.7, for each stratum of Abelian differentials. We postpone the more intricate analysis of the connected components to Section 5.

Let $\mathcal{H}(k_1, k_2, \dots, k_m)$ be a stratum of Abelian differentials. Recall that the number of odd k_i must itself be even, say $2l$. Furthermore, if g is the genus of the underlying topological surface, we have $2g - 2 = \sum_{i=1}^m k_i$.

4.1. Constructing a surface. We start by constructing a square-tiled surface. First, we ensure that we land in the stratum $\mathcal{H}(k_1, k_2, \dots, k_m)$. We start out with a long horizontal square-tiled surface with some large number $y^2 - g + 1$ of squares and opposite side identifications, see Figure 2.4. The surface obtained by identifying the sides is a torus, and there are no




FIGURE 2.4. A horizontal square-tiled surface.

singularities of the flat structure. We can add an angle of 4π to some marked point by inserting a vertical strip of $y_i + 1$ square tiles, as in Figure 2.5. We




FIGURE 2.5. Inserting a vertical strip of squares creates a cone point with angle 6π out of a marked point.

treat the $y_i \geq 1$ as variables that we will need to specify later on. This operation can be repeated in order to add an integer multiple of 4π to the angle around any cone point or marked point. For example, Figure 2.6 indicates how to insert another vertical strip of square tiles in order to add 4π to the cone angle around a cone point with angle 6π . Iterating this procedure, we can reach all strata with even multiplicities.




FIGURE 2.6. A vertical strip of squares can be inserted in order to add another 4π to the angle around a cone point.




FIGURE 2.7. Inserting an L-shaped square-tiled surface creates two cone angles of 4π out of one marked point.

In order to create odd multiplicities, we insert an L-shaped square-tiled surface with $y_i + 1$ tiles, $y_i \geq 2$, as shown in Figure 2.7. This creates two cone points of angle 4π , which is multiplicity one. Recall that there must be an even number $2l$ of odd multiplicities k_i , so we can repeat this step l times to have the right number of odd multiplicities, and then successively add two to the multiplicities by inserting vertical strips as above, until we reach the stratum $\mathcal{H}(k_1, k_2, \dots, k_m)$. Following this procedure, we need to add a total of l L-shapes and $g - l - 1$ vertical strips.

4.2. Calculating the polynomial. The square-tiled surface we construct in Section 4.1 naturally decomposes into horizontal and vertical annuli that are one square wide. Let X be the intersection matrix for the core curves α_i of the horizontal annuli and the core curves β_j of the vertical annuli. We index the rows by horizontal curves and the columns by

vertical curves. We now describe the matrix XX^\top . Since the curves α_i and β_j pairwise intersect in a tree-like fashion, we use the following way of looking at the computation. The i -th diagonal coefficient equals the number of vertical curves intersecting the i -th horizontal curve α_i . Furthermore, an off-diagonal ij -th coefficient is equal to 1 if there exists a vertical curve intersecting both horizontal curves α_i and α_j . Otherwise, it equals 0.

In order to write down the matrix XX^\top , we quickly recall our construction. We have one horizontal curve that we start with. It intersects y^2 vertical curves. We further have one horizontal curve for each L-shaped surface we inserted, of which there are l in total. These curves respectively intersect y_i vertical curves, for $i = 1, \dots, l$, and are linked to the starting horizontal curve via an intersecting vertical curve.

For example, if we insert two L-shaped surfaces with $y_1 + 1$ and $y_2 + 1$ tiles, respectively, we obtain the matrix

$$XX^\top = \begin{pmatrix} y^2 & 1 & 1 \\ 1 & y_1 & 0 \\ 1 & 0 & y_2 \end{pmatrix}$$

with characteristic polynomial obtained by developing the first column of the matrix $tI - XX^\top$:

$$\begin{aligned} (t - y^2)(t - y_1)(t - y_2) - (t - y_2) - (t - y_1) &= \\ &= -y^2 \prod_{i=1}^2 (t - y_i) + t \prod_{i=1}^2 (t - y_i) - \sum_{i=1}^2 \prod_{j \neq i} (t - y_j). \end{aligned}$$

It is straightforward to generalise the last form of the characteristic polynomial to an arbitrary number l of inserted L-shaped surfaces.

Conveniently, the form of the characteristic polynomial turns out to be basically the same even if we insert vertical strips, but this needs a more careful calculation. We first describe the coefficients of the matrix XX^\top we get from inserting vertical strips: for each vertical surface we insert, we get another y_i horizontal curves, all intersecting a single vertical curve that also intersects the starting horizontal curve. Here, i runs from $l + 1$ to $g - 1$. We present the matrix using parameters $b, b_i \in \mathbb{R}$. These parameters are helpful in the proof of Lemma 2.12, and later in Section 5.3. For the purpose of the calculation of XX^\top in this section, we simply have $b = b_i = 1$ for all i . We write $\mathbf{b}_{n \times m}$ for the $n \times m$ matrix with all entries equal to $b \in \mathbb{R}$. In case $n = m$, we simplify and write \mathbf{b}_n .

DEFINITION 2.10. For parameters $b, b_i \in \mathbb{R}$, $i = 1, \dots, l$, we consider the matrix

$$XX^\top = \left(\begin{array}{cccc|cccc} y^2 & b_1 & \cdots & b_l & \mathbf{1}_{1 \times y_{l+1}} & \mathbf{1}_{1 \times y_{l+2}} & \cdots & \mathbf{1}_{1 \times y_{g-1}} \\ b_1 & y_1 & & & & & & \\ \vdots & & \ddots & & & & & \\ b_l & & & y_l & & & & \\ \hline \mathbf{1}_{y_{l+1} \times 1} & & & & \mathbf{1}_{y_{l+1}} & & & \\ \mathbf{1}_{y_{l+2} \times 1} & & & & & \mathbf{1}_{y_{l+2}} & & \\ \vdots & & & & & & \ddots & \\ \mathbf{1}_{y_{g-1} \times 1} & & & & & & & \mathbf{1}_{y_{g-1}} \end{array} \right).$$

For the characteristic polynomial of the matrix XX^\top , we have the following result.

LEMMA 2.11. The characteristic polynomial of XX^\top equals

$$p(t, y, \mathbf{y}) = t^a \left(-y^2 \prod_{i=1}^{g-1} (t - y_i) + t \prod_{i=1}^{g-1} (t - y_i) - \sum_{i=1}^{g-1} c_i \prod_{j \neq i} (t - y_j) \right),$$

where $a = \sum_{i=l+1}^{g-1} (y_i - 1)$, $c_{l+1} = y_{l+1} b^2$, $c_i = y_i$ for $i \geq l+2$ and $c_i = b_i^2$ otherwise.

PROOF. This calculation is slightly tedious, but obtained in a fairly straightforward manner by developing the first column of $(tI - XX^\top)$. We begin by observing that the determinants of the $y_i \times y_i$ matrices

$$tI_{y_i} - \mathbf{1}_{y_i} = \begin{pmatrix} t-1 & -1 & \cdots & -1 \\ -1 & t-1 & & \\ \vdots & & \ddots & \\ -1 & & & t-1 \end{pmatrix},$$

$$M_{y_i}(t) = \begin{pmatrix} -1 & -1 & \cdots & -1 \\ -1 & t-1 & & \\ \vdots & & \ddots & \\ -1 & & & t-1 \end{pmatrix}$$

are respectively given by the polynomials $t^{y_i-1}(t - y_i)$ and $-t^{y_i-1}$. The former calculation follows by inspecting the eigenvalues of the matrix $\mathbf{1}_{y_i}$, and the latter is derived by solving the equation

$$\det(tI_{y_i} - \mathbf{1}_{y_i}) = t \det(tI_{y_i-1} - \mathbf{1}_{y_i-1}) + \det(M_{y_i}(t)).$$

We note that changing the diagonal coefficient (-1) of the matrix $M_{y_i}(t)$ with some other diagonal coefficient $(t-1)$ does not change the determinant.

Now, by developing the first column of $(tI - XX^\top)$, we get that the characteristic polynomial of XX^\top has the following summands. The first summand (obtained by deleting the first row and the first column when developing) equals

$$(t - y^2) \prod_{i=1}^l (t - y_i) \prod_{i=l+1}^{g-1} \det(tI_{y_i} - \mathbf{1}_{y_i}) = t^a (t - y^2) \prod_{i=1}^{g-1} (t - y_i),$$

where $a = \sum_{i=l+1}^{g-1} (y_i - 1)$. The rest of the summands are obtained as follows. Assume that in the development we delete the first column and the k -th row, where $k \geq 2$. We have to take the determinant of the matrix obtained by deleting the k -th row of the matrix

$$\left(\begin{array}{ccc|cccc} -b_1 & \cdots & -b_l & -\mathbf{1}_{1 \times y_{l+1}} & -\mathbf{1}_{1 \times y_{l+2}} & \cdots & -\mathbf{1}_{1 \times y_{g-1}} \\ t - y_1 & & & & & & \\ & \ddots & & & & & \\ & & t - y_l & & & & \\ \hline & & & tI_{y_{l+1}} - \mathbf{1}_{y_{l+1}} & & & \\ & & & & tI_{y_{l+2}} - \mathbf{1}_{y_{l+2}} & & \\ & & & & & \ddots & \\ & & & & & & tI_{y_{g-1}} - \mathbf{1}_{y_{g-1}} \end{array} \right).$$

After switching adjacent rows (a total of $k - 2$ times) to move the first row to be the $k - 1$ st one, the matrix obtained is almost of block diagonal form and we can read off the determinant. For the rows $k = 2, \dots, l + 1$, we obtain the summand

$$\begin{aligned} (-b_{k-1})(-1)^{1+k+k-2} & \left(\prod_{j \neq k-1, 1 \leq j \leq l} (t - y_j) \prod_{i=l+1}^{g-1} \det(tI_{y_i} - \mathbf{1}_{y_i}) \right) (-b_{k-1}) = \\ & = -b_{k-1}^2 t^a \prod_{j \neq k-1} (t - y_j). \end{aligned}$$

For the rows $k > l + y_{l+1}$, we obtain summands of the form

$$\begin{aligned} (-1)(-1)^{1+k+k-2} & \left(\prod_{j=1}^l (t - y_j) \prod_{j \neq i, l+1 \leq j \leq g-1} \det(tI_{y_j} - \mathbf{1}_{y_j}) \right) \det(M_{y_i}(t)) = \\ & = -t^a \prod_{j \neq i} (t - y_j). \end{aligned}$$

Here, we assume for the calculation that the k -th row intersects the diagonal block $tI_{y_i} - \mathbf{1}_{y_i}$, where $i \geq l + 2$. There are a total of y_i summands of this type. If the k -th row intersects the block $tI_{y_{l+1}} - \mathbf{1}_{y_{l+1}}$, the corresponding constant vectors of the first row and the first column have coefficients $b \in \mathbb{R}$. In this case, we obtain y_{l+1} times the summand

$$-b^2 t^a \prod_{j \neq l+1} (t - y_j).$$

Adding all summands, we finally obtain the polynomial

$$t^a \left((t - y^2) \prod_{i=1}^{g-1} (t - y_i) - \sum_{i=1}^{g-1} c_i \prod_{j \neq i} (t - y_j) \right),$$

where $a = \sum_{i=l+1}^{g-1} (y_i - 1)$, $c_{l+1} = y_{l+1} b^2$, $c_i = y_i$ for $i \geq l + 2$ and $c_i = b_i^2$ otherwise. \square

LEMMA 2.12. *Let $k \geq 1$, and let $y_i, c_i \in \mathbb{Z}$ for $i = 1, \dots, k$ such that all y_i are pairwise distinct and all c_i are positive. Then the polynomial*

$$p(t, y) = -y^2 \prod_{i=1}^k (t - y_i) + t \prod_{i=1}^k (t - y_i) - \sum_{i=1}^k c_i \prod_{j \neq i} (t - y_j)$$

is irreducible in $\mathbb{Z}[t, y]$.

PROOF. We regard the polynomial $p(t, y) \in \mathbb{Z}[t, y] \cong (\mathbb{Z}[t])[y]$ as a polynomial of degree two in the variable y , with coefficients in $\mathbb{Z}[t]$. We note that the coefficient of y^2 and the constant coefficient $p(t, 0)$ are relatively prime in $\mathbb{Z}[t]$. This follows from the observation that the roots of the coefficient of y^2 are exactly the y_i , while none of those numbers is a root of the constant coefficient. Indeed, we have

$$p(y_i, 0) = -c_i \prod_{j \neq i} (y_i - y_j) \neq 0.$$

This implies that the only possibility to factor $p(t, y)$ is by writing it as a product of two factors linear in the variable y . To rule this out, we apply Eisenstein's criterion as follows. The constant coefficient $p(t, 0)$ has a simple root: the Perron–Frobenius eigenvalue of a matrix of the form XX^\top as in Definition 2.10, where we set $l = g - 1 = k$, $b_i = \sqrt{c_i}$ and $y = 0$. Now, let $q(t) \in \mathbb{Z}[t]$ be the irreducible factor of $p(t, 0) \in \mathbb{Z}[t]$ containing this root. Then $q(t)$ divides the constant coefficient $p(t, 0)$ but $q(t)^2$ does not. Furthermore, $q(t)$ does not divide the coefficient of y^2 since otherwise it would have a root in common with the constant coefficient $p(t, 0)$. Eisenstein's criterion

now implies that $p(t, y)$ can not be factored into a product of two factors with positive degree in the variable y . \square

4.3. Main results for strata. We are now ready to prove the analogues of Theorem 1.7 and Theorem 1.6 for strata of Abelian differentials.

THEOREM 2.13. *Every number $1 \leq d \leq g$ is realised as the degree of the trace field of a product of two affine multitwists on a surface in every stratum of Abelian differentials on Riemann surfaces of genus g .*

PROOF. Let $\mathcal{H}(k_1, k_2, \dots, k_m)$ be a stratum of Abelian differentials. We use the surface constructed in Section 4.1. By the Thurston–Veech construction [65, 68], there exists a flat structure on it, obtained by changing the side lengths of the rectangles, such that the multitwists T_α and T_β have affine representatives, and such that the degree of the trace field is given by the degree of the Perron–Frobenius eigenvalue μ^2 of XX^\top .

Let $2 \leq d \leq g$ be the some degree of a trace field we want to construct. Set $g - d + 1$ of the $g - 1$ parameters y_i equal to 2 and all others > 2 and pairwise distinct. In this way, the characteristic polynomial of XX^\top can be factored as $(t - 2)^{g-d} p(t, y)$, where the polynomial $p(t, y)$ is of degree d in the variable t and with pairwise distinct y_i . In particular, Lemma 2.12 implies that $p(t, y)$ is irreducible as a polynomial in $\mathbb{Z}[t, y]$. Now, by Hilbert’s irreducibility theorem [28], there are infinitely many integer specifications of y such that the resulting polynomial is irreducible in $\mathbb{Z}[t]$. For $|y|$ large enough, all these specifications can be realised geometrically as in Section 4.1, since we start with $y^2 - g + 1$ squares in the construction. In particular, for every such y , we obtain an Abelian differential with trace field of degree d . \square

THEOREM 2.14. *Every even number $2 \leq 2d \leq 2g$ is realised as the degree of a product of two affine multitwists on a surface in every stratum of Abelian differentials on Riemann surfaces of genus g .*

PROOF. In the proof of Theorem 2.13, we have constructed examples with Perron–Frobenius eigenvalue μ^2 of XX^\top having degree d by letting exactly $g - d + 1$ of the parameters y_i equal to 2. For these examples, we now bound $\sigma(2I + \Omega)$ in order to apply Theorem 2.8 to $T_\alpha \circ T_\beta$. Let Ω' be the matrix obtained from Ω by deleting all the rows and all the columns corresponding to y or the $g - 1 - (g - d + 1) = d - 2$ parameters y_i that are not set equal to 2. We have

$$\sigma(\Omega + 2I) \geq \sigma(\Omega' + 2I) - (d - 1).$$

By construction, Ω' is the adjacency matrix of a forest consisting of path graphs (some of which might be of length zero). In particular, one directly verifies that $2I + \Omega'$ is positive definite. We get

$$\sigma(\Omega + 2I) \geq \sigma(\Omega' + 2I) - (d - 1) = n + m - 2d + 2 > n + m - 2d.$$

Furthermore, one directly checks that the matrix $\Omega + 2I$ has a negative eigenvalue as soon as $y > 4$, which we are allowed to assume. This implies

$$n + m > \sigma(\Omega + 2I) + \text{null}(\Omega + 2I).$$

Theorem 2.8 applies and the mapping class $T_\alpha \circ T_\beta$ is pseudo-Anosov with stretch factor λ of degree $2d$. \square

REMARK 2.15. The mapping classes we construct above are positive arborescent and so can all be obtained by capping off monodromies of certain fibred links called positive arborescent Hopf plumbings. This relation is discussed for example in the background chapter of the authors' PhD thesis [35]. The pseudo-Anosov stretch factors therefore appear as the dominating roots of the Alexander polynomials of these links. It is conceivable that our argument, or at least a portion thereof, could be replaced by a careful analysis of these Alexander polynomials using the skein relation. However, the calculations we present here can readily be applied to our examples in Section 5, which are not necessarily obtained from arborescent Hopf plumbings anymore.

5. Connected components of strata

In this section, we study the connected components of strata of Abelian differentials. After recalling the classification of the connected components, we first analyse to which connected components our examples from Section 4 belong. We then construct examples covering all remaining connected components, finally proving Theorem 1.6 and Theorem 1.7 in their full generality.

5.1. Classification of connected components of strata. The connected components of the strata of the moduli space of Abelian differentials are classified by Kontsevich and Zorich [27]. There are at most three connected components, and the classification uses two topological invariants that we describe now.

- (1) *Hyperellipticity.* For $g \geq 2$, the strata $\mathcal{H}(2g - 2)$ and $\mathcal{H}(g - 1, g - 1)$ have a component that consists entirely of hyperelliptic Riemann

surfaces, where the hyperelliptic involution permutes the two zeros (when there are two).

- (2) *Parity of the spin structure.* If the degrees of the singularities of a stratum are all even, then one can define a spin structure, or equivalently a quadratic form q on the first homology group. The parity of this spin structure (or the Arf invariant of the form) is a topological invariant.

REMARK 2.16. If a translation surface belongs to a hyperelliptic component $\mathcal{H}^{hyp}(2g-2)$ or $\mathcal{H}^{hyp}(g-1, g-1)$ and admits a cylinder decomposition, then all cylinders are fixed by the hyperelliptic involution, and each of them contains exactly two fixed points in its interior. Since the total number of fixed points is $2g+2$, this observation can be used to show that a translation surface does not belong to a hyperelliptic component.

We will use the topological definition of the spin structure (see [27, Section 3.1] for details) to have an effective way to compute its parity in terms of the Arf invariant of q . Since the flat metric (X, ω) has trivial holonomy, outside of finite number of singularities, we have a well-defined horizontal direction. Consider a smooth simple closed oriented curve α on X which does not contain any singularities. The total change of the angle between the tangent vector to α and the tangent vector to the horizontal is equal to $2\pi \cdot \text{ind}(\alpha)$, where $\text{ind}(\alpha) \in \mathbb{Z}$. If we choose any symplectic basis $(a_i, b_i)_{i=1, \dots, g}$ of $H_1(X; \mathbb{Z}/2\mathbb{Z})$, then the parity of the spin structure is [27, Equation (4)]:

$$(1) \quad \Phi(\omega) = \sum_{i=1}^g q(a_i)q(b_i) \pmod{2},$$

where $q(\alpha) = \text{ind}(\alpha) + 1$ for an oriented smooth path α . Together with the formula $q(\alpha + \beta) = q(\alpha) + q(\beta) + i(\alpha, \beta)$ for any $\alpha, \beta \in H_1(X; \mathbb{Z}/2\mathbb{Z})$, it is easy to calculate the parity of the spin structure given in any (non symplectic) basis of the first homology.

Next, we explain concretely how to compute $\Phi(\omega)$, where (X, ω) is obtained from the construction in Section 4. Observe that (X, ω) belongs to a non hyperelliptic component if $g > 2$: if $(X, \omega) \in \mathcal{H}^{hyp}(2g-2)$, then the number of cylinders we have inserted is $g-1$. By Remark 2.16 they contribute to $2g-2$ fixed points of the hyperelliptic involution (located on the $2g-2$ horizontal core curves), say $p_1, p'_1, \dots, p_{g-1}, p'_{g-1}$. There are two more fixed points q, q' on the horizontal core curve of the long cylinder \mathcal{C}

we start with, and one fixed point on its boundary, say q'' , that is on the same vertical closed curve as q' . The last fixed point is the singularity. On the other hand, each inserted cylinder should have two fixed points on its *vertical* core curve: one is p_i , the other one is $p_i'' \in \mathcal{C}$. Thus necessarily we have $p_i'' = q$ for all $i = 1, \dots, g-1$. This is possible only if $g-1 = 1$. Finally, for $(X, \omega) \in \mathcal{H}^{\text{hyp}}(g-1, g-1)$ the situation is similar.

5.2. Nonhyperelliptic components, spin 1. Consider (X, ω) obtained from the construction in Section 4 when all k_i are even. As a basis of the first homology $H_1(X, \mathbb{Z}/2\mathbb{Z})$, we take horizontal curves $\gamma_0, \dots, \gamma_{g-1}$, where γ_0 is the horizontal curve that we start with, and γ_i is in the i th vertical cylinder, and vertical curves $\eta_0, \dots, \eta_{g-1}$, where η_0 crosses γ_0 only once, and η_i is the core curve of the i th vertical cylinder for $i > 0$. By construction, for every i, j

$$i(\gamma_0, \eta_j) = 1, \quad i(\gamma_i, \eta_j) = \delta_{ij} \text{ for } i > 0, \quad \text{and } i(\gamma_i, \gamma_j) = i(\eta_i, \eta_j) = 0.$$


We can thus form a symplectic basis as follows:

$$\begin{cases} a_1 = \gamma_0, & b_1 = \eta_0 \\ a_i = \gamma_{i-1}, & b_i = \eta_{i-1} - \eta_0 \text{ for } i \neq 0 \end{cases}$$

Clearly $\text{ind}(\gamma_i) = \text{ind}(\eta_i) = 0$. Substituting in Equation (1), we conclude:

$$\begin{aligned} \Phi(\omega) &= 1 + \sum_{i=2}^g q(\gamma_{i-1})q(\eta_{i-1} - \eta_0) = \\ &= 1 + \sum_{i=2}^g (q(\eta_{i-1}) + q(\eta_0) + i(\eta_0, \eta_{i-1})) \equiv 1 \pmod{2}. \end{aligned}$$

5.3. Reaching the nonhyperelliptic component $\mathcal{H}(2k_1, \dots, 2k_m)$, spin $0, m > 1$. We now use a slightly different model defined as follows. Start with the surface depicted in Figure 2.8, with a long horizontal cylinder made out of $y^2 - g + 1$ squares. It belongs to $\mathcal{H}^{\text{hyp}}(2, 2)$. Its spin structure is 0 as we can check directly, or by using the formulae in [27, Corollary 5]. We can insert $g-3 \geq 1$ vertical strips of $y_i + 1$ square tiles (for $i = 3, \dots, g-1$) as in Section 4 in order to add zeros of even multiplicities and to reach the stratum $\mathcal{H}(2k_1, \dots, 2k_m)$ where $\sum 2k_i = 2g-2$. This construction does not change the spin structure as we can see on the computation below. We let γ_0 the horizontal core curve in the long cylinder, and $\gamma_1, \dots, \gamma_{g-1}$ the other horizontal core curves contained in the i th cylinder. Similarly, we let η_i for $i = 0, \dots, g-1$ the vertical core curves: η_0 is the core curve of the

FIGURE 2.8. A surface in $\mathcal{H}^{\text{hyp}}(2, 2)$ (with even spin structure).

vertical cylinder with label A_0 and η_i is the core curve of the i th vertical cylinder for $i > 0$. We have for every i, j

$$\begin{aligned} i(\gamma_0, \eta_i) &= 1 \text{ for } i \neq 1 \text{ and } i(\gamma_0, \eta_1) = 2, \\ i(\gamma_i, \eta_i) &= \delta_{ij} \text{ for } i > 0, \\ i(\gamma_i, \gamma_j) &= i(\eta_i, \eta_j) = 0. \end{aligned}$$

We can thus form a symplectic basis of $H_1(S; \mathbb{Z}/2\mathbb{Z})$ as follows:

$$\begin{cases} a_1 = \gamma_0, & b_1 = \eta_0 \\ a_2 = \gamma_1, & b_2 = \eta_1 \\ a_i = \gamma_{i-1}, & b_i = \eta_{i-1} - \eta_0 \text{ for } i > 2 \end{cases}$$

By using Equation (1) this leads to

$$\begin{aligned} \Phi(\omega) &= 1 + 1 + \sum_{i=3}^g q(\gamma_{i-1})q(\eta_{i-1} - \eta_0) = \\ &= 1 + 1 + \sum_{i=3}^g (q(\eta_{i-1}) + q(\eta_0) + i(\eta_0, \eta_{i-1})) \equiv 0 \pmod{2}. \end{aligned}$$

We now compute the degree of the trace field. In order to write down the matrix XX^\top , we apply the strategy described in Section 4.2. Observe that the horizontal curve that we start with crosses $y^2 - g + 1 + g - 3 = y^2 - 2$ squares. More precisely it intersects $y^2 - 4$ vertical curves once and one vertical curve twice. We obtain the following matrix, where $\mathbf{b}_{n \times m}$ stands

for the $n \times m$ matrix with all entries equal to $b \in \mathbb{Z}$:

$$XX^\top = \left(\begin{array}{c|cccc} y^2 & \mathbf{2}_{1 \times y_1} & \mathbf{1}_{1 \times y_2} & \cdots & \mathbf{1}_{1 \times y_{g-1}} \\ \mathbf{2}_{y_1 \times 1} & \mathbf{1}_{y_1 \times y_1} & & & \\ \mathbf{1}_{y_2 \times 1} & & \mathbf{1}_{y_2 \times y_2} & & \\ \vdots & & & \ddots & \\ \mathbf{1}_{y_{g-1} \times 1} & & & & \mathbf{1}_{y_{g-1} \times y_{g-1}} \end{array} \right).$$

From Lemma 2.11 with $l = 0$, we see that the characteristic polynomial of XX^\top equals $t^a \cdot p(t, y, \mathbf{y})$, where

$$p(t, y, \mathbf{y}) = -y^2 \prod_{i=1}^{g-1} (t - y_i) + t \prod_{i=1}^{g-1} (t - y_i) - \sum_{i=1}^{g-1} c_i \prod_{j \neq i} (t - y_j),$$

for $a = \sum_{i=1}^{g-1} (y_i - 1)$, $c_1 = 4y_1$ and $c_i = y_i$ if $i \geq 2$. From Lemma 2.12, we deduce that $p(t, y, \mathbf{y})$ is irreducible in $\mathbb{Z}[t, y]$ given that all $y_i \in \mathbb{N}$ are pairwise distinct (here our parameter b in Definition 2.10 equals 2). As before, we can factor out $(t - 2)^{g-d}$ and obtain an irreducible polynomial of degree d by setting $g - d + 1$ of the $g - 1$ parameters y_i equal to 2. We can then apply the same strategy than the proof of Theorem 2.13 to get the result.

COROLLARY 2.17. *Every number $1 \leq d \leq g$ is realised as the degree of the trace field of a product of two affine multitwists on a surface in every non-hyperelliptic connected component with spin 0 of a stratum, except $\mathcal{H}(2g-2)$, of Abelian differentials on Riemann surfaces of genus g .*

We further apply the same strategy to realise all even degrees as stretch factors. One can copy the proof of Theorem 2.14 word for word and obtain the following result.

COROLLARY 2.18. *Every even number $2 \leq 2d \leq 2g$ is realised as the degree of a product of two affine multitwists on a surface in every nonhyperelliptic connected component with spin 0 of a stratum, except $\mathcal{H}(2g - 2)$, of Abelian differentials on Riemann surfaces of genus g .*

5.4. Reaching the nonhyperelliptic component of $\mathcal{H}(2g-2)$, with spin 0.

5.4.1. *Degree $d = 2$.* We start with the model presented in Figure 2.8, and insert $g - 3$ vertical cylinders ($g > 3$) with parameters $y_1 = 2$ and $y_i = 1$ for $i > 1$ (see also Figure 2.9). The number of squares in grey color equals exactly $y^2 - 2 - 3 - (g - 3) = y^2 - g - 2$. The surface belongs to $\mathcal{H}(2, 2g - 4)$.

Since it can be continuously deformed to the surface in Figure 2.8 with spin 0, it also has spin 0. Now we collapse all the grey squares. The resulting surface belongs to the stratum $\mathcal{H}^{\text{non hyp}}(2g-2)$. Again this continuous deformation does not change the parity of the spin structure.




FIGURE 2.9. A surface in $\mathcal{H}^{\text{non hyp}}(2, 2g-4)$ for $g > 3$. If we collapse the handle (in grey color) we obtain a surface in $\mathcal{H}^{\text{non hyp}}(2g-2)$.

Following the computation in the previous subsection, we now obtain the $(g+1) \times (g+1)$ intersection matrix (recall $y^2 - g - 2 = 0$)

$$XX^\top = \left(\begin{array}{c|cccc} g+2 & 2 & 2 & 1 & \cdots & 1 \\ \hline 2 & 1 & 1 & & & \\ 2 & 1 & 1 & & & \\ 1 & & & 1 & & \\ \vdots & & & & \ddots & \\ 1 & & & & & 1 \end{array} \right).$$

By Lemma 2.11, the characteristic polynomial of XX^\top equals

$$t^a \left(-(g+2) \prod_{i=1}^{g-1} (t-y_i) + t \prod_{i=1}^{g-1} (t-y_i) - \sum_{i=1}^{g-1} c_i \prod_{j \neq i} (t-y_j) \right),$$

where $a = 1$, $c_1 = 4y_1$ and $c_i = y_i$ for $i \geq 2$. Thus the polynomial is

$$\begin{aligned} t(t-1)^{g-3} & (-(g+2)(t-2)(t-1) + t(t-2)(t-1) - 8(t-1) - (g-2)(t-2)) = \\ & = t^2(t-1)^{g-3}(t^2 - t \cdot (g+5) + 2g+2). \end{aligned}$$

In particular, the degree of the trace field is either one or two. The discriminant of $t^2 - t \cdot (g+5) + 2g+2$ is $D = (g+5)^2 - 8 \cdot (g+1) = g^2 + 2g + 17$. We see that $(g+1)^2 < D < (g+5)^2$. If the degree of the trace field is one then D is a square, and one of the following three cases holds:

- (1) $D = (g + 2)^2$. Then $g^2 + 2g + 17 = g^2 + 4g + 4$. Solving in g we find $2g = 13$ which is a contradiction;
- (2) $D = (g + 3)^2$. Then $g^2 + 2g + 17 = g^2 + 6g + 9$. Solving in g we find $6g = 1$ which is a contradiction;
- (3) $D = (g + 4)^2$. Then $g^2 + 2g + 17 = g^2 + 4g + 4$. Solving in g we find $g = 2$ which is again a contradiction with $g > 3$.

This implies that D is not a square and hence the degree of the trace field must be two.

5.4.2. *Degree $2 < d \leq g$.* We consider the modified version of our construction as depicted in Figure 2.10. When $g > 3$ the surface is not hyperelliptic.




FIGURE 2.10. A surface in the even spin non hyperelliptic connected component of $\mathcal{H}(2g - 2)$ for $g > 3$.

For the computation of the spin structure, we consider the “obvious” core curves γ_i and η_i (for $i = 0, \dots, g - 1$) of the horizontal and vertical cylinders. It forms a (non symplectic) basis of the homology:

$$\begin{aligned} i(\gamma_0, \eta_i) &= 1 \text{ for } i \neq g - 1 \text{ and } i(\gamma_0, \eta_{g-1}) = 0, \\ i(\gamma_i, \eta_i) &= \delta_{ij} \text{ for } i = 0, \dots, g - 1 \\ i(\gamma_i, \gamma_j) &= i(\eta_i, \eta_j) = 0. \end{aligned}$$

We can thus form a symplectic basis of $H_1(S; \mathbb{Z}/2\mathbb{Z})$ as follows:

$$\begin{cases} a_1 = \gamma_0, & b_1 = \eta_0 \\ a_i = \gamma_{i-1}, & b_i = \eta_{i-1} - \eta_0 \text{ for } i = 2, \dots, g - 1 \\ a_g = \gamma_{g-1}, & b_g = \eta_{g-1} - b_{g-1} = \eta_{g-1} - \eta_{g-2} + \eta_0 \end{cases}$$

Equation (1) reads

$$\Phi(\omega) = q(\gamma_0)q(\eta_0) + \sum_{i=2}^{g-1} q(\gamma_{i-1})q(\eta_{i-1} - \eta_0) + q(\gamma_{g-1})q(\eta_{g-1} - \eta_{g-2} + \eta_0)$$

Since $q(\eta_{i-1} - \eta_0) = q(\eta_{i-1}) + q(\eta_0) + i(\eta_{i-1}, \eta_0) = 1 + 1 + 0 = 0 \pmod{2}$, the sum with the $g - 2$ terms vanishes. For the last term, a direct computation leads to

$$\begin{aligned} q(\eta_{g-1} - \eta_{g-2} + \eta_0) &= q(\eta_{g-1}) + q(\eta_{g-2}) + q(\eta_0) + \\ &\quad + i(\eta_{g-1}, \eta_{g-2}) + i(\eta_{g-1}, \eta_0) + i(\eta_{g-2} + \eta_0) = \\ &= 1 + 1 + 1 + 0 + 0 + 0 = 1 \pmod{2}. \end{aligned}$$

Finally we get $\Phi(\omega) = 1 + 0 + 1 = 0 \pmod{2}$.

The intersection matrix (with the parameters $y_{g-2} = y_{g-1} = 1$) is

$$XX^\top = \left(\begin{array}{c|ccc|cc} y^2 & \mathbf{1}_{1 \times y_1} & \cdots & \mathbf{1}_{1 \times y_{g-3}} & 1 & 0 \\ \mathbf{1}_{y_1 \times 1} & \mathbf{1}_{y_1} & & & 0 & 0 \\ \vdots & & \ddots & & \vdots & \vdots \\ \mathbf{1}_{y_{g-3} \times 1} & & & \mathbf{1}_{y_{g-3}} & 0 & 0 \\ 1 & 0 & \cdots & 0 & 2 & 1 \\ 0 & 0 & \cdots & 0 & 1 & 1 \end{array} \right).$$

By developing along the last column, its characteristic polynomial equals

$$\begin{aligned} (2) \quad (t-1) &\left((t-2)p(t, y, \mathbf{y}) - t^a \prod_{i=1}^{g-3} (t-y_i) \right) - p(t, y, \mathbf{y}) = \\ &p(t, y, \mathbf{y})(t^2 - 3t + 1) - t^a(t-1) \prod_{i=1}^{g-3} (t-y_i) = \\ &= t^a \left(-y^2(t^2 - 3t + 1) \prod_{i=1}^{g-3} (t-y_i) + (t^3 - 3t^2 + 1) \prod_{i=1}^{g-3} (t-y_i) \right. \\ &\quad \left. - (t^2 - 3t + 1) \sum_{i=1}^{g-3} c_i \prod_{j \neq i} (t-y_j) \right), \end{aligned}$$

where $p(t, y, \mathbf{y})$ is the degree $g - 2$ polynomial in Lemma 2.11, with the parameters $a = \sum_{i=1}^{g-3} (y_i - 1)$ and $c_i = y_i$ for all $1 \leq i \leq g - 3$. Following the same line of proof we used for Lemma 2.12, we show

LEMMA 2.19. *The polynomial*

$$-y^2(t^2-3t+1) \prod_{i=1}^{g-3} (t-y_i) + (t^3-3t^2+1) \prod_{i=1}^{g-3} (t-y_i) - (t^2-3t+1) \sum_{i=1}^{g-3} y_i \prod_{j \neq i} (t-y_j)$$

is irreducible in $\mathbb{Z}[t, y]$ given that all $y_i \in \mathbb{N}$ are distinct.

PROOF. We follow the proof of Lemma 2.12. We note that the polynomial has degree two in the variable y with no non trivial common factor between the coefficient of y^2 and the constant coefficient. By the Perron–Frobenius theorem, there is a simple irreducible factor of the constant coefficient. Thus Eisenstein’s criterion applies in $(\mathbb{Z}[t])[y]$. \square

Again, the proofs of Theorems 2.13 and 2.14 carry over and provide the results of Corollaries 2.17 and 2.18 also for the stratum $\mathcal{H}(2g-2)$.

5.5. Reaching the hyperelliptic components of $\mathcal{H}(2g-2)$ and of $\mathcal{H}(g-1, g-1)$. We start by constructing a square-tiled surface. Pick a long horizontal square-tiled cylinder made of $2n+1$ squares with identifications A_0, A_1, \dots, A_n and A'_1, \dots, A'_n as depicted in Figure 2.11. We then add a stair case template, made of k steps, using a total of $2k$ squares. Finally we insert a long vertical square-tiled cylinder with some large number y of squares and identifications C_1, \dots, C_y as in Figure 2.11. We treat y as a variable that we will need to specify later on. This creates a surface $X_{n,k,y}$. Similarly, one can construct a surface $Y_{n,k,y}$ by collapsing one square corresponding to the label A_0 .

LEMMA 2.20. *The genus of $X_{n,k,y}$ and $Y_{n,k,y}$ is $g = n + k + 2$. Moreover $X_{n,k,y}$ belongs to the hyperelliptic connected component of $\mathcal{H}(g-1, g-1)$ while $Y_{n,k,y}$ belongs to the hyperelliptic connected component of $\mathcal{H}(2g-2)$.*

PROOF OF LEMMA 2.20. Clearly the two square-tiled surfaces are hyperelliptic: the involution fixes the $k+2$ horizontal cylinders. By inspecting the gluing, one sees that $X_{n,k,y}$ has two zeros, each of order $g-1$. The cone angle at each zero is $g \cdot 2\pi$. Since the total number of squares contributing to the cone angle is $2n+2k+2+2$, we get $(2n+2k+4) \cdot 2\pi = g \cdot 2\pi + g \cdot 2\pi$. Hence, $g = n + k + 2$.

Similarly, $Y_{n,k,y}$ has one zero, of order $2g-2$ and cone angle $(2g-1) \cdot 2\pi$. Now the total number of squares contributing to the cone angle is one less, namely $2n+2k+2+1$. Thus, $(2n+2k+3) \cdot 2\pi = (2g-1) \cdot 2\pi$. \square

along the last column that $p_k(t, y) = -y^2(q_k(t) + q_{k-1}(t)) + tq_k(t)$. We now claim that the roots of q_k and q_{k-1} are pairwise distinct and simple.

PROOF OF THE CLAIM. We note that B_{k-1} is obtained from B_k by deleting the last row and the last column. Interlacing results for real symmetric matrices tell us that the eigenvalues of B_k and B_{k-1} interlace. This means that if $\lambda_1 \leq \dots \leq \lambda_{k+1}$ are the eigenvalues of B_k and if $\mu_1 \leq \dots \leq \mu_k$ are the eigenvalues B_{k-1} , then we have

$$\lambda_i \leq \mu_i \leq \lambda_{i+1}$$

for all $1 \leq i \leq k$. The crucial point is that in our case these inequalities are strict, which can be proved as follows. We first note that the matrix B_k is clearly symmetric and positive definite (for $\alpha \geq 1$). This implies that all its leading principal minors are positive. Now, Theorem 7 by Gantmacher and Krein [9] states that a tridiagonal matrix with positive coefficients on the main diagonal and the adjacent diagonals is oscillatory if and only if all the leading principal minors are positive, implying that B_k is oscillatory. In turn, Theorem 6.5 by Ando [2] states that for oscillatory matrices, all the interlacing inequalities are strict. That is, if $\lambda_1 \leq \dots \leq \lambda_{k+1}$ are the eigenvalues of B_k and if $\mu_1 \leq \dots \leq \mu_k$ are the eigenvalues B_{k-1} , then we have

$$\lambda_i < \mu_i < \lambda_{i+1}$$

for all $1 \leq i \leq k$. In particular, the eigenvalues of B_k and the eigenvalues of B_{k-1} are pairwise distinct and simple. \square

We now finish the proof the lemma. Let $F \neq t$ be an irreducible factor of q_k . Since the roots of q_k are simple, F^2 is not a factor of tq_k . If F is a factor of $q_k + q_{k-1}$ then q_k and q_{k-1} share a common root, which is not possible by the claim. Hence, by Eisenstein's criterion, $p_k(t, y)$ is irreducible when regarded as a polynomial in the variable y and so can not be factored in the form $(ay+b)(cy+d)$. So, if there is a factorisation of $p_k(t, y)$, then one of the factors must have degree zero in the variable y . But such a factorisation cannot exist, since $q_k(t) + q_{k-1}(t)$ and $tq_k(t)$ are relatively prime in $\mathbb{Z}[t]$. Indeed, since the roots of $q_k(t)$ and $q_{k-1}(t)$ are distinct, the only possible common factor of tq_k and $q_k + q_{k-1}$ is t . But $p_k(0, y)$ is the determinant of XX^\top and equals $y^2 \cdot (\alpha - 1) \neq 0$. This proves the lemma. \square

THEOREM 2.22. *For any hyperelliptic connected component \mathcal{C} of the stratum $\mathcal{H}(2g - 2)$, every number $1 \leq d \leq g - 1$ is realised as the degree of the trace field of a product of two affine multitwists on a surface in \mathcal{C} .*

PROOF OF THEOREM 2.22. Since the case $d = 1$ is clear by considering square-tiled surfaces, let us assume $d \geq 2$ and set $k = d - 2 \geq 0$. We construct a surface $X_{n,k,y}$ or $Y_{n,k,y}$ where $n = g - d = g - k - 2 \geq 0$, see Lemma 2.20. If $d < g$ then $n \neq 0$ and $\alpha \neq 1$. If $d = g$, that is, if $n = 0$, then by assumption we consider only $Y_{n,k,y} \in \mathcal{H}(g - 1, g - 1)$ so that $\alpha = 4n + 2 = 2 \neq 1$. Thus Lemma 2.21 applies and the characteristic polynomial of XX^\top , viewed as a polynomial in $\mathbb{Z}[y, t]$ is irreducible. Then by Hilbert's irreducibility theorem [28], there are infinitely many specifications of y so that the resulting polynomial is irreducible as a polynomial in the variable t . Note that all specifications can be realised geometrically. Indeed, one can choose $y > 0$ by symmetry. In particular, applying the Thurston–Veech construction, there exists a product of two multitwists on the surface of genus $g = n + d$ in the desired connected component. \square

We also prove the analogous theorem for degrees of stretch factors.

THEOREM 2.23. *For any hyperelliptic connected component \mathcal{C} of the stratum $\mathcal{H}(g - 1, g - 1)$, every even number $2 \leq 2d \leq 2g$ is realised as the degree of the stretch factor of a product of two affine multitwists on a surface in \mathcal{C} .*

For any hyperelliptic connected component \mathcal{C} of $\mathcal{H}(2g - 2)$, every even number $2 \leq 2d \leq 2g - 2$ is realised as the degree of the stretch factor of a product of two affine multitwists on a surface in \mathcal{C} .

PROOF. We use the same examples as in the proof of Theorem 2.22. We first deal with the case $d = 2$ by taking the specific example $y = 1$. In this case, we have $k = 0$ and we get $XX^\top = \begin{pmatrix} \alpha & 1 \\ 1 & 1 \end{pmatrix}$. We obtain

$$\mu^2 = \frac{\alpha + 1 + \sqrt{\alpha^2 - 2\alpha + 5}}{2},$$

which is an algebraic number of degree two over \mathbb{Q} . Indeed, we have

$$\alpha^2 > \alpha^2 - 2\alpha + 5 > (\alpha - 1)^2$$

in case $\alpha \neq 1, 2$, so this number is not a square and μ^2 is not rational. Neither is it in case $\alpha = 2$, by direct calculation, and the case $\alpha = 1$ is not needed.

We are now ready to apply Theorem 2.8. Let Ω' be the matrix obtained from Ω by deleting the row and the column corresponding to the cylinder with $2n + 1$ or $2n + 2$ squares. We have

$$\sigma(\Omega + 2I) \geq \sigma(\Omega' + 2I) - 1.$$

By construction, Ω' is the adjacency matrix of a forest consisting of path graphs, so that $2I + \Omega'$ is positive definite. We get

$$\sigma(\Omega + 2I) \geq \dim(\Omega) - 2 > \dim(\Omega) - 4.$$

Theorem 2.8 applies and the mapping class $T_\alpha \circ T_\beta$ is pseudo-Anosov with stretch factor λ of degree $2d = 4$.

For the case $d \geq 3$, we take the examples as in the proof of Theorem 2.22, without specialising y . Let Ω' be the matrix obtained from Ω by deleting the rows and the columns corresponding to the horizontal cylinder with $2n + 1$ or $2n + 2$ squares, and to the vertical cylinder with $y + 1$ squares. We have

$$\sigma(\Omega + 2I) \geq \sigma(\Omega' + 2I) - 2.$$

By construction, Ω' is the adjacency matrix of a forest consisting of path graphs, so that $2I + \Omega'$ is positive definite. We get

$$\sigma(\Omega + 2I) \geq \dim(\Omega) - 4 > \dim(\Omega) - 2d.$$

Again, Theorem 2.8 applies and the mapping class $T_\alpha \circ T_\beta$ is pseudo-Anosov with stretch factor λ of degree $2d$. \square

5.6. Reaching the hyperelliptic component of $\mathcal{H}(2g - 2)$ with degree g . Take the stair case model with a “long” stair made of y^2 squares, see Figure 2.12. The $g \times g$ matrix is

$$XX^\top = \left(\begin{array}{c|cccc} y^2 & 1 & & & \\ \hline 1 & 2 & 1 & & \\ & 1 & \ddots & \ddots & \\ & & \ddots & \ddots & 1 \\ & & & 1 & 2 & 1 \\ & & & & 1 & 1 \end{array} \right).$$

Let $p_g(t, y)$ be the characteristic polynomial of XX^\top . We will use the characteristic polynomial $q_g(t)$ of the $g \times g$ matrix B_g obtained from $2I_g + \text{Adj}(A_g)$ by adding -1 to the last diagonal entry, where $\text{Adj}(A_g)$ is the adjacency matrix of the path graph with g vertices.




FIGURE 2.12. A stair case template in the hyperelliptic component of $\mathcal{H}(2g - 2)$.

By developing the determinant of $t\text{Id}_g - XX^\top$ along the first column we get

$$p_g(t, y) = -y^2 q_{g-1}(t) + t q_{g-1}(t) - q_{g-2}(t).$$

We now claim that the polynomials $q_{g-1}(t)$ and $t q_{g-1}(t) - q_{g-2}(t)$ are relatively prime. Using the same argument as in Lemma 2.21, we get that the matrix B_g is oscillatory, and hence the roots of p_{g-1} and p_{g-2} are all simple and pairwise distinct. Since the minimal polynomial of the Perron–Frobenius eigenvalue of B_g is a simple irreducible factor F of the polynomial $t q_{g-1}(t) - q_{g-2}(t)$ that is not also a factor of $q_{g-1}(t)$, Eisenstein’s criterion applies and $p_g(t, y)$ is irreducible. Thus there are infinitely many specifications of $y > 0$ such that $p_g(t, y) \in \mathbb{Z}[t]$ is irreducible. This yields the degree $d = g$ for the hyperelliptic component of $\mathcal{H}(2g - 2)$ for any $g > 1$.

Using this model, it is now straightforward to adapt the proofs of Theorem 2.22 and Theorem 2.23 to construct examples where the trace field is of degree g and the stretch factor is of degree $2g$.

This finishes the case distinction of connected components of strata, and therefore the proofs of Theorems 1.6 and 1.7.

6. Proof of Theorem 1.3

The goal of this section is to realise every positive integer $d \leq 3g - 3$ as the multicurve intersection degree of a pair of multicurves $\alpha, \beta \subset S$ on S_g for $g \geq 3$ in such a way that the multicurves α and β consist of components that are separating or that come in bounding pairs, where for each bounding pair one of the curves is a component of α and the other

is a component of β . This guarantees that $T_\alpha \circ T_\beta^{-1}$ is an element of the Torelli group $\mathcal{I}(S_g)$. Indeed, on the homology level, the action of a Dehn twist along a curve γ is given by $[\delta] \mapsto [\delta] + i(\delta, \gamma)[\gamma]$, where $i(\cdot, \cdot)$ is the algebraic intersection form on $H_1(S_g; \mathbb{Z})$. It follows that a Dehn twist along a separating curve is an element of $\mathcal{I}(S_g)$, but also each composition of Dehn twists $T_{\alpha_i} \circ T_{\beta_j}^{-1}$ for a bounding pair α_i and β_j . Since all the individual Dehn twists T_{α_i} commute, as well as all the individual Dehn twists T_{β_j} , it follows that if the multicurves α and β consist of components that are separating or that come in bounding pairs, where for each bounding pair one of the curves is a component of α and the other is a component of β , the mapping class $T_\alpha \circ T_\beta^{-1}$ is an element of $\mathcal{I}(S_g)$, and hence so is the mapping class $T_\alpha^n \circ T_\beta^{-n}$. Theorem 1.4 then provides Theorem 1.3 in the case $g \geq 3$. For the case $g = 2$ we note that the statement is proved for $d = 3g - 3 = 3$ in Section 3, and for $d \leq 2 = g$ it is proved in Section 4 above.

We start with the maximal degree $3g - 3$ and then discuss how to adapt the construction in order to realise smaller degrees.

6.1. Multicurve intersection degree $3g - 3$. We start by realising, on the surface of genus $g \geq 2$ with one boundary component, a pair of filling multicurves α and β such that $\chi_{X \times X^\top} \in \mathbb{Z}[t]$ is irreducible and of degree $3g - 2$, in such a way that the multicurves α and β consist of components that are separating or that come in bounding pairs, where for each bounding pair one of the curves is a component of α and the other is a component of β . The construction is done by induction on the genus $g \geq 2$.

6.1.1. *For $g = 2$ with one boundary component.* We consider the two multicurves α and β shown in Figure 2.13, with y parallel copies of the component β_1 . We first note that the components α_1 and α_3 are separating. Furthermore, the components α_2 and α_4 have their counterparts in the multicurve β with which they each form a bounding pair. Finally, the component of β of which there are y parallel copies and the component of β drawn in light blue in Figure 2.13 are separating.

We first note that the components $\alpha_1, \beta_1, \alpha_3$ and β_3 are separating. Furthermore, the components α_2 and α_4 have their counterparts β_2 and β_4 , respectively, in the multicurve β with which they each form a bounding pair.




FIGURE 2.13. Two multicurves α and β on the surface of genus two with one boundary component. There are y parallel copies of the component β_1 , and the separating component β_3 is drawn in light blue.

We directly calculate

$$XX^\top = \begin{pmatrix} 84 + 16y & 40 + 8y & 40 & 16 \\ 40 + 8y & 20 + 4y & 20 & 8 \\ 40 & 20 & 20 & 8 \\ 16 & 8 & 8 & 4 \end{pmatrix},$$

and it is a direct check (by the computer) that the characteristic polynomial of XX^\top is irreducible if $y = 2$ or $y = 3$. This finishes the case $g = 2$ with one boundary component.

6.1.2. *For $g > 2$ and one boundary component.* In order to increase the genus by one, we glue a surface of genus one with two boundary components as follows. On this surface, we consider the two multicurves α and β shown in Figure 2.14. We directly calculate

$$XX^\top = \begin{pmatrix} 16y + 4 & 8y \\ 8y & 4y \end{pmatrix} =: C_y,$$

and $\chi_{XX^\top}(t) = t^2 - (20y + 4)t + 16y$ with discriminant $16 \cdot (25y^2 + 6y + 1)$, which is never a square. Indeed, we have

$$(5y)^2 = 25y^2 < 25y^2 + 6y + 1 < 25y^2 + 10y + 1 = (5y + 1)^2.$$




FIGURE 2.14. Two multicurves α (in red) and β (in blue) on the surface of genus one with two boundary components. The multicurve β has y parallel copies of its separating component.

In particular, the polynomial $\chi_{XX\tau}$ is irreducible for positive integers y .

For the inductive step, let $g \geq 2$. Assume we have constructed on the surface of genus g with one boundary component a pair of multicurves α', β' such that the characteristic polynomial $\chi_{XX\tau} \in \mathbb{Z}[t]$ is irreducible and of degree $3g - 2$, in such a way that the multicurves α and β consist of components that are separating or that come in bounding pairs, where for each bounding pair one of the curves is a component of α and the other is a component of β . Further, assume that α'_1 is a simple closed curve that encircles all the handles of the surface, except for the rightmost one. Then, we take such a model surface and glue to its boundary a surface of genus one with two boundary components, as shown in Figure 2.14, and add two new curves α_0 and β_0 to the multicurves. The curve α_0 encircles all the handles of the newly formed surface, except for the rightmost one, and the curve β_0 runs along the glued boundary components, and twice intersects α_0 but no other component of α , see Figure 2.15.




FIGURE 2.15.

The proof of irreducibility is now exactly the same as in the non-Torelli case in Section 3. The only thing we need to check is that the multicurves α

and β consist of components that are separating or that come in bounding pairs, where for each bounding pair one of the curves is a component of α and the other is a component of β . But this is clearly the case, since all the curves we add in the inductive step are separating or come as a bounding pair.

6.1.3. *The closed case for $g \geq 4$.* The last step is to make the surfaces closed. We simply glue together two pieces of genus g', g'' , for $g' + g'' = g$, and one boundary component together along their boundaries. The same argument as in the inductive step provides irreducible characteristic polynomials of degree

$$3g' - 2 + 3g'' - 2 + 1 = 3g - 3.$$

6.1.4. *The closed case for $g = 3$.* We need a different argument. In this case, we start with the surface of genus two and one boundary component depicted in Figure 2.13, and close it off to the left by gluing a surface of genus one with one boundary component, see Figure 2.16. We add the




FIGURE 2.16. Two multicurves α and β on the surface of genus three. There are two new components of α when compared to Figure 2.13: a nonseparating component (red) that we call α_5 and a separating component (orange) that we call α_6 . Similarly, there are two new components of β : a separating component (blue) that we call β_5 and a nonseparating component (light blue) that we call β_6 .

curves α_5, α_6 , two parallel copies of β_5 and finally the curve β_6 . We note that β_5 and α_6 are separating, and β_6 and α_5 form a bounding pair. To read off the matrix XX^T for this pair of multicurves, we note two things. Firstly, the component α_5 intersects the curves β_j the same number of times as α_1 , except for the component β_5 of which there are two parallel copies.

Secondly, the component α_6 intersects each curve β_j twice the number of times as α_5 , except for β_6 . Using this, it is straightforward to calculate (ordering the curves as $\alpha_6, \alpha_5, \alpha_1, \alpha_2, \alpha_3, \alpha_4$)

$$XX^\top = \begin{pmatrix} 500 & 248 & 232 & 112 & 80 & 32 \\ 248 & 124 & 116 & 56 & 40 & 16 \\ 232 & 116 & 116 & 56 & 40 & 16 \\ 112 & 56 & 56 & 28 & 20 & 8 \\ 80 & 40 & 40 & 20 & 20 & 8 \\ 32 & 16 & 16 & 8 & 8 & 4 \end{pmatrix},$$

which is then checked to have irreducible characteristic polynomial.

6.2. Multicurve intersection degrees $d < 3g - 3$. We now show how to modify our construction from Section 6.1 in order to realise multicurve intersection degrees smaller than the maximal multicurve intersection degree $3g - 3$. We need new building blocks to construct our surfaces.

Block 1. Our first block is obtained from the surface depicted in Figure 2.13, simply by dropping the component α_3 . A direct verification yields that for $y = 1, 2$ the characteristic polynomial of XX^\top is irreducible and of degree 3.

Block 2. Our second block is obtained from the surface depicted in Figure 2.17. The characteristic polynomial of the matrix XX^\top for the multicurves α and β is irreducible and of degree 1. Versions of this block with distinct characteristic polynomial can be obtained by taking y parallel copies of β .




FIGURE 2.17. Two separating and filling curves α (red) and β (blue) on the surface of genus two with one boundary component.

Block 3. Take a surface as depicted in Figure 2.18. We denote the red multicurve by α and the blue multicurve by β . The multicurve α has a total of $k + 1$ separating components: one for each of the handles that separates the handle, and one that separates all the handles. We denote the component of α that separates all the handles of the surface in Figure 2.18 by α_1 , and we denote the other separating components of α by $\alpha_2, \alpha_4, \dots, \alpha_{2k}$ from left to right. Finally, the remaining nonseparating components of α are $\alpha_3, \alpha_5, \dots, \alpha_{2k+1}$ from left to right. The multicurve β consists of a




FIGURE 2.18. A surface of genus k with two boundary components, as well as two multicurves α (in red) and β (in blue). The separating components of β can have several parallel copies: the separating components β_1, \dots, β_k running through the handles have y_1, \dots, y_k copies, and the separating component β_0 to the right of the left-most handle has $y^2 - k - 4(y_1 + \dots + y_k)$ copies.

separating and a nonseparating component in each handle, as well as one separating curve β_0 to the right of the left-most handle. All the separating components of β can have multiple parallel copies, see Figure 2.18. In this situation, we have

$$XX^\top = \left(\begin{array}{c|ccc} 4y^2 & v_{y_1}^\top & v_{y_2}^\top & \cdots & v_{y_k}^\top \\ v_{y_1} & C_{y_1} & 0 & & \\ v_{y_2} & 0 & C_{y_2} & & \\ \vdots & & & \ddots & \\ v_{y_k} & & & & C_{y_k} \end{array} \right), \quad C_{y_i} = \begin{pmatrix} 16y_i + 4 & 8y_i \\ 8y_i & 4y_i \end{pmatrix},$$

$$v_{y_i} = \begin{pmatrix} 16y_i + 4 \\ 8y_i \end{pmatrix}.$$

By Remark 2.6, $\chi_{XX^\top}(t, y) \in \mathbb{Z}[t][y]$ is irreducible. By Hilbert's irreducibility theorem [28], there are infinitely many specifications of y such

that $y^2 - k - 4(y_1 + \cdots + y_k) > 0$ and such that $\chi_{XX^\top}(t) \in \mathbb{Z}[t]$ is irreducible and of degree $2k + 1$. Note that we can drop the separating components $\alpha_2, \alpha_4, \dots, \alpha_{2k}$ of α winding around one handle one by one in order to decrease the degree, reducing a 2-by-2 block to a 1-by-1 block, consisting of the coefficient $4y_i$, for each component dropped in this way. If all the y_i are chosen pairwise distinct, Remark 2.6 guarantees that the polynomial $\chi_{XX^\top}(t, y) \in \mathbb{Z}[t][y]$ is irreducible. We can in this way construct all degrees $k + 1 \leq d \leq 2k + 1$ for the surface of genus k and 2 boundary components.

6.2.1. *Realising multicurve intersection degrees $3g - 6 \leq d < 3g - 3$.*

Using a block of type 1 or 2 instead of our standard starting surface depicted in Figure 2.13, we can reduce the multicurve intersection degree by 1 or 3, respectively. Since we use such block on both sides of the surface in our construction, this gives the possibility to reduce the degree by any among the numbers 1, 2, 3, 4 or 6. In particular, we can clearly realise the multicurve intersection degrees $3g - 4, 3g - 5$ and $3g - 6$. This argument works for every $g \geq 4$.

In case of $g = 3$, we need a separate argument. The idea is to copy our example of maximal degree from Figure 2.16, but leave out first α_1 and then also α_3 . If we drop α_1 , which amounts to deleting the third row and column, we obtain a new matrix XX^\top which is directly verified to have an irreducible characteristic polynomial. If we now drop also α_3 , which amounts to deleting the second-last row and column, we obtain a new matrix XX^\top which again is verified to have irreducible characteristic polynomial. Since we have only dropped separating components, we have not changed the fact that the multicurves α and β consist of components that are separating or that come in bounding pairs, where for each bounding pair one of the curves is a component of α and the other is a component of β .

We have therefore realised $d = 5, 4$ for $g = 3$. For $g = 3$, the case $d = 3$ equals the case $d = 3g - 6$, which is treated below.

6.2.2. *Realising multicurve intersection degrees $g \leq d \leq 3g - 6$.* We start by constructing a surface of genus $g - 2$ with two boundary components, which we then close off in a second step.

Using surfaces of the type depicted in Figure 2.14 and applying the inductive step procedure, we can construct a surface of genus $g - 2 \geq 1$ and two boundary components, as well as filling multicurves α and β with

intersection degree $3(g-2)-1 = 3g-7$. Using at some point in the inductive procedure a block of type 3 of genus $k \leq g-2$, as depicted in Figure 2.18, we can reduce the degree by up to $2k-2 \leq 2g-6$, realising multicurve intersection degrees from $g-1$ to $3g-7$ on the surface of genus $g-2$ with two boundary components. Now we close the surface, as depicted in Figure 2.19, adding the new components α_0 and β_0 to the multicurves α and β , respectively.




FIGURE 2.19. Two separating curves α_0 and β_0 . There are ρ parallel copies of β_0 .

We obtain the matrix

$$XX^T = \left(\begin{array}{c|ccc} 64\rho + 16a_1 & 4a_1 & \dots & 4a_n \\ \hline 4a_1 & & & \\ \vdots & & A & \\ 4a_n & & & \end{array} \right),$$

where A is the matrix XX^T before adding the curves α_0 and β_0 . Since we have $a_1 = 4a$, we can set $\rho = y^2 - a$ to have the top-left coefficient $64y^2$, which is exactly the form of the matrix in Lemma 2.4. Finishing the argument as usual, we can realise the multicurve intersection degrees $g \leq d \leq 3g-6$ for $g \geq 3$.

6.2.3. Realising multicurve intersection degrees $1 \leq d < g$. Realising multicurve intersection degree one is clearly achieved by taking a pair of separating filling curves on the surface S_g .

For $2 \leq d < g$, let us define $f = g-1-d$. We start with a surface block of type 3 of genus $g-2$, where we deleted all the components of α that are separating. We also remove the component of β in the middle of Figure 2.18. Furthermore, we let the $f+1 \leq g-2$ first of the parameters y_i be equal to 1. Then we close off the surface as in the previous case, adding one

component α_0 to α and one component β_0 to β , compare with Figure 2.19. Assume there are ρ parallel copies of β_0 . We get

$$XX^\top = \begin{pmatrix} 64(\rho - g + 2) + 256\delta & 32y_1 & 32y_2 & \cdots & 32y_{g-2} \\ & 32y_1 & 4y_1 & & \\ & 32y_2 & & 4y_2 & \\ & \vdots & & & \ddots \\ & 32y_{g-2} & & & 4y_{g-2} \end{pmatrix},$$

where $\delta = y_1 + \cdots + y_{g-2}$. We choose ρ such that $64(\rho - g + 2) + 256\delta = 64y^2$. To simplify the calculations, we let $z_i = 4y_i$ for $i = 1, \dots, g - 2$. The matrix becomes

$$XX^\top = \begin{pmatrix} 64y^2 & 8z_1 & 8z_2 & \cdots & 8z_{g-2} \\ & 8z_1 & z_1 & & \\ & 8z_2 & & z_2 & \\ & \vdots & & & \ddots \\ & 8z_{g-2} & & & z_{g-2} \end{pmatrix}.$$

By Lemma 2.11, the characteristic polynomial of XX^\top equals

$$p(t, y, \mathbf{z}) = -64y^2 \prod_{i=1}^{g-2} (t - z_i) + t \prod_{i=1}^{g-2} (t - z_i) - \sum_{i=1}^{g-2} 64z_i^2 \prod_{j \neq i} (t - z_j).$$

If all the z_i are pairwise distinct, this polynomial is irreducible as a polynomial in t, y by Lemma 2.12. However, we chose that the first $f + 1$ coefficients y_1, \dots, y_{f+1} are equal to 1 and the other $y_i \neq 1$ and pairwise distinct. In particular, the polynomial $p(t, y)$ factors as $(t - 4)^f \tilde{p}(t, y)$, where $\tilde{p}(t, y)$ is of degree $g - 1 - f = d$ in the variable t and with pairwise distinct z_i . In particular, Lemma 2.12 implies that $\tilde{p}(t, y) \in \mathbb{Z}[t, y]$ is irreducible. Hilbert's irreducibility theorem [28] guarantees the existence of infinitely many positive specifications of y such that the resulting polynomial is irreducible in $\mathbb{Z}[t]$.

This finishes the proof of Theorem 1.3.

We end this section with a proof of Theorem 1.5.

PROOF OF THEOREM 1.5. For every $g \geq 3$ and every positive integer $1 \leq d \leq 3g - 3$, we have constructed a pair of filling multicurves α and β , with a parameter y , such that $\chi_{XX^\top}(t, y) \in \mathbb{Z}[t, y]$ is irreducible. By Remark 2.2, we may run the same argument to show that also the polynomial $\chi_{XX^\top}(t^2, y) \in \mathbb{Z}[t, y]$ is irreducible. By Hilbert's irreducibility theorem [28], we find infinitely many specifications of y such that the

polynomial $\chi_{XX^\top}(t^2) \in \mathbb{Z}[t]$ is irreducible of degree $2d$. The leading eigenvalue μ of the matrix XX^\top is a root of a characteristic polynomial $\chi_{XX^\top}(t)$, so $\chi_{XX^\top}(t^2) \in \mathbb{Z}[t]$ is the minimal polynomial of $\sqrt{\mu}$. Hence, the multicurve bipartite degree of the pair of multicurves α and β equals $2d$.


For $g = 2$, we use the example constructed in Section 3 for $d = 3$ and the examples from Section 4 for $d = 1, 2$. Similarly to Remark 2.2, one can run the same proof as for Lemma 2.12 to show that $\chi_{XX^\top}(t^2, y) \in \mathbb{Z}[t, y]$ is irreducible. \square

7. The case $g = 2$ of Theorems 1.1 and 1.2

While Theorem 1.3 implies Theorems 1.1 and 1.2 in the case $g \geq 3$, we still have to give a proof for the genus two surface, for which we now present ad hoc examples. For this purpose, we use the flipper software [4] and start with the genus two surface with one puncture $S_{2,1}$, see the figure below. In flipper, mapping classes are defined via Dehn twists along the curves a, b, c, d, e, f .

Consider the separating curve γ depicted in blue. By the chain relation, we know that $T_\gamma = (T_a \circ T_f)^6$. We consider the following three conjugates of T_γ :

- (1) $T_1 = (T_f \circ T_a \circ T_b) \circ T_\gamma \circ (T_f \circ T_a \circ T_b)^{-1}$,
- (2) $T_2 = (T_c \circ T_b) \circ T_\gamma \circ (T_c \circ T_b)^{-1}$,
- (3) $T_3 = (T_a \circ T_b) \circ T_\gamma \circ (T_a \circ T_b)^{-1}$.



Obviously $T_\gamma, T_1, T_2, T_3 \in \mathcal{I}(S_{2,1})$. For each even degree $d \in \{2, 4, 6\}$, we exhibit a word in the above elements. We check that the mapping class is pseudo-Anosov with singularity pattern $(1, 1, 1, 1; 0)$, and we compute its stretch factor by using flipper [4]. The vector $(1, 1, 1, 1; 0)$ means that the invariant foliations have four 3-pronged type singularities, and a 2-pronged one at the puncture.

| pseudo-Anosov mapping class $[f] \in \text{Mod}(S_{2,1})$ | minimal polynomial of $\lambda(f)$ |
|---|--|
| $T_\gamma \circ T_1 \circ T_2^{-1}$ | $t^2 - 66t + 1$ |
| $T_\gamma \circ T_1 \circ T_2$ | $t^4 - 72t^3 + 110t^2 - 72t + 1$ |
| $T_\gamma \circ T_1 \circ T_2 \circ T_3$ | $t^6 - 266t^5 + 143t^4 - 204t^3 + 143t^2 - 266t + 1$ |

In order to obtain elements in $\text{Mod}(S_2)$, we appeal to the forgetful map. Since our examples do not have a 1-pronged singularity at the puncture, we can fill it in order to get a pseudo-Anosov mapping class with the same stretch factor in $\text{Mod}(S_2)$, see [21, Lemma 2.6] for details. This completes the proof of Theorem 1.2 for $g = 2$. Now for each example, the minimal polynomial of $\lambda(f)$ is reciprocal and has degree $2d$ (for $d = 1, 2, 3$ respectively). Hence the minimal polynomial of $\lambda(f) + \lambda(f)^{-1}$ has degree d as required, and Theorem 1.1 is proved as well.

8. Explicit pseudo-Anosov maps

In [42, Problem 10.4], Margalit asks for explicit examples of pseudo-Anosov maps with specific stretch factor degrees. Our construction of multicurves allows us to do so for small genera.

More precisely, we now provide an almost explicit construction of multicurves with intersection degrees $1 < d \leq 3g - 3$ for which the pseudo-Anosov mapping class $T_\alpha \circ T_\beta$ has stretch factor degree $2d$.

Aided by the computer, we can subsequently find completely explicit examples of multicurves and therefore entirely explicit pseudo-Anosov mapping classes on S_g arising from the Thurston–Veech construction realising the maximal stretch factor degree $6g - 6$, for all genera $g \leq 201$.

In a first step, we build upon our examples in Section 3 to realise also all trace field degrees $d < 3g - 3$ on S_g for $g \geq 2$.

8.1. Multicurve intersection degrees $d < 3g - 3$. Recall that in Section 3, we constructed pairs of multicurves for which the Thurston–Veech construction realises the multicurve intersection degree $3g - 3$, for $g \geq 2$. In a first step, we build upon our examples from Section 3 to realise also all possible smaller multicurve intersection degrees $1 \leq d < 3g - 3$ on S_g for $g \geq 2$.

We need a new building block for our surfaces, see Figure 2.20.

We denote the red multicurve by α and the blue multicurve by β . The multicurve α has $k + 1$ separating components: one for each of the handles that separates the handle, and one that separates all the handles. We denote the component of α that separates all the handles of the surface in Figure 2.20 by α_1 , and we denote the other separating components of α by $\alpha_2, \alpha_4, \dots, \alpha_{2k}$ from left to right. Finally, the remaining nonseparating components of α are $\alpha_3, \alpha_5, \dots, \alpha_{2k+1}$ from left to right.




FIGURE 2.20. A surface of genus k with two boundary components, as well as two multicurves α (in red) and β (in blue). Some components of β have several parallel copies, as indicated by y_1, \dots, y_k and $y^2 - k$.

In this situation, we have

$$XX^\top = \left(\begin{array}{c|cccc} 4y^2 & v^\top & v^\top & \cdots & v^\top \\ \hline v & B_{y_1} & 0 & & \\ v & 0 & B_{y_2} & & \\ \vdots & & & \ddots & \\ v & & & & B_{y_k} \end{array} \right), \quad B_{y_i} = \begin{pmatrix} 4 & 2 \\ 2 & y_i \end{pmatrix}, \quad v = \begin{pmatrix} 4 \\ 2 \end{pmatrix}.$$

Let $p_{y_i}(t) = t^2 - (4 + y_i)t + 4(y_i - 1)$ be the characteristic polynomial of B_{y_i} . We know from Section 3 that p_{y_i} is irreducible if $y \geq 12$. So, choosing all $y_i \geq 12$ pairwise distinct, Remark 2.6 guarantees that the polynomial $\chi_{XX^\top}(t, y) \in \mathbb{Z}[t][y]$ is irreducible. By Hilbert's irreducibility theorem [28], there are infinitely many specifications of y such that $y^2 - k > 0$ and such that $\chi_{XX^\top}(t) \in \mathbb{Z}[t]$ is irreducible and of degree $2k + 1$.

Case 1: $2g \leq d < 3g - 3$. Assume we want to realise the multicurve intersection degree $3g - 3 - f$ for $0 < f \leq g - 3$. Let $k = f + 2 \leq g - 1$. Start the inductive procedure as in Section 3 with the surface from Figure 2.20 as a starting point, adding $g - 1 - k$ more handles. The exact same argument yields a surface of genus $g - 1$ with two boundary components, and a characteristic polynomial $\chi_{XX^\top} \in \mathbb{Z}[t]$ that is irreducible and of degree $2k + 1 + 3(g - 1 - k) = 3g - 3 - k + 1$. Closing up the surface exactly as in Section 3 yields $3g - 3 - k + 2 = 3g - 3 - f$ as a multicurve intersection degree on the closed orientable surface of genus g .

Case 2: $g < d < 2g$. Assume we want to realise the multicurve intersection degree $2g - f$ for $0 < f \leq g - 1$. Take the surface depicted in Figure 2.20 for $k = g - 1$. Now, remove f of the separating curve $\alpha_2, \dots, \alpha_{2g-2}$. This slightly modifies the matrix XX^\top : f of the 2-by-2 blocks on the diagonal are now 1-by-1 blocks, with the single coefficient y_i . Nevertheless, since all the y_i are chosen pairwise distinct, Remark 2.6 guarantees that the polynomial $\chi_{XX^\top}(t, y) \in \mathbb{Z}[t][y]$ is irreducible. We note that for the coefficients y_i in the 1-by-1 blocks, any positive integer can be chosen. By Hilbert's irreducibility theorem [28], there are infinitely many specifications of y such that $\chi_{XX^\top}(t) \in \mathbb{Z}[t]$ is irreducible and of degree $2g - 1 - f$. Closing up the surface as in Section 3 yields the multicurve intersection degree $2g - f$ on the closed orientable surface of genus g .

Case 3: $1 \leq d \leq g$. This is the case we have already dealt with in Section 4.

8.2. Even degree stretch factors. We now show that in our construction of multicurves of Section 8.1, the degree of the stretch factor of $T_\alpha \circ T_\beta$ equals two over the trace field. It uses our first nonsplitting criterion, Theorem 2.8.

Recall that $\sigma(A)$ and $\text{null}(A)$ denote the signature and the nullity, respectively, of the matrix A .

THEOREM 2.24. *Let α and β a pair of multicurves described in Section 8.1, realising a multicurve intersection degree $1 \leq d \leq 3g - 3$. Then the mapping class $T_\alpha \circ T_\beta$ is pseudo-Anosov with stretch factor λ of degree $2d$.*

For the case $1 \leq d \leq g$, this is shown in Section 4.

PROOF OF THEOREM 2.24. According to Theorem 2.8, all there is to show is

$$(3) \quad \dim(\Omega) > \sigma(\Omega + 2I) + \text{null}(\Omega + 2I) > \dim(\Omega) - 2d.$$

We now make a case distinction depending on d .

Case 1: $2g \leq d \leq 3g - 3$. We consider the submatrix Ω' of Ω that is obtained by deleting all the rows and columns corresponding to components of the multicurve α that have been added during the inductive step or closing up of the surface. Furthermore, if $d < 3g - 3$, we also remove the component of α encircling multiple handles of the starting surface, that is, the surface depicted in Figure 2.20.

A base change by a permutation matrix brings $\Omega' + 2I$ into block diagonal form with $g - 1$ blocks corresponding to genus one surface pieces as depicted in Figure 2.1, and a block of the form $2I$. For a block of the former type, and for $y > 4$, we directly calculate that the nullity is zero and the signature equals the dimension of the block minus two. Already, this implies that certainly the signature of $\Omega + 2I$ is not equal to its dimension, and it only remains to verify the lower bound in (3).

By construction, if the genus equals $g \geq 2$, we have $g - 1$ surface pieces as in Figure 2.1. This in particular implies that $\sigma(\Omega') = \dim(\Omega') - 2g + 2$. Furthermore, we have $\dim(\Omega) - \dim(\Omega') = d - 2g + 2$. The latter equality follows from that fact that the number of components of α in our construction is exactly d , and there are two components per surface pieces as in Figure 2.1. We now calculate

$$\begin{aligned} \sigma(\Omega + 2I) &\geq \sigma(\Omega') - (\dim(\Omega) - \dim(\Omega')) \\ &= (\dim(\Omega') - 2g + 2) - (d - 2g + 2) \\ &= \dim(\Omega') - d \\ &> \dim(\Omega) - 2d, \end{aligned}$$

which implies (3), so we are done for this case.

Case 2: $g < d < 2g$. We consider the submatrix Ω' of Ω that is obtained by deleting two rows and two columns corresponding to components of the multicurve α : the one corresponding to the component encircling multiple handles in Figure 2.20 and the one obtained from closing the surface. Recall that we have removed $f = 2g - d$ separating curves $\alpha_2, \dots, \alpha_{2g-2}$.

A base change by a permutation matrix brings $\Omega' + 2I$ into block diagonal form with $g - 1 - f$ blocks corresponding to surface pieces as in Figure 2.1, f blocks corresponding to surface pieces as in Figure 2.1 but with the separating component of α removed, and a block of the form $2I$.

For a block of the first type, and for $y > 4$, recall from the previous case that the nullity is zero and the signature equals the dimension of the block minus two. For a block of the second type, the sum of the nullity and the signature equals the dimension of the block if $y \leq 3$, and it equals the dimension of the block minus two if $y > 3$. We may assume that for at least one block of the second type, we have $y = 3$. This is enough to ensure that $\dim(\Omega) > \sigma(\Omega + 2I) + \text{null}(\Omega + 2I)$, so again we only need to verify the lower bound in (3).

By construction, for genus $g \geq 2$, we have $g - 1$ surface pieces as in Figure 2.1. Having at least one piece with $y \leq 3$ implies

$$\sigma(\Omega') > \dim(\Omega') - 2g + 2.$$

Furthermore, we have $\dim(\Omega) - \dim(\Omega') = 2$. We now calculate

$$\begin{aligned} \sigma(\Omega + 2I) &\geq \sigma(\Omega') - (\dim(\Omega) - \dim(\Omega')) \\ &> (\dim(\Omega') - 2g + 2) - 2 \\ &= \dim(\Omega') - 2g \\ &= \dim(\Omega) - 2g + 2 \geq \dim(\Omega) - 2d, \end{aligned}$$

which implies (3) also in the case $g < d < 2g$, so we are done. \square

8.3. Explicit examples with stretch factor degree $6g - 6$. We conclude this section by giving explicit computations supporting a conjecture on the irreducibility of the characteristic polynomials constructed in Section 3 for specific values of y . In the inductive step of Section 3, one uses a map

$$\phi_k : \begin{array}{ccc} M_k(\mathbb{Z}) \times \mathbb{Z} & \longrightarrow & M_{k+3}(\mathbb{Z}) \\ (C, y) & \mapsto & \left(\begin{array}{c|c|c} 4y^2 & * & * \\ * & C & \\ * & & A \end{array} \right), \quad \text{with } A = \begin{pmatrix} 4 & 2 \\ 2 & 12 \end{pmatrix}. \end{array}$$

For $g > 1$ we inductively construct the $(3g - 1) \times (3g - 1)$ matrix M_g with the maps ϕ_{3i-1} for $i = 1, \dots, g - 1$:

$$M_g = \phi_{3(g-1)-1}(\phi_{3(g-2)-1}(\dots \phi_{3 \cdot 2-1}(\phi_{3 \cdot 1-1}(B, y^{(1)}), y^{(2)}), \dots, y^{(g-2)}), y^{(g-1)}),$$

with $B = \begin{pmatrix} 4 & 2 \\ 2 & 13 \end{pmatrix}$ and suitable parameters $y^{(i)}$ given by Hilbert's irreducibility theorem. The condition $y^2 > \frac{1}{4}c_{11} + 1$ appearing in the construction is obviously equivalent to $(y^{(i)})^2 > (y^{(i-1)})^2 + 1$. Finally, following Section 3 the matrix XX^\top for the multicurves α and β on the closed surface of genus $g + 1$ takes the form

$$\widetilde{M}_g = \left(\begin{array}{c|c} y^2 & * \\ * & M_g \end{array} \right)$$

with the condition $y^2 > \frac{1}{4}(M_g)_{11} = (y^{(g)})^2$.

By computer assistance, one immediately checks the following proposition.

PROPOSITION 2.25. *For any $1 < g \leq 200$, if we set $y^{(i)} = i + 1$ for indices $i = 1, \dots, g - 1$, then the characteristic polynomial χ_{M_g} is irreducible over \mathbb{Q} . Moreover, for $y = g + 1$, $\chi_{\widetilde{M}_g}$ is irreducible over \mathbb{Q} .*

Together with Theorem 2.24, this gives explicit examples of pseudo-Anosov maps realising the upper bound $6g - 6$ in Theorem 1.3 for every genus $1 < g \leq 201$. We don't know whether χ_{M_g} and $\chi_{\widetilde{M}_g}$ are actually irreducible for every $g > 200$ with the parameters $y^{(i)} = i + 1$ chosen as in Proposition 2.25.

CHAPTER 3

Algebraic characterisation of stretch factors

In this chapter, we give the proof of Theorem 1.8, characterising up to powers the algebraic integers that arise as stretch factors in the Thurston–Veech construction. The material is copied and adapted from our joint work with Pankau [38].

1. Perron and bi-Perron numbers

In this section, we prove a result about the trace fields of Perron and bi-Perron numbers. Recall that a Perron number λ is a real algebraic integer > 1 all of whose Galois conjugates have modulus in the open interval $(0, \lambda)$, except for λ itself. The following statement is given in the proof of Lemma 8.2 of Strenner [62].

LEMMA 3.1. *Let λ be a Perron number of degree l , and let $\lambda_1, \dots, \lambda_l$ be its Galois conjugates. Then for all positive integers k , λ^k is also of degree l and $\lambda_1^k, \dots, \lambda_l^k$ are its Galois conjugates.*

The property of Perron numbers highlighted by this lemma is a key ingredient to proving the following proposition.

PROPOSITION 3.2. *Let λ be a Perron number. Then, we have one of the following.*

(1) *If $[\mathbb{Q}(\lambda) : \mathbb{Q}]$ is odd, then for all k we have*

$$\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda^k + \lambda^{-k}).$$

(2) *If $[\mathbb{Q}(\lambda) : \mathbb{Q}]$ is even, then for all k we have*

$$\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda^{2k+1} + \lambda^{-(2k+1)})$$

and

$$\mathbb{Q}(\lambda^2 + \lambda^{-2}) = \mathbb{Q}(\lambda^{2k} + \lambda^{-2k})$$

with $\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda^2 + \lambda^{-2})$ if and only if $-\lambda^{-1}$ is not a Galois conjugate of λ .

We note that it is possible for the stretch factor λ of an orientation-reversing pseudo-Anosov map to have $-\lambda^{-1}$ as a Galois conjugate. The following example describes an instance of this phenomenon in the simplest case of the torus.

EXAMPLE 3.3. The golden ratio $\varphi = \frac{1+\sqrt{5}}{2}$ is a bi-Perron number with minimal polynomial $t^2 - t - 1$. By definition, $-\varphi^{-1} = \frac{1-\sqrt{5}}{2}$ is a Galois conjugate of φ , and we have that $\mathbb{Q}(\varphi) = \mathbb{Q}(\varphi + \varphi^{-1}) = \mathbb{Q}(\sqrt{5})$ but we also have $\mathbb{Q}(\varphi^2 + \varphi^{-2}) = \mathbb{Q}$. While the golden ratio is not the stretch factor of an orientation-preserving Anosov map of the torus, it is the stretch factor of an orientation-reversing one: indeed, the spectral radius of the matrix $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ is the golden ratio.

We now prove Proposition 3.2, which will be important in the proof of Theorem 1.8 below.

PROOF OF PROPOSITION 3.2. We start by noting that since λ is a Perron number, then Lemma 3.1 tells us that $[\mathbb{Q}(\lambda) : \mathbb{Q}] = [\mathbb{Q}(\lambda^k) : \mathbb{Q}]$ for all positive integers k . This immediately implies that $\mathbb{Q}(\lambda) = \mathbb{Q}(\lambda^k)$ for all k , since the former is a field extension of the latter.

We now prove part (1) of the proposition by assuming that $[\mathbb{Q}(\lambda) : \mathbb{Q}]$ is odd. Since $\mathbb{Q}(\lambda^k)$ is a field extension of $\mathbb{Q}(\lambda^k + \lambda^{-k})$, and since λ^k is a root of the polynomial $t^2 - (\lambda^k + \lambda^{-k})t + 1$, then we must have that

$$(*) \quad [\mathbb{Q}(\lambda^k) : \mathbb{Q}(\lambda^k + \lambda^{-k})] = 1 \text{ or } 2.$$

Now this degree cannot equal 2 because $\mathbb{Q}(\lambda) = \mathbb{Q}(\lambda^k)$, hence $[\mathbb{Q}(\lambda^k) : \mathbb{Q}]$ is odd, so by the tower theorem for field extensions, none of the intermediate extensions can have even degree. Hence, $\mathbb{Q}(\lambda) = \mathbb{Q}(\lambda^k) = \mathbb{Q}(\lambda^k + \lambda^{-k})$ for all positive integers k . In particular, we have that $\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda^k + \lambda^{-k})$ for all k .

The proof of part (2) of the proposition will be broken into two steps. The first step will be to prove that the equality $\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda^2 + \lambda^{-2})$ holds if and only if $-\lambda^{-1}$ is not a Galois conjugate of λ . We do this by proving the contrapositive. The second step will be to prove the general equalities in the statement by considering the cases for when $-\lambda^{-1}$ is a Galois conjugate or not.

It is important to note that $-\lambda^{-1}$ can only be a Galois conjugate in the even degree case since if it is a conjugate, then for any other conjugate μ ,

then so is $-\mu^{-1}$. Hence the minimal polynomial has an even number of roots.

We start by assuming that $[\mathbb{Q}(\lambda) : \mathbb{Q}]$ is even. Now, because of the inclusions $\mathbb{Q}(\lambda^2 + \lambda^{-2}) \subseteq \mathbb{Q}(\lambda + \lambda^{-1}) \subseteq \mathbb{Q}(\lambda) = \mathbb{Q}(\lambda^2)$ then we have the following tower, with the possible degree of each extension listed:

$$\begin{array}{c} \mathbb{Q}(\lambda) = \mathbb{Q}(\lambda^2) \\ \left| \begin{array}{c} 1 \text{ or } 2 \\ \mathbb{Q}(\lambda + \lambda^{-1}) \end{array} \right. \\ \left| \begin{array}{c} 1 \text{ or } 2 \\ \mathbb{Q}(\lambda^2 + \lambda^{-2}) \end{array} \right. \end{array}$$

We start by assuming that $\mathbb{Q}(\lambda + \lambda^{-1}) \neq \mathbb{Q}(\lambda^2 + \lambda^{-2})$, and show that then $-\lambda^{-1}$ must be a Galois conjugate of λ . Since these fields are not equal, then we immediately have that the top extension must be degree 1 because we know from (*) that the degree of $\mathbb{Q}(\lambda^2)$ over $\mathbb{Q}(\lambda^2 + \lambda^{-2})$ is 1 or 2. Therefore, we have $\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda) = \mathbb{Q}(\lambda^2)$, and the tower collapses to

$$\begin{array}{c} \mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda) = \mathbb{Q}(\lambda^2) \\ \left| \begin{array}{c} 2 \\ \mathbb{Q}(\lambda^2 + \lambda^{-2}) \end{array} \right. \end{array}$$

Hence, $t^2 - (\lambda^2 + \lambda^{-2})t + 1$ is the minimal polynomial for λ^2 over the field $\mathbb{Q}(\lambda^2 + \lambda^{-2})$. Therefore, we see that the non-identity automorphism $\phi \in \text{Gal}(\mathbb{Q}(\lambda^2)/\mathbb{Q}(\lambda^2 + \lambda^{-2}))$ maps λ^2 to λ^{-2} . Hence, $[\phi(\lambda)]^2 = \lambda^{-2}$, and we get $\phi(\lambda) = \pm\lambda^{-1}$. Note, we are allowed to apply ϕ to λ in this case since $\mathbb{Q}(\lambda) = \mathbb{Q}(\lambda^2)$. Now, it cannot be the case that $\phi(\lambda) = \lambda^{-1}$ because this would imply that $\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda^2)$ is fixed by $\text{Gal}(\mathbb{Q}(\lambda^2)/\mathbb{Q}(\lambda^2 + \lambda^{-2}))$, which contradicts the definition of the Galois group. Therefore, we have $\phi(\lambda) = -\lambda^{-1}$, which implies that $-\lambda^{-1}$ is a Galois conjugate of λ .

Now, running the argument in reverse, if $-\lambda^{-1}$ is a Galois conjugate of λ then there exists a \mathbb{Q} -automorphism ϕ of $\mathbb{Q}(\lambda)$ such that $\phi(\lambda) = -\lambda^{-1}$. This immediately implies that $\mathbb{Q}(\lambda^2 + \lambda^{-2})$ is fixed by ϕ but $\mathbb{Q}(\lambda + \lambda^{-1})$ is not, therefore, we have $\mathbb{Q}(\lambda + \lambda^{-1}) \neq \mathbb{Q}(\lambda^2 + \lambda^{-2})$.

We now generalise the argument and prove the equalities in statement (2) of the proposition. Suppose that $-\lambda^{-1}$ is a Galois conjugate of λ . Then,

the automorphism ϕ that interchanges λ and $-\lambda^{-1}$ fixes $\mathbb{Q}(\lambda^{2k} + \lambda^{-2k})$ for all positive integers k . This implies that

$$[\mathbb{Q}(\lambda^{2k}) : \mathbb{Q}(\lambda^{2k} + \lambda^{-2k})] = 2$$

for all k . But $\mathbb{Q}(\lambda^{2k} + \lambda^{-2k})$ is a subfield of $\mathbb{Q}(\lambda^2 + \lambda^{-2})$ for all k so

$$\mathbb{Q}(\lambda^2 + \lambda^{-2}) = \mathbb{Q}(\lambda^{2k} + \lambda^{-2k})$$

for all k . On the other hand $\mathbb{Q}(\lambda^{2k+1} + \lambda^{-(2k+1)})$ is not fixed for any k , hence

$$\mathbb{Q}(\lambda^{2k+1} + \lambda^{-(2k+1)}) = \mathbb{Q}(\lambda^{2k+1}) = \mathbb{Q}(\lambda) = \mathbb{Q}(\lambda + \lambda^{-1})$$

for all k .

Now, in the case where $-\lambda^{-1}$ is not a conjugate of λ , above arguments immediately imply that $\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda^2 + \lambda^{-2})$. Therefore, both of these fields must either equal $\mathbb{Q}(\lambda) = \mathbb{Q}(\lambda^k)$ for all positive integers k , or for no k . Suppose that the following equality holds for all k :

$$\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda^2 + \lambda^{-2}) = \mathbb{Q}(\lambda^k).$$

Then, it must hold that $\mathbb{Q}(\lambda^{2k} + \lambda^{-2k}) = \mathbb{Q}(\lambda^{2k+1} + \lambda^{-(2k+1)}) = \mathbb{Q}(\lambda^k)$ for all k , because if equality fails for some k , then there would have to exist a $\mathbb{Q}(\lambda^n + \lambda^{-n})$ -automorphism ϕ (where $n = 2k$ or $n = 2k + 1$) that interchanges λ^n with λ^{-n} . Thus, $\phi(\lambda) = \lambda^{-1}$, since it cannot equal $-\lambda^{-1}$ by assumption. Hence, $\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda)$ is a fixed field, which is a contradiction. Therefore, for all k we have

$$\mathbb{Q}(\lambda^{2k} + \lambda^{-2k}) = \mathbb{Q}(\lambda^{2k+1} + \lambda^{-(2k+1)}) = \mathbb{Q}(\lambda^k).$$

Finally, if the fields $\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda^2 + \lambda^{-2})$ do not equal $\mathbb{Q}(\lambda) = \mathbb{Q}(\lambda^k)$ for any k , then since $\mathbb{Q}(\lambda^k + \lambda^{-k})$ is a subfield of $\mathbb{Q}(\lambda + \lambda^{-1})$ for all k , (*) implies that $[\mathbb{Q}(\lambda^k) : \mathbb{Q}(\lambda^k + \lambda^{-k})] = 2$ for all k , therefore

$$\mathbb{Q}(\lambda + \lambda^{-1}) = \mathbb{Q}(\lambda^2 + \lambda^{-2}) = \mathbb{Q}(\lambda^k + \lambda^{-k})$$

for all k . □

The following result is the key to the construction of a geometric situation that corresponds to the power of a bi-Perron number. It was used by Pankau to show that every Salem number has a power that is the stretch factor of a pseudo-Anosov homeomorphism arising from the Thurston–Veech construction [48]. We now show that the proof works almost identically for

bi-Perron numbers all of whose Galois conjugates are contained in $\mathbb{S}^1 \cup \mathbb{R}$. This extension is also presented in Pankau's PhD thesis [49].

PROPOSITION 3.4. *Let λ be a bi-Perron number all of whose Galois conjugates are contained in $\mathbb{S}^1 \cup \mathbb{R}$. Then there exists a positive integer k so that $\lambda^k + \lambda^{-k}$ equals the spectral radius of a positive symmetric integer matrix.*

PROOF. We very closely follow Pankau's proof in [48]. For the convenience of the reader, we summarise the key steps and mention where we have to pay attention because our setting is slightly more general than in the original argument.

Let λ be a bi-Perron number all of whose Galois conjugates lie in $\mathbb{S}^1 \cup \mathbb{R}$. Then $\lambda + \lambda^{-1}$ is a totally real Perron number. Indeed, $\mathbb{Q}(\lambda + \lambda^{-1})$ is a subfield of $\mathbb{Q}(\lambda)$, so every embedding of $\mathbb{Q}(\lambda + \lambda^{-1})$ into \mathbb{C} is the restriction of an embedding of $\mathbb{Q}(\lambda)$ into \mathbb{C} . In particular, each Galois conjugate of $\lambda + \lambda^{-1}$ is of the form $\lambda_i + \lambda_i^{-1}$, where λ_i is a Galois conjugate of λ . Since $\lambda_i \in \mathbb{S}^1 \cup \mathbb{R}$, it follows that $\lambda_i + \lambda_i^{-1} \in \mathbb{R}$.

Let $f(t)$ be the minimal polynomial of $\lambda + \lambda^{-1}$, and denote by n its degree. Without loss of generality, we assume that $-\lambda^{-1}$ is not a Galois conjugate of λ . Indeed, if $-\lambda^{-1}$ is a Galois conjugate of λ , we can simply run the argument for λ^2 .

Step 1. By a result of Estes [13], there exists a rational symmetric matrix Q of size $(n+e) \times (n+e)$ with characteristic polynomial $f(t)(t-1)^e$, where e equals 1 or 2.

Step 2. By conjugation with an element in $O(n+e, \mathbb{Q})$ and possibly a small perturbation, we may assume that the eigenvector of the matrix Q for the eigenvalue $\lambda + \lambda^{-1}$ is positive, compare with the discussion starting with Proposition 5.2 in [48].

Step 3. Define the matrix

$$\mathcal{M} = \begin{pmatrix} Q & -I \\ I & 0 \end{pmatrix}.$$

We now describe the characteristic polynomial of \mathcal{M} . In the proof of Proposition 5.3 in [48], it is shown that μ is an eigenvalue for \mathcal{M} with eigenvector $(\mathbf{v}, \mu^{-1}\mathbf{v})^\top$ exactly if $\mu + \mu^{-1}$ is an eigenvalue for Q with eigenvector \mathbf{v} . Hence, the characteristic polynomial of \mathcal{M} equals $t^n f(t+t^{-1})(t^2-t+1)^e$. We

note the following discrepancy with Proposition 5.3 in [48]: if the characteristic polynomial $g(t)$ of λ is not reciprocal, then the polynomial $t^n f(t + t^{-1})$ equals $g(t)g^*(t)$, where $g^*(t) = t^n g(t^{-1})$. On the other hand, if $g(t)$ is reciprocal, which is the case exactly if λ has a Galois conjugate on the unit circle (for example if λ is a Salem number), then $t^n f(t + t^{-1})$ equals $g(t)$. In any case, the characteristic polynomial of \mathcal{M} has integer coefficients and $\det(\mathcal{M}) = 1$.

Step 4. By Proposition 5.4 in [48], for any positive integer k , $\mathcal{M}^k + \mathcal{M}^{-k}$ is a block diagonal matrix with two blocks \mathcal{Q}_k . Here, \mathcal{Q}_k is a rational symmetric matrix with characteristic polynomial $f_k(t)(t-a)^e$, where $f_k(t)$ is the minimal polynomial of $\lambda^k + \lambda^{-k}$ and a is among the numbers $-2, -1, 1, 2$. The proof does not depend on whether the characteristic polynomial $g(t)$ of λ is reciprocal or not. Also, by the discussion right above Proposition 5.5 in [48], the eigenspaces of Q and \mathcal{Q}_k agree. In particular, the eigenvector \mathbf{v} for the eigenvalue $\lambda^k + \lambda^{-k}$ of \mathcal{Q}_k is positive.

Step 5. We now prove that the matrix \mathcal{Q}_k is positive for k large enough. We write $\mathbf{e}_i = c_i \mathbf{v} + \mathbf{w}_i$ for every basis vector \mathbf{e}_i , where \mathbf{w}_i is a fixed vector (independent of k) in the orthogonal complement of \mathbf{v} , and $c_i > 0$. Since \mathbf{w}_i lies in the orthogonal complement to \mathbf{v} , it is a linear combination of eigenvectors of \mathcal{Q}_k other than \mathbf{v} . In particular, the modulus of every coefficient of $\mathcal{Q}_k \mathbf{w}_i$ is bounded from above by $|\lambda_2^k + \lambda_2^{-k}| \cdot \|\mathbf{w}_i\|_\infty$, where $\lambda_2 + \lambda_2^{-1}$ is the second-largest root in modulus of $f(t)$. Now, since λ is a bi-Perron number and $-\lambda^{-1}$ is not among its Galois conjugates, the ratio between $\lambda^k + \lambda^{-k}$ and $\lambda_2^k + \lambda_2^{-k}$ becomes arbitrarily large when k tends to infinity. Therefore,

$$\mathcal{Q}_k \mathbf{e}_i = c_i (\lambda^k + \lambda^{-k}) \mathbf{v} + \mathcal{Q}_k \mathbf{w}_i$$

becomes positive for large k , since $c_i > 0$ and \mathbf{v} is a positive vector.

Step 6. For large enough k , the matrix \mathcal{M}^k has integer coefficients by Proposition 5.5 in [48]. Hence, also \mathcal{M}^{-k} has integer coefficients for large enough k , since $\det(\mathcal{M}) = 1$. In particular, also \mathcal{Q}_k has integer coefficients for large enough k . This finishes the proof that for k large enough, the number $\lambda^k + \lambda^{-k}$ equals the spectral radius of a positive symmetric integer matrix \mathcal{Q}_k . \square

2. Coxeter transformations and the Thurston–Veech construction

2.1. The Coxeter transformation. Coxeter groups are abstract generalisations of reflection groups. They admit a presentation encoded in a graph with weighted edges, the so-called Coxeter diagram, and they are linear by Tits’ representation. As we can single out the only input we need from the theory of Coxeter groups in Lemma 3.5 below, we do not give the definitions and instead refer to Bourbaki’s classic [5].

In case the underlying graph of a Coxeter diagram is bipartite, there is a well-defined conjugacy class of matrices obtained via Tits’ representation, the so-called bipartite Coxeter transformation, see, for example, McMullen [44]. By a result of A’Campo, the spectrum of this matrix is contained in $\mathbb{S}^1 \cup \mathbb{R}_{>0}$, and determines, for example, whether the group is finite [1]. All we need for our purposes is the following formula relating the spectra of the Coxeter adjacency matrix and the bipartite Coxeter transformation. We do not give the definition of the Coxeter adjacency matrix, but simply note that in our case of Coxeter diagrams with simple edges (which, in the language of Coxeter groups, means that every edge is of weight 3), the Coxeter adjacency matrix equals the ordinary adjacency matrix of the underlying abstract graph.

LEMMA 3.5. *Let Ω be the adjacency matrix of a finite bipartite graph with simple edges, understood as a Coxeter diagram Γ with edge weights equal to 3. Then the eigenvalues λ_i of the bipartite Coxeter transformation associated with Γ are related to the eigenvalues α_i of Ω by*

$$\alpha_i^2 - 2 = \lambda_i + \lambda_i^{-1}.$$

PROOF. This is exactly what is shown in the proof of Proposition 5.3 of McMullen’s article [44]. \square

2.2. Pseudo-Anosov stretch factors and the Thurston–Veech construction. Lemma 3.6 lists the main consequence of the Thurston–Veech construction we need in the present context.

LEMMA 3.6. *For a bi-Perron number $\lambda > 1$, the following are equivalent.*

- (1) *The number λ is the stretch factor of a pseudo-Anosov homeomorphism arising from the Thurston–Veech construction.*

(2) *The number λ is the spectral radius of a product of matrices*

$$\begin{pmatrix} 1 & r \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ -r & 1 \end{pmatrix},$$

where r is the spectral radius of an adjacency matrix of a finite bipartite graph with simple edges.

PROOF OF LEMMA 3.6. The Thurston–Veech construction [65, 68] directly implies the statement for the case where r is the spectral radius of a geometric intersection matrix of multicurves α and β that intersect minimally and fill a surface S . It therefore suffices to show that the set of numbers that appear as spectral radii of such intersection matrices equals the set of numbers that appear as the spectral radii of adjacency matrices of finite bipartite graphs with simple edges.

Given two multicurves α and β that intersect minimally and fill a surface S , their geometric intersection matrix is, by definition, a symmetric nonnegative integer matrix. In the proof of Proposition 2.1 of Hoffman [23], it is shown that any spectral radius of such a matrix is also the spectral radius of an adjacency matrix of a finite bipartite graph with simple edges. This proves the first direction.

Conversely, let A be the adjacency matrix of a finite bipartite graph with simple edges. Since the matrix A is conjugate to a diagonal block matrix with each block on the diagonal corresponding to a connected component of the graph, we can restrict to a block realising the spectral radius of A . In other words, we assume without loss of generality that the graph is connected.

Note that if the connected finite bipartite graph we consider has only a single vertex, then this implies $r = 0$. In particular, the only product of matrices we get in Lemma 3.6 is the identity matrix, which has spectral radius 1. This case is irrelevant, since we restrict ourselves to bi-Perron numbers $\lambda > 1$. We can therefore assume that the connected finite bipartite graph has at least two vertices. It is now straightforward to abstractly construct a closed surface S filled by two multicurves α and β that have the matrix A as their geometric intersection matrix. In order to do so, take two collections of annuli K_i and K'_j that are in one-to-one correspondence with the vertices of the bipartite graph, respecting the bipartition. For each edge of the bipartite graph, locally identify the annuli K_i and K'_j corresponding to the endpoints of the edge along a common square whose

boundary alternatingly belongs to the boundary of K_i and K'_j . We glue such that the orientation of the annuli is respected and their core curves intersect once, see Figure 3.1 for an example of such a glueing corresponding to the complete bipartite graph $K_{2,3}$ on two and three vertices.




FIGURE 3.1. Two collections of annuli, horizontal and vertical, obtained by identifying the boundary of the rectangles as indicated by the letters a, b, c, d, e . After identifying the squares labelled s_1, \dots, s_6 pairwise by translations, the intersection matrix of the core curves of the annuli equals the adjacency matrix of the complete bipartite graph $K_{2,3}$ on two and three vertices. In order to obtain the adjacency matrix of a subgraph, simply omit some of the identifications.

So far, we have constructed a compact surface with boundary. To finish, glue a disc along each boundary component to obtain a closed surface S . By construction, the core curves of the annuli K_i and K'_j define two multicurves α and β , respectively, filling S and with geometric intersection matrix A . Furthermore, the multicurves α and β must intersect minimally, since simple closed curves with zero or one point of intersection always minimise the number of intersections within their respective isotopy classes. \square

2.3. Proof of Theorem 1.8. We prove the following implications: (a) implies (c) implies (b) implies (a).

(a) *implies* (c): Let λ be a bi-Perron number all of whose Galois conjugates are contained in $\mathbb{S}^1 \cup \mathbb{R}$. By Proposition 3.4, there exists a positive symmetric integer matrix M that has $\lambda^k + \lambda^{-k}$ as its spectral radius, for some positive integer k . In the proof of Proposition 2.1 of Hoffman [23], it

is shown that any number that is the spectral radius of a positive symmetric integer matrix is also the spectral radius of an adjacency matrix of a finite bipartite graph with simple edges. In particular, $\lambda^k + \lambda^{-k}$ is the spectral radius of an adjacency matrix Ω of a bipartite graph Γ with simple edges. By Lemma 3.5, the spectral radius x of the bipartite Coxeter transformation associated with Γ equals λ^{2k} . Indeed, we have $(\lambda^k + \lambda^{-k})^2 - 2 = x + x^{-1}$, which yields $\lambda^{2k} + \lambda^{-2k} = x + x^{-1}$ and hence $x = \lambda^{2k}$, as $x \mapsto x + x^{-1}$ is a strictly monotonic function on $[1, \infty)$.

(c) *implies (b)*: In the above implication, we have seen that λ^{2k} is the spectral radius of a bipartite Coxeter transformation associated with a bipartite Coxeter diagram with simple edges if and only if $\lambda^k + \lambda^{-k}$ is the spectral radius of an adjacency matrix Ω of a finite bipartite graph Γ with simple edges. We now use Lemma 3.6 for the matrix product

$$\begin{pmatrix} 1 & \lambda^k + \lambda^{-k} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -(\lambda^k + \lambda^{-k}) & 1 \end{pmatrix} = \begin{pmatrix} 1 - (\lambda^k + \lambda^{-k})^2 & \lambda^k + \lambda^{-k} \\ -(\lambda^k + \lambda^{-k}) & 1 \end{pmatrix},$$

the trace of which equals $2 - (\lambda^k + \lambda^{-k})^2$. In particular, the eigenvalues must satisfy the equation $-t - t^{-1} = (\lambda^k + \lambda^{-k})^2 - 2 = \lambda^{2k} + \lambda^{-2k}$. Hence, the eigenvalue with larger modulus is $-\lambda^{2k}$ and so λ^{2k} is the spectral radius of the matrix product. By Lemma 3.6, the number λ^{2k} is the stretch factor of a pseudo-Anosov homeomorphism arising from the Thurston–Veech construction.

(b) *implies (a)*: Assume that λ^k is the stretch factor of a pseudo-Anosov homeomorphism arising from the Thurston–Veech construction. By a result of Hubert and Lanneau, the associated trace field $\mathbb{Q}(\lambda^k + \lambda^{-k})$ is totally real [25].

We first consider the case where $-\lambda^{-1}$ is not a Galois conjugate of λ . From Proposition 3.2, we know that $\mathbb{Q}(\lambda + \lambda^{-1})$ equals $\mathbb{Q}(\lambda^k + \lambda^{-k})$ for all positive integers k . Hence, if the field $\mathbb{Q}(\lambda^k + \lambda^{-k})$ is totally real, then obviously so must be $\mathbb{Q}(\lambda + \lambda^{-1})$, and all Galois conjugates of $\lambda + \lambda^{-1}$ must be real. We note that all Galois conjugates of λ are roots of the polynomial $t^{\deg(p)}p(t+t^{-1})$, where $p(t)$ is the minimal polynomial of $\lambda + \lambda^{-1}$. In particular, all Galois conjugates λ_i of λ must satisfy $\lambda_i + \lambda_i^{-1} \in \mathbb{R}$ and so $\lambda_i \in \mathbb{S}^1 \cup \mathbb{R}$.

In the case where $-\lambda^{-1}$ is a Galois conjugate of λ , Proposition 3.2 shows that one out of $\mathbb{Q}(\lambda + \lambda^{-1})$ and $\mathbb{Q}(\lambda^2 + \lambda^{-2})$ equals $\mathbb{Q}(\lambda^k + \lambda^{-k})$. In the former case, we are done by the above argument. In the latter case, the same

argument gives that all Galois conjugates λ_i^2 of λ^2 are contained in $\mathbb{S}^1 \cup \mathbb{R}$. Hence, all Galois conjugates λ_i of λ are contained in $\mathbb{S}^1 \cup \mathbb{R} \cup i\mathbb{R}$. We are done by the observation that no Galois conjugate of λ can be totally imaginary. Indeed, assume λ_i is such a Galois conjugate. Then also $\overline{\lambda_i}$ is a Galois conjugate of λ , and we have $\lambda_i^2 = \overline{\lambda_i^2}$. As an irreducible integer polynomial has no multiple zeroes, this implies $\deg(\lambda^2) < \deg(\lambda)$, a contradiction by Lemma 3.1. This finishes the proof of Theorem 1.8.

REMARK 3.7. Our proof strategy of cyclically showing (a) implies (c) implies (b) implies (a) allows us to single out Lemma 3.5 as the only input needed on Coxeter transformations. We note that while we do so fairly implicitly, one can explicitly compare bipartite Coxeter transformations with the elements obtained via the Thurston–Veech construction using a product of exactly two multitwists, see Section 8 of Leininger [33], thus providing a more conceptual proof of (c) implies (b).

3. Galois conjugates of bi-Perron numbers

For the Galois conjugates of a bi-Perron number, we have the following result; the statement is different from the one of Theorem 1.8 in that we only have to use squares for the characterisation, and we only need Coxeter diagrams that are trees.

THEOREM 3.8. *For a Galois conjugate λ of a bi-Perron number, the following are equivalent.*

- (a) *All Galois conjugates of λ are contained in $\mathbb{S}^1 \cup \mathbb{R}$.*
- (b) *The number λ^2 is an eigenvalue of a Coxeter transformation associated with a tree.*

We note that the bi-Perron number in the statement might not be the spectral radius of the Coxeter transformation. Furthermore, we do not include a statement concerning stretch factors, since in the setting of the Thurston–Veech construction we cannot assure that λ is actually a Galois conjugate of a stretch factor, but only an eigenvalue of the action induced on the first homology of the surface by a pseudo-Anosov homeomorphism.

Again, no result of the generality of Theorem 3.8 can be obtained without taking squares: by a result of A’Campo [1], a Coxeter transformation associated with a tree has no negative real eigenvalue, except for possibly -1 .

3.1. Proof of Theorem 3.8. We prove (a) implies (b) implies (a).

(a) *implies (b)*: Let λ be a Galois conjugate of a bi-Perron number all of whose Galois conjugates are contained in $\mathbb{S}^1 \cup \mathbb{R}$. Then $\lambda + \lambda^{-1}$ is a totally real algebraic integer, so by a theorem of Salez [54], $\lambda + \lambda^{-1}$ is an eigenvalue of an adjacency matrix Ω of a finite tree Γ . By Lemma 3.5, the eigenvalues ρ_i of the bipartite Coxeter transformation associated with Γ seen as a bipartite Coxeter diagram with simple edges are related to the eigenvalues α_i of Ω by

$$\alpha_i^2 - 2 = \rho_i + \rho_i^{-1}.$$

By plugging in $\lambda + \lambda^{-1}$ for α_i we see that

$$\lambda^2 + \lambda^{-2} = \rho_i + \rho_i^{-1}.$$

Hence we have $\rho_i = \lambda^2$, that is, λ^2 is an eigenvalue of the Coxeter transformation associated with Γ .

(b) *implies (a)*: This follows from the result that all the eigenvalues of the Coxeter transformation of a tree are contained in $\mathbb{S}^1 \cup \mathbb{R}_{>0}$, due to A'Campo [1]. In particular, we have that all Galois conjugates of λ^2 are contained in $\mathbb{S}^1 \cup \mathbb{R}_{>0}$. Now, since λ is a Perron number, if $\lambda_1, \dots, \lambda_l$ are the Galois conjugates of λ , then $\lambda_1^2, \dots, \lambda_l^2$ are the Galois conjugates of λ^2 by Lemma 3.1. Hence, all Galois conjugates of λ lie in $\mathbb{S}^1 \cup \mathbb{R}$.

Spectral radii of integer matrices

In this chapter, we prove Theorems 1.13 and 1.14 on spectral radii of integer matrices, as well as Theorem 1.16 on the comparison of the minimal spectral radii of reciprocal and skew-reciprocal matrices. The content is copied, adapted and consolidated from [36, 37].

1. The clique polynomial

In this section, we review parts of McMullen's technique using the curve graph and its clique polynomial in order to single out minimal spectral radii among nonnegative matrices. We try to keep the discussion as concise as possible and refer to the original article [45] and the references therein for a more complete discussion.

Let Γ be a directed graph. A *simple closed curve* in Γ is the union of directed edges describing a closed directed loop in Γ that visits every vertex at most once. The *curve graph* G of Γ is obtained as follows: there is a vertex for every simple closed curve in Γ , and two vertices are connected by an edge if and only if the corresponding simple closed curves have no vertex of Γ in common. Each vertex of G is given a weight describing the number of edges contained in the simple closed curve.

A subset K of the vertices of G is a *clique* if the subgraph induced by K is complete. The *clique polynomial* of G is defined to be

$$Q(t) = \sum_K (-1)^{|K|} t^{w(K)},$$

where we also allow $K = \emptyset$, and $w(K)$ is the sum of all weights of vertices in K .

With a nonnegative square matrix A of dimension $n \times n$, we associate a directed graph Γ_A that has n vertices and directed edges between the vertices according to the coefficients of A . Let G_A be the associated curve graph and let $Q_A(t)$ be its clique polynomial. By a well-known result in graph theory, the characteristic polynomial of A is the reciprocal of $Q_A(t)$,

that is, $\chi_A(t) = t^n Q_A(t^{-1})$. In particular, the spectral radius of A equals the inverse of the smallest modulus among the roots of $Q_A(t)$.

1.1. McMullen's classification of graphs with small growth.

McMullen defines a minimal growth rate $\lambda(G)$ for graphs G . We refer to McMullen's original article [45] for more details. The only statement we need for our purposes is that if A is a nonnegative matrix of dimension $n \times n$ and with spectral radius $\rho(A)$, then $\lambda(G_A)$ is a lower bound for the normalised spectral radius $\rho(A)^n$.

| G | $\lambda(G)$ |
|--|------------------------------|
| nA_1 $\circ \cdots \circ$ | n |
| A_2^* $\circ \text{---} \circ \quad \circ$ | 4 |
| A_2^{**} $\circ \text{---} \circ \quad \circ \quad \circ$ | $3 + 2\sqrt{2} \approx 5.82$ |
| A_2^{***} $\circ \text{---} \circ \quad \circ \quad \circ \quad \circ$ | $4 + 2\sqrt{2} \approx 7.46$ |
| A_3^* $\circ \text{---} \circ \text{---} \circ \quad \circ$ | $3 + 2\sqrt{2} \approx 5.82$ |
| Y^* $\begin{array}{c} \circ \quad \circ \\ \diagdown \quad / \\ \circ \\ \diagup \quad \diagdown \\ \circ \end{array} \quad \circ$ | $4 + 2\sqrt{2} \approx 7.46$ |

FIGURE 4.1. Some graphs and their minimal growth rates.

We are interested in irreducible matrices A . This means that the directed graph Γ_A is strongly connected, which in turn implies that the associated curve graph G_A has complement G'_A that is connected. Here, the complement G'_A is defined to be the graph with the same vertex set as G_A but the complementary edges. There are very few curve graphs G with G' connected and minimal growth rate $\lambda(G) < 8$. Our argument is based on the following classification due to McMullen [45].

THEOREM 4.1 (Theorem 1.6 in McMullen [45]). *The graphs G with G' connected and $1 < \lambda(G) < 8$ are given by*

$$A_2^*, A_2^{**}, A_2^{***}, A_3^*, Y^* \text{ and } nA_1,$$

For $2 \leq n \leq 7$.

This result tells us that if we want to describe all irreducible nonnegative matrices with normalised spectral radius < 8 , all we have to do is check among those whose associated curve graph is among the ones shown in Figure 4.1. In fact, in Section 2 we split the proof of Theorems 1.13 and 1.14 into four propositions, dealing with the four distinct cases. In each case, the first thing we do is to realise the proposed minimal normalised spectral radius, which (except in one case for $g = 3$) turns out to be < 8 . Then McMullen's classification result applies and we only have to check the curve graphs given in Theorem 4.1 to finish the proof.

2. Skew-reciprocity and minimal spectral radii

In this section, we prove Theorems 1.13 and 1.14. We break down the proof into four separate propositions, distinguishing between the irreducible and the primitive case, as well as the case of even and odd g .

The condition of skew-reciprocity poses slightly different constraints on the coefficients of the polynomial than reciprocity. First of all, we note that if the roots of a polynomial $f \in \mathbb{Z}[t]$ are invariant under the transformation $\lambda \mapsto -\lambda^{-1}$, then we have $f(t) = \pm t^{\deg(f)} f(-t^{-1})$. This entails the following constraints.

LEMMA 4.2. *Let $f \in \mathbb{Z}[t]$ be a monic skew-reciprocal polynomial of degree $2g$. Then we have the following conditions on the coefficients of f :*

- (1) *the moduli of the coefficients of t^d and t^{2g-d} agree. More precisely,*
- (2) *if g is even and $f(0) = 1$, the coefficients of t^d and t^{2g-d} agree for even d and differ by a sign for odd d ,*
- (3) *if g is even and $f(0) = -1$, the coefficients of t^d and t^{2g-d} agree for odd d and differ by a sign for even d . In particular, the middle coefficient of f vanishes,*
- (4) *if g is odd and $f(0) = 1$, the coefficients of t^d and t^{2g-d} agree for even d and differ by a sign for odd d . In particular, the middle coefficient of f vanishes,*
- (5) *if g is odd and $f(0) = -1$, the coefficients of t^d and t^{2g-d} agree for odd d and differ by a sign for even d .*

PROOF. Let $f(t) = a_{2g}t^{2g} + \dots + a_0$ be a skew-reciprocal polynomial. The polynomial relation $f(t) = \pm t^{2g}f(-t^{-1})$ given by skew-reciprocity translates to the relation

$$a_d = \pm(-1)^{2g-d}a_{2g-d} = \pm(-1)^d a_{2g-d}$$

for each pair coefficients a_d and a_{2g-d} . Clearly, the coefficients are symmetric up to a possible sign that alternates between $+1$ and -1 as we change the index d of the coefficient by one. In particular, the sign is the same for all even d and it is the same for all odd d . Now recall that $f(t)$ is monic, that is, $a_{2g} = 1$. In this case, $f(0) = 1$ means that the coefficients a_d and a_{2g-d} agree for $d = 0$ and hence all even d , and they differ by a sign for odd d . Similarly, $f(0) = -1$ means that the coefficients differ by a sign for $d = 0$ and hence for all even d , and they agree for even d . Finally, the middle coefficient a_g needs to be zero if a_d and a_{2g-d} differ by a sign for all d with the same parity as g . This distinction yields the four different cases (2)–(5) described in the statement of Lemma 4.2. \square

EXAMPLE 4.3. To see the main proof ideas applied to the simplest nontrivial example, we now determine which spectral radii are obtained by skew-reciprocal matrices $A \in \text{GL}_{2g}(\mathbb{Z})$ with curve graph $G_A = 2A_1$, which is the graph with two isolated vertices. In this case $Q(t) = 1 - t^a - t^b$, where a and b are the weights of the vertices. The polynomial $Q(t)$ is of degree $2g$, and without loss of generality we assume $b = 2g$. The only way to have the moduli of the coefficients symmetrically distributed as in (1) of Lemma 4.2 is if $a = g$. Therefore, the only clique polynomial we possibly obtain is $Q(t) = 1 - t^g - t^{2g}$. Hence, the only characteristic polynomial we possibly obtain is $t^{2g} - t^g - 1$. We make the following observations:

- (1) if A is primitive, then $g = 1$. Indeed, otherwise the characteristic polynomial is a polynomial in t^g with $g > 1$. Such a polynomial cannot be the characteristic polynomial of a primitive matrix. On the other hand, for $g = 1$, the polynomial $t^2 - t - 1$ is the minimal polynomial of the golden ratio, realised as the characteristic polynomial of the matrix $\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$;
- (2) for g is odd, the polynomial $t^{2g} - t^g - 1$ is the characteristic polynomial of a nonnegative irreducible matrix in $\text{GL}_{2g}(\mathbb{Z})$, namely a

standard companion matrix. For example, $t^6 - t^3 - 1$ is the characteristic polynomial of the matrix

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

which is irreducible;

- (3) for even g , we note that since the constant coefficient of the characteristic polynomial of A is negative, we are in case (3) of Lemma 4.2. In particular, the coefficient of t^g should be zero instead of -1 . This means that for even g , there are no skew-reciprocal matrices $A \in \mathrm{GL}_{2g}(\mathbb{Z})$ with curve graph $2A_1$.

In summary, we obtain that the following spectral radii can be realised for matrices $A \in \mathrm{GL}_{2g}(\mathbb{Z})$ with curve graph $2A_1$:

- (i) among primitive skew-reciprocal matrices $A \in \mathrm{GL}_{2g}(\mathbb{Z})$, only the golden ratio is realised, for $g = 1$. For $g > 1$, there are no primitive skew-reciprocal matrices with curve graph $2A_1$;
- (ii) among nonnegative irreducible skew-reciprocal $A \in \mathrm{GL}_{2g}(\mathbb{Z})$, the largest root of the polynomial $t^{2g} - t^g - 1$ is realised, for odd g . For even g , there are no nonnegative irreducible skew-reciprocal matrices $A \in \mathrm{GL}_{2g}(\mathbb{Z})$ with curve graph $2A_1$.

2.1. The irreducible case.

2.1.1. The case of odd g .

PROPOSITION 4.4. *Let $g \geq 1$ odd. Among all skew-reciprocal nonnegative irreducible matrices $A \in \mathrm{GL}_{2g}(\mathbb{Z})$, the minimal spectral radius > 1 is realised by the largest root λ_{2g} of the polynomial $t^{2g} - t^g - 1$.*

PROOF. Let $A \in \mathrm{GL}_{2g}(\mathbb{Z})$ be a nonnegative skew-reciprocal matrix. Then its square A^2 is a reciprocal matrix. In particular, by McMullen's result on minimal normalised spectral radii for nonnegative reciprocal matrices [45] we know that its normalised spectral radius must be at least φ^4 . Therefore, the normalised spectral radius of A must be at least φ^2 , which incidentally equals $(\lambda_{2g})^{2g}$. In order to finish the proof, it therefore suffices to realise the polynomial $t^{2g} - t^g - 1$ as the characteristic polynomial of a

nonnegative irreducible matrix in $\text{GL}_{2g}(\mathbb{Z})$. This is straightforward, as it can be achieved by a standard companion matrix, see (2) in Example 4.3 for the case $g = 3$. \square

REMARK 4.5. It actually follows from McMullen's classification that the polynomial $t^{2g} - t^g - 1$ is the unique characteristic polynomial that can appear for a matrix that minimises the spectral radius. Indeed, only the graph $2A_1$ can appear as curve graph, with clique polynomial $1 - t^a - t^b$. For this polynomial to be skew-reciprocal we must either have $a = 2g$ and $b = g$ or $b = 2g$ and $a = g$. Both cases yield our candidate polynomial.

2.1.2. *The case of even g .* We note that the above proof does not work for even g : if g is even, then the polynomial $t^{2g} - t^g - 1$ is not skew-reciprocal, as noted in Example 4.3. We instead have the following minimisers.

PROPOSITION 4.6. *Let $g \geq 2$ even. Among all skew-reciprocal nonnegative irreducible matrices $A \in \text{GL}_{2g}(\mathbb{Z})$, the minimal spectral radius > 1 is realised by the largest root λ_{2g} of the polynomial $t^{2g} - t^{g+1} - t^{g-1} - 1$.*

PROOF. The largest root λ_{2g} of the polynomial $t^{2g} - t^{g+1} - t^{g-1} - 1$ is clearly realised as the spectral radius of a nonnegative irreducible matrix $A \in \text{GL}_{2g}(\mathbb{Z})$. Indeed, again we can achieve this by a standard companion matrix. We now note that $(\lambda_{2g})^{2g}$ is a descending sequence converging to $3 + 2\sqrt{2}$ and starting at $\varphi^4 < 7$ for $g = 2$. In particular, we can finish the proof by showing that for even $g \geq 2$, λ_{2g} minimises the spectral radius among all skew-reciprocal nonnegative irreducible matrices $A \in \text{GL}_{2g}(\mathbb{Z})$ that have one of the graphs in Figure 4.1 except A_2^{***} or Y^* as their curve graph.

- (1) $G = 2A_1$. As noted in Example 4.3, there are no nonnegative irreducible skew-reciprocal matrices $A \in \text{GL}_{2g}(\mathbb{Z})$ with g even and $2A_1$ as their curve graph.
- (2) $G = 3A_1$. In this case $Q(t) = 1 - t^a - t^b - t^c$. Without loss of generality we assume $c = 2g$. If we want the moduli of the coefficients symmetrically distributed, we are left with the options

$$1 - t^{g-d} - t^{g+d} - t^{2g}$$

for $0 \leq d \leq g$. The case $d = g$ is ruled out as the resulting polynomial is a monomial. The case $d = 0$ is ruled out by Lemma 4.2, as above. It follows that $1 < d < g$. By Proposition 3.2 in [39], we know that the largest root of the reciprocal

polynomial $t^{2g} - t^{g+d} - t^{g-d} - 1$ is a strictly increasing function of d . So, the smallest spectral is obtained for $d = 1$, resulting in our candidate polynomial $t^{2g} - t^{g+1} - t^{g-1} - 1$.

- (3) $G = 4A_1$ or $G = 6A_1$. As for $G = 2A_1$, the number of terms is odd. The only way to have the moduli of the coefficients symmetrically distributed is to have a middle coefficient, a contradiction to Lemma 4.2.
- (4) $G = 5A_1$. In this case $Q(t) = 1 - t^a - t^b - t^c - t^d - t^e$. Without loss of generality $e = 2g$. As in the case $G = 3A_1$, there must be at least two paired terms of power $\neq 0, g$. We assume without loss of generality that $0 < a < g < b = 2g - a < 2g$. Since the polynomial reciprocal to $Q(t)$ is realised by a standard companion matrix, we can delete its coefficients that correspond to the terms $-t^c$ and $-t^d$ and obtain a matrix with strictly smaller spectral radius and characteristic polynomial $t^{2g} - t^b - t^a - 1$, where the a, b and g satisfy $0 < a < g < b = 2g - a < 2g$. We have shown in the case $G = 3A_1$ that the minimal spectral radius obtained by a matrix with such a characteristic polynomial is our candidate λ_{2g} .
- (5) $G = A_2^*$. In this case, $Q(t) = 1 - t^a - t^b - t^c + t^{a+b}$, where c is the weight on the isolated vertex of A_2^* . Since there are three terms not paired with the constant term, we deduce there must be a nonvanishing middle coefficient. By Lemma 4.2, this means that the leading coefficient of $Q(t)$ is positive, so we have $a + b = 2g$ and $c = g$. We note that the possibilities that remain are of the form $1 - t^{g-d} - t^g - t^{g+d} + t^{2g}$, which are reciprocal. Theorem 7.3 by McMullen [45] then provides that the normalised spectral radius is $\geq \varphi^4 = \lambda_4$ and $> \lambda_{2g}$ for $g > 2$.
- (6) $G = A_2^{**}$. In this case, $Q(t) = 1 - t^a - t^b - t^c - t^d + t^{a+b}$, where c and d are the weights on the isolated vertices of A_2^{**} . If the leading coefficient of $Q(t)$ is positive, then $Q(t)$ is reciprocal. By McMullen's analysis of the curve graph A_2^{**} for reciprocal weights, a normalised spectral radius in this case must be $\geq (2 + \sqrt{3})^2 > \lambda_{2g}$. It remains to consider the case where the leading coefficient of $Q(t)$ is negative. Without loss of generality, we assume $c = 2g$. We have the following conditions on the other parameters a, b, d .
- $a + b < 2g$ and hence $a + b \leq 2g - 2$. Indeed, t^{a+b} appears with a positive sign and must be paired with a term $-t^x$ with

negative sign, for $x = a, b$ or d . In particular, Lemma 4.2 implies that $a + b$ must be even.

- either $d + a = 2g$ or $d + b = 2g$. We assume without loss of generality that $d + b = 2g$. Thus, $1 \leq a < b, d < a + b \leq 2g - 2$.

Now consider the directed graph $\Gamma_{a,b,d}$, where a weight w on an




FIGURE 4.2. The directed graph $\Gamma_{a,b,d}$.

edge indicates $w - 1$ additional vertices placed on the edge. We note that the clique polynomial of the curve graph of $\Gamma_{a,b,d}$ is $Q(t)$. Furthermore, deleting the edge of length 1 in $\Gamma_{a,b,d}$ that forms the simple closed curve of length a strictly decreases the spectral radius of the associated adjacency matrix. Furthermore, the new curve graph is $3A_1$ and the new clique polynomial is obtained by removing the terms $-t^a$ and t^{a+b} , and hence skew-reciprocal. This is a case we have already dealt with.

- (7) $G = A_3^*$. In this case, $Q(t) = 1 - t^a - t^b - t^c - t^d + t^{a+b} + t^{b+c}$, where d is the weight of the isolated vertex and b is the weight of the vertex of degree two of A_3^* . Since the number of summands is odd, the middle term must have a nonvanishing coefficient. By Lemma 4.2, the leading coefficient of $Q(t)$ is positive, and we assume without loss of generality that $b + c = 2g$. Note that since the coefficients of t^b and t^c have the same sign, b and c need to be even. We now distinguish three cases: either $a + b = g$, $a = g$, or $d = g$.

- if $a + b = g$, then $Q(t) = 1 - t^a - t^b + t^g - t^{2g-a} - t^{2g-b} + t^{2g}$, which is reciprocal. By McMullen's analysis of the curve graph A_3^* for reciprocal weights [45], the normalised spectral radius must either be $> 12.5 > \varphi^4$, or $Q(t)$ is among the examples arising from A_2^* , a case we have dealt with already;

- if $a = g$, then $d + a + b = 2g$. Since b and g are even, so must be $a + b = g + b$, and hence also d . By Lemma 4.2, the coefficients of t^d and of t^{a+b} should have the same sign, a contradiction;
- if $d = g$, we have $2a + b = 2g$, and hence $a \leq g - 1$. Let $\Gamma'_{a,b,c}$ be defined as in Figure 4.3, where the edge of weight $g - a - 1$




FIGURE 4.3. The directed graph $\Gamma'_{a,b,c}$.

is contracted in case $a = g - 1$. We note that the clique polynomial of the curve graph of $\Gamma'_{a,b,c}$ is exactly our $Q(t)$, where $d = g$. Deleting the edge of length one that forms the simple closed curve of length a strictly decreases the spectral radius of the associated adjacency matrix. The clique polynomial we obtain after this deletion of an edge is of the form $1 - t^b - t^g - t^{2g-b} + t^{2g}$, a case we have dealt with in our study of $G = A_2^*$.

□

REMARK 4.7. Our proof of Proposition 4.6 actually shows that for $g \neq 2$, the polynomial $t^{2g} - t^{g+1} - t^{g-1} - 1$ is the unique characteristic polynomial that can appear for a matrix minimising the spectral radius. Except if $g = 2$, where we have a second possibility (appearing in (5)) for the characteristic polynomial: $t^4 - 3t^2 + 1$. In the case $g = 2$, both minimising polynomials are divisible by the minimal polynomial of the golden ratio.

REMARK 4.8. In the cases (4), (6) and (7) of the proof of Proposition 4.6, we construct irreducible matrices and reduce some of their coefficients in order to obtain irreducible matrices with strictly smaller spectral radii that belong to other cases we have already dealt with. For a quicker proof of Proposition 4.6, we could use the monotonicity property for the

spectral radius formulated by McMullen [45] on the level of the weighted curve graph. However, this monotonicity is not strict in general. In particular, this proof strategy seems to fail to provide the uniqueness of the minimising characteristic polynomials described in Remark 4.7.

2.2. The primitive case.

2.2.1. The case of even g .

PROPOSITION 4.9. *Let $g \geq 2$ even. Among all skew-reciprocal primitive matrices $A \in \text{GL}_{2g}(\mathbb{Z})$, the minimal spectral radius > 1 is realised by the largest root μ_{2g} of the polynomial $t^{2g} - t^{g+1} - t^{g-1} - 1$.*

PROOF. By Proposition 4.6, we know that μ_{2g} is actually the minimal spectral radius among all nonnegative irreducible matrices. In order to prove the result, it is enough to show that the polynomial $t^{2g} - t^{g+1} - t^{g-1} - 1$ is the characteristic polynomial of a primitive matrix. This is the case. Indeed, we can take the standard companion matrix for the polynomial and draw its directed adjacency graph. We directly see that there are directed cycles of length $g - 1, g + 1$ and $2g$. In order to show that the matrix is primitive, it suffices to show that their common greatest divisor is 1. Let n be a positive integer that divides both $2g$ and $g + 1$. Since g is even, $g + 1$ is odd and so n has to be odd as well. Now since n is odd and divides $2g$, it divides g . We have that n divides both g and $g + 1$ and therefore $n = 1$. This finishes the proof. \square

2.2.2. The case of odd g .

PROPOSITION 4.10. *Let $g \geq 3$ odd. Among all skew-reciprocal primitive matrices $A \in \text{GL}_{2g}(\mathbb{Z})$, the minimal spectral radius > 1 is realised by the largest root μ_{2g} of the polynomial $t^{2g} - t^{g+2} - t^{g-2} - 1$.*

PROOF FOR $g \geq 5$. We take a standard companion matrix to realise the largest root μ_{2g} of the polynomial $t^{2g} - t^{g+2} - t^{g-2} - 1$ as a spectral radius of a nonnegative matrix in $\text{GL}_{2g}(\mathbb{Z})$. Furthermore, the associated directed graph has simple closed curves of lengths $2g, g + 2, g - 2$, which have greatest common divisor 1. Indeed, since g is odd so is $g + 2$, so if n divides both $2g$ and $g + 2$, then it must be odd itself and hence divide g . But then, since n divides both g and $g + 2$, it must divide 2. But n being odd implies $n = 1$. This shows that the companion matrix we constructed is primitive.

We now note that $(\mu_{2g})^{2g}$ is a descending sequence converging to $3 + 2\sqrt{2}$, with first values $\mu_6 \approx 8.19$ and $\mu_{10} \approx 6.42$. The example in the case $g = 3$ is

too large to be covered by McMullen's classification, and we give a separate argument covering this case below. For $g \geq 5$ odd, we can proceed as before, and finish the proof by showing that μ_{2g} minimises the spectral radius among all skew-reciprocal primitive matrices $A \in \text{GL}_{2g}(\mathbb{Z})$ with one of the graphs in Figure 4.1 except A_2^{***} or Y^* as their curve graph.

- (1) $G = 2A_1$. As we noted in Example 4.3, there exist no primitive skew-reciprocal matrices $A \in \text{GL}_{2g}(\mathbb{Z})$ with $g > 1$ and $2A_1$ as their curve graph.
- (2) $G = 3A_1$. In this case, $Q(t) = 1 - t^a - t^b - t^c$. Without loss of generality, we assume that $c = 2g$, which implies $a = 2g - b$ if we want symmetrically distributed coefficients. We have multiple possibilities for a :
 - $a = g$. In this case, $Q(t) = 1 - 2t^g - t^{2g}$, which is not primitive.
 - $a = g - 1$. In this case $Q(t) = 1 - t^{g-1} - t^{g+1} - t^{2g}$. Lemma 4.2 implies that this polynomial is not skew-reciprocal. Indeed, for it to be skew-reciprocal, the coefficients of t^{g+1} and t^{g-1} would have to differ by a sign since $g - 1$ is even;
 - $a = g - 2$. This case gives exactly our candidate polynomial with largest root μ_{2g} ;
 - $a < g - 2$. By Proposition 3.2 in [39], we know that the largest root of the reciprocal polynomial $t^{2g} - t^{g+d} - t^{g-d} - 1$ is a strictly increasing function of d . In particular, the spectral radii we obtain for $a < g - 2$ are strictly larger than our candidate.
- (3) $G = 4A_1$. In this case, $Q(t) = 1 - t^a - t^b - t^c - t^d$. We realise the reciprocal of $Q(t)$ as the characteristic polynomial of a standard companion matrix. Since there are five terms, there must be a middle coefficient. Deleting this middle coefficient amounts to decreasing a coefficient of the companion matrix from 1 to 0, strictly reducing the spectral radius. After this modification, the polynomial is among the examples we have already dealt with in the case $G = 3A_1$.
- (4) $G = 5A_1$ or $G = 6A_1$. This case can be dealt with in the same way as the case $G = 4A_1$. We delete the coefficients of a pair of terms whose powers add to $2g$ (in the case of $G = 5A_1$) and additionally the middle coefficient (in the case of $G = 6A_1$).
- (5) $G = A_2^*$. In this case, $Q(t) = 1 - t^a - t^b - t^c + t^{a+b}$, where c is the weight on the isolated vertex of A_2^* . Since there are five terms,

there must be a middle coefficient, which by Lemma 4.2 implies that the leading coefficient is negative. We therefore must have $c = 2g$ and we can assume without loss of generality that $b = g$ to get a polynomial of the form

$$Q(t) = 1 - t^a - t^g + t^{a+g} - t^{2g},$$

and in particular $2a + g = 2g$. But this implies that $g = 2a$ is even, a contradiction.

- (6) $G = A_2^{**}$. In this case, $Q(t) = 1 - t^a - t^b - t^c - t^d + t^{a+b}$, where c and d are the weights on the isolated vertices of A_2^{**} . If the leading coefficient of $Q(t)$ is positive, then $a + b = 2g$ and the resulting polynomial is reciprocal. By McMullen's analysis of the curve graph A_2^{**} for reciprocal weights [45], a normalised spectral radius in this case must be $\geq (2 + \sqrt{3})^2 > \mu_{2g}$. It remains to consider the case of a negative leading coefficient. Without loss of generality, we assume $c = 2g$. In order to have symmetrically distributed moduli of the coefficients, we must either have $2a + b = 2g$ and $b + d = 2g$ or $2b + a = 2g$ and $a + d = 2g$. Both cases imply that d is even, and hence so must be b (in the former case) or a (in the latter case). In both cases, we get a contradiction to Lemma 4.2, which states that the coefficients must differ by a sign for even powers.
- (7) $G = A_3^*$. In this case, $Q(t) = 1 - t^a - t^b - t^c - t^d + t^{a+b} + t^{b+c}$. Since there are seven terms, there must be a nonvanishing middle coefficient. By Lemma 4.2, this can only happen if the leading coefficient is negative. We must have $d = 2g$ and get

$$Q(t) = 1 - t^a - t^b - t^c + t^{a+b} + t^{b+c} - t^{2g}.$$

We distinguish cases depending on which term has power g :

- if one among a, b and c equals g , we have $a, b, c \leq g$ and furthermore $a + b, b + c > g$. This implies that $a + b, b + c$ and two among a, b, c are even by Lemma 4.2. But then clearly all among a, b, c are even, and hence is g , a contradiction;
- if $a + b = g$, then $a, b < g$ and hence $c, b + c > g$. Also $b + c < 2g$ so we must have $a + c = 2g = 2b + c$. In particular, $a = 2b$ is even, and hence so must be c . This contradicts Lemma 4.2, which states that coefficients must differ by a sign for terms

with even powers. The argument for the case $b + c = g$ is obtained by switching a and c .

□

PROOF FOR $g = 3$. The candidate polynomial $t^6 - t^5 - t - 1$ has maximal real root $\mu_6 \approx 1.4196 > \sqrt{2}$, so we only need to check other polynomials with roots bounded from above by this number, and bounded from below by $\sqrt{2}$. Indeed, our proof in the case $g \geq 5$ shows that there is no spectral radius $< 8^{\frac{1}{6}} = \sqrt{2}$ among skew-reciprocal primitive matrices $A \in \text{GL}_6(\mathbb{Z})$. We now distinguish cases depending on the determinant of A .

Case 1: $\det(A) = 1$. In this case, the characteristic polynomial must have a factor $(t^2 + 1)$. The reason for this is that the eigenvalues of a skew-reciprocal matrix come in groups:

- if $\lambda \notin \mathbb{R}$, $\lambda \neq \pm i$ is an eigenvalue, then so are $-\lambda^{-1}$, $\bar{\lambda}$ and $-\bar{\lambda}^{-1}$. These four roots of the characteristic polynomial contribute $+1$ to the determinant,
- if $\lambda \in \mathbb{R}$, $\lambda \neq 0$ is an eigenvalue, then so is $-\lambda^{-1}$. These two roots of the characteristic polynomial contribute -1 to the determinant,
- if $\lambda = \pm i$ is an eigenvalue, then so is $\bar{\lambda} = -\lambda$. These two roots contribute $+1$ to the determinant.

For $g = 3$, the only way for determinant $+1$ is if the last case appears at least once. This implies that the polynomial is divisible by $(t - i)(t + i) = t^2 + 1$. By Perron–Frobenius theory, we know that A has at least two real roots. In particular, the first case cannot occur and the spectral radius is a totally real algebraic integer with at most two Galois conjugates of modulus > 1 . If it is not an integer, it is an algebraic integer of degree at least two. In particular, the Mahler measure of its minimal polynomial is at least φ^2 by Corollary 1' of Schinzel [55]. Since at most two Galois conjugates have modulus > 1 , it follows that the modulus of the larger root is bounded from below by $\varphi \approx 1.61 > \mu_6$.

Case 2: $\det(A) = -1$. We first rule out the case where all eigenvalues are real. The spectral radius is an algebraic integer of degree at most six that is maximal in modulus among all its Galois conjugates. Skew-reciprocity of A and the fact that the minimal polynomial has constant coefficient ± 1 imply that at most half of the Galois conjugates of the spectral radius can have modulus > 1 . Again, Schinzel's Corollary 1' in [55] implies that the spectral radius is bounded from below by $\varphi \approx 1.61 > \mu_6$.

In the remaining case, the spectral radius ρ of A is of degree six and has four non-real Galois conjugates $\lambda, -\lambda^{-1}, \bar{\lambda}$ and $-\bar{\lambda}^{-1}$. Let the characteristic polynomial of A be given by

$$\begin{aligned} P(t) &= t^6 + at^5 + bt^4 + ct^3 - bt^2 + at - 1 \\ &= (t - \rho)(t + \rho^{-1})(t - \lambda)(t + \lambda^{-1})(t - \bar{\lambda})(t + \bar{\lambda}^{-1}). \end{aligned}$$

We get the following estimates for the coefficients a, b and c .

- Since $\rho < 1.42$, we have $|\rho - \rho^{-1}| < 0.72$. For the coefficient a , we get

$$\begin{aligned} |a| &\leq |\rho - \rho^{-1}| + |(\lambda + \bar{\lambda}) - (\lambda^{-1} + \bar{\lambda}^{-1})| \\ &< 0.72 + 2|\operatorname{Re}(\lambda) - \operatorname{Re}(\lambda^{-1})| = 0.72 + 2\left|\operatorname{Re}(\lambda) - \frac{\operatorname{Re}(\lambda)}{|\lambda|^2}\right| \\ &= 0.72 + 2|\operatorname{Re}(\lambda)| \left(1 - \frac{1}{|\lambda|^2}\right) < 0.72 + 1.44 < 3, \end{aligned}$$

where we used $|\operatorname{Re}(\lambda)| \leq |\lambda| < 1.42$ in the second to last inequality. Up to replacing $P(t)$ by $P(-t)$, we may assume that $a \in \{-2, -1, 0\}$.

- Since $|\lambda| \leq \rho < 1.42$, we have $|\lambda - \lambda^{-1}| < 2.13$ and $|\bar{\lambda} - \bar{\lambda}^{-1}| < 2.13$. We calculate

$$c = -2a + (\rho^{-1} - \rho)(\lambda^{-1} - \lambda)(\bar{\lambda}^{-1} - \bar{\lambda}),$$

where

$$|(\rho^{-1} - \rho)(\lambda^{-1} - \lambda)(\bar{\lambda}^{-1} - \bar{\lambda})| < 0.72 \cdot (2.13)^2 < 3.2.$$

In particular, $c \in \{-2a - 3, \dots, -2a + 3\}$.

- We have

$$b = -3 + (\rho^{-1} - \rho)(\lambda^{-1} - \lambda + \bar{\lambda}^{-1} - \bar{\lambda}) + (\lambda^{-1} - \lambda)(\bar{\lambda}^{-1} - \bar{\lambda}),$$

where

$$\begin{aligned} |(\rho^{-1} - \rho)(\lambda^{-1} - \lambda + \bar{\lambda}^{-1} - \bar{\lambda}) + (\lambda^{-1} - \lambda)(\bar{\lambda}^{-1} - \bar{\lambda})| \\ < 0.72 \cdot 1.44 + (2.13)^2 < 5.58. \end{aligned}$$

In particular, $b \in \{-8, \dots, 2\}$.

There are now $3 \cdot 7 \cdot 11 = 231$ remaining polynomials to check. Listing them all as well as their roots, it is a quick check to see which ones among them have a real root with modulus between 1.41 and 1.42; only three polynomials

remain. Among these three polynomials, only our candidate $t^6 - t^5 - t - 1$ remains if we insist that the real root with modulus between 1.41 and 1.42 be maximal in modulus among all the roots of the polynomial. \square

REMARK 4.11. Again we have shown that for $g \geq 3$, the polynomial

$$t^{2g} - t^{g+2} - t^{g-2} - 1$$

is the unique characteristic polynomial that can appear for a matrix minimising the spectral radius. The case $g = 3$ is not covered by McMullen's classification but our ad-hoc argument rules out all other possibilities for characteristic polynomials: while we gave ourselves the liberty to replace $P(t)$ by $P(-t)$ during the proof, we note that the root of $t^6 + t^5 + t - 1$ that is maximal in modulus is real and negative. Therefore, the polynomial $t^6 + t^5 + t - 1$ is not the characteristic polynomial of a primitive matrix.

3. A reformulation of the question of Schinzel and Zassenhaus

The goal of this section is to prove Theorem 1.16.

Recall that a polynomial $f \in \mathbb{Z}[t]$ of even degree $2d$ is called reciprocal if we have $f(t) = \pm t^{2d} f(t^{-1})$, and skew-reciprocal if $f(t) = \pm t^{2d} f(-t^{-1})$.

LEMMA 4.12. *Let $f \in \mathbb{Z}[t]$ be a monic skew-reciprocal polynomial of degree 2^{i+1} with $|\overline{f}| > 1$. Then either $f(t) = g(t^2)$, where $g(t)$ is a reciprocal polynomial of degree 2^i , or f has a nonreciprocal irreducible factor.*

PROOF. Let $f \in \mathbb{Z}[t]$ be a monic skew-reciprocal polynomial of degree 2^{i+1} such that $|\overline{f}| > 1$.

If a polynomial $f \in \mathbb{Z}[t]$ of degree 2^{i+1} is both reciprocal and skew-reciprocal, we have $f(t) = g(t^2)$, where $g(t)$ is a reciprocal polynomial of degree 2^i .

If f is not reciprocal, it must have at least one nonreciprocal irreducible factor. \square

Let λ_i and $\tilde{\lambda}_i$ be the smallest houses larger than 1 among all monic integer reciprocal and skew-reciprocal polynomials of degree 2^i , respectively. Furthermore, let $r_i = 2^i \log(\lambda_i)$ and $s_i = 2^i \log(\tilde{\lambda}_i)$.

Our proof is based on the following theorem of Breusch [7].

THEOREM 4.13 (Breusch [7]). *The Mahler measure of any integer non-reciprocal irreducible polynomial other than t is greater than 1.179.*

LEMMA 4.14. *For $i \geq 1$, we have $s_{i+1} \geq \min\{r_i, \log(1.179)\}$.*

PROOF. Let $f \in \mathbb{Z}[t]$ be a monic skew-reciprocal polynomial of degree 2^{i+1} such that $\overline{|f|} > 1$. We use Lemma 4.12 to distinguish two cases. If $f(t) = g(t^2)$, where $g(t)$ is a reciprocal polynomial of degree 2^i , we have $\overline{|f|^2} = \overline{|g|}$. It follows that $2^{i+1}\log\overline{|f|} = 2^i\log\overline{|g|} \geq r_i$. On the other hand, if f has a nonreciprocal irreducible factor, then Theorem 4.13 implies $2^{i+1}\log\overline{|f|} \geq \log(1.179)$. \square

LEMMA 4.15. *The answer to the question of Schinzel and Zassenhaus is positive exactly if the sequence $\{r_i\}$ is bounded strictly away from zero.*

PROOF. By Theorem 4.13, the question of Schinzel and Zassenhaus is equivalent to the same question restricted to reciprocal polynomials. Furthermore, any reciprocal polynomial $f(t)$ can be multiplied by $(t+1)^k$, where k is at most the degree of $f(t)$, so that it becomes reciprocal of degree 2^i , for some $i \geq 1$, keeping its house. This means that the question of Schinzel and Zassenhaus is equivalent to the same question for (not necessarily irreducible) reciprocal polynomials of degree 2^i . The statement of the lemma now follows from the fact that $\{r_i\} = \{2^i \log(\lambda_i)\}$ is strictly bounded away from zero exactly if there exists a constant c such that $\lambda_i > 1 + \frac{c}{2^i}$ for all i . \square

PROOF OF THEOREM 1.16. For one direction, we assume there exists a sequence $\left\{\frac{q_{N_j}}{q_{n_j}}\right\}$, where $0 < n_j < N_j$, that converges to zero. If $f(t)$ is a reciprocal polynomial of even degree, then $f(t^2)$ is a skew-reciprocal polynomial. This implies $s_i \leq r_{i-1}$. In particular, we have

$$r_{N_j} = r_{n_j} \prod_{i=n_j+1}^{N_j} \frac{r_i}{r_{i-1}} \leq r_{n_j} \prod_{i=n_j+1}^{N_j} \frac{r_i}{s_i} \leq 4 \log(\varphi) \frac{q_{N_j}}{q_{n_j}},$$

where φ is the golden ratio. In the last inequality, we use $r_{n_j} \leq r_1 = 4 \log(\varphi)$. The numbers r_{N_j} converge to 0 as $j \rightarrow \infty$, giving a negative answer to the question of Schinzel and Zassenhaus by Lemma 4.15.

For the other direction, assume the set $\left\{\frac{q_N}{q_n} \in \mathbb{R} : n, N \in \mathbb{N}, n < N\right\} \subset \mathbb{R}$ is bounded away from zero.

Claim. $r_N \geq \frac{\log(1.179)q_N}{\max\{q_1, \dots, q_{N-1}\}}$.

We admit the claim for a moment. By our assumption, there is a constant bounding all fractions $\frac{q_N}{q_n}$ with $0 < n < N$ away from zero. In particular, by the claim, there exists a constant bounding r_N strictly away from

zero for all N . This is equivalent to a positive answer to the question of Schinzel and Zassenhaus by Lemma 4.15.

We now prove the claim by induction on N .

Base case: For $N = 2$, we verify

$$r_2 = \frac{r_2}{s_2} \cdot s_2 \geq \frac{q_2}{q_1} \min\{r_1, \log(1.179)\} = \frac{\log(1.179)q_2}{\max\{q_1\}},$$

where the inequality is due to Lemma 4.14, and the equality on the right follows from $r_1 = 4 \log(\varphi)$, which is larger than $\log(1.179)$.

Inductive step: We again use Lemma 4.14. We have

$$r_{N+1} = \frac{r_{N+1}}{s_{N+1}} \cdot s_{N+1} \geq \frac{q_{N+1}}{q_N} \min\{r_N, \log(1.179)\}.$$

Using the induction hypothesis on r_N , this yields

$$\begin{aligned} r_{N+1} &\geq \frac{q_{N+1}}{q_N} \min \left\{ \frac{\log(1.179)q_N}{\max\{q_1, \dots, q_{N-1}\}}, \log(1.179) \right\} \\ &= \log(1.179)q_{N+1} \min \left\{ \frac{1}{\max\{q_1, \dots, q_{N-1}\}}, \frac{1}{q_N} \right\} \\ &= \frac{\log(1.179)q_{N+1}}{\max\{q_1, \dots, q_N\}}, \end{aligned}$$

which completes the inductive step. \square

4. Symplectic vs. antisymplectic matrices

Our Theorem 1.16 can also be formulated for definitions of reciprocity and skew-reciprocity that dictate the sign in the defining equation. For example, as is done often in the literature, we could have defined a polynomial $f \in \mathbb{Z}[t]$ of degree $2d$ to be *reciprocal* if $f(t) = t^{2d}f(t^{-1})$. Similarly, we could have defined a polynomial $f \in \mathbb{Z}[t]$ of even degree $2d$ to be *skew-reciprocal* if $f(t) = (-1)^d t^{2d} f(-t^{-1})$.

This is the route we took originally in the article [36]. It turns out that the monic polynomials of this kind are exactly the characteristic polynomials of integer symplectic and antisymplectic matrices, respectively, and we prove the exact same Theorem 1.16 in this setting [36]. This provides yet another reformulation of the question of Schinzel and Zassenhaus, this time in terms of comparing the minimal spectral radii of symplectic and antisymplectic integer matrices of given sizes. In turn, these are exactly the matrices arising from the actions induced on first homology by orientation-preserving and orientation-reversing mapping classes, respectively.

Minimal stretch factors on nonoriented surfaces

In this chapter, we discuss our results on minimal pseudo-Anosov stretch factors under the additional hypothesis of an orientable invariant foliation for nonorientable surfaces (Theorem 1.18) and for orientation-reversing maps of orientable surfaces (Theorem 1.21). We also prove Theorem 1.24 on the Galois conjugates of stretch factors in these two settings. This material is copied and adapted from our joint work with Strenner [39].

1. Construction of pseudo-Anosov maps on nonorientable surfaces

In this and the next section, we use Penner's construction to construct pseudo-Anosov mapping classes. We briefly recall Penner's construction below, stating it in a way that works both for orientable and for nonorientable surfaces. For more details, see [50, Section 4] or [62, Section 2].

In Penner's construction, we have a collection of two-sided simple closed curves $C = \{c_1, \dots, c_n\}$ that fill the surface (the complement of the curves is a union of disks and once-punctured disks), that pairwise intersect minimally, and that are *marked inconsistently*. This means that there is a small regular neighborhood $N(c_i)$ for each curve c_i and an orientation of each annulus $N(c_i)$ such that the orientation of $N(c_i)$ and $N(c_j)$ are different at each intersection whenever $i \neq j$. Penner showed that any product of the Dehn twists T_{c_i} is pseudo-Anosov assuming that

- each twist T_{c_i} is right-handed according to the orientation of $N(c_i)$,
- each twist T_{c_i} is used in the product only with positive powers,
- each twist T_{c_i} is used in the product at least once.

Note that if the surface is oriented, then the above conditions in Penner's construction say that the collection of curves is a union of two multicurves α and β , and the Dehn twists along the curves in α are all right-handed, whereas the Dehn twists along the curves in β are all left-handed with respect to the orientation of the surface.

We will present the construction of our examples as follows. First we define the rotationally symmetric graphs that will be the intersection graphs of the collections of curves. Then we describe the rotationally symmetric surfaces and curves on these surfaces whose intersection matrices realise the given graphs. Finally, we define our mapping classes as a composition of a Dehn twist and a rotation.

1.1. The graphs. Let k and n be integers of different parity such that $n \geq 3$ and $1 \leq k \leq n - 1$. Let $G_{n,k}$ be the graph whose vertices are the vertices of a regular n -gon and every vertex v is connected to the k vertices that are the farthest away from v in the cyclic order of the vertices.




FIGURE 5.1. The graphs $G_{9,2}$, $G_{6,3}$ and $G_{10,5}$.

1.2. The surfaces. For each $G_{n,k}$, we will construct a nonorientable surface $\Sigma_{n,k}$ that contains a collection of curves with intersection graph $G_{n,k}$. To construct $\Sigma_{n,k}$, start with a disk with one crosscap. By this, we mean that we cut a smaller disk out of the disk and identify the antipodal points of the boundary of the small disk. We indicate this identification with a cross inside the small disk, see Figure 5.2. The resulting surface is homeomorphic to the Möbius strip.


Next, we consider $2n$ disjoint intervals on the boundary of the disk and label the intervals with integers from 1 to n so that each label is used exactly twice. In the cyclic order, the labels are $1, s, 2, s + 1, \dots, n, s + n$, where we have $s = \frac{n+k+3}{2}$ and all labels are understood modulo n .

For each label, the corresponding two intervals are connected by a twisted strip, as on Figure 5.2.

LEMMA 5.1. *The Euler characteristic of $\Sigma_{n,k}$ is $-n$.*

PROOF. The disk with a crosscap has zero Euler characteristic (it is homeomorphic to a Möbius strip), and each attached twisted strip has contribution -1 . \square

LEMMA 5.2. *The number of boundary components of $\Sigma_{n,k}$ is $\gcd(n, k)$.*

FIGURE 5.2. The surface $\Sigma_{10,5}$ and the curve c_{10} .

PROOF. We will show that the number of boundary components of $\Sigma_{n,k}$ is the same as the number of orbits of the dynamical system $x \mapsto x + n - k$ in the group $\mathbb{Z}/2n\mathbb{Z}$. The number of such orbits is $\gcd(n - k, 2n) = \gcd(k, n)$, since $n - k$ is odd.

To prove our claim, we identify $\mathbb{Z}/2n\mathbb{Z}$ with the $2n$ intervals in the cyclic order. We claim that the right endpoint of the interval at position i lies on the same boundary component as the right endpoint of the interval at position $i + n - k$. One can see this by induction. In the case $k = n - 1$, the cyclic order of labels is $1, 1, \dots, n, n$, so the twisted strips identify the right endpoint of every interval with the right endpoint of the next interval. When $k = n - 3$, the cyclic order is $1, n, 2, 1, \dots, n, n - 1$, in which case every third right endpoint is on the same boundary component, and so on. \square

PROPOSITION 5.3. *The surface $\Sigma_{n,k}$ is homeomorphic to the nonorientable surface of genus $n - \gcd(k, n) + 2$ with $\gcd(k, n)$ boundary components.*

PROOF. The Euler characteristic of the nonorientable surface of genus g with b boundary components is $2 - g - b$. By Lemma 5.1 and Lemma 5.2, we obtain the equation $2 - g - \gcd(k, n) = -n$. Rearranging, we finally obtain $g = n - \gcd(k, n) + 2$. \square

1.3. The curves. We construct a two-sided curve c_i for each of the labels $i = 1, \dots, n$ as follows. Each curve consists of two parts. One part of each curve is the core of the strip corresponding to the label. The other part is an arc inside the disk that passes through the crosscap and connects the corresponding two intervals. The curve c_{10} is shown on Figure 5.2.

Note that every pair of curves intersects either once or not at all. The curves c_i and c_j are disjoint if and only if the two i labels and the two j labels *link* in the cyclic order. In other words, if the two i labels separate the two j labels.

LEMMA 5.4. *The intersection graph of the curves c_i on $\Sigma_{n,k}$ is $G_{n,k}$.*

PROOF. We prove the lemma by induction. If $k = n - 1$, then $s = 1$, so the cyclic order is $1, 1, 2, 2, \dots, n, n$. Since the no two labels link, all pairs of curves intersect and the intersection graph is the complete graph $G_{n,n-1}$.

Now suppose k is decreased by 2. Then s is decreased by 1, and we obtain the cyclic order $1, n, 2, 1, 3, 2, \dots, n, n - 1$. As a consequence, 1 becomes linked with 2 and n . Hence the intersection graph is indeed $G_{n,k}$.

It is easy to see that every time k is decreased by two each label is linked with two more labels, hence the intersection graph is always $G_{n,k}$. \square

LEMMA 5.5. *The curves c_i can be marked so that all intersections are inconsistent.*

PROOF. Choose markings for the c_i which are invariant under the rotational symmetry, see Figure 5.3. The marking of the curves is indicated by the coloring as follows. Consider the orientable surface obtained by removing the crosscap and cutting the strips attached to the disk in the middle. Choose an orientation of this surface. Then color the arcs composing the curves using red and blue depending on whether the orientation of the neighborhood of the curve matches the orientation of the surface or not. Note that the color of a curve changes when it goes through the crosscap or the middle of a strip.

Since blue and red meets at every intersection, the marking is inconsistent. \square




FIGURE 5.3. A collection of filling inconsistently marked curves.

1.4. The mapping classes. Denote by r the rotation of $\Sigma_{n,k}$ by one click in the clockwise direction. Define the mapping class

$$f_{n,k} = r \circ T_{c_1}$$

where T_{c_1} is a Dehn twist about the curve c_1 . (There are two possible directions for the Dehn twist, but either choice works for our purposes.) Note that

$$f_{n,k}^n = T_{c_n} \circ \cdots \circ T_{c_1},$$

so $f_{n,k}^n$ arises from Penner's construction. In particular, $f_{n,k}^n$ is pseudo-Anosov and so is $f_{n,k}$.

We remark that for $k = n - 1$, the mapping class $f_{n,k}$ coincides with the nonorientable Arnoux-Yoccoz mapping class h_{n-1} , described as a product of a Dehn twist and a finite order mapping class in joint work with Strenner in [40].

PROPOSITION 5.6. *The stretch factor of $f_{n,k}$ is the largest root of*

$$t^n - t^{n-r} - t^{n-r-1} - \dots - t^{r+1} - t^r - 1,$$

where $r = \frac{n-k+1}{2}$.

PROOF. To compute the stretch factor, we use Penner's approach in the section titled "An upper bound by example" in [51]. Penner constructed an invariant bigon track by smoothing out the intersections of the curves c_i . Each c_i defines a characteristic measure μ_i on this bigon track, defined by assigning 1 to the branches traversed by c_i and zero to the rest. The cone generated by the μ_i is invariant under both T_{c_1} and r , hence it contains the unstable foliation, and the stretch factor is given by the largest eigenvalue of the action of $r \circ T_{c_1}$ on this cone. The rotation r acts by a permutation matrix and the matrix corresponding to T_{c_1} is the sum of the identity matrix and matrix obtained by the intersection matrix $i(C, C)$ by zeroing out all rows except the first row. The product of these two matrices takes the following form:

$$\begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 \end{pmatrix}$$

This particular matrix belongs to $f_{10,5}$.

This matrix is the companion matrix of the polynomial in the statement of the proposition. Hence the characteristic polynomial of this matrix is indeed that polynomial. \square

Our next goal is to determine the singularity structure of the mapping classes $f_{n,k}$. For this, first we need a lemma.

Consider the complementary regions of the curves $\{c_1, \dots, c_n\}$. There are two types of regions depending on whether a region contains a boundary component of $\Sigma_{n,k}$ (type 1) or not (type 2). A region of type 1 is an annulus that is bounded by a boundary component β of $\Sigma_{n,k}$ on one side and by a polygonal path consisting of arcs of the curves c_i on the other side. The shaded region on Figure 5.3 illustrates a region of type 1.

LEMMA 5.7. *The length of these polygonal paths is $\frac{4n}{\gcd(k,n)}$.*

PROOF. This follows from the observation that every point in the orbit in $\mathbb{Z}/2n\mathbb{Z}$ corresponding to the boundary component β (see the proof of Lemma 5.2) has two associated arcs. Since the number of orbits is $\gcd(k, n)$, the length of each orbit is $\frac{2n}{\gcd(k,n)}$, and hence the length of the polygonal path is twice this quantity. \square

PROPOSITION 5.8. *The pseudo-Anosov mapping class $f_{n,k}$ has $\gcd(k, n)$ singularities, one for each boundary component. The number of prongs of each singularity is $\frac{2n}{\gcd(k,n)}$.*

PROOF. Each complementary region of the curves $\{c_1, \dots, c_n\}$ contains either one singularity or none. The number of prongs of a singularity equals the number of cusps of the bigon track obtained by the smoothing process that are contained in the same region as the singularity. If the number of cusps is 2, then the region does not contain a singularity. If the number of cusps is $k > 2$, then it contains a k -pronged singularity.

Regions of type 2 are rectangles (bounded by four subarcs of the curves c_i), and hence contain two cusps. So they do not correspond to singularities.

The lengths of the polygonal paths bounding regions of type 1 are $\frac{4n}{\gcd(k,n)}$ by Lemma 5.7, so the number of cusps in these regions is $\frac{2n}{\gcd(k,n)}$. Therefore the singularities have that many prongs. By Lemma 5.2, the number of such regions is $\gcd(k, n)$, so that is also the number of the singularities. \square

As a corollary of Propositions 5.6, 5.8 and 5.3, we have the following.

COROLLARY 5.9. *There exist pseudo-Anosov mapping classes with an orientable invariant foliation on the surfaces N_g with the data below. All of these examples belong to the family $f_{n,k}$ for the n and k shown in the table.*

| g | n | k | $\lambda(f_{n,k})$ | <i>minimal polynomial</i> | <i>singularity type</i> |
|--------|-----|-----|--------------------|--|-------------------------|
| 4^* | 3 | 2 | 1.83929 | $t^3 - t^2 - t - 1$ | (6) |
| 5^* | 6 | 3 | 1.51288 | $t^4 - t^3 - t^2 + t - 1$ | (4,4,4) |
| 6^* | 5 | 2 | 1.42911 | $t^5 - t^3 - t^2 - 1$ | (10) |
| 7^* | 10 | 5 | 1.42198 | $t^6 - t^5 - t^3 + t - 1$ | (4,4,4,4,4) |
| 8^* | 7 | 2 | 1.28845 | $t^7 - t^4 - t^3 - 1$ | (14) |
| 9 | 8 | 3 | 1.35680 | $t^8 - t^5 - t^4 - t^3 - 1$ | (16) |
| 10^* | 9 | 2 | 1.21728 | $t^9 - t^5 - t^4 - 1$ | (18) |
| 11 | 12 | 3 | 1.22262 | $\frac{t^{12}-t^7-t^6-t^5-1}{t^2+t+1}$ | (8,8,8) |
| 12^* | 11 | 2 | 1.17429 | $t^{11} - t^6 - t^5 - 1$ | (22) |
| 13 | 22 | 11 | 1.27635 | $t^{12} - t^{11} - t^6 + t - 1$ | (4 ¹¹) |
| 14^* | 13 | 2 | 1.14551 | $t^{13} - t^7 - t^6 - 1$ | (26) |
| 15 | 14 | 3 | 1.18750 | $t^{14} - t^8 - t^7 - t^6 - 1$ | (28) |
| 16^* | 15 | 2 | 1.12488 | $t^{17} - t^9 - t^8 - 1$ | (30) |
| 17 | 18 | 3 | 1.14259 | $\frac{t^{18}-t^{10}-t^9-t^8-1}{t^2+t+1}$ | (12,12,12) |
| 18^* | 17 | 2 | 1.10938 | $t^{19} - t^{10} - t^9 - 1$ | (34) |
| 19 | 18 | 5 | 1.20514 | $t^{18} - t^{11} - t^{10} - t^9 - t^8 - t^7 - 1$ | (36) |
| 20^* | 19 | 2 | 1.09730 | $t^{23} - t^{12} - t^{11} - 1$ | (38) |

(4¹¹ means that there are 11 singularities with 4 prongs.)


In each genus, the family $f_{n,k}$ contains several examples. In the table above, we have listed only the example with the smallest stretch factor. In the starred cases, we will be able to certify that the given stretch factors are not only minimal in the family $f_{n,k}$ but among all pseudo-Anosov maps with an orientable invariant foliation.

2. Orientation-reversing pseudo-Anosov mapping classes on odd genus surfaces

In this section, we construct an orientation-reversing pseudo-Anosov mapping class with small stretch factor on every odd genus orientable surface. The construction is analogous to the construction in the previous section, but simpler. As in the previous section, we separate the construction of the surfaces, the curves and finally the mapping classes.

2.1. The surfaces. For every $k \geq 2$, consider the surface Σ_k obtained by chaining together $2k$ annuli in a cycle as on Figure 5.4.

PROPOSITION 5.10. *The number of boundary components of Σ_k is 4 if k is even and 2 if k is odd.*

FIGURE 5.4. The surface Σ_k .

PROOF. The boundary of Σ_k is composed of $8k$ arcs, 4 arcs for each annulus. Our goal is to determine which of them belong to the same boundary component.

Denote by r the rotation of Σ_k by one click. By tracing the boundary, one can see that every boundary point x lies on the same boundary component as $r^4(x)$. Moreover, the path between x and $r^4(x)$ traverses each of the 4 types of arcs exactly once. Therefore it suffices to pick any boundary point x and determine into how many equivalence classes the set $\{x, r(x), \dots, r^{2k-1}(x)\}$ falls apart. The number of such equivalence classes is 4 if k is even and 2 if k is odd. \square

PROPOSITION 5.11. *The surface Σ_k is homeomorphic to $S_{k-1,4}$ if k is even and $S_{k,2}$ if k is odd.*

PROOF. We have $\chi(\Sigma_k) = 2k$. From the equation $\chi = 2 - 2g - b$, where g is the genus and b is the number of boundary components, it follows that $g = k + 1 - \frac{b}{2}$. The statement now follows from Proposition 5.10. \square

As a consequence, the construction only produces odd genus examples.

2.2. The curves. From now on, suppose that k is even. Consider the set $C = \{c_1, \dots, c_{2k}\}$ of core curves of the $2k$ annuli. Our numbering

will differ from the standard cyclic numbering; we will explain this shortly. As in Section 1.3, any rotationally symmetric marking of the curves is an inconsistent marking.

The intersection graph of C is a cycle of length $2k$. We draw this cycle as on Figure 5.5: the vertices are the vertices of a regular polygon and every vertex is connected to the two vertices that are the second furthest in the cyclic order. We number the curves according to the cyclic orientation induced by this picture.




FIGURE 5.5. Our unusual way of numbering the curves. For example, the curve c_1 intersects c_k and c_{k+2} , not c_2 and c_{2k} .

2.3. The mapping classes. Denote by r the rotation of Σ_k (see Figure 5.4) by one click in the clockwise direction. Since c_i and $r(c_i)$ intersect for all i , the rotation r induces a rotation of the cycle on Figure 5.5 by $k-1$ clicks. So r^{k-1} rotates the cycle by $(k-1)^2 = k^2 - 2k + 1$ clicks, which is congruent to 1 modulo $2k$ if k is even. Therefore r^{k-1} induces rotating the cycle on Figure 5.5 by one click (in the clockwise direction, assuming that we have chosen the numbering of the curves accordingly). In particular, we have $r^{k-1}(c_{i+1}) = c_i$.

We are now ready to define the mapping class:

$$\psi_k = r^{k-1} \circ T_{c_1}.$$

Note that

$$\psi_k^{2k} = T_{c_{2k}} \circ \cdots \circ T_{c_1},$$

so ψ_k^{2k} arises from Penner's construction. In particular, ψ_k^{2k} is pseudo-Anosov and so is ψ_k . Note that while ψ_k^{2k} is orientation-preserving, ψ_k is orientation-reversing. This follows from the fact that T_{c_1} is orientation-preserving and r is orientation-reversing.

PROPOSITION 5.12. *The stretch factor of ψ_k is the largest root of*

$$t^{2k} - t^{k+1} - t^{k-1} - 1.$$

PROOF. The proof is similar to the proof of Proposition 5.6. We have

$$i(C, C) = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 \end{pmatrix}, \quad M = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \end{pmatrix}$$

where M is the matrix of the action of ψ_k on the cone of measures (the product of a permutation matrix and the sum of the identity matrix and the first row of $i(C, C)$). The matrices above illustrate the case $k = 4$. The matrix M is the companion matrix of the polynomial in the proposition. \square

PROPOSITION 5.13. *The pseudo-Anosov mapping class ψ_k has four k -pronged singularities.*

PROOF. By Proposition 5.10 and its proof, each of the four boundary components of Σ_k consists of $2k$ arcs if k is even. There is a prong for every second corner of the boundary path, therefore there are k prongs for each singularity. \square

COROLLARY 5.14. *There exist orientation-reversing pseudo-Anosov mapping classes with orientable invariant foliations on the surfaces S_g with the data below. All of these examples belong to the family ψ_k for the k shown in the table.*

| g | k | $\lambda(\psi_k)$ | largest root of | singularity type |
|-----|-----|-------------------|---|------------------|
| 1 | 2 | 1.61803 | $t^2 - t - 1 = \frac{t^4 - t^3 - t - 1}{t^2 + 1}$ | no singularities |
| 3 | 4 | 1.25207 | $t^8 - t^5 - t^3 - 1$ | (4,4,4,4) |
| 5 | 6 | 1.15973 | $t^{12} - t^7 - t^5 - 1$ | (6,6,6,6) |
| 7 | 8 | 1.11707 | $t^{16} - t^9 - t^7 - 1$ | (8,8,8,8) |
| 9 | 10 | 1.09244 | $t^{20} - t^{11} - t^9 - 1$ | (10,10,10,10) |
| 11 | 12 | 1.07638 | $t^{24} - t^{13} - t^{11} - 1$ | (12,12,12,12) |

PROOF. The statement follows from Propositions 5.12, 5.13 and 5.11. The reason we have no singularities in the genus 1 case is that by Proposition 5.13 the “singularities” have two prongs, so they are not actually singularities. \square

We remark that in the genus 1, there is an alternative, simpler construction that yields the same stretch factor. Consider the matrix

$$M = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$$

with determinant -1 . The corresponding linear map $\mathbb{R}^2 \rightarrow \mathbb{R}^2$ maps \mathbb{Z}^2 to \mathbb{Z}^2 , hence it descends to an Anosov diffeomorphism f of the torus $\mathbb{R}^2/\mathbb{Z}^2$, as we described in the teaser at the very beginning of this habilitation thesis. Its stretch factor is the largest root of the $t^2 - t - 1$, the characteristic polynomial of M .

3. Restrictions on polynomials

Pseudo-Anosov stretch factors are roots of integral polynomials. The properties of these integral polynomials are similar, but slightly different depending on whether a pseudo-Anosov mapping class is an orientation-preserving or orientation-reversing mapping class on an orientable surface or a mapping class on a nonorientable surface. In this section, we discuss these properties for nonorientable surfaces and orientation-reversing mapping classes.

Recall that a polynomial $p(t)$ of degree n is reciprocal if $p(t) = \pm t^n p(t^{-1})$, in other words, when its coefficients are the same in reverse order up to sign. Analogously, recall that $p(t)$ is skew-reciprocal if $p(t) = \pm t^n p(-t^{-1})$.

PROPOSITION 5.15. *Let $\psi : N_g \rightarrow N_g$ be a pseudo-Anosov map with a transversely orientable invariant foliation on the closed nonorientable surface N_g of genus g . Then its stretch factor λ is a root of a (not necessarily irreducible) polynomial $p(t) \in \mathbb{Z}[t]$ with the following properties:*

- (1) $\deg(p) = g - 1$.
- (2) $p(t)$ is monic and its constant coefficient is ± 1 .
- (3) The absolute values of the roots of $p(t)$ other than λ lie in the open interval (λ^{-1}, λ) . In particular, $p(t)$ is not reciprocal or skew-reciprocal.
- (4) $p(t)$ is reciprocal mod 2.

PROOF. Note that exactly one of the stable and unstable foliations is transversely orientable (otherwise the surface itself would be orientable). We will assume that it is the stable foliation, otherwise we replace ψ by its inverse.

Consider the action $\psi^* : H^1(N_g; \mathbb{R}) \rightarrow H^1(N_g; \mathbb{R})$ defined by pullback on cohomology with real coefficients. Since the stable foliation is transversely orientable, it is represented by a closed real 1-form, that is, an element of $H^1(N_g; \mathbb{R})$. The stable foliation \mathcal{F}^s is the one whose leaves are contracting and hence the surface is expanding in the transverse direction. Therefore the measure of a transverse arc in the pullback $\psi^*(\mathcal{F}^s)$ is λ times its measure in \mathcal{F}^s . Hence \mathcal{F}^s is an eigenvector of the map ψ^* with eigenvalue λ or $-\lambda$.

Let $p(x)$ be the characteristic polynomial of ψ^* . Note that since we have $\dim(H^1(N_g, \mathbb{R})) = g - 1$, we also have $\deg(p) = g - 1$, proving (1).

The polynomial $p(t)$ has integral coefficients, since ψ^* restricts to an action $H^1(N_g; \mathbb{Z}) \rightarrow H^1(N_g; \mathbb{Z})$. This restriction is invertible, since the action of $\text{Mod}(N_g)$ on $H^1(N_g; \mathbb{Z})$ is a group representation, so the determinant of ψ^* is ± 1 . Therefore the constant coefficient of $p(t)$ is ± 1 . Also, as a characteristic polynomial, $p(t)$ is monic. This proves (2).

It is a standard fact from the theory of orientation-preserving pseudo-Anosov mapping classes on orientable surfaces that stretch factors are strictly maximal among their Galois conjugates. Moreover, in case the invariant foliations are orientable, the spectral radius of the action induced on the first homology equals the stretch factor, and every other eigenvalue is strictly smaller, see, for example, [46, Theorem 5.3 (1)]. Applying this result to the orientation-preserving lift $\tilde{\psi} : S_{g-1} \rightarrow S_{g-1}$ of ψ to the orientable double cover S_{g-1} of N_g , we obtain that any root λ' of $p(t)$ other than $\pm\lambda$ satisfies $|\lambda'| < |\lambda|$. Applying the same theorem for $\tilde{\psi}^{-1}$, we conclude that any root λ' of $p(t)$ other than $\pm\lambda^{-1}$ satisfies $|\lambda'^{-1}| < |\lambda|$. Therefore absolute values of the roots of $p(t)$ other than $\pm\lambda$ and possibly $\pm\lambda^{-1}$ lie in the open interval (λ^{-1}, λ) . However, it was shown in the proof of [63, Proposition 2.3] that if λ or $-\lambda$ is a root of $p(t)$, then λ^{-1} and $-\lambda^{-1}$ cannot be roots of $p(t)$, hence mentioning the edge case $\pm\lambda^{-1}$ in the previous sentence is not necessary.

If $p(t)$ was reciprocal, then λ and λ^{-1} or $-\lambda$ and $-\lambda^{-1}$ would have to be roots. If it was skew-reciprocal, then λ and $-\lambda^{-1}$ or $-\lambda$ and λ^{-1} would

have to be roots. As we have just shown, these scenarios are impossible, because $\pm\lambda^{-1}$ is not a root of $p(t)$. This proves (3).

The fact that $p(t)$ is reciprocal mod 2 was shown in [63, Proposition 4.2]. This justifies (4).

Finally, notice that we have not guaranteed that λ is a root of $p(t)$ —we have only shown that either λ or $-\lambda$ is a root. If it is $-\lambda$, then the polynomial $p(-t)$ or $-p(-t)$ satisfies all the required properties. \square

We call a $2n \times 2n$ matrix A *anti-symplectic* if the corresponding linear transformation sends the standard symplectic form on \mathbb{R}^{2n} to its negative. Formally, this can be written as

$$AJA^T = -J$$

where $J = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix}$ and I is the $n \times n$ identity matrix.

PROPOSITION 5.16. *The characteristic polynomial $p(t)$ of a $2n \times 2n$ anti-symplectic matrix is skew-reciprocal.*

PROOF. Let A be an $2n \times 2n$ anti-symplectic matrix. Since A^2 is symplectic, we have $\det(A^2) = 1$ and $\det(A) = \pm 1$. Since $\det(J) = 1$, we have

$$\begin{aligned} p(t) &= \det(A - tI) = \det(AJ - tJ) = \det(AJ + tAJA^T) \\ &= \det(A) \det(J) \det(I + tA^T) = \pm \det(I + tA) \\ &= \pm t^{2n} \det(A + t^{-1}I) = \pm t^{2n} p(-t^{-1}), \end{aligned}$$

hence $p(t)$ is skew-reciprocal. \square

The proof above is a straightforward modification of the standard proof that the characteristic polynomials of symplectic matrices are reciprocal.

PROPOSITION 5.17. *Let $\psi : S_g \rightarrow S_g$ be an orientation-reversing pseudo-Anosov map with transversely orientable invariant foliations. Then its stretch factor λ is a root of a (not necessarily irreducible) polynomial $p(t) \in \mathbb{Z}[t]$ with the following properties:*

- (1) $\deg(p) = 2g$.
- (2) $p(t)$ is monic and its constant coefficient is $(-1)^g$.
- (3) $p(t) = (-1)^g t^{2g} p(-t^{-1})$.
- (4) $p(-\lambda^{-1}) = 0$.
- (5) *The absolute values of the roots of $p(t)$ other than λ and $-\lambda^{-1}$ lie in the open interval (λ^{-1}, λ) .*

PROOF. Let $p(t)$ be the characteristic polynomial of the induced action $\psi_* : H_1(S_g) \rightarrow H_1(S_g)$. Clearly, (1) holds. Similarly to the proof of Proposition 5.15, we obtain (5) by a reduction to the known statement for orientation-preserving pseudo-Anosov mapping classes. This time, we directly obtain (5) by applying [46, Theorem 5.3] to the square of ψ .

An orientation-reversing homeomorphism sends the intersection form on $H_1(S_g)$ to its negative. Proposition 5.16 implies that $p(t) = \pm t^{2g}p(-t^{-1})$. To decide which sign is right, we only need to compute the sign of the constant coefficient of $p(t)$. If the constant coefficient is 1, then the sign is positive. If the constant coefficient is -1 , then the sign is negative. To put this in another way, we have

$$(4) \quad p(t) = p(0)t^{2g}p(-t^{-1}).$$

For orientation-preserving homeomorphisms, the action on homology is symplectic, hence its determinant is $+1$. It follows that, for fixed g , the determinant is either $+1$ for all orientation-reversing homeomorphisms of S_g or -1 for all orientation-reversing homeomorphisms of S_g . It is sufficient to check only one homeomorphism to decide which one. For example, consider the reflection $i : S_g \rightarrow S_g$ about the plane containing the curves b_i on Figure 5.6.




FIGURE 5.6. The standard homology basis for S_2 .

The curves $\{a_1, \dots, a_g, b_1, \dots, b_g\}$ form a homology basis. We have that $i(b_i) = b_i$ and $i(a_i) = -a_i$ for all i . So the matrix of A is a diagonal matrix whose diagonal contains 1s and -1 s, g of each. Hence the determinant is $(-1)^g$ and we have shown (2).

Item (3) follows from equation (4) and the fact $p(0) = (-1)^g$ we have just shown.

Either λ or $-\lambda$ is a root of $p(t)$. If it is $-\lambda$, we may replace $p(t)$ with $p(-t)$; the previously proven properties remain true. Finally, (4) follows from (3) by setting $t = \lambda$. \square

We are now ready to prove Theorem 1.24. We emphasise that, unlike in the previous propositions, in this theorem we are not assuming that the surface is closed or that the pseudo-Anosov mapping class has a transversely orientable invariant foliation.

PROOF OF THEOREM 1.24. An irreducible polynomial $p(t) \in \mathbb{Z}[t]$ with a (complex) root α on the unit circle is reciprocal. This holds because then also α^{-1} is a root of $p(t)$, therefore α is a root of the polynomial $t^d p(t^{-1})$, where d is the degree of $p(t)$. But the minimal polynomial is unique up constant factor, so $t^d p(t^{-1}) = \pm p(t)$. Hence $p(t)$ is indeed reciprocal. So if a stretch factor λ has a Galois conjugate on the unit circle, then the minimal polynomial of λ is reciprocal and λ^{-1} is also a root of the minimal polynomial.

However, by [63, Proposition 2.3], λ and λ^{-1} are not Galois conjugates if λ is a stretch factor of a pseudo-Anosov map (possibly with no orientable invariant foliations) on a nonorientable surface (possibly with punctures). This completes the proof in the case when the pseudo-Anosov map is supported on a nonorientable surface.

We now prove the orientation-reversing case. If our surface is closed, then, by Proposition 5.17, λ and λ^{-1} are not Galois conjugates if λ is a stretch factor of an orientation-reversing pseudo-Anosov map with orientable invariant foliations. If the foliations are not orientable, we can lift the map to the orientation double cover of the foliations to obtain an orientation-reversing pseudo-Anosov map with orientable invariant foliations and with the same stretch factor. Therefore λ and λ^{-1} are not Galois conjugates in this case, either.

If our surface has punctures, then we can fill in the punctures after making the foliations orientable to obtain a pseudo-Anosov map with the same stretch factor on a closed surface, reducing to the closed case discussed in the previous paragraph. This completes the proof in the case when the pseudo-Anosov map is orientation-reversing. \square

4. Elimination of polynomials

In this section we first prove bounds on the sum of k th powers of roots of a polynomial when the absolute values of the roots are bounded by some positive real number $r > 1$. These bounds are improved versions of Lemma A.1. of [29], using the special properties of the polynomials in Propositions 5.15 and 5.17.

Then we describe how we use this lemma and Propositions 5.15 and 5.17 in order to systematically narrow down the set of possible minimal polynomials of the minimal stretch factors.

4.1. Power sum bounds. We begin by proving two elementary lemmas.

LEMMA 5.18. *Suppose $r > 1$ and $r^{-1} \leq a_1, \dots, a_d \leq r$ are positive real numbers such that $a_1 \cdots a_d = 1$. Then*

$$\sum_{i=1}^d a_i \leq \begin{cases} n(r + r^{-1}) & \text{if } d = 2n \text{ is even} \\ n(r + r^{-1}) + 1 & \text{if } d = 2n + 1 \text{ is odd.} \end{cases}$$

PROOF. The function $x \mapsto x + x^{-1}$ is increasing on the interval $x \geq 1$. So if there are $i \neq j$ so that $r^{-1} < a_i, a_j < r$, then we can increase the sum by moving a_i and a_j away from each other by keeping their product unchanged, until at least one of them is r^{-1} or r . After every such operation, the number of a_i that are equal to r^{-1} or r increases. So eventually we get to a point where at most one a_i is not r^{-1} or r . When $d = 2n$, no such a_i can exist, and exactly half of the a_i equal r , the other half r^{-1} , otherwise their product would not be 1. When $d = 2n + 1$, exactly one such a_i exists, it equals 1 and exactly half of the remaining a_i equal r , the other half r^{-1} . \square

LEMMA 5.19. *Suppose $r > 1$ and a_1, \dots, a_d are positive real numbers such that $r^{-1} \leq a_1 \leq \dots \leq a_d \leq r$ and $a_1 \cdots a_d = 1$ and $a_1 \geq a_d^{-1}$. Then*

$$a_d - \sum_{i=1}^{d-1} a_i \geq \begin{cases} \min\{2 - 2n, -(n-2)r - nr^{-1}\} & \text{if } d = 2n \text{ is odd.} \\ \min\{1 - 2n, -(n-2)r - 1 - nr^{-1}\} & \text{if } d = 2n + 1 \text{ is odd.} \end{cases}$$

Moreover, the inequalities are strict if $a_1 > a_d^{-1}$.

PROOF. Similarly to the proof of Lemma 5.18, our approach is to change the numbers a_1, \dots, a_{d-1} to increase $\sum_{i=1}^{d-1} a_i$ as much as possible while keeping the hypotheses true. Whenever there are $i \neq j$ such that we

have $a_d^{-1} < a_i, a_j < a_d$, we push a_i and a_j apart until at least one of them equals a_d^{-1} or a_d . The end result is the same as before, so we have

$$a_d - \sum_{i=1}^{d-1} a_i \geq a_d - ((n-1)a_d + na_d^{-1}) = -(n-2)a_d - na_d^{-1}.$$

when $d = 2n$ and

$$a_d - \sum_{i=1}^{d-1} a_i \geq a_d - ((n-1)a_d + 1 + na_d^{-1}) = -(n-2)a_d - 1 - na_d^{-1}$$

when $d = 2n + 1$. Since the function $x \mapsto -(n-2)x - nx^{-1}$ is concave, its minimum on the interval $[1, r]$ is taken at one of the endpoints.

The inequalities in the case $a_1 > a_d^{-1}$ are strict, since in the optimal distribution there has to be an a_i that takes the value a_d^{-1} . \square

Now we apply Lemmas 5.18 and 5.19 for roots of polynomials.

COROLLARY 5.20. *Suppose $P(t)$ is a monic polynomial of degree d with constant coefficient ± 1 . Let z_1, \dots, z_d be the roots of $P(t)$ and let*

$$p_k = z_1^k + \dots + z_d^k$$

the k th power sum of the roots.

Suppose there is a root $\lambda > 1$ such that all the other roots have absolute values in the interval $[\lambda^{-1}, \lambda]$. For any $r > \lambda$, we have

$$\min\{2 - 2n, -(n-2)r^k - nr^{-k}\} \leq p_k \leq n(r^k + r^{-k})$$

if $d = 2n$ is even and

$$\min\{1 - 2n, -(n-2)r^k - 1 - nr^{-k}\} \leq p_k \leq n(r^k + r^{-k}) + 1$$

if $d = 2n + 1$ is odd.

Moreover, strict inequality holds in the lower bound when no eigenvalue equals λ^{-1} .

PROOF. Let z_1, \dots, z_d be the roots of $P(t)$ and let $a_i = |z_i|^k$ for every i . Note that $a_1 \cdots a_d = 1$ and $r^{-k} \leq a_1, \dots, a_d \leq r^k$. Assuming $a_1 \leq \dots \leq a_d$, we have $a_d = \lambda^k$. Since

$$a_d - \sum_{i=1}^{d-1} a_i \leq p_k = z_1^k + \dots + z_{d-1}^k + \lambda^k \leq \sum_{i=1}^d a_i,$$

the bounds follow from Lemma 5.18 and Lemma 5.19 \square

4.2. Newton's formulas. In this section, we recall Newton's formulas that relate power sums of the roots to the coefficients of the polynomial.

We will use the notation

$$P(t) = t^d - c_1 t^{d-1} - \cdots - c_{d-1} t \pm 1$$

for the coefficients of monic polynomials of degree d . As in the statement of Corollary 5.20, we denote by p_k the k th power sum of the roots of $P(t)$.

Newton's formulas relating power sums and symmetric polynomials can be stated either as

$$(5) \quad p_k = -c_1 p_{k-1} - c_2 p_{k-2} - \cdots - c_{k-1} p_1 - k c_k$$

or as

$$(6) \quad c_k = \frac{-c_1 p_{k-1} - c_2 p_{k-2} - \cdots - c_{k-1} p_1 - p_k}{k}$$

for all $1 \leq k \leq d-1$.

As Lanneau and Thiffeault point out in Section A.1 of [29], is it more computationally efficient to bound the power sums p_i and using Newton's formulas to compute the coefficients c_k from the p_i than to bound the coefficients directly. This is because many scenarios get ruled out just because the numerator in equation (6) is not divisible by k .

4.3. The polynomial elimination algorithm. We give a lower bound on the minimal stretch factor $\delta^+(N_g)$ by a systematic elimination of polynomials. We describe this process below. In order to illustrate the effect of each step in the algorithm, we give the number of candidate polynomials left after each step when $g = 12$ (when the degree is 11).

ALGORITHM 5.21. *Let $d \geq 4$, $g = d-1$ and $r > 1$ such that $\delta^+(N_g) < r$. Perform the following steps in order to obtain a small set of polynomials of degree g that include one polynomial whose root is $\delta^+(N_g)$:*

- (1) *Compute the possible values of p_1, \dots, p_{d-1} using the bounds given by Corollary 5.20. For $d = 11$, the total number of combinations is $20 \cdot 20 \cdot 21 \cdot 23 \cdot 24 \cdot 27 \cdot 30 \cdot 34 \cdot 38 \cdot 43 = 10,641,541,131,648,000$.*
- (2) *Compute the coefficients c_1, \dots, c_{d-1} using Equation 6, keeping only the cases when all c_i are integers. $57,643,952$ cases remain.*
- (3) *Discard all cases where the polynomial is not reciprocal mod 2. $1,808,922$ cases remain.*

- (4) Try ± 1 for the constant coefficient. We now doubled the number of cases to 3,617,844.
- (5) Consider the reciprocal polynomial $P^*(t) = \pm t^d P(t^{-1})$ (with the sign chosen so that the polynomial is monic), and use Equation 5 to compute the power sums p_1^*, \dots, p_{d-1}^* of this polynomial from the reversed sequence $\pm c_{d-1}, \dots, \pm c_1$ of coefficients, where the signs here depend on the sign chosen in the previous step. Discard the cases that do not satisfy the bounds of Corollary 5.20. 5075 cases remain.
- (6) Test the remaining polynomials by Newton's method for finding roots. Start with the upper bound for the Perron root. Since the polynomial is increasing and convex in $[\lambda, \infty)$, we should get a decreasing sequence of x -values larger than λ . Discard the cases when this fails. Stop when two consecutive x -values are very close to each other. 421 cases remain.
- (7) Discard the polynomials where the largest eigenvalue in absolute value is not real. 86 cases remain.
- (8) Discard the cases where the multiplicity of the largest eigenvalue is larger than 1. 54 cases remain.
- (9) Discard the cases when there is a root with absolute value less than or equal to λ^{-1} . 33 cases remain.
- (10) Discard the cases where the largest eigenvalue is larger than our upper bound. 1 case remains.

In practice, the first three steps are implemented in a more sophisticated way. Our computers cannot handle as many as 10,641,541,131,648,000 cases, so the implementation does not actually consider all those combinations. It first chooses a value for p_1 and sets $c_1 = -p_1$ by (6). Then it has to choose a value for p_2 of the same parity as p_1 , since $c_2 = -\frac{c_1 p_1 - p_2}{2}$. Similarly, the value chosen for p_3 is then determined mod 3, therefore a huge number of combinations for the p_i are never considered. Also, once more than half of the c_i are computed, we obtain additional constraints on the p_i , since our polynomial has to be reciprocal mod 2. Since these divisibility checks are done early, and not after the whole polynomial is constructed, a huge number of cases gets eliminated early.

The idea of bounding the coefficients using power sums as in steps (1) and (2) instead of using symmetric polynomials to express the coefficients in terms of the roots is due to Lanneau and Thiffeault [29]. As the numbers

suggest, these are the steps that are responsible for bringing the size of the set of possible polynomials down from an astronomical size to one that is approachable by computers.

Step (3) is special to nonorientable surfaces and is also crucial. Without this step, not only would the searching process be much slower, but there are quite a few polynomials that pass all the other tests but this is the only step that eliminates them. Perhaps this step is the main reason for why we do not need Lefschetz number tests unlike Lanneau and Thiffeault in the orientable case.

Step (5) is also special to nonorientable surfaces, since in the orientable case the polynomials are reciprocal, so the reciprocal polynomial does not contain any additional information. This was one of the last tests we added, and this reduced the running time of the algorithm for the $d = 11$ case from several hours a few minutes. The reason this works so effectively is that in most of the coefficient sequences at this point, the last few coefficients (c_{d-1} , c_{d-2} , etc.) are much bigger than the required bounds.

Step (6) is another computationally inexpensive test that quickly eliminates a large fraction of the polynomials. The idea of this test is also due to Lanneau and Thiffeault.

The most computationally expensive part is computing the roots. We only compute the roots after Step (6), only in 421 cases. So in terms of total time, actually steps (1)–(6) take more than 99% of the running time.

We use a very similar algorithm in order to give a lower bound for the minimal stretch factor $\delta_{rev}^+(S_g)$ among orientation-reversing pseudo-Anosov maps. The difference is that we use the properties from Proposition 5.17 instead of the ones from Proposition 5.15. Our implementation of these algorithms can be found at

<https://github.com/b5strbal/polynomial-filtering>.

4.4. Minimal stretch factors. We are now ready to single out the minimal stretch factor $\delta^+(N_g)$ among pseudo-Anosov homeomorphisms with an orientable invariant foliation for certain nonorientable closed surfaces N_g . Theorem 1.18 is a direct consequence of Corollary 5.9 and Proposition 5.22 below.

PROPOSITION 5.22. *Let g and r be as in one of the rows in the table below. Let f be a pseudo-Anosov mapping class with an orientable invariant foliation on N_g whose stretch factor λ is smaller than r . Then λ must be a root of the polynomial shown in the table.*

In the cases where no polynomial is given, we have indicated to how many polynomials we were able to restrict the list of candidate polynomials.

| g | r | Polynomial candidates | largest root |
|-----|---------|-----------------------------|--------------|
| 4 | 1.84 | $t^3 - t^2 - t - 1$ | 1.83929 |
| 5 | 1.52 | $t^4 - t^3 - t^2 + t - 1$ | 1.51288 |
| 6 | 1.43 | $t^5 - t^3 - t^2 - 1$ | 1.42911 |
| 7 | 1.422 | $t^6 - t^5 - t^3 + t - 1$ | 1.42198 |
| 8 | 1.2885 | $t^7 - t^4 - t^3 - 1$ | 1.28845 |
| 9 | 1.3568 | 18 candidates | |
| 10 | 1.2173 | $t^9 - t^5 - t^4 - 1$ | 1.21728 |
| 11 | 1.22262 | 5 candidates | |
| 12 | 1.1743 | $t^{11} - t^6 - t^5 - 1$ | 1.17429 |
| 13 | 1.2764 | 288 candidates | |
| 14 | 1.14552 | $t^{13} - t^7 - t^6 - 1$ | 1.14551 |
| 15 | 1.1875 | 84 candidates | |
| 16 | 1.1249 | $t^{15} - t^8 - t^7 - 1$ | 1.12488 |
| 17 | 1.1426 | 16 candidates | |
| 18 | 1.10939 | $t^{17} - t^9 - t^8 - 1$ | 1.10938 |
| 20 | 1.09731 | $t^{19} - t^{10} - t^9 - 1$ | 1.09730 |

PROOF. The proof consists of running Algorithm 5.21 and is computer-assisted. However, we will prove the proposition by hand in genus 4. We follow the first four steps explicitly, then (since only a handful of polynomials remain) we finish the proof with an ad hoc but simple argument.

Step (1): we have

$$-1 = \min\{-1, 1.84 - 1 - 1/1.84\} < p_1 < 1.84 + 1 + 1/1.84 \approx 3.38,$$

therefore the possible values for p_1 are 0, 1, 2 and 3. We have

$$-1 = \min\{-1, 1.84^2 - 1 - 1/1.84^2\} < p_2 < 1.84^2 + 1 + 1/1.84^2 \approx 4.68,$$

so the possible values for p_2 are 0, 1, 2, 3 and 4.

Step (2): By (6), we have $c_1 = -p_1$ and $c_2 = \frac{p_1^2 - p_2}{2}$, therefore p_1 and p_2 have the same parity. Hence the possible pairs are (0, 0), (0, 2), (0, 4), (1, 1), (1, 3), (2, 0), (2, 2), (2, 4), (3, 1) and (3, 3).

Step (3): The pair $(p_1, \frac{p_1^2 - p_2}{2})$ also has the same parity, since our polynomial is reciprocal mod 2. That leaves the choices $(0, 0)$, $(0, 4)$, $(1, 3)$, $(2, 0)$, $(2, 4)$, $(3, 3)$ for (p_1, p_2) .

Step (4): We construct the list of possible polynomials.

- (1) $t^3 \pm 1$
- (2) $t^3 - 2t \pm 1$
- (3) $t^3 - t^2 - t \pm 1$
- (4) $t^3 - 2t^2 + 2t \pm 1$
- (5) $t^3 - 2t^2 \pm 1$
- (6) $t^3 - 3t^2 + 3t \pm 1$

The polynomial has to be irreducible, since the degree of a stretch factor on a nonorientable surface is at least three [62, Proposition 8.7.]. The polynomials where neither 1 nor -1 are roots are $t^3 - t^2 - t - 1$, $t^3 - 2t^2 + 2t + 1$, $t^3 - 2t^2 - 1$ and $t^3 - 3t^2 + 3t + 1$. The second and fourth polynomial do not have a positive real root, and the third polynomial has a root that is approximately 2.2. That leaves us with $t^3 - t^2 - t - 1$. \square

We have stopped at genus 20 because of computational difficulties. The genus 18 case took about half a day to run on a single computer. In the genus 20 case the algorithm took about a day to complete when run parallel on 30 computers. We estimate that the genus 22 case would need to run for a few months on the same cluster of computers.

In the odd genus cases, the issue is not the running time, but the fact that our tests are not good enough to eliminate all polynomials that should be eliminated. In the hope of dealing with more odd genus cases, we have also implemented the Lefschetz number tests used by Lanneau and Thiffeault [29, Section 2.3]. These tests help eliminate a large percentage of the remaining polynomials, but, unfortunately, not all. Table 5.1 below shows the polynomials that we could not eliminate in the genus 9, 11 and 13 cases, in addition to the polynomials that we have constructed in Corollary 5.9.

Most of the polynomials that we are not able to eliminate are products of polynomials that appear in some lower genus and cyclotomic polynomials. We think that these polynomials should be possible to eliminate, but we do not know how. In particular, we think that in the genus 9 and 11 cases all three remaining polynomials could be eliminated, and we conjecture that the examples constructed in Corollary 5.9 are the minimal stretch factor examples, compare with Conjecture 1.20.

| g | Polynomial | Stratum | Stretch factor |
|-----|--------------------------------|---------|----------------|
| 9 | $(t^7 - t^4 - t^3 - 1)(t - 1)$ | (4^7) | 1.28845 |
| | $(t^7 - t^5 - t^2 - 1)(t - 1)$ | (4^7) | 1.30740 |
| 11 | $(t^9 - t^5 - t^4 - 1)(t - 1)$ | (4^9) | 1.21728 |
| 13 | 18 polynomials | | |

TABLE 5.1. The polynomials and possible strata in genus 9, 11 that we cannot rule out using Algorithm 5.21 and Lefschetz arguments. The notation a^b means an orbit of length b consisting of a -pronged singularities.

Similarly, in the orientation-reversing case, Theorem 1.21 follows directly from Corollary 5.14 and Proposition 5.23 below.

PROPOSITION 5.23. *Let g and r be as in one of the rows in the table below. Let f be an orientation-reversing pseudo-Anosov mapping class with orientable invariant foliations on S_g whose stretch factor λ is smaller than r . Then λ must be a root of the polynomial shown in the table.*

| g | r | Polynomial candidates | largest root |
|-----|--------|--------------------------------|--------------|
| 1 | 1.62 | $t^2 - t - 1$ | 1.61803 |
| 2 | 1.62 | $t^2 - t - 1$ | 1.61803 |
| 3 | 1.253 | $t^8 - t^5 - t^3 - 1$ | 1.25207 |
| 4 | 1.253 | $t^8 - t^5 - t^3 - 1$ | 1.25207 |
| 5 | 1.16 | $t^{12} - t^7 - t^5 - 1$ | 1.15973 |
| 6 | 1.16 | $t^{12} - t^7 - t^5 - 1$ | 1.15973 |
| 7 | 1.1171 | $t^{16} - t^9 - t^7 - 1$ | 1.11707 |
| 8 | 1.1171 | $t^{16} - t^9 - t^7 - 1$ | 1.11707 |
| 9 | 1.0925 | $t^{20} - t^{11} - t^9 - 1$ | 1.09244 |
| 10 | 1.0925 | $t^{20} - t^{11} - t^9 - 1$ | 1.09244 |
| 11 | 1.0764 | $t^{24} - t^{13} - t^{11} - 1$ | 1.07638 |

PROOF. Analogously to the proof of Proposition 5.22, the proof of this statement is also computer-assisted. The algorithm used is a slight modification of Algorithm 5.21 as mentioned at the end of Section 4.3.

The polynomials in the table for $g \geq 3$ are not irreducible: they are products of $t^2 + 1$ and an irreducible factor. When $g \geq 3$ is odd, the polynomial we get as a result of the elimination process is this irreducible factor. When $g \geq 4$ is even, then the only polynomial left is the product of that irreducible factor and $t^2 - 1$. In either case, the stretch factor has to be a root of the irreducible factor, therefore a root of the polynomials in the table. The reasons for why we have not listed the irreducible factors in the

table is that they have many more terms and they do not show such a clear pattern as the polynomials in the table. Moreover, the polynomials in the table appear also in Corollary 5.14.

Similarly, in genus 2, the polynomial remaining after the elimination process is $(t^2 - t - 1)(t^2 - 1)$, so the stretch factor would have to be a root of $t^2 - t - 1$. \square

By using the Lefschetz number arguments of Lanneau and Thiffeault, we think it is possible to show that in genus 2 the only remaining polynomial, namely $(t^2 - t - 1)(t^2 - 1)$, cannot actually be the characteristic polynomial of the action on the first homology for an orientation-reversing pseudo-Anosov map with orientable invariant foliations. This would imply the strict inequality $\delta_{rev}^+(S_2) > \delta_{rev}^+(S_1)$. Since Proposition 5.23 shows a very clear pattern, we conjecture that the stretch factor candidates in Proposition 5.23 cannot be realised for any even genus, leading to Conjecture 1.23.

Minimal stretch factors in Penner's construction

The goal of this chapter is to give proofs of Theorems 1.26 and 1.27 on minimal stretch factors in Penner's construction on nonorientable closed surfaces. The content is copied and adapted from our joint work with Strenner [41].

1. Nonorientable surfaces

We start by giving three simple observations concerning the nonorientable case of Penner's construction, which will be used later on. Lemma 6.1 hints at why searching for the minimal stretch factor among pseudo-Anosov mapping classes arising from Penner's construction is more complicated on nonorientable surfaces than on orientable ones: the intersection graph of the curves used in the construction always contains at least one cycle, while for the minimising examples on closed orientable surfaces, it is a path [34].

LEMMA 6.1. *If a collection of curves $\{c_i\}$ as in Penner's construction fills a nonorientable surface, then their intersection graph is not bipartite.*

PROOF. Let S be a nonorientable closed surface, and let $\{c_i\}$ be a collection of curves as in Penner's construction that fill S . Recall that there exist homeomorphisms φ_{c_i} of regular neighbourhoods of the curves c_i to the standard annulus such that at each intersection point the pullback orientations disagree. If the intersection graph of the curves $\{c_i\}$ were bipartite, we could simply switch the orientation of the regular neighbourhoods of the the curves corresponding to one set of the bipartition to obtain a situation in which at each intersection point, the orientations of the regular neighbourhoods agree. In particular, as the curves $\{c_i\}$ are assumed to fill the surface S , we could extend this consistent orientation to an orientation of the surface S . \square

LEMMA 6.2. *Let S be any surface, and let $\{c_1, \dots, c_l\}$ be a collection of two-sided curves in S that intersect inconsistently and with the pattern of a cycle of odd length l . Then, a small regular neighbourhood Σ_0 of the union of the curves c_i is homeomorphic to N_{l+1} minus a disc.*

In particular, a collection of two-sided curves $\{c_i\}$ that intersect inconsistently and with the pattern of an odd cycle can only fill a nonorientable closed surface of even genus. Indeed, applying Lemma 6.2 to a collection of curves $\{c_1, \dots, c_l\}$ that in addition fill a closed surface S , we directly obtain the following statement.

COROLLARY 6.3. *Let S be a closed surface, and let $\{c_1, \dots, c_l\}$ be a collection of two-sided curves in S that intersect inconsistently and with the pattern of a cycle of odd length l . If the collection of curves $\{c_1, \dots, c_l\}$ fills S , then S is homeomorphic to N_{l+1} .*

PROOF OF LEMMA 6.2. Let $\{c_1, \dots, c_l\}$ be a collection of two-sided curves that intersect inconsistently and with the pattern of a cycle of odd length l . We want to show that the boundary of a small regular neighbourhood Σ_0 of the union of the curves c_i has exactly one boundary component. The statement then follows directly from the fact that Σ_0 is homotopy equivalent to a wedge of $l + 1$ circles, and hence has Euler characteristic $-l$.

Consider the surface Σ_1 obtained from Σ_0 by removing a *square*: the intersection of the annulus neighbourhoods of c_1 and c_l . The surface Σ_1 is homeomorphic to the surface obtained by chaining together l annuli and removing a square from the first and last annuli as in Figure 6.1. The




FIGURE 6.1. The surface Σ_1 . The letters A, B, C and D indicate how the strands of the boundary $\partial\Sigma_1$ connect.

boundary of the first and last annuli each have four arcs on the boundary of Σ_1 . In $\partial\Sigma_0 \cap \partial\Sigma_1$, the four arcs on the first annulus are connected to the four arcs on the last annulus as shown on Figure 6.1.

To reverse the process and construct the surface Σ_0 from Σ_1 , we need to glue $\partial\Sigma_1 \setminus \partial\Sigma_0$ to a square. Since the curves c_i are assumed to intersect inconsistently, there are two ways to do this, see Figure 6.2. We can see that




FIGURE 6.2. Glueing together the leftmost annulus (depicted vertically) and the rightmost annulus (depicted horizontally) from Figure 6.1 so that the orientations do not agree on the intersection.

in each case, all arcs get identified to a single boundary component. \square

A subgraph Γ' of the intersection graph Γ is *induced* if it contains all edges of Γ that connect pairs of vertices of Γ' .

LEMMA 6.4. *Let S be a surface filled by a collection of curves $\{c_i\}$ satisfying the hypotheses of Penner's construction. If the intersection graph Γ contains a cycle of odd length l as an induced subgraph, then S is nonorientable and its genus is greater than or equal to $l + 1$.*

PROOF. Let $c_{i_1}, \dots, c_{i_l} \subset S$ be the curves corresponding to the induced cycle of length l . This means that two consecutive curves c_{i_j} and $c_{i_{j+1}}$ intersect once, where the index j is taken (mod l). Since the cycle of length l is an induced subgraph of Γ , there are no other intersections between curves c_{i_j} . By Lemma 6.2, a regular neighbourhood U of the union of the curves c_{i_1}, \dots, c_{i_l} is not orientable, and hence neither is S . Furthermore, U has exactly one boundary component, and $\chi(U) = -l$. Thus, the surface S has the nonorientable closed surface of genus $l + 1$ as a connected summand, and, in particular, is of genus at least $l + 1$ itself. \square

2. Stretch factor theory of the cycle

The goal of this section is to describe the stretch factors arising from Penner's construction using curves with an odd cycle as their intersection graph, such as in Example 6.5 below.

EXAMPLE 6.5. Take an odd number l of annuli and glue them together to form a circle. Insert two half-twists in each annulus in order to produce a nonorientable surface. This is depicted in Figure 1.2 on the left for $l = 5$. Finally, glue in a disc along the boundary component (there is only one boundary component) to obtain a closed nonorientable surface, which by a direct Euler characteristic count is shown to be of genus $g = l + 1$. Number the core curves c_i of the annuli in the clockwise fashion. It is not hard to see that one can find homeomorphisms φ_{c_i} from regular neighbourhoods of the curves c_i to the standard annulus so that the curves c_i intersect inconsistently. Furthermore, there are no bigons, since each pair of curves c_i and c_j intersects at most once. It follows that the collection of core curves c_i satisfies the hypotheses of Penner's construction. The intersection graph is a cycle of length l .

Let C_l be a cycle of length l encoding the intersection of curves used in Penner's construction: to each curve c_i corresponds a vertex v_i of C_l . We now study mapping classes defined by a word w in the Dehn twists T_{c_i} so that every twist T_{c_i} appears exactly once. To every such word, we associate an acyclic orientation of C_l : an edge between v_i and v_j is directed from v_i to v_j if T_{c_i} occurs in w before T_{c_j} and vice-versa. The *flow difference* of an acyclic orientation of the cycle C_l is the number of edges oriented in the clockwise sense minus the number of edges oriented in the anticlockwise sense.

LEMMA 6.6. *Let w and w' be two words in the Dehn twists T_{c_i} so that every twist appears exactly once in each of them. If w and w' induce acyclic orientations of C_l with the same flow difference, then the matrices $\rho(w)$ and $\rho(w')$ from Penner's construction are conjugate.*

PROOF. Let W be the set of words in the Dehn twists T_{c_i} so that every twist appears exactly once. By a result of Shi [58], there exists a one-to-one correspondence between acyclic orientations of the cycle and words in W up to the commutation relation of Dehn twists (which commute exactly if the defining curves do not intersect). Moreover, two acyclic orientations of the cycle are connected by a sequence of source-to-sink operations if and only if they have the same flow difference by a result of Pretzel [52]. Here, a *source-to-sink operation* denotes the process of making a source of the directed graph into a sink by switching the orientations of all adjacent edges. By Shi's correspondence, on the level of the words, making a source into

a sink or vice-versa translates to a conjugation by the Dehn twist along the corresponding curve. In particular, the two matrices $\rho(w)$ and $\rho(w')$ associated with two pseudo-Anosov mapping classes arising from Penner's construction are conjugate if the words w and w' induce acyclic orientations of C_l with the same flow difference. \square

Since conjugate matrices have the same eigenvalues, we only have to study one standard representative for each flow difference. By symmetry, we also have to consider only the absolute value of the flow difference.

2.1. A formula for the stretch factor. The goal of this section is to show that for a cycle of fixed length, the stretch factor of Dehn twist products is a strictly increasing function of the absolute value of the flow difference. This follows from the Propositions 6.7 and 6.8 below, which give the means to directly compute the stretch factor given the length of the cycle and the flow difference.

Let $C = \{(x, y) \in \mathbb{R}^2 : y > 0, |x| < y\}$. Furthermore, define the function $f : C \rightarrow \mathbb{R}_{>0}$ by mapping (x, y) to the largest real solution of the equation $t - t^{\frac{y+x}{2y}} - t^{\frac{y-x}{2y}} - 1 = 0$.

PROPOSITION 6.7. *The function f is well-defined and*

- (1) *is 0-homogeneous, and, in fact, only depends on $|\frac{x}{y}|$,*
- (2) *is continuous and strictly increasing in $|\frac{x}{y}|$.*

PROPOSITION 6.8. *For a tuple $(d, l) \in \mathbb{Z}^2 \cap C$ such that $d \equiv l \pmod{2}$, the value $f(d, l)$ equals the stretch factor of the Penner mapping classes with flow difference d on the cycle of length l .*

PROOF OF PROPOSITION 6.7. Notice that

$$t - t^{\frac{y+x}{2y}} - t^{\frac{y-x}{2y}} - 1 = t - t^{\frac{1}{2} + \frac{x}{y}} - t^{\frac{1}{2} - \frac{x}{y}} - 1 = t - t^{\frac{1}{2} + |\frac{x}{y}|} - t^{\frac{1}{2} - |\frac{x}{y}|} - 1.$$

This proves (1), assuming that f is well-defined. Define

$$h(t, s) = t - t^{\frac{1}{2} + s} - t^{\frac{1}{2} - s} - 1$$

for $0 < s < \frac{1}{2}$. For every s , $h(1, s) = -2$. Furthermore, $\partial_t h(t, s) > 0$ for all $t > 1$. It follows that for any fixed s , the function $h(\cdot, s)$ has exactly one real zero > 1 . This shows that f is well-defined. Furthermore, $\partial_t h(t, s)$ depends continuously on s , therefore so does the real zero > 1 of the function $h(\cdot, s)$. This proves the first part of (2). In order to see the second part of (2), notice that $\partial_s h(t, s) < 0$. This implies that the real zero > 1 of the function $h(\cdot, s)$ is strictly increasing in s . \square

EXAMPLE 6.9 (Twist and click homeomorphisms). Let $l, c \in \mathbb{N}$ be natural numbers such that $c < l$ and $\gcd(c, l) = 1$. Let Σ_l be the surface obtained by thickening a collection of l curves with the l -cycle as their intersection graph, so that between any two intersections there is a half-twist. This is depicted for $l = 5$ in Figure 1.2 on the left. Consider the mapping class $\phi_{l,c}$ obtained by a Dehn twist along one of the curves composed with a c -fold click, that is, a rotation of the symmetric surface Σ_l by an angle $c \cdot \frac{2\pi}{l}$. The l -th power of such a mapping class $\phi_{l,c}$ arises from Penner's construction using the core curves of the annuli with the l -cycle as their intersection graph, and every curve gets twisted along exactly once. Since $\phi_{l,c}^l$ is pseudo-Anosov by Penner's construction, so is $\phi_{l,c}$ by the classification of surface homeomorphisms and the stretch factor of $\phi_{l,c}$ is the l -th root of the stretch factor of $\phi_{l,c}^l$.

The following lemma describes the stretch factor of the twist and click mapping classes introduced in Example 6.9. For $c = 2$, the result is also stated by Strenner and the author in [39]. In this case the absolute value of the flow difference is 1. The proofs are basically identical.

LEMMA 6.10. *Let a be the smallest natural number with $ac \equiv 1 \pmod{l}$. Then, the stretch factor of $\phi_{l,c}$ is given by the largest real root of the polynomial $t^l - t^{l-a} - t^a - 1$.*

PROOF. The mapping class $\phi_{l,c}^l$ is pseudo-Anosov and arises from Penner's construction. Furthermore, the associated matrix $\rho(\phi_{l,c}^l)$ in Penner's construction equals the action on the first homology of the surface induced by $\phi_{l,c}^l$. To see this, choose the collection of the core curves of the annuli as a basis for the first homology, oriented invariantly under rotation. From this it follows that $\phi_{l,c}^l$ has an orientable invariant foliation, and hence so does $\phi_{l,c}$. In particular, the stretch factor of $\phi_{l,c}$ is given by the spectral radius of its action induced on the first homology of the surface, which we describe now. Number the core curves of the twisted bands in the following way. The first curve c_1 is the one along which we do a Dehn twist in the definition of $\phi_{l,c}$. The second curve c_2 is the image of c_1 under rotation of Σ_l by an angle $-c \cdot \frac{2\pi}{l}$. The third curve c_3 is the image of c_2 under rotation of Σ_l by an angle $-c \cdot \frac{2\pi}{l}$, and so on. As a basis for the first homology $H_1(\Sigma_l; \mathbb{R})$, we choose the homology classes of c_1, c_2, \dots, c_l . We obtain that the rotation r of Σ_l by an angle $c \cdot \frac{2\pi}{l}$ acts by a permutation matrix, sending c_i to c_{i-1} , where the indices are taken $(\text{mod } l)$. Furthermore, the Dehn twist T_{c_1} acts

as the identity on the curves c_i for $i \neq a, l - a$, and adds the curve c_1 to the curves c_a and c_{l-a} . The product of these matrix actions is a companion matrix for the polynomial $t^l - t^{l-a} - t^a - 1$. For example, for $l = 5$ and $c = 1$, we have $a = 1$ and

$$(T_{c_1})_* = \begin{pmatrix} 1 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}, r_* = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \end{pmatrix},$$

and

$$(\phi_{5,1})_* = r_* \cdot (T_{c_1})_* = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 \end{pmatrix},$$

so $(\phi_{5,1})_*$ has characteristic polynomial $t^5 - t^4 - t - 1$. \square

LEMMA 6.11. *Let a be the smallest positive integer with $ac \equiv 1 \pmod{l}$. Then, the flow difference of $\phi_{l,c}^l$ is $l - 2a$.*

PROOF. We identify the elements of $\mathbb{Z}/l\mathbb{Z}$ with the vertices of the cycle C_l of length l . We consider the sequence $0, c, 2c, \dots, (l-1)c \in \mathbb{Z}/l\mathbb{Z}$, which is the sequence in which $\phi_{l,c}^l$ twists along the curves (corresponding to elements of $\mathbb{Z}/l\mathbb{Z}$). In order to determine the flow difference of $\phi_{l,c}^l$, it suffices to know for each pair of adjacent elements $k, k+1 \in \mathbb{Z}/l\mathbb{Z}$ which element appears first in the sequence. Indeed, if $k \in \mathbb{Z}/l\mathbb{Z}$ appears first in the sequence, then the edge connecting the k th and the $k+1$ st vertex is oriented towards the $k+1$ st vertex, and vice versa.

Assume for a moment that a is minimal so that $ac \equiv \pm 1 \pmod{l}$. Then, the residue classes $0, c, 2c, \dots, (a-1)c \in \mathbb{Z}/l\mathbb{Z}$ are pairwise nonadjacent. Now, the next residue class in the sequence is $ac \equiv 1 \equiv 0 + 1 \pmod{l}$. Each element that occurs in the sequence after ac can also be obtained by adding 1 to an element that occurred already before in the sequence. We deduce that we obtain $l - a$ edges pointing in the clockwise direction and a edges pointing in the anticlockwise direction. This yields a flow difference of $(l - a) - a = l - 2a$.

If a is not the minimal natural number so that $ac \equiv \pm 1 \pmod{l}$, then we have $l - a < a$ and $(l - a)c \equiv -1 \pmod{l}$. We can repeat the same argument, but the direction of each edge is switched. We obtain $l - (l - a)$ edges pointing in the anticlockwise direction and $l - a$ edges pointing in the clockwise direction. This yields a flow difference of $l - a - (l - (l - a)) = l - 2a$. \square

PROOF OF PROPOSITION 6.8. Assume for a moment that l is odd. We first reduce to the case $\gcd(d, l) = 1$. For this, assume for a moment $\gcd(d, l) > 1$. We have

$$f(d, l) = f\left(\frac{d}{\gcd(d, l)}, \frac{l}{\gcd(d, l)}\right)$$

by 0-homogeneity of f . Note that a Penner mapping class of flow difference $\frac{d}{\gcd(d, l)}$ on the cycle of length $\frac{l}{\gcd(d, l)}$ is covered $\gcd(d, l)$ -fold by a Penner mapping class with flow difference d on the cycle of length l . It therefore suffices to prove the statement for $\gcd(d, l) = 1$.

In the twist and click mapping classes for a fixed odd length l , as c runs through the numbers smaller than l with $\gcd(c, l) = 1$, also the corresponding a runs through the numbers smaller than l with $\gcd(a, l) = 1$. Therefore, the numbers $l - 2a$ run through the odd numbers of absolute value smaller than l with $\gcd(l - 2a, l) = 1$. In particular, we obtain every flow difference d with $\gcd(d, l) = 1$ as an l -th power of a twist and click example. By Lemma 6.10 and 6.11, the stretch factor of the Penner mapping classes with flow difference d on the cycle of length l is the l -th power of the largest real root of the polynomial $t^l - t^{l-a} - t^a - 1$, where $a = (l - d)/2$. Equivalently, the stretch factor equals the largest real solution of the equation

$$t - t^{\frac{l+d}{2l}} - t^{\frac{l-d}{2l}} - 1 = 0,$$

which finishes the proof in the case where l is odd.

Now, let l be even. The proof of this case is similar, the main difficulty being that by dividing both l and d by $\gcd(d, l)$, it is possible to break the condition $d \equiv l \pmod{2}$. This time, we reduce our argument to the case $\gcd(d, l) = 2$ and $d \not\equiv l \pmod{4}$. Indeed, this is exactly the case where the condition $d \equiv l \pmod{2}$ does not hold anymore after dividing both l and d by 2. Notice that any other case either reduces to this one or a case where l is odd, by a covering argument as above. It therefore suffices to prove the statement for $\gcd(d, l) = 2$ and $d \not\equiv l \pmod{4}$.

As in the argument for odd l , we again use the twist and click mapping classes from Example 6.9. The only difference is that in this case, the numbers $l - 2a$ run through the even numbers of absolute value smaller than l with $\gcd(l - 2a, l) = 2$ and $l - 2a \not\equiv l \pmod{4}$. Indeed, since a is odd, we obtain $l - 2a \not\equiv l \pmod{4}$. On the other hand, since we get all a with $\gcd(a, l) = 1$ by varying c with $\gcd(c, l) = 1$, we obtain all flow differences d with $\gcd(d, l) = 2$ and $d \not\equiv l \pmod{4}$ by an l -th power of a twist and click example. Here, we again use Lemma 6.11 to argue that $d = l - 2a$. As in the case of odd l , the statement follows from Lemma 6.10. \square

REMARK 6.12. For odd l and $c = 2$, the twist and click mapping class $\phi_{l,c}$ conjecturally minimises the stretch factor among pseudo-Anosov mapping classes with an orientable invariant foliation on the nonorientable closed surface of genus $l + 1$. This has been shown for even genus up to 20 in our joint work with Strenner [39]. Adding the mapping classes $\phi_{l,c}$ for other c to the picture as in Proposition 6.7 exhibits a strong similarity with theory of the normalised stretch factor on a fibred face of the Thurston norm ball [17, 18]. Indeed, we expect many of the mapping classes $\phi_{l,c}$ to lie in a common fibred cone.

3. Even genus minimal stretch factors

The goal of this section is to single out the minimal stretch factor examples among mapping classes arising from Penner's construction on a closed nonorientable surface of even genus. We will often use the following lemma to obtain lower bounds for the stretch factor of mapping classes arising from Penner's construction. It was implicitly used already in the case of orientable surfaces by the author [34].

LEMMA 6.13. *Let ϕ be a mapping class arising from Penner's construction using a collection of curves $\{c_i\}$. If the intersection graph of the curves $\{c_i\}$ contains a tree Γ (possibly with multiple edges between two vertices) as a subgraph, then*

$$\lambda(\phi) \geq \frac{2 + \alpha^2 + \sqrt{4\alpha^2 + \alpha^4}}{2},$$

where α is the largest eigenvalue of the adjacency matrix of Γ .

PROOF. Let ϕ be a mapping class arising from Penner's construction using a collection of curves $\{c_i\}$, and let the tree Γ be a subgraph of the intersection graph of the curves $\{c_i\}$. For two matrices A and B of the

same dimensions, we write $A \leq B$ if $a_{ij} \leq b_{ij}$ for all i, j . Recall that the spectral radius of nonnegative matrices is monotonic under " \leq ", see, for example, [6]. We may therefore assume that ϕ is a product of Dehn twists T_{c_i} so that every curve c_i gets twisted along exactly once. Let ϕ_Γ be the subproduct of Dehn twists T_{c_i} along exactly those curves c_i which correspond to the vertices of Γ . We have that $\lambda(\phi)$ is an upper bound for the spectral radius of $\rho(\phi_\Gamma)$. The spectral radius of $\rho(\phi_\Gamma)$ is in turn an upper bound for the Penner stretch factor $\lambda(\Gamma)$ associated with the subgraph Γ and its induced order of twisting. Note that by a result of Steinberg, the order of twisting does not change the conjugacy class, since Γ is a tree [61]. It follows that $\lambda(\Gamma)$ is independent of the Dehn twist product order on Γ . In particular, we may calculate $\lambda(\Gamma)$ as the stretch factor of a product of two multitwists, in which case the Thurston–Veech construction yields

$$\lambda(\Gamma) + \lambda(\Gamma)^{-1} - 2 = \alpha^2,$$

where α is the largest eigenvalue of the adjacency matrix of Γ , see [65, 68]. Solving this equation for $\lambda(\Gamma)$ yields the result. \square

Let φ_l be the mapping class defined by the l th power of the twist-and-click mapping class $\phi_{l,2}$, where l is an odd natural number. By Lemma 6.11, if $c = 2$, then $a = \frac{l+1}{2}$ and the absolute value of the flow difference associated with φ_l equals $|l - 2a| = 1$. Both Lemma 6.15 and Lemma 6.14 follow readily from Propositions 6.7 and 6.8.

LEMMA 6.14. *Among pseudo-Anosov mapping classes arising from Penner's construction using curves with an odd l -cycle as their intersection graph, the mapping class φ_l has minimal stretch factor.*

PROOF. By Proposition 6.8, the stretch factor of the Penner mapping classes with flow difference d on the cycle of length l equals $f(d, l)$. By Proposition 6.7, the function $f(d, l)$ is strictly increasing in $|\frac{d}{l}| = \frac{|d|}{l}$. This means that for a cycle of fixed length l , the stretch factor is a strictly increasing function of the absolute value of the flow difference. In particular, the stretch factor is minimised for the minimal absolute value of the flow difference, which for a cycle of odd length l is 1. \square

LEMMA 6.15. *We have $\lambda(\varphi_{l+2}) < \lambda(\varphi_l)$.*

PROOF. By Proposition 6.8, we have that the stretch factor of $\lambda(\varphi_j)$ is $f(1, j)$. By Proposition 6.7, the function $f(1, j)$ is strictly increasing in $|\frac{1}{j}| = \frac{1}{j}$ and hence strictly decreasing in j . \square

We are now ready to describe the Penner mapping classes of minimal stretch factor on nonorientable closed surfaces of even genus. Note that we only have to consider nonorientable surfaces of genus at least four since the mapping class group of the Klein bottle is finite and thus does not contain pseudo-Anosov elements.

THEOREM 6.16. *The mapping class φ_l has the minimal stretch factor among pseudo-Anosov mapping classes arising from Penner's construction for a nonorientable closed surface of even genus $g = l + 1$.*

PROOF. Let N_{l+1} be the nonorientable closed surface of genus $l + 1$, which we assume to be an even number. We know that there exists the mapping class φ_l on N_{l+1} , with stretch factor $\lambda(\varphi_l)$. Furthermore, let ϕ be any mapping class on N_{l+1} arising from Penner's construction. As before, we are allowed to assume that every curve used for the construction of ϕ gets twisted along exactly once. We distinguish cases depending on the intersection graph of the curves used in the construction of ϕ .

Case 1: the intersection graph contains a double edge. Let c_1 and c_2 be two curves that intersect at least twice. Since a bipartite family of curves which intersect inconsistently cannot fill a nonorientable surface, there must be at least one other curve c_3 intersecting either c_1 or c_2 . In particular, the intersection graph of the curves $\{c_i\}$ contains the tree Γ with three vertices, one double edge and one simple edge as a subgraph, depicted in Figure 6.3 on the left. The adjacency matrix of this tree has maximal eigenvalue $\sqrt{5}$ and we use Lemma 6.13 to conclude

$$\lambda(\phi) \geq \frac{7 + 3\sqrt{5}}{2} \approx 6.854.$$

This number is larger than the stretch factor of any mapping class φ_l by the values given in Table 6.1 and the monotonicity due to Lemma 6.15.




FIGURE 6.3.

Case 2: the intersection graph contains an odd cycle of length $k \leq l$: In this case, by an argument similar to the argument used to prove Lemma 6.13, the stretch factor is always bounded from below by the stretch factor of

a pseudo-Anosov mapping class arising from Penner's construction using curves that intersect with the pattern of an odd cycle of length $k \leq l$. In particular, Lemmas 6.14 and 6.15 imply $\lambda(\phi) \geq \lambda(\varphi_k) \geq \lambda(\varphi_l)$.

| cycle length | flow difference | stretch factor |
|--------------|-----------------|-----------------|
| 3 | 1 | ≈ 6.222 |
| 5 | 1 | ≈ 5.961 |
| 5 | 3 | ≈ 7.520 |
| 7 | 1 | ≈ 5.895 |
| 7 | 3 | ≈ 6.529 |
| 7 | 5 | ≈ 8.841 |

TABLE 6.1. Some stretch factors for short odd cycles.

Case 3: the intersection graph only contains odd cycles of length $k > l$: Take an odd cycle of minimal length $k > l$ among odd cycles. This cycle is necessarily an induced subgraph of the intersection graph. Otherwise, the intersection graph would either have to contain a double edge (which we may rule out by Case 1) or an edge connecting two nonadjacent vertices of the cycle, which implies the existence of an odd cycle of length $< k$. Hence, by Lemma 6.4, the genus of the surface N_{l+1} is bounded from below by $k + 1 > l + 1 = g$, a contradiction. \square

3.1. A proof of Theorem 1.27 and almost a proof of Theorem 1.26. By Theorem 6.16 and Proposition 6.8, we know that for even genus g , the minimal stretch factor $\delta_P(N_g)$ equals the largest real solution of the equation

$$t - t^{\frac{g}{2g-2}} - t^{\frac{g-2}{2g-2}} - 1 = 0.$$

Setting $g = 2k$ yields exactly the statement of Theorem 1.27.

We are now ready to show Theorem 1.26, except for the existence of the limit $\lim_{k \rightarrow \infty} \delta_P(N_{2k+1})$ for nonorientable closed surfaces of odd genus.

THEOREM 6.17. *For the minimal stretch factor $\delta_P(N_g)$ among pseudo-Anosov mapping classes arising from Penner's construction on the nonorientable closed surface of genus g , the limit $\lim_{k \rightarrow \infty} \delta_P(N_{2k})$ exists, and*

- (a) $\lim_{k \rightarrow \infty} \delta_P(N_{2k}) = 3 + 2\sqrt{2}$,
- (b) $\liminf_{k \rightarrow \infty} \delta_P(N_{2k+1}) > 3 + 2\sqrt{2}$.

PROOF. By Theorem 1.27, we know that for even g , $\delta_P(N_g)$ equals the largest real solution of the equation

$$t - t^{\frac{g}{2g-2}} - t^{\frac{g-2}{2g-2}} - 1 = 0.$$

As $g \rightarrow \infty$, this solution converges to the largest real solution of the equation

$$t - 2t^{\frac{1}{2}} - 1 = 0,$$

which is $3 + 2\sqrt{2}$, the square of the silver ratio. This proves the existence of the limit $\lim_{k \rightarrow \infty} \delta_P(N_{2k})$ and the exact value in (a).

In order to prove (b), we show that a Penner stretch factor on a nonorientable surface of odd genus is bounded from below by $3 + 2\sqrt{2} + \delta$, for $\delta = \frac{1}{10}$. We can use similar steps as in the proof of Theorem 6.16 and the values in Table 6.1 to reduce the argument to the case where the intersection graph contains an induced cycle of length > 9 . By Corollary 6.3, the corresponding curves cannot fill the surface, since it is of odd genus. Hence, the intersection graph must contain at least one more vertex connecting to the induced cycle. In particular, it contains a subgraph of the form depicted in Figure 6.3 on the right. In this case, we use Lemma 6.13 to obtain that the stretch factor is bounded from below by 5.946. \square

In order to show that the limit $\lim_{k \rightarrow \infty} \delta_P(N_{2k+1})$ exists, we need a better grip on the actual minimal Penner stretch factors $\delta_P(N_g)$ for odd g . To this end, we study odd cycles with an extra vertex in the next section.

4. Stretch factor theory of the enriched cycle

Let P_l be the *enriched cycle of length l* , that is, the l -cycle with an additional vertex connecting to exactly one vertex of the cycle. In order to deal with closed nonorientable surfaces of odd genus, we have to study these examples systematically. The goal of this section is to prove the following analogues of Lemma 6.14 and Lemma 6.15 for enriched cycles. For l odd, let μ_l be the stretch factor arising from Penner's construction using curves that have P_l as their intersection graph and a Dehn twist product with flow difference 1. By the *flow difference* of an enriched cycle we just mean the flow difference of the induced cycle obtained by removing the extra vertex. As in Lemma 6.6, there is exactly one conjugacy class for each flow difference.

LEMMA 6.18. *The minimal stretch factor arising from Penner's construction using curves that have P_l as their intersection graph is μ_l .*

LEMMA 6.19. *We have $\mu_l \geq \mu_{l+2}$.*

In order to prove Lemma 6.18 and Lemma 6.19, we will study fibred link representatives and the Perron–Frobenius eigenvectors of the matrices

associated with the mapping classes arising from Penner's construction, respectively.

4.1. Fibred links. An oriented compact surface Σ (with oriented boundary) embedded in \mathbb{S}^3 is a *fibre surface* if its interior $\overset{\circ}{\Sigma}$ is the fibre of a locally-trivial fibre bundle $p : \mathbb{S}^3 \setminus \partial\Sigma \rightarrow \mathbb{S}^1$. In this case, the oriented boundary $\partial\Sigma$ is called a *fibred link*. Such a fibration is determined by a mapping class of Σ up to conjugation, the *monodromy* of the fibration.

Given two oriented surfaces embedded in \mathbb{S}^3 , it is possible to obtain new oriented surfaces by *plumbing*, that is, glueing the two surfaces (which are separated by an oriented embedded sphere \mathbb{S}^2) together along a square (which is contained in the sphere \mathbb{S}^2) whose boundary arcs alternatingly belong to the boundary of one surface or the other, see Figure 6.4 for an example. We assume that in the plumbing square, the orientations of both surfaces and the sphere agree.




FIGURE 6.4. A fibre surface which is obtained by successive plumbing of Hopf bands to the standard disc. The plumbing square for each of the plumbings is coloured grey.

By a result of Stallings, a plumbing Σ of two fibre surfaces Σ_1 and Σ_2 is again a fibre surface [60]. Furthermore, we now assume Σ_1 to be on the negative side (the “inside”) of the sphere \mathbb{S}^2 and Σ_2 on the positive side (the “outside”) of the sphere \mathbb{S}^2 . Then, the monodromy ϕ of the plumbing is given by the composition $\phi_1 \circ \phi_2$ of the two monodromies ϕ_1 and ϕ_2 of the plumbing summands Σ_1 and Σ_2 , extended to Σ by the identity on $\Sigma \setminus \Sigma_1$ and $\Sigma \setminus \Sigma_2$, respectively. For this to make sense, recall that mapping classes of surfaces with boundary are assumed to fix the boundary pointwise.

The positive Hopf band and the negative Hopf band are fibre surfaces and their monodromies are a positive Dehn twist and a negative Dehn twist along the core curve, respectively. This fact, as well as Stallings' result, is accessibly explained by Baader and Graf, who interpret the concept of

fibredness in terms of elastic cords [3]. By Stallings' result, a successive plumbing of Hopf bands yields a product of Dehn twists along the core curves of the Hopf bands plumbed. In this way, it is possible to represent certain mapping classes as monodromies of fibred links.

4.2. Realising Penner maps as fibred link monodromies. With the preceding discussion on fibred link monodromies, it is clear what we should do in order to obtain mapping classes that arise from Penner's construction as monodromies of fibred links: plumb positive and negative Hopf bands such that the core curves of the positive Hopf bands do not intersect among themselves and likewise for the negative Hopf bands. For instance, Figure 6.5 depicts two fibre surfaces. Both are obtained from the closed standard disc, which is situated in the middle, by consecutive plumbing of Hopf bands. There is a total of three positive and three negative Hopf bands plumbed in alternating fashion. In this way, we obtain fibred links whose monodromy is a product of Dehn twists along curves which intersect each other with the pattern of a cycle. Furthermore, the monodromy is a pseudo-Anosov mapping class arising from Penner's construction by using the core curves of the plumbed positive Hopf bands as one multicurve and the core curves of the plumbed negative Hopf bands as the other multicurve. It is straightforward to see that we are able to represent any order of Dehn twists by varying the order of plumbing, and, in particular, every flow difference.




FIGURE 6.5. Fibred links realising flow difference 0 (on the left) and flow difference 2 (on the right).

REMARK 6.20. By the process of representing Penner mapping classes as monodromies of fibred links, we only obtain orientable surfaces. However, we

can still study the stretch factors of Penner mapping classes on nonorientable surfaces by first lifting to the orientable double cover. For example, if we want to study the stretch factor of a Penner mapping class given by an odd cycle of length l with a flow difference d , we can lift it to a Penner mapping class on an orientable surface, with intersection graph the cycle of length $2l$ and with flow difference $2d$. The stretch factors of the original Penner mapping class and its lift agree and the latter can be represented by a fibred link monodromy. In this context, we recall that a Penner stretch factor depends only on the intersection graph and the twist order, so it suffices to represent this information and not the actual Penner mapping classes.

4.3. The Alexander polynomial of fibred links. The Alexander polynomial Δ_L of an oriented link L is defined recursively by the skein relation

$$(SR) \quad \Delta_{L_+} = \Delta_{L_-} + \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right) \Delta_{L_0},$$

and the initial condition $\Delta_U = 1$, where U is the unknot. For a fixed crossing, the links L_+ , L_- and L_0 correspond to the positive version of the crossing, the negative version of the crossing and the orientation-preserving smoothing of the crossing, see Figure 6.6. The important property in our context is




FIGURE 6.6. The links L_+ , L_- and L_0 used in the skein relation are obtained by local adjustments at a crossing.

that for fibred links, the Alexander polynomial equals the characteristic polynomial of the action on the first homology of the fibre surface induced by the monodromy, up to a normalisation factor, see, for example, [8]. The normalisation factor equals $\sqrt{t}^{b_1(L)}$ with a possible sign -1 to make the leading coefficient positive. This follows from the fact that the characteristic polynomial of a matrix of size $b_1(L)$ is of degree $b_1(L)$ and with leading coefficient $+1$, while the highest power of t appearing with nonzero coefficient in the Alexander polynomial of a fibred link is $\sqrt{t}^{b_1(L)}$, where $b_1(L)$ is the first Betti number of the fibre surface for L . We would like to stress that

while the Alexander polynomial is often defined up to powers of the variable and up to sign, the skein-theoretic definition we use here gives a well-defined Laurent polynomial in \sqrt{t} .

4.4. A proof of Lemma 6.18. Let $\Delta_{d,l}$ be the Alexander polynomial of the fibred link realisation of the mapping class arising via Penner’s construction on curves that intersect with the pattern of a cycle of even length $l = p + q$, where every curve is twisted along exactly once and the twist order yields flow difference $d = p - q$, where we may assume $d \geq 0$. The following proposition contains the key result on Alexander polynomials of fibred link realisations.

PROPOSITION 6.21. *In the above notation, we have*

$$\Delta_{d,l} - \Delta_{d+2,l} = \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right) \left(\sqrt{t}^{d+1} - \sqrt{t}^{-(d+1)} \right).$$

Our proof of Proposition 6.21 relies on the skein relation for the Alexander polynomial. We split it up into several separate statements. For a natural number d , let H_d be the link obtained by the closure of d stacked copies of the braid depicted on the left in Figure 6.7. On the right in Figure 6.7, the link H_d is shown for $d = 4$.




FIGURE 6.7. The braid building block for the links H_d (on the left) and the link H_4 (on the right).

LEMMA 6.22. *Let $p \geq q, p' \geq q'$ be strictly positive natural numbers such that $p + q = p' + q' = l$. Let $d = p - q$ and $d' = p' - q'$. Then, we have*

$$\Delta_{d,l} - \Delta_{d',l} = \Delta_{H_d} - \Delta_{H_{d'}}.$$

PROOF. Let L_d and $L_{d'}$ be two links representing flow differences d and d' , respectively, on a cycle of length $l = p + q = p' + q'$, consisting

of l Hopf bands plumbed to a closed standard disc. We consider diagrams of L_d and $L_{d'}$ as described in Section 4.2 and Figure 6.5, and apply the skein relation to a crossing of a twist of one of the plumbed bands. The change from a positive crossing to a negative crossing or vice-versa manifestly untwists the band. The smoothing of the crossing as in the link L_0 of the skein relation cuts the band. The resulting link is a plumbing of Hopf bands along a path. In this case, whether a band passes over another one or vice-versa does not change the link up to isotopy. In particular, the L_0 -terms in the skein relation for L_d and $L_{d'}$ agree. Hence, $\Delta_{d,l} - \Delta_{d',l}$ equals the difference of the Alexander polynomials of the links L_d and $L_{d'}$, but with one Hopf band untwisted. This argument can be repeated for each of the Hopf bands, which finally yields $\Delta_{d,l} - \Delta_{d',l} = \Delta_{H_d} - \Delta_{H_{d'}}$, since the link L_d with all Hopf bands untwisted is exactly the link H_d . Indeed, the link L_d with all Hopf bands untwisted can be divided into p sectors that resemble a positive half-twist on three strands and q sectors that resemble a negative half-twist on three strands, with the orientation of the middle strand reversed, compare with Figure 6.8. Since a positive and a negative half-twist cancel each




FIGURE 6.8. A link isotopy supported in bounded sectors.

other (which is very well perceivable in Figure 6.8 on the right), we are left with $d = p - q$ positive half-twists, that is, the link H_d . \square

LEMMA 6.23. *For even d , we have*

$$\Delta_{H_d} = \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right) \left(\Delta_{T_{2,d}} - 2 \sum_{i=1}^{\frac{d}{2}} \Delta_{T_{2,2i}} \right),$$

where $T_{2,2i}$ denotes the $(2, 2i)$ -torus link.

PROOF. The idea is to subsequently use the skein relation of the Alexander polynomial on all crossings where the middle strand of H_d passes below

an other strand, starting from the highest such crossing and proceeding to the lowest. This allows for a computation after finitely many steps until the middle strand corresponds to a split component. The crossing changes in the skein relation simplifies the linking of the middle strand with the other strands. This is depicted in Figure 6.9. We now show that the L_0 -




FIGURE 6.9.

smoothings accumulate $\Delta_{T_{2,2i}}$ -summands with a coefficient $-\left(\sqrt{t} - \frac{1}{\sqrt{t}}\right)$. Assume that we have already changed the highest $k \geq 0$ crossings where the middle strand of H_d passes below an other strand. We now describe what happens when we smooth the $k + 1$ st crossing as in the L_0 -part of the skein relation. Explicitly drawing the diagrams reveals that if k is even, we obtain a torus link $T_{2,k}$ and if k is odd, we obtain a torus link $T_{2,k+1}$. Here, we also recall that $d = p - q$ is even. This is depicted in Figure 6.10 for $d = 6$ and $k = 3$. In the diagram for H_d , there are d undercrossings of




FIGURE 6.10.

the middle strand. After we have changed all the crossings, the link is split and has Alexander polynomial 0. Hence, the Alexander polynomial of H_d is

the sum over all d undercrossings of the Alexander polynomial of the torus link obtained by the corresponding L_0 -smoothing, as described above, with a coefficient

$$-\left(\sqrt{t} - \frac{1}{\sqrt{t}}\right).$$

This yields the desired result. \square

The following lemma is a standard fact on the Alexander polynomial of torus links. Using the skein relation for the Alexander polynomial, its verification is straightforward.

LEMMA 6.24.

$$\Delta_{T_{2,2i}} = \sum_{j=0}^{i-1} (-1)^j \left(\sqrt{t}^{2i-1-2j} - \sqrt{t}^{-(2i-1-2j)} \right).$$

PROOF OF PROPOSITION 6.21. By Lemma 6.22, we have

$$\Delta_{d,l} - \Delta_{d+2,l} = \Delta_{H_d} - \Delta_{H_{d+2}}.$$

On the other hand, using first Lemma 6.23 and then Lemma 6.24, we obtain

$$\begin{aligned} \Delta_{H_d} - \Delta_{H_{d+2}} &= \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right) (\Delta_{T_{2,d+2}} + \Delta_{T_{2,d}}) \\ &= \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right) \left(\sqrt{t}^{d+1} - \sqrt{t}^{-(d+1)} \right), \end{aligned}$$

which is what we wanted to show. \square

LEMMA 6.25. *The leading coefficient of $\Delta_{d,l}$ is $+1$.*

PROOF. By the exact same skein relations as in the proof of Lemma 6.22, we have that $\Delta_{d,l}$ is a sum of Δ_{H_d} and summands of Alexander polynomials of plumbings of Hopf bands of alternating kind along a path, with varying coefficients from the skein relation. A careful inspection reveals that the leading coefficient of $\Delta_{d,l}$ equals the leading coefficient of the Alexander polynomial of the longest such path starting and ending with a negative Hopf band. With a recursion on the length of such a path, one can show that this leading coefficient is $+1$. \square

We will show that for the enriched cycle, the stretch factor is a monotonic function of the absolute value of the flow difference. In particular, this implies that for an odd enriched cycle of length l , the minimal Penner stretch factor is obtained by the example with flow difference 1.

PROOF OF LEMMA 6.18. We will compare the stretch factors of Penner mapping classes obtained by twisting along curves which intersect like an

enriched odd cycle, when we vary the flow difference associated with the order of twisting. A concrete surface (of genus 7, that is, $l = 5$) for which we can build such an example is shown in Figure 1.2 on the right. We first lift the mapping class to the double cover orienting the surface. By doing this, we double the length of the cycle and the flow difference. Furthermore, there are now two extra vertices connecting to the cycle at opposite ends, one corresponding to a curve along which we twist positively and one corresponding to a curve along which we twist negatively. This lifted mapping class has the same stretch factor as the Penner mapping class we started with.

We find a fibred link representative of the lifted mapping class by taking the usual representative for the cycle and plumbing two extra Hopf bands H_1 and H_2 , such as in Figure 6.11 in the top middle. We assume H_1 to be a positive Hopf band and H_2 to be a negative Hopf band. Let $L(d, l)$ be the enriched fibred link representative, where l is the even length of the cycle, and $d \geq 0$ is the flow difference.

We use the skein relation (SR) for the enriched fibred link representative $L(d, l) = L(d, l)_+$ at a positive crossing of H_1 . This yields

$$\Delta_{L(d, l)} = \Delta_{L(d, l)_-} + \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right) \Delta_{L(d, l)_0}.$$

The links $L_+ = L(d, l)$, $L_- = L(d, l)_-$ and $L_0 = L(d, l)_0$ appearing in the skein relation are depicted in Figure 6.11. We note that L_- is given by a plumbing of Hopf bands along a tree. In particular, the Alexander polynomial Δ_{L_-} does not depend on the order of twisting, and, in particular, does not depend on the flow difference. From this, we deduce

$$\Delta_{L(d, l)} - \Delta_{L(d+2, l)} = \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right) (\Delta_{L(d, l)_0} - \Delta_{L(d+2, l)_0}).$$

The links $L = L(d, l)_0$ and $L' = L(d+2, l)_0$ are fibred link representatives of enriched cycles of even length l and flow difference $d \geq 0$ and $d + 2$, respectively. Furthermore, the extra band H_2 corresponding to the extra vertex which is negative. We can get rid of the band H_2 by another use of the skein relation, similarly to the skein relation we used to get rid of H_1 , but we have to take care of the change in sign of the band. Using the skein relation on a negative crossing of H_2 , which is negative, we obtain

$$\Delta_L = \Delta_{L_+} - \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right) \Delta_{L_0}$$




FIGURE 6.11. The links appearing in the skein relation (SR) for the fibred link representations of enriched cycles, compare also with Figure 6.6. The crossing used in the skein relation is one of the crossings of the band H_1 .

and

$$\Delta_{L'} = \Delta_{L'_+} - \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right) \Delta_{L'_0},$$

respectively. We note that again, the links L_+ and L'_+ are given by a plumbing of Hopf bands along a forest. In particular, we have $\Delta_{L_+} = \Delta_{L'_+}$,

and this yields

$$\Delta_L - \Delta_{L'} = - \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right) \left(\Delta_{L_0} - \Delta_{L'_0} \right),$$

where Δ_{L_0} equals $\Delta_{d,l}$ and $\Delta_{L'_0}$ equals $\Delta_{d+2,l}$. Applying Proposition 6.21, this gives

$$\begin{aligned} \Delta_{L(l,d)} - \Delta_{L(l,d+2)} &= - \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right)^2 \left(\Delta_{L_0} - \Delta_{L'_0} \right) \\ &= - \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right)^3 \left(\sqrt{t}^{d+1} - \sqrt{t}^{-(d+1)} \right). \end{aligned}$$

Now, let $\chi_{d,l}$ and $\chi_{d+2,l}$ be the characteristic polynomials of the action induced on the first homology by the monodromies of the fibred link representatives $L(d+2,l)$ and $L(d+2,l)$, respectively.

We note that the leading coefficient of both $\Delta_{L(d,l)}$ and $\Delta_{L(d+2,l)}$ is -1 . Indeed, we have used the skein relation on one positive and one negative crossing to go from $\Delta_{L(d,l)}$ and $\Delta_{L(d+2,l)}$ to $\Delta_{d,l}$ and $\Delta_{d+2,l}$, respectively, which have leading coefficient $+1$. Tracking the sign of the leading coefficient through the two skein relations yields that it switches. In particular, the leading coefficients of $\Delta_{L(d,l)}$ and $\Delta_{L(d+2,l)}$ are -1 , since the leading coefficients of $\Delta_{d,l}$ and $\Delta_{d+2,l}$ are $+1$ by Lemma 6.25. This means that to normalise the Alexander polynomials $\Delta_{L(d,l)}$ and $\Delta_{L(d+2,l)}$ to the characteristic polynomials $\chi_{d,l}$ and $\chi_{d+2,l}$ of the action induced on first homology by the monodromies of $L(d,l)$ and $L(d+2,l)$, respectively, we have to multiply by $-(\sqrt{t})^{l+2}$. This yields a difference of

$$\begin{aligned} \chi_{d,l} - \chi_{d+2,l} &= \left(\sqrt{t} \right)^{l+2} \left(\sqrt{t} - \frac{1}{\sqrt{t}} \right)^3 \left(\sqrt{t}^{d+1} - \sqrt{t}^{-(d+1)} \right) \\ &= (t-1)^3 \left(t^{\frac{l+d}{2}} - t^{\frac{l-d-2}{2}} \right). \end{aligned}$$

Clearly, this difference is strictly positive for any real number $t > 1$. In particular, when evaluated at real numbers strictly greater than 1, the characteristic polynomial of the action corresponding to the flow difference $d \geq 0$ is strictly greater than the characteristic polynomial of the action corresponding to the flow difference $d+2$. This implies that the largest real root is strictly greater for the characteristic polynomial of the action corresponding to the flow difference $d+2$, and thus proves the claim. \square

4.5. A monotonicity criterion and a proof of Lemma 6.19. Let ϕ be a mapping class obtained from Penner's construction using curves with

intersection graph Γ , and such that every curve gets twisted along exactly once. We further assume Γ to have only simple edges. The intersection graph Γ is equipped with an acyclic orientation given by the order in which the curves used in Penner's construction get twisted along. In this context, the matrix $\rho(\phi)$ associated with ϕ in Penner's construction (Theorem 1.25) only depends on the intersection graph Γ and the acyclic orientation. In particular, the same is true for the stretch factor, and we consider both $\rho(\phi)$ and $\lambda(\phi)$ as a function of the intersection graph Γ with its acyclic orientation. We write $\rho(\Gamma)$ and $\lambda(\Gamma)$, respectively.

Let $y \in \mathbb{R}^n$ with coefficients y_i (corresponding to the vertices v_i of Γ) be a Perron–Frobenius eigenvector of the matrix $\rho(\Gamma)$. This is a vector whose entries satisfy the following set of equations. For each vertex v_i of Γ , let v_{o_1}, \dots, v_{o_r} be the vertices of Γ connected to v_i by an edge pointing away from v_i . Similarly, let v_{i_1}, \dots, v_{i_q} be the vertices of Γ connected to v_i by an edge pointing towards v_i . Then the matrix $\rho(\Gamma)$ defined in Penner's construction acts on an arbitrary vector $x \in \mathbb{R}^n$ as follows:

$$(P) \quad \begin{aligned} (\rho(\phi)x)_i &= x_i \\ &+ x_{o_1} + \cdots + x_{o_r} \\ &+ (\rho(\phi)x)_{i_1} + \cdots + (\rho(\phi)x)_{i_q}. \end{aligned}$$

In other words, to a weight we add all the weights adjacent in the graph, and we do this to all the weights in the order given by the acyclic orientation. In particular, the Perron–Frobenius eigenvector y of $\rho(\Gamma)$ satisfies the equation

$$(PF) \quad \begin{aligned} \lambda y_i &= y_i \\ &+ y_{o_1} + \cdots + y_{o_r} \\ &+ \lambda(y_{i_1} + \cdots + y_{i_q}). \end{aligned}$$

PROPOSITION 6.26. *Assume that locally around a vertex v_i the acyclic orientation of Γ looks like a source-sink path as in Figure 6.12 on the left. Assume furthermore that the coefficient y_i of the Perron–Frobenius eigenvector y corresponding to v_i is smaller than or equal to the coefficients y_{i-2} and y_{i+2} corresponding to the two outer vertices. Then, locally prolonging the path by two vertices as shown in Figure 6.12 on the right yields an acyclically oriented graph Γ' with associated Penner stretch factor $\lambda'(\Gamma)$ such that $\lambda'(\Gamma) \leq \lambda(\Gamma)$.*




FIGURE 6.12.

The following lemma is a standard description of the Perron–Frobenius eigenvalue of a Perron–Frobenius matrix.

LEMMA 6.27. *Let A be a Perron–Frobenius matrix of size $n \times n$, and let λ be its Perron–Frobenius eigenvalue. Then*

$$\lambda = \min \left(\max_{x_i \neq 0} \left(\frac{(Ax)_i}{x_i} \right) \right),$$

where the minimum is taken over all nonnegative vectors $x \in \mathbb{R}^n \setminus \{0\}$.

PROOF OF PROPOSITION 6.26. Let $\rho(\Gamma')$ be the matrix with dimension $(n+2) \times (n+2)$ associated with the acyclically oriented graph Γ' by Penner’s construction. We will describe a nonnegative vector $x \in \mathbb{R}^{n+2}$ such that for each coefficient x_i , we have

$$\frac{(\rho(\Gamma')x)_i}{x_i} \leq \lambda(\Gamma).$$

It then follows from Lemma 6.27 that the Perron–Frobenius eigenvalue $\lambda(\Gamma')$ associated with Γ' is bounded from above by $\lambda(\Gamma)$.

Let $x \in \mathbb{R}^{n+2}$ be the vector with entries as shown in Figure 6.13 on the right, where we assume that all the entries corresponding to vertices outside the local picture are equal to the corresponding entry y_j of the Perron–Frobenius eigenvector y of $\rho(\Gamma)$.

Except for the five middle vertices, all the entries corresponding to the vertices on the right satisfy the exact same equations (P), and thus satisfy also the respective equations (PF). In particular, the corresponding entries x_j satisfy

$$\frac{(\rho(\Gamma')x)_j}{x_j} = \lambda.$$

We still have to show

$$(*) \quad \frac{(\rho(\Gamma')x)_j}{x_j} \leq \lambda$$




FIGURE 6.13. The graphs Γ on the left and the graph Γ' on the right, obtained by a local modification around the vertex v_i .

for the entries $x_{i-1}, x_i, x_{i+1}, x_{n+1}$ and x_{n+2} corresponding to the five middle vertices. Using the equations (P) and (PF) as well as the assumption

$$y_i \leq \min(y_{i-2}, y_{i+2}),$$

we now verify this by direct computation. We first note that

$$\begin{aligned} (\rho(\Gamma')x)_{n+1} &= x_{n+1} + x_{i-1} + x_i \\ &= y_i + y_{i-1} + \min(y_{i-1}, y_{i+1}) \\ &\leq \lambda(\Gamma)y_i = \lambda(\Gamma)x_{n+1}, \end{aligned}$$

which proves (*) for x_{n+1} . The analogue computation for x_{n+2} also yields

$$(\rho(\Gamma')x)_{n+2} \leq \lambda(\Gamma)x_{n+2}.$$

Similarly, for x_{i-1} , we have

$$\begin{aligned} (\rho(\Gamma')x)_{i-1} &= x_{i-1} + (\rho(\Gamma')x)_{n+1} + (\rho(\Gamma')x)_{i-2} \\ &= y_{i-1} + (\rho(\Gamma')x)_{n+1} + (\rho(\Gamma')x)_{i-2} \\ &\leq y_{i-1} + \lambda(\Gamma)y_i + \lambda(\Gamma)y_{i-2} \\ &= \lambda(\Gamma)y_{i-1} = \lambda(\Gamma)x_{i-1}, \end{aligned}$$

which proves (*) for x_{i-1} . The analogue computation for x_{i+1} yields

$$(\rho(\Gamma')x)_{i+1} \leq \lambda(\Gamma)x_{i+1}.$$

Finally, for x_i , we have

$$\begin{aligned}
(\rho(\Gamma')x)_i &= x_i + (\rho(\Gamma')x)_{n+1} + (\rho(\Gamma')x)_{n+2} \\
&\leq \min(y_{i-1}, y_{i+1}) + 2\lambda(\Gamma)y_i \\
&\leq \min(y_{i-1}, y_{i+1}) + \lambda(\Gamma)y_i + \lambda(\Gamma)\min(y_{i-2}, y_{i+2}) \\
&\leq \min(\lambda(\Gamma)y_{i-1}, \lambda(\Gamma)y_{i+1}) \\
&= \lambda(\Gamma)\min(y_{i-1}, y_{i+1}) = \lambda(\Gamma)x_i,
\end{aligned}$$

which proves (*) for x_i and finishes the proof. \square

We are now ready to show that the sequence of stretch factors (μ_l) of the Penner mapping classes associated with the enriched cycle of odd length l and flow difference 1 is nowhere increasing.

PROOF OF LEMMA 6.19. For $l = 3, 5, 7, 9, 11, 13$ we simply check the statement by hand (on a computer) and notice that $\mu_{13} < 6.13$. The results of the calculation are given in Table 6.2. We now proceed by induction

| l | μ_l |
|-----|-----------------|
| 3 | ≈ 6.996 |
| 5 | ≈ 6.452 |
| 7 | ≈ 6.277 |
| 9 | ≈ 6.194 |
| 11 | ≈ 6.148 |
| 13 | ≈ 6.120 |

TABLE 6.2. Some values of μ_l .

on the length l of the cycle. Assume we have shown the statement up to cycles of length $\leq 2n - 1$. We want to show $\mu_{2n-1} \geq \mu_{2n+1}$. The idea is to apply Proposition 6.26 to the sink whose corresponding entry of the Perron–Frobenius eigenvector is minimal.

Let P_{2n-1} be the enriched cycle of length $2n - 1$, acyclically oriented as in Figure 6.14, where the edges which are not displayed are oriented in alternating fashion. Clearly, the absolute value of the flow difference equals 1, so the associated Penner stretch factor $\Gamma(P_{2n-1})$ is μ_{2n-1} . Now, let $y \in \mathbb{R}^{2n}$ be the Perron–Frobenius eigenvector for the matrix $\rho(P_{2n-1})$, where its i th entry y_i corresponds to the vertex v_i of P_{2n-1} . We are interested in the minimal entry of y . It is a direct observation that if there is an edge pointing from v_j to v_k , then $y_j \leq y_k$. This follows from the equation (PF). In particular, the minimal entry of y corresponds to a source.




FIGURE 6.14.

If the minimal entry y_i of y corresponds to $v_i \neq v_1$, then we can apply Proposition 6.26. This yields a Penner mapping class on the enriched cycle of length $2n + 1$ with the same flow difference and smaller stretch factor. In particular, we have $\mu_{2n-1} \geq \mu_{2n+1}$ and we are done.

Now assume the minimal entry of y is y_1 . Consider Figure 6.15, which




FIGURE 6.15.

describes an acyclically oriented tree Γ and thus a Penner stretch factor $\lambda(\Gamma)$ and an associated matrix $\rho(\Gamma)$. Furthermore, Figure 6.15 describes a non-negative vector $x \in \mathbb{R}^9$, with entries given by the indicated vertex weights. Using the equations (P) and (PF) and the assumption that y_1 is minimal among the entries of y , one can show that, for any $1 \leq j \leq 9$,

$$(**) \quad \frac{(\rho(\Gamma)x)_j}{x_j} \leq \mu_{2n-1} < 6.13.$$

This can be verified very similar to the calculations in the proof of Lemma 6.19. For all entries except the one with weight y_{2n-1} , the inequality (**) follows

very directly from a comparison with the equation (PF) for the corresponding entry of the Perron–Frobenius eigenvector y . To show (**) for the entry with weight y_{2n-1} , it is necessary to make use of the assumption that y_1 is smaller than y_{2n-2} .

Proving (**) for all entries x_j of the vector x yields a contradiction, since the spectral radius of $\rho(\Gamma)$ can be calculated directly and is strictly larger than 6.13. \square

5. Odd genus minimal stretch factors

Let l be an odd natural number and let ψ_l be the mapping class arising from Penner’s construction using curves with an enriched l -cycle P_l as their intersection graph, and with flow difference 1. The stretch factor of ψ_l equals μ_l . We will use the values of μ_l calculated in Table 6.2. We are now ready to show that the mapping classes ψ_l minimise the stretch factor among mapping classes arising from Penner’s construction on nonorientable surfaces of odd genus. Note that we have to show this for genus greater than or equal to 5, since the genus 3 nonorientable closed surface does not admit pseudo-Anosov mapping classes.

THEOREM 6.28. *The mapping class ψ_l minimises the stretch factor among mapping classes arising from Penner’s construction on the nonorientable closed surface of odd genus $g = l + 2$.*

PROOF. Let N_{l+2} be the nonorientable closed surface of genus $l + 2$, which we assume to be even. We know that there exists the mapping class ψ_l on N_{l+2} , with stretch factor $\mu_l = \lambda(\psi_l)$. Let ϕ be any mapping class on N_{l+2} arising from Penner’s construction, where we assume that every curve used for the construction of ϕ gets twisted along exactly once. Exactly as in the proof of Theorem 6.16, we distinguish cases depending on the intersection graph of the curves used in the construction of ϕ .

Case 1: the intersection graph contains a double edge. We use the same argument as in Case 1 of Theorem 6.16. Let c_1 and c_2 be two curves that intersect at least twice. There must be at least one other curve c_3 intersecting either c_1 or c_2 , and the intersection graph of the curves $\{c_i\}$ contains the tree Γ with three vertices, one double edge and one simple edge as a subgraph, depicted in Figure 6.3 on the left. As in the proof of Theorem 6.16, we obtain $\lambda(\phi) \geq \frac{7+3\sqrt{5}}{2} \approx 6.854$. Note that $\mu_5 \approx 6.452$ but $\mu_3 \approx 6.996$, so the argument works for genus at least 7. In order to accommodate genus 5

in the argument, we need to consider also slightly larger subgraphs than the one used in Case 1 of the proof of Theorem 6.16: there must be at least one other edge, since the intersection graph must contain an odd cycle by Lemma 6.1. More precisely, the intersection graph actually contains one of the four graphs shown in Figure 6.16 as a subgraph. All Penner mapping




FIGURE 6.16.

classes associated with one of those four graphs can directly be shown to have stretch factor bounded from below by 7. Hence, by monotonicity of the spectral radius of nonnegative matrices under “ \leq ”, we can exclude double edges of the intersection graph also for genus 5.

Case 2: the intersection graph contains an odd cycle of length $k \leq l$: We can use the same argument as in the proof of Theorem 6.16, this time invoking the monotonicity lemmas we have proved for enriched cycles. Since an odd cycle cannot fill a nonorientable surface of odd genus by Corollary 6.3, we may assume that the intersection graph contains an enriched cycle of length $k \leq l$ as an induced subgraph. In particular, the stretch factor $\lambda(\phi)$ is bounded from below by the stretch factor of a pseudo-Anosov mapping class arising from Penner’s construction using curves that intersect with the pattern of an enriched odd cycle of length $k \leq l$. In particular, Lemmas 6.18 and 6.19 directly imply $\lambda(\phi) \geq \lambda(\psi_k) \geq \lambda(\psi_l)$.

Case 3: the intersection graph only contains odd cycles of length $k > l$: Take an odd cycle of minimal length $k > l$ among odd cycles. Exactly as in Case 3 of the proof of Theorem 6.16, we may assume this cycle is in fact an induced subgraph of the intersection graph. Hence, by Lemma 6.4, the genus of the surface is bounded from below by $k + 1 > l + 2 = g$, a contradiction. \square

We can now complete the proof of Theorem 1.26.

PROOF OF THEOREM 1.26. All statements of Theorem 1.26 are implied by Theorem 6.17, except for the existence of the limit $\lim_{k \rightarrow \infty} \delta_P(N_{2k+1})$. To show that this limit indeed exists, the only thing we have to note is that the

sequence $\delta_P(N_{2k+1})$ is not increasing in k . But this is a direct consequence of Lemma 6.19, since $\delta_P(N_{2k+1}) = \lambda(\psi_{2k-1}) = \mu_{2k-1}$ by Theorem 6.28. \square

Minimal stretch factors in the fully-punctured case

In this chapter, we prove Theorems 1.29 and 1.31 on the minimal normalised stretch factors of orientation-reversing pseudo-Anosov maps in the fully-punctured setting. The material is copied and adapted from joint work with Lanneau and Tsang [32].

1. Background

In this section, we recall some background material from [22]. The main goal is to explain Theorem 7.1, which is essentially a summary of results by Hironaka and Tsang in [22, Section 3] but allowing for orientation-reversing maps.

1.1. Standardly embedded train tracks. A *train track* τ on a finite-type surface S is an embedded finite graph with a partition of half-edges incident to each vertex into two nonempty subsets. We refer to the data of the partition as the *smoothing*, and incorporate this data by choosing some tangent line at each vertex v , and arranging the half-edges in the two subsets to be tangent to the two sides of the line.

A *boundary component* of τ is a boundary component of a complementary region of τ in S . In this paper, it will always be the case that τ is a deformation retract of S , hence the boundary components of τ are in canonical one-to-one correspondence with the punctures of S .

Let $\partial\tau = \partial_I\tau \sqcup \partial_O\tau$ be a partition of the boundary components of τ into a nonempty set of *inner boundary components* and a nonempty set of *outer boundary components* respectively.

A train track τ is *standardly embedded* (with respect to $(\partial_I\tau, \partial_O\tau)$) if its set of edges $E(\tau)$ can be partitioned into a set of *infinitesimal edges* $E_{\text{inf}}(\tau)$ and a set of *real edges* $E_{\text{real}}(\tau)$, such that:

- The smoothing at each vertex is defined by separating the infinitesimal edges and the real edges.
- The union of infinitesimal edges is a disjoint union of cycles, which we call the *infinitesimal polygons*.

- The infinitesimal polygons are exactly the inner boundary components of τ .

Let τ and τ' be train tracks on S . A *train track map* is a map $f : S \rightarrow S$ that sends vertices of τ to vertices of τ' and smooth edge paths of τ to smooth edge paths of τ' .

Two examples of train track maps are the *subdivision move* on an edge e , and the *elementary folding move* on a pair of edges (e_1, e_2) that are adjacent on one side of a vertex. We pictorially recall the definitions of these in Figure 7.1 top and bottom, and refer the reader to [22, Definitions 3.14 and 3.15] for the precise descriptions.




FIGURE 7.1. Top: The subdivision move on e . Bottom: The elementary folding move on (e_1, e_2) .

The *transition matrix* of a train track map f is the nonnegative integer matrix $f_* \in M_{E(\tau') \times E(\tau)}(\mathbb{Z}_{\geq 0})$ whose (e', e) -entry is given by the number of times $f(e)$ passes through e' .

THEOREM 7.1. *Let $f : S \rightarrow S$ be a fully-punctured pseudo-Anosov map with at least two puncture orbits. Then f is homotopic to a train track map on a standardly embedded train track τ . This train track map, which we will denote by f as well, has the following properties:*

- f is a composition $\sigma f_n \cdots f_1$, where each f_i is a subdivision move or an elementary folding move, and σ is an isomorphism of train tracks.
- The transition matrix of f is of the form

$$f_* = \begin{bmatrix} P & * \\ 0 & f_*^{\text{real}} \end{bmatrix}$$

where P is a permutation matrix and f_*^{real} is a $|\chi(S)|$ -by- $|\chi(S)|$ primitive matrix whose spectral radius equals the stretch factor of f .

PROOF. This is essentially shown over [22, Propositions 3.10, 3.12, 3.13, 3.21]. Now, strictly speaking, [22] only deals with orientation-preserving pseudo-Anosov maps, but this hypothesis is not used in the proofs of these propositions. In the following, we give a quick summary of these proofs for the reader's convenience and to demonstrate that they carry through in the orientation-reversing setting.

Fix a partition $\mathcal{X} = \mathcal{X}_I \sqcup \mathcal{X}_O$ of the set of punctures of S into two nonempty f -invariant subsets. We first construct a partition of S into rectangles: For each $x \in \mathcal{X}_O$, let σ_x^u be the unstable star at x for which each of its prongs has μ^s -length 1. Then for each $x \in \mathcal{X}_I$, we construct a stable star σ_x^s at x by extending the stable prongs at x until it bumps into some σ_x^u . Finally, for each $x \in \mathcal{X}_O$, we extend σ_x^u by extending the prongs of σ_x^u until it bumps into some σ_x^s . It is straightforward to check that each complementary region of $\bigcup_{x \in \mathcal{X}_I} \sigma_x^s \cup \bigcup_{x \in \mathcal{X}_O} \sigma_x^u$ is a rectangle.

From this partition, we can construct the standardly embedded train track τ by taking its set of vertices to be in one-to-one correspondence with pairs of adjacent prongs of the stable stars σ_x^s . The infinitesimal edges of τ are taken to be in one-to-one correspondence with the prongs of σ_x^s , with their endpoints at the two vertices corresponding to the two pairs the prong lies in. The real edges of τ are taken to be in one-to-one correspondence with the rectangles of the partition, with their endpoints at the two vertices corresponding to pairs on which the two stable sides of the rectangle lies along. We illustrate a local picture of this construction in Figure 7.2.




FIGURE 7.2. A local example of the construction going from the partition by rectangles to the standardly embedded train track.

By construction, τ is a deformation retract of S , hence we can compute $\chi(S) = \chi(\tau) = |V| - (|E_{\text{inf}}| + |E_{\text{real}}|) = |V| - (|V| + |E_{\text{real}}|) = -|E_{\text{real}}|$.

Observe that $f^{-1}(\bigcup_{x \in \mathcal{X}_I} \sigma_x^s) \supset \bigcup_{x \in \mathcal{X}_I} \sigma_x^s$, $f^{-1}(\bigcup_{x \in \mathcal{X}_O} \sigma_x^u) \subset \bigcup_{x \in \mathcal{X}_O} \sigma_x^u$. We can choose an increasing sequence of stable stars interpolating from the star $\bigcup_{x \in \mathcal{X}_I} \sigma_x^s$ to $f^{-1}(\bigcup_{x \in \mathcal{X}_I} \sigma_x^s)$ and a decreasing sequence of unstable stars interpolating from $\bigcup_{x \in \mathcal{X}_O} \sigma_x^u$ to $f^{-1}(\bigcup_{x \in \mathcal{X}_O} \sigma_x^u)$, so that at each stage the stable and unstable stars partition S into rectangles. By repeating our construction above, we get a sequence of train tracks that differ by subdivision and elementary folding moves. We illustrate an example of a stage of this sequence in Figure 7.3. Since the last train track in this sequence is constructed from $f^{-1}(\bigcup_{x \in \mathcal{X}_I} \sigma_x^s) \cup f^{-1}(\bigcup_{x \in \mathcal{X}_O} \sigma_x^u)$, f induces a train track isomorphism from this train track back to τ . This shows the first item in the statement.




FIGURE 7.3. As we expand the stable stars and shrink the unstable stars, the train track undergoes subdivision and elementary folding moves.

By construction, f maps each infinitesimal edge to a single infinitesimal edge, hence if we list the infinitesimal edges in front of the real edges, the upper and lower left blocks of the transition matrix are as claimed in the second item of the statement.

It remains to show the claimed properties of the lower right block f_*^{real} . To show primitiveness, notice that the (e', e) -entry of f_*^k is given by the number of times $f^k(e)$ passes through e' . Using the fact that each leaf of ℓ^u is dense in S , one can deduce that for each e' there exists $k_{e'}$ such that the (e', e) -entry of $(f_*^{\text{real}})^{k_{e'}}$ is positive for each e . This implies that $(f_*^{\text{real}})^{\prod k_{e'}}$ is positive, hence f_*^{real} is primitive.

To find the spectral radius of f_* , we define an eigenvector u by taking its e -entry to be the μ^u -measure of the corresponding rectangle R . By definition, the e -entry of $f_*^{\text{real}}u$ is the $f_*\mu^u$ -measure of R , which is λ times the e -entry of u . Hence u is a λ -eigenvector of f_*^{real} . Since u is positive, by the Perron–Frobenius theorem, λ is the spectral radius of f_*^{real} . \square

For future convenience, we will refer to a standardly embedded train track τ as in Theorem 7.1 as a *f-invariant* standardly embedded train track, and we will refer to the block f_*^{real} as the *real transition matrix*.

2. Skew-reciprocity up to cyclotomic factors

Recall that a polynomial $p \in \mathbb{Z}[t]$ is *reciprocal* if its roots are invariant under the transformation $t \mapsto t^{-1}$. A polynomial $q \in \mathbb{Z}[t]$ is *skew-reciprocal* if its roots are invariant under the transformation $t \mapsto -t^{-1}$. A square matrix is said to be *reciprocal/skew-reciprocal* if its characteristic polynomial is reciprocal/skew-reciprocal, respectively.

Hironaka–Tsang [22] show that if f is an orientation-preserving pseudo-Anosov map, then the transition matrix f_* preserves a skew-symmetric bilinear form ω , called the Thurston symplectic form. Despite its name, ω is in general degenerate, but by showing that f_* acts on the radical $\text{rad}(\omega)$ by a cyclotomic, thus reciprocal matrix, one can conclude that f_* is reciprocal nevertheless.

Now if instead f is orientation-reversing, then f_* would send ω to $-\omega$. One might be tempted to conclude from this that f_* would be a skew-reciprocal matrix. However, this is incorrect since it fails to account for the action of f_* on $\text{rad}(\omega)$. What turns out to be true instead is that f_* satisfies a weaker property as follows:

DEFINITION 7.2. A polynomial $r \in \mathbb{Z}[t]$ is *skew-reciprocal up to cyclotomic factors* if its roots, with possibly the exception of roots of unity, are invariant under the transformation $t \mapsto -t^{-1}$. Equivalently, the polynomial r is the product of any number of cyclotomic factors and a skew-reciprocal

polynomial. A square matrix is *skew-reciprocal up to cyclotomic factors* if its characteristic polynomial is skew-reciprocal up to cyclotomic factors.

In Section 2.1, we will explain why f_* and f_*^{real} are skew-reciprocal up to cyclotomic factors. In Section 2.2, we will show a simple necessary condition regarding the coefficients of a polynomial that is skew-reciprocal up to cyclotomic factors.

2.1. The Thurston symplectic form. Let τ be a train track. The *weight space* $\mathcal{W}(\tau)$ of τ is the subspace of $\mathbb{R}^{E(\tau)}$ consisting of elements (w_e) that satisfy $\sum_{e \in E_v^1} w_e = \sum_{e \in E_v^2} w_e$ at every vertex v , where $E_v^1 \sqcup E_v^2$ is the smoothing at v .

DEFINITION 7.3. Let $w, w' \in \mathcal{W}(\tau)$. We define

$$\omega(w, w') = \sum_{v \in V(\tau)} \sum_{e_1 \text{ left of } e_2} (w_{e_1} w'_{e_2} - w_{e_2} w'_{e_1})$$

where the second summation is taken over all pairs of edges (e_1, e_2) on a side of v for which e_1 is on the left of e_2 . Then ω is clearly a skew-symmetric bilinear form on $\mathcal{W}(\tau)$. We call ω the *Thurston symplectic form*.

PROPOSITION 7.4. *Let f be a fully-punctured pseudo-Anosov map with at least two puncture orbits. Let τ be a f -invariant standardly embedded train track. Then there exists a linear map T and a permutation matrix $P \in M_{V(\tau) \times V(\tau)}(\mathbb{R})$ which fit into the commutative diagram*

$$\begin{array}{ccccccc} 0 & \longrightarrow & \mathcal{W}(\tau) & \longrightarrow & \mathbb{R}^{E(\tau)} & \xrightarrow{T} & \mathbb{R}^{V(\tau)} \\ & & \downarrow f_* & & \downarrow f_* & & \downarrow P \\ 0 & \longrightarrow & \mathcal{W}(\tau) & \longrightarrow & \mathbb{R}^{E(\tau)} & \xrightarrow{T} & \mathbb{R}^{V(\tau)} \end{array}$$

Moreover, if f is orientation-reversing, then $f_* : \mathcal{W}(\tau) \rightarrow \mathcal{W}(\tau)$ sends ω to $-\omega$.

PROOF. The first statement follows from [22, Proposition 4.2]. Again, strictly speaking Hironaka and Tsang [22] deal with orientation-preserving maps, but this hypothesis is not used in the proof of this proposition. We give a quick summary of the proof to demonstrate this.

For each vertex v , we define a linear map $T_v : \mathbb{R}^{E(\tau)} \rightarrow \mathbb{R}$ by the following formula: $T_v(w) = \sum_{e \in E_{v,\text{real}}} w_e - \sum_{e \in E_{v,\text{inf}}} w_e$. The map T is then defined by $T(w) = (T_v(w))$. By definition, $\mathcal{W}(\tau) = \ker T$.

We then define

$$P_{v',v} = \begin{cases} 1, & \text{if } f \text{ sends } v \text{ to } v' \\ 0, & \text{otherwise.} \end{cases}$$

It is straightforward to check that $PT = Tf_*$.

For the moreover statement, we use the factorization of f as $\sigma f_n \cdots f_1$ from Theorem 7.1. By [22, Lemmas 4.4 and 4.5], each subdivision move and elementary folding move preserves ω . Meanwhile, the train track isomorphism σ is orientation-reversing if and only if f is orientation-reversing. In this case, σ sends ω to $-\omega$. \square

We can further separate the action of f_* on $\mathcal{W}(\tau)$ into two parts by considering the radical of ω .

Let c be an even-pronged boundary component c of τ . Label the sides of ∂c by I_1, \dots, I_n in a cyclic order. The *radical element* r_c is the element of $\mathcal{W}(\tau)$ where we assign to each edge on I_k a weight of $(-1)^k$.

PROPOSITION 7.5. *Let f be a fully-punctured pseudo-Anosov map with at least two puncture orbits. Let τ be a f -invariant standardly embedded train track. Let $\partial_{\text{even}}\tau$ be the set of even-pronged punctures. Then there exists a signed permutation matrix P which fits into the commutative diagram*

$$\begin{array}{ccccccc} \mathbb{R}^{\partial_{\text{even}}\tau} & \longrightarrow & \mathcal{W}(\tau) & \longrightarrow & \mathcal{W}(\tau)/\text{rad}(\omega) & \longrightarrow & 0 \\ & & \downarrow P & & \downarrow f_* & & \\ \mathbb{R}^{\partial_{\text{even}}\tau} & \longrightarrow & \mathcal{W}(\tau) & \longrightarrow & \mathcal{W}(\tau)/\text{rad}(\omega) & \longrightarrow & 0 \end{array}$$

Moreover, if f is orientation-reversing, then $f_* : \mathcal{W}(\tau)/\text{rad}(\omega) \rightarrow \mathcal{W}(\tau)/\text{rad}(\omega)$ is skew-reciprocal.

PROOF. [22, Proposition 5.5] states that $\text{rad}(\omega) = \text{span}\{r_c\}$. Meanwhile, observe that $f_*(r_c) = \pm r_{f(c)}$. This implies that f_* acts on $\mathbb{R}^{\partial_{\text{even}}\tau}$ via a signed permutation matrix P .

For the moreover statement, ω descends to a symplectic form on the quotient $\mathcal{W}(\tau)/\text{rad}(\omega)$. The map $f_* : \mathcal{W}(\tau)/\text{rad}(\omega) \rightarrow \mathcal{W}(\tau)/\text{rad}(\omega)$ is anti-symplectic with respect to this symplectic form, hence by [39, Proposition 4.2], this matrix is skew-reciprocal. \square

PROPOSITION 7.6. *Let f be an orientation-reversing and fully-punctured pseudo-Anosov map with at least two puncture orbits. Let τ be a f -invariant standardly embedded train track. Then the real transition matrix f_*^{real} is skew-reciprocal up to cyclotomic factors.*

PROOF. The only roots of the characteristic polynomial of a signed permutation matrix are roots of unity. From Proposition 7.5, we deduce that the induced map $f_* : \mathcal{W}(\tau) \rightarrow \mathcal{W}(\tau)$ is skew-reciprocal up to cyclotomic factors. From Proposition 7.4, it follows that $f_* : \mathbb{R}^{E(\tau)} \rightarrow \mathbb{R}^{E(\tau)}$ is skew-reciprocal up to cyclotomic factors. Finally, from the second item in Theorem 7.1, it follows that f_*^{real} is skew-reciprocal up to cyclotomic factors. \square

REMARK 7.7. We exhibit an example that shows that Proposition 7.6 becomes false if we omit ‘up to cyclotomic factors’. Let S be the surface $(\mathbb{R}^2 \setminus (\frac{1}{2}\mathbb{Z})^2) / \mathbb{Z}^2$, which is a torus with four punctures. The matrix $\begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$ induces an orientation-reversing pseudo-Anosov map f on S . It is straightforward to check that f has two puncture orbits. The standardly embedded train track τ illustrated in Figure 7.4 top is a f -invariant train track. In the basis (a, b, c, d) as indicated in Figure 7.4 bottom left, one computes

$$f_*^{\text{real}} = \begin{bmatrix} 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

The characteristic polynomial of f_*^{real} is $t^4 - t^2 - 2t - 1 = (t^2 - t - 1)(t^2 + t + 1)$. Hence in this example, f_*^{real} is skew-reciprocal up to cyclotomic factors but not skew-reciprocal.

2.2. A parity condition. Recall that a polynomial

$$p(t) = a_n t^n + \dots + a_1 t + a_0 \in \mathbb{Z}[t]$$

is reciprocal if and only if $p(t) = \pm t^n p(t^{-1})$. This polynomial equation transforms to the equation $a_d = \pm a_{n-d}$ for all pairs of symmetrically distributed coefficients a_d and a_{n-d} .

Similarly, a polynomial $q(t) = b_m t^m + \dots + b_1 t + b_0 \in \mathbb{Z}[t]$ is skew-reciprocal if and only if $q(t) = \pm t^m q(-t^{-1})$, which in turn implies the equation $b_d = \pm (-1)^{m-d} b_{m-d}$ on the level of coefficients.

In comparison, the coefficients of a polynomial that is skew-symmetric up to cyclotomic factors may not be symmetric up to signs. For example, consider the polynomial $(t^2 - t - 1)(t^2 + t + 1) = t^4 - t^2 - 2t - 1$ as in Remark 7.7.




FIGURE 7.4. An example that shows that Proposition 7.6 becomes false if we omit ‘up to cyclotomic factors’

Nevertheless, we have the following lemma, which states a symmetry condition that holds for any product of a reciprocal and a skew-reciprocal polynomial.

LEMMA 7.8. *Suppose $r(t) = p(t)q(t) = c_k t^k + \dots + c_1 t + c_0$, where the first factor $p \in \mathbb{Z}[t]$ is a reciprocal polynomial and the second factor $q \in \mathbb{Z}[t]$ is a skew-reciprocal polynomial. Then $c_d + c_{k-d}$ is even for all $d = 0, \dots, k$.*

PROOF. Considering the coefficients of both p and q in $\mathbb{Z}/2\mathbb{Z}$, we have the congruences $a_d \equiv a_{n-d} \pmod{2}$ and $b_d \equiv b_{m-d} \pmod{2}$. Hence the coefficients of $r = pq$ also satisfy $c_d \equiv c_{k-d} \pmod{2}$, which means exactly that $c_d + c_{k-d}$ is even. □

Clearly, Lemma 7.8 is only a rough necessary condition for polynomials that are skew-reciprocal up to cyclotomic factors. In Section 3 we will need to refine our analysis in order to obtain an asymptotic classification result for small normalised spectral radii of matrices $A \in \text{GL}_n(\mathbb{Z})$ that are primitive and skew-reciprocal up to cyclotomic factors.

3. Analysis of small curve graphs

In this section, we will show that if $n \geq 4$ then for any primitive matrix $A \in \text{GL}_n(\mathbb{Z})$ that is skew-reciprocal up to cyclotomic factors, we have $\rho(A)^n \geq 3 + 2\sqrt{2}$. Together with Theorem 7.1 and Proposition 7.6, this implies the inequality in Theorem 1.29. The starting point is the following proposition.

PROPOSITION 7.9. *Let $A \in \text{GL}_n(\mathbb{Z})$ be a primitive matrix with spectral radius $\rho(A) > 1$ such that $\rho(A)^n < 3 + 2\sqrt{2}$. Then the reciprocal of the characteristic polynomial of A has one of the following forms:*

$$\begin{aligned}
 (2A_1) \quad & Q(t) = 1 - t^a - t^b \\
 (3A_1) \quad & Q(t) = 1 - t^a - t^b - t^c \\
 (4A_1) \quad & Q(t) = 1 - t^a - t^b - t^c - t^d \\
 (5A_1) \quad & Q(t) = 1 - t^a - t^b - t^c - t^d - t^e \\
 (A_2^*) \quad & Q(t) = 1 - t^a - t^b - t^c + t^{a+b}
 \end{aligned}$$

for suitable positive parameters $a, b, c, d, e \in \mathbb{Z}$.

PROOF. This result is contained in [45]. We briefly indicate how to deduce this exact statement from the technology in that paper.

To a matrix $A \in \text{GL}_n(\mathbb{Z})$ we associate a directed graph Γ that has n vertices and directed edges between vertices according to the coefficients of the matrix A . The associated curve graph G_A has one vertex for every simple closed curve in Γ , where a simple closed curve is a directed closed loop that visits every vertex at most once. Further, two vertices of G_A are connected by an edge if and only if the corresponding simple closed curves have no vertex of Γ in common. Every vertex v of G_A has a weight $w(v)$ describing the number of edges in the associated simple closed curve. It is a well-known result from graph theory that the clique polynomial $Q(t)$ of the weighted graph G_A is the reciprocal of the characteristic polynomial of the matrix $A : Q(t) = t^n \chi_A(t^{-1})$.

It suffices to describe the possible clique polynomials that can arise for the matrix A . This is the content of McMullen’s classification of curve graphs with small growth rate. Indeed, by [45, Theorem 1.6], the only curve graphs that can arise for the matrix A are the graphs nA_1 for $n = 2, \dots, 5$ or A_2^* . To be precise, McMullen’s result lists all possible curve graphs that can arise for $\rho(A)^n < 8$. We then further ruled out those graphs where we cannot

have $\rho(A)^n < 3 + 2\sqrt{2}$. The possible clique polynomials for the remaining curve graphs are given in the statement of the proposition. \square

Proposition 7.9 gives us a list of five different forms of polynomials that we need to consider. We will analyse these one-by-one below. Throughout the analysis we write $\varphi = \frac{1+\sqrt{5}}{2}$ for the golden ratio and $\sigma = 1 + \sqrt{2}$ for the silver ratio.

Case $2A_1$. In this case $Q(t) = 1 - t^a - t^b$. Without loss of generality, we may assume $b = n$. This yields characteristic polynomials of the form

$$P(t) = t^n - t^a - 1.$$

Observe that n must be even, since otherwise by skew-reciprocity up to cyclotomic factors, $P(t)$ must have an odd number of roots (with multiplicity) on the unit circle, thus $P(t)$ must have 1 or -1 as a root, yet $P(\pm 1)$ is an odd integer so this is impossible. We write $n = 2g$, $g \geq 2$. The only way for $P(t)$ to satisfy the parity condition of Lemma 7.8 is if $a = g$. But then the polynomial is a polynomial in t^g and A would not be primitive.

REMARK 7.10. Here we used the assumption that $n \geq 4$ to have $g \geq 2$, thus obtain a contradiction to A being primitive. When $n = 2$, we can have $P(t) = t^2 - t - 1$, whose normalised largest root is $\varphi^2 < 3 + 2\sqrt{2}$. In fact, this characteristic polynomial is realised by the orientation-reversing fully-punctured pseudo-Anosov map obtained as the action of $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ on the 4-punctured sphere, which has two puncture orbits.

Case $3A_1$. In this case $Q(t) = 1 - t^a - t^b - t^c$, and we may assume $c = n$. This yields characteristic polynomials of the form $P(t) = t^n - t^a - t^b - 1$. There are two ways to satisfy the parity condition of Lemma 7.8: either $a + b = n$ or $a = b$. Hence the characteristic polynomial has one of the following two forms:

- (1) $P(t) = t^n - t^{\frac{n}{2}+d} - t^{\frac{n}{2}-d} - 1$, where $0 \leq d < \frac{n}{2}$, or
- (2) $P(t) = t^n - 2t^a - 1$, where $0 \leq a < n$.

In the former case, we claim that the largest root of $P(t)$ is strictly increasing in d .

PROOF OF CLAIM. We introduce the function

$$h(x, s) = x^2 - x^{1+s} - x^{1-s} - 1$$

defined on the domain $C = \{(x, s) \in \mathbb{R}^2 : x > 1, 0 \leq s < 1\}$. Observe that $x = \lambda^{\frac{n}{2}}$ is the largest real root of $h(x, s)$ with $s = \frac{2d}{n} < 1$. For any fixed s , notice that

$$\partial_x^2 h(x, s) = 2 - (1 + s)sx^{s-1} + (1 - s)sx^{-s-1} > 0.$$

Thus $\partial_x h(x, s)$ is strictly increasing in x . Furthermore, we have the divergence $\lim_{x \rightarrow \infty} \partial_x h(x, s) = +\infty$ and $h(1, s) = -2 < 0$, hence the equation $h(x, s) = 0$ has a unique real solution $x(s) > 1$ depending continuously on s . Meanwhile, $\partial_s h(x, s) \leq 0$ on C , where equality holds if and only if $s = 0$. Since $\partial_x h(x(s), s) > 0$ for any s , we have $x'(s) \geq 0$, equality holds if and only if $s = 0$. Thus $x(s)$ is strictly increasing in s . \square

In particular, the minimal largest root is attained at $d = 0$, that is, when $P(t) = t^n - 2t^{\frac{n}{2}} - 1$ has normalised largest root $3 + 2\sqrt{2}$.

In the latter case, if $P(t)$ is skew-reciprocal, then we must have n even and $a = \frac{n}{2}$, which would reduce to the former case. Hence we can assume that $P(t)$ is skew-reciprocal up to cyclotomic factors but not skew-reciprocal. Then there must exist a root of unity ξ that is a zero: $\xi^n - 2\xi^a = 1$. The only way for this to happen is if $\xi^n = \xi^a = -1$. Since A is primitive, n and a are coprime, hence $\xi = -1$ and n and a are odd. Writing the polynomial as $P(t) = t^a(t^{n-a} - 1) - (t^a + 1)$, we compute that the quotient of $P(t)$ by $t + 1$ to be

$$R(t) = t^a(t^{n-a-1} - t^{n-a-2} + \dots - 1) - (t^{a-1} - t^{a-2} + \dots + 1)$$

Observe that $R(t)$ has no roots of unity as a root. Indeed, by the reasoning above, the only possible root is $\xi = -1$, yet $R(-1) = (n - a) - a = n - 2a$ is an odd integer. Meanwhile, by inspecting the coefficients of $R(t)$,

$$R(t) = \begin{cases} t^{n-1} - t^{n-2} + \dots + t - 1 & \text{if } a \geq 3 \\ t^{n-1} - t^{n-2} + t^{n-3} - \dots + t^2 - t - 1 & \text{if } a = 1 \end{cases}.$$

we see that $R(t)$ is not skew-reciprocal, giving us a contradiction.

REMARK 7.11. In the analysis above, we use the assumption $n \geq 4$ to deduce that $R(t)$ is not skew-reciprocal. When $n = 3$ and $a = 1$, the polynomial $R(t) = t^2 - t - 1$ is actually skew-reciprocal. Correspondingly, the characteristic polynomial $P(t) = t^3 - 2t - 1$ is skew-reciprocal up to cyclotomic factors, and has normalised largest root $\varphi^3 < 3 + 2\sqrt{2}$. It is not clear to us at the moment whether this characteristic polynomial is attained

by any orientation-reversing fully-punctured pseudo-Anosov map with at least two puncture orbits.

Case 4A₁. In this case $Q(t) = 1 - t^a - t^b - t^c - t^d$, and we may assume that $d = n$. This yields characteristic polynomials of the form

$$P(t) = t^n - t^a - t^b - t^c - 1.$$

As in case 2A₁, n must be even. Let us write $n = 2g$.

There are three coefficients other than the ones associated to t^{2g} and -1 . The parity condition of Lemma 7.8 implies that one of them must be the middle coefficient: say $b = g$. There are two ways to satisfy the parity condition of Lemma 7.8: either $a + c = 2g$ or $a = c$. Hence the characteristic polynomial has one of the following two forms:

- (1) $P(t) = t^{2g} - t^{g+d} - t^g - t^{g-d} - 1$, where $0 \leq d < g$, or
- (2) $P(t) = t^{2g} - 2t^a - t^g - 1$, where $0 < a < 2g$.

In the former case, we claim that the largest root of $P(t)$ is strictly increasing in d .

PROOF OF CLAIM. We introduce the function

$$h(x, s) = x^2 - x^{1+s} - x - x^{1-s} - 1$$

defined on the domain $C = \{(x, s) \in \mathbb{R}^2 : x > 1, 0 \leq s < 1\}$. Observe that $x = \lambda^g$ is the largest real root of $h(x, s)$ with $s = \frac{d}{g} < 1$. For any fixed s , notice that

$$\partial_x^2 h(x, s) = 2 - (1 + s)sx^{s-1} + (1 - s)sx^{-s-1} > 0.$$

Thus $\partial_x h(x, s)$ is strictly increasing in x . Again, $\lim_{x \rightarrow \infty} \partial_x h(x, s) = +\infty$ and $h(1, s) = -3 < 0$, hence the equation $h(x, s) = 0$ has a unique real solution $x(s) > 1$ depending continuously on s . Meanwhile, $\partial_s h(x, s) \leq 0$ on C , where equality holds if and only if $s = 0$. Since $\partial_x h(x(s), s) > 0$ for any s , we have $x'(s) \geq 0$, equality holds if and only if $s = 0$. Thus $x(s)$ is strictly increasing in s . \square

In particular, the minimal largest root λ is attained at $d = 0$, that is, when $P(t) = t^{2g} - 3t^g - 1$ has normalised largest root $\left(\frac{3 + \sqrt{13}}{2}\right)^2 > 3 + 2\sqrt{2}$.

In the latter case, if $P(t)$ is skew-reciprocal, then we must have $a = g$, which belongs to the former case. Hence we can assume that $P(t)$ is skew-reciprocal up to cyclotomic factors but not skew-reciprocal. Then there

must exist a root of unity ξ that is a zero: $\xi^{2g} - \xi^g - 1 = 2\xi^a$. We claim that the only way this can happen is if ξ is a sixth root of unity.

To see this, we write $\xi^g = e^{i\theta}$ and take the norm of both sides of the equation

$$\begin{aligned} 4 &= (\cos 2\theta - \cos \theta - 1)^2 + (\sin 2\theta - \sin \theta)^2 \\ &= (\cos^2 2\theta + \cos^2 \theta + 1 + 2 \cos \theta - 2 \cos 2\theta - 2 \cos 2\theta \cos \theta) \\ &\quad + (\sin^2 2\theta + \sin^2 \theta - 2 \sin 2\theta \sin \theta) \\ &= 3 + 2 \cos \theta - 2 \cos 2\theta - 2 \cos(2\theta - \theta) \\ &= 3 - 2 \cos 2\theta. \end{aligned}$$

Thus $\cos 2\theta = -\frac{1}{2}$ and $\theta \in \{\frac{\pi}{3}, \frac{2\pi}{3}, \frac{4\pi}{3}, \frac{5\pi}{3}\}$. For each of the four possibilities for θ , we have:

| θ | ξ^g | ξ^{2g} | $\xi^a = \frac{\xi^{2g} - \xi^g - 1}{2}$ |
|------------------|-----------------------|-----------------------|--|
| $\frac{2\pi}{3}$ | $e^{\frac{2i\pi}{3}}$ | $e^{\frac{4i\pi}{3}}$ | $e^{\frac{4i\pi}{3}}$ |
| $\frac{4\pi}{3}$ | $e^{\frac{4i\pi}{3}}$ | $e^{\frac{2i\pi}{3}}$ | $e^{\frac{2i\pi}{3}}$ |
| $\frac{\pi}{3}$ | $e^{\frac{i\pi}{3}}$ | $e^{\frac{2i\pi}{3}}$ | -1 |
| $\frac{5\pi}{3}$ | $e^{\frac{5i\pi}{3}}$ | $e^{\frac{4i\pi}{3}}$ | -1 |

Suppose ξ is a primitive q^{th} root of unity. For the first two cases, since $\xi^{3g} = \xi^{3a} = 1$, we have $q | \gcd(3g, 3a) = 3$, i.e. $q = 3$. In turn, we have $a = 3A - k$ and $g = 3G + k$ for $k = 1$ or 2 . Similarly for the last two cases, since $\xi^{6g} = \xi^{2a} = 1$, we have $q | \gcd(6g, 2a) \leq 6$. But ξ^g is a primitive sixth root of unity, so in fact $q = 6$. In turn, we have $a = 6A + 3$ and $g = 6G + k$ for $k = \pm 1$.

In particular, this reasoning shows that the first two and the last two cases are mutually disjoint. Correspondingly, $P(t)$ is in one of two possible forms:

- (i) $P(t) = t^{6G+2k} - 2t^{3A-k} - t^{3G+k} - 1$, where $k = 1$ or 2 , and where we have $1 \leq A < 2G + k$, or
- (ii) $P(t) = t^{12G+2k} - 2t^{6A+3} - t^{6G+k} - 1$, where $k = \pm 1$, and $0 \leq A < 2G$.

For case (i), we claim that if $P(t)$ is skew-reciprocal up to cyclotomic factors but not skew-reciprocal, then $k = 2$, $G = 0$, and $A = 1$, that is, $P(t) = t^4 - t^2 - 2t - 1 = (t^2 - t - 1)(t^2 + t + 1)$, compare with Remark 7.7, which has normalised largest root $\varphi^4 > 3 + 2\sqrt{2}$.

PROOF OF CLAIM. Assume $P(t)$ is skew-reciprocal up to cyclotomic factors, but not skew-reciprocal. We reasoned above that it has a cubic root of

unity as a root. We will compute the quotient $R(t)$ of $P(t)$ by $t^2 + t + 1$.

$$P(t) = 2(t^3 - 1) \frac{(t^{6G+2k} - t^{3A-k})}{t^3 - 1} - (t^{6G+2k} + t^{3G+k} + 1).$$

For the first term, we get

$$2(t^3 - 1) \frac{(t^{6G+2k} - t^{3A-k})}{t^3 - 1} = 2t^{3A-k} (t^3 - 1) (t^{6G-3A+3k-3} + \dots + t^3 + 1).$$

On the other hand, for the second term, we get

$$\begin{aligned} t^{6G+2k} + t^{3G+k} + 1 &= (t^{6G+2k} - t^{3G+2k}) - (t^{3G} - 1) \\ &\quad + (t^{3G+2k} + t^{3G+k} + t^{3G}) \\ &= t^{3G+2k} (t^{3G} - 1) - (t^{3G} - 1) + t^{3G} (t^{2k} + t^k + 1) \\ &= (t^{3G} - 1) (t^{3G+2k} - 1) + t^{3G} (t^{2k} + t^k + 1) \\ &= (t^3 - 1) (t^{3G-3} + \dots + t^3 + 1) (t^{3G+2k} - 1) \\ &\quad + t^{3G} (t^{2k} + t^k + 1). \end{aligned}$$

Hence

$$\begin{aligned} R(t) &= 2t^{3A-k} (t - 1) (t^{6G-3A+3k-3} + \dots + t^3 + 1) \\ &\quad - (t - 1) (t^{3G-3} + \dots + t^3 + 1) (t^{3G+2k} - 1) - t^{3G} \frac{t^{2k} + t^k + 1}{t^2 + t + 1}. \end{aligned}$$

We now check that $R(t)$ has no roots of unity as a root. Indeed, by the reasoning above, if ξ is a root then it is a third root of unity. However, we compute

$$\begin{aligned} R(\xi) &= \begin{cases} 2(1 - \xi^2)(2G - A + 1) - 3G - 1 & \text{if } k = 1 \\ 2(\xi^2 - \xi)(2G - A + 2) + 3\xi G + 2\xi & \text{if } k = 2 \end{cases} \\ &\neq 0. \end{aligned}$$

Thus we only need to show that $R(t)$ is not skew-reciprocal, except for the specified case.

For $k = 1$,

$$\begin{aligned} R(t) &= 2 \sum_{i=A}^{2G} (t^{3i} - t^{3i-1}) - t^{3G} - \sum_{i=0}^{G-1} (t^{3i+1} - t^{3i}) (t^{3G+2} - 1) \\ &= t^{6G} - t^{6G-1} + \dots + t - 1 \end{aligned}$$

is not skew-reciprocal.

For $k = 2$ and $G > 0$,

$$\begin{aligned} R(t) &= 2 \sum_{i=A}^{2G+1} (t^{3i-1} - t^{3i-2}) - (t^2 - t + 1)t^{3G} - \sum_{i=0}^{G-1} (t^{3i+1} - t^{3i})(t^{3G+4} - 1) \\ &= \begin{cases} t^{6G+2} - t^{6G+1} + \dots + t - 1 & \text{if } A > 1 \\ t^{6G+2} - t^{6G+1} + 0t^{6G} + \dots + 2t^2 - t - 1 & \text{if } A = 1 \end{cases} \end{aligned}$$

is not skew-reciprocal. If $G = 0$, then $A = 1$ and $R(t) = t^2 - t - 1$ is actually skew-reciprocal. \square

In case (ii), we claim that $P(t)$ is never skew-reciprocal up to cyclotomic factors.

PROOF OF CLAIM. We follow the same strategy as in case (i). We first compute the quotient of $P(t)$ by $t^2 - t + 1$.

$$P(t) = t^{12G+2k} - 2t^{6A+3} - t^{6G+k} - 1 = (t^{12G+2k} - t^{6G+k} + 1) - 2(t^{6A+3} + 1).$$

For the first term:

$$\begin{aligned} t^{12G+2k} - t^{6G+k} + 1 &= (t^{12G+2k} - t^{6G+2k}) - (t^{6G} - 1) \\ &\quad + (t^{6G+2k} - t^{6G+k} + t^{6G}) \\ &= (t^{6G} - 1)(t^{6G+2k} - 1) + t^{6G+k-1}(t^{1+k} - t + t^{1-k}) \\ &= (t^{6G-3} - \dots + t^3 - 1)(t^3 + 1)(t^{6G+2k} - 1) \\ &\quad + t^{6G+k-1}(t^2 - t + 1). \end{aligned}$$

For the second term:

$$t^{6A+3} + 1 = (t^{6A} - \dots - t^3 + 1)(t^3 + 1).$$

Hence

$$\begin{aligned} R(t) &= (t^{6G-3} - \dots + t^3 - 1)(t + 1)(t^{6G+2k} - 1) + t^{6G+k-1} \\ &\quad - 2(t^{6A} - \dots - t^3 + 1)(t + 1). \end{aligned}$$

We then check that $R(t)$ has no root of unity as a root. The only possible root ξ satisfies $\xi^3 = -1$, and

$$\begin{aligned} R(\xi) &= \begin{cases} 6\xi G - \xi - 2(2A + 1)(\xi + 1) & \text{if } k = -1 \\ 6G + 1 - 2(2A + 1)(\xi + 1) & \text{if } k = 1 \end{cases} \\ &\neq 0 \end{aligned}$$

Finally we observe that

$$\begin{aligned} R(t) &= \sum_{i=0}^{2G-1} (-1)^{i+1} (t^{3i+1} + t^{3i}) (t^{6G+2k} - 1) + t^{6G+k-1} \\ &\quad - 2 \sum_{i=0}^{2A} (-1)^i (t^{3i+1} + t^{3i}) \\ &= \begin{cases} t^{12G+2k-2} + t^{12G+2k-3} - \dots - t - 1 \\ t^{12G-4} - t^{12G-5} - 2t^{12G-6} - \dots + 0t^2 - t - 1, \end{cases} \end{aligned}$$

where the first case corresponds to $k = 1$ or $A < 2G - 1$ and the second case corresponds to $k = -1$ and $A = 2G - 1$, is not skew-reciprocal. \square

Case 5A₁. In this case, $Q(t) = 1 - t^a - t^b - t^c - t^d - t^e$, and we may assume $e = n$. This yields characteristic polynomials of the form

$$P(t) = t^n - t^a - t^b - t^c - t^d - 1.$$

To satisfy the parity condition of Lemma 7.8, the characteristic polynomial has one of the three possible forms

- (1) $P(t) = t^n - t^{\frac{n}{2}+a} - t^{\frac{n}{2}+b} - t^{\frac{n}{2}-b} - t^{\frac{n}{2}-a} - 1$, for $0 \leq a, b < \frac{n}{2}$, or
- (2) $P(t) = t^n - 2t^a - t^{\frac{n}{2}+b} - t^{\frac{n}{2}-b} - 1$, for $0 < a < n$, $0 \leq b < \frac{n}{2}$, or
- (3) $P(t) = t^n - 2t^a - 2t^b - 1$, for $0 < a, b < n$.

In the first case, we claim that the largest root of $P(t)$ is strictly increasing in a and b .

PROOF OF CLAIM. We introduce the function

$$h(x, s, u) = x^2 - x^{1+s} - x^{1+u} - x^{1-s} - x^{1-u} - 1$$

defined on the domain $C = \{(x, s, u) \in \mathbb{R}^3 : x > 1, 0 \leq s, u < 1\}$. We observe that $x = \lambda^{\frac{n}{2}}$ is the largest real root of $h(x, s, u)$ with $s = \frac{2a}{n} < 1$ and $u = \frac{2b}{n} < 1$. For any fixed (s, u) and $x > 1$,

$$\partial_x^3 h(x, s, u) = (1+s)s(1-s)(x^{s-2} - x^{-s-2}) + (1+u)u(1-u)(x^{u-2} - x^{-u-2}) > 0.$$

Hence $\partial_x^2 h(x, s, u)$ is strictly increasing in x . Moreover, we directly calculate $\partial_x h(1, s, u) = 2 - (1+s) - (1+u) - (1-s) - (1-u) = -2 < 0$ and also $\partial_x h(x, s, u) > 0$ for large enough $x > 1$. Thus $\partial_x h(\cdot, s, u)$ has a unique zero $x_0 > 1$, and $h(\cdot, s, u)$ is decreasing on $(1, x_0)$ and increasing on (x_0, ∞) . Since $h(1, s, u) = -4 < 0$ and $\lim_{x \rightarrow \infty} h(x, s, u) = +\infty$ this proves that the equation $h(x, s, u) = 0$ has a unique real solution $x(s, u) > x_0 > 1$, depending continuously on (s, u) . Also, $\partial_s h(x, s, u) \leq 0$ and $\partial_u h(x, s, u) \leq 0$,

with equality if and only if $s = 0$ and $u = 0$ respectively, hence by the implicit function theorem $x(s, u)$ is strictly increasing in s and u . \square

In particular, the minimal largest root λ of $P(t)$ is attained at $a = b = 0$, when $P(t) = t^n - 4t^{\frac{n}{2}} - 1$ has normalised largest root $(2 + \sqrt{5})^2 > 3 + 2\sqrt{2}$.

In the second case, we claim that the largest root of $P(t)$ is strictly increasing in a and b .

PROOF OF CLAIM. We introduce the function

$$h(x, s, u) = x^2 - 2x^s - x^{1+u} - x^{1-u} - 1$$

defined on the domain $C = \{(x, s, u) \in \mathbb{R}^3 : x > 1, 0 < s < 2, 0 \leq u < 1\}$. Observe that $x = \lambda^{\frac{n}{2}}$ is the largest real root of $h(x, s, u)$ with $s = \frac{2a}{n}$ and $u = \frac{2b}{n}$. We further separate C into $C_1 = \{(x, s, u) \in C : 0 < s \leq 1\}$ and $C_2 = \{(x, s, u) \in C : 1 \leq s < 2\}$.

For any fixed $(s, u) \in C_1$,

$$\partial_x^2 h(x, s, u) = 2 + 2s(1 - s)x^{s-2} - (1 + u)ux^{u-1} + (1 - u)ux^{-u-1} > 0.$$

Thus $\partial_x h(x, s, u)$ is strictly increasing in x . Furthermore, we have divergence $\lim_{x \rightarrow \infty} \partial_x h(x, s, u) = +\infty$ and $h(1, s, u) = -4 < 0$, hence the equation $h(x, s, u) = 0$ has a unique real solution $x(s, u) > 1$ depending continuously on s, u . Furthermore, $\partial_s h(x, s, u) < 0$ and $\partial_u h(x, s, u) \leq 0$, with equality if and only if $u = 0$, hence by the implicit function theorem $x(s, u)$ is strictly increasing in s and u in C_1 .

For any fixed $(s, u) \in C_2$,

$$\partial_x^3 h(x, s, u) = 2s(s - 1)(2 - s)x^{s-3} + (1 + u)u(1 - u)(x^{u-2} - x^{-u-2}) > 0.$$

Hence $\partial_x^2 h(x, s, u)$ is strictly increasing in x . Moreover, we directly calculate $\partial_x h(1, s, u) = 2 - 2s - (1 + u) - (1 - u) = -2s < 0$ and $\partial_x h(x, s, u) > 0$ for large enough $x > 1$. Thus $\partial_x h(\cdot, s, u)$ has a unique zero $x_0 > 1$, and $h(\cdot, s, u)$ is decreasing on $(1, x_0)$ and increasing on (x_0, ∞) . Since $h(1, s, u) = -4 < 0$ and $\lim_{x \rightarrow \infty} h(x, s, u) = +\infty$ this proves that the equation $h(x, s, u) = 0$ has a unique real solution $x(s, u) > x_0 > 1$, depending continuously on (s, u) . Furthermore, $\partial_s h(x, s, u) < 0$ and $\partial_u h(x, s, u) \leq 0$, with equality if and only if $u = 0$, hence by the implicit function theorem $x(s, u)$ is strictly increasing in s and u in C_2 . \square

In particular, the minimal largest root is attained at $a = 1, b = 0$, when $P(t) = t^n - 2t^{\frac{n}{2}} - 2t - 1$. We can compare its root λ to the largest real root

α of $t^n - 2t^{\frac{n}{2}} - 1$. We find $P(\alpha) = \alpha^n - 2\alpha^{\frac{n}{2}} - 2\alpha - 1 = -2\alpha < 0$ thus $\lambda > \alpha$. In particular, $\lambda^n > \alpha^n = 3 + 2\sqrt{2}$ as desired.

In the last case, if $P(t) = t^n - 2t^a - 2t^b - 1$ is skew-reciprocal, then we reduce to the first case, so we can assume that it is skew-reciprocal up to cyclotomic factors but not skew-reciprocal. Suppose $P(t)$ has primitive q^{th} root of unity $\xi = e^{i\theta}$ as a root, namely

$$(7) \quad \xi^n = 2\xi^a + 2\xi^b + 1.$$

Taking the norm of each side of this equation, we get

$$\begin{aligned} 1 &= (2 \cos a\theta + 2 \cos b\theta + 1)^2 + (2 \sin a\theta + 2 \sin b\theta)^2 \\ &= 1 + 4(\cos^2 a\theta + \cos^2 b\theta + \cos a\theta + \cos b\theta + 2 \cos a\theta \cos b\theta) \\ &\quad + 4(\sin^2 a\theta + \sin^2 b\theta + 2 \sin a\theta \sin b\theta) \\ &= 9 + 8(\cos a\theta \cos b\theta + \sin a\theta \sin b\theta) + 4 \cos a\theta + 4 \cos b\theta \\ &= 9 + 8 \cos(a-b)\theta + 8 \cos \frac{(a+b)\theta}{2} \cos \frac{(a-b)\theta}{2} \\ &= 1 + 16 \cos^2 \frac{(a-b)\theta}{2} + 8 \cos \frac{(a+b)\theta}{2} \cos \frac{(a-b)\theta}{2} \\ &= 1 + 8 \cos \frac{(a-b)\theta}{2} \left(2 \cos \frac{(a-b)\theta}{2} + \cos \frac{(a+b)\theta}{2} \right) \\ &= 1 + 8 \cos \frac{(a-b)\theta}{2} \cdot 2 \left(\cos \frac{a\theta}{2} \cos \frac{b\theta}{2} + \sin \frac{a\theta}{2} \sin \frac{b\theta}{2} \right) \\ &\quad + 8 \cos \frac{(a-b)\theta}{2} \left(\cos \frac{a\theta}{2} \cos \frac{b\theta}{2} - \sin \frac{a\theta}{2} \sin \frac{b\theta}{2} \right) \\ &= 1 + 8 \cos \frac{(a-b)\theta}{2} \left(3 \cos \frac{a\theta}{2} \cos \frac{b\theta}{2} + \sin \frac{a\theta}{2} \sin \frac{b\theta}{2} \right) \end{aligned}$$

Thus either $\cos \frac{(a-b)\theta}{2} = 0$ or $3 \cos \frac{a\theta}{2} \cos \frac{b\theta}{2} + \sin \frac{a\theta}{2} \sin \frac{b\theta}{2} = 0$.

Since

$$\begin{aligned} \xi^{a-b} + 1 &= (1 + \cos(a-b)\theta) + i \sin(a-b)\theta \\ &= 2 \cos \frac{(a-b)\theta}{2} \left(\cos \frac{(a-b)\theta}{2} + i \sin \frac{(a-b)\theta}{2} \right), \end{aligned}$$

we have $\cos \frac{(a-b)\theta}{2} = 0$ if and only if $\xi^{a-b} = -1$ and $\xi^n = 1$. This is in turn equivalent to q being even, $a-b$ being an odd multiple of $\frac{q}{2}$, and n being a multiple of q . Also, in this case, $\frac{(a-b)\theta}{2}$ is an odd multiple of $\frac{\pi}{2}$ and we have $\tan \frac{a\theta}{2} \tan \frac{b\theta}{2} = -1$.

Meanwhile, $3 \cos \frac{a\theta}{2} \cos \frac{b\theta}{2} + \sin \frac{a\theta}{2} \sin \frac{b\theta}{2} = 0$ if and only if

- $\cos \frac{a\theta}{2} = 0$ and $\sin \frac{b\theta}{2} = 0$ or

- $\cos \frac{b\theta}{2} = 0$ and $\sin \frac{a\theta}{2} = 0$ or
- $\tan \frac{a\theta}{2} \tan \frac{b\theta}{2} = -3$.

In the first two cases, we have $\xi^{a-b} = -1$ and $\xi^n = 1$, which as pointed out above, is equivalent to q being even, $a - b$ being an odd multiple of $\frac{q}{2}$, and n being a multiple of q .

This reasoning shows that a root $\xi = e^{i\theta}$ of $P(t)$ that is a primitive q^{th} root of unity is exactly one of the following two types:

- (i) If q is even, $a - b$ is an odd multiple of $\frac{q}{2}$, and n is a multiple of q .
- (ii) If $\tan \frac{a\theta}{2} \tan \frac{b\theta}{2} = -3$.

Suppose there are no roots of type (ii). We claim that in this case $P(t)$ cannot be skew-reciprocal up to cyclotomic factors.

PROOF OF CLAIM. There is at least one root of type (i). In particular n is even. Let $Q = \gcd(a - b, \frac{n}{2})$ and let $n = 2NQ$. Without loss of generality suppose $a > b$ and let $a - b = AQ$.

We first compute the quotient $R(t)$ of $P(t)$ by $t^Q + 1$:

$$\begin{aligned} P(t) &= t^{2NQ} - 2t^{AQ+b} - 2t^b - 1 \\ &= (t^{2NQ} - 1) - 2t^b(t^{AQ} + 1). \end{aligned}$$

Hence

$$R(t) = (t^{(2N-1)Q} - \dots + t^Q - 1) - 2t^b(t^{(A-1)Q} - \dots - t^Q + 1).$$

We then check that $R(t)$ has no root of unity as a root. The only possible roots satisfy $\xi^Q = \pm 1$, but

$$\begin{aligned} R(\xi) &= \begin{cases} -2N - 2\xi^b A & \text{if } \xi^Q = -1 \\ -2\xi^b & \text{if } \xi^Q = 1 \end{cases} \\ &\neq 0. \end{aligned}$$

Here we used the fact that $\gcd(A, N) = 1$ to see that the top expression is nonzero.

Finally, we check that $R(t)$ is not skew-reciprocal. If $Q > 1$, then due to primitivity, $b \not\equiv 0 \pmod{Q}$ and the coefficients of the two terms occupy disjoint sets of degrees, so $R(t)$ cannot be skew-reciprocal unless Q is even and $\frac{(A-1)Q}{2} + b = \frac{(2N-1)Q}{2}$. But in this case we have $a + b = n$ and we reduce to case (1) tackled above. If $Q = 1$, then $R(t)$ has odd degree and cannot be skew-reciprocal. □

Hence we can suppose that there is a root $\xi = e^{i\theta}$ of type (ii). Say ξ is a primitive q^{th} root of unity, then we have $\tan \frac{a\pi}{q} \tan \frac{b\pi}{q} = -3$. If q were odd, then using Galois conjugation, we find another root of unity ξ^2 as a root of $P(t)$. Noting that ξ^2 must be of type (ii) as well, this leads to another equation

$$(8) \quad \tan \frac{2a\pi}{q} \tan \frac{2b\pi}{q} = -3.$$

We claim that if (8) is satisfied for every type (ii) root, then $P(t)$ cannot be skew-reciprocal up to cyclotomic factors.

PROOF OF CLAIM. Let us write $\alpha = \tan \frac{a\pi}{q}$ and $\beta = \tan \frac{b\pi}{q}$. Then we have $\alpha\beta = -3$ and $\frac{2\alpha}{1-\alpha^2} \frac{2\beta}{1-\beta^2} = -3$. Hence

$$4 = (1 - \alpha^2)(1 - \beta^2) = 1 - (\alpha + \beta)^2 + 2\alpha\beta + (\alpha\beta)^2 = 1 - (\alpha + \beta)^2 - 6 + 9.$$

This leads to $(\alpha + \beta)^2 = 0$, so $\alpha = -\beta = \pm\sqrt{3}$. Without loss of generality suppose $\tan \frac{a\pi}{q} = \alpha = \sqrt{3}$ and $\tan \frac{b\pi}{q} = \beta = -\sqrt{3}$. Then $\xi^a = e^{\frac{2ia\pi}{q}} = e^{\frac{2i\pi}{3}}$ and $\xi^b = e^{\frac{2ib\pi}{q}} = e^{\frac{4i\pi}{3}}$, and $\xi^n = 2\xi^a + 2\xi^b + 1 = -1$.

Since $\xi^{3a} = \xi^{3b} = \xi^{2n} = 1$, $q | \gcd(3a, 3b, 2n) \leq 6 \gcd(a, b, n) = 6$. In fact, since ξ^a and ξ^b are third roots of unity while ξ^n is a second root of unity, we have $q = 6$. In turn, $a = 6A + 2$, $b = 6B + 4$, and $n = 6N + 3$. In particular, since n is odd, $P(t)$ cannot have any roots of type (i).

We first compute the quotient $R(t)$ of $P(t)$ by $t^2 - t + 1$:

$$\begin{aligned} P(t) &= t^{6N+3} - 2t^{6A+2} - 2t^{6B+4} - 1 \\ &= (t^{6N} - 1)t^3 - 2(t^{6A} - 1)t^2 - 2(t^{6B} - 1)t^4 - (2t^4 - t^3 + 2t^2 + 1) \end{aligned}$$

Hence

$$\begin{aligned} R(t) &= (t^{6N-3} - \dots + t^3 - 1)(t+1)t^3 - 2(t^{6A-3} - \dots + t^3 - 1)(t+1)t^2 \\ &\quad - 2(t^{6B-3} - \dots + t^3 - 1)(t+1)t^4 - (2t^2 + t + 1) \end{aligned}$$

We then check that $R(t)$ has no root of unity as a root. By the reasoning above, the only possible roots ξ satisfy $\xi^3 = -1$, but

$$\begin{aligned} R(\xi) &= 2N(\xi + 1) - 4A(\xi + 1)\xi^2 - 4B(\xi + 1)\xi^4 - (2\xi^2 + \xi + 1) \\ &= 2N(\xi + 1) - 4A(\xi - 2) - 4B(-2\xi + 1) - (3\xi - 1) \\ &= (2N - 4A + 8B - 3)\xi + (2N + 8A - 4B + 1) \neq 0 \end{aligned}$$

since N, A, B are integers.

Finally, since $R(t)$ is of odd degree, it cannot be skew-reciprocal. \square

Thus q is even for at least some type (ii) root $\xi = e^{\frac{2i\pi}{q}}$. For this value of q , we set $q = 2^m u$ with u odd and $m \geq 1$. Since $u - 2$ and $q = 2^m u$ are coprime, $P(t)$ has the root of unity ξ^{u-2} as a root, necessarily of type (ii). We thus have a new relation $\tan a(u - 2)\frac{\pi}{q} \tan b(u - 2)\frac{\pi}{q} = -3$.

An elementary calculation gives $\tan a(u - 2)\frac{\pi}{q} = \tan\left(\frac{a\pi}{2^m} - \frac{2a\pi}{q}\right)$ and $\tan b(u - 2)\frac{\pi}{q} = \tan\left(\frac{b\pi}{2^m} - \frac{2b\pi}{q}\right)$. If $m = 1$ and a, b are even, one has $\tan\left(\frac{a\pi}{2^m} - \frac{2a\pi}{q}\right) = -\tan\frac{2a\pi}{q}$ and $\tan\left(\frac{b\pi}{2^m} - \frac{2b\pi}{q}\right) = -\tan\frac{2b\pi}{q}$. If this is the case for every type (ii) root with an even value of q , then we are reduced to the previous situation (8) and we again run into a contradiction. Hence we can assume that either $m = 1$ and at least one of a, b is odd, or $m > 1$.

If $m = 1$, a is odd, and b is even, one has $\tan\left(\frac{a\pi}{2^m} - \frac{2a\pi}{q}\right) = \frac{1}{\tan\frac{2a\pi}{q}}$. This gives

$$\frac{-\tan\frac{2b\pi}{q}}{\tan\frac{2a\pi}{q}} = -3,$$

or equivalently

$$3\frac{2\alpha}{1 - \alpha^2} = \frac{2\beta}{1 - \beta^2} = \frac{-\frac{6}{\alpha}}{1 - \frac{9}{\alpha^2}} = \frac{-6\alpha}{\alpha^2 - 9}$$

$$\frac{\alpha^2 - 9}{\alpha^2 - 1} = 1$$

which is absurd. Symmetrically, we also get a contradiction if $m = 1$, a is even, and b is odd.

If $m = 1$ and a, b are odd, we have

$$\tan\frac{2a\pi}{q} \tan\frac{2b\pi}{q} = -\frac{1}{3},$$

or equivalently

$$-3\frac{2\alpha}{1 - \alpha^2} = \frac{1 - \beta^2}{2\beta} = \frac{1 - \frac{9}{\alpha^2}}{-\frac{6}{\alpha}} = \frac{\alpha^2 - 9}{-6\alpha}$$

$$36\alpha^2 = (\alpha^2 - 9)(1 - \alpha^2)$$

$$\alpha^4 + 26\alpha^2 + 9 = 0$$

which is also absurd since α is a real number.

Hence we can assume $m > 1$. In this case we will arrive at a contradiction by considering the number fields $\mathbb{Q}(e^{\frac{2i\pi}{2^m u}})$ and $\mathbb{Q}(e^{\frac{2i\pi}{2^{m-1}u}})$. The degrees of these number fields are $\varphi(2^m u) = 2^{m-1}\varphi(u)$ and $\varphi(2^{m-1}u) = 2^{m-2}\varphi(u)$, respectively, hence the latter is strictly contained in the former.

Now by primitiveness, at least one of n, a, b is odd. By grouping up the terms in (7) involving an odd power of ξ , we would be able to express $\xi = e^{\frac{2i\pi}{2^m u}}$ as an element of $\mathbb{Q}(e^{\frac{2i\pi}{2^{m-1}u}})$, leaving us with a contradiction, unless if n is even, a and b are odd, and $\xi^n - 1 = 0$ and $\xi^a + \xi^b = 0$. But then ξ would be of type (i), contradicting what we have assumed above.

Case A_2^* . In this case $Q(t) = 1 - t^a - t^b - t^c + t^{a+b}$, where c is the weight on the isolated vertex. As in case $2A_1$, n must be even. Let us write $n = 2g$. There are two cases to consider: either $a + b = 2g$ or $c = 2g$.

In the former case, the characteristic polynomial is of the form

$$P(t) = t^{2g} - t^{2g-a} - t^c - t^a + 1.$$

The only way to satisfy the parity condition of Lemma 7.8 is if $c = g$. In this case, the polynomial is actually reciprocal. In particular, the normalised spectral radius of the matrix A is bounded from below by $\varphi^4 > 3 + 2\sqrt{2}$ by McMullen's analysis [45].

In the latter case, the characteristic polynomial is of the form

$$P(t) = t^{2g} - t^a - t^b + t^{a+b-2g} - 1.$$

There are three coefficients, so to satisfy the parity condition of Lemma 7.8 at least one of them must be the middle coefficient: either $a = g$, $b = g$ or $a + b = 3g$.

(i) If $a = g$, then we get characteristic polynomials of the form

$$P(t) = t^{2g} - t^b - t^g + t^{b-g} - 1.$$

In order to satisfy the parity condition of Lemma 7.8, we must have $(b - g) + b = 2g$, that is, g is even and $b = \frac{3g}{2}$. If $g \geq 4$ then $P(t)$ is a polynomial in $t^{\frac{g}{2}}$, hence cannot be the characteristic polynomial of a primitive matrix. If $g = 2$ then one can check that the polynomial $P(t) = t^4 - t^3 - t^2 + t - 1$ is not skew-reciprocal up to cyclotomic factors directly.

(ii) If $b = g$, we can repeat the same argument as for $a = g$.

(iii) If $a + b = 3g$, then we get characteristic polynomials of the form

$$P(t) = t^{2g} - t^a - t^b + t^g - 1.$$

The only way to satisfy the parity condition of Lemma 7.8 is if we have $a = b$, that is, if g is even and $a = b = \frac{3g}{2}$. If $g \geq 4$ then $P(t)$ is a polynomial in $t^{\frac{g}{2}}$, hence cannot be the characteristic polynomial

of a primitive matrix. If $g = 2$ then

$$P(t) = t^4 - 2t^3 + t^2 - 1 = (t^2 - t - 1)(t^2 - t + 1),$$

thus $\rho(A)^n = \varphi^4 > 3 + 2\sqrt{2}$.

4. Sharpness of Theorem 1.29

In this section, we demonstrate the asymptotic sharpness of Theorem 1.29.

Given an integer $k \geq 2$, we let Σ_k be the surface

$$(S^1 \times (-1, 1)) \setminus \left\{ (\exp \frac{ai\pi}{k}, 0) \mid a \in \mathbb{Z}/2k \right\}.$$

That is, we puncture $2k$ points from an open annulus in a cyclic fashion. See Figure 7.5. Note that $\chi(\Sigma_k) = -2k$.




FIGURE 7.5. The surface Σ_k and several maps on Σ_k .

We define

$$p_k = \begin{cases} k + 1 & \text{if } k \text{ is even} \\ k + 2 & \text{if } k \text{ is odd} \end{cases}$$

Note that p_k and $2k$ are relatively prime, hence we can set q_k to be an inverse of $p_k \pmod{2k}$.

We set up some maps on Σ_k which we will use to define f_k : We denote by a_n be the arc $\left\{ (\exp i\theta, 0) \mid \theta \in (\frac{nq_k\pi}{k}, \frac{(nq_k+1)\pi}{k}) \right\}$, and denote by σ_n the positive half-twist around a_n . We denote by t the translation by q_k units, that is, $t(\exp i\theta, r) = (\exp i(\theta + \frac{q_k\pi}{k}), r)$. Note that t maps each arc a_n to a_{n+1} . Finally, we denote by η the reflection across $S^1 \times \{0\}$, that is, the map $\eta(\exp i\theta, r) = (\exp i\theta, -r)$, see Figure 7.5.

We then define f_k to be the composition $\eta t \sigma_0$.

PROPOSITION 7.12. *The map $f_k : \Sigma_k \rightarrow \Sigma_k$ is an orientation-reversing fully-punctured pseudo-Anosov map with at least two puncture orbits. The stretch factor of f_k is the largest root of*

$$\begin{cases} t^{2k} - t^{k+1} - t^{k-1} - 1 & \text{if } k \text{ is even} \\ t^{2k} - t^{k+2} - t^{k-2} - 1 & \text{if } k \text{ is odd.} \end{cases}$$

PROOF. It is clear that f_k is orientation-reversing and has at least two puncture orbits. To show the rest of the properties, we will pass to a double branched cover.

Consider the double cover of Σ_k determined by sending loops around each puncture $(\exp \frac{ai\pi}{k}, 0)$ to $1 \in \mathbb{Z}/2$ and sending loops of the type $S^1 \times \{r\}$ to $0 \in \mathbb{Z}/2$. We let $\widehat{\Sigma}_k$ be the surface obtained by filling in the punctures lying above each $(\exp \frac{ai\pi}{k}, 0)$. Note that each arc a_n lifts to a closed curve c_n in $\widehat{\Sigma}_k$. It is convenient to think of $\widehat{\Sigma}_k$ as a chain of $2k$ bands with cores c_n . See Figure 7.6. Note that for each n , c_n intersects each of c_{n+p_k} and c_{n-p_k} once, and intersects those curves only.

Each positive half-twist σ_n on Σ_k lifts to a positive Dehn twist $\widehat{\sigma}_n$ around c_n . The translation t lifts to the translation \widehat{t} that sends each curve c_n to c_{n+1} . The reflection η lifts to a reflection $\widehat{\eta}$ which preserves each c_n . Thus f_k lifts to the map $\widehat{f}_k = \widehat{\eta}\widehat{t}\widehat{\sigma}_0$. It suffices to show that \widehat{f}_k is a fully-punctured pseudo-Anosov map with stretch factor as recorded in Proposition 7.12.

Our proof for this is a copy of the arguments in Chapter 5. In fact, for k even, it is easy to see that \widehat{f}_k is exactly¹ the map denoted by ψ_k in Section 2 of Chapter 5.

Let τ_k be the train track obtained by taking the union of the c_n and smoothing according to their orientations. Note that τ_k is a deformation retract of $\widehat{\Sigma}_k$. Also, note that each c_n is carried by τ_k , that is, each of c_n determines an element of the weight space $\mathcal{W}(\tau_k)$, which we denote by c_n as well. Observe that the set $\{c_n\}$ is linearly independent, and $(f_k)_*$ preserves the subspace U spanned by $\{c_n\}$. In fact, under the basis (c_1, \dots, c_{2k}) , the action of $(f_k)_*$ is represented by the matrix $P + N$ where

$$P_{ij} = \begin{cases} 1 & \text{if } i = j + 1 \\ 0 & \text{otherwise} \end{cases}$$

¹Note that the surface denoted by $\widehat{\Sigma}_k$ here is the surface denoted by Σ_k in Section 2 of Chapter 5.




FIGURE 7.6. The double branched cover $\widehat{\Sigma}_k$ of Σ_k .

and

$$N_{ij} = \begin{cases} 1 & \text{if } (i, j) = (1, p_k) \text{ or } (1, -p_k) \\ 0 & \text{otherwise} \end{cases}$$

where we regard the values of the indices mod $2k$.

It is straightforward to check that $P + N$ is primitive, thus \widehat{f}_k is a fully-punctured pseudo-Anosov map and τ_k fully carries the unstable foliation of the map \widehat{f}_k .

Using the fact that $P + N$ is a companion matrix, one computes the characteristic polynomial of $P + N$ to be

$$t^{2k} - t^{p_k} - t^{2k-p_k} - 1 = \begin{cases} t^{2k} - t^{k+1} - t^{k-1} - 1 & \text{if } k \text{ is even} \\ t^{2k} - t^{k+2} - t^{k-2} - 1 & \text{if } k \text{ is odd.} \end{cases}$$

Thus the spectral radius of $P + N$ is the largest root of this polynomial. Applying the Perron–Frobenius theorem we conclude that the stretch factor of \widehat{f}_k is this largest root as well. \square

The normalised stretch factor $P_k = \lambda(f_k)^{2k}$ of f_k is in turn the largest root of the equation

$$\begin{cases} P_k - (P_k^{\frac{1}{2k}} + P_k^{-\frac{1}{2k}})P_k^{\frac{1}{2}} - 1 = 0 & \text{if } k \text{ is even} \\ P_k - (P_k^{\frac{1}{k}} + P_k^{-\frac{1}{k}})P_k^{\frac{1}{2}} - 1 = 0 & \text{if } k \text{ is odd.} \end{cases}$$

As $k \rightarrow \infty$, both families of equations converge to $P - 2P^{\frac{1}{2}} - 1 = 0$, whose largest root is $3 + 2\sqrt{2}$. Thus we conclude that $\lim_{k \rightarrow \infty} P_k = 3 + 2\sqrt{2}$.

5. Proof of Theorem 1.31

We explain the proof of Theorem 1.31 in this section. The strategy is much simpler than that of Theorem 1.29 and essentially consists of determining, among the orientation-preserving fully-punctured pseudo-Anosov maps that have normalised stretch factor strictly less than φ^4 , which ones admit an orientation-reversing square root.

The set of normalised stretch factors of such maps is determined by Tsang [67, Table 1] and consists of

- (1) four quadratic stretch factors, φ^2 , $\frac{4+\sqrt{12}}{2}$, $\frac{5+\sqrt{21}}{2}$, σ^2 , given by linear Anosov homeomorphisms on the punctured torus $S_{1,1}$ with Euler characteristic -1 ;
- (2) two normalised stretch factors (Lehmer's number)⁹ ≈ 4.311 , realised on $S_{5,1}$, and $|\text{LT}_{1,2}|^3 \approx 5.107$, realised on $S_{2,1}$.

We immediately see that φ and σ are normalised stretch factors of fully-punctured orientation-reversing pseudo-Anosov maps, given, respectively, by matrices $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 2 & 1 \\ 1 & 0 \end{pmatrix} \in \text{GL}_2(\mathbb{Z})$ on $S_{1,1}$. Meanwhile, the normalised stretch factor φ^2 is realised on the punctured sphere $S_{0,4}$ (see Remark 7.10).

For the other two quadratic stretch factors $\frac{4+\sqrt{12}}{2}$ and $\frac{5+\sqrt{21}}{2}$, one can check that the minimal polynomial of their square roots are, respectively, $x^4 - 4x^2 + 1$ and $x^4 - 5x^2 + 1$. Since the degree is not 2, they cannot be the normalised stretch factor of an orientation-reversing Anosov map on $S_{1,1}$.

On the other hand, Lehmer's number, respectively $|\text{LT}_{1,2}|$, is the largest real root of the polynomial $t^{10} + t^9 - t^7 - t^6 - t^5 - t^4 - t^3 + t + 1$, respectively the polynomial $t^4 - t^3 - t^2 - t + 1$. They are in particular Salem numbers (all roots different from $\lambda^{\pm 1}$ have modulus one). By Theorem 1.24, the Galois conjugates of the stretch factor of an orientation-reversing pseudo-Anosov map cannot lie on the unit circle.

Finally, the fact that the spectrum contains a dense subset of $[\sigma^2, \infty)$ follows from the examples in Section 4 (see the discussion in [67, Section 6.1]).

6. Discussion and further questions

As mentioned in the introduction, Theorem 1.29 is motivated by the work of Hironaka and Tsang [22]. However, the statement of Theorem 1.29 is less satisfying than the results in [22] in two aspects:

- Hironaka and Tsang [22] determine the actual value of the minimal normalised stretch factor on surfaces with Euler characteristic $-2k$, for each k . Meanwhile, Theorem 1.29 only determines the asymptotic value of the minimal values in the orientation-reversing case, as $k \rightarrow \infty$.
- Hironaka and Tsang [22] show that the minimal normalised stretch factors on surfaces with odd Euler characteristic are greater than those on surfaces with even Euler characteristic. It is conceivable that this remains true in the orientation-reversing case, especially since we can only prove that Theorem 1.29 is sharp for surfaces with even Euler characteristic, but for now we do not know for sure.

Regarding the first point, we conjecture the following.

CONJECTURE 7.13. *Let $f : S \rightarrow S$ be an orientation-reversing fully-punctured pseudo-Anosov map on a finite-type orientable surface. If we suppose $-\chi(S) = 2k \geq 4$ and if f has at least two puncture orbits, then the normalised stretch factor of f is greater than or equal to the largest real root of*

$$\begin{cases} t^{2k} - t^{k+1} - t^{k-1} - 1 & \text{if } k \text{ is even} \\ t^{2k} - t^{k+2} - t^{k-2} - 1 & \text{if } k \text{ is odd.} \end{cases}$$

In other words, we conjecture that the examples we demonstrated in Section 4 attain the minimal normalised stretch factor among orientation-reversing fully-punctured pseudo-Anosov maps defined on surfaces with even Euler characteristic and with at least two puncture orbits.

Theoretically, one should be able to verify Conjecture 7.13, and, separately, address the second point above, by following the approach in this paper. The additional work lies in extending the analysis in Section 3 of this chapter to a larger collection of curve graphs.

The main deterrent for doing so is the anticipated complexity of the analysis. We currently do not have many tools for identifying polynomials that are skew-reciprocal up to cyclotomic factors. Our parity condition Lemma 7.8 is a rough necessary condition that helps reduce the analysis into finitely many families, but the number of parameters of these families grows as the number of terms in the characteristic polynomial increases. The characteristic polynomial in case $5A_1$ has 6 terms and the analysis is already quite involved. For the curve graphs A_2^{**} and A_3^* , which have the

next largest minimal growth rate, the characteristic polynomial will have 6 and 7 terms respectively, and if one repeats our strategy on these cases the analysis is expected to be even more unwieldy.

Worse still, for $k = 3$, one would have to extend the analysis to all curve graphs with growth rate < 8.186 . We currently do not have a list of such graphs, since McMullen [45] only compiled all graphs with minimum growth rate ≤ 8 . Note that this issue also arises in Chapter 4 where we have to employ additional techniques to compute the minimal spectral radius for primitive skew-reciprocal 6-by-6 matrices.

In conclusion, it would be wise to develop a deeper understanding of polynomials that are skew-reciprocal up to cyclotomic factors before attempting to generalise our analysis by brute force.

We now turn to address Theorem 1.31. An obvious future direction regarding this theorem is to understand the behavior of normalised stretch factors in the range (φ^2, σ^2) . We make the following conjecture.

CONJECTURE 7.14. *The minimal accumulation point of the set of normalised stretch factors of orientation-reversing fully-punctured pseudo-Anosov maps is σ^2 .*

In other words, there are only finitely many isolated values for the normalised stretch factor in the range (φ^2, σ^2) .

Note that if Conjecture 7.14 is true, then it would imply that Theorem 1.29 and Corollary 1.30 remain true without the ‘with at least two puncture orbits’ hypothesis.

Our current approach is inadequate towards verifying this conjecture because there are infinitely many fully-punctured orientation-preserving pseudo-Anosov maps with normalised stretch factor in the range (φ^4, σ^4) . It would take an infinite amount of time to check whether each of these admit an orientation-reversing square root, at least without some new ideas.

Another approach might be to directly generalise the ideas in [67] to the orientation-reversing case. In [67], Tsang showed that the mapping torus of an orientation-preserving fully-punctured pseudo-Anosov map with normalised stretch factor $\leq \varphi^4$ admits a veering triangulation with an explicitly bounded number of tetrahedra. Now, veering triangulations can only exist on orientable 3-manifolds, but there is a generalisation of the notion of veering triangulations to so-called *veering branched surfaces*, and the latter can exist on non-orientable 3-manifolds. We refer to [66] for details.

One should be able to generalise Tsang's ideas [67] to show that the mapping torus of a fully-punctured orientation-reversing pseudo-Anosov map with normalised stretch factor $\leq \sigma^2$ admits a veering branched surface with an explicitly bounded number of triple points. If this number is reasonably small, one could then attempt to generate a list of all veering branched surfaces up to that number of triple points and study the maps that arise as monodromies of fiberings of the corresponding 3-manifolds.

Bibliography

- [1] N. A'Campo: *Sur les valeurs propres de la transformation de Coxeter*, Invent. Math. **33** (1976), no. 1, 61–67.
- [2] T. Ando: *Totally positive matrices*, Linear Algebra Appl. **90** (1987), 165–219.
- [3] S. Baader, C. Graf: *Fibred links in S^3* , Expo. Math. **34** (2016), no. 4, 423–435.
- [4] M. Bell, *flipper* (computer software), <https://pypi.org/project/flipper/>, 2013–2021, version 2 July 2021.
- [5] N. Bourbaki: *Groupes et algèbres de Lie*, Chapitres IV, V, VI, Hermann, 1968, Paris.
- [6] A. E. Brouwer, W. H. Haemers: *Spectra of graphs*, Springer, New York, 2012.
- [7] R. Breusch: *On the distribution of the roots of a polynomial with integral coefficients*, Proc. Amer. Math. Soc. **2** (1951), 939–941.
- [8] G. Burde, H. Zieschang, M. Heusener: *Knots*, De Gruyter, Berlin, 2014.
- [9] F. Gantmakher, M. Krein: *Sur les matrices complètement non négatives et oscillatoires*, Compos. Math. **4** (1937), 445–476.
- [10] K. Calta, T. A. Schmidt: *Infinitely many lattice surfaces with special pseudo-Anosov maps*, J. Mod. Dyn. **7** (2013), no. 2, 239–254.
- [11] J.-H. Cho, J.-Y. Ham: *The minimal dilatation of a genus-two surface*, Experiment. Math. **17** (2008), no. 3, 257–267.
- [12] V. Dimitrov: *A proof of the Schinzel-Zassenhaus conjecture on polynomials*, preprint (2019), <https://arxiv.org/abs/1912.12545>
- [13] D. R. Estes: *Eigenvalues of symmetric integer matrices*, J. Number Theory **42** (1992), no. 3, 292–296.
- [14] B. Farb, D. Margalit: *A primer on mapping class groups*, Princeton University Press, Princeton, 2012.
- [15] A. Fathi: *Démonstration d'un théorème de Penner sur la composition des twists de Dehn*, Bull. Soc. Math. France **120** (1992), no. 4, 467–484.
- [16] A. Fathi, F. Laudenbach, V. Poénaru: *Travaux de Thurston sur les surfaces*, Astérisque, no. 66–67, Société Mathématique de France, Paris 1979.
- [17] D. Fried: *Fibrations over S^1 with pseudo-Anosov monodromy*, in *Travaux de Thurston sur les surfaces*, volume 66-67 of Astérisque, Soc. Math. France, Paris, 1979.
- [18] D. Fried: *Flow equivalence, hyperbolic systems and a new zeta function for flows*, Comment. Math. Helvetici **57** (1982), 237–259.
- [19] D. Fried: *Growth rate of surface homeomorphisms and flow equivalence*, Ergodic Theory Dynam. Systems **5** (1985), no. 4, 539–563.
- [20] E. Gutkin, C. Judge: *Affine mappings of translation surfaces: Geometry and arithmetic*, Duke Math. J. **103** (2000), no. 2, 191–213.

- [21] E. Hironaka, E. Kin, *A family of pseudo-Anosov braids with small dilatation*, Algebr. Geom. Topol. **6** (2006), 699–738.
- [22] E. Hironaka, C. C. Tsang: *Standardly embedded train tracks and pseudo-Anosov maps with minimum expansion factor*, Groups Geom. Dyn. **19** (2025), no. 4, 1263–1317.
- [23] A. J. Hoffman: *On limit points of spectral radii of non-negative symmetric integral matrices*, Graph Theory Appl., Proc. Conf. Western Michigan Univ. 1972, Lect. Notes Math. **303** (1972), 165–172.
- [24] J. Hubbard, H. Masur: *Quadratic differentials and foliations*, Acta Math. **142** (1979), 221–274.
- [25] P. Hubert, E. Laneeau: *Veech groups without parabolic elements*, Duke Math. J. **133** (2006), no. 2, 335–346.
- [26] R. Kenyon, J. Smillie: *Billiards on rational-angled triangles*, Comment. Math. Helv. **75** (2000), no. 1, 65–108.
- [27] M. Kontsevich, A. Zorich: *Connected components of the moduli space of Abelian differentials with prescribed singularities*, Invent. Math. **153** (2003), no. 3, 631–678.
- [28] S. Lang: *Le théorème d'irréductibilité de Hilbert*, Semin. Bourbaki 12 (1959/1960), Exp. No. 201, 13 pp., 1960.
- [29] E. Laneeau, J.-L. Thiffeault: *On the minimum dilatation of pseudo-Anosov homomorphisms on surfaces of small genus*, Ann. Inst. Fourier **61** (2011), no. 1, 105–144.
- [30] E. Laneeau, L. Liechti: *Trace field degrees of Abelian differentials*, Comment. Math. Helv. **99** (2024), no. 1, 81–110.
- [31] E. Laneeau, L. Liechti: *Trace field degrees in the Torelli group*, preprint (2024), accepted for publication in J. Eur. Math. Soc. (2025).
- [32] E. Laneeau, L. Liechti, C. C. Tsang: *Minimal stretch factors of orientation-reversing fully-punctured pseudo-Anosov maps*, Trans. Am. Math. Soc. **378** (2025), no. 4, 2943–2968.
- [33] C. J. Leininger: *On groups generated by two positive multi-twists: Teichmüller curves and Lehmer's number*, Geom. Topol. **8** (2004), 1301–1359.
- [34] L. Liechti: *Minimal dilatation in Penner's construction*, Proc. Amer. Math. Soc. **145** (2017), no. 9, 3941–3951.
- [35] L. Liechti: *On the spectra of mapping classes and the 4-genera of positive knots*, 2017, Dissertation, Universität Bern.
- [36] L. Liechti: *On the arithmetic and the geometry of skew-reciprocal polynomials*, Proc. Amer. Math. Soc. **147** (2019), no. 12, 5131–5139.
- [37] L. Liechti: *On the minimal spectral radii of skew-reciprocal integer matrices*, New York J. Math **30** (2024), 307–322.
- [38] L. Liechti, J. Pankau: *The geometry of bi-Perron numbers with real or unimodular Galois conjugates*, Int. Math. Res. Not. (2022), no. 23, 18526–18540.
- [39] L. Liechti, B. Strenner: *Minimal pseudo-Anosov stretch factors on nonoriented surfaces*, Algebr. Geom. Topol. **20** (2020), no. 1, 451–485.
- [40] L. Liechti, B. Strenner, *The Arnoux–Yoccoz mapping classes via Penner's construction*, Bull. Soc. Math. France **148** (2020), no. 3, 383–397.
- [41] L. Liechti, B. Strenner: *Minimal Penner dilatations on nonorientable surfaces*, J. Topol. Anal. **13** (2021), no. 1, 187–218.

- [42] D. Margalit: *Problems, Questions, and Conjectures about Mapping Class Groups*, Breadth in contemporary topology, 157–186, Proc. Sympos. Pure Math., 102, Amer. Math. Soc., Providence, RI, 2019.
- [43] J. McKee, C. Smyth: *Around the unit circle–Mahler measure, integer matrices and roots of unity*, Springer, Cham, 2021.
- [44] C. T. McMullen: *Coxeter groups, Salem numbers and the Hilbert metric*, Publ. Math. Inst. Hautes Études Sci. **95** (2002), 151–183.
- [45] C. T. McMullen: *Entropy and the clique polynomial*, J. Topol. **8** (2015), 184–212.
- [46] C. T. McMullen: *Billiards and Teichmüller curves on Hilbert modular surfaces*, J. Amer. Math. Soc. **16** (2003), no. 4, 857–885.
- [47] M. R. Murty: *Polynomials assuming square values*, Number theory and discrete geometry. Proceedings of the international conference in honour of Professor R. P. Bambah, Chandigarh, India, November 30–December 3, 2005, pages 155–163. Mysore: Ramanujan Mathematical Society, 2008.
- [48] J. Pankau: *Salem number stretch factors and totally real number fields arising from Thurston’s construction*, Geom. Topol. **24**, no. 4, 1695–1716.
- [49] J. Pankau: *On stretch factors of pseudo-Anosov maps*, PhD thesis, University of California Santa Barbara, 2018.
- [50] R. C. Penner: *A construction of pseudo-Anosov homeomorphisms*, Trans. Amer. Math. Soc. **310** (1988), no. 1, 179–197.
- [51] R. C. Penner: *Bounds on least dilatations*, Proc. Amer. Math. Soc. **113** (1991), no. 2, 443–450.
- [52] O. Pretzel: *On reorienting graphs by pushing down maximal vertices*, Order **3** (1986), no. 2, 135–153.
- [53] R. M. Robinson: *Intervals containing infinitely many sets of conjugate algebraic integers*, Studies in Mathematical Analysis and Related Topics: Essays in Honor of George Pólya, Stanford 1962, pp. 305–315.
- [54] J. Salez: *Every totally real algebraic integer is a tree eigenvalue*, J. Comb. Theory, Ser. B **111** (2015), 249–256.
- [55] A. Schinzel: *Addendum to the paper ‘On the product of the conjugates outside the unit circle of an algebraic number’*, Acta Arith. **26** (1975), 329–331.
- [56] A. Schinzel, H. Zassenhaus: *A refinement of two theorems of Kronecker*, Michigan Math. J. **12** (1965), no. 1, 81–85.
- [57] I. Schur: *Über die Verteilung der Wurzeln bei gewissen algebraischen Gleichungen mit ganzzahligen Koeffizienten*, Math. Z. **1** (1918), no. 4, 377–402.
- [58] J.-Y. Shi: *The enumeration of Coxeter elements*, J. Algebraic Combin. **6** (1997), no. 2, 161–171.
- [59] H. Shin, B. Strenner: *Pseudo-Anosov mapping classes not arising from Penner’s construction*, Geom. Topol. **19** (2015), no. 6, 3645–3656.
- [60] J. Stallings: *Constructions of fibred knots and links*, Algebraic and Geometric Topology, 55–60, Proc. Sympos. Pure Math. **32** (1978) Amer. Math. Soc., Providence, R.I.
- [61] R. Steinberg: *Finite reflection groups*, Trans. Amer. Math. Soc. **91** (1959), 493–504.

- [62] B. Strenner: *Algebraic degrees of pseudo-Anosov stretch factors*, *Geom. Funct. Anal.* **27** (2017), no. 6, 1497–1539.
- [63] B. Strenner: *Lifts of pseudo-Anosov homeomorphisms of nonorientable surfaces have vanishing SAF invariant*, *Math. Res. Lett.* **25** (2018), no. 2, 677–685.
- [64] O. Teichmüller: *Extremale quasikonforme Abbildungen und quadratische Differentiale*, *Abh. Preuß. Akad. Wiss.*, no. 22, 1–197 (1940).
- [65] W. Thurston: *On the geometry and dynamics of diffeomorphisms of surfaces*, *Bull. Am. Math. Soc.* **19** (1988), no. 2, 417–431.
- [66] C. C. Tsang: *Veering branched surfaces, surgeries, and geodesic flows*, *New York J. Math.* **29** (2023), 1425–1495.
- [67] C. C. Tsang: *On the set of normalized dilatations of fully-punctured pseudo-Anosov maps*, preprint (2023), <https://arxiv.org/abs/2306.10245>.
- [68] W. A. Veech: *Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards*, *Invent. Math.* **97** (1989), no. 3, 553–583.
- [69] Mehdi Yazdi: *Lower bound for the Perron–Frobenius degrees of Perron numbers*, *Ergodic Theory Dyn. Syst.* **41** (2021), no. 4, 1264–1280.