

Bader, Philipp (2026) *Translation lengths of pseudo-Anosovs.* PhD thesis.

https://theses.gla.ac.uk/86383/

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses
https://theses.gla.ac.uk/
research-enlighten@glasgow.ac.uk

Translation lengths of pseudo-Anosovs

by

Philipp Bader

A thesis submitted in fulfilment of the requirements
for the degree of

Doctor of Philosophy

at the

School of Mathematics & Statistics
College of Science & Engineering
University of Glasgow



July 2026

Στη μαμά μου

Abstract

Pseudo-Anosov mapping classes of a surface serve as a bridge between dynamics and flat geometry. From a dynamical point of view, one can study intrinsic invariants, such as the stretch factor, or consider properties of actions of pseudo-Anosovs on different spaces. On the other hand, there is a naturally assigned flat surface to each pseudo-Anosov. On this surface, the pseudo-Anosov can be viewed as an affine map. This makes the investigation of flat surfaces, such as translation surfaces, crucial in the theory of pseudo-Anosovs.

In this work, we study pseudo-Anosov mapping classes via their actions on Teichmüller space, the curve graph and the pants graph and focus on measuring the translation lengths of these actions. In particular, we discuss results concerning the minimal translation lengths and present comparison theorems. We then explore the connection between pseudo-Anosovs and flat surfaces with an emphasis on square-tiled surfaces. After reviewing the relation between square-tiled surfaces, filling pairs of curves, and pseudo-Anosovs, we use a group-theoretic result on commutators of permutations to derive statements about these notions. We end with an analysis of geminal square-tiled surfaces and the corresponding pillowcase-tiled surfaces.

Contents

Introduction **1**

1 Preliminaries **10**
1.1 Teichmüller and moduli space 10
1.2 The classification of mapping classes 13
1.3 Translation surfaces 18
1.4 The curve graph and the pants graph 27
1.5 Translation lengths and pseudo-Anosov actions 30

2 Minimal hyperelliptic translation lengths **40**
2.1 Relevant notions 43
2.2 Topological Rauzy-Veech induction 47
2.3 Some properties of f_γ and V_γ 52
2.4 Translation lengths in hyperelliptic components 60

3 Infinite multiplicity of $l_\mathcal{C}$ **71**
3.1 Modification of Penner's example 73
3.2 Construction of infinite multiplicity 77

4 Commutators of n-cycles **86**
4.1 Proof of Theorem 4.1 87

5 Filling pairs, square-tiled surfaces and pseudo-Anosov stretch factors **99**
5.1 Filling pairs of curves and relation with STS 102
5.2 Affine pseudo-Anosov maps 107
5.3 Proofs of the results 108
5.4 Generalisations and future directions 114

6 Pants graph translation lengths **117**
6.1 Statements of the results 118
6.2 The map from pants graph to curve graph 120
6.3 Translation lengths of pseudo-Anosovs 121
6.4 The map from Teichmüller space to pants graph 125
6.5 Future directions 128

7 Geminal square-tiled surfaces **130**
7.1 Preliminaries 132
7.2 Normality and cyclicity of geminal covers 136
7.3 Pillowcase-tiled surfaces 145
7.4 Open questions 153

Bibliography **155**

Acknowledgements

First, I would like to express my gratitude to my supervisors, Vaibhav Gadre and Tara Brendle. Vaibhav, thank you for your enthusiasm for mathematics, which has always been motivating and inspiring to me. Tara, thank you for all your encouragement and support. I would also like to thank my examiners, Jim Belk and Pascal Hubert, for their interest in this thesis, the careful reading, and helpful comments.

A special thank you goes to Paul Apisa and Gabriela Brown for welcoming me during my research visit in Madison and making it such an enjoyable experience. The mathematical conversations we had significantly broadened my knowledge and enriched my research. I want to thank Luke Jeffreys for hosting me in Bristol multiple times and for all the insightful discussions as well as fun activities we had together. I am grateful to Sebastian Hensel and Panos Papadopoulos for many exchanges of ideas over the years, both before and during the PhD. Thank you to Ryan Lamy, whose mathematical curiosity taught me many valuable lessons. I also thank Mehdi Yazdi for very helpful conversations regarding part of this thesis. Lastly, it has been a pleasure to work alongside my academic brother Riccardo Giannini; you were the first person I met in Glasgow and immediately made me feel at home.

I am deeply grateful to the entire Glasgow School of Mathematics for making my PhD such a valuable and memorable chapter of my life. From mathematical discussions, seminars and reading groups, to light-hearted conversations at lunch breaks and social activities beyond the department, such as parkrun trips, football games or dinner evenings, I would like to thank everyone among the many people who are too numerous to name, and feel incredibly fortunate to have had the opportunity to study and live in Glasgow. I am especially thankful to Isacco and JJ for having become "PhD that matter" and for spending some of the most emotional moments of our lives together.

I would never have reached this stage without the support of my family. I would like to thank my uncle Rainer and my aunt Sabine for all our calls during this time and for always cheering me up. I am thankful to my cousins Vanessa, Kleio and Athina for the joyful moments we shared, from DnD sessions to our trips together. A heartfelt thank you to Konstantinos, who has always been an invaluable friend throughout my life. I am grateful to have been guided by my grandparents and my aunt Leoli, who cannot be with us in this moment but will always accompany me. Finally, I want to express my deepest gratitude to my parents, who have taught me everything and have always been there for me; without your unwavering support, I would not be where I am today.

To finish, I would like to highlight my gratefulness to Giulia. Thank you for being by my side, for always believing in me, and for sharing so many beautiful moments together.

Author's declaration

I declare that, except where explicit reference is made to the contribution of others, this dissertation is the result of my own work and has not been submitted for any other degree at the University of Glasgow or any other institution.

Introduction

The theory of surfaces is a long-studied, rich branch of mathematics that connects various fields, such as geometry, topology and dynamics. Originating in the 19^{th} century through the work of Gauß and Riemann, it remains a thriving area of research. We want to briefly touch upon the general theory of surfaces while exploring the vast, interesting problems that arise and simultaneously introducing some of the important notions for this work.

We consider a closed, connected, oriented surface of genus g, which we denote by S_g. Throughout, we usually omit the word "connected", but a surface is always assumed to be connected, unless stated otherwise. Oftentimes, one can equip a surface with a hyperbolic metric and consider geometric questions, such as finding the lengths of geodesics under this metric, cf. [Mir08], [MP19]. A hyperbolic metric on a surface S_g is equivalent to a choice of a Riemann surface structure on S_g. The space of all such structures is the *moduli space* $\mathcal{M}_g$ which is of interest to topologists as well as algebraic geometers, cf. [DM69], [Kon92]. To study $\mathcal{M}_g$, one can consider its (orbifold) universal cover $\mathcal{T}_g$, the *Teichmüller space*. The deck group of this cover is the *mapping class group* $\mathrm{Mod}(S_g)$, which is the group of orientation-preserving homeomorphisms of S_g up to isotopy. This discrete group is of independent interest to geometric group theorists, cf. [BKMM12], [BBF14]. The mapping class group is weakly hyperbolic, as can be seen by its action on the hyperbolic *curve graph* $\mathcal{C}(S_g)$; a graph whose vertices correspond to isotopy classes of curves on S_g. This action can moreover be used to show finite generation of $\mathrm{Mod}(S_g)$ (see [FM12, Section 4.3], or [BF04], [Kor04] for further results on finite generation).

The elements of the mapping class group have been classified into three categories by Thurston [Thu88]. Every $f \in \mathrm{Mod}(S_g)$ is either *periodic*, *reducible* or *pseudo-Anosov*. The latter are elements of infinite order that do not preserve any finite set of curves on S_g. Equivalently, pseudo-Anosovs can be characterised as elements that stretch and contract

two foliations on the surface by some *stretch factor* $\lambda \in \mathbb{R}$. The study of numbers that can arise as stretch factors of pseudo-Anosovs has sparked many interesting questions to attract number theorists to the field, cf. [Hir09], [Pan20].

Closely related to the study of Teichmüller and moduli space is the study of the space of translation surface structures on S_g. A *translation surface* can be described by a polygon in the plane whose sides are identified by translations to yield a closed surface. Understanding the topology of spaces of translation surfaces, particularly in the context of computing their fundamental group, is an active area of current research, cf. [CS23], [Gia24]. The group $\mathrm{GL}_2^+(\mathbb{R})$ acts on polygons by linear transformations and consequently on the space of translation surfaces. This action is where the dynamics come into play. In particular, the action contains the so-called *Teichmüller flow*, which is used to understand the geodesics of moduli space, cf. [AGY06], [Via06].

A *square-tiled surface* is a special type of translation surface which, as the name suggests, is built out of squares. The study of square-tiled surfaces is interesting in its own right, as we will see throughout, but also of importance in the bigger picture. For instance, counts on the number of square-tiled surfaces with certain properties contain information about the volume of spaces of translation surfaces, cf. [Zor02], [DGZZ20].

In this work, we study pseudo-Anosov elements of the mapping class group along with square-tiled surfaces. In particular, we consider actions of the mapping class group $\mathrm{Mod}(S_g)$ on various spaces, such as the Teichmüller space $\mathcal{T}(S_g)$, the curve graph $\mathcal{C}(S_g)$ or the *pants graph* $\mathcal{P}(S_g)$ and focus on the *translation lengths* of pseudo-Anosov elements under these actions – a measure of how much the pseudo-Anosov translates the respective space on average. To define the translation lengths, we equip $\mathcal{T}(S_g)$ with the Teichmüller metric and $\mathcal{C}(S_g)$ and $\mathcal{P}(S_g)$ with the shortest path metric, respectively. We also illustrate the connection between pseudo-Anosovs and square-tiled surfaces, and portray how examining either yields information about the other. We end with an analysis of a special type of square-tiled surfaces, called *geminal*.

In what follows, we present our main results and outline the structure of this work.

Let S_g be a closed, oriented surface of genus $g \geqslant 2$ and $\mathrm{Mod}(S_g)$ be the mapping class group. For a pseudo-Anosov $f \in \mathrm{Mod}(S_g)$, let $\lambda(f)$ denote its stretch factor. For a metric space (X, d) on which f acts as an isometry, we let

$$l_X(f) \coloneqq \liminf_{n\to\infty} \frac{d(x, f^n(x))}{n},$$

where $x \in X$, denote the stable translation length of f. The stretch factor of f is compa-

rable to its Teichmüller space translation length via the equality $l_{\mathcal{T}}(f) = \log(\lambda(f))$. The set of stretch factors over all pseudo-Anosovs in $\mathrm{Mod}(S_g)$ admits a minimum, i.e.

$$\lambda_g \coloneqq \min\{\lambda(f) \mid f \in \mathrm{Mod}(S_g) \text{ pseudo-Anosov}\}$$

exists and is strictly greater than 1. Equivalently, the analogous quantity for Teichmüller translation lengths

$$l_g^{\mathcal{T}} \coloneqq \min\{l_{\mathcal{T}}(f) \mid f \in \mathrm{Mod}(S_g) \text{ pseudo-Anosov}\}$$

exists and is strictly greater than 0. In fact, the asymptotic behaviour of these quantities has been determined. We have that $\lambda_g \asymp e^{\frac{1}{g}}$ and $l_g^{\mathcal{T}} \asymp \frac{1}{g}$ [Pen91]. Computing the explicit minimum is a big open question in the field.

Every pseudo-Anosov arises as an affine homeomorphism for some translation surface. Each translation surface has a certain number of singularities, and each singularity has a certain order. The moduli space of all translation surfaces can be written as a union of strata, where each stratum is defined by prescribing the number and order of singularities of translation surfaces in that stratum. The connected components of strata have been classified [KZ03]. For fixed genus g, there are exactly two connected components consisting entirely of hyperelliptic translation surfaces; one is a component of the stratum with a single singularity of order $2g-2$, and the other of the stratum with two singularities, both of order $g-1$.

Boissy–Lanneau consider the set of pseudo-Anosovs that arise as affine homeomorphisms for some translation surface in a hyperelliptic stratum component [BL22]. We call a pseudo-Anosov in this set a hyperelliptic pseudo-Anosov. Boissy–Lanneau determine the minimal Teichmüller translation length of hyperelliptic pseudo-Anosovs, i.e. they compute the precise value of

$$l_{g,\mathrm{hyp}}^{\mathcal{T}} \coloneqq \min\{l_{\mathcal{T}}(f) \mid f \in \mathrm{Mod}(S_g) \text{ hyperelliptic pseudo-Anosov}\},$$

and determine a pseudo-Anosov attaining the minimum. Let f_g be this pseudo-Anosov. In particular, it holds that

$$l_{g,\mathrm{hyp}}^{\mathcal{T}} = l_{\mathcal{T}}(f_g) \geqslant \log(\sqrt{2}),$$

for any genus g.

We consider the same set of pseudo-Anosovs as Boissy–Lanneau and obtain an analogous result for the translation length on the curve graph $\mathcal{C}(S_g)$. Let $l_{g,\mathrm{hyp}}^{\mathcal{C}}$ denote the minimal curve graph translation length over all hyperelliptic pseudo-Anosovs. We compute the

asymptotic behaviour of the minimum $l^{\mathcal{C}}_{g,\mathrm{hyp}}$ and show that it is asymptotically attained by f_g.

Theorem (A). *The minimal curve graph translation length over all hyperelliptic pseudo-Anosovs satisfies*

$$l^{\mathcal{C}}_{g,\mathrm{hyp}} \asymp \frac{1}{g}.$$

Furthermore, for the minimiser of the Teichmüller translation length f_g, we have

$$\frac{1}{16g-12} \leqslant l_{\mathcal{C}}(f_g) \leqslant \frac{1}{g-1}.$$

We observe that the behaviour of the curve graph translation length of hyperelliptic pseudo-Anosovs is different to the one of Teichmüller space: while the minimum $l^{\mathcal{T}}_{g,\mathrm{hyp}}$ is bounded from below by a positive constant, independently of the genus, the minimum $l^{\mathcal{C}}_{g,\mathrm{hyp}}$ tends to 0 as the genus increases. We proceed with comparison results between Teichmüller space and curve graph translation lengths, where we don't restrict to hyperelliptic pseudo-Anosovs anymore.

Theorem (B). *For any $g \geqslant 3$, there exists a sequence $(f_n)_{n=1}^{\infty} \subset \mathrm{Mod}(S_g)$ of pseudo-Anosovs with*

$$\lim_{n\to\infty} l_{\mathcal{T}}(f_n) = \infty \text{ and } l_{\mathcal{C}}(f_n) \leqslant \frac{1}{g-1} \text{ for all } n \in \mathbb{N}.$$

It follows that there exists a sequence $(h_g)_{g=2}^{\infty}$, where h_g is a pseudo-Anosov of a genus g surface with

$$l_{\mathcal{T}}(h_g) \to \infty \text{ and } l_{\mathcal{C}}(h_g) \to 0,$$

as $g \to \infty$.

The theorem asserts in particular that the difference between the Teichmüller space and the curve graph translation length of a pseudo-Anosov is unbounded. The proof is constructive, i.e. the pseudo-Anosovs f_n and h_g are given explicitly.

Two conjugate pseudo-Anosovs have the same translation length. We say that the *multiplicity* of a number $\lambda \in \mathbb{R}$ is the number of conjugacy classes of pseudo-Anosovs with translation length λ. If we fix the genus, then in the case of Teichmüller space translation length, the multiplicity is always finite, i.e. for any $\lambda \in \mathbb{R}$, there exist at most finitely many pairwise non-conjugate pseudo-Anosovs with Teichmüller translation length equal to λ [FM12, Theorem 14.9]. As a corollary of the above theorem, we obtain that multiplicities can be infinite for the curve graph.

Corollary (C). *For any $g \geqslant 3$, there exists $q \in \mathbb{Q}$ and infinitely many pairwise non-conjugate pseudo-Anosovs in* $\mathrm{Mod}(S_g)$ *with curve graph translation length equal to q.*

The proof of this corollary is non-constructive. The point is essentially that the sequence $(f_n)_{n=1}^{\infty}$, constructed in Theorem B, admits a subsequence that satisfies the corollary. However, there is no way to specify the subsequence exactly. If a pseudo-Anosov satisfies certain requirements, we are also able to produce an explicit construction of infinitely many pseudo-Anosovs of equal curve graph translation length.

Theorem (D). *Let $f \in \mathrm{Mod}(S_g)$ be a pseudo-Anosov with $l_{\mathcal{C}}(f) = m$, where $m \in \mathbb{N}$, and assume that there is a non-separating curve α attaining the translation length, i.e.*

$$d_{\mathcal{C}}(\alpha, f^k(\alpha)) = km,$$

for all $k \in \mathbb{Z}$. Then, for any pseudo-Anosov h of the surface $S_g \setminus \alpha$, there exists $N_h \in \mathbb{N}$ such that for $n \geqslant N_h$, each $f_{h,n} := h^n \circ f$ satisfies $l_{\mathcal{C}}(f_{h,n}) = l_{\mathcal{C}}(f)$.

Apart from Teichmüller space and curve graph, we also focus on translation lengths of pseudo-Anosovs acting on the pants graph $\mathcal{P}(S_g)$ – a graph whose vertices correspond to pants decompositions of the surface. As remarked above, the minimal Teichmüller space translation length $l_g^{\mathcal{T}}$ behaves asymptotically like $\frac{1}{g}$. For the curve graph, the analogous quantity is smaller; from [GT11], we know that

$$l_g^{\mathcal{C}} := \min\{l_{\mathcal{C}}(f) \mid f \in \mathrm{Mod}(S_g) \text{ pseudo-Anosov}\} \asymp \frac{1}{g^2}.$$

We establish asymptotic bounds on the minimal pants graph translation length.

Theorem (E). *Let $l_g^{\mathcal{P}}$ be the minimal pants graph translation length over all pseudo-Anosovs in* $\mathrm{Mod}(S_g)$. *There exists $C > 0$ such that*

$$\frac{1}{g} \lesssim l_g^{\mathcal{P}} \lesssim C.$$

The proof of this theorem relies on considering a map

$$p : \mathcal{P}(S_g) \to \mathcal{C}(S_g),$$

which allows us to compare translation lengths on the pants graph to those on the curve graph. The map p sends a pants decomposition to one of the curves it contains.

In a similar fashion, the systole map

$$\mathrm{sys}\colon \mathcal{T}(S_g) \to \mathcal{C}(S_g),$$

which sends a hyperbolic surface to one of its systoles (i.e. shortest curves), allows us to compare Teichmüller space and curve graph translation lengths. The systole map is coarsely Lipschitz, and Gadre–Hironaka–Kent–Leininger compute that the optimal Lipschitz constant behaves asymptotically like $\frac{1}{\log(g)}$ [GHKL13]. We establish the coarsely Lipschitz property for the map p and produce asymptotic bounds for its optimal Lipschitz constant.

Theorem (F). *The map $p\colon \mathcal{P}(S_g) \to \mathcal{C}(S_g)$ is coarsely Lipschitz. Let μ_g be the optimal Lipschitz constant. It holds that*

$$\frac{1}{g^2} \lesssim \mu_g \lesssim \frac{1}{g}.$$

Finally, we consider a map

$$q\colon \mathcal{T}(S_g) \to \mathcal{P}(S_g),$$

defined by Brock [Bro03], which relates Teichmüller space and the pants graph. As a corollary of the above theorem, we obtain:

Corollary (G). *The map q is coarsely Lipschitz, and its optimal Lipschitz constant ν_g satisfies*

$$\nu_g \gtrsim \frac{g}{\log(g)}.$$

We now turn our attention to square-tiled surfaces. A square-tiled surface can be described combinatorially in terms of permutations. In fact, the data of two permutations, which encode how to glue the squares together in the horizontal and vertical directions, respectively, is enough to determine a square-tiled surface. Furthermore, the commutator of these two permutations contains information on the stratum the square-tiled surface lies in. Motivated by this observation, we prove the following purely group-theoretic result. Let S_n be the symmetric group and A_n the alternating group on n elements. Let $C(n) \subset S_n$ be the conjugacy class of n-cycles. Finally, let $[\tau, \pi]$ denote the commutator of two permutations $\tau, \pi \in S_n$.

Theorem (H). *The map*

$$C(n) \times C(n) \to A_n, \ (\tau, \pi) \mapsto [\tau, \pi]$$

is surjective for all $n \geqslant 6$.

We note that this result also follows from a paper by Vavpetič, and thank Goulnara Arzhantseva for pointing out the reference to us [Vav16].

Using Theorem H on commutators, we are able to present slick proofs of results on square-tiled surfaces and stretch factors of pseudo-Anosovs that are mostly folklore in the field, or also follow from [Jef21]. These results are stated precisely in Chapter 5. In particular, we show that in any stratum there exist $(1,1)$-square-tiled surfaces with any number of squares above the minimum necessary. A $(1,1)$-square tiled surface is one whose corresponding permutations are both n-cycles, hence the relation to Theorem H. From this, it follows that for any $n \geqslant 2g-1$, there is a pair of filling curves on S_g, i.e. two curves whose complementary components are all contractible, such that these two curves intersect in exactly n points. Regarding stretch factors, for a fixed stratum $\mathcal{H}$, we let $\Lambda_{\mathcal{H}}$ denote the set of stretch factors over all pseudo-Anosovs that are affine homeomorphisms on a translation surface in $\mathcal{H}$. We can show that for any finite collection of strata $\mathcal{H}_1, \ldots, \mathcal{H}_n$, the intersection of the corresponding stretch factor sets is infinite, i.e.

$$\left| \bigcap_{i=1}^{n} \Lambda_{\mathcal{H}_i} \right| = \infty.$$

Finally, as remarked above, the multiplicity of a stretch factor is finite. However, we sketch a proof that the supremum of the set of multiplicities is infinite. Concretely, let $\Lambda_g \subset \mathbb{R}$ be the set of numbers that arise as stretch factors of a pseudo-Anosov in $\mathrm{Mod}(S_g)$. For $\lambda \in \Lambda_g$, let m_λ be the number of pairwise non-conjugate pseudo-Anosovs with stretch factor λ. We have that for any $g \geqslant 1$,

$$\sup_{\lambda \in \Lambda_g} m_\lambda = \infty.$$

In the final chapter of this work, we study *geminal* translation surfaces, which were introduced by Apisa–Wright [AW22]. On a translation surface, there is a notion of a flat cylinder. A geminal translation surface is characterised by the property that cylinders come in parallel pairs. The theory of these surfaces plays an important role in the classification of orbit closures of the $\mathrm{GL}_2^+(\mathbb{R})$-action on the moduli space of translation surfaces [AW23].

Let T_4 be the torus constructed by gluing together four squares, which is defined in Chapter 7. For a geminal square-tiled surface, there is a natural covering map to T_4. Apisa–Wright posed the problem of whether this covering is always normal and, in that case, if the deck group is cyclic [AW22, Problem 8.16]. It has been shown that this covering is not always normal [AB25]. We show that, when the covering is normal, the deck group is indeed cyclic. In particular, we obtain:

Theorem (I). *Let X be a geminal square-tiled surface and $p\colon X \to T_4$ the natural covering map. Then, we have the following equivalence:*

$$p \text{ is normal} \iff p \text{ is cyclic.}$$

Shortly before publication of this thesis, we became aware that the direction normality implies cyclicity in the above theorem was also proved in a group-theoretic way in [The22]. We present an independent, more geometrically flavoured proof.

Outline. We point out that after reading Chapter 1, which introduces the main notions and provides the necessary definitions, all other chapters are designed so that they can be read independently. For that reason, some of the main definitions appear multiple times throughout this work. Each chapter has its own introduction, which contains the main results of that chapter. All the aforementioned theorems can be found in the respective chapters, either stated exactly as above or with more details. Some of the chapters end with a section on open questions or possible future directions.

In Chapter 1, we introduce the main notions and give a general overview of the necessary theory. In particular, we define Teichmüller space, curve graph and pants graph, as well as translation surfaces and square-tiled surfaces. We also discuss pseudo-Anosov mapping classes and their translation lengths.

In Chapter 2, we discuss the construction of pseudo-Anosovs via Rauzy-Veech induction and prove Theorem A on the minimal curve graph translation length of hyperelliptic pseudo-Anosovs.

In Chapter 3, we construct the sequences of pseudo-Anosovs that appear in Theorem B. To do that, we modify Penner's famous example of a pseudo-Anosov [Pen91]. We then prove Corollary C and Theorem D on the infinite multiplicity of curve graph translation lengths. For the proof of Theorem D, we use the theory of subsurface projections developed by Masur–Minsky [MM00]. Chapters 2 and 3 contain work from a paper of the author [Bad25b].

In Chapter 4, we prove the group-theoretic Theorem H on the commutator map. This chapter contains work from a preprint by the author, which is not intended for publication [Bad25a].

In Chapter 5, we explore the connection between square-tiled surfaces and filling sets of curves, and discuss how these can be used to construct pseudo-Anosovs. We then prove the results about these notions mentioned above using the theorem from Chapter 4.

In Chapter 6, we focus on pants graph translation lengths and prove Theorems E and F, as well as Corollary G.

In Chapter 7, we discuss geminal square-tiled surfaces and prove Theorem I. Since geminal square-tiled surfaces are closely related to so-called 1-cylinder pillowcase-tiled surfaces, we include a discussion of the latter and highlight connections between the two.

Chapter 1

Preliminaries

§ 1.1 | Teichmüller and moduli space

In this section, we provide the main definitions of *Teichmüller* and *moduli space*, as well as of the *mapping class group* of a surface. In particular, we see that the mapping class group acts on Teichmüller space and that moduli space can be described as the quotient of this action.

There are many approaches to defining Teichmüller space. Here, we present two equivalent ones: one as the space of *marked Riemann surfaces* and the other as the space of *marked hyperbolic surfaces*. We start by introducing the main notions.

Definition 1.1. *A* Riemann surface X *is a topological surface together with a maximal atlas of charts to the complex plane* $\mathbb{C}$ *such that the chart transitions are holomorphic.*

Definition 1.2. *A* hyperbolic surface X *is a quotient of the hyperbolic plane* $\mathbb{H}$ *by a discrete subgroup of* $\mathrm{Isom}^+(\mathbb{H})$ *acting freely on* $\mathbb{H}$*. We shall denote this subgroup by* Γ*.*

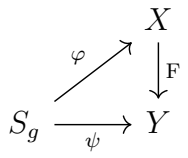
The genus of a Riemann or hyperbolic surface X is defined to be the genus of the underlying topological surface. A hyperbolic surface inherits a metric from $\mathbb{H}$ which we call the hyperbolic metric on X. Both Riemann and hyperbolic surfaces have a natural orientation coming from the canonical orientations of $\mathbb{C}$ and $\mathbb{H}$, respectively.

It turns out that for closed surfaces of genus $\geqslant 2$, the above two definitions are equivalent, i.e. any Riemann surface is a hyperbolic surface and vice versa. To see this, note that the *Uniformisation Theorem* implies that the universal cover of such a Riemann surface X is the unit disc $D = \{z \in \mathbb{C} \,:\, |z| < 1\}$. Hence, X can be described as the quotient of D by a discrete group Γ acting freely, via biholomorphisms on D. We can identify $\mathbb{H}$ with D via the Poincaré disc model. The crucial point now is that the group of biholomorphisms of D

is isomorphic to the orientation-preserving isometry group of $\mathbb{H}$. This group is $\mathrm{PSL}_2(\mathbb{R})$ acting by Möbius transformations. Therefore, the group Γ is a subgroup of $\mathrm{PSL}_2(\mathbb{R})$ and acts by isometries on $\mathbb{H}$, which implies that X is a hyperbolic surface. The reverse implication works the same way.

From now on, we will use the terms Riemann and hyperbolic surface interchangeably. Let S_g be a closed, connected, oriented topological surface of genus g. Unless otherwise specified, we assume that $g \geqslant 2$. A *marking* of S_g is an orientation-preserving homeomorphism $\varphi : S_g \to X$, where X is some Riemann surface. One can think of a marking as a way of equipping S_g with a Riemann surface structure (or with a hyperbolic metric) by pulling back the structure from X via φ. In this regard, the homeomorphism φ tells us how to identify S_g with X. We want to think of the pulled back structures as being the same if this identification is made by two isotopic homeomorphisms. This idea motivates the following definition of Teichmüller space.

Definition 1.3. *The Teichmüller space of S_g is the space of all pairs (X, φ) where X is a Riemann surface and φ a marking, up to the following equivalence relation: two pairs (X, φ) and (Y, ψ) are identified if there exists a biholomorphism $\mathrm{F} : X \to Y$ such that the diagram*



commutes up to isotopy, i.e. $\mathrm{F} \circ \varphi \simeq \psi$.

We denote Teichmüller space by $\mathcal{T}(S_g)$ and refer to its elements as marked Riemann surfaces (or marked hyperbolic surfaces). We can use *Teichmüller's Theorem* (see [FM12, Section 11.4]) to equip $\mathcal{T}(S_g)$ with a metric - and consequently a topology. We call this metric the *Teichmüller metric* and denote it by $d_{\mathcal{T}}$. We will go into more detail on the construction of $d_{\mathcal{T}}$ in Section 1.3, where translation surfaces are introduced.

The *moduli space* $\mathcal{M}_g$ is the space of all closed Riemann surfaces of genus g up to biholomorphism, or equivalently, the space of all closed hyperbolic surfaces of genus g up to orientation-preserving isometry. There is a natural surjection from Teichmüller space to moduli space by forgetting the marking, i.e. a map

$$\mathcal{T}(S_g) \twoheadrightarrow \mathcal{M}_g, \ (X, \varphi) \mapsto X.$$

We equip $\mathcal{M}_g$ with the quotient topology. Our next goal is to describe an action of a group on $\mathcal{T}(S_g)$ such that the above map becomes the quotient map of that group action. For

this, we define the *mapping class group.*

Definition 1.4. *The* mapping class group *of S_g is the group of all orientation-preserving self-homeomorphisms of S_g up to isotopy, i.e. up to the equivalence relation $f \sim g$ if and only if f is isotopic to g.*

We denote the mapping class group by $\mathrm{Mod}(S_g)$. Note that another way of describing the mapping class group is to start with the group of orientation-preserving homeomorphisms $\mathrm{Homeo}^+(S_g)$ and collapse each connected component to a point. Here, we equip $\mathrm{Homeo}^+(S_g)$ with the compact-open topology. In particular, if $\mathrm{Homeo}_0(S_g)$ denotes the component of the identity, we have

$$\mathrm{Mod}(S_g) = \mathrm{Homeo}^+(S_g)/\mathrm{Homeo}_0(S_g).$$

This implies that the mapping class group is a discrete group. Furthermore, there is a natural action of $\mathrm{Mod}(S_g)$ on Teichmüller space by changing the marking. We define:

$$\mathrm{Mod}(S_g) \times \mathcal{T}(S_g) \to \mathcal{T}(S_g), \quad (f, (X, \varphi)) \mapsto f \cdot (X, \varphi) \coloneqq (X, \varphi \circ f^{-1}).$$

In the above, there is a slight abuse of notation. Since $f \in \mathrm{Mod}(S_g)$ is only a class of homeomorphisms, the composition $\phi \circ f^{-1}$ is only defined up to isotopy. Due to the definition of Teichmüller space however, this is sufficient to guarantee that the action is well-defined. Finally, note that we have to use the inverse in order to obtain the desired equality $(fh) \cdot (X, \varphi) = f \cdot h \cdot (X, \varphi)$ for $f, h \in \mathrm{Mod}(S_g)$ and $(X, \varphi) \in \mathcal{T}(S_g)$.

Lemma 1.5. *The quotient of the* $\mathrm{Mod}(S_g)$*-action on Teichmüller space is the moduli space, i.e.* $\mathcal{T}(S_g)/\mathrm{Mod}(S_g) = \mathcal{M}_g$.

Proof. Since the action of $\mathrm{Mod}(S_g)$ only changes the marking, but not the Riemann surface of the pair (X, φ), the forgetful map $\mathcal{T}(S_g) \to \mathcal{M}_g$ factors through the quotient. We need to show that $\mathcal{T}(S_g)/\mathrm{Mod}(S_g) \to \mathcal{M}_g$ is injective. This follows from the fact that for any two markings $\varphi, \psi : S_g \to X$, there exists an element $f \in \mathrm{Mod}(S_g)$ with $f \cdot (X, \varphi) = (X, \psi)$. Indeed, set f to be the class of the homeomorphism $\psi^{-1} \circ \varphi$ and the claim follows. □

One can show that the action of the mapping class group is properly discontinuous, meaning that the set $\{f \in \mathrm{Mod}(S_g) \mid f \cdot K \cap K \neq \varnothing\}$ is finite for any compact $K \subset \mathcal{T}(S_g)$. However, it is not a free action. In fact, $f \cdot (X, \varphi) = (X, \varphi)$ implies that f^{-1}, and hence also f, is a biholomorphism with respect to the pulled back structure of X under φ. Since, by a theorem of Hurwitz, the biholomorphism group of a closed Riemann surface is finite,

we obtain that all the point stabilisers of our action are finite.

An important result to mention is that $\mathcal{T}(S_g)$ is homeomorphic to $\mathbb{R}^{6g-6}$. The proof of this statement uses classical hyperbolic geometry and can be found in detail in [FM12]. To obtain an idea about why the result is true, consider the following construction: choose $3g-3$ simple closed curves on S_g that split the surface into a disjoint union of 3-punctured spheres, so-called *pairs of pants* (see Figure 1.12 in Section 1.4 below). One can show that the hyperbolic lengths of these curves, together with so-called twist parameters that describe how the pairs of pants are glued together, determine the hyperbolic metric completely. Hence, for each curve we require two real parameters, i.e. $6g-6$ parameters in total to describe a point in Teichmüller space.

As by the above result Teichmüller space is contractible, its topology is not very interesting to study. However, understanding the Teichmüller metric $d_{\mathcal{T}}$ with which it is equipped has raised a number of fascinating questions. Below, we will describe $d_{\mathcal{T}}$ in more detail and discuss interesting problems concerning it.

We end this section by remarking that the homeomorphism of $\mathcal{T}(S_g)$ to $\mathbb{R}^{6g-6}$ turns Teichmüller space naturally into a manifold. The action of $\mathrm{Mod}(S_g)$ being properly discontinuous implies that $\mathcal{M}_g$ can be equipped with a structure as an orbifold whose singular points correspond to the Riemann surfaces X with non-trivial biholomorphism group.

§ 1.2 | The classification of mapping classes

The aim of this section is to present a classification of mapping classes. In particular, we define three types of mapping classes and show that each $f \in \mathrm{Mod}(S_g)$ is one of those three types. Historically, this discussion was initiated by Nielsen and later completed by Thurston, which is why the classification is usually referred to as the *Nielsen-Thurston classification* (see [Thu88] for more information). The main idea is to imitate the classification of isometries of the hyperbolic plane into elliptic, parabolic and hyperbolic. The latter can be done by analysing the action of an isometry f on the compactification of the hyperbolic plane and realising that f has either a unique fixed point in $\mathbb{H}$, or that there is no fixed point in $\mathbb{H}$, in which case f either fixes a unique point in the boundary $\partial\mathbb{H}$, or exactly two points in $\partial\mathbb{H}$. The three cases then exactly correspond to the classes of elliptic, parabolic or hyperbolic isometries.

Similarly, the classification of mapping classes is based on studying how $\mathrm{Mod}(S_g)$ acts on Teichmüller space $\mathcal{T}(S_g)$. We start by defining the three types of mapping classes.

Definition 1.6. *A mapping class $f \in \mathrm{Mod}(S_g)$ is called* periodic *if it has finite order, i.e. if there is an $n \in \mathbb{N}$ such that $f^n = \mathrm{id}$.*

The *Nielsen realisation theorem* states that any periodic mapping class can be represented by an isometry of some hyperbolic metric, i.e. for a periodic $f \in \mathrm{Mod}(S_g)$ there exists a homeomorphism f' in the isotopy class of f and an element $(X, \varphi) \in \mathcal{T}(S_g)$ such that f' is an isometry with respect to the pulled back metric of X under φ. This can be used to show that the maximal order of a periodic element in $\mathrm{Mod}(S_g)$ is $4g + 2$. Furthermore, such an element is always realisable as follows: take a regular $4g+2$-gon and glue opposite parallel edges via a translation. The result is a genus g surface. Rotating the polygon clockwise by $\frac{2\pi}{4g+2}$ yields a homeomorphism of order $4g + 2$.

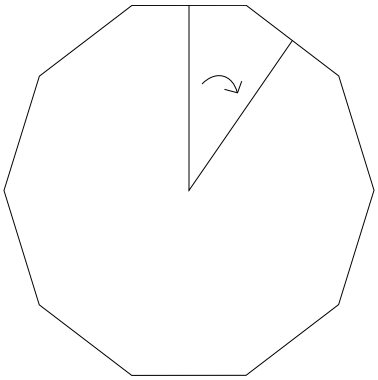


Figure 1.1: The rotation of a 10-gon yields a periodic element of order 10 in $\mathrm{Mod}(S_2)$.

By a *curve* α on S_g we will mean an isotopy class of essential, simple, closed curves. In other words, one can think of α as a non-trivial element of $\pi_1(S_g)$ that can be represented by a loop that doesn't intersect itself. Note that since both $f \in \mathrm{Mod}(S_g)$ and a curve α are only defined up to isotopy, it makes sense to speak about the curve $f(\alpha)$. For two curves α and β, we define their intersection number $i(\alpha, \beta)$ as the minimal number of intersections across all representatives.

Definition 1.7. *A mapping class $f \in \mathrm{Mod}(S_g)$ is called* reducible*, if there exists a finite set of curves $A = \{\alpha_1, ..., \alpha_n\}$ such that $i(\alpha_j, \alpha_k) = 0$ for any $j, k \in \{1, ..., n\}$ and such that f acts as a permutation on the set A.*

If f is reducible and A is a set satisfying the above definition, then A is called a *reduction system* for f. Note that there is always some power of f fixing each curve in a reduction system. The reason for the name *reducible* is because one can cut at the curves of a reduction system to obtain a possibly disconnected surface whose components have lower genus. Observe that f induces a homeomorphism of this new surface, which permutes the components. Hence, in order to understand f, we need to understand its action on

surfaces of lower genus, which means we have *reduced* the problem to a lower complexity one.

For the third type of mapping classes, we first need to introduce the notion of a *measured foliation* of a surface. For us, a foliation will mean a collection of disjoint 1-dimensional submanifolds (called *leaves*) of S_g, together with finitely many points $p_1, ..., p_n \in S_g$ (called *singularities*) such that S_g is the union of all the leaves and singularities and such that locally the following picture arises:

- For any $p \in S_g$ that is not a singularity, there exists a neighbourhood of p and a chart to $\mathbb{R}^2$ such that the foliation in this chart looks like a foliation of $\mathbb{R}^2$ by lines parallel to the x-axis.
- For any singularity $p_i \in S_g$, there exists a $k_i \geqslant 3$, a neighbourhood of p_i and a chart to $\mathbb{R}^2$ such that the foliation in this chart looks like a *k-pronged saddle*.

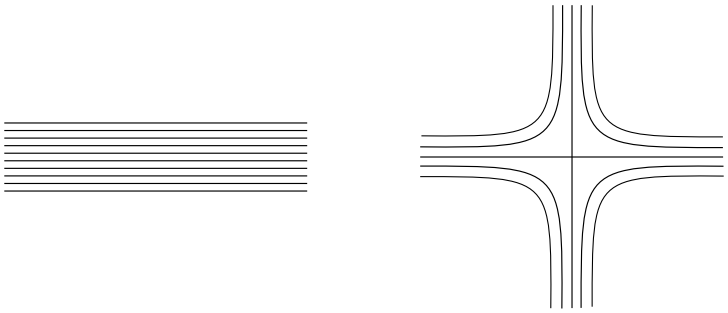


Figure 1.2: Foliated neighbourhood of a regular point and a 4-pronged saddle.

We usually denote a foliation on S_g by $\mathcal{F}$. A *transverse measure* μ to the foliation $\mathcal{F}$ is a function that assigns a positive real number to each arc transverse to the foliation. We require in particular that the arc misses the singularities of the foliation and that it is transverse to each leaf of $\mathcal{F}$ at every point of its interior. We also require that μ is invariant under leaf-preserving isotopy and regular with respect to the Lebesgue measure (see [FM12, Section 11.2] for details).

We call the pair $(\mathcal{F}, \mu)$ a *measured foliation*. A homeomorphism φ of S_g acts on the set of foliations by mapping a foliation $\mathcal{F}$ to its image $\varphi(\mathcal{F})$. Furthermore, if μ is a transverse measure on $\mathcal{F}$, then $\varphi \cdot \mu$, defined by $(\varphi \cdot \mu)(\alpha) := \mu(\varphi^{-1}(\alpha))$, is a transverse measure on $\varphi(\mathcal{F})$. We say that φ fixes a measured foliation if $(\varphi(\mathcal{F}), \varphi \cdot \mu) = (\mathcal{F}, \mu)$. Furthermore, we say that φ stretches or contracts a measured foliation, if $(\varphi(\mathcal{F}), \varphi \cdot \mu) = (\mathcal{F}, \lambda\mu)$ where $\lambda \in \mathbb{R}_+$ is either > 1 or < 1 respectively. Here, $\lambda\mu$ simply means the measure that multiplies the μ-measure of each arc by λ.

Two measured foliations are called transverse if the singularities of the foliations coincide

and if their leaves are transverse at all points. We are now ready to define the last class of mapping classes, the so-called *pseudo-Anosov* ones.

Definition 1.8. *A mapping class $f \in \mathrm{Mod}(S_g)$ is called* pseudo-Anosov, *if there exists a homeomorphism φ in the isotopy class of f, two transverse measured foliations $(\mathcal{F}^u, \mu_u)$, $(\mathcal{F}^s, \mu_s)$ and a real number $\lambda > 1$ such that*

$$(\varphi(\mathcal{F}^u), \varphi \cdot \mu_u) = (\mathcal{F}^u, \lambda\mu_u) \quad \textit{and} \quad (\varphi(\mathcal{F}^s), \varphi \cdot \mu_s) = (\mathcal{F}^s, \lambda^{-1}\mu_s).$$

We call $\mathcal{F}^s$ and $\mathcal{F}^u$ the *stable* and *unstable* foliation and say that f contracts the stable and stretches the unstable foliation by λ. We refer to this λ as the *stretch factor* of the pseudo-Anosov.

As discussed at the start of this section, the following theorem asserts that every element in the mapping class group belongs to one of the three classes defined above.

Theorem 1.9 (Nielsen-Thurston classification)**.** *Let $f \in \mathrm{Mod}(S_g)$. Then f is either periodic, reducible or pseudo-Anosov. Furthermore, no pseudo-Anosov is periodic or reducible.*

Proof. A proof can be found in [FM12, Section 13.6]. □

Note that while the subset of pseudo-Anosov mapping classes is disjoint from the one of periodic or reducible ones, there do exist mapping classes that are both periodic and reducible.

Examples 1.10. Here are three examples of mapping classes that are both periodic and reducible, only reducible or only periodic.

- The *hyperelliptic involution* ι is an order 2 mapping class given by rotating the surface by 180 degrees as in Figure 1.3. The two red curves in the figure form a reduction system. Therefore, ι is both periodic and reducible.

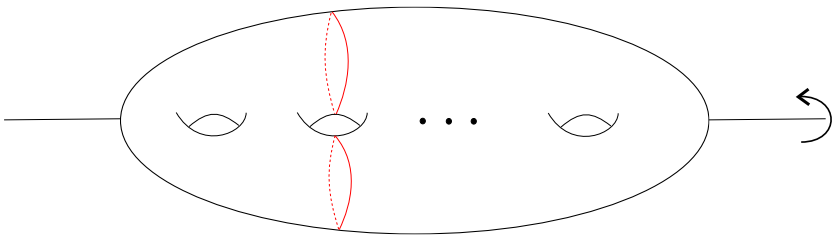


Figure 1.3: The hyperelliptic involution.

- A *Dehn twist* T_α about a curve α is a mapping class represented by a homeomorphism which in an annular neighbourhood $S^1 \times [0,1]$ of α is given by the map $(x,t) \mapsto (e^{2\pi i t}x\,,\, t)$. Figure 1.4 shows the effect of a Dehn twist locally and Figure 1.5 describes the effect of the Dehn twist T_α on a curve intersecting α non-trivially. The curve α itself is fixed by T_α, but no power of T_α fixes any curve β with $i(\alpha,\beta) \neq 0$. Hence, T_α is reducible but not periodic.

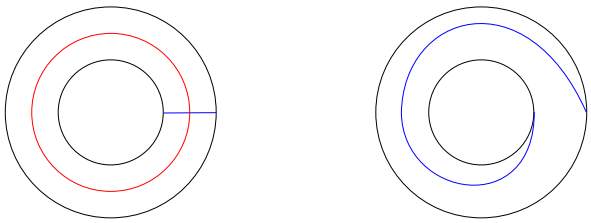


Figure 1.4: A Dehn twist locally in an annular neighbourhood of the curve.

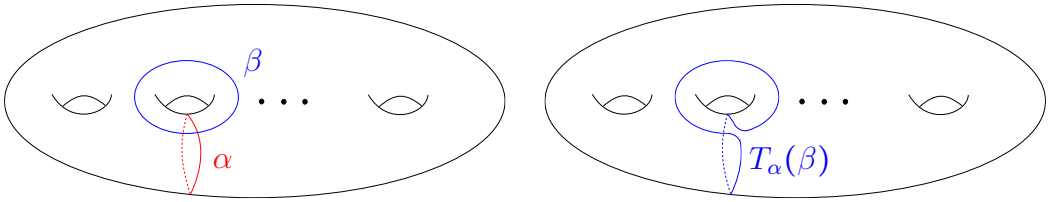


Figure 1.5: The effect of a Dehn twist.

- The aforementioned rotation of a $4g+2$-gon gives an example of a periodic but not reducible mapping class. To see that the rotation is not reducible, note that in the orbit of any curve α there is another curve intersecting α non-trivially.

Keeping in mind that a pseudo-Anosov is never periodic or reducible, we could have also used the following alternative definition:

A *pseudo-Anosov* $f \in \mathrm{Mod}(S_g)$ is an infinite-order element that does not fix any finite set of disjoint curves.

In this regard, the powerful statement proved by Thurston is that any such element admits a description in terms of foliations as above. In order to prove Theorem 1.9, Thurston constructed a boundary for Teichmüller space whose points can be thought of as equivalence classes of measured foliations. This boundary is homeomorphic to a $6g-7$ dimensional sphere, and the union of Teichmüller space and its boundary to a $6g-6$ dimensional compact ball. The action of the mapping class group extends to one on this ball, and we can hence apply Brouwer's Fixed Point Theorem. By studying whether the fixed point lies inside Teichmüller space or on the boundary, Thurston was

able to determine the classification of Theorem 1.9. For example, it follows from the Nielsen realisation theorem that periodic mapping classes are exactly the ones having a fixed point in $\mathcal{T}(S_g)$. The action of a pseudo-Anosov, which resembles the action of a hyperbolic isometry on $\mathbb{H}$, is particularly interesting and will be discussed in more detail in Section 1.5 when introducing translation lengths.

§ 1.3 | Translation surfaces

In this section, we introduce the concepts of a *translation* and a *half-translation surface* (see [AM24] for an introduction to this theory). We provide a geometric as well as an algebraic description of these notions and see how they can be related to the study of Teichmüller space and the Teichmüller metric. Finally, we focus on *square-tiled surfaces*, a special case of translation surfaces, and present a combinatorial description in terms of permutations.

In order to present the geometric definition of a translation surface, consider the following construction.

For $k \in \mathbb{N}$, take k copies $H_1, ..., H_k$ of the closed upper half-plane $\{z \in \mathbb{C} \,|\, \mathrm{Im}(z) \geqslant 0\}$ and glue the H_i along their boundaries $\{z \in \mathbb{C} \,|\, \mathrm{Im}(z) = 0\}$ as follows: Glue the half-line $[0, \infty)$ of H_i to the half-line $(-\infty, 0]$ of H_{i+1} via $x \mapsto -x$ for all $i = 1, \dots, k-1$ and analogously the half-line $[0, \infty)$ of H_k to the half-line $(-\infty, 0]$ of H_1. Let $\overline{H}_k$ denote the resulting space, which we refer to as the *half-plane construction* of order k. One can picture the space $\overline{H}_k$ as a k-pronged saddle (see Figure 1.2 above) where all the lines of the foliation are straight lines.

Definition 1.11. *A* translation surface X *is a closed surface together with a finite collection of points* $\{p_1, \dots, p_n\}$ *and a maximal atlas of charts on* $X \setminus \{p_1, \dots, p_n\}$ *with values in* $\mathbb{C}$ *whose transition maps are translations, i.e. of the form* $z \mapsto z + c$ *for some* $c \in \mathbb{C}$. *Furthermore, we require that around every* p_i *there is a* $k_i \in \mathbb{N}$ *and a neighbourhood of* p_i *that admits a homeomorphism to* $\overline{H}_{2k_i+2}$ *which is a local isometry everywhere except for* p_i.

Definition 1.12. *A* half-translation surface X *is a closed surface together with a finite collection of points* $\{p_1, \dots, p_n\}$ *and a maximal atlas of charts on* $X \setminus \{p_1, \dots, p_n\}$ *with values in* $\mathbb{C}$ *whose transition maps are half-translations, i.e. of the form* $z \mapsto \pm z + c$ *for some* $c \in \mathbb{C}$. *Furthermore, we require that around every* p_i *there is an* $l_i \geqslant -1$ *and a neighbourhood of* p_i *that admits a homeomorphism to* $\overline{H}_{l_i+2}$ *which is a local isometry everywhere except for* p_i.

We remark that an atlas of charts whose transition functions are (half-)translations defines

a flat metric on $X \setminus \{p_1, ..., p_n\}$, which justifies that we can speak about isometries in the above definitions. Moreover, we call the points p_i *singularities* of the (half-)translation surface X and note that the cone angle around some p_i is given by $2\pi(k_i + 1)$ in the translation surface case and $\pi(l_i + 2)$ in the half-translation surface case. The *Gauß–Bonnet theorem* implies that

$$k_1 + ... + k_n = 2g - 2 \text{ and } l_1 + ... + l_n = 4g - 4,$$

where g is the genus of X. Finally, note that since (half-)translations are orientation-preserving, every (half-)translation surface carries a natural orientation coming from the standard orientation of the complex plane.

We now want to give a purely algebraic definition of a translation surface. Let X be a Riemann surface and let $T^*_{1,0}X$ denote the holomorphic cotangent bundle of X. An *abelian differential* (or *holomorphic* 1*-form*) is a holomorphic section of $T^*_{1,0}X$.

Definition 1.13. *A* translation surface *is a pair* (X, ω) *where* X *is a Riemann surface and* ω *a non-zero abelian differential.*

Similarly, one can also describe a half-translation surface algebraically. Let X be a Riemann surface. Let $S(T^*_{1,0}(X))$ denote the symmetric square of the holomorphic cotangent bundle. A *quadratic differential* is a section of $S(T^*_{1,0}(X))$.

Definition 1.14. *A* half-translation surface *is a pair* (X, q) *where* X *is a Riemann surface and* q *a non-zero meromorphic quadratic differential with at most simple poles.*

These definitions sound quite different from the ones given at the start of the section, but turn out to be equivalent. We refer to definitions 1.11 and 1.12 as the geometric definitions and to 1.13 as well as 1.14 as the algebraic definitions.

Lemma 1.15. *The geometric and algebraic definitions of a translation and a half-translation surface, respectively, are equivalent.*

Proof. We will sketch the proof and refer the reader to [Wri15, Proposition 1.6] for details.

Starting with a Riemann surface X and an abelian differential ω, one can show that around every point where ω is non-zero, there exists a neighbourhood U and a local chart $z : U \to \mathbb{C}$ such that $\omega = dz$ on U. For a quadratic differential q, the same is true with dz replaced by dz^2. One can now check that the transition maps of these charts are translations in the case of ω, and half-translations in the case of q. Hence, we constructed an atlas of charts on $X \setminus Z(\omega)$, respectively $X \setminus Z(q)$, where $Z(\cdot)$ denotes the finite set of

zeros of the differentials, with the desired property. Around a zero of ω, there is a chart z such that $\omega = z^k dz$ and around a zero of q, one such that $q = z^l dz^2$. This can be used to show that a neighbourhood of any zero is locally isometric to the desired half-plane construction.

On the other hand, starting with a collection of charts as in the geometric definitions, we observe that since (half-)translations are holomorphic maps, our charts equip $X \setminus \{p_1, \dots, p_n\}$ with a Riemann surface structure. Furthermore, by setting $\omega = dz$ (or $q = dz^2$) on every chart z, we obtain an abelian (or quadratic) differential on $X \setminus \{p_1, \dots, p_n\}$. Finally, using the structure around a singularity p_i, one has to show that there is a unique chart z_i with $z_i(p_i) = 0$, such that the collection of charts together with the z_i turn X into a Riemann surface and such that setting $\omega = z_i^{k_i} dz_i$ around p_i extends the abelian differential to all of X. For a quadratic differential, the same argument with $q = z_i^{l_i} dz_i^2$ instead works. Here, k_i and l_i are the integers coming from the condition on a neighbourhood of the singularities in the geometric definitions. □

By the geometric definitions, it is evident that every translation surface is also a half-translation surface. Accordingly, every abelian differential can be seen as a quadratic differential. Indeed, if ω is an abelian differential, then $\omega \otimes \omega$ defines a quadratic differential, and we view abelian differentials as a subset of quadratic ones in this way.

The singularities of a (half-)translation surface correspond to the zeros of the differential in the algebraic description, and the k_i (or l_i) that determine the cone angles are exactly the order of the corresponding zero. In the algebraic description, we obtain the equalities $\sum_{i=1}^{n} k_i = 2g - 2$ and $\sum_{i=1}^{n} l_i = 4g - 4$ from the *Riemann–Roch theorem.*

Two translation surfaces (X, ω) and (Y, η) are *isomorphic* if there is a biholomorphism $f : X \to Y$ such that $f^*(\eta) = \omega$, i.e. the abelian differential η pulls back to ω under f. Note that an isomorphism of translation surfaces necessarily preserves the number and order of singularities. We define an isomorphism of half-translation surfaces in the analogous way. Let $\mathcal{H}_g$ and $\mathcal{Q}_g$ denote the spaces of all translation, respectively half-translation surfaces of genus g up to isomorphism.

For any partition of $2g-2$, i.e. any tuple of natural numbers $(k_1, \dots, k_n)$ with $k_1 + \dots + k_n = 2g - 2$, let $\mathcal{H}(k_1, \dots, k_n)$ denote the subspace of $\mathcal{H}_g$ of all translation surfaces with n singularities of orders given by the k_i. We call such a subspace a *stratum* and note that $\mathcal{H}_g$ is the disjoint union of all the strata corresponding to a partition of $2g-2$. Similarly, $\mathcal{Q}_g$ decomposes as a disjoint union of analogously defined strata corresponding to a partition of $4g - 4$. Note that in this case, we allow a partition to have entries equal to -1.

The algebraic definition of a (half-)translation surface shows that there is a natural map from $\mathcal{H}_g$, or $\mathcal{Q}_g$, to the moduli space of Riemann surfaces $\mathcal{M}_g$. Namely, the map that forgets the differential, i.e. sends an element (X, ω), or (X, q), to X.

We now want to look at some examples of translation surfaces. A way to construct a translation surface is by drawing a finite number of polygons in the complex plane such that the edges of the polygons are paired so that

- each pair consists of two parallel edges of equal length;
- the two paired edges are on opposite sides of the respective polygons.

Here, being on opposite sides means that the normal unit vector pointing into the interior of the polygon for one of the two edges is the negative of the corresponding vector for the other paired edge.

We glue the paired edges together by the appropriate translation. The second bullet point above ensures that the resulting surface is orientable. Since the polygons are embedded in the complex plane, there is a natural way to obtain charts whose chart transitions are translations at all points except possibly at the corners of the polygons. The corners will all have cone angles that are multiples of 2π, and we hence obtain a translation surface according to the geometric definition with the corners of the polygons corresponding to the singularities. See Figure 1.6 for two such examples. The numbers in the figure indicate the pairing of the edges that should be glued by translation.

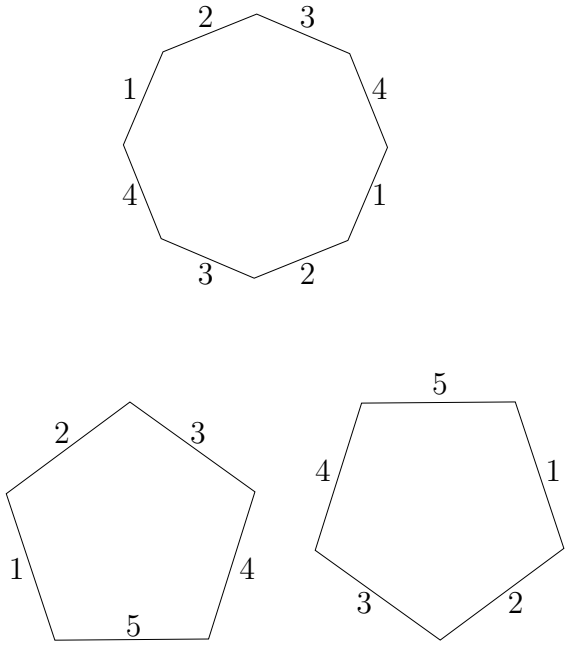


Figure 1.6: Two translation surfaces.

In fact, one can show that any translation surface can be cut into polygons along straight lines connecting singularities. This implies that we can represent any translation surface by polygons in the above way. Two translation surfaces given as collections of polygons are isomorphic if one can be deformed into the other by cutting along straight lines and regluing the resulting pieces along paired edges, as in Figure 1.7. The figure shows an example of a torus, where parallel edges are supposed to be glued together.

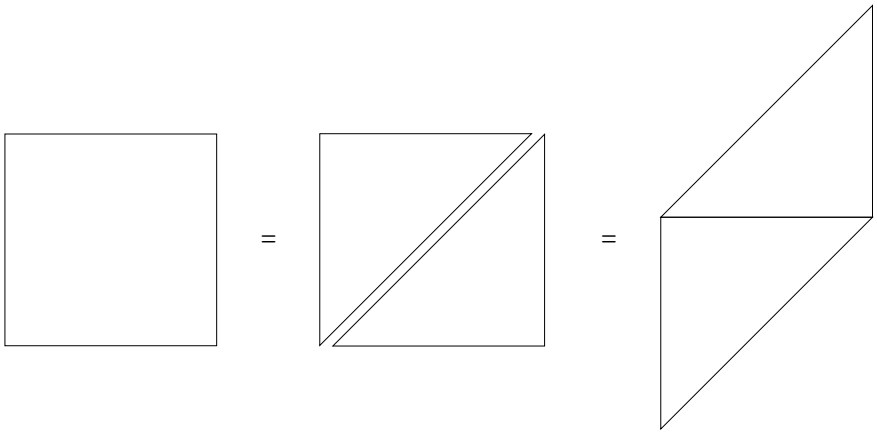


Figure 1.7: The cut and reglue procedure.

The discussion about obtaining translation surfaces via polygons works analogously in the half-translation surface case. Here, we are allowed to glue parallel edges by half-translations, see Figure 1.8. In the figure, the numbered edges are glued by translation, while the arrows indicate the sides glued by a half-translation. The only subtlety to remark is that in this case, the edges glued by translation should still be on opposite sides of the polygons, while the edges glued by half-translation should be on the same side of the polygons.

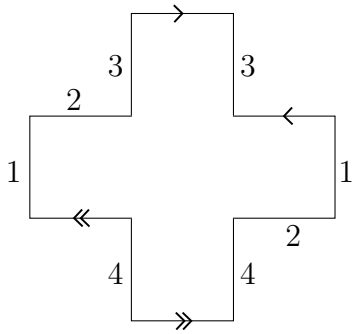


Figure 1.8: A half-translation surface.

We end with the definition of a *cylinder* of a (half-)translation surface which will be used in Chapter 7.

Definition 1.16. *A* cylinder *of a (half-)translation surface is the image of an isometric embedding of* $\mathbb{R}/(c\mathbb{Z}) \times (0, h)$*, which is not the restriction of an isometric embedding of a larger cylinder. The parameters* c *and* h *are called the circumference and height of the cylinder.*

§ 1.3.1 | Quadratic differentials and the Teichmüller metric

We now want to explore the connection between Teichmüller space and quadratic differentials and in particular understand the role of the latter with regard to the Teichmüller metric.

Let X be a Riemann surface of genus g and let $\mathcal{Q}(X)$ be the space of holomorphic quadratic differentials on X. The first observation we make is that $\mathcal{Q}(X)$ is a complex vector space of dimension $3g - 3$ [FM12, Theorem 11.7]. Hence, $\mathcal{Q}(X)$ is a real $6g - 6$ dimensional vector space, matching the dimension of $\mathcal{T}(S_g)$. We want to think of $\mathcal{Q}(X)$ as the tangent space of a point (X, φ) in Teichmüller space. We remark that usually $\mathcal{Q}(X)$ is identified with the cotangent space, but for our purposes it will be more straightforward to think of elements in $\mathcal{Q}(X)$ as tangent vectors.

Recall that each $0 \neq q \in \mathcal{Q}(X)$ is a half-translation surface and can hence be drawn as a collection of polygons in the plane. The one-parameter group

$$\left\{ \begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix} : \ t \in \mathbb{R} \right\}$$

acts by linear maps on $\mathbb{R}^2$ and thus on the collection of polygons. This action preserves the properties that paired edges have, such as being parallel. Hence, for any t the resulting collection of polygons again defines a half-translation surface and in particular a Riemann surface which we denote by X_t. Furthermore, note that if we choose a marking $\varphi : S_g \to X$ for the original Riemann surface $X = X_0$, there is a natural way to carry over this marking to a marking φ_t for X_t. We conclude that for each non-zero $q \in \mathcal{Q}(X)$ and marking φ, we obtain a line $\mathbb{R} \to \mathcal{T}(S_g)$ passing through the point (X, φ) at time 0. It turns out that all these lines are embeddings of $\mathbb{R}$ and we define them to be the geodesics of Teichmüller space.

Teichmüller's Theorem guarantees that for any two points (X, φ), $(Y, \psi) \in \mathcal{T}(S_g)$ there is a unique geodesic passing through both of these points. Let q_X be a quadratic differential on X such that it defines the geodesic passing through (X, φ) at time 0 for which $(Y, \psi) = (X_t, \varphi_t)$ for some $t > 0$. We define the Teichmüller metric to be

$$d_{\mathcal{T}}((X, \varphi), (Y, \psi)) := t.$$

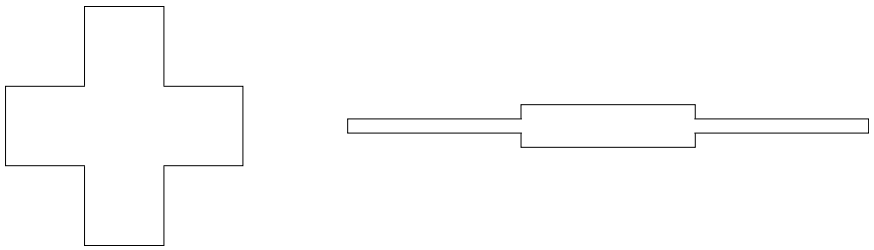


Figure 1.9: The effect of the $\begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}$-action on a translation surface.

Due to Teichmüller's theorem and the theory of quadratic differentials, $d_{\mathcal{T}}$ is a well-defined metric. Furthermore, $d_{\mathcal{T}}$ is a complete metric on Teichmüller space.

We end this subsection by remarking that the action of the one-parameter group described above can be generalised to an action of the group $\mathrm{GL}_2^+(\mathbb{R})$ on the space of quadratic or abelian differentials. Indeed, given any (half-)translation surface (X, ω) and a matrix $A \in \mathrm{GL}_2^+(\mathbb{R})$, we can define $A \cdot (X, \omega)$ as the (half-)translation surface arising from letting A act on a polygon representation of (X, ω). Note that this action preserves the number and orders of the singular points of X, which implies that we have an action of $\mathrm{GL}_2^+(\mathbb{R})$ on any stratum $\mathcal{H}(k_1, \ldots, k_n)$ or $\mathcal{Q}(k_1, \ldots, k_n)$. Studying the subspaces that arise as (closures of) orbits of this action has led to many interesting questions and groundbreaking results, some of which we mention in Chapter 7. The stabiliser of (X, ω) under the action is usually denoted by $\mathrm{SL}(X, \omega)$ and referred to as the *Veech group* of (X, ω). The Veech group is always a subgroup of $\mathrm{SL}_2(\mathbb{R})$. In the case when $\mathrm{SL}(X, \omega)$ is a lattice, i.e. has finite covolume, we call the corresponding orbit of (X, ω) a *Teichmüller curve*. The condition of obtaining a Teichmüller curve is equivalent to requiring the orbit to be closed.

There is a different but equivalent way of defining the Veech group. Let (X, ω) be a translation surface with singularities $p_1, \ldots, p_n$. An *affine homeomorphism* of (X, ω) is a homeomorphism of X that sends singularities to singularities and is locally affine on $X \setminus \{p_1, \ldots, p_n\}$. Here, locally affine means that every point has a chart neighbourhood on which the homeomorphism is given by a map of the form $x \mapsto Ax + b$ where $A \in \mathrm{GL}_2(\mathbb{R})$ with $\det(A) = \pm 1$ and $b \in \mathbb{R}^2$. Note that since chart transitions are translations, the matrix A is the same in every chart. Hence, we have a well-defined map

$$d \colon \mathrm{Aff}^+(X, \omega) \to \mathrm{SL}_2(\mathbb{R}),$$

where $\mathrm{Aff}^+(X, \omega)$ is the group of orientation-preserving affine homeomorphisms. The Veech group can be defined as the image of this map, i.e. $\mathrm{SL}(X, \omega) = \mathrm{image}(d)$. The kernel of the above map is the group $\mathrm{Aut}(X, \omega)$ of so-called *translation automorphisms*.

Overall, for any translation surface (X, ω), we obtain a short exact sequence of the form

$$1 \to \mathrm{Aut}(X, \omega) \to \mathrm{Aff}^+(X, \omega) \to \mathrm{SL}(X, \omega) \to 1.$$

For example, for (X, ω) given by the unit square with opposite sides identified, i.e. the *standard torus*, we have

$$1 \to \mathbb{R}^2/\mathbb{Z}^2 \to \mathrm{Aff}^+(X, \omega) \to \mathrm{SL}_2(\mathbb{Z}) \to 1,$$

and every orientation-preserving affine homeomorphism is of the form $x \to Ax + b$ with $A \in \mathrm{SL}_2(\mathbb{Z})$ and $b \in \mathbb{R}^2$. See [HS06] for more information on Veech groups and affine homeomorphisms.

§ 1.3.2 | Square-tiled surfaces

We turn our attention to *square-tiled surfaces*. These are examples of translation surfaces which are built by gluing squares together.

Let $(C_i)_{i=1}^n$ be a finite collection of unit squares embedded pairwise disjointly into $\mathbb{C}$ such that one square is standardly embedded as $[0, 1] \times [0, 1]$ and every other embedding differs from the standard one by some translation.

A *square-tiled surface* is the quotient space of $\bigcup_{i=1}^n C_i$ arising from a choice of gluing of sides of the C_i such that

- each side is glued to exactly one parallel side by translation,
- if a side is on the left (resp. top) of a square, then it is glued to a side on the right (resp. bottom) of a square.

Note that these conditions are analogous to the ones required for the general polygon construction of translation surfaces. Every square-tiled surface is a closed, orientable surface that comes with a subdivision into squares. As defined above, a square-tiled surface is not necessarily connected. For example, gluing opposite sides of each square C_i yields a disjoint union of n tori. We will usually only consider connected square-tiled surfaces. We will use the abbreviation STS for a square-tiled surface. By construction, a STS is a translation surface.

Since every square of a STS admits a natural map to the torus by identifying opposite sides, a STS admits a map to the torus which is a branched cover, branched over the image point of all the corners of the C_i. Conversely, every finite-degree branched cover of the torus with branching over exactly one point can be described as a STS.

There is a purely combinatorial way to describe square-tiled surfaces which will be of use. Let X be a STS built out of n squares $C_1, \ldots, C_n$. We assign two permutations in the symmetric group S_n to X as follows:

For $i \in \{1, \ldots, n\}$, let $\pi_h(i)$ be the index describing the square that the right side of C_i is glued to. This yields a permutation $\pi_h \in S_n$ which we call the *horizontal permutation*. Similarly, for $i \in \{1, \ldots, n\}$, let $\pi_v(i)$ be the index describing the square that the bottom side of C_i is glued to. This yields a permutation $\pi_v \in S_n$ which we call the *vertical permutation*. The permutations π_h and π_v depend on the choice of numbering of the squares. A different choice yields conjugate permutations.

Conversely, given two permutations in S_n we can construct a STS by assigning to each $i \in \{1, \ldots, n\}$ a square and gluing sides as the permutations suggest. Hence, square-tiled surfaces with n squares are in one-to-one correspondence with the choice of two permutations in S_n up to simultaneous conjugation. This can also be described as the set

$$\mathrm{Hom}(\mathbb{F}_2, S_n)/S_n,$$

where $\mathbb{F}_2$ is the free group on two generators and S_n acts on the set of homomorphisms by conjugation. This is a special case of a more general discussion of branched covers of the torus (see, for example, [EO01, Section 2.1]).

The combinatorial description of a STS also helps in determining the stratum of the STS. As the following lemma shows, the commutator of the two permutations carries all the information needed. We denote the conjugacy class of permutations with m cycles of lengths $k_1, \ldots, k_m$ by $C(k_1, \ldots, k_m)$.

Lemma 1.17. *Let X be a square-tiled surface defined by the permutations $\pi_h, \pi_v \in S_n$. Let $C(k_1, \ldots, k_m)$ be the conjugacy class that contains the commutator $[\pi_h, \pi_v]$. Then X is an element of the stratum $\mathcal{H}(k_1 - 1, \ldots, k_m - 1)$.*

Proof. For a permutation $\sigma \in S_n$ and $i \in \{1, \ldots, n\}$, we say that *doing one step* of σ at i leads us to j, if $\sigma(i) = j$. Note that doing one step of the commutator $[\pi_h, \pi_v] = \pi_h^{-1}\pi_v^{-1}\pi_h\pi_v$ at any square i of X leads to a possibly different square j where we end up by starting at the original square i and walking the sequence down-right-up-left. The bottom-right corner of square i is the same point p as the bottom-right corner of square j, and the path traced out by going from i to j shows that the cone angle around p is at least 2π. In fact, the cone angle is exactly 2π if and only if $i = j$, and bigger otherwise. Continuing to do one step in the commutator at j gets us potentially to another square and so on. We continue this procedure until we are back at square i, i.e. until the cycle of $[\pi_h, \pi_v]$ has

closed up. Each step adds another 2π to the cone angle, and we conclude that p is a point with cone angle equal to $2\pi k$, where k is the length of the cycle of $[\pi_h, \pi_v]$ containing i. Recalling the definition of a translation surface, this means that p is a singularity of order $k-1$. Note that in the case where $i = j$, and therefore the cone angle equal to 2π, the point p is actually a regular point, or equivalently a singularity of order 0.

Overall, we see that for each cycle of the commutator, we obtain a singularity of order one less than the length of that cycle, which is what we wanted to show. □

Remark 1.18. The above lemma, together with the Gauß-Bonnet formula gives a way to compute the genus of a STS. Given X as in the lemma and letting $C(k_1, \ldots, k_m)$ be the conjugacy class of the corresponding commutator, we have

$$g(X) = \frac{k_1 + \cdots + k_m - m}{2} + 1.$$

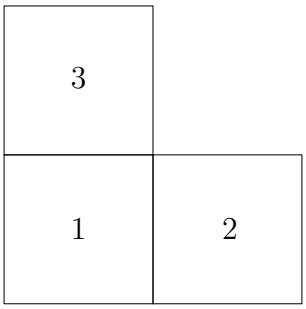


Figure 1.10: The *L-shaped* square-tiled surface of genus 2 corresponding to the permutations $\pi_h = (1\ 2)\,(3)$ and $\pi_v = (1\ 3)\,(2)$. Note that $[\pi_h, \pi_v] = (1\ 2\ 3)$ which implies that the L-shape is in $\mathcal{H}(2)$.

One reason square-tiled surfaces are of interest is that one can derive information about the volume of strata of translation surfaces by counting square-tiled surfaces (see, for example, [Zor02]). Here, the volume is the so-called *Masur–Veech volume* defined by equipping a stratum with the local *period coordinates* and pulling back the Lebesgue measure under these coordinates. Square-tiled surfaces then correspond to integral points under these local coordinates.

The Veech group of a STS is a finite index subgroup of $\mathrm{SL}(2,\mathbb{Z})$ which is a lattice in $\mathrm{SL}(2,\mathbb{R})$. Consequently, the orbit of any STS is closed and defines a Teichmüller curve.

§ 1.4 | The curve graph and the pants graph

In this section, we present the definitions of two graphs that are naturally associated to a surface. These are the *curve graph* and the *pants graph*. The mapping class group acts

on both. Intuitively, these graphs can be thought of as discrete models for Teichmüller space, and we will see how they are related to one another.

Before getting to the main part of this section, we introduce the notions of a *coarsely Lipschitz map* and of the *asymptotic behaviour* of sequences.

A map $\varphi\colon X \to Y$ between metric spaces is *coarsely Lipschitz* if there exist $C, K \geqslant 0$ such that

$$d_Y(\varphi(x_1), \varphi(x_2)) \leqslant C d_X(x_1, x_2) + K$$

for all $x_1, x_2 \in X$. The *optimal Lipschitz constant* of φ is the infimum over all $C \geqslant 0$ for which there is some $K \geqslant 0$ satisfying the coarsely Lipschitz condition.

We say that two sequences $(a_n)_{n\in\mathbb{N}}, (b_n)_{n\in\mathbb{N}} \subset \mathbb{R}$ have the same *asymptotic behaviour*, denoted $a_n \asymp b_n$, if there is a $C > 1$ such that $\frac{a_n}{b_n} \in [\frac{1}{C}, C]$ for all $n \in \mathbb{N}$. For us, the sequences will usually depend on the genus, i.e. we will consider sequences of the form (a_g) with $g \in \mathbb{N}_{\geqslant 2}$.

The curve graph and the pants graph can be defined more generally for surfaces with punctures or boundary, but we will only be dealing with closed surfaces. Let S_g be a closed, oriented surface of genus g. Recall that by *curve* we refer to an isotopy class of essential, simple, closed curves on S_g. Below, we will sometimes be imprecise when distinguishing between curves and their representatives, but the context of the discussion will always make sure that no confusion can arise.

Definition 1.19. *The* curve graph $\mathcal{C}(S_g)$ *is a graph whose vertices correspond to curves. Two vertices are joined by an edge if and only if they correspond to curves for which there exist representatives that are disjoint.*

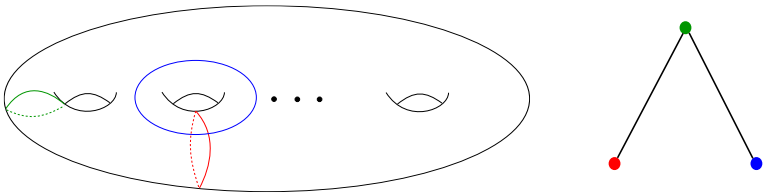


Figure 1.11: A small portion of the curve graph and the corresponding curves.

We equip $\mathcal{C}(S_g)$ with the shortest path metric, where we impose each edge to have length 1. We denote the metric by $d_{\mathcal{C}}$. The curve graph is a locally infinite, infinite diameter, hyperbolic graph [MM99, Theorem 1.1]. A mapping class f acts on the vertices of $\mathcal{C}(S_g)$ by sending a curve α to its image under f, i.e. $f \cdot \alpha = f(\alpha)$. Since this action preserves the property of having disjoint representatives of curves, we conclude that $\mathrm{Mod}(S_g)$ acts on the curve graph by isometries.

We now want to see how Teichmüller space and the curve graph are related. The *systole* of a hyperbolic surface X is a geodesic of minimal length. Note that a systole is in general not unique. The *systole map* is a function

$$\mathrm{sys}\colon \mathcal{T}(S_g) \to \mathcal{C}(S_g)$$

that takes a point in Teichmüller space to one of its systoles. Concretely, if (X, φ) is a point in $\mathcal{T}(S_g)$ and γ a systole, then $\mathrm{sys}(X, \varphi) = \varphi^{-1}(\gamma)$. One could also think of this map as a set-valued map, i.e. one that takes a point in $\mathcal{T}(S_g)$ to the set of all its systoles. This set is always a finite set of disjoint curves, i.e. corresponds to a finite set of diameter $\leqslant 1$ in the curve graph. For our purposes, we picture the systole map as picking a curve in the corresponding set for all elements of Teichmüller space. Since we will be interested in coarse properties of the map, it won't matter which curves are picked.

Masur–Minsky show that the systole map is coarsely Lipschitz, i.e. there are $C, K > 0$ such that

$$d_{\mathcal{C}}(\mathrm{sys}(X, \varphi), \mathrm{sys}(Y, \psi)) \leqslant C\, d_{\mathcal{T}}((X, \varphi), (Y, \psi)) + K,$$

for all $(X, \varphi), (Y, \psi) \in \mathcal{T}(S_g)$ [MM99]. One can then ask what the optimal Lipschitz constant, denoted by κ_g, is. The answer to this question is given in [GHKL13]. The authors determine the asymptotic behaviour of the optimal constants and show that $\kappa_g \asymp \frac{1}{\log(g)}$.

A *pants decomposition* of S_g is a collection of curves which have representatives that are pairwise disjoint and such that each complementary region is a three-punctured sphere, i.e. a *pair of pants*, as in Figure 1.12. Note that a pants decomposition has to consist of $3g - 3$ curves and that these curves split the surface into $2g - 2$ pairs of pants.

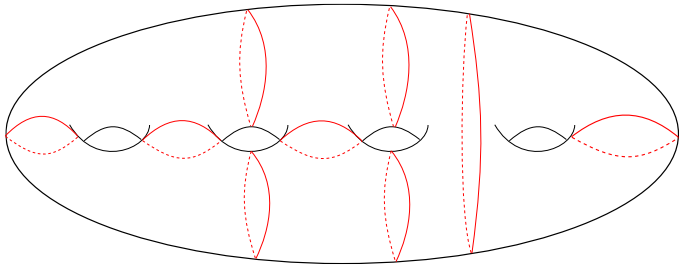


Figure 1.12: A pants decomposition of a genus 4 surface.

Let $\{\alpha_1, \ldots, \alpha_{3g-3}\}$ be a pants decomposition. For each curve α_i, after cutting S_g at all other α_j with $j \neq i$, the curve α_i lies in a subsurface which is either a once-punctured torus or a four-punctured sphere. In the first case, we can replace α_i by a curve α_i' intersecting α_i exactly once and obtain a new pants decomposition $\{\alpha_1, \ldots, \alpha_i', \ldots, \alpha_{3g-3}\}$. In the sec-

ond case, we replace α_i by a curve intersecting it exactly twice and again obtain a new pants decomposition as before. We refer to these two replacement moves as *elementary moves.*

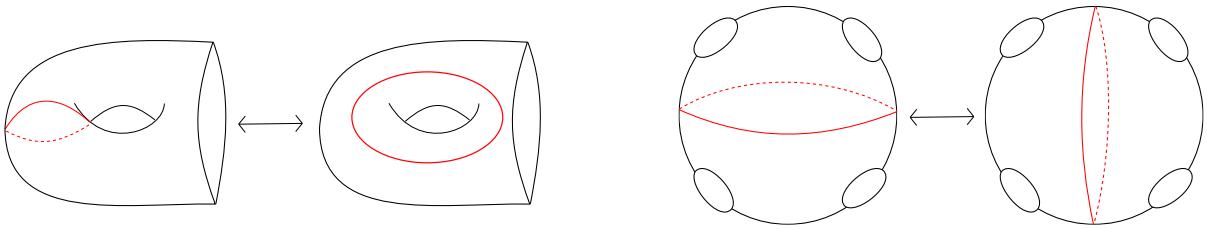


Figure 1.13: Examples of the two elementary moves.

Definition 1.20. *The pants graph $\mathcal{P}(S_g)$ is the graph whose vertices correspond to pants decompositions and edges to elementary moves.*

As in the curve graph case, we equip the pants graph with the shortest path metric where each edge has length 1 and denote it by $d_{\mathcal{P}}$.

There is a natural map from the pants graph to the curve graph, which takes a pants decomposition to one of its curves. We denote this map by

$$p : \mathcal{P}(S_g) \to \mathcal{C}(S_g).$$

In Chapter 6, Lemma 6.6, we show that this map is coarsely Lipschitz.

Finally, we remark that the pants graph $\mathcal{P}(S_g)$ and Teichmüller space $\mathcal{T}(S_g)$ equipped with the so-called *Weil-Petersson metric* are quasi-isometric as shown by Brock [Bro03, Theorem 1.1]. We go into more detail about the Weil-Petersson metric in Chapter 6.

§ 1.5 | Translation lengths and pseudo-Anosov actions

The purpose of this section is to discuss the action of the mapping class group $\mathrm{Mod}(S_g)$ on the spaces introduced above. In particular, we want to see how the different types of elements, i.e. periodic, reducible and pseudo-Anosov, act on Teichmüller space $\mathcal{T}(S_g)$, the curve graph $\mathcal{C}(S_g)$, or the pants graph $\mathcal{P}(S_g)$. In order to quantify the action of an element of $\mathrm{Mod}(S_g)$ on either space, we introduce notions of *translation lengths* of isometries of metric spaces. We will focus on the so-called *stable translation length*, which is often also referred to as the asymptotic translation length. The most interesting case turns out to be the one of pseudo-Anosovs, and we shall discuss it in more detail.

Let (X, d) be a metric space and f an isometry of X.

Definition 1.21. *The* translation length *of f is defined as*

$$\mathcal{L}(f) \coloneqq \inf_{x \in X} d(x, f(x)).$$

Definition 1.22. *The* stable translation length *is given by*

$$l(f) \coloneqq \liminf_{n \to \infty} \frac{d(x, f^n(x))}{n},$$

where $x \in X$ is any point.

In fact, since $(d(x, f^n(x)))_{n \in \mathbb{N}}$ is a subadditive sequence, the sequence $(\frac{d(x, f^n(x))}{n})_{n \in \mathbb{N}}$ converges, so we could replace $\liminf_{n \to \infty}$ by $\lim_{n \to \infty}$ in the definition, but this will not play a role in the rest of this work.

Both of the above notions are a measure of how much an isometry translates the space. Intuitively, while the translation length can be thought of as the minimal displacement, the stable translation length quantifies how much the space moves on average. A priori, it is not clear that the definition of the latter does not depend on the choice of the point $x \in X$.

Lemma 1.23. *The stable translation length is well defined.*

Proof. Let $y \neq x$ be any other point in X. Then, by the triangle inequality

$$d(x, f^n(x)) \leq d(x, y) + d(y, f^n(y)) + d(f^n(y), f^n(x)),$$

for any $n \in \mathbb{N}$. Now, since f is an isometry of X, we have $d(x, y) = d(f^n(x), f^n(y))$ and therefore

$$d(x, f^n(x)) \leq 2d(x, y) + d(y, f^n(y)).$$

Dividing by n and taking the lim inf on both sides yields

$$\liminf_{n \to \infty} \frac{d(x, f^n(x))}{n} \leqslant \liminf_{n \to \infty} \frac{d(y, f^n(y))}{n}.$$

By reversing the roles of x and y, one obtains the opposite inequality, which altogether shows that the definition of the stable translation length is independent of the choice of the point x. □

Note that both translation lengths are, by definition, non-negative real numbers. If f has a fixed point, then both $\mathcal{L}(f)$ and $l(f)$ are 0. We also have the following observation:

Lemma 1.24. *For any isometry f of X, it holds that $l(f) \leqslant \mathcal{L}(f)$.*

Proof. Let $x \in X$ and $n \in \mathbb{N}$. Then $d(x, f^n(x)) \leqslant nd(x, f(x))$ by the triangle inequality

and the fact that f is an isometry. Hence, $l(f) \leqslant d(x, f(x))$. Since x was arbitrary, we conclude $l(f) \leqslant \mathcal{L}(f)$. □

Assume now that the space X is either $\mathcal{T}(S_g)$ equipped with the Teichmüller metric $d_{\mathcal{T}}$ or one of the graphs $\mathcal{C}(S_g)$ or $\mathcal{P}(S_g)$ equipped with the shortest path metric $d_{\mathcal{C}}$ or $d_{\mathcal{P}}$. The mapping class group acts on all of the above while preserving the corresponding metric, and we can hence interpret each element $f \in \mathrm{Mod}(S_g)$ as an isometry of any of these spaces. Thus, we obtain translation lengths for all cases, which we denote by $\mathcal{L}_{\mathcal{T}}(f), \mathcal{L}_{\mathcal{C}}(f), \mathcal{L}_{\mathcal{P}}(f)$ and $l_{\mathcal{T}}(f), l_{\mathcal{C}}(f), l_{\mathcal{P}}(f)$ respectively.

We remark that the *extended mapping class group* $\mathrm{Mod}^{\pm}(S_g)$, i.e. the analogously defined group where orientation-reversing homeomorphisms are also allowed, is in fact isomorphic to the isometry groups of all three spaces $\mathcal{T}(S_g), \mathcal{C}(S_g)$ and $\mathcal{P}(S_g)$ (see [FM12, Section 12.1], [Iva97, Theorem 1], and [Mar04, Theorem 1]). Hence, the discussion below covers essentially all the isometries of the spaces considered.

Here, we discuss translation lengths for the actions on Teichmüller space and the curve graph, while we refer the reader to Chapter 6 for a discussion about the pants graph.

Teichmüller space. We first analyse the case where $X = \mathcal{T}(S_g)$. We have already discussed that if f is a periodic mapping class, then f fixes a point in Teichmüller space. It follows that $\mathcal{L}_{\mathcal{T}}(f) = l_{\mathcal{T}}(f) = 0$. If f is reducible, then we can find a maximal reduction system, i.e. a maximal set of disjoint curves that f preserves. Some power of f will then fix each curve in the reduction system and act on the complementary components as either a periodic or a pseudo-Anosov mapping class. If there are pseudo-Anosov components, then $\mathcal{L}_{\mathcal{T}}(f) > l_{\mathcal{T}}(f) > 0$. Otherwise, $\mathcal{L}_{\mathcal{T}}(f) = l_{\mathcal{T}}(f) = 0$. In any case, $\mathcal{L}_{\mathcal{T}}(f)$ is not realised, i.e. there does not exist $x \in \mathcal{T}(S_g)$ such that $\mathcal{L}_{\mathcal{T}}(f) = d_{\mathcal{T}}(x, f(x))$ (see [FM12, Section 13.6] for details).

In the case of f being pseudo-Anosov, there are two transverse foliations which are fixed by f. This was used by Thurston to argue that f fixes exactly two points on the boundary of Teichmüller space and consequently that there is a unique geodesic $\gamma \subset \mathcal{T}(S_g)$ on which f acts by translation. Hence, by choosing $x \in \gamma$, we obtain

$$l_{\mathcal{T}}(f) = \liminf_{n\to\infty} \frac{d_{\mathcal{T}}(x, f^n(x))}{n} = \liminf_{n\to\infty} \frac{n\, d_{\mathcal{T}}(x, f(x))}{n} = d_{\mathcal{T}}(x, f(x)).$$

It now also follows that $\mathcal{L}_{\mathcal{T}}(f) = l_{\mathcal{T}}(f) = d_{\mathcal{T}}(x, f(x)) > 0$ for any $x \in \gamma$, since

$$d_{\mathcal{T}}(x, f(x)) = l_{\mathcal{T}}(f) \leqslant \mathcal{L}_{\mathcal{T}}(f) \leqslant d_{\mathcal{T}}(x, f(x)).$$

Curve graph. One of the main breakthrough results in the study of the curve graph is that $\mathcal{C}(S_g)$ is a hyperbolic space [MM99]. Therefore, its isometries can be categorised as *elliptic*, *parabolic* and *hyperbolic* in the sense of Gromov [Gro87]. Here, an elliptic isometry f is one where each orbit $\{f^n(x) \mid x \in \mathcal{C}(S_g), n \in \mathbb{Z}\}$ is bounded. A parabolic isometry has a unique limit point in the Gromov boundary to which every orbit converges. Finally, a hyperbolic isometry has exactly two limit points to which every orbit converges.

If f is periodic or reducible, then there is a finite set of curves fixed by f. It follows that f is an elliptic isometry and $l_{\mathcal{C}}(f) = 0$. If f fixes a curve, then $\mathcal{L}_{\mathcal{C}}(f) = 0$. Otherwise, $\mathcal{L}_{\mathcal{C}}(f) = 1$ unless f is elliptic but not reducible, in which case we don't know if $\mathcal{L}_{\mathcal{C}}(f)$ could be larger.

Bowditch showed that the action of a pseudo-Anosov f on the curve graph is similar to the one on Teichmüller space. There is a geodesic $\gamma \subset \mathcal{C}(S_g)$ and a power $n \in \mathbb{N}$ such that f^n fixes γ setwise and acts on it by translation [Bow08]. It follows that f is a hyperbolic isometry. Furthermore, $l_{\mathcal{C}}(f^n) = \mathcal{L}_{\mathcal{C}}(f^n) = d_{\mathcal{C}}(x, f^n(x)) \in \mathbb{N}$ for any $x \in \gamma$. From this, it also follows that $\mathcal{L}_{\mathcal{C}}(f)$ is an integer $\geqslant 1$ and that $l_{\mathcal{C}}(f) = \frac{l_{\mathcal{C}}(f^n)}{n}$ is a rational number.

We conclude by remarking that the above discussion shows that the curve graph does not admit any parabolic isometries.

We summarise the above discussion in the following Table 1.1.

	periodic	**reducible**	**pseudo-Anosov**
On $\mathcal{T}$	$\exists$ fixed point	depends on pA components	preserves geodesic
$\mathcal{L}_{\mathcal{T}}$	0 *realised*	$\begin{cases} 0, & \text{if there is no pseudo-Anosov component} \\ > 0, & \text{if there is a pseudo-Anosov component} \end{cases}$ *not realised*	> 0 *realised*
$l_{\mathcal{T}}$	0	$\begin{cases} 0, & \text{if there is no pseudo-Anosov component} \\ > 0, & \text{if there is a pseudo-Anosov component} \end{cases}$	$= \mathcal{L}_{\mathcal{T}} > 0$
On $\mathcal{C}$	*elliptic*	*elliptic*	*hyperbolic*
$\mathcal{L}_{\mathcal{C}}$	$\in \mathbb{N}$ *realised*	$\in \{0, 1\}$ *realised*	$\in \mathbb{N}$ *realised*
$l_{\mathcal{C}}$	0	0	$\in \mathbb{Q}_{>0}$

Table 1.1: Translation lengths of types of mapping classes on $\mathcal{T}(S_g)$ and $\mathcal{C}(S_g)$.

§ 1.5.1 | An example of a pseudo-Anosov

We consider an example of a mapping class of the torus $T = S_1$. In the case of a torus, foliations don't have singular points. Picturing the torus as a square with opposite sides identified, one can see that up to isotopy, every foliation is given by foliating the square by parallel straight lines. If the angle between a side of a square and a line of the foliation is rational, then all leaves are embedded circles. If the angle is irrational, then all leaves are embedded real lines and each leaf is dense in T.

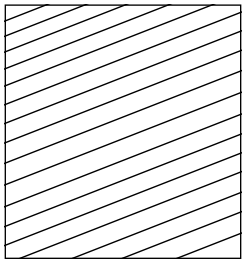


Figure 1.14: A foliation of the torus.

A pseudo-Anosov (in fact, in the case of a torus just called *Anosov*) homeomorphism is one for which there are two such transverse foliations which are stretched or contracted, respectively. Note that the foliations necessarily have to consist of dense leaves.

Every homeomorphism of T induces an isomorphism on homology $H_1(T) \cong \mathbb{Z}^2$, which induces a map $\mathrm{Mod}(T) \to \mathrm{Aut}(H_1(T)) \cong \mathrm{GL}_2(\mathbb{Z})$. Since the mapping class group consists only of orientation-preserving homeomorphisms, this map has image inside $\mathrm{SL}_2(\mathbb{Z})$. In fact, it turns out that this map induces an isomorphism between $\mathrm{Mod}(T)$ and $\mathrm{SL}_2(\mathbb{Z})$. Another way to think of this is again picturing the torus as a square and letting a matrix of $\mathrm{SL}_2(\mathbb{Z})$ act on this square. The image is a parallelogram which can be cut along straight lines connecting the corners and reglued along sides that are identified until obtaining the original square. For example, the matrix

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

would lead to the reverse procedure in Figure 1.7 above. Mapping a point of T to the corresponding point after acting by the matrix and performing the cut and reglue procedure, we obtain a homeomorphism. Hence, every $A \in \mathrm{SL}_2(\mathbb{Z})$ induces a mapping class of T and every mapping class is of this form. In more abstract language, the fact that we are able to cut and reglue back to the original square corresponds exactly to the fact that $\mathrm{SL}_2(\mathbb{Z})$ is the Veech group of the square torus.

Consider the mapping class f that corresponds to the matrix

$$A = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z}).$$

In fact, since

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

it can be seen that the mapping class f is the composition of the two Dehn twists $T_\alpha \circ T_\beta^{-1}$, where α and β are the standard horizontal and vertical curves on the torus as in Figure 1.15.

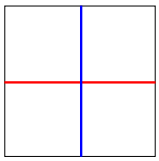


Figure 1.15: The horizontal and vertical curves on the torus in red and blue, respectively.

Let $\varphi = \frac{1}{2}(1 + \sqrt{5})$ be the golden ratio and $\lambda = \frac{1}{2}(3 + \sqrt{5})$. Note that

$$v_1 := (\varphi, 1), \; v_2 := (-\varphi^{-1}, 1)$$

are the eigenvectors of A corresponding to the eigenvalues λ and λ^{-1}.

Define the foliations $\mathcal{F}_1, \mathcal{F}_2$ on T as the straight line foliations on the square torus parallel to v_1 and v_2, respectively. Since the vectors are orthogonal, the corresponding foliations are transverse. Furthermore, since v_1, v_2 are the eigenvectors of A, the mapping class f preserves the two foliations and stretches $\mathcal{F}_1$, respectively contracts $\mathcal{F}_2$, by the factor λ. Hence, f is Anosov with stretch factor λ.

The Teichmüller space of the torus is isometric to the hyperbolic plane $\mathbb{H}$ scaled by the factor $\frac{1}{2}$, i.e. we can identify $\mathcal{T}(T)$ with the upper half-plane $\mathbb{H}$ and for any two points $x, y \in \mathbb{H}$, we have

$$d_{\mathcal{T}}(x, y) = \frac{1}{2} d_{\mathbb{H}}(x, y).$$

The mapping class group acts on $\mathcal{T}(T)$ by Möbius transformations, i.e. an element

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z}) \cong \mathrm{Mod}(T)$$

acts by the isometry

$$z \mapsto \frac{az+b}{cz+d},$$

for $z \in \mathbb{H}$ (see [FM12, proof of Theorem 11.20]).

In particular, the mapping class f acts by the map

$$z \mapsto \frac{2z+1}{z+1}$$

which is a hyperbolic isometry that fixes the points φ, $-\varphi^{-1} \in \partial\mathbb{H}$ and acts on the geodesic connecting the two fixed points by translation. Hence, we know from the general discussion above that

$$l_{\mathcal{T}}(f) = \mathcal{L}_{\mathcal{T}}(f) = \frac{1}{2}\, d_{\mathbb{H}}(z, f(z)),$$

for any point z on the geodesic connecting φ and $-\varphi^{-1}$.

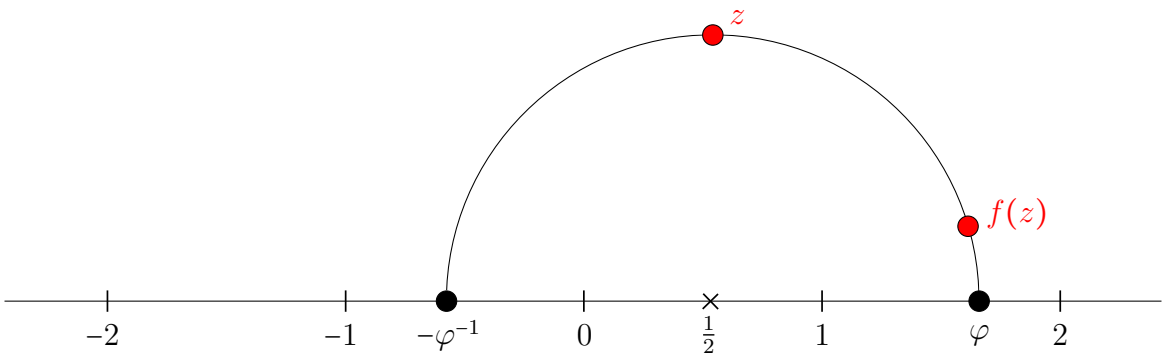


Figure 1.16: The geodesic in $\mathbb{H}$ on which f acts by translation.

Note that $z := \frac{1}{2}(1+\sqrt{5}\,i)$ is such a point. We compute:

$$f(z) = \frac{1+\sqrt{5}\,i+1}{\frac{1}{2}(1+\sqrt{5}\,i)+1} = \frac{1}{7}(11+\sqrt{5}\,i),$$

and finally

$$\frac{1}{2}\, d_{\mathbb{H}}(z, f(z)) = \log(x+\sqrt{x^2+1}),$$

with

$$x = \frac{\|z - f(z)\|}{2\sqrt{\mathrm{Im(z)}\,\mathrm{Im(f(z))}}},$$

where $\|\cdot\|$ stands for the Euclidean norm. Doing the calculations, one obtains that the

Teichmüller (stable) translation length of f is given by

$$l_{\mathcal{T}}(f) = \mathcal{L}_{\mathcal{T}}(f) = \log(\frac{1}{2}(3+\sqrt{5})) = \log(\lambda).$$

This is in fact no coincidence, as it turns out that for any genus g, and any pseudo-Anosov $f \in \mathrm{Mod}(S_g)$, the translation length $l_{\mathcal{T}}(f) = \mathcal{L}_{\mathcal{T}}(f)$ is always given by the logarithm of the stretch factor of f.

We now turn our attention to the curve graph $\mathcal{C}(T)$. Since on a torus there are no two different, disjoint curves, we need to modify the definition of the curve graph in this case. The vertices of $\mathcal{C}(T)$ still correspond to curves on T, but two vertices are joined by an edge if and only if there are representatives that intersect exactly once.

Thinking of the torus once more as a square with opposite sides identified, any curve on T can be represented by a straight line of rational slope. More formally, we can identify the torus with $\mathbb{R}^2/\mathbb{Z}^2$, pick a primitive vector

$$v = \begin{pmatrix} p \\ q \end{pmatrix} \in \mathbb{Z}^2,$$

and draw a straight line in $\mathbb{R}^2$ from the origin to v. This line descends to a closed curve of slope $\frac{p}{q}$ on the torus, and every curve of the torus arises in this way.

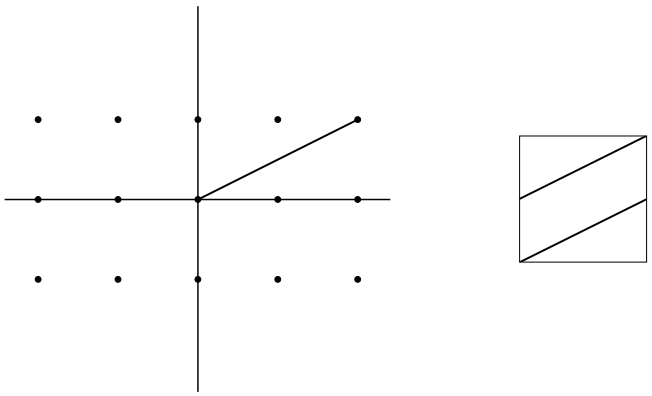


Figure 1.17: A straight line to $v = \begin{pmatrix} 2 \\ 1 \end{pmatrix}$ and the corresponding curve of slope 2 on T.

Thus, the vertices of the curve graph correspond to reduced fractions $\frac{p}{q} \in \mathbb{Q}$ together with the fraction $\frac{1}{0}$ describing the horizontal straight line on T, or equivalently to primitive vectors in $\mathbb{Z}^2$. The number of intersections of two curves of slopes $\frac{p}{q}$ and $\frac{r}{s}$ is given by the

absolute value of the determinant of the matrix

$$\begin{pmatrix} p & r \\ q & s \end{pmatrix}.$$

Hence, two vertices of the curve graph are joined by an edge if and only if $|ps - rq| = 1$, and we see that $\mathcal{C}(T)$ is exactly the *Farey graph*.

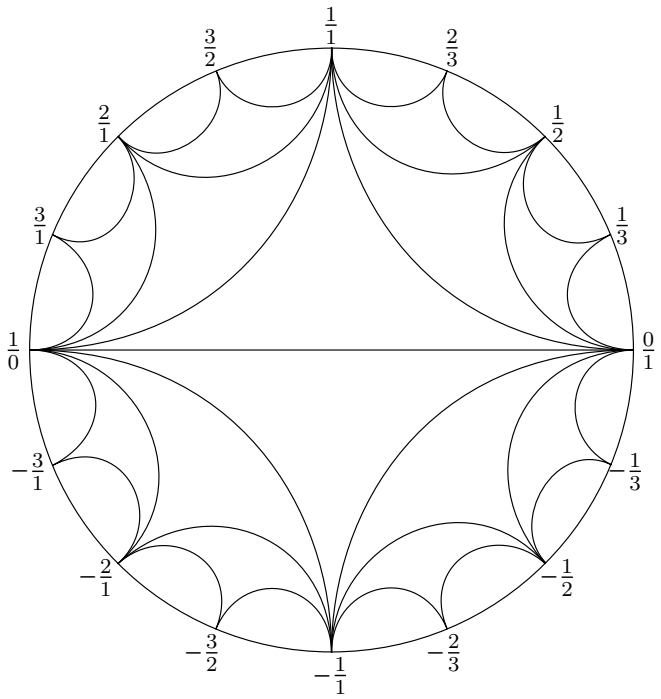


Figure 1.18: A portion of the Farey graph.

Since the mapping class f acts linearly on the square as the matrix

$$\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix},$$

a curve $\frac{p}{q}$ is mapped by f to the curve $\frac{2p+q}{p+q}$.

As we expect from an Anosov mapping class, we can see that f doesn't fix any curve, and hence obtain $\mathcal{L}_{\mathcal{C}}(f) > 0$. Since f sends the curve $\frac{1}{0}$ to the curve $\frac{2}{1}$, and

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

i.e. $d_{\mathcal{C}}(\frac{1}{0}, \frac{2}{1}) = 1$, we obtain

$$\mathcal{L}_{\mathcal{C}}(f) = 1.$$

For the stable curve graph translation length, it follows that $l_{\mathcal{C}}(f) \leqslant 1$. In order to compute $l_{\mathcal{C}}(f)$ precisely, we need to iterate f on some curve. Picking again $\frac{1}{0}$, we get

$$\frac{1}{0} \xrightarrow{f} \frac{2}{1} \xrightarrow{f} \frac{5}{3} \xrightarrow{f} \frac{13}{8} \xrightarrow{f} \cdots,$$

i.e. f^n sends the curve $\frac{1}{0}$ to the curve $\frac{F_{2n+1}}{F_{2n}}$, where F_i is the i^{th} *Fibonacci number*.

Using the definition of the Farey graph, one can check that the sequence of these fractions defines a geodesic in $\mathcal{C}(T)$, and so f preserves this geodesic and acts on it by translation by 1. In particular, we conclude that

$$l_{\mathcal{C}}(f) = \liminf_{n\to\infty} \frac{d_{\mathcal{C}}(f^n(\frac{1}{0}), \frac{1}{0})}{n} = \liminf_{n\to\infty} \frac{n}{n} = 1 = \mathcal{L}_{\mathcal{C}}(f).$$

In total, we've computed the following translation lengths of the Anosov f :

$$l_{\mathcal{T}}(f) = \mathcal{L}_{\mathcal{T}}(f) = \log(\lambda), \text{ and } l_{\mathcal{C}}(f) = \mathcal{L}_{\mathcal{C}}(f) = 1.$$

We saw that in both $\mathcal{T}(T)$ and $\mathcal{C}(T)$, the mapping class f preserves a geodesic and acts on it by translation. While for Teichmüller space this is always the case for a pseudo-Anosov, it is a special feature of our particular f in the case of the curve graph. In our example, the two spaces can be identified with the hyperbolic plane $\mathbb{H}$ and the Farey graph, respectively. Since the Farey graph embeds naturally into the hyperbolic plane where the vertices are aligned along the boundary $\partial\mathbb{H}$ as in Figure 1.18, we can think of the curve graph as a discrete model of Teichmüller space.

Iterating f on some point of the geodesic f preserves, defines a sequence which converges to the golden ratio $\varphi \in \partial\mathbb{H}$. For Teichmüller space, we see this from the description of the Möbius transformation corresponding to f, and for the curve graph by the fact that the sequence of ratios of Fibonacci numbers converges to φ.

Chapter 2

Minimal hyperelliptic translation lengths

In this chapter, we prove results on the translation lengths of pseudo-Anosovs that are affine homeomorphisms with respect to hyperelliptic translation surfaces. In particular, we compute the asymptotic behaviour of the minimal curve graph translation length of such pseudo-Anosovs. We start by recalling the relevant notions defined in Chapter 1 and present the main results.

Let S_g be a closed, oriented surface of genus $g \geq 2$. The *mapping class group* $\mathrm{Mod}(S_g)$ is the group of orientation-preserving homeomorphisms up to isotopy. $\mathrm{Mod}(S_g)$ admits natural actions on the *Teichmüller space* $\mathcal{T}(S_g)$ and on the *curve graph* $\mathcal{C}(S_g)$. The former is the space of marked hyperbolic structures on S_g, while the latter is a graph whose vertices correspond to isotopy classes of essential, simple, closed curves and its edges to disjointness of these curves. $\mathcal{T}(S_g)$ admits a metric called the *Teichmüller metric*, which resembles the hyperbolic metric in many properties. For example, the isometries of $\mathcal{T}(S_g)$ can be split into three categories, just like isometries of hyperbolic space are classified as being elliptic, parabolic or hyperbolic. The action of the mapping class group on Teichmüller space was used by Thurston to categorise mapping classes. The Nielsen-Thurston classification states that any $f \in \mathrm{Mod}(S_g)$ is either periodic, reducible or pseudo-Anosov.

Pseudo-Anosovs act on $\mathcal{T}(S_g)$ in a similar way to a hyperbolic isometry: for any pseudo-Anosov f, there exists a bi-infinite geodesic in $\mathcal{T}(S_g)$ such that f fixes this geodesic setwise and acts on it by translation. Denote the length of this translation by $l_{\mathcal{T}}(f)$. Using the definition of the Teichmüller metric, one can see that $l_{\mathcal{T}}(f) = \log \lambda(f)$, where $\lambda(f)$ is the stretch factor of the pseudo-Anosov. Hence, measuring the translation length on Teichmüller space is equivalent to determining the stretch factor of f. A big open question concerning the study of mapping classes is which numbers arise as a stretch

factor of a pseudo-Anosov. It is known that for fixed genus g, the set of all λ which are stretch factors for some pseudo-Anosov of S_g is a closed, discrete subset of $\mathbb{R}$. In particular, this set has a minimum, and it is an open problem to find the precise value of this minimum or the pseudo-Anosov attaining it. It is known, however, that the minimum has asymptotic behaviour like $\frac{1}{g}$. All of the above can be found in more detail in [FM12, Chapter 14].

The curve graph $\mathcal{C}(S_g)$ is a Gromov-hyperbolic graph, and the action of pseudo-Anosovs on the curve graph is by hyperbolic isometries [MM99, Proposition 4.6]. Here, a hyperbolic isometry is one for which the limit set of any orbit consists of exactly two points in the Gromov-boundary (as defined in [Gro87, page 209]). In fact, Bowditch showed that, similarly to the picture in Teichmüller space, for a pseudo-Anosov f there exists a bi-infinite geodesic in $\mathcal{C}(S)$ such that some power f^k preserves this geodesic and acts on it by translation [Bow08]. As before, we can measure the length of this translation, call it $l_{\mathcal{C}}(f^k)$ and set $l_C(f) \coloneqq \frac{l_{\mathcal{C}}(f^k)}{k}$ to be the stable translation length of f on the curve graph. In Section 2.1, we give an equivalent definition of $l_{\mathcal{C}}(f)$ which justifies the name *stable* curve graph translation length. From Bowditch's work, it follows that the $l_{\mathcal{C}}(f)$ are always rational. As is true for stretch factors, the set of all numbers arising as stable curve graph translation lengths of pseudo-Anosovs of a fixed genus g surface has a minimum. Again, it is an open problem to find the specific value of this minimum and the pseudo-Anosov attaining it. It is known that the minimum is of order $\frac{1}{g^2}$ [GT11].

Regarding the problem of finding the exact minima, one can relax the question by restricting to a subset of all pseudo-Anosovs and try to find the minimal stretch factor or stable curve graph translation length over this subset. Boissy–Lanneau restrict themselves to the set of pseudo-Anosovs that are affine with respect to a translation surface in a hyperelliptic stratum component. They proceed to determine the minimal stretch factor over this set and explicitly construct a pseudo-Anosov that attains this minimum [BL22]. Let f_g be this pseudo-Anosov for the respective genus g surface. In particular, Boissy–Lanneau show that $\lambda(f_g) \geqslant \sqrt{2}$ for any g.

Our first goal is to compute the stable curve graph translation length of the f_g. We show:

Theorem 2.1. *For all $g \geq 2$, it holds that*

$$\frac{1}{16g-12} \leqslant l_C(f_g) \leqslant \frac{1}{g-1}.$$

We proceed with showing that, in fact, the minimal stable curve graph translation length over all pseudo-Anosovs in a hyperelliptic component behaves like $l_{\mathcal{C}}(f_g)$. In the following,

we write $f(n) \asymp h(n)$ for two functions $f, h : \mathbb{N}_{\geq 2} \to \mathbb{R}$ if

$$\frac{f(n)}{h(n)} \in [\frac{1}{C}, C]$$

for some $C > 1$ and all n.

Theorem 2.2. *Let $l^{\mathcal{C}}_{g,\mathrm{hyp}}$ be the minimal stable curve graph translation length over all pseudo-Anosovs that are affine with respect to a genus g translation surface in a hyperelliptic component. Then:*

$$l^{\mathcal{C}}_{g,\mathrm{hyp}} \asymp \frac{1}{g}.$$

Remark 2.3. We remark that while the stable curve graph translation length of f_g is of the same order as $l^{\mathcal{C}}_{g,\mathrm{hyp}}$, it remains unclear whether f_g is the actual minimiser of the stable curve graph translation length over all pseudo-Anosovs in a hyperelliptic component, i.e. whether $l^{\mathcal{C}}_{g,\mathrm{hyp}} = l_{\mathcal{C}}(f_g)$.

Boissy–Lanneau's result can be interpreted as saying that pseudo-Anosovs in hyperelliptic components have a large stretch factor. Theorem 2.2 says that the stable curve graph translation lengths of such pseudo-Anosovs are also large. The minimal stable curve graph translation length over all pseudo-Anosovs is of order $\frac{1}{g^2}$ and can hence certainly not be attained by a pseudo-Anosov in a hyperelliptic component. However, Theorem 2.2 also shows that after restricting to the pseudo-Anosovs in hyperelliptic components, the minimal stable curve graph translation length behaves differently from the minimal stretch factor. While the minimal stretch factor stays bounded from below by the constant $\sqrt{2}$, the minimal stable curve graph translation length behaves like $\frac{1}{g}$ and, in particular, tends to 0 as we increase the genus.

This raises the question of how the two translation lengths $l_{\mathcal{T}}(f)$ and $l_{\mathcal{C}}(f)$ are related in general. We investigate this question further in Chapter 3.

Outline. In Section 2.1, we introduce the basic notions needed for this chapter. In particular, we recall the definitions of pseudo-Anosovs and their stable curve graph translation length and present some concepts of the theory of train tracks which will be useful for the proof of the lower bound in Theorem 2.1. In Section 2.2, we present Rauzy-Veech induction, which is used to define the maps f_g. However, since we are interested in $l_{\mathcal{C}}(f_g)$, which is a topological invariant, we omit the discussion of flat structures – which is normally used in order to define mapping classes through Rauzy-Veech induction – and present a purely topological version of it. In Section 2.3, we discuss some properties of the notions defined in Section 2.2 that will be useful for the proofs. Finally, in Section 2.4, we define the pseudo-Anosovs f_g and prove Theorems 2.1 and 2.2.

§ 2.1 | Relevant notions

§ 2.1.1 | Pseudo-Anosovs and translation lengths

Let S_g be a closed, oriented, genus g surface. If not specified otherwise, we always assume that $g \geq 2$. The mapping class group of S_g is the group of orientation-preserving self-homeomorphisms of S_g up to isotopy. We denote this group by $\mathrm{Mod}(S_g)$.

Definition 2.4. *A homeomorphism φ of S_g is called* pseudo-Anosov *if there are two measured foliations $(\mathcal{F}^s, \mu_s), (\mathcal{F}^u, \mu_u)$ and $\lambda > 1$ such that*

- *F^s and F^u are transverse,*
- $\varphi \cdot (\mathcal{F}^s, \mu_s) = (\mathcal{F}^s, \lambda^{-1}\mu_s)$,
- $\varphi \cdot (\mathcal{F}^u, \mu_u) = (\mathcal{F}^u, \lambda\mu_u)$.

For a pseudo-Anosov φ, we call $\mathcal{F}^s, \mathcal{F}^u$ the stable and unstable foliation, respectively, and λ the stretch factor (or the dilatation) of φ. A mapping class $f \in \mathrm{Mod}(S_g)$ is called pseudo-Anosov, if there exists a representative homeomorphism φ of f that is pseudo-Anosov. For a given pseudo-Anosov $f \in \mathrm{Mod}(S_g)$, there are different pseudo-Anosov homeomorphisms representing f; however, they are all conjugate to each other. Hence, the stretch factor of f is well-defined. We sometimes write $\lambda(f)$ instead of just λ to emphasise the pseudo-Anosov the stretch factor corresponds to. See [FM12, Part 3] for more information on pseudo-Anosov theory.

For a general metric space (X, d) and an isometry f of X, the *translation length* of f is defined as

$$\inf_{x \in X} d(x, f(x)).$$

The *stable translation length* is given by

$$\liminf_{n \to \infty} \frac{d(x, f^n(x))}{n},$$

where $x \in X$ is any point. By Lemma 1.23, the stable translation length is well defined.

Let $\mathcal{T}(S_g)$ denote the Teichmüller space of S_g. The mapping class group acts by isometries on $\mathcal{T}(S_g)$ with respect to the Teichmüller metric. For a pseudo-Anosov f, its Teichmüller space translation length and stable translation length are the same and given by $\log(\lambda(f))$ (see [FM12, Chapter 14] for details). We denote the (stable) Teichmüller space translation length of a pseudo-Anosov f by $l_{\mathcal{T}}(f)$.

Let $\mathcal{C}(S_g)$ denote the curve graph of S_g. This is the graph with vertices corresponding to

isotopy classes of essential simple closed curves, where two vertices are joined by an edge if there exist representative curves that are disjoint. Since homeomorphisms preserve the property of being an essential simple closed curve as well as the disjointness of two such curves, the mapping class group $\mathrm{Mod}(S_g)$ acts on the curve graph $\mathcal{C}(S_g)$ by graph automorphisms. We equip $\mathcal{C}(S_g)$ with the shortest path metric $d_{\mathcal{C}}$ where each edge has length 1. The mapping class group action is an action by isometries with respect to $d_{\mathcal{C}}$. A pseudo-Anosov f doesn't fix any finite set of curves on the surface, and therefore its curve graph translation length is always greater than or equal to 1. It is more interesting to consider the stable translation length.

For $f \in \mathrm{Mod}(S_g)$, the *stable curve graph translation length* of f is given by

$$l_{\mathcal{C}}(f) = \liminf_{n\to\infty} \frac{d_{\mathcal{C}}(\alpha, f^n(\alpha))}{n},$$

where $\alpha \in \mathcal{C}(S_g)$ is any vertex. From now on, when it comes to the curve graph, we will only consider the stable translation length.

§ 2.1.2 | Train tracks

In this section, we introduce the notion of train tracks and discuss their importance for pseudo-Anosov maps. As before, let S_g denote a closed, oriented surface of genus g. For the following definition, we equip S_g with a smooth structure.

Definition 2.5. *A* train track τ *on* S_g *is an embedded graph such that each edge is a smooth path and at each vertex all adjacent edges are mutually tangent.*

Vertices of τ are usually referred to as *switches* and edges of τ as *branches*. The tangency condition yields a splitting of the set of branches adjacent to a switch into two sets, the incoming and outgoing branches. Figure 2.1 shows part of a train track; the switch on the left has two incoming and one outgoing branch, while the switch on the right has one incoming and three outgoing branches. We note that the choice of which branches are called incoming and which are called outgoing is irrelevant and could be reversed. In what follows, we mention some important notions we make use of. For a more detailed introduction to train tracks, see [PH92].

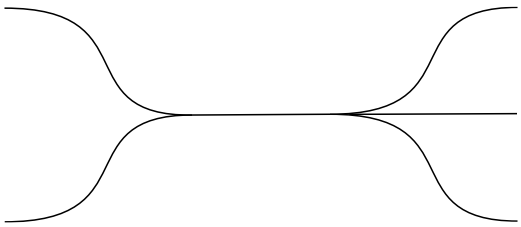


Figure 2.1: Part of a train track.

We say that a train track is *large* if all components of $S_g \setminus \tau$ are polygons.

A *train route* is a smooth path in τ. Hence, a train route passing through a switch can only pass from an incoming to an outgoing branch or the other way around. A curve γ on S is *carried* by τ if it is isotopic to a closed train route. A train track is *recurrent* if every branch is contained in a closed train route. A train track is *transversely recurrent* if for every branch there exists a simple closed curve intersecting the branch efficiently, i.e. there are no bigons between the branch and the curve. Finally, a train track is *birecurrent* if it is both recurrent and transversely recurrent.

A *measure* μ on τ is an assignment of a non-negative real number (also called a *weight*) to each branch of τ such that at each switch the sum of weights of the incoming branches equals the sum of weights of the outgoing ones. Denote by $P(\tau)$ the set of measures on τ. Note that a curve carried by τ induces a counting measure on τ, assigning to each branch the number of times the curve passes over this branch. Hence, we can think of such a curve as an element of $P(\tau)$. Let $\mathrm{int}(P(\tau))$ denote the set of measures that are positive on each branch.

A *diagonal extension* of τ is a train track τ' which consists of the same switches and branches as τ and has possibly some extra branches, called *diagonals*, which start and terminate in cusps of some complementary region of $S_g \setminus \tau$ and are not homotopic to any of the already existing branches. Figure 2.2 shows an example of a diagonal. Let $E(\tau)$ be the set of diagonal extensions of τ and $PE(\tau)$ the union of the sets of measures of all diagonal extensions. Furthermore, let $\mathrm{int}(PE(\tau))$ be the set of measures in $PE(\tau)$ that are positive on each branch of τ.

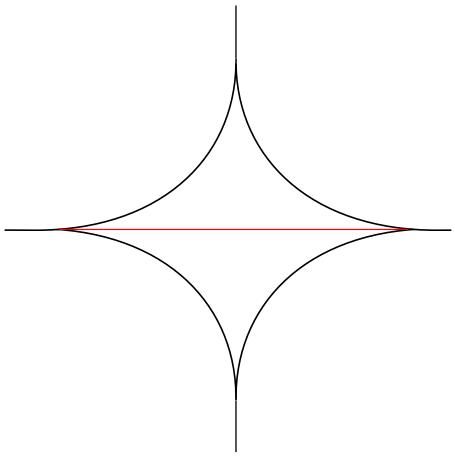


Figure 2.2: An example of a diagonal in a train track.

The next lemma, often referred to as the *Nesting Lemma*, is crucial in our computation of the lower bound in Theorem 2.1. In the following, one should think of the sets $PE(\tau)$ and

$\mathrm{int}(PE(\tau))$ as subsets of the curve graph $\mathcal{C}(S_g)$, i.e. only consider the measures coming from curves, and $\mathcal{N}_1$ denotes the 1-neighbourhood of a subset of $\mathcal{C}(S_g)$.

Lemma 2.6 (Nesting lemma, [MM99])**.** *Let τ be a large recurrent train track. Then*

$$\mathcal{N}_1(\mathrm{int}(PE(\tau))) \subset PE(\tau).$$

Proof. The Nesting Lemma is due to Masur–Minsky [MM99, Lemma 4.4]. They prove it under the stronger assumption that the train track is birecurrent. The weakening of this to require that the train track be recurrent only is established in [GT11, Lemma 3.2]. □

A train track σ is *carried* by the train track τ if there is an isotopy of S_g that takes every train route of σ to a train route of τ. We write $\sigma \prec \tau$. A train track is *invariant* for a homeomorphism f if $f(\tau) \prec \tau$. In this case, we can define a matrix V_f as follows.

Let n be the number of branches of τ. Let $b_1, ..., b_n$ be a numbering of the branches and V_f be the $n \times n$ matrix whose (i, j)-entry counts how often the branch $f(b_j)$ passes over the branch b_i. Hence, the j^{th} column of V_f describes the train route $f(b_j)$. We call V_f a *train track matrix* for f. Note that V_f depends on the invariant train track τ as well.

Train tracks play an important role in the theory of pseudo-Anosovs, because for every pseudo-Anosov, there exists an invariant train track. The corresponding train track matrix is a Perron-Frobenius matrix, and its leading eigenvalue is the stretch factor of the pseudo-Anosov [FM12, Chapter 15.2]. We will use an invariant train track and its corresponding train track matrix to compute the lower bound in Theorem 2.1.

§ 2.1.3 | Translation surfaces and affine pseudo-Anosovs

We recall the necessary definitions of the terms that come up in Theorem 2.2.

A *translation surface* is a pair (X, ω) where X is a Riemann surface and ω a non-zero holomorphic 1-form on X.

The name translation surface is justified by the fact that a translation surface admits a natural atlas of charts away from the zeros of ω whose transition functions are translations (see, for example, [Wri15]). The zeros of ω are usually referred to as *singularities*.

A homeomorphism f of a translation surface (X, ω) is called *affine* if it sends singularities to singularities and is affine on the natural charts. Note that this is well-defined since the chart transitions are translations, so in particular, affine maps.

A Riemann surface X is called *hyperelliptic* if it admits an involution, i.e. an order 2 biholomorphism $\iota : X \to X$, such that $X/\langle\iota\rangle \cong \mathbb{P}^1$, where $\mathbb{P}^1$ denotes the Riemann sphere.

A translation surface (X, ω) is *hyperelliptic* if the underlying Riemann surface X is hyperelliptic.

Given a translation surface (X, ω), let $x_1, ..., x_n$ be the zeros of ω and let $k_1, ..., k_n$ be the orders of the x_i respectively. The Riemann-Roch Theorem ensures that $k_1 + ... + k_n = 2g-2$, where g is the genus of X. Thus, for any given tuple $(k_1, ..., k_n)$ with $k_1 + ... + k_n = 2g-2$, one can define the space $\mathcal{H}(k_1, ..., k_n)$ to be the space of all translation surfaces with exactly n singularities of orders $k_1, ..., k_n$. These spaces are referred to as *strata* of translation surfaces. Kontsevich–Zorich classify the connected components of these strata [KZ03]. In particular, the classification shows that the two strata $\mathcal{H}(2g-2)$ and $\mathcal{H}(g-1, g-1)$ contain a connected component that consists entirely of hyperelliptic translation surfaces. These connected components are referred to as *hyperelliptic components*.

If S_g is a topological closed, oriented, genus g surface, then we say that $f \in \mathrm{Mod}(S_g)$ is affine for a translation surface in a hyperelliptic component (or simply f is in a hyperelliptic component) if there exists a representative φ of f, a translation surface (X, ω) in a hyperelliptic component, a homeomorphism $\psi : S_g \to X$ and an affine homeomorphism φ_X of (X, ω) such that the following diagram commutes:

$$\begin{array}{ccc} S_g & \xrightarrow{\varphi} & S_g \\ \downarrow{\scriptstyle \psi} & & \downarrow{\scriptstyle \psi} \\ X & \xrightarrow{\varphi_X} & X. \end{array}$$

§ 2.2 | Topological Rauzy-Veech induction

In this section, we present the method of using Rauzy-Veech induction to obtain mapping classes. This method relies on the theory of translation surfaces. Since we are interested in the stable curve graph translation length, which is a topological invariant of a pseudo-Anosov, we present a purely topological version of Rauzy-Veech induction that omits the discussion of translation surfaces.

§ 2.2.1 | Rauzy-Veech induction on permutations

An *interval exchange map* is a self-map of the interval $I = [0, 1]$, for which there is a subdivision of I into subintervals such that the map permutes these. Consequently, such a map is determined by a permutation $\pi \in S_n$, describing how to permute the subintervals, and a "length vector" $\lambda \in \mathbb{R}^n_+$ with $\lambda_1 + ... + \lambda_n = 1$, defining the length of each subinterval.

Given an interval exchange map, Rauzy-Veech induction is a method to produce a new interval exchange map out of the given one. An interval exchange map together with some extra data gives rise to a translation surface. Therefore, the study of translation

surfaces is closely related to that of interval exchange maps, and Rauzy-Veech induction can be used to define homeomorphisms that are affine with respect to a certain translation surface. We refer the reader to [Yoc10] for an introduction to this theory. Boissy–Lanneau use Rauzy-Veech induction to define pseudo-Anosov maps that are affine with respect to some translation surface in the hyperelliptic stratum component of $\mathcal{H}(2g-2)$ and to find the pseudo-Anosov with minimal stretch factor out of all of these [BL22]. For a given g, we denote the pseudo-Anosov that attains this minimum by f_g. Boissy–Lanneau proceed to do the same with the hyperelliptic stratum component of $\mathcal{H}(g-1,g-1)$. However, the stretch factor of f_g is less than the minimal one for $\mathcal{H}(g-1,g-1)$, and hence it is the minimal one over all hyperelliptic components. This is why we focus on f_g. We point out that the techniques presented in the computation of $l_{\mathcal{C}}(f_g)$ below can be used to compute the stable curve graph translation length of any pseudo-Anosov defined in terms of Rauzy-Veech induction.

In this work, our goal is to compute the stable curve graph translation length of the f_g. In order to define the pseudo-Anosovs f_g (as in [BL22, Section 5]), we use Rauzy-Veech induction. Since we are interested in $l_{\mathcal{C}}(f_g)$, which is a purely topological invariant, we omit the discussion of translation surfaces and present a more topological version of Rauzy-Veech induction. In particular, instead of an interval exchange map, i.e., a permutation plus a length vector, we only require a permutation.

Let $\pi \in S_n$ be a permutation. We say that π is *irreducible*, if $\pi(\{1,...,k\}) = \{1,...,k\}$ implies $k = n$. Throughout, we assume that all permutations are irreducible. Note that this ensures that $\pi(n) \neq n$, which is necessary in order to define the Rauzy-Veech induction below.

We want to work with so-called labelled permutations. Let $\mathcal{A}$ be a finite alphabet with n elements and $\pi_t : \mathcal{A} \to \{1,...,n\}$ and $\pi_b : \mathcal{A} \to \{1,...,n\}$ bijections, such that $\pi = \pi_b \circ \pi_t^{-1}$. We call the pair (π_t, π_b) a *labelled permutation* and say that this labelled permutation is irreducible if the corresponding π is irreducible. Here, the subscripts t and b stand for "top" and "bottom", respectively. Note that π does not determine π_t and π_b uniquely. In other words, after fixing an alphabet $\mathcal{A}$, we have a surjective, but not injective map from the set of labelled permutations to S_n.

As an example, consider $\pi = (1\ 2\ 3) \in S_3$ and $\mathcal{A} = \{A, B, C\}$. One choice for π_t, π_b would be $\pi_t(A) = 1$, $\pi_t(B) = 3$, $\pi_t(C) = 2$ and $\pi_b(A) = 2$, $\pi_b(B) = 1$, $\pi_b(C) = 3$. We usually represent our labelled permutation as a matrix

$$\begin{pmatrix} \pi_t^{-1}(1) & ... & \pi_t^{-1}(n) \\ \pi_b^{-1}(1) & ... & \pi_b^{-1}(n) \end{pmatrix},$$

which in our example gives

$$\begin{pmatrix} A & C & B \\ B & A & C \end{pmatrix}.$$

Given a labelled permutation (π_t, π_b), we construct a surface $X(\pi_t, \pi_b)$ by taking a regular $2n$-gon, labelling its top sides according to π_t and its bottom sides according to π_b and gluing the sides with the same label in the way they are oriented in Figure 2.3. Here, the gluings are just topological and not necessarily by translations. In order to distinguish between the polygon and the resulting surface after side identifications, we denote the polygon by $P(\pi_t, \pi_b)$. Figure 2.3 shows an example for the labelled permutation

$$\begin{pmatrix} A & B & C & D \\ D & A & C & B \end{pmatrix}.$$

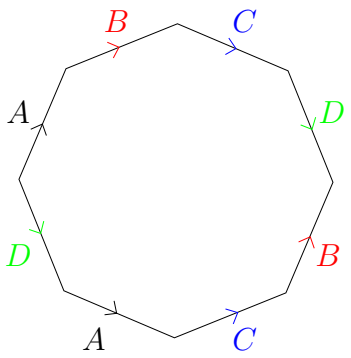


Figure 2.3: An example of the $2n$-gon construction for $n = 4$.

We now define Rauzy-Veech induction. Let (π_t, π_b) be a labelled permutation with $\pi = \pi_b \circ \pi_t^{-1}$. Rauzy-Veech induction consists of two purely combinatorial ways to obtain a new labelled permutation from the given one. The two ways are called top, resp. bottom Rauzy-Veech induction, which we denote by $\mathcal{R}_t$ and $\mathcal{R}_b$ respectively.

- Top Rauzy-Veech induction $\mathcal{R}_t$.
  Let $\alpha := \pi_t^{-1}(n)$ and $\beta := \pi_b^{-1}(n)$. In this case, we say that the top is the winner and the bottom is the loser, or equivalently, α is the winner, whereas β is the loser. We refer to the Rauzy-Veech induction in this case as being of type t.
  We define the Rauzy-Veech induction $\mathcal{R}_t(\pi_t, \pi_b)$ as (π_t', π_b'), where $\pi_t' = \pi_t$ and π_b' differs from π_b by moving β to the right of α and translating everything afterwards to the right by one. For example

$$\begin{pmatrix} A & B & C & D \\ D & C & B & A \end{pmatrix} \xrightarrow{\mathcal{R}_t} \begin{pmatrix} A & B & C & D \\ D & A & C & B \end{pmatrix}.$$

- Bottom Rauzy-Veech induction $\mathcal{R}_b$.
  Let $\alpha := \pi_t^{-1}(n)$ and $\beta := \pi_b^{-1}(n)$. In this case, we say that the bottom is the winner and the top the loser, or equivalently, β is the winner, whereas α is the loser. We refer to the Rauzy-Veech induction in this case as being of type b.
  We define the Rauzy-Veech induction $\mathcal{R}_b(\pi_t, \pi_b)$ as (π_t', π_b'), where $\pi_b' = \pi_b$ and π_t' differs from π_t by moving α to the right of β and translating everything afterwards to the right by one. For example

$$\begin{pmatrix} A & B & C & D \\ D & C & B & A \end{pmatrix} \xrightarrow{\mathcal{R}_b} \begin{pmatrix} A & D & B & C \\ D & C & B & A \end{pmatrix}.$$

We define a graph called the *labelled Rauzy diagram* $\mathcal{D}_n$. The vertices of this graph correspond to irreducible labelled permutations of fixed length n, and we connect two vertices $(\pi_t, \pi_b), (\pi_t', \pi_b')$ by a with t resp. b labelled, oriented edge (from (π_t, π_b) to (π_t', π_b')) if $\mathcal{R}_t(\pi_t, \pi_b) = (\pi_t', \pi_b')$ resp. $\mathcal{R}_b(\pi_t, \pi_b) = (\pi_t', \pi_b')$.

The labelled Rauzy diagram is, in general, not a connected graph, but we always restrict ourselves to a connected component. We will be interested in the connected component of the labelled permutation

$$\begin{pmatrix} \alpha_1 & \alpha_2 & \dots & \alpha_n \\ \alpha_n & \alpha_{n-1} & \dots & \alpha_1 \end{pmatrix}.$$

For example, if $n = 3$, the connected component of

$$\begin{pmatrix} A & B & C \\ C & B & A \end{pmatrix}$$

is

$$_{\mathrm{t}}\circlearrowleft \begin{pmatrix} A & C & B \\ C & B & A \end{pmatrix} \underset{b}{\overset{b}{\rightleftarrows}} \begin{pmatrix} A & B & C \\ C & B & A \end{pmatrix} \underset{t}{\overset{t}{\rightleftarrows}} \begin{pmatrix} A & B & C \\ C & A & B \end{pmatrix} \circlearrowright_{\mathrm{b}}.$$

Finally, we want to assign a matrix to every edge of the labelled Rauzy diagram. Let e be an edge of $\mathcal{D}_n$. Let (π_t, π_b) be the initial vertex of e. Since e corresponds to a Rauzy-Veech move (either t or b), it determines a winner-loser pair (α, β). Let $E_{\alpha,\beta}$ be the matrix with a single non-zero entry equal to 1 in the (α, β) position and let $V_{\alpha,\beta} = Id + E_{\alpha,\beta}$. We assign to e the matrix $V_{\alpha,\beta}$.

§ 2.2.2 | Construction of mapping classes through Rauzy-Veech induction

Let (π_t, π_b) be a labelled permutation, $P(\pi_t, \pi_b)$ the corresponding polygon and $X(\pi_t, \pi_b)$ the corresponding surface. We want to interpret Rauzy-Veech induction on (π_t, π_b) in a geometric way, i.e. relate $X(\pi_t, \pi_b)$ and $X(\mathcal{R}(\pi_t, \pi_b))$. Consider the following cut-and-

paste process on $P(\pi_t, \pi_b)$.

Let T be the triangle consisting of the top side corresponding to $\alpha := \pi_t^{-1}(d)$, the bottom side corresponding to $\beta := \pi_b^{-1}(d)$ and the straight line joining their two distinct endpoints, see Figure 2.4. Cut the triangle T and glue it back by either gluing the top side α to the bottom side corresponding to α or gluing the bottom side β to the top side corresponding to β. After having glued back T, we choose an ambient isotopy of the plane to isotope the resulting polygon back to being a regular $2n$-gon, see Figure 2.5.

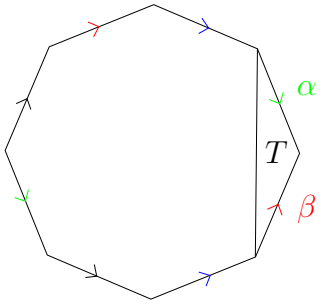


Figure 2.4: The triangle T.

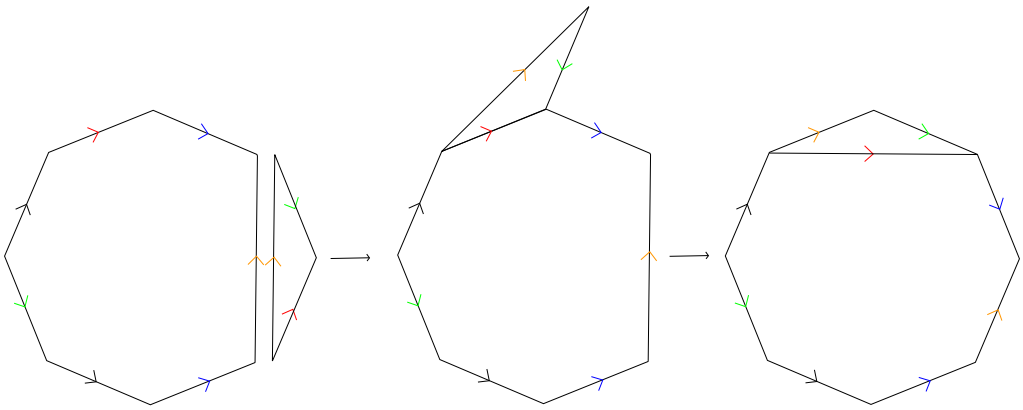


Figure 2.5: The cut-and-paste construction for a bottom move.

We refer to the case where we glue the α's as a top move and to the case where we glue the β's as a bottom move. In both cases, the new polygon has two new unlabelled sides, which we label by α in the case of a top move and by β in the case of a bottom move.

In total, we obtain that performing a top or bottom move to $P(\pi_t, \pi_b)$ as described above yields a new $2n$-gon together with a labelled permutation that describes the gluing of the sides, where this new labelled permutation is exactly obtained by Rauzy-Veech induction of type t, respectively b on (π_t, π_b). After gluing sides according to the new permutation, we obtain the surface $X(\mathcal{R}(\pi_t, \pi_b))$.

In summary, we now have a way to assign a surface to a given labelled permutation

and interpret a Rauzy-Veech move (equivalently, an outgoing edge in the labelled Rauzy diagram) as a cut-and-paste move of a triangle. The surface obtained after the cut-and-paste move corresponds to the resulting labelled permutation. We now want to use this interpretation to define mapping classes.

Consider an edge path γ in the labelled Rauzy diagram, so that its endpoints define the same unlabelled permutation. We call such a path an *allowed path*. In other words, an allowed path γ is a finite sequence of t's and b's such that its starting point (π_t, π_b) and endpoint (π'_t, π'_b) satisfy $\pi_b \circ \pi_t^{-1} = \pi'_b \circ \pi_t'^{-1}$. Since the unlabelled permutations are the same, the associated surfaces $X(\pi_t, \pi_b)$ and $X(\pi'_t, \pi'_b)$ admit a natural homeomorphism $\varphi : X(\pi_t, \pi_b) \to X(\pi'_t, \pi'_b)$, which just comes by identifying the $2n$-gons $P(\pi_t, \pi_b)$ and $P(\pi'_t, \pi'_b)$ via the identity map. We refer to φ as the *change of labelling*. There is a second homeomorphism from $X(\pi_t, \pi_b)$ to $X(\pi'_t, \pi'_b)$ that is given by tracking the moves of γ, i.e. sending each point to the corresponding point after each cut-and-paste step. We denote this homeomorphism by ψ. Note that at each cut-and-paste step we have chosen an ambient isotopy of the plane to obtain a regular $2n$-gon (cf. right arrow in Figure 2.5). Any other choice of isotopies would result in an isotopic homeomorphism ψ'. For us, it is enough to define ψ up to isotopy, since we only need a well-defined mapping class. We refer to ψ as the *cut-and-paste homeomorphism*.

For an allowed path γ, we let $f_\gamma := \psi^{-1} \circ \varphi \in \mathrm{Mod}(X(\pi_t, \pi_b))$, where (π_t, π_b) is the starting point of γ and φ, ψ are as above.

To f_γ we associate a matrix as follows: Let $e_1, \ldots, e_k$ be the edges of γ ordered from first to last, i.e. the starting point of e_1 is (π_t, π_b) and the endpoint of e_k is (π'_t, π'_b). Let V_{e_i} be the corresponding matrices as discussed at the end of Section 2.2.1. Finally, let P be the permutation matrix that is 1 in the (α, β) entry if and only if $\varphi^{-1}(\alpha) = \beta$, or equivalently $\varphi(\beta) = \alpha$, and 0 otherwise. Here, $\varphi(\beta) = \alpha$ is to be understood as φ maps the side labelled β in the $2n$-gon to the side labelled α in the other $2n$-gon. We let $V_\gamma := V_{e_1} \cdot \ldots \cdot V_{e_k} \cdot P$.

Note that V_γ is a matrix with only non-negative entries. We say that V_γ is *primitive* if there exists a $k \in \mathbb{N}$ such that V_γ^k is positive, i.e. consists of strictly positive entries. In [Vee82], it is shown that whenever the matrix V_γ is primitive, the corresponding mapping class f_γ is a pseudo-Anosov. Furthermore, the spectral radius of V_γ is the stretch factor of f_γ.

§ 2.3 | Some properties of f_γ and V_γ

Let γ be an allowed path in the labelled Rauzy diagram. In this section, we want to understand how the matrix V_γ describes the map f_γ. Our first goal is to show that there

is an invariant train track τ_γ for f_γ such that V_γ is the corresponding train track matrix.

Let (π_t, π_b) be the starting point of γ and let $P_\gamma := P(\pi_t, \pi_b)$, $X_\gamma := X(\pi_t, \pi_b)$ be the corresponding polygon and surface. Let τ_γ be the train track on X_γ which consists of a single switch in the middle of P_γ and n branches, each of which goes from the switch to a bottom side of P_γ, comes back from the corresponding top side of P_γ and connects back to the switch. For the labelled permutation

$$\begin{pmatrix} A & B & C & D \\ D & A & C & B \end{pmatrix},$$

an example can be seen in Figure 2.6.

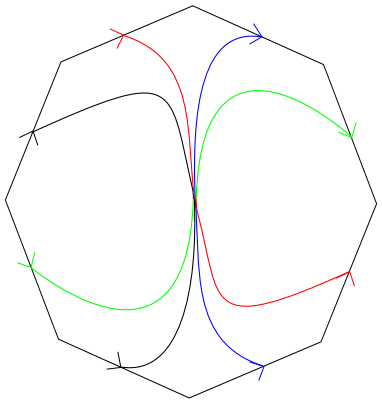


Figure 2.6: The train track τ_γ.

Label each branch of τ_γ with the same alphabet used to define the labelled permutations π_t, π_b, i.e. call a branch α if it connects the sides of P_γ labelled α. We show:

Lemma 2.7. *τ_γ is an invariant train track for f_γ and the corresponding train track matrix is V_γ.*

Proof. Let (π_t', π_b') be the endpoint of the path γ. Recall that with the notation of Section 2.2.2, $f_\gamma = \psi^{-1} \circ \varphi$, where $\varphi, \psi : X_\gamma \to X(\pi_t', \pi_b')$ are the change of labelling and cut-and-paste homeomorphism, respectively.

Hence, in order to determine $f_\gamma(\tau_\gamma)$, we need to first apply φ, and then perform all the cut-and-paste moves backwards. The map φ maps τ_γ to a train track on $X(\pi_t', \pi_b')$ which when drawn on $P(\pi_t', \pi_b')$ looks exactly the same as τ_γ. However, the labels of the branches may have changed. This is exactly captured by the permutation matrix P defined in Section 2.2.2. Any branch α of τ_γ is mapped under φ to the branch $\varphi(\alpha) =: \beta$ and the α-column of P has exactly one non-zero entry equal to 1 in the (β, α) position.

For the cut-and-paste moves, we analyse a single b-move (a t-move is analogous) and then argue by induction. So, assume we apply the inverse of a b-move with a winner-loser pair (α, β). This corresponds to cutting the triangle consisting of the top sides α and β and a straight line connecting them and gluing the α-side of it to the bottom side labelled α, (compare Figure 2.7, where $(\alpha, \beta) = (B, C)$. In the middle polygon of the figure, we labelled the "new" B-edge). The effect on our train track is now easy to observe. The α-branch will still be the α-branch in the new polygon, while the β-branch can be isotoped to run exactly once over itself and once over the α-branch. All other branches will remain the same. This is exactly captured by the matrix $V_{\alpha,\beta}$. For any $\alpha' \neq \beta$, the α'-column consists of a single non-zero entry equal to 1 in the (α', α') position. The β-column has two non-zero entries equal to 1 in the positions (β, β) and (α, β).

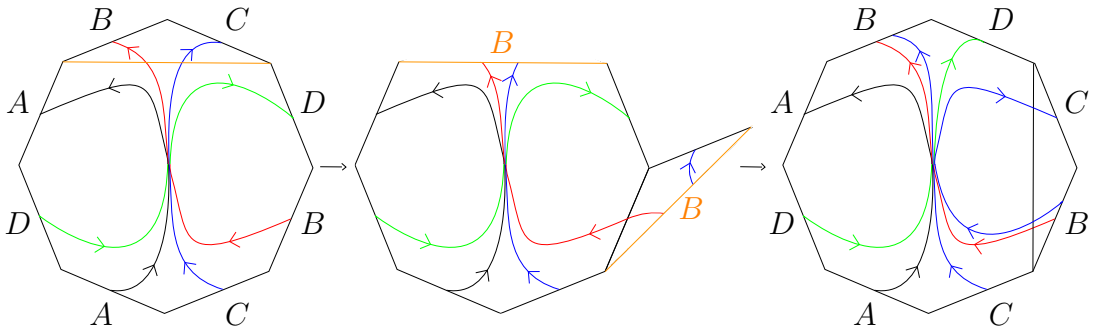


Figure 2.7: The inverse of a b-move with its effect on the train track.

Hence, we conclude inductively that after applying φ and each cut-and-paste move reversed, our newly obtained train track $f_\gamma(\tau_\gamma)$ is carried by τ_γ. Furthermore, if $e_1, ..., e_k$ are the edges of γ, then the α-column of the matrix

$$V_{e_1} \cdot ... \cdot V_{e_k} \cdot P = V_\gamma$$

describes the image of the branch $f_\gamma(\alpha)$ for any α, i.e. V_γ is the train track matrix. □

The above lemma will be used in the proof of the lower bound in Theorem 2.1. For the proof of the upper bound, it will be more convenient to observe what happens to the sides of the polygon P_γ. Note that these sides do not necessarily define closed curves in X_γ. However, we will see that for the pseudo-Anosovs we will be interested in, the sides do define closed curves. In general, the sides only define paths on X_γ, but a concatenation of these paths defines a closed curve, and this in turn can be used to apply a similar method as the one presented in Section 2.4 to compute upper bounds of the stable curve graph translation length in other examples.

Since $l_{\mathcal{C}}(f_\gamma) = l_{\mathcal{C}}(f_\gamma^{-1})$, it does not matter if we use f_γ or f_γ^{-1} to compute the translation

length. In fact, we want to study the action of f_γ^{-1} on the sides of P_γ instead of the one of f_γ. We make this choice because $f_\gamma^{-1} = \varphi^{-1} \circ \psi$ (compare the notation of Section 2.2.2), and it is slightly easier to apply the cut-and-paste homeomorphism ψ instead of applying its inverse. The reason we didn't consider f_γ^{-1} to begin with is that the train track τ_γ is not carried by f_γ^{-1} but only by f_γ.

Let α be a side of the polygon P_γ. We think of α as a side of P_γ as well as a path in X_γ interchangeably as long as there is no ambiguity.

Lemma 2.8. *The path $f_\gamma^{-1}(\alpha)$ can be homotoped so that it lies only on the sides of the polygon P_γ. Furthermore, the (α, β) entry of V_γ counts exactly how often $f_\gamma^{-1}(\alpha)$ passes over β.*

Proof. Since $f_\gamma^{-1} = \varphi^{-1} \circ \psi$, we need to first apply the cut-and-paste moves to P_γ and then the inverse of the change of labelling.

Consider a single edge e of γ corresponding to a bottom cut-and-paste move with winner-loser pair (α, β). The case of a top move is analogous.

Any side which does not correspond to the winner side remains a side of the new polygon with the same label (compare Figure 2.5). The side α gets mapped to a line connecting two corners of the new polygon, which can be homotoped to pass exactly once over the new side labelled α and the side labelled β (compare Figure 2.8). This is captured by the rows of the matrix $V_{\alpha,\beta}$. For any $\beta' \neq \alpha$, the β'-row of $V_{\alpha,\beta}$ has a single non-zero entry equal to 1 in the (β', β') position, while the α-row has two non-zero entries equal to 1 in positions (α, α) and (α, β).

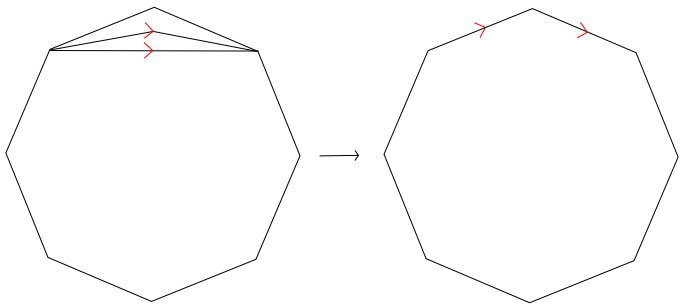


Figure 2.8: The homotopy of the side α after the cut-and-paste move.

Hence, inductively, we obtain that after applying all cut-and-paste moves, the image of the sides of P_γ can be homotoped to lie on the sides of $P(\pi_t', \pi_b')$, and letting $e_1, \ldots, e_k$ be the edges of γ, the α-row of the matrix $V_{e_1} \cdot \ldots \cdot V_{e_k}$ describes the image of α. Note that the order we multiply the matrices together is the same as the one we use in the proof of

Lemma 2.7, even though we apply ψ instead of ψ^{-1}. This is because here sides should be thought of as row vectors, so we are multiplying matrices from the right, while previously the branches of the train track had to be thought of as column vectors, so we multiplied matrices from the left.

Finally, we need to apply φ^{-1}. This just permutes the sides of the polygon. The columns of the matrix P encode the permutation according to φ. Hence, in order to encode φ^{-1} in terms of rows of the matrix, we need to multiply with $(P^{-1})^T$. Since P is a permutation matrix, its transpose is its inverse, so $(P^{-1})^T = P$.

We conclude that f_γ^{-1} maps sides of P_γ to paths that can be homotoped to lie only on sides of P_γ and that the rows of the matrix $V_{e_1} \cdot \ldots \cdot V_{e_k} \cdot P = V_\gamma$ encode how often each path $f_\gamma^{-1}(\alpha)$ passes over each side. □

Lemma 2.8 shows that the rows of the matrix V_γ describe the image of the corresponding sides of P_γ under f_γ^{-1}. Note, that if we orient the sides as in Figure 2.3, then for any α, we have that $f_\gamma^{-1}(\alpha)$ traverses each side β in the given orientation. This is because the gluing of the triangle for a Rauzy-Veech move preserves this orientation.

Lemma 2.9. *Let α be a side of P_γ that is never a winner for any edge of γ. Then the α-row of V_γ just has a single non-zero entry which is equal to* 1.

Proof. A cut-and-paste move corresponding to a Rauzy-Veech move maps any side of the starting polygon that is not a winner to some side of the resulting polygon. Hence, if α is never a winner, the image of α after applying all the edges of γ will be a single side β. According to Lemma 2.8, the α-row of V_γ is consequently 0 in every entry except for (α, β) where it is 1. □

Lemma 2.9 tells us that if α is never a winner in any step of the path γ, then the side corresponding to it in P_γ, resp. the path in X_γ, gets mapped under f_γ^{-1} onto a single other side, resp. path. This will be crucial in the computation of the upper bound of the stable curve graph translation length in Section 2.4, since sides with this property are easy to control.

§ 2.3.1 | The flip move

In order to define the pseudo-Anosovs f_g, apart from Rauzy-Veech moves, we will need to consider the so-called *flip move* f. Similarly to the Rauzy-Veech moves, given a labelled

permutation, the flip move produces a new labelled permutation in the following way:

$$\begin{pmatrix} \pi_t^{-1}(1) & \pi_t^{-1}(2) & \dots & \pi_t^{-1}(n) \\ \pi_b^{-1}(1) & \pi_b^{-1}(2) & \dots & \pi_b^{-1}(n) \end{pmatrix} \xrightarrow{f} \begin{pmatrix} \pi_b^{-1}(n) & \pi_b^{-1}(n-1) & \dots & \pi_b^{-1}(1) \\ \pi_t^{-1}(n) & \pi_t^{-1}(n-1) & \dots & \pi_t^{-1}(1) \end{pmatrix}.$$

For example, we have

$$\begin{pmatrix} A & C & B \\ B & A & C \end{pmatrix} \xrightarrow{f} \begin{pmatrix} C & A & B \\ B & C & A \end{pmatrix}.$$

We can also give a geometric meaning to the flip move. Let (π_t, π_b) be a labelled permutation and $P = P(\pi_t, \pi_b)$ the corresponding polygon. We define the flip move on P as the rotation of P by 180 degrees. This defines a polygon P' which yields a new labelled permutation (π_t', π_b') that is exactly the one obtained by the flip move on (π_t, π_b).

We define the *labelled augmented Rauzy diagram* to be the labelled Rauzy diagram with extra edges corresponding to the flip move. We label these extra edges by f. Note that the labelled augmented Rauzy diagram might have fewer connected components than the labelled Rauzy diagram. As before, given a path γ in the labelled augmented Rauzy diagram, we say that γ is allowed if the starting point and endpoint define the same unlabelled permutation, and in an analogous way, we obtain a corresponding mapping class f_γ of X_γ.

We again assign a matrix V_γ to γ in the same way as before, where, when we pass through an edge corresponding to a flip move, we don't change our matrix. So, V_γ is again a non-negative integer matrix.

Example 2.10. Let

$$(\pi_t, \pi_b) = \begin{pmatrix} A & B & C \\ C & A & B \end{pmatrix}$$

and consider the path $\gamma = f$ in the labelled augmented Rauzy diagram starting at (π_t, π_b) and consisting of a single flip move. Since

$$\begin{pmatrix} A & B & C \\ C & A & B \end{pmatrix} \xrightarrow{f} \begin{pmatrix} B & A & C \\ C & B & A \end{pmatrix},$$

we see that starting and ending permutation of γ define the same unlabelled permutation, namely (1 2 3). Hence, we obtain a mapping class f_γ and a corresponding matrix V_γ. In this case, V_γ is equal to the permutation matrix, since the flip move doesn't contribute

anything by definition, i.e.

$$V_\gamma = P = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

The mapping class f_γ is an involution of the torus, since in terms of polygons it corresponds to a 180-degree rotation of a hexagon with sides glued as imposed by the permutations.

In the above example, the matrix V_γ was of order 2, so in particular not primitive. Here, we show:

Lemma 2.11. *Let γ be a path in the labelled augmented Rauzy diagram. Let f_γ and V_γ be the corresponding mapping class and matrix. If V_γ is a primitive matrix, then f_γ is pseudo-Anosov.*

Proof. Assume that γ consists of n edges out of which $1 \leqslant k \leqslant n$ are flip moves. The main observation for the proof is that geometrically, a flip move followed by a Rauzy-Veech move is the same as an analogous Rauzy-Veech move on the left side of the polygon followed by a flip move. More precisely, if we denote cut-and-paste moves (as defined in Section 2.2) by t^R, b^R and the analogously defined cut-and-paste moves on the left side of the polygon by t^L, b^L, we have

$$ft^R = b^L f, \text{ and } fb^R = t^L f.$$

These equations should be read as saying: applying a top (resp. bottom) move on the right followed by a flip move is the same as first applying a flip move followed by a bottom (resp. top) move on the left.

Hence, we can write our cut-and-paste homeomorphism ψ corresponding to f_γ as a composition of cutting and pasting on the right and on the left, followed by k flip moves. Since the flip moves are of order two, depending on the parity of k, we are either left with one or with zero flip moves.

Now, analogously to the proof of Lemma 2.7, we have to check that the train track τ_γ is carried by f_γ and that the train track matrix is given by V_γ. If we only have to do cut-and-paste moves on the right, we already know this from Lemma 2.7. Assume we have to do a cut-and-paste move on the left. We look at the case where we have to do a top move t^L, (the case b^L is analogous). This means that our path in the labelled augmented Rauzy diagram had an edge corresponding to a b^R move, which we substituted

by $ft^Lf = ft^Lf^{-1} = b^R$. Hence, we have to check that a t^L move with winner-loser pair (α, β) has the same effect as a b^R move with winner-loser pair (α, β). The action on the train track by the latter is given by the matrix $V_{\alpha,\beta}$.

Performing a top move on the left with winner-loser pair (α, β) means that we cut the triangle on the left and glue its α-side to the corresponding α-side on the bottom (compare Figure 2.9, which is the flipped version of Figure 2.5).

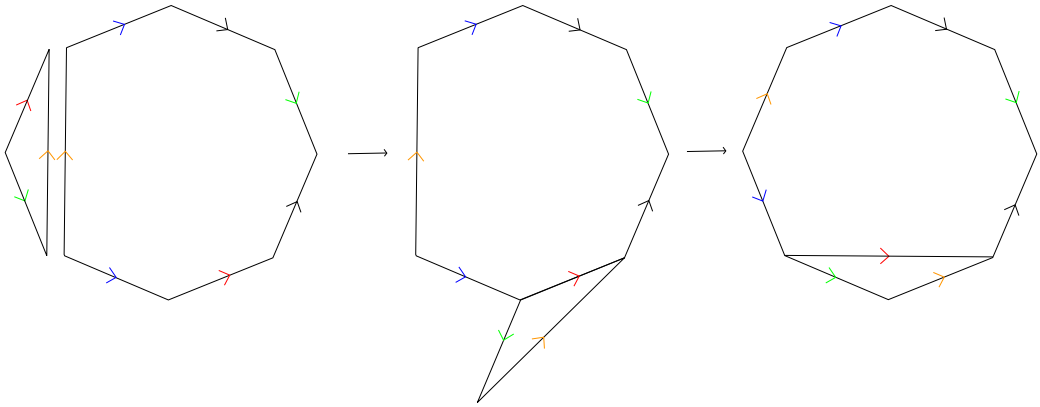


Figure 2.9: A top move on the left.

In order to see the effect on the train track, we have to apply the inverse of this move and observe that every branch of the train track remains the same, except for the β-branch that now runs exactly once over itself and once over the α-branch (compare Figure 2.10, the flipped version of Figure 2.7, where $(\alpha, \beta) = (B, C)$). This is captured by the matrix $V_{\alpha,\beta}$, which is what we wanted to show.

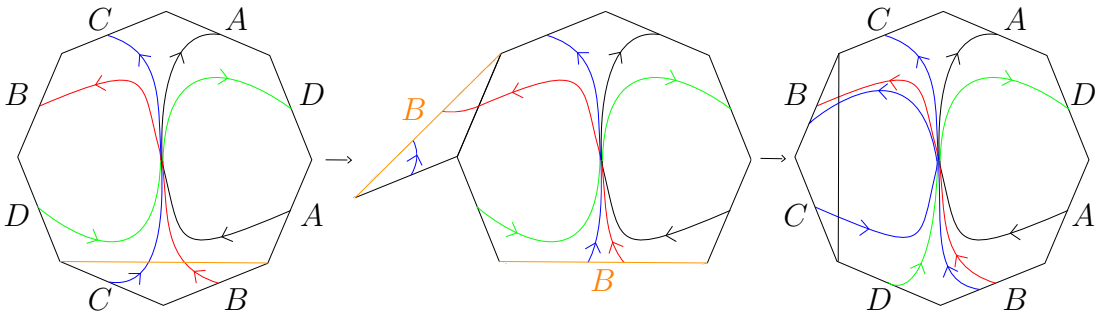


Figure 2.10: The inverse of a left t-move with its effect on the train track.

Note that if, after performing all left and right cut-and-paste moves, ψ does consist of a flip move, this only changes the orientation of the train track. The α-branch of the train track will still be the α-branch after the flip for every α. Hence, the train track τ_γ is carried by f_γ and the train track matrix is given by V_γ.

It is a standard fact that if a large, birecurrent train track is carried by a mapping class and the corresponding train track matrix is primitive, then the mapping class is pseudo-Anosov [Pen88, Corollary 3.2]. The train track τ_γ is large, since its complement is a union of polygons. Furthermore, τ_γ is recurrent since every branch itself is a closed train route. Finally, each branch of τ_γ is intersected efficiently by a simple closed curve that can be built as a concatenation of sides of the corresponding polygon P_γ. This concludes the proof. □

§ 2.4 | Translation lengths in hyperelliptic components

In this section, we prove Theorems 2.1 and 2.2. We start by defining the pseudo-Anosov maps f_g, which is due to Boissy–Lanneau. We are not going to prove that the f_g are indeed in a hyperelliptic component, but instead refer the reader to [BL12] or [BL22]. When defining the f_g, we use a different labelling of the permutations Boissy–Lanneau use. We then proceed with proving Theorem 2.1. We split the proof into the proof for the upper bound and the lower bound of $l_{\mathcal{C}}(f_g)$ since the proofs use different techniques. Finally, we prove Theorem 2.2.

§ 2.4.1 | The pseudo-Anosovs f_g

First, we recall the construction of the f_g which can be found in [BL12], as well as [BL22].

Consider the labelled permutation

$$(\pi_t^g, \pi_b^g) = \begin{pmatrix} \alpha_1 & \alpha_2 & \alpha_3 & \dots & \alpha_g & \alpha_{g+1} & \alpha_{g+2} & \dots & \alpha_{2g} \\ \alpha_{2g} & \alpha_{g-1} & \alpha_{g-2} & \dots & \alpha_1 & \alpha_{2g-1} & \alpha_{2g-2} & \dots & \alpha_g \end{pmatrix},$$

and the path $\gamma_g = ftb^g$ in the labelled augmented Rauzy diagram starting at (π_t^g, π_b^g), i.e. γ_g is the path obtained by first applying g times the type b-move, followed by once the type t-move and finally once the flip move f. Note that we write ftb^g, i.e. use the convention to read from right to left. Furthermore, we write $x \cdot (\pi_t, \pi_b)$ to denote the endpoint of any path x in the labelled augmented Rauzy-diagram starting at a labelled permutation (π_t, π_b).

After applying b a number of k times ($k \in \{1, \dots, g\}$) to our starting permutation, we end up at the permutation

$$b^k \cdot (\pi_t^g, \pi_b^g) = \begin{pmatrix} \alpha_1 & \alpha_2 & \dots & \alpha_g & \alpha_{2g-k+1} & \dots & \alpha_{2g} & \alpha_{g+1} & \dots & \alpha_{2g-k} \\ \alpha_{2g} & \alpha_{g-1} & \dots & \alpha_1 & \alpha_{2g-1} & \dots & \alpha_{2g-k} & \alpha_{2g-k-1} & \dots & \alpha_g \end{pmatrix}.$$

Whenever we write "..." in the above permutation, we mean to continue the sequence

by adding or subtracting one to the subscript as appropriate. We observe that after g moves of type b we obtain $b^g \cdot (\pi_t^g, \pi_b^g) = (\pi_t^g, \pi_b^g)$. Hence, we traced out a loop in the labelled augmented Rauzy diagram and can therefore already obtain a mapping class corresponding to this loop. However, this mapping class wouldn't be a pseudo-Anosov. We continue the path by applying t.

From the above, we obtain

$$tb^g \cdot (\pi_t^g, \pi_b^g) = t \cdot (\pi_t^g, \pi_b^g) = \begin{pmatrix} \alpha_1 & \alpha_2 & \dots & \alpha_{g+1} & \alpha_{g+2} & \dots & \alpha_{2g} \\ \alpha_{2g} & \alpha_g & \dots & \alpha_1 & \alpha_{2g-1} & \dots & \alpha_{g+1} \end{pmatrix}.$$

Finally, we apply f to obtain

$$\gamma_g \cdot (\pi_t^g, \pi_b^g) = f \cdot \begin{pmatrix} \alpha_1 & \alpha_2 & \dots & \alpha_{g+1} & \alpha_{g+2} & \dots & \alpha_{2g} \\ \alpha_{2g} & \alpha_g & \dots & \alpha_1 & \alpha_{2g-1} & \dots & \alpha_{g+1} \end{pmatrix} =$$

$$= \begin{pmatrix} \alpha_{g+1} & \dots & \alpha_{2g-1} & \alpha_1 & \dots & \alpha_g & \alpha_{2g} \\ \alpha_{2g} & \dots & \alpha_{g+2} & \alpha_{g+1} & \dots & \alpha_2 & \alpha_1 \end{pmatrix}.$$

Comparing the starting point and endpoint of γ_g, we see that the two labelled permutations define the same unlabelled permutation, and therefore γ_g defines a mapping class f_g on the surface $X_g := X(\pi_t^g, \pi_b^g)$. One easily checks by induction that the surface X_g is indeed a surface of genus g.

Boissy–Lanneau show that f_g is pseudo-Anosov and that its stretch factor is bounded below by $\sqrt{2}$ for all g [BL12]. This is done by computing the matrix V_{γ_g}, showing that it is primitive and computing its spectral radius. Since Boissy–Lanneau use a different labelling, the matrix they obtain differs from the one we compute below by conjugation by a permutation matrix.

§ 2.4.2 | Proof of Theorem 2.1

Our main goal is to prove the following:

Theorem. *For all $g \geqslant 2$, it holds that*

$$\frac{1}{16g-12} \leqslant l_C(f_g) \leqslant \frac{1}{g-1}.$$

The remaining part of this section is devoted to proving the theorem. Since $l_C(f_g) = l_C(f_g^{-1})$, we can use either f_g or its inverse for the proof. We use f_g^{-1} for the proof of the upper bound and f_g for the proof of the lower bound.

The first thing to note is that on the surface $X(\pi_t^g, \pi_b^g)$, all polygon sides $\alpha_1, ..., \alpha_{2g}$ are closed curves. This follows from the equivalent statement in the following lemma.

Lemma 2.12. *All the corners of the polygon $P(\pi_t^g, \pi_b^g)$ define the same point in $X(\pi_t^g, \pi_b^g)$.*

Proof. Consider the corner on the left of the top side corresponding to α_1. Since in $X(\pi_t^g, \pi_b^g)$ the top and bottom sides corresponding to α_1 are identified, the above corner defines the same point as the left corner of the bottom side corresponding to α_1. This, in turn, is the same corner as the right corner of the bottom side corresponding to α_2. From the gluing, we move to the right corner of the top side corresponding to α_2, which is the same as the left corner of the top side corresponding to α_3. Iterating this process shows that all the left corners of the top sides corresponding to α_i with i odd define the same point in $X(\pi_t^g, \pi_b^g)$.

After reaching the left corner of the top side corresponding to $2g-1$ and gluing to the bottom, we end up at the right corner of the bottom side corresponding to α_1, which is the same as the left corner of the top side corresponding to α_2. From there, iterating again shows that we also identify all the left corners of the top sides of the α_i with i even. Hence, all the corners get identified in $X(\pi_t^g, \pi_b^g)$. □

Since the α_i are closed curves, we can use them to find bounds on the stable curve graph translation length of f_g^{-1}. The ones that are easier to control are the ones that, seen as sides of the polygon, get mapped to another side of the polygon.

Lemma 2.13. *The only two curves that are winners for some edge of γ_g are α_g and α_{2g}.*

Proof. For the first b-move in γ_g, the winner is $(\pi_b^g)^{-1}(2g) = \alpha_g$. Since after a b-move, π_b^g remains unchanged, we get that for all the following b-moves, the winner is still α_g. Finally, we apply the t-move to $b^g \cdot (\pi_t^g, \pi_b^g) = (\pi_t^g, \pi_b^g)$, so the winner is $(\pi_t^g)^{-1}(2g) = \alpha_{2g}$. □

The above, together with Lemma 2.9, implies that all other curves α_i with $i \notin \{g, 2g\}$ are mapped by f_g^{-1} onto some other α_j, where $j \in \{1, ..., 2g\}$. Therefore, for any curve α_i with $i \notin \{g, 2g\}$, we can read off its image just by looking at the labelled permutation that is the endpoint of γ.

In particular, if we write $f_g^{-1} = \varphi^{-1} \circ \psi$ as in Section 2.2.2, then for $i \notin \{g, 2g\}$, ψ takes α_i to the side α_i of $X(\gamma_g(\pi_t^g, \pi_b^g))$. Hence, $f_g^{-1}(\alpha_i) = \varphi^{-1}(\alpha_i)$. It remains to study the relabelling homeomorphism φ.

Let $(\widetilde{\pi}_t^g, \widetilde{\pi}_b^g) := \gamma_g \cdot (\pi_t^g, \pi_b^g)$ be the endpoint of γ_g. Since on the level of polygons φ just maps $P(\pi_t^g, \pi_b^g)$ onto $P(\widetilde{\pi}_t^g, \widetilde{\pi}_b^g)$ as the identity, we can read off what φ does on the sides of the polygons by just looking at the labelled permutations, i.e.

$$\varphi(\alpha_i) = (\widetilde{\pi}_t^g)^{-1} \circ \pi_t^g(\alpha_i), \text{ or equivalently } \varphi(\alpha_i) = (\widetilde{\pi}_b^g)^{-1} \circ \pi_b^g(\alpha_i).$$

We summarise what we showed in the following lemma.

Lemma 2.14. *For any $i \notin \{g, 2g\}$, the pseudo-Anosov f_g^{-1} maps the curve α_i to the curve*

$$\varphi^{-1}(\alpha_i) = (\pi_t^g)^{-1} \circ \widetilde{\pi}_t^g(\alpha_i).$$

From the starting and ending points of γ, we read off that

$$(\pi_t^g)^{-1} \circ \widetilde{\pi}_t^g(\alpha_i) = \begin{cases} \alpha_{g+i-1} & \text{if } i = 1, ..., g \\ \alpha_{i-g} & \text{if } i = g+1, ..., 2g-1 \\ \alpha_{2g} & \text{if } i = 2g. \end{cases}$$

We are now ready to prove the upper bound for Theorem 2.1.

Proof of upper bound in Theorem 2.1. The strategy we want to use for proving that a pseudo-Anosov f satisfies $l_{\mathcal{C}}(f) \leqslant \frac{k}{n}$ for some $k, n \in \mathbb{N}$ is the following: We find a curve α such that

$$d_{\mathcal{C}}(\alpha, f^n(\alpha)) \leqslant k.$$

Then, by the triangle inequality and the fact that f acts as an isometry on the curve graph, it follows that for any $j \geqslant 1$

$$\frac{d_{\mathcal{C}}(\alpha, f^{jn}(\alpha))}{j} \leqslant \frac{j \cdot d_{\mathcal{C}}(\alpha, f^n(\alpha))}{j} \leqslant k,$$

and so $l_{\mathcal{C}}(f^n) \leqslant k$. It is now easy to see that then $l_{\mathcal{C}}(f) \leqslant \frac{k}{n}$ (see [GT11, Lemma 2.2]).

Consider f_g as defined at the beginning of the section. We want to apply the aforementioned strategy to f_g^{-1} and the curve α_{2g-1}.

From Lemma 2.14 and the equation after, we see that $f_g^{-1}(\alpha_{2g-1}) = \alpha_{g-1}$. Iterating this

process yields

$$\begin{aligned} f_g^{-1}(\alpha_{2g-1}) &= \alpha_{g-1} \\ f_g^{-2}(\alpha_{2g-1}) &= \alpha_{2g-2} \\ f_g^{-3}(\alpha_{2g-1}) &= \alpha_{g-2} \\ f_g^{-4}(\alpha_{2g-1}) &= \alpha_{2g-3} \\ &\vdots \\ f_g^{-(2g-2)}(\alpha_{2g-1}) &= \alpha_g. \end{aligned}$$

More formally, we have $f_g^{-k}(\alpha_{2g-1}) = \alpha_{g-\frac{k+1}{2}}$ for all odd k between 1 and $2g-3$ and $f_g^{-k}(\alpha_{2g-1}) = \alpha_{2g-\frac{k+2}{2}}$ for all even k between 2 and $2g-2$.

Note that after $2g-2$ iterates of f_g^{-1}, we hit the curve α_g for which Lemma 2.14 doesn't apply anymore, so we can't use it to say something about $f_g^{-(2g-1)}(\alpha_{2g-1})$. Note furthermore, that the orbit of α_{2g-1} under f_g^{-1} traces out all the curves α_i with $i \notin \{g, 2g\}$ before hitting the curve α_g, so out of all the choices we could make in order to apply Lemma 2.14, this gives the best upper bound.

Since α_{2g-1} and α_g are both sides of the polygon $P(\pi_t^g, \pi_b^g)$, the curves α_{2g-1} and α_g of the surface $X(\pi_t^g, \pi_b^g)$ intersect in one point, which is the image of the corners of $P(\pi_t^g, \pi_b^g)$. So, the geometric intersection number of the two curves is at most 1.

If the geometric intersection number is 0, then $d_{\mathcal{C}}(\alpha_{2g-1}, \alpha_g) = 1$. If the geometric intersection number is 1, then by thickening up the curves, one can find a neighbourhood of their union that is homeomorphic to a torus with one boundary component β. This β is a closed curve that is disjoint from both α_{2g-1} and α_g, and it is furthermore essential, because the genus of our surface is $\geqslant 2$. It follows that $d_{\mathcal{C}}(\alpha_{2g-1}, \alpha_g) = 2$.

In any case, we have

$$d_{\mathcal{C}}(\alpha_{2g-1}, f_g^{-(2g-2)}(\alpha_{2g-1})) = d_{\mathcal{C}}(\alpha_{2g-1}, \alpha_g) \leqslant 2.$$

It follows that

$$l_{\mathcal{C}}(f_g) = l_{\mathcal{C}}(f_g^{-1}) \leqslant \frac{2}{2g-2} = \frac{1}{g-1}.$$

□

It remains to prove the lower bound of Theorem 2.1. This requires different techniques from the ones used in the proof of the upper bound. It is done by using an invariant train

track for f_g and the Nesting Lemma presented in Section 2.1.2.

Proof of lower bound in Theorem 2.1. Let τ be the train track on X_g as defined in Section 2.3 (compare Figure 2.6). The train track τ is large, since its complement is a $(4g-2)$-gon. Furthermore, τ is recurrent since every branch itself is a closed train route. Finally, in Lemma 2.7 (and the proof of Lemma 2.11) we showed that τ is invariant for f_g and the corresponding train track matrix is given by V_{γ_g}.

We have to compute V_{γ_g}. Since the sequence of winner-loser pairs of γ_g is given by

$$(\alpha_g, \alpha_{2g}), (\alpha_g, \alpha_{2g-1}), ..., (\alpha_g, \alpha_{g+1}), (\alpha_{2g}, \alpha_g),$$

we obtain

$$V_{\gamma_g} = V_{\alpha_g,\alpha_{2g}} \cdot V_{\alpha_g,\alpha_{2g-1}} \cdot ... \cdot V_{\alpha_g,\alpha_{g+1}} \cdot V_{\alpha_{2g},\alpha_g} \cdot P,$$

where P is the permutation matrix as defined in Section 2.2.2.

With the convention that the α_i are ordered according to their indices, we compute

$$V_{\alpha_g,\alpha_{2g}} \cdot V_{\alpha_g,\alpha_{2g-1}} \cdot ... \cdot V_{\alpha_g,\alpha_{g+1}} = \begin{pmatrix} Id & A \\ 0 & Id \end{pmatrix},$$

where each block is a $g \times g$ block and A is a matrix whose entries in the last row are all 1 and every other entry is 0. Furthermore, we have

$$V_{\alpha_{2g},\alpha_g} = \begin{pmatrix} Id & 0 \\ B & Id \end{pmatrix},$$

where B has a single non-zero entry equal to 1 in its bottom right corner and

$$P = \begin{pmatrix} 0_{g\times g-1} & Id_{g\times g} & 0_{g\times 1} \\ Id_{g-1\times g-1} & 0_{g-1\times g} & 0_{g-1\times 1} \\ 0_{1\times g-1} & 0_{1\times g} & 1_{1\times 1} \end{pmatrix},$$

where the subscripts indicate the size of each block.

In total, we obtain

$$V_{\gamma_g} = \begin{pmatrix} Id & A \\ 0 & Id \end{pmatrix}\begin{pmatrix} Id & 0 \\ B & Id \end{pmatrix} P = \begin{pmatrix} Id + B & A \\ B & Id \end{pmatrix} P = \begin{pmatrix} A_{g\times g-1} & Id_{g\times g} + B_{g\times g} & B_{g\times 1} \\ Id_{g-1\times g-1} & 0_{g-1\times g} & 0_{g-1\times 1} \\ 0_{1\times g-1} & B_{1\times g} & 1_{1\times 1} \end{pmatrix}.$$

In the above, the subscripts indicate the size of the matrices. $A_{g\times g-1}$ is a matrix whose entries in the last row are all 1 and every other entry is 0. $B_{g\times g}$ is the same as B and $B_{g\times 1}, B_{1\times g}$ have all entries 0 except for the $(g,1)$ or $(1,g)$ entry respectively, which is equal to 1.

Since the $(2g,2g)$ entry on the diagonal of V_{γ_g} is positive, we can use [Tsa09, Proposition 2.4] to argue that $V_{\gamma_g}^{4g}$ is a positive matrix. In fact, a more thorough inspection of V_{γ_g} shows that already its $(4g-4)^{\text{th}}$ power is a positive matrix, but we omit the details since we only need to obtain a bound of order $\frac{1}{g}$.

The fact that $V_{\gamma_g}^{4g}$ is a positive matrix implies that given a curve γ carried by τ, the curve $f_g^{4g}(\gamma)$ passes over every branch of τ. Given instead a curve γ carried by some diagonal extension of τ, then [GT11, Lemma 5.2] implies that $f_g^K(\gamma)$ with $K = 6(2g-2)+4g = 16g-12$ passes over every branch of τ.

Now, for $n \geqslant 1$ let $\tau_n := f_g^{Kn}(\tau)$ and set $\tau = \tau_0$. Then, the above observation shows that

$$PE(\tau_n) \subset int(PE(\tau_{n-1})).$$

From the Nesting Lemma (Section 2.1.2) we obtain the following sequence of inclusions:

$$\cdots PE(\tau_n) \subset int(PE(\tau_{n-1})) \subset \mathcal{N}_1(int(PE(\tau_{n-1}))) \subset PE(\tau_{n-1}) \subset int(PE(\tau_{n-2}))\cdots$$

Finally, choose a curve δ which is not carried by any extension of τ but such that $\gamma := f_g^K(\delta)$ is. Then for $n \geqslant 1$, we get that $f_g^{Kn}(\gamma)$ is an element of $PE(\tau_n)$. The above sequence of inclusions now implies that $d_{\mathcal{C}}(\gamma, f_g^{Kn}(\gamma)) \geqslant n$ for any $n \geqslant 1$. We conclude:

$$l_{\mathcal{C}}(f_g^K) = \liminf_{n\to\infty} \frac{d_{\mathcal{C}}(\gamma, f_g^{Kn}(\gamma))}{n} \geqslant 1,$$

which implies

$$l_{\mathcal{C}}(f_g) = \frac{l_{\mathcal{C}}(f_g^K)}{K} \geqslant \frac{1}{K} = \frac{1}{16g-12}.$$

□

§ 2.4.3 | Proof of Theorem 2.2

This Section is devoted to the proof of Theorem 2.2.

Let $g \geqslant 2$ and

$$l^{\mathcal{C}}_{g,\text{hyp}} := \min\{\, l_{\mathcal{C}}(f) \mid f \text{ is a pseudo-Anosov in a hyperelliptic component of genus } g\}.$$

In order to prove Theorem 2.2, we have to prove that $l^{\mathcal{C}}_{g,\mathrm{hyp}} \asymp \frac{1}{g}$. Hence, it suffices to find an upper and a lower bound for $l^{\mathcal{C}}_{g,\mathrm{hyp}}$ that are of order $\frac{1}{g}$. From Theorem 2.1, together with the fact that the pseudo-Anosov f_g is in a hyperelliptic component, it follows immediately that $l^{\mathcal{C}}_{g,\mathrm{hyp}} \leqslant l_{\mathcal{C}}(f_g) \leqslant \frac{1}{g-1}$. Hence, it remains to find a lower bound. For that, we will show that any pseudo-Anosov in a hyperelliptic component of genus g has stable curve graph translation length greater than or equal to $\frac{1}{16g-10}$. The argument we use is very similar to the proof of the lower bound in Theorem 2.1. We start with some preliminary observations that are necessary for the proof.

Let f be a pseudo-Anosov in a hyperelliptic component. Then, there exists a translation surface (X,ω) in a hyperelliptic component, a representative φ of f and a corresponding pseudo-Anosov φ_X on X (compare Section 2.1.3). Note that $l_{\mathcal{C}}(f) = l_{\mathcal{C}}(\varphi_X)$. Since X is a hyperelliptic Riemann surface, it admits a hyperelliptic involution ι. Following the notation of Boissy–Lanneau, we let $\{\varphi_X, \iota\circ\varphi_X\} = \{\varphi^+, \varphi^-\}$, where φ^+ preserves the orientation of the associated foliations and φ^- reverses it (see [BL22] for details). The important property for our purposes is given by the following lemma:

Lemma 2.15. *It holds that* $l_{\mathcal{C}}(\varphi^+) = l_{\mathcal{C}}(\varphi^-)$.

Proof. We have to show that $l_{\mathcal{C}}(\varphi_X) = l_{\mathcal{C}}(\iota\circ\varphi_X)$. Since ι is a hyperelliptic involution, there is a curve α that is fixed by ι. Furthermore, conjugating ι by φ_X yields a hyperelliptic involution of X which – since ι is unique – has to be equal to ι. Thus, ι commutes with φ_X (cf. [BL12, Lemma 2.3]). It follows that

$$(\iota\circ\varphi_X)^n(\alpha) = \varphi_X^n\circ\iota^n(\alpha) = \varphi_X^n(\alpha)$$

for any $n\in\mathbb{N}$, and therefore

$$l_{\mathcal{C}}(\iota\circ\varphi_X) = \liminf_{n\to\infty}\frac{d_{\mathcal{C}}(\alpha,(\iota\circ\varphi_X)^n(\alpha))}{n} = \liminf_{n\to\infty}\frac{d_{\mathcal{C}}(\alpha,\varphi_X^n(\alpha))}{n} = l_{\mathcal{C}}(\varphi_X).$$

□

Let $n\in\{2g, 2g+1\}$. Let $\mathcal{A}$ be an alphabet consisting of n letters $\alpha_1, \dots, \alpha_n$ and let

$$(\pi_t,\pi_b) = \begin{pmatrix}\alpha_1 & \alpha_2 & \dots & \alpha_{n-1} & \alpha_n\\ \alpha_n & \alpha_{n-1} & \dots & \alpha_2 & \alpha_1\end{pmatrix}.$$

Throughout this section, we refer to (π_t,π_b) as the *central permutation*. Consider the connected component of the labelled Rauzy diagram of (π_t,π_b). We call this component $\mathcal{D}_{(\pi_t,\pi_b)}$. Rauzy showed that this connected component is isomorphic to the so-called

unlabelled Rauzy diagram of $\pi = \pi_b \circ \pi_t^{-1}$ (see [Rau79] or [BL12]). This means in particular that no two different labelled permutations in $\mathcal{D}_{(\pi_t,\pi_b)}$ define the same unlabelled permutation.

Any mapping class defined through Rauzy-Veech induction on $\mathcal{D}_{(\pi_t,\pi_b)}$ is affine for a hyperelliptic component. More concretely, if $n = 2g$, then it's affine for $\mathcal{H}(2g-2)$, and if $n = 2g+1$, then for $\mathcal{H}(g-1, g-1)$ (see, for example [BL12]).

Given the central permutation, we can perform t-moves to it and obtain a loop in $\mathcal{D}_{(\pi_t,\pi_b)}$ consisting of $n-1$ vertices corresponding to

$$t^m \cdot (\pi_t, \pi_b) = \begin{pmatrix} \alpha_1 & \alpha_2 & \alpha_3 & \dots & \alpha_{m+1} & \alpha_{m+2} & \alpha_{m+3} & \dots & \alpha_n \\ \alpha_n & \alpha_m & \alpha_{m-1} & \dots & \alpha_1 & \alpha_{n-1} & \alpha_{n-2} & \dots & \alpha_{m+1} \end{pmatrix}$$

for $m = 1, \dots, n-1$. We refer to this loop in the labelled Rauzy diagram as the *central loop*.

We are now ready to prove Theorem 2.2:

Proof of Theorem 2.2. Let f be in a hyperelliptic component of genus g and let φ^+, φ^- be as constructed above. In order to prove Theorem 2.2, it suffices to show that $l_{\mathcal{C}}(\varphi^+) = l_{\mathcal{C}}(\varphi^-) \geqslant \frac{1}{16g-10}$.

Boissy–Lanneau prove the following statement:

If φ^+ cannot be obtained by Rauzy-Veech induction on $\mathcal{D}_{(\pi_t,\pi_b)}$, then φ^- is obtained by a path in the augmented labelled Rauzy diagram that starts at a permutation in the central loop and consists of a single flip move which corresponds to the last edge of the path (cf. [BL22, Theorem 4.1]).

From this, we obtain that at least one of φ^+, φ^- will be of the form f_γ (compare the notation of Section 2.2) where γ is one of two cases:

- γ is a closed path in $\mathcal{D}_{(\pi_t,\pi_b)}$.
- γ is a path in the augmented labelled Rauzy diagram starting at some $t^m \cdot (\pi_t, \pi_b)$ and whose last edge is the only edge with label f.

Hence, it suffices to analyse the stable curve graph translation length of f_γ in the two cases.

Recall from Section 2.3 that we assign a matrix V_γ to f_γ, which is a train track matrix for a natural train track carried by f_γ. We want to have a power k such that V_γ^k becomes

a positive matrix. Having found such a k, we can use [GT11, Lemma 5.2] and argue analogously to the proof of the lower bound in Theorem 2.1 to conclude $l_{\mathcal{C}}(f_\gamma) \geqslant \frac{1}{6(2g-2)+k}$.

We are left with the task of finding the power k. Recall, that since V_γ is primitive, it suffices to find a non-zero diagonal entry. Then, we can apply the result of Tsai to obtain k [Tsa09, Proposition 2.4].

We analyse the two cases of the path γ separately:

In the first case, γ is a closed path in $\mathcal{D}_{(\pi_t,\pi_b)}$. So, V_γ is a product of matrices $V_{\alpha,\beta}$. In particular, we don't have to multiply by a permutation matrix at the end, because the starting point and endpoint of γ are the same labelled permutation. Since the $V_{\alpha,\beta}$ are non-negative matrices with all diagonal entries equal to 1, all diagonal entries of V_γ are positive.

In the second case, γ is a path that starts at $t^m \cdot (\pi_t, \pi_b)$ for some $m \in \{1, \dots, n-1\}$. The edges of γ are all t or b edges, except for the last one which is an f edge. The last vertex of γ defines the same unlabelled permutation as the starting point $t^m \cdot (\pi_t, \pi_b)$. This implies that the second to last vertex of γ is a labelled permutation in $\mathcal{D}_{(\pi_t,\pi_b)}$ which defines the same unlabelled permutation as $f \cdot t^m \cdot (\pi_t, \pi_b)$. Since

$$t^m \cdot (\pi_t, \pi_b) = \begin{pmatrix} \alpha_1 & \alpha_2 & \alpha_3 & \dots & \alpha_{m+1} & \alpha_{m+2} & \alpha_{m+3} & \dots & \alpha_n \\ \alpha_n & \alpha_m & \alpha_{m-1} & \dots & \alpha_1 & \alpha_{n-1} & \alpha_{n-2} & \dots & \alpha_{m+1} \end{pmatrix},$$

we have

$$f \cdot t^m \cdot (\pi_t, \pi_b) = \begin{pmatrix} \alpha_{m+1} & \dots & \alpha_{n-2} & \alpha_{n-1} & \alpha_1 & \dots & \alpha_{m-1} & \alpha_m & \alpha_n \\ \alpha_n & \dots & \alpha_{m+3} & \alpha_{m+2} & \alpha_{m+1} & \dots & \alpha_3 & \alpha_2 & \alpha_1 \end{pmatrix}.$$

The labelled permutation

$$(\pi_t', \pi_b') := t^{n-m-1} \cdot (\pi_t, \pi_b) = \begin{pmatrix} \alpha_1 & \alpha_2 & \alpha_3 & \dots & \alpha_{n-m} & \alpha_{n-m+1} & \alpha_{n-m+2} & \dots & \alpha_n \\ \alpha_n & \alpha_{n-m-1} & \alpha_{n-m-2} & \dots & \alpha_1 & \alpha_{n-1} & \alpha_{n-2} & \dots & \alpha_{n-m} \end{pmatrix}$$

is an element of $\mathcal{D}_{(\pi_t,\pi_b)}$ which defines the same unlabelled permutation as $f \cdot t^m \cdot (\pi_t, \pi_b)$. Since there is a unique such permutation, the second-to-last vertex of γ has to be (π_t', π_b'). Hence, the endpoint of γ is

$$f \cdot (\pi_t', \pi_b') = \begin{pmatrix} \alpha_{n-m} & \dots & \alpha_{n-2} & \alpha_{n-1} & \alpha_1 & \dots & \alpha_{n-m-2} & \alpha_{n-m-1} & \alpha_n \\ \alpha_n & \dots & \alpha_{n-m+2} & \alpha_{n-m+1} & \alpha_{n-m} & \dots & \alpha_3 & \alpha_2 & \alpha_1 \end{pmatrix}.$$

From the above, we obtain that V_γ is a product of matrices $V_{\alpha,\beta}$ followed by multiplication with a permutation matrix P which encodes the relabelling between $t^m \cdot (\pi_t, \pi_b)$ and

$f \cdot (\pi_t', \pi_b')$. Notice that the n^{th} entry in the top row of the two labelled permutations is the same, namely α_n. This implies that the (n, n)-entry of P is equal to 1. Since any product of matrices $V_{\alpha,\beta}$ has positive entries on the diagonal, it follows that V_γ has a positive entry on the diagonal, namely the entry (n, n).

We showed that in both cases the matrix V_γ has a positive entry on the diagonal. Since V_γ is an $n \times n$ matrix with either $n = 2g$ or $n = 2g + 1$, using [Tsa09, Proposition 2.4], we obtain that V_γ^{4g+2} is a positive matrix. Hence, we have found the power $k = 4g + 2$ and conclude that $l_{\mathcal{C}}(f_\gamma) \geqslant \frac{1}{6(2g-2)+4g+2} = \frac{1}{16g-10}$.

Since this lower bound is of order $\frac{1}{g}$, this concludes the proof. □

Chapter 3

Infinite multiplicity of $l_{\mathcal{C}}$

We begin this chapter by examining the relationship between Teichmüller and curve graph translation lengths of pseudo-Anosovs. In particular, we construct a sequence of pseudo-Anosov mapping classes whose Teichmüller translation lengths diverge to infinity, while the curve graph translation lengths remain small. As a corollary of this result, we obtain that there exist curve graph translation lengths with infinite multiplicity. We then proceed to give an explicit way of constructing infinitely many non-conjugate pseudo-Anosovs with the same curve graph translation length. In what follows, we use the same notation as in the previous chapter.

Using the systole map from $\mathcal{T}(S_g)$ to $\mathcal{C}(S_g)$, which is coarsely Lipschitz, one can show that $l_{\mathcal{C}}(f) \leqslant K \cdot l_{\mathcal{T}}(f)$ for any $f \in \mathrm{Mod}(S_g)$, where the constant K depends only on the genus of S_g and can be chosen to be approximately $\frac{1}{\log(g)}$ (see [GHKL13] for details). In particular, this shows that small Teichmüller space translation length (or stretch factor) implies small stable curve graph translation length, but raises the question about the opposite. The f_g from Theorem 2.1 are an example of a sequence of pseudo-Anosovs such that $l_{\mathcal{T}}(f_g)$ is uniformly bounded away from 0, but $l_C(f_g) \to 0$ as $g \to \infty$. One can ask if an even better statement is possible, i.e. if there exists a sequence (h_g) with h_g a pseudo-Anosov of S_g such that $l_{\mathcal{T}}(h_g) \to \infty$ and $l_{\mathcal{C}}(h_g) \to 0$, as $g \to \infty$. Another question would be what happens when we fix the genus, i.e. can we find a sequence of pseudo-Anosovs $(f_n)_{n=1}^{\infty} \subset \mathrm{Mod}(S_g)$ with $l_{\mathcal{T}}(f_n) \to \infty$ as $n \to \infty$ while $l_{\mathcal{C}}(f_n) \leqslant C$ for some constant $C > 0$ and all n. Note that after fixing the genus, this is the best we can ask for since there is a minimal stable curve graph translation length, so we can't have that $l_{\mathcal{C}}(f_n)$ approaches 0.

We provide affirmative answers to the questions above by showing:

Theorem 3.1. *For any $g \geqslant 3$, there exists a sequence $(f_n)_{n=1}^{\infty} \subset \mathrm{Mod}(S_g)$ of pseudo-Anosovs with*

$$\lim_{n\to\infty} l_{\mathcal{T}}(f_n) = \infty \text{ and } l_{\mathcal{C}}(f_n) \leqslant \frac{1}{g-1} \text{ for all } n \in \mathbb{N}.$$

We remark that in Theorem 3.1, the constant bounding the curve graph translation length can even be chosen to be of order $\frac{1}{g^2}$. This will be discussed in the proof of the Theorem. However, it is enough to state the Theorem as above in order to obtain the following:

Corollary 3.2. *There exists a sequence $(h_g)_{g=2}^{\infty}$, where h_g is a pseudo-Anosov of a genus g surface with*

$$l_{\mathcal{T}}(h_g) \to \infty \text{ and } l_{\mathcal{C}}(h_g) \to 0,$$

as $g \to \infty$.

Theorem 3.1 gives another interesting consequence about stable curve graph translation lengths, namely that there are such lengths with infinite multiplicity:

Corollary 3.3. *For any $g \geqslant 3$, there exists $q \in \mathbb{Q}$ such that there are infinitely many non-conjugate pseudo-Anosovs in $\mathrm{Mod}(S_g)$ with stable curve graph translation length q.*

This is different from Teichmüller space translation length, where it is known that some $x \in \mathbb{R}$ is attained as $l_{\mathcal{T}}(f)$ by at most finitely many non-conjugate pseudo-Anosovs f. We remark that the result of Corollary 3.3 is folklore among experts in the field. For example, the pseudo-Anosovs constructed by Watanabe can also be used to obtain the corollary [Wat20].

The result asserts the existence of a number $q \in \mathbb{Q}$ with infinite multiplicity with respect to stable curve graph translation lengths. However, given a pseudo-Anosov f, it is not clear whether $l_{\mathcal{C}}(f)$ is attained by an infinite number of non-conjugate pseudo-Anosovs. We end this work by presenting a method that, for certain pseudo-Anosovs f, constructs infinitely many non-conjugate pseudo-Anosovs f_i with $l_{\mathcal{C}}(f_i) = l_{\mathcal{C}}(f)$. In particular, we prove:

Theorem 3.4. *Let f be a pseudo-Anosov with $l_{\mathcal{C}}(f) = m$, where $m \in \mathbb{N}$, and assume that there is a non-separating curve α attaining the translation length, i.e.*

$$d_{\mathcal{C}}(\alpha, f^k(\alpha)) = km,$$

for all $k \in \mathbb{Z}$. Then, for any pseudo-Anosov h of the surface $S_g \setminus \alpha$, there exists $N_h \in \mathbb{N}$ such that for $n \geqslant N_h$, each $f_{h,n} := h^n \circ f$ satisfies $l_{\mathcal{C}}(f_{h,n}) = l_{\mathcal{C}}(f)$.

The property of non-conjugacy of the $f_{h,n}$ is discussed in Section 3.2. By the surface $S_g \setminus \alpha$ in the above theorem, we mean the surface of genus $g-1$ with two boundary components resulting from removing an open, annular neighbourhood of α from S_g. Note that the pseudo-Anosov h fixes the boundary of $S_g \setminus \alpha$ pointwise and can hence be extended to a homeomorphism of S_g, i.e. when we write $h \circ f$ we think of h as an element of $\mathrm{Mod}(S_g)$. Of course, h seen as an element of $\mathrm{Mod}(S_g)$ is not a pseudo-Anosov but is reducible, since it fixes the curve α.

As described in the previous chapters, every pseudo-Anosov admits a power with an invariant geodesic axis in the curve graph. Hence, if the axis consists of non-separating curves, then this power belongs to the set of pseudo-Anosovs to which Theorem 3.4 applies. Another example of a pseudo-Anosov that satisfies the prerequisites of the Theorem is one constructed as a product of large powers of Dehn twists about two filling curves. Such a pseudo-Anosov has an invariant geodesic axis [BK24]. Again, if this axis consists of non-separating curves, then Theorem 3.4 applies.

It is not clear whether the conditions in Theorem 3.4 can be relaxed or whether a similar construction can be used for all pseudo-Anosovs; the question of whether all stable curve graph translation lengths have infinite multiplicity remains open.

To prove Theorem 3.1, we modify Penner's famous example of a pseudo-Anosov [Pen91]. For the proof of Theorem 3.4, we use a result about subsurface projections from Masur–Minsky [MM00].

§ 3.1 | Modification of Penner's example

In this section, we prove Theorem 3.1 as well as Corollaries 3.2 and 3.3, which follow directly from the theorem. For the proof of Theorem 3.1, we first construct the f_n and show that they are pseudo-Anosov. We proceed with finding bounds for their stretch factors and stable curve graph translation lengths, respectively, to prove the claims of Theorem 3.1. The idea is to modify Penner's example of a pseudo-Anosov in order to obtain the desired properties.

For $g \geqslant 3$, consider a genus g surface S_g as in Figure 3.1.

Let ρ be the order g homeomorphism that rotates the surface in Figure 3.1 anti-clockwise and let a, b, c be the curves depicted. Let T_γ denote the right-handed Dehn twist about a curve γ. Penner shows that the mapping class $\rho \circ T_c \circ T_a^{-1} \circ T_b$ is pseudo-Anosov (see [Pen91]). We define $f_n \in \mathrm{Mod}(S_g)$ by $f_n := \rho \circ T_c \circ T_a^{-1} \circ T_b^n$.

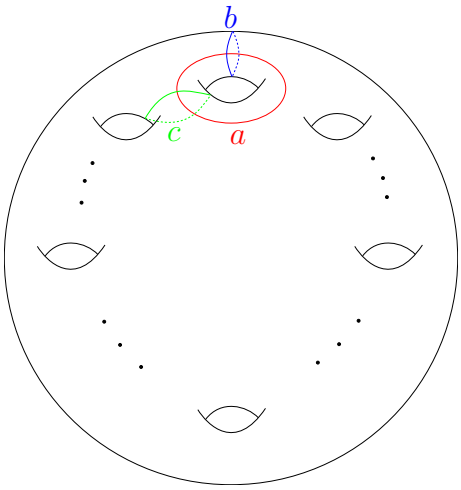


Figure 3.1: The surface S_g.

Lemma 3.5. *$f_n = \rho \circ T_c \circ T_a^{-1} \circ T_b^n$ is pseudo-Anosov.*

Proof. The curves $\rho^i(a), \rho^i(b), \rho^i(c)$ for $i = 1, \dots, g$ fill the surface. Note that

$$f_n^g = (\rho \circ T_c \circ T_a^{-1} \circ T_b^n)^g = \prod_{i=1}^{g} T_{\rho^i(c)} \circ T_{\rho^i(a)}^{-1} \circ T_{\rho^i(b)}^n.$$

From [Pen88, Theorem 3.1], with $\mathcal{C} = \{\rho^i(b), \rho^i(c) \,|\, i = 1, \dots, g\}$ and $\mathcal{D} = \{\rho^i(a) \,|\, i = 1, \dots, g\}$, it follows that f_n^g is pseudo-Anosov. Assume f_n were periodic or reducible. Then, some power f_n^m of f_n fixes a curve. Hence f_n^{mg} fixes that curve too, which is absurd since f_n^g is pseudo-Anosov. So, f_n has to be pseudo-Anosov. □

Lemma 3.6. *For all n, it holds that $\lambda(f_n) \geqslant n^{\frac{1}{g}}$, and consequently $l_{\mathcal{T}}(f_n) \geqslant \log(n^{\frac{1}{g}})$.*

Proof. The proof is analogous to the one Penner presents in [Pen91, Section "An upper bound by example"]. There, Penner defines a train track which is invariant for f_n. By analysing the space of measures arising as counting measures of the curves

$$a, b, c, \rho(a), \rho(b), \rho(c), \dots, \rho^{g-1}(a), \rho^{g-1}(b), \rho^{g-1}(c)$$

on this train track, Penner constructs a $3g \times 3g$-matrix M which counts how often each curve passes over any other curve and shows that the spectral radius of M is the stretch factor of the pseudo-Anosov f_1. More precisely, if we denote the above curves in the above order by γ_i, with $i = 1, \dots, 3g$, then the (i, j) entry of M counts how often $f_1(\gamma_j)$ passes over γ_i, i.e. the j^{th} column describes the image of γ_j.

Using the same argument, we obtain that $\lambda(f_n)$ can be computed as the spectral radius

of the matrix

$$M_n = \begin{pmatrix} 0 & 0 & 0 & \cdots & 0 & Id \\ A_n & B_n & 0 & \cdots & 0 & C_n \\ 0 & Id & 0 & \cdots & 0 & 0 \\ 0 & 0 & Id & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & Id & 0 \end{pmatrix},$$

where all blocks are 3×3 matrices and M_n is a $3g \times 3g$ matrix. Here, Id denotes the identity matrix and

$$A_n = \begin{pmatrix} n+1 & 1 & 1 \\ n & 1 & 0 \\ n+1 & 1 & 2 \end{pmatrix}, \; B_n = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \; C_n = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Penner's example is exactly the case where $n = 1$. Computing the g^{th} iterate of the matrix for $g \geqslant 4$ yields

$$M_n^g = \begin{pmatrix} A_n & B_n & 0 & 0 & \cdots & 0 & 0 & 0 & C_n \\ C_nA_n & D_n + C_nB_n & B_nA_n & 0 & \cdots & 0 & 0 & 0 & C_n^2 \\ 0 & C_n & D_n & B_nA_n & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & C_n & D_n & B_nA_n \\ B_nA_n & 0 & 0 & 0 & \cdots & 0 & 0 & C_n & D_n \end{pmatrix},$$

where $D_n := A_n + B_nC_n$. This computation is analogous to [Pen91].

In the case $g = 3$, one obtains

$$M_n^g = \begin{pmatrix} A_n & B_n & C_n \\ C_nA_n & D_n + C_nB_n & B_nA_n + C_n \\ B_nA_n & C_n & D_n \end{pmatrix}.$$

By the Collatz-Wielandt formula, we obtain that the spectral radius of M_n^g is at least as big as the lowest row sum, which in both cases is given by the second row and equals $n+1$. Denoting the spectral radius of M_n by $\rho(M_n)$, we have

$$\lambda(f_n) = \rho(M_n) = \rho(M_n^g)^{\frac{1}{g}} \geqslant (n+1)^{\frac{1}{g}} \geqslant n^{\frac{1}{g}}.$$

□

Lemma 3.7. *For any n, it holds that $l_{\mathcal{C}}(f_n) \leqslant \frac{1}{g-1}$.*

Proof. Consider the curve $\gamma := \rho(b)$. Since it is disjoint from a, b, c, we have $f_n(\gamma) = \rho(\gamma) = \rho^2(b)$. Iteratively, we obtain

$$f_n^{g-1}(\gamma) = \rho^{g-1}(\gamma) = \rho^g(b) = b.$$

Since b is disjoint from γ, we get that $d_{\mathcal{C}}(\gamma, f_n^{g-1}(\gamma)) = 1$, which implies that

$$l_{\mathcal{C}}(f_n) \leqslant \frac{d_{\mathcal{C}}(\gamma, f_n^{g-1}(\gamma))}{g-1} = \frac{1}{g-1}.$$

□

Lemma 3.6 shows that the stretch factors of the f_n tend to infinity, while Lemma 3.7 implies that their stable curve graph translation length is bounded from above by the constant required. This finishes the proof of Theorem 3.1.

Gadre–Tsai show that the stable curve graph translation length of f_1, i.e. Penner's original example, can be bounded from above not just by $\frac{1}{g-1}$ but even by $\frac{4}{g^2+g-4}$ (see [GT11, proof of Theorem 6.1]). This is done by computing more iterations of the curve $\rho(b)$. In particular, Gadre–Tsai argue that an iterate of $\rho(b)$ is contained in a neighbourhood of a union of some of the curves in the set $\{\rho^i(a), \rho^i(b), \rho^i(c) \,|\, i = 1, \dots, g\}$, and as long as this union is not over all the curves of the set, one can find a curve disjoint from both $\rho(b)$ and its iterate. So, the proof boils down to controlling the neighbourhoods for a high enough iterate. Since f_n differs from f_1 only by twisting more often about b, the iterates of $\rho(b)$ under f_n are contained in the same neighbourhoods as the iterates of $\rho(b)$ under f_1. Hence, looking at f_n instead of f_1 does not affect Gadre–Tsai's argument, and we obtain $l_{\mathcal{C}}(f_n) \leqslant \frac{4}{g^2+g-4}$. We leave out the details, since our bound of $\frac{1}{g-1}$ is enough in order to prove Corollaries 3.2 and 3.3.

Corollary (3.2)**.** *There exists a sequence $(h_g)_{g=2}^{\infty}$, where h_g is a pseudo-Anosov of a genus g surface, with*

$$l_{\mathcal{T}}(h_g) \to \infty \text{ and } l_{\mathcal{C}}(h_g) \to 0,$$

as $g \to \infty$.

Proof. For $g = 2$, let h_g be an arbitrary pseudo-Anosov. For $g \geqslant 3$, let $n_g := g^g$ and let $h_g := f_{n_g}$, where f_{n_g} is as constructed at the beginning of Section 3.1 for the corresponding genus g.

From Lemma 3.6 we obtain $l_{\mathcal{T}}(h_g) \geqslant \log(g)$, and from Lemma 3.7 we get that $l_{\mathcal{C}}(h_g) \leqslant \frac{1}{g-1}$. It follows that

$$\lim_{g\to\infty} l_{\mathcal{T}}(h_g) = \infty \text{ and } \lim_{g\to\infty} l_{\mathcal{C}}(h_g) = 0.$$

□

Corollary (3.3)**.** *For any $g \geqslant 3$, there exists $q \in \mathbb{Q}$ such that there are infinitely many non-conjugate pseudo-Anosovs in* $\mathrm{Mod}(S_g)$ *with stable curve graph translation length q.*

Proof. Let $L = \{l_{\mathcal{C}}(f) \,|\, f \in \mathrm{Mod}(S_g) \text{ pseudo-Anosov}\}$ be the set of all stable curve graph translation lengths of pseudo-Anosovs on S_g. Bowditch proves that there is a power m such that every pseudo-Anosov $f \in \mathrm{Mod}(S_g)$ preserves a geodesic in the curve graph after being raised to the power m (see [Bow08, Theorem 1.4]). Hence, the set L is contained in the set $\{\frac{n}{m} \,|\, n \in \mathbb{N}\}$. In particular, this implies that for any constant C, the set $\{x \in L \,|\, x \leqslant C\}$ is finite.

From Theorem 3.1, we know that there is a sequence of pseudo-Anosovs f_n with $\lambda(f_n) \to \infty$ as $n \to \infty$, and furthermore $l_{\mathcal{C}}(f_n) \leqslant C$ for all n, where C can be taken to be any constant above $\frac{1}{g-1}$. After having to possibly pass to a subsequence, we can assume that all the $\lambda(f_n)$ are pairwise different, and so all the f_n are pairwise non-conjugate.

Finally, since the set $\{l_{\mathcal{C}}(f_n) \,|\, n \in \mathbb{N}\}$ is a finite set by the discussion above, there must be a q that is attained by infinitely many of the f_n. □

§ 3.2 | Construction of infinite multiplicity

In this section, we prove Theorem 3.4. The proof relies on results about subsurface projections by Masur–Minsky [MM00]. We start by recalling some of the theory.

Let S_g be a closed, oriented, genus g surface and Y an incompressible, connected subsurface. For a set X, let $\mathcal{P}(X)$ denote the set of finite subsets of X. Masur–Minsky construct subsurface projection maps $\pi_Y : \mathcal{C}(S_g) \to \mathcal{P}(\mathcal{C}(Y))$. The way to think about these maps is the following: Given a curve in $\mathcal{C}(S_g)$, look at its intersection with Y. This intersection is a collection of arcs. Turn these arcs into closed curves by piecing them together with the boundary components of Y. This yields a finite union of curves which we define as the image of π_Y of the curve we started with (see [MM00, Section 2] for details). Note that the image of a curve is the empty set if and only if its intersection with Y is empty.

We write d_Y for the distance in $\mathcal{C}(Y)$. This should not be confused with $d_{\mathcal{C}}$ which as before denotes the distance in $\mathcal{C}(S_g)$. The crucial result of Masur–Minsky for this work is [MM00, Theorem 3.1]. We state it here for completeness.

Theorem ([MM00]). *Let Y be a subsurface of S_g as above. Let G be a geodesic in $\mathcal{C}(S_g)$ such that every vertex of G has non-empty subsurface projection to Y. There exists a constant M depending only on the genus of S_g such that* $\operatorname{diam}(\pi_Y(G)) \leqslant M$.

In other words, the above theorem says that every geodesic in $\mathcal{C}(S_g)$, whose vertices have non-trivial image under π_Y, projects to a uniformly bounded set. We'll make use of this as follows: Assume we are given curves $\beta_1, \beta_2 \in \mathcal{C}(S_g)$ with $d_Y(\pi_Y(\beta_1), \pi_Y(\beta_2)) > M$. Then, any geodesic joining β_1 and β_2 in $\mathcal{C}(S_g)$ has to pass through a curve that doesn't intersect the subsurface Y. In particular, we will apply the above theorem to $Y = S_g \setminus \alpha$, where α is the curve from Theorem 3.4. Here, we make use of the fact that α is non-separating, so that $S_g \setminus \alpha$ is connected. Note that for $Y = S_g \setminus \alpha$, the only curve in S_g with empty subsurface projection to Y is the curve α. Hence, in this case, we can conclude that the geodesic connecting β_1 and β_2 has to pass through α.

From now on, we omit to write the subsurface projection π_Y when measuring the distance in $\mathcal{C}(Y)$, i.e. if β_1, β_2 are in $\mathcal{C}(S_g)$, then we simply write $d_Y(\beta_1, \beta_2)$ instead of $d_Y(\pi_Y(\beta_1), \pi_Y(\beta_2))$.

We recall the setup of Theorem 3.4. Let $f \in \operatorname{Mod}(S_g)$ be a pseudo-Anosov and assume there are curves α_i, $i \in \mathbb{Z}$, such that $f(\alpha_i) = \alpha_{i+1}$ and $d_{\mathcal{C}}(\alpha_i, \alpha_j) = |i - j|\, m$ for all $i, j \in \mathbb{Z}$ and some fixed $m \in \mathbb{N}$. Note that this ensures that $l_{\mathcal{C}}(f) = m$. Furthermore, note that the condition on α in Theorem 3.4 guarantees the existence of the α_i as the iterates of α under f. We use the convention that $\alpha = \alpha_1$.

From now on, let $Y = S_g \setminus \alpha$ and h be a pseudo-Anosov of Y. Since h is pseudo-Anosov, we know that $d_Y(\beta, h^k(\beta)) \to \infty$ as $k \to \infty$ for any curve β in Y. Choose $N_h \in \mathbb{N}$ big enough such that for all $n \geqslant N_h$ we have

$$d_Y(\alpha_2, h^n(\alpha_2)) > 2M + d_Y(\alpha_0, \alpha_2), \tag{3.1}$$

where M denotes the constant from the Masur-Minsky theorem. The reason for this choice will become apparent in the proof of Lemma 3.8 below.

Fix some $n \geqslant N_h$ and define $F = f_{h,n} := h^n \circ f$. By abuse of notation, from now on we simply write h instead of h^n. Hence, we have $F = h \circ f$ and our goal is to compute $l_{\mathcal{C}}(F)$.

In order to compute the stable curve graph translation length of F, we need to iterate F on some curve. We choose the curve α_0. In the following, we use the notation $h_i := f^{i-1} h f^{-(i-1)}$ for all $i \geqslant 1$. In particular, $h_1 = h$. We obtain:

$$F(\alpha_0) = h_1 \circ f(\alpha_0) = \alpha_1,$$

as well as

$$F^2(\alpha_0) = F(\alpha_1) = h_1(\alpha_2).$$

Iterating this process yields

$$F^k(\alpha_0) = h_1 f \dots h_1 f\,(\alpha_0) = h_1 \underbrace{f\, h_1 f^{-1}}_{h_2} \underbrace{f^2 h_1 f^{-2}}_{h_3} f^3 h_1 \dots \underbrace{f^{k-1} h_1 f^{-(k-1)}}_{h_{k-1}} f^k\,(\alpha_0) =$$

$$= h_1 h_2 h_3 \dots h_{k-1}\,(\alpha_k).$$

Hence, to prove Theorem 3.4, we have to compute the distance in the curve graph between α_0 and $h_1 h_2 h_3 \dots h_{k-1}\,(\alpha_k)$. We formulate this as a Lemma 3.8 below. The following Figure 3.2 is a schematic drawing of the vertices of the curve graph we are interested in and should help the reader visualise the distances appearing in the proof of Lemma 3.8. Note that f "acts on" Figure 3.2 by translating horizontally.

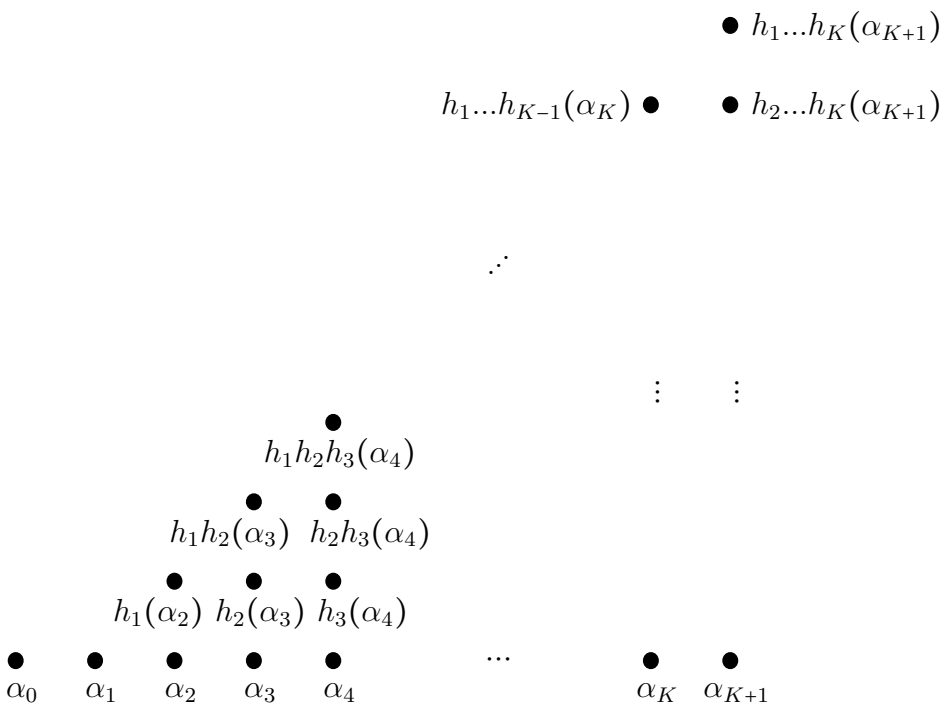


Figure 3.2: Schematic diagram of part of the curve graph.

Lemma 3.8. *It holds that*

$$d_{\mathcal{C}}(\alpha_0,\, h_1 h_2 h_3 \dots h_{k-1}\,(\alpha_k)) = km,$$

for any $k \in \mathbb{N}$.

Proof. We prove this statement by induction over k. The base case $k = 1$ holds since $d_{\mathcal{C}}(\alpha_0, \alpha_1) = m$ by assumption. We also compute the case $k = 2$ explicitly since it provides an insight into how the induction step works. For this case, we want to calculate $d_{\mathcal{C}}(\alpha_0, h_1(\alpha_2))$. Since $d_{\mathcal{C}}(\alpha_0, \alpha_1) = m$ and

$$d_{\mathcal{C}}(\alpha_1, h_1(\alpha_2)) = d_{\mathcal{C}}(h_1(\alpha_1), h_1(\alpha_2)) = d_{\mathcal{C}}(\alpha_1, \alpha_2) = m,$$

we obtain that $d_{\mathcal{C}}(\alpha_0, h_1(\alpha_2)) \leqslant 2m$. Furthermore, using equation (3.1) we obtain

$$d_Y(\alpha_0, h_1(\alpha_2)) \geqslant d_Y(\alpha_2, h_1(\alpha_2)) - d_Y(\alpha_0, \alpha_2) > 2M > M.$$

It follows from the Masur–Minsky theorem that any geodesic between α_0 and $h_1(\alpha_2)$ in $\mathcal{C}(S_g)$ has to pass through the curve α_1. We obtain:

$$d_{\mathcal{C}}(\alpha_0, h_1(\alpha_2)) \geqslant d_{\mathcal{C}}(\alpha_0, \alpha_1) + d_{\mathcal{C}}(\alpha_1, h_1(\alpha_2)) = m + m = 2m.$$

Assume now that for all natural numbers k up to some $K \in \mathbb{N}$, the claim

$$d_{\mathcal{C}}(\alpha_0, h_1 h_2 h_3 \dots h_{k-1}(\alpha_k)) = km$$

holds. We have to show the analogous claim for $K + 1$.

Note that $f^{-1}(h_2 \dots h_K(\alpha_{K+1})) = h_1 \dots h_{K-1}(\alpha_K)$. It follows that

$$d_{\mathcal{C}}(\alpha_1, h_1 \dots h_K(\alpha_{K+1})) = d_{\mathcal{C}}(\alpha_1, h_2 \dots h_K(\alpha_{K+1})) = d_{\mathcal{C}}(\alpha_0, h_1 \dots h_{K-1}(\alpha_K)) = Km,$$

and hence

$$d_{\mathcal{C}}(\alpha_0, h_1 \dots h_K(\alpha_{K+1})) \leqslant d_{\mathcal{C}}(\alpha_0, \alpha_1) + d_{\mathcal{C}}(\alpha_1, h_1 \dots h_K(\alpha_{K+1})) \leqslant m + Km = (K+1)m.$$

For the other inequality, we want to use the subsurface projection to Y. Our goal is to show that the distance of the projections of α_0 and $h_1 \dots h_K(\alpha_{K+1})$ is sufficiently big.

First, note that

$$d_{\mathcal{C}}(h_1(\alpha_2),\, h_1 \dots h_K(\alpha_{K+1})) = d_{\mathcal{C}}(\alpha_2,\, h_2 \dots h_K(\alpha_{K+1})),$$

and by applying f^{-1} to both sides the above is equal to

$$d_{\mathcal{C}}(\alpha_1,\, h_1 \dots h_{K-1}(\alpha_K)) = d_{\mathcal{C}}(\alpha_1,\, h_2 \dots h_{K-1}(\alpha_K)),$$

which by applying f^{-1} again is equal to

$$d_{\mathcal{C}}(\alpha_0,\, h_1 \dots h_{K-2}(\alpha_{K-1})),$$

which by the induction hypothesis equals $(K-1)m$. From the calculations above, we also know that

$$d_{\mathcal{C}}(\alpha_1,\, h_1 \dots h_K(\alpha_{K+1}))$$

equals Km.

We now want to argue that a geodesic between $h_1(\alpha_2)$ and $h_1 \dots h_K(\alpha_{K+1})$ doesn't pass through α_1. So, assume such a geodesic did pass through α_1. Then:

$$d_{\mathcal{C}}(h_1(\alpha_2),\, h_1 \dots h_K(\alpha_{K+1})) \geqslant d_{\mathcal{C}}(h_1(\alpha_2), \alpha_1) + d_{\mathcal{C}}(\alpha_1, h_1 \dots h_K(\alpha_{K+1})) =$$

$$= m + Km = (K+1)m.$$

But then $(K-1)m = d_{\mathcal{C}}(h_1(\alpha_2),\, h_1 \dots h_K(\alpha_{K+1})) \geqslant (K+1)m$ which is a contradiction. Hence, no geodesic between $h_1(\alpha_2)$ and $h_1 \dots h_K(\alpha_{K+1})$ passes through α_1. In particular, this means that any such geodesic consists of vertices with non-empty subsurface projection image to Y. So, we can use the result of Masur–Minsky to conclude that a projection to $\mathcal{C}(Y)$ of such a geodesic has an image whose diameter is bounded by M. In particular, we obtain

$$d_Y(h_1(\alpha_2),\, h_1 \dots h_K(\alpha_{K+1})) \leqslant M.$$

Using equation (3.1) again, it follows that

$$d_Y(\alpha_2, h_1 \dots h_K(\alpha_{K+1})) \geqslant d_Y(\alpha_2, h_1(\alpha_2)) - d_Y(h_1 \dots h_K(\alpha_{K+1}),\, h_1(\alpha_2)) >$$

$$> 2M + d_Y(\alpha_0, \alpha_2) - M = M + d_Y(\alpha_0, \alpha_2),$$

and finally,

$$d_Y(\alpha_0, h_1 \dots h_K(\alpha_{K+1})) \geqslant d_Y(\alpha_2, h_1 \dots h_K(\alpha_{K+1})) - d_Y(\alpha_0, \alpha_2) >$$

$$> M + d_Y(\alpha_0, \alpha_2) - d_Y(\alpha_0, \alpha_2) = M.$$

This inequality, together with the Masur-Minsky theorem, shows that any geodesic from α_0 to $h_1 \dots h_K(\alpha_{K+1})$ has to pass through α_1. For the distance of the two points, this implies:

$$d_{\mathcal{C}}(\alpha_0, h_1 \dots h_K(\alpha_{K+1})) \geqslant d_{\mathcal{C}}(\alpha_0, \alpha_1) + d_{\mathcal{C}}(\alpha_1, h_1 \dots h_K(\alpha_{K+1})) =$$

$$= m + Km = (K+1)m.$$

□

We are now ready to prove Theorem 3.4:

Proof of Theorem 3.4. It only remains to show that $F = h \circ f$ satisfies $l_{\mathcal{C}}(F) = l_{\mathcal{C}}(f) = m$. By definition of the stable curve graph translation length, we have

$$l_{\mathcal{C}}(F) = \liminf_{k\to\infty} \frac{d_{\mathcal{C}}(\alpha_0, F^k(\alpha_0))}{k}.$$

Above, we showed that

$$F^k(\alpha_0) = h_1 \dots h_{k-1}(\alpha_k),$$

and in Lemma 3.8 we computed that

$$d_{\mathcal{C}}(\alpha_0, h_1 \dots h_{k-1}(\alpha_k)) = km.$$

Hence, we conclude

$$l_{\mathcal{C}}(F) = \liminf_{k\to\infty} \frac{km}{k} = m.$$

□

Remark 3.9. For a pseudo-Anosov f that satisfies the condition of Theorem 3.4, we now have a way of constructing infinitely many pseudo-Anosovs different from f with the same stable curve graph translation length. Note that we can even change the curve $\alpha = \alpha_1$ and use any other curve of the α_k in the construction. More precisely, for any curve α_k and any pseudo-Anosov $h \in \mathrm{Mod}(S_g \setminus \alpha_k)$, there are infinitely many pseudo-Anosovs of the form $h^n \circ f$ for big enough n that satisfy $l_{\mathcal{C}}(h^n \circ f) = l_{\mathcal{C}}(f)$.

The remark specifies how to obtain pseudo-Anosovs not equal to f with the same stable curve graph translation length. However, this would not be as interesting if all the newly obtained pseudo-Anosovs were conjugate to each other. In the following lemma, we ensure that this does not happen when we add an extra assumption to f.

Lemma 3.10. *Let $g \geqslant 3$. Assume that there is a symplectic basis of homology of the form $\alpha, \beta_1, \dots, \beta_{2g-1}$, where α is any of the α_k, such that the image under the symplectic representation of f with respect to this basis is a non-negative matrix M. Then, there is a pseudo-Anosov h of $S_g \setminus \alpha$ such that the sequence $h^n \circ f$ contains infinitely many non-conjugate mapping classes.*

Proof. Let $\alpha, \beta_1, \dots, \beta_{2g-1}$ be a symplectic basis of $H_1(S_g)$, where α and β_1 intersect exactly once. Let M be the matrix defined by the action of f on homology with respect to the above basis, and assume that M is non-negative.

Let $S_g \setminus \alpha$ be the surface with two boundary components that results from cutting S_g at α. Let S' be the closed surface of genus $g-1$ that results from collapsing the two boundary components of $S_g \setminus \alpha$ to points p_1, p_2 respectively. Note that $\beta_1, \ldots, \beta_{2g-1}$ is a basis of $H_1(S', \{p_1, p_2\})$, where β_1 is now an arc connecting the two points p_1 and p_2, and $\beta_2, \ldots, \beta_{2g-1}$ is a basis of $H_1(S')$. A homeomorphism fixing the boundary of $S_g \setminus \alpha$ is pseudo-Anosov if and only if the corresponding homeomorphism of $S' \setminus \{p_1, p_2\}$ is pseudo-Anosov.

Let $D \subset S'$ be a small disc around β_1 that does not intersect any of the other β_i. Let h' be a homeomorphism of S' that fixes D and acts on $H_1(S')$ with respect to the basis $\beta_2, \ldots, \beta_{2g-1}$ by a primitive, symplectic matrix B. Here, by primitive matrix, we mean a non-negative matrix that admits a power whose entries are all positive. Such h' exists, as the symplectic representation is surjective and we can always isotope a homeomorphism to fix a given disc.

The homeomorphism h' fixes in particular the points p_1, p_2. The induced action of h' on $H_1(S', \{p_1, p_2\})$ with respect to the basis $\beta_1, \ldots, \beta_{2g-1}$ is then given by the matrix

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

where 1 is a single entry, and 0 stands for the zero row and column vector with $2g-2$ entries, respectively.

Let h_{pA} be a pseudo-Anosov of $S' \setminus \{p_1, p_2\}$ which acts trivially on $H_1(S', \{p_1, p_2\})$ and such that h' does not map the set of the two invariant foliations of h_{pA} to itself. For example, one can construct such h_{pA} as a point-pushing map along a curve based at p_1 that fills $S' \setminus \{p_2\}$ and is null-homologous in $H_1(S')$. Then h_{pA} acts trivially on the required homology group and Kra's construction implies that h_{pA} is pseudo-Anosov [FM12, Theorem 14.6]. The condition on the foliations of h_{pA} can then be achieved by appropriate conjugation.

A standard north-south dynamics argument now shows that composing h' with a sufficiently high power of h_{pA} yields a pseudo-Anosov (see, for example, [Iva92, Exercise 3(a), page 102]). Call this pseudo-Anosov h. Note that since h_{pA} acts by the identity on $H_1(S', \{p_1, p_2\})$, the pseudo-Anosov h acts in the same way as h'.

We can also view h as a pseudo-Anosov of $S_g \setminus \alpha$, and, after gluing the two boundary curves of $S_g \setminus \alpha$ back together, as a (reducible) mapping class of S_g. Seen as a mapping class of S_g, h acts on $H_1(S_g)$ with respect to the basis $\alpha, \beta_1, \ldots, \beta_{2g-1}$ by the matrix

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

where the 1's are single entries and the 0's are either also single entries or stand for the zero row or column vectors of dimension $2g-2$. Note that the 0 in the second entry of the first row describes the number of twists used when gluing the boundaries of $S_g \setminus \alpha$ back together. Hence, in this entry we could achieve any $k \in \mathbb{Z}$, but for convenience we choose to glue with no twisting.

Computing powers of the above matrix shows that $h^n \circ f$ acts as

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & B^n \end{pmatrix} M.$$

Write

$$M = \begin{pmatrix} m_{11} & m_{12} & m_{r_1} \\ m_{21} & m_{22} & m_{r_2} \\ m_{c_1} & m_{c_2} & M' \end{pmatrix},$$

where m_{r_1}, m_{r_2} and m_{c_1}, m_{c_2} are row and column vectors of dimension $2g-2$, respectively, and M' is a $(2g-2) \times (2g-2)$ matrix.

Then the trace of the matrix describing the homology action of $h^n \circ f$ is

$$\operatorname{tr}\left(\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & B^n \end{pmatrix} \begin{pmatrix} m_{11} & m_{12} & m_{r_1} \\ m_{21} & m_{22} & m_{r_2} \\ m_{c_1} & m_{c_2} & M' \end{pmatrix} \right) = m_{11} + m_{22} + \operatorname{tr}(B^n M').$$

Since M is non-negative, M' is non-negative. As $g \geqslant 3$ and M' is a $(2g-2) \times (2g-2)$ block of an invertible matrix, M' is non-zero. Because B is primitive, the entries of B^n tend to infinity as n gets larger. We conclude that

$$\operatorname{tr}(B^n M') \xrightarrow{n \to \infty} \infty,$$

and consequently the trace of the matrices describing the action of $h^n \circ f$ on $H_1(S_g)$ tends to infinity as $n \to \infty$. Since conjugate mapping classes have the same trace when acting on homology, this shows that there is an infinite subsequence n_k such that the pseudo-Anosovs $h^{n_k} \circ f$ are pairwise non-conjugate. □

The above lemma applies, for example, to pseudo-Anosovs f in the Torelli group. We believe that the construction in Theorem 3.4 can be used to construct infinitely many non-conjugate pseudo-Anosovs even without the homology assumption of the lemma. Note that Remark 3.9 says that we can use any triple (α_k, h, n), where h is a pseudo-Anosov of $S_g \setminus \alpha_k$ and n is a sufficiently high power, to build the suitable pseudo-Anosovs $h^n \circ f$. In the lemma, by adding the homology assumption on f, we found infinitely many non-conjugate pseudo-Anosovs only by varying n. Hence, by varying α_k and h as well, one might be able to find infinitely many non-conjugate pseudo-Anosovs without having to impose the extra homology assumption. A thorough study of the conjugacy classes of the maps $h^n \circ f$ that depend on the triple (α_k, h, n) would be interesting.

Chapter 4

Commutators of n-cycles

In this chapter, we show a purely group-theoretic result about commutators of permutations. This result will later be used in Chapter 5 to prove statements about filling pairs of curves, square-tiled surfaces, and pseudo-Anosov stretch factors. We begin by introducing some background together with the main notions.

The commutator subgroup of a group is defined as the subgroup generated by all commutators. A product of two commutators need not be a commutator itself. Hence, it makes sense to ask whether the commutator subgroup of a given group consists entirely of commutators.

In 1951, Ore proved that the above holds for the finite symmetric groups by showing that every element of the alternating group can be written as a commutator in the symmetric group [Ore51]. Furthermore, Ore proceeded to prove that for $n \geqslant 5$, every element in the alternating group A_n is already a commutator in the alternating group itself. This led to the so-called Ore conjecture, which states that every element in a finite, non-abelian, simple group is a commutator. Over the years, the conjecture was established for more simple groups and finally completely proven in 2010 by Liebeck–O'Brien–Shalev–Tiep [LOST09].

In this work, we reinforce Ore's original result by showing that for $n \geqslant 6$, every element in the alternating group A_n can be written as a commutator of two n-cycles. In the case where n is odd, n-cycles are elements of A_n themselves.

We remark that our result is very closely related to a paper by Vavpetič, who also focuses on the problem of showing that elements in A_n can be written as a product of two cycles [Vav16]. However, since in that paper the length of the cycles is not specified, one would need to do some further analysis to derive our result from [Vav16]. Below, we present a proof of our theorem that is independent of the work by Vavpetič.

We start by setting up some notation:

Let S_n be the symmetric group on n elements. For $\pi \in S_n$ we say that π has cycle type $(k_1, \dots, k_m)$ if π consists of m cycles that have length $k_1, \dots, k_m$ respectively. Note that $k_1 + \dots + k_m = n$. The conjugacy classes of S_n are determined by cycle type, i.e. two permutations $\pi, \tau \in S_n$ are conjugate if and only if they have the same cycle type. For $k_1, \dots, k_m$ with $k_1 + \dots + k_m = n$, we denote the corresponding conjugacy class by $C(k_1, \dots, k_m) \subset S_n$. We refer to elements in $C(n)$ as *n-cycles*.

Denote by $[\tau, \pi] := \tau^{-1}\pi^{-1}\tau\pi$ the commutator of $\tau, \pi \in S_n$. Every commutator is an even permutation and hence an element of A_n. It is a classical result that the commutator subgroup of S_n, i.e. the group generated by all commutators, is A_n for all $n \geqslant 2$.

Here, for $n \geqslant 6$ we prove the stronger result that every element of A_n is a commutator of two n-cycles.

Theorem 4.1. *The map*

$$C(n) \times C(n) \to A_n, \ (\tau, \pi) \mapsto [\tau, \pi]$$

is surjective for all $n \geqslant 6$.

Remark 4.2. For $n = 2$, Theorem 4.1 is also true since A_2 is the trivial group. For $n \in \{3, 4, 5\}$ however, it fails. For $n = 3$, we have that $C(3) = \{\sigma, \sigma^{-1}\}$ with $\sigma = (1\,2\,3)$ and so any commutator of 3-cycles is the identity. However, $\sigma, \sigma^{-1} \in A_3$. For $n = 4$ and $n = 5$, one can check that $C(2, 2) \subset A_4$ and $C(1, 2, 2) \subset A_5$ respectively are not in the image.

§ 4.1 | Proof of Theorem 4.1

Let $\sigma_n := (1\ 2\ \dots\ n) \in C(n)$. We refer to σ_n as the *shift permutation*. Let $k_1, \dots, k_m \in \mathbb{N}$ with $k_1 + \dots + k_m = n$. The conjugacy class $C(k_1, \dots, k_m)$ is a subset of A_n if and only if the tuple $(k_1, \dots, k_m)$ contains an even number of even numbers. We refer to such conjugacy classes as *even conjugacy classes*.

We denote the commutator map from Theorem 4.1 by f. Note that it suffices to show that one element of each even conjugacy class is in the image of f. For if $\rho \in \text{image}(f)$ and ρ' is any other element in the same conjugacy class as ρ, then there exist $\tau, \pi \in C(n)$ and $\phi \in S_n$ with

$$f(\tau, \pi) = \rho \text{ and } \rho' = \phi\rho\phi^{-1},$$

and hence we have $\tau' := \phi\tau\phi^{-1}$, $\pi' := \phi\pi\phi^{-1} \in C(n)$ with

$$f(\tau', \pi') = [\phi\tau\phi^{-1}, \phi\pi\phi^{-1}] = \phi[\tau, \pi]\phi^{-1} = \phi\rho\phi^{-1} = \rho',$$

and so $\rho' \in \text{image}(f)$.

The above observation shows that it suffices to find examples of permutations $\tau, \pi \in C(n)$ whose commutator is contained in $C(k_1, \dots, k_m)$ for every even conjugacy class $C(k_1, \dots, k_m)$. If $[\tau, \pi] \in C(k_1, \dots, k_m)$, then by simultaneously conjugating τ and π we can assume that $\tau = \sigma_n$, where σ_n is the shift permutation. Hence, in order to prove the main theorem, we want to prove the following reformulation of it:

Claim: For every $n \geqslant 6$ and every even conjugacy class $C(k_1, \dots, k_m) \subset S_n$, there exists a $\pi \in C(n)$ with $[\sigma_n, \pi] \in C(k_1, \dots, k_m)$.

To prove the claim, we construct explicit examples in certain conjugacy classes, which we will call *irreducible* and show that it suffices to construct these examples. The latter follows from the subsequent lemmata.

In the proofs, for two permutations α and β on different elements, we use the notation $(\alpha)(\beta)$ to indicate the permutation that does both at the same time. For example, if $\alpha = (1\,2\,3)$ and $\beta = (4\,5)(6)$, then $(\alpha)(\beta) = (1\,2\,3)(4\,5)(6)$.

Lemma 4.3. *Let $n \geqslant 2$ and $\pi \in C(n)$ with $\pi(1) = 2$ such that $[\sigma_n, \pi] \in C(k_1, \dots, k_m)$. Then there exists $\pi_+ \in C(n+1)$ with $\pi_+(1) = 2$ such that*

$$[\sigma_{n+1}, \pi_+] \in C(k_1, \dots, k_m, 1).$$

Proof. If $n = 2$, then we necessarily have that $\pi = (1\ 2)$, and $[\sigma_2, \pi] \in C(1, 1)$. Defining $\pi_+ := (1\ 2\ 3)$, gives $[\sigma_3, \pi_+] \in C(1, 1, 1)$.

Now, let $n \geqslant 3$. Since π is an n-cycle with $\pi(1) = 2$, we can write $\pi = (1\ 2\ i\ \dots\ j)$ where $i, j \in \{3, \dots, n\}$ and between i and j we could have any order of the remaining numbers.

Define $\pi_+ := (1\ \ 2\ \ i\ \ \dots\ \ j\ \ n+1) \in C(n+1)$ where between i and j, π_+ is the same as π.

Let $\alpha := [\sigma_n, \pi] \in C(k_1, \dots, k_m)$. We claim that $[\sigma_{n+1}, \pi_+] = (\alpha)(n+1) \in C(k_1, \dots, k_m, 1)$. To see this, first note that

$$\sigma_{n+1}^{-1}\pi_+^{-1}\sigma_{n+1}\pi_+(n+1) = \sigma_{n+1}^{-1}\pi_+^{-1}\sigma_{n+1}(1) = \sigma_{n+1}^{-1}\pi_+^{-1}(2) = \sigma_{n+1}^{-1}(1) = n+1.$$

For any $k \in \{1, \dots, n\}$, we need to check that $[\sigma_n, \pi](k) = [\sigma_{n+1}, \pi_+](k)$. Since π and π_+,

as well as σ_n and σ_{n+1}, only differ from each other by a single element that was added to the respective cycle, the above immediately follows for all k apart from two cases, namely $k = j$ and k such that $\pi(k) = n$.

For these remaining two cases, we carry out the computations. For $k = j$ we obtain:

$$[\sigma_{n+1}, \pi_+](k) = \sigma_{n+1}^{-1}\pi_+^{-1}\sigma_{n+1}(n+1) = \sigma_{n+1}^{-1}\pi_+^{-1}(1) = \sigma_{n+1}^{-1}(n+1) = n,$$

$$[\sigma_n, \pi](k) = \sigma_n^{-1}\pi^{-1}\sigma_n(1) = \sigma_n^{-1}\pi^{-1}(2) = \sigma_n^{-1}(1) = n.$$

For k with $\pi(k) = n$ we obtain:

$$[\sigma_{n+1}, \pi_+](k) = \sigma_{n+1}^{-1}\pi_+^{-1}\sigma_{n+1}(n) = \sigma_{n+1}^{-1}\pi_+^{-1}(n+1) = \sigma_{n+1}^{-1}(j) = j-1,$$

$$[\sigma_n, \pi](k) = \sigma_n^{-1}\pi^{-1}\sigma_n(n) = \sigma_n^{-1}\pi^{-1}(1) = \sigma_n^{-1}(j) = j-1.$$

□

Lemma 4.4. *Let $n_1, n_2 \geqslant 2$ and $n := n_1 + n_2$. Let $\pi_1 \in C(n_1)$, $\pi_2 \in C(n_2)$ with $\pi_1(1) = \pi_2(1) = 2$ such that $[\sigma_{n_1}, \pi_1] \in C(k_1, \ldots, k_m)$ and $[\sigma_{n_2}, \pi_2] \in C(l_1, \ldots, l_s)$. Then there exists $\pi \in C(n)$ with $\pi(1) = 2$ such that*

$$[\sigma_n, \pi] \in C(k_1, \ldots, k_m, l_1, \ldots, l_s).$$

Proof. Analogously to the proof of the previous lemma, we write $\pi_1 = (1\ 2\ i_1\ \ldots\ j_1)$, $\pi_2 = (1\ 2\ i_2\ \ldots\ j_2)$, where $i_1, j_1 \in \{3, \ldots, n_1\}$ and $i_2, j_2 \in \{3, \ldots, n_2\}$. In the case where n_1 or n_2 is equal to 2, the corresponding set $\{n_1, \ldots, 3\}$ or $\{n_2, \ldots, 3\}$ is understood to be the empty set, and the corresponding permutation is $(1\ 2)$.

Let $\alpha := [\sigma_{n_1}, \pi_1] \in C(k_1, \ldots, k_m)$, $\beta := [\sigma_{n_2}, \pi_2] \in C(l_1, \ldots, l_s)$ and define

$$\pi := (1\ \ 2\ \ i_1\ \ \ldots\ \ j_1\ \ n_1+1\ \ n_1+2\ \ n_1+i_2\ \ \ldots\ \ n_1+j_2).$$

Here, between i_1 and j_1, π is the same as π_1 and between n_1+i_2 and n_1+j_2, π is the same as π_2 after adding n_1 to each element. By definition $\pi \in C(n)$, and we want to show that $[\sigma_n, \pi] = (\alpha)(\beta + n_1) \in C(k_1, \ldots, k_m, l_1, \ldots, l_s)$, where the notation $\beta + n_1$ means to add n_1 to each element of β.

In particular, we want to show that $[\sigma_n, \pi](k) = [\sigma_{n_1}, \pi_1](k)$ for every $k \in \{1, \ldots, n_1\}$ and $[\sigma_n, \pi](l + n_1) = [\sigma_{n_2}, \pi_2](l) + n_1$ for every $l \in \{1, \ldots, n_2\}$. By the definition of π and σ_n, this is straightforward for all k and l apart from two cases respectively, namely $k = j_1$

and k with $\pi_1(k) = n_1$, as well as $l = j_2$ and l with $\pi_2(l) = n_2$. For these cases, we carry out the computations:

For $k = j_1$ we obtain:

$$[\sigma_n, \pi](k) = \sigma_n^{-1}\pi^{-1}\sigma_n(n_1 + 1) = \sigma_n^{-1}\pi^{-1}(n_1 + 2) = \sigma_n^{-1}(n_1 + 1) = n_1,$$

$$[\sigma_{n_1}, \pi_1](k) = \sigma_{n_1}^{-1}\pi_1^{-1}\sigma_{n_1}(1) = \sigma_{n_1}^{-1}\pi_1^{-1}(2) = \sigma_{n_1}^{-1}(1) = n_1.$$

For k with $\pi_1(k) = n_1$ we obtain:

$$[\sigma_n, \pi](k) = \sigma_n^{-1}\pi^{-1}\sigma_n(n_1) = \sigma_n^{-1}\pi^{-1}(n_1 + 1) = \sigma_n^{-1}(j_1) = j_1 - 1,$$

$$[\sigma_{n_1}, \pi_1](k) = \sigma_{n_1}^{-1}\pi_1^{-1}\sigma_{n_1}(n_1) = \sigma_{n_1}^{-1}\pi_1^{-1}(1) = \sigma_{n_1}^{-1}(j_1) = j_1 - 1.$$

For $l = j_2$ we obtain:

$$[\sigma_n, \pi](l + n_1) = \sigma_n^{-1}\pi^{-1}\sigma_n(1) = \sigma_n^{-1}\pi^{-1}(2) = \sigma_n^{-1}(1) = n = n_2 + n_1,$$

$$[\sigma_{n_2}, \pi_2](l) = \sigma_{n_2}^{-1}\pi_2^{-1}\sigma_{n_2}(1) = \sigma_{n_2}^{-1}\pi_2^{-1}(2) = \sigma_{n_2}^{-1}(1) = n_2.$$

For l with $\pi_2(l) = n_2$ we obtain:

$$[\sigma_n, \pi](l + n_1) = \sigma_n^{-1}\pi^{-1}\sigma_n(n) = \sigma_n^{-1}\pi^{-1}(1) = \sigma_n^{-1}(j_2 + n_1) = j_2 - 1 + n_1,$$

$$[\sigma_{n_2}, \pi_2](l) = \sigma_{n_2}^{-1}\pi_2^{-1}\sigma_{n_2}(n_2) = \sigma_{n_1}^{-1}\pi_1^{-1}(1) = \sigma_{n_1}^{-1}(j_2) = j_2 - 1.$$

□

Regarding the above lemmata, given an even conjugacy class $C(k_1, \ldots, k_m)$, we want to split the tuple $(k_1, \ldots, k_m)$ into smaller pieces for which we can construct the desired permutations for the corresponding conjugacy classes. A straightforward idea would be to split it into odd singletons and pairs of even numbers, construct permutations for each corresponding conjugacy class and use Lemma 4.4 to obtain a permutation for the whole conjugacy class. More precisely, this means that we would only need to construct permutations π whose commutator with σ_n is contained in $C(n)$ for n odd or $C(2k, 2l)$ for $k, l \geqslant 1$. However, this does not work, since for $C(3)$ or $C(2, 2)$ there does not exist such a desired permutation. Hence, we need to consider more conjugacy classes. The following list will be referred to as *irreducible conjugacy classes*:

$C(n)$	$n \geqslant 5$ odd
$C(2k, 2l)$	$k, l \geqslant 1$, except $k = l = 1$
$C(n, 3)$	$n \geqslant 1$ odd
$C(2k, 2l, 3)$	$k, l \geqslant 1$
$C(n, 2, 2)$	$n \geqslant 3$ odd
$C(2k, 2, 2, 2)$	$k \geqslant 1$

Lemma 4.5. *Let $n \geqslant 6$ and let $C(k_1, \ldots, k_m) \subset S_n$ be an even conjugacy class. Assume that for every irreducible conjugacy class C there exists a π_C with $\pi_C(1) = 2$ such that its commutator with the corresponding shift permutation lies in C. Then there exists $\pi \in C(n)$ with $[\sigma_n, \pi] \in C(k_1, \ldots, k_m)$.*

Proof. In this proof whenever we say *a permutation for the conjugacy class $C(l_1, \ldots, l_s)$*, we mean an m-cycle τ, where $m := l_1 + \ldots + l_s$, with the properties $\tau(1) = 2$ and $[\sigma_m, \tau] \in C(l_1, \ldots, l_s)$. Furthermore, we say that $C(l_1, \ldots, l_s)$ *contains a number l* if $l_j = l$ for some $j = 1, \ldots, s$.

Our goal is to show that there is a π for $C(k_1, \ldots, k_m)$. If $C(k_1, \ldots, k_m)$ does not contain the numbers $1, 2$ and 3, then by the assumption we can find permutations for the irreducible conjugacy classes $C(k_i)$ with k_i odd and $C(k_j, k_{j'})$ with $k_j, k_{j'}$ even. Then, by Lemma 4.4 we can also find a π for $C(k_1, \ldots, k_m)$.

So, we need to consider the cases where $C(k_1, \ldots, k_m)$ contains $1, 2$ or 3. First, assume that $C(k_1, \ldots, k_m)$ contains 1, but does not contain 2 or 3. Let $k_i = 1$ for some i. Then, if we can find a permutation for $C(k_1, \ldots, k_{i-1}, k_{i+1}, \ldots, k_m)$, by Lemma 4.3 we can also find one for $C(k_1, \ldots, k_m)$. Inductively, by this process we can remove all numbers equal to 1 and be left with the previous case, so again we can find a π for $C(k_1, \ldots, k_m)$.

Now, assume that $C(k_1, \ldots, k_m)$ contains 3, but does not contain 1 or 2. Consider the permutations (1 2 4 3 6 5) and (1 2 4 9 5 7 8 6 3). Their commutator with the corresponding shift permutations is contained in $C(3, 3)$ and $C(3, 3, 3)$ respectively. So, by Lemma 4.4, we can find permutations for any conjugacy class of the form $C(3, \ldots, 3)$, that is, for any conjugacy class containing only and at least two 3's. Therefore, if $C(k_1, \ldots, k_m)$ contains more than one 3, we can find a permutation for the conjugacy class $C(3, \ldots, 3)$ with the number of 3's that $C(k_1, \ldots, k_m)$ contains and, using Lemma 4.4, reduce to the case of having no 3's. If $C(k_1, \ldots, k_m)$ contains exactly one 3, then it contains at least either one more odd number k_i or a pair of even numbers $k_j, k_{j'}$. From the assumption, we can build a permutation for the irreducible conjugacy classes $C(k_i, 3)$ or $C(k_j, k_{j'}, 3)$ and, by

Lemma 4.4, reduce to the case of no 3's again.

Now, assume that $C(k_1, \ldots, k_m)$ contains 2, but does not contain 1 or 3. Consider the permutations (1 2 5 4 7 8 6 3) and (1 2 8 12 9 6 7 3 5 10 11 4). Their commutator with the corresponding shift permutations is contained in $C(2,2,2,2)$ and $C(2,2,2,2,2,2)$ respectively. So, by Lemma 4.4, we can find permutations for any conjugacy class of the form $C(2, \ldots, 2)$, that is, for any conjugacy class containing only and at least two pairs of 2's. If $C(k_1, \ldots, k_m)$ contains an even number of 2's which is greater or equal to four, we can find a permutation for the conjugacy class $C(2, \ldots, 2)$ with the number of 2's that $C(k_1, \ldots, k_m)$ contains and using Lemma 4.4 reduce to the case of having no 2's. If $C(k_1, \ldots, k_m)$ contains an odd number of 2's which is greater or equal to five, then there is at least one more even number k_j, and we can find permutations for the conjugacy class $C(2, \ldots, 2)$, where the number of 2's is one less than the number of 2's contained in $C(k_1, \ldots, k_m)$, and for $C(2, k_j)$ since it is irreducible. Again, using Lemma 4.4, we can reduce to the case of having no 2's. If $C(k_1, \ldots, k_m)$ contains exactly one or exactly three 2's, then there is at least one other even number k_j, and by the assumption we can find permutations for $C(2, k_j)$ or $C(2,2,2,k_j)$ and hence reduce to having no 2's. If $C(k_1, \ldots, k_m)$ contains exactly two 2's, then there is either another odd number k_i or another pair of even numbers $k_j, k_{j'}$. Note that $k_j \neq 2 \neq k_{j'}$. By the assumption, we can find permutations for the irreducible classes $C(k_i, 2, 2)$ or $C(k_j, 2)$ and $C(k_{j'}, 2)$. We conclude that also in this case we can reduce to having no 2's.

Now, assume that $C(k_1, \ldots, k_m)$ contains 1 and 3, but not 2. Consider the permutation (1 2 4 3). Its commutator with σ_4 lies in $C(3,1)$. Let k be the minimum of the number of 1's and the number of 3's contained in $C(k_1, \ldots, k_m)$. By Lemma 4.4, we can find a permutation in the conjugacy class $C(3, \ldots, 3, 1, \ldots, 1)$ with exactly k 3's and k 1's and reduce to the case of having either only 1's or only 3's left.

Now, assume that $C(k_1, \ldots, k_m)$ contains 1 and 2, but not 3. Consider the permutation (1 2 6 4 5 3). Its commutator with σ_6 lies in $C(2,2,1,1)$. If $C(k_1, \ldots, k_m)$ contains numbers not equal to 1 or 2, by the previous cases we can find a permutation for the conjugacy class containing those numbers and all the 2's. Then by Lemma 4.3, we can also find a permutation for $C(k_1, \ldots, k_m)$. Otherwise, $C(k_1, \ldots, k_m)$ consists only of 1's and 2's. If there are more than two 2's, we already saw that we can find a permutation for $C(2, \ldots, 2)$ and then add the 1's by Lemma 4.3. If there are exactly two 2's, then there have to be at least two 1's since $n \geqslant 6$. In that case, we use the permutation for $C(2,2,1,1)$ constructed above and add the remaining 1's with Lemma 4.3.

Now, assume that $C(k_1, \ldots, k_m)$ contains 2 and 3, but not 1.

Consider the permutations (1 2 4 6 7 5 3), (1 2 4 6 7 9 10 5 8 3) and (1 2 7 5 3 11 10 4 6 9 8). The commutator of these three permutations with the corresponding shift permutations is contained in $C(3,2,2)$, $C(3,3,2,2)$ and $C(3,2,2,2,2)$ respectively. Using these permutations, together with the ones for the conjugacy classes $C(2,\dots,2)$, $C(3,\dots,3)$ and Lemma 4.4, we can construct permutations for every conjugacy class of the form $C(2,\dots,2,3,\dots,3)$ for any even number of 2's and any number of 3's. Hence, if $C(k_1,\dots,k_m)$ contains an even number of 2's, we build the permutation for the corresponding conjugacy class $C(2,\dots,2,3,\dots,3)$ and by Lemma 4.4 are able to reduce to the case without 2's or 3's. If $C(k_1,\dots,k_m)$ contains an odd number of 2's, then there is at least one more even number $k_j \neq 2$, so we use permutations for $C(2,\dots,2,3,\dots,3)$, where the number of 2's is one less than the number of 2's contained in $C(k_1,\dots,k_m)$, and for $C(2,k_j)$ in order to reduce to the case without 2's and 3's.

Finally, assume that $C(k_1,\dots,k_m)$ contains 1, 2 and 3. Without loss of generality we can assume that $k_{m'} = k_{m'+1} = \dots = k_m = 1$ for some m', i.e.

$$C(k_1,\dots,k_m) = C(k_1,\dots,k_{m'-1},1,\dots,1).$$

By the above discussion, we can construct a permutation for $C(k_1,\dots,k_{m'-1})$. Hence, Lemma 4.3 yields a permutation π for $C(k_1,\dots,k_m)$. □

With the above Lemma 4.5 in hand, it remains to show that every irreducible conjugacy class is contained in the image of the commutator map. For that, we explicitly construct permutations π whose commutator with the shift permutation lies in the given irreducible conjugacy class. The only constraint on π we need to fulfil is that $\pi(1) = 2$.

§ 4.1.1 | Permutations for the irreducible conjugacy classes

For each irreducible conjugacy class below, we define a permutation $\pi \in C(n)$ for the corresponding n such that $[\sigma_n, \pi]$ is contained in that conjugacy class. All permutations we define satisfy $\pi(1) = 2$, so together with the above lemmata, this section concludes the proof of Theorem 4.1.

In the following, whenever we write $\dots^{+d}\dots$ or $\dots^{-d}\dots$ for some $d \in \mathbb{N}$ in a cycle of a permutation, we mean to continue the cycle by adding or subtracting d consecutively until reaching the next number in the permutation. For example $(1\ 2\ \dots^{+2}\dots\ 8\ 7) = (1\ 2\ 4\ 6\ 8\ 7)$. Whenever we write $\dots^{(+d,-d')}\dots$ for some $d, d' \in \mathbb{N}$, we mean to continue by adding d, then subtracting d' and so on until reaching the next number. For example $(1\ 2\ \dots^{(+3,-1)}\dots\ 9) = (1\ 2\ 5\ 4\ 7\ 6\ 9)$.

$\boxed{C(n) \text{ for } n \geqslant 5 \text{ odd}}$

Let $\pi := (1\ \ 2\ \ 4\ \ 5\ \ \overset{+2}{\dots}\ \ n\ \ n-1\ \ \overset{-2}{\dots}\ \ 6\ \ 3)$. We compute:

$$[\sigma_n, \pi] = \begin{cases} (1\ \ 5\ \ \overset{+4}{\dots}\ \ n-4\ \ n-1\ \ \overset{-4}{\dots}\ \ 4\ \ 7\ \ \overset{+4}{\dots}\ \ n-2\ \ 2\ \ 3\ \ n\ \ n-3\ \ \overset{-4}{\dots}\ \ 6), \\ \text{if } n \equiv 1 \bmod 4, \\ \\ (1\ \ 5\ \ \overset{+4}{\dots}\ \ n-2\ \ 2\ \ 3\ \ n\ \ n-3\ \ \overset{-4}{\dots}\ \ 4\ \ 7\ \ \overset{+4}{\dots}\ \ n-4\ \ n-1\ \ \overset{-4}{\dots}\ \ 6), \\ \text{if } n \equiv 3 \bmod 4. \end{cases}$$

In any case, we have $[\sigma_n, \pi] \in C(n)$.

$\boxed{C(2k, 2l) \text{ for } k, l \geqslant 1 \text{ except } k = l = 1}$

We subdivide this case into $k = l$, $k = l + 1$ and $k > l + 1$.

For $k = l$ (and $l \geqslant 4$), let

$$\pi = (1\ \ 2\ \ 5\ \ 3\ \ 4\ \ \overset{(+3,-1)}{\dots}\ \ 2l+2\ \ 2l+6\ \ 2l+4\ \ 2l+5\ \ \overset{(+3,-1)}{\dots}\ \ 4l\ \ 4l-1).$$

We compute:

$$[\sigma_{4l}, \pi] = \begin{cases} (1\ \ 4\ \ \overset{+4}{\dots}\ \ 2l\ \ 2l+5\ \ \overset{+4}{\dots}\ \ 4l-3\ \ 4l-2\ \ \overset{-4}{\dots}\ \ 2l+6\ \ 2l+3\ \ \overset{-4}{\dots}\ \ 3) \\ (2\ \ \overset{+4}{\dots}\ \ 2l+2\ \ 2l+7\ \ \overset{+4}{\dots}\ \ 4l-1\ \ 4l\ \ \overset{-4}{\dots}\ \ 2l+4\ \ 2l+1\ \ \overset{-4}{\dots}\ \ 5), \\ \text{if } l \equiv 0 \bmod 2, \\ \\ (1\ \ 4\ \ \overset{+4}{\dots}\ \ 2l+2\ \ 2l+7\ \ \overset{+4}{\dots}\ \ 4l-3\ \ 4l-2\ \ \overset{-4}{\dots}\ \ 2l+4\ \ 2l+1\ \ \overset{-4}{\dots}\ \ 3) \\ (2\ \ \overset{+4}{\dots}\ \ 2l\ \ 2l+5\ \ \overset{+4}{\dots}\ \ 4l-1\ \ 4l\ \ \overset{-4}{\dots}\ \ 2l+6\ \ 2l+3\ \ \overset{-4}{\dots}\ \ 5), \\ \text{if } l \equiv 1 \bmod 2. \end{cases}$$

Counting the elements in the two cycles of $[\sigma_{4l}, \pi]$ shows that in both cases $[\sigma_{4l}, \pi] \in C(2l, 2l) = C(2k, 2l)$. For $l \in \{2, 3\}$, the respective permutations $(1\ 2\ 4\ 5\ 8\ 6\ 7\ 3)$ and $(1\ 2\ 5\ 3\ 4\ 7\ 6\ 9\ 8\ 12\ 10\ 11)$ give the desired results $C(4, 4)$ and $C(6, 6)$.

For $k = l + 1$ (and $l \geqslant 4$) let

$$\pi = (1\ \ 2\ \ 5\ \ 3\ \ 4\ \ \overset{(+3,-1)}{\dots}\ \ 2l+2\ \ 2l+6\ \ 2l+4\ \ 2l+5\ \ \overset{(+3,-1)}{\dots}\ \ 4l\ \ 4l-1\ \ 4l+2\ \ 4l+1).$$

Since this permutation is obtained by the one in the $k = l$ case by adding $4l+2, 4l+1$ at the end, which doesn't change the $(+3,-1)$ pattern, a similar computation to the one before shows $[\sigma_{4l+2}, \pi] \in C(2l+2, 2l) = C(2k, 2l)$. For $l \in \{1,2,3\}$, the respective permutations $(1\ 2\ 5\ 4\ 6\ 3), (1\ 2\ 5\ 3\ 4\ 7\ 6\ 10\ 8\ 9)$ and $(1\ 2\ 5\ 3\ 4\ 7\ 6\ 9\ 8\ 12\ 10\ 11\ 14\ 13)$ give the desired results $C(4,2), C(6,4)$ and $C(8,6)$.

For $k > l+1$ (and $l \geqslant 2$), let

$$\pi = (1\ \ 2\ \ 5\ \ 3\ \ 4\ \overset{(+3,-1)}{\therefore}\ 4l{-}4\ \ 4l\ \ 4l{-}2\ \ 4l{-}1\ \ 4l{+}1\ \ 4l{+}4\ \ 4l{+}2\ \ 4l{+}3\ \overset{(+3,-1)}{\therefore}\ 2k{+}2l\ \ 2k{+}2l{-}1).$$

We compute:

$$[\sigma_{2k+2l}, \pi] = \begin{cases} (1\ \ 4\ \overset{+4}{\therefore}\ 4l-4\ \ 4l-2\ \ 4l-5\ \overset{-4}{\therefore}\ 3) \\ (2\ \overset{+4}{\therefore}\ 4l-6\ \ 4l-1\ \overset{+4}{\therefore}\ 2k+2l-1\ \ 2k+2l\ \overset{-4}{\therefore}\ 4l+4 \\ 4l+1\ \overset{+4}{\therefore}\ 2k+2l-3\ \ 2k+2l-2\ \overset{-4}{\therefore}\ 4l+2\ \ 4l\ \ 4l-3\ \overset{-4}{\therefore}\ 5), \\ \text{if } k+l \equiv 0 \bmod 2, \\ \\ (1\ \ 4\ \overset{+4}{\therefore}\ 4l-4\ \ 4l-2\ \ 4l-5\ \overset{-4}{\therefore}\ 3) \\ (2\ \overset{+4}{\therefore}\ 4l-6\ \ 4l-1\ \overset{+4}{\therefore}\ 2k+2l-3\ \ 2k+2l-2\ \overset{-4}{\therefore}\ 4l+4 \\ 4l+1\ \overset{+4}{\therefore}\ 2k+2l-1\ \ 2k+2l\ \overset{-4}{\therefore}\ 4l+2\ \ 4l\ \ 4l-3\ \overset{-4}{\therefore}\ 5), \\ \text{if } k+l \equiv 1 \bmod 2. \end{cases}$$

Again counting the elements in the two cycles of $[\sigma_{2k+2l}, \pi]$ shows that in both cases $[\sigma_{2k+2l}, \pi] \in C(2k, 2l)$. The remaining case is $k > l+1$ and $l = 1$. If we let $n := 2k+2$, this corresponds to $C(n-2, 2)$ for $n \geqslant 8$. For this (if $n \geqslant 10$), let

$$\pi = (1\ \ 2\ \ 4\ \ 5\ \ 8\ \ 7\ \ 10\ \overset{+2}{\therefore}\ n\ \ n-1\ \overset{-2}{\therefore}\ 9\ \ 6\ \ 3),$$

and compute:

$$[\sigma_n,\pi]=\begin{cases}(4\ 8)\\ (1\ \ 5\ \ 10\ \ \therefore^{+4}\ \ n-2\ \ 2\ \ 3\ \ n\ \ n-3\ \ \therefore^{-4}\ \ 9\ \ 7\ \ 12\ \ \therefore^{+4}\ \ n-4\ \ n-1\ \ \therefore^{-4}\ \ 11\ \ 6),\\ \text{if } n\equiv 0 \bmod 4,\\ \\ (4\ 8)\\ (1\ \ 5\ \ 10\ \ \therefore^{+4}\ \ n-4\ \ n-1\ \ \therefore^{-4}\ \ 9\ \ 7\ \ 12\ \ \therefore^{+4}\ \ n-2\ \ 2\ \ 3\ \ n\ \ n-3\ \ \therefore^{-4}\ \ 11\ \ 6),\\ \text{if } n\equiv 2 \bmod 4.\end{cases}$$

Hence, $[\sigma_n,\pi]\in C(n-2,2)$. Finally, for $n=8$, the permutation $(1\ 2\ 4\ 6\ 8\ 7\ 5\ 3)$ gives the desired conjugacy class $C(6,2)$.

$\boxed{C(n,3) \text{ for } n\geqslant 1 \text{ odd}}$

For $n=1$, $\pi=(1\ 2\ 4\ 3)$ and for $n=3$, $\pi=(1\ 2\ 4\ 3\ 6\ 5)$ give the desired result. For $n\geqslant 5$, let $\pi=(1\ \ 2\ \ 4\ \ 5\ \ 7\ \ 8\ \ \therefore^{+2}\ \ n+3\ \ n+2\ \ \therefore^{-2}\ \ 9\ 6\ 3)$. We compute:

$$[\sigma_{n+3},\pi]=\begin{cases}(1\ 5\ 6)\,(2\ \ 3\ \ n+3\ \ n\ \ \therefore^{-4}\ \ 9\ \ 4\ \ \therefore^{+4}\ \ n-1\ \ n+2\ \ \therefore^{-4}\ \ 7\ \ 10\ \ \therefore^{+4}\ \ n+1)\\ \text{if } n\equiv 1 \bmod 4,\\ \\ (1\ 5\ 6)\,(2\ \ 3\ \ n+3\ \ n\ \ \therefore^{-4}\ \ 7\ \ 10\ \ \therefore^{+4}\ \ n-1\ \ n+2\ \ \therefore^{-4}\ \ 9\ \ 4\ \ \therefore^{+4} n+1)\\ \text{if } n\equiv 3 \bmod 4.\end{cases}$$

In any case, we have $[\sigma_{n+3},\pi]\in C(n,3)$.

$\boxed{C(2k,2l,3) \text{ for } k,l\geqslant 1}$

Let $n:=2k+2l$ and $\pi\in C(n)$ be any of the permutations constructed in the $C(2k,2l)$ case above such that $[\sigma_n,\pi]\in C(2k,2l)$. Then π is of the form $(1\ 2\ \ldots\ i\ n\ \ldots\ j)$, where $i,j\in\{3,\ldots,n-1\}$. Define $\tau\in C(n+5)$ by

$$\tau:=(1\ \ 2\ \ \ldots\ \ i\ \ n\ \ \ldots\ \ j\ \ n+1\ \ n+3\ \ n+2\ \ n+5\ \ n+4),$$

where τ is the same as π from 1 to j. Note that whatever π is, the commutator $[\sigma_{n+5}, \tau]$ contains the cycle $(n+1 \;\; n+4 \;\; n+5)$. For any $m \in \{1, \ldots, n\} \setminus \{i, j\}$ it is straightforward that $[\sigma_{n+5}, \tau](m) = [\sigma_n, \pi](m)$. Furthermore, we also have $[\sigma_{n+5}, \tau](i) = j - 1 = [\sigma_n, \pi](i)$. For j we have $[\sigma_n, \pi](j) = n$, whereas the commutator $[\sigma_{n+5}, \tau]$ contains a cycle of the form $(\ldots \;\; j \;\; n+2 \;\; n+3 \;\; n \;\; \ldots)$. Since $[\sigma_n, \pi] \in C(2k, 2l)$ and one can check that the element n is contained in the $2k$ cycle of $[\sigma_n, \pi]$ for any permutation π constructed above, we conclude that $[\sigma_{n+5}, \tau] \in C(2k+2, 2l, 3)$.

Using the permutations for the $C(2k, 2l)$ case above and the method just described, we cover every case except for $C(4, 2, 3)$ and $C(2k, 2k, 3)$ for $k \geqslant 1$. The permutation $(1\ 2\ 4\ 5\ 8\ 7\ 9\ 6\ 3)$ gives the desired result for $C(4, 2, 3)$. For $C(2k, 2k, 3)$ (and $k \geqslant 4$) consider

$$\pi = (1 \;\; 2 \;\; 6 \;\; 3 \;\; 4 \;\; 10 \;\; 5 \;\; 8 \;\; 9 \;\; 7 \;\; 12 \;\; \overset{(-1,+3)}{\ldots} \;\; 4k-3 \;\; 4k-1 \;\; 4k+1 \;\; 4k \;\; 4k+3 \;\; 4k+2).$$

We compute

$$\begin{aligned}[\sigma_{4k+3}, \pi] = (1 \;\; 5 \;\; 7 \;\; 13 \;\; \overset{+4}{\ldots} \;\; 4k-3 \;\; 4k \;\; 4k+1 \;\; 4k-2 \;\; \overset{-4}{\ldots} \;\; 10)\\ (2 \;\; 8 \;\; 3 \;\; 9 \;\; 4 \;\; 11 \;\; \overset{+4}{\ldots} \;\; 4k-5 \;\; 4k-4 \;\; \overset{-4}{\ldots} \;\; 12 \;\; 6)\\ (4k-1 \;\; 4k+2 \;\; 4k+3),\end{aligned}$$

and see that $[\sigma_{4k+3}, \pi] \in C(2k, 2k, 3)$.

Finally, the permutations $(1\ 2\ 4\ 6\ 7\ 5\ 3)$ and $(1\ 2\ 4\ 6\ 5\ 3\ 7\ 9\ 8\ 11\ 10)$ as well as $(1\ 2\ 6\ 3\ 4\ 10\ 5\ 8\ 9\ 7\ 11\ 13\ 12\ 15\ 14)$ give the desired result for the conjugacy classes $C(2, 2, 3), C(4, 4, 3)$ and $C(6, 6, 3)$ respectively.

$\boxed{C(n, 2, 2) \text{ for } n \geqslant 3 \text{ odd}}$

For $n = 3$, $\pi = (1\ 2\ 4\ 6\ 7\ 5\ 3)$ gives the desired result. For $n \geqslant 5$, we let

$$\pi = (1 \;\; 2 \;\; 4 \;\; n+2 \;\; n+4 \;\; n+3 \;\; 5 \;\; 6 \;\; \overset{+2}{\ldots} \;\; n+1 \;\; n \;\; \overset{-2}{\ldots} \;\; 7 \;\; 3).$$

We compute:

$$[\sigma_{n+4},\pi] = \begin{cases} (2\ n+2)\,(4\ n+3) \\ (1\ \ 6\ \ \overset{+4}{\dots}\ \ n-3\ \ n\ \ \overset{-4}{\dots}\ \ 5\ \ 8\ \ \overset{+4}{\dots}\ \ n-1\ \ 3\ \ n+4\ \ n+1\ \ n-2\ \ \overset{-4}{\dots}\ \ 7) \\ \text{if } n \equiv 1 \bmod 4, \\ \\ (2\ n+2)\,(4\ n+3) \\ (1\ \ 6\ \ \overset{+4}{\dots}\ \ n-1\ \ 3\ \ n+4\ \ n+1\ \ n-2\ \ \overset{-4}{\dots}\ \ 5\ \ 8\ \ \overset{+4}{\dots}\ \ n-3\ \ n\ \ \overset{-4}{\dots}\ \ 7) \\ \text{if } n \equiv 3 \bmod 4. \end{cases}$$

In any case, we have $[\sigma_{n+4},\pi] \in C(n,2,2)$.

$\boxed{C(2k,2,2,2) \text{ for } k \geqslant 1}$

For $k = 1$, $\pi = (1\ 2\ 5\ 4\ 7\ 8\ 6\ 3)$ and for $k = 2$, $\pi = (1\ 2\ 4\ 8\ 9\ 7\ 5\ 6\ 10\ 3)$ give the desired result. For $k \geqslant 3$, we distinguish whether $k \equiv 0 \pmod 2$ or $k \equiv 1 \pmod 2$. In the first case, let $\pi = (1\ \ 2\ \ 5\ \ 4\ \ 7\ \ 9\ \ 10\ \ \overset{+2}{\dots}\ \ n-2\ \ n-3\ \ \overset{-2}{\dots}\ \ 11\ \ 8\ \ n-1\ \ n\ \ 6\ \ 3)$. In the second case, let $\pi = (1\ \ 2\ \ 5\ \ 4\ \ 7\ \ \overset{+2}{\dots}\ \ n-3\ \ n-2\ \ \overset{-2}{\dots}\ \ 8\ \ n-1\ \ n\ \ 6\ \ 3)$. Here, $n := 2k+6$.

We compute:

$$[\sigma_n,\pi] = \begin{cases} (1\ 5)\,(2\ n-1)\,(3\ n) \\ (4\ \ 10\ \ \overset{+4}{\dots}\ \ n-4\ \ 7\ \ 8\ \ n-2\ \ n-5\ \ \overset{-4}{\dots}\ \ 9\ \ 12\ \ \overset{+4}{\dots}\ \ n-6\ \ n-3\ \ \overset{-4}{\dots}\ \ 11\ \ 6) \\ \text{if } k \equiv 0 \bmod 2, \\ \\ (1\ 5)\,(2\ n-1)\,(3\ n) \\ (4\ \ 9\ \ \overset{+4}{\dots}\ \ n-3\ \ 7\ \ \overset{+4}{\dots}\ \ n-5\ \ n-4\ \ \overset{-4}{\dots}\ \ 8\ \ n-2\ \ \overset{-4}{\dots}\ \ 6) \\ \text{if } k \equiv 1 \bmod 2. \end{cases}$$

In any case, we have $[\sigma_n,\pi] \in C(2k,2,2,2)$.

Chapter 5

Filling pairs, square-tiled surfaces and pseudo-Anosov stretch factors

In this chapter, we display the connection between sets of filling curves and square-tiled surfaces and present how these concepts can be used to construct pseudo-Anosovs. We then use Theorem 4.1 on commutators of n-cycles in the symmetric group presented in Chapter 4 to derive four theorems regarding the aforementioned notions. Theorems 1 and 2 below also follow from Luke Jeffrey's work [Jef21]. Theorem 3 is novel. Theorem 4 is known to experts in the field, but to the author's knowledge, it has not been written up anywhere in this form. Our main goal is to portray how the purely group-theoretic result on commutators allows us to create slick new proofs of geometric statements.

We start by stating Theorems 1, 2, 3 and 4 and providing some background information. Sections 5.1 and 5.2 contain formal definitions of the main notions appearing in the theorems. In Section 5.3, we present proofs for Theorems 1, 2 and 3, and sketch proof ideas for Theorem 4.

Given a closed, oriented, genus g surface S_g and a filling pair of curves α, β of S_g, a natural question is how many times α and β intersect, i.e. what their geometric intersection number $i(\alpha, \beta)$ is. It is known that for every genus other than 2, the minimal intersection number of two such curves is $2g - 1$, and examples achieving this minimum have been constructed (see, for example, [AH15]). In genus 2, the minimal intersection number of a filling pair of curves is 4 (see [AH15, Theorem 2.17]). We show that every number above the minimum can be achieved as the intersection number of a filling pair of curves. In particular, we prove:

Theorem (1).

- *Let $g \geqslant 3$. Then for all $n \geqslant 2g-1$ there exists a filling pair of curves (α, β) on a genus g surface with $i(\alpha, \beta) = n$.*
- *For $g = 2$, there exists such a filling pair for all $n \geqslant 4$.*
- *For $g = 1$, there exists such a filling pair for all $n \geqslant 1$.*

In fact, the proof of Theorem 1 is constructive and the filling pairs we obtain satisfy $i(\alpha, \beta) = \hat{i}(\alpha, \beta)$, where $\hat{i}$ denotes the algebraic intersection number. Hence, one could replace geometric with algebraic intersection number in the above theorem.

The question regarding the intersection of filling pairs is related to the question about the number of squares needed to build a square-tiled surface with a single horizontal and a single vertical cylinder. Such a square-tiled surface will be called a $(1,1)$-square-tiled surface. We will see the connection between such square-tiled surfaces and filling pairs of curves in Section 5.1.1.

A square-tiled surface is, in particular, a translation surface. Therefore, one can ask which stratum of translation surfaces it is contained in. The question we are interested in is the following: given a stratum $\mathcal{H}$, with how many squares can one build a $(1,1)$-square-tiled surface X with $X \in \mathcal{H}$? As for the filling pair question, the minimal such number of squares needed has been found, and examples realising this minimum have been provided by Jeffreys [Jef21]. In fact, Jeffreys' work is even more subtle as it doesn't just provide examples for each stratum, but even for each connected component of a stratum. Here, we don't make this refinement. We show that for a given stratum, one can construct a $(1,1)$-square-tiled surface in that stratum with any number of squares above the minimum necessary.

Theorem (2).

- *Let $g \geqslant 3$ and $(k_1, \ldots, k_m)$ such that $k_1 + \ldots + k_m = 2g-2$. Then for all $n \geqslant 2g-2+m$ there exists a $(1,1)$-square-tiled surface $X \in \mathcal{H}(k_1, \ldots, k_m)$ with exactly n squares.*
- *For $g = 2$, there exist such $(1,1)$-square-tiled surfaces $X \in \mathcal{H}(2)$ for all $n \geqslant 4$ and $X \in \mathcal{H}(1,1)$ for all $n \geqslant 6$.*
- *For $g = 1$, there exists such a $(1,1)$-square-tiled surface $X \in \mathcal{H}(0)$ for all $n \geqslant 1$.*

Filling pairs of curves can be used to construct pseudo-Anosovs with stretch factors related to the intersection number of the curves. Furthermore, such a pseudo-Anosov is affine with respect to a translation surface corresponding to the filling pair. We will

be interested in seeing whether pseudo-Anosovs that are affine with respect to different strata of translation surfaces can have the same stretch factor.

In particular, we are interested in the following sets: for a stratum of translation surfaces $\mathcal{H}$, let $\Lambda_{\mathcal{H}}$ denote the set of stretch factors of pseudo-Anosovs that are affine with respect to some translation surface in $\mathcal{H}$. We refer to the sets $\Lambda_{\mathcal{H}}$ as *affine stretch factor sets*. Since the sets of affine pseudo-Anosovs are disjoint, one might expect that the affine stretch factor sets are as well. However, we show that the affine stretch factor sets have a big overlap; in other words, the same stretch factors appear for many pseudo-Anosovs, no matter which stratum they are affine for or which genus g surface they live on.

For a stratum $\mathcal{H} = \mathcal{H}(k_1, \ldots, k_m)$, we define $|\mathcal{H}| := k_1 + \ldots + k_m + m$, and show:

Theorem (3). *Let $n \in \mathbb{N}$. Then the set*

$$\bigcap_{|\mathcal{H}| \leq n} \Lambda_{\mathcal{H}}$$

is infinite. Equivalently, any finite intersection of affine stretch factor sets is infinite.

Finally, we focus on the multiplicity of stretch factors. Let Λ_g be the set of all $\lambda > 1$ that arise as a stretch factor of a pseudo-Anosov on S_g. All conjugates of a given pseudo-Anosov are also pseudo-Anosov mapping classes and have the same stretch factor. Hence, if we let $PA(g)$ denote the set of conjugacy classes of pseudo-Anosovs in $\mathrm{Mod}(S_g)$, we obtain a well-defined map

$$PA(g) \to \Lambda_g$$

taking a pseudo-Anosov to its stretch factor. The cardinality of the preimage of a given $\lambda \in \Lambda_g$ under the above map is called the *multiplicity* of λ and will be denoted by m_λ. Each m_λ is known to be a finite number [FM12, Theorem 14.9]. Here, we pose the question of how big m_λ can be, i.e. of determining the value of

$$m_g := \sup_{\lambda \in \Lambda_g} m_\lambda.$$

We show that for any genus, this supremum is infinite, i.e. on every S_g there exist stretch factors that are attained by arbitrarily many pairwise non-conjugate pseudo-Anosovs.

Theorem (4). *For any $g \geqslant 1$, it holds that*

$$m_g = \infty.$$

§ 5.1 | Filling pairs of curves and relation with STS

Let S_g be a closed, oriented surface of genus g. Unlike in Chapter 1, in this section we use the word *curve* to refer to an essential, simple, closed curve γ on S_g and not to its isotopy class. For two curves γ_1, γ_2, we write $i(\gamma_1, \gamma_2)$ to denote their geometric intersection number, i.e. the minimal intersection number over all γ_1', γ_2' where γ_1' and γ_2' are isotopic to γ_1 and γ_2 respectively. If γ_1 and γ_2 attain this minimum, we say that they are in *minimal position*.

We say that a set of curves $C = \{\gamma_1, \dots, \gamma_n\}$ *fills* the surface if the curves are pairwise in minimal position and the complement of C is a union of topological disks. In the case where $C = \{\alpha, \beta\}$ consists of two elements, we say that the two curves form a *filling pair*.

Let $\{\alpha, \beta\}$ be a filling pair of curves. Let $n \coloneqq i(\alpha, \beta)$ be their geometric intersection number. The intersection number has to satisfy the inequality $n \geqslant 2g - 1$. Indeed, this follows by an Euler characteristic argument:

Define a graph G on S_g by letting the intersection points of α and β be the vertices and the arcs of α and β connecting the intersection points be the edges. Then G consists of n vertices and has valence 4 at every vertex since exactly two arcs of α and two arcs of β come out of every intersection point. The number of edges of G is therefore $\frac{4n}{2} = 2n$. Applying the formula for the Euler characteristic and using the fact that there is at least one face yields

$$2 - 2g \geqslant n - 2n + 1,$$

which implies $n \geqslant 2g - 1$.

§ 5.1.1 | Connection to square-tiled surfaces

Recall that a square-tiled surface was defined as a collection of squares appropriately glued together as outlined in Section 1.3.2. Recall further that a STS can be described in terms of two permutations $\pi_h, \pi_v \in S_n$ where n is the number of squares of the STS.

We now explore the relationship between square-tiled surfaces and filling curves. Let X be a square-tiled surface consisting of n squares. We can construct closed curves on X by starting at the midpoint of a square and going horizontally to the right until the curve closes up. Call all the resulting curves $\alpha_1, \dots, \alpha_k$. Note that the number k is the number of cycles of the permutation π_h and that the number of squares some α_i passes through is given by the length of the corresponding cycle. Analogously, we construct curves $\beta_1, \dots, \beta_l$ where instead of moving right, we move downwards. As before, l is given by the number of cycles of π_v, and the length of the cycles determines the number of squares each β_i

passes through.

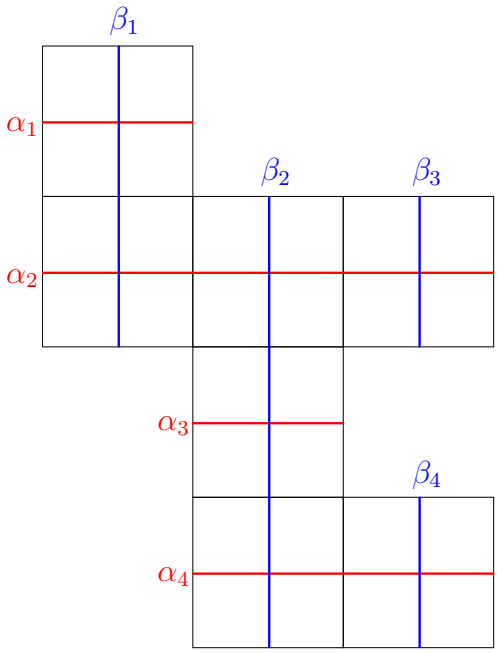


Figure 5.1: A square-tiled surface and its corresponding set of filling curves.

It might happen that one of the α_i (or β_i) is isotopic to an α_j (or β_j). If this happens, then we can reduce the number of squares by collapsing the two isotopic curves to a single curve as demonstrated in Figure 5.2. From now on, we assume that all α_i are pairwise non-isotopic and that the same is true for the β_i.

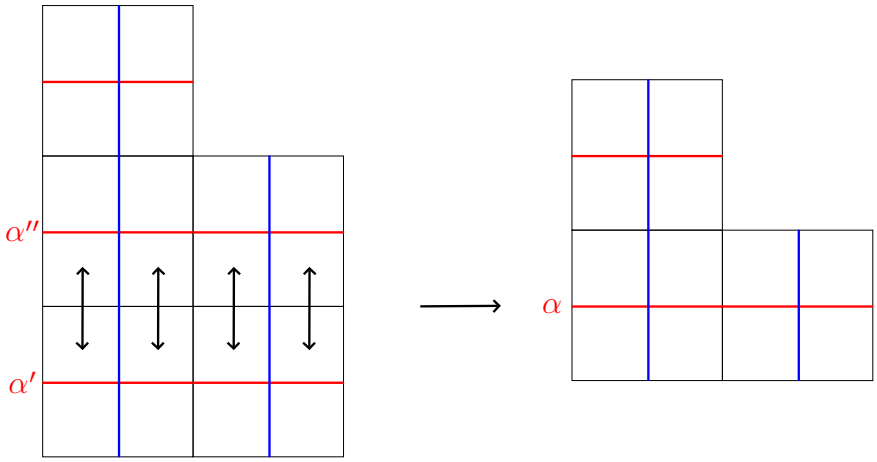


Figure 5.2: The two isotopic curves α' and α'' are collapsed to a single curve α.

Let $A := \{\alpha_1, \ldots, \alpha_k\}$ and $B := \{\beta_1, \ldots, \beta_l\}$. We have $i(\alpha_i, \alpha_j) = 0 = i(\beta_r, \beta_s)$ for any i, j, r, s. Furthermore, the set $A \cup B$ fills the surface since the complementary regions are all topological disks, each of which contains at most one singularity of the STS. The

number of disks is equal to the number of cycles of the commutator $[\pi_h, \pi_v]$. In the special case where both $A = \{\alpha\}$ and $B = \{\beta\}$ consist of a single curve, we call the square-tiled surface a $(1,1)$-*square-tiled surface*. The terminology comes from the fact that, in this case, there is exactly one horizontal and one vertical cylinder. The condition defining a $(1,1)$-STS is equivalent to requiring π_h and π_v to be n-cycles. In this case, the intersection number $i(\alpha,\beta) = n$ is the same as the number of squares of X. In fact, with an appropriate choice of orientation of the curves, the algebraic intersection number $\hat{i}(\alpha,\beta)$ also equals n in this case (see [FM12, Section 1.2.3] for a definition of the algebraic intersection number).

Conversely, we can start with a filling pair of curves α, β with $i(\alpha,\beta) = n$ and construct the graph G as above. Let Γ be the dual graph of G. Then Γ consists of n faces which are all squares. The number of edges of Γ is $2n$, and the number of vertices is $n + 2 - 2g$. Since all the faces of Γ are squares, we obtain a square tiling of the surface. Now, think of the faces of Γ as squares embedded in the plane in a way that α traverses each square horizontally and β traverses each square vertically. The graph Γ then tells us how to glue the squares together to obtain the original surface. In order for the squares to be glued in a way satisfying the definition of a STS, and hence yielding a translation surface, it is necessary and sufficient that $i(\alpha,\beta) = |\hat{i}(\alpha,\beta)|$. Otherwise, some squares will have to be glued by a half-translation, and we obtain a half-translation surface.

§ 5.1.2 | Example: $i(\alpha,\beta) = 3$

The goal of this subsection is to portray the connection between filling pairs and square tilings with the help of an example. This will consist of showing that genus 2 behaves differently from other genera: while for any $g \neq 2$ there is a filling pair of curves with intersection number $2g-1$, on a genus 2 surface, the minimal intersection number of a filling pair is four, instead of three. Here, we show that three is impossible, while Section 5.3 below provides a proof that intersection number four (and any number above) is possible.

Proposition 5.1. *There is no filling pair of curves α, β on a genus 2 surface with*

$$i(\alpha,\beta) = 3.$$

Proof. Let S be a closed, oriented surface and α, β two curves with $i(\alpha,\beta) = 3$ that fill S. The goal is to show that S has to be a torus.

Let Γ be the graph dual to the curves α, β as defined above. Γ yields a square tiling of S which consists of three squares. Let C_1, C_2, C_3 be these three squares. Each C_i contains exactly one arc of α and one arc of β that intersect once inside C_i. Embed each C_i in the

plane such that α is horizontal and β vertical. The graph Γ specifies how to glue the C_i together in order to reconstruct the surface S. Since α is a closed curve, without loss of generality, we can assume that the three squares are glued together in a horizontal line as shown in Figure 5.3. Note that the left side of C_1 is glued to the right side of C_3 by a translation (as opposed to a half-translation) since S is orientable.

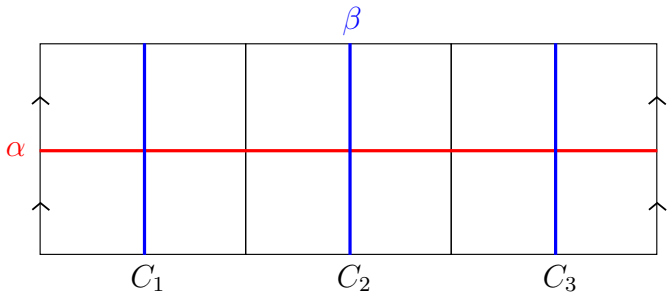


Figure 5.3: The square tiling of S.

It remains to analyse how the three top and three bottom sides of the C_i are glued together. Again, due to the orientability of S, if two top sides are glued together, they must be glued by a half-translation, while if a top side is glued to a bottom side, it must be by a translation. We encode the gluing by a labelled permutation of the form

$$\begin{pmatrix} A & B & C \\ C & A & B \end{pmatrix},$$

where the letters indicate the pairing of sides that are glued together. In the above permutation, we would glue all sides by a translation, but we could also have a permutation of the form

$$\begin{pmatrix} A & A & B \\ B & C & C \end{pmatrix},$$

where we would glue the sides labelled by B by a translation and the sides labelled by A and C by a half-translation.

Note that since β is a closed curve, we cannot have a letter appearing right below itself, since this would create a closed curve that intersects α only once. Keeping this in mind, up to relabelling, we are left with the following possibilities:

$$\begin{pmatrix} A & A & B \\ B & C & C \end{pmatrix}, \begin{pmatrix} A & A & B \\ C & B & C \end{pmatrix}, \begin{pmatrix} A & B & A \\ B & C & C \end{pmatrix}, \begin{pmatrix} A & B & A \\ C & C & B \end{pmatrix},$$

$$\begin{pmatrix} A & B & B \\ C & A & C \end{pmatrix}, \begin{pmatrix} A & B & C \\ C & A & B \end{pmatrix}, \begin{pmatrix} A & B & B \\ C & C & A \end{pmatrix}, \begin{pmatrix} A & B & C \\ B & C & A \end{pmatrix}.$$

Now, observe that we can cut the square C_1 and reglue it to the right of C_3 without changing the (half-)translation surface structure. This cutting and regluing move changes our labelled permutation by deleting the two leftmost entries of both rows and rewriting them as rightmost entries, for example

$$\begin{pmatrix} A & A & B \\ B & C & C \end{pmatrix} \longrightarrow \begin{pmatrix} A & B & A \\ C & C & B \end{pmatrix}.$$

After doing all such possible moves and possibly relabelling, we see that we are left with the permutations:

$$\begin{pmatrix} A & A & B \\ B & C & C \end{pmatrix}, \begin{pmatrix} A & A & B \\ C & B & C \end{pmatrix}, \begin{pmatrix} A & B & C \\ C & A & B \end{pmatrix}, \begin{pmatrix} A & B & C \\ B & C & A \end{pmatrix}.$$

It is easy to see that the latter two permutations yield a torus by looking at Figure 5.3 and observing the following. If we glue the left side of C_1 to the right side of C_3 as indicated, we obtain a cylinder. Using any of the gluings from the latter two permutations above says that we have to glue the whole of the top boundary of this cylinder to the whole of its bottom boundary. Since the cyclic order in which the letters appear in the top and bottom rows of the permutations is the same, we obtain a torus.

We now analyse the case

$$\begin{pmatrix} A & A & B \\ B & C & C \end{pmatrix},$$

which corresponds to the half-translation surface in Figure 5.4.

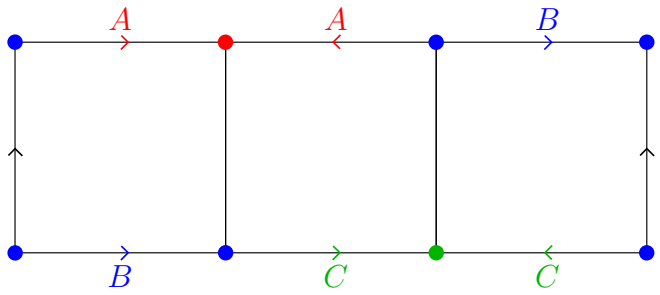


Figure 5.4: The half-translation surface defined by the permutation $\begin{pmatrix} A & A & B \\ B & C & C \end{pmatrix}$.

As can be seen from the figure, there are three singularities on this half-translation surface with cone angles π, π and 4π respectively. This implies that the permutation

$$\begin{pmatrix} A & A & B \\ B & C & C \end{pmatrix},$$

yields a surface in the stratum $\mathcal{Q}(-1,-1,2)$. From Section 1.3, we know that the genus of such a surface is given by the formula

$$4g - 4 = -1 - 1 + 2 = 0,$$

which means that $g = 1$ and we are again in the case of a torus.

The permutation

$$\begin{pmatrix} A & A & B \\ C & B & C \end{pmatrix}$$

can be analysed analogously and yields a torus as well.

We conclude that in any case, the surface S is a torus. In particular, only the torus admits a filling pair of curves with intersection number three, but a genus 2 surface does not. □

§ 5.2 | Affine pseudo-Anosov maps

In this section, we define the *affine stretch factor set* of a stratum of translation surfaces and describe the *Thurston construction* of pseudo-Anosovs.

Recall from Section 2.1.3 that a homeomorphism f of S_g is said to be affine for some translation surface if it can be identified with a homeomorphism of that translation surface which is affine. We say that f is affine for a stratum $\mathcal{H}$, if it is affine for some translation surface in that stratum. We denote the set of all pseudo-Anosovs that are affine for $\mathcal{H}$ by $PA(\mathcal{H})$. For a pseudo-Anosov f, we denote its stretch factor by λ_f. Finally, we define

$$\Lambda_{\mathcal{H}} := \{\lambda \in \mathbb{R} \mid \lambda = \lambda_f \text{ for some } f \in PA(\mathcal{H})\},$$

and call $\Lambda_{\mathcal{H}}$ the *affine stretch factor set* for $\mathcal{H}$.

The following theorem, due to Thurston, is a powerful tool in order to construct pseudo-Anosov mapping classes.

A finite set $A = \{a_1, \dots, a_n\}$ whose elements a_i are curves on the surface S_g is called a *multicurve* if all a_i are pairwise disjoint. We write T_A for the mapping class given by twisting about all the a_i and call T_A a *multitwist*. Note that the order in which we twist about the a_i does not matter since they are disjoint. Thurston proved the following:

Theorem 5.2 (Thurston's construction)**.** *Let A, B be multicurves on S_g such that $A \cup B$ fills S_g. Let $\langle T_A, T_B \rangle$ denote the subgroup of* $\mathrm{Mod}(S_g)$ *generated by T_A, T_B. Then, there exists a real number μ and a representation $\rho\colon \langle T_A, T_B \rangle \to \mathrm{PSL}_2(\mathbb{R})$ defined by*

$$T_A \mapsto \begin{pmatrix} 1 & -\sqrt{\mu} \\ 0 & 1 \end{pmatrix}, \quad T_B \mapsto \begin{pmatrix} 1 & 0 \\ \sqrt{\mu} & 1 \end{pmatrix},$$

such that $f \in \langle T_A, T_B \rangle$ is pseudo-Anosov if and only if $\rho(f)$ is hyperbolic. Furthermore, the stretch factor of such f is given by the larger of the two eigenvalues of $\rho(f)$.

Proof. See [FM12, Theorem 14.1]. □

Remark 5.3. The real number μ in Theorem 5.2 can be computed as follows. Let $A = \{a_1, \dots, a_k\}$, $B = \{b_1, \dots, b_l\}$ and let N be the $k \times l$ matrix defined by $N_{rs} = i(a_r, b_s)$. One can show that NN^t is a primitive matrix, so the Perron-Frobenius theorem applies [FM12, Theorem 14.2]. The real number μ is the spectral radius of NN^t.

When we say larger eigenvalue in the theorem above, we mean the eigenvalue > 1 of the representative of $\rho(f)$ in $\mathrm{SL}_2(\mathbb{R})$ with positive eigenvalues. We remark that a matrix in $\mathrm{SL}_2(\mathbb{R})$ is hyperbolic if and only if its trace is of absolute value greater than 2.

In fact, the representation ρ from the above theorem also determines which mapping classes of $\langle T_A, T_B \rangle$ are periodic or reducible, but we won't need to make use of this.

In Section 5.1.1, we discussed the connection of filling pairs and square-tiled surfaces, or the more general square-tiled half-translation surfaces. In an analogous way, a pair of multicurves that together fill S_g yields a square tiling of the surface, which in turn gives rise to a square-tiled (half-)translation surface. Let $\mathcal{H}$ (or $\mathcal{Q}$) denote the stratum of this (half-)translation surface. Then, every mapping class in $\langle T_A, T_B \rangle$ is affine for $\mathcal{H}$ (or $\mathcal{Q}$) (see the proof of [FM12, Theorem 14.1]). Hence, Thurston's construction allows us to explicitly build pseudo-Anosovs for which we know the stretch factor as well as the stratum they are affine for.

§ 5.3 | Proofs of the results

The proofs of all four Theorems 1, 2, 3 and 4 rely on Theorem 4.1, the result about commutators proved in Chapter 4. The idea is that in all proofs we want to build $(1,1)$-square-tiled surfaces. These surfaces then give rise to filling pairs of curves, which in turn can be used for Thurston's construction of pseudo-Anosovs. Since $(1,1)$-square-tiled surfaces are specified by two permutations which are n-cycles, we are able to make use of Theorem 4.1.

We quickly recall the theorem. Let S_n be the symmetric group and A_n the alternating group on n elements. Let $C(n)$ denote the conjugacy class of n-cycles. Denote by

$$[\pi, \sigma] := \pi^{-1}\sigma^{-1}\pi\sigma$$

the commutator of $\pi, \sigma \in S_n$.

In Theorem 4.1, we showed that the map

$$C(n) \times C(n) \to A_n,\ (\pi, \sigma) \mapsto [\pi, \sigma]$$

is surjective for all $n \geqslant 6$.

In each of the following proofs, we make use of this result.

§ 5.3.1 | Proof of Theorem 1

Theorem (1).

- *Let $g \geqslant 3$. Then for all $n \geqslant 2g-1$ there exists a filling pair of curves (α, β) on a genus g surface with $i(\alpha, \beta) = n$.*
- *For $g = 2$, there exists such a filling pair for all $n \geqslant 4$.*
- *For $g = 1$, there exists such a filling pair for all $n \geqslant 1$.*

Proof. Let $n \geqslant 5$ be odd. Then there exist $\pi, \sigma \in C(n) \subset S_n$ with $[\pi, \sigma] \in C(n)$. For $n \geqslant 7$ this follows from Theorem 4.1 and for $n = 5$ one can take $\sigma = (1\,2\,3\,4\,5)$ and $\pi = (1\,2\,4\,5\,3)$ so that $[\sigma, \pi] = (1\,4\,2\,3\,5) \in C(5)$.

The two permutations π and σ give rise to a $(1,1)$-square-tiled surface X. Since by Lemma 1.17 the cycle type of the commutator determines the stratum of the STS, we obtain that X lies in the stratum $\mathcal{H}(n-1)$. Hence, the genus g of X satisfies $2g-2 = n-1$, i.e. $g = \frac{n+1}{2}$ or $n = 2g-1$. Furthermore, following Section 5.1.1, X admits a filling pair of curves α, β which are the horizontal and vertical curve on X and satisfy $i(\alpha, \beta) = n$. So, for $g \geqslant 3$, we constructed a filling pair of curves α, β on a genus g surface with $i(\alpha, \beta) = 2g-1$.

Let now $g \geqslant 3$ be fixed and $n > 2g-1$. Then we have $n \geqslant 6$ and hence there exist permutations $\sigma, \pi \in C(n)$ with $[\sigma, \pi] \in C(2g-1, 1, \ldots, 1)$, where the number of 1's is equal to $n-2g+1$. As before, this yields a $(1,1)$-square-tiled surface X which in turn yields a pair of filling curves α, β with $i(\alpha, \beta) = n$. In this case, we obtain $X \in \mathcal{H}(2g-2)$ which implies that X has genus g. This finishes the proof for $g \geqslant 3$.

Now, consider $g = 2$. For $n = 4$, the permutations $\sigma = (1\,2\,3\,4)$ and $\pi = (1\,2\,4\,3)$ yield

a $(1,1)$-square-tiled surface X with a filling pair of curves α, β with $i(\alpha,\beta) = 4$. Since $[\sigma,\pi] = (1\,3\,4)\,(2) \in C(3,1)$, we get $X \in \mathcal{H}(2)$, so X is a genus 2 surface. For $n = 5$, the permutations $\sigma = (1\,2\,3\,4\,5)$ and $\pi = (1\,2\,4\,3\,5)$ satisfy $[\sigma,\pi] \in C(3,1,1)$ and hence yield a genus 2 surface with a filling pair of curves with intersection number 5. For $n \geqslant 6$, Theorem 4.1 implies that there exist permutations $\sigma, \pi \in C(n)$ with $[\sigma,\pi] \in C(3,1,\ldots,1)$ where the number of 1's is $n-3$. These permutations yield a surface with the desired filling pair.

Finally, consider $g = 1$. The torus S_1 is homeomorphic to the product of two circles $S^1 \times S^1$. Let α, β be the two standard curves on the torus coming from the inclusion of the two copies of S^1. The two curves fill the torus and satisfy $i(\alpha,\beta) = 1$. Let T_β denote the Dehn twist about β. The pair $(\alpha, T_\beta^n(\alpha))$ is a filling pair with $i(\alpha, T_\beta^n(\alpha)) = n$ for any $n \geqslant 1$. □

§ 5.3.2 | Proof of Theorem 2

Theorem (2).

- *Let $g \geqslant 3$ and $(k_1, \ldots, k_m)$ such that $k_1 + \ldots + k_m = 2g-2$. Then for all $n \geqslant 2g-2+m$ there exists a $(1,1)$-square-tiled surface $X \in \mathcal{H}(k_1, \ldots, k_m)$ with exactly n squares.*
- *For $g = 2$, there exist such $(1,1)$-square-tiled surfaces $X \in \mathcal{H}(2)$ for all $n \geqslant 4$ and $X \in \mathcal{H}(1,1)$ for all $n \geqslant 6$.*
- *For $g = 1$, there exists such a $(1,1)$-square-tiled surface $X \in \mathcal{H}(0)$ for all $n \geqslant 1$.*

Proof. First, let $k_1, \ldots, k_m$ be such that $k_1 + \ldots + k_m = 2g-2$ and $2g-2+m \geqslant 6$. Let $n \geqslant 2g-2+m$, so in particular $n \geqslant 6$ and we can apply Theorem 4.1. There exist permutations $\sigma, \pi \in C(n) \subset S_n$ with $[\sigma,\pi] \in C(k_1+1, \ldots, k_m+1, 1, \ldots, 1)$ where the number of 1's is equal to $n-2g+2-m$. Hence, by Lemma 1.17, π and σ give rise to a $(1,1)$-square-tiled surface X with exactly n squares that is contained in the stratum $\mathcal{H}(k_1, \ldots, k_m)$.

It remains to check the case where $2g-2+m \leq 5$. This corresponds to the strata $\mathcal{H}(4), \mathcal{H}(2), \mathcal{H}(1,1)$ and $\mathcal{H}(0)$.

The case of $\mathcal{H}(4)$: for $n = 5$, the two permutations $\sigma = (1\,2\,3\,4\,5)$ and $\pi = (1\,2\,4\,5\,3)$ satisfy $[\sigma,\pi] = (1\,4\,2\,3\,5) \in C(5)$. Therefore, σ and π give rise to a $(1,1)$-square-tiled surface $X \in \mathcal{H}(4)$ consisting of 5 squares. For $n \geqslant 6$, Theorem 4.1 says that there exist $\sigma, \pi \in C(n)$ with $[\sigma,\pi] \in C(5,1,\ldots,1)$ where the number of 1's is $n-5$. These σ and π give rise to a $(1,1)$-square-tiled surface as desired.

The case of $\mathcal{H}(2)$: for $n = 4$, the permutations $\sigma = (1\,2\,3\,4)$ and $\pi = (1\,2\,4\,3)$ satisfy $[\sigma,\pi] = (1\,3\,4)\,(2) \in C(3,1)$ and hence give rise to a $(1,1)$-square-tiled surface in $\mathcal{H}(2)$

with 4 squares. For $n = 5$, the permutations $\sigma = (1\,2\,3\,4\,5)$ and $\pi = (1\,2\,4\,3\,5)$ satisfy $[\sigma, \pi] \in C(3,1,1)$ and hence give rise to the desired $(1,1)$-square-tiled surface. For $n \geqslant 6$, it follows from Theorem 4.1 that permutations with the desired property exist.

The case of $\mathcal{H}(1,1)$: we only claim that one can build $(1,1)$-square-tiled surfaces in $\mathcal{H}(1,1)$ with $n \geqslant 6$ squares, so the result follows from Theorem 4.1.

The case of $\mathcal{H}(0)$: for $n = 1$, a single square with opposite sides identified gives the desired square-tiled surface. For $n \geqslant 2$, consider the permutations $\sigma = \pi = (0\,1\,2 \ldots n-1) \in C(n)$. Then $[\sigma, \pi] = [\sigma, \sigma] = \mathrm{id} \in C(1, \ldots, 1)$. Therefore, σ and π give rise to a $(1,1)$-square-tiled surface $X \in \mathcal{H}(0)$ with n squares. □

§ 5.3.3 | Proof of Theorem 3

Recall that for a stratum $\mathcal{H} = \mathcal{H}(k_1, \ldots, k_m)$ we let $|\mathcal{H}| := k_1 + \ldots + k_m + m$.

Theorem (3). *Let $n \in \mathbb{N}$. Then the set*

$$\bigcap_{|\mathcal{H}| \leq n} \Lambda_{\mathcal{H}}$$

is infinite. Equivalently, any finite intersection of affine stretch factor sets is infinite.

Proof. Let $n \geqslant 6$. It suffices to show the Theorem for such n since this implies the statement for smaller n as well. Theorem 2 implies that there is a $(1,1)$-square-tiled surface with exactly n squares $X_{\mathcal{H}} \in \mathcal{H}$ for every stratum $\mathcal{H}$ with $|\mathcal{H}| \leq n$. On each $X_{\mathcal{H}}$ there is a filling pair of curves $\alpha_{\mathcal{H}}, \beta_{\mathcal{H}}$ with $i(\alpha_{\mathcal{H}}, \beta_{\mathcal{H}}) = n$. Thus, we can apply the Thurston construction (cf. Theorem 5.2 and the subsequent remark) to each such pair and obtain representations $\rho_{\mathcal{H}} : \langle T_{\alpha_{\mathcal{H}}}, T_{\beta_{\mathcal{H}}} \rangle \to \mathrm{PSL}_2(\mathbb{R})$ with

$$T_{\alpha_{\mathcal{H}}} \mapsto \begin{pmatrix} 1 & -n \\ 0 & 1 \end{pmatrix}, \quad T_{\beta_{\mathcal{H}}} \mapsto \begin{pmatrix} 1 & 0 \\ n & 1 \end{pmatrix}.$$

The important property is that all these representations have the same image. Call the image $\mathcal{I} \subset \mathrm{PSL}_2(\mathbb{R})$. Following Section 5.2, a hyperbolic matrix $A \in \mathcal{I}$ corresponds to at least one pseudo-Anosov $f_{\mathcal{H}} \in \langle T_{\alpha_{\mathcal{H}}}, T_{\beta_{\mathcal{H}}} \rangle$ that is affine for $\mathcal{H}$. The stretch factor of $f_{\mathcal{H}}$ is λ, the larger of the two eigenvalues of A. Hence,

$$\lambda \in \bigcap_{|\mathcal{H}| \leq n} \Lambda_{\mathcal{H}}.$$

Then for any power k, the mapping class $f_{\mathcal{H}}^k$ is a pseudo-Anosov that is affine for $\mathcal{H}$ with

stretch factor λ^k. Hence,

$$\{\lambda^k \mid k \geqslant 1\} \subset \bigcap_{|\mathcal{H}| \leq n} \Lambda_{\mathcal{H}}$$

which implies that the latter set is infinite. □

Remark 5.4. The proof shows that by analysing the hyperbolic matrices in $\mathcal{I}$, one can find many more stretch factors in $\bigcap_{|\mathcal{H}| \leq n} \Lambda_{\mathcal{H}}$, which are not just powers of λ.

§ 5.3.4 | Proof of Theorem 4

Recall that $m_g := \sup_{\lambda \in \Lambda_g} m_\lambda$, where m_λ is the multiplicity of a stretch factor λ.

Theorem (4). *For any $g \geqslant 1$, it holds that*

$$m_g = \infty.$$

Remark 5.5. In the torus case $g = 1$, we have that $\mathrm{Mod}(S_1) \cong \mathrm{SL}_2(\mathbb{Z})$ and m_λ is the number of conjugacy classes of hyperbolic matrices with eigenvalue λ. Counts of these have been studied by number theorists before (see, for instance, [CCC80]), thus yielding a proof of Theorem 4 in the $g = 1$ case. One can then use a branched cover to lift mapping classes of some finite-index subgroup of $\mathrm{Mod}(S_1)$ to a higher-genus surface S_g (cf. [FM12, Section 14.1.1]). In particular, a hyperbolic matrix lifts to a pseudo-Anosov whose stretch factor is an eigenvalue of the matrix. Hence, analysing conjugacy classes of finite-index subgroups of $\mathrm{SL}_2(\mathbb{Z})$ leads to a proof of Theorem 4.

We want to present a proof sketch which uses filling pairs of curves and the Thurston construction. In particular, we make use of Theorem 1, which in turn relies on Theorem 4.1. The main idea is that we can build pseudo-Anosovs with the same stretch factor by either taking a filling pair with low intersection number and twisting a large amount of times about the curves, or taking a filling pair with high intersection number and only twisting a small amount of times about the curves.

Proof sketch of Theorem 4. Let $g \geqslant 1$ and $N \in \mathbb{N}$. Our goal is to find a $\lambda \in \Lambda_g$ with $m_\lambda \geqslant N$. Since N was chosen arbitrarily, this would be enough to finish the proof.

If $g \geqslant 3$, let $n_1 := 2g - 1$. If $g = 2$, let $n_1 := 4$. If $g = 1$, let $n_1 := 3$. Furthermore, define integers

$$n_1 < n_2 < \ldots < n_N$$

by $n_{i+1} := n_i^2 - 2$ for all $i = 1, \ldots, N - 1$.

By Theorem 1, we know that for each n_i there is a filling pair of curves (α_i, β_i) on S_g

with $i(\alpha_i, \beta_i) = n_i$. Define

$$f_i := T_{\alpha_i} \circ T_{\beta_i} \in \mathrm{Mod}(S_g).$$

By the Thurston construction, the f_i are pseudo-Anosovs and their stretch factors are given by the larger eigenvalue of the matrices

$$\begin{pmatrix} 1 & -n_i \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ n_i & 1 \end{pmatrix} = \begin{pmatrix} 1 - n_i^2 & -n_i \\ n_i & 1 \end{pmatrix} = \begin{pmatrix} n_i^2 - 1 & n_i \\ -n_i & -1 \end{pmatrix}.$$

Note that the last equality holds since we are working in $\mathrm{PSL}_2(\mathbb{R})$.

Hence, for the stretch factors, we obtain

$$\lambda_{f_i} = \frac{1}{2}(n_i^2 - 2 + n_i\sqrt{n_i^2 - 4}).$$

Define $h_i := f_i^{2^{N-i}} \in \mathrm{Mod}(S_g)$, i.e.

$$h_N = f_N, \quad h_{N-1} = f_{N-1}^2, \quad h_{N-2} = f_{N-2}^4, \quad \dots \quad , h_1 = f_1^{2^{N-1}}.$$

As powers of pseudo-Anosovs, the h_i are pseudo-Anosovs as well. Furthermore, their stretch factors are

$$\lambda_{h_i} = \lambda_{f_i}^{2^{N-i}} = \frac{1}{2^{2^{N-i}}}(n_i^2 - 2 + n_i\sqrt{n_i^2 - 4})^{2^{N-i}}.$$

We want to show that all the stretch factors of the h_i are the same:

For arbitrary $n \in \mathbb{N}$, we have

$$\frac{1}{2^2}(n^2 - 2 + n\sqrt{n^2 - 4})^2 = \frac{1}{4}(n^4 - 4n^2 + 4 + 2(n^2 - 2)n\sqrt{n^2 - 4} + n^2(n^2 - 4)) =$$

$$= \frac{1}{2}(n^4 - 4n^2 + 2 + (n^2 - 2)n\sqrt{n^2 - 4}) = \frac{1}{2}((n^2 - 2)^2 - 2 + (n^2 - 2)\sqrt{(n^2 - 2)^2 - 4}).$$

This shows that for any $i \in \{1, \dots, N-1\}$, we have $\lambda_{f_i}^2 = \lambda_{f_{i+1}}$. It follows that

$$\lambda_{h_i} = \lambda_{f_i}^{2^{N-i}} = \lambda_{f_{i+1}}^{2^{N-i-1}} = \lambda_{f_{i+2}}^{2^{N-i-2}} = \cdots = \lambda_{f_N}^{2^{N-i-(N-i)}} = \lambda_{f_N} = \lambda_{h_N}.$$

Hence, letting $\lambda := \lambda_{h_N}$, we have now constructed N pseudo-Anosovs $h_1, \dots, h_N \in \mathrm{Mod}(S_g)$ with stretch factor λ.

In order to conclude that $m_\lambda \geqslant N$ and finish the proof, it only remains to show that the h_i are pairwise non-conjugate. We only sketch a possible approach to proving this. Note that, in fact, it would be enough to show that there is a subset of the h_i which are pairwise non-conjugate as long as the size of this subset tends to infinity with N.

Let $i \neq j \in \{1, \dots, N\}$. Then $h_i = T_{\alpha_i} T_{\beta_i} \cdots T_{\alpha_i} T_{\beta_i}$ and $h_j = T_{\alpha_j} T_{\beta_j} \cdots T_{\alpha_j} T_{\beta_j}$ where the products $T_{\alpha_i} T_{\beta_i}$ and $T_{\alpha_j} T_{\beta_j}$ appear 2^{N-i} and 2^{N-j} times respectively in the equalities. Note that conjugating h_j by any $f \in \mathrm{Mod}(S_g)$ yields

$$f h_j f^{-1} = f T_{\alpha_j} f^{-1} f T_{\beta_j} f^{-1} f \cdots f^{-1} f T_{\alpha_j} f^{-1} f T_{\beta_j} f^{-1} = T_{f(\alpha_j)} T_{f(\beta_j)} \cdots T_{f(\alpha_j)} T_{f(\beta_j)}.$$

As α_j, β_j, the curves $f(\alpha_j), f(\beta_j)$ also form a filling pair with intersection number n_j. Hence, to finish the proof, it suffices to show the following conjecture:

Conjecture. Let (α, β) and (γ, δ) be filling pairs of curves on S_g with $i(\alpha, \beta) \neq i(\gamma, \delta)$. Let $U_1 := \langle T_\alpha T_\beta \rangle$, $U_2 := \langle T_\gamma T_\delta \rangle \subset \mathrm{Mod}(S_g)$, which are two subgroups, both isomorphic to $\mathbb{Z}$ that consist entirely of pseudo-Anosovs (and the identity). Then, it holds that

$$U_1 \cap U_2 = \{\mathrm{id}\}.$$

We remark that the conjecture is stronger than what is needed to finish the proof, but it is of independent interest as a property of mapping class groups. □

§ 5.4 | Generalisations and future directions

We discussed the correspondence between filling pairs of curves and $(1,1)$-square-tiled surfaces. The latter are characterised by a pair of n-cycles. In general, a STS gives rise to two multicurves $A = \{\alpha_1, \dots, \alpha_k\}$ and $B = \{\beta_1, \dots, \beta_l\}$. Such a STS is defined by two permutations $\pi_h, \pi_v \in S_n$ where n is the number of squares of the STS and

$$\pi_h \in C(i_1, \dots, i_k),\ \pi_v \in C(j_1, \dots, j_l),$$

for some partitions $i_1 + \dots + i_k = n = j_1 + \dots + j_l$.

We say that X is a (k,l)-STS if it consists of exactly k horizontal and l vertical cylinders. A generalisation of Theorem 2 can be obtained by answering the following question.

Question 5.6. Given a stratum $\mathcal{H}(k_1, \dots, k_m)$, for which $n \in \mathbb{N}$ does there exist a (k,l)-square-tiled surface $X \in \mathcal{H}(k_1, \dots, k_m)$ with exactly n squares?

To answer this question, one would have to derive a similar statement about commutators as in Chapter 4. This time, however, instead of analysing a single map, one needs to consider all the maps

$$C(i_1, \dots, i_k) \times C(j_1, \dots, j_l) \to A_n,\ (\pi, \tau) \mapsto [\pi, \tau],$$

for any partitions $i_1 + \dots + i_k = n = j_1 + \dots + j_l$. Determining the images of all such

commutator maps would give a precise answer to the above question.

Another topic to address would be the question of how many square-tiled surfaces with a given property exist. Theorem 2 only concerns the existence of $(1,1)$-STS in a given stratum with a given number of squares. One can consider a more refined question:

Question 5.7. Given $\mathcal{H}(k_1, \dots, k_m)$ and a number $n \in \mathbb{N}$, how many $(1,1)$-square-tiled surfaces $X \in \mathcal{H}(k_1, \dots, k_m)$ exist that have exactly n squares?

Of course, we could consider the even more general problem and replace $(1,1)$ by (k,l) in the above question, but that seems out of reach at the moment.

The most interesting case in Question 5.7 is when n is the minimal number for which there exists a $(1,1)$-STS with n squares in $\mathcal{H}(k_1, \dots, k_m)$, i.e. $n = k_1 + \dots + k_m + m$. Aougab–Menasco–Nieland consider the minimal stratum $\mathcal{H}(2g-2)$ and show that there exist at least $(g-2)!$ many $(1,1)$-STS $X \in \mathcal{H}(2g-2)$ with exactly $2g-1$ squares [AMN22, Theorem 1.2]. It would be interesting if one could obtain similar results in other strata as well.

Answering the two above questions would also yield information about filling pairs of (multi)curves. For example, Aougab–Menasco–Nieland's result implies that there are at least $(g-2)!$ distinct mapping class group orbits of filling pairs of curves whose intersection number is $2g-1$. Since these filling pairs are built out of square-tiled surfaces, they have the property that their geometric intersection number equals the absolute value of their algebraic intersection number.

In general, if one starts with a filling pair of multicurves, the dual graph yields a square tiling of the surface, which either yields a STS, i.e. a translation surface, or a half-translation surface. In order to obtain a general result about filling pairs of multicurves, one would also have to answer questions similar to the above ones about half-translation surfaces that are built out of squares.

Regarding Theorem 3, we showed that for any finite collection of strata, there is a number $\lambda \in \mathbb{R}$ that is contained in all the corresponding stretch factor sets. Since this implies that all the powers of λ are also contained in all the stretch factor sets, we concluded that the intersection of these sets is infinite. As noted in the remark after the proof of Theorem 3, we can propose the task of constructing more stretch factors in the intersection.

We say that the real numbers $\lambda_1, \dots, \lambda_k$ are multiplicatively independent if

$$\lambda_1^{n_1} \cdot \dots \cdot \lambda_k^{n_k} = 1$$

implies $n_1 = \dots = n_k = 0$ for any integers $n_1, \dots, n_k$.

Question 5.8. Let $\mathcal{H}_1, \ldots, \mathcal{H}_n$ be strata of translation surfaces. How many multiplicatively independent stretch factors are contained in $\bigcap_{i=1}^{n} \Lambda_{\mathcal{H}_i}$?

Finally, one can also put number-theoretic restrictions on the stretch factors considered and ask the above question with these extra restrictions. It is known that stretch factors are *bi-Perron numbers* [Fri85, Theorem 1]. A prominent open question in the field is whether every bi-Perron number (or at least a power of it) is the stretch factor of some pseudo-Anosov [Mar19, Question 10.3]. Other classes of numbers that have been studied with regard to stretch factors are the so-called *Pisot* or *Salem numbers* - which are both examples of Perron numbers (see, for example [Pan20]).

Chapter 6

Pants graph translation lengths

In this chapter, we discuss translation lengths of pseudo-Anosovs acting on the pants graph. In particular, we are interested in the minimal such translation length and want to see how these translation lengths compare to the ones on Teichmüller space and on the curve graph.

Let S_g be a closed, oriented surface of genus $g \geqslant 2$. Recall that the pants graph $\mathcal{P}(S_g)$ was defined as the graph whose vertices correspond to pants decompositions of the surface, and whose edges correspond to elementary moves (see Section 1.4). Contrary to the curve graph, the pants graph is not a hyperbolic graph [BF06, Theorem 4.1]. It is, however, connected and of infinite diameter. Connectedness was shown in [HLS00], and infinite diameter follows, for example, from [Bro03].

Throughout this chapter, we use the same notation as in Chapter 1. In particular, $\mathrm{Mod}(S_g)$ denotes the mapping class group of S_g. We will consider the Teichmüller space $\mathcal{T}(S_g)$, the curve graph $\mathcal{C}(S_g)$ and the pants graph $\mathcal{P}(S_g)$. The Teichmüller metric will be denoted by $d_{\mathcal{T}}$ and the shortest path metric on the above graphs by $d_{\mathcal{C}}$ and $d_{\mathcal{P}}$ respectively. The (stable) translation length of a mapping class $f \in \mathrm{Mod}(S_g)$ acting on either of the above spaces will be denoted by $l_{\mathcal{T}}(f), l_{\mathcal{C}}(f)$ or $l_{\mathcal{P}}(f)$ respectively.

The same argument as for Teichmüller space or the curve graph shows that the translation length $l_{\mathcal{P}}(f)$ is 0 if f is a periodic mapping class. For a pseudo-Anosov f, we have the analogous statement as in the curve graph case: there exists a power of f that preserves a geodesic and translates along it [Irm15]. Consequently, $l_{\mathcal{P}}(f)$ is positive and rational. For a reducible mapping class, like in the Teichmüller case, it depends on whether the mapping class has pseudo-Anosov components or not. If it does, then $l_{\mathcal{P}}(f) > 0$, if it does not, then $l_{\mathcal{P}}(f) = 0$.

In what follows, we restrict to translation lengths of pseudo-Anosovs. We are mostly interested in what the minimal translation length is, i.e. the value of

$$l_g^{\mathcal{P}} := \inf \{ l_{\mathcal{P}}(f) \mid f \in \mathrm{Mod}(S_g) \text{ pseudo-Anosov} \}.$$

To compare Teichmüller and curve graph translation lengths, we saw that it is useful to consider the systole map, which is a coarsely Lipschitz map between the two spaces:

$$\mathrm{sys} : \mathcal{T}(S_g) \to \mathcal{C}(S_g).$$

Gadre–Hironaka–Kent–Leininger showed that for the optimal Lipschitz constant κ_g of the systole map, it holds that $\kappa_g \asymp \frac{1}{\log(g)}$ [GHKL13].

We use a similar strategy to compare pants graph translation lengths to the ones of $\mathcal{T}(S_g)$ and $\mathcal{P}(S_g)$ by considering maps

$$p : \mathcal{P}(S_g) \to \mathcal{C}(S_g) \text{ and } q : \mathcal{T}(S_g) \to \mathcal{P}(S_g),$$

which we show are coarsely Lipschitz. The map p is the map that sends a pants decomposition to one of its curves, introduced in Section 1.4. The map q is the one used by Brock to show that $(\mathcal{P}(S_g), d_{\mathcal{P}})$ is quasi-isometric to $(\mathcal{T}(S_g), d_{WP})$, where d_{WP} stands for the Weil-Petersson metric. We describe the map q in Section 6.4.

§ 6.1 | Statements of the results

Let $p : \mathcal{P}(S_g) \to \mathcal{C}(S_g)$ be the map that sends a pants decomposition to one of its curves and μ_g be its optimal Lipschitz constant, i.e.

$$\mu_g := \inf \{ C > 0 \mid p \text{ is } (C, K)\text{-coarsely Lipschitz for some } K > 0 \}.$$

Let $l_g^{\mathcal{P}}$ be the minimal pseudo-Anosov pants graph translation length as defined above.

Our main goal in this section is to produce bounds on the optimal Lipschitz constant μ_g of the map p and on the minimal translation length $l_g^{\mathcal{P}}$. In the following theorems, the symbol $\lesssim$ stands for "less than or equal in asymptotic behaviour", i.e. for two sequences $(a_n), (b_n)$, we have $a_n \lesssim b_n$, if $a_n \asymp b_n$ or $\frac{a_n}{b_n} \to 0$ as $n \to \infty$.

Theorem 6.1. *The map $p : \mathcal{P}(S_g) \to \mathcal{C}(S_g)$ is coarsely Lipschitz. The optimal Lipschitz constant μ_g satisfies*

$$\frac{1}{g^2} \lesssim \mu_g \lesssim \frac{1}{g}.$$

For the minimal translation length, we have:

Theorem 6.2. *There is a constant $C > 0$ such that the minimal pseudo-Anosov pants graph translation length $l_g^{\mathcal{P}}$ satisfies*

$$\frac{1}{g} \lesssim l_g^{\mathcal{P}} \lesssim C.$$

In fact, the two theorems are intertwined since knowing the value of μ_g gives us a way to compare pants graph translation lengths with curve graph ones. Namely, in Lemma 6.7, we show that for any $f \in \mathrm{Mod}(S_g)$, it holds that

$$l_{\mathcal{C}}(f) \leqslant \mu_g \cdot l_{\mathcal{P}}(f).$$

Gadre–Tsai showed that the minimal pseudo-Anosov curve graph translation length $l_g^{\mathcal{C}}$ is asymptotic to $\frac{1}{g^2}$. Therefore, the above inequality implies that the upper bound for μ_g in Theorem 6.1 is equivalent to the lower bound of $l_g^{\mathcal{P}}$ in Theorem 6.2 and vice versa.

We observe that Theorem 6.2 shows that the asymptotic behaviour of the minimal pants graph translation length is strictly bigger than the analogous one in the case of the curve graph and at least as big as in the case of Teichmüller space. It would be interesting if equality with the upper bound holds, i.e. if $l_g^{\mathcal{P}} \asymp C$, as this would mean that contrary to Teichmüller space or the curve graph, the minimal translation lengths on the pants graph do not tend to 0 as the genus increases.

Brock showed that the pants graph is quasi-isometric to Teichmüller space equipped with the Weil-Petersson metric [Bro03, Theorem 1.1]. Let

$$q : \mathcal{T}(S_g) \to \mathcal{P}(S_g)$$

be the map Brock considers. We describe this map in detail in Section 6.4 and show that it is coarsely Lipschitz when $\mathcal{T}(S_g)$ is equipped with the Teichmüller metric.

Lemma 6.3. *There exist constants $C_g, K_g \geqslant 0$ such that*

$$d_{\mathcal{P}}(q(X,\varphi), q(Y,\psi)) \leqslant C_g \, d_{\mathcal{T}}((X,\varphi),(Y,\psi)) + K_g$$

for all $(X,\varphi),(Y,\psi) \in \mathcal{T}(S_g)$.

Finally, we say that two maps $f_1, f_2 : X \to Y$ between two metric spaces are *coarsely equivalent* and write $f_1 \approx f_2$ if there is a $K \geqslant 0$ such that

$$d_Y(f_1(x), f_2(x)) \leqslant K$$

for all $x \in X$.

We will observe that the above defined maps sys, p and q fit into a coarsely commuting diagram.

Lemma 6.4. *The following is a diagram of coarsely equivalent maps*

$$\begin{array}{ccc} \mathcal{T}(S_g) & & \\ \downarrow{\scriptstyle q} & \searrow^{\text{sys}} & \\ \mathcal{P}(S_g) & \xrightarrow{p} & \mathcal{C}(S_g), \end{array}$$

i.e. $p \circ q \approx \text{sys}$.

Let ν_g be the optimal Lipschitz constants of the map q. We obtain:

Corollary 6.5. *The above coarsely commuting diagram implies that*

$$\kappa_g \leqslant \nu_g \cdot \mu_g.$$

In particular,

$$\nu_g \gtrsim \frac{g}{\log(g)}.$$

Outline. In Section 6.2, we discuss the map p, show that it is coarsely Lipschitz and prove the upper bound in Theorem 6.1. As mentioned above, this simultaneously shows the lower bound in Theorem 6.2. In Section 6.3, we prove the upper bound in Theorem 6.2 by constructing an explicit example of a pseudo-Anosov. Finally, in Section 6.4, we discuss the quasi-isometry between the pants graph and Teichmüller space with the Weil-Petersson metric and prove Lemmata 6.3 and 6.4, as well as Corollary 6.5.

§ 6.2 | The map from pants graph to curve graph

Let $\mathcal{X} = \{\alpha_1, \ldots, \alpha_{3g-3}\} \in \mathcal{P}(S_g)$ be a pants decomposition. The map

$$p : \mathcal{P}(S_g) \to \mathcal{C}(S_g)$$

is defined by $p(\mathcal{X}) = \alpha_i \in \mathcal{C}(S_g)$ for some i. As was the case for the systole map, we could instead view p as a set-valued map that sends a pants decomposition to the set of its curves in the curve graph. Note that this set forms a complete subgraph as all the curves

are pairwise disjoint. However, since we are only interested in the coarse behaviour of p, we can think of it as mapping $\mathcal{X}$ to one of its curves α_i, and the choice of i won't matter.

Recall that a map $f : X \to Y$ between two metric spaces is coarsely Lipschitz if there exist $C, K > 0$ such that

$$d_Y(f(x_1), f(x_2)) \leqslant C d_X(x_1, x_2) + K$$

for all $x_1, x_2 \in X$.

Lemma 6.6. *The map $p : \mathcal{P}(S_g) \to \mathcal{C}(S_g)$ is coarsely Lipschitz.*

Proof. We show that the constants C and K in the definition of coarsely Lipschitz can be chosen to be $\frac{2}{3g-4}$ and 2, respectively. The main observation is that whenever we perform $3g-4$ elementary moves, i.e. have two pants decompositions $\mathcal{X}, \mathcal{Y}$ with $d_{\mathcal{P}}(\mathcal{X}, \mathcal{Y}) \leqslant 3g-4$, then their images satisfy $d_{\mathcal{C}}(p(\mathcal{X}), p(\mathcal{Y})) \leqslant 2$. This is because there are $3g-3$ curves in a pants decomposition, hence performing $3g-4$ moves will leave at least one curve unchanged, say this curve is $x_1 \in \mathcal{X} \cap \mathcal{Y}$. Hence,

$$d_{\mathcal{C}}(p(\mathcal{X}), p(\mathcal{Y})) \leqslant d_{\mathcal{C}}(p(\mathcal{X}), x_1) + d_{\mathcal{C}}(x_1, p(\mathcal{Y})) \leqslant 1 + 1 = 2.$$

Now, let $\mathcal{X}, \mathcal{Y}$ be two arbitrary pants decompositions and let $n := d_{\mathcal{P}}(\mathcal{X}, \mathcal{Y})$. Write $n = d(3g-4) + r$ with $r < 3g-4$. Let $\mathcal{X} = \mathcal{X}_0, \mathcal{X}_1, ..., \mathcal{X}_{d+1} = \mathcal{Y}$ be a sequence of pants decompositions with $d_{\mathcal{P}}(\mathcal{X}_i, \mathcal{X}_{i+1}) \leqslant 3g-4$ for all i. We obtain

$$d_{\mathcal{C}}(p(\mathcal{X}), p(\mathcal{Y})) \leqslant \sum_{i=0}^{d} d_{\mathcal{C}}(p(\mathcal{X}_i), p(\mathcal{X}_{i+1})) \leqslant 2(d+1) = 2d + 2 \leqslant \frac{2n}{3g-4} + 2,$$

which proves the claim. □

Note that the proof implies that the optimal Lipschitz constant μ_g satisfies

$$\mu_g \leqslant \frac{2}{3g-4},$$

i.e. $\mu_g \lesssim \frac{1}{g}$.

§ 6.3 | Translation lengths of pseudo-Anosovs

The mapping class group $\mathrm{Mod}(S_g)$ acts on both $\mathcal{P}(S_g)$ and $\mathcal{C}(S_g)$ by isometries.

Recall that for an isometry f of a metric space X we defined the (stable) translation

length as

$$l(f) \coloneqq \liminf_{n\to\infty} \frac{d_X(f^n(x), x)}{n},$$

where $x \in X$ is any element.

For any $f \in \mathrm{Mod}(S_g)$, we can measure the translation lengths $l_{\mathcal{P}}(f)$ and $l_{\mathcal{C}}(f)$ corresponding to the action on the pants and the curve graph, respectively. The two translation lengths are related by the Lipschitz constant μ_g.

Lemma 6.7. *For any $f \in \mathrm{Mod}(S_g)$, it holds that*

$$l_{\mathcal{C}}(f) \leqslant \mu_g \cdot l_{\mathcal{P}}(f).$$

Proof. Let $\varepsilon > 0$. Then there are $C, K > 0$ with $C < \mu_g + \varepsilon$ such that

$$d_{\mathcal{C}}(p(\mathcal{X}), p(\mathcal{Y})) \leqslant C d_{\mathcal{P}}(\mathcal{X}, \mathcal{Y}) + K,$$

for all $\mathcal{X}, \mathcal{Y} \in \mathcal{P}(S_g)$.

Let $f \in \mathrm{Mod}(S_g)$. Note that the map p is *coarsely equivariant* in the sense that

$$d_{\mathcal{C}}(h \cdot p(\mathcal{X}),\, p(h \cdot \mathcal{X})) \leqslant 1$$

for all $h \in \mathrm{Mod}(S_g)$. It follows that

$$d_{\mathcal{C}}(f^n(p(\mathcal{X})), p(\mathcal{X})) \leqslant d_{\mathcal{C}}(p(f^n(\mathcal{X})), p(\mathcal{X})) + 1 \leqslant C d_{\mathcal{P}}(f^n(\mathcal{X}), \mathcal{X}) + K + 1$$

for any $\mathcal{X} \in \mathcal{P}(S_g)$. Dividing by n and taking the lim inf on both sides yields

$$l_{\mathcal{C}}(f) \leqslant C \cdot l_{\mathcal{P}}(f).$$

Since $\varepsilon > 0$ was arbitrary, we obtain the claim of the lemma. □

We remark that knowing that $l_{\mathcal{C}}(f) > 0$ for any pseudo-Anosov f, the above lemma can be used to show that also $l_{\mathcal{P}}(f) > 0$.

For the minimal pseudo-Anosov pants graph translation length $l_g^{\mathcal{P}}$, Lemma 6.7, together with the bound on μ_g we obtained, implies

$$l_g^{\mathcal{P}} \gtrsim \frac{l_g^{\mathcal{C}}}{\mu_g} \gtrsim \frac{g}{g^2} = \frac{1}{g}.$$

In order to prove the upper bound for $l_g^{\mathcal{P}}$, we construct explicit examples of pseudo-Anosovs.

Proposition 6.8. *For any $g \geqslant 3$, there exists a pseudo-Anosov $F_g \in \mathrm{Mod}(S_g)$ and a constant $C > 0$ such that $l_{\mathcal{P}}(F_g) \leqslant C$.*

Proof. Consider the surface S_g as pictured in Figure 6.1. Let $\rho \in \mathrm{Mod}(S_g)$ be the order $g-1$ mapping class that arises through anti-clockwise rotation of S_g. Let $\alpha_1, \alpha_2, \alpha_3$ and β be the curves pictured in Figure 6.1. For $i \in \{1,2,3\}$, let $A_i := \{\rho^n(\alpha_i) \,|\, n = 1, \dots, g-1\}$, and let $B := \{\rho^n(\beta) \,|\, n = 1, \dots, g-1\}$. Finally, let $A := A_1 \cup A_2 \cup A_3$ and note that A and B are multicurves whose union fills the surface.

We define $F_g \in \mathrm{Mod}(S_g)$ to be the composition

$$F_g := \rho \circ T_\beta^{-1} \circ T_{\alpha_1} \circ T_{\alpha_2} \circ T_{\alpha_3}.$$

We have that

$$F_g^{g-1} = \prod_{n=1}^{g-1} T_{\rho^n(\beta)}^{-1} \circ T_{\rho^n(\alpha_1)} \circ T_{\rho^n(\alpha_2)} \circ T_{\rho^n(\alpha_3)},$$

which is pseudo-Anosov due to [Pen88, Theorem 3.1] applied to the sets A and B. Hence, also F_g is pseudo-Anosov.

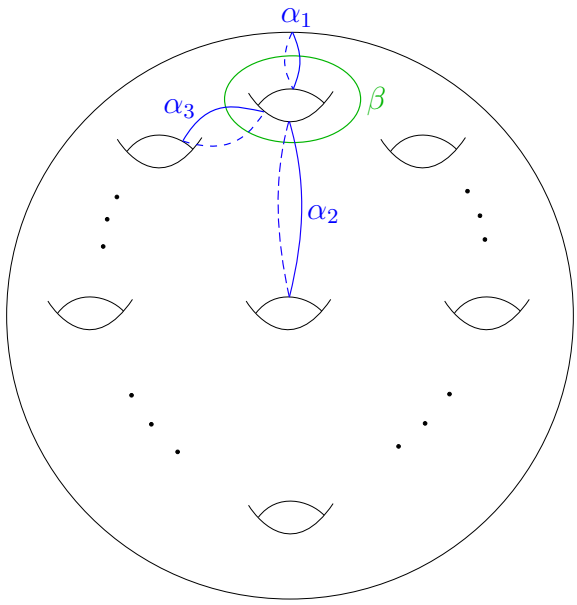


Figure 6.1: The pseudo-Anosov $F_g = \rho \circ T_\beta^{-1} \circ T_{\alpha_1} \circ T_{\alpha_2} \circ T_{\alpha_3}$.

We want to compute the pants graph translation length of F_g. Let $\gamma_1, \gamma_2, \gamma_3$ be the curves portrayed in Figure 6.2 and $\Gamma := \{\gamma_1, \gamma_2, \gamma_3\}$. Consider the pants decomposition

$$\mathcal{X}_g = \bigcup_{n=1}^{g-1} \rho^n(\Gamma) \in \mathcal{P}(S_g),$$

which can also be seen in Figure 6.2. Note that $\gamma_1 = \alpha_1$ and $\gamma_2 = \alpha_2$.

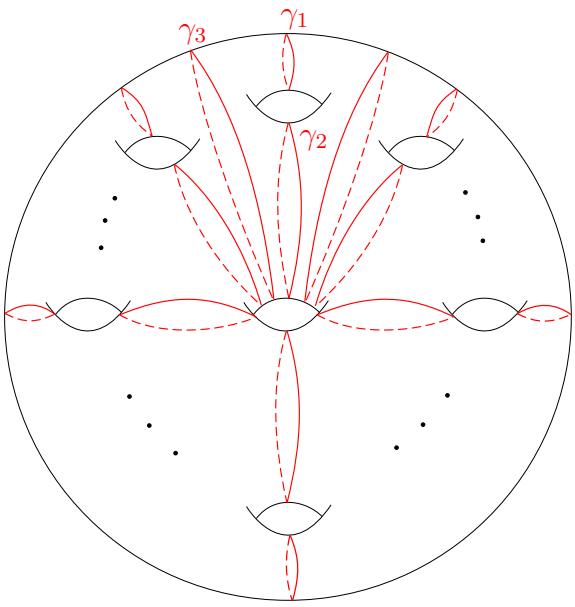


Figure 6.2: The pants decomposition $\mathcal{X}_g$.

When applying F_g to the pants decomposition $\mathcal{X}_g$, any curve $\gamma \in \mathcal{X}_g \setminus \Gamma$ gets mapped to $\rho(\gamma)$. The curves $\gamma_1, \gamma_2, \gamma_3$ are the only ones affected by the twists. Hence, the pants decomposition $F_g(\mathcal{X}_g)$ differs from $\mathcal{X}_g$ by three curves which are contained in a subsurface homeomorphic to a torus with three boundary components, bounded by the curves γ_1, γ_2 and $\rho(\gamma_3)$. Concretely,

$$F_g(\mathcal{X}_g) = (\mathcal{X}_g \setminus \rho(\Gamma)) \cup \{\rho(T_\beta^{-1}(\gamma_1)), \rho(T_\beta^{-1}(\gamma_2)), \rho(T_\beta^{-1}(T_{\alpha_3}(\gamma_3)))\}.$$

Since the pants graph of a torus with three boundary components is connected [HLS00, Theorem D], there is a number of elementary moves $C > 0$ one can perform inside the subsurface to modify the three new curves back to the original ones, i.e. $d_{\mathcal{P}}(\mathcal{X}_g, F_g(\mathcal{X}_g)) = C$.

It now follows that

$$l_{\mathcal{P}}(F_g) = \liminf_{n\to\infty} \frac{d_{\mathcal{P}}(\mathcal{X}_g, F_g^n(\mathcal{X}_g))}{n} \leqslant \liminf_{n\to\infty} \frac{n \cdot C}{n} = C.$$

□

The important observation in the above proposition is that the constant C is independent of the genus g. Hence, it follows that

$$l_g^{\mathcal{P}} \leqslant C,$$

for all $g \geqslant 3$, and in particular $l_g^{\mathcal{P}} \lesssim C$. Using Lemma 6.7 again, we also obtain

$$\mu_g \gtrsim \frac{l_g^{\mathcal{C}}}{l_g^{\mathcal{P}}} \gtrsim \frac{1}{g^2}.$$

This finishes the proofs of Theorems 6.1 and 6.2.

§ 6.4 | The map from Teichmüller space to pants graph

The Weil-Petersson metric d_{WP} on Teichmüller space $\mathcal{T}(S_g)$ is – in contrast to the Teichmüller metric $d_{\mathcal{T}}$ – a Riemannian metric. It is induced by a Hermitian inner product on the vector space $\mathcal{Q}(X)$ of quadratic differentials of a Riemann surface X, which is identified with the cotangent space of $\mathcal{T}(S_g)$ at X. Concretely, this inner product is defined as

$$\langle q_1, q_2 \rangle := \int_X \frac{q_1 \bar{q}_2}{\rho^2},$$

where $\rho(z)|dz|$ is the hyperbolic metric on X (see [Wol09] for details on the Weil-Petersson metric). As the Teichmüller metric, the Weil-Petersson metric descends to a metric on moduli space. While $d_{\mathcal{T}}$ is complete, d_{WP} is not.

Definition 6.9. *Two metric spaces $(X, d_X), (Y, d_Y)$ are called* quasi-isometric *if there is a map $f : X \to Y$ and $A \geqslant 1$, as well as $B, C \geqslant 0$, such that*

$$\frac{1}{A} d_X(x_1, x_2) - B \leqslant d_Y(f(x_1), f(x_2)) \leqslant A d_X(x_1, x_2) + B$$

for all $x_1, x_2 \in X$, and such that for any $y \in Y$, there exists $x \in X$ with

$$d_Y(y, f(x)) \leqslant C.$$

The map f satisfying the above conditions is called a quasi-isometry.

As with coarsely Lipschitz maps, we refer to the infimum over all $A \geqslant 1$ for which there exists $B \geqslant 0$ such that the inequalities in the definition are satisfied as the *optimal constant* for the quasi-isometry.

We now want to describe a map

$$q : \mathcal{T}(S_g) \to \mathcal{P}(S_g),$$

which is a quasi-isometry between the two spaces equipped with d_{WP} and $d_{\mathcal{P}}$, respectively.

For a simple closed curve γ on S_g and a point $(X, \varphi) \in \mathcal{T}(S_g)$, let $l_X^{\varphi}(\gamma)$ denote the hyperbolic length of the curve $\varphi(\gamma)$ on X. Bers proved that for any genus $g \geqslant 2$, there is a constant b_g such that for any $(X, \varphi) \in \mathcal{T}(S_g)$ there exists a pants decomposition

$\{\gamma_1, \ldots, \gamma_{3g-3}\} \in \mathcal{P}(S_g)$ with $l_X^{\varphi}(\gamma_i) \leqslant b_g$ for all $i = 1, \ldots, 3g-3$ [Ber85, Theorem 2]. Intuitively, this says that for any hyperbolic surface, we can find a pants decomposition consisting of "short" curves, where by short we mean of length $\leqslant b_g$.

For a given $\mathcal{X} \in \mathcal{P}(S_g)$, let $U(\mathcal{X}) \subset \mathcal{T}(S_g)$ be the set of points for which all curves in $\mathcal{X}$ are short. By Bers' theorem, we know that

$$\bigcup_{\mathcal{X} \in \mathcal{P}(S_g)} U(\mathcal{X}) = \mathcal{T}(S_g).$$

Let $r : \mathcal{P}(S_g) \to \mathcal{T}(S_g)$ be the map that sends $\mathcal{X}$ to some point in $U(\mathcal{X})$. Brock shows that this map is a quasi-isometry for the Weil-Petersson metric on Teichmüller space [Bro03].

Let

$$q : \mathcal{T}(S_g) \to \mathcal{P}(S_g)$$

be a quasi-isometry inverse to r, i.e. a map that sends a point (X, φ) to a pants decomposition $\mathcal{X}$ for which $(X, \varphi) \in U(\mathcal{X})$.

We know that q is a quasi-isometry with respect to the Weil-Petersson metric and want to show that q is coarsely Lipschitz with respect to the Teichmüller metric. For that, we use [Lin74, Theorem in Section 3], which states that

$$d_{WP}((X, \varphi), (Y, \psi)) \leqslant \alpha_g \cdot d_{\mathcal{T}}((X, \varphi), (Y, \psi)),$$

where α_g is the square root of the area of any hyperbolic genus g surface and $(X, \varphi), (Y, \psi) \in \mathcal{T}(S_g)$ are arbitrary. In particular, we have that $\alpha_g \asymp \sqrt{g}$.

Since q is a quasi-isometry, there exist $A_g \geqslant 1, B_g \geqslant 0$ with

$$d_{\mathcal{P}}(q(X, \varphi), q(Y, \psi)) \leqslant A_g d_{WP}((X, \varphi), (Y, \psi)) + B_g,$$

for all $(X, \varphi), (Y, \psi) \in \mathcal{T}(S_g)$.

Putting the two inequalities together, we obtain

$$d_{\mathcal{P}}(q(X, \varphi), q(Y, \psi)) \leqslant A_g \alpha_g d_{\mathcal{T}}((X, \varphi), (Y, \psi)) + B_g,$$

which finishes the proof of Lemma 6.3.

Now, consider the diagram

$$\begin{array}{ccc} \mathcal{T}(S_g) & & \\ \downarrow{\scriptstyle q} & \searrow^{\text{sys}} & \\ \mathcal{P}(S_g) & \xrightarrow{p} & \mathcal{C}(S_g). \end{array}$$

We want to show that this diagram coarsely commutes, i.e. that there is a $K_g \geqslant 0$ such that

$$d_{\mathcal{C}}(\text{sys}(X,\varphi), p \circ q(X,\varphi)) \leqslant K_g,$$

for all $(X,\varphi) \in \mathcal{T}(S_g)$.

Let $(X,\varphi) \in \mathcal{T}(S_g)$ and let $\alpha := \text{sys}(X,\varphi)$ be a systole and $\beta := p \circ q(X,\varphi)$ a curve in a short pants decomposition. Note that $l^{\varphi}_X(\alpha) \leqslant l^{\varphi}_X(\beta) \leqslant b_g$, where b_g is the Bers constant. The Collar Lemma in hyperbolic geometry (see, for instance, [FM12, Lemma 13.6]) implies that the intersection number of α and β is bounded from above by a constant that depends only on their lengths. Since both lengths are $\leqslant b_g$, we can uniformly bound $i(\alpha,\beta)$ by a constant M_g depending only on the genus of S_g. Finally, we can bound the curve graph distance of two curves in terms of their intersection number (see, for example, [Hem01, Lemma 2.1]) and obtain

$$d_{\mathcal{C}}(\alpha,\beta) \leqslant 2\log_2(i(\alpha,\beta)) + 2 \leqslant 2\log_2(M_g) + 2.$$

Setting $K_g := 2\log_2(M_g) + 2$ gives what we wanted and finishes the proof of Lemma 6.4.

We are now ready to prove Corollary 6.5.

Proof of Corollary 6.5. Let κ_g, μ_g, ν_g be the optimal Lipschitz constants of the maps sys, p, q respectively. For any $\varepsilon > 0$, there are $C_p, C_q, B_p, B_q > 0$ with $C_p < \mu_g + \varepsilon$, $C_q < \nu_g + \varepsilon$, such that the pairs $(C_p, B_p), (C_q, B_q)$ satisfy the coarsely Lipschitz condition for p, q respectively.

Let $(X,\varphi), (Y,\psi) \in \mathcal{T}(S_g)$. Then, since $\text{sys} \approx p \circ q$, there is $K_g \geqslant 0$ such that

$$d_{\mathcal{C}}(\text{sys}(X,\varphi), \text{sys}(Y,\psi)) \leqslant d_{\mathcal{C}}(p \circ q(X,\varphi), p \circ q(Y,\psi)) + 2K_g.$$

Since p and q are coarsely Lipschitz, we obtain

$$d_{\mathcal{C}}(p \circ q(X,\varphi), p \circ q(Y,\psi)) \leqslant C_p C_q d_{\mathcal{T}}((X,\varphi),(Y,\psi)) + C_p B_q + B_p.$$

Putting everything together yields

$$d_{\mathcal{C}}(\text{sys}(X,\varphi), \text{sys}(Y,\psi)) \leqslant C_p C_q d_{\mathcal{T}}((X,\varphi),(Y,\psi)) + C_p B_q + B_p + 2K_g,$$

which shows that $\kappa_g \leqslant C_p C_q$. Since $\varepsilon > 0$ was arbitrary, this implies

$$\kappa_g \leqslant \mu_g \nu_g.$$

Finally, we know that $\kappa_g \asymp \frac{1}{\log(g)}$ and saw that $\mu_g \lesssim \frac{1}{g}$, which implies that

$$\nu_g \gtrsim \frac{\kappa_g}{\mu_g} \gtrsim \frac{g}{\log(g)}.$$

□

Remark 6.10. The discussion in this section shows that the optimal quasi-isometry constant A_g of the map q with respect to the Weil-Petersson metric satisfies

$$A_g \gtrsim \frac{\nu_g}{\alpha_g} \gtrsim \frac{g}{\log(g)\sqrt{g}} = \frac{\sqrt{g}}{\log(g)}.$$

§ 6.5 | Future directions

Theorems 6.1 and 6.2 provide bounds on the quantities μ_g and $l_g^{\mathcal{P}}$. The next natural step is to improve these bounds in order to determine the precise asymptotic behaviour of these constants.

Question 6.11. What is the asymptotic behaviour of μ_g and of $l_g^{\mathcal{P}}$?

There are two approaches to this question. The first is to conduct a more careful analysis of the map p to improve the upper bound of μ_g found in Lemma 6.6. The most optimistic one could hope for is showing that pants decompositions of distance g^2 contain curves of distance bounded by a constant. This would imply that $\mu_g \asymp \frac{1}{g^2}$ and consequently, $l_g^{\mathcal{P}} \asymp C$.

The second approach would be to try to improve the upper bound of $l_g^{\mathcal{P}}$ by either studying the example of Proposition 6.8 more thoroughly or coming up with other examples of pseudo-Anosovs of small pants graph translation length. In this case, the most optimistic outcome would be to find pseudo-Anosovs whose pants graph translation length is asymptotic to $\frac{1}{g}$. This would imply that $l_g^{\mathcal{P}} \asymp \frac{1}{g}$ and $\mu_g \asymp \frac{1}{g}$.

We remark that finding a lower bound greater than $\frac{1}{g^2}$ for the curve graph translation length $l_{\mathcal{C}}(F_g)$, where F_g is the pseudo-Anosov from Proposition 6.8, would also improve the lower bound for μ_g.

Similarly to the question above, we can ask about the other constants appearing in this chapter.

Question 6.12. What is the asymptotic behaviour of ν_g and of the optimal quasi-isometry constant in Brock's map?

As we saw, the optimal Lipschitz constants are related to the corresponding translation lengths. For example, we have

$$\frac{l_{\mathcal{C}}(f)}{l_{\mathcal{T}}(f)} \leqslant \kappa_g, \text{ or } \frac{l_{\mathcal{C}}(f)}{l_{\mathcal{P}}(f)} \leqslant \mu_g,$$

for any $f \in \mathrm{Mod}(S_g)$.

A sequence of pseudo-Anosovs $f_g \in \mathrm{Mod}(S_g)$ are called $(\mathcal{T},\mathcal{C})$-*ratio-optimising* if

$$\frac{l_{\mathcal{C}}(f_g)}{l_{\mathcal{T}}(f_g)} \asymp \kappa_g \asymp \frac{1}{\log(g)}.$$

Aougab–Taylor, as well as Jeffreys, show the ubiquity of $(\mathcal{T},\mathcal{C})$-ratio-optimising pseudo-Anosovs ([AT17] and [Jef21]). In particular, Aougab–Taylor construct such pseudo-Anosovs arbitrarily deep in the Johnson filtration, while Jeffreys focuses on finding such pseudo-Anosovs in every connected component of a stratum of abelian differentials.

Note that $(\mathcal{T},\mathcal{C})$-ratio-optimising pseudo-Anosovs cannot achieve the minimal translation length on both Teichmüller space and the curve graph, since in that case the ratio would be $\frac{1}{g}$ instead of $\frac{1}{\log(g)}$.

In an analogous way, we can define $(\mathcal{P},\mathcal{C})$-*ratio-optimising* pseudo-Anosovs f_g by requiring

$$\frac{l_{\mathcal{C}}(f_g)}{l_{\mathcal{P}}(f_g)} \asymp \mu_g.$$

Question 6.13. Is there a sequence of $(\mathcal{P},\mathcal{C})$-ratio-optimising pseudo-Anosovs $f_g \in \mathrm{Mod}(S_g)$ with $l_{\mathcal{C}}(f_g) \asymp \frac{1}{g^2}$ and $l_{\mathcal{P}}(f_g) \asymp L_g^{\mathcal{P}}$?

Conversely, instead of looking for ratio-optimising pseudo-Anosovs, one can ask how big the gap between $l_{\mathcal{C}}(f)$ and $l_{\mathcal{T}}(f)$ or $l_{\mathcal{P}}(f)$ for a pseudo-Anosov f can become. In Chapter 3, it was shown that the difference between $l_{\mathcal{T}}(f)$ and $l_{\mathcal{C}}(f)$ can be arbitrarily large.

We ask for an analogous statement with Teichmüller space replaced by the pants graph.

Question 6.14. Is there a sequence of pseudo-Anosovs $(f_n) \in \mathrm{Mod}(S_g)$ such that $l_{\mathcal{C}}(f_n) \asymp \frac{1}{g^2}$ and $l_{\mathcal{P}}(f_n) \to \infty$ as $n \to \infty$?

Chapter 7

Geminal square-tiled surfaces

Let $\mathcal{H} = \mathcal{H}(k_1, \ldots, k_m)$ be a stratum of translation surfaces of genus g, determined by the equality $2g-2 = k_1 + \ldots + k_m$. Recall from Section 1.3 that there is an action of $\mathrm{GL}_2^+(\mathbb{R})$ on $\mathcal{H}$. This action contains information about the Teichmüller flow on moduli space, while it can also be used to answer questions about dynamics of billiard flows (see, for example, [MT02] or [Smi00]). Trying to understand the $\mathrm{GL}_2^+(\mathbb{R})$-action has caught the attention of many mathematicians and has created a vast field of research. A particular problem that has been studied extensively concerns the classification of orbit closures of the $\mathrm{GL}_2^+(\mathbb{R})$-action on $\mathcal{H}$. As Zorich puts it, the most optimistic hope would be to have a *magic table* which lists all possible orbit closures and contains all the information on their relevant invariants, such as volume, Lyapunov exponents, etc. (see [Zor06] for details on these notions).

For every stratum, there is a natural choice of local coordinates, called period coordinates. An *affine invariant submanifold* is a subspace which is locally given by linear equations in these period coordinates. A breakthrough result from 2015 by Eskin–Mirzakhani–Mohammadi states that orbit closures are affine invariant submanifolds [EMM15]. A year later, Filip proceeded to show that orbit closures are algebraic varieties [Fil16]. Even though these results give us a lot of information about the structure of orbit closures, we are still far from a complete classification.

The *rank* of an orbit closure is an integer invariant computed as follows. One can view the tangent space of the orbit closure at a point (X, ω) as a subspace of the relative cohomology group $H^1(X, Z(\omega); \mathbb{C})$, where $Z(\omega)$ is the set of zeros of ω. Projecting to the absolute cohomology group $H^1(X; \mathbb{C})$ then yields a subspace, which turns out to have even dimension. The rank is given as half of the dimension of this subspace, which is an integer between 1 and g. The two extremes of a $\mathrm{GL}_2^+(\mathbb{R})$-orbit are on the one hand that

its closure is a whole component of the stratum $\mathcal{H}$, in which case the rank is g, or on the other hand that the orbit is already closed, i.e. a Teichmüller curve, in which case the rank is 1. Already the classification of Teichmüller curves turns out to be a very difficult challenge and has only been completed in genus 2 up to now (see [McM22] for a survey on known results about the classification of Teichmüller curves).

Another possibility for an orbit closure is to consist entirely of translation surfaces that arise as branched covers of lower-genus translation surfaces. Mirzakhani conjectured informally that any orbit closure of rank $\geqslant 2$ is either a whole component of a stratum or a locus of branched covers. Apisa–Wright later formalised the idea of consisting entirely of such covers in what they call being a *full locus of covers* of a stratum of abelian or quadratic differentials [AW24]. Although the conjecture turned out to be false as stated above, the new examples constructed to disprove it all have rank equal to 2 [MMW17].

In 2023, Apisa–Wright showed that Mirzakhani's conjecture does hold if the rank is large enough. In particular, the two authors prove that every orbit closure of rank at least $\frac{g}{2} + 1$ is either a connected component of a stratum, or the locus of all holonomy double covers of a stratum of quadratic differentials [AW23]. A main tool in the proof of this result is the study of so-called *geminal* orbit closures. These are orbit closures in which any cylinder on any translation surface in the orbit closure is either *free* or has a *twin* - hence the name geminal (see [AW22] for a concrete definition).

In this chapter, we focus on geminal square-tiled surfaces, which we define in Section 7.1. Apisa–Wright conjectured that every geminal STS is a normal cover of the *four-square torus* with cyclic deck group. This conjecture can be phrased in purely group theoretic terms and was done so in [AW22] as Problem 8.16. The conjecture was disproved by Abdalla–Brown, who constructed an infinite family of non-cyclic 1-cylinder pillowcase-tiled surfaces [AB25]. The translation covers of these surfaces yield examples of non-cyclic geminal STS.

Our main goal is to show that a geminal STS admits a cover to the four-square torus, which is normal if and only if it is cyclic. We state this as Theorem 7.7 and formulate the equivalent statement in a more group-theoretic language as Theorem 7.8 below. The result can be interpreted as saying that one part of the Apisa–Wright conjecture is true, namely, if we assume that the cover is normal, then it has to be cyclic. In particular, the examples by Abdalla–Brown cannot be normal covers.

§ 7.1 | Preliminaries

Recall that a square-tiled surface is a finite-degree branched cover of the torus T, branched over one point p. As such, a STS corresponds to a homomorphism

$$\rho : \pi_1(T \setminus \{p\}) \cong \mathbb{F}_2 \to S_n,$$

where n is the number of preimages of any point in $T \setminus \{p\}$, or equivalently the number of squares of the STS. The homomorphism is given by lifting a based loop of $T \setminus \{p\}$ to the STS and recording the permutation of preimages of the basepoint it induces. We refer to the homomorphism ρ as the *monodromy map*. Throughout, we identify $\pi_1(T \setminus \{p\})$ with $\mathbb{F}_2 = \langle a, b \rangle$ by mapping the standard horizontal and vertical loops, oriented as in Figure 7.1, to a and b, respectively. Note that under this identification, the homomorphism ρ maps the two generators a and b of $\mathbb{F}_2$ to the permutations π_h and π_v of the STS as described in Section 1.3.2.

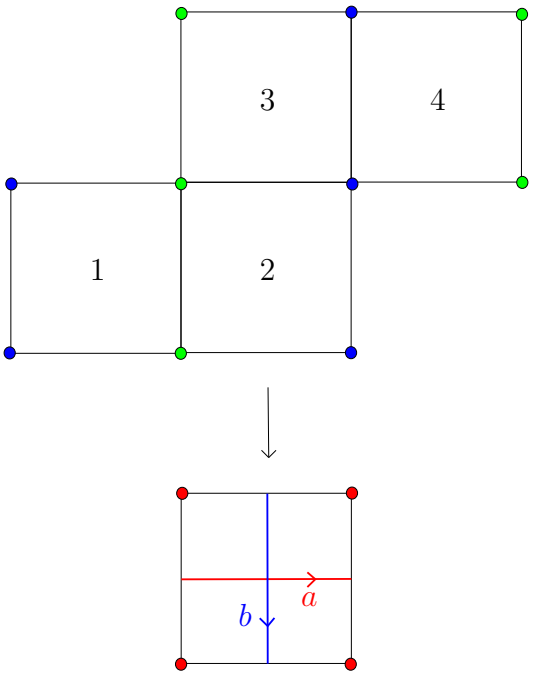


Figure 7.1: A square-tiled surface as a branched cover of the torus corresponding to the map with $\rho(a) = (1\ 2)(3\ 4)$, $\rho(b) = (1)(2\ 3)(4)$. In the figure, it is implicit that parallel opposite sides are glued together.

Throughout this chapter, we make the convention to read elements in the fundamental group from right to left, i.e. the word ba, for example, corresponds to the concatenation of first the horizontal and then the vertical loop. This is done to coincide with the usual way permutations are multiplied together, so to work with $\rho(ba) = \rho(b)\rho(a)$, instead of $\rho(ba) = \rho(a)\rho(b)$.

Let X be a STS with n squares given by the map $\rho\colon \mathbb{F}_2 \to S_n$, and let $X \to T$ be the corresponding branched cover. Throughout this chapter, we assume that X is connected. Let $B = \{p_1, \dots, p_m\} \subset X$ be the set of preimages of $p \in T$. We call B the set of *branch points* on X even though some of the points might be unramified. Note that by Lemma 1.17, the number m is given by the number of cycles of the commutator $[\rho(a), \rho(b)]$.

For $k \in \{1, \dots, n\}$, let $\mathrm{Fix}(k)$ be the subgroup of S_n of all permutations π with $\pi(k) = k$. Note that all the subgroups $\mathrm{Fix}(k)$ are conjugate to each other.

Lemma 7.1. *The fundamental group of $X \setminus B$ is given by the preimage of $\mathrm{Fix}(k)$ under ρ, i.e.*

$$\pi_1(X \setminus B) \cong \rho^{-1}(\mathrm{Fix}(k)),$$

for any $k \in \{1, \dots, n\}$. Furthermore, the cover $X \to T$ is normal if and only if

$$\rho^{-1}(\mathrm{Fix}(1)) = \rho^{-1}(\mathrm{Fix}(2)) = \dots = \rho^{-1}(\mathrm{Fix}(n)) = \ker(\rho).$$

In that case, the deck group of the cover is given by $\mathbb{F}_2/\ker(\rho)$.

Proof. The branched cover $X \to T$ induces an unbranched cover $X \setminus B \to T \setminus \{p\}$. Fix a basepoint $p_0 \in T \setminus \{p\}$. It is standard covering space theory that after fixing a preimage of p_0 as a basepoint in $X \setminus B$, the fundamental group of $X \setminus B$ is isomorphic to the subgroup of $\pi_1(T \setminus \{p\}) \cong \mathbb{F}_2$ of loops that lift to loops. If we choose the basepoint in $X \setminus B$ to be inside square k, then these loops are exactly the elements of $\mathbb{F}_2$ that map to a permutation fixing k, i.e. they form the subgroup $\rho^{-1}(\mathrm{Fix}(k))$. Note that any path connecting two different choices of basepoint in $X \setminus B$ maps to a loop in $T \setminus \{p\}$ which induces a conjugation of the corresponding subgroups. In particular, all subgroups $\rho^{-1}(\mathrm{Fix}(k))$ for $k = 1, \dots, n$ are isomorphic.

By definition, the branched cover $X \to T$ is normal if and only if the unbranched cover $X \setminus B \to T \setminus \{p\}$ is normal. The latter is true if $\pi_1(X \setminus B)$ is a normal subgroup of $\pi_1(T \setminus \{p\})$, which holds if and only if $\rho^{-1}(\mathrm{Fix}(k)) = \ker(\rho)$ for some, hence any $k \in \{1, \dots, n\}$. It is again standard covering space theory that the deck group G of a normal cover fits into a short exact sequence of the form

$$1 \to \pi_1(X \setminus B) \to \pi_1(T \setminus \{p\}) \xrightarrow{\rho} G \to 1,$$

where $G \cong \mathrm{im}(\rho) \cong \mathbb{F}_2/\ker(\rho)$. □

From now on, we identify $\pi_1(X \setminus B)$ with $\rho^{-1}(\mathrm{Fix}(1))$. In the following remark, we denote this group by Π.

Remark 7.2. If $X \to T$ is not a normal cover, then the deck group is given by $N(\Pi)/\Pi$ where $N(\Pi)$ is the normaliser, i.e. the largest subgroup of $\mathbb{F}_2$ in which Π is normal [Hat02, Proposition 1.39].

We say that a cover is *cyclic* if the image of the monodromy ρ is a cyclic group. The following lemma says that in our context, cyclic covers are automatically normal.

Lemma 7.3. *If the image of $\rho\colon \mathbb{F}_2 \to S_n$ is a cyclic group, then the cover $X \to T$ is normal.*

Proof. Since X is assumed to be connected, we must necessarily have that $\mathrm{im}(\rho)$ is generated by an n-cycle. In this case, $\mathrm{Fix}(k) \cap \mathrm{im}(\rho) = \{\mathrm{id}\}$ for any $k \in \{1, \dots, n\}$. It follows that for all k, we have $\rho^{-1}(\mathrm{Fix}(k)) = \ker(\rho)$, and therefore, by Lemma 7.1, that the cover is normal. □

In fact, the above proof shows that in the setting of Lemma 7.3, the deck group has to be $\mathbb{Z}/n\mathbb{Z}$.

We now look at an example of a STS that will be important for the rest of this chapter.

Example 7.4. The *four-square torus* X is the STS corresponding to the map

$$\rho\colon \mathbb{F}_2 \to S_4,\ a \mapsto (1\ 2)\ (3\ 4),\ b \mapsto (1\ 3)\ (2\ 4).$$

The set of branch points B on X of the cover $X \to T$ consists of four elements, which can be seen in Figure 7.2. Since X is topologically a torus, the fundamental group Π of $X \setminus B$ is isomorphic to $\mathbb{F}_5$ which we identify with the subgroup $\rho^{-1}(\mathrm{Fix}(1)) \leqslant \mathbb{F}_2$. The latter is generated by the elements $a^2, b^2, [b^{-1}, a], [b, a], [a^{-1}, b^{-1}] \in \mathbb{F}_2$. The cover is normal, and the deck group is given by $\mathrm{im}(\rho) \cong \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$.

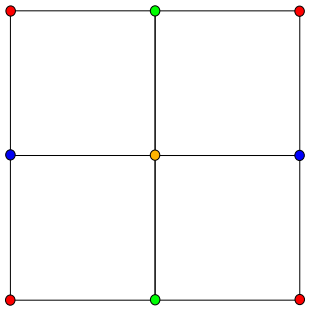


Figure 7.2: The four-square torus.

Definition 7.5. *A square-tiled surface X is* geminal *if it is made out of $4d$ squares for some $d \in \mathbb{N}$ and the corresponding $\rho : \mathbb{F}_2 \to S_{4d}$ satisfies the following condition:*

For every $x \in \mathbb{F}_2$ that is part of a size-two generating set, i.e. for which there exists $y \in \mathbb{F}_2$ with $\mathbb{F}_2 = \langle x, y \rangle$, it holds that $\rho(x) \in C(2d, 2d) \subset S_{4d}$.

Note that the four-square torus from example 7.4 is geminal. In fact, every geminal X covers the four-square torus as we will see below.

The condition we imposed on a geminal STS can be interpreted geometrically as saying that cylinders come in pairs. This is because every cylinder on T is described by an element $x \in \mathbb{F}_2$ that is part of a size-two generating set, and the condition on x implies that the corresponding cylinder lifts to two parallel cylinders on X of equal height and circumference.

We call an element $x \in \mathbb{F}_2$ as in Definition 7.5 a *primitive* element. Note that primitive elements in $\mathbb{F}_2$ must map to primitive elements in $\mathbb{Z}^2$ under the abelianisation map.

Example 7.6. The *Eierlegende Wollmilchsau (EW)* is an example of a geminal STS with 8 squares as shown in Figure 7.3. In fact, it can be checked that it is the only geminal STS with 8 squares. Furthermore, the EW is a double cover of the four-square torus, as we will see below. The name "Eierlegende Wollmilchsau" is German and describes a fictitious farm animal that has many favourable features, such as being able to lay eggs as well as produce milk and wool. Mathematicians chose this name for this STS because it serves as a counterexample in many instances in the study of translation surfaces (see, for example, [HS08]).

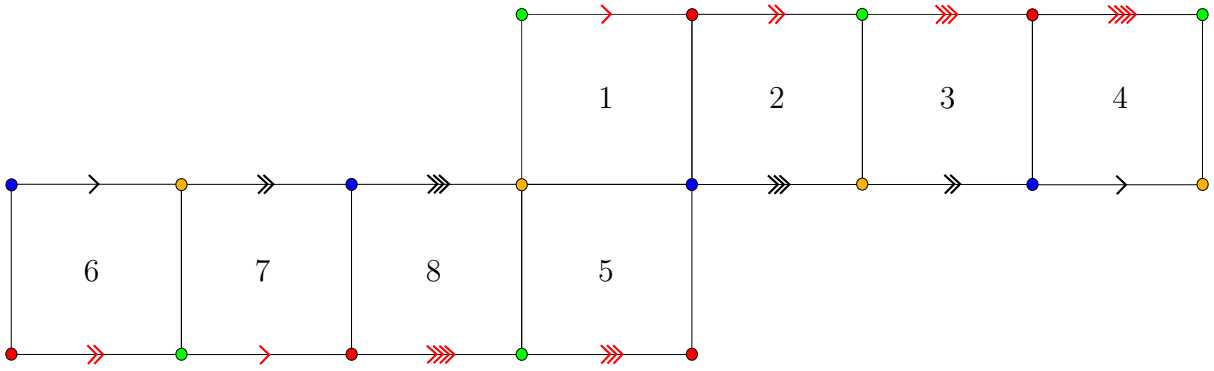


Figure 7.3: The Eierlegende Wollmilchsau corresponds to the monodromy map with $\rho(a) = (1\ 2\ 3\ 4)(5\ 6\ 7\ 8)$, $\rho(b) = (1\ 5\ 3\ 7)(2\ 8\ 4\ 6)$. The EW lies in the stratum $\mathcal{H}(1,1,1,1)$ and has genus 3.

§ 7.1.1 | Description via graphs

Recall from Section 5.1 that we can assign a graph G to a STS. The way this was done was to define the graph G by requiring that each vertex should correspond to a square of the STS, and two vertices should be joined by an edge if the corresponding squares are adjacent. If we further label the edges by whether they correspond to a horizontal or a vertical adjacency of the corresponding squares and direct them to point to the right, resp. bottom, of a square, then the graph contains all the information of the STS.

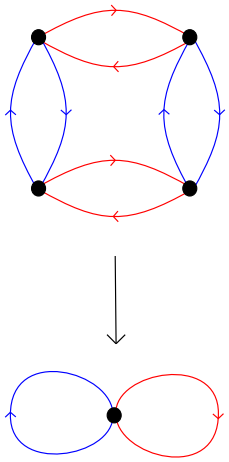


Figure 7.4: The covering between the graphs of the four-square torus and the standard torus.

Let now G_T be the graph corresponding to the standard torus T, which is a two-petaled rose. Let X be any other STS. Then, its graph G_X is an unbranched cover of G_T defined by the corresponding map $\rho\colon \mathbb{F}_2 \to S_n$. Indeed, note that G_T and G_X are homotopy equivalent to $T \setminus \{p\}$ and $X \setminus B$ respectively, where $\{p\}$ and B are the branching sets as defined above. Hence, $\pi_1(G_T) \cong \mathbb{F}_2$ and $\pi_1(G_X) \cong \rho^{-1}(\mathrm{Fix}(1))$. It follows that the lemmata in Section 7.1 hold analogously for the spaces G_T and G_X. The reason we choose to think about square-tiled surfaces as graphs from now on is that we can restrict ourselves to thinking about unbranched covers, and it becomes easier to visualise examples.

§ 7.2 | Normality and cyclicity of geminal covers

From now on, we omit writing the word *branched* and refer to the relevant branched covers simply as *covers*.

Let T be the standard torus and T_4 be the four-square torus from Example 7.4. Recall that any STS X, in particular also T_4, comes with a cover to T. In this section, we want to prove our main result about geminal square-tiled surfaces:

Theorem 7.7. *Let X be a geminal square-tiled surface. The cover $X \to T$ factors through a cover*

$$p\colon X \to T_4,$$

for which the following are equivalent:

1) p is normal,

2) p is cyclic.

Theorem 7.7 answers the second part of the Apisa–Wright conjecture (formulated as Problem 8.16 in [AW22]) affirmatively. The fact that for any geminal STS X there is a cover $p\colon X \to T_4$, can be interpreted as saying that T_4 is the smallest geminal STS.

Theorem 7.8 below is a purely group-theoretic description of our main result in terms of the corresponding homomorphism $\rho\colon \mathbb{F}_2 \to S_{4d}$. The proof of the theorem is geometric in flavour and uses the graph description of square-tiled surfaces from the end of the last section.

Let $\mathbb{F}_2 = \langle a, b\rangle$ and

$$\rho_4\colon \mathbb{F}_2 \to \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z},$$

be given by $a \mapsto (1,0)$ and $b \mapsto (0,1)$. We think of ρ_4 as the homomorphism corresponding to T_4, where we interpret $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ as the subgroup of S_4 generated by $(1\,2)\,(3\,4)$ and $(1\,3)\,(2\,4)$. Let $K \coloneqq \ker(\rho_4)$ and note that K is isomorphic to the fundamental group of T_4 minus the branching locus, or equivalently to $\pi_1(G_4)$, where G_4 is the graph corresponding to T_4.

Let now $d \in \mathbb{N}$ and X be a geminal STS with corresponding monodromy $\rho\colon \mathbb{F}_2 \to S_{4d}$. We will see that we can split the set $\{1, \dots, 4d\}$ into four sets of order d, such that the action of K under ρ preserves these four sets. Let $\bar{\rho}\colon K \to S_d$ be the restriction of the action of K to any of the four sets. This map corresponds to the monodromy map of the cover $X \to T_4$.

Theorem 7.8. *It holds that*

$$\rho^{-1}(\mathrm{Fix}(k)) \leqslant K,$$

for all $k \in \{1, \dots, 4d\}$. Furthermore, the following are equivalent:

1) $\rho^{-1}(\mathrm{Fix}(1)) \trianglelefteq K$

2) $\bar{\rho}\colon K \to S_d$ has cyclic image.

Remark 7.9. The first part of the theorem, namely that $\rho^{-1}(\mathrm{Fix}(k))$ is a subgroup of K, exactly corresponds to the fact that the cover $X \to T$ factors through a cover $X \to T_4$.

The two equivalent conditions say that this cover is normal if and only if it is cyclic. The proof below furthermore implies that if we are in the case of a cyclic, normal cover, then $\rho^{-1}(\mathrm{Fix}(k))$ is a normal subgroup of K for all $k \in \{1, \dots, 4d\}$. Moreover, the set

$$\{\rho^{-1}(\mathrm{Fix}(k)) \,|\, k \in \{1, \dots, 4d\}\}$$

consists of four elements, all of which are conjugate subgroups of $\mathbb{F}_2$, and

$$K/\rho^{-1}(\mathrm{Fix}(k)) \cong \mathbb{Z}/d\mathbb{Z}$$

for any k.

The rest of this section is devoted to the proof of Theorem 7.8.

The property of being geminal implies that $\rho(a)$ and $\rho(b)$ are in the conjugacy class $C(2d, 2d)$. We refer to such elements as $(2d, 2d)$-cycles. Up to renumbering, we may assume that

$$\rho(a) = (1\ 2\ 3\ \dots\ 2d)\ (2d+1\ 2d+2\ \dots\ 4d).$$

With this in mind, we now want to describe what $\rho(b)$ looks like.

Lemma 7.10. *The two cycles of the $(2d, 2d)$-cycle $\rho(b)$ alternate between the two cycles of $\rho(a)$. Furthermore, without loss of generality, we can assume that one cycle of $\rho(b)$ consists entirely of odd numbers and the other entirely of even ones.*

Proof. Denote the two cycles of $\rho(a)$ by c_1^a and c_2^a, where

$$c_1^a = (1\ 2\ 3\ \dots\ 2d) \quad \text{and} \quad c_2^a = (2d+1\ 2d+2\ \dots\ 4d).$$

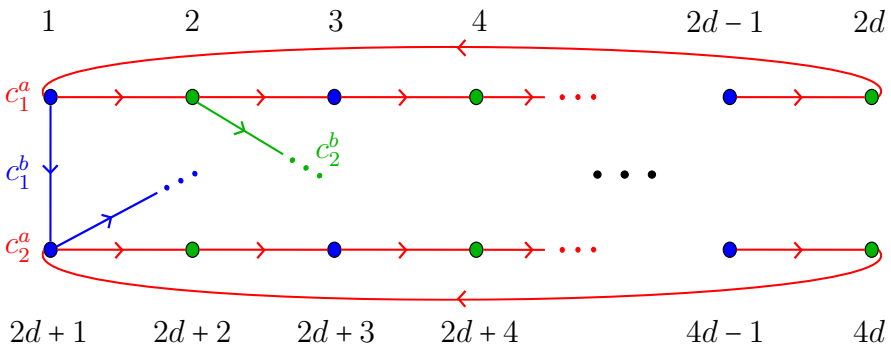


Figure 7.5: A schematic drawing of the cycles of $\rho(a)$ and $\rho(b)$.

Similarly, let c_1^b, c_2^b be the two $2d$-cycles of $\rho(b)$. Assume that the cycles of $\rho(b)$ don't alternate between the cycles of $\rho(a)$. Then, without loss of generality, there exist $i, j \in \{1, \dots, 2d\}$ such that $c_1^b(i) = j$. Let k be the power such that $(c_1^a)^k(i) = j$. Now, consider the element $b^{-1}a^k \in \mathbb{F}_2$. This element is primitive and should hence map to a $(2d, 2d)$-cycle under ρ. However,

$$\rho(b^{-1}a^k)(i) = \rho(b^{-1})((c_1^a)^k(i)) = (c_1^b)^{-1}(j) = i,$$

which means that $\rho(b^{-1}a^k)$ fixes i and can therefore not be a $(2d, 2d)$-cycle. This shows that $\rho(b)$ has to alternate between c_1^a and c_2^a.

Up to relabelling, we may assume that $\rho(b)(1) = 2d + 1$. We claim that with this convention, it follows that c_1^b consists entirely of odd numbers and c_2^b entirely of even ones. Assume this was not the case. Then, since the cycles of $\rho(b)$ alternate between the ones of $\rho(a)$, a cycle of $\rho(b)$ needs to necessarily contain two numbers that are consecutive in a cycle of $\rho(a)$, i.e. there exists a power k and some $i \in \{1, \dots, 4d\}$ such that

$$\rho(b)^k(i) = \rho(a)(i).$$

But then the element $a^{-1}b^k$ is a primitive element whose image under ρ fixes i. This leads to the desired contradiction, which implies the claim and finishes the proof. □

With the above lemma in hand, we are ready to prove the first part of the theorem. Recall that X is our geminal STS, T_4 the four-square torus and T the standard torus.

Lemma 7.11. *The standard cover $X \to T$ factors through the cover $T_4 \to T$, or equivalently there is an inclusion of subgroups $\rho^{-1}(\mathrm{Fix}(1)) \leqslant K \leqslant \mathbb{F}_2$.*

Proof. Let G_X, G_{T_4} and G_T be the graphs corresponding to the square-tiled surfaces X, T_4 and T as described in Section 7.1.1 and let $p_X : G_X \to G_T$, $p_{T_4} : G_{T_4} \to G_T$ be the corresponding covering maps. It suffices to construct a covering map $p : G_X \to G_{T_4}$ such that $p_{T_4} \circ p = p_X$, since this implies that

$$\rho^{-1}(\mathrm{Fix}(1)) \cong \pi_1(G_X) \leqslant K \cong \pi_1(G_{T_4}) \leqslant \mathbb{F}_2 \cong \pi_1(G_T),$$

which in turn implies that there exists a cover $X \to T_4$ as desired.

Lemma 7.10 implies that the graph G_X must be of the form as in Figure 7.5. In particular, G_X consists of $4d$ vertices which we number $1, \dots, 4d$, such that:

- There are edges forming the ordered cycles $1, 2, 3 \ldots, 2d$ and $2d+1, 2d+2 \ldots, 4d$ - corresponding to $\rho(a)$;
- There are edges forming cycles containing the vertices $1, 3, 5, \ldots, 4d{-}1$ and $2, 4, 6, \ldots, 4d$ which alternate between the two cycles of the first bullet point - corresponding to $\rho(b)$.

Recall the graph G_{T_4} from Figure 7.4. G_{T_4} consists of 4 vertices, which we now number by $1, 2, 3$ and 4 such that:

- There are two edges joining 1 and 2 as well as two edges joining 3 and 4 - corresponding to $\rho_4(a)$;
- There are two edges joining 1 and 3 as well as two edges joining 2 and 4 - corresponding to $\rho_4(b)$.

Define $p : G_X \to G_{T_4}$ by sending

- each odd vertex in $\{1, \ldots, 2d\}$ to 1,
- each even vertex in $\{1, \ldots, 2d\}$ to 2,
- each odd vertex in $\{2d+1, \ldots, 4d\}$ to 3,
- each even vertex in $\{2d+1, \ldots, 4d\}$ to 4.

Since, as described in Section 7.1.1, the graphs are directed, the map p is determined on the edges by the above assignment of vertices.

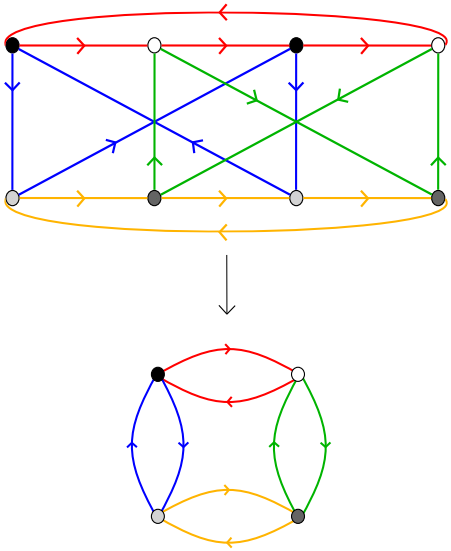


Figure 7.6: The cover p in the case of the Eierlegende Wollmilchsau.

One can now see that each point on G_{T_4} has exactly d preimages under p. In particular, the preimage of an edge is a disjoint union of d edges, and the preimage of a small neighbourhood of a vertex - a four-prong - is a disjoint union of d homeomorphic copies of that neighbourhood. Hence, p is a covering, and it is straightforward to check that $p_{T_4} \circ p = p_X$. In fact, on the vertices the equality holds since they are all mapped to a single vertex by p_X and p_{T_4}, and on the edges it holds because p sends cycles corresponding to $\rho(a), \rho(b)$ to cycles corresponding to $\rho_4(a), \rho_4(b)$, respectively. □

Finally, we are ready to prove the main theorem. In the proof, we use the same notation as in the proofs of the last two lemmata.

Proof of Theorem 7.8. It only remains to show the equivalence of 1) and 2) as stated in the Theorem. Let X be a geminal STS and G_X be the corresponding graph. It will be helpful to keep Figure 7.5 in mind during the proof, which is a schematic drawing of the graph G_X. Let $p : G_X \to G_{T_4}$ be the covering constructed in the proof of Lemma 7.11. The covering corresponds to an inclusion of fundamental groups

$$\rho^{-1}(\mathrm{Fix}(1)) \leqslant K.$$

Here, we have chosen the vertices labelled by 1 in G_X and G_{T_4} to act as basepoints. Note that a different choice would lead to the inclusion

$$\rho^{-1}(\mathrm{Fix}(k)) \leqslant K,$$

where $\rho^{-1}(\mathrm{Fix}(k))$ is conjugate to $\rho^{-1}(\mathrm{Fix}(1))$ by an element in $\mathbb{F}_2$ and K is unchanged since it is a normal subgroup of $\mathbb{F}_2$.

Consider the monodromy map $\rho : \mathbb{F}_2 \to S_{4d}$ corresponding to X. Restricting ρ to K corresponds to looking at the permutations induced by lifts of loops of G_{T_4} under the cover p. Such a permutation will preserve the sets of preimages of the four vertices of G_{T_4} which are

$$A_{\mathrm{odd}} := p^{-1}(1) = \{1, 3, 5, \dots, 2d-1\},\ A_{\mathrm{even}} := p^{-1}(2) = \{2, 4, \dots, 2d\},$$

and

$$B_{\mathrm{odd}} := p^{-1}(3) = \{2d+1, 2d+3, \dots, 4d-1\},\ B_{\mathrm{even}} := p^{-1}(4) = \{2d+2, 2d+4 \dots, 4d\}.$$

Hence, for the restriction, we obtain

$$\rho|_K : K \longrightarrow \underbrace{\mathrm{Sym}(A_{\mathrm{odd}}) \times \mathrm{Sym}(A_{\mathrm{even}}) \times \mathrm{Sym}(B_{\mathrm{odd}}) \times \mathrm{Sym}(B_{\mathrm{even}})}_{\cong (S_d)^4} \leqslant S_{4d}.$$

Projecting to the first factor $\mathrm{Sym}(A_{\mathrm{odd}}) \cong S_d$, we obtain a map

$$\bar{\rho}\colon K \to S_d$$

which can be thought of as the monodromy map of the cover $p\colon G_X \to G_{T_4}$. Again, a different choice of factor to project to corresponds to a different choice of basepoint and would lead to a monodromy map with isomorphic image.

1) $\Rightarrow$ 2): If $\rho^{-1}(\mathrm{Fix}(1))$ is a normal subgroup of K, then p is a normal cover. In particular, the deck group acts transitively on the set A_{odd} and has to preserve the cycle c_a^1 which connects these points (cf. Figure 7.5). This implies that the deck group has to be a cyclic group of order d. On the other hand, the deck group is isomorphic to the image of the monodromy map $\bar{\rho}\colon K \to S_d$. Hence, we conclude that

$$\mathrm{im}(\bar{\rho}) \cong \mathbb{Z}/d\mathbb{Z}.$$

2) $\Rightarrow$ 1): This direction holds essentially by the same argument as in Lemma 7.3. Assume that the image of $\bar{\rho}$ is cyclic. Since G_X is connected, the image has to necessarily be $\mathbb{Z}/d\mathbb{Z}$. Furthermore, as in Lemma 7.3, it follows that

$$\rho^{-1}(\mathrm{Fix}(1)) = \ker(\bar{\rho}),$$

which means that $\rho^{-1}(\mathrm{Fix}(1))$ is a normal subgroup of K.

□

Remark 7.12. To define the map $\bar{\rho}$ in the above proof, the choice of projecting to $\mathrm{Sym}(A_{\mathrm{odd}})$ instead of $\mathrm{Sym}(A_{\mathrm{even}}), \mathrm{Sym}(B_{\mathrm{odd}})$ or $\mathrm{Sym}(B_{\mathrm{even}})$ was arbitrary. Indeed, any other choice would also describe the monodromy of the cover, but with a different choice of basepoint. In particular, for a normal cover, any choice would lead to a cyclic image. However, we emphasise that this does in general not imply that the map $\rho|_K\colon K \to S_{4d}$ has cyclic image, as the following example illustrates.

Example 7.13 (Ornithorynque). Let X be the STS with 12 squares defined by $\rho\colon \mathbb{F}_2 \to S_{12}$ with

$$\rho(a) = (1\ 2\ 3\ 4\ 5\ 6)\,(7\ 8\ 9\ 10\ 11\ 12),\ \ \rho(b) = (1\ 7\ 3\ 11\ 5\ 9)\,(2\ 10\ 6\ 12\ 4\ 8).$$

This STS is often called the *Ornithorynque*, and is geminal [AW22, Example 2.7]. Furthermore, X is a normal cover of T_4 with deck group $\mathbb{Z}/3\mathbb{Z}$, generated by rotating the vertices in the top circle in Figure 7.7 by +2 (mod 6), the vertices in the bottom circle by −2 (mod 6) and mapping the edges accordingly.

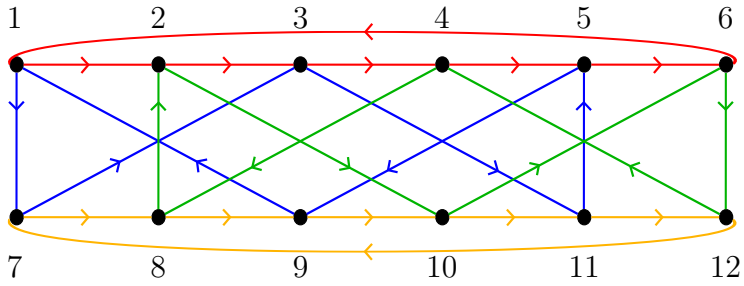


Figure 7.7: The graph G_X of the Ornithorynque.

The four sets described in the proof of Theorem 7.8 are $\{1,3,5\},\{2,4,6\},\{7,9,11\}$ and $\{8,10,12\}$. The map $\bar{\rho}\colon K \to S_3$ was defined by looking at the action of K on $\{1,3,5\}$. Equivalently, we could define a $\bar{\rho}'$ by looking at the action on $\{2,4,6\}$. Both $\bar{\rho}$ and $\bar{\rho}'$ have cyclic image, generated by $(1\ 3\ 5)$ and $(2\ 4\ 6)$, respectively. However, the image of $\rho|_K$ is not cyclic, as

$$\rho(a^2) = (1\ 3\ 5)\,(2\ 4\ 6)\,(7\ 9\ 11)\,(8\ 10\ 12),$$

but

$$\rho(ab^2a) = (1)(3)(5)\,(2\ 6\ 4)\,(7\ 11\ 9)\,(8)(10)(12).$$

Essentially, the reason an element $x \in K$ does in general not act on all four of the above 3-element sets by elements of the same order is that x describes different loops when based at the four vertices of G_{T_4}, which are non-homotopic.

§ 7.2.1 | Summary

We end this section with a quick summary of the results and implications of the above proofs.

Let X be a geminal STS with $4d$ squares for some $d \in \mathbb{N}$ and let T_4 be the four-square torus and T be the standard torus. There is a tower of covers

$$\begin{array}{c} X \\ \downarrow p \\ T_4 \\ \downarrow \\ T, \end{array}$$

and the cover p is normal if and only if it is cyclic. The cover $T_4 \to T$ is a degree 4 normal cover with deck group $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$. However, the cover $X \to T$ is in general not normal, even if p is. On the other hand, if $X \to T$ is normal, then so is p. The degree of p is d, while the degree of $X \to T$ is $4d$.

We described the above covers in terms of the fundamental groups of the spaces minus the respective branching sets and the monodromy maps to the corresponding symmetric groups. In particular, the covers correspond to the diagram

$$\begin{array}{ccc} \rho^{-1}(\mathrm{Fix}(1)) = \bar{\rho}^{-1}(\mathrm{Fix}(1)) & & \\ \downarrow & & \\ \mathbb{F}_5 & \xrightarrow{\bar{\rho}} & S_d \\ & \searrow^{\rho} & \uparrow \\ & & S_d \times S_d \times S_d \times S_d \\ & & \downarrow \\ \mathbb{F}_2 & \xrightarrow{\rho} & S_{4d}. \end{array}$$

Here, all vertical arrows pointing downward are inclusions. $\mathbb{F}_5$ sits inside $\mathbb{F}_2$ as the subgroup generated by $a^2, b^2, [b^{-1}, a], [b, a], [a^{-1}, b^{-1}]$, and the inclusion of $S_d \times S_d \times S_d \times S_d$ into S_{4d} is induced by the inclusions of the sets

$$\{1, 3, 5, \ldots, 2d-1\},\ \{2, 4, \ldots, 2d\},\ \{2d+1, 2d+3, \ldots, 4d-1\},\ \{2d+2, 2d+4 \ldots, 4d\}$$

into

$$\{1, \ldots, 4d\}.$$

The map $S_d \times S_d \times S_d \times S_d \to S_d$ is the projection to the first factor.

The map ρ is the defining map of the STS X where we have chosen the generator a to map to $(1\ 2\ 3\ \ldots\ 2d)\ (2d+1\ 2d+2\ \ldots\ 4d) \in C(2d, 2d)$. Every element of $\mathbb{F}_5$ maps under ρ to an element in $C(d, d, d, d)$ which preserves the sets

$$\{1, 3, 5, \ldots, 2d-1\},\ \{2, 4, \ldots, 2d\},\ \{2d+1, 2d+3, \ldots, 4d-1\},\ \{2d+2, 2d+4 \ldots, 4d\}.$$

Hence, we can view $\rho|_{\mathbb{F}_5}$ as a map to $S_d \times S_d \times S_d \times S_d$. The map $\bar{\rho}$ is the concatenation of $\rho|_{\mathbb{F}_5}$ and the projection to the first factor. The groups $\rho^{-1}(\mathrm{Fix}(1))$ and $\bar{\rho}^{-1}(\mathrm{Fix}(1))$ are equal and correspond to the fundamental group of $X \setminus B$, where B is the set of branch points.

The cover p being normal is equivalent to $\rho^{-1}(\mathrm{Fix}(1))$ being a normal subgroup of $\mathbb{F}_5$ which in turn is equivalent to

$$\rho^{-1}(\mathrm{Fix}(1)) = \ker(\bar{\rho}),$$

or to

$$\rho^{-1}(\mathrm{Fix}(1)) = \rho^{-1}(\mathrm{Fix}(3)) = \rho^{-1}(\mathrm{Fix}(5)) = \ \ldots\ = \rho^{-1}(\mathrm{Fix}(2d-1)).$$

Theorem 7.8 implies that if this is the case, then

$$\mathbb{F}_5/\rho^{-1}(\mathrm{Fix}(1)) \cong \mathbb{Z}/d\mathbb{Z}.$$

§ 7.3 | Pillowcase-tiled surfaces

The discussion of geminal square-tiled surfaces is closely related to so-called 1-*cylinder pillowcase-tiled surfaces*. In what follows, we define pillowcase-tiled surfaces and discuss their connection to the above.

The *pillowcase* P is a half-translation surface given by gluing the sides of two squares as in Figure 7.8. Note that since we glue the top and bottom sides by a half-translation, the

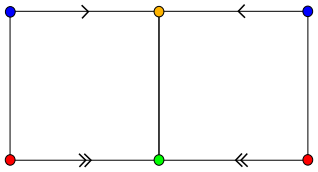


Figure 7.8: The pillowcase.

pillowcase is not a square-tiled surface under our definition, but it is a half-translation surface. The four points corresponding to the corners of the squares in Figure 7.8 are singularities with cone angle π. Hence, P is an element of the stratum $\mathcal{Q}(-1,-1,-1,-1)$ of quadratic differentials. The genus of P is zero, so the underlying topological space is a sphere.

A *pillowcase-tiled surface* Y, also abbreviated to PTS, is a finite-degree cover of the pillowcase P branched over (at most) the four singularities. If the degree of the cover is $d \in \mathbb{N}$, then Y can be represented by $2d$ squares, glued together in an appropriate way.

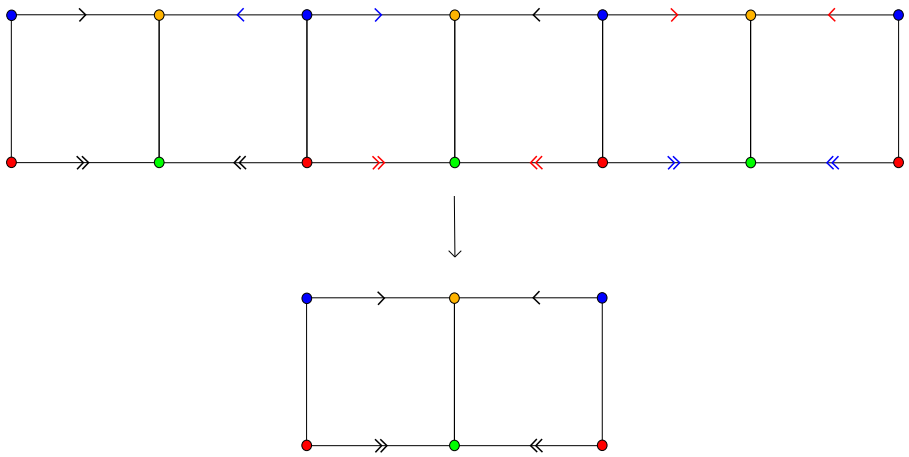


Figure 7.9: An example of a pillowcase-tiled surface.

Similarly to square-tiled surfaces, we can define a pillowcase-tiled surface Y via permutations. If we let B be the set of the four singularities on P, then Y induces a monodromy homomorphism $\rho\colon \pi_1(P \setminus B) \cong \mathbb{F}_3 \to S_d$. Conversely, every such homomorphism defines a PTS, and two homomorphisms define the same PTS if they differ by conjugation on S_d. Since a homomorphism from $\mathbb{F}_3$ is determined by the image of the three generators, we conclude that pillowcase-tiled surfaces with $2d$ squares are in one-to-one correspondence with

$$(S_d)^3/S_d,$$

where S_d acts by simultaneous conjugation on each of the three factors.

The four-square torus T_4 is a double cover of the pillowcase P, where the deck group $\mathbb{Z}/2\mathbb{Z}$ is generated by a 180-degree rotation. In fact, for any half-translation surface Y, there

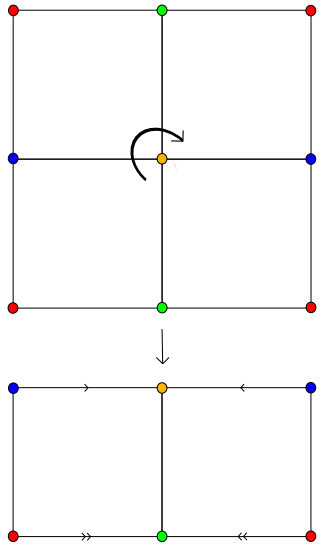


Figure 7.10: The four-square torus covers the pillowcase.

exists a canonical double cover X which is a translation surface, called the *translation cover*, or *holonomy double cover*, of Y (see [AM24, Section 2.6.4]). We briefly describe the construction of a translation cover.

Start with a half-translation surface Y given as a collection of polygons. Denote this collection by A. Consider another collection of the same polygons rotated by 180 degrees. Denote this collection by $\bar{A}$. We will use the collection $A \cup \bar{A}$ to build a translation surface X. If, to obtain Y, two sides of A were glued by a translation, we glue the same two sides of A as well as the corresponding ones of $\bar{A}$ by a translation. If, to obtain Y, two sides α, β of A were glued by a half-translation, then let $\bar{\alpha}, \bar{\beta}$ be the corresponding two sides of $\bar{A}$ and glue α to $\bar{\beta}$ and β to $\bar{\alpha}$ by a translation. We have now obtained a translation surface X. The natural map coming from mapping $A \cup \bar{A}$ to A by the identity on A and a 180

degree rotation on $\bar{A}$ defines a double cover from X to Y. The branch points correspond to the singularities of Y of cone angle an odd multiple of π.

The four-square torus is the translation cover of the pillowcase. Furthermore, given any PTS Y, its translation cover X is a STS which necessarily covers the four-square torus. We obtain a commutative diagram of covers

$$\begin{array}{ccc} X & \xrightarrow{2} & Y \\ \downarrow{\scriptstyle d} & & \downarrow{\scriptstyle d} \\ T_4 & \xrightarrow{2} & P, \end{array} \tag{7.1}$$

where the numbers on the arrows stand for the degree of the corresponding covers. As double covers, the horizontal arrows are normal covers.

Lemma 7.14. *If the cover $Y \to P$ in the above diagram is normal, then so is the cover $X \to T_4$.*

Proof. Diagram 7.1 gives rise to the following inclusion of fundamental groups

$$\begin{array}{ccc} \pi_1(X^*) & \longrightarrow & \pi_1(Y^*) \\ \downarrow & & \downarrow \\ \pi_1(T_4^*) & \longrightarrow & \pi_1(P^*). \end{array}$$

where the * indicates the removal of the corresponding branching set. The horizontal inclusions are of index 2, so in particular normal. Note that if we view each group as a subgroup of $\pi_1(P^*)$, then we can write

$$\pi_1(X^*) = \pi_1(Y^*) \cap \pi_1(T_4^*).$$

Assuming that the inclusion $\pi_1(Y^*) \hookrightarrow \pi_1(P^*)$ is normal, we want to show that $\pi_1(X^*) \hookrightarrow \pi_1(T_4^*)$ is normal.

So, let $x \in \pi_1(X^*)$ and $t \in \pi_1(T_4^*)$. Then, $txt^{-1} \in \pi_1(T_4^*)$ and furthermore, $txt^{-1} \in \pi_1(Y^*)$ since $x \in \pi_1(Y^*)$ and $\pi_1(Y^*) \trianglelefteq \pi_1(P^*)$. Overall, we have

$$txt^{-1} \in \pi_1(Y^*) \cap \pi_1(T_4^*) = \pi_1(X^*),$$

which shows that $\pi_1(X^*)$ is normal in $\pi_1(T_4^*)$. □

As described in the above proof, the diagram 7.1 gives rise to a diagram of fundamental groups which are isomorphic to the following groups:

$$\begin{array}{ccc} \mathbb{F}_{4d+1} & \longrightarrow & \mathbb{F}_{2d+1} \\ \downarrow & & \downarrow \\ \mathbb{F}_5 & \longrightarrow & \mathbb{F}_3. \end{array}$$

The horizontal arrows have cokernel $\mathbb{Z}/2\mathbb{Z}$. We can describe the bottom map in terms of generators. Let $x_1, \ldots, x_5$ be the generators of $\mathbb{F}_5$ which correspond to the elements $a^2, b^2, [b^{-1}, a], [b, a], [a^{-1}, b^{-1}] \in \mathbb{F}_2 \cong \pi_1(T \setminus p)$ as in Example 7.4. One can think of the element a as following a directed red edge in Figure 7.11 and the element b as following a directed blue edge. Let x, y, z be the generators of $\mathbb{F}_3$ corresponding to the images of ba, ab^{-1} and $b^{-1}a^{-1}$ respectively. Again, Figure 7.11 can be used to visualise that this is indeed a generating set. The map $\mathbb{F}_5 \to \mathbb{F}_3$ in the above diagram is given by

$$x_1 \mapsto yx, \quad x_2 \mapsto y^{-1}z^{-1}, \quad x_3 \mapsto y^{-1}zyx, \quad x_4 \mapsto zx, \quad x_5 \mapsto z^{-1}x^{-1}.$$

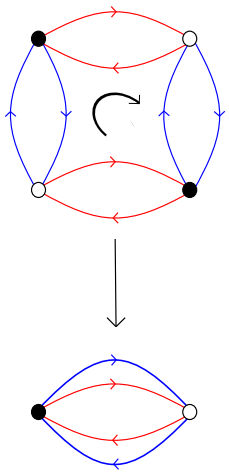


Figure 7.11: The graph picture of the translation cover of the pillowcase.

A 1-*cylinder* pillowcase-tiled surface is a PTS for which the closure of every cylinder is the whole surface. The following definition is an equivalent, purely group-theoretic description of 1-cylinder pillowcase-tiled surfaces, which is due to a theorem by Abdalla–Brown [AB25, Theorem 2]. Recall that a PTS is determined by a triple of permutations $(\pi_1, \pi_2, \pi_3) \in S_d$. We define an action of the braid group B_3 on $(S_d)^3$ by

$$\sigma_1 \cdot (\pi_1, \pi_2, \pi_3) = (\pi_2,\ \pi_2^{-1}\pi_1\pi_2,\ \pi_3), \quad \sigma_2 \cdot (\pi_1, \pi_2, \pi_3) = (\pi_1,\ \pi_3,\ \pi_3^{-1}\pi_2\pi_3),$$

where σ_1, σ_2 are the standard generators of B_3.

In the following definition, we chose the generating set $\{x, y, z\}$ of $\pi_1(P \setminus B) \cong \mathbb{F}_3$ to be as in Figure 7.12. In the above description, this set corresponds to the images of $ba, ab^{-1}, b^{-1}a^{-1}$.

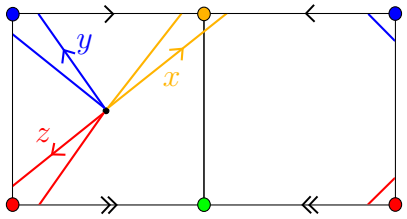


Figure 7.12: The choice of generators x, y, z for Definition 7.15.

Definition 7.15. *A* 1-cylinder pillowcase-tiled surface *is a PTS defined by a map*

$$\rho : \mathbb{F}_3 = \langle x, y, z \rangle \to S_d$$

with the property that if $(\pi_1, \pi_2, \pi_3) = g \cdot (\rho(x), \rho(y), \rho(z))$ *for any* $g \in B_3$*, then* $\pi_2\pi_1 \in C(d)$.

The connection between 1-cylinder PTS and geminal STS is given by the following lemma.

Lemma 7.16. *The translation cover of a* 1*-cylinder PTS* Y *is a geminal STS* X.

Proof. It is easiest to see the statement in terms of cylinders. The condition of a 1-cylinder PTS Y translates geometrically to the property that the closure of any cylinder is the whole surface Y [AB25]. We sketch how this can be seen. The Veech group of the pillowcase P is $\mathrm{SL}_2(\mathbb{Z})$, and it acts transitively on the cylinder directions of P. The condition that $\rho(y)\rho(x) \in C(d)$ means that the horizontal cylinder lifts to a unique cylinder in Y whose closure is all of Y. The action of the braid group can be thought of as moving the horizontal direction to any other cylinder direction and repeating the same argument. In fact, we could have used the natural action of $\mathrm{SL}_2(\mathbb{Z})$ instead. However, using the isomorphism $\mathrm{PSL}_2(\mathbb{Z}) \cong B_3/Z(B_3)$, Abdalla–Brown show that one can equivalently consider the action of B_3 described above [AB25, Lemma 6].

The condition in the definition of a geminal STS X translates geometrically to the property that for any cylinder direction, there are exactly two parallel cylinders whose closure is all of X. This can be seen using the fact that the Veech group of the torus T is again $\mathrm{SL}_2(\mathbb{Z})$ and acts transitively on cylinder directions T. The latter are described by elements of $\mathbb{F}_2$ that are part of a two-element generating set. The images of these elements under the monodromy map being $(2d, 2d)$-cycles means that each cylinder lifts to two cylinders whose closure is all of X.

Let now Y be a 1-cylinder PTS, X its translation cover and C_1 a cylinder on X. The image C of C_1 in Y is a cylinder, necessarily with closure equal to Y. From the construction of a translation cover, we see that a cylinder lifts to two parallel cylinders. So, the preimage

of C in X consists of C_1 and a parallel cylinder C_2. From Y being the closure of C, it now follows that the closure of $C_1 \cup C_2$ is all of X and therefore that X is geminal. □

Our main Theorem on geminal square-tiled surfaces now yields the analogous statement for 1-cylinder pillowcase-tiled surfaces:

Corollary 7.17. *The cover $Y \to P$, where Y is a 1-cylinder PTS and P is the pillowcase, is normal if and only if it is cyclic.*

Proof. The direction "cyclic implies normal" follows as in Lemma 7.3, since Y and P are connected.

For the other direction, assume the cover $Y \to P$ is normal. Then, by the Lemmata 7.16 and 7.14, there is an associated geminal STS X with normal cover $X \to T_4$. Theorem 7.8 now implies that $X \to T_4$ is cyclic. Since $X \to T_4$ is the pullback cover of $Y \to P$ under the double cover $T_4 \to P$, the deck group of $Y \to P$ is isomorphic to the one of $X \to T_4$, hence also cyclic. □

Corollary 7.17 says in particular that given any monodromy

$$\rho : \mathbb{F}_3 \to S_d$$

as in Definition 7.15 with cyclic image, we get a PTS Y which is a normal cover of the pillowcase with deck group $\mathbb{Z}/d\mathbb{Z}$. The translation cover X of Y is a geminal STS which is a normal cover of the four-square torus with deck group $\mathbb{Z}/d\mathbb{Z}$. Y consists of $2d$ squares, while X consists of $4d$. Abdalla–Brown show that 1-cylinder PTS only exist if d is odd [AB25, Theorem 1]. This implies that not every normal, cyclic geminal STS arises as the translation cover of a 1-cylinder PTS. For instance, the EW from Example 7.6 is such a geminal STS. Since the EW is a degree 2 cover of T_4, it would have to be the translation cover of a 1-cylinder PTS of degree 2, which does not exist.

The main result of Abdalla–Brown states that there are examples of 1-cylinder PTS that are not cyclic – hence not normal – covers of the pillowcase [AB25, Theorem 1]. In fact, such examples exist for any odd degree $d \geqslant 5$. The translation cover of such a PTS is an example of a non-cyclic geminal STS.

Example 7.18. We present an example of a non-cyclic geminal STS arising as the translation cover of a PTS coming from the construction of Abdalla–Brown. Let Y be the PTS of degree 5 defined by the monodromy

$$\rho(x) = (1\ 3\ 2\ 5\ 4),\ \rho(y) = (1\ 5\ 3\ 2\ 4),\ \rho(z) = (1\ 5\ 4\ 3\ 2),$$

where x, y, z is the generating set of Figure 7.12. The PTS Y is portrayed in Figure 7.13. In the figure, the Greek and Latin letters indicate the gluing of the sides, which are all glued by half-translations. One can check that the paths x, y and z lift to paths in Y that indeed induce the permutations on the squares labelled $1, 2, 3, 4, 5$ as specified above.

Since the product $\rho(z)\rho(y)\rho(x) = \mathrm{id}$, it follows that Y satisfies the definition of a 1-cylinder PTS. In fact, for any d, if a PTS is defined by three d-cycles whose product is the identity permutation, then it is a 1-cylinder PTS (see [AB25, proof of Theorem 1]). Since $\rho(y) \notin \langle \rho(x) \rangle$, the subgroup of S_5 generated by $\rho(x), \rho(y), \rho(z)$ is not cyclic and therefore Y is not a normal cover of P.

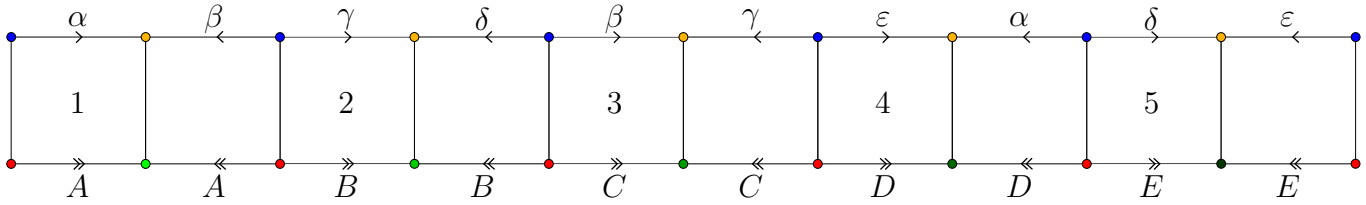


Figure 7.13: The non-cyclic 1-cylinder pillowcase-tiled surface Y.

Let X be the translation cover of Y, depicted in Figure 7.14. In this figure, all the sides are glued by translations. From Lemma 7.16, we know that X is geminal. From Figure 7.14, we see that we can describe X in terms of the monodromy $\rho_X : \mathbb{F}_2 = \langle a, b \rangle \to S_{20}$, given by

$$\rho_X(a) = (1\ 2\ 3\ 4\ 5\ 6\ 7\ 8\ 9\ 10)\,(11\ 12\ 13\ 14\ 15\ 16\ 17\ 18\ 19\ 20),$$

$$\rho_X(b) = (1\ 11\ 5\ 17\ 3\ 19\ 9\ 13\ 7\ 15)\,(2\ 12\ 8\ 16\ 10\ 14\ 4\ 20\ 6\ 18).$$

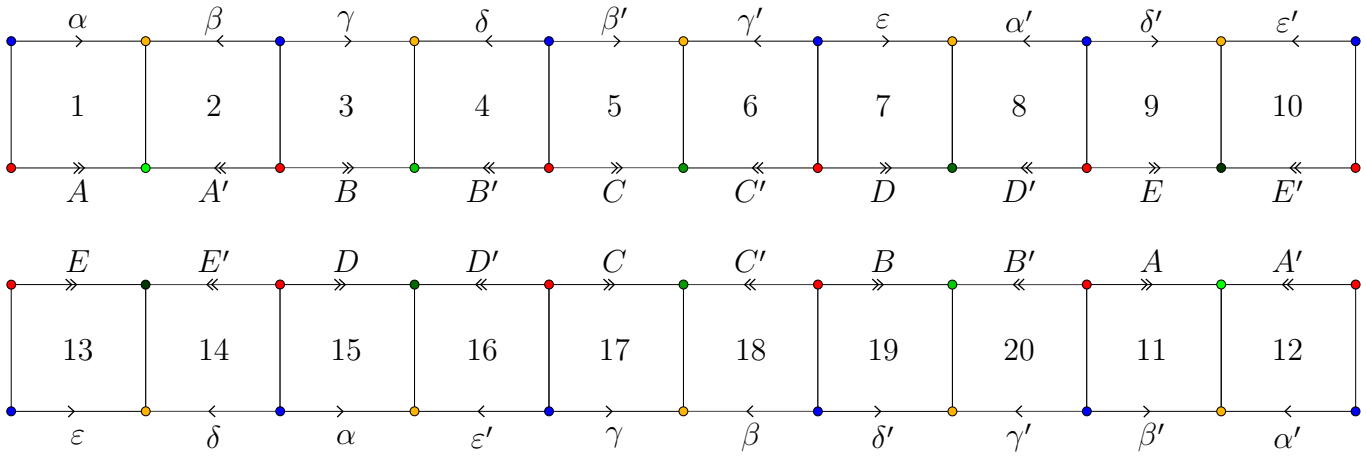


Figure 7.14: The non-cyclic geminal square-tiled surface X.

Hence, as a cover of T_4, the geminal STS X is given by the map $\bar{\rho}_X : \mathbb{F}_5 \to \mathrm{Sym}(\{1,3,5,7,9\})$,

$$\bar{\rho}_X(x_1) = (1\ 3\ 5\ 7\ 9),\ \bar{\rho}_X(x_2) = (1\ 5\ 3\ 9\ 7),$$

$$\bar{\rho}_X(x_3) = (1\ 3\ 7\ 5\ 9),\ \bar{\rho}_X(x_4) = \mathrm{id},\ \bar{\rho}_X(x_5) = (1\ 5\ 7\ 9\ 3),$$

where the x_i are the generators of $\mathbb{F}_5$ introduced in Example 7.4. Since the subgroup of $\mathrm{Sym}(\{1,3,5,7,9,\}) \cong S_5$ generated by the above permutations is not cyclic, X is an example of a geminal STS whose cover to T_4 is non-cyclic and non-normal.

The non-normality of the cover can also be seen geometrically. Consider Figure 7.15, which shows a part of the graph corresponding to X. If X were a normal cover of T_4, there should be a deck transformation mapping the vertex labeled 1 to the vertex labeled 3. Such a deck transformation has to rotate the top cycle by two vertices and rotate the bottom cycle by some amount. Observe, however, that the bottom cycle can not be rotated in a way such that the resulting deck transformation preserves all the arrows between the two cycles. This implies that the desired deck transformation does not exist and therefore that X is not normal.

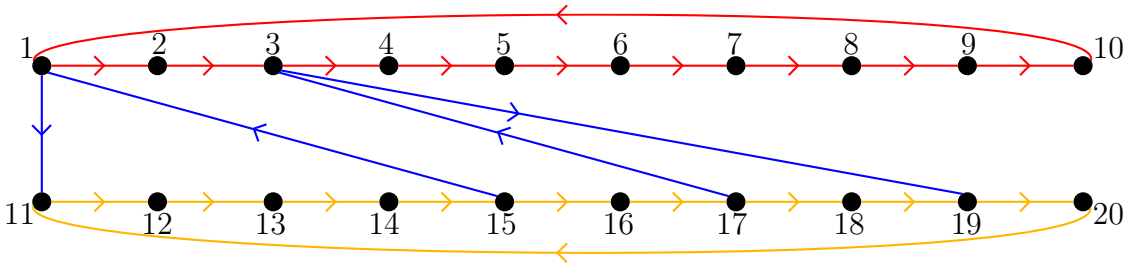


Figure 7.15: Part of the graph corresponding to X.

Remark 7.19. As in the above example, any non-cyclic 1-cylinder PTS Y constructed by Abdalla–Brown is defined by a monodromy given by three d-cycles whose product is the identity. This means that the preimages of three of the singularities of P consist of a single point of cone angle πd, while the preimage of the fourth singularity consists of d points of cone angle π. Hence, Y lies in the stratum $\mathcal{Q}(-1^d, (d-2)^3)$, where the exponents indicate the number of singularities of the given order. It follows that the genus of Y is $\frac{d-1}{2}$. In particular, one obtains such a Y in any genus $\geqslant 2$.

For the translation cover X, the cone points of angle π of Y lift to regular points while the ones of cone angle πd lift to points of cone angle $2\pi d$. Hence, $X \in \mathcal{H}((d-1)^3)$ and $g(X) = \frac{3d-1}{2}$. In particular, the construction only yields examples of non-cyclic geminal STS in certain genera. In our concrete example above, we obtained a genus 2 surface $Y \in \mathcal{Q}(-1^5, 3^3)$ and a genus 7 surface $X \in \mathcal{H}(4,4,4)$.

Example 7.20. Translation covers of non-cyclic 1-cylinder pillowcase-tiled surfaces only yield examples of non-cyclic geminal square-tiled surfaces whose cover to the four-square torus is of odd degree. In order to construct an example of an even-degree, non-cyclic geminal STS, we can start with an odd-degree, non-cyclic geminal STS X and consider the double cover $\tilde{X}$ that arises as the pullback of the Eierlegende Wollmilchsau. Concretely, if $p\colon X \to T_4$ and $p'\colon \mathrm{EW} \to T_4$ are the corresponding covers, then

$$\tilde{X} = \{(x, e) \in X \times \mathrm{EW} \,|\, p(x) = p'(e)\} \subset X \times \mathrm{EW},$$

and $\tilde{X}$ fits into the commutative diagram

$$\begin{array}{ccc} \tilde{X} & \xrightarrow{d} & \mathrm{EW} \\ \big\downarrow{\scriptstyle 2} & & \big\downarrow{\scriptstyle 2} \\ X & \xrightarrow{d} & T_4, \end{array}$$

where the numbers on the arrows indicate the degrees of the covers, and d is odd. By construction, $\tilde{X}$ is a non-cyclic, even-degree geminal STS.

§ 7.4 | Open questions

We end this chapter with some questions that arise from the discussion. We saw that a geminal square-tiled surface X covers the four-square torus T_4 and that this cover is normal if and only if it is cyclic. In order to construct non-normal, non-cyclic covers, we started with a 1-cylinder pillowcase-tiled surface Y that is a non-cyclic cover of the pillowcase P and looked at its translation cover X. Since such Y only exist if the cover to P has odd degree, we only obtain examples of non-cyclic geminal STS of odd degree. Taking a suitable double cover of such a STS then yielded examples of even-degree, non-cyclic geminal square-tiled surfaces. Overall, we now have examples for non-cyclic geminal square-tiled surfaces X whose corresponding cover $X \to T_4$ has either degree d or $2d$ with $d \geqslant 5$ odd. We can ask about the degrees we haven't covered yet.

Question 7.21. For $k \in \mathbb{N}$, is every geminal square-tiled surface of degree $4k$ normal and cyclic? What happens in the cases of degree $d = 2, 3, 6$?

It is not hard to check that for $d = 2$, the EW is the only geminal STS, which implies that all degree 2 geminal STS are normal and cyclic. We expect the same to hold for the other two low-degree examples, i.e., $d = 3$ and $d = 6$, although we have not carried out a detailed analysis. To answer the degree $4k$ question, one would need to either classify all such geminal square-tiled surfaces or find a non-cyclic example that doesn't arise from pillowcase-tiled surfaces.

As noted in the remark at the end of the last section, we are only able to construct non-cyclic geminal STS in certain strata and certain genera.

Question 7.22. Do there exist non-cyclic geminal square-tiled surfaces in every genus (above a certain number)? In which strata other than $\mathcal{H}(d-1^3)$ do non-cyclic geminal square-tiled surfaces exist?

Of course, it is also interesting to ask the above questions for 1-cylinder pillowcase-tiled surfaces or to drop the non-cyclic assumption and ask for cyclic covers instead.

As we saw, a translation cover of a 1-cylinder PTS is a geminal STS. We can ask if in odd degrees every geminal STS arises in this way.

Question 7.23. Does every odd-degree geminal square-tiled surface admit a double cover to a 1-cylinder pillowcase-tiled surface?

Finally, Abdalla–Brown's construction is up to now the only known construction to obtain non-cyclic 1-cylinder PTS, and hence non-cyclic geminal STS. Maybe slightly informal questions one can ask are the following.

Question 7.24. Are there other constructions apart from Abdalla–Brown's one that yield non-cyclic 1-cylinder pillowcase-tiled surfaces? Is there a direct construction of non-cyclic geminal square-tiled surfaces that does not use the intermediate step of a pillowcase-tiled surface?

Bibliography

[AB25] Malak Abdalla and Gabriela Brown. An infinite family of non-cyclic 1-cylinder pillowcase-tiled surfaces. *arXiv:2510.18255*, 2025.

[AGY06] Artur Avila, Sébastien Gouëzel, and Jean-Christophe Yoccoz. Exponential mixing for the Teichmüller flow. *Publications mathématiques de l'IHÉS 104, 143-211*, 2006.

[AH15] Tarik Aougab and Shinnyih Huang. Minimally intersecting filling pairs on surfaces. *Algebraic & Geometric Topology 15, 903-932*, 2015.

[AM24] Jayadev S. Athreya and Howard Masur. *Translation surfaces*. American Mathematical Society, Graduate Studies in Mathematics 242, 2024.

[AMN22] Tarik Aougab, William Menasco, and Mark Nieland. Origamis associated to minimally intersecting filling pairs. *Pacific Journal of Mathematics 317, No. 1, 1-20*, 2022.

[AT17] Tarik Aougab and Samuel Taylor. Pseudo-Anosovs optimizing the ratio of teichmüller to curve graph translation length. In *In the tradition of Ahlfors-Bers, VII*. Contemporary Mathematics 696, 17-28, 2017.

[AW22] Paul Apisa and Alex Wright. Generalizations of the Eier-legende–Wollmilchsau. *Cambridge Journal of Mathematics 10, 859-933*, 2022.

[AW23] Paul Apisa and Alex Wright. High rank invariant subvarieties. *Annals of Mathematics 198, 657-726*, 2023.

[AW24] Paul Apisa and Alex Wright. Reconstructing orbit closures from their boundaries. *Memoirs of the American Mathematical Society 298, No. 1487*, 2024.

[Bad25a] Philipp Bader. Commutators of n-cycles in the symmetric group. *Not intended for publication*, 2025. URL: `https://arxiv.org/abs/2509.22749`.

[Bad25b] Philipp Bader. Comparing Teichmüller and curve graph translation lengths. *Algebraic & Geometric Topology (to appear)*, 2025+. URL: `https://arxiv.org/abs/2501.16563`.

[BBF14] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. Constructing group actions on quasi-trees and applications to mapping class groups. *Publications mathématiques de l'IHÉS 122, 1-64*, 2014.

[Ber85] Lipman Bers. An inequality for riemann surfaces. *Differential Geometry and Complex Analysis, 87-93*, 1985.

[BF04] Tara Brendle and Benson Farb. Every mapping class group is generated by 6 involutions. *Journal of Algebra 278, 187-198*, 2004.

[BF06] Jeffrey Brock and Benson Farb. Curvature and rank of Teichmüller space. *American Journal of Mathematics 128, 1-22*, 2006.

[BK24] Hyungryul Baik and Changsub Kim. On translation lengths of pseudo-Anosov maps on the curve graph. *Bulletin of the Korean Mathematical Society 61, No. 3, 585–595*, 2024.

[BKMM12] Jason Behrstock, Bruce Kleiner, Yair Minsky, and Lee Mosher. Geometry and rigidity of mapping class groups. *Geometry & Topology 16, 781-888*, 2012.

[BL12] Corentin Boissy and Erwan Lanneau. Pseudo-Anosov homeomorphisms on translation surfaces in hyperelliptic components have large entropy. *Geometric and Functional Analysis 22, 74–106*, 2012.

[BL22] Corentin Boissy and Erwan Lanneau. Length spectrum of hyperelliptic components. *Journal of the European Mathematical Society 24, No. 6, 1839-1892*, 2022.

[Bow08] Brian H. Bowditch. Tight geodesics in the curve complex. *Inventiones mathematicae 171, 281–300*, 2008.

[Bro03] Jeffrey F. Brock. The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores. *Journal of the American Mathematical Society 16, No. 3, 495-535*, 2003.

[CCC80] S. Chowla, J. Cowles, and M. Cowles. On the number of conjugacy classes in $\mathrm{SL}(2,\mathbb{Z})$. *Journal of number theory 12, 372-377*, 1980.

[CS23] Aaron Calderon and Nick Salter. Framed mapping class groups and the monodromy of strata of abelian differentials. *Journal of the European Mathematical Society 25, 4719-4790*, 2023.

[DGZZ20] Vincent Delecroix, Elise Goujard, Peter Zograf, and Anton Zorich. Contribution of one-cylinder square-tiled surfaces to Masur-Veech volumes. In *Quelques aspects de la théorie des systèmes dynamiques : un hommage à Jean-Christophe Yoccoz (Volume I)*. Astérisque, 2020.

[DM69] Pierre Deligne and David Mumford. The irreducibility of the space of curves of given genus. *Publications mathématiques de l'IHÉS 36, 75-109*, 1969.

[EMM15] Alex Eskin, Maryam Mirzakhani, and Amir Mohammadi. Isolation, equidistribution and orbit closures for the $\mathrm{SL}(2,\mathbb{R})$ action on moduli space. *Annals of Mathematics 182, 673-721*, 2015.

[EO01] Alex Eskin and Andrei Okounkov. Asymptotics of numbers of branched coverings of a torus and volumes of moduli spaces of holomorphic differentials. *Inventiones mathematicae 145, 59-103*, 2001.

[Fil16] Simion Filip. Splitting mixed hodge structures over affine invariant manifolds. *Annals of Mathematics 183, 681-713*, 2016.

[FM12] Benson Farb and Dan Margalit. *A primer on mapping class groups*. Princeton University Press, Princeton Mathematical Series 49, 2012.

[Fri85] David Fried. Growth rate of surface homeomorphisms and flow equivalence. *Ergodic Theory and Dynamical Systems 5, 539-563*, 1985.

[GHKL13] Vaibhav Gadre, Eriko Hironaka, Richard P. Kent, and Christopher Leininger. Lipschitz constants to curve complexes. *Mathematical Research Letters 20, No. 4, 647-656*, 2013.

[Gia24] Riccardo Giannini. Monodromy Kernels for Strata of Translation Surfaces. *International Mathematical Research Notices 2024, No. 10, 8716-8733*, 2024.

[Gro87] Mikhael Gromov. Hyperbolic groups. *Essays in Group Theory (S.M. Gersten), MSRI Publications 8, 75-263*, 1987.

[GT11] Vaibhav Gadre and Chia-Yen Tsai. Minimal pseudo-Anosov translation lengths on the complex of curves. *Geometry & Topology 15, 1297-1312*, 2011.

[Hat02] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.

[Hem01] John Hempel. 3-manifolds as viewed from the curve complex. *Topology 40, 631-657*, 2001.

[Hir09] Eriko Hironaka. What is Lehmer's number? *Notices of the American Mathematical Society 56, No. 3, 374-375*, 2009.

[HLS00] Allen Hatcher, Pierre Lochak, and Leila Schneps. On the Teichmüller tower of mapping class groups. *Journal für reine und angewandte Mathematik 521, 1-24*, 2000.

[HS06] Pascal Hubert and Thomas Schmidt. An introduction to Veech surfaces. In *Handbook of dynamical systems 1B, 501-526*. Amsterdam: Elsevier, 2006.

[HS08] Frank Herrlich and Gabriela Schmithüsen. An extraordinary origami curve. *Mathematische Nachrichten 281, No. 2, 219-237*, 2008.

[Irm15] Ingrid Irmer. Stable lengths on the pants graph are rational. *New York Journal of Mathematics 21, 1153-1168*, 2015.

[Iva92] Nikolai V. Ivanov. *Subgroups of Teichmüller modular groups*. American Mathematical Society, Translations of Mathematical Monographs 115, 1992.

[Iva97] Nikolai V. Ivanov. Automorphisms of complexes of curves and of Teichmüller spaces. *International Mathematical Research Notices 14, 651-666*, 1997.

[Jef21] Luke Jeffreys. Single-cylinder square-tiled surfaces and the ubiquity of ratio-optimising pseudo-Anosovs. *Transactions of the American Mathematical Society 374, No. 8, 5739-5781*, 2021.

[Kon92] Maxim Kontsevich. Intersection Theory on the Moduli Space of Curves and the Matrix Airy Function. *Communications in Mathematical Physics 147, 1-23*, 1992.

[Kor04] Mustafa Korkmaz. Generating the surface mapping class group by two elements. *Transactions of the American Mathematical Society 357, No. 8, 3299-3310*, 2004.

[KZ03] Maxim Kontsevich and Anton Zorich. Connected components of the moduli spaces of abelian differentials with prescribed singularities. *Inventiones mathematicae 153, 631-678*, 2003.

[Lin74] Michele Linch. A comparison of metrics on Teichmüller space. *Proceedings of the American Mathematical Society 43, No. 2*, 1974.

[LOST09] Martin Walter Liebeck, Eamonn Anthony O'Brien, Aner Shalev, and Pham Huu Tiep. The ore conjecture. *Journal of the European Mathematical Society 12, 939–1008*, 2009.

[Mar04] Dan Margalit. Automorphisms of the pants complex. *Duke Mathematical Journal 121, No. 3, 457-479*, 2004.

[Mar19] Dan Margalit. Problems, Questions, and Conjectures about Mapping Class Groups. In *Breadth in Contemporary Topology*, Proceedings of Symposia in Pure Mathematics 102, 157-186. American Mathematical Society, 2019.

[McM22] Curtis McMullen. Billiards and Teichmüller curves. *Bulletin of the American Mathematical Society 60, No. 2*, 2022.

[Mir08] Maryam Mirzakhani. Growth of the number of simple closed geodesics on hyperbolic surfaces. *Annals of Mathematics 168, 97-125*, 2008.

[MM99] Howard Masur and Yair Minsky. Geometry of the complex of curves I: Hyperbolicity. *Inventiones Mathematicae 138, 103–149*, 1999.

[MM00] Howard Masur and Yair Minsky. Geometry of the complex of curves II: Hierarchical structure. *Geometric and Functional Analysis 10, 902-974*, 2000.

[MMW17] Curtis McMullen, Ronen Mukamel, and Alex Wright. Cubic curves and totally geodesic subvarieties of moduli space. *Annals of Mathematics 185, 957-990*, 2017.

[MP19] Maryam Mirzakhani and Bram Petri. Lengths of closed geodesics on random surfaces of large genus. *Commentarii Mathematici Helvetici 94, No. 4, 869-889*, 2019.

[MT02] Howard Masur and Serge Tabachnikov. Rational billiards and flat structures. In *Handbook of dynamical systems 1A, 1015-1089*. Elsevier Science, 2002.

[Ore51] Oystein Ore. Some remarks on commutators. *Proceedings of the American Mathematical Society 2, 307-314*, 1951.

[Pan20] Joshua Pankau. Salem number stretch factors and totally real fields arising from Thurston's construction. *Geometry & Topology 24, 1695-1716*, 2020.

[Pen88] Robert C. Penner. A construction of pseudo-Anosov homeomorphisms. *Proceedings of the American Mathematical Society 310, No. 1, 179-197*, 1988.

[Pen91] Robert C. Penner. Bounds on least dilatations. *Proceedings of the American Mathematical Society 113, No. 2, 443-450*, 1991.

[PH92] Robert Penner and John Harer. *Combinatorics of train tracks*. Princeton University Press, Annals of Mathematics Studies 125, 1992.

[Rau79] Gérard Rauzy. Échanges d'intervalles et transformations induites. *Acta Arithmetica 34, No. 4, 315-328*, 1979.

[Smi00] John Smillie. The dynamics of billiard flows in rational polygons. In *Dynamical Systems*. Encyclopedia of Math. Sciences 100, 360-382, 2000.

[The22] Andrea Hannah Thevis. *On the geometry of p-origamis and beyond*. PhD thesis, Universität des Saarlandes, 2022.

[Thu88] William Thurston. On the geometry and dynamics of diffeomorphisms of surfaces. *Bulletin of the American Mathematical Society 19, No. 2*, 1988.

[Tsa09] Chia-Yen Tsai. The asymptotic behavior of least pseudo-Anosov dilatations. *Geometry & Topology 13, 2253–2278*, 2009.

[Vav16] Aleš Vavpetič. Commutators of cycles in permutation groups. *Ars Mathematica Contemporanea 10, 67-77*, 2016.

[Vee82] William A. Veech. Gauss measures for transformations on the space of interval exchange maps. *Annals of Mathematics 115, No. 2, 201-242*, 1982.

[Via06] Marcelo Viana. Ergodic theory of interval exchange maps. *Revista Matemática Complutense 19, No. 1, 7-100*, 2006.

[Wat20] Yohsuke Watanabe. Pseudo–Anosov mapping classes from pure mapping classes. *Transactions of the American Mathematical Society 373, No. 1, 419–434*, 2020.

[Wol09] Scott Wolpert. The Weil-Petersson metric geometry. In *Handbook of Teichmüller theory II*. EMS Publishing House, 2009.

[Wri15] Alex Wright. Translation surfaces and their orbit closures: An introduction for a broad audience. *EMS Surveys in Mathematical Sciences 2, No. 1, 63-108*, 2015.

[Yoc10] Jean-Christophe Yoccoz. Interval exchange maps and translation surfaces. In *Homogeneous Flows, Moduli Spaces and Arithmetic*, Clay Mathematics Proceedings 10, 1-69. American Mathematical Society, 2010.

[Zor02] Anton Zorich. Square Tiled Surfaces and Teichmüller Volumes of the Moduli Spaces of Abelian Differentials. In *Rigidity in Dynamics and Geometry, 459-471*. Springer Berlin Heidelberg, 2002.

[Zor06] Anton Zorich. Flat surfaces. In *Frontiers in Number Theory, Physics, and Geometry I, 437-583*. Springer Berlin, 2006.