A short introduction to translation surfaces, Veech surfaces, and Teichmüller dynamics

D. Massart

July 24, 2021

Abstract

We review the different notions about translation surfaces which are necessary to understand McMullen's classification of $\mathrm{GL}_2^+(\mathbb{R})$ -orbit closures in genus two. In Section 2 we recall the different definitions of a translation surface, in increasing order of abstraction, starting with cutting and pasting plane polygons, ending with Abelian differentials. In Section 3 we define the moduli space of translation surfaces and explain its stratification by the type of zeroes of the Abelian differential, the local coordinates given by the relative periods, its relationship with the moduli space of complex structures and the Teichmüller geodesic flow. In Part II we introduce the $\mathrm{GL}_2^+(\mathbb{R})$ -action, and define the related notions of Veech group, Teichmüller disk, and Veech surface. In Section 8 we explain how McMullen classifies $\mathrm{GL}_2^+(\mathbb{R})$ -orbit closures in genus 2: you have orbit closures of dimension 1 (Veech surfaces, of which a complete list is given), 2 (Hilbert modular surfaces, of which again a complete list is given), and 3 (the whole moduli space of complex structures). In the last section we review some recent progress in higher genus.

keywords: translation surface, Veech surface, Abelian differential, moduli space
AMS classification : 37D40, 30F60, 32G15

Contents

I	What is a translation surface?	3
1	Introduction	3
2	Three different definitions of a translation surface	3
2.1	The most hands-on definition: polygons with identifications	3
2.2	Main examples of translation surfaces	6
2.3	Definition through an atlas	8
2.4	Definition of a translation surface as a holomorphic differential	10
2.5	Definition of a half-translation surface as a quadratic differential	14
2.6	Dynamical system point of view	15

3	Moduli space	16
3.1	Strata	17
3.2	Period coordinates	18
3.3	Dimension of the strata	19
3.4	Quadratic differentials	19
3.4.1	Dimension of strata of quadratic differentials	19
3.5	Another look at the dimension of $\mathcal{H}_g$	20
4	The Teichmüller geodesic flow	21
4.1	Teichmüller's theorem	21
4.2	Masur's criterion	22
II	Orbits of the $GL_2^+(\mathbb{R})$-action	22
5	Veech groups	23
5.1	The Veech group of the square torus is $SL_2(\mathbb{Z})$	24
5.2	Veech groups of the regular polygons	24
5.3	Veech groups of square-tiled surfaces	25
6	Teichmüller disks	27
6.1	Example: the Teichmüller disk of the torus	28
6.1.1	Hyperbolic metric vs. Teichmüller metric	30
6.1.2	Hyperbolic geodesics	30
6.1.3	Horocycles	31
6.1.4	Relationship between the asymptotic behaviour of hyperbolic geodesics, and the dynamic behaviour of Euclidean geodesics	32
6.2	Examples of higher genus Teichmüller disks	32
6.2.1	Three-squared surfaces	32
6.2.2	Regular polygons	33
7	Veech surfaces	34
7.1	The Smillie-Weiss theorem	35
7.2	The Veech alternative	36
7.2.1	Examples of Veech surfaces	36
8	Classification of orbits in genus two	36
8.1	Square-tiled surfaces	36
8.2	A homological detour	37
8.2.1	The tautological subspace	37
8.2.2	The trace field	38
8.2.3	The Jacobian torus	40
8.2.4	Real multiplication	41

Part I

What is a translation surface?

1 Introduction

Translation surfaces are a great pedagogical tool at almost every level of education. At the elementary, or secondary, level, you are impeded by the lack of embedding into 3-space, but billiards provide plenty of access points into geometry. At the undergraduate level, they help in topology, to explain quotient spaces, or in differential geometry, to explain atlases and transition maps, or in complex analysis, with meromorphic functions. At the graduate level, the doors of abstraction burst open, and in a few light steps you get to the enchanted gardens of moduli spaces, Teichmüller theory, and elementary algebraic geometry. The beauty of the subject lies in the interplay between the basic (cutting and pasting little pieces of paper) and the abstract (moduli spaces of Riemann surfaces).

There are already many wonderful introductions to translation surfaces (for instance, [60], [13], [18], [38], or [58]), and this one is in no way meant as a substitute to any of them, but it could be used as a stepping stone. It is specifically aimed at beginning graduate students. Mostly, it consists of what, given two (or five) hours and a blackboard, I would tell a student looking for a thesis subject and eager to know what translation surfaces are about.

I have tried to include (sketches of) proofs of two types of results: first, those that are too elementary to be included in the aforementioned introductory papers, but can still cause a good deal of head-scratching; and some, not uncommon in the subject, that are of the “takes genius to see it, but not that hard once you know it” variety.

I thank Erwan Lanneau for many useful conversations, the editor of the present volume, and Pablo Montealegre for their careful reading of the manuscript, and Smail Cheboui for letting me use some of the beautiful drawings of [5] (the ugly ones are mine).

2 Three different definitions of a translation surface

2.1 The most hands-on definition: polygons with identifications

Let P be a polygon in the Euclidean plane, not necessarily convex, maybe not even connected, but with its sides pairwise parallel and of equal length. Glue each side to a parallel side of equal length. There are several ways of doing this, so let’s be a bit more careful.

First, let us make the convention that we orient the sides of a polygon so that the interior of the polygon lies to the left.

Now let us identify the sides in pairs, in such a way that the resulting surface is orientable (see Figure 1 for a non-orientable example). One way to think of it is to

imagine that the upside of the polygon is painted red, while the downside is painted black. The surface is orientable if it has one red side, and one black side. This means that each edge is glued to another edge in such a way that the colors match, so the arrows on the edges do not match.

Let us assume, furthermore, that we glue each edge to a parallel edge of equal length. Let n be the number of edges. Thus there exists some permutation σ of $\{1, \dots, n\}$ such that $v_{\sigma(i)} = \pm v_i$, for all $i = 1, \dots, n$. Since we must oppose the arrows when gluing, there are two possible cases: either $v_{\sigma(i)} = -v_i$, in which case the two edges are glued by translation (see Figure 1, top right), or $v_{\sigma(i)} = v_i$, in which case the edges are glued with a half-turn (see Figure 1, bottom).

Formally, you define an equivalence relation $\sim$ on P , and you are considering the quotient space $X = P / \sim$. Although the polygon itself may not be connected, we require the resulting quotient space to be connected, just because if it is not, we study its connected components separately. What you get is a compact, orientable manifold of dimension two (a surface, for short).

Figure 1: Here the arrows indicate the gluing, not the orientation of the boundary, so arrows must match under the gluing. Note that in the case of the sphere, the edges are not glued by translation, but with a half-turn.

Here is why the quotient space is a manifold.

What we need is to find, for each $x \in X$, a neighborhood U_x of x in X , and a chart $\phi_x : U_x \rightarrow \mathbb{C}$, such that for any x, y in X , if $U_x \cap U_y \neq \emptyset$, then $\phi_x \circ \phi_y^{-1}$ preserves whatever structure it is that you want your manifold to come with (if you want a differentiable manifold, you require $\phi_x \circ \phi_y^{-1}$ to be differentiable, if you want a complex manifold, you require it to be holomorphic, and so on).

If x is the image in X of an interior point $\bar{x}$ of P , then x has a neighborhood U_x in X which is the homeomorphic image in X of an open neighborhood $\bar{U}_x$ of $\bar{x}$ in P (hence in $\mathbb{C}$). We take the inverse quotient map $U_x \rightarrow \bar{U}_x$ as the chart ϕ_x .

Shorter and somewhat sloppier version: the quotient map, restricted to the interior of P , is a homeomorphism. Identify x with its pre-image. Take a neighborhood U_x of x which is contained in the interior of P , and take the identity as a chart.

If x is the image in X of an interior point of an edge of P , then x has only two pre-images in P , and x has a neighborhood in X whose pre-image in P is the reunion of two half-disks of equal radius and parallel diameters. If the two edges are glued by translation, we define a chart by the identity on one of the half-disks, and a translation on the other half-disk. If the two edges are glued with a half-turn, we define a chart by the identity on one of the half-disks, and a translation composed with a half-turn on the other half-disk.

If x is the image in X of a vertex of P , then all pre-images in P of x are vertices $P_1, \dots, P_k$ of P , because we glue edges of equal length, and x has a neighborhood U_x in X whose pre-image in P is a union of circular sectors S_i , $i = 1, \dots, k$, with radius r and straight boundaries e_i and f_i , so that f_i is parallel to e_{i+1} , and f_k is parallel to e_1 . The angle θ_i of the sector S_i is the angle between two adjacent (at P_i) edges of P . Note that since f_k is parallel to e_1 , the angles θ_i sum to a multiple of π , say $p\pi$. Furthermore, if the edges f_k and e_1 are glued by a translation, the angles θ_i sum to a multiple of 2π . If the edges f_k and e_1 are glued with a half-turn, the angles θ_i sum to an odd multiple of π . Denote

$$\Theta_i = \sum_{j=1}^{i-1} \theta_j.$$

We define a chart which takes each S_i to a circular sector with vertex at 0, by

$$P_i + \rho \exp i(\theta + \Theta_i) \mapsto \rho \exp i \frac{2}{p}(\theta + \Theta_i), \text{ for } 0 \leq \rho \leq r, 0 \leq \theta \leq \theta_i.$$

We say a point in X which is the image of an interior point of P is of type I, a point in X which is the image of an interior point of an edge of P is of type II, and a point in X which is the image of a vertex of P is of type III.

If x and y are both of type I, and $U_x \cap U_y \neq \emptyset$, then $\phi_x \circ \phi_y^{-1}$ is the identity, which preserves every structure imaginable.

If x and y are both of type I or II, and $U_x \cap U_y \neq \emptyset$, then $\phi_x \circ \phi_y^{-1}$ is either the identity, or a translation, or a "half-translation" $z \mapsto -z + c$, all of which preserve almost every structure imaginable.

If x is of type III and y is of type I, II, or III and $U_x \cap U_y \neq \emptyset$, then $\phi_x \circ \phi_y^{-1}$ is holomorphic (it is essentially a branch of the $p/2$ -th root). If both x and y are of type III, we choose U_x and U_y so that $U_x \cap U_y = \emptyset$.

Thus we have defined a complex structure on X (an atlas with holomorphic transition maps), but this does not tell the whole story. Denote $\Sigma = \{x_1, \dots, x_k\}$ the set of all images in X of the vertices of P , then if all edges are glued by translation, $X \setminus \Sigma$ has an atlas whose transition maps are translations, whence the name **translation surface**. Among the three surfaces in Figure 1, only the torus is a translation surface.

If some pair of edges is glued with a half turn, $X \setminus \Sigma$ has an atlas whose transition maps take the form $z \mapsto \pm z + c$, in which case X is called a **half-translation surface**. The sphere in Figure 1 is a half-translation surface. The Klein bottle is not orientable, so it is neither a translation nor a half-translation surface.

Points of type III in X with a total angle $> 2\pi$ will hereafter be called **singularities**. While being of type III depends on the polygon, being a singularity only depends on the quotient space X .

The fact that translations and the $z \mapsto -z$ map preserve almost every structure imaginable on the plane (complex, metric, you name it) entails that translation and half-translation surfaces come with a lot of structure, and part of the appeal of this theory is that you can view it from so many different angles.

2.2 Main examples of translation surfaces

Before proceeding further, let us introduce our favorite examples. Probably the first thing that comes to mind after seeing the definition is an even-sided, regular polygon.

The surface obtained by identifying opposite sides of a square is a torus.

The surface obtained by identifying opposite sides of a regular hexagon is also a torus, as may be seen by computing the Euler characteristic: one two-cell (the hexagon itself), three edges (one for each pair of opposite edges of the hexagon), and two vertices (if the vertices of the hexagon are cyclically numbered, all even (resp. odd)-numbered vertices are identified into one point), so the Euler characteristic is zero.

For $n > 1$, the surface obtained by identifying opposite sides of a regular $4n$ -gon has genus n , with all vertices identified into one point, with a total angle $(4n - 2)\pi$. The surface obtained by identifying opposite sides of a regular $4n + 2$ -gon has genus n , with two points of type III, both with angle $2n\pi$.

Another surface of interest is the double $(2n + 1)$ -gon, which, from the combinatorial viewpoint, is the same as the regular $4n$ -gon, so it has genus n , with all vertices identified into one point, with a total angle $(4n - 2)\pi$. See the double pentagon on Figure 2.

Our second main example is a generalization of the square: imagine that instead of just one square, you have a finite collection of same-sized squares, each of which has its edges labeled “right”, “left”, “top”, “bottom”, and glue each right (resp. left) edge with a left (resp. right) edge, and each top (resp. bottom) edge with a bottom (resp. top) edge. Such surfaces are called **square-tiled**, or sometimes **origami**. We have already seen the one-squared surface, which is a torus. Two-squared surfaces are also tori (compute the Euler characteristic). The first examples of genus 2 are three-squared (see Figure

Figure 2: the double pentagon: different polygons, same surface

3). They have only one point of type III, with total angle 6π . A four-squared, genus 2 surface, with two points of type III, each with total angle 4π , is shown in Figure 4.

Note that according to common usage, square-tiled surfaces are translation surfaces, so the Klein bottle and the sphere in Figure 1 are not square-tiled surfaces, even though they are actually tiled by squares.

Unlike regular polygons, which are somewhat few and far between, square-tiled surfaces are swarming all over the place. In fact, any translation surface may be approximated, in a sense that will be made precise later, by square-tiled surfaces, just because any line may be approximated uniformly by a stair-shaped broken line made of horizontal and vertical segments.

A useful feature of square-tiled surfaces is that they are covers of the square torus, ramified over one point. Of course there is nothing special about the square here, each regular polygon comes with its own family of ramified covers.

It is usually not a good idea to try to visualize the surface X in space, if only because, having non-positive curvature, it does not embed isometrically in $\mathbb{R}^3$. Computing the Euler characteristic is the way to go. It is interesting, however, when it lets you visualize a decomposition of the surface into flat cylinders. Figures 5, 6, and 7 show a topological embedding into $\mathbb{R}^3$ of the surface on the left of Figure 3, called $St(3)$ after [46]. Figures 8, 9, and 10 show a topological embedding into $\mathbb{R}^3$ of the surface of Figure 4, called $St(4)$ after [46].

Figure 3: Two different three-squared surfaces of genus two, with $St(3)$ on the left

2.3 Definition through an atlas

What if we took the property we have just proved as a definition? Let us say that a translation surface is a compact manifold X of dimension two, such that there exists a finite subset $\Sigma = \{x_1, \dots, x_k\}$, called the singular set of X , and an atlas of X , such that the transition map between any two charts whose domains do not meet Σ is a translation, and the transition map between a chart whose domain meets the singular set, and a chart whose domain does not, is $z \mapsto z^k$ for some $k \in \mathbb{N}$. Is that an equivalent definition? Meaning, from this data, can we extract a polygon with identifications, so that when we perform the identifications, we get back the translation atlas?

Well, among the many structures translations preserve, there is the Euclidean metric. Therefore we can equip our surface $X \setminus \Sigma$ with a Riemannian metric which is locally Euclidean, so its geodesics are locally straight lines. The Riemannian metric may not extend to the whole of X , but the distance function does. In particular we can draw geodesics between any two singularities (elements of Σ).

Now, draw geodesics between singularities, as many (but finitely many) of them as you like, as long as the connected components of the complement in X of those geodesics are simply connected. Being simply connected, they may be developed to the plane, to polygons. Then, apply the construction of Subsection 2.1 to this polygon (remember, I never said the polygon from which X is glued should be connected).

The atlas we obtain from gluing back may not be the same atlas we started with,

Figure 4: $St(4)$, a four-squared surface of genus 2, with two singularities

Figure 5: Identifications for the surface $St(3)$, stage 1

but they share a common maximal atlas. So, with the usual polite fiction of a manifold as a maximal atlas, the new definition is equivalent to the first one.

Figure 6: Identifications for the surface $St(3)$, stage 2

Although we shall try to ignore it as much as possible, we have to mention the following annoying

Fact 2.1. *Given a homeomorphism f of the surface X , and an atlas $(U_i, \phi_i)_i$ on X , the charts ϕ_i may be pre-composed with f , thus giving a new atlas $(U_i, \phi_i \circ f)_i$, which in general is not compatible with the first one.*

This fact is annoying because it only makes a difference if you tag each point in X with a label and are interested in tracking each individual label. The usual way around it is through another polite fiction: decide two maximal atlases are equivalent if they may be deduced from each other by pre-composition with a homeomorphism (sometimes it is convenient to restrict to homeomorphisms which are isotopic to the identity map). Then define your structure as this most unfathomable object: an equivalence class of maximal atlases.

2.4 Definition of a translation surface as a holomorphic differential

We have seen that a translation surface (as in our first definition) has a complex structure, which is essentially a protractor (a way to measure angles between tangent vectors at any point). But it actually has much more: a graduated ruler (since we can measure distances), and a compass, whose needle points North, wherever you are, except at the singular set. This is because our polygon is a subset of the Euclidean plane, so we can choose any direction we want as the North, and since all identifications are translations

Figure 7: Identifications for the surface $St(3)$, stage 3

Figure 8: Identifications for the surface $St(4)$, stage 1

Figure 9: Identifications for the surface $St(4)$, stage 2

(except at the singular set), this goes down to the quotient space X . At a vertex of P , the needle of the compass gets a little dizzy, but not in the same way as an actual compass at the magnetic pole of the Earth would: while the latter has infinitely many directions to choose from, our compass only has finitely many. If the angle around the singular point x_i is $2k\pi$, then there are, so to say, k different northes at x_i .

Formally, the package (protractor, ruler, compass) is called a **holomorphic, or Abelian, differential**, and usually denoted ω . If you have a complex manifold X , a holomorphic differential is just, for every $x \in X$, the choice of a complex linear map $l(x)$ from the tangent space to X at x , to $\mathbb{C}$, with the condition that $l(x)$ depends homomorphically on x . Our X 's have complex dimension one, and the only complex linear maps from $\mathbb{C}$ to itself are the maps $z \mapsto \lambda z$, with $\lambda \in \mathbb{C}$. When we think of the identity as a holomorphic differential on $\mathbb{C}$, we denote it dz , so we will usually think of a holomorphic differential ω on X as $f(z)dz$, with f holomorphic, in local charts, with the usual compatibility requirement that if two chart domains intersect, and ω reads $f(z)dz$ in one chart, and $g(z)dz$ in the other, and $T(z)$ is the transition map, then $f(T(z))T'(z) = g(z)$ (f and g should be thought of as derivatives, since $f(z)dz$ is a differential, so we just apply the chain rule for derivatives).

If you would like to see it put another way: if $x \in X$, and v is a tangent vector to X at x , and we have two different charts ϕ and ψ at x , such that ω reads $f(z)dz$ in the chart ϕ , and $g(z)dz$ in the chart ψ , then $\omega(x).v$ reads $f(\phi(x))\phi'(x).v$ in the chart ϕ , and $g(\psi(x))\psi'(x).v$ in the chart ψ . Since the evaluation of ω does not depend on the chart, we have $f(\phi(x))\phi'(x) = g(\psi(x))\psi'(x)$. Now let T be the transition map $\phi \circ \psi^{-1}$, then, setting $z = \psi(x)$, we have $f(T(z))\phi'(\psi^{-1}(z)) = g(z)\psi'(\psi^{-1}(z))$, which is $f(T(z))T'(z) = g(z)$.

Figure 10: Identifications for the surface $St(4)$, stage 3

Not every compact manifold has a non-zero holomorphic differential, for instance, if you try and extend dz to the sphere $\mathbb{CP}^1$, you get a double pole at infinity, so the needle of your compass will act tipsy at infinity, like at the North Pole. Translation surfaces are different from $\mathbb{CP}^1$: the holomorphic differential dz in the plane goes down to a holomorphic differential on the quotient space X , albeit with zeroes. The order of the zeroes is given by the angle at the singularities, as follows from the chain rule: if g is a chart around a singular point, and f is a regular chart, then the transition map is $z \mapsto z^k$, so $kf(z^k)z^{k-1} = g(z)$, in particular g has a zero of order $k - 1$ at the singular point.

Assume ω reads $g(z)dz$ in some chart (U, ϕ) at $x_0 \in X$, with $\phi(x_0) = 0$ and $g(0) \neq 0$, so we say that x_0 is a regular point of ω . Then we may take a primitive of ω as a new chart at x_0 :

$$\begin{aligned} \psi : V &\longrightarrow \mathbb{C} \\ x &\longmapsto \int_{x_0}^x \omega \end{aligned}$$

where V is some neighborhood of x_0 , with $V \subset U$. The map ψ is a local biholomorphism because ω does not vanish at x_0 . Let G be the local primitive of g , defined in V , such that $G(0) = 0$. Note that $\psi(x) = G(\phi(x))$, so the transition map $\psi \circ \phi^{-1}$ between ϕ and ψ is G . Thus, by the chain rule, assuming ω reads $f(z)dz$ in the chart ψ , we have

$$f(\psi \circ \phi^{-1}(z)) G'(z) = g(z)$$

so f is constant $= 1$, that is, ω reads dz in the chart ψ .

Now assume we have two different charts, at different, regular, points x and y , whose domains intersect, and assume that in both charts ω reads dz . Let ϕ be the transition map between the two charts. Then, by the chain rule, we have $\phi' = 1$, that is, ϕ is a translation. Therefore, from a holomorphic differential, we recover a translation atlas, thus proving the equivalence of our three definitions.

2.5 Definition of a half-translation surface as a quadratic differential

Now assume that instead of a translation surface, we have a half-translation surface. What could play the role of a holomorphic differential in that case? Well, the $z \mapsto -z$ map does not preserve the linear form dz , but it does preserve the quadratic form dz^2 .

If you have a complex manifold X , a **holomorphic quadratic differential** is to a holomorphic differential what a quadratic form is to a linear form, so it is, for every $x \in X$, the choice of a complex-valued quadratic form $l(x)$ from the tangent space to X at x , to $\mathbb{C}$, with the condition that $l(x)$ depends holomorphically on x . Our X 's have complex dimension one, and the only complex-valued quadratic forms from $\mathbb{C}$ to itself are the maps $z \mapsto \lambda z^2$, with $\lambda \in \mathbb{C}$. When we think of the square map as a quadratic differential on $\mathbb{C}$, we denote it dz^2 , so we will usually think of a holomorphic quadratic differential ω on X as $f(z)dz^2$, with f holomorphic, in local charts, with the usual compatibility requirement that if two chart domains intersect, and ω reads $f(z)dz^2$ in one chart, and $g(z)dz^2$ in the other, and $\phi(z)$ is the transition map, then $f(\phi(z))(\phi'(z))^2 = g(z)$ (f and g should be thought of as squares of derivatives). Actually we should be allowing simple poles for f as well, for instance the half-translation sphere in Figure 1 has four simple poles, but we shall not dwell on that.

Note that given a holomorphic differential ω , written $f(z)dz$ in local coordinates, the formula $f(z)^2 dz^2$ yields a quadratic holomorphic differential, so the set of holomorphic differentials (modulo identification of a differential with its polar opposite) identifies with a subset of the set of quadratic holomorphic differentials. It is a proper subset, however, because not every holomorphic function is a square.

While not every quadratic differential is a square, every quadratic differential becomes a square in a suitable double cover of X , by the classical Riemann surface construction of

the square root function. For instance, the half-translation sphere in Figure 1 is covered, twofold, by a 4-square torus, in fact the involution of the covering is the hyperelliptic involution of the torus. It is often convenient to deal with translation surfaces only, taking double covers when need be.

2.6 Dynamical system point of view

So far we have seen a translation surface as a geometric, or complex-analytic, object, but it is interesting to view it as a dynamical system.

We have seen that the North is well-defined on a translation surface X , outside the singularities; so we have a local flow ϕ_t on X , usually called the **vertical flow**, defined by “walk North for t miles”. The flow is complete (i.e well-defined for any $t \in \mathbb{R}$) if we remove the finite, hence negligible, union of orbits which end in a singularity (which we call **singular orbits**).

Of course we can multiply our Abelian differential by $e^{i\theta}$ for any θ , thus turning the vertical by θ , so in fact we have a family of flows indexed by $\theta \in [0, 2\pi]$.

We ask the usual dynamical questions: are there periodic orbits? if so, can we enumerate them? are there invariant measures not supported on periodic orbits? If so, how many of them?

In the case of the square torus, the questions are readily answered: by taking the first return map on a closed transversal, we see that the dynamical properties of the flow in a given direction θ are those of a rotation on the circle. That is, when $\theta/\pi \in \mathbb{Q}$, every orbit is periodic, with the same period.

When $\theta/\pi \in \mathbb{R} \setminus \mathbb{Q}$, every orbit is equidistributed, that is, the proportion of its time it spends in a given interval is the proportion of the circle this interval occupies. In particular the flow in the direction θ is uniquely ergodic, that is, it supports a unique invariant measure, up to a scaling factor.

In the case of square-tiled surfaces, the answer is equally satisfying. Recall that any square-tiled surface X is a ramified cover of the square torus, so the flow in the direction θ on X projects to the flow in the direction θ on the square torus. Therefore, if $\theta/\pi \in \mathbb{Q}$, every orbit of the flow in the direction θ on X projects to a closed orbit in the square torus, so it must be periodic itself. The period need not be the same for all orbits, though, for instance the vertical flow on $X = St(3)$ has orbits of period 1 and 2 (see Figure 3). In fact $X = St(3)$, minus the set of singular orbits, is the reunion of two cylinders, one made with orbits of length 2, and the other made of orbits of length 1.

The answer when $\theta/\pi \in \mathbb{R} \setminus \mathbb{Q}$ is a bit trickier, because measures cannot be projected, they may only be lifted, and it is unclear why every invariant measure in X should be the lift of an invariant measure in the torus, but as in the torus case, orbits are equidistributed.

This dichotomy between directions which are **completely periodic**, meaning that the surface decomposes into cylinders of periodic orbits, and directions which are uniquely ergodic, is called the **Veech dichotomy**. In the next sections we are going to see a powerful criterion for unique ergodicity, and apply it to find a larger class of surfaces which satisfy the Veech dichotomy.

In general it is hard to prove that a given surface (say, the double pentagon) satisfies the Veech dichotomy. On the other hand, it is easy to find a surface which does not satisfy the Veech dichotomy (see Figure 11).

In [29] an example is given of a translation surface with a direction θ such that the flow in the direction θ is minimal (every orbit is dense) but not uniquely ergodic, that is, it supports several distinct invariant measures.

Figure 11: Take the surface $St(3)$ and shear the top edge to the right by x , keeping the same gluing pattern. The shaded square projects to the surface as a cylinder of vertical periodic geodesics of length 1, while the rest projects to a cylinder where every vertical geodesic is dense, if x is irrational.

3 Moduli space

Now that we know what a translation structure on a surface is, we would like to know what it means for two translation structures to be the same, or almost the same. Sameness is easy to define, if hard to visualize, through atlases: we say two translation structures (as atlases) are equivalent if they share a common (equivalence class of) maximal translation atlas. Observe that a necessary condition for sameness is that the number of singular points, and the angle around each singular point, are the same.

It is a bit harder to see sameness from polygons, because two polygons which look very different may yield the same surface after identifications, see Figure 2. The correct

definition of sameness is that two polygons yield the same translation surface if one may be cut and pasted into the other, provided the pieces are only re-arranged by translation (no rotation or flip allowed).

It is not always easy to tell when two polygons do *not* yield the same surface, either. The two polygons of Figure 3 are really different surfaces, because the one on the right is foliated, in the vertical direction, by closed geodesics of length 3, while the one on the left is not.

3.1 Strata

The previous discussion suggests lumping together all translation surfaces with the same number of singularities, and equal angles around each singularities. The set of all such translation surfaces is called a **stratum** (term coined by Veech [55], to be explained later). It is usually denoted $\mathcal{H}(k_1, \dots, k_n)$, where $2\pi(k_i + 1)$ is the angle around the i -th singular point. For instance $\mathcal{H}(2)$ is the stratum of surfaces with only one singular point, of angle 6π , while $\mathcal{H}(1, 1)$ is the stratum of surfaces with two singular points, each with angle 4π .

Note that if two translation surfaces lie in the same stratum, they must have the same genus. To see this, let us assume, to begin with, that X is obtained by identifying the sides of a connected $2N$ -gon, with the vertices identifying into the n singularities of an Abelian differential ω .

Then the Euler characteristic $\chi(X)$ of X is $1 - N + n$.

Let us say the angle around each singularity x_i is $2(k_i + 1)\pi$, then the sum of the interior angles of the polygon equals the sum of the angles around the singularities, that is,

$$\sum_{i=1}^n 2(k_i + 1)\pi = (2N - 2)\pi,$$

whence

$$\sum_{i=1}^n k_i = N - 1 - n = -\chi(X) = 2\text{genus}(X) - 2.$$

Now, assume that X is obtained by identifying the sides of a $2N$ -gon with p connected components, each with N_i vertices, for $i = 1, \dots, p$. Then

$$\chi(X) = p - \frac{1}{2} \sum_{i=1}^p N_i + n,$$

and the sum of the interior angles of the polygons is

$$\sum_{i=1}^p (N_i - 2)\pi = \sum_{i=1}^n 2(k_i + 1)\pi$$

whence

$$\sum_{i=1}^n k_i = \frac{1}{2} \sum_{i=1}^p N_i - p - n = -\chi(X).$$

For instance, the stratum $\mathcal{H}(0)$ consists of all flat tori. In genus 2, we have the strata $\mathcal{H}(2)$ or $\mathcal{H}(1, 1)$, and those are the only strata of genus 2, because the only way the number 2 can be partitioned is as $1 + 1$ or $2 + 0$. The surfaces in Figures 2 and 3 lie in $\mathcal{H}(2)$, while the surface in Figure 4 lies in $\mathcal{H}(1, 1)$.

The regular $4n$ -gon, and the double $(2n + 1)$ -gon, lie in the stratum $\mathcal{H}(2n - 2)$. The regular $(4n + 2)$ -gon lies in the stratum $\mathcal{H}(n - 1, n - 1)$. On the other hand, square-tiled surfaces are dense in every stratum.

Now we want to define what it means for translation surfaces to be close, that is, we are looking for a topology on the set of translation surfaces. In fact we can do even better: we may find local coordinates on a given stratum, thus making the stratum (almost) a manifold.

3.2 Period coordinates

Fix a basis for the $\mathbb{Z}$ -module $H_1(X, \Sigma, \mathbb{Z})$. Such a basis may be chosen as (the relative homology classes of) $\alpha_1, \dots, \alpha_{2g}, c_1, \dots, c_{n-1}$, where g is the genus of X , $\alpha_1, \dots, \alpha_{2g}$ are simple closed curves (based at x_1) which generate the absolute homology $H_1(X, \mathbb{Z})$, and for each $i = 1, \dots, n - 1$, c_i is a simple arc joining x_1 to x_{i+1} .

Definition 3.1. *The **period coordinates** of the holomorphic differential ω , with respect to the basis $\alpha_1, \dots, \alpha_{2g}, c_1, \dots, c_{n-1}$, are the $2g + n - 1$ complex numbers*

$$\int_{\alpha_1} \omega, \dots, \int_{\alpha_{2g}} \omega, \int_{c_1} \omega, \dots, \int_{c_{n-1}} \omega.$$

The complex numbers $\int_{\alpha_1} \omega, \dots, \int_{\alpha_{2g}} \omega$ are called the **absolute periods** of ω , because the homology classes $[\alpha_1], \dots, [\alpha_{2g}]$ live in the absolute homology $H_1(X, \mathbb{Z})$, while the complex numbers $\int_{c_1} \omega, \dots, \int_{c_{n-1}} \omega$ are called the **relative periods** of ω , because the homology classes $[c_1], \dots, [c_{n-1}]$ live in the relative homology $H_1(X, \Sigma, \mathbb{Z})$.

Why this is actually a set of local coordinates on the stratum of X is a theorem of [53] (see also [13], section 2.3, and also [59], 55). Let us briefly explain the idea. Assume two holomorphic differential ω_1 and ω_2 lie in the same stratum and have the same period coordinates. Assume that ω_1 and ω_2 are close enough, in some sense, that they can be considered as smooth (not necessarily holomorphic) differential forms on the same surface X (seen as a differentiable manifold), with the same homology basis $\alpha_1, \dots, \alpha_{2g}, c_1, \dots, c_{n-1}$. Each holomorphic differential ω_1 and ω_2 gives you a Euclidean metric on X , for which geodesic representatives of $\alpha_1, \dots, \alpha_{2g}, c_1, \dots, c_{n-1}$ can be found. Cut X open along these geodesics, in each of the two Euclidean metrics, you get a pair of $2(2g + 2n - 1)$ -gons whose sides are represented by the period coordinates of the holomorphic differentials. Since ω_1 and ω_2 have the same period coordinates, the two polygons are isometric, by a translation, so ω_1 and ω_2 give the same translation structure.

Example: assume (X, ω) is a square-tiled surface. Then the homology basis may be chosen so that the curves $\alpha_1, \dots, \alpha_{2g}, c_1, \dots, c_{n-1}$ lie in the sides of the squares. Thus the period coordinates lie in $\mathbb{Z}[i]$. Conversely, if all period coordinates lie in $\mathbb{Z}[i]$, then X may be cut into a plane polygon, all of whose vertices lie in $\mathbb{Z}[i]$, so X is actually a square-tiled surface.

3.3 Dimension of the strata

Thus we know the (complex) dimension of the stratum $\mathcal{H}(k_1, \dots, k_n)$ is $2g + n - 1$. The largest stratum is the one with $2g - 2$ singularities of index 1, its dimension is $4g - 3$. The smallest stratum is the one with only one singularity of index $2g - 2$, its dimension is $2g$. The strata are not disconnected from each other, each stratum but the largest one lies in the closure of one or several of the larger ones: for instance, if we have a sequence (X, ω_m) of translation surfaces in the stratum $\mathcal{H}(k_1, \dots, k_n)$, and $\int_{c_1} \omega_m \rightarrow 0$ when $m \rightarrow \infty$, then the limit surface lies in $\mathcal{H}(k_1 + k_2, k_3, \dots, k_n)$, because the singularities x_1 and x_2 merge in the limit. This is the reason they are called strata, because they are nested in each other's closure, like matriochka. The union of all strata, with the topology given by the period coordinates, is called the **moduli space of translation surfaces of genus g** , denoted $\mathcal{H}_g$. It contains the largest (or principal) stratum, denoted $\mathcal{H}(1^{2g-2})$, as a dense open subset, so its dimension is $4g - 3$.

While the global topology of $\mathcal{H}_g$ is mysterious, we have a simple compactness criterion (**Mumford's compactness theorem**, [41]) for subsets of $\mathcal{H}_g$: a sequence X_n of elements of $\mathcal{H}_g$, of uniformly bounded area, goes to infinity if and only if there exist closed geodesics in X_n whose lengths go to zero.

3.4 Quadratic differentials

The theory of translation surfaces becomes much more interesting when you think about it in connection with Teichmüller theory.

The set of complex structures on surfaces of genus g is a well-studied object, going back to Riemann (see [42], reprinted in [43], pp.88-144, or, for the faint of heart, [16]). It is worth noting that Riemann actually got there when studying Abelian differentials.

It is called the **moduli space of complex structures of genus g** , and denoted $\mathcal{M}_g$. It has a topology, but actually it has much more structure than that, in fact it is, except at some special points which correspond to very symmetrical surfaces, a complex manifold of dimension $3g - 3$, provided $g > 1$. The genus one case is special and will be discussed in more detail in Section 6.1. The cotangent space to moduli space at a given X is the vector space of holomorphic quadratic differentials (see [16], Theorem 6.6.1).

Recall that we have just computed the dimension of $\mathcal{H}_g$, which we found to be $4g - 3$, while the cotangent bundle of $\mathcal{M}_g$ has dimension $6g - 6$, so the squares of Abelian differentials have codimension $2g - 3$ in the space of quadratic differentials.

3.4.1 Dimension of strata of quadratic differentials

Following [23], we denote $\mathcal{Q}_g$ the space of all quadratic differentials of genus g , and $\mathcal{Q}(k_1, \dots, k_n)$ the set of quadratic differentials with n zeroes, of respective orders $k_1, \dots, k_n$, which are not squares.

Computations identical to those of Subsection 3.1 show that $\mathcal{Q}_g$ stratifies as the union of $\mathcal{Q}(k_1, \dots, k_n)$, with $k_1 + \dots + k_n = 4g - 4$.

It is interesting to compute the dimension of the strata of non-square quadratic differentials, because there is one notable difference with the Abelian case. In the case of the Abelian stratum $\mathcal{H}(k_1, \dots, k_n) \subset \mathcal{H}_g$, we cut our surface into a $2(2g+n-1)$ -gon, and basically said that sides being pairwise equal, you need $2g+n-1$ complex parameters to determine an element of $\mathcal{H}(k_1, \dots, k_n)$. Let us try to apply the same argument to quadratic differentials.

Let us take an element of $\mathcal{Q}(k_1, \dots, k_n)$, and cut it into a polygon, whose vertices correspond to the singularities of the quadratic differential. Let us assume for simplicity that the polygon is connected. Then the sum of the interior angles of the polygon equals the sum of the angles around the singularities, which is $\sum_{i=1}^n (k_i + 2)\pi$. Thus the polygon has $\sum_{i=1}^n (k_i + 2) + 2 = 4g - 4 + 2n + 2$ edges. Since the edges are pairwise equal, this gives us $2g + n - 1$ complex parameters. But then, the largest stratum, with $n = 4g - 4$ and $k_i = 1, i = 1, \dots, 4g - 4$, would have dimension $6g - 5$, instead of $6g - 6$ as befits the cotangent bundle of $\mathcal{M}_g$.

To understand this, we must take a closer look at how we identify the sides of a polygon. Recall that we made the convention that we orient the sides of a polygon so that the interior of the polygon lies to the left. With this convention, a necessary and sufficient condition for plane vectors $v_1, \dots, v_n$ to be the oriented edges of a polygon is that $v_1 + \dots + v_n = 0$.

Next, recall that we identify the sides in pairs, in such a way that the resulting surface is orientable, so each edge is glued to another edge in such a way that the arrows on the edges do not match.

Recall, furthermore, that we glue each edge to a parallel edge of equal length, so there exists some permutation σ of $\{1, \dots, n\}$ such that $v_{\sigma(i)} = \pm v_i$, for all $i = 1, \dots, n$. Since we must oppose the arrows when gluing, there are two possible cases: either $v_{\sigma(i)} = -v_i$, in which case the two edges are glued by translation, or $v_{\sigma(i)} = v_i$, in which case the edges are glued with a half-turn.

In the former case, we may choose v_i freely without being constrained by the relation $v_1 + \dots + v_n = 0$, while in the latter case, it imposes a condition on v_i :

$$-2v_i = v_1 + \dots + \hat{v}_i + \dots + \hat{v}_{\sigma(i)} + \dots + v_n$$

where the hats mean that $v_{\sigma(i)}$ and v_i are omitted in the sum.

So, when some pair of sides is glued with a half-turn, this pair of sides is determined by the others, so we lose one complex parameter. Now, saying that the quadratic differential is not a square is precisely saying that some pair of sides is glued with a half-turn, since we have seen that when every side is glued by translation, we get an Abelian differential. Therefore, the dimension of the stratum $\mathcal{Q}(k_1, \dots, k_n)$ is $2g + n - 2$.

3.5 Another look at the dimension of $\mathcal{H}_g$

Here is another way to understand the fact that $\dim \mathcal{H}_g = 4g - 3$. A translation structure, or holomorphic differential, consists of the following data: a complex structure on a surface X , and a complex-valued differential form, holomorphic for the given complex

structure. As we have seen, a complex structure is determined by $3g - 3$ complex parameters.

Besides, the vector space of holomorphic (for a given complex structure) differentials has complex dimension g . This is because, given a complex structure X , by the Hodge theorem, the complex vector space $H^1(X, \mathbb{C})$, which has dimension $2g$, is the direct sum of two isomorphic summands, the subspace of holomorphic differentials, and the subspace of anti-holomorphic differentials. Thus a translation structure is determined by $3g - 3 + g = 4g - 3$ complex parameters, and $\mathcal{H}_g$ may be viewed as a real vector bundle over $\mathcal{M}_g$, with fiber $H^1(X, \mathbb{R})$.

4 The Teichmüller geodesic flow

4.1 Teichmüller's theorem

Take two complex structures X_1 and X_2 of genus g . Then Teichmüller's theorem says there exists a holomorphic quadratic differential q on X_1 , and a number $t \in \mathbb{R}$, such that, assuming for simplicity that q is the square of a holomorphic differential ω , the 1-differential form with real part $e^t \operatorname{Re} \omega$ and imaginary part $e^{-t} \operatorname{Im} \omega$ is holomorphic with respect to the complex structure X_2 .

Then one may define, at least locally, a distance function by $d(X_1, X_2) := |t|$, and this distance comes with geodesics (shortest paths): if $s \in [0, t]$, and X_s is the complex structure which makes the complex differential with real part $e^s \operatorname{Re} \omega$ and imaginary part $e^{-s} \operatorname{Im} \omega$ holomorphic, then $d(X_1, X_s) = |s|$, so the path $s \mapsto X_s$ is a geodesic.

Teichmüller's distance has an infinitesimal expression, like a Riemannian metric, and more precisely, it is a Finsler metric (a Finsler metric is to a Riemannian metric what a norm is to a Euclidean norm). This was discovered by Teichmüller, see [49], p.26, translated in [50], or just [16], Theorem 6.6.5.

A quadratic differential ω induces a Riemannian (flat except at the singularities) metric on X , in particular it comes with a volume form, so it makes sense to evaluate the total volume of X with respect to q , denoted $\operatorname{Vol}(X, q)$. The map $(X, q) \mapsto \operatorname{Vol}(X, q)$, restricted to the tangent space to $\mathcal{M}_g$ at X , is a norm: it is 1-homogeneous, positive except at $\omega = 0$, and satisfies the triangle inequality. The first two properties are immediate, the last one may warrant a short proof.

First, observe that if the quadratic differential q is $f(z)dz^2$ in some chart, then the volume form of q , seen as a flat metric, is $\frac{i}{2} |f(z)| dz \wedge d\bar{z}$, because $dz \wedge d\bar{z} = -2idx \wedge dy$. This expression is coordinate-invariant, because if q is $g(z)dz^2$ in some other chart, and T is the transition map, then $f(z) = g(T(z))T'(z)^2$, so

$$|f(z)| dz \wedge d\bar{z} = |g(T(z))| T'(z) \overline{T'(z)} dz \wedge d\bar{z} = |g(T(z))| (T'(z) dz) \wedge (\overline{T'(z)} d\bar{z})$$

so the total volume of q may be calculated in local coordinates. Now, assume we have two quadratic differentials $q_1 = f_1(z)dz^2$ and $q_2 = f_2(z)dz^2$, so $q_1 + q_2 = (f_1 + f_2)dz^2$, the triangle inequality in $\mathbb{C}$ yields $|f_1 + f_2| \leq |f_1| + |f_2|$, and, integrating over X , we get $\operatorname{Vol}(X, q_1 + q_2) \leq \operatorname{Vol}(X, q_1) + \operatorname{Vol}(X, q_2)$.

This norm endows $\mathcal{M}_g$ with a metric: if we have a C^1 path $X(t)$ in $\mathcal{M}_g$, its derivative $\dot{X}(t)$ is a quadratic differential, holomorphic at $X(t)$, and we just say that the velocity of X at time t is $\text{Vol}(X, \dot{X}(t))$. This is not a Riemannian metric, because the norm is not Euclidean, the specific name is Finsler metric. In any case it induces a geodesic flow on the cotangent bundle of $\mathcal{M}_g$: start at a complex structure X , in the direction given by a quadratic differential ω , and go to X_t such that the complex quadratic differential with real part $e^t \text{Re } \omega$ and imaginary part $e^{-t} \text{Im } \omega$ is holomorphic with respect to the complex structure X_t .

Most interesting to us here is the fact that $\mathcal{H}_g$, which is a submanifold of the cotangent bundle of $\mathcal{M}_g$, is invariant by the geodesic flow. In fact, we shall see, in Part II, that $\mathcal{H}_g$ is invariant by a much more interesting action.

4.2 Masur's criterion

In dynamical systems, we have a saying, which goes "sow in the parameter space, reap in the phase space". This is particularly relevant to translation surfaces, because the parameter space $\mathcal{H}_g$ comes with so much structure: a geodesic flow, invariant measures, affine coordinates... Masur's criterion (possibly the single most useful result in the theory) is a striking example.

Theorem 4.1 (Masur's criterion, [28]). *Given an Abelian differential ω , the orbit of ω under the Teichmüller geodesic flow is recurrent if and only if the vertical flow of ω is uniquely ergodic.*

The idea is very clearly explained in [39]. We'll see several applications of Masur's criterion, here is one:

Theorem 4.2 ([21]). *Given an Abelian differential ω , for almost every θ , the vertical flow of $e^{i\theta}\omega$ is uniquely ergodic.*

Once we have Masur's criterion, it is easy to get a weak version of Theorem 4.2: for almost every Abelian differential ω , for almost every θ , the vertical flow of $e^{i\theta}\omega$ is uniquely ergodic. The proof is basically just Poincaré's Recurrence Theorem. Of course, to apply it we need an invariant measure of full support and finite total volume, so we apply it, not on $\mathcal{H}_g$, but on the subset $\mathcal{H}_g^1 \subset \mathcal{H}_g$ of Abelian differentials of area 1. It is a theorem by Masur and Veech ([27], [52]) that the Lebesgue measure induced on $\mathcal{H}_g^1$ by the period coordinates has finite total volume.

Part II

Orbits of the $GL_2^+(\mathbb{R})$ -action

Every translation surface may be viewed as a polygon in the Euclidean plane with parallel sides of equal length pairwise identified. The group $GL_2^+(\mathbb{R})$ acts linearly on polygons,

mapping pairs of parallel sides of equal length to pairs of parallel sides of equal length, so it acts on translation surfaces.

Let us take a basis $\mathcal{B} = (\alpha_1, \dots, \alpha_{2g}, c_1, \dots, c_{n-1})$ of $H_1(X, \Sigma, \mathbb{Z})$. If we cut X into a polygon along the curves $\alpha_1, \dots, \alpha_{2g}, c_1, \dots, c_{n-1}$, the period coordinates of (X, ω) in the basis $\mathcal{B}$ are the sides of the polygon (identifying a complex number and its affix). So, if $A \in \mathrm{GL}_2^+(\mathbb{R})$, the period coordinates, in the basis $\mathcal{B}$, of the surface $A.(X, \omega)$, are the complex numbers

$$A. \int_{\alpha_1} \omega, \dots, A. \int_{\alpha_{2g}} \omega, A. \int_{c_1} \omega, \dots, A. \int_{c_{n-1}} \omega,$$

where A acting on a complex numbers means that it acts on its affix. The matrix A induces a homeomorphism (which is actually a diffeomorphism outside the singularities), from the translation surface (X, ω) , to the surface $A.(X, \omega)$.

First, let us observe that strata are invariant under $\mathrm{GL}_2^+(\mathbb{R})$. This is because acting by an element of $\mathrm{GL}_2^+(\mathbb{R})$ on a polygon does not change the order in which sides are identified, which determines the singularities, and the angles around the singularities.

The $\mathrm{SL}_2(\mathbb{R})$ -action contains the geodesic flow, in the following sense: the geodesic, with respect to the Teichmüller metric, which has initial position X , and initial velocity ω , is the orbit of (X, ω) under the diagonal subgroup of $\mathrm{SL}_2(\mathbb{R})$,

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

It is usually preferred to deal with the $\mathrm{GL}_2^+(\mathbb{R})$ -action rather than with the geodesic flow, for the following reason, which is a theorem by Eskin and Mirzakhani [10], henceforth referred to as the Magic Wand, after [61]. While the invariant sets for the geodesic flow may be wild (for instance, Teichmüller disks contain geodesic laminations, which are locally the product of a Cantor set with an interval), the closed invariant sets of the $\mathrm{GL}_2^+(\mathbb{R})$ -action are always nice submanifolds (again, ignoring singularities which occur at very symmetrical surfaces), locally defined by affine equations in the period coordinates. Any attempt at informally describing the proof, assuming the author could do it, would be as long as this paper. A recurring theme is the analogy with Ratner's theorems.

Now, by a result by Masur and (independantly) Veech ([27, 52]), the geodesic flow is ergodic, with respect to a measure of full support in $\mathcal{H}_g$, so almost every orbit is dense; consequently, almost every $\mathrm{GL}_2^+(\mathbb{R})$ -orbit is dense. In fact, any stratum supports a full-support, ergodic measure. This means finding interesting (meaning: other than strata closures and $\mathcal{H}_g$ itself), closed invariant subsets won't be easy.

5 Veech groups

It may happen that for some polygon P and some element A of $\mathrm{GL}_2^+(\mathbb{R})$, both P and $A.P$, after identifications, are just the same translation surface X . This happens when $A.P$ can be cut into pieces, and the pieces re-arranged, by translations, into A .

The **Veech group** of (X, ω) is the subgroup of $\mathrm{GL}_2^+(\mathbb{R})$ which preserves X . Since the Veech group must preserve volumes, it is a subgroup of $\mathrm{SL}_2(\mathbb{R})$, and it turns out to be a Fuchsian group (see [18], Lemma 2). We sometimes denote it $SL(X, \omega)$ after McMullen.

Take a matrix A in $SL(X, \omega)$, then it defines a homeomorphism from X to itself, which we again denote A for simplicity. Let A_* be the linear automorphism of $H_1(X, \Sigma, \mathbb{Z})$ induced by A . Then the period coordinates of (X, ω) , in a basis $\mathcal{B}$ of $H_1(X, \Sigma, \mathbb{Z})$, are exactly the period coordinates of $A.(X, \omega)$ in the basis $(A_*)^{-1}(\mathcal{B})$. Thus, modulo a change of basis, the set of periods is invariant under the Veech group.

Note that $SL(X, \omega)$ is never cocompact, by the following argument: a translation surface always has a closed geodesic, for topological reasons; up to a rotation, we may assume this geodesic is horizontal. Let l be its length. Applying the Teichmüller geodesic flow for a time t , the length of the closed geodesic becomes $e^{-t}l$, so the Teichmüller geodesic goes to infinity in the moduli space, therefore $\mathrm{SL}_2(\mathbb{R})/SL(X, \omega)$ cannot be compact.

The next best thing to being co-compact is to have finite co-volume, Veech groups of finite co-volume are going to play an important part. Note that for a Fuchsian group Γ to have finite co-volume, its ends at infinity must be finitely many cusps (see Section 4.2 of [20]).

Now let us see some examples.

5.1 The Veech group of the square torus is $\mathrm{SL}_2(\mathbb{Z})$

Figure 12 shows that the matrices

$$T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \text{ and } S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix},$$

which generate $\mathrm{SL}_2(\mathbb{Z})$, lie in the Veech group of the square torus. Conversely, every element of the Veech group preserve the set of periods, so, in the case of the torus, it must preserve $\mathbb{Z}[i]$, hence it must lie in $\mathrm{SL}_2(\mathbb{Z})$.

5.2 Veech groups of the regular polygons

It is obvious that the rotation by π/n is in the Veech group of the regular $2n$ -gon. If you view the double $(2n+1)$ -gon as a star, as in Figure 2, you see that the rotation by $2\pi/(2n+1)$ is in the Veech group of the double $(2n+1)$ -gon.

Figure 13 explains how to find a parabolic element in the Veech group of the regular n -gon: first, the regular n -gon may be cut and pasted into a slanted stair-shape. Then the matrix

$$\begin{pmatrix} 1 & -2 \cot \frac{\pi}{n} \\ 0 & 1 \end{pmatrix}$$

brings the slanted stair-shape to its mirror image with respect to the vertical axis, and the mirror image may be cut and pasted back to a regular n -gon.

Figure 12: Action of T and S on the torus $\mathbb{T}^2$

Figure 13: The octagon is cut and pasted into a slanted stair shape

In fact, Veech proved in [54] that the Veech groups of the aforementioned surfaces are generated by the rotation, and by the parabolic element.

5.3 Veech groups of square-tiled surfaces

The Veech group of a square-tiled surface, if we assume that the squares are copies of the unit square in $\mathbb{R}^2$, is a subgroup of $\mathrm{SL}_2(\mathbb{Z})$, because any element of the Veech group must preserve the set of periods, hence it must preserve $\mathbb{Z}[i]$. In fact it has finite index in $\mathrm{SL}_2(\mathbb{Z})$ (see [14]).

Figure 15 shows why the parabolic element T^2 lies in the Veech group of the surface $St(3)$. Figure 14 shows why the parabolic element T does not lie in the Veech group of the surface $St(3)$, in fact it takes $St(3)$ to the surface on the right in Figure 3. The

Veech group of $St(3)$ is generated by T^2 and S , the rotation of order 4. It has index 3 in $SL_2(\mathbb{Z})$, the right cosets are those of id , T and its transpose.

In [45] an algorithm is given to compute the Veech group of any square-tiled surface. It is implemented in the Sagemath package [9].

Figure 14: Action of T on $St(3)$

6 Teichmüller disks

The orbit of (X, ω) under $\mathrm{GL}_2^+(\mathbb{R})$, or, more properly, its projection to the Teichmüller space $\mathcal{T}_g$, is called the **Teichmüller disk** of (X, ω) . The concept (under the name “complex geodesic”, which is still used sometimes) originates in [49], § 121. Here is why

Figure 15: Action of T^2 on $St(3)$

it is called a disk.

Observe that if ω is a holomorphic differential on a Riemann surface X , and $A \in \mathrm{GL}_2^+(\mathbb{R})$ is a similitude, then $A.\omega$ is still holomorphic with respect to the complex structure of X , so if we are looking at the projection to the Teichmüller space $\mathcal{T}_g$, the action of $\mathrm{GL}_2^+(\mathbb{R})$ factors through the quotient space of $\mathrm{GL}_2^+(\mathbb{R})$ by similitudes. Recall that said quotient space is the hyperbolic disk $\mathbb{H}^2$.

The **Teichmüller curve** of (X, ω) is the quotient of the hyperbolic disk $\mathbb{H}^2$ by $SL(X, \omega)$, which is a Fuchsian group. The fact that $SL(X, \omega)$ is a Fuchsian group tells us that this quotient is a hyperbolic surface, possibly with some conical singularities at symmetrical surfaces (any symmetry of X is an element of $SL(X, \omega)$ which fixes X).

Since the $\mathrm{GL}_2^+(\mathbb{R})$ -action contains the Teichmüller geodesic flow as a subgroup action, Teichmüller curves, which are $\mathrm{GL}_2^+(\mathbb{R})$ -orbits, are invariant under the geodesic flow. Furthermore, Teichmüller curves, as hyperbolic manifolds, are isometrically embedded in $\mathcal{M}_g$ (see [44]).

Figure 14 shows that the two surfaces of Figure 3 lie in the same Teichmüller curve.

6.1 Example: the Teichmüller disk of the torus

The Teichmüller disk of the torus is special because it is also the Teichmüller *space* of the torus.

A torus, as a complex curve, is $\mathbb{C}/\Lambda$, where $\Lambda = \mathbb{Z}u \oplus \mathbb{Z}v$ is a lattice in $\mathbb{R}^2$, u and v being two non-colinear vectors in $\mathbb{R}^2$. Since $\mathbb{Z}u \oplus \mathbb{Z}v = \mathbb{Z}v \oplus \mathbb{Z}u$, we may assume that (u, v) is a positive basis of $\mathbb{R}^2$. Thus the matrix M whose columns are the coordinates of u and v , in that order, lies in $\mathrm{GL}_2^+(\mathbb{R})$. The Abelian differential we consider on $\mathbb{C}/\Lambda$ is $dz = dx + idy$. A matrix $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{GL}_2^+(\mathbb{R})$ acts on dz by pull-back :

$$\forall \begin{pmatrix} x \\ y \end{pmatrix} \in \mathbb{R}^2, A^*dz \begin{pmatrix} x \\ y \end{pmatrix} = dz \left(A. \begin{pmatrix} x \\ y \end{pmatrix} \right) = dz \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix} = ax + by + i(cx + dy)$$

so $A^*dz = a.dx + b.dy + i(c.dx + d.dy)$. The definition of the pull-back means that the complex 1-form A^*dz on $\mathbb{C}/\Lambda$, which is non-holomorphic unless A happens to be $\mathbb{C}$ -linear, becomes the Abelian differential dz on the torus $\mathbb{C}/A.\Lambda$, where $A.\Lambda = \mathbb{Z}A.u \oplus \mathbb{Z}A.v$.

Let (e_1, e_2) be the canonical basis of $\mathbb{R}^2$. The straight segments $\{te_i : 0 \leq t \leq 1\}$, for $i = 1, 2$, become closed curves in $\mathbb{C}/\mathbb{Z}^2 = \mathbb{C}/\mathbb{Z}e_1 \oplus \mathbb{Z}e_2$. We again denote e_1 and e_2 , respectively, the homology classes of those closed curves. For any torus $\mathbb{C}/\Lambda$, with $\Lambda = \mathbb{Z}u \oplus \mathbb{Z}v$, $M \in \mathrm{GL}_2^+(\mathbb{R})$ being the matrix whose columns are the coordinates of u and v , in that order, M may be viewed as a diffeomorphism from $\mathbb{C}/\mathbb{Z}^2$ to $\mathbb{C}/\Lambda$. This allows us to consider (e_1, e_2) as a basis of $H_1(\mathbb{C}/\Lambda, \mathbb{R})$: e_1 (resp. e_2) is the homology class of the closed curve $\{tu, \text{ resp. } tv : 0 \leq t \leq 1\}$ in $\mathbb{C}/\Lambda$.

We compute period coordinates in the basis (e_1, e_2) , and since $\omega = dz$, the period coordinates of $\mathbb{C}/\Lambda$ are $z(u)$ and $z(v)$, where $z(u)$ is the complex number whose affix is u . So, considering periods as vectors in $\mathbb{R}^2$, $\mathrm{GL}_2^+(\mathbb{R})$ acts on periods linearly on the left.

If $\phi : \mathbb{R}^2 \longrightarrow \mathbb{R}^2$ is a similarity (a non-zero $\mathbb{C}$ -linear map), then $\mathbb{C}/\Lambda$ and $\mathbb{C}/\phi(\Lambda)$ are bi-holomorphic. So, if we are interested in the projection of the Teichmüller disk to the moduli space of complex structures, we might as well consider the quotient of $GL_2^+(\mathbb{R})$ by the subgroup H of similarities, acting on the left. Beware that $GL_2^+(\mathbb{R})$ must then act on the quotient from the right.

Say two matrices $M_1, M_2 \in GL_2^+(\mathbb{R})$ are equivalent under H if there exists $P = \begin{pmatrix} a & c \\ -c & a \end{pmatrix} \in H$, such that $PM_1 = M_2$. We use the following representative of an equivalence class:

if $u = \begin{pmatrix} a \\ c \end{pmatrix}, v = \begin{pmatrix} b \\ d \end{pmatrix}, M = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in GL_2^+(\mathbb{R})$, then

$$\frac{1}{a^2 + c^2} \begin{pmatrix} a & c \\ -c & a \end{pmatrix} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \frac{1}{a^2 + c^2} \begin{pmatrix} 1 & ab + cd \\ 0 & ad - bc \end{pmatrix}.$$

Observe that the matrix $\frac{1}{a^2 + c^2} \begin{pmatrix} a & c \\ -c & a \end{pmatrix}$, viewed as a map from $\mathbb{C}$ to $\mathbb{C}$, is $z \mapsto \frac{1}{z(u)}z$. Geometrically, this means that the complex number which represents (the biholomorphism class of) $\mathbb{C}/\Lambda$, where $\Lambda = \mathbb{Z}u \oplus \mathbb{Z}v$, is just $[z(u) : z(v)] \in \mathbb{C}P^1$, so the canonical representative is $\left[1 : \frac{z(v)}{z(u)}\right] \in \mathbb{C}P^1$.

Since (u, v) is a positive basis of $\mathbb{R}^2$, $\frac{z(v)}{z(u)}$ lies in the upper half-plane $\mathbb{H}^2$, so we have a bijection

$$\Psi : H \backslash GL_2^+(\mathbb{R}) \longrightarrow \mathbb{H}^2.$$

Take an element of $H \backslash GL_2^+(\mathbb{R})$, represented by a matrix $\begin{pmatrix} 1 & x \\ 0 & y \end{pmatrix}$, and an element $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ of $GL_2^+(\mathbb{R})$. Then, recalling that $GL_2^+(\mathbb{R})$ acts on the quotient from the right,

$$M := \begin{pmatrix} 1 & x \\ 0 & y \end{pmatrix} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} a + cx & b + dx \\ cy & dy \end{pmatrix}.$$

Set $z = x + iy$, so the affixes of the columns of the matrix M are $a + cz$ and $b + dz$, then the map Ψ takes the equivalence class of the matrix M to the point $\frac{dz+b}{cz+a}$ in $\mathbb{H}^2$. This means that the map Ψ is equivariant, with respect to the right action of $GL_2^+(\mathbb{R})$ on right cosets, and the right action of $GL_2^+(\mathbb{R})$ on $\mathbb{H}^2$ defined by

$$\mathbb{H}^2 \times GL_2^+(\mathbb{R}) \longrightarrow \mathbb{H}^2 \tag{1}$$

$$\left(z, \begin{pmatrix} a & b \\ c & d \end{pmatrix}\right) \longmapsto \frac{dz + b}{cz + a}. \tag{2}$$

6.1.1 Hyperbolic metric vs. Teichmüller metric

So far we have identified the Teichmüller disk of the torus with the hyperbolic plane. We already know that the Veech group of the square torus is $\mathrm{SL}_2(\mathbb{Z})$, so the Teichmüller curve of the torus is the *modular curve* $\mathbb{H}^2/\mathrm{SL}_2(\mathbb{Z})$ (more on which can be found in [1]).

The hyperbolic plane, as a set, is not particularly interesting; what is interesting about it is the hyperbolic metric. We shall now see that Teichmüller's metric, on the Teichmüller space of the torus viewed as the hyperbolic plane, is precisely the hyperbolic metric, up to a multiplicative factor, which we shall ignore.

We shall use the fact that the hyperbolic metric, up to a multiplicative factor, is the only Finsler metric on $\mathbb{H}^2$ invariant under the action of $\mathrm{SL}_2(\mathbb{R})$. This is because $\mathrm{SL}_2(\mathbb{R})$ acts transitively on the tangent bundle of $\mathbb{H}^2$, so once we know the length of one tangent vector at some point, we know the length of every tangent vector, at every point.

By the equivariance of the map Ψ , and since we are not interested in the multiplicative factor, all we have to check is the invariance of the Teichmüller metric under the action of $\mathrm{SL}_2(\mathbb{R})$.

Now recall the Teichmüller norm, on the fiber, over X , of the holomorphic quadratic differential bundle, is just the total area. Since $\mathrm{SL}_2(\mathbb{R})$ preserves area, it preserves the Teichmüller metric. Thus we have proved that the Teichmüller metric is a multiple of the hyperbolic metric on $\mathbb{H}^2$. This was originally observed in § 5 of [49].

6.1.2 Hyperbolic geodesics

Let us investigate a little bit the relationship between hyperbolic geodesics and flat tori. The subgroup

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

of $\mathrm{SL}_2(\mathbb{R})$ is mapped by Ψ to the geodesic

$$\begin{aligned} \mathbb{R} &\longrightarrow \mathbb{H}^2 \\ t &\longmapsto e^{-2t}i. \end{aligned}$$

For any $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{R})$,

$$\mathcal{G}A = \left\{ \begin{pmatrix} e^t a & e^t b \\ e^{-t} c & e^{-t} d \end{pmatrix} : t \in \mathbb{R} \right\}$$

is mapped by Ψ to the geodesic

$$\begin{aligned} \mathbb{R} &\longrightarrow \mathbb{H}^2 \\ t &\longmapsto \frac{e^{2t}ab + e^{-2t}cd + i(ad - bc)}{e^{2t}a^2 + e^{-2t}c^2} \end{aligned}$$

whose endpoints at infinity are $\frac{b}{a}$ and $\frac{d}{c}$.

Here is a geometric interpretation of the endpoints. Consider the complex differential form $\eta := A^*dz = a.dx + b.dy + i(c.dx + d.dy)$. Then, for any $t \in \mathbb{R}$,

$$(g_t A)^* dz = e^t \Re \eta + i e^{-t} \Im \eta,$$

that is, $e^t \Re \eta + i e^{-t} \Im \eta$ is holomorphic on the torus $\mathbb{C}/\Lambda$, where $\Lambda = g_t.A.\mathbb{Z}^2$. So η^2 is the quadratic differential given by Teichmüller's theorem (see 4.1), associated with the geodesic Ψg_t , whose real part is contracted along the geodesic, while its imaginary part is expanded. The endpoints at infinity of the geodesic are the (reciprocals of the opposites of the) slopes of the respective kernels of the real and imaginary parts of η .

6.1.3 Horocycles

The subgroup

$$\mathcal{N} := \left\{ n_t = \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix} : t \in \mathbb{R} \right\}$$

of $\mathrm{SL}_2(\mathbb{R})$ is mapped by Ψ to the horocycle

$$\begin{aligned} \mathbb{R} &\longrightarrow \mathbb{H}^2 \\ t &\longmapsto i + t. \end{aligned}$$

For any $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{R})$,

$$\mathcal{N}A = \left\{ \begin{pmatrix} a+tc & b+td \\ c & d \end{pmatrix} : t \in \mathbb{R} \right\}$$

is mapped by Ψ to the horocycle

$$\begin{aligned} \mathbb{R} &\longrightarrow \mathbb{H}^2 \\ t &\longmapsto \frac{(a+tc)(b+td) + cd + i(ad-bc)}{(a+tc)^2 + c^2} \end{aligned}$$

whose endpoint at infinity is $\frac{d}{c}$, and whose apogee is $\frac{d}{c} + i\frac{ad-bc}{c^2}$.

The pull-back, by $n_t A$, of the real 1-form dy , is $c.dx + d.dy$, because

$$(n_t A)^* dy \begin{pmatrix} x \\ y \end{pmatrix} = dy \left(n_t A \begin{pmatrix} x \\ y \end{pmatrix} \right) = dy \begin{pmatrix} (a+tc)x + (b+td)y \\ cx + dy \end{pmatrix} = cx + dy,$$

in particular $(n_t A)^* dy$ is constant along the horocycle $\Psi(\mathcal{N}A)$. The point at infinity of the horocycle is (minus the reciprocal of) the slope of the kernel of the constant 1-form.

The reason for the annoying recurrence of "minus the reciprocal of..." is that the matrix $\begin{pmatrix} d & b \\ c & a \end{pmatrix}$ is conjugate, by $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$, which acts on the hyperbolic plane by $z \mapsto -z$, to the inverse matrix $A^{-1} = \begin{pmatrix} d & -b \\ -c & a \end{pmatrix}$.

If $\frac{d}{c} \in \mathbb{Q}$, then all the closed geodesics whose velocity vectors are tangent to the kernel of $c.dx + d.dy$ are closed, and of length $c^2 + d^2$. So the torus $n_t A$, for $t \in \mathbb{R}$, may be seen as a cylinder of girth $c^2 + d^2$, with the boundaries identified. When t varies in $\mathbb{R}$, the identification varies but the closed geodesics remain fixed.

6.1.4 Relationship between the asymptotic behaviour of hyperbolic geodesics, and the dynamic behaviour of Euclidean geodesics

The set of slopes of closed geodesics in $\mathbb{C}/\mathbb{Z}^2$ is $\mathbb{Q}$. The Veech group $SL_2(\mathbb{Z})$ acts transitively on $\mathbb{Q}$. Given a hyperbolic geodesic $t \mapsto \gamma(t)$ in the modular surface $\mathbb{H}^2/SL_2(\mathbb{Z})$, the following four points are equivalent:

- $\gamma(t)$ goes to infinity in $\mathbb{H}^2/SL_2(\mathbb{Z})$, when $t \rightarrow \infty$
- the lifts of γ to $\mathbb{H}^2$ have a rational endpoint when $t \rightarrow \infty$
- all orbits of the vertical flow of the Abelian differential associated with γ are closed
- γ is invariant under the parabolic subgroup of $SL_2(\mathbb{Z})$ generated by $\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$.

6.2 Examples of higher genus Teichmüller disks

6.2.1 Three-squared surfaces

The Veech group of the surface $St(3)$ is an index 3 subgroup of $SL_2(\mathbb{Z})$. It has a fundamental domain which is an ideal triangle in $\mathbb{H}^2$, with ideal vertices $-1, 1, \infty$. The parabolic transformation T^2 identifies the two vertical boundaries, and the rotation S identifies the two halves of the semi-circular boundary. The point i , which corresponds to the identity matrix, that is, to the surface $St(3)$ itself, is a singular point: a conical point with angle π . This is because it is invariant by S , which acts as the involution $z \mapsto -\frac{1}{z}$ on the hyperbolic plane. The other two three-squared translation surfaces, which correspond to the matrices T and ST , are both (since they induce the same complex structure) mapped by Ψ , as defined in Subsection 6.1, to the point $1+i$ (which is identified with $1-i$ by T^2).

The Teichmüller curve of $St(3)$ has two cusps, one at ∞ , and the other at ± 1 (which become the same point in the Teichmüller curve). Since $St(3)$ is square-tiled, the set of slopes of closed geodesics is $\mathbb{Q}$. The difference with the torus is that the Veech group does not act transitively on $\mathbb{Q}$, it has two orbits, one for each cusp. The orbit of ∞ consists of all fractions $\frac{p}{q}$, with $p \not\equiv q \pmod{2}$. The orbit of ± 1 consists of all fractions $\frac{p}{q}$, with $p \equiv q \pmod{2}$.

The first cusp is called a two-cylinder cusp, while the second cusp is called a one-cylinder cusp, for the following reason. If a geodesic, in the Teichmüller curve, escapes to infinity in the two-cylinder cusp, then any lift of this geodesic to $\mathbb{H}^2$ corresponds, by Teichmüller's theorem, to a quadratic differential, the trajectories of whose real part are closed, and decompose $St(3)$ into two cylinders of closed geodesics. For instance, the

vertical direction admits two cylinders, one obtained by identifying the green boundaries in Figure 5 is made of vertical closed geodesics of length 2, and the other, obtained by identifying the blue boundaries in Figure 5 is made of vertical closed geodesics of length 1. If a geodesic, in the Teichmüller curve, escapes to infinity in the one-cylinder cusp, then any lift of this geodesic to $\mathbb{H}^2$ corresponds, by Teichmüller's theorem, to a quadratic differential, the trajectories of whose real part are closed, pairwise homotopic, and of equal length. For instance, all geodesics in the 45° direction, except those that hit the singular point, are closed and of length $3\sqrt{2}$.

See [6] for more on this Teichmüller curve. Note that the three-squared Teichmüller curve has genus zero, topologically it is a sphere minus two points, one for each cusp. The genus, and the number of cusps, of Teichmüller curves of square-tiled surfaces may be arbitrarily large. See [17] for the stratum $\mathcal{H}_2$; given a square-tiled surface in any stratum, the genus and number of cusps of its Teichmüller curve may be computed with [9].

6.2.2 Regular polygons

We have seen that the Veech group of the regular polygons is generated by a parabolic element, and a rotation. The Veech group of the double $2n + 1$ -gon has a fundamental domain which looks like Figure 16, left. It has only one cusp, at infinity in Figure 16. This is reflected in the fact that the Veech group acts transitively on the set of directions of closed geodesics. It is an n -cylinder cusp, meaning that in any direction of closed geodesic, the surface decomposes into n cylinders.

The point $(0, 1)$ in Figure 16, left, is the image by Ψ of the identity matrix, that is, the double $2n + 1$ -gon. In the Teichmüller curve it is singular, more precisely it is a conical point with angle $\frac{2\pi}{2n+1}$. This reflects the fact that the double $2n + 1$ -gon is invariant by a rotation of order $2n + 1$.

The corner of coordinates $(\pm \cot \frac{\pi}{2n+1}, y)$, is the image by Ψ , when $n = 2$, of the golden L (see [8]), and when $n > 2$, of stair-shaped polygons (see [26]). Those polygons are invariant by S , which acts as the involution $z \mapsto -\frac{1}{z}$ on the hyperbolic plane, this is why the corner becomes a conical point with angle π in the Teichmüller curve.

The main difference between the Teichmüller curve of the double $2n + 1$ -gon, and Teichmüller curve of the regular $4n$ -gon, is that the latter has two cusps, at ∞ and $(\pm \cot \frac{\pi}{4n}, 0)$ on Figure 16, right. See [47] for more on the Teichmüller curve of the regular octagon, where the fundamental domain is drawn in the hyperbolic disk rather than in the hyperbolic half-plane.

The Teichmüller curve of the octagon, however, is different from the three-square Teichmüller curve, in that both cusps are two-cylinders, meaning that if a geodesic, in the Teichmüller curve, escapes to infinity in any cusp, then any lift of this geodesic to $\mathbb{H}^2$ corresponds, by Teichmüller's theorem, to a quadratic differential, the trajectories of whose real part are closed, and decompose the octagon into two cylinders of closed geodesics.

Then, you may ask, how do we distinguish between the cusps? First, look at the horizontal direction in the octagon (Figure 13). We leave it to the reader to check that

the cylinder made up with the triangles 5 and 6 has height 1 and length $1 + \sqrt{2}$, while the cylinder made up with the triangles 1, 2, 3 and 4 has height $\frac{1}{\sqrt{2}}$ and length $2 + \sqrt{2}$.

The **module** of a cylinder is the ratio of its height to its length. So we see that for the horizontal direction, the ratio of the modules of the cylinders is 2. What is important about this ratio of modules, is that it is invariant by $\mathrm{GL}_2^+(\mathbb{R})$, simply because linear maps, while they may not preserve length, preserve ratios of length of collinear segments. Now we leave it to the reader to check that in the direction of the short red diagonal, one cylinder is made up of the triangles 1 and 2, while the other is made of triangles 3, 4, 5, 6, and the ratio of their modules is not 2. This means that no element of the Veech group takes the horizontal direction to the direction of the short red diagonal, so the Veech group, acting on the set of directions of closed geodesics, has at least two orbits, therefore there are at least two cusps.

7 Veech surfaces

When facing a daunting dynamical system, and looking for closed invariant sets, the first thing to do is to look for *small* closed invariant sets: if your dynamical system is an $\mathbb{R}$ -action (a flow), you are going to look for fixed points, or periodic orbits. In the case of an action by a larger group, you are going to look for closed (topologically speaking) orbits. If the $\mathrm{GL}_2^+(\mathbb{R})$ -orbit of a holomorphic differential (X, ω) is closed in $\mathcal{H}_g$, we say

Figure 16: Fundamental domains for the Veech groups of the double $2n + 1$ -gon, on the left, and the $4n$ -gon, on the right. The number y is such that the circular boundary meets the vertical boundary perpendicularly. The angle at the point $(0, 1)$ is $\frac{2\pi}{2n+1}$ (left), and $\frac{\pi}{2n}$ (right).

(X, ω) is a **Veech surface**.

7.1 The Smillie-Weiss theorem

What Veech was after was Veech groups with finite co-volume. Later it was proved in [48] that having a closed Teichmüller disk is equivalent to having a finite-covolume Veech group.

What follows, while very far from a proof, is meant to convey a bit of the idea.

First, let us assume that for some Abelian differential (X, ω) , $\Gamma = SL(X, \omega)$ has finite co-volume in $\mathbb{H}^2$. Then $SL_2(\mathbb{R})/\Gamma$ consists of a compact part K and finitely many cusps $C_1, \dots, C_n$ (see [20]). Therefore $K.(X, \omega)$ is compact, so it is closed in $\mathcal{H}_g$, and $C_i.(X, \omega)$ is closed in $\mathcal{H}_g$ unless there exists a sequence A_n of matrices in C_i , which converges to infinity in $SL_2(\mathbb{R})$, and such that $A_n.(X, \omega)$ does not go to infinity in $\mathcal{H}_g$. But recall that every cusp of $\mathbb{H}^2/\Gamma$ is stabilized by a parabolic element of the Veech group, which in turns yields a decomposition of X into cylinders of closed geodesics. When A_n goes to infinity, since $A_n \in C_i$, the length of the closed geodesics must go to zero, which means that $A_n.(X, \omega)$ leaves every compact set of $\mathcal{H}_g$.

Conversely, let us assume that for some Abelian differential (X, ω) , the orbit $GL_2^+(\mathbb{R}).(X, \omega)$ is closed in $\mathcal{H}_g$. First, let us observe that since $GL_2^+(\mathbb{R}).(X, \omega)$ is closed in $\mathcal{H}_g$, the embedding of $SL_2(\mathbb{R})/\Gamma$ into $\mathcal{H}_g$ is proper (the inverse image of any compact set is compact). This can be seen from the Magic Wand Theorem of [10]: since $GL_2^+(\mathbb{R}).(X, \omega)$ is closed, it must be an embedded submanifold, in particular, it is properly embedded. This means that the ends at infinity of $GL_2^+(\mathbb{R}).(X, \omega)$ in $\mathcal{H}_g$ are exactly the images in $\mathcal{H}_g$ of the ends at infinity of $SL_2(\mathbb{R})/\Gamma$.

Now let us ask ourselves, what kind of ends at infinity a closed Teichmüller disk may have?

The first two examples that come to mind are a hyperbolic funnel (an end at infinity of a quotient of the hyperbolic disk by a hyperbolic element), and a cusp (the thin end a quotient of the hyperbolic disk by a parabolic element). One difference between the two is that the former has infinite volume, while the latter has not. Another difference is that geodesics that go to infinity in a cusp are all asymptotic to each other, while there exists a point x in $\mathbb{H}^2$, and an open interval I of directions such that any geodesic going from x with direction in I goes to infinity in the funnel. Now we apply Theorem 4.2: among those directions, there must be a uniquely ergodic one, let us call it θ (meaning that the vertical flow of $e^{i\theta}\omega$ is uniquely ergodic). But then, by Masur's Criterion, geodesics corresponding to uniquely ergodic directions do not go to infinity. This contradiction shows that the ends at infinity of a Teichmüller disk are cusps.

Of course, we also have to prove that there are but finitely many cusps. Note that by [19], a Veech group may be infinitely generated, but in that case it is not the Veech group of a Veech surface. Here we draw the Magic Wand again: orbits closures are algebraic subvarieties of $\mathcal{H}_g$, so orbit closures of dimension one cannot have infinitely many cusps. Now, since cusps have finite volume, a closed Teichmüller disk is the union of a compact part, and finitely many cusps, so it has finite volume.

7.2 The Veech alternative

The Veech alternative (see [54]) is another great example of sowing in the parameter space and reaping in the phase space. It says that every direction on a Veech surface is either uniquely ergodic, or completely periodic. This means that if (X, ω) is a Veech surface, for every $\theta \in \mathbb{R}$, the vertical flow of $(X, e^{i\theta}\omega)$ is either uniquely ergodic, or every orbit which does not hit a singularity is periodic. Beware this does not mean the flow itself is periodic, for the periods of the orbits may not be commensurable.

Let us give some flavour of the proof. Assume the vertical flow on a Veech surface (X, ω) is not uniquely ergodic, so the orbit of ω under the Teichmüller geodesic flow goes to infinity. Now, since (X, ω) is a Veech surface, the only ends at infinity of its Teichmüller disk are cusps; and a geodesic which goes to infinity in a cusp must be invariant by the parabolic element which generates the cusp.

7.2.1 Examples of Veech surfaces

The first examples found, by Veech ([54]), are the surfaces obtained by identifying opposite sides of a regular $2n$ -gon.

Once people got interested in Veech surfaces, it was immediately realized (see [14]) that square-tiled surfaces are Veech surfaces, just because their Veech groups are (conjugate to) subgroups of $\mathrm{SL}_2(\mathbb{Z})$, which has finite co-volume.

In general, finding the Veech surfaces in a given stratum is a hard problem. However, in genus two, it was solved by McMullen in [31–36], see Subsection 8.2.

8 Classification of orbits in genus two

The Magic Wand, magic as it is, does not tell us everything about the orbit closures. For instance, in most strata, all we know is that

- by ergodicity, almost every orbit is dense
- the union of all orbits of square-tiled surfaces (each of which is closed) is dense.

When are there orbit closures of intermediate dimension? Are there closed orbits (i.e. Veech surfaces) besides those of square-tiled surfaces? Speaking of square-tiled surfaces, how do they distribute in Teichmüller disks? For instance, could it be that each square-tiled surface is the only square-tiled surface in its Teichmüller disk?

The answers are known in genus two: we have a complete list of non-square-tiled Veech surfaces ([4], [32]), a complete list of orbit closures of intermediate dimension ([33]), and, in the stratum $\mathcal{H}(2)$, a complete list of orbits of *primitive* (more on this later) square-tiled surfaces ([17]).

8.1 Square-tiled surfaces

Observe that a square may be subdivided into little rectangles, and then we may use some matrix in $\mathrm{GL}_2(\mathbb{R})$ to square the rectangles, so a square-tiled surface may be square-tiled

in many ways.

Following [17], we say a square-tiled surface is primitive if it is not obtained from another square-tiled surface by subdividing squares into rectangles and then applying a matrix to square the rectangles.

Then [17] says that for $n > 4$, there are exactly two Teichmüller disks of primitive n -square surfaces $\mathcal{H}(2)$. For $n = 3, 4$, there is only one Teichmüller disk of primitive n -square surfaces. Each Teichmüller disk contains finitely many primitive n -square surfaces, the exact number is given by the index of the Veech group in $\mathrm{SL}_2(\mathbb{Z})$. In genus two it turns out that this index is never 1, however, there exist square-tiled surfaces of higher genus whose Veech group is $\mathrm{SL}_2(\mathbb{Z})$ (see “Eierlegende Wollmilchsau” and “Ornythorynque” in [13]). Even in genus two the situation is not fully understood, for instance, we don’t know how many Teichmüller disks of primitive n -square surfaces there are in $\mathcal{H}(1, 1)$.

8.2 A homological detour

8.2.1 The tautological subspace

Since the complex-valued 1-form ω is holomorphic, the real-valued 1-forms $\Re(\omega)$ and $\Im(\omega)$ are closed, by the Cauchy-Riemann relations. We denote S the 2-dimensional subspace of $H^1(X, \mathbb{R})$ generated by the cohomology classes of $\Re(\omega)$ and $\Im(\omega)$.

Recall that $H^1(X, \mathbb{R})$ has a symplectic structure: let $[X] \in H^2(X, \mathbb{R})$ be the fundamental class, then

$$\forall a, b \in H^1(X, \mathbb{R}), \exists c \in \mathbb{R}, a \wedge b = c[X].$$

The symplectic 2-form is then defined as $(a, b) \mapsto c$. It is Poincaré dual to the intersection form (the bilinear form Int on $H_1(X, \mathbb{Z})$ such that $\mathrm{Int}(h, k)$ is the total intersection, counted with signs, of any two representatives of h and k in transverse position) on $H_1(X, \mathbb{R})$, that is, if $P : H_1(X, \mathbb{R}) \rightarrow H^1(X, \mathbb{R})$ is the Poincaré map, then

$$\forall h, k \in H_1(X, \mathbb{R}), P(h) \wedge P(k) = \mathrm{Int}(h, k)[X].$$

The subspace S is symplectic, meaning that the 2-form $\cdot \wedge \cdot$, restricted to $S \times S$, is non-degenerate. Indeed, recall that in an appropriate chart, $\omega = dz = dx + idy$, so

$$\Re(\omega) \wedge \Im(\omega) = dx \wedge dy.$$

The symplectic orthogonal $S^\perp$ of S is the set of cohomology classes c such that $c \wedge \Re(\omega) = c \wedge \Im(\omega) = 0$. We leave it as an exercise for the reader that in a symplectic vector space, the symplectic orthogonal of a symplectic subspace is also symplectic.

Now let us take a look at how the Veech group Γ acts on S . Let γ be an element of Γ , and let $\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{R})$ be the matrix of the derivative of γ , seen as a diffeomorphism of X minus its singular set. Let (x, y) be a point in X , in local coordinates, and let

(u, v) be a tangent vector to X at (x, y) . Then

$$\begin{aligned}\gamma^*(dx)_{(x,y)} \cdot (u, v) &= dx(d\gamma_{(x,y)} \cdot (u, v)) = dx(au + bv, cu + dv) = au + bv \\ \gamma^*(dy)_{(x,y)} \cdot (u, v) &= dy(d\gamma_{(x,y)} \cdot (u, v)) = dy(au + bv, cu + dv) = cu + dv\end{aligned}$$

so the action of Γ by pull-back on S , endowed with the basis $[dx], [dy]$, is given by the same matrix as the linear action of γ on $\mathbb{R}^2$. For this reason S is sometimes called the **tautological subspace** of $H^1(X, \mathbb{R})$.

8.2.2 The trace field

The **trace field** $K(\Gamma)$ of a subgroup Γ of $\mathrm{SL}_2(\mathbb{R})$ is the subfield of $\mathbb{R}$ generated by the traces of the elements of Γ . We are going to see that

Lemma 8.1. *When Γ is the Veech group of some translation surface of genus two, $K(\Gamma) = \mathbb{Q}[\sqrt{d}]$ for some $d \in \mathbb{N}^*$.*

Of course, for most surfaces, the Veech group is trivial, so the trace field is $\mathbb{Q}$. The trace field is most interesting when the Veech group is large, especially for Veech surfaces.

Figure 17: On the left, the surface $L(a, 1)$. The sides e_1 and f_1 have length 1, while the sides e_2 and f_2 have length $a - 1$. The surface decomposes into two cylinders of horizontal closed geodesics, the lower one of width a and height 1, and the upper one of width 1 and height $a - 1$. On the right, the top rectangle in the surface $L(a, 1)$ has been translated to the left. Sides are still identified by vertical or horizontal translations. The resulting surface lies in the stratum $\mathcal{H}(1, 1)$, like $St(4)$, but it has the same absolute periods as $L(a, 1)$. The closed curves with respect to which the periods are computed are drawn in red.

It is proved in [14] that for a Veech surface of any genus, having $\mathbb{Q}$ as its trace field is equivalent to being in the Teichmüller disk of a square-tiled surface. In the genus two case, this means that if d is a square, then X is parallelogram-tiled (that is, X lies in the Teichmüller disk of a square-tiled surface).

Take $\gamma \in \Gamma$, and consider the induced automorphism γ^* of $H^1(X, \mathbb{R})$. Since γ^* is induced by a homeomorphism of X , it preserves the integer lattice $H^1(X, \mathbb{Z})$ of $H^1(X, \mathbb{R})$, and the same goes for $(\gamma^*)^{-1}$. Thus the endomorphism $\gamma^* + (\gamma^*)^{-1}$ of $H^1(X, \mathbb{R})$ may be represented by an integer matrix, in particular its eigenvalues are algebraic numbers.

But recall that γ^* , restricted to S , is just γ , so $\gamma^* + (\gamma^*)^{-1}$, restricted to S , is just $\gamma + \gamma^{-1} = \text{Tr}(\gamma) \cdot \text{Id}$. Therefore the trace of γ is a double eigenvalue of $\gamma^* + (\gamma^*)^{-1}$, so the square of the minimal polynomial χ of $\text{Tr}(\gamma)$ divides the characteristic polynomial of $\gamma^* + (\gamma^*)^{-1}$, hence χ has degree at most two.

We have just proved that for every $\gamma \in \Gamma$, there exists $d \in \mathbb{N}^*$ such that $\text{Tr}(\gamma) \in \mathbb{Q}[\sqrt{d}]$. The same argument applies to any linear combination of traces of elements of Γ . We claim that d may be chosen independantly of γ . If it were not the case, say, $\text{Tr}(\gamma) \in \mathbb{Q}[\sqrt{d}]$ and $\text{Tr}(\gamma') \in \mathbb{Q}[\sqrt{d'}]$ for some $\gamma' \neq \gamma$ and $d' \neq d$, then we could find a non-quadratic element in $K(\Gamma)$ (think, for instance, that $\sqrt{2} + \sqrt{3}$ is not quadratic). Therefore, $K(\Gamma) = \mathbb{Q}[\sqrt{d}]$ for some $d \in \mathbb{N}^*$.

Example. Consider the surface $L(a, 1)$ on the left of Figure 17, where $a = \frac{1+\sqrt{d}}{2}$. For $d = 5$, this surface lies in the Teichmüller curve of the double pentagon, it is usually called the Golden L (see [7, 8]). For $d = 2$, it can be shown to lie in the Teichmüller curve of the regular octagon (see 47).

The matrix $A := \begin{pmatrix} 1 & 4a \\ 0 & 1 \end{pmatrix}$ acts on the lower cylinder, whose height is 1 and whose width is a , as the 4-th power of the horizontal Dehn twist, because

$$A \cdot \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix} + 4 \begin{pmatrix} a \\ 0 \end{pmatrix},$$

and it acts on the upper cylinder, whose height is $a - 1$ and whose width is 1, as the $(d - 1)$ -th power of the horizontal Dehn twist, because

$$A \cdot \begin{pmatrix} 0 \\ a - 1 \end{pmatrix} = \begin{pmatrix} 4a(a - 1) \\ a - 1 \end{pmatrix} = \begin{pmatrix} 0 \\ a - 1 \end{pmatrix} + (d - 1) \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

Thus the matrix A lies in the Veech group of $L(a, 1)$. The same arguments apply to the vertical cylinders, so the matrix $\begin{pmatrix} 1 & 0 \\ 4a & 1 \end{pmatrix}$ also lies in the Veech group, and so does their product $\begin{pmatrix} 1 + 16a^2 & 4a \\ 4a & 1 \end{pmatrix}$, whose trace is $2 + 16a^2 = 2 + 4d + 8\sqrt{d}$, hence the trace field of $L(a, 1)$ is $\mathbb{Q}[\sqrt{d}]$.

8.2.3 The Jacobian torus

For any manifold X , the **Jacobian torus** $Jac(X)$ of X is the quotient of $H^1(X, \mathbb{R})$ by the integer lattice $H^1(X, \mathbb{Z})$. When X is a Riemann surface, $H^1(X, \mathbb{R})$ identifies with the space of Abelian differentials, by sending an Abelian differential to the cohomology class of its real part, and $H^1(X, \mathbb{Z})$ acts by translations, which are holomorphic, so $Jac(X)$ is actually a complex manifold.

The Jacobian torus also has a symplectic structure, given by the wedge pairing, and the symplectic structure is compatible with the complex structure, in the following way: any complex subspace in $H^1(X, \mathbb{R})$ is also a symplectic subspace. If E is any complex subspace of $H^1(X, \mathbb{R})$, then the symplectic orthogonal $E^\perp$ is also a complex subspace. For instance, the canonical subspace S and its symplectic orthogonal are both complex subspaces of $H^1(X, \mathbb{R})$.

An **endomorphism** of $Jac(X)$ is a $\mathbb{C}$ -linear endomorphism of $H^1(X, \mathbb{R})$, which preserves the integer lattice, so it quotients to a self-map of $Jac(X)$, and is self-adjoint with respect to the symplectic form. Such a thing actually exists, here is an example.

Assume X is a translation surface of genus two, and γ is an element of the Veech group of X , such that its trace is irrational (hence quadratic since the genus is two). Think of the example we have just seen on the surface $L(a, 1)$.

Then both γ^* and $(\gamma^*)^{-1}$ are $\mathbb{R}$ -linear endomorphisms of $H^1(X, \mathbb{R})$ which preserve the integer lattice, so the same goes for $\gamma^* + (\gamma^*)^{-1}$. Recall that $T := \gamma^* + (\gamma^*)^{-1}$, restricted to S , is just $Tr(\gamma).Id$, so $t := Tr(\gamma)$ is a double eigenvalue of T . Thus the Galois conjugate $\bar{t}$ of t is also a double eigenvalue of T . But both γ^* and $(\gamma^*)^{-1}$ are symplectic, so, since they preserve S , they must preserve the symplectic orthogonal $S^\perp$. Therefore, since the dimension of $S^\perp$ is two (recall that the genus of X is two), we get that T , restricted to $S^\perp$, is $\bar{t}.Id$. This proves that T is $\mathbb{C}$ -linear, since it may be expressed by the matrix $\begin{pmatrix} t & 0 \\ 0 & \bar{t} \end{pmatrix}$ in some complex basis.

This is the part that breaks down in higher genus, and the reason why McMullen's results in genus two have been dubbed miraculous (see [60]). Let us see exactly how it breaks down. Assume X is a translation surface of genus $g > 2$, with a trace field of degree g (the maximal possible degree). Let γ be an element of the Veech group of X , whose trace is an algebraic integer of degree g . Then the $g - 1$ Galois conjugates of $Tr(\gamma)$ are double eigenvalues of $T = \gamma^* + (\gamma^*)^{-1}$, so $S^\perp$ decomposes as a direct sum of 2-dimensional real subspaces, on each of which T acts by homothety. The trouble is, those 2-dimensional real subspaces have no reason to be complex subspaces, even though $S^\perp$ is a complex subspace, so we cannot conclude that T is $\mathbb{C}$ -linear. In some cases (see [24]) T happens to be $\mathbb{C}$ -linear, so McMullen's methods may be applied.

To prove that T is an endomorphism of $Jac(X)$, we still have to prove that T is self-adjoint with respect to the symplectic form. Since γ is a homeomorphism of X , γ^* preserves the symplectic form, so the adjoint of γ^* is $(\gamma^*)^{-1}$. Similarly the adjoint of $(\gamma^*)^{-1}$ is γ^* , so $T = \gamma^* + (\gamma^*)^{-1}$ is self-adjoint.

8.2.4 Real multiplication

Let X be a translation surface of genus two, with trace field $\mathbb{Q}[\sqrt{d}]$ for some non-square $d \in \mathbb{N}$. Let γ be an element of the Veech group of X , with irrational trace, and let $T = \gamma^* + (\gamma^*)^{-1}$.

The endomorphisms of $Jac(X)$ form a ring, and the subring generated by Id and T is isomorphic to a finite index subring of the ring of integers of $\mathbb{Q}[\sqrt{d}]$, the isomorphism being the trace of the restriction to S . Of course there is another isomorphism, which is the trace of the restriction to $S^\perp$.

For any translation surface, when the ring of endomorphisms of $Jac(X)$ has a subring isomorphic to (some finite index subring of) the integer ring of some number field K , we say $Jac(X)$ has **real multiplication** by K . So, what we have proved so far is

Lemma 8.2. *The Jacobian of any non-square-tiled Veech surface of genus two has real multiplication by some $\mathbb{Q}[\sqrt{d}]$, with $d \in \mathbb{N}$ non-square.*

For instance, the double pentagon has real multiplication by $\mathbb{Q}[\sqrt{5}]$, while the octagon has real multiplication by $\mathbb{Q}[\sqrt{2}]$. More generally, for any non-square $d \in \mathbb{N}$, the surface $L(a, 1)$, with $a = \frac{1+\sqrt{d}}{2}$, has real multiplication by $\mathbb{Q}[\sqrt{d}]$.

Now, what is the point of all this, you may ask? The point is that there is a way to know which surfaces have Jacobians with real multiplication, and the stroke of genius was to think of looking at the problem this way.

Let us denote W_d the projection to the moduli space $\mathcal{M}_2$ of the set of translation surfaces of genus two whose Jacobian has real multiplication by $\mathbb{Q}[\sqrt{d}]$.

Theorem 8.3. *For $d \in \mathbb{N}$ non-square, W_d is a two-dimensional algebraic subvariety of $\mathcal{M}_2$.*

Proof. We will focus on the dimension, and refer the reader to [31] for the algebraicity statement. Let us forget about Jacobians for a moment, and consider the set of all complex tori of dimension two (also known, poetically, as principally polarized Abelian varieties). This set may be viewed as the set of all complex structures on $\mathbb{R}^4/\mathbb{Z}^4$, which are compatible with the symplectic structure on $\mathbb{Z}^4$, that is, denoting $(\cdot, \cdot)$ the canonical scalar product and $\cdot \wedge \cdot$ the symplectic structure, for any a, b , $a \wedge b = (a, b)$.

Among all tori, let us consider those whose endomorphism ring contains a copy of (a finite index subring of) the integer ring of $\mathbb{Q}[\sqrt{d}]$. Let ϕ be a generator of said subring, which then is the $\mathbb{Z}$ -module of rank two $\mathbb{Z}[\phi] = \{a + b\phi : a, b \in \mathbb{Z}\}$. The embedding of $\mathbb{Z}[\phi]$ into the ring of endomorphisms of $\mathbb{Z}^4$ makes $\mathbb{Z}^4$ a $\mathbb{Z}[\phi]$ -module, necessarily of rank two. Tensorizing by $\mathbb{Q}$, we may view $\mathbb{Q}^4$ as a two-dimensional vector space over $\mathbb{Q}[\sqrt{d}] = \mathbb{Q}[\phi]$.

Let $\bar{\phi}$ be the Galois conjugate of ϕ , so $P_\phi = (X - \phi)(X - \bar{\phi})$ is the minimal polynomial of ϕ , and let T_ϕ be the endomorphism of $\mathbb{Q}^4$ associated with ϕ . Then $P_\phi(T_\phi) = 0$. Since T preserves $\mathbb{Z}^4$, $T_\phi \neq \phi Id$, so T_ϕ has both ϕ and $\bar{\phi}$ as eigenvalues, which entails that, as

an endomorphism of a two-dimensional $\mathbb{Q}[\phi]$ -vector space, T_ϕ is diagonalizable. Both eigenspaces have dimension one over $\mathbb{Q}[\phi]$, which is dimension two over $\mathbb{Q}$. Since T is self-adjoint, the eigenspaces are mutually orthogonal, and they are symplectic, for if they were not, since their dimension is two, then the restriction of the symplectic form to either of them would vanish, and since they are orthogonal, the symplectic form on $\mathbb{Z}^4$ would be zero, a contradiction.

The complex structure on $\mathbb{R}^4/\mathbb{Z}^4$ is then completely determined by the requirement that the eigenspaces be complex lines, and by the choice of a complex structure for each eigenspace, compatible with the symplectic structure. Each choice is determined by one complex parameter (recall our discussion of the Teichmüller space, which is also the Teichmüller disk, of the torus), so the complex structure on $\mathbb{R}^4/\mathbb{Z}^4$ is determined by two complex parameters.

Next, observe that almost all complex tori are Jacobians: essentially, those who are not Jacobians of regular surfaces, are Jacobians of degenerate surfaces, for instance a bouquet of two 1-dimensional tori. So the set of Jacobians of regular surfaces has dimension two over $\mathbb{C}$. □

McMullen's result actually says much more: the image, in the Teichmüller space $\mathcal{T}_2$, of the set of all translation surfaces of genus two whose Jacobian has real multiplication by $\mathbb{Q}[\sqrt{d}]$, for $d \in \mathbb{N}$ non-square, is an holomorphically embedded copy of $\mathbb{H}^2 \times \mathbb{H}^2$. Its image W_d in the moduli space $\mathcal{M}_2$ is called a **Hilbert modular surface**.

Now we'd like to understand how Hilbert modular surfaces behave with respect to the $\mathrm{SL}_2(\mathbb{R})$ -action, and the strata.

Lemma 8.4. *For any non-square $d \in \mathbb{N}$, W_d is $\mathrm{SL}_2(\mathbb{R})$ -invariant.*

Proof. Let (X, ω) be a translation surface of genus two whose Jacobian has real multiplication by $\mathbb{Q}[\sqrt{d}]$, for $d \in \mathbb{N}$ non-square, let T be an endomorphism of the Jacobian of X , and let γ be an element of $\mathrm{SL}_2(\mathbb{R})$. Then $(\gamma^*)^{-1}$ conjugates T with an endomorphism of the Jacobian of $\gamma.(X, \omega)$, so $\gamma.(X, \omega)$ has real multiplication by $\mathbb{Q}[\sqrt{d}]$. □

Recall that we denote by P the canonical projection $\mathcal{H}_2 \rightarrow \mathcal{M}_2$.

Lemma 8.5. *The intersection $P(\mathcal{H}(2)) \cap W_d$ is not empty.*

Proof. Consider the surface $L(a, 1)$, with $a = \frac{1+\sqrt{d}}{2}$. □

Lemma 8.6. *The intersection $P(\mathcal{H}(1, 1)) \cap W_d$ is not empty.*

Proof. Consider the surface on the right of Figure 17, obtained from $L(a, 1)$ by shifting the upper cylinder to the left. This surface has the same absolute periods as $L(a, 1)$, hence it has the same Jacobian, since Jacobians ignore relative periods. Hence it lies in W_d . □

By Lemmata 8.4 and 8.5, $P(\mathcal{H}(2)) \cap W_d$ is a reunion of $\mathrm{SL}_2(\mathbb{R})$ -orbits, so its dimension is at least one. If its dimension were two, since two is the dimension of W_d , and the latter is connected as a quotient of $\mathbb{H}^2 \times \mathbb{H}^2$, then $\mathcal{H}(2) \cap W_d$ would be the whole of W_d . But this is impossible by Lemma 8.6. So the dimension of $P(\mathcal{H}(2)) \cap W_d$ is one, and since it is an algebraic subvariety of W_d , it must be a finite union of closed orbits. There, we have found Veech surfaces in $\mathcal{H}(2)$, with trace field $\mathbb{Q}[\sqrt{d}]$ for any non-square $d \in \mathbb{N}$, and proved that there are only finitely many of them, for each d . We have also proved that the surface $L(a, 1)$ is a Veech surface for $a = \frac{1+\sqrt{d}}{2}$, just by computing its trace field.

Again, McMullen's results actually say much more: the intersection $P(\mathcal{H}(2)) \cap W_d$ contains exactly one orbit, except when $d \equiv 1 \pmod{8}$, $d \neq 9$, in which case there are exactly two (see [32]). The intersection $P(\mathcal{H}(1, 1)) \cap W_d$ contains no closed orbit, except when $d = 5$, in which case it contains exactly one, the Teichmüller curve of the regular decagon (see [34, 35]). An orbit closure is either the orbit itself, in which case it is a Teichmüller curve, of which we have a complete list, or W_d for some d , or a whole stratum (see [33]). Note that in genus two, both strata project surjectively to the moduli space $\mathcal{M}_2$ (see [12], III.7.5, Corollary 1), so an orbit which is dense in its stratum projects to a dense subset of $\mathcal{M}_2$.

9 What is known in higher genus?

First, a bit of vocabulary: a Veech surface is said to be **geometrically primitive** if it is not a ramified cover of a Veech surface of lower genus. For instance, square-tiled surfaces, other than tori, are not geometrically primitive since they cover the torus. A Veech surface is said to be **algebraically primitive** if its trace field has maximal possible degree, which is the genus of the surface. Veech's family (the regular polygons) and McMullen's family in genus 2 are both algebraically and geometrically primitive, as well as Ward's (see [57]).

The general philosophy of Veech surface hunters seems to be that algebraically and geometrically primitive Veech surfaces are scarce, and you should expect at most finitely many of them in a given stratum of genus > 2 . In [30] this very statement is proved, for the minimal stratum of each genus > 2 .

In [3] a family of Veech surfaces is found, which generalizes Veech's family. In [36], [24], the authors find a way around the fact that McMullen's techniques do not generalize in higher genus, and discover a family of Veech surfaces in genera 3 and 4, which generalizes McMullen's family, in the sense that their trace fields are quadratic (so they are not algebraically primitive). In [37], [11], another family of Veech surfaces in genus 4 is found, again with quadratic trace fields.

References

- [1] P. Arnoux, Le codage du flot géodésique sur la surface modulaire. Enseign. Math. (2) 40 (1994), no. 1-2, 29-48.
- [2] D. Aulicino, D. Nguyen, A. Wright, Classification of higher rank orbit closures in $\mathcal{H}^{odd}(4)$. J. Eur. Math. Soc. (JEMS) 18 (2016), no. 8, 1855-1872.
- [3] I. Bouw and M. Möller, Teichmüller curves, triangle groups, and Lyapunov exponents, Ann. of Math. 172 (2010), 139-185.
- [4] K. Calta, Veech surfaces and complete periodicity in genus two. J. Amer. Math. Soc. 17 (2004), no. 4, 871-908.
- [5] S. Cheboui, Intersection algébrique sur les surfaces à petits carreaux, Ph.D. Thesis, Université de Montpellier, 2021
- [6] S. Cheboui, A. Kessi, D. Massart, Algebraic intersection in a family of Teichmüller disks, preprint, <https://arxiv.org/abs/2007.1084>
- [7] D. Davis, D. Fuchs, S. Tabachnikov, Periodic trajectories in the regular pentagon. Mosc. Math. J. 11 (2011), no. 3, 439-461, 629.
- [8] D. Davis, S. Lelièvre, Periodic paths on the pentagon, double pentagon and golden L. Preprint arXiv:1810.11310
- [9] V. Delecroix et al., Surface Dynamics - SageMath package, Version 0.4.1, <https://doi.org/10.5281/zenodo.3237923>, 2019
- [10] A. Eskin, M. Mirzakhani, Invariant and stationary measures for the $SL(2, \mathbb{R})$ action on moduli space. Publ. Math. Inst. Hautes Études Sci. 127 (2018), 95-324
- [11] A. Eskin, C. McMullen, R. Mukamel, A. Wright, Billiards, quadrilaterals and moduli spaces. J. Amer. Math. Soc. 33 (2020), no. 4, 1039-1086.
- [12] H. Farkas, I. Kra, Riemann surfaces. Second edition. Graduate Texts in Mathematics, 71. Springer-Verlag, New York, 1992
- [13] G. Forni, C. Matheus, Introduction to Teichmüller theory and its applications to dynamics of interval exchange transformations, flows on surfaces and billiards. J. Mod. Dyn. 8 (2014), no. 3-4, 271-436.
- [14] E. Gutkin, C. Judge, Affine mappings of translation surfaces: geometry and arithmetic. Duke Math. J. 103 (2000), no. 2, 191-213.
- [15] J. Hubbard, H. Masur, Quadratic differentials and foliations. Acta Math. 142 (1979), no. 3-4, 221-274.

- [16] J. Hubbard, Teichmüller theory and applications to geometry, topology, and dynamics. Vol. 1. Teichmüller theory. Matrix Editions, Ithaca, NY, 2006.
- [17] P. Hubert, S. Lelièvre, Prime arithmetic Teichmüller discs in $\mathcal{H}(2)$. Israel J. Math. 151 (2006), 281–321.
- [18] P. Hubert, T. Schmidt, An introduction to Veech surfaces, Handbook of dynamical systems. Vol. 1B, 501–526, Elsevier B. V., Amsterdam, 2006. (short version of the lecture notes from the summer school at Luminy, 2003, <https://www.math.uchicago.edu/~masur/hs.pdf>)
- [19] P. Hubert, T. Schmidt, Infinitely generated Veech groups. Duke Math. J. 123 (2004), no. 1, 49–69.
- [20] S. Katok, Fuchsian Groups, University of Chicago Press (1992)
- [21] S. Kerckhoff, H. Masur, J. Smillie, Ergodicity of billiard flows and quadratic differentials. Ann. of Math. (2) 124 (1986), no. 2, 293–311.
- [22] M. Kontsevich, A. Zorich, Connected components of the moduli spaces of Abelian differentials with prescribed singularities. Invent. Math. 153 (2003), no. 3, 631–678.
- [23] E. Lanneau, Promenade dynamique et combinatoire dans les espaces de Teichmüller, mémoire d’habilitation, <https://www-fourier.ujf-grenoble.fr/~lanneau/articles/hdr-lanneau.pdf>
- [24] E. Lanneau, D. Nguyen, Teichmüller curves generated by Weierstrass Prym eigenforms in genus 3 and genus 4. J. Topol. 7 (2014), no. 2, 475–522.
- [25] E. Lanneau, D. Nguyen, A. Wright, Finiteness of Teichmüller curves in non-arithmetic rank 1 orbit closures. Amer. J. Math. 139 (2017), no. 6, 1449–1463.
- [26] E. Lanneau, D. Massart, Algebraic intersection in regular polygons, in progress
- [27] H. Masur, Interval exchange transformations and measured foliations. Ann. of Math. (2) 115 (1982), no. 1, 169–200.
- [28] H. Masur, Hausdorff dimension of the set of nonergodic foliations of a quadratic differential. Duke Math. J. 66 (1992), no. 3, 387–442.
- [29] H. Masur, S. Tabachnikov, Rational billiards and flat structures. Handbook of dynamical systems, Vol. 1A, 1015–1089, North-Holland, Amsterdam, 2002.
- [30] C. Matheus, A. Wright, Hodge-Teichmüller planes and finiteness results for Teichmüller curves. Duke Math. J. 164 (2015), no. 6, 1041–1077.
- [31] C. McMullen, Billiards and Teichmüller curves on Hilbert modular surfaces. J. Amer. Math. Soc. 16 (2003), no. 4, 857–885.

- [32] C. McMullen, Teichmüller curves in genus two: Discriminant and spin. *Math. Ann.* 333 (2005), pp. 87–130.
- [33] C. McMullen, Dynamics of $SL_2(\mathbb{R})$ over moduli space in genus two. *Ann. of Math.* (2) 165 (2007), no. 2, 397–456.
- [34] C. McMullen, Teichmüller curves in genus two: the decagon and beyond. *J. Reine Angew. Math.* 582 (2005), 173–199.
- [35] C. McMullen, Teichmüller curves in genus two: torsion divisors and ratios of sines. *Invent. Math.* 165 (2006), no. 3, 651–672.
- [36] Prym varieties and Teichmüller curves. *Duke Math. J.* 133 (2006), no. 3, 569–590.
- [37] C. McMullen, R. Mukamel, A. Wright, Cubic curves and totally geodesic subvarieties of moduli space. *Ann. of Math.* (2) 185 (2017), no. 3, 957–990.
- [38] M. Möller, Affine groups of flat surfaces, in *Handbook of Teichmüller Theory. Volume II*, European Math. Soc. (EMS), Zürich, 2009, 369–387.
- [39] T. Monteil, Introduction to the theorem of Kerckhoff, Masur and Smillie Oberwolfach Report, No. 17 (2010), 989–994, <https://lipn.univ-paris13.fr/~monteil/papiers/fichiers/kms.pdf>
- [40] R. Mukamel, Fundamental domains and generators for lattice Veech groups. *Comment. Math. Helv.* 92 (2017), no. 1, 57–83.
- [41] D. Mumford, A remark on Mahler’s compactness theorem. *Proceedings of the American Mathematical Society* 28, (1971), no. 1, 289–294.
- [42] B. Riemann, Theorie der Abel’schen Functionen. *J. Reine Angew. Math.* 54 (1857), 115–155.
- [43] B. Riemann, *Gesammelte mathematische Werke, wissenschaftlicher Nachlass und Nachträge*. Springer-Verlag, Berlin; BSB B. G. Teubner Verlagsgesellschaft, Leipzig, 1990.
- [44] H. Royden, Invariant metrics on Teichmüller space. *Contributions to analysis* (a collection of papers dedicated to Lipman Bers), pp. 393–399. Academic Press, New York, 1974.
- [45] G. Schmithüsen, An algorithm for finding the Veech group of an origami. *Experiment. Math.* 13 (2004), no. 4, 459–472.
- [46] G. Schmithüsen, Examples for Veech groups of origamis. *The geometry of Riemann surfaces and abelian varieties*, 193–206, *Contemp. Math.*, 397, Amer. Math. Soc., Providence, RI, 2006.

- [47] J. Smillie, C. Ulcigrai, Geodesic flow on the Teichmüller disk of the regular octagon, cutting sequences and octagon continued fractions maps. Dynamical numbers—interplay between dynamical systems and number theory, 29–65, Contemp. Math., 532, Amer. Math. Soc., Providence, RI, 2010.
- [48] J. Smillie, B. Weiss, Characterizations of lattice surfaces. *Invent. Math.* 180 (2010), no. 3, 535–557.
- [49] O. Teichmüller, Extremale quasikonforme Abbildungen und quadratische Differentiale. *Abh. Preuss. Akad. Wiss. Math.-Nat. Kl.* 1939, (1940). no. 22, 197 pp.
- [50] O. Teichmüller, Extremal quasiconformal mappings and quadratic differentials. Translated from the German by Guillaume Thret. IRMA Lect. Math. Theor. Phys., 26, Handbook of Teichmller theory. Vol. V, 321–483, Eur. Math. Soc., Zrich, 2016.
- [51] W. Thurston, On the geometry and dynamics of diffeomorphisms of surfaces. *Bull. Amer. Math. Soc. (N.S.)* 19 (1988), no. 2, 417–431.
- [52] W. Veech, Gauss measures for transformations on the space of interval exchange maps. *Ann. of Math. (2)* 115 (1982), no. 1, 201–242.
- [53] W. Veech, The Teichmller geodesic flow. *Ann. of Math. (2)* 124 (1986), no. 3, 441–530.
- [54] W. Veech, Teichmller curves in moduli space, Eisenstein series and an application to triangular billiards. *Invent. Math.* 97 (1989), no. 3, 553–583.
- [55] W. Veech, Moduli spaces of quadratic differentials. *J. Analyse Math.* 55 (1990), 117–171.
- [56] Ya. Vorobets, Plane structures and billiards in rational polygons: the Veech alternative. *Russian Math. Surveys* 51 (1996), no. 5, 779–817
- [57] C. Ward, Calculation of Fuchsian groups associated to billiards in a rational triangle. *Ergodic Theory Dynam. Systems* 18 (1998), no. 4, 1019–1042.
- [58] A. Wright, Translation surfaces and their orbit closures: an introduction for a broad audience. *EMS Surv. Math. Sci.* 2 (2015), no. 1, 63–108.
- [59] A. Zemljakov, A. Katok, Topological transitivity of billiards in polygons. *Mat. Zametki* 18 (1975), no. 2, 291–300.
- [60] A. Zorich, Flat surfaces. *Frontiers in number theory, physics, and geometry. I*, 437–583, Springer, Berlin, 2006.
- [61] A. Zorich, Le thorme de la baguette magique de A. Eskin et M. Mirzakhani. *Gaz. Math.* No. 142 (2014), 39–54.

Adress:

Daniel Massart, Institut Montpelliérain Alexander Grothendieck, Université de Montpellier, France

email: daniel.massart@umontpellier.fr