FLAT STRUCTURES AND COMPLEX STRUCTURES IN TEICHMÜLLER THEORY

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Joshua Paul Bowman

August 2009

FLAT STRUCTURES AND COMPLEX STRUCTURES IN TEICHMÜLLER THEORY

Joshua Paul Bowman, Ph.D.

Cornell University 2009

We consider canonical invariants of flat surfaces and complex structures, including the combinatorics of Delaunay triangulations and boundary strata of the Siegel half-plane. These objects have been previously considered by various other authors; we provide fresh perspectives on how they arise naturally, develop some new results on their geometric structure, and give explicit examples of applications. We also study an important infinite family of flat surfaces, and extend this family by adding a surface of infinite genus, the study of whose affine structure leads to interesting new examples of dynamical and geometric behavior.

BIOGRAPHICAL SKETCH

Joshua Paul Bowman was born in Dallas, TX, in 1977. His family moved to Memphis, TN, in 1980, and he attended Harding Academy of Memphis from 1982 to 1995, participating in band, chorus, and theatre. In 1999, he obtained a Bachelor of Arts degree *summa cum laude* from St. Olaf College in Northfield, MN, where he majored in mathematics and music and performed with the St. Olaf Band, the St. Olaf Orchestra, and various choirs, including the Viking Chorus, the Chapel Choir, and Cantorei. Following a year living in Minneapolis, MN, he joined the Peace Corps in July 2000 and served as a Volunteer for two years, teaching secondary mathematics in Kérouané, Guinea, West Africa. During this time he founded a club for students to engage in science and mathematics outside of the classroom and sought out materials for science classes; this latter project led to the dissemination of laboratory materials to six *lycées* around the country.

Bowman began his graduate studies at Cornell University in 2003, and obtained his master's degree in 2006. His research interests began with combinatorial geometries, then shifted to hyperbolic geometry, from which he was led to Teichmüller theory. He recently completed a translation of *The Scientific Legacy of Henri Poincaré* for the AMS. Among his activities in Ithaca have been contradance, singing with the Cornell Chorale and the choir of Christ Chapel, hosting "music-listening parties" to share his interest in contemporary composition, and leading Bible studies with the Graduate Christian Fellowship. He participated in the annual Math Department Spring Concert each of the past six years, first performing with a barbershop quintet, and organizing the event in 2008 and 2009.

In 2008, Bowman became engaged to Hannah Newfield-Plunkett, who graduated from Cornell in May 2009 with degrees in mathematics *magna cum laude* and English *summa cum laude*. They will be married on August 1, 2009.

To my father, who talked about math with me at dinner.

To my mother, who supported all my creative and academic activities.

To Hannah, who fills my life with joy.

ACKNOWLEDGEMENTS

I am grateful to my advisor John H. Hubbard for his guidance and instruction and to my committee members John Smillie, Bill Thurston, and Cliff Earle for being accessible and willing to discuss the points of my work I found sticky. Thanks to Barbara Hubbard for being welcoming and encouraging. Thanks to the many friends I have made among the grad students in the math department for mathematical conversations, potluck dinners, and trips to the Dairy Bar. Thanks as well to the numerous authors and colleagues from whom I gained much clarity.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER 1

PRELIMINARIES

1.1 Flat surfaces, locally Euclidean surfaces, and differentials

The field of flat surfaces is a playground for a range of subjects from elementary Euclidean and hyperbolic geometry to complex analysis and dynamical systems. It has its origins in Teichmüller's theory of extremal quasi-conformal maps and the theories of Nielsen and Anosov on surface homeomorphisms. Until the 1970s, the perspective was primarily that of complex analysis. In the 1980s, it began to shift to a more elementary geometric viewpoint, with dynamical questions playing a larger role. The main connection between these two perspectives is the relative ease with which a surface can be endowed with a conformal structure. Among the conformal maps of $\mathbb{C}$ are the Euclidean isometries. Thus an ordinary "flat" geometric surface, such as a polyhedron or the usual torus $\mathbb{R}^2/\mathbb{Z}^2$, has a canonical conformal structure associated to it. The Euclidean structure is then specified by a differential, holomorphic with respect to the conformal structure. Even apparent singularities for the geometric structure, such as the vertices of a polyhedron or more generally any "cone-type" points, are shown to be nothing more than zeroes or "at worst simple" poles for the differential. (See [46] for a careful exposition.)

As Teichmüller's work shows, in the moduli theory of Riemann surfaces, preference is given to *quadratic differentials*—that is, tensors that are locally of the form $f(z)\,\mathrm{d}z^2$. From the geometric perspective, these may be constructed from polygons in $\mathbb{R}^2$ whose edges are identified by either translation or rotation by π.

There are several classes of surfaces that are interesting to consider, and the terminology to describe them is not quite standard. We will therefore take *locally Euclidean surface* to mean a surface without boundary formed from Euclidean polygons by gluing their edges via isometries. If the polygons are located in $\mathbb{R}^2$ and all of the edge identifications are by translations, then we call the result a *translation surface*—in the complex analytic realm, this corresponds to a Riemann surface carrying an abelian differential. Most of our results will be concentrated on a class intermediate between these two—we will take *flat surface* to mean either a translation surface or a Riemann surface carrying a quadratic differential, which from the elementary point of view can be constructed from polygons in $\mathbb{R}^2$ by gluing edges via translations or central symmetries (rotations by π). We will mention, in §4.4 and only briefly, a class of surfaces that encompasses the flat surfaces (but not more general locally Euclidean surfaces), called *homothety surfaces*.

1.2 Actions of linear groups and projective linear groups

The spaces of translation surfaces and flat surfaces admit actions by $\mathrm{SL}_2(\mathbb{R})$ and $\mathrm{PSL}_2(\mathbb{R})$, respectively, via post-composition with charts of the translation or flat structure. The action by $\mathrm{SO}_2(\mathbb{R})$ is trivial on the level of complex structures, and merely rotates the directions on each surface. This fact is used to identify the $\mathrm{PSL}_2(\mathbb{R})$ orbit of a flat surface with the unit tangent bundle of the upper half-plane $\mathbb{H}$, with the underlying space $\mathbb{H}$ corresponding to a *Teichmüller disk* in the moduli space of Riemann surfaces.

The (*projective*) *Veech group* of a flat surface is the subgroup of $\mathrm{PSL}_2(\mathbb{R})$ consisting of elements that send the surface to an isometric flat surface (i.e., the

stabilizer of the flat surface, although one needs to pay attention to the marking in many applications). Equivalently, it is the group of derivatives of homeomorphisms that are affine with respect to the flat structure. (See [15] for a clear exposition, as well as an indication of connections with the complex analytic viewpoint.)

There is great interest in knowing the Veech group, or better yet the affine group, of a given flat surface; it is generically trivial, but many important surfaces appear with "large" Veech groups. Two of the best-known constructions of surfaces with non-trivial affine groups have been provided by Thurston (later refined by Veech), in his classification of surface automorphisms, and by a paper of Arnoux and Yoccoz. We will present these constructions later in this work.

CHAPTER 2

RESULTS OF THIS THESIS

Here we state the main results of this work.

Part I: Delaunay triangulations of flat surfaces

In the first part, we consider the data of Delaunay triangulations of flat surfaces. These data determine a canonical partition of the Teichmüller space of flat structures, which leads to a tessellation of the disk (hyperbolic plane) attached to each flat surface. These tessellations have been considered by W. Veech in a preprint [48] and recently by A. Broughton and C. Judge [8]. We provide independent proofs of Veech's results, and draw new consequences from the construction. Specifically, we have the following:

Theorem 2.1. *Each flat surface generates a tessellation of the upper half-plane that is invariant under the Veech group Γ of the surface. If Γ is a lattice, then it has a finite-index subgroup with a fundamental domain composed of tiles of the corresponding tessellation. The squares of the edge lengths and angle cotangents appearing in the tessellation lie in the holonomy field of the generating flat surface.*

Moreover, we conjecture the following: *Each tile of the tessellation associated to a flat surface is contained in a hyperbolic triangle, possibly with vertices at infinity, and therefore has area at most π.* Apart from the finiteness of the area of the tiles, this would be the first universal bound on the geometry of these tiles. We expect a proof of this result to be forthcoming.

Part II: Complex structures and odd cohomology

In the second part, we address a question raised by J. Hubbard: Is it possible to characterize the Jacobian (or some sub-abelian variety) of a Riemann surface using purely topological data? This question remains open, in general. The specific context of Hubbard's question involves a complex structure J_0 on the odd cohomology of a Riemann surface X_0 (with respect to an involution that appears as part of the construction), and asks whether this complex structure coincides with that on the odd part of the Jacobian of X_0. We answer this question in the negative, and provide a characterization of J_0 in the case that a certain polynomial appearing in the construction is irreducible.

In these investigations, we found it useful to develop and augment some of the theory of the Siegel half-plane $\mathfrak{H}$ in as coordinate-free a way as possible. We obtain the following results:

Theorem 2.2. *The Siegel half-plane $\mathfrak{H}$ of complex dimension $(n^2+n)/2$ has, in the Lagrangian Grassmannian variety of a certain related complex vector space, a canonical boundary which is stratified by real manifolds indexed by $1 \leq p \leq n$ and having dimension $n^2+n-(p^2+p)/2$. A geodesic for the Siegel metric on $\mathfrak{H}$ has two endpoints in $\partial\mathfrak{H}$, which must lie in the same stratum.*

Using these boundary strata, we are able to distinguish certain complex-analytic maps into $\mathfrak{H}$, and to conclude:

Theorem 2.3. *The complex structure J_0 on $H^1(X_0,\mathbb{R})^-$ belongs to a family $\{J_t\}$ of complex structures that does not coincide with those arising from any Teichmüller disk having non-elementary Veech group. In particular, J_0 does not coincide with the complex structure arising from the Jacobian of X_0.*

The family $\{J_t\}$ will be given explicitly, and shown to extend to a maximal holomorphically immersed disk in the Siegel half-plane of $H^1(X_0, \mathbb{R})^-$.

Part III: Examples

In the third part, we provide examples of iso-Delaunay tessellations, primarily in the case of surfaces with lattice Veech groups, which lead to independent verifications of what their Veech groups are.

Finally, we carry out an in-depth study of an exceptional family of translation surfaces, those provided by P. Arnoux and J.-C. Yoccoz. We prove:

Theorem 2.4. *The Arnoux–Yoccoz surfaces converge metrically on compact subsets of the complements of their singular points to a surface of infinite genus carrying a natural* 1*-form of finite area. The Veech group of this limit surface is isomorphic to* $\mathbb{Z} \times \mathbb{Z}/(2)$.

Part I

Flat structures

CHAPTER 3

DELAUNAY WEIGHTS OF PURE SIMPLICIAL COMPLEXES

3.1 Delaunay triangulations in Euclidean space

Definition 3.1. Let τ be a connected, pure k-dimensional simplicial complex embedded in k-dimensional Euclidean space $\mathbb{R}^k$ such that its vertex set τ^0 is discrete, and let $\overline{\tau} \subset \mathbb{R}^k$ be its underlying set of points. The *empty circumsphere* condition states that, for each simplex $T \in \tau$, no point of τ^0 is contained in the interior of the sphere circumscribed around T. If τ satisfies the empty circumsphere condition and $\overline{\tau}$ is convex, then τ is called a *Delaunay triangulation* of its vertex set.

Note that τ is not required to be finite, but it must be locally finite. Given a discrete set of points $\tau^0 \subset \mathbb{R}^k$, there are several ways of proving the existence of a Delaunay triangulation τ of τ^0, one of which we will give below. The uniqueness (up to certain trivial exchanges) is obtained by comparing the Delaunay triangulation with the dual Voronoi construction, which we here omit. A key observation before proceeding further is Delaunay's famous lemma [14], which reduces the empty circumsphere condition to one that can be checked "locally", i.e., on adjacent pairs of facets.

Proposition 3.2 (Delaunay lemma)**.** *Let τ be as in the previous definition, and assume $\overline{\tau}$ is convex. Then τ satisfies the empty circumsphere condition if and only if every subcomplex consisting of two adjacent facets of τ does.*

Delaunay's proof is elementary and extremely geometric. First, recall that the *power* of a point $P \in \mathbb{R}^k$ with respect to a sphere $S \subset \mathbb{R}^k$ is the following invariant:

let a secant line to S through P be drawn, let P' and P'' be the points of intersection of this line with S, and compute the scalar product $\langle P' - P, P'' - P \rangle$. If the line is taken to be the diameter through P, then this reduces to $|P - O|^2 - \rho^2$, where O is the center of S and ρ is its radius. Now Delaunay's argument consists of looking at the linear segment $[AB]$ from an arbitrary vertex $A \in \tau^0$ to a point B in the interior of an arbitrary facet $T \in \tau$, such that the segment only intersects k-faces and $(k-1)$-faces of τ. Let $T_0, T_1, \ldots, T_n = T$ be the sequence of facets of τ that are intersected by $[AB]$ while moving from A to B, and for each $0 \leq i \leq n$, let a_i be the power of A with respect to the circumsphere of T_i. The assumption that pairs of adjacent simplices satisfy the empty circumsphere condition implies that the sequence $\{a_i\}$ is non-decreasing, and because $a_0 = 0$, we have $a_n \geq 0$. Therefore A is not contained in the interior of the circumsphere of T.

Note that, if two simplices in $\mathbb{R}^k$ share a facet, then each has a "free" vertex, and the property that one of these free vertices is not contained in the circumsphere of the other simplex is symmetric: the common facet determines a hyperplane H in $\mathbb{R}^k$, and the two free vertices must be on opposite sides of H (otherwise, the interiors of the simplices would overlap). By considering the possible relative positions of two spheres that intersect along a sphere of one lower dimension, we see that one free vertex is contained in the second sphere if and only if the second free vertex is contained in the first sphere.

Thus the entire situation reduces to determining whether one point in $\mathbb{R}^k$ is contained on or outside the sphere determined by another set of $k+1$ points. This is checked by means of the *insphere test*. Let T be a simplex in $\mathbb{R}^k$ with vertex set $\{v_0, v_1, \ldots, v_k\}$, ordered so that $\{v_1 - v_0, \ldots, v_k - v_0\}$ is a direct basis of $\mathbb{R}^k$, and let S be the circumsphere of T.

Proposition 3.3 (Insphere test)**.** *Given any point v' of $\mathbb{R}^k$, let Δ be the $(k+2) \times (k+2)$ determinant*

$$\Delta = \begin{vmatrix} (v_0)_1 & (v_0)_2 & \cdots & (v_0)_k & |v_0|^2 & 1 \\ (v_1)_1 & (v_1)_2 & \cdots & (v_1)_k & |v_1|^2 & 1 \\ \vdots & \vdots & & \vdots & \vdots & \vdots \\ (v_k)_1 & (v_k)_2 & \cdots & (v_k)_k & |v_k|^2 & 1 \\ v'_1 & v'_2 & \cdots & v'_k & |v'|^2 & 1 \end{vmatrix}.$$

$$\textit{Then } \Delta \textit{ is } \begin{cases} \textit{negative} & \textit{if and only if } v' \textit{ lies inside } S; \\ 0 & \textit{if and only if } v' \textit{ lies on } S; \\ \textit{positive} & \textit{if and only if } v' \textit{ lies outside } S. \end{cases}$$

The proof proceeds by considering the paraboloid $\mathcal{P}$ with equation $x_{k+1} = x_1^2 + \cdots + x_k^2$ in $\mathbb{R}^{k+1}$. Each v_i projects vertically to $\tilde{v}_i = ((v_i)_1, \ldots, (v_i)_k, |v_i|^2) \in \mathcal{P}$, and the $k+1$ points $\{\tilde{v}_0, \ldots, \tilde{v}_k\}$ are contained in a unique hyperplane $\mathcal{H}$. The intersection of $\mathcal{H}$ with $\mathcal{P}$ is an ellipsoid, which is the vertical projection to $\mathcal{H}$ of the sphere $S \subset \mathbb{R}^k \subset \mathbb{R}^{k+1}$. Because the orientations on $\mathcal{H}$ and $\mathbb{R}^k$ induced by the upward vertical direction in $\mathbb{R}^{k+1}$ correspond to each other via the vertical projection $\mathbb{R}^k \to \mathcal{H}$, $\{\tilde{v}_1 - \tilde{v}_0, \ldots, \tilde{v}_k - \tilde{v}_0\}$ is a direct basis for $\mathcal{H}$. Now, if $\tilde{v}'$ is the vertical projection of v' to $\mathcal{P}$, it follows that the $(k+1) \times (k+1)$ determinant

$$\begin{vmatrix} \tilde{v}_1 - \tilde{v}_0 & \cdots & \tilde{v}_k - \tilde{v}_0 & \tilde{v}' - \tilde{v}_0 \end{vmatrix},$$

which equals Δ, is positive if and only if $\tilde{v}'$ lies above $\mathcal{H}$ (in which case v' is outside of S), negative if and only if $\tilde{v}'$ lies below $\mathcal{H}$ (in which case v' lies inside of S), and 0 if and only if $\tilde{v}'$ lies on $\mathcal{H}$ (in which case v' lies on S).

This method of proof also leads to what is perhaps the simplest proof of the existence of a Delaunay triangulation for a discrete set $\tau^0 \subset \mathbb{R}^k$: the vertical

projection of τ^0 to $\mathcal{P}$ is again a discrete set $\tilde{\tau}^0$, and its convex hull is therefore a (possibly unbounded) polyhedron $\mathcal{C}$. If $\mathcal{C}$ is unbounded, we set $\tilde{\tau} = \partial\mathcal{C}$, and if $\mathcal{C}$ is bounded, then we obtain $\tilde{\tau}$ from $\partial\mathcal{C}$ by deleting any cells whose interiors lie over the interiors of other cells in $\mathcal{C}$. If any of the facets of $\tilde{\tau}$ are not simplices, they may be triangulated, making $\tilde{\tau}$ into a simplicial complex. The vertical projection τ of $\tilde{\tau}$ to $\mathbb{R}^k$ is an embedding, and by the proof of the incircle test, it is a Delaunay triangulation of its vertex set τ^0.

We now wish to convert the incircle test, which relies on a set of coordinates, into a purely metric statement, which will allow us to extend the definition of Delaunay triangulations to simplicial complexes that are Euclidean on each facet, but are not necessarily embedded. For the remainder of this paper, we will only concern ourselves with the case of two-dimensional complexes whose underlying space is a manifold, which we will call *simplicial surfaces*.

3.2 Cotangents and Delaunay weights

The expression for the cotangent of an angle formed by a pair of vectors $\xi, \eta \in \mathbb{R}^2$ will be useful throughout what follows:

$$\cot \angle(\xi, \eta) = \frac{\langle \xi, \eta \rangle}{|\xi\ \eta|}, \tag{3.1}$$

where $\langle \xi, \eta \rangle$ denotes the inner product of ξ and η, and $|\xi\ \eta|$ is the determinant of ξ and η, i.e., the standard 2-form on $\mathbb{R}^2$ evaluated on ξ and η. Note also that if α, β, γ are the angles in a triangle, then they satisfy

$$\cot\alpha \cot\beta + \cot\beta \cot\gamma + \cot\gamma \cot\alpha = 1. \tag{3.2}$$

The surface of equation $xy + yz + zx = 1$ is a hyperboloid H of two sheets, one corresponding to the cotangents of three angles that sum to π and the other corresponding to the cotangents of three angles that sum to 2π. The metric induced on this hyperboloid by the quadratic form $-(xy + yz + zx)$ makes each sheet isometric to the standard hyperbolic plane (e.g., the upper half-plane model), and this isometry can be realized concretely by sending the point z in the upper half-plane to the cotangents of the angles in the triangle $(0, 1, z)$. This metric will appear later as the Teichmüller metric on the bundle of quadratic differentials.

Definition 3.4. Let T_1 and T_2 be a pair of Euclidean triangles, joined along a common edge E. Let α and β be the angles opposite E. We call

$$w(E) = \frac{\cot \alpha + \cot \beta}{2}$$

the *Delaunay weight* of E (suppressing its dependence on T_1 and T_2 in the terminology, but context will always make abundantly clear what determines the weights on E). We will call an edge in a simplicial surface τ *Delaunay* if its Delaunay weight is non-negative, and we call τ a Delaunay triangulation of $\overline{\tau}$ if all edges are Delaunay.

Here is the motivation for the above definition. We may assume that T_1 and T_2 lie in the same plane, and by classical geometry they satisfy the empty circumcircle condition if and only if the sum of the angles opposite F is less than or equal to π. This condition is transcendental in the edge lengths, however, and we would like to give an algebraic condition. If α and β are angles strictly between 0 and π, then we have

$$\alpha + \beta \leq \pi \iff \cot \alpha + \cot \beta \geq 0,$$

and the equality is also an if-and-only-if statement (see Figure 3.1). The 1/2 in

Definition 3.4 is a normalizing factor that will make some later results look more natural; it is also consistent with the notation of [5].

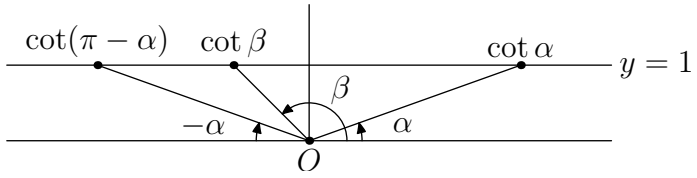

Figure 3.1: Proof that $\alpha + \beta \leq \pi \iff \cot\alpha + \cot\beta \geq 0$.

Definition 3.5. Let T be a Euclidean triangle and E a distinguished edge of T. The *modulus of T with respect to E* is

$$\mu_E(T) = \frac{\text{altitude of } T \text{ from vertex opposite } E}{\text{length of } E} = \frac{2 \cdot \text{area of } T}{(\text{length of } E)^2}$$

For many purposes, it is more useful to use the inverse of the modulus of a triangle (e.g., it is additive along cylinders, since all the triangles have the same height; it appears naturally in descriptions of the space of triangles up to similarity). If $\xi, \eta \in \mathbb{R}^2$ form a direct basis and $\zeta = \xi - \eta$, then the triangle T with ξ, η, ζ as side vectors has inverse modulus

$$\mu_E^{-1}(T) = \frac{\langle \zeta, \zeta \rangle}{|\xi\ \eta|} \tag{3.3}$$

with respect to the side E along which ζ lies. We note that the tangent space to the hyperboloid H (mentioned following equation (3.2)) at a point $(\cot\alpha, \cot\beta, \cot\gamma)$ corresponding to a triangle T with angles α, β, γ can be described as the kernel of the 1×3 matrix whose entries are the inverse moduli of T with respect to each of its sides; this follows from the relation

$$\mu_E^{-1}(T) = \cot\alpha + \cot\beta$$

where E is the edge of T between the angles α and β.

Lemma 3.6. *Let T_1 and T_2 be Euclidean triangles in $\mathbb{R}^2$ joined along a common edge E. Then a sufficient condition for E to be a Delaunay edge is*

$$\mu_E^{-1}(T_1)\,\mu_E^{-1}(T_2) \leq 4.$$

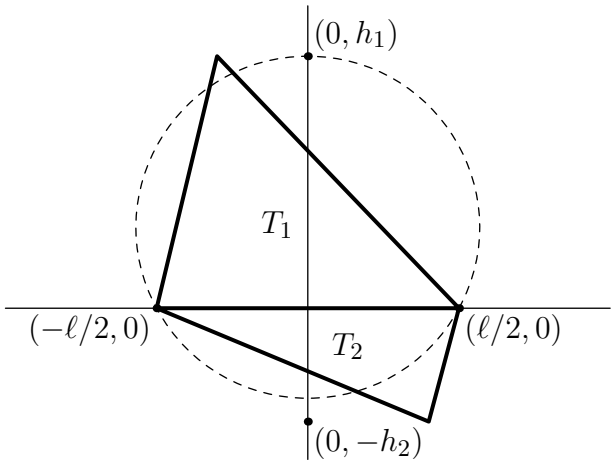

Figure 3.2: An inequality involving inverse moduli of joined triangles.

Proof. We may assume without loss of generality that E is the segment along the x-axis from $(-\ell/2, 0)$ to $(\ell/2, 0)$, T_1 lies above the x-axis, and T_2 lies below the x-axis (Figure 3.2). Let h_1 be the altitude of T_1 and h_2 be the altitude of T_2, both taken with respect to E. Then the Delaunay weight of E is minimized when the "free" vertices of T_1 and T_2 lie on the y-axis. The center of the circle $\mathcal{C}$ through $(-\ell/2)$, $(\ell/2)$, and $(0, h_1)$ is located at $(0, \frac{{h_1}^2 - (\ell/2)^2}{2h_1})$, and the radius of $\mathcal{C}$ is $\frac{{h_1}^2 + (\ell/2)^2}{2h_1}$. In order for $(0, -h_2)$ to lie outside of the interior of $\mathcal{C}$, we must have $h_1 h_2 \geq (\ell/2)^2$. This inequality is equivalent to the one in the statement of the result, and if it is satisfied, then E is *a fortiori* a Delaunay edge for T_1 and T_2. □

3.3 Delaunay triangulations of flat surfaces

Definition 3.7. Let (X, q) be a flat surface and Z be a discrete subset of X that includes the singularities of q. A (q, Z)-*triangle* on X is the image of a 2-simplex on X, embedded on its interior, whose edges are geodesic with respect to $|q|^{1/2}$, whose vertices lie in Z, and whose interior contains no points of Z. A (q, Z)-*triangulation* of X is a simplicial structure on X, all of whose facets are (q, Z)-triangles. We denote by $T(X, q, Z)$ the set of (q, Z)-triangulations of X. A (q, Z)-triangulation of X is *Delaunay* if all of its edges are Delaunay.

For simplicity, we will almost always drop the (q, Z) prefix and refer simply to "triangles", "triangulations", and "Delaunay triangulations" of X.

The principle concern is the existence and uniqueness of Delaunay triangulations. These have been established for compact half-translation surfaces (or locally Euclidean surfaces more generally) by a series of results over the last two decades.

- In [45] (originally circulated as a preprint starting c. 1987), Thurston sketched the construction of Delaunay triangulations for locally Euclidean structures on a sphere.
- Masur and Smillie [29] proved the existence of a Delaunay triangulation for any compact locally Euclidean surface by dualizing the construction of Voronoi cells and applied properties of the Delaunay triangulation of a surface to get sharp estimates in their study of non-ergodic directions.
- Rivin [38] studied triangulations of simplicial surfaces by attaching weights to the edges, which he called *dihedral angles* (these weights are similar to our Delaunay weights, but they are defined simply by adding the angles opposite

an edge of the triangulation, and not the cotangents of these angles); he described the space of locally Euclidean surfaces for which a given triangulation is Delaunay as a polytope in the space of functions on the edge set of the triangulation.

- Veech [50] first described a partition of the Teichmüller space of flat surfaces by the combinatorial structure of the Delaunay triangulations of its points and showed that this partition is equivariant with respect to the action of the mapping class group, analogously to Penner's partition of the decorated Teichmüller space [35].
- Indermitte, Liebling, Troyanov, and Clémençon [24] considered "edge flips" between various triangulations (and also give an application of Delaunay triangulations and Voronoi cells to modeling biological growth).
- Bobenko and Springborn [5] took up the work of Rivin and Indermitte et al, and completed the proof of the uniqueness of the Delaunay triangulation of a surface, using a method similar to that of Delaunay's original proof (see the discussion following Proposition 3.2); they also use what we have called Delaunay edge weights to define an intrinsic Laplacian on the vertex set of a simplicial surface.

We summarize these results in the following proposition:

Proposition 3.8. *Let (X, q) be a compact flat surface and Z a non-empty discrete subset of X containing the zeroes of q. Then X has a Delaunay triangulation, unique up to exchanges of edges with Delaunay weight* 0*. Any triangulation of X may be transformed into a Delaunay triangulation by a finite sequence of edge flips, at each step exchanging an edge having negative Delaunay weight for a transverse edge of positive Delaunay weight.*

CHAPTER 4

ISO-DELAUNAY TESSELLATIONS

4.1 Pre-geodesic conditions on points of $\mathrm{SL}_2(\mathbb{R})$

We denote the half-plane model of the hyperbolic plane by $\mathbb{H}$, and we use $z = u+iv$ for the canonical coordinate on $\mathbb{H}$. Geodesics for the Poincaré metric $|\mathrm{d}z|/v$ can be either lines or circles perpendicular to the u-axis. These two kinds of geodesics can be unified algebraically: any geodesic γ is the set of points satisfying an equation of the form

$$a(u^2 + v^2) + bu + c = 0 \qquad \text{with } b^2 - 4ac > 0.$$

The latter condition ensures that the geodesic has exactly two endpoints in $\partial\mathbb{H} = \mathbb{R} \cup \{\infty\}$:

- if $a \neq 0$, then the endpoints are the (real) solutions to $au^2 + bu + c = 0$;
- if $a = 0$, then the endpoints are ∞ and $-c/b$.

In this model, the isometry group $\mathrm{PSL}_2(\mathbb{R})$ acts by fractional linear transformations:

$$\left[\begin{pmatrix} a & b \\ c & d \end{pmatrix}\right] \cdot z = \frac{az+b}{cz+d}.$$

We make this into a right action by defining $z \cdot [A] = [A]^{-1} \cdot z$. The stabilizer of i under this right action is $\mathrm{SO}_2(\mathbb{R})/\{\pm\mathrm{id}\}$. The projection $P : \mathrm{PSL}_2(\mathbb{R}) \to \mathbb{H}$ given by $P([A]) = i \cdot [A]$ canonically identifies $\mathrm{PSL}_2(\mathbb{R})$ with the unit tangent bundle of $\mathbb{H}$; each right coset of $\mathrm{SO}_2(\mathbb{R})/\{\pm\mathrm{id}\}$ is sent to the set of unit vectors at a point of $\mathbb{H}$. We choose a canonical representative of each coset via the QR-decomposition

in $\mathrm{SL}_2(\mathbb{R})$:

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} v^{-1/2} & -uv^{-1/2} \\ 0 & v^{1/2} \end{pmatrix} \qquad \text{where}$$

$$\cos\theta = \frac{a}{\sqrt{a^2+c^2}}, \quad \sin\theta = \frac{c}{\sqrt{a^2+c^2}}, \quad u = -\frac{ab+cd}{a^2+c^2}, \quad \text{and} \quad v = \frac{1}{a^2+c^2}. \tag{4.1}$$

With this notation, $P([A]) = u + iv$, and the second matrix in the factorization of A is the canonical representative of the coset to which A belongs.

We now characterize subsets in $\mathrm{PSL}_2(\mathbb{R})$ of the form $P^{-1}(\gamma)$, where γ is a geodesic in $\mathbb{H}$.

Definition 4.1. A *pre-geodesic condition* on points $A \in \mathrm{SL}_2(\mathbb{R})$ is one that can be expressed by an equation of the form $\langle A\xi_1, A\xi_2\rangle = 0$ for some $\xi_1, \xi_2 \in \mathbb{R}^2$ such that $|\xi_1\ \xi_2| \neq 0$.

Note that, when ξ_1 and ξ_2 are fixed, the function $\langle A\xi_1, A\xi_2\rangle$ on $\mathrm{SL}_2(\mathbb{R})$ is invariant under the involution $A \mapsto -A$, and therefore it is also a well-defined function on $\mathrm{PSL}_2(\mathbb{R})$.

Lemma 4.2. *For every geodesic γ in $\mathbb{H}$, there exist $\xi_1, \xi_2 \in \mathbb{R}^2$ such that*

$$P^{-1}(\gamma) = \big\{[A] \in \mathrm{PSL}_2(\mathbb{R}) \mid \langle A\xi_1, A\xi_2\rangle = 0\big\}.$$

Conversely, every set of this form, with $|\xi_1\ \xi_2| \neq 0$, projects via P to a geodesic in $\mathbb{H}$.

Proof. We show the second claim first. Let $\xi_1 = (x_1, y_1)$ and $\xi_2 = (x_2, y_2)$. Then, using the QR-decomposition for points $A \in \mathrm{SL}_2(\mathbb{R})$ and the notation of (4.1), we have

$$\langle A\xi_1, A\xi_2\rangle = v^{-1}\big(x_1x_2 - (x_1y_2 + x_2y_1)\,u + y_1y_2(u^2+v^2)\big).$$

The discriminant of the equation $\langle A\xi_1, A\xi_2\rangle = 0$ is then

$$(x_1y_2 + x_2y_1)^2 - 4x_1x_2y_1y_2 = x_1^2y_2^2 - 2x_1x_2y_1y_2 + x_2^2y_1^2 = |\xi_1\ \xi_2|^2,$$

and so the set of solutions to this equation projects to a geodesic in $\mathbb{H}$ provided $|\xi_1\ \xi_2| \neq 0$.

Now suppose γ is given, and let $r_1, r_2 \in \mathbb{R} \cup \{\infty\}$ be its endpoints. If one of these, say r_2, equals ∞, then take $\xi_1 = (r_1, 1)$ and $\xi_2 = (1, 0)$. If both lie in $\mathbb{R}$, then take $\xi_1 = (r_1, 1)$ and $\xi_2 = (r_2, 1)$. Then the set of solutions to $\langle A\xi_1, A\xi_2\rangle = 0$ projects to γ via P. □

Thus the endpoints of the geodesic $P(\{[A] \mid \langle A\xi_1, A\xi_2\rangle = 0\})$ are the cotangents of the angles formed by ξ_1 and ξ_2 with the positive x-axis; that is, we recover $S^1 = \mathbb{RP}^1$ as the boundary of $\mathbb{H}$. We note that the description we have given for geodesics in $\mathbb{H}$ is closely related to the description of $\mathbb{H}$ as the space of inner products on $\mathbb{R}^2$, modulo scaling: for any $A \in \mathrm{SL}_2(\mathbb{R})$, $A^\top A$ is symmetric and positive definite, and therefore it defines an inner product on $\mathbb{R}^2$. This inner product depends only on the right coset of $\mathrm{SO}_2(\mathbb{R})$ to which A belongs (see (4.1)). The equation $\langle A\xi_1, A\xi_2\rangle = 0$ is satisfied by those A for which ξ_1 and ξ_2 are orthogonal in the corresponding inner product.

Example 4.3 (Another example of a pre-geodesic condition on points of $\mathrm{SL}_2(\mathbb{R})$)**.** Given two non-colinear $\xi_1, \xi_2 \in \mathbb{R}^2$, require that $A\xi_1$ and $A\xi_2$ have the same length. This situation is described by the equation $\langle A\xi_1, A\xi_1\rangle = \langle A\xi_2, A\xi_2\rangle$, which is equivalent to the condition $\langle A(\xi_1 + \xi_2), A(\xi_1 - \xi_2)\rangle = 0$.

It is worth considering what meaning a "solution" to $\langle A\xi, A\xi\rangle = 0$ could have when ξ is a non-zero vector in $\mathbb{R}^2$. To make sense of this equation, we consider

more generally equations of the form $\langle A\xi, A\xi\rangle = r$, where $r \geq 0$. For $r > 0$, the solutions to $\langle A\xi, A\xi\rangle = r$ project to a *horocycle* in $\mathbb{H}$: setting $\xi = (x; y)$ and using the QR-decomposition again, we have

$$\langle A\xi, A\xi\rangle = r \iff x^2 - 2xy\,u + y^2(u^2 + v^2) = rv$$

which if $y = 0$ is the equation of a horizontal line, and if $y \neq 0$ is the equation of a Euclidean circle tangent to the real axis at $x/y = \cot\xi$. If we fix ξ and let $r \to 0$, then in the case $y = 0$ the intersection of the horizontal line with the imaginary axis tends to ∞, and in the case $y \neq 0$ the radius of the circle tends to 0. We have thus proved the following.

Lemma 4.4. *Let $\xi \in \mathbb{R}^2$ be non-zero, and let $c \in \partial\mathbb{H}$ denote the cotangent of the angle ξ forms with the horizontal direction. Let $\{A_n\}$ be a sequence in $\mathrm{SL}_2(\mathbb{R})$. Then the following are equivalent:*

1. *$P([A_n]) \to c$ in the horoball topology as $n \to \infty$;*
2. *$\|A_n\xi\| \to 0$ as $n \to \infty$.*

Recall that *horoball topology* on $\overline{\mathbb{H}}$ is generated by open sets of $\mathbb{H}$ and, for each point c of the boundary $\partial\mathbb{H}$, the union of $\{c\}$ with any open horoball centered at c. Note that this is a refinement of the topology on $\overline{\mathbb{H}}$ induced by its inclusion into $\mathbb{CP}^1$. In particular, the topology induced by the horoball topology on $\partial\mathbb{H}$ as a subspace is discrete.

At this point, we could examine the sets of solutions to general equations of the form $\langle A\xi_1, A\xi_2\rangle = r$—these correspond to sets that project via P to Euclidean circles that intersect the real axis (curves of constant curvature between 0 and 1)—for which the pre-geodesic and "pre-horocyclic" conditions are limiting cases, but we will not need them in the future, and so we set aside this study.

We can obtain "new" pre-geodesic conditions by taking certain "linear combinations" of "old" ones. We obtain the following by expanding and factoring the discriminant.

Lemma 4.5. *Let $\xi_1, \xi_2, \eta_1, \eta_2 \in \mathbb{R}^2$ and $a, b \in \mathbb{R}$. Then the equation*

$$a \langle A\xi_1, A\xi_2 \rangle + b \langle A\eta_1, A\eta_2 \rangle = 0$$

defines a pre-geodesic condition on points of $\mathrm{SL}_2(\mathbb{R})$ if and only if

$$a^2 |\xi_1\ \xi_2|^2 + b^2 |\eta_1\ \eta_2|^2 + 2ab\big(|\xi_1\ \eta_1| \cdot |\eta_2\ \xi_2| + |\xi_1\ \eta_2| \cdot |\eta_1\ \xi_2|\big) > 0.$$

We close this section by considering inequalities of the form

$$\langle A\xi_1, A\xi_2 \rangle \geq 0.$$

If $|\xi_1\ \xi_2| = 0$, then this inequality is either always or never satisfied, depending on the sign of $\langle \xi_1, \xi_2 \rangle$. If $|\xi_1\ \xi_2| \neq 0$, then the solutions to this inequality project to a *hyperbolic half-plane* (as opposed to the Poincaré half-plane, which is all of $\mathbb{H}$). Such a half-plane is equivalent to an *oriented* geodesic: the half-plane is the set of points "to the left" as we move along the geodesic in the direction of its orientation. Thus, if $|\xi_1\ \xi_2| > 0$, we call $\cot \xi_1$ the *first* endpoint and $\cot \xi_2$ the *second* endpoint of the corresponding geodesic.

4.2 Iso-Delaunay half-planes for convex quadrilaterals

Let $Q = P_1P_2P_3P_4$ be a simple quadrilateral in the plane, oriented counterclockwise. We allow two adjacent sides to be colinear, but in this section we will use the phrase "strictly convex" when we wish to exclude this possibility (i.e., "Q is

strictly convex" means that Q equals the convex hull of $\{P_1, P_2, P_3, P_4\}$, and not of any proper subset). For clarity, we will also assume Q has non-empty interior, or equivalently that the vertices of Q do not all lie in a single line. We say Q is *cyclic* if it is inscribable in a circle (see Figure 4.1).

Lemma 4.6. *The condition "$A(Q)$ is cyclic" is a pre-geodesic condition on $A \in \mathrm{SL}_2(\mathbb{R})$ if and only if Q is strictly convex.*

In a certain sense, this is geometrically obvious: if $A(Q)$ is cyclic, then it is in particular strictly convex. But the image of a non-convex quadrilateral by a linear map can never be strictly convex. Conversely, if Q is strictly convex, the natural equation that arises apparently has the form of a pre-geodesic condition on points of $\mathrm{SL}_2(\mathbb{R})$. This is what we will verify.

Proof. Set $\xi_1 = P_2 - P_1$, $\xi_2 = P_3 - P_2$, $\xi_3 = P_4 - P_3$, and $\xi_4 = P_1 - P_4$. The statement that Q is strictly convex is equivalent to $|\xi_i\ \xi_{i+1}| > 0$ for all i, where the indices are to be taken modulo 4. Note that Q can have at most one angle that equals or exceeds π, so at most one of these inequalities can fail to be satisfied. Relabeling (cyclically) if necessary, we may assume that $|\xi_1\ \xi_2| > 0$ and $|\xi_3\ \xi_4| > 0$.

The condition for $A(Q)$ to be cyclic is $\cot \angle(A\xi_1, A\xi_2) + \cot \angle(A\xi_3, A\xi_4) = 0$, or

$$|\xi_3\ \xi_4|\,\langle A\xi_1, A\xi_2\rangle + |\xi_1\ \xi_2|\,\langle A\xi_3, A\xi_4\rangle = 0, \tag{4.2}$$

where we have used the fact that $A \in \mathrm{SL}_2(\mathbb{R})$ to extract it from the determinants. By Lemma 4.5, Equation (4.2) defines a pre-geodesic condition if and only if

$$|\xi_1\ \xi_2| \cdot |\xi_3\ \xi_4| + |\xi_1\ \xi_3| \cdot |\xi_4\ \xi_2| + |\xi_1\ \xi_4| \cdot |\xi_3\ \xi_2| > 0.$$

Substituting $\xi_4 = -(\xi_1 + \xi_2 + \xi_3)$ and $\xi_3 = -(\xi_1 + \xi_2 + \xi_4)$ in the middle term, we

obtain

$$|\xi_1\ \xi_2| \cdot |\xi_3\ \xi_4| + (|\xi_1\ \xi_2| + |\xi_1\ \xi_4|) \cdot (|\xi_1\ \xi_2| + |\xi_3\ \xi_2|) + |\xi_1\ \xi_4| \cdot |\xi_3\ \xi_2| > 0,$$

which simplifies (after collecting a few terms and dividing by 2) to

$$|\xi_4\ \xi_1| \cdot |\xi_2\ \xi_3| > 0.$$

Because both factors on the left cannot simultaneously be negative, this inequality holds if and only if Q is strictly convex. This completes the proof. □

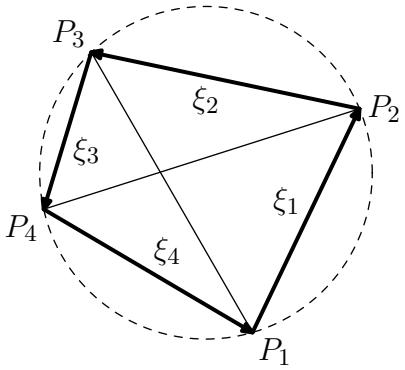

Figure 4.1: A cyclic quadrilateral: both diagonals have Delaunay weight zero.

Corollary 4.7. *Suppose Q is strictly convex, and let D be a diagonal of Q. If $w_A(D)$ denotes the Delaunay weight of $A(D)$ in $A(Q)$, then the equation $w_A(D) = 0$ is a pre-geodesic condition on $A \in \mathrm{SL}_2(\mathbb{R})$.*

This is simply a restatement of Lemma 4.6. We also will need a more detailed study of the various possibilities for a diagonal of a quadrilateral to have non-negative Delaunay weight.

Corollary 4.8. *Let T_1 and T_2 be triangles in $\mathbb{R}^2$ with disjoint interiors, joined by a common edge E. Let $\mathbb{H}_E$ be the projection to $\mathbb{H}$ via P of the set defined by $w_A(E) \geq 0$.*

1. *If $T_1 \cup T_2$ is a strictly convex quadrilateral, then $\mathbb{H}_E$ is a closed hyperbolic half-plane.*

2. *If $T_1 \cup T_2$ is a non-convex polygon, then $\mathbb{H}_E = \mathbb{H}$.*

3. *If $T_1 \cup T_2$ is a triangle, then one side E_1 of T_1 is aligned with a side E_2 of T_2; in this case, $\mathbb{H}_E = \mathbb{H}$ and the equation $w_A(E) = 0$ corresponds to the point of $\partial\mathbb{H}$ that is the cotangent of the common angle that E_1 and E_2 form with the x-axis.*

Remark 4.9. If Q is already cyclic, as in Figure 4.1, the endpoints of the geodesic in $\mathbb{H}$ described by Lemma 4.6 are given by the directions of the bisectors of the angles formed by the intersection of the diagonals of Q. I thank Chris Judge for pointing this out to me. In our approach, this can be seen as follows. Given any two non-zero vectors $\xi, \eta \in \mathbb{R}^2$, a non-zero vector σ *bisects* the angle $\angle(\xi, \eta)$ if $\cot \angle(\xi, \sigma) = \cot \angle(\sigma, \eta)$, which equation can be rearranged as

$$|\eta\ \sigma| \langle \xi, \sigma \rangle + |\xi\ \sigma| \langle \eta, \sigma \rangle = 0. \tag{4.3}$$

If we assume further that $|\xi\ \eta| \neq 0$, then (4.3) is equivalent to the condition

$$|\eta\ \sigma|^2 \langle \xi, \xi \rangle - |\xi\ \sigma|^2 \langle \eta, \eta \rangle = 0. \tag{4.4}$$

(This equivalence derives from the facts that $|\eta\ \sigma|\,\xi - |\xi\ \sigma|\,\eta$ is colinear with σ and that (4.3) is equivalent to $|\eta\ \sigma|\,\xi + |\xi\ \sigma|\,\eta$ being orthogonal to σ.) By Lemma 4.5, if the vectors ξ, η, σ satisfy $|\xi\ \sigma| \cdot |\eta\ \sigma| \neq 0$, then "$A\sigma$ bisects the angle $\angle(A\xi, A\eta)$" is a pre-geodesic condition on A. We need to show that, for appropriate values of the three vectors, this is equivalent to a cyclicity condition involving the side vectors of a quadrilateral.

Suppose we have as before a simple quadrilateral Q with side vectors

$\xi_1, \xi_2, \xi_3, \xi_4$, and Q is cyclic. For $\theta \in S^1$ and $\varepsilon > 0$, consider the matrix

$$A_{\theta,\varepsilon} = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} 1+\varepsilon & 0 \\ 0 & 1-\varepsilon \end{pmatrix} \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix},$$

whose invariant directions are $\pm(\cos\theta, \sin\theta)$ and $\pm(-\sin\theta, \cos\theta)$. (The middle factor in the expression for $A_{\theta,\varepsilon}$ is just the first-order Taylor expansion of $\exp\left(\begin{smallmatrix} \varepsilon & 0 \\ 0 & -\varepsilon \end{smallmatrix}\right)$.) Using the fact that Q is cyclic, we obtain

$$|\xi_3\ \xi_4|\,\langle A_{\theta,\varepsilon}\xi_1, A_{\theta,\varepsilon}\xi_2\rangle + |\xi_1\ \xi_2|\,\langle A_{\theta,\varepsilon}\xi_3, A_{\theta,\varepsilon}\xi_4\rangle = 2\varepsilon\big(|\xi_3\ \xi_4|\,\langle R\xi_1, \xi_2\rangle + |\xi_1\ \xi_2|\,\langle R\xi_3, \xi_4\rangle\big), \tag{4.5}$$

where $R = \begin{pmatrix} \cos^2\theta - \sin^2\theta & 2\cos\theta\sin\theta \\ 2\cos\theta\sin\theta & \sin^2\theta - \cos^2\theta \end{pmatrix}$. We are looking for values of θ such that the coefficient of ε in (4.5) is zero.

The diagonals of Q are given by the vectors $\xi_1 + \xi_2$ and $\xi_2 + \xi_3$. Let σ be a unit vector that bisects the angle $\angle(\xi_1 + \xi_2, \xi_2 + \xi_3)$; that is, $\|\sigma\| = 1$ and (from (4.4))

$$|\xi_2 + \xi_3\ \sigma|^2\,\|\xi_1 + \xi_2\|^2 - |\xi_1 + \xi_2\ \sigma|^2\,\|\xi_2 + \xi_3\|^2 = 0.$$

After replacing one copy of $\xi_1 + \xi_2$ with $-(\xi_3 + \xi_4)$ and one copy of $\xi_2 + \xi_3$ with $-(\xi_1 + \xi_4)$ in each term, expanding, and recollecting terms, this condition becomes

$$\begin{aligned} 0 \ &= |\xi_1\ \sigma|\,\langle \xi_1, \sigma\rangle\,\big(|\xi_3\ \xi_2| + |\xi_4\ \xi_2| + |\xi_4\ \xi_3|\big) \\ &\quad + |\xi_2\ \sigma|\,\langle \xi_2, \sigma\rangle\,\big(|\xi_3\ \xi_1| + |\xi_3\ \xi_4| + |\xi_4\ \xi_1|\big) \\ &\quad + |\xi_3\ \sigma|\,\langle \xi_3, \sigma\rangle\,\big(|\xi_1\ \xi_4| + |\xi_2\ \xi_4| + |\xi_2\ \xi_1|\big) \\ &\quad + |\xi_4\ \sigma|\,\langle \xi_4, \sigma\rangle\,\big(|\xi_1\ \xi_3| + |\xi_2\ \xi_3| + |\xi_1\ \xi_2|\big). \end{aligned}$$

It can be shown by a direct but somewhat lengthy computation that, if $\sigma = \binom{\cos\theta}{\sin\theta}$, then this expression equals the coefficient of ε in (4.5). Therefore the angle bisectors of the diagonals of Q are the invariant directions that determine the geodesic.

4.3 Consequences of a construction by Veech

Let C be an ordinary Euclidean cylinder with boundary components $\partial_1 C$ and $\partial_2 C$, and let Z_1 and Z_2 be non-empty finite subsets of $\partial_1 C$ and $\partial_2 C$, respectively. We will construct the Delaunay triangulation of C with respect to $Z = Z_1 \cup Z_2$, i.e., we endow C with the structure of a simplicial complex τ in such a way that the 0-cells of τ are points of Z, the connected components of $\partial_1 C - Z_1$ and $\partial_2 C - Z_2$ are 1-cells of τ, and each of the remaining 1-cells of τ is a segment connecting a point of Z_1 and a point of Z_2 and having non-negative Delaunay weight in τ.

For each pair of adjacent points P, P' in Z_1 or Z_2, let $\sigma(P, P')$ be the directed segment that is perpendicular to $\partial_1 C$ and $\partial_2 C$ and has its first endpoint at the midpoint of $[P, P']$. The segments $\sigma(P, P')$ partition C into rectangles. We call one of these rectangles *oriented* if the directed segments point in opposite directions, and *non-oriented* otherwise. The union of a non-oriented rectangle with an adjacent oriented rectangle yields a new oriented rectangle. An oriented rectangle is *direct* if the orientation induced by the directed edges is direct, and *indirect* otherwise. If two segments with opposite directions are superimposed, we say they form a *degenerate* rectangle. Define the *inverse modulus* of an oriented rectangle R to be $\mu^{-1}(R) = \pm \ell / h$, where h is the distance from $\partial_1 C$ to $\partial_2 C$ and ℓ is the remaining dimension of R, taken with a positive sign if R is direct and with a negative sign if R is indirect.

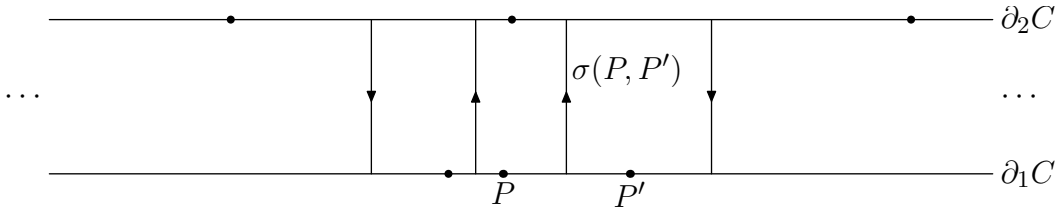

Figure 4.2: A direct, a non-oriented, and an indirect rectangle

Lemma 4.10. *Let $P_1P_1'P_2P_2'$ be a quadrilateral, oriented counterclockwise, such that $[P_1P_1']$ is parallel to $[P_2P_2']$, and let R be the oriented rectangle formed by $\sigma(P_1, P_1')$, $\sigma(P_2, P_2')$, and the lines supporting $[P_1P_1']$ and $[P_2P_2']$. Then the Delaunay weight of $[P_1, P_2]$ equals the inverse modulus of R.*

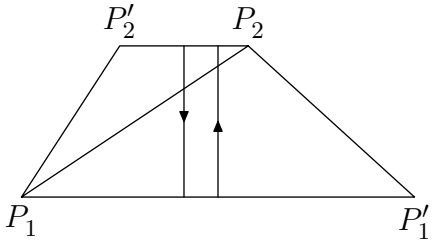

Proof. Assume the lines (P_1P_1') and (P_2P_2') are horizontal, as in the figure above. Let h be the distance between these two lines, and let the x-coordinates of P_1, P_1', P_2, P_2' be x_1, x_1', x_2, x_2', respectively. Then the Delaunay weight of the edge $[P_1P_2]$ is

$$\frac{1}{2}\left(\frac{x_1' - x_2}{h} + \frac{x_1 - x_2'}{h}\right) = \frac{1}{h}\left(\frac{x_1 + x_1'}{2} - \frac{x_2 + x_2'}{2}\right),$$

which is the inverse modulus of R. □

Now we complete the above construction by adding one edge for each rectangle of the partition. Each oriented rectangle—direct or indirect—arises from a trapezoid as in Lemma 4.10 and determines which diagonal of the trapezoid should be chosen. Any non-oriented rectangle belongs to a maximal sequence of adjacent such rectangles, which is flanked by two oppositely-oriented rectangles; these latter two determine a pair of edges with a common endpoint, which should be connected to all vertices on the opposite boundary component in between the pair of edges. A degenerate rectangle arises from an isosceles trapezoid, for which both diagonals have zero Delaunay weight, and either can be chosen to complete the triangulation.

Lemma 4.11. *Let (X, q) be a compact half-translation surface, let $Z \subset X$ be a discrete subset containing the singular points of q, and let $\theta \in \mathbb{RP}^1$. Then θ is a periodic direction of (X, q) if and only if there exist $\tau \in T(X, q, Z)$ and a sequence $\{A_n\} \subset \mathrm{SL}_2(\mathbb{R})$ such that $P([A_n]) \to \cot\theta$ in the horoball topology and τ is a Delaunay triangulation of $[A_n] \cdot (X, q)$ for all n.*

For completeness, we include the proof of this lemma here, although it is entirely due to Veech [48]. (We have made a slightly more general statement that does not affect the details of the proof in any way.)

Proof. By applying $\left[\left(\begin{smallmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{smallmatrix}\right)\right]$, we may assume that θ is horizontal.

First, suppose the horizontal direction of (X, q) is periodic. Using the construction from the start of this section, compute the Delaunay triangulation of each horizontal cylinder. Since each triangle of these trianglations has a horizontal edge, by contracting the horizontal direction we may make the inverse moduli of these triangles as small as we like (see Equation (3.3) on page 13). In particular, because there are finitely many triangles, we may contract the horizontal direction sufficiently that all these inverse moduli are smaller than 2. Then by Lemma 3.6 all of the horizontal saddle connections are in the Delaunay triangulation of such surfaces. If, while contracting the horizontal direction, we expand in the vertical direction, then the partition of the horizontal cylinders into oriented and non-oriented rectangles does not change (although the inverse moduli of the rectangles vary, they do not change sign), and so the Delaunay edges crossing each cylinder also remain so.

Next, suppose $P([A_n]) \to \infty$ in the horoball topology, and $\tau \in T(X, q, Z)$ is a Delaunay triangulation for every $[A_n] \cdot (X, q)$. We wish to show that every triangle

of τ has a horizontal edge; then these triangles can be joined along their non-horizontal edges to give a cylinder decomposition of (X, q). Write $\tau = \tau_\theta \sqcup \tau_\theta'$, where τ_θ is the set of faces of τ that have a horizontal edge, and τ_θ' is its complement. Suppose $\tau_\theta' \neq \varnothing$. For all sufficiently large n, each element of τ_θ has two angles that measure $> \pi/4$, and each element of τ_θ' has an angle that measures $> 3\pi/4$. Let $F_0 \in \tau_\theta'$, let α_0 be its largest angle, and let E_0 be the edge opposite α_0. E_0 cannot join F_0 to an element of τ_θ, since for any $F \in \tau_\theta$, two of the angle of F are larger than $\pi/4$, and the remaining angle is opposite a vertical edge, while E_0 cannot be vertical. Hence E_0 joins F_0 to some $F_1 \in \tau_\theta'$. The largest angle α_1 of F_1, also being $> 3\pi/4$, cannot be opposite E_0, and so the edge E_1 opposite α_1 satisfies $\text{length}(E_1) > \text{length}(E_0)$. Continuing inductively, we construct a sequence of edges $E_0, E_1, E_2, \ldots$, with $\text{length}(E_{i+1}) > \text{length}(E_i)$ for all i. But the set of edges of τ is finite, so this is a contradiction. We conclude $\tau_\theta' = \varnothing$, that is, $\tau = \tau_\theta$. □

Here are some other consequences of the above construction.

Corollary 4.12. *Let (X, q) be a compact flat surface, and assume that $\theta \in \mathbb{RP}^1$ is a periodic direction of (X, q). Then there is a horoball neighborhood B of $\cot\theta$ such that, whenever $P([A]) \in B$, any Delaunay triangulation of $[A] \cdot (X, q)$ contains all saddle connections in the direction θ.*

Proof. The inverse modulus of a triangle depends only on its base and its height. Therefore, for all sufficiently "small" horoball neighborhoods of $\cot\theta$, the inverse moduli of the triangles obtained by applying Veech's construction to the direction θ are all < 2. By the above reasoning, every Delaunay triangulation of points in this region must include the saddle connections in the direction θ. □

Corollary 4.13. *Let C be a horizontal cylinder with distinguished non-empty finite subsets of the two components of ∂C, and let τ be its Delaunay triangulation.*

Let w^+ (resp. w^-) be the smallest Delaunay weight of the edges of τ forming an obtuse (resp. acute) angle with ∂C, measured in the direct sense from the boundary component. Then the Delaunay triangulation of C remains unchanged by $\left(\begin{smallmatrix} 1 & t \\ 0 & 1 \end{smallmatrix}\right)$ for $-w^- \leq t \leq w^+$.

We will deduce other consequences in the next section.

4.4 Iso-Delaunay tessellations

Definition 4.14. Let Σ be a set of convex, finite-area (but not necessarily compact) polygons in $\mathbb{H}$. Then Σ is a *tessellation* of $\mathbb{H}$ if the elements of Σ cover $\mathbb{H}$, and whenever σ_1 and σ_2 are distinct elements of Σ, either $\sigma_1 \cap \sigma_2$ is empty or it is either a side or a vertex of both σ_1 and σ_2. If Σ is a tessellation, then the elements of Σ are called the *tiles*. The *edges* and *vertices* of a tessellation are the sides and vertices of its tiles.

Definition 4.15. Given a tessellation Σ of $\mathbb{H}$, an isometry $f : \mathbb{H} \to \mathbb{H}$ is an *automorphism* of Σ if $\{f(\sigma) \mid \sigma \in \Sigma\} = \Sigma$. $\mathrm{Aut}(\Sigma)$ is the group of all such automorphisms; $\mathrm{Aut}^+(\Sigma)$ is the subgroup of orientation-preserving automorphisms.

Note that $\mathrm{Aut}^+(\Sigma)$ has index at most 2 in $\mathrm{Aut}(\Sigma)$.

Lemma 4.16. *Let Σ be a tessellation of $\mathbb{H}$. Then $\mathrm{Aut}(\Sigma)$ is a discrete group of isometries. In particular, $\mathrm{Aut}^+(\Sigma)$ is a Fuchsian group.*

Proof. It is equivalent to show that $\mathrm{Aut}(\Sigma)$ does not have non-identity elements that are arbitrarily close to the identity. First observe that any finite-area polygon has a discrete stabilizer in $\mathrm{Isom}(\mathbb{H})$. Let $p \in \mathbb{H}$, $\sigma \in \Sigma$ be such that p is an interior

point of σ and p is not fixed by any element in the stabilizer of σ. Then there exists some $\varepsilon(p) > 0$ such that any element $f \in \mathrm{Aut}(\Sigma)$ moves p by at least $\varepsilon(p)$: either f stabilizes σ, in which case the above observation applies, or it moves σ to another tile of Σ, in which case p is moved by at least twice its distance to the boundary of σ. Because p cannot be moved arbitrarily small amounts, $\mathrm{Aut}(\Sigma)$ acts discretely on $\mathbb{H}$. □

Any triangulation $\tau \in T(X, q, Z)$ determines a closed, possibly empty, geodesically convex region in $\mathbb{H}$: for each edge $E \in \tau$, define

$$\mathbb{H}_E = P\big(\{[A] \in \mathrm{PSL}_2(\mathbb{R}) \mid w_A(E) \geq 0\}\big)$$

and define the *tile* corresponding to τ as

$$\mathbb{H}_\tau = \bigcap_{E \in \tau} \mathbb{H}_E. \tag{4.6}$$

Example 4.17. Let $\Lambda \subset \mathbb{C}$ be a lattice, and let $X = \mathbb{C}/\Lambda$ be the corresponding torus, carrying the 1-form ω induced by $\mathrm{d}z$. Let Z be the one-point set consisting of the image of Λ on X. Any (q, Z)-triangulation of X then consists of a symplectic basis for the homology of X, and their sum; these partition X into two congruent triangles, with their corresponding edges E_1, E_2, E_3 glued. Let $\mathcal{E} = \{E_1, E_2, E_3\}$, and let $w : \mathcal{T}(X) \to \mathbb{R}^{\mathcal{E}}$ be the weight map. Because each edge is opposite two congruent angles in a triangle, the image of w lies in the hyperboloid of equation $x_1x_2 + x_2x_3 + x_3x_1 = 1$, and the weights are all positive precisely when the triangles composing X have no obtuse angle, i.e., when $x_i \geq 0$ for all i, which condition determines an ideal triangle (see Figure 4.3). Because $\mathrm{SL}_2(\mathbb{Z})$ acts transitively on the symplectic bases of $H_1(X, \mathbb{R})$ and each basis determines an ideal triangle, the regions $\mathbb{H}_\tau$ arising from the triangulations $\tau \in T(X, \omega, Z)$ determine an $\mathrm{SL}_2(\mathbb{Z})$-invariant tessellation of $\mathbb{H}$ by ideal triangles, i.e., the Farey tessellation.

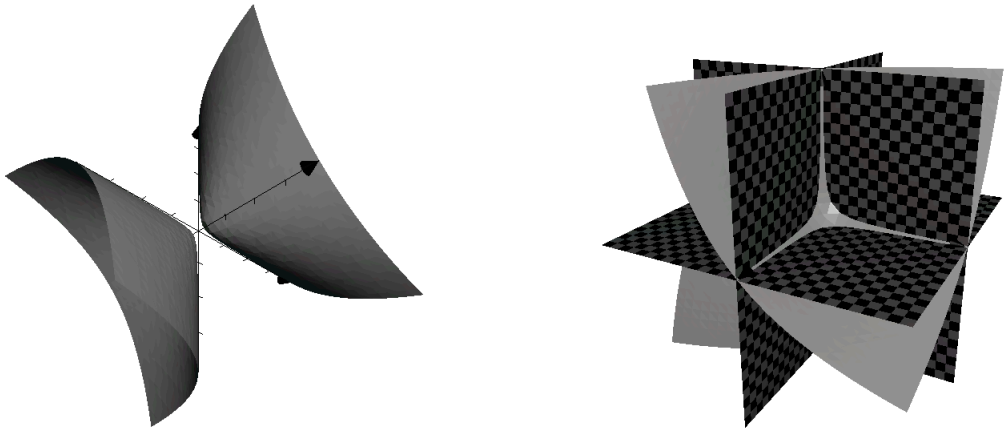

Figure 4.3: LEFT: The hyperboloid $x_1x_2 + x_2x_3 + x_3x_1 = 1$. RIGHT: The ideal triangle cut out by the coordinate planes.

Because any $\tau \in T(X, q, Z)$ has only finitely many edges, (4.6) is a finite intersection, and therefore its boundary is piecewise geodesic. We conjecture the following bound on its geometry: $\mathbb{H}_\tau$ *is contained in an ideal triangle in* $\mathbb{H}$*, and hence has area bounded by* π. Veech shows that the tiles have finite area by applying a result of Vorobets [52] that a compact flat surface has only countably many saddle connections. From our conjectured bound and another observation by Veech, we obtain that origami provide the extreme case of the Farey tessellation.

Corollary 4.18. $\mathbb{H}_\tau$ *has area* π *if and only if* (X, q, Z) *is affinely equivalent to an origami for which* Z *consists of all the corners of the square tiles.*

Definition 4.19. Let (X, q, Z) be a pointed compact flat surface, and set

$$\Sigma(X, q, Z) = \{\mathbb{H}_\tau \mid \tau \in T(X, q, Z),\ \mathbb{H}_\tau \text{ has non-empty interior}\}.$$

We call $\Sigma(X, q, Z)$ the *iso-Delaunay tessellation* of $\mathbb{H}$ arising from (X, q, Z).

It follows from our previous work that:

Theorem 4.20. $\Sigma(X, q, Z)$ *is a tessellation of* $\mathbb{H}$*, in the sense of Definition 4.15.*

Remark 4.21. Other tessellations of $\mathbb{H}$ can be defined by other pre-geodesic conditions associated to flat surfaces, as has been done, for example, by J. Smillie and B. Weiss [43] using the condition in Example 4.3: each tile of their tessellation is associated to a saddle connection E of the flat surface, and consists of point in $\mathbb{H}$ for which E has the shortest length among all saddle connections.

The following observation is a result of the fact that the Delaunay triangulation(s) of a surface depends only on its metric structure, and not on the marking.

Proposition 4.22. *$\Sigma(X, q, Z)$ is $\Gamma(X, q, Z)$-invariant.*

Here is a consequence of Corollary 4.12.

Theorem 4.23. *If $\Gamma(X, q, Z)$ is a lattice in $(\mathrm{P})\mathrm{SL}_2(\mathbb{R})$, then it contains a finite-index subgroup having a fundamental domain composed of (finitely many) tiles of $\Sigma(X, q, Z)$.*

Proof. Let $\Gamma = \Gamma(X, q, Z)$. Then $\mathbb{H}/\Gamma$ has finite area, and therefore finitely many cusps. Each of these has a neighborhood that only intersects tiles of $\Sigma(X, q, Z)$; after removing these tiles, the remainder is compact. Because the boundaries of the tiles of $\Sigma(X, q, Z)$ consist of finitely many geodesic segments, no sequence of tiles can accumulate on the boundary of any fixed tile. Therefore, each tile has an open neighborhood that does not completely contain any other tile. These neighborhoods cover the compact portion of the fundamental domain, and therefore finitely many of them cover it. □

When the holonomy field (see [25]) of a flat surface is known, we can also obtain some elementary number-theoretic restrictions on the corresponding tessellation.

Proposition 4.24. *Let $\mathbb{K} \subseteq \mathbb{R}$ be the holonomy field of (X, q, Z). Then the cusps of $\Sigma(X, q, Z)$ lie in $\mathbb{K} \cup \{\infty\}$. Moreover, each geodesic supporting an edge of a tile in $\Sigma(X, q, Z)$ has an equation with coefficients in $\mathbb{K}$, and therefore an endpoint of any such geodesic, as well as the cotangent of any angle in a tile of $\Sigma(X, q, Z)$ lies in at most a quadratic extension of $\mathbb{K}$.*

Proof. A cusp of $\Sigma(X, q, Z)$ corresponds to a periodic direction on (X, q, Z), whose (co-)slope must lie in $\mathbb{K} \cup \{\infty\}$ (cf. [18]). The endpoints of a geodesic supporting an edge of $\Sigma(X, q, Z)$ are solutions to a (linear or) quadratic equation with coefficients in $\mathbb{K}$. Likewise, the cotangent of the angle between two Poincaré geodesics in $\mathbb{H}$ may be obtained as a square root of a rational expression in the coefficients of the geodesics. □

On the necessity of (± 1)-holonomy for finite area of tiles

An example of a *homothety surface* shows that several of the results of this section depend not only on the compactness of X, but also on the fact that q yields a well-defined area on X. Many of the definitions for homothety surfaces carry over naturally from the case of locally Euclidean surfaces, and so we omit them.

Let $h > 1$, and let X be the torus that is the quotient of $\mathbb{C} - \{0\}$ by the homothety $z \mapsto hz$, with the induced homothety structure $\mathfrak{h}$. Homotheties form a normal subgroup of all affine maps, and so in general the space of homothety surfaces also admits an action by $\mathrm{GL}_2(\mathbb{R})$. In the case of $(X, \mathfrak{h})$, this action is trivial: all elements of $\mathrm{GL}_2(\mathbb{R})$ commute with the transition map $z \mapsto hz$. Thus we can say that the "Veech group" of $(X, \mathfrak{h})$ is all of $\mathrm{GL}_2(\mathbb{R})$. We mark four points of X, $Z = \{(\pm 1, 0), (0, \pm 1)\}$ and triangulate $(X, \mathfrak{h}, Z)$ as in Figure 4.4. This

triangulation is Delaunay in the sense we have given, as angles are scale-invariant.

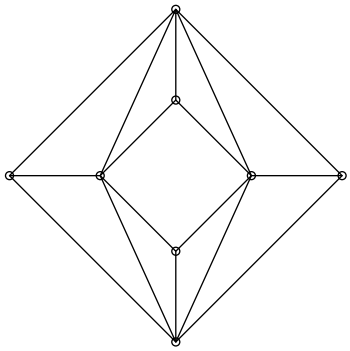

Figure 4.4: A homothety structure on the torus with a corresponding affine triangulation; the inner boundary is glued to the outer by a central homothety with scaling factor $h > 1$.

There are only two distinct conditions that arise from the requirement that edges remain Delaunay:

- each diagonal of a trapezoid already has weight 0 and remains so when the trapezoids are stretched in the directions of slope ± 1, thus leading to the condition $|z| \geq 1$;
- each edge of slope ± 1 yields the condition $|z| \leq h$;
- the remaining edges are diagonals of non-convex quadrilaterals, hence they do not create any constraints.

The region described therefore has infinite area.

Several people I have spoken with have said they thought about this surface at one point or another. The following behaviors of "geodesics" (linear trajectories) are interesting:

- *every* direction has *two* closed trajectories;

- any trajectory that is not periodic *accumulates* on the parallel closed trajectories, one in forward time and the other in backward time.

This also explains the failure of Veech's argument for the iso-Delaunay region to have finite area—it depends on a result of Vorobets [52] that at most countably many directions can have periodic trajectories.

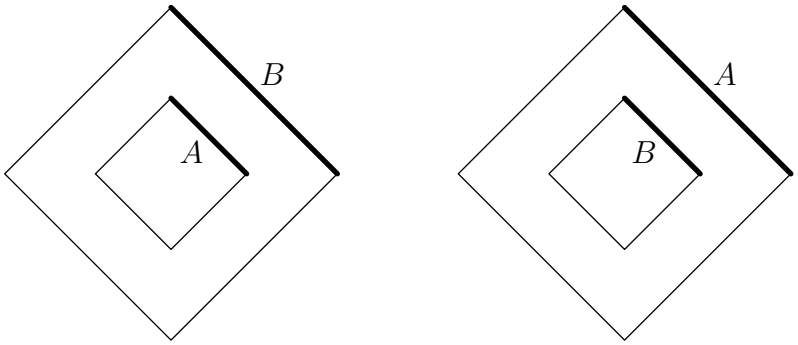

Figure 4.5: A genus 2 surface formed from two copies of the surface in Figure 4.4.

Another useful property of this surface is that it can be used to construct homothety surfaces whose affine group contains a reducible mapping class *that is not contained in the affine group of any translation or flat surface.* Start with two copies of the torus described above. Slit each of them and sew them together as shown in Figure 4.5. Then applying the linear map $\left(\begin{smallmatrix} 1 & 0 \\ 0 & h \end{smallmatrix}\right)$ induces an affine automorphism which is a composition of Dehn twists around disjoint curves, some direct, some indirect. This latter property is what rules out the possibility of this homeomorphism being contained in the affine group of a flat surface: when such elements are the product of Dehn twists around disjoint curves, they must all be in the same direction.

Part II

Complex structures

CHAPTER 5

STRATA OF $\mathrm{GR}_{\mathbb{C}}(\mathbb{C} \otimes V)$ AND COMPLEX STRUCTURES ON V

The Grassmannian variety of a complex vector space admits a natural stratification when the vector space is equipped with a real structure. In this chapter we describe these strata.

Given a function $f : X \to Y$, $\operatorname{gr} f$ denotes its graph in the space $X \times Y$.

5.1 A lemma on direct sum splittings

Lemma 5.1. *Let V be a vector space. Given A and B in $\mathrm{GL}(V)$, set $V' = \ker(A - B)$ and $V'' = \ker(A + B)$. If $A^2 = B^2$, then V' and V'' are invariant subspaces for both A and B. If $A^{-1}B = B^{-1}A$, then $V = V' \oplus V''$.*

Proof. Suppose $A^2 = B^2$. Then because B is nonsingular, we have

$$\ker(A + B)A = \ker B(B + A) = \ker(A + B),$$

and therefore V'' is invariant under A. By similar arguments, V'' is invariant under B, and V' is invariant under both A and B.

Now suppose $A^{-1}B = B^{-1}A$. To split V into the sum $V' \oplus V''$, we want to find a projection $P : V \to V'$ such that $\mathrm{id} - P$ is a projection $V \to V''$. That is, we want to solve the system

$$\begin{cases} (A - B)P = 0 \\ (A + B)(\mathrm{id} - P) = 0 \\ P^2 = P \end{cases}$$

for P. The first equation yields $AP = BP$. Substituting into the second equation, we obtain $A + B - 2AP = 0$, from which

$$P = \frac{1}{2}(\mathrm{id} + A^{-1}B) \qquad \text{and} \qquad \mathrm{id} - P = \frac{1}{2}(\mathrm{id} - A^{-1}B).$$

The third equation is then satisfied because $(A^{-1}B)^2 = \mathrm{id}$. □

Example 5.2. If A is any involution on V, then its (± 1)-eigenspaces sum to all of V. Any isomorphism G from V to its real dual space $V^\top$ induces an involution on $\mathrm{Hom}(V)$ by $A \mapsto G^{-1}A^\top G$; for example, if G gives rise to an inner product, then this is the transpose or adjoint operator, and Lemma 5.1 becomes the statement that every matrix can be uniquely written as the sum of a symmetric matrix and an anti-symmetric matrix.

Example 5.3. If J_1 and J_2 are complex structures on V that commute, then the subspaces defined by the equations $J_1 = J_2$ and $J_1 = -J_2$ are complementary complex vector spaces with respect to both J_1 and J_2. If J is any complex structure on V, then two commuting complex structures on $\mathrm{Hom}(V)$ are given by pre-composition and post-composition (right and left multiplication, respectively) by J. The subspace of $\mathrm{Hom}(V)$ on which these coincide is the space of complex-linear maps, and the subspace on which they differ is the space of complex-antilinear maps (with respect to J in both cases).

5.2 Complex conjugation

A complex vector space is a real vector space $\mathfrak{V}$ with a linear map $\mathfrak{J} : \mathfrak{V} \to \mathfrak{V}$ such that $\mathfrak{J}^2 = -\mathrm{id}$, or in other words a real vector space with a faithful action by $\mathbb{C}$ that distributes over vector addition. The real vector space $\mathbb{C}$ has a canonical

complex structure, which we shall denote m_i, and the canonical real structure $x + iy \mapsto x - iy$, which we shall denote $\bar{\cdot}$. In order to have a real structure on the complex vector space $(\mathfrak{V}, \mathfrak{J})$, i.e., an involution conj : $\mathfrak{V} \to \mathfrak{V}$ that extends the action of $\mathbb{C}$ on $\mathfrak{V}$ to include complex conjugation, we must have a distinguished real subspace $V \subset \mathfrak{V}$ such that V is the $(+1)$-eigenspace of conj and $\mathfrak{J}V$ is the (-1)-eigenspace of conj. Clearly conj and V determine each other, and so there is no loss of generality in assuming that $(\mathfrak{V}, \mathfrak{J}, \text{conj})$ is simply $(\mathbb{C} \otimes_{\mathbb{R}} V, m_i, \bar{\cdot})$, where the operations in the latter triple act on the first coordinate of $\mathbb{C} \otimes_{\mathbb{R}} V$. Hereafter we write $\mathbb{C} \otimes V = \mathbb{C} \otimes_{\mathbb{R}} V$.

5.3 Intersections of subspaces of $\mathbb{C} \otimes V$

Let V be a real n-dimensional vector space, and let $\nu : V \to \mathbb{C} \otimes V$ be the canonical inclusion $\nu(v) = 1 \otimes v$. We denote by $\mathcal{G} = \mathrm{Gr}_{\mathbb{C}}(\mathbb{C} \otimes V)$ the Grassmannian variety of complex subspaces of $\mathbb{C} \otimes V$ and by $\mathcal{G}_p = \mathrm{Gr}_{\mathbb{C}}(p, \mathbb{C} \otimes V)$ the connected component of p-dimensional subspaces. Observe that, if W is a complex subspace of $\mathbb{C} \otimes V$, then so is $\overline{W} = \{\overline{x} \mid x \in W\}$.

Lemma 5.4. *If $W \in \mathcal{G}$, then $\dim_{\mathbb{C}}(W \cap \overline{W}) = \dim_{\mathbb{R}}(W \cap \nu(V))$.*

Proof. The subspace $W' = \{x + \overline{x} \mid x \in W \cap \overline{W}\}$ is fixed pointwise by complex conjugation. Moreover, $iW' = \{x - \overline{x} \mid x \in W \cap \overline{W}\}$ because $i(x + \overline{x}) = ix - \overline{ix}$ and $W \cap \overline{W}$ is invariant under m_i. Therefore $W \cap \overline{W} = W' \oplus iW'$ as real vector spaces, and $W' = W \cap \nu(V)$. Because $\dim_{\mathbb{R}} W' = \dim_{\mathbb{R}} iW'$, we conclude both are equal to $\dim_{\mathbb{C}}(W \cap \overline{W})$. □

For $0 \leq p \leq n$, we define $\delta_p : \mathcal{G}_p \to \mathbb{Z}$ by $\delta_p(W) = \dim_{\mathbb{C}}(W \cap \overline{W})$ and we set

$\mathcal{G}_p^q = \delta_p^{-1}(q)$. Note that, in order for $\mathcal{G}_p^q$ to be non-empty, we must have $q \leq p$ and $q \geq \max\{0, 2p - n\}$. The $\mathcal{G}_p^q$ are the natural strata of $\mathcal{G}_p$.

Example 5.5. Suppose $V = \mathbb{R}^2$. Then $\mathbb{C} \otimes V = \mathbb{C}^2$ canonically. The Grassmannian component $\mathcal{G}_1$ is simply the complex projective line $\mathbb{CP}^1$, i.e., the Riemann sphere, consisting of complex lines in $\mathbb{C}^2$ parametrized by their slopes. The stratum $\mathcal{G}_1^1$ is the real projective line $\mathbb{RP}^1 \subset \mathbb{CP}^1$, i.e., the extended real axis, consisting of those complex lines in $\mathbb{C}^2$ that are complexifications of lines in $\mathbb{R}^2$. The remaining non-trivial stratum $\mathcal{G}_1^0(\mathbb{R}^2)$ has two components: the upper half-plane and the lower half-plane in $\mathbb{C}$.

The following theorem is the goal of this section:

Theorem 5.6. *Each non-empty stratum $\mathcal{G}_p^q$ is a smooth submanifold of $\mathcal{G}$ with real dimension $q(n-q) + 2(n-p)(p-q)$. The closure of $\mathcal{G}_p^q$ in $\mathcal{G}_p$ is $\bigcup_{q' \geq q} \mathcal{G}_p^{q'}$.*

In the following chapter, we will specialize to the case of strata that form a boundary to the Siegel half-plane. This special case has already been introduced by Friedland and Freitas [17] (see also Freitas's thesis), but much of the structure of the strata can be obtained from this more general context. For example, it is immediate that the action of $\mathrm{GL}(V)$ on $\mathcal{C}(V)$ extends continuously to an action on the boundary, as is the following result.

Lemma 5.7. *The natural action of $\mathrm{GL}(V)$ on $\mathcal{G}$ preserves each stratum $\mathcal{G}_p^q$. The action is transitive when restricted to a stratum.*

Proof. The action of $\bar{\cdot}$ commutes with the action of $\mathrm{GL}(V)$, because they act on different factors of $\mathbb{C} \otimes V$. This implies the first claim.

To show the second claim, let $W_1, W_2 \in \mathcal{G}_p^q$, and let $W_1' = W_1 \cap \nu(V)$, $W_2' = W_2 \cap \nu(V)$. Choose $\mathbb{R}$-bases B_1' and B_2' for W_1' and W_2', respectively, and complete these to $\mathbb{C}$-bases B_1 and B_2 of W_1 and W_2. Each B_i may be chosen so that the real and imaginary parts of its elements form an $\mathbb{R}$-linearly independent set: this follows from

$$2(p-q) + q = 2p - q \leq n$$

by the assumption that $\mathcal{G}_p^q$ is non-empty. We can therefore choose an element of $\mathrm{GL}(V)$ that sends B_1 to B_2, by acting on the real and imaginary parts. □

Remark 5.8. Lemma 5.7 says that each stratum is an orbit of the reductive group $\mathrm{GL}(V)$ acting on a complex Grassmannian variety. This means that several of the results in this section could be obtained by a "Matsuki-type" correspondence; I thank Allen Knutson for pointing this out to me. We will continue to give concrete proofs tailored to the situation at hand, however.

Lemma 5.9. *Let $\mathcal{G}_p^q$ be a nonempty stratum. Then the product $\mathcal{G}_q^q \times \mathcal{G}_{p-q}^0$ is a rank $q(p-q)$ complex affine bundle over $\mathcal{G}_p^q$ via the projection $(E, F) \mapsto E + F$.*

Proof. Let $(E, F) \in \mathcal{G}_q^q \times \mathcal{G}_{p-q}^0$. Our first claim is that $E \cap F = \{0\}$. Indeed, if $\alpha \otimes v \in F$, then $\overline{\alpha}(\alpha \otimes v) = |\alpha|^2 \otimes v$ is also in F; this element is fixed by conjugation on $\mathbb{C} \otimes V$, hence also contained in $\overline{F}$, which implies $\alpha \otimes v = 0$. Therefore $E + F$ is an element of $\mathcal{G}_p$. Moreover, it lies in $\mathcal{G}_p^q$ because $(E+F) \cap (\overline{E+F}) = E$. Conversely, if $W \in \mathcal{G}_p^q$, then $W \cap \overline{W}$ is a point in $\mathcal{G}_q^q$, and any complement to $W \cap \overline{W}$ in W is a point in $\mathcal{G}_{p-q}^0$. Therefore the map is surjective.

Given $W \in \mathcal{G}_p^q$, $W \cap \overline{W}$ is the unique maximal subspace of W fixed by conjugation. Therefore the fiber over W consists of all pairs $(W \cap \overline{W}, F)$, where F is a complement to $W \cap \overline{W}$ in W. Once we choose a particular point $(W \cap \overline{W}, F_0)$ in

the fiber, the whole fiber is parametrized by $\mathrm{Hom}_{\mathbb{C}}(F_0, W \cap \overline{W})$, with the parameterization $A \mapsto (W \cap \overline{W}, \mathrm{gr}\, A)$.

To get a local vector bundle structure, observe that a smooth section $s : U \to \mathcal{G}^0_{p-q}$ defined on a small open set U in $\mathcal{G}^q_p$ induces a section (perhaps on a smaller open set) $W \mapsto (W \cap \overline{W}, s(W))$. The fiber over W is then identified with $\mathrm{Hom}_{\mathbb{C}}(s(W), W \cap \overline{W})$. □

From this result (and its proof), we can compute the tangent space to each stratum:

Lemma 5.10. *Let $W \in \mathcal{G}^q_p$, where $\mathcal{G}^q_p$ is a nonempty stratum. Set $W' = \nu(V) \cap W$. Then*

$$T_W \mathcal{G}^q_p \cong \mathrm{Hom}_{\mathbb{R}}(W', V/W') \oplus \mathrm{Hom}_{\mathbb{C}}(W/(W \cap \overline{W}), (\mathbb{C} \otimes V)/W).$$

Proof. Recall that any choice of a complement E to W in $\mathbb{C} \otimes V$ yields an identification $T_W \mathcal{G}_p \cong \mathrm{Hom}(W, E) \cong \mathrm{Hom}_{\mathbb{C}}(W, (\mathbb{C} \otimes V)/W)$ (see for example [51, §10.1]); this is the basis of our analysis.

Write $W = W_1 \oplus W_2$, where $W_1 = \mathbb{C} \otimes W' = W \cap \overline{W}$, and W_2 is a complement to W_1 in W. The tangent space to $\mathcal{G}^q_q \times \mathcal{G}^0_{p-q}$ at (W_1, W_2) is $\mathrm{Hom}_{\mathbb{R}}(W', V/W') \oplus \mathrm{Hom}_{\mathbb{C}}(W_2, (\mathbb{C} \otimes V)/W_2)$. But $W_2 \cong W/W_1$, and the tangent space to the fiber over W at (W_1, W_2) is $\mathrm{Hom}_{\mathbb{C}}(W_2, W_1)$; the quotient of $\mathrm{Hom}_{\mathbb{C}}(W_2, (\mathbb{C} \otimes V)/W_2)$ by $\mathrm{Hom}_{\mathbb{C}}(W_2, W_1)$ is $\mathrm{Hom}_{\mathbb{C}}(W_2, (\mathbb{C} \otimes V)/W)$. The result follows. □

The identification in Lemma 5.9 can in fact be shown to be canonical (i.e., independent of all choices in the proof), but we will not need this fact.

Proof of Theorem 5.6. The (real) dimension of $\mathcal{G}_p^q$ follows from Lemma 5.9. It is embedded as a smooth submanifold of $\mathcal{G}_p$. because it is an orbit of $\mathrm{GL}(V)$. Its closure is shown to be as stated by letting subspaces of various dimensions approach their conjugates. This also shows that δ_p is an upper semicontinuous function. $\square$

5.4 The manifold of complex structures on a real vector space

Now let V be a real $2n$-dimensional vector space. Define the inclusion $\nu : V \to \mathbb{C} \otimes V$ and the strata $\mathcal{G}_p^q$ as before. Our focus will be on the component $\mathcal{G}_n$ of $\mathrm{Gr}_{\mathbb{C}}(\mathbb{C} \otimes V)$. Set $\mathcal{U}(V) = \mathcal{G}_n^0$, the open stratum of $\mathcal{G}_n$. When $K \in \mathcal{U}(V)$, we have the canonical splitting $\mathbb{C} \otimes V = K \oplus \overline{K}$, and so the tangent space to $\mathcal{G}_n$ at K may be identified with $\mathrm{Hom}_{\mathbb{C}}(K, \overline{K})$. We recall that the exponential map $\mathrm{Hom}_{\mathbb{C}}(K, \overline{K}) \to \mathcal{G}_n$ is given by $A \mapsto \mathrm{gr}\, A$. We will see that $\mathcal{U}(V)$ is naturally isomorphic to the space of complex structures on V, and that the remaining strata in $\mathcal{G}_n$ form a natural boundary to the space of complex structures; the points of this boundary have strong geometric meaning.

The *manifold of complex structures on* V is the subvariety $\mathcal{C}(V)$ of $\mathrm{Hom}(V)$ defined by

$$\mathcal{C}(V) = \{J \in \mathrm{Hom}(V) \mid J^2 = -\mathrm{id}\}.$$

It is a smooth manifold because $\mathrm{GL}(V)$ acts transitively on it by conjugation. Each point $J \in \mathcal{C}(V)$ splits $\mathrm{Hom}(V)$ (see Example 5.3) into a sum of J-linear and

J-antilinear maps—which we denote by

$$\mathrm{Hom}_J(V) = \{A \in \mathrm{Hom}(V) \mid AJ = JA\} \quad \text{and}$$
$$\mathrm{Hom}_{\overline{J}}(V) = \{A \in \mathrm{Hom}(V) \mid AJ = -JA\}$$

—and makes $\mathrm{Hom}(V)$ itself into a complex vector space, with the complex structure given by post-composition by J, which we denote $\mathscr{L}_J$; $\mathrm{Hom}_J(V)$ and $\mathrm{Hom}_{\overline{J}}(V)$ are complex subspaces. We have $\dim_{\mathbb{C}} \mathrm{Hom}_J(V) = \dim_{\mathbb{C}} \mathrm{Hom}_{\overline{J}}(V) = n^2$. Differentiating the condition $J^2 + \mathrm{id} = 0$ shows that $\mathrm{Hom}_{\overline{J}}(V)$ is the tangent space to $\mathcal{C}(V)$ at J. Because the complex structures $\mathscr{L}_J$ vary smoothly with J, they endow $\mathcal{C}(V)$ with an almost-complex structure. The following result shows that this almost-complex structure on $\mathcal{C}(V)$ is integrable.

Theorem 5.11. *$\mathcal{C}(V)$ and $\mathcal{U}(V)$ are isomorphic as almost-complex manifolds.*

The basic correspondence between points in these two manifolds is often used in Hodge-theoretic situations (see for example [20, App. A.4] or [51]). This result in itself is therefore not new, but it seems to have remained "folklore" knowledge (cf. [4]). Consequently, we will sketch the connection between $\mathcal{C}(V)$ and $\mathcal{U}(V)$, then more carefully extract some results that will be useful in later chapters.

Given $J \in \mathcal{C}(V)$, let $k(J)$ be the real $2n$-dimensional subspace

$$k(J) = \{1 \otimes Jv - i \otimes v \mid v \in V\} \subset \mathbb{C} \otimes V.$$

From the equality $\overline{k(J)} = k(-J) = \{1 \otimes Jv + i \otimes v \mid v \in V\}$ we conclude $k(J) \cap \overline{k(J)} = \{0\}$. That $k(J)$ and $\overline{k(J)}$ are complex subspaces of $\mathbb{C} \otimes V$ can be checked directly, or by an appeal to Lemma 5.1. (Note that $k(J)$ and $\overline{k(J)}$ are also defined by the respective equations $ix = -Jx$ and $ix = Jx$.) Thus we have a function $k : \mathcal{C}(V) \to \mathcal{U}(V)$.

To obtain an inverse map, we must associate to each $K \in \mathcal{U}(V)$ a unique complex structure $j(K)$ on V. If $K \in \mathcal{U}(V)$, then $\nu(V)$ is transverse to both K and $\overline{K}$, and the (real) dimension of all three is equal. Hence the projection of $\nu(V)$ onto either component of the sum $K \oplus \overline{K}$ is a linear isomorphism. We denote by $\kappa : V \to \overline{K}$ the composition of ν and the projection $\mathbb{C} \otimes V \to \overline{K}$. Then

$$j(K) = \kappa^{-1} \circ m_i \circ \kappa$$

is a complex structure on V. This defines a function $j : \mathcal{U}(V) \to \mathcal{C}(V)$.

The action of $\mathrm{GL}(V)$ on $\mathcal{C}(V)$, as previously observed, is by conjugation. The action on $\mathcal{U}(V)$ is the canonical one induced on $\mathcal{G}_n$ by the action on the second factor of $\mathbb{C} \otimes V$. These actions are conjugated by j and k since, for $A \in \mathrm{GL}(V)$, we have

$$k(AJA^{-1}) = \{1 \otimes AJA^{-1}v - i \otimes v \mid v \in V\} = \{1 \otimes AJv - i \otimes Av \mid v \in V\} = A(k(J)).$$

As could be expected at this point, we have:

Lemma 5.12. *The maps j and k are inverse $\mathrm{GL}(V)$-equivariant bijections.*

We also have natural interpretations of complex-linear and complex-antilinear maps of V in terms of maps between subspaces of $\mathbb{C} \otimes V$, which shed light on the how the tangent spaces to $\mathcal{C}(V)$ and $\mathcal{U}(V)$ relate to each other.

Lemma 5.13. *Let $J \in \mathcal{C}(V)$, and set $K = k(J)$. Then we have the following equalities:*

$$\mathrm{Hom}_{\mathbb{C}}(K) = \mathrm{Hom}_{\mathbb{C}}(\overline{K}) = \mathrm{Hom}_J(V)$$

$$\mathrm{Hom}_{\mathbb{C}}(K, \overline{K}) = \mathrm{Hom}_{\mathbb{C}}(\overline{K}, K) = \mathrm{Hom}_{\overline{J}}(V)$$

when elements of the latter sets, acting on $\mathbb{C} \otimes V$, are restricted to the appropriate subspaces.

Proof. Let $A \in \mathrm{Hom}_J(V)$. Then K and $\overline{K}$ are invariant subspaces of A, because

$$A(1 \otimes Jv \pm i \otimes v) = 1 \otimes AJv \pm i \otimes Av = 1 \otimes J(Av) \pm i \otimes Av \qquad \text{for all } v \in V.$$

Similarly, if $A \in \mathrm{Hom}_{\overline{J}}(V)$, then A sends K to $\overline{K}$ (as well as $\overline{K}$ to K). Thus we have the inclusions $\mathrm{Hom}_J(V) \subset \mathrm{Hom}_{\mathbb{C}}(K)$ and $\mathrm{Hom}_{\overline{J}}(V) \subset \mathrm{Hom}_{\mathbb{C}}(K, \overline{K})$. A dimension count shows that the inclusions must be surjections, which proves the equalities. □

Lemma 5.14. *If λ is an eigenvalue of $A \in \mathrm{Hom}_{\overline{J}}(V)$, then $-\lambda$ is also an eigenvalue of A, and J interchanges the corresponding eigenspaces.*

Proof. Let $A \in \mathrm{Hom}_{\overline{J}}(V)$, and suppose $v \in \mathbb{C} \otimes V$ satisfies $Av = \lambda v$. Then $AJv = -JAv = -\lambda Jv$, from which the result follows. □

By composing k with the canonical charts at points of $\mathcal{U}(V)$, we get canonical charts on $\mathcal{C}(V)$. The next lemma gives explicit formulas for these charts and their inverses, as well as domains on which they are defined.

Lemma 5.15. *Let $J_0 \in \mathcal{C}(V)$, and set $K_0 = k(J_0)$. If $J \in \mathcal{C}(V)$ is such $J_0 J$ does not have 1 as an eigenvalue, then $k(J)$ is the graph in $K_0 \oplus \overline{K_0} = \mathbb{C} \otimes V$ of*

$$A = (J_0 - J)(J_0 + J)^{-1} \quad \in \mathrm{Hom}_{\overline{J_0}}(V) = \mathrm{Hom}_{\mathbb{C}}(K_0, \overline{K_0}).$$

Conversely, if $A \in \mathrm{Hom}_{\overline{J_0}}(V)$ does not have 1 as an eigenvalue, then

$$j(\mathrm{gr}(A)) = J_0(\mathrm{id} - A)(\mathrm{id} + A)^{-1} \quad \in \mathcal{C}(V).$$

Proof. By definition, $k(J) = \{1 \otimes Jv - i \otimes v \mid v \in V\}$. Since $k(J) \subset K_0 \oplus \overline{K_0}$, for every $v \in V$ there exist unique v' and v'' in V such that

$$1 \otimes Jv - i \otimes v = (1 \otimes J_0 v' - i \otimes v') + (1 \otimes J_0 v'' + i \otimes v'').$$

This leads to the system of equations

$$\begin{cases} Jv = J_0(v' + v'') \\ v = v' - v'' \end{cases}.$$

Solving for v'' in terms of v', we find

$$v'' = (J_0 - J)(J_0 + J)^{-1}v',$$

which proves the first result.

If $A \in \mathrm{Hom}_{\overline{J_0}}(V)$ does not have 1 as an eigenvalue, then by Lemma 5.14 neither does it have -1 as an eigenvalue, and therefore $\mathrm{id} + A$ is invertible. The second result now follows from the first by solving the first equation for J. □

Example 5.16. Suppose $V = \mathbb{R}^2$ and J_0 is the standard complex structure $\left[\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right]$. Each point $\left[\begin{smallmatrix} a & -(a^2+1)/b \\ b & -a \end{smallmatrix}\right]$ in $\mathcal{C}(V)$ is identified with $a + bi \in \mathbb{C}$, and $T_{J_0}\mathcal{C}(V)$ is the space of constant Beltrami forms $\alpha\,\mathrm{d}\overline{z}/\mathrm{d}z$, $\alpha \in \mathbb{C}$. The formulas in Lemma 5.15 become the following Möbius transformations:

$$z \mapsto (i - z)/(i + z) \qquad \text{and} \qquad w \mapsto i(1 - w)/(1 + w),$$

which exchange the upper half-plane $\mathbb{H}$ with $\mathbb{D}$. The only point J in $\mathcal{U}(V)$ such that $J_0 J$ has 1 as an eigenvalue is $-J_0$, corresponding to $-i$, which is sent to ∞ by the first map above.

We record here the condition for a map from a domain $U \subset \mathbb{C}$ into $\mathcal{C}(V)$ to be holomorphic. If z is a coordinate on U, let D' and D'', respectively, denote differentiation with respect to $\mathrm{Re}\, z$ and $\mathrm{Im}\, z$. Suppose $J(z) : \zeta \mapsto J_\zeta$ is a family of complex structures parameterized by $\zeta \in U$. Then the Cauchy–Riemann condition for $J(z)$ to be holomorphic is

$$(D' + J_\zeta D'')J(z) = 0. \tag{5.1}$$

5.5 Strata of $\mathcal{G}_n(V)$ in local coordinates

At this point it is clear that the dimension of the $(+1)$-eigenspace of an element of $\mathrm{Hom}_{\overline{J}}(V)$ is important, and we suspect that it relates to the strata of $\mathcal{G}_n(V)$. The following lemma makes this relationship precise.

Lemma 5.17. *Let $J \in \mathcal{C}(V)$ and set $K = k(J)$. If $A \in \mathrm{Hom}_{\overline{J}}(V) = \mathrm{Hom}_{\mathbb{C}}(K, \overline{K})$, then*

$$\delta_n(\mathrm{gr}(A)) = \dim_{\mathbb{R}} \ker(\mathrm{id} - A) = \dim_{\mathbb{R}} \ker(\mathrm{id} + A).$$

Proof. For $1 \otimes Jv - 1 \otimes v \in k(J)$ and $A \in \mathrm{Hom}_{\overline{J}}(A)$, we have

$$(1 \otimes Jv - 1 \otimes v) + A(1 \otimes Jv - 1 \otimes v) = 1 \otimes J(\mathrm{id} - A)v - i \otimes (\mathrm{id} + A)v$$

This element of $\mathrm{gr}\, A$ is fixed by conjugation if $v \in \ker(\mathrm{id} + A)$ and in the (-1)-eigenspace of conjugation if $v \in \ker(\mathrm{id} - A)$. From this observation the result follows. □

The boundary of $\mathcal{C}(V)$ is therefore sent by the chart centered at J_0 to the affine subvariety of $\mathrm{Hom}_{\overline{J_0}}(V)$ defined by any of the following:

$$\det(\mathrm{id} - A) = 0, \qquad \det(\mathrm{id} + A) = 0, \qquad \det(\mathrm{id} - A^2) = 0.$$

The principle stratum $\mathcal{G}_n^1$ of $\partial\mathcal{C}(V)$ is a dense subset. More generally, as a corollary to Theorem 5.6, we have

Corollary 5.18. *The dimension of $\mathcal{G}_n^p \subset \partial\mathcal{C}(V)$ is $2n^2 - p^2$.*

Given $J_0 \in \mathcal{C}(V)$, the eigenvectors of $-JJ_0 = J^{-1}J_0$ for J in the domain of the chart given in Lemma 5.15 coincide with those of the image element, and the

eigenvalues are related by a naturally arising fractional linear transformation: if v is an eigenvector of $A = (J_0 - J)(J_0 + J)^{-1}$ with corresponding eigenvalue λ, then

$$-JJ_0v = (\mathrm{id} - A)^{-1}(\mathrm{id} + A)v = (\mathrm{id} - A)^{-1}(1 + \lambda)v = \frac{1 + \lambda}{1 - \lambda}v. \qquad (5.2)$$

This relation will be particularly useful in studying local versions of the Siegel half-plane in the next chapter.

CHAPTER 6

THE GEOMETRY OF $\mathfrak{H}$ AND ITS BOUNDARY

In this chapter, we introduce a symplectic structure on a vector space and place the associated Siegel half-plane in the context of the previous chapter, thereby recovering much of its classical geometry, while also refining some of the results and proofs.

6.1 Linear maps from a vector space to its dual spaces

Let V be a finite-dimensional real vector space. We denote by $V^\top = \mathrm{Hom}(V, \mathbb{R})$ the real dual of V. If W is another real vector space, then any linear map $A : V \to W$ has a dual map, called its *transpose*, $A^\top : W^\top \to V^\top$, defined by $A^\top \alpha = \alpha A$. Recall that V is canonically isomorphic to its double dual space $V^{\top\top}$ by sending $v \in V$ to the evaluation map $ev_v : \alpha \mapsto \alpha v$. Hence the transpose of a map $V \to V^\top$ is again a map $V \to V^\top$.

Any linear map $B : V \to V^\top$ induces a bilinear form b on V, defined by $b(v, w) = (Bw)v$, which is non-degenerate precisely when B is an isomorphism. We call $B : V \to V^\top$ *symmetric* if $B^\top = B$ and *anti-symmetric* if $B^\top = -B$. A *pseudo-Euclidean structure on* V is a symmetric linear isomorphism $G : V \to V^\top$, and a *symplectic structure on* V is an anti-symmetric linear isomorphism $\Sigma : V \to V^\top$.

A linear map $B : V \to V^\top$ is *positive semi-definite*, written $B \geq 0$, if $(Bv)v \geq 0$ for all $v \in V$. B is *positive definite*, written $B > 0$, if $(Bv)v > 0$ for all $v \neq 0$; in particular, such a map is an isomorphism. A *Euclidean structure* is a positive definite pseudo-Euclidean structure.

Given a pseudo-Euclidean structure G, a linear map $A : V \to V$ is called *self-adjoint* or *symmetric with respect to* G if $(GA)^\top = GA$. The bilinear form induced on V by a symplectic structure is called a *symplectic form*, and the bilinear form induced by a (pseudo-)Euclidean structure is called a (*pseudo-*)*inner product*.

An element of $\mathrm{Hom}(V)$ is said to be *diagonalizable* if its eigenspaces span V. For clarity and later reference, we restate the spectral theorem for finite-dimensional vector spaces.

Theorem 6.1 (Spectral theorem)**.** *A linear map $V \to V$ is diagonalizable if and only if it is symmetric with respect to some Euclidean structure on V.*

Proof. The "if" part is the usual spectral theorem. The "only if" part is by construction: suppose $A \in \mathrm{Hom}(V)$ is diagonalizable. Choose an eigenbasis for A, and let G be a Euclidean structure for which this basis is orthogonal. Then $(GA)^\top = GA$. □

The set $\mathrm{Hom}(V, \mathbb{C})$ of real linear maps $V \to \mathbb{C}$ is equipped with a canonical complex structure, which is post-composition by m_i. Now suppose that V is equipped with a complex structure $J : V \to V$, $J^2 = -\mathrm{id}$. The *conjugate space of* V, denoted $\overline{V}$, is V equipped with the complex structure $-J$. We are interested in the complex dual space of V, which we denote by $V^* = \mathrm{Hom}_{\mathbb{C}}(V, \mathbb{C})$, and also in the *dual conjugate* and *conjugate dual* spaces of V, respectively denoted $(\overline{V})^*$ and $\overline{(V^*)}$.

The spaces $(\overline{V})^*$ and $\overline{(V^*)}$ are canonically isomorphic; both consist of complex-antilinear maps $V \to \mathbb{C}$, and there is an explicit complex-linear isomorphism between the two, given by $\alpha \mapsto \overline{\alpha}$. The evaluation map $v \mapsto ev_v$ again yields

canonical isomorphisms of V with its double complex dual, double dual conjugate, and double conjugate dual spaces.

The transpose operator is defined as before and sends a linear map $A : V \to W$ to the linear map $A^\top : W^* \to V^*$; it enjoys the same linearity properties as A (i.e., if A is complex-linear, then so is $A^\top$, and likewise for complex-antilinearity). In addition, we now also have a *conjugate transpose* operator, which is defined by $A^*\alpha = \overline{\alpha}A$. The conjugate transpose of A may be viewed in a number of ways, but most usefully as a map $\overline{(W^*)} \to (\overline{V})^*$, in which case it again enjoys the same linearity properties as A. Hence the conjugate transpose of a map $V \to (\overline{V})^*$ is again a map $V \to (\overline{V})^*$.

A *pseudo-Hermitian structure on* V is a complex-linear isomorphism $H : V \to (\overline{V})^*$ such that $H^* = H$. This is equivalent to the conditions that $\operatorname{Re} H$ be a pseudo-Euclidean structure on V and that $\operatorname{Im} H$ be a symplectic structure on V. A linear map $A : V \to V$ is called *self-adjoint* or *Hermitian* if $(HA)^* = HA$. A *Hermitian structure* is a pseudo-Hermitian structure whose real part is a Euclidean structure.

Example 6.2. The complex manifold $\mathcal{C}(V)$ carries a canonical pseudo-Hermitian structure $\mathcal{H}$ on its tangent bundle, defined on each tangent space $T_J\mathcal{C}(V) = \operatorname{Hom}_{\overline{J}}(V)$ by

$$(A, B) \mapsto (\mathcal{H}_J B)A = \operatorname{tr} BA - i \operatorname{tr} BJA = \operatorname{tr} AB + i \operatorname{tr} AJB.$$

The symmetry of $\operatorname{Re} \mathcal{H}_J$ is a standard property of the trace. The anti-symmetry of $\operatorname{Im} \mathcal{H}_J$ arises from the definition of $\operatorname{Hom}_{\overline{J}}(V)$ as the space of J-antilinear maps. $\mathcal{H}$ is clearly invariant under the action of $\mathrm{GL}(V)$ by conjugation, since, on the level of tangent spaces to $\mathcal{C}(V)$, the action of $\mathrm{GL}(V)$ is the adjoint action, which preserves traces.

We will soon see subspaces of $\mathrm{Hom}_{\overline{J}}(V)$ on which $\mathcal{H}_J$ is a Hermitian structure.

6.2 Compatible complex structures

Three important subspaces of $\mathrm{Hom}(V)$ play starring roles in what follows. Suppose G is a pseudo-Euclidean structure, Σ is a symplectic structure, and J is a complex structure on V. Then we have the classical orthogonal, symplectic, and complex-linear groups

$$\begin{aligned}
\mathrm{O}_G(V) &= \{A \in \mathrm{GL}(V) \mid A^\top G A = G\},\\
\mathrm{Sp}_\Sigma(V) &= \{A \in \mathrm{GL}(V) \mid A^\top \Sigma A = \Sigma\},\\
\mathrm{GL}_J(V) &= \{A \in \mathrm{GL}(V) \mid AJ = JA\}.
\end{aligned}$$

For the sake of simplicity, we will drop the subscripts of the first two when clarity permits. The Lie algebras of these groups are

$$\begin{aligned}
\mathfrak{so}(V) &= \{A \in \mathrm{Hom}(V) \mid (GA)^\top = -GA\},\\
\mathfrak{sp}(V) &= \{A \in \mathrm{Hom}(V) \mid (\Sigma A)^\top = \Sigma A\},\\
\mathrm{Hom}_J(V) &= \{A \in \mathrm{Hom}(V) \mid AJ = JA\}.
\end{aligned}$$

Each of these has a natural complement in $\mathrm{Hom}(V)$. We have already seen that $\mathrm{Hom}_{\overline{J}}(V)$ complements $\mathrm{Hom}_J(V)$ (see previous chapter). By Example 5.2, a complement to $\mathfrak{sp}(V)$ is $\{A \in \mathrm{Hom}(V) \mid (\Sigma A)^\top = -\Sigma A\}$, and a complement to $\mathfrak{so}(V)$ is $\{A \in \mathrm{Hom}(V) \mid (GA)^\top = GA\}$. This latter is precisely the space of maps on V that are symmetric with respect to G.

The interesting equation to study is $G = \Sigma J$; this is (almost) the *compatibility condition*. Note that when two of G, Σ, J are given, there is at most one possibility

for the remaining one that is an object of the appropriate type. Note also that when G, Σ, and J satisfy this relation, we obtain the fourth kind of classical group, a *unitary group*, as the intersection of any two of the above groups:

$$\mathrm{U}(V) = \mathrm{O}(V) \cap \mathrm{Sp}(V) = \mathrm{O}(V) \cap \mathrm{GL}_J(V) = \mathrm{Sp}(V) \cap \mathrm{GL}_J(V);$$

The corresponding pseudo-Hermitian structure on V is $H = G + i\Sigma$. The Lie algebra of $\mathrm{U}(V)$ is

$$\mathfrak{u}(V) = \mathfrak{so}(V) \cap \mathfrak{sp}(V) = \mathfrak{so}(V) \cap \mathrm{Hom}_J(V) = \mathfrak{sp}(V) \cap \mathrm{Hom}_J(V).$$

One wonders, given a fixed Σ, for which $J \in \mathcal{C}(V)$ is ΣJ a pseudo-Euclidean structure? The answer is precisely the ones that are themselves symplectic. Using the above notation, this statement becomes:

$$\forall\ J \in \mathcal{C}(V), \quad J \in \mathfrak{sp}(V) \iff J \in \mathrm{Sp}(V).$$

We write $\mathcal{C}_\Sigma(V) = \mathcal{C}(V) \cap \mathrm{Sp}(V)$. The tangent space to $\mathrm{Sp}(V)$ at $J \in \mathcal{C}_\Sigma(V)$ is the space of operators $A : V \to V$ that are symmetric with respect to ΣJ, i.e,

$$T_J\mathrm{Sp}(V) = \{A \in \mathrm{Hom}(V) \mid (\Sigma J A)^\top = \Sigma J A\},$$

which shows, firstly, that the dimension of $\mathrm{Sp}(V)$ is $2n^2 + n$, and, secondly, that the normal space to $\mathrm{Sp}(V)$ at J can be identified with $\mathfrak{so}_{\Sigma J}(V)$. The tangent space to $\mathcal{C}_\Sigma(V)$ at J is of course $\mathrm{Hom}_{\overline{J}}(V) \cap \mathfrak{sp}(V)$; it is moreover a J-invariant subspace of $T_J\mathcal{C}(V)$, which shows that $\mathcal{C}_\Sigma(V)$ itself is a complex manifold.

The connected components of $\mathcal{C}_\Sigma(V)$ are indexed by the signature of the quadratic form associated to ΣJ. We say that $J \in \mathcal{C}_\Sigma(V)$ is *compatible* with Σ if ΣJ is a Euclidean structure. The set of all complex structures compatible with Σ forms the *Siegel half-plane* $\mathfrak{H}$:

$$\mathfrak{H} = \{J \in \mathcal{C}_\Sigma(V) \mid \Sigma J > 0\}.$$

(For a discussion of compatible complex structures, see for example [11, Part V].) The action of $\mathrm{GL}(V)$ on $\mathcal{C}(V)$ restricts to an action of $\mathrm{Sp}(V)$ on $\mathcal{C}_\Sigma(V)$; the stabilizer of each point is the corresponding unitary group. This action preserves the connected components of $\mathcal{C}_\Sigma(V)$ (because $\mathrm{Sp}(V)$ is connected), hence in particular $\mathrm{Sp}(V)$ acts on $\mathfrak{H}$.

Proposition 6.3. *If $J \in \mathfrak{H}$, then the pseudo-Hermitian structure $\mathcal{H}_J$ defined in Example 6.2 restricts to a Hermitian structure $T_J\mathfrak{H} \subset T_J\mathcal{C}(V)$.*

Proof. We only need to show that $(\mathcal{H}_J A)A = \operatorname{tr} A^2 > 0$ for all $A \neq 0$. This follows from the spectral theorem, because A is symmetric with respect to ΣJ, and therefore all of its eigenvalues are real. □

Therefore $\mathfrak{H}$ carries a canonical $\mathrm{Sp}(V)$-invariant Riemannian metric, the *Siegel metric*. In §6.4, we will again have occasion to apply the spectral theorem when we describe the geodesics for this metric (and a family of related metrics). First, however, we examine equivalent descriptions of $\mathfrak{H}$, including its image by the logarithmic map and how a choice of coordinates leads to the description by complex matrices.

6.3 Isotropic subspaces and $\Lambda(\mathbb{C} \otimes V)$

Let V be a real vector space of dimension $2n$, and fix a symplectic structure Σ on V. A subspace W of V is *isotropic* if $(\Sigma w)v = 0$ for all $v, w \in W$. Given $A \in \mathrm{Sp}(V)$, it is well-known that for each eigenvalue λ of A, the corresponding eigenspace E_λ is isotropic if $\lambda \neq 1$, and more generally its symplectic complement is the sum of all eigenspaces $E_{\lambda'}$ where λ' is an eigenvalue of A and $\lambda\lambda' \neq 1$. An isotropic subspace

$W \subset V$ is *Lagrangian* if $\dim W = n$. The set of Lagrangian subspaces of V forms the *Lagrangian Grassmannian* $\Lambda(V) \subset \mathrm{Gr}_{\mathbb{R}}(n, V)$. The symplectic structure on V extends to a symplectic structure on $\mathbb{C} \otimes V$, and so the same definitions apply in the complexified case, as well.

Let $\mathcal{U}(V) \subset \mathcal{G}_n$ be as in the previous chapter, along with the functions $j : \mathcal{U}(V) \to \mathcal{C}(V)$ and $k : \mathcal{C}(V) \to \mathcal{U}(V)$.

Proposition 6.4. *The image $\mathcal{U}_\Sigma(V)$ of $\mathcal{C}_\Sigma(V)$ in $\mathcal{G}_n$ by k is $\mathcal{U}(V) \cap \Lambda(\mathbb{C} \otimes V)$.*

Proof. Let $J \in \mathcal{C}(V)$. For any $v, w \in V$, we have

$$(\Sigma(1 \otimes Jw - i \otimes w))(1 \otimes Jv - i \otimes v) = (\Sigma Jw)Jv - (\Sigma w)v - i((\Sigma Jw)v + (\Sigma w)Jv)$$

The real part of this expression vanishes for all $v, w \in V$ if and only if $J \in \mathrm{Sp}(V)$. The imaginary part vanishes for all $v, w \in V$ if and only if $J \in \mathfrak{sp}(V)$. Because the conditions $J \in \mathrm{Sp}(V)$ and $J \in \mathfrak{sp}(V)$ are equivalent, the result is shown. □

Proposition 6.5. *Let $J \in \mathcal{C}_\Sigma(V)$, and let $L \subset V$ be any subspace. Then L is isotropic if and only if $L \perp_{\Sigma J} JL$. In particular, if $L \in \Lambda(V)$, then V splits into the orthogonal sum $L \oplus JL$.*

Proof. The first claim follows immediately from the equality $(\Sigma w)v = ((\Sigma J)w)Jv$. The second claim follows from a dimension count. □

We write $\Lambda_p = \Lambda(\mathbb{C} \otimes V) \cap \mathcal{G}_n^p$. The smallest stratum $\Lambda_n \cong \Lambda(V)$ is known to be the Shilov boundary of $\mathfrak{H}$ (that is, it is the smallest set $\mathfrak{S}$ of boundary points such that every harmonic function on $\mathfrak{H}$ attains a maximum on $\mathfrak{H} \cup \mathfrak{S}$). Recently [17], it was shown that $\overline{\mathfrak{H}} \subset \Lambda(\mathbb{C} \otimes V)$ is the Busemann 1-compactification of $\mathfrak{H}$. We can now compute the dimensions of all strata (cf. Corollary 5.18).

Theorem 6.6. *Λ_p is a smooth real manifold of dimension $n^2 + n - (p^2 + p)/2$.*

Proof. This simply requires lightly modifying the proofs of Lemmata 5.9 and 5.10, applying the symmetry of elements of the tangent space with respect to a Euclidean structure. □

6.4 Geodesics in $\mathfrak{H}$

Throughout this section, whenever λ is an eigenvalue of a linear operator, we use E_λ to denote to corresponding eigenspace.

Lemma 6.7. *If $A, B \in \mathrm{Hom}(V, V^\top)$ are both positive definite, then all of the eigenvalues of $A^{-1}B$ are positive.*

Proof. Suppose λ is an eigenvalue of $A^{-1}B$ with corresponding eigenvector $v \neq 0$. Then $Bv = \lambda Av$. Thus we have $(Bv)v = \lambda(Av)v$, and because $A > 0$ and $B > 0$, also $\lambda > 0$. □

Theorem 6.8. *Let J_0 and J_1 be in $\mathfrak{H}$. Then*

1. *$-J_1J_0$ is diagonalizable, and all of its eigenvalues are positive.*
2. *For any eigenvalue $\lambda \neq 1$, J_0 and J_1 interchange E_λ and $E_{1/\lambda}$. If 1 is an eigenvalue, then E_1 is invariant under both J_0 and J_1.*
3. *$(-J_1J_0)^t$ is in $\mathrm{Sp}(V)$ for all $t \in \mathbb{R}$.*

Proof. First observe that $-J_1J_0 = -J_1\Sigma^{-1}\Sigma J_0 = (\Sigma J_1)^{-1}(\Sigma J_0)$. The first claim now follows from the spectral theorem and Lemma 6.7.

An eigenspace E_λ is the kernel of $-J_1J_0-\lambda\cdot\mathrm{id}$. This map factors as $J_1(\lambda J_1-J_0)$, and because J_1 is non-singular, E_λ is also the kernel of $\lambda J_1 - J_0$. Suppose $v \in E_\lambda$. Then $(J_1 - \lambda J_0)J_0 v = -J_1(\lambda J_1 - J_0)v = 0$, and therefore J_1 maps E_λ to $E_{1/\lambda}$. Because $-J_0J_1$ has the same eigenspaces as $-J_1J_0$, the same argument shows that J_1 maps E_λ to $E_{1/\lambda}$. In particular, if 1 is an eigenvalue, then E_1 is invariant under J_0 and J_1.

By part (1), $(-J_1J_0)^t$ is defined on each E_λ by $w \mapsto \lambda^t w$. We need to show that this map is symplectic. Suppose $\lambda, \lambda' \in \mathscr{E}$, and $v \in E_\lambda, w \in E_{\lambda'}$. Then

$$\big(\Sigma(-J_1J_0)^t w\big)(-J_1J_0)^t v = \Sigma({\lambda'}^t w)\lambda v = (\lambda\lambda')^t(\Sigma w)v.$$

If $\lambda\lambda' \neq 1$, part (2) shows that both sides of this equality are zero. If $\lambda\lambda' = 1$, then the equality shows that Σ is preserved on $E_\lambda \oplus E_{1/\lambda}$ (or on E_1, if $\lambda = \lambda' = 1$). Because the E_λs sum to V, this shows that $(-J_1J_0)^t$ is symplectic for all $t \in \mathbb{R}$. $\square$

We obtain as a corollary the following extremely well-known properties:

Corollary 6.9. $\mathrm{Sp}(V)$ *acts transitively on* $\mathfrak{H}$ *by conjugation.* $\mathfrak{H}$ *is path-connected.*

Proof. Let $J_0, J_1 \in \mathfrak{H}$. Part (3) of Theorem 6.8 implies in particular that $\sqrt{-J_1J_0} \in \mathrm{Sp}(V)$. We will show that $\sqrt{-J_1J_0}J_0\sqrt{-J_0J_1} = J_1$. The inverse of $\sqrt{-J_1J_0}$ is $\sqrt{-J_0J_1}$, because taking inverses of linear transformations commutes with taking square roots (when both exist). It suffices to show that J_2 equals $\sqrt{-J_1J_0}J_0\sqrt{-J_0J_1}$ on $E_\lambda \oplus E_{1/\lambda}$ for each $\lambda \neq 1$, since $J_1 = J_0$ on E_1 if 1 is an eigenvalue. On $E_\lambda \oplus E_{1/\lambda}$, J_1 restricts to $(1/\lambda)J_0 \oplus \lambda J_0$, $\sqrt{-J_0J_1}$ restricts to $\lambda^{-1/2}\mathrm{id} \oplus \lambda^{1/2}\mathrm{id}$, and $\sqrt{-J_1J_0}$ restricts to $\lambda^{1/2}\mathrm{id} \oplus \lambda^{-1/2}\mathrm{id}$. J_0 interchanges E_λ and $E_{1/\lambda}$. Therefore J_1 equals the composition of $\sqrt{-J_0J_1}$, J_0, and $\sqrt{-J_1J_0}$. (See Figure 6.1.)

It follows that, given $J_0, J_1 \in \mathfrak{H}$, $J_t = (-J_1J_0)^{t/2}J_0(-J_0J_1)^{t/2}$ is a path from J_0 to J_1. □

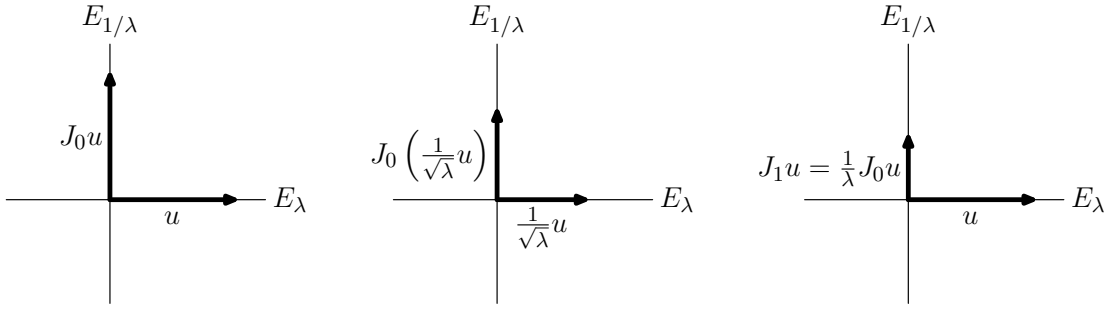

Figure 6.1: A geodesic motion in $\mathfrak{H}$.

Theorem 6.8 also suggests a family of metrics on $\mathfrak{H}$: given $J_1, J_2 \in \mathfrak{H}$, let $\mathcal{E}$ be the set of eigenvalues of $-J_2J_1$. Then define

$$d_\alpha(J_1, J_2) = \frac{1}{2}\Big(\sum_{\lambda\in\mathcal{E}} \dim E_\lambda \,|\log\lambda|^\alpha\Big)^{1/\alpha} \quad (1 \le \alpha < \infty).$$

(Recall that E_λ denotes the eigenspace corresponding to the eigenvalue λ.)

Proposition 6.10. *For every $1 \le \alpha < \infty$, d_α is a metric on $\mathfrak{H}$.*

As observed by Freitas–Friedland, given $J_0, J_1 \in \mathfrak{H}$, the path $t \mapsto (-J_1J_0)^{t/2}J_0(-J_0J_1)^{t/2}$ is geodesic for every d_α, and it is the unique geodesic from J_0 to J_1 if $1 < \alpha < \infty$. The Siegel metric is d_2.

Lemma 6.11. *Let $J_0, J_1 \in \mathfrak{H}$, and let m be the multiplicity of 1 as an eigenvalue of $-J_1J_0$. Set $J_t = (-J_1J_0)^{t/2}J_0(-J_0J_1)^{-t/2}$ and $K_{\pm\infty} = \lim_{t\to\pm\infty} k(J_t) \in \mathcal{G}$. Then $\delta(K_{\pm\infty}) = n - m/2$.*

Proof. Let E_0 be the subspace on which $-J_1J_0$ restricts to the identity, and let E_+ (resp. E_-) be the sum of the eigenspaces with corresponding eigenvalues greater

(resp. less) than 1. Set $F_0 \subset \mathbb{C} \otimes V$ to be the image of E_0 under the isomorphism $V \to k(J_0)$. As $t \to \pm\infty$, $k(J_t)$ limits to $F_0 \oplus (\mathbb{C} \otimes E_\pm)$. Because $\dim_{\mathbb{R}} E_\pm = (2n - m)/2$, the result follows. □

Corollary 6.12. *The endpoints of a geodesic in $\mathfrak{H}$ for any metric d_α, $\alpha > 1$, lie in the same stratum Λ_p of $\partial\mathfrak{H}$.*

6.5 Local coordinates on $\mathfrak{H}$

We have already shown that $\mathfrak{H}$ is a complex manifold. As is well-known, it is in fact biholomorphic to a bounded, contractible, open region in a complex vector space. We present this as follows:

Proposition 6.13. *Let $J_0 \in \mathfrak{H}$. Then the canonical chart k_0 on $\mathcal{C}(V)$ at J_0 includes all of $\mathfrak{H}$ in its domain and sends $\mathfrak{H}$ to an open, contractible, bounded domain of $T_{J_0}\mathfrak{H}$.*

Proof. There are several pieces to prove.

First, given any other $J \in \mathfrak{H}$, all eigenvalues of $J_0 J$ are negative by Theorem 6.8, hence in particular $J_0 J$ does not have 1 as an eigenvalue. Therefore, by Lemma 5.15, all of $\mathfrak{H}$ lies in the domain of the canonical chart at J_0.

Secondly, we show that any symplectic J maps to $\mathrm{Hom}_{\overline{J_0}}(V) \cap \mathfrak{sp}(V)$ under k_0. That is, we want to determine the condition on $A \in \mathrm{Hom}_{\overline{J_0}}(V)$ such that ${k_0}^{-1}(A)$ is symplectic. Because ${k_0}^{-1}(A)$ is a complex structure, we have

$$J_0(\mathrm{id} - A)(\mathrm{id} + A)^{-1} = (\mathrm{id} - A)^{-1}(\mathrm{id} + A)J_0.$$

(Note that power series in A commute.) The requirement $(\Sigma {k_0}^{-1}(A))^\top = \Sigma {k_0}^{-1}(A)$ is equivalent to each of the following:

$$(\mathrm{id} + A^\top)^{-1}(\mathrm{id} - A^\top)\Sigma J_0 = \Sigma(\mathrm{id} - A)^{-1}(\mathrm{id} + A)J_0,$$
$$(\mathrm{id} - A^\top)\Sigma(\mathrm{id} - A) = (\mathrm{id} + A^\top)\Sigma(\mathrm{id} + A),$$
$$-A^\top\Sigma - \Sigma A = A^\top\Sigma + \Sigma A,$$
$$(\Sigma A)^\top = \Sigma A.$$

Thirdly, $k_0(\mathfrak{H})$ is open because it is a component of the complement in $\mathrm{Hom}_{\overline{J_0},\Sigma}(V)$ of the zero set of $\det(\mathrm{id} - A)$.

Fourthly, we show that $k_0(\mathfrak{H})$ is bounded. Each eigenvalue of $A = k_0(J)$, $J \in \mathfrak{H}$, is related to an eigenvalue of $-JJ_0$ by Equation (5.2) on page 50. Because all the eigenvalues of $-JJ_0$ are negative, the eigenvalues of A must lie between -1 and 1. This gives a bounded condition on A (since its eigenvalues are real).

Finally, we show that $k_0(\mathfrak{H})$ is contractible. Since all the eigenvalues of $A \in k_0(\mathfrak{H})$ have absolute value less than 1, the same holds for tA, $t \leq 1$. Therefore $k_0(\mathfrak{H})$ deformation retracts onto the origin. □

It is worthwhile to consider how to recover the more standard definition of $\mathfrak{H}$ as the set of symmetric $n \times n$ complex matrices with positive definite imaginary part. This classical representation of $\mathfrak{H}$ has many advantages, among them the simplicity of its definition, but it fails to include within its scope the entire boundary of $\mathfrak{H}$ that we wish to study. It will be useful to have an easy set of conditions to check, however, for the application we make in the next chapter, and more generally the rest of this section can serve as reference for those who wish to know how the different representations are related.

Choose $L \in \Lambda(V)$ and a basepoint $J_0 \in \mathfrak{H}$. Then, by Lemma 6.5, $V = L \oplus J_0 L$, and therefore any other $J \in \mathcal{C}(V)$ can be expressed in block-matrix form as follows:

$$\begin{pmatrix} J_{11} & J_{12} \\ J_{21} & J_{22} \end{pmatrix} : L \oplus J_0 L \to L \oplus J_0 L.$$

The equation $J^2 = -\mathrm{id}$ translates to the conditions

$$\begin{cases} J_{11}^2 + J_{12} J_{21} = -\mathrm{id}_L \\ J_{22}^2 + J_{21} J_{12} = -\mathrm{id}_{J_0 L} \\ J_{11} J_{12} + J_{12} J_{22} = 0 \\ J_{21} J_{11} + J_{22} J_{21} = 0 \end{cases} . \tag{6.1}$$

Now we determine the conditions to ensure $J \in \mathfrak{H}$. Lemma 6.5 implies that $J_{21} = \mathrm{proj}_{J_0 L} J$ is an isomorphism from L to $J_0 L$, hence invertible. In this case, the system (6.1) is equivalent to

$$J_{12} = -(J_{11}^2 + \mathrm{id}_L) J_{21}^{-1} \qquad \text{and} \qquad J_{22} = -J_{21} J_{11} J_{21}^{-1}. \tag{6.2}$$

Σ itself can be written as $\left(\begin{smallmatrix} 0 & \Sigma_{12} \\ -\Sigma_{12}^\top & 0 \end{smallmatrix}\right)$, where $\Sigma_{12} \in \mathrm{Hom}(J_0 L, L^\top)$ is an isomorphism, because L and $J_0 L$ are both in $\Lambda(V)$ (again by Lemma 6.5). For J to be in $\mathcal{C}_\Sigma(V)$, ΣJ must be symmetric, which translates to

$$\begin{cases} (\Sigma_{12} J_{21})^\top = \Sigma_{12} J_{21} \\ \Sigma_{12}^\top J_{12} = J_{12}^\top \Sigma_{12} \\ \Sigma_{12} J_{22} = -J_{11}^\top \Sigma_{12} \end{cases} . \tag{6.3}$$

Combining the second equation in (6.2) with the first equation in (6.3), the final equation in (6.3) becomes

$$(\Sigma_{12} J_{21} J_{11})^\top = \Sigma_{12} J_{21} J_{11}. \tag{6.4}$$

Lastly, we need $\Sigma J > 0$. Clearly we must have $\Sigma_{12}J_{21} > 0$, because this is the restriction of ΣJ to L. But this condition is also sufficient: if $u \oplus v \in L \oplus J_0 L$, then by setting $u' = J_{21}^{-1}v$, we can reduce the computation of $(\Sigma J(u \oplus v))(u \oplus v)$ to a computation in L. We get

$$\begin{aligned}
&(\Sigma_{21}J_{21}u)u - (\Sigma_{12}v)J_{12}v + (\Sigma_{12}J_{22}v)u - (\Sigma_{12}v)J_{11}u \\
&= (\Sigma_{12}J_{21}u)u + (\Sigma_{12}J_{21}u')({J_{11}}^2 + \mathrm{id}_{L_0})u' - (\Sigma_{12}J_{21}J_{11}u')u - (\Sigma_{12}J_{21}u')J_{11}u \\
&= (\Sigma_{12}J_{21}u)u + (\Sigma_{12}J_{21}J_{11}u')J_{11}u' + (\Sigma_{12}J_{21}u')u' - 2(\Sigma_{12}J_{21}J_{11}u')u \\
&= (\Sigma_{12}J_{21}u')u' + (\Sigma_{12}J_{21}(u - J_{11}u'))(u - J_{11}u').
\end{aligned}$$

Both of the terms in this final sum are non-negative. If $v \neq 0$, then the first term is positive, and if $v = 0$ but $u \neq 0$, the second term is positive. Thus we obtain the "bilinear relations":

Proposition 6.14. *If* $J = \left(\begin{smallmatrix} J_{11} & J_{12} \\ J_{21} & J_{22} \end{smallmatrix}\right)$ *is any element of* $\mathrm{Hom}(V)$*, then necessary and sufficient conditions to have* $J \in \mathfrak{H}$ *are* $\Sigma_{12}J_{21} > 0$*,* $(\Sigma_{12}J_{21})^\top = \Sigma_{12}J_{21}$*, and the equations of* (6.2) *and* (6.4)*.*

Given $J = \left(\begin{smallmatrix} J_{11} & J_{12} \\ J_{21} & J_{22} \end{smallmatrix}\right) \in \mathfrak{H}$, the functions $X = J_{11}J_{21}^{-1}$ and $Y = J_{21}^{-1}$ therefore both map $J_0 L \to L$ and satisfy the conditions

$$\Sigma_{12}Y^{-1} > 0, \qquad (\Sigma_{12}Y^{-1})^\top = \Sigma_{12}Y^{-1}, \qquad \text{and} \qquad (\Sigma_{12}X^{-1})^\top = \Sigma_{12}X^{-1}.$$

(These expressions have meaning because Σ_{12} is an isomorphism $J_0 L \to L^\top$. Note also that the latter equation is indeed equivalent to (6.4), because J_{11} is symmetric with respect to $\Sigma_{12}J_{21}$ if and only if J_{11}^{-1} is.) Conversely, given such maps $X, Y : J_0 L \to L$, the function

$$\begin{pmatrix} XY^{-1} & -(XY^{-1}X + Y) \\ Y^{-1} & -Y^{-1}X \end{pmatrix} \in \mathrm{Hom}(V)$$

lies in $\mathfrak{H}$. After choosing an orthonormal basis for L (with respect to ΣJ_0), and taking the image basis in $J_0 L$, the above expressions define a bijection between $\mathfrak{H}$ and the appropriate space of matrices, of the form $X + iY$. Under this correspondence, the action of $\mathrm{Sp}(V)$ on $\mathfrak{H}$ (by conjugation) becomes an action by "generalized fractional linear transformations": $\left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right) : Z \mapsto (AZ + B)(CZ + D)^{-1}$.

6.6 Torelli space and the period map

Given a compact Riemann surface X of genus $g \geq 1$, let σ be the symplectic form on $H_1(X, \mathbb{Z})$ given by algebraic intersection of curves. Let $\Omega(X) \cong \mathbb{C}^g$ denote the space of abelian differentials on X. Given $\omega \in \Omega(X)$ and a smooth simple closed curve γ, we call $\mathrm{per}_\gamma(\omega) = \int_\gamma \omega$ the period of ω along γ. The value of this integral depends only on the homology class $[\gamma] \in H_1(X, \mathbb{Z})$, and therefore the period map $\mathrm{per} : [\gamma] \mapsto \mathrm{per}_\gamma$ is a canonical injection $H_1(X, \mathbb{Z}) \to \Omega(X)^*$. The image of per in $\Omega(X)^*$, which we will simply denote $\mathrm{per}(X)$, is clearly a lattice, and so the quotient $\mathrm{Jac}(X) = \Omega(X)^*/\mathrm{per}(X)$ is a complex torus, called the Jacobian of X. $\mathrm{Jac}(X)$ is characterized by choosing a symplectic basis for $H_1(X, \mathbb{Z})$ and computing its period matrix, as described below.

We need to recall some general facts about dual bases. Suppose W is a complex vector space of dimension d, and $\mathscr{W} = \{w_1, \ldots, w_d\}$ is a basis for W. Then we get a canonical evaluation map $ev_{\mathscr{W}} : W^* \to \mathbb{C}^d$, whose jth component is ev_{w_j}. The dual basis to $\mathscr{W}$ is the basis $\mathscr{W}^* = \{{w_1}^*, \ldots, {w_d}^*\}$ for W^* such that ${w_j}^*(w_k) = \delta_{jk}$. (Note that the element ${w_j}^*$ depends on the entire basis $\mathscr{W}$, not just the element w_j.) If $\mathscr{W}' = \{{w_1}', \ldots, {w_d}'\}$ is any other basis for W, then $ev_{\mathscr{W}'}(\mathscr{W}^*) = [{w_j}^*({w_k}')]_{j,k=1}^d$ is the change-of-basis matrix that converts the coefficients of a vector in the basis

$\mathscr{W}'$ to its coefficients in the basis $\mathscr{W}$. Observe that $ev_{\mathscr{W}}((\mathscr{W}')^*) = ev_{\mathscr{W}'}(\mathscr{W}^*)^{-1}$.

With notation as in the previous paragraphs, choosing a basis for $\Omega(X)$ induces, by composition of $ev_{\mathscr{W}}$ with the period map, a linear map $\Pi_{\mathscr{W}} : H_1(X, \mathbb{Z}) \to \mathbb{C}^g$, as in the diagram below:

$$\begin{array}{ccc} H_1(X,\mathbb{Z}) & \xrightarrow{\text{per}} & \Omega(X)^* \\ & \searrow^{\Pi_{\mathscr{W}}} & \downarrow{\scriptstyle ev_{\mathscr{W}}} \\ & & \mathbb{C}^g \end{array}$$

Let $\mathscr{A}$ and $\mathscr{B}$ be subsets of $H_1(X, \mathbb{Z})$ of size g, each spanning a Lagrangian subspace in $H_1(X, \mathbb{R}) = \mathbb{R} \otimes_{\mathbb{Z}} H_1(X, \mathbb{Z})$, and together forming a symplectic basis for $H_1(X, \mathbb{Z})$ (for example, start with a canonical system of curves on X—say $\{\alpha_1, \ldots, \alpha_g\}$ and $\{\beta_1, \ldots, \beta_g\}$—and take the images of these in $H_1(X, \mathbb{Z})$). Note that $\text{per}(\mathscr{A})$ and $\text{per}(\mathscr{B})$ are both $\mathbb{C}$-bases for $\Omega(X)^*$; we identify $\mathscr{A}$ and $\mathscr{B}$ with their embedded images in $\Omega(X)^*$. The Riemann period matrix of X with respect to the basis $\mathscr{A} \cup \mathscr{B}$ is defined to be $\Pi = \Pi_{\mathscr{A}^*}(\mathscr{B})$. If we write out what this means, we find

$$\Pi_{jk} = \int_{\beta_j} {\alpha_k}^*,$$

in accord with the usual definition of the period matrix. More generally, given any (symplectic) basis $\mathscr{W}$ of $\Omega(X)$, the $g \times 2g$ matrix $\begin{bmatrix} \Pi_{\mathscr{W}}(\mathscr{A}) & \Pi_{\mathscr{W}}(\mathscr{B}) \end{bmatrix}$ is called a full period matrix for X; this is the matrix of $\Pi_{\mathscr{W}}$ with respect to $\mathscr{A} \cup \mathscr{B}$ and $\mathscr{W}$. To obtain the standard Riemann matrix from a full period matrix, compute $\Pi = \big(\Pi_{\mathscr{W}}(\mathscr{A})^{-1}\big)\Pi_{\mathscr{W}}(\mathscr{B})$.

6.7 Example: the KFT family

As an example of applying the results of this chapter to a family of Riemann surfaces, we consider the KFT family of [39]. This is the family of genus 3 sur-

faces whose automorphism groups contain the symmetric group on four letters—including Klein's quartic curve, the quartic Fermat surface, and the tetrahedron (that is, the six lines determined by four points in projective space), hence the name. Rodríguez and Gonzáles-Aguilera compute the period matrices of these surfaces as having the form

$$Z = \tau Z_0, \qquad \text{where} \qquad Z_0 = \begin{pmatrix} 3 & -1 & -1 \\ -1 & 3 & -1 \\ -1 & -1 & 3 \end{pmatrix} \quad \text{and} \quad \tau \in \mathbb{H}$$

These are precisely the matrices stabilized by the image of a certain faithful representation $S_4 \to \mathrm{Sp}_6(\mathbb{R})$ (see [39] for details). By work of Silhol in [41], this family of surfaces is a Teichmüller curve generated by a quadratic differential—in fact, an origami, which exhibits the S_4 symmetry of the family as its group of isometries.

If $Z = \tau Z_0$, then $X = xZ_0$ and $Y = yZ_0$, where $\tau = x + iy$. Let I_3 be the 3×3 identity matrix. Then the complex structure in $\mathfrak{H}_3$ corresponding to Z is

$$J_\tau = \frac{1}{y} \begin{pmatrix} xI_3 & -(x^2+y^2)\, Z_0 \\ {Z_0}^{-1} & -xI_3 \end{pmatrix}$$

Taking the product of two of these, we get

$$J_{\tau_2} J_{\tau_1} = \frac{1}{y_1 y_2} \begin{pmatrix} (x_1 x_2 - {x_2}^2 - {y_2}^2) I_3 & [x_1({x_2}^2 + {y_2}^2) - x_2({x_1}^2 + {y_1}^2)] Z_0 \\ (x_1 - x_2) {Z_0}^{-1} & (x_1 x_2 - {x_1}^2 - {y_1}^2) I_3 \end{pmatrix}$$

The determinant of $-J_{\tau_2} J_{\tau_1} - \lambda \cdot \mathrm{id}$ is

$$\left(\lambda^2 - \frac{1}{y_1 y_2} \big((x_1 - x_2)^2 + {y_1}^2 + {y_2}^2 \big) \lambda + 1 \right)^3$$

the discriminant of whose inner quadratic polynomial is always non-negative. Therefore to move between two distinct points of KFT along a geodesic in $\mathfrak{H}_3$ requires uniform motion in three two-dimensional symplectic subspaces.

CHAPTER 7

ODD COHOMOLOGY

7.1 Orientation covers of generic quadratic differentials

Let X be a fixed surface of genus $g \geq 2$, and let $\pi : \tilde{X} \to X$ be a degree 2 covering branched over $4g-4$ points. The genus of $\tilde{X}$ is then $\tilde{g} = 4g-3$. The sheet exchange $\tau : \tilde{X} \to \tilde{X}$ induces an involution on the first cohomology group of $\tilde{X}$, splitting it into a $(+1)$-eigenspace $H^1(\tilde{X}, \mathbb{R})^+$ and a (-1)-eigenspace $H^1(\tilde{X}, \mathbb{R})^-$. We call elements of $H^1(\tilde{X}, \mathbb{R})^+$ *even cohomology classes* and elements of $H^1(\tilde{X}, \mathbb{R})^-$ *odd cohomology classes*, both with respect to τ. If $\tilde{X}$ is given the structure of a differentiable manifold, then a representative form of an even (resp. odd) cohomology class is called an *even* (resp. *odd*) *form* on $\tilde{X}$. Each cohomology class in $H^1(X, \mathbb{R})$ pulls back by π to an element of $H^1(\tilde{X}, \mathbb{R})^+$, and likewise any even cohomology class on $\tilde{X}$ descends to a cohomology class on X. Hence $\dim_{\mathbb{R}} H^1(\tilde{X}, \mathbb{R})^+ = 2g$, which implies $\dim_{\mathbb{R}} H^1(\tilde{X}, \mathbb{R})^- = 6g - 6$.

When X is a Riemann surface and q is a quadratic differential on X with simple zeroes, then the above situation arises naturally by taking $\pi : \tilde{X} \to X$ to be the double cover of X branched at the zeroes of q, and endowing $\tilde{X}$ with the conformal structure that makes π conformal. Since $\pi \circ \tau = \pi$, τ is also conformal. (Locally over a zero of q, π looks like $z \mapsto z^2$ and τ looks like $z \mapsto -z$.) Then $\pi^* q$ is the square of an abelian differential $\omega_q \in \Omega(\tilde{X})$, with ω_q well-defined up to a choice of sign. In this case, ω_q is called a *square root* of q, and the pair $(\tilde{X}, \omega_q)$ is called the *orientation cover* of (X, q). If $\mathcal{Q}(X)$ denotes the space of quadratic differentials on X, then the map $\mathcal{Q}(X) \to \Omega(\tilde{X})$ defined by $q' \mapsto {\omega_q}^{-1} \pi^* q'$ is an injection whose image is $\Omega(\tilde{X})^-$, the space of odd abelian differentials on $\tilde{X}$ (which is isomorphic as

a real vector space to $H^1(\tilde{X}, \mathbb{R})^-$). As in the topological case considered previously, abelian differentials on X pullback via π to differentials in $\Omega(\tilde{X})^+$, the space of even abelian differentials on $\tilde{X}$.

If Y is any Riemann surface of genus $g(Y) \geq 1$ and $\{\omega_1, \ldots, \omega_{g(Y)}\}$ is a basis for $\Omega(Y)$, then $\{[\operatorname{Re}\omega_1], [\operatorname{Im}\omega_1], \ldots, [\operatorname{Re}\omega_{g(Y)}], [\operatorname{Im}\omega_{g(Y)}]\}$ is a basis for $H^1(Y, \mathbb{R})$. $\Omega(Y)$ is a complex vector space, and the corresponding complex structure on $H^1(Y, \mathbb{R})$ sends $[\operatorname{Re}\omega]$ to $[\operatorname{Im}\omega]$ for any $\omega \in \Omega(Y)$. Obviously, if $\Omega(Y)$ has a splitting into even and odd forms, then this way of finding a basis for $H^1(Y, \mathbb{R})$ preserves the splitting, and we obtain a basis of the even and odd cohomology of Y; in particular, $\omega \in \Omega(Y)^-$ if and only if $[\operatorname{Re}\omega] \in H^1(Y, \mathbb{R})^-$, or equivalently $[\operatorname{Im}\omega] \in H^1(Y, \mathbb{R})^-$.

7.2 The Thurston–Veech construction

Next we describe a special case of the Thurston–Veech [44, 49] (or "bouillabaisse") construction of quadratic differentials and an associated complex structure on the odd cohomology of a surface. Let X be a compact orientable surface of genus $g \geq 2$. The largest number of disjoint, pairwise non-homotopic curves on X is then $3g - 3$ (corresponding to "splitting the surface into pairs of pants"). We call such a collection a *maximal multicurve* on X. Take two multicurves $\mathscr{A}$ and $\mathscr{B}$ on X such that no element of $\mathscr{A}$ is homotopic to an element of $\mathscr{B}$. Assume, moreover, that the union of all the curves cuts X into simply-connected pieces. Encode the intersections of elements of $\mathscr{A}$ and $\mathscr{B}$ in a $(3g-3) \times (3g-3)$ matrix M, which can also be thought of as the matrix of a linear transformation $\mathbb{R}^{\mathscr{A}} \to \mathbb{R}^{\mathscr{B}}$.

For each intersection of an element in $\mathscr{A}$ and $\mathscr{B}$, we construct a metric rectangle

on X. Its dimensions are dictated by the entries of an eigenvector of $M^\top M : \mathbb{R}^{\mathscr{A}} \to \mathbb{R}^{\mathscr{A}}$, as follows: Topological considerations show that some power of $M^\top M$ has all positive entries, which means that it is a Perron–Frobenius matrix, and therefore its largest eigenvalue λ_{PF} has multiplicity 1 and a corresponding eigenvector v with all positive entries. $u = \lambda_{\mathrm{PF}}^{-1/2} M v$ is then an eigenvector of $MM^\top$ with corresponding eigenvalue λ_{PF}. A rectangle crossed by $\alpha \in \mathscr{A}$ and $\beta \in \mathscr{B}$ has its height given by the α component of v and its width given by the β component of u.

Actually, there is a real parameter of choice in the construction: we could take $u = e^t \lambda_{\mathrm{PF}}^{-1/2} M v$ for any real t, and we would then find $v = e^{-t} \lambda_{\mathrm{PF}}^{-1/2} M^\top u$. We thus obtain a family of Riemann surfaces X_t, each carrying a specific quadratic differential q_t.

7.3 A question about abelian varieties

J. Hubbard posed a question regarding these surfaces. Suppose the following:

1. all of the components of the complement of $\mathscr{A} \cup \mathscr{B}$ are either hexagons or rectangles;
2. the intersection matrix M is invertible.

The first condition implies that every q_t has only simple zeroes. So, as in the previous section, we take an orientation cover $(\tilde{X}_t, \omega_t)$ of each (X_t, q_t), chosen continuously. We can identify all these topologically with a branched double cover $\tilde{X} \to X$. The second condition implies that the lifts of the elements of $\mathscr{A} \cup \mathscr{B}$ to $\tilde{X}$ form a basis for $H_1(\tilde{X}, \mathbb{R})^-$. We even have the symplectic form on $H_1(\tilde{X}, \mathbb{R})$ in

the coordinates of $\tilde{\mathscr{A}}$ and $\tilde{\mathscr{B}}$:

$$\Sigma = \begin{bmatrix} 0 & -2M^\top \\ 2M & 0 \end{bmatrix}.$$

These lifts are *purely topological*, and hence we do not need to consider a subscript on either the homology classes or the symplectic form. The corresponding symplectic form on cohomology is $\Sigma^\top = -\Sigma$.

J. Hubbard posed a question dealing with a family of complex structures on the odd cohomology of $\tilde{X}_t$; the idea is to turn cohomology classes supported on $\tilde{\mathscr{A}}$ into classes supported on $\tilde{\mathscr{B}}$. We define, for all $t \in \mathbb{R}$, in the basis dual to $\tilde{\mathscr{A}} \cup \tilde{\mathscr{B}}$:

$$J_t = \begin{bmatrix} 0 & -(M^\top M)^{-\frac{1}{2}-\frac{t}{\log \lambda_{\mathrm{PF}}}} M^\top \\ M(M^\top M)^{-\frac{1}{2}+\frac{t}{\log \lambda_{\mathrm{PF}}}} & 0 \end{bmatrix}, \qquad t \in \mathbb{R}. \tag{7.1}$$

(It is important to note that $M^\top M$ is a positive definite symmetric matrix, so all of its eigenvalues are real and positive.) The question is: does this complex structure coincide with the "natural" complex structure described previously?

This question may be phrased in terms of analytic curves in the Siegel half-plane $\mathfrak{H}$. After choosing the data of a pair of multi-curves $\mathscr{A}$ and $\mathscr{B}$ on X, there is a free real parameter in the Thurston–Veech construction of flat surfaces, and varying this parameter produces a family of surfaces X_t that are related by the Teichmüller geodesic flow. The Jacobians of the corresponding double covers $\tilde{X}_t$ split into even and odd parts; each of these parts is again an abelian variety, polarized (but not principally so) by the restriction of the symplectic form on $H_1(\tilde{X}, \mathbb{R})$. The odd parts of the Jacobians of the $\tilde{X}_t$ lie in the Siegel upper half-plane $\mathfrak{H} \subset \mathcal{C}(H^1(\tilde{X}, \mathbb{R})^-)$. The complex structure proposed by Hubbard also extends to a one real-parameter family of complex structures. We shall show that all of these complex structures also trace out a curve $\mathfrak{H}$. The question is whether these two

curves coincide.

The main result of this chapter is as follows:

Theorem 7.1. *The family J_t extends to a holomorphically immersed maximal disk in $\mathfrak{H}$ that does not coincide with the disk arising from any Teichmüller disk having non-trivial Veech group. In the case that the Thurston–Veech construction produces an irreducible characteristic polynomial, J_t is obtained by acting independently on the subspace of $H^1(\tilde{X}, \mathbb{R})^-$ spanned by $\{[\operatorname{Re}\omega_q], [\operatorname{Im}\omega_q]\}$ and each of its Galois conjugates via a diagonal action.*

7.4 Complex structures on the odd cohomology of bouillabaisse surface covers

We again use $\mathscr{E}(M^\top M)$ to denote the set of eigenvalues of $M^\top M$, and λ_{PF} to denote its Perron–Frobenius eigenvalue.

Using the results of chapter 6, we can examine how the complex structures (7.1) relate to each other geometrically by looking at the following product:

$$-J_{t_2}J_{t_1} = \begin{bmatrix} (M^\top M)^{\frac{t_1-t_2}{\log\lambda_{\mathrm{PF}}}} & 0 \\ 0 & (M^\top M)^{\frac{t_2-t_1}{\log\lambda_{\mathrm{PF}}}} \end{bmatrix}.$$

The spectrum of this matrix is $\{\lambda^{\pm(t_1-t_2)/\lambda_{\mathrm{PF}}} \mid \lambda \in \mathscr{E}(M^\top M)\}$. The eigenspace corresponding to $\lambda^{\pm(t_1-t_2)/\lambda_{\mathrm{PF}}}$ is $E_{\lambda^{\pm 1}}$.

Lemma 7.2. *We have the following limits in $\mathrm{Gr}_{\mathbb{C},n}(H_1(\tilde{X}, \mathbb{C}))$:*

$$\lim_{t\to+\infty} k(J_t) = \mathbb{C}\otimes V_- \quad \textit{and} \quad \lim_{t\to-\infty} k(J_t) = \mathbb{C}\otimes V_+,$$

where V_+ is the sum of the E_λ with $\lambda > 1$ and V_- is the sum of the E_λ with $\lambda < 1$.

Proof. Two defining equations for $k(J_t)$ in $\mathbb{C} \otimes (\tilde{\mathscr{A}} \cup \tilde{\mathscr{B}})$ are

$$x = -i(M^\top M)^{-\frac{1}{2} - \frac{t}{\log \lambda_{\mathrm{PF}}}} M^\top y \qquad \text{and} \qquad y = iM(M^\top M)^{-\frac{1}{2} + \frac{t}{\log \lambda_{\mathrm{PF}}}} x,$$

where $x \in \mathbb{C} \otimes \tilde{\mathscr{A}}$ and $y \in \mathbb{C} \otimes \tilde{\mathscr{B}}$. We recall that $M^\top$ sends eigenvectors of $MM^\top$ to eigenvectors of $M^\top M$ with the same eigenvalues. Now by looking at the appropriate equation on each $E_\lambda \oplus E_{1/\lambda}$, we conclude the desired result. □

Let $\mathbb{B} = \{z \in \mathbb{C} \mid |\operatorname{Im} z| < \pi/2\}$; $\mathbb{B}$ is conformally equivalent to the open unit disk via the map $z \mapsto \tanh(z/2)$. We will construct an analytic family of complex structures on $H_1(\tilde{X}, \mathbb{R})^-$, varying with a parameter t in $\mathbb{B}$. To simplify notation, we introduce the function, for real t,

$$\Theta(t) = \frac{t}{\log \lambda_{\mathrm{PF}}} \log M^\top M.$$

This function yields a matrix whose leading eigenvalue is t; we will be most concerned with behaviors near $t = \pm\pi/2$.

Now we extend our family of complex structures to a complex disk. For each $t = t' + it'' \in \mathbb{B}$, define J_t in the basis $\tilde{\mathscr{A}} \cup \tilde{\mathscr{B}}$ by

$$J_{t'+it''} = \begin{bmatrix} -\tan \Theta(t'') & -\sec \Theta(t'')(M^\top M)^{-\frac{1}{2} - \frac{t'}{\log \lambda}} M^\top \\ M(M^\top M)^{-\frac{1}{2} + \frac{t'}{\log \lambda}} \sec \Theta(t'') & M(\tan \Theta(t''))M^{-1} \end{bmatrix} \tag{7.2}$$

In order to show that this is well-defined for all $t \in \mathbb{B}$, we need to show that the spectrum of $M^\top M$ is bounded below by $1/\lambda$. Recall that all of the eigenvalues of $M^\top M$ are positive.

Lemma 7.3.

1. *For all $t \in \mathbb{B}$, $J_t \in \mathfrak{H}$.*
2. *J_t depends analytically on t.*

Proof.

(1) Using the trigonometric identity $\tan^2\theta - \sec^2\theta = -1$ and the fact that all real powers of a positive definite matrix are positive definite, direct computation shows that J_t satisfies the conditions of Proposition 6.14 for any $t \in \mathbb{B}$. (Note that $\tilde{\mathscr{A}}$ and $\tilde{\mathscr{B}}$ are Lagrangian subspaces of $H_1(\tilde{X}, \mathbb{R})^-$, and that $J_t\,\tilde{\mathscr{A}} = \tilde{\mathscr{B}}$ for all t.)

(2) Let D' denote differentiation with respect to $t' = \operatorname{Re} t$ and D'' denote differentiation with respect to $t'' = \operatorname{Im} t$. Recall that at any $J \in \mathcal{C}(V)$, the complex structure on $T_J\mathcal{C}(V)$ is given by left-multiplication by J. Now the family J_t is seen by direct computation to satisfy equation (5.1) on page 48. □

Lemma 7.4. *The disk $J_t : \mathbb{B} \to \mathfrak{H}$ extends continuously to a the boundary $\partial\mathbb{B} = \{z \in \mathbb{C} \mid |\operatorname{Im} z| = \pi/2\}$, and $\delta(k(J_t)) = 1$ for all $t \in \partial\mathbb{B}$.*

Proof. For each $t = t' + it'' \in \mathbb{B}$, let

$$A_t = (\mathrm{id} + J_{t'}J_t)(\mathrm{id} - J_{t'}J_t)^{-1} = 2(\mathrm{id} - J_{t'}J_t)^{-1} - \mathrm{id} \in \operatorname{Hom}_{\overline{J_{t'}}}(H_1(\tilde{X}, \mathbb{R})^-).$$

Throughout these computations, I represents the $(3g-3)\times(3g-3)$ identity matrix. When quotients appear, their values (as matrices) are well-defined because the denominator is invertible and the expressions involved commute; we write them as quotients to save horizontal space and for ease of reading.

$$J_{t'}J_{t'+it''} = \begin{bmatrix} -\sec\Theta(t'') & -\tan\Theta(t'')(M^\top M)^{-\frac{1}{2}-\frac{t'}{\log\lambda}}M^\top \\ -M(M^\top M)^{-\frac{1}{2}+\frac{t'}{\log\lambda}}\tan\Theta(t'') & -M\sec\Theta(t'')M^{-1} \end{bmatrix}$$

$$2(\mathrm{id} - J_{t'} J_{t'_i t''})^{-1} =$$
$$\begin{bmatrix} I & -\dfrac{\sin\Theta(t'')}{I+\cos\Theta(t'')}(M^\top M)^{-\frac{1}{2}-\frac{t'}{\log\lambda}}M^\top \\ -M(M^\top M)^{-\frac{1}{2}+\frac{t'}{\log\lambda}}\dfrac{\sin\Theta(t'')}{I+\cos\Theta(t'')} & I \end{bmatrix}$$

We can find the value of δ on for $|\operatorname{Im} t| = \pi/2$ by computing the dimension of the kernel of $\mathrm{id} - A_t$ when $t'' = \pm\pi/2$. By the previous calculation and the definition of A_t, we have immediately

$$\mathrm{id} - A_{t'\pm\pi/2} =$$
$$\begin{bmatrix} I & \dfrac{\sin\Theta(\pm\pi/2)}{I+\cos\Theta(\pm\pi/2)}(M^\top M)^{-\frac{1}{2}-\frac{t'}{\log\lambda}}M^\top \\ M(M^\top M)^{-\frac{1}{2}+\frac{t'}{\log\lambda}}\dfrac{\sin\Theta(\pm\pi/2)}{I+\cos\Theta(\pm\pi/2)} & I \end{bmatrix}.$$

This matrix is row equivalent to

$$\begin{aligned} &\begin{bmatrix} I & \dfrac{\sin\Theta(\pm\pi/2)}{I+\cos\Theta(\pi/2)}(M^\top M)^{-\frac{1}{2}-\frac{t'}{\log\lambda}}M^\top \\ 0 & I - M\dfrac{\sin^2\Theta(\pi/2)}{I+\cos\Theta(\pi/2))^2}M^{-1} \end{bmatrix} \\ &= \begin{bmatrix} I & \dfrac{\sin\Theta(\pm\pi/2)}{I+\cos\Theta(\pi/2)}(M^\top M)^{-\frac{1}{2}-\frac{t'}{\log\lambda}}M^\top \\ 0 & I - M\left(I - \dfrac{I-\cos^2\Theta(\pi/2)}{(I+\cos\Theta(\pi/2))^2}\right)M^{-1} \end{bmatrix} \\ &= \begin{bmatrix} I & \dfrac{\sin\Theta(\pi/2)}{I+\cos\Theta(\pi/2)}(M^\top M)^{-\frac{1}{2}-\frac{t'}{\log\lambda}}M^\top \\ 0 & 2M\dfrac{\cos\Theta(\pi/2)}{I+\cos\Theta(\pi/2)}M^{-1} \end{bmatrix}. \end{aligned}$$

The nullity of this last matrix equals the nullity of the $(2,2)$-block, which in turn equals the nullity of $\cos\Theta(\pi/2)$. Because $M^\top M$ is Perron–Frobenius, its largest eigenvalue is strictly greater in absolute value than its remaining eigenvalues, which means the same is true of $\Theta(t'')$ for all $t'' > 0$. Thus $\cos\Theta(\pi/2)$ vanishes on the eigenspace of $\Theta(\pi/2)$ corresponding to the eigenvalue $\pi/2$, and this is the entire kernel. Therefore the nullity is 1. □

An important feature of this family is that, for all real t, J_t coincides with the Hodge complex structure on the subspace spanned by ω_t. That is,

Lemma 7.5. *For all $t \in \mathbb{R}$, $J_t \operatorname{Re} \omega_t = \operatorname{Im} \omega_t$.*

Proof. Direct computation. □

Theorem 7.6. *J_t is a maximal disk in $\mathfrak{H}$ that does not arise from a Teichmüller disk.*

We will lean on the following observation from [23]: *given two directions on a flat surface that are affinely equivalent, if one is periodic, then both are, and they have the same number of cylinders, with the same height and width data, up to scaling.*

Proof. Suppose $J_t : \mathbb{B} \to \mathfrak{H}$ were the image of a Teichmüller disk. Then, because its action along the subspace spanned by $\operatorname{Re} \omega_t$ and $\operatorname{Im} \omega_t$ coincides with the Teichmüller flow along this curve, J_t must be the Teichmüller disk generated by ω_t. By the Thurston–Veech construction, the affine group of $(\widetilde{X}_t, q_t)$ includes many pseudo-Anosov elements, none of which stabilize the horizontal or vertical directions of q_t. Let θ_1 and θ_2 be the images of the horizontal and vertical directions by any of these elements.

Because θ_1 and θ_2 are affinely equivalent to the horizontal and vertical directions, they have systems of curves $\mathscr{A}'$ and $\mathscr{B}'$ with the same intersection properties as those used to construct q_t. We could therefore begin with this system of binding curves to construct the same disk J_t. The directions on the boundary of $\mathbb{B}$ corresponding to θ_1 and θ_2 should then have the same degeneracy properties (i.e., they should land in the same stratum of $\partial\mathfrak{H}$) as the horizontal and vertical directions.

However, as calculated above, $\delta(\pm\infty) > 1$, while δ is constantly 1 on the boundary of $\mathbb{B}$. This is a contradiction, and we conclude that the family $\{J_t\}$ does not arise from a Teichmüller disk. $\square$

Remark 7.7. When the characteristic polynomial of $M^\top M$ is irreducible, the various roots (all of which are real) determine various cohomology classes on $\widetilde{X}$. These are the "Galois conjugates" of $\operatorname{Re}\omega_t$ and $\operatorname{Im}\omega_t$. It is with respect to the basis of $H^1(\widetilde{X},\mathbb{R})^-$ formed by these classes that J_t takes the form $\left(\begin{smallmatrix} 0 & -I \\ I & 0 \end{smallmatrix}\right)$. This may be seen in one of two ways. First, the result of Lemma 7.5 is purely algebraic, and so it applies also to these Galois conjugate classes. Second, as observed at the beginning of this section and by an application of Theorem 6.8, the eigenspaces of $M^\top M$ dictate the behavior of J_t.

Part III

Examples

CHAPTER 8

SAMPLE ISO-DELAUNAY TESSELLATIONS

In this chapter, we provide some examples of iso-Delaunay tessellations, primarily for surfaces known to have lattice Veech groups.

8.1 Genus 2 surfaces from L-shaped tables

K. Calta and C. McMullen provided the classification of non-arithmetic genus 2 surfaces with lattice Veech groups in [9] and [30, 32, 31, 34], respectively. McMullen mentions in particular surfaces that arise from the Katok–Zemljakov billiard construction [53] applied to "L-shaped tables", i.e., rectangles from which a rectangle has been removed from one corner. M. Bainbridge has computed the Euler characteristics of the resulting Teichmüller curves [3]. Here we show the tessellations from a particular case of this series, that of a square with side length $a = (1 + \sqrt{d})/2$, from which a corner square has been cut, leaving two sides of length 1.

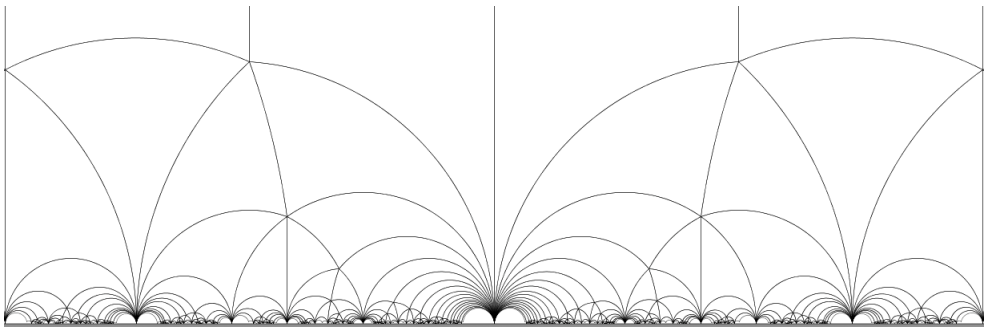

Figure 8.1: $d = 3$

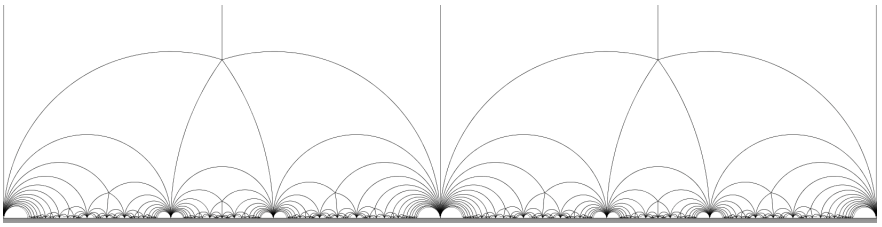

Figure 8.2: $d = 5$

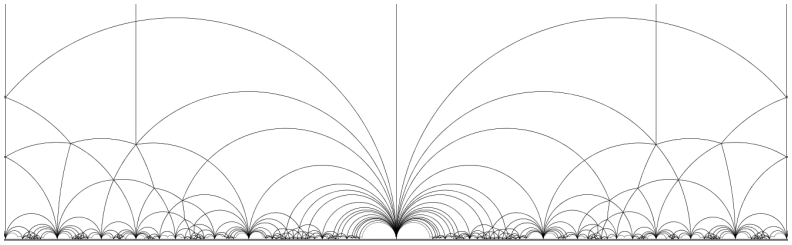

Figure 8.3: $d = 13$

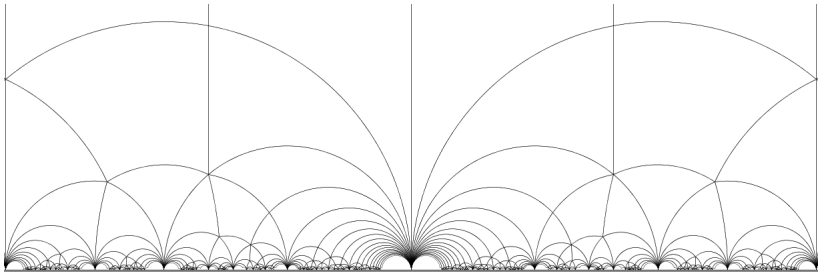

Figure 8.4: $d = 17$

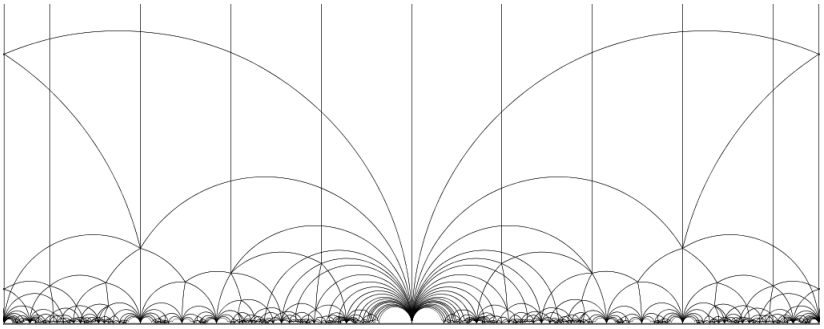

Figure 8.5: $d = 37$

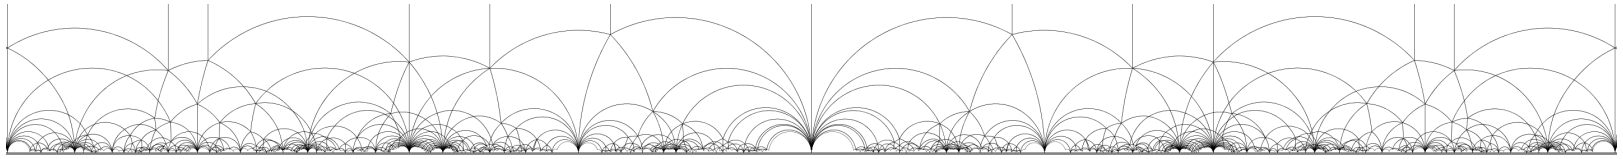

Figure 8.6: $d = 6$

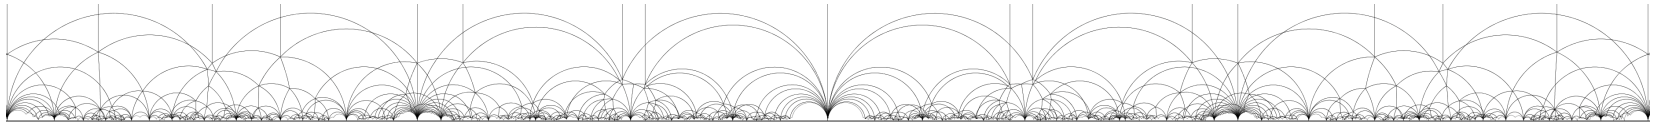

Figure 8.7: $d = 10$

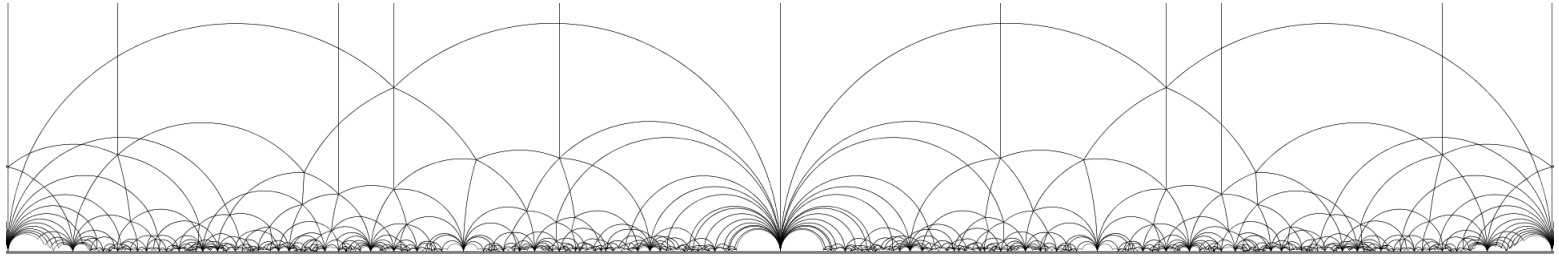

Figure 8.8: $d = 15$

8.2 Surfaces from rational triangles

We use the notation $T(p, q, r)$ to mean the triangle with angles $\pi(p/n)$, $\pi(r/n)$, and $\pi(q/n)$, where $n = p + q + r$, $S(p, q, r)$ to mean the surface obtained from applying the billiard construction to $T(p, q, r)$, and $\Sigma(p, q, r)$ to mean the corresponding iso-Delaunay tessellation.

In [48], Veech computed the iso-Delaunay tessellation for the surface that arises from billiards in a regular n-gon, and called these the "bicuspid surfaces", because they (and their covers) are precisely the surfaces for which the tessellations consist exclusively of tiles having two ideal vertices. (See for example the case $d = 5$ in the previous section, the so-called "golden table" which is affinely equivalent to the surface consisting of a pair of polygons, cf. [30].)

To date, four "exceptional" triangles have been discovered, not belonging to any general family. Three are acute, and by a theorem of Kenyon and Smillie [25] (modulo a number-theoretic conjucture, proved by Puchta [37]) they and Veech's examples are the only acute triangles whose surfaces have a lattice Veech group. These three have been related to the exceptional Coxeter–Dynkin diagrams E6, E7, and E8 by C. Leininger [27]. The fourth is discussed in P. Hooper's thesis (the relevant section is available separately as a preprint [19]).

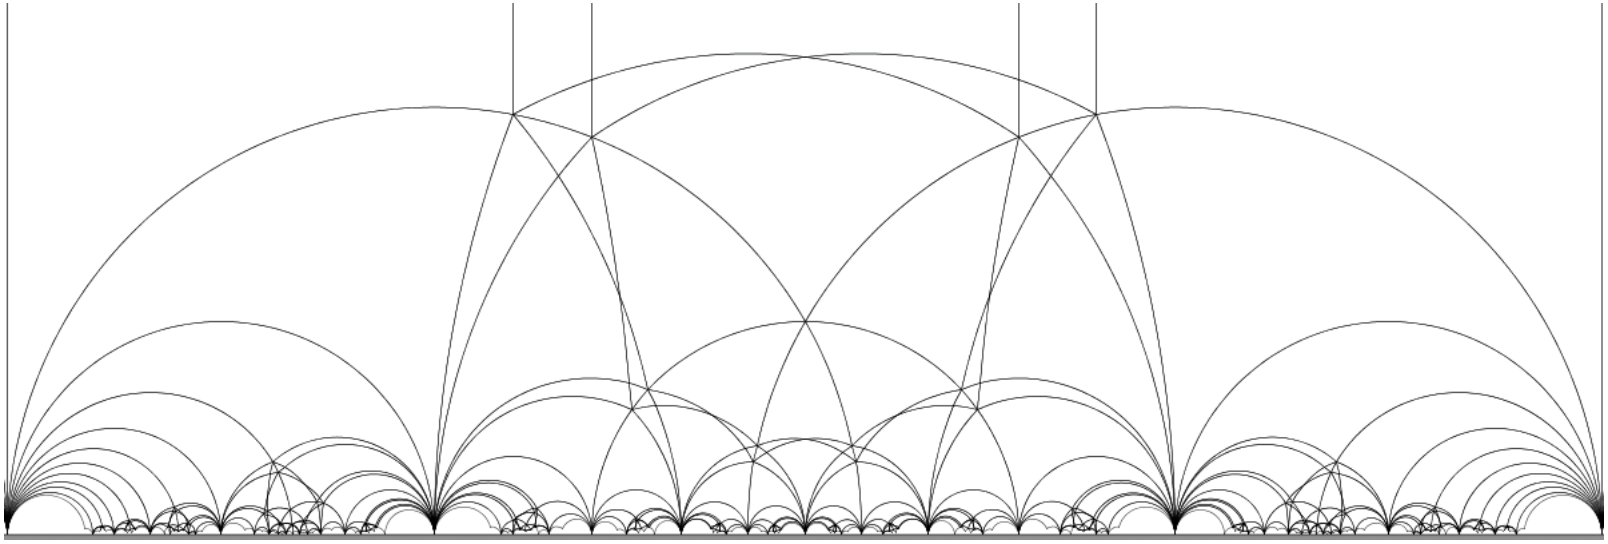

Figure 8.9: $\Sigma(3,4,5)$: genus 3. Discovered by Veech [49]. Associated to the Coxeter–Dynkin diagram E6.

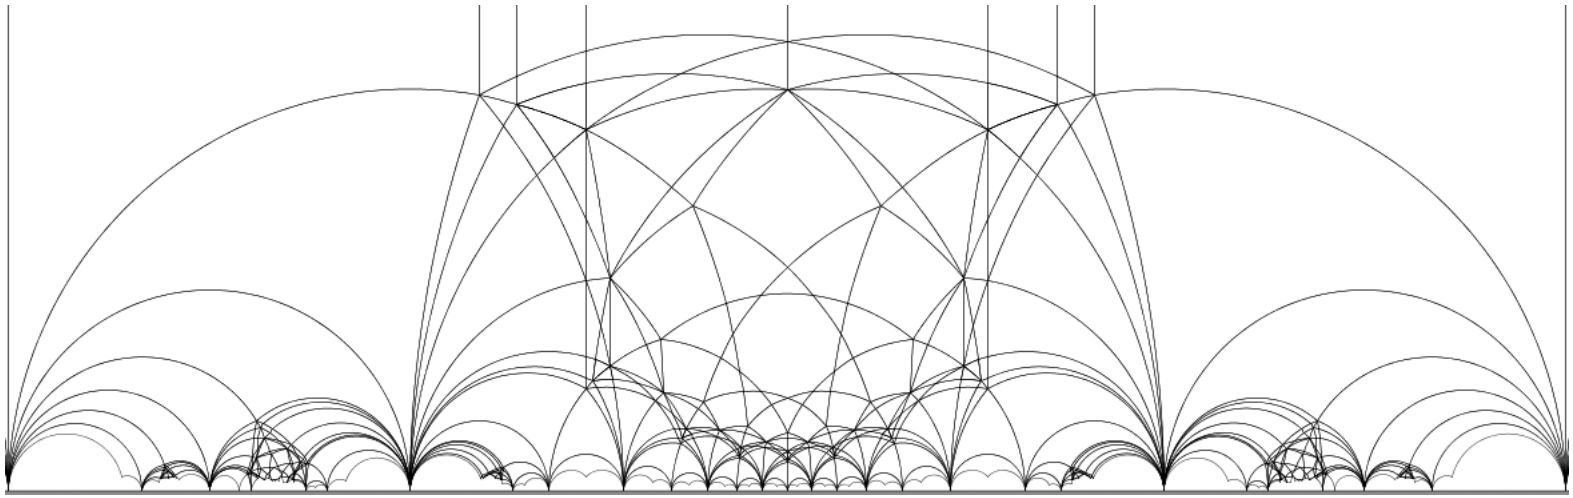

Figure 8.10: $\Sigma(2,3,4)$: genus 3. Discovered by Kenyon–Smillie [25]. Associated to the Coxeter–Dynkin diagram E7.

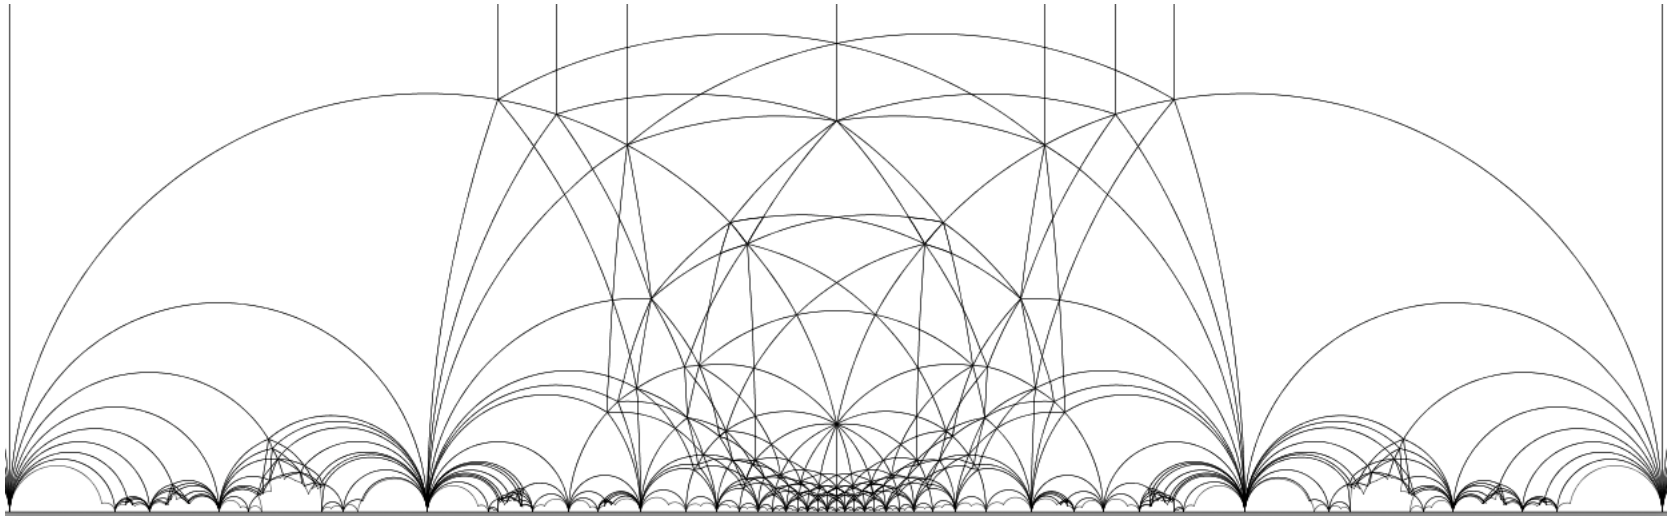

Figure 8.11: $\Sigma(3,5,7)$: genus 4. Discovered by Vorobets [52]. Associated to the Coxeter–Dynkin diagram E8.

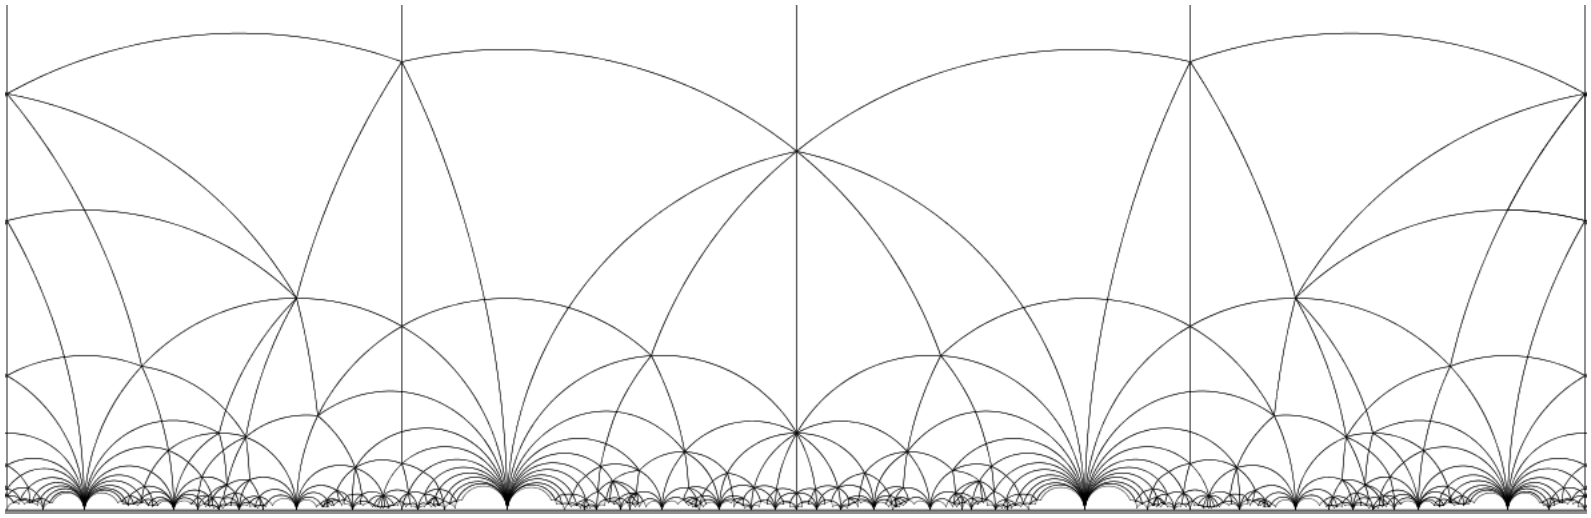

Figure 8.12: $\Sigma(1, 4, 7)$: genus 4. Discovered by Hooper [19] and McMullen [33].

(Note: the foregoing images are not all presented at the same scale; in several cases the point i, corresponding to the original surface, is not even included, but its location may be discerned by finding a reflection symmetry in a circle orthogonal to the central axis and taking the intersection of this circle with the axis.)

CHAPTER 9

AN EXCEPTIONAL SET OF EXAMPLES: THE ARNOUX–YOCCOZ SURFACES

Introduction: from the golden ratio to the geometric series

From our calculus courses, we know that the infinite geometric series $\frac{1}{2}+\frac{1}{4}+\frac{1}{8}+\cdots$ converges to 1. Indeed, using the summation formula $\sum_{k=1}^{\infty} x^k = x/(1-x)$, we find that $\frac{1}{2}$ is the unique solution to the equation $\sum_{k=1}^{\infty} x^k = 1$. From even earlier in our lives, perhaps, we recall that the equation $x + x^2 = 1$ has a unique positive solution, whose inverse is the golden ratio. The expression $x + x^2$ may be viewed as a partial geometric series, which can be extended to n terms: $x + \cdots + x^n$.

The positive solutions to the equations $x + \cdots + x^n = 1$ for $n \geq 3$ are instrumental in creating a certain family of measured foliations on surfaces, which were introduced by P. Arnoux and J.-C. Yoccoz in 1981 [2]. It was shown in 2005 by P. Hubert and E. Lanneau [22] that the Arnoux–Yoccoz examples do not arise from the Thurston–Veech construction (see chapter 7). In this chapter we will present the surfaces constructed by Arnoux and Yoccoz and give explicit triangulations, then use these to prove certain properties common to all these surfaces. We will also see that the family can be extended to include the cases $n = 2$ and $n = \infty$. These extreme cases will turn out to be exceptional in their construction—the first corresponds to a singular surface and the second to a surface of infinite type—but we hope that the self-similarity property that the golden ratio and the geometric series share with all of the other examples (see §9.1 and §9.4) will illuminate the entire sequence of surfaces for the reader.

9.1 Interval exchange maps

In this section we review the algebraic numbers and interval exchange maps involved in the construction of the Arnoux–Yoccoz translation surfaces. Given any $g \geq 2$, the polynomial

$$x^g + x^{g-1} + \cdots + x - 1 \tag{9.1}$$

has a unique positive root, since its values at 0 and 1 are -1 and $g-1$, respectively, and its derivative is positive for all positive x. We denote the positive root of (9.1) simply as α, suppressing its dependence on g. Arnoux and Yoccoz showed that the inverse of α is a Pisot number, which means that α is in fact the only root of (9.1) that lies within the unit disk. Hubert and Lanneau showed that, if g is even, then (9.1) has one negative root, and if g is odd, then α is the only real root. We add to these properties the following:

Lemma 9.1. *For each $g \geq 2$, the positive root α of (9.1) satisfies*

$$\frac{1}{2^{g+2}} < \alpha - \frac{1}{2} < \frac{1}{2^{g+1}} \tag{9.2}$$

and therefore it converges to $1/2$ exponentially fast as $g \to \infty$.

We will make use of this convergence in §9.4.

Proof. To obtain the lower bound, we will show that, when $r = 1/2 + 1/2^{g+2}$, the polynomial (9.1) evaluated at r is negative. This is equivalent to

$$\frac{1 - r^{g+1}}{1 - r} < 2, \qquad \text{or} \qquad \left(1 + \frac{1}{2^{g+1}}\right)^{g+1} > 1,$$

which is true for all $g \geq 2$. The upper bound is obtained similarly. □

Arnoux and Yoccoz [2] introduced an interval exchange map—i.e, a piecewise isometry of an interval that is bijective and has only finitely many points

of discontinuity—based on the geometric properties of α. First, the unit interval is subdivided into g intervals of lengths α, α^2, ..., α^g. Each of these subintervals is divided in half, and the halves are exchanged within each subinterval. Finally the entire unit interval is divided into half, and these two halves are exchanged. We denote the total process f_g (see Figure 9.1). We will occasionally be interested in the behavior of f_g and its iterates on the endpoints of the subintervals, so for specificity we restrict the map to $[0, 1)$ and assume that the left endpoint of each piece is carried along. The key feature of f_g is its self-similarity:

Proposition 9.2 (Arnoux–Yoccoz)**.** *Let $\tilde{f}_g$ be the interval exchange map induced on $[0, \alpha)$ by the first return map of f_g. Then f_g is conjugate to $\tilde{f}_g$.*

The proof uses an explicit piecewise affine map $h_g : [0, 1) \to [1, \alpha)$, defined as follows:

$$h_g(x) = \begin{cases} \alpha x + \frac{\alpha+\alpha^{g+1}}{2}, & x \in \left[0, \frac{1-\alpha^g}{2}\right) \\ \alpha x - \frac{\alpha-\alpha^{g+1}}{2}, & x \in \left[\frac{1-\alpha^g}{2}, 1\right) \end{cases}$$

which satisfies $f_g = h_g^{-1} \circ \tilde{f}_g \circ h_g$. In §9.4, we will show similar kinds of results for certain exchanges on infinitely many subintervals.

In their original paper, citing work of G. Levitt, Arnoux and Yoccoz state that, for a given interval exchange map:

> *... on peut construire une suspension canonique, et l'on sait que toute suspension possédant les mêmes singularités (en type et en nombre) que cette suspension canonique lui est homéomorphe par un homéomorphisme préservant la mesure transverse du feuilletage.*

(The "canonical suspension" is a closed surface with a measured foliation along with a closed curve transverse to the foliation on which the first return map of

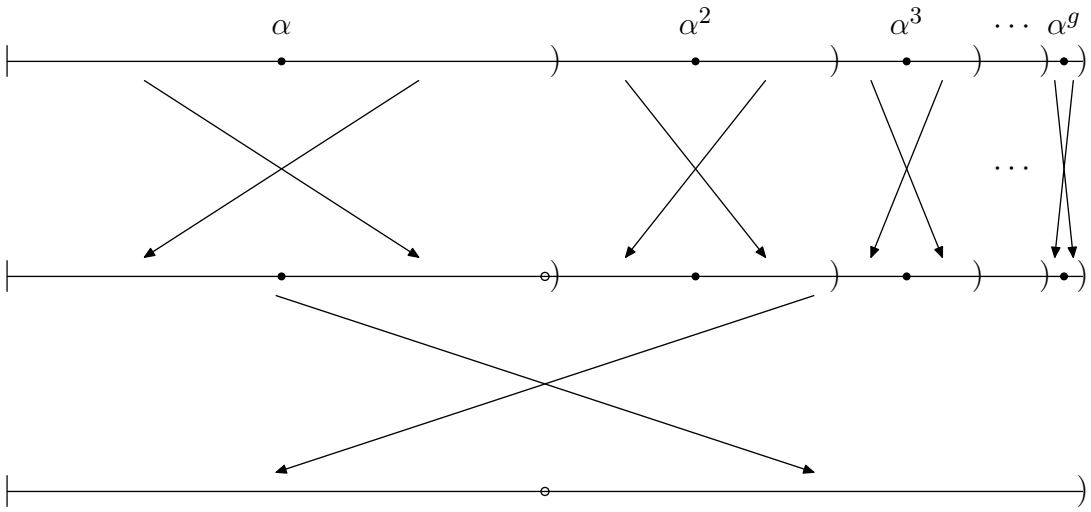

Figure 9.1: The interval exchange f_g as a composition of two involutions.

the foliation induces the given interval exchange map.) They then use this result and the self-similarity of f_g to demonstrate the existence of a pseudo-Anosov homeomorphism ψ_g on a surface of genus g such that the expansion constant of ψ_g is $1/\alpha$. Fortunately, in a separate article [1], Arnoux gives an explicit description of the canonical suspension of f_3 and illustrates ψ_3. In this chapter we will present the generalization of Arnoux's construction to all genera and exploit these presentations to make further conclusions about the Arnoux–Yoccoz surfaces.

9.2 Steps and slits

Fix $g \geq 3$. In this section, we will present the genus g Arnoux–Yoccoz surface (X_g, ω_g) by generalizing Arnoux's presentation of (X_3, ω_3). Starting with a unit square, we carve out a "staircase" in the upper right-hand corner, with the widths of the steps, from left to right, given by α, α^2, $\ldots$, α^g, and the distances between the steps, going down, given by α^g, α^{g-1}, $\ldots$, α. We further slit this square along

several vertical segments σ_1, σ_2, ..., σ_g. The slits are made starting along the bottom edge of the square at points whose x-coordinates are images by f_g of 0 and the left-hand endpoints of the intervals $[\alpha^i, \alpha^{i+1})$ $(1 \leq i \leq g-1)$.

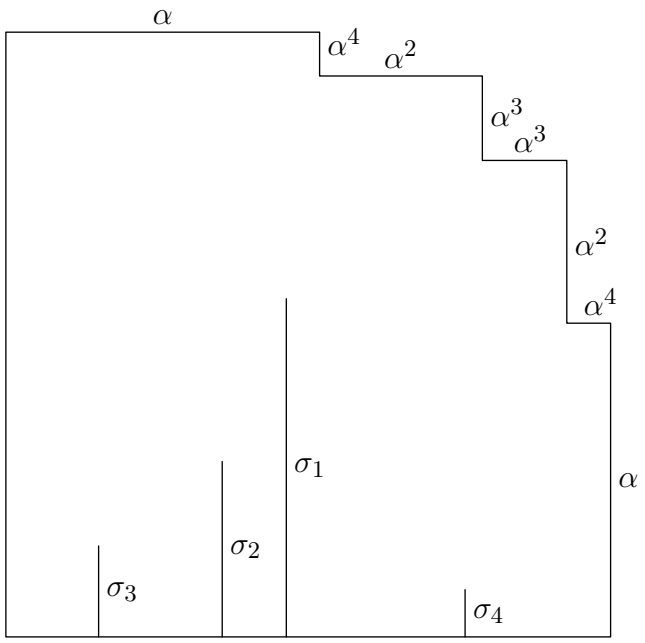

Figure 9.2: The steps and slits for the genus 4 Arnoux–Yoccoz surface

Now we wish to provide appropriate gluings for the surface to have an affine self-map. These identifications are as follows:

- The tops of the steps are glued to the bottom of the unit square according to the interval exchange f_g.
- The vertical edge of the bottommost step, having length α, is identified with the bottom portion of the leftmost vertical edge.
- The remaining top portion of the leftmost edge of the square, having length $1-\alpha$, is identified with the bottom portion to the left of σ_1.
- The vertical edge of the step having height α^i $(2 \leq i \leq g)$ is identified with the bottom portion to the right of the segment σ_{i-1}.

- The remaining top portion to the right of each segment σ_i $(1 \leq i \leq g-1)$ is identified with the left of the segment σ_{i+1}.
- The right-hand side of σ_g is identified with the left-hand side of the top of σ_1.

There is a one real-parameter family of surfaces that satisfy the gluings given above; the easiest parameter to vary is $|\sigma_g|$. We want to single out a value for this parameter so that the surface admits a pseudo-Anosov affine map. The required condition is described by the equation $\alpha(1+|\sigma_g|) = (1-\alpha)+|\sigma_g|$, which says that the length of σ_1 is α times the sum of the length of σ_g and the length of the left edge of the square (i.e., 1). Solving this equation, we find $|\sigma_g| = (2\alpha-1)/(1-\alpha)$, which determines the lengths of the remaining slits.

The pseudo-Anosov homeomorphism $\psi_g : X_g \to X_g$ expands the horizontal foliation of ω_g by a factor of $1/\alpha$ and contracts the vertical foliation by a factor of α. It permutes the vertical segments in a predictable manner: for each i from 1 to $g-1$, ψ_g sends σ_i to σ_{i+1}, and also sends the union of σ_g with the left-hand edge of the initial square to σ_1. The step of height α^i is also sent to the step of height α^{i+1} $(1 \leq i \leq g-1)$.

9.3 Triangulations

Again fix $g \geq 3$. In this section we construct the surface (X_g, ω_g) from $4g$ triangles. Begin with the points $P_0, \ldots, P_g, Q_0, \ldots, Q_g$ in $\mathbb{R}^2$, chosen as follows (see

Figure 9.3):

$$P_0 = \left(\frac{1-\alpha^g}{2}, \frac{\alpha^2}{1-\alpha}\right), \qquad Q_0 = \left(-\frac{\alpha^g}{2}, \alpha\right),$$
$$P_1 = \left(-\frac{\alpha^{g-1}+\alpha^g}{2}, \frac{\alpha-\alpha^2+\alpha^3}{1-\alpha}\right), \qquad P_g = \left(1+\frac{\alpha-\alpha^g}{2}, \frac{3\alpha-1-\alpha^2}{1-\alpha}\right),$$
$$P_i = \left(\frac{\alpha-\alpha^i}{1-\alpha}, \frac{\alpha}{1-\alpha}\right) \qquad \text{for } i = 2, \ldots, g-1,$$
$$Q_i = \left(\frac{2\alpha-\alpha^i-\alpha^{i+1}}{2(1-\alpha)}, \frac{\alpha-\alpha^{g-i+2}}{1-\alpha}\right) \qquad \text{for } i = 1, \ldots, g.$$

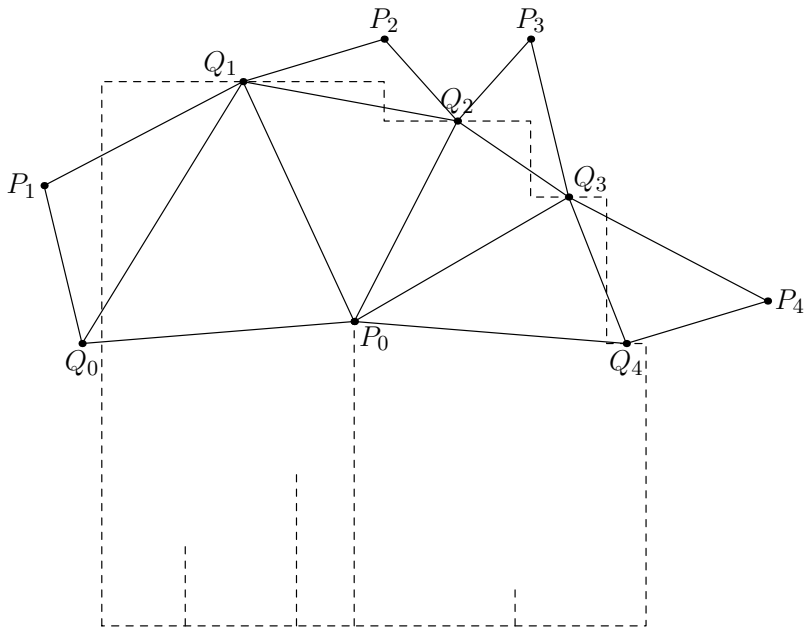

Figure 9.3: The points $P_0, \ldots, P_4, Q_0, \ldots, Q_4$ relative to (X_4, ω_4)'s staircase

For $i = 1, \ldots, g$, set $T_i = P_0 Q_i Q_{i-1}$ and $T_{g+i} = P_i Q_{i-1} Q_i$. For $i = 1, \ldots, 2g$, let T_i' be the reflection of T_i in the horizontal axis. Glue the T_is along their common boundaries, and likewise for the T_i's. Then each remaining "free" edge is a translation of another; we glue each such pair of edges:

- P_0Q_0 is paired with $P_0'Q_g'$, and $P_0'Q_0'$ is paired with P_0Q_g.
- P_1Q_1 is paired with $P_g'Q_{g-1}'$, and $P_1'Q_1'$ is paired with P_gQ_{g-1}.

- P_1Q_0 is paired with $P_{g-1}Q_{g-1}$, and $P'_1Q'_0$ is paired with $P'_{g-1}Q'_{g-1}$.
- P_gQ_g is paired with Q_1P_2, and $P'_gQ'_g$ is paired with $Q'_1P'_2$.
- For $i = 2, \ldots, g-2$, P_iQ_i is paired with $Q'_iP'_{i+1}$ and $P'_iQ'_i$ is paired with Q_iP_{i+1}.

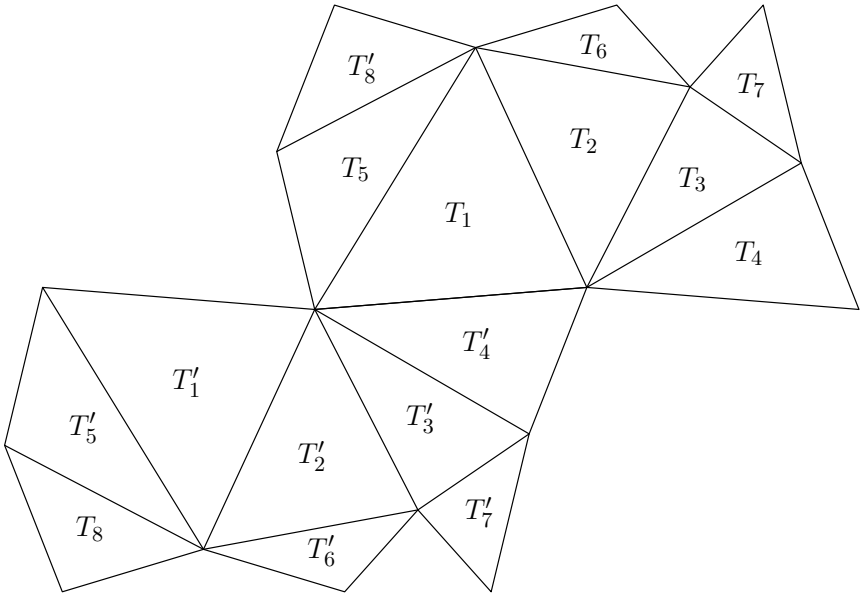

Figure 9.4: The triangles comprising (X_4, ω_4)

(See Figure 9.4.) All of the P_is and Q'_is are identified to become a cone point, and likewise for all of the Q_is and P'_is. An Euler characteristic computation shows that the resulting surface has genus g, and because the two cone points are symmetric, they each have a cone angle of $2g\pi$. One can verify the following result directly by checking that the surface we have constructed from triangles is isometric to the staircase presentation.

Proposition 9.3. *The T_is and T'_is induce a triangulation of (X_g, ω_g).*

Corollary 9.4. $\mathrm{Aff}(X_g, \omega_g)$ *contains a fixed-point free, orientation-reversing involution ρ_g, which commutes with ψ_g, and whose derivative is reflection in the x-axis.*

The existence of this symmetry occurs for a completely general reason: f_g is conjugate to its inverse by the following "rotation" of the unit interval:

$$r(x) = \begin{cases} x + \frac{1}{2}, & x \in [0, \frac{1}{2}) \\ x - \frac{1}{2}, & x \in [\frac{1}{2}, 1) \end{cases}$$

By the reasoning invoked in §9.1, the surface obtained from (X_g, ω_g) by applying complex conjugation to the charts of ω_g (which is a suspension of f_g^{-1}, and therefore of f_g) is translation equivalent to (X_g, ω_g) itself, which yields the existence of ρ_g.

Corollary 9.5. *The compact non-orientable surface of Euler characteristic* $1 - g$ *admits a pseudo-Anosov homeomorphism whose invariant foliations have one singular point and whose expansion constant has degree* g*.*

This corollary generalizes a result from the original paper by Arnoux and Yoccoz, in which the surface (X_3, ω_3) is constructed by two separate methods: first by a lifting a pseudo-Anosov homeomorphism of $\mathbb{RP}^2$ to the non-orientable surface of Euler characteristic -2 and then to genus 3, and second by the method of suspending an interval exchange map.

Corollary 9.6. *If* $g \geq 4$*, then* X_g *is not hyperelliptic.*

Proof. Every abelian differential on a hyperelliptic surface is odd with respect to the hyperelliptic involution. If, for some $g \geq 4$, X_g were hyperelliptic, then there would have to be an isometry of (X_g, ω_g) with derivative $-\mathrm{id}$. Such an isometry would have to preserve the Delaunay triangulation of (X_g, ω_g). But no other triangle is isometric to T_1, so such an isometry does not exist. □

9.4 A limit surface: $(X_\infty, \omega_\infty)$

Lemma 9.1 implies that each triangle that appears in the construction of some (X_g, ω_g) has a "limiting position"; from these we can construct a "limit surface" of infinite genus. To be precise, we obtain a non-compact translation surface $(X_\infty, \omega_\infty)$, where X_∞ has infinite genus, whose metric completion is the one-point compactification of X_∞. In a sense, the two cone points of the (X_g, ω_g), $g < \infty$, have "collapsed" into each other, leaving an essential singularity at which all of the "curvature" of the space $(X_\infty, \omega_\infty)$ is concentrated. We shall briefly address in §9.5 the nature of singularities on this surface. See Figure 9.5 for the definition of this surface; ω_∞ is, as usual, the 1-form induced on the quotient by $\mathrm{d}z$ in the plane. A *critical trajectory* of $(X_\infty, \omega_\infty)$ is a geodesic trajectory that leaves every compact subset of X_∞. A *saddle connection* of $(X_\infty, \omega_\infty)$ is a geodesic trajectory (of finite length) that leaves every compact subset of X_∞ in both directions.

Theorem 9.7. *X_∞ is a Riemann surface of infinite genus with one end, and ω_∞ is an abelian differential of finite area on X_∞ without zeroes on X_∞.* $\mathrm{Aff}(X_\infty, \omega_\infty)$ *includes an orientation-reversing isometric involution ρ_∞ without fixed points on X_∞ and a pseudo-Anosov homeomorphism ψ_∞ with expansion constant* 2. *These two elements commute.*

Proof. (In this paragraph, we follow the method of proof used by R. Chamanara in [12].) That X_∞ is a Riemann surface is evident, as are the claims about ω_∞. The fact that X_∞ has infinite genus can be deduced from the existence of a set of pairwise non-homotopic simple closed curves $\{\gamma'_n, \gamma''_n\}_{n\in\mathbb{N}}$, where γ'_n (respectively, γ''_n) connects the midpoints of the edges labelled C'_n (respectively, C''_n), and each γ'_n intersects only γ''_n (and vice versa). To show that X_∞ has only one topological end, we construct a sequence of compact subsurfaces with boundary. Let K_g be

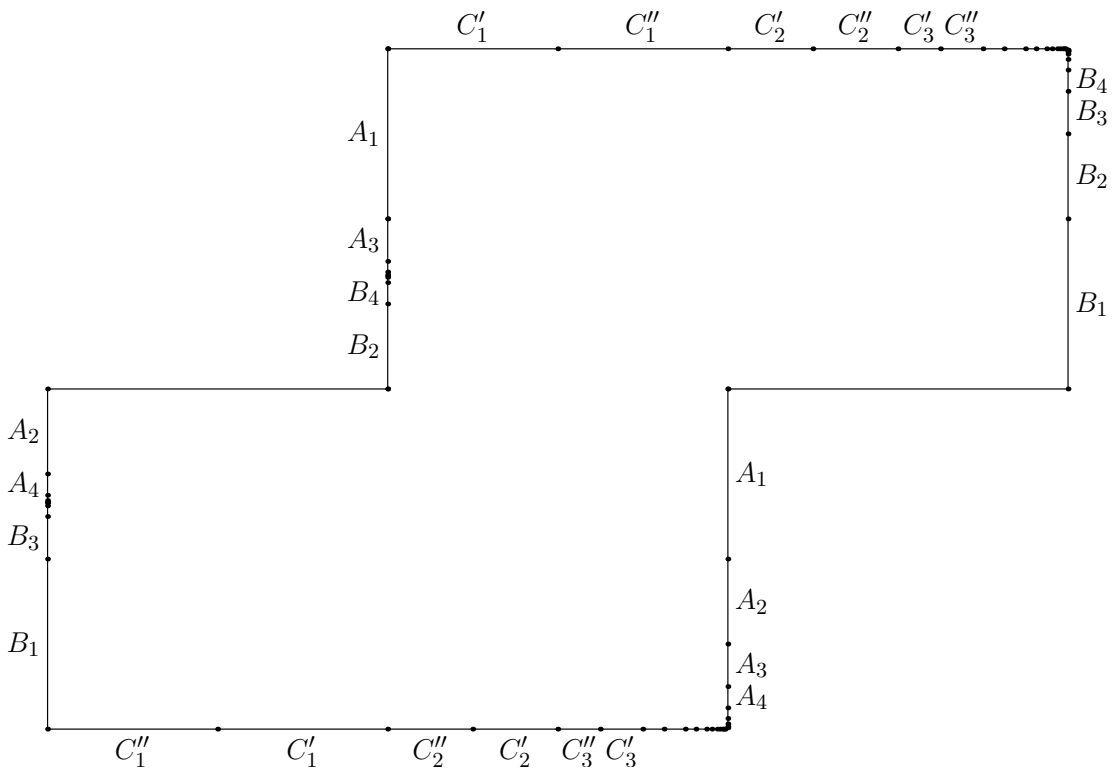

Figure 9.5: The surface $(X_\infty, \omega_\infty)$. Each pair of edges with the same label is identified by translation. The length of each A_n, B_n, C_n', or C_n'' is $1/2^{n+1}$.

the complement of the union of the open squares having side length $1/2^{g+1}$ and centered at the endpoints of the segments A_n, B_n, C_n', C_n''. These K_g satisfy $K_g \subset K_{g+1}$ and $\bigcup K_g = X_\infty$, and the complement of each K_g has one component. Therefore by definition X_∞ has one topological end.

The orientation-reversing affine map ρ_∞ is visible in Figure 9.5 as a glide-reflection in a horizontal axis with translation length $1/2$. It sends the interior of the upper rectangle to the interior of the lower rectangle, each edge labeled A_n to an edge labeled B_n, and each edge labeled C_n' to an edge labeled C_n''. Therefore it has no fixed points.

Now we demonstrate the pseudo-Anosov affine map ψ_∞. Let R be the central

rectangle in Figure 9.5, and let S_1 and S_2 be the squares in the lower left and upper right, respectively. Expand R horizontally by a factor of 2, and contract R vertically by a factor of $1/2$ to obtain $\psi_\infty(R)$. Do the same with the rectangle R' which is the union of S_1 and S_2 (the top edge of S_1 is glued to the bottom edge of S_2) to obtain $\psi_\infty(R')$. Take $\psi_\infty(R)$ and lay it over S_1 and the lower half of R, and lay $\psi_\infty(R')$ over S_2 and the top half of R. This affine map is compatible with all identifications. That ψ_∞ and ρ_∞ commute may be checked directly. □

The pseudo-Anosov map $\psi_\infty : X_\infty \to X_\infty$ is a variant of the well-studied baker map, and thus $(X_\infty, \omega_\infty)$ is an alternate infinite-genus realization of this map, which was demonstrated on a "hyperelliptic" infinite-genus surface by Chamanara–Gardiner–Lakic [13]. The topological type of X_∞ is that of a "Loch Ness monster" and is therefore related to the surfaces described in [36], although the flat structure of ω_∞ does not fall into the class of surfaces studied there.

Let us make precise the notion of $(X_\infty, \omega_\infty)$ as a "limit" of (X_g, ω_g). We establish canonical piecewise-affine embeddings $\iota_g : K_g \to X_g$, where the K_g are the subsurfaces defined in the proof of Theorem 9.7, in such a way that $\iota_g^* |\omega_g|$ converges to $|\omega_\infty|$ on compact subsets of X_∞ as $g \to \infty$. (Here $|\omega_n|$ indicates the metric induced on X_n by ω_n, $3 \le n \le \infty$.) In fact, each ι_g will be defined on an open set U_g containing K_g and dense in X_∞.

For each $3 \le g < \infty$, let U_g be the surface obtained from Figure 9.5 by making all identifications up through index $\lfloor g/2 \rfloor$ for the A_is and B_is, and all identifications up through index $\lfloor (g-1)/2 \rfloor$ for the C_i's and the C_i''s. (Here and elsewhere $x \mapsto \lfloor x \rfloor$ denotes the "floor" function.) Retract the union of the triangles

$$\left\{ \left(\frac{1}{2}, \frac{1}{2}\right), \left(1, \frac{2^{\lfloor g/2 \rfloor} - 1}{2^{\lfloor g/2 \rfloor}}\right), (1,1) \right\} \quad \text{and} \quad \left\{ \left(\frac{1}{2}, \frac{1}{2}\right), (1,1), \left(\frac{2^{\lfloor (g-1)/2 \rfloor} - 1}{2^{\lfloor (g-1)/2 \rfloor}}, 1\right) \right\}$$

onto the triangle $\{(1/2, 1/2), (1, 1 - 1/2^{\lfloor g/2 \rfloor}), (1 - 1/2^{\lfloor (g-1)/2 \rfloor}, 1)\}$ by a homeomorphism, affine on each of the original triangles. Now a surface of genus g with two punctures can be created directly by identifying the "free" edge of this triangle with one of the "free" segments on the leftmost edges of the polygon.

Figure 9.6 shows the outlines of the first few surfaces in the sequence (X_g, ω_g). By adjusting the positions of the triangles in the upper right and upper left corners (e.g., removing the triangles labelled $T_{2g-\lfloor g/2 \rfloor}$ through T_{2g}, in addition to their mirror images, and regluing them along their longest edges in the appropriate location), one finds that there is a piecewise-affine map ι_g carrying U_g to X_g. Moreover, because $U_{g-1} \subset U_g$, ι_g restricts to an embedding of U_{g-1}, as well.

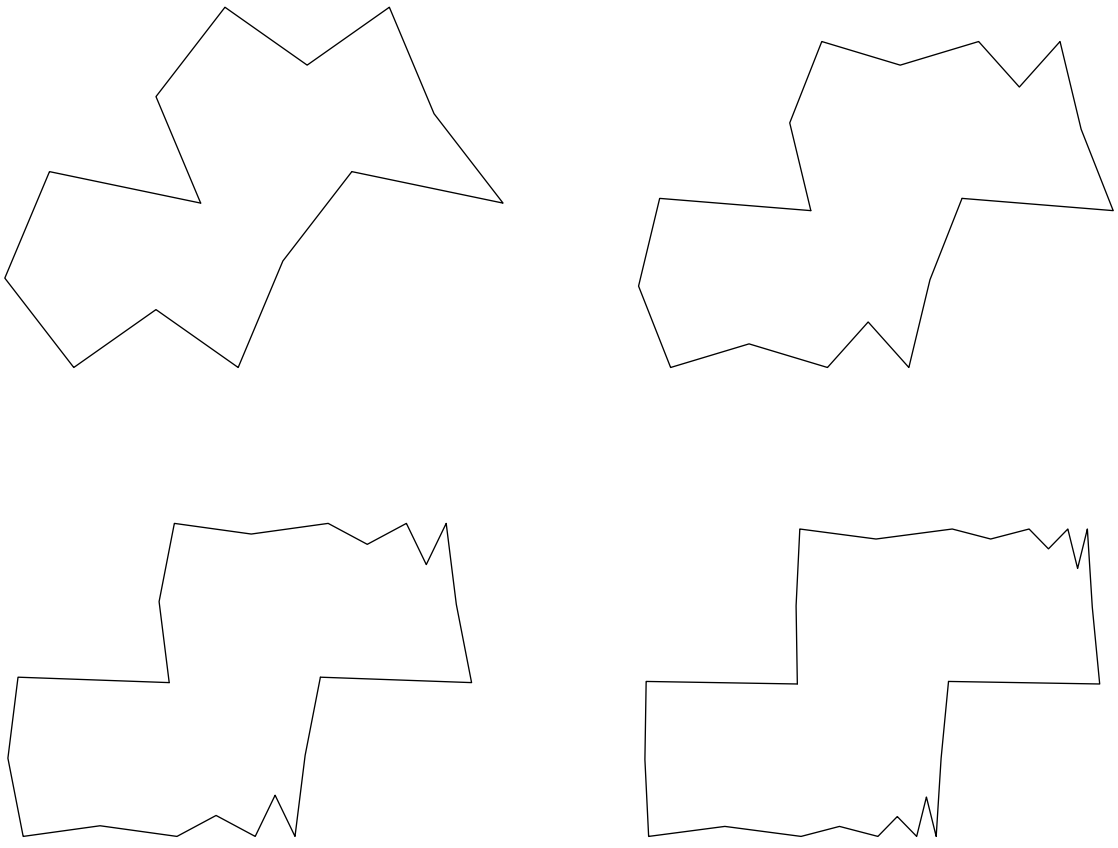

Figure 9.6: Outlines of the surfaces (X_g, ω_g) for $g = 3, 4, 5, 6$

Theorem 9.8. *The metrics $\iota_g^* |\omega_g|$ converge to $|\omega_\infty|$ uniformly on compact subsets of X_∞.*

Proof. Any compact $K \subset X_\infty$ is contained in some U_n. For any pair of points $P', P'' \in K$, the ratio of the distance from P' to P'' in each of the metrics $\iota_g^*|\omega_g|$ and $|\omega_\infty|$ is bounded by the quasi-conformal constants and the Jacobian determinants of the maps ι_g, which are uniformly bounded over all of K. As these constants approach 1, so do the ratios of lengths over K, uniformly. □

9.5 The affine group of $(X_\infty, \omega_\infty)$

In this section we will explore some of the geometry and dynamics of $(X_\infty, \omega_\infty)$, culminating in a proof of the following:

Theorem 9.9. $\mathrm{Aff}(X_\infty, \omega_\infty) \cong \mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ *is generated by ψ_∞ and ρ_∞.*

Let us revise our definition of "interval exchange map" to include injective maps from an interval to itself that are upper semicontinuous piecewise isometries. (This keeps with the "continuous at left endpoints" convention, although we may lose the property of bijectivity, as we shall see.) Then the vertical foliation of $(X_\infty, \omega_\infty)$ induces an interval exchange map $f_\infty : [1, 0) \to [1, 0)$, which can also be defined as follows: first, swap the two halves of each interval $[\frac{2^n-1}{2^n}, \frac{2^{n+1}-1}{2^{n+1}})$, then swap $[1, 1/2)$ with $[1/2, 1)$.

We can encode f_∞ symbolically as follows: if we do not allow the binary expansion of a number to terminate with only 1s, then each number in $[0, 1)$ has a unique binary expansion. Use these to identify $[0, 1)$ with the set $\mathfrak{B} \subset (\mathbb{F}_2)^{\mathbb{N}}$ consisting of

sequences that do not terminate with only 1s. Given a sequence $a = a_0a_1a_2\cdots$, we obtain $f_\infty(a)$ as follows:

1. find the first $i \in \mathbb{N}$ such that $a_i = 0$, and replace a_{i+1} with $a_{i+1} + 1$;
2. replace a_0 with $a_0 + 1$.

The inverse map f_∞^{-1} simply reverses these two steps. Both f_∞ and f_∞^{-1} are bijections. We remark that the first return map of f_∞ on either $[0, 1/2)$ or $[1/2, 1)$ is simply the restriction of f_∞^2 to the respective interval.

To aid our study at this point, we use the map r defined in §9.3 along with the following:

$$h'(x) = \frac{x}{2}, \qquad h''(x) = (r \circ h')(x) = \frac{x}{2} + \frac{1}{2},$$

$$h_\infty(x) = (h' \circ r)(x) = \begin{cases} \frac{1}{2}\left(x + \frac{1}{2}\right), & x \in [0, \frac{1}{2}) \\ \frac{1}{2}\left(x - \frac{1}{2}\right), & x \in [\frac{1}{2}, 1) \end{cases}$$

In terms of binary expansions, we can describe the effects of these functions on a sequence $a \in \mathfrak{B}$ as follows:

- r replaces a_0 with $a_0 + 1$;
- h' appends a 0 to the beginning of the sequence;
- h'' appends a 1 to the beginning of the sequence;
- h_∞ replaces a_0 with $a_0 + 1$ and appends a 0 to the beginning of the sequence.

The formalism of encoding these maps to act on infinite binary sequences makes immediate the following result.

Lemma 9.10. *Let f_∞, r, h', h'', and h_∞ act on $\mathfrak{B}$ as above. Then:*

- *r conjugates f_∞ to f_∞^{-1}.*
- *h' conjugates $f_\infty^2|_{[0,\,1/2)}$ to f_∞^{-1}.*
- *h'' conjugates $f_\infty^2|_{[1/2,\,1)}$ to f_∞.*
- *h_∞ conjugates $f_\infty^2|_{[0,\,1/2)}$ to f_∞.*

Proof. We will prove the second claim. It is equivalent to show that $f_\infty^2 h' f_\infty = h'$. Let $a = a_0 a_1 a_2 \cdots$ be a sequence in $\mathfrak{B}$, and let $i_0 \geq 0$ be the first value for which $a_{i_0} = 0$. Then $(h' f_\infty(a))_0 = 0$, $(h' f_\infty(a))_1 = a_0 + 1$, $(h' f_\infty(a))_{i_0+2} = a_{i_0+1} + 1$, and $(h' f_\infty(a))_{i+1} = a_i$ for all other i. Applying f_∞ to $h' f_\infty(a)$ results in $(1, a_0, \ldots, a_{i_0-1}, 0, a_{i_0+1} + 1, a_{i_0+2}, \ldots)$. Now $i_0 + 1$ is the first index i such that $(f_\infty h' f_\infty a)_i = 0$. Applying f_∞ again replaces $(f_\infty h' f_\infty a)_{i_0+2}$ with a_{i+1} and changes the leading 1 to a 0, so that $f_\infty^2 h' f_\infty(a) = h'(a)$.

The proofs of the other claims are similar; in fact, the first claim is trivial, while the latter two claims follow from the first two. □

As a caveat regarding exchanges of infinitely many intervals, we describe the interval exchange F_∞ induced on a vertical segment by the horizontal foliation. We use the horizontal flow in the positive x-direction, in which case F_∞ has the following effect on $\mathfrak{B}$: for each sequence a,

1. find the least $i > 0$ such that $a_i \neq a_0$;
2. replace a_{i-1-2j} with $a_{i-1-2j} + 1$ for all $0 \leq j \leq \lfloor i/2 \rfloor$.

Note that this algorithm fails to define F_∞ on the zero sequence $\bar{0}$; we will see momentarily that $1/3$ does not have a preimage by F_∞, and so we can define

$F_\infty(0) = 1/3$ without compromising the injectivity or semicontinuity of F_∞. The inverse F_∞^{-1} acts on $\mathfrak{B}$ as follows: for each sequence a,

1. find the least $i > 0$ such that $a_i = a_{i-1}$;
2. replace a_{i-1-2j} with $a_{i-1-2j} + 1$ for all $0 \leq j \leq \lfloor i/2 \rfloor$.

This algorithm fails for *two* points in $\mathfrak{B}$, namely $\overline{01} = 1/3$ and $\overline{10} = 2/3$; these have no pre-images by F_∞. Hence we can "fix" F_∞ by defining $F_\infty(0)$ to be either $1/3$ or $2/3$, but the choice is arbitrary. In either case, F_∞ will still not have all of $\mathfrak{B}$ as its image. The points $1/3$ and $2/3$ do form an *attracting cycle* for F_∞^{-1}, however. The special role of $1/3$ and $2/3$ will be useful to keep in mind.

Let $\mathfrak{D} \subset [0, 1)$ denote the set of dyadic rationals in $[0, 1)$—that is, the set of rational numbers of the form $n/2^m$ for some $n, m \in \mathbb{Z}$. $\mathfrak{D}$ sits inside $\mathfrak{B}$ as the set of sequences that are eventually 0. For each $x \in [0, 1)$, let $\mathcal{O}^\pm(x)$ be the orbit of x under $f_\infty^{\pm 1}$.

Lemma 9.11. $\mathfrak{D} = \mathcal{O}^\pm(0) \sqcup \mathcal{O}^\pm(1/2)$.

Another way to state this result is that the union of the forward and backward orbits of a sequence $a \in \mathfrak{D}$ is entirely determined by the parity of the number of 1s in the sequence a. We call $\mathrm{TM}(a) = \sum a_i \in \mathbb{F}_2$ the *Thue–Morse function*: for any particular $a \in \mathfrak{D}$, this sum is finite, and $\mathrm{TM}(a)$ is invariant under f_∞ because two digits are changed from a to $f_\infty(a)$. We also define the *index of* a to be the smallest natural number $\mathrm{Ind}(a) \in \mathbb{N}$ such that $a_i = 0$ for all $i > \mathrm{Ind}(a)$. (Recall that our sequences in $\mathfrak{B}$ start with a_0, and so $\mathrm{Ind}(\bar{0}) = \mathrm{Ind}(1\bar{0}) = 0$.) We will show

that the following table determines which orbit contains $a \in \mathfrak{D} - \{\bar{0}, 1\bar{0}\}$:

		TM(a)	
		0	1
Ind(a)	even	$\mathcal{O}^-(0)$	$\mathcal{O}^+(1/2)$
	odd	$\mathcal{O}^+(0)$	$\mathcal{O}^-(1/2)$

(9.3)

One consequence of the proof will be a quick algorithm for computing the exact value of $n \in \mathbb{Z}$ so that $f_\infty^n(0) = a$ or $f_\infty^n(1/2) = a$.

Proof of Lemma 9.11. Let H be the semigroup of functions $\mathfrak{B} \to \mathfrak{B}$ consisting of words in h' and h''. The map from H to $\mathfrak{B}$ defined by $w \mapsto w(\bar{0})$ induces a set-theoretic bijection between $\mathfrak{D}$ and the quotient of H by the relation $w \sim wh'$. Throughout the proof, we will use the equivalence $\mathfrak{D} \leftrightarrow H/\sim$, by which $(a_0, a_1, \dots, a_{\mathrm{Ind}(a)}, 0, \dots)$ corresponds to the equivalence class of $\eta_0\eta_1\cdots\eta_{\mathrm{Ind}(a)}$, with

$$\eta_i = \begin{cases} h' & \text{if } a_i = 0 \\ h'' & \text{if } a_i = 1 \end{cases}.$$

In particular, $\eta_{\mathrm{Ind(a)}} = h''$ if $\mathrm{Ind}(a) \geq 1$.

Let $a \in \mathfrak{D}$. We proceed by induction on $\mathrm{Ind}(a)$. Direct computation shows that

$$h''h''(\bar{0}) = f_\infty h'(\bar{0}) = f_\infty(\bar{0}) \qquad \text{and} \qquad h'h''(\bar{0}) = f_\infty^{-1}h''(\bar{0}) = f_\infty^{-1}(1\bar{0}),$$

and therefore if $\mathrm{Ind}(a) = 1$, a is in the union of the orbits of $\bar{0}$ and $1\bar{0}$. Now suppose $\mathrm{Ind}(a) \geq 2$, and let $w = \eta_0\eta_1\cdots\eta_{\mathrm{Ind}(a)-1}h''$ be the corresponding word in $\mathfrak{H}$. Using the above computations, we can rewrite the effect of w on $\bar{0}$ in the following way:

$$w(\bar{0}) = \begin{cases} \eta_0\eta_1\cdots\eta_{\mathrm{Ind}(a)-2} f_\infty h'(\bar{0}) & \text{if } \eta_{\mathrm{Ind}(a)-1} = h'' \\ \eta_0\eta_1\cdots\eta_{\mathrm{Ind}(a)-2} f_\infty^{-1} h''(\bar{0}) & \text{if } \eta_{\mathrm{Ind}(a)-1} = h' \end{cases}.$$

From Lemma 9.10, we have

$$f_\infty^2 h' = h' f_\infty^{-1} \qquad \text{and} \qquad f_\infty^2 h'' = h'' f_\infty.$$

These relations allow us to move f_∞ to the far left of the word, each time exchanging a power of f_∞ for a power whose absolute value is twice as great, which means we have expressed a as $f_\infty^n(b)$, where $\mathrm{Ind}(b) < \mathrm{Ind}(a)$. Here $|n| = 2^{\mathrm{Ind}(a)-1}$, and the sign of n is determined by the number of 0s among $a_0, \ldots, a_{\mathrm{Ind}(a)-1}$. By induction, we have shown that every point of $\mathfrak{D}$ lies in the union of the orbits of $\bar{0}$ and $1\bar{0}$.

Because $\mathrm{TM}(a)$ is invariant under f_∞, $\bar{0}$ and $1\bar{0}$ are not in the same orbit, and therefore $\mathfrak{D}$ is a disjoint union of these two orbits. □

Remark 9.12. We will need a bit more information about the points of discontinuity of f_∞. These correspond precisely to sequences of the form $11\cdots11\bar{0}$ or $11\cdots1101\bar{0}$ (the initial number of 1s may be zero). From the information in (9.3), we see that the forward and backward orbits of both 0 and 1/2 each contain infinitely many such points.

Lemma 9.13. *Saddle connections are dense in the vertical foliation of* $(X_\infty, \omega_\infty)$. *Every vertical critical trajectory is a saddle connection.*

Proof. Let $x \in \mathfrak{D}$, and consider the point $(x, 0)$ on the boundary of the unit square. If x is not already a point of discontinuity of f_∞, then by Lemma 9.11 and Remark 9.12, there exist $m, n > 0$ such that $f_\infty^{-m}(x)$ and $f_\infty^n(x)$ are points of discontinuity of f_∞. Because f_∞ is determined by the vertical flow, this means there is a vertical saddle connection passing through $(x, 0)$ and connecting $(f_\infty^{-m}(x), 0)$ to $(f_\infty^n(x), 0)$. If x is a point of discontinuity of f_∞, then so is $f_\infty(x)$, and there is a vertical saddle connection from $(x, 0)$ to $(f_\infty(x), 1)$. □

The proof shows, moreover, that the union of the vertical critical trajectories contains precisely those points that have representatives in Figure 9.5 with a dyadic rational x-coordinate.

For clarity in the proof of the next lemma, we introduce the notion of a *germ of a singularity* on a locally Euclidean surface M. This is a sequence $\mathfrak{g}$ of open sets $U_0 \supset U_1 \supset U_1 \supset \cdots$ such that $\bigcap U_i = \varnothing$, and, for each i, $\partial U_i \subset M$ is a connected 1-manifold (i.e., homeomorphic to either S^1 or $\mathbb{R}$) and either U_i is simply connected, or U_i is conformally equivalent to a punctured disk, or $\pi_1(U_i)$ is infinitely generated. In addition, we require that if any $U_i \in \mathfrak{g}$ is simply connected, then for every $\ell > 0$, there exists $\varepsilon > 0$ such that U_i contains an embedded curve of constant curvature $1/\varepsilon$ and length ℓ; we say in this case that $\mathfrak{g}$ is the germ of an *infinite-angle* singularity. In the case that each U_i is a punctured disk, we say that $\mathfrak{g}$ is the germ of a *cone-type* singularity. In the remaining case, we call $\mathfrak{g}$ the germ of an *essential*, or *end-type*, singularity. A germ of a singularity is *regular* if either it is of cone-type or, for some $U_i \in \mathfrak{g}$ and some ε, there is isometric embedding of $\mathbb{R}$ into U_i as a curve with constant curvature $1/\varepsilon$. (In [36], the authors call a flat surface *tame* if all of its singularities are regular.) An infinite-angle singularity that is not regular is a *spire*. If $\gamma : (0, T) \to M$ (for some $0 < T \leq \infty$) is a critical trajectory on M (i.e., $\gamma(t)$ leaves every compact subset of M as $t \to 0$), then γ *emanates* from a germ of a singularity $\mathfrak{g}$ if, for every i and every $\delta < T$, $U_i \in \mathfrak{g}$ contains points of $\gamma((0, \delta))$.

A great deal of theory about germs of singularities remains to be established, but for our purposes here we only need the following observation:

Proposition 9.14. *An affine homeomorphism of a locally Euclidean surface M sends a germ of a singularity to a germ of a singularity of the same type.*

Lemma 9.15. *The vertical direction of $(X_\infty, \omega_\infty)$ is not affinely equivalent to any other direction on $(X_\infty, \omega_\infty)$.*

Proof. Let $\mathcal{F}_v$ be the vertical foliation of $(X_\infty, \omega_\infty)$, and let $\mathcal{F}_\theta$ be the foliation in some other direction θ. Assume there exists some $\varphi \in \mathrm{Aff}(X_\infty, \omega_\infty)$ that sends θ to the vertical direction. Let L be the critical leaf of $\mathcal{F}_\theta$ emanating from $(0, 2/3)$ in Figure 9.5. Then $\varphi(L)$ must be a critical trajectory in the vertical direction, which means it must be a saddle connection. By composing φ with some power of ψ_∞ and ρ_∞, if necessary, we may assume $\varphi(L)$ is the saddle connection L_0 from $(0, 0)$ to $(0, 1/2)$.

Let $\mathfrak{g}$ be the germ of a spire singularity from which L_0 emanates (e.g., take each $U_i \in \mathfrak{g}$ to be the union of open semi-circles of decreasing radius centered at the right endpoints of the segments C_n'). By the previous proposition, φ^{-1} must send $\mathfrak{g}$ to the germ of a spire singularity from which L emanates. However, no such germ exists: in order to permit arbitrarily long curves of constant non-zero curvature to intersect L arbitrarily close to $(0, 2/3)$, the open sets needed must have infinitely generated fundamental group. This gives us a contradiction, from which we conclude the desired result. □

Now we are ready to prove the main theorem of this section.

Proof of Theorem 9.9. By Lemma 9.15, any affine homeomorphism φ of $(X_\infty, \omega_\infty)$ must preserve the vertical direction. Because it must preserve the set of saddle connections, and the lengths of the vertical saddle connections are all powers of 2, the derivative of φ must act on the vertical direction by multiplication by $\pm 2^n$ for some $n \in \mathbb{Z}$. By composing φ with a power of ψ_∞ and ρ_∞, if necessary, we may assume that φ is orientation-preserving and the derivative of φ is the identity

in the vertical direction. Note that, because the area of $(X_\infty, \omega_\infty)$ is finite, the derivative of φ must lie in $\mathrm{SL}_2(\mathbb{R})$, which implies that its only eigenvalue is 1.

Thus φ is either a translation automorphism or a parabolic map. The latter is impossible because $(X_\infty, \omega_\infty)$ does not have any cylinders in the vertical direction. The existence of non-trivial translation automorphisms is ruled out directly, for example by observing that each vertical saddle connection has only one other of the same length (its image by ρ_∞), and no translation automorphism can carry one to the other. Therefore the original map φ was a product of a power of ψ_∞ and ρ_∞, and the result is proved. $\square$

APPENDIX A

FROM THE TOP: $G = 1$, $G = 2$

We have already extended the family of Arnoux–Yoccoz surfaces (X_g, ω_g) to the index $g = \infty$. In this appendix we indicate what happens we when extend the construction to create (X_1, ω_1) and (X_2, ω_2) so that the sequence (X_g, ω_g) is defined for all indices $1 \leq g \leq \infty$.

The defining equation for α in the case $g = 1$ is $\alpha = 1$. The corresponding surface is a torus, formed from the unit square by the usual top-bottom and left-right identifications. Hence $(X_1, \omega_1) = (\mathbb{C}/(\mathbb{Z} \oplus i\mathbb{Z}), \mathrm{d}z)$ and ψ_1 is the identity map.

In the case $g = 2$, we get the equation $\alpha + \alpha^2 = 1$, which means that $\alpha = (\sqrt{5} - 1)/2$ is the inverse of the golden ratio. Beginning with the unit square, a single square of side length $1 - \alpha = \alpha^2$ is removed from the upper right corner. Two slits are made, one from $(\alpha/2, 0)$ to $(\alpha/2, 1)$ and the other from $((1+\alpha)/2, 0)$ to $((1 + \alpha)/2, \alpha)$, thereby cutting the square into three separate pieces. After the usual identifications are made, following the procedure of §9.2, the result is a disconnected pair of tori. This is to be expected: the corresponding interval exchange map f_2 is reducible. Viewed on the circle $[0, 1]/\{0 \sim 1\}$, it splits into two interval exchanges, each of which swaps a pair of segments whose lengths are in the golden ratio. The pair of tori taken together admits a pseudo-Anosov homeomorphism ψ_2 with expansion constant $1/\alpha = (1 + \sqrt{5})/2$, which in the process exchanges the components.

Genus 2 is not entirely absent in this picture, however. If we shorten the height of the first slit to $1 - \varepsilon$ and that of the second slit to $\alpha - \varepsilon$, then the

same identifications are possible, and we obtain a connected sum of the two tori, resulting in two cone points of angle 4π. As $\varepsilon \to 0$, the two cone points collapse into a single point, which becomes a marked point on each of the two tori. Thus we can think of (X_2, ω_2) as a degenerate genus 2 surface.

Because (X_2, ω_2) is not connected, we adopt the convention that the affine group $\mathrm{Aff}(X_2, \omega_2)$ only consists of affine self-maps whose derivative is constant. The orientation-reversing map $\rho_2 \in \mathrm{Aff}(X_2, \omega_2)$ exchanges the components. By composing any $\varphi \in \mathrm{Aff}(X_2, \omega_2)$ with ρ_2 or ψ_2, if necessary, we may assume that φ is orientation-preserving and also preserves the components of X_2. The orientation-preserving affine group of a torus with a marked point is $\mathrm{SL}_2(\mathbb{Z})$; in this special case, the derivative homomorphism is an isomorphism. Thus, to compute the remainder of $\mathrm{Aff}(X_2, \omega_2)$, we wish to find the intersection of the affine groups of the two components. Let

$$M_1 = \begin{pmatrix} 1 & -\alpha \\ \alpha & 1 \end{pmatrix} \qquad \text{and} \qquad M_2 = \begin{pmatrix} \alpha & -1 \\ 1 & \alpha \end{pmatrix}.$$

Following a certain normalization, the two components of X_2 have the columns of M_1 and M_2 for their respective homology bases. Then we want to determine

$$(M_1 \cdot \mathrm{SL}_2(\mathbb{Z}) \cdot M_1^{-1}) \cap (M_2 \cdot \mathrm{SL}_2(\mathbb{Z}) \cdot M_2^{-1})$$

or, equivalently, $(M_2^{-1} M_1 \cdot \mathrm{SL}_2(\mathbb{Z}) \cdot M_1^{-1} M_2) \cap \mathrm{SL}_2(\mathbb{Z})$. We have

$$M_1^{-1} M_2 = (M_2^{-1} M_1)^\top = \frac{\alpha}{2-\alpha} \begin{pmatrix} 2 & -1 \\ 1 & 2 \end{pmatrix}$$

and we want to find the quadruples of integers (X, Y, Z, W) with $XW - YZ = 1$ such that the following is in $\mathrm{SL}_2(\mathbb{Z})$:

$$M_2^{-1} M_1 \begin{pmatrix} X & Y \\ Z & W \end{pmatrix} M_1^{-1} M_2 = \frac{1}{5} \begin{pmatrix} 4X + 2(Y+Z) + W & 4Y + 2(W-X) - Z \\ 4Z + 2(W-X) - Y & 4W - 2(Y+Z) + X \end{pmatrix}.$$

That is, each of the entries in the final product must be congruent to 0 modulo 5. This is a necessary and sufficient condition. All four entries yield the same linear condition $X + 3Y + 3Z + 4W \equiv 0 \bmod 5$, which is satisfied in particular if $X \equiv W \equiv 1$ and $Y \equiv Z \equiv 0 \bmod 5$. Thus the Veech group of (X_2, ω_2) contains a copy of the principle 5-congruence subgroup of $\mathrm{SL}_2(\mathbb{Z})$; therefore it is a lattice in $\mathrm{SL}_2(\mathbb{R})$.

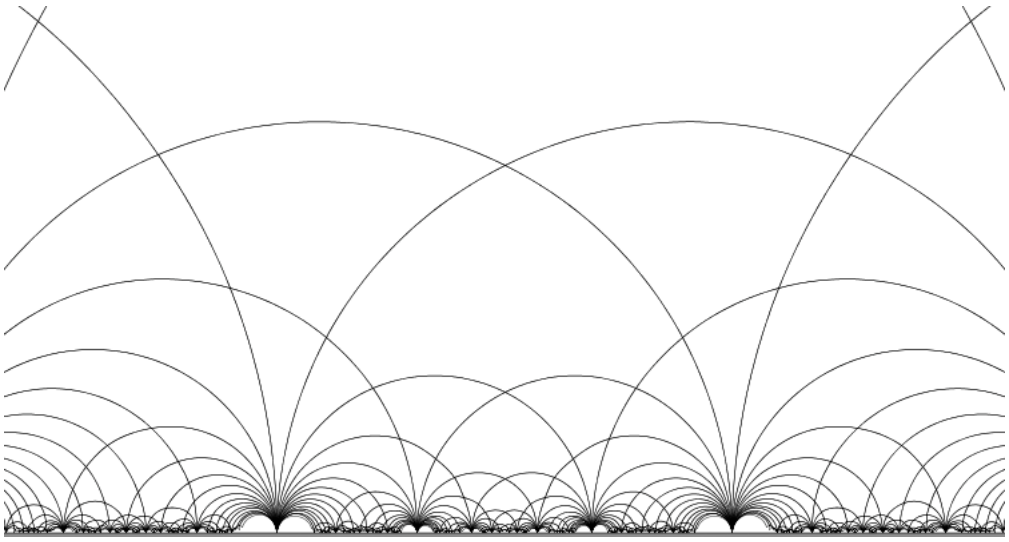

Figure A.1: Iso-Delaunay tessellation for (X_2, ω_2), the "genus 2 Arnoux–Yoccoz surface"; note that this is actually a superposition of two Farey tessellations, one corresponding to each component of X_2. (See Example 4.17.)

APPENDIX B

EQUATIONS FOR THE $G = 3$ SURFACE AND RELATED SURFACES

The contents of this appendix are to appear as an article [6] in the proceedings of the 2008 Ahlfors–Bers Colloquium. To avoid conflicts of notation within this appendix, we will refer to what has been denoted (X_3, ω_3) simply as $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$. The Veech group and $\mathrm{SL}_2(\mathbb{R})$-orbit closure of this surface were studied by Hubert–Lanneau–Möller in [23].

B.1 Delaunay polygons of the genus 3 Arnoux–Yoccoz surface

Let $\alpha \approx 0.543689$ be the real root of the polynomial $x^3 + x^2 + x - 1$. Let S_0 be the square with vertex set $\{(0,0), (\alpha^2, \alpha), (\alpha^2 - \alpha, \alpha^2 + \alpha), (-\alpha, \alpha^2)\}$, and let T_0 be the trapezoid with vertex set $\{(0,0), (1-\alpha, 1-\alpha), (1-\alpha-\alpha^2, 1), (-\alpha, \alpha^2)\}$. We form a flat surface $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$ from two copies of S_0 and four copies of T_0: reflecting S_0 across either a horizontal or vertical axis yields the same square S_1 (up to translation); we denote by $T_{1,0}$, $T_{0,1}$, and $T_{1,1}$ the reflections of T_0 across a vertical axis, across a horizontal axis, and across both, respectively. (In fact, $T_{1,0}$, $T_{0,1}$, and $T_{1,1}$ are all rotations of T_0 by multiples of $\pi/2$, but this description via reflections will be invariant under horizontal and vertical scaling, i.e., the Teichmüller geodesic flow.) Identify the long base of T_0 with the long base of $T_{1,1}$, as well as their short bases; do the same with $T_{1,0}$ and $T_{0,1}$. Each remaining side of a trapezoid is parallel to exactly one side of S_0 or S_1; identify by translations those sides which are parallel. (See Figure B.1.)

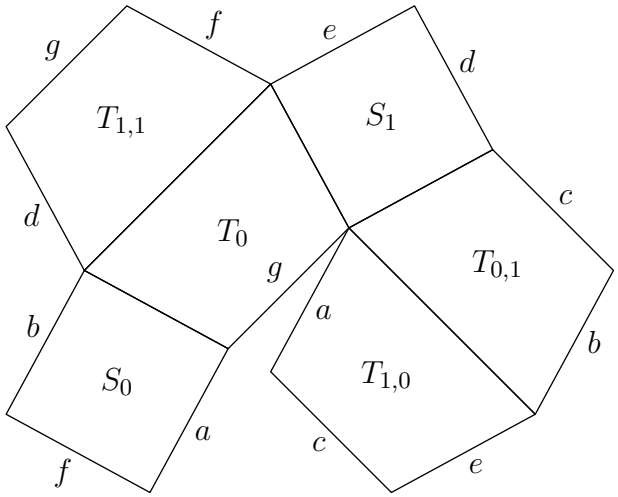

Figure B.1: The decomposition of $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$ into its Delaunay polygons. Edges with the same label are identified by translation.

The resulting flat surface $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$ has genus 3 and two singularities each with cone angle 6π. The images of S_0 and T_0 are the Delaunay polygons of ω_{AY}. X_{AY} is hyperelliptic; the hyperelliptic involution $\tau : X_{\mathrm{AY}} \to X_{\mathrm{AY}}$ is evident in Figure B.1 as rotation by π around the centers of the squares and the midpoints of the edges joining two trapezoids; these six points together with the cone points are therefore the Weierstrass points of the surface. (See Figure B.3 for the quotient of X_{AY} by τ.) Moreover, ω_{AY} is odd with respect to τ, i.e., $\tau^*\omega_{\mathrm{AY}} = -\omega_{\mathrm{AY}}$.

The pseudo-Anosov diffeomorphism ψ_{AY} constructed by Arnoux–Yoccoz scales the surface by a factor of $1/\alpha$ in the horizontal direction and by α in the vertical direction. In Figure B.2 we show the result of applying this affine map to Figure B.1, along with the new Delaunay edges. Two of the trapezoids—having the orientations of $T_{1,1}$ and $T_{1,0}$—are clearly visible; the squares and the other two trapezoids are constructed from the remaining triangles.

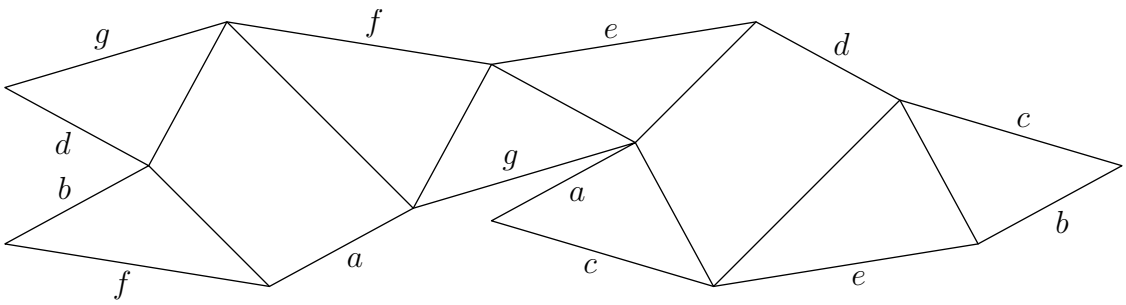

Figure B.2: The result of applying the Arnoux–Yoccoz pseudo-Anosov diffeomorphism to ω_{AY}. The original shapes of S_0 and T_0 and their copies can be reconstructed by matching edges.

B.2 X_{AY} as a cover of $\mathbb{RP}^2$

The reflections applied to S_0 and T_0 in §B.1 induce a pair of orientation-reversing involutions without fixed points on X_{AY}. These can be visualized (as in Figure B.1) as "glide-reflections", one along a horizontal axis and the other along a vertical axis. Both exchange S_0 and S_1. Let σ_1 be the involution that exchanges T_0 and $T_{1,0}$; i.e., its derivative is reflection across the horizontal axis. Let σ_2 be the involution that exchanges T_0 and $T_{0,1}$; i.e., its derivative is reflection across the vertical axis. The product of σ_1 and σ_2 is the hyperelliptic involution τ, and neither sends any point of X_{AY} to its image by τ. They therefore descend to a single involution σ on $\mathbb{CP}^1$ without fixed points. The quotient of $\mathbb{CP}^1$ by σ is homeomorphic to $\mathbb{RP}^2$.

In fact, the presence of σ_1 and σ_2 is implicit in the work of Arnoux–Yoccoz. The original paper [2] begins with a measured foliation of $\mathbb{RP}^2$ with one "tripod" (a singular point of valence three) and three "thorns" (singular points of valence 1), which is then lifted to the genus 3 example we have described. In Figure B.3, right, we illustrate $\mathbb{RP}^2$ as the quotient of X_{AY} by the group generated by σ_1 and σ_2. In Figure B.3, left, we see $\mathbb{CP}^1$, on which σ acts again by a "glide-reflection", which

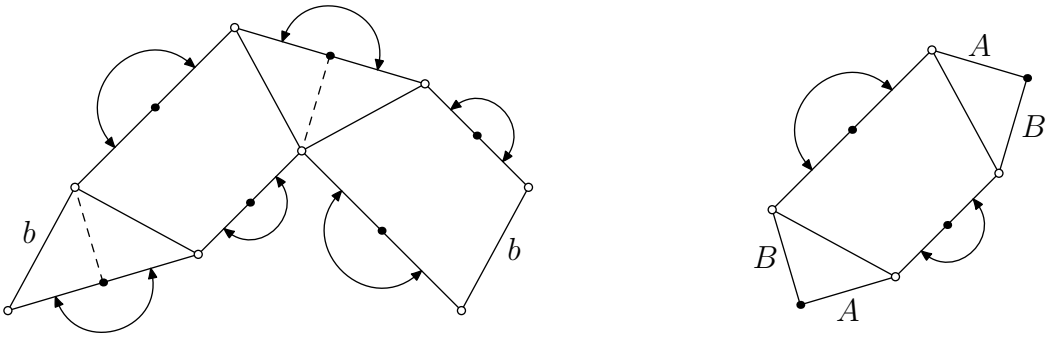

Figure B.3: Quotient surfaces of X_{AY}. LEFT: $\mathbb{CP}^1$ as the quotient of X_{AY} by $\langle\tau\rangle$. The edges marked b are identified by translation. RIGHT: $\mathbb{RP}^2$ as the quotient of X_{AY} by the group $\langle\sigma_1, \sigma_2\rangle$. Edges with the same label are identified by a glide reflection along either a horizontal or a vertical axis. Dashed lines in the left picture indicate preimages of the segments labeled B on the right.

is the sheet exchange for the cover $\mathbb{CP}^1 \to \mathbb{RP}^2$. In both pictures we have drawn vertices that become tripods as open circles, and the vertices that become thorns as filled-in circles. The vertical foliation of the surface on the right of Figure B.3 is the starting point of [2].

In §B.3 we will show that both $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$ and another affinely equivalent surface have *real structures* (orientation-reversing involutions whose fixed-point set is 1-dimensional) that are not evident in the original construction. These additional structures will allow us to write equations for the surfaces and fit them into families of flat surfaces with a common group of isometries. In §B.4 we will transfer these results to genus 2 quadratic differentials. In §B.5 we will conclude by showing that we have found all the surfaces that are obtained by applying the geodesic flow to $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$ and have real structures.

B.3 Two families of surfaces

B.3.1 Labeling the Weierstrass points of X_{AY}

As before, we denote the hyperelliptic involution of X_{AY} by τ, and we let σ_1 and σ_2 be the involutions described in §B.2, with $\sigma : \mathbb{CP}^1 \to \mathbb{CP}^1$ the involution covered by both σ_1 and σ_2.

The purpose of this section is to show the following.

Theorem B.1. *The surface $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$ belongs to a family $(X_{t,u}, \omega_{t,u})$, with $t > 1$ and $u > 0$, such that $X_{t,u}$ has the equation*

$$y^2 = x(x-1)(x-t)(x+u)(x+tu)(x^2+tu), \tag{B.1}$$

and $\omega_{t,u}$ is a multiple of $x\,\mathrm{d}x/y$ on $X_{t,u}$.

Each of the surfaces in Theorem B.1 has a pair of real structures ρ_1 and ρ_2 whose product is again τ, and which therefore descend to a single real structure ρ on $\mathbb{CP}^1$. Any product of the form $\rho_i\sigma_j$ $(i, j \in \{1, 2\})$ is a square root of τ, and therefore the group generated by $\{\sigma_1, \sigma_2, \rho_1, \rho_2\}$ is the dihedral group of the square. We will exhibit these isometries in our presentation of $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$. In §B.3.3 we will look at surfaces in this family that have additional symmetries.

Let $\varpi : X_{\mathrm{AY}} \to \mathbb{CP}^1$ be the degree 2 map induced by τ, i.e., $\varpi \circ \tau = \varpi$. We can normalize ϖ so that the zeroes of ω_{AY} are sent to 0 and ∞, and the midpoint of the short edge between T_0 and $T_{1,1}$ is sent to 1. We wish to find the images of the remaining Weierstrass points, so that we can write an affine equation for X_{AY} in the form $y^2 = P(x)$, where P is a degree 7 polynomial with roots at 0

and 1. Hereafter we assume that ϖ is the restriction to X_{AY} of the coordinate projection $(x, y) \mapsto x$. Consequently, we may consider each Weierstrass point as either a point $(w, 0)$ that solves $y^2 = P(x)$ or simply as a point w on the x-axis.

Each of the real structures ρ_1 and ρ_2 has a fixed-point set with three components: in one case, say ρ_1, the real components are the line of symmetry shared by T_0 and $T_{1,1}$, and the two bases of $T_{1,0}$ and $T_{0,1}$. The fixed-point set of ρ_2 is then the union of the corresponding lines in the orthogonal direction. Because ρ_1 and ρ_2 fix the points 0, 1, and ∞, ρ fixes the real axis; therefore ρ is simply complex conjugation.

With this normalization, the involution σ on $\mathbb{CP}^1$ exchanges 0 and ∞ and preserves the real axis; therefore σ has the form $x \mapsto -r/\overline{x}$ for some real $r > 0$.

Let $s = (s, 0)$ be the center of S_0. Then $\rho_1(s) = \rho_2(s) = \sigma_1(s) = \sigma_2(s)$ is the center of S_1, which implies $\rho(s) = \sigma(s)$, i.e., $\overline{s} = -r/\overline{s}$. The solutions to this equation are $\pm i\sqrt{r}$. By considering the location of the fixed-point sets of ρ_1 and ρ_2, we see that the image of S_0 by ϖ lies in the upper half-plane; therefore $s = i\sqrt{r}$, and $-i\sqrt{r}$ is the center of S_1.

Let t be the midpoint of the long edge of T_0. Applying σ_1 or σ_2 shows that the midpoint of the long edge of $T_{1,0}$ is at $-r/t$.

We already know that 1 is the center of the short edge of T_0. Since the short edge of $T_{0,1}$ is the image of this edge by σ_1 or σ_2, the midpoint of this edge must be at $-r$.

To simplify notation, let us make the substitution $u = r/t$, so that $r = tu$ (hence σ has the form $\sigma(x) = -tu/\overline{x}$). Thus X_{AY} has the equation (B.1) for some

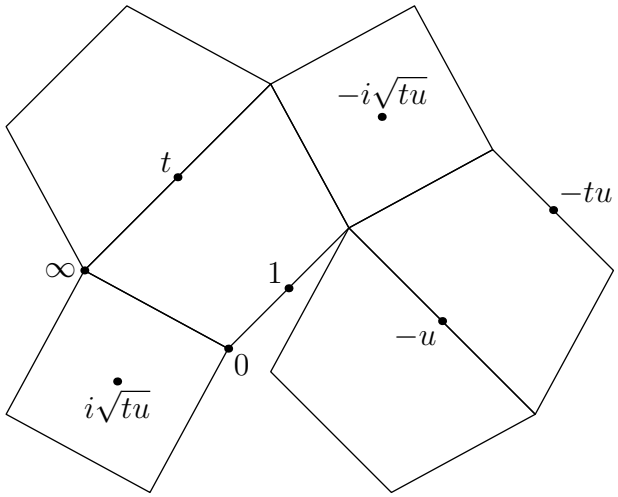

Figure B.4: The Weierstrass points of X_{AY}, following normalization ($t > 1$, $u > 0$). The real structures ρ_1 and ρ_2 appear as reflections in the lines of slope ± 1.

$(t, u) = (t_{\mathrm{AY}}, u_{\mathrm{AY}})$. Furthermore, ω_{AY} is the square root of a quadratic differential on $\mathbb{CP}^1$ with simple zeroes at 0 and ∞ and simple poles at 1, t, $-u$, $-tu$, and $\pm i\sqrt{tu}$. There is therefore some complex constant c such that

$$\omega_{\mathrm{AY}}{}^2 = \varpi^* \left(\frac{cx}{(x-1)(x-t)(x+u)(x+tu)(x^2+tu)} \, \mathrm{d}x^2 \right),$$

i.e., $\omega_{\mathrm{AY}} = \pm\sqrt{c}\, x \, \mathrm{d}x/y$. This establishes Theorem B.1.

B.3.2 Integral equations

To find t_{AY} and u_{AY} requires solving a system of equations involving hyperelliptic integrals, which we establish in this section using relative periods of ω_{AY}. Choose a square root of

$$f_{t,u}(x) = \frac{x}{(x-1)(x-t)(x+u)(x+tu)(x^2+tu)}$$

in the open first quadrant such that its extension $\sqrt{f_{t,u}(x)}$ to the complement of $\{1, t, i\sqrt{tu}\}$ in the closed first quadrant is positive on the open interval $(0, 1)$. Let

η_0 be the Delaunay edge between S_0 and T_0; $\varpi(\eta_0)$ is then a curve from 0 to ∞ in the first quadrant. Integrating $\sqrt{cf_{t,u}(x)}\,\mathrm{d}x$ on the portion of the first quadrant below $\varpi(\eta_0)$ will then give a conformal map to half of T_0. We will be interested in integrals along the real axis.

The vector from 0 to 1 along the short side of T_0 is $\frac{1}{2}(1-\alpha)(1+i)$, while the the line of symmetry of T_0 from 1 to t gives the vector $\frac{1}{2}(\alpha+\alpha^2)(-1+i)$. Observe that

$$i \cdot (1-\alpha)(1+i) = \alpha \cdot (\alpha+\alpha^2)(-1+i),$$

and therefore

$$i\int_0^1 \sqrt{cf_{t,u}(x)}\,\mathrm{d}x = \alpha\int_1^t \sqrt{cf_{t,u}(x)}\,\mathrm{d}x. \tag{B.2}$$

Similarly, the vector from t to ∞ along the long side of T_0 is $\dfrac{1}{2}(1-\alpha^2)(-1-i)$, and because $1-\alpha^2 = (1+\alpha)(1-\alpha)$, we have

$$-(1+\alpha)\int_0^1 \sqrt{cf_{t,u}(x)}\,\mathrm{d}x = \int_t^\infty \sqrt{cf_{t,u}(x)}\,\mathrm{d}x. \tag{B.3}$$

In both equations we can cancel out the c, which was ever only a global (complex) scaling factor anyway. Now bring i under the square root on the right-hand side of (B.2) in order to make the radicand positive. We thus obtain from (B.2) and (B.3) the system of (real) integral equations

$$\begin{cases} \displaystyle\int_0^1 \sqrt{f_{t,u}(x)}\,\mathrm{d}x = \alpha\int_1^t \sqrt{-f_{t,u}(x)}\,\mathrm{d}x \\ \displaystyle(1+\alpha)\int_0^1 \sqrt{f_{t,u}(x)}\,\mathrm{d}x = -\int_t^\infty \sqrt{f_{t,u}(x)}\,\mathrm{d}x \end{cases} \tag{B.4}$$

whose solution is the desired pair $(t_{\mathrm{AY}}, u_{\mathrm{AY}})$. Using numeric methods, we find

$$t_{\mathrm{AY}} \approx 1.91709843377 \qquad \text{and} \qquad u_{\mathrm{AY}} \approx 2.07067976690.$$

We conjecture that t_{AY} and u_{AY} lie in some field of small degree over $\mathbb{Q}(\alpha)$.

B.3.3 Other exceptional surfaces in this family

An examination of the geometric arguments in §B.3.1 and an application of the principle of continuity to t and u show the following:

Theorem B.2. *Every $(X_{t,u}, \omega_{t,u})$ as in Theorem B.1 can be formed by replacing T_0 in the description from §B.1 with an isosceles trapezoid T, S_0 with the square built on a leg of T, and the copies of T_0 with the rotations of T by $\pi/2$.*

The placement of t and u on $\mathbb{R}$ determines the shape of the trapezoid T, and any isosceles trapezoid may be obtained by an appropriate choice of t and u. In this section, we examine certain shapes that give $(X_{t,u}, \omega_{t,u})$ extra symmetries and determine the corresponding values of t and u. We continue to use τ to denote the hyperelliptic involution of $X_{t,u}$.

Suppose that T is a rectangle. Then there are two orthogonal closed trajectories, running parallel to the sides of T and connecting the centers $\pm i\sqrt{tu}$ of the squares, and either of these can be made into the fixed-point set of a real structure on $X_{t,u}$. The product of these two real structures is again τ, so they descend to a single real structure on $\mathbb{CP}^1$. This real structure exchanges 0 with ∞ and fixes $\pm i\sqrt{tu}$, so it must be inversion in the circle $|x|^2 = tu$. It also exchanges 1 with t, which implies $1 \cdot t = tu$, i.e., $u = 1$. The remaining parameter t is determined by solving the single integral equation

$$\int_0^1 \sqrt{\frac{x}{(x^2-1)(x^2-t^2)(x^2+t)}}\,\mathrm{d}x = \mu \int_1^t \sqrt{\frac{-x}{(x^2-1)(x^2-t^2)(x^2+t)}}\,\mathrm{d}x$$

where 2μ is the ratio of the width of T to its height. Recall that an *origami*, also called a *square-tiled surface*, is a flat surface that covers the square torus with at most one branch point (cf. [40, 16, 54]). By looking at rational values of μ, we have the following result:

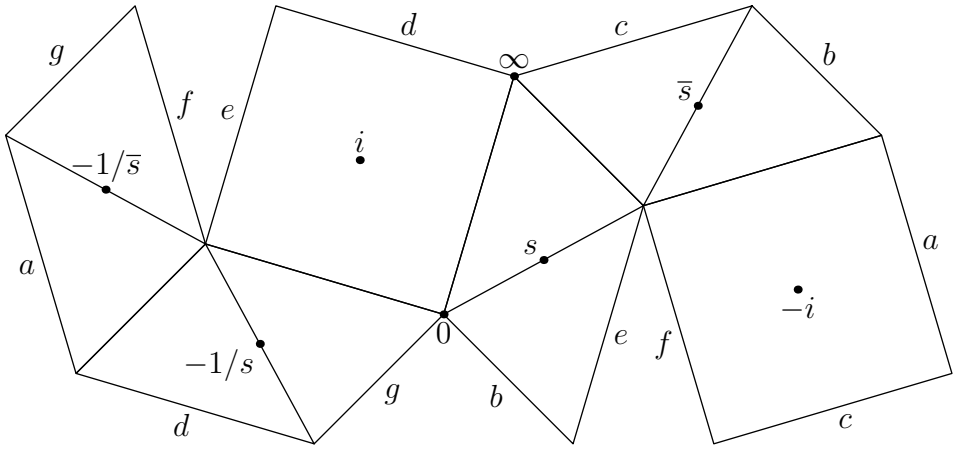

Figure B.5: Another surface in the $\mathrm{GL}_2(\mathbb{R})$-orbit of ω_{AY} with additional real structures. Edges with the same label are identified.

Corollary B.3. *The family* $(X_{t,1}, \omega_{t,1})$ *contains a dense set of origamis.*

These are not the only $(X_{t,u}, \omega_{t,u})$ that are origamis, however. If T is a trapezoid whose legs are orthogonal to each other, then $(X_{t,u}, \omega_{t,u})$ is again an origami.

B.3.4 Second family of surfaces

Conjugating ρ_1 by the pseudo-Anosov element ψ_{AY} guarantees the existence of another orientation-reversing involution in the affine group of ω_{AY}. This element fixes a point "half-way" (in the Teichmüller metric, for instance) between ω_{AY} and its image by ψ_{AY}, lying in the Teichmüller disk of $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$. This surface can be found either by scaling the vertical direction by $\sqrt{\alpha}$ and the horizontal direction by $1/\sqrt{\alpha}$ or, to keep our coordinates in the field $\mathbb{Q}(\alpha)$, just by scaling the horizontal by $1/\alpha$. This surface, which we will denote $(X'_{\mathrm{AY}}, \omega'_{\mathrm{AY}})$, is shown in Figure B.5, along with its Delaunay polygons.

Theorem B.4. *The surface* $(X'_{\mathrm{AY}}, \omega'_{\mathrm{AY}})$ *belongs to a family* (X_s, ω_s)*, with* $\operatorname{Im} s > 0$

and $s \neq i$, such that X_s has the equation

$$y^2 = x(x^2+1)(x-s)(x-\overline{s})(x+1/s)(x+1/\overline{s}), \tag{B.5}$$

and ω_s is a multiple of $x\,\mathrm{d}x/y$ on X_s.

Again, we have two real structures ρ_1' and ρ_2' whose product is the hyperelliptic involution τ. Each of these only has one real component, however: the union of the sides of the parallelograms running parallel to the axis of reflection. The only Weierstrass points that lie on these components are 0 and ∞; the remaining Weierstrass points are the centers of the squares and of the parallelograms. We again let ρ' be the induced real structure on $\mathbb{CP}^1$ and assume that it fixes the real axis (this we can do because we have only fixed the positions of two points on $\mathbb{P}^1$), so that the remaining Weierstrass points come in conjugate pairs.

The fixed-point free involutions σ_1 and σ_2 from §B.2 again preserve the union of the real loci of ρ_1' and ρ_2', and therefore they descend to a fixed-point free involution σ of the form $x \mapsto -r/\overline{x}$, with r real and positive. We have one more free real parameter for normalization, so we can assume $r = 1$. This implies that the centers of the squares are at $\pm i$. Let s be the center of one of the parallelograms; then applying ρ_1' and σ_1 shows that the remaining Weierstrass points are $\overline{s}$, $1/s$, and $1/\overline{s}$. Using developing vectors again, we can find equations that define s, in a manner analogous to finding (B.4).

As an analogue to Theorem B.2, we have:

Theorem B.5. *Every (X_s, ω_s) as in Theorem B.4 can be formed from a parallelogram P, a square built on one side of P, the rotation of P by $\pi/2$, and the images of P and its rotation by reflection across their remaining sides.*

The shape of P is determined by the value of s. If $s = \frac{1}{2}(\sqrt{3} + i)$, then P becomes a square, and we obtain one of the "escalator" surfaces in [28]. More generally, if s is any point of the unit circle, then P is a rectangle, and inversion in the unit circle corresponds to another pair of real structures on X, which are the reflections across the axes of symmetry of P. By considering those rectangular P whose side lengths are rationally related, we have as before:

Corollary B.6. *The family $(X_{e^{i\theta}}, \omega_{e^{i\theta}})$ (with $0 < \theta < \pi/2$) contains a dense set of origamis.*

Another origami appears when P is composed of a pair of right isosceles triangles so that s lies not on the hypotenuse, but on a leg of each.

B.4 Quadratic differentials and periods on genus 2 surfaces

We do not know how to compute the rest of the periods for $X_{t,u}$ or X_s, apart from those of $\omega_{t,u}$ or ω_s, respectively. In this section, however, we consider the periods of certain related genus 2 surfaces, which demonstrate remarkable relations.

Let X be any hyperelliptic genus 3 surface with an abelian differential ω that is odd with respect to the hyperelliptic involution and has two double zeroes. The pair (X, ω) has a corresponding pair (Ξ, q), where Ξ is a genus 2 surface and q is a quadratic differential on Ξ with four simple zeroes. Geometrically, the correspondence may be described as follows: two of the zeroes of ω are at Weierstrass points of X, hence (X, ω^2) covers a flat surface $(\mathbb{CP}^1, \widetilde{q})$ where $\widetilde{q}$ has six poles and two simple zeroes (Figure B.3). Then (Ξ, q) is the double cover of $(\mathbb{CP}^1, \widetilde{q})$ branched at the poles of $\widetilde{q}$. In our cases, the genus 2 surface may be

obtained by cutting along opposite sides of one of the squares in Figure B.1 or B.5, then regluing each of these via a rotation by π to the free edge provided by cutting along the other (cf. [26, 47]).

First we consider the family $(X_{t,u}, \omega_{t,u})$ and the related genus 2 flat surfaces $(\Xi_{t,u}, q_{t,u})$. To be explicit, the defining expressions of both types of surfaces are:

$$X_{t,u} : y^2 = x(x-1)(x-t)(x+u)(x+tu)(x^2+tu), \qquad \omega_{t,u} = \frac{x\,\mathrm{d}x}{y};$$
$$\Xi_{t,u} : y^2 = (x-1)(x-t)(x+u)(x+tu)(x^2+tu), \qquad q_{t,u} = \frac{x\,\mathrm{d}x^2}{y^2}.$$

The order 4 rotation $\rho_1\sigma_1$ of $X_{t,u}$ persists on $\Xi_{t,u}$. Following R. Silhol [42], we find a new parameter a, depending on t and u, so that the Riemann surface

$$\Xi_a : y^2 = x(x^2-1)(x-a)(x-1/a)$$

is isomorphic to $\Xi_{t,u}$. Doing so simply requires a change of coordinates in x, namely

$$\Phi(x) = i\sqrt{tu}\frac{(x-1)}{(x+tu)}.$$

Then $\Phi(1) = 0$, $\Phi(-tu) = \infty$, and $\Phi(\pm i\sqrt{tu}) = \mp 1$. The images of t and u by Φ are

$$a = i\sqrt{\frac{u}{t}}\frac{(t-1)}{(u+1)} \qquad \text{and} \qquad \frac{1}{a} = i\sqrt{\frac{t}{u}}\frac{(1+u)}{(1-t)}.$$

Because $t > 1$ and $u > 0$, a lies on the positive imaginary axis and $1/a$ lies on the negative imaginary axis. The involution ρ becomes reflection across the imaginary axis. The images of 0 and ∞ by Φ are

$$\Phi(0) = \frac{i}{\sqrt{tu}} \qquad \text{and} \qquad \Phi(\infty) = \frac{\sqrt{tu}}{i},$$

so the image of $q_{t,u}$ on Ξ_a is a scalar multiple of

$$\left(x - \frac{i}{\sqrt{tu}}\right)\left(x - \frac{\sqrt{tu}}{i}\right)\frac{\mathrm{d}x^2}{y^2} = \left(x^2 + i\left(\frac{tu-1}{\sqrt{tu}}\right)x + 1\right)\frac{\mathrm{d}x^2}{y^2}$$

These calculations imply that, for each pair (t_0, u_0), there is a one-parameter family of surfaces $(\Xi_{t,u}, q_{t,u})$ such that $\Xi_{t,u}$ is isomorphic to Ξ_{t_0,u_0} while $q_{t,u}$ and q_{t_0,u_0} represent different differentials on the abstract Riemann surface.

Now we apply the same analysis to the second family. This time we are moving from (X_s, ω_s) to (Σ_s, q_s), as defined below:

$$X_s : y^2 = x(x^2+1)(x-s)(x-\overline{s})(x+1/s)(x+1/\overline{s}), \qquad \omega_s = \frac{x\,\mathrm{d}x}{y};$$
$$\Sigma_s : y^2 = (x^2+1)(x-s)(x-\overline{s})(x+1/s)(x+1/\overline{s}), \qquad q_s = \frac{x\,\mathrm{d}x^2}{y^2}.$$

We change coordinates in x using

$$\Psi(x) = i\left(\frac{x-s}{sx+1}\right)$$

so that $\Psi(s) = 0$, $\Psi(-1/s) = \infty$, and $\Psi(\pm i) = \mp 1$. This time we get the curve $y^2 = x(x^2-1)(x-a)(x-1/a)$, where

$$a = \Phi(\overline{s}) = \frac{2\,\mathrm{Im}\,s}{1+|s|^2} \qquad \text{and} \qquad \frac{1}{a} = \Phi\left(-\frac{1}{\overline{s}}\right) = \frac{1+|s|^2}{2\,\mathrm{Im}\,s}.$$

Here we have $0 < a < 1$ and $1/a > 1$; ρ' becomes inversion in the unit circle. The points 0 and ∞ on Σ_s become $\Phi(0) = -is$ and $\Phi(\infty) = i/s$. Again, we find just a one-parameter family of genus 2 Riemann surfaces, each carrying a one-parameter family of quadratic differentials corresponding to distinct surfaces X_s.

In [42], it is shown that the full period matrix for any of the surfaces Ξ_a can be expressed in terms of a single parameter, thanks to the fourfold symmetry of the surface. This parameter is the ratio of $\int_{-1}^{0}\varphi$ and $\int_{0}^{1/a}\varphi$, where $\varphi = \frac{\mathrm{d}x}{y} - \frac{x\,\mathrm{d}x}{y}$. This ratio is real precisely when a lies on the positive imaginary axis, as in our first family, and in these cases the period matrix of Ξ_a is purely imaginary.

B.5 Final remarks

The involutions we have exhibited also act on the Teichmüller disk generated by $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$, and their effects can be seen via the *iso-Delaunay tessellation* shown in Figure B.6. Each element of Γ acts on $\mathbb{H}$ by an isometry, preserving or reversing orientation according to the sign of its determinant. Figure B.6 is symmetric with respect to the central axis (the imaginary axis in $\mathbb{C}$); both σ_1 and σ_2 yield elements of Γ that reflect $\mathbb{H}$ across this axis. The hyperbolic element of Γ corresponding to ψ_{AY} fixes the points 0 and ∞ in $\partial\mathbb{H}$ and translates points along the imaginary axis by $z \mapsto z/\alpha^2$. A sequence of concentric circles is visible in the tessellation; these are related by ψ_{AY}, and one is the unit circle, so their radii are all powers of $1/\alpha^2 \approx 3.38$.

There are two kinds of distinguished points on the central axis: ones where two geodesics meet and ones where three geodesics meet. The latter are those whose corresponding surface is isometric to $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$, while the former correspond to $(X'_{\mathrm{AY}}, \omega'_{\mathrm{AY}})$. The real structures ρ_1 and ρ_2 (resp. ρ'_1 and ρ'_2) yield an element of Γ that reflects $\mathbb{H}$ across the unit circle (resp. across the circle $|z| = 1/\alpha$). The order 4 rotations of $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$ and $(X'_{\mathrm{AY}}, \omega'_{\mathrm{AY}})$ are thus visible as the order 2 rotations of $\mathbb{H}$ around these distinguished points.

If any other flat surface on the central axis had real structures, then its symmetries, too, would have to induce a reflection of $\mathbb{H}$ that preserves the tessellation. No such point exists; therefore we have described all the surfaces within the orbit of $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$ under the geodesic flow that demonstrate additional symmetries.

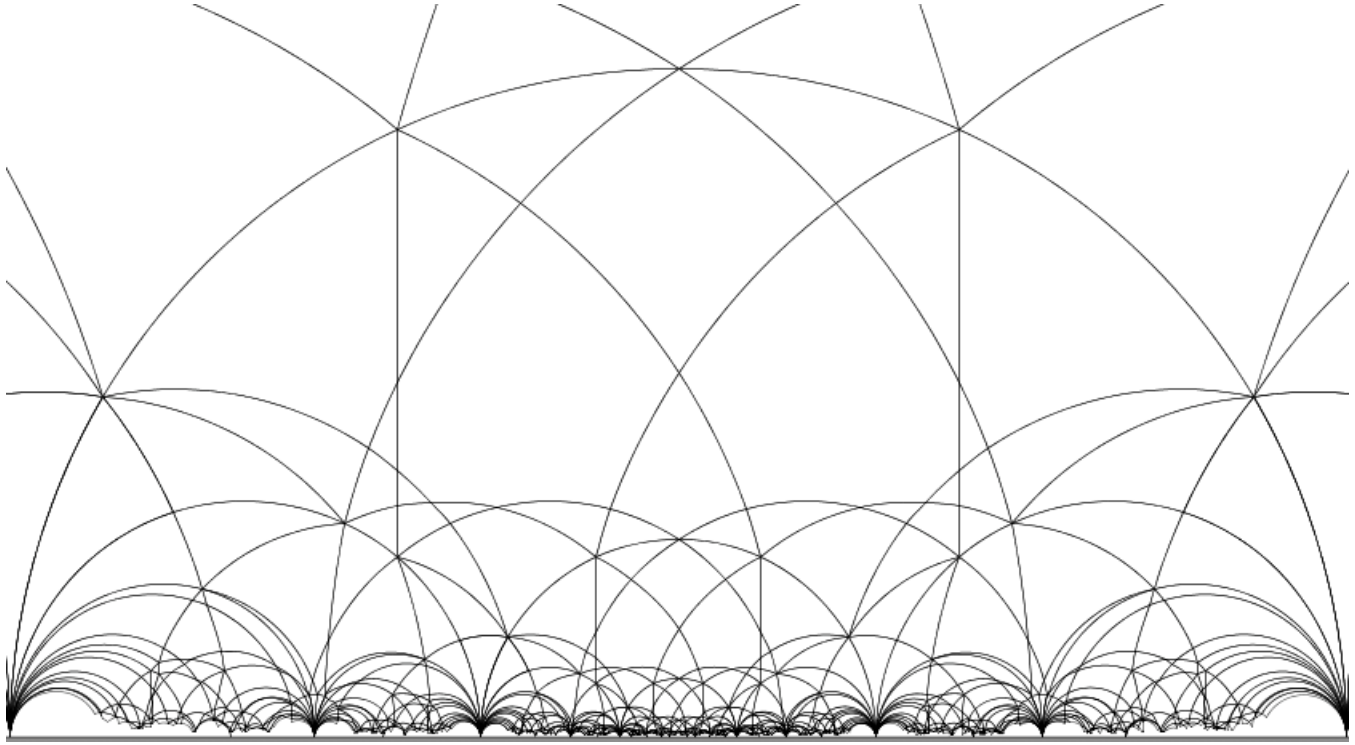

Figure B.6: The iso-Delaunay tessellation of $\mathbb{H}$ arising from $(X_{\mathrm{AY}}, \omega_{\mathrm{AY}})$

BIBLIOGRAPHY

[1] Pierre Arnoux. Un exemple de semi-conjugaison entre un échange d'intervalles et une translation sur le tore. *Bull. Soc. Math. France*, 116(4):489–500, 1988.

[2] Pierre Arnoux and Jean-Christophe Yoccoz. Construction de difféomorphismes pseudo-Anosov. *C. R. Acad. Sc. Paris Sr. I Math.*, 292:75–78, 1981.

[3] Matt Bainbridge. Euler characteristics of Teichmüller curves in genus two. *Geometry & Topology*, 11:1887–2073, 2007.

[4] Christina Birkenhake. Complex structures contained in classical groups. *J. Lie Theory*, 8(1):139–152, 1998.

[5] Alexander I. Bobenko and Boris A. Springborn. A discrete Laplace–Beltrami operator for simplicial surfaces. *Discrete Comput. Geom.*, 38(4):740–756, 2007.

[6] Joshua P. Bowman. Orientation-reversing involutions of the genus 3 Arnoux–Yoccoz surface and related surfaces. To appear, *Proceedings of the Fourth Ahlfors–Bers Colloquium*.

[7] Joshua P. Bowman. Teichmüller geodesics, Delaunay triangulations, and Veech groups. To appear in *Proceedings of the International Workshop on Teichmüller theory and moduli problems*.

[8] Allen Broughton and Christopher Judge. Sub-conics in translation surfaces. In preparation.

[9] Kariane Calta. Veech surfaces and complete periodicity in genus two. *J. Amer. Math. Soc.*, 17:871–908, 2004.

[10] Kariane Calta and John Smillie. Algebraically periodic translation surfaces. *J. Mod. Dyn.*, 2(2):209–248, 2008.

[11] Ana Cannas da Silva. *Lectures on symplectic geometry*, volume 1764 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2001.

[12] Reza Chamanara. Affine automorphism groups of surfaces of infinite type. In *In the tradition of Ahlfors and Bers, III*, volume 355 of *Contemp. Math.*, pages 123–145. Amer. Math. Soc., Providence, RI, 2004.

[13] Reza Chamanara, Frederick P. Gardiner, and Nikola Lakic. A hyperelliptic realization of the horseshoe and baker maps. *Ergodic Theory Dynam. Systems*, 26(6):1749–1768, 2006.

[14] Boris N. Delaunay. Sur la sphère vide. *Bull. Acad. Sci. USSR*, (6):793–800, 1934.

[15] Clifford J. Earle and Frederick P. Gardiner. Teichmüller disks and Veech's $\mathcal{F}$-structures. *Contemp. Math.*, 201:165–189, 1997.

[16] Alex Eskin and Andrei Okounkov. Asymptotics of numbers of branched coverings of a torus and volumes of moduli spaces of holomorphic differentials. *Inv. Math.*, 145(1):59–103, 2001.

[17] Shmuel Friedland and Pedro J. Freitas. Revisiting the Siegel upper half plane. I and II. *Linear Algebra Appl.*, 376:19–67, 2004.

[18] Eugene Gutkin and Chris Judge. Affine mappings of translation surfaces: Geometry and arithmetic. *Duke Math. J.*, 103:191–213, 2000.

[19] W. Patrick Hooper. Another Veech triangle. Preprint, available online.

[20] John H. Hubbard. *Teichmüller theory and applications to geometry, topology, and dynamics. Vol. 1: Teichmüller theory.* Matrix Editions, Ithaca, NY, 2006.

[21] John H. Hubbard and Howard Masur. Quadratic differentials and foliations. *Acta Math.*, 142:221–274, 1979.

[22] Pascal Hubert and Erwan Lanneau. Veech groups without parabolic elements. *Duke Math. J.*, 133(2):335–346, 2006.

[23] Pascal Hubert, Erwan Lanneau, and Martin Möller. The Arnoux–Yoccoz Teichmüller disc. *Geom. Funct. Anal.*, 18(6):1988–2016, 2009.

[24] C. Indermitte, Th. M. Liebling, M. Troyanov, and H. Clémençon. Voronoi diagrams on piecewise flat surfaces and an application to biological growth. *Theoret. Comput. Sci.*, 263(1-2):263–274, 2001. Combinatorics and computer science (Palaiseau, 1997).

[25] Richard Kenyon and John Smillie. Billiards in rational-angled triangles. *Comment. Math. Helv.*, 75:65–108, 2000.

[26] Erwan Lanneau. Hyperelliptic components of the moduli spaces of quadratic differentials with prescribed singularities. *Comm. Math. Helv.*, 79:471–501, 2004.

[27] Christopher J. Leininger. On groups generated by two positive multi-twists: Teichmüller curves and Lehmer's number. *Geom. and Top.*, 8:1301–1359, 2004.

[28] Samuel Lelièvre and Robert Silhol. Multi-geodesic tessellations, fractional Dehn twists and uniformization of algebraic curves. arXiv:math/0702374v1.

[29] Howard Masur and John Smillie. Hausdorff dimension of sets of nonergodic measured foliations. *Ann. Math.*, 134:455–543, 1991.

[30] Curtis T. McMullen. Billiards and Teichmüller curves on Hilbert modular surfaces. *J. Amer. Math. Soc.*, 16:857–885, 2003.

[31] Curtis T. McMullen. Teichmüller curves in genus two: Discriminant and spin. *Math. Ann.*, 333:87–130, 2005.

[32] Curtis T. McMullen. Teichmüller curves in genus two: The decagon and beyond. *J. reine angew. Math.*, 582:173–200, 2005.

[33] Curtis T. McMullen. Prym varieties and Teichmüller curves. *Duke Math. J.*, 133:569–590, 2006.

[34] Curtis T. McMullen. Teichmüller curves in genus two: Torsion divisors and ratios of sines. *Inv. math.*, 165:651–672, 2006.

[35] Robert C. Penner. The decorated Teichmüller space of punctured surfaces. *Comm. Math. Phys.*, 113(2):299–339, 1987.

[36] Piotr Przytycki, Gabriela Schmithüsen, and Ferran Valdez. Veech groups of Loch Ness monsters. arXiv:math/0906.5268v1.

[37] Jan-Christoph Puchta. On triangular billiards. *Comment. Math. Helv.*, 76(3):501–505, 2001.

[38] Igor Rivin. Euclidean structures on simplicial surfaces and hyperbolic volume. *Ann. Math.*, 139:553–580, 1994.

[39] Rubí E. Rodríguez and Víctor González-Aguilera. Fermat's quartic curve, Klein's curve and the tetrahedron. In *Extremal Riemann surfaces (San Francisco, CA, 1995)*, volume 201 of *Contemp. Math.*, pages 43–62. Amer. Math. Soc., Providence, RI, 1997.

[40] Gabriela Schmithüsen. An algorithm for finding the Veech group of an origami. *Experiment. Math.*, 13(4):459–472, 2004.

[41] Robert Silhol. Hyperbolic lego and algebraic curves in genus 2 and 3. In *Complex manifolds and hyperbolic geometry (Guanajuato, 2001)*, volume 311 of *Contemp. Math.*, pages 313–334. Amer. Math. Soc., Providence, RI, 2002.

[42] Robert Silhol. Genus 2 translation surfaces with an order 4 automorphism. In *The geometry of Riemann surfaces and abelian varieties*, volume 397 of *Contemp. Math.*, pages 207–213. Amer. Math. Soc., Providence, RI, 2006.

[43] John Smillie and Barak Weiss. Characterizations of lattice surfaces. To appear in *Inv. Math.*

[44] William P. Thurston. On the geometry and dynamics of diffeomorphisms of surfaces. *Bull. A.M.S.*, 19:417–431, 1988.

[45] William P. Thurston. Shapes of polyhedra and triangulations of the sphere. In *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 511–549 (electronic). Geom. Topol. Publ., Coventry, 1998.

[46] Marc Troyanov. Les surfaces euclidiennes à singularités coniques. *Enseign. Math.*, 32:79–94, 1986.

[47] Sergey Vasilyev. Genus two Veech surfaces arising from general quadratic differentials. arXiv:math/0504180.

[48] William A. Veech. Bicuspid F-structures and Hecke groups. Preprint.

[49] William A. Veech. Teichmüller curves in moduli space, Eisenstein series, and an application to triangular billiards. *Inv. Math.*, 97:553–583, 1989.

[50] William A. Veech. Delaunay partitions. *Topology*, 36(1):1–28, 1997.

[51] Claire Voisin. *Hodge Theory and Complex Algebraic Geometry. I*, volume 76 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, English edition, 2007. Translated from the French by L. Schneps.

[52] Yaroslav B. Vorobets. Planar structures and billiards in rational polygons: the Veech alternative. *Russ. Math. Surv*, 51:779–817, 1996.

[53] A. N. Zemljakov and A. B. Katok. Topological transitivity of billiards in polygons. *Mat. Zametki*, 18(2):291–300, 1975.

[54] Anton Zorich. Square tiled surfaces and Teichmüller volumes of the moduli spaces of abelian differentials. In *Rigidity in dynamics and geometry (Cambridge, 2000)*, pages 459–471. Springer, Berlin, 2002.