Geometry & Topology 24 (2020) 373–408

$\mathrm{GL}_2\mathbb{R}$–invariant measures in marked strata: generic marked points, Earle–Kra for strata and illumination

PAUL APISA

We show that nontrivial $\mathrm{GL}(2,\mathbb{R})$–invariant point markings for translation surfaces in strata of abelian differentials exist only when the translation surface belongs to a hyperelliptic component. As an application, we establish constraints on sections of the universal curve restricted to orbifold covers of subvarieties of the moduli space of Riemann surfaces that contain a Teichmüller disk. We also solve the finite blocking problem for generic translation surfaces.

30F60, 32G15, 37F30

1 Introduction

Let $\mathcal{QM}_g$ be the moduli space of quadratic differentials on genus g Riemann surfaces. The space admits a $\mathrm{GL}(2,\mathbb{R})$ action generated by Teichmüller geodesic flow and complex scalar multiplication. The space $\mathcal{QM}_g$ also admits a $\mathrm{GL}(2,\mathbb{R})$–invariant stratification given by specifying the number of zeros and poles of the quadratic differentials and their orders of vanishing.

A point of $\mathcal{QM}_g$ will be denoted by (X,q), where X is a Riemann surface and q is a quadratic differential. A point (X,q) will be called generic if its $\mathrm{GL}(2,\mathbb{R})$ orbit is dense in the connected component of the stratum containing it. We are interested in studying collections of points P on quadratic differentials (X,q). We will always assume that P does not contain any zeros or poles of q. Say that P is generic if the orbit closure of $(X,q;P)$ is as big as possible, ie

$$\dim_{\mathbb{C}} \overline{\mathrm{GL}(2,\mathbb{R})\cdot(X,q;P)} = |P| + \dim_{\mathbb{C}} \overline{\mathrm{GL}(2,\mathbb{R})\cdot(X,q)}.$$

A nongeneric point is called a periodic point.

A translation surface, which we will denote by (X,ω), is a Riemann surface X together with a holomorphic one-form ω. The one-form determines a quadratic differential by taking its square. The $\mathrm{GL}(2,\mathbb{R})$ action on the space of quadratic differentials preserves

Published: 25 March 2020 DOI: 10.2140/gt.2020.24.373

the locus of translation surfaces. We will now characterize the nongeneric collections of points on generic translation surfaces.

Theorem 1.1 *If (X, ω) is a generic translation surface of genus at least two and P is a nongeneric collection of points then (X, ω) belongs to a hyperelliptic component of a stratum and P contains either a Weierstrass point or two points exchanged by the hyperelliptic involution.*

Remark 1.2 In Apisa and Wright [3], it is shown by separate methods that if (X, q) is a generic quadratic differential that does not belong to a stratum of abelian differentials then the conclusion of Theorem 1.1 holds. Neither result implies the other.

Holomorphic sections of the universal curve In Hubbard [14], it was shown that holomorphic sections of the universal curve over $\mathrm{Teich}_{g,n}$ — the Teichmüller space of genus g Riemann surfaces with n punctures — exist only when $(g, n) = (2, 0)$, in which case the sections mark fixed points of the hyperelliptic involution. Earle and Kra [6] generalized this result by allowing the sections to mark punctured points. They showed that in this more general setting, the only new sections that could arise were in genus two and are given by taking a punctured point and either marking it or its image under the hyperelliptic involution. The following result gives a constraint on holomorphic sections of finite orbifold covers of the universal curve that explicitly identifies a finite collection of points that may be marked by holomorphic sections.

Let C be a subvariety of $\mathcal{M}_{g,n}$ and let Γ be a torsionfree finite-index subgroup of the mapping class group $\mathrm{Mod}_{g,n}$. Let $C(\Gamma)$ be the preimage of C on $\mathrm{Teich}_{g,n}/\Gamma$.

Theorem 1.3 *If C contains the Teichmüller disk generated by the quadratic differential (X, q), then any holomorphic section of the universal curve over $C(\Gamma)$ marks a point p on X that is a pole, zero or periodic point of (X, q).*

Remark 1.4 By Eskin, Filip and Wright [7, Theorem 1.5], there are finitely many periodic points on (X, q) if and only if its holonomy double cover is not a torus cover. Therefore, for such quadratic differentials, Theorem 1.3 shows that a holomorphic section of the universal curve can only mark finitely many points and explicitly describes which points may be marked.

Two applications of Theorems 1.1 and 1.3 Apart from $\mathcal{M}_{g,n}$ itself, other examples of algebrogeometrically interesting varieties that contain a Teichmüller disk are the theta-null divisor (see Müller [27] and Grushevsky and Zakharov [13]), the antiramification locus (see Farkas and Verra [10]) and the Weierstrass divisor (see Cukierman [5]). More examples are listed in Mullane [26] and the Kodaira dimensions of many such loci are computed in Gendron [12]. Each of these examples is the projection of a nonhyperelliptic component $\mathcal{H}$ of some stratum of abelian differentials to $\mathcal{M}_{g,n}$. Therefore, the following is immediate from Theorems 1.1 and 1.3.

Corollary 1.5 *If $\mathcal{C}$ is any of the preceding loci and X is a Riemann surface in $\mathcal{C}$, then there is an abelian differential ω such that (X,ω) belongs to $\mathcal{H}$ and the only points on X that may be marked by a holomorphic section of the universal curve defined on $\mathcal{C}$ are zeros of ω.*

Remark 1.6 In particular, this corollary recovers a classical fact that for any finite-index torsionfree subgroup Γ of the mapping class group, there are no holomorphic sections of the universal curve defined over $\mathrm{Teich}_{g,n}/\Gamma$ unless $g=2$.

Another application of Theorem 1.3 is to sections of the universal curve defined over Hilbert modular surfaces in genus two. In [20; 21; 22; 23], McMullen showed that the closures of Teichmüller disks generated by an abelian differential on a genus two Riemann surface is $\mathcal{M}_2$, is contained in a locus of torus covers, or is a Hilbert modular surface. The periodic points on a torus cover (X,ω) are exactly the preimages of torsion points on the torus. It was shown in Apisa [1] that for any other abelian differential (X,ω) in genus two — except those whose orbit closures project to E_5, ie the Hilbert modular surface parametrizing curves whose Jacobian admits real multiplication by the maximal order in $\mathbb{Q}[\sqrt{5}]$ — that the only periodic points are Weierstrass points.

Corollary 1.7 *The only holomorphic sections of the universal curve defined over genus two Hilbert modular surfaces (excluding E_5) mark Weierstrass points or zeros of the eigenform.*

Remark 1.8 Up to the hyperelliptic involution there is one additional point that may be marked on the universal curve above E_5 and it is described in Kumar and Mukamel [16].

Finite blocking on translation surfaces One motivation for studying marked points on translation surfaces is rational billiards. Given a rational billiard table and two points p and q one may ask whether there is a billiard trajectory from p to q. More ambitiously, one may seek to ascertain whether there is a finite collection of points that block all shots from p to q. To study rational billiards, one often applies the unfolding construction of Katok and Zemlyakov [30] to turn the billiard table into a translation surface where billiard trajectories correspond to straight lines on the translation surface. The previously posed problem then becomes the finite blocking problem: given two points p and q on a translation surface is there a finite collection of points that intersects every straight line from p to q? We will continue to assume that neither p nor q coincide with a singularity of the flat metric.

Theorem 1.9 *A generic translation surface contains a pair of finitely blocked points if and only if it belongs to a hyperelliptic component, in which case the pair consists of a point and its image under the hyperelliptic involution.*

Proof of Theorem 1.9 given Theorem 1.1 By Apisa and Wright [3, Theorem 3.15] and Theorem 1.1, finitely blocked pairs of points only occur when the generic translation surfaces belongs to a hyperelliptic component; and in this case, the pair of points consists of either two Weierstrass points or two points exchanged by the hyperelliptic involution and in both cases the finite blocking set is the collection of Weierstrass points.

Since the collection of translation surfaces represented by strictly convex $2n$–gons with opposite sides identified is open, nonempty and $\mathrm{GL}(2,\mathbb{R})$–invariant in hyperelliptic components of strata, its complement consists of nongeneric translation surfaces. Thus, the generic translation surface may be represented by a strictly convex $2n$–gon with opposite sides identified. The Weierstrass points are the midpoints of the polygon and its edges (and the vertices when n is even). By convexity a Weierstrass point is at most finitely blocked from itself and so the pairs of points finitely blocked from each other are exactly the ones containing a point and its image under the hyperelliptic involution. □

For more on the finite blocking problem, see Lelièvre, Monteil and Weiss [17]. For applications to billiards, see Mirzakhani and Wright [25, Sections 6–7], Apisa and Wright [3, Section 3] and Apisa [1].

Organization The proof of Theorem 1.3 is independent of the rest of the paper and appears in Section 3. The outline of the proof of Theorem 1.1 is given in Section 6 and reduced to two more technical results that are established in the subsequent two sections. The two main tools used in the proof of Theorem 1.1 are the construction of generic horizontally and vertically periodic translation surfaces in every component of every stratum of holomorphic one-forms in Section 4 and results in Section 5 that constrain the positions of periodic points using cylinders.

Acknowledgement The author thanks Alex Eskin for suggesting the problem and thanks Alex Eskin, Alex Wright and Curt McMullen for their insightful comments. He thanks Ronen Mukamel for suggesting the parallel between the main theorems and the work of Hubbard and of Earle and Kra. He thanks Matt Bainbridge for suggesting the connection to the problem of classifying holomorphic sections. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1144082 and upon work supported by the National Science Foundation under Award No. 1803625. The author gratefully acknowledges their support.

2 Background

Let $\mathcal{M}_{g,n}$ denote the moduli space of connected orientable genus g surfaces with n punctures. Let $\Omega\mathcal{M}_{g,n}$ be the moduli space of holomorphic one-forms (equivalently abelian differentials) on Riemann surfaces of genus g with n punctures. Notice that a collection of polygons in $\mathbb{C}$ — with sides identified by translation to form a genus g surface with n punctures — produces an element of $\Omega\mathcal{M}_{g,n}$ by defining a one-form that is equal to dz on each polygon, where z is the coordinate on $\mathbb{C}$. Such a surface is called a translation surface. Every element of $\Omega\mathcal{M}_{g,n}$ may be written as a translation surface. If X is a Riemann surface and ω a holomorphic one-form, then the flat metric on (X,ω) refers to the metric induced on X from its presentation as a translation surface.

As mentioned in the introduction, $\Omega\mathcal{M}_g$ admits a $\mathrm{GL}(2,\mathbb{R})$ action and a $\mathrm{GL}(2,\mathbb{R})$–invariant stratification. In terms of the polygonal presentation of translation surfaces, the $\mathrm{GL}(2,\mathbb{R})$ action may be described as the action induced by the $\mathrm{GL}(2,\mathbb{R})$ action on the polygons after making the identification of $\mathbb{C}$ with $\mathbb{R}^2$ via $x+iy \to (x,y)$.

Recall that the stratification of $\Omega\mathcal{M}_g$ was given by fixing the orders and multiplicities of the zeros of the one-forms. Given a one-form ω on a Riemann surface X we will let Σ denote its zero set. In the polygonal presentation, the points of Σ are the cone points of the flat metric (ie the points where the circumference of a circle of radius ϵ is not $2\pi\epsilon$ for sufficiently small ϵ). Local coordinates on strata are given by period coordinates, which is a coordinate system that records the periods of a basis of relative homology $H_1(X,\Sigma;\mathbb{Z})$. It follows that the tangent space at (X,ω) of a stratum of abelian differentials containing (X,ω) is $H^1(X,\Sigma;\mathbb{C})$.

If $|\Sigma| = 1$, then the change of coordinates between period coordinate charts is given by changing one homology basis with another and hence is an element of $\mathrm{Sp}(2g,\mathbb{Z})$.

More generally, when $|\Sigma|$ is arbitrary, the change of coordinates for period coordinates remains a constant volume-preserving linear function. This endows strata of $\Omega\mathcal{M}_g$ with a linear structure and a well-defined Lebesgue measure; see Zorich [31] for more details.

Lebesgue measure on strata of $\Omega\mathcal{M}_g$, also called Masur–Veech measure, induces a finite $\mathrm{SL}(2,\mathbb{R})$–invariant measure on $\mathcal{U}$, the locus of unit-area translation surfaces. This measure is ergodic with respect to the $\mathrm{SL}(2,\mathbb{R})$ action on $\mathcal{U}$ by Masur [18] and Veech [28]. Notice that it is necessary to restrict to $\mathcal{U}$ so that the total mass of the Masur–Veech measure is finite. We must also replace the $\mathrm{GL}(2,\mathbb{R})$ action by an $\mathrm{SL}(2,\mathbb{R})$ action in order to preserve the set $\mathcal{U}$. One might now inquire whether there is a classification of ergodic $\mathrm{SL}(2,\mathbb{R})$–invariant measures on $\mathcal{U}$.

2.1 Affine invariant submanifolds and equivariant maps between them

Definition 2.1 An affine invariant submanifold is a closed subset of a stratum of $\Omega\mathcal{M}_g$, which, in local period coordinates, can be expressed as the vanishing locus of a collection of real linear equations.

The arguments of Masur and Veech mentioned in the previous paragraph show that Lebesgue measure on affine invariant submanifolds restricted to $\mathcal{U}$ produces a finite ergodic $\mathrm{SL}(2,\mathbb{R})$–invariant measure. Eskin and Mirzakhani [8] showed that this is the only way to produce such measures on $\mathcal{U}$. Using this result, Eskin, Mirzakhani and Mohammadi [9] showed that the $\mathrm{GL}(2,\mathbb{R})$ orbit closure of any holomorphic one-form in a stratum is an affine invariant submanifold.

Definition 2.2 Given an affine invariant submanifold $\mathcal{M}$, let $\mu_{\mathcal{M}}$ be Lebesgue measure on $\mathcal{M}\cap\mathcal{U}$. Given a $\mathrm{GL}(2,\mathbb{R})$–equivariant measurable map $f\colon \mathcal{M}\to\mathcal{N}$ between two affine invariant submanifolds, define the pushforward of $\mathcal{M}$, denoted by $f_*\mathcal{M}$, to be the affine invariant submanifold represented by the ergodic $\mathrm{SL}(2,\mathbb{R})$–invariant measure given by $f_*\mu_{\mathcal{M}}$.

Notice that $f(\mathcal{M})$ and $f_*\mathcal{M}$ coincide up to sets of measure zero. Specifically, $f_*\mathcal{M}\setminus f(\mathcal{M})$ has $f_*\mu_{\mathcal{M}}$–measure zero since its preimage lies in the complement of $\mathrm{supp}(\mu_{\mathcal{M}})=\mathcal{M}\cap\mathcal{U}$. Conversely, $f(\mathcal{M})\setminus f_*\mathcal{M}$ lies outside the support of $f_*\mu_{\mathcal{M}}$ and hence has $f_*\mu_{\mathcal{M}}$–measure zero. Thus, its preimage on $\mathcal{M}$ has measure zero.

Example 2.3 Let $\mathcal{M}_{g,n}$ be the moduli space of genus g surfaces with n labeled marked points. Consider the affine invariant submanifold $\mathcal{M}$ in $\Omega\mathcal{M}_{2,2}$ consisting of holomorphic one-forms with a double zero and with two distinct marked points, marking the two punctures, that are exchanged by the hyperelliptic involution. Let $\mathcal{H}$ be the stratum of $\Omega\mathcal{M}_{2,1}$ consisting of abelian differentials with one double zero and one distinct marked point that records the location of the puncture. There is a forgetful map $f\colon \mathcal{M} \to \mathcal{H}$ that forgets the second marked point. Notice that $f_*\mathcal{M}$ is the collection of all holomorphic one-forms with a double zero and with a marked point that does not coincide with the double zero. However, $f(\mathcal{M})$ is the collection of all holomorphic one-forms with a double zero and with a marked point that does not coincide with any Weierstrass point. This illustrates how $f_*\mathcal{M}$ and $f(\mathcal{M})$ may differ up to sets of measure zero.

Finally we note the following.

Theorem 2.4 (Filip [11]) *Affine invariant submanifolds are algebraic subvarieties of $\Omega\mathcal{M}_g$.*

2.2 The tangent space of an affine invariant submanifold

As shown above, the tangent space at (X,ω) of a stratum of abelian differentials containing (X,ω) is $H^1(X,\Sigma;\mathbb{C})$. Therefore, the tangent space of affine invariant submanifolds containing a point (X,ω) is a subspace of $H^1(X,\Sigma;\mathbb{C})$. Since affine invariant submanifolds are locally linear, the tangent subspace is locally constant in period coordinates.

Let $p\colon H^1(X,\Sigma;\mathbb{C}) \to H^1(X;\mathbb{C})$ be the projection from relative to absolute cohomology. Notice that the kernel of the projection has size at most $|\Sigma|-1$.

Theorem 2.5 (Avila, Eskin and Möller [4]) *If $\mathcal{M}$ is an affine invariant submanifold containing a translation surface (X,ω), then $p(T_{(X,\omega)}\mathcal{M})$ is complex symplectic.*

2.3 Cylinder deformations

Definition 2.6 A cylinder c in a translation surface (X,ω) is a maximal isometrically embedded Euclidean cylinder that contains no points of Σ in its interior (the isometric embedding is with respect to the flat metric on (X,ω)). The cylinder is foliated by isotopic curves — whose homotopy class γ_c is called the core curve — that generate the fundamental group of the cylinder. The distance between the two boundaries of c is called the height, h_c, of the cylinder. The condition that a cylinder is maximal just means that h_c is as big as possible. The modulus of a cylinder is defined to be $h_c/\left|\int_{\gamma_c}\omega\right|$. Notice that there are two choices (each of which is an inverse of the other) for the core curve of c. Whenever a collection of parallel cylinders is specified we will tacitly assume that their core curves all have periods whose arguments, when taken as complex numbers, are identical.

Masur showed in [19] that every translation surface contains a cylinder. A fundamental tool in the sequel will be cylinder deformations.

Definition 2.7 Suppose that (X,ω) is a translation surface in an affine invariant submanifold $\mathcal{M}$. Two cylinders C_1 and C_2 on (X,ω) will be said to be $\mathcal{M}$–equivalent if C_1 and C_2 are parallel at all surfaces in an open neighborhood of (X,ω) in $\mathcal{M}$. A maximal collection of equivalent cylinders on (X,ω) will be called an $\mathcal{M}$–equivalence class. Given a collection $\mathcal{C}$ of cylinders whose periods have argument $e^{i\theta}$, the standard shear, $\sigma_{\mathcal{C}}$, and the standard dilation, $a_{\mathcal{C}}$, are the cohomology classes

$$\sigma_{\mathcal{C}} := e^{i\theta}\sum_{c\in\mathcal{C}} h_c\gamma_c^* \quad \text{and} \quad a_{\mathcal{C}} := i e^{i\theta}\sum_{c\in\mathcal{C}} h_c\gamma_c^*.$$

These cohomology classes belong to $H^1(X,\Sigma;\mathbb{C})$, where Σ is the zero set of ω. Recall that this cohomology group is the tangent space at (X,ω) of the stratum containing (X,ω). Since this stratum is linear we can travel along the line in the direction of the $\sigma_{\mathcal{C}}$ (resp. $a_{\mathcal{C}}$). In the case that the cylinders in $\mathcal{C}$ are horizontal, this line corresponds to the path given by applying the matrix $\begin{pmatrix}1&t\\0&1\end{pmatrix}$ $\left(\text{resp. } \begin{pmatrix}1&0\\0&1+t\end{pmatrix}\right)$, for $t\in\mathbb{R}$, to the cylinders in $\mathcal{C}$ while fixing the complement of $\mathcal{C}$. This justifies the terminology "standard shear" (resp. dilation). If $\mathcal{C}$ is a nonhorizontal collection of cylinders whose core curves are oriented so that their periods have argument $e^{i\theta}$, then flowing along $\sigma_{\mathcal{C}}$ (resp. $a_{\mathcal{C}}$) corresponds to shearing (resp. dilating the heights of) the cylinders in $\mathcal{C}$.

Example 2.8 Let $\mathcal{H}$ be a connected component of a stratum of abelian differentials. Let C_1 and C_2 be two parallel cylinders on a translation surface (X,ω) in $\mathcal{H}$ with core curves γ_1 and γ_2 respectively. Suppose without loss of generality that C_1 and C_2 are horizontal. If γ_1 and γ_2 are homologous, then $\int_{\gamma_1}\eta = \int_{\gamma_2}\eta$ for any nearby translation surface $(Y,\eta)\in\mathcal{H}$ and so the two cylinders are $\mathcal{H}$–parallel. Conversely, if γ_1 and γ_2 are not homologous then there is a simple closed curve γ whose oriented intersection number with γ_1 is 1 and whose oriented intersection number with γ_2 is zero. Since γ^* belongs to $H^1(X,\Sigma;\mathbb{C})$, which is the tangent space to $\mathcal{H}$, we may travel along a line in the $i\cdot\gamma^*$ direction. This keeps γ_2 horizontal while making γ_1 nonhorizontal. This example shows that two cylinders are $\mathcal{H}$–equivalent if and only if their core curves are homologous.

Theorem 2.9 (Wright [29, Corollary 3.4]) *Suppose that (X,ω) is a horizontally periodic translation surface whose $\mathrm{GL}(2,\mathbb{R})$ orbit closure is $\mathcal{M}$. Let $(C_1,\dots,C_n)$ be an enumeration of the horizontal cylinders and suppose that cylinder C_i has modulus m_i and core curve (oriented from left to right) γ_i of length c_i for $i=1,\dots,n$. Let $W\subseteq\mathbb{Q}^n$ be the subset of rational homogeneous linear relations that the moduli $(m_i)_{i=1}^n$ satisfy; ie $w\in W$ if and only if $w\cdot m=0$. If $(v_i)_{i=1}^n\in\mathbb{C}^n$ belongs to $W^\perp$, then*

$$\sum_{i=1}^n c_i v_i \gamma_i^* \in T_{(X,\omega)}\mathcal{M}.$$

Theorem 2.10 [29, Lemma 4.11] *If $\mathcal{C}$ is an $\mathcal{M}$–equivalence class of horizontal cylinders on a translation surface (X,ω) contained in an affine invariant submanifold $\mathcal{M}$ then the standard shear and the standard dilation are both contained in $T_{(X,\omega)}\mathcal{M}$.*

2.4 The boundary of an affine invariant submanifold

In Mirzakhani and Wright [24], a partial compactification of any affine invariant submanifold is defined as follows. A sequence (X_n,ω_n) of translation surfaces in a component of a stratum $\mathcal{H}$ converges to (Y,η) — where Y is a compact Riemann surface with finitely many connected components and η is a holomorphic one-form on each component together with a collection of marked points (which are considered part of the zero set Σ of η by thinking of the marked points as zeros of order zero) — if there is a nested collection of open neighborhoods $U_n\subseteq Y$ such that:

(1) $\bigcap_n U_n=\Sigma$.

(2) There are maps $g_n\colon Y\setminus U_n\to X_n$ that are diffeomorphisms onto their image and such that $g_n^*\omega_n$ converges to ω in the compact open topology.

(3) The injectivity radius of all points $p\in X_n$ — ie the supremum of the radius of balls (in the flat metric) centered at p that are embedded and do not contain points of Σ — not in the image of g_n goes to zero uniformly in n.

Given an affine invariant submanifold $\mathcal{M}$, let $\overline{\mathcal{M}}^{\mathrm{MW}}$ denote its partial compactification, which is given by $\mathcal{M}$ together with the union of all its limit points (as defined above). The $\mathrm{GL}(2,\mathbb{R})$ action on $\mathcal{M}$ extends continuously to $\overline{\mathcal{M}}^{\mathrm{MW}}$.

Theorem 2.11 (Mirzakhani and Wright [24, Corollary 1.2]) *Suppose that $\mathcal{M}$ is an affine invariant submanifold and (X_n,ω_n) is a sequence of points in $\mathcal{M}$ that converges to a (possibly disconnected) translation surface (Y,η) in the boundary of $\mathcal{M}$. The orbit closure of any component of (Y,η) has dimension smaller than $\dim\mathcal{M}$.*

2.5 Marked points on affine invariant submanifolds

Finally, the fundamental object of study in the sequel will be the following.

Definition 2.12 Let $\mathcal{M}$ be an affine invariant submanifold in a connected component of a stratum of abelian differentials $\mathcal{H}$. Fix a positive integer n. Let $\mathcal{H}(0^n)$ be the collection of translation surfaces in $\mathcal{H}$ together with n distinct marked points that do not coincide with zeros of the underlying abelian differential. The space $\mathcal{H}(0^n)$ is called a component of a marked stratum of abelian differentials. Let $\pi\colon \mathcal{H}(0^n) \to \mathcal{H}$ be the forgetful map. Define $\mathcal{M}(0^n)$ to be the preimage of $\mathcal{M}$ under π.

The forgetful map π extends to an equivariant map $\overline{\mathcal{M}(0^n)}^{\mathrm{MW}} \to \overline{\mathcal{M}}^{\mathrm{MW}}$ by forgetting the marked points. Let $\overline{\mathcal{M}(0^n)}$ be the preimage of $\mathcal{M}$ in $\overline{\mathcal{M}(0^n)}^{\mathrm{MW}}$ under the forgetful map. Notice that $\overline{\mathcal{M}(0^n)}$ is the partial compactification of $\mathcal{M}(0^n)$, where marked points are allowed to coincide with each other and with points of Σ.

Extending this definition to quadratic differentials introduces a minor complication because the Mirzakhani–Wright compactification is specifically defined for affine invariant submanifolds in strata of abelian differentials. Nevertheless, every quadratic differential has a canonical branched cover — called the holonomy double cover — so that the pullback of the quadratic differential under this cover is the square of an abelian differential.

Definition 2.13 Let $\mathcal{M}$ be an affine invariant submanifold in a connected component of a stratum of quadratic differentials $\mathcal{Q}$. Let $\widetilde{\mathcal{Q}}$ denote the collection of holonomy double covers of quadratic differentials in $\mathcal{Q}$, which is itself an affine invariant submanifold. On each surface in $\widetilde{\mathcal{Q}}$ we use the term "holonomy involution" to refer to the involution whose quotient gives the holonomy double cover.

As before, fix a positive integer n and let $\mathcal{Q}(0^n)$ denote the collection of n distinct marked points that do not coincide with zeros or poles of the quadratic differential. The locus of holonomy double covers $\widetilde{\mathcal{Q}(0^n)}$ is the $\mathrm{GL}(2,\mathbb{R})$–invariant subset of $\widetilde{\mathcal{Q}}(0^{2n})$ consisting of surfaces in $\widetilde{\mathcal{Q}}$ together with $2n$ distinct points that do not coincide with zeros and which project to exactly n points under the holonomy involution. Let $\mathcal{C}$ be the closure of $\widetilde{\mathcal{Q}(0^n)}$ in $\overline{\widetilde{\mathcal{Q}}(0^{2n})}$ (this uses Definition 2.12 with $\widetilde{\mathcal{Q}}$ as the affine invariant submanifold). Define the points in $\overline{\mathcal{Q}(0^n)}$ to be the set of marked surfaces in $\mathcal{C}$ after quotienting by the holonomy involution. This creates a bijection between $\mathcal{C}$ and $\overline{\mathcal{Q}(0^n)}$. We give $\overline{\mathcal{Q}(0^n)}$ the topology that makes this bijection a homeomorphism.

The forgetful map from $\mathcal{Q}(0^n)$ to $\mathcal{Q}$ extends to one between $\overline{\mathcal{Q}(0^n)}$ and $\mathcal{Q}$. Let $\overline{\mathcal{M}(0^n)}$ (resp. $\mathcal{M}(0^n)$) be the preimage of $\mathcal{M}$ under the forgetful map whose domain is $\overline{\mathcal{Q}(0^n)}$ (resp. $\mathcal{Q}(0^n)$). The fibers of the forgetful map with domain $\overline{\mathcal{M}(0^n)}$ are all collections of n marked points, where the points are now allowed to coincide with each other, zeros and poles.

Lemma 2.14 *Let $\mathcal{M}$ be an affine invariant submanifold and let $\mathcal{N}$ be an affine invariant submanifold in $\mathcal{M}(0)$. Let $\pi\colon \mathcal{M}(0) \to \mathcal{M}$ be the map that forgets the marked points. If $\pi_*\mathcal{N} = \mathcal{M}$ then $\pi(\mathcal{N})$ is an open dense subset of $\mathcal{M}$ and the fiber over a generic point in $\mathcal{M}$ is nonempty. Moreover, π is open.*

Proof Since all fibers of the forgetful map $\pi\colon \overline{\mathcal{M}(0)} \to \mathcal{M}$ are compact, π is proper and its image is closed. The image of π has full measure in $\pi_*\mathcal{N}$, which is $\mathcal{M}$ by assumption. If $\mathcal{N} = \mathcal{M}(0)$ then the result is immediate, so suppose instead that $\mathcal{N}$ and $\mathcal{M}$ have the same dimension. Notice that all surfaces in $\overline{\mathcal{M}(0)}$ are connected by construction. As a reminder, $\overline{\mathcal{M}(0)}$ is the partial compactification defined in Definition 2.12, which is a subset of the Mirzakhani–Wright partial compactification. Therefore, by Mirzakhani and Wright [24, Corollary 1.2] — see our Theorem 2.11 — the components of the boundary of $\overline{\mathcal{N}}^{\mathrm{MW}} \cap \overline{\mathcal{M}(0)}$ have strictly smaller dimension than $\mathcal{M}$ and hence their pushfoward $\mathcal{C}$ cannot be $\mathcal{M}$ by Sard's theorem. Therefore, $\pi(\mathcal{N})$ contains the complement of $\mathcal{C}$, which is an open dense set in $\mathcal{M}$. Since $\pi\colon \mathcal{N} \to \mathcal{M}$ is a finite holomorphic map between equidimensional varieties (by Filip [11]; see our Theorem 2.4), it is open. □

Definition 2.15 If $\mathcal{M}$ is an affine invariant submanifold then let $\mathcal{M}^{\mathrm{ord}}(0^n)$ be the finite cover of $\mathcal{M}(0^n)$ where the marked points are labeled. Let $\pi_k\colon \mathcal{M}(0^n) \to \mathcal{M}(0^{n-1})$ be the map that forgets the k^{th} marked point for $k \in \{1, \dots, n\}$.

Lemma 2.16 *Let $\mathcal{N}$ be an affine invariant submanifold in $\mathcal{M}(0^n)$ that pushes forward to $\mathcal{M}$ under the map π that forgets all marked points. If $(X, \omega; P)$ is generic in $\mathcal{M}(0^n)$, then there is an open neighborhood U of $(X, \omega; P)$ in $\mathcal{N}$ such that $\pi(U)$ is an open set around (X, ω) in $\mathcal{M}$.*

Proof Without loss of generality, we will work on $\mathcal{M}^{\mathrm{ord}}(0^n)$. Since each π_k restricted to $\mathcal{N}$ is open by Lemma 2.14 and since π is equal to $\pi_1 \circ \cdots \circ \pi_n$, its restriction to $\mathcal{N}$ is open as well. □

3 Holomorphic sections over varieties containing a Teichmüller disk and proof of Theorem 1.3

Throughout this section we will make the following assumption.

Assumption 3.1 Let Γ be a torsionfree finite-index subgroup of the mapping class group. Let C be a complex analytic subvariety of $\mathrm{Teich}_{g,n}/\Gamma$, where $\mathrm{Teich}_{g,n}$ is the Teichmüller space of a closed genus g surface with n punctures, with $3g-3+n>0$. Let $\pi\colon \mathcal{C}_{g,n} \to \mathrm{Teich}_{g,n}/\Gamma$ be the universal curve. Suppose additionally that C contains a Teichmüller disk generated by the quadratic differential (X,q), where X is a Riemann surface and q a quadratic differential on X. Suppose that $\mathcal{Q}$ is the stratum of quadratic differentials to which (X,q) belongs and let $\mathcal{M}$ be the orbit closure of (X,q) in $\mathcal{Q}$.

Lemma 3.2 *Every holomorphic section of the universal curve over C induces a $\mathrm{GL}(2,\mathbb{R})$–equivariant section of the forgetful map from $\overline{\mathcal{M}(0)}$ to $\mathcal{M}$.*

Proof Let $s\colon C \to \mathcal{C}_{g,n}$ be a holomorphic section of the restriction of π to $\pi^{-1}(C)$. Let $\iota\colon \mathbb{D} \to C$ be the inclusion of the Teichmüller disk into C. The inclusion is an isometry with respect to the underlying Kobayashi hyperbolic metrics. Since $\iota = \pi \circ s$ and π and s are contractions in the Kobayashi hyperbolic metrics, it follows that s restricted to the embedded Teichmüller disk in C is an isometry in the Kobayashi metrics. Therefore, $s(\iota(\mathbb{D}))$ is a Teichmüller disk in $\mathcal{C}_{g,n}$.

Sufficiently close Riemann surfaces X_1 and X_2 contained in $\iota(\mathbb{D})$ are joined by a dilatation-minimizing homeomorphism given by a geodesic in $\iota(\mathbb{D})$. By Teichmüller's theorem, this homeomorphism is unique up to pre- and post-composition with a conformal automorphism. However, since Γ is torsionfree there are no such automorphisms in Γ that fix X_1 or X_2. If γ is the geodesic from X_1 to X_2 in $\iota(\mathbb{D})$, then $s(\gamma)$ is a geodesic of the same length in $\mathcal{C}_{g,n}$ and hence corresponds to a homeomorphism with the same dilatation. By uniqueness this path must correspond to the same homeomorphism and so if Teichmüller geodesic flow along (X_1,q_1) produces the geodesic γ, Teichmüller geodesic flow along $(s(X),q)$ produces $s(\gamma)$.

Let $\widetilde{s}\colon \mathcal{M} \to \overline{\mathcal{M}(0)}$ be the section of the forgetful map from $\overline{\mathcal{M}(0)}$ to $\mathcal{M}$ given by sending a quadratic differential (X,q) to $(s(X),q)$. The argument above shows that this map is $\mathrm{GL}(2,\mathbb{R})$–equivariant on Teichmüller disks since it is equivariant under complex scalar multiplication and Teichmüller geodesic flow. Since $\mathcal{M}$ is foliated by $\mathrm{GL}(2,\mathbb{R})$–invariant Teichmüller disks the claim follows. □

Remark 3.3 The proof of Lemma 3.2 only uses the hypothesis that Γ is torsionfree, not that it is finite-index.

Proof of Theorem 1.3 Let s be a holomorphic section of the restriction of π to $\pi^{-1}(C)$. Since it is a continuous section, it follows that $s(C)$, and hence $\widetilde{s}(\mathcal{M})$, is closed. By Lemma 3.2, $\widetilde{s}(\mathcal{M})$ is closed and $\mathrm{GL}(2,\mathbb{R})$–invariant and therefore it is an affine invariant submanifold by Eskin, Mirzakhani and Mohammadi [9]. Notice that this application of Eskin, Mirzakhani and Mohammadi uses the fact that Γ is finite-index since this implies that the Lebesgue measure of the collection of unit-area half-translation surfaces in $\mathcal{M}$ is finite. Since $\widetilde{s}(\mathcal{M})$ is an affine invariant submanifold that does not coincide with $\mathcal{M}(0)$, it follows that the point that s marks above X is a periodic point, zero or pole of (X,q). □

Remark 3.4 The same proof shows that measurable equivariant sections of the forgetful map from $\overline{\mathcal{M}}(0)$ to $\mathcal{M}$ only mark periodic points, zeros or poles. In the measurable setting, the section is used to push forward Lebesgue measure to a measure on $\overline{\mathcal{M}(0)}$, which must be Lebesgue measure on an affine invariant submanifold by Eskin and Mirzakhani [8]. The details are omitted.

4 Explicit translation surfaces in every component of every stratum

We will construct explicit generic translation surfaces in each connected component of every stratum of abelian differentials. The connected components were classified in [15].

Theorem 4.1 (Kontsevich and Zorich [15]) *All strata are connected except:*

- *For $g > 3$, $\mathcal{H}(2g-2)$ has three connected components characterized by odd spin, even spin and hyperellipticity.*
- *For odd $g > 3$, $\mathcal{H}(g-1,g-1)$ has three connected components characterized by odd spin, even spin and hyperellipticity.*
- *For even $g > 3$, $\mathcal{H}(g-1,g-1)$ has two connected components characterized by hyperellipticity and nonhyperellipticity.*
- *For $g > 3$, $\mathcal{H}(2k_1,\dots,2k_n)$ has two connected components characterized by odd and even spin (excluding the case $\mathcal{H}(g-1,g-1)$ for odd $g > 3$, which, as mentioned above, has three components).*
- *$\mathcal{H}(4)$ and $\mathcal{H}(2,2)$ each have two connected components, a hyperelliptic and an odd one.*

To distinguish which connected component of a stratum a specific translation surface belongs to, we will use the following criterion.

Theorem 4.2 [15, Corollary 2] *Let $\mathcal{H}$ be a stratum of abelian differentials. For each connected component C of the minimal stratum there is a unique component of $\mathcal{H}$ that contains C in its closure.*

We are now in a position to create horizontally and vertically periodic translation surfaces in each component of each stratum. First, we establish a convention:

Convention for figures We will often use polygons, all of whose edges will be vertical or horizontal, to represent translation surfaces using the following two conventions. The edge of a polygon will mean a line segment in the boundary of the polygon that connects two vertices and has no vertex in its interior.

(1) The intersection of a dotted line and an edge is a vertex of the polygon.

(2) If a pair of unmarked vertical (resp. horizontal) edges contain interior points that can be connected by a horizontal (resp. vertical) line that lies in the interior of the polygon then they are identified.

Under this convention the two translation surfaces in Figure 1 are identical.

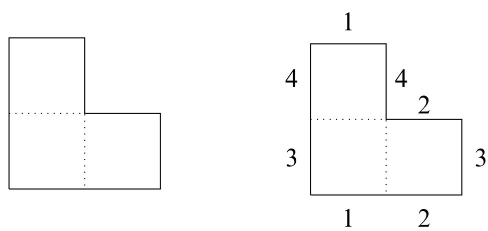

Figure 1: Equivalent representations of the same translation surface.

Using separatrix diagrams, Kontsevich and Zorich produce surfaces that belong to each component of the minimal stratum; see [15, Figure 4]. The surfaces are represented as translation surfaces in Figure 2.

Proposition 4.3 (genericity criterion) *Suppose that (X,ω) is a translation surface in a component $\mathcal{H}$ of a stratum of abelian differentials of genus g surfaces with n zeros and no marked points. If (X,ω) has $g+n-1$ horizontal cylinders whose moduli satisfy no rational linear relation, then (X,ω) is generic.*

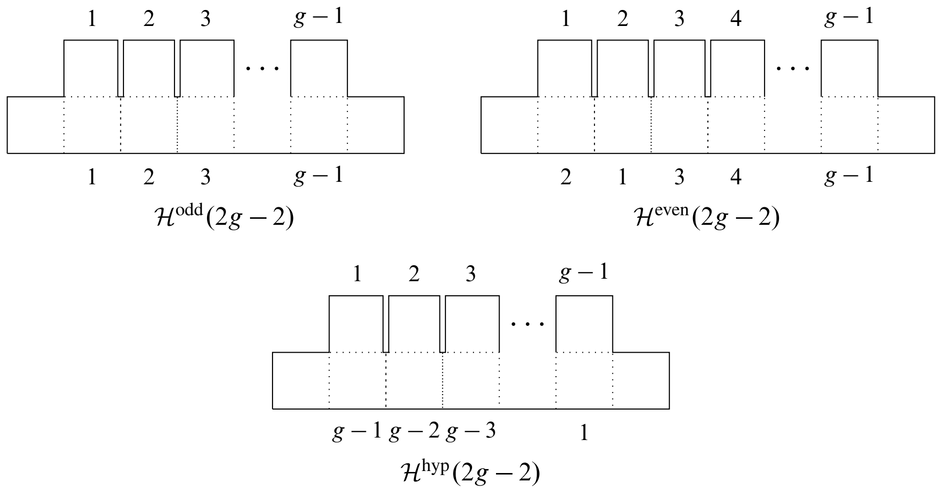

Figure 2: Surfaces in each component of the minimal stratum.

Proof Let $\mathcal{M}$ be the orbit closure of (X,ω). It suffices to show that $\mathcal{M}$ coincides with a component of $\mathcal{H}$. By Wright [29, Corollary 3.4] — see our Theorem 2.9 — since the moduli satisfy no rational linear relation the tangent space of $\mathcal{M}$ at (X,ω) includes $\{\gamma_1^*,\dots,\gamma_{g+n-1}^*\}$, where the γ_i are core curves oriented from left to right of the horizontal cylinders and γ_i^* denotes the dual cohomology class under the intersection pairing. Recall that the identification of $T_{(X,\omega)}\mathcal{M}$ with a subspace of $H^1(X,\Sigma;\mathbb{C})$ for Σ the zero set of ω was described in Section 2.2. The dual cohomology classes span a complex vector space of dimension $g+n-1$.

Let $p\colon T_{(X,\omega)}\mathcal{M}\to H^1(X,\mathbb{C})$ be the projection from relative to absolute cohomology. By Avila, Eskin and Möller [4] — see our Theorem 2.5 — the image of the projection is a complex symplectic vector space. The kernel of the projection has (complex) dimension at most $n-1$. Since the projection of $\{\gamma_1^*,\dots,\gamma_{g+n-1}^*\}$ spans an isotropic subspace, which has dimension at most g, it follows that the kernel of p has dimension exactly $n-1$ and that the projection of $\{\gamma_1^*,\dots,\gamma_{g+n-1}^*\}$ spans a Lagrangian subspace. Since $p(T_{(X,\omega)}\mathcal{M})$ is complex symplectic it follows that p is a surjection with maximal-dimensional kernel. It follows that $T_{(X,\omega)}\mathcal{M}$ is isomorphic to $H^1(X,\Sigma;\mathbb{C})$ and hence that $\mathcal{M}$ has full dimension. Since $\mathcal{M}$ is open and closed it must coincide with a component of $\mathcal{H}$. □

Generic surfaces in $\mathcal{H}^{\mathrm{hyp}}(2g-2)$ and $\mathcal{H}^{\mathrm{hyp}}(g-1,g-1)$ It is straightforward to verify that the translation surfaces in Figure 3 are in the indicated components; see for example [2, Section 2]. The genericity criterion (Proposition 4.3) implies that the translation surfaces are generic provided that all moduli of horizontal cylinders satisfy no rational linear relation.

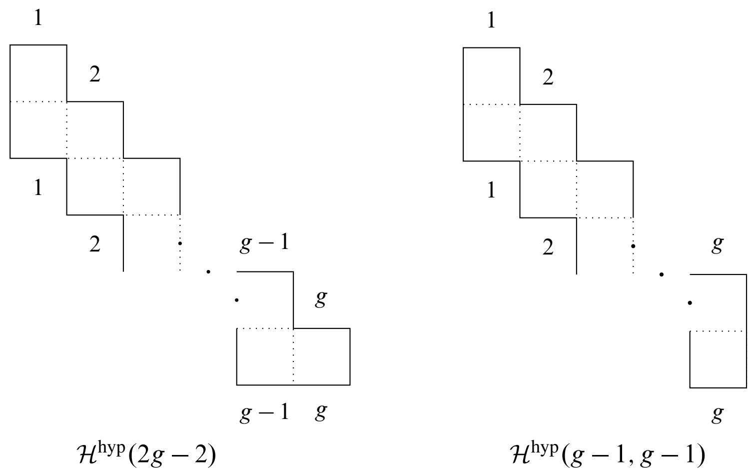

Figure 3: Hyperelliptic translation surfaces.

Generic surfaces in even components of strata Proceed as follows to find a surface in $\mathcal{H} := \mathcal{H}^{\mathrm{even}}(2k_1, \dots, 2k_n)$. Start with the surface in Figure 4, set $g = 1 + \sum_i k_i$ and $S := \{a_0, a_{k_1}, a_{k_1+k_2}, \dots, a_{k_1+\dots+k_{n-1}}\}$. Then collapse every saddle connection in $\{a_1, \dots, a_{g-2}\} \setminus S$. The resulting surface belongs to the even component since further collapsing the saddle connections in $S \setminus \{a_0\}$ gives a path in the stratum whose endpoint is a surface in the even minimal component. (Specifically it will be the surface in the top right of Figure 2 after we perform a half Dehn-twist in horizontal cylinders C_1 and C_2, which has the effect of changing the labels from the order $(1,2)$ (resp. $(3,4)$) on the top boundary of C_1 (resp. C_2) to $(2,1)$ (resp. $(4,3)$).)

We want to specify a translation surface (X,ω) in $\mathcal{H}$ constructed by the process described in the preceding paragraph. To do so, we must first specify the lengths of the saddle connections in Figure 4. Let $(\alpha_i)_{i=0}^{g-2} \cup (x_i)_{i=1}^{g+1} \cup (y_i)_{i=0}^{g-1}$ be a collection of positive real numbers with the property that all rational linear relations satisfied by these numbers are generated by the relations $x_1 = x_4$ and $x_2 = x_3$. Now we will set the horizontal saddle connections labeled by an integer $i \in \{1, \dots, g+1\}$ to have length x_i, the heights of the cylinders C_i (resp. C) to have length y_i for $i > 0$ (resp. length y_0) and the saddle connections labeled a_i to have length 0 if they were collapsed by the procedure described in the previous paragraph and α_i otherwise.

By construction the moduli of the vertical (resp. horizontal) cylinders of (X,ω) satisfy no rational linear relation and so by the genericity criterion (Proposition 4.3), (X,ω) is generic.

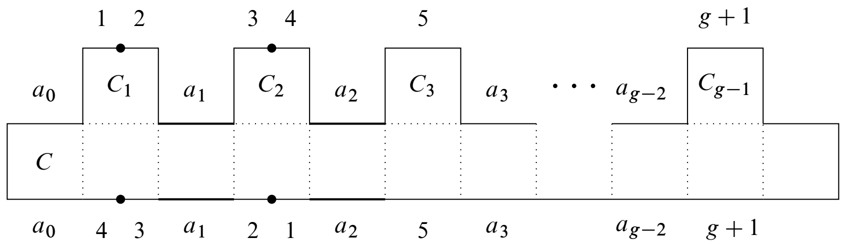

Figure 4: $\mathcal{H}^{\text{even}}(2,\dots,2)$ with labeled horizontal cylinders.

The following lemma will be needed in the final two sections in order to apply the main technical result of the paper, Lemma 5.5. A reader who seeks to understand the motivation for proving the following result is encouraged to read the statement of Lemma 5.5 first.

Lemma 4.4 *If $g>3$ then on (X,ω) there are at least two $\mathcal{H}$–equivalence classes of vertical cylinders $\mathcal{D}_1$ and $\mathcal{D}_2$ such that $C\cap\mathcal{D}_i$ is connected for $i=1,2$ and such that the complement of $\mathcal{D}_1\cup\mathcal{D}_2$ in C consists of two disjoint rectangles of different horizontal lengths.*

Remark 4.5 In the special case that the complement of $\mathcal{D}_1\cup\mathcal{D}_2$ in C is connected, we will take one rectangle to be an "empty rectangle" whose horizontal length is zero.

Remark 4.6 The restriction that $g>3$ is natural since when $g=3$ the even components of strata coincide with hyperelliptic components.

Proof Let D_i be the unique vertical cylinder passing through C_i for $i\in\{2,\dots,g+1\}$. Let D_1 (resp. D_2) be the vertical cylinder that passes through saddle connection 1 (resp. 2). Finally, let A_i be the vertical cylinder passing through the saddle connection labeled a_i (provided that this saddle connection has not been collapsed in the construction of (X,ω)).

By Example 2.8, two cylinders are $\mathcal{H}$–equivalent if and only if their core curves are homologous. This shows that the $\mathcal{H}$–equivalence classes of vertical cylinders are

$$\{D_1,D_2\},\quad \{D_3\},\quad \dots,\quad \{D_{g-1}\},\quad \{A_1\},\quad \{A_0,A_2,\dots,A_{g-2}\}.$$

If $g > 4$, then let $\mathcal{D}_1 := \{D_3\}$ and $\mathcal{D}_2 = \{D_4\}$. These are two $\mathcal{H}$–equivalence classes that have connected intersection with C. The complement of $\mathcal{D}_1 \cup \mathcal{D}_2$ in C consists of two rectangles, one of which that has length α_3 or zero (depending on whether or not saddle connection a_3 was collapsed) and the other of which is a positive real number that is not rationally linearly related to α_3 by construction. Suppose now that $g = 4$.

Suppose first that the horizontal saddle connection labeled a_1 is uncollapsed. Set $\mathcal{D}_1 = \{A_1\}$ and $\mathcal{D}_2 = \{D_3\}$. The complement of $\mathcal{D}_1 \cup \mathcal{D}_2$ in C is a disjoint union of two rectangles. These rectangles have horizontal lengths $|a_2| + x_1 + x_2$ and $|a_0| + x_1 + x_2$, where $|\cdot|$ denotes flat length. These two lengths are distinct by construction.

Suppose now that a_1 is collapsed. Set $\mathcal{D}_1 = \{D_1, D_2\}$ and $\mathcal{D}_2 = \{D_3\}$. Again the complement of $\mathcal{D}_1 \cup \mathcal{D}_2$ in C is a disjoint union of two rectangles whose lengths are $|a_0|$ and $|a_2|$, which again are distinct by construction. □

Generic surfaces in remaining components To find generic surfaces in all other connected components of the remaining strata we glue together copies of the surfaces in Figure 5 along the horizontal cylinders that intersect all vertical cylinders. By the genericity criterion (Proposition 4.3) whenever the vertical cylinders have rationally unrelated moduli this surface is generic.

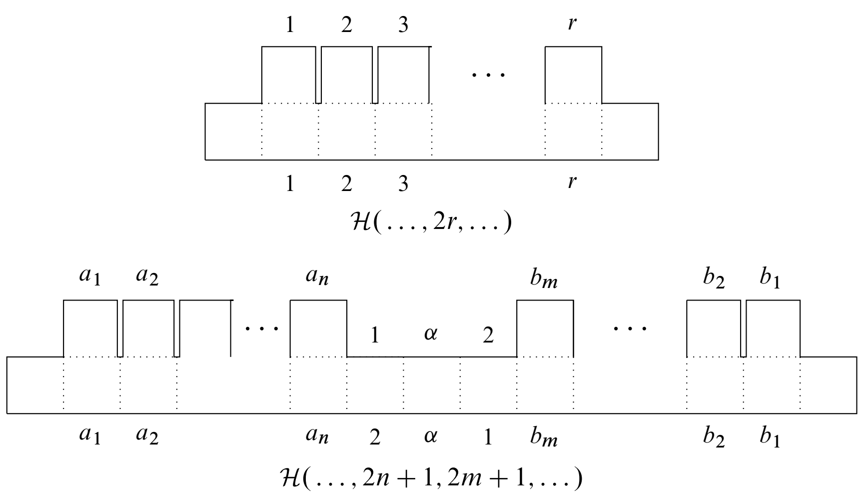

Figure 5: Surfaces in $\mathcal{H}^{\text{odd}}$, $\mathcal{H}^{\text{nonhyp}}$ and connected strata.

Definition 4.7 Each of the nonhyperelliptic surfaces constructed in this subsection contain a horizontal cylinder that intersects every vertical cylinder. This cylinder will be called the central horizontal cylinder.

Lemma 4.8 *Let (X, ω) be a surface constructed by the above process. Let C be the central horizontal cylinder. There are $\mathcal{H}$–equivalence classes of cylinders $\mathcal{D}_1$ and $\mathcal{D}_2$ that contain one cylinder intersecting C exactly once. The complement of $\mathcal{D}_1$ and $\mathcal{D}_2$ in C consists of two disjoint rectangles whose horizontal lengths are distinct for generic choices of lengths of the horizontal saddle connections.*

Proof Suppose first that the (X, ω) has two zeros of even order. Then the surface contains two copies of the surface in Figure 5, top. It is sufficient to take the two vertical cylinders that pass through the horizontal saddle connection labeled 1. Similarly, if there are four zeros of odd order we have two copies of the surface in Figure 5, bottom, and we take the two vertical cylinders that pass through the horizontal saddle connection α.

If there is one zero of even order and two zeros of odd order, then we have a surface like the one in Figure 5, top, and one like the one in Figure 5, bottom, and we take the vertical cylinder passing through the saddle connection labeled 1 on the first and the one passing through the horizontal saddle connection labeled α on the second.

If there is only one zero, then it is of even order and we take the two vertical cylinders that pass through 1 and 2. If there are only two zeros and both are of odd order then we take the vertical cylinder passing through α and the vertical cylinder passing through a_1 or b_1 (whichever exists).

As in Lemma 4.4, two cylinders are $\mathcal{H}$–equivalent if and only if their core curves are homologous. All cylinders selected in this proof have no distinct $\mathcal{H}$–equivalent cylinders. (To see this, note that the constructed surface is made up of vertically periodic subsurfaces as in Figure 5. On each subsurface no two vertical cylinders have homologous core curves. Moreover, if two vertical cylinders have homologous core curves then they must belong to the same subsurface.)

All cylinders selected in this proof intersect C exactly once, and therefore it is immediate that $C \setminus (\mathcal{D}_1 \cup \mathcal{D}_2)$ is a union of two disjoint rectangles. By construction, one of these rectangles completely contains a vertical cylinder. By altering the height (ie horizontal length) of this cylinder we see that the two rectangles do not have generically identical horizontal lengths. □

5 Marked points in cylinders

In this section we will prove results about marked points and cylinders that form the technical core of the paper. We will make the following standing assumption:

Assumption 5.1 $\mathcal{M}$ is an affine invariant submanifold and (X,ω) is a translation surface whose $\mathrm{GL}(2,\mathbb{R})$ orbit closure is $\mathcal{M}$.

Lemma 5.2 *Let P be a collection of distinct points on (X,ω) and suppose that $\mathcal{M}'$ is the orbit closure of $(X,\omega;P)$. If $\mathcal{C}$ is an $\mathcal{M}$–equivalence class of cylinders on (X,ω), then $\mathcal{C}'$ is an $\mathcal{M}'$–equivalence class, where $\mathcal{C}'$ contains the cylinders in $\mathcal{C}$ divided into subcylinders by the points in P.*

Proof First we will show that any two cylinders in $\mathcal{C}'$ are $\mathcal{M}'$ equivalent. Let C_i be two cylinders in $\mathcal{C}'$ and let γ_i be their core curves for $i=1,2$. By assumption, there is a neighborhood U of (X,ω) in $\mathcal{M}$ on which γ_1 and γ_2 are collinear. Let U' be a preimage of the U in $\mathcal{M}'$ on which C_1 and C_2 persist as cylinders. In this neighborhood, γ_1 and γ_2 must remain collinear and hence $\mathcal{M}'$–equivalent.

It remains to show that if C_1 and C_2 are $\mathcal{M}'$–equivalent cylinders with core curves γ_1 and γ_2, then γ_1 and γ_2 must be collinear on a neighborhood of (X,ω) in $\mathcal{M}$. By Lemma 2.16 there is a neighborhood of $(X,\omega;P)$ in $\mathcal{M}'$ that projects to an open neighborhood of (X,ω) in $\mathcal{M}$. Let U' be such a neighborhood of $(X,\omega;P)$ in $\mathcal{M}'$ where γ_1 and γ_2 are collinear, and let U be its image under the forgetful map. Since γ_1 and γ_2 are collinear on U' they are collinear on U and hence are homotopic to core curves of $\mathcal{M}$–equivalent cylinders. □

Definition 5.3 Given a cylinder C in a translation surface we say that the height of the cylinder is the distance between the two boundaries in the flat metric. Suppose that p is a marked point contained in a cylinder C. Let h_C be the height of the cylinder, and let h_p be the distance from the point to one of the two boundary curves of C. We say that p lies at rational height in C if the ratio h_p/h_C is rational.

Lemma 5.4 (rational height lemma) *Let $\mathcal{C}$ be an $\mathcal{M}$–equivalence class of cylinders. Suppose that any two cylinders in $\mathcal{C}$ have a rational ratio of moduli. If a periodic point belongs to the interior of a cylinder in $\mathcal{C}$ then it lies at rational height.*

Proof Let p be a periodic point contained in the interior of a cylinder in $\mathcal{C}$. Let $\mathcal{M}'$ be the orbit closure of $(X,\omega;p)$. Let $\mathcal{C}'$ be the collection of subcylinders on $(X,\omega;p)$ into which $\mathcal{C}$ is divided. By Lemma 5.2, $\mathcal{C}'$ is an $\mathcal{M}'$ equivalence class. Let $\sigma_{\mathcal{C}'}$ be the standard shear on $\mathcal{C}'$. Since the cylinders in $\mathcal{C}$ have a rational ratio of moduli, the flow along $\sigma_{\mathcal{C}}$ is periodic. Suppose, for a contradiction, that p does not have rational height. In this case, the flow along $\sigma_{\mathcal{C}'}$ is not periodic and so the orbit closure of $(X,\omega;p)$ contains $(X,\omega;q)$, where q is any point in C along the core curve of C that intersects p.

Let γ_1 and γ_2 be the two core curves of the cylinders into which p divides C. Since p may be moved along the core curve of C while fixing all cylinders in (X,ω), it follows that the tangent space of $\mathcal{M}'$ at $(X,\omega;p)$ contains the deformation $\gamma_1^* \smallsetminus \gamma_2^*$. Let U' be a neighborhood as in Lemma 2.16 of $(X,\omega;p)$ in $\mathcal{M}'$ on which the cylinders in $\mathcal{C}$ persist. This neighborhood projects to a neighborhood U of (X,ω) in $\mathcal{M}$. Since the tangent space contains the deformation $\gamma_1^* \smallsetminus \gamma_2^*$ the fiber of the projection from $\mathcal{M}'$ to $\mathcal{M}$ that forgets marked points has real dimension at least one. Therefore, the dimension of $\mathcal{M}'$ is strictly larger than the dimension of $\mathcal{M}$, which contradicts the assumption that p is a periodic point. □

Lemma 5.5 *Let (X,ω) be a generic translation surface in an affine invariant submanifold $\mathcal{M}$. Let C be a horizontal cylinder, and let $\mathcal{D}_1$ and $\mathcal{D}_2$ be two vertical distinct $\mathcal{M}$–equivalence classes of cylinders such that:*

1. *The intersection of $\mathcal{D}_i$ with the interior of C is connected and nonempty for $i=1,2$.*
2. *Any cylinder in the $\mathcal{M}$–equivalence class $\mathcal{C}$ of C has a modulus that is an integer multiple of the modulus of C.*

If p is an $\mathcal{M}$–periodic point in the interior of C, then up to relabeling $\mathcal{D}_1$ and $\mathcal{D}_2$, the point p is at the center of the rectangle given by the intersection of $\mathcal{D}_1$ and C. Furthermore, removing $\mathcal{D}_1$ and $\mathcal{D}_2$ divides C into two rectangles of equal size.

Proof Let $\mathcal{M}'$ be the orbit closure of $(X,\omega;p)$. Let $\mathcal{C}$ be the collection of horizontal cylinders equivalent to C on (X,ω).

For simplicity we begin by applying a element of $\mathrm{GL}(2,\mathbb{R})$ so that C has unit height. By the rational height lemma (Lemma 5.4) the marked point p lies at rational height h in C. Since we have normalized the height of C to be one this means that h is rational; in particular suppose that $h=n/m$, where n and m are coprime positive integers.

Part 1 We may assume that p belongs to $\mathcal{D}_1$ Suppose not. By Lemma 5.2, since p is not contained in $\mathcal{D}_1$, the standard shear $a_{\mathcal{D}_1}$ is tangent to $\mathcal{M}'$. Traveling in the $a_{\mathcal{D}_1}$ direction in $\mathcal{M}'$ from $(X,\omega;p)$ widens $\mathcal{D}_1$ while fixing the part of the translation surface (and marked point) in the complement $\mathcal{D}_1$. Travel in this direction until the intersection of $\mathcal{D}_1$ and C accounts for at least $(m-1)/m$ proportion of the area of C. Let (Y,η) be the new translation surface.

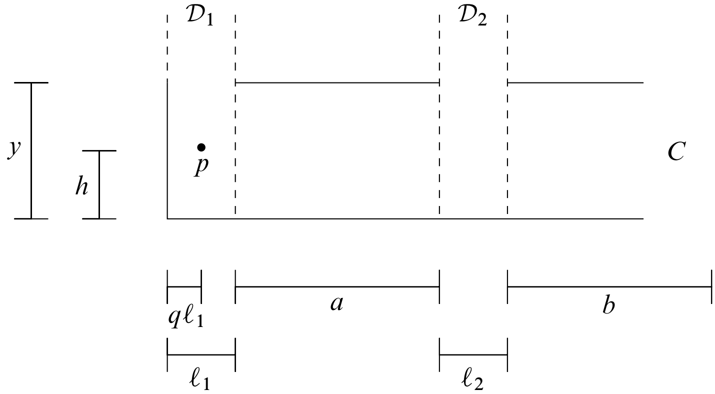

Figure 6: Lemma 5.5 shows that $a = b$ and that, after scaling so C has unit height, $q = h = \frac{1}{2}$.

Part 1a (Y, η) may be taken to be generic in $\mathcal{M}$ We formed (Y, η) by traveling in the $a_{\mathcal{D}_1}$ direction from (X, ω) in $\mathcal{M}$. Let ℓ be the segment joining (X, ω) to (Y, η) in $\mathcal{M}$. Each proper affine invariant submanifold (of which there are only countably many by Eskin, Mirzakhani and Mohammadi [9]) contained in $\mathcal{M}$ intersects ℓ in a closed set. If a neighborhood U of (Y, η) had the property that every point in $U \cap \ell$ was contained in some proper affine invariant submanifold, then by the Baire category theorem there would be a proper affine invariant submanifold $\mathcal{N}$ such that $\mathcal{N} \cap \ell$ had interior in $U \cap \ell$. Since affine invariant submanifolds are linear this would imply that all of ℓ was contained in $\mathcal{N}$, which contradicts the fact that ℓ contains a translation surface (X, ω) that is generic in $\mathcal{M}$. Therefore, we may assume that the point (Y, η) was chosen to be generic in $\mathcal{M}$.

Part 1b The hypotheses of the lemma continue to hold on (Y, η) and the fiber of $\mathcal{M}'$ over (Y, η) contains $(Y, \eta; p)$ where p belongs to $\mathcal{D}_1$ Traveling in the $a_{\mathcal{D}_1}$ direction keeps p fixed in the complement of $\mathcal{D}_1 \cup \mathcal{D}_2$ and keeps the heights of the cylinders in $\mathcal{C}$ constant. By Wright [29, Theorem 1.9], the ratio of lengths of core curves of cylinders in $\mathcal{C}$ are constant, and so the condition on moduli of cylinders in $\mathcal{C}$ continues to hold.

Letting $\mathcal{C}'$ be the collection of cylinders on $(Y, \eta; p)$ that project to cylinders in $\mathcal{C}$ we have that $\mathcal{C}'$ is an $\mathcal{M}'$–equivalence class by Lemma 5.2. Travel from $(Y, \eta; p)$ in the $\sigma_{\mathcal{C}'}$ direction until one complete Dehn twist has been performed in C. The resulting unmarked translation surface is (Y, η) since all cylinders in $\mathcal{C}$ have moduli that are integer multiples of the modulus of C. Since $\mathcal{D}_1 \cap C$ is connected and accounts for at least $(m-1)/m$ of the area of C, it follows that the marked point p now belongs to $\mathcal{D}_1$ and we have shown that $(Y,\eta;p)$ belongs to $\mathcal{M}'$.

Part 1c If the conclusion of the lemma holds on (Y,η), it does so on (X,ω) also To pass from (X,ω) to (Y,η), we first traveled along $a_{\mathcal{D}_1}$ and then traveled along $\sigma_{\mathcal{C}'}$ to perform one Dehn twist in C. If the conclusion of the lemma holds on (Y,η), then traveling in the opposite direction along $\sigma_{\mathcal{C}'}$ to perform the opposite Dehn twist in C moves p to the midpoint of the rectangle $\mathcal{D}_2\cap C$. Let $\mathcal{D}_1'$ be the collection of subcylinders into which the cylinders in $\mathcal{D}_1$ are divided by p. Traveling along $a_{\mathcal{D}_1'}$ back to (X,ω) keeps the complement of $\mathcal{D}_1$ fixed and so the conclusion of the lemma held on $(X,\omega;p)$, as desired. Therefore, we may suppose without loss of generality (by replacing (X,ω) with (Y,η)) that p belongs to $\mathcal{D}_1$.

Part 2 Determining the position of p Suppose now that p belongs to $\mathcal{D}_1$. As before, $a_{\mathcal{D}_2}$ is tangent to $\mathcal{M}'$. Travel from $(X,\omega;p)$ in the $a_{\mathcal{D}_2}$ direction in $\mathcal{M}'$ until the intersection of $\mathcal{D}_2$ and C accounts for at least $(m-1)/m$ proportion of the area of C. Without loss of generality, we may replace $(X,\omega;p)$ with the resulting marked translation surface.

Let ℓ_i be the horizontal length of the rectangle $C\cap\mathcal{D}_i$ for $i=1,2$. The complement of $\mathcal{D}_1\cup\mathcal{D}_2$ in C is two disjoint rectangles. Let a (resp. b) be the horizontal length of the rectangle to the right (resp. left) of $\mathcal{D}_1\cap C$; see Figure 6. Let ℓ be the length of the core curve of C. Let $q\in[0,1]$ be chosen so that p is a distance of $q\ell_1$ from the left boundary of $\mathcal{D}_1\cap C$. Let $\mathcal{D}_1'$ be the collection of subcylinders into which p divides $\mathcal{D}_1$ on $(X,\omega;p)$.

Travel in the $a_{\mathcal{D}_1'}$ direction from $(X,\omega;p)$ so that the length of the core curve of C increases by s and then travel in the $\sigma_{\mathcal{C}'}$ direction to perform exactly one Dehn twist in C (we will assume that the core curves of horizontal cylinders are oriented from left to right; this specifies that we are applying the matrix $\begin{pmatrix} 1 & (\mathrm{Mod}(C))^{-1} \\ 0 & 1 \end{pmatrix}$ to the cylinders in $\mathcal{C}'$, where $\mathrm{Mod}(C)$ denote the modulus of C). The distance of the marked point from the left-hand boundary of $\mathcal{D}_2\cap C$ is

$$h(\ell+s)-(1-q)(\ell_1+s)-a.$$

Traveling back along the $a_{\mathcal{D}_1'}$ direction returns to the unmarked surface (X,ω) while leaving the position of the marked point fixed in the complement of $\mathcal{D}_1$. Since p is a periodic point, the fiber of the forgetful map from $\mathcal{M}'$ to $\mathcal{M}$ over (X,ω) is finite.

Therefore, $h(\ell+s)-(1-q)(\ell_1+s)-a$ is constant as a function of s. In other words, $h=(1-q)$. If we sheared C in the other direction we would have by symmetry that $h=q$ and so $q=h=\frac{1}{2}$. By symmetry, after shearing the marked point into $\mathcal{D}_2$ the distance from the left-hand boundary of $\mathcal{D}_2\cap C$ is $\frac{1}{2}\ell_2$, ie

$$\tfrac{1}{2}(\ell-\ell_1)-a=\tfrac{1}{2}\ell_2.$$

Since $\ell=\ell_1+a+\ell_2+b$ we see that $a=b$, as desired. □

6 Proof of Theorem 1.1

Throughout this section, we make the following assumption.

Assumption 6.1 Let $\mathcal{M}$ be a nonempty affine invariant submanifold in $\mathcal{H}(0^n)$, where $\mathcal{H}$ is an unmarked stratum of abelian differentials. Suppose that $\mathcal{M}$ contains marked points on a translation surface (X,ω) that is generic in $\mathcal{H}$ and that is one of the surfaces constructed in Section 4. Finally, suppose after passing to a finite cover that the marked points are labeled as $\{p_1,\dots,p_n\}$. Let $\pi_k\colon \mathcal{H}^{\mathrm{ord}}(0^n)\to\mathcal{H}^{\mathrm{ord}}(0^{n-1})$ be the map that forgets the k^{th} marked point for $k\in\{1,\dots,n\}$.

Theorem 6.2 *Periodic points exist on (X,ω) if and only if $\mathcal{H}$ is hyperelliptic, in which case they are Weierstrass points.*

Theorem 6.3 *If $n\ge 2$ and $(\pi_k)_*\mathcal{M}=\mathcal{H}^{\mathrm{ord}}(0^{n-1})$ for all $k\in\{1,\dots,n\}$ then $\mathcal{H}$ is hyperelliptic, $n=2$, and the fiber of $\mathcal{M}$ over (X,ω) contains all pairs of distinct points exchanged by the hyperelliptic involution.*

If $\mathcal{H}$ is hyperelliptic, then given a pair $\{i,j\}$ of integers in $\{1,\dots,n\}$, let $\mathcal{H}_{ij}$ denote the subset of $\mathcal{H}(0^n)$ where p_i and p_j are exchanged by the hyperelliptic involution. If $i=j$, this will mean that p_i is a fixed point of the hyperelliptic involution. We will prove the following strengthening of Theorem 1.1.

Theorem 6.4 *If $\mathcal{M}$ is as in Assumption 6.1 then either $\mathcal{M}=\mathcal{H}(0^n)$ or the stratum $\mathcal{H}$ is hyperelliptic and there is a subset S of pairs of integers in $\{1,\dots,n\}$ such that $\mathcal{M}=\bigcap_{\{i,j\}\in S}\mathcal{H}_{ij}$.*

Proof of Theorem 6.4 given Theorem 6.2 and Theorem 6.3 Fix $\mathcal{H}$ and proceed by induction on n. The $n=1$ case is Theorem 6.2. Suppose now that $n>1$ and that $\mathcal{M}\neq\mathcal{H}(0^n)$.

Suppose first that for some $k \in \{1, \dots, n\}$, $(\pi_k)_*\mathcal{M}$ has dimension equal to $\dim \mathcal{M} - 1$. Suppose without loss of generality after relabeling that $k = 1$. By the induction hypothesis, $\mathcal{H}$ is hyperelliptic and $(\pi_1)_*\mathcal{M} = \bigcap_{\{i,j\}\in S} \mathcal{H}_{ij}(0^{n-1})$, where S is some subset of pairs of integers in $\{2, \dots, n\}$. It follows that $\mathcal{M}$ is contained in $\bigcap_{\{i,j\}\in S} \mathcal{H}_{ij}(0^n)$ and therefore coincides with it since both manifolds are connected, closed and of the same dimension.

Suppose now that for all $k \in \{1, \dots, n\}$, $(\pi_k)_*\mathcal{M}$ has the same dimension as $\mathcal{M}$ and that it does not coincide with $\mathcal{H}(0^{n-1})$. By the induction hypothesis, for each k, there is a subset S_k of pairs of integers in $\{1, \dots, n\} \setminus \{k\}$ such that $(\pi_k)_*\mathcal{M} = \bigcap_{\{i,j\}\in S_k} \mathcal{H}_{ij}$. The number of elements of S_k is exactly the codimension of $(\pi_k)_*\mathcal{M}$ in $\mathcal{H}(0^{n-1})$. Suppose without loss of generality that $\{1, \ell_1\}$ is contained in S_{ℓ_2}. Then this pair cannot be contained in S_1 and so $S_1 \cup S_{\ell_2}$ contains at least as many elements as the codimension of $\mathcal{M}$ in $\mathcal{H}(0^n)$. Hence, $\mathcal{M} = \bigcap_{\{i,j\}\in S_1 \cup S_{\ell_2}} \mathcal{H}_{ij}(0^n)$ since both manifolds are connected, closed and of the same dimension.

The only case that remains is when $(\pi_k)_*\mathcal{M} = \mathcal{H}^{\mathrm{ord}}(0^{n-1})$ for all $k \in \{1, \dots, n\}$, which follows from Theorem 6.3. □

Proof of Theorem 1.1 Let (Y, η) be a generic translation surface in a component $\mathcal{H}$ of a stratum of abelian differentials. Let P be a nongeneric collection of n points on (Y, η) and let $\mathcal{M}$ be the orbit closure of $(Y, \eta; P)$ in $\mathcal{H}(0^n)$. By assumption, $\mathcal{M} \neq \mathcal{H}(0^n)$. By Lemma 2.14, for a generic surface (X, ω) as in Assumption 6.1, there is a collection of points Q — which are necessarily nongeneric — such that $(X, \omega; Q) \in \mathcal{M}$. By Theorem 6.4, it follows that $\mathcal{H}$ is hyperelliptic and that there is a subset S of pairs of integers in $\{1, \dots, n\}$ such that $\mathcal{M} = \bigcap_{\{i,j\}\in S} \mathcal{H}_{ij}$. Since P is nongeneric it follows that it must either contain a Weierstrass point or two points exchanged by the hyperelliptic involution. □

7 Proof of Theorem 6.3

Assumption 6.1 will remain in effect for this section. Assume too that $\mathcal{M}$ is a proper affine invariant submanifold in $\mathcal{H}^{\mathrm{ord}}(0^n)$ where $n \geq 2$ and suppose that $(\pi_k)_*\mathcal{M} = \mathcal{H}^{\mathrm{ord}}(0^{n-1})$ for all $k \in \{1, \dots, n\}$. If $\mathcal{H}$ is nonhyperelliptic then let C be the central horizontal cylinder and if $\mathcal{H}$ is hyperelliptic let C be any horizontal cylinder that intersects two vertical cylinders. By Lemmas 4.4 and 4.8, there are two $\mathcal{H}$–equivalence classes of vertical cylinders $\mathcal{D}_1$ and $\mathcal{D}_2$ whose intersection with C is connected and nonempty.

Lemma 7.1 *There are two marked points. If one marked point lies in a cylinder in* $\mathcal{D}_1$ *or* $\mathcal{D}_2$*, the other one lies in that cylinder as well.*

Proof Let $P = \{p_1, \dots, p_{n-1}\}$ be a collection of $n-1$ points where p_1 lies in $\mathcal{D}_1$ and divides the cylinders in $\mathcal{D}_1$ into subcylinders whose moduli admit no rational homogeneous linear relation. Suppose too that the points $p_2, \dots, p_{n-1}$ belong to $\mathcal{D}_2$ and divide it into vertical subcylinders whose moduli also admit no rational homogeneous linear relation. The genericity criterion (Proposition 4.3) implies that $(X, \omega; P)$ is generic in $\mathcal{H}(0^{n-1})$. By Lemma 2.14, there is a point p_n such that $(X, \omega; P \cup \{p_n\})$ belongs to $\mathcal{M}$.

Let $\mathcal{D}$ be an equivalence class in $\{\mathcal{D}_1, \mathcal{D}_2\}$ that does not contain p_n. Since $\mathcal{D}$ is its own equivalence class in $\mathcal{H}$ and since it is divided into subcylinders whose moduli admit no rational homogeneous linear relation, it follows that each subcylinder in $\mathcal{D}$ may be sheared while fixing the rest of the surface. Phrased differently, in the fiber of $\mathcal{M}$ over (X, ω) any marked point in $\mathcal{D}$ may be moved freely while fixing all other points in P. However, if p_k is a point (for $k \in \{1, \dots, n-1\}$) that belongs to $\mathcal{D}$, then the fiber of $\pi_k\colon \mathcal{M} \to \mathcal{H}^{\mathrm{ord}}(0^{n-1})$ is one-dimensional and, by assumption, $(\pi_k)_* \mathcal{M} = \mathcal{H}^{\mathrm{ord}}(0^{n-1})$. This implies that $\mathcal{M} = \mathcal{H}^{\mathrm{ord}}(0^n)$, which is a contradiction. Therefore, there are no points belonging to $\mathcal{D}$ and so $n = 2$.

For the final statement, suppose, for a contradiction, that $\{p_1, p_2\}$ is a fiber of $\mathcal{M}$ over (X, ω) under the map that forgets marked points and suppose too that p_1 belongs to $\mathcal{D}$ for $\mathcal{D} \in \{\mathcal{D}_1, \mathcal{D}_2\}$, but that p_2 does not. Since the map π_2 that forgets the second point is open by Lemma 2.16, there is a nearby surface $(X, \omega; p_1', p_2')$ in $\mathcal{M}$ where p_1' divides the cylinders in $\mathcal{D}$ into subcylinders whose moduli admit no rational homogeneous linear relation and p_2' remains outside of $\mathcal{D}$. This contradicts the previous paragraph. □

Because $\pi_2\colon \mathcal{M} \to \mathcal{H}(0)$ is a finite holomorphic map, we see that given a point $(X, \omega; p_1, p_2) \in \mathcal{M}$ we move p_1 to a new point p_1' (at least locally) and there will be a unique nearby point $(X, \omega; p_1', p_2') \in \mathcal{M}$. Since the equations that define the affine invariant submanifold $\mathcal{M}$ have real coefficients, if p_1 moves horizontally (resp. vertically), so does p_2.

Now suppose without loss of generality that p_1 lies at irrational height in $\mathcal{D}_1$ and at irrational height in C. Let p_2 be a point such that $(X, \omega; p_1, p_2) \in \mathcal{M}$. By Lemma 7.1, p_2 must also lie in the interior $\mathcal{D}_1 \cap C$, which is a rectangle. Moving p_1 to the left we see that p_1 reaches the left boundary of $\mathcal{D}_1 \cap C$ at the same moment that p_2 reaches the vertical boundary. Now reversing direction and moving p_1 to the right we see again that p_1 and p_2 arrive at the vertical boundary of $\mathcal{D}_1$ at the same moment. This implies either that one point lies above the other and both move at the same speed in the same direction (horizontally) or that both points at some point were on opposite boundaries of $\mathcal{D}_1$ and move in opposite directions (horizontally) at the same speed. The first case cannot occur, since if it does we may simply shear the central horizontal cylinder and find two marked points that do not lie above each other and that still move in the same direction at the same speed. The same argument applied to the vertical direction shows that when one point moves at unit speed in the v direction, the other point moves at unit speed in the $-v$ direction. Coupled with the fact that the points arrive at the $\mathcal{D}_1 \cap C$ boundary at the same times we have that when $\mathcal{H}$ is hyperelliptic the two points are exchanged by the hyperelliptic involution.

We will now show that $\mathcal{H}$ must be hyperelliptic. The present situation is pictured in Figure 7. If $\mathcal{H}$ is not hyperelliptic then we may ensure that $a < b$ (by Lemmas 4.4 and 4.8). If p_2 moves to the right at unit speed, then p_1 moves to the left at unit speed and hence p_2 arrives in the interior of $\mathcal{D}_2$ before p_1, contradicting Lemma 7.1.

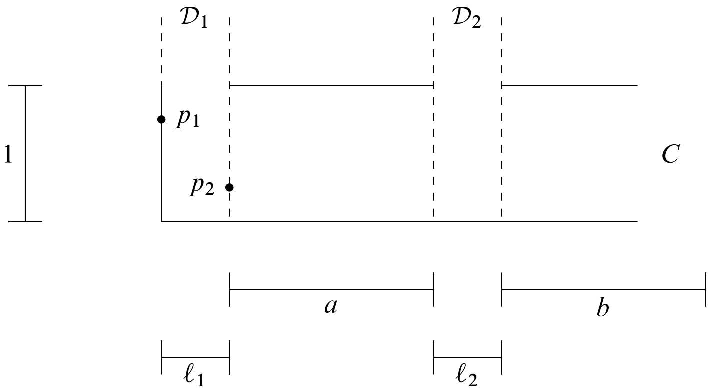

Figure 7: Two marked points in the central horizontal cylinder.

8 Proof of Theorem 6.2

Throughout this section we will make the following assumption.

Assumption 8.1 Let $\mathcal{H}$ be a stratum of abelian differentials and let (X, ω) be a generic translation surface in $\mathcal{H}$ constructed in Section 4. Let $\mathcal{M}$ be an affine invariant submanifold properly contained in $\mathcal{H}(0)$. By Lemma 2.14, the fiber in $\mathcal{M}$ over (X,ω) is nonempty and any point p in the fiber is a periodic point.

Definition 8.2 A cylinder C is called $\mathcal{H}$–free if $\{C\}$ is an $\mathcal{H}$–equivalence class. This is equivalent to the condition that there is no other parallel cylinder C with a core curve that is homologous to the core curve of C.

Proposition 8.3 *The periodic points on $\mathcal{H}^{\text{hyp}}(2g-2)$ and $\mathcal{H}^{\text{hyp}}(g-1,g-1)$ are exactly the Weierstrass points.*

Proof Let $\mathcal{H}$ be either $\mathcal{H}^{\text{hyp}}(2g-2)$ or $\mathcal{H}^{\text{hyp}}(g-1,g-1)$. The surface (X,ω) is pictured again in Figure 8.

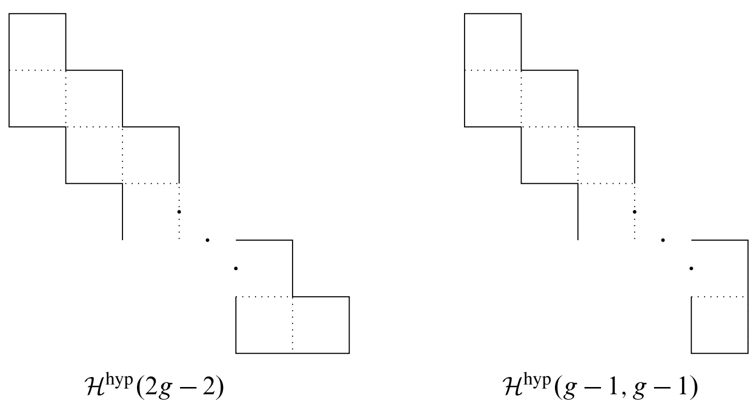

Figure 8: Hyperelliptic translation surfaces.

Each horizontal and vertical cylinder is $\mathcal{H}$–free. By Lemma 5.5 if p is a periodic point in (X,ω) that lies in the interior of a horizontal cylinder that intersects two vertical cylinders, it is automatically a Weierstrass point. The same holds if it lies in the interior of a vertical cylinder that intersects two horizontal ones. We may therefore assume (up to exchanging each instance of the word "horizontal" for the word "vertical" and vice versa) that p lies in the interior of a horizontal cylinder that intersects only one vertical cylinder and on the boundary of a vertical cylinder. By Lemma 5.2, we may shear this cylinder and remain in $\mathcal{M}$. Shearing the horizontal cylinder so as to perform one complete Dehn twist moves the periodic point into the interior of a vertical cylinder that intersects two horizontal cylinders and so we have that p is a Weierstrass point by Lemma 5.5. □

Assumption 8.4 Assume now that $\mathcal{H}$ is nonhyperelliptic and let C be the central horizontal cylinder in (X,ω).

Proposition 8.5 *A periodic point on (X,ω) must lie on the boundary of C and in the interior of a vertical cylinder that is not $\mathcal{H}$–free.*

Proof Let p be a periodic point in (X,ω). By Lemmas 4.4, 4.8 and 5.5 the periodic point cannot lie in the interior of C. We will proceed now by cases based on the containment of p in vertical cylinders.

Case 1 ***p* is contained in a vertical cylinder *V* that is *$\mathcal{H}$*–free, is contained in *C* and only intersects the core curve of *C* once** In this case, Lemma 5.2 implies that we may shear V so as to perform one complete Dehn twist and remain in $\mathcal{M}$. This moves p to a periodic point in the interior of C, which is a contradiction.

Case 2 ***p* is contained in a vertical cylinder *V* that is *$\mathcal{H}$*–free and is contained in *C*** By the previous case, V must intersect the core curve of C at least twice. By construction of the surfaces in Section 4 the situation must be as depicted in Figure 9. The marked point is then contained in a $\mathcal{H}$–free cylinder (drawn in dashed lines). Using Lemma 5.2 to shear this cylinder to perform one complete Dehn twist we see that p may be moved to a periodic point in the interior of C, which is a contradiction.

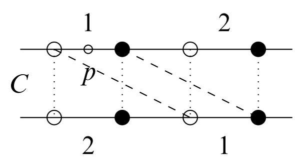

Figure 9: The translation surface in Case 2.

Case 3 ***p* is contained in a *$\mathcal{H}$*–free vertical cylinder *V* that is not contained in *C*** By construction of the surfaces in Section 4, V only intersects two horizontal cylinders, C and H, and core curves intersect exactly once; see Figure 10. Notice that C is $\mathcal{H}$–free since by construction there is a vertical cylinder V' contained in it. The core curve of V' cannot intersect any other horizontal cylinder and so C is not homologous to (equivalently $\mathcal{H}$–equivalent to) any other horizontal cylinder. Similarly, H is $\mathcal{H}$–free since it is not equivalent to C and since applying the matrix $\left(\begin{smallmatrix}1&1\\t&0\end{smallmatrix}\right)$, for small t, to V makes H and C nonhorizontal while fixing all other horizontal cylinders (two equivalent cylinders must remain parallel in a neighborhood of the original surface).

Applying Lemma 5.5, with $\mathcal{D}_1 = \{C\}$ and $\mathcal{D}_2 = \{H\}$ we see that if p lies in the interior of V, then there must be a periodic point in C, which is a contradiction. If p does not lie in the interior of V, then by Lemma 5.2 we may shear H to perform one complete Dehn twist while fixing the remainder of the translation surface and remaining in $\mathcal{M}$. This shear moves p to the interior of V and so we are done.

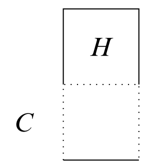

Figure 10: The vertical cylinder V.

Case 4 p is contained in a vertical cylinder that is not contained in C and that is not $\mathcal{H}$–free By construction of the surfaces in Section 4, $\mathcal{H}$ is an even component of a stratum of abelian differentials. Moreover, either p is contained on the boundary of a vertical cylinder, as in Figure 11, or is contained in one of the cylinders passing through the saddle connections labeled $\{1, \dots, 4\}$ in Figure 11.

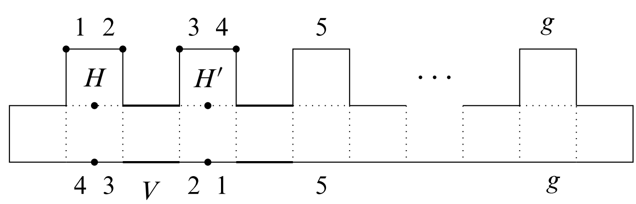

Figure 11: Points on the boundary of C and on the boundary of a vertical cylinder.

Let H and H' be the indicated horizontal cylinders in Figure 11, which are $\mathcal{H}$–free. Suppose without loss of generality that p is contained in the horizontal cylinder H or its boundary. By Lemma 5.2 we may shear the cylinders H and H' to arrive at the surface in Figure 12. Let D be the diagonal cylinder with dashed boundary that passes through the horizontal saddle connection labeled 1. By Lemma 5.2 we may shear H if necessary to perform one complete Dehn twist and move p into the interior of D while remaining in $\mathcal{M}$. By Lemma 5.5 — where D is intersected by the equivalence classes $\{H\}$ and $\{C\}$ — it follows that there is a periodic point contained in the interior of C. By Lemma 5.2, we may shear H and H' while fixing the remainder of the translation surface to return to (X, ω) with a periodic point p in the interior of C, a contradiction.

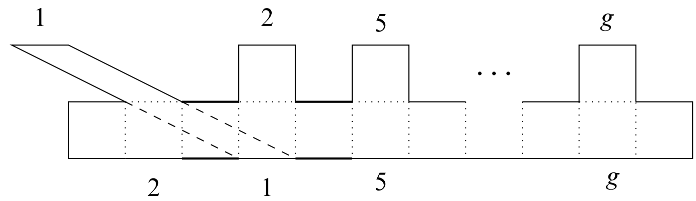

Figure 12: Moving potentially periodic points into the interior of C.

Therefore, p lies on the boundary of C and is contained in the interior of a vertical cylinder which is not $\mathcal{H}$–free. □

Corollary 8.6 *The only component of $\mathcal{H}(2g-2)$ with periodic points is $\mathcal{H}^{\mathrm{hyp}}(2g-2)$.*

Proof Let $\mathcal{H}$ be a minimal nonhyperelliptic stratum of abelian differentials. If $\mathcal{M}$ is an affine invariant submanifold properly contained in $\mathcal{H}(0)$ that pushes forward to $\mathcal{H}$ under the forgetful map, then its fiber over (X,ω) is nonempty by Lemma 2.14. If $(X,\omega;p)$ is an element of the fiber then p is contained in a vertical cylinder that is contained in C and that is not $\mathcal{H}$–free. By construction of the surfaces in Section 4 there are no such cylinders and so we are done. □

Proof of Theorem 6.2 Let $\mathcal{H}$ be a component of a stratum of abelian differentials with at least two zeros. Proceed by induction on $\dim_{\mathbb{C}} \mathcal{H}$. The result has already been established for hyperelliptic components (Proposition 8.3), which establishes the base case and allows us to assume without loss of generality that $\mathcal{H}$ is not a hyperelliptic component. Assume that p is a periodic point on (X,ω). By Proposition 8.5, p is contained on the boundary of C and in a vertical cylinder V that is contained in C and that is not $\mathcal{H}$–free.

Case 1 ***C* contains an *$\mathcal{H}$*–free vertical cylinder *W*** By construction of the surfaces in Section 4, there are only two types of $\mathcal{H}$–free vertical cylinders contained in C, which are pictured in Figure 13.

By Lemma 5.2, if (X,ω) contains an $\mathcal{H}$–free cylinder W we may travel in $\mathcal{M}$ in the direction of the standard shear a_W to shrink the horizontal cross curve of W until it vanishes. This degenerates $\mathcal{M}$ to the boundary of $\mathcal{H}(0)$. Notice that in both cases this degeneration causes two distinct zeros to collide on the boundary, but, by considering Euler characteristic, no genus is lost. By Mirzakhani and Wright [24, Corollary 1.2], the resulting translation surface has an orbit closure of strictly smaller dimension than $\mathcal{M}$.

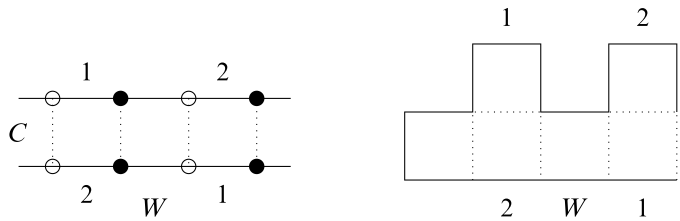

Figure 13: The two types of $\mathcal{H}$–free cylinders W in C on a subsurface of (X, ω) in components with two zeros of odd order (left) and a subsurface of (X, ω) in even components (right).

By the genericity criterion (Proposition 4.3), the boundary translation surface remains generic in the component of the stratum to which it belongs, which necessarily has complex dimension one less than $\mathcal{M}$. Therefore, p remains a periodic point.

By the induction hypothesis, p is a Weierstrass point and the boundary translation surface belongs to a hyperelliptic component. In particular, p must lie halfway across V on the boundary of C, dividing V into two subcylinders of equal modulus, which we will call V_1 and V_2. By construction, on $(X, \omega; p)$ the only rational linear homogeneous equation that holds on moduli of cylinders equivalent to V is that the moduli of V_1 and V_2 are equal. By Wright [29] — in particular Theorems 2.9 and 2.10 — and Lemma 5.2, V may be sheared on $(X, \omega; p)$ while remaining in $\mathcal{M}$ and fixing the remainder of the translation surface. Shearing so as to perform one complete Dehn twist moves p into the interior of C on (X, ω), which contradicts Proposition 8.5.

Notice that as a corollary of this step, $\mathcal{H}$ does not contain two zeros of odd order.

Case 2 The surface belongs to an even component Let H and H' be the horizontal cylinders labeled in Figure 11. By Lemma 5.2, the result of shearing them while fixing the remainder of the surface remains in $\mathcal{M}$. Therefore, we shear them to find the surface in $\mathcal{M}$ depicted in Figure 14, which contains a vertical cylinder that contains H and H' and that passes through them exactly once.

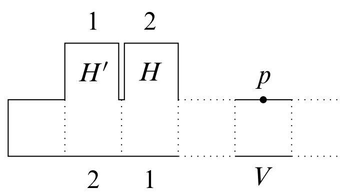

Figure 14: A translation surface in the even component.

By Lemma 5.2, applying the standard dilation $a_{\{H\}}$ to H causes its vertical cross curve to vanish and passes to a surface on the boundary of $\mathcal{M}$. When the cross curve vanishes, the boundary translation surface (shown in Figure 15) has the zero of order $2k+2$ (where k is a positive integer) on the boundary of H split into two zeros, one of order 1 and one of order $2k-1$.

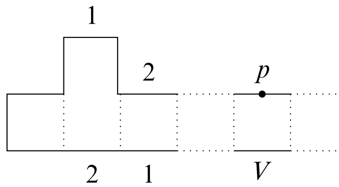

Figure 15: The boundary translation surface (Y, η).

Let $\mathcal{M}'$ be the orbit closure of $(Y, \eta; p)$. By the genericity criterion (Proposition 4.3), (Y, η) remains generic. By Mirzakhani and Wright [24, Corollary 1.2], the dimension of $\mathcal{M}'$ must be strictly smaller than the dimension of $\mathcal{M}$. It follows that p is a periodic point on (Y, η). However, no such points exist on translation surfaces of genus greater than two with a zero of odd order. Thus, (X, ω) cannot belong to an even component.

Case 3 The surface does not belong to an even component The marked point is contained in one of the two configurations shown in Figure 16, where V is adjacent on the right to a vertical cylinder W that contains a horizontal cylinder H (both shown in the figure).

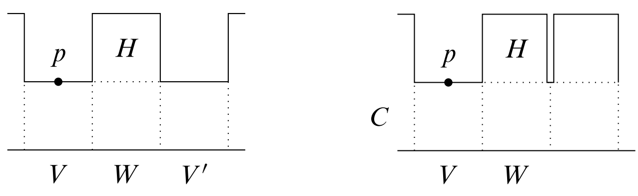

Figure 16: Two possible configurations: the periodic point borders an order 2 zero (left) and the periodic point borders a zero of order 4 or more (right).

By Lemma 5.2, applying the standard dilation $a_{\{H\}}$ to H causes its vertical cross curve to vanish and passes to a surface on the boundary of $\mathcal{M}$. The underlying translation surface moves from $\mathcal{H}(2k, \dots)$ to $\mathcal{H}(0, 2k-2, \dots)$ (see Figure 17).

Let $(Y, \eta; \{p\} \cup Q)$ be the boundary translation surface where Q are the marked points that arise in the degeneration. Let $\mathcal{M}'$ be the orbit closure of $(Y, \eta; \{p\} \cup Q)$. The

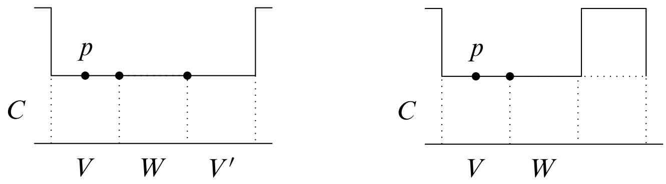

Figure 17: The result of collapsing H: the surface collapses to $\mathcal{H}(0,0,\dots)$ (left) and the surfaces collapses to $\mathcal{H}(0,2k-2,\dots)$ (right).

genericity criterion (Proposition 4.3) implies that $(Y,\eta;Q)$ is generic in the stratum $\mathcal{H}'$ that contains it. By Mirzakhani and Wright [24, Corollary 1.2], $\mathcal{M}'$ is an affine invariant submanifold that is properly contained in $\mathcal{H}'(0)$. The induction hypothesis implies that (Y,η) belongs to a hyperelliptic component and that either p is a Weierstrass point or is exchanged with a point in Q under the hyperelliptic involution.

Consider first the configuration in Figure 17, right. Since $V\cup W$ must be fixed by the hyperelliptic involution and since W may be made arbitrarily long horizontally we see that p is neither a Weierstrass point nor a point exchanged under the hyperelliptic involution with a point in Q.

Consider now the configuration in Figure 17, left, and let V_a and V_b be the left and right subcylinders that p splits V into. We see that p must be exchanged under the hyperelliptic involution with the rightmost point in Q and so V_a and V' have identical moduli. Repeating the argument with the vertical cylinder W' that V borders on the left shows that (X,ω) must contain the subsurface shown in Figure 18 and satisfy the property that the modulus of V' is the same as the modulus of V_a and the modulus of V'' the same as the modulus of V_b.

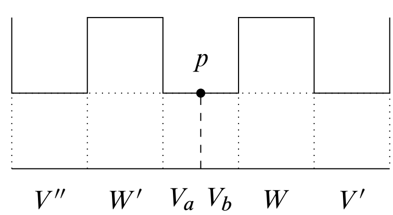

Figure 18: The surface $(X,\omega;p)$.

Letting $\mathrm{Mod}(D)$ denote the modulus of a cylinder D we see that on this surface,

$$\mathrm{Mod}(V)=\mathrm{Mod}(V_a)+\mathrm{Mod}(V_b)=\mathrm{Mod}(V')+\mathrm{Mod}(V''),$$

which is a rational linear homogeneous relation on moduli satisfied by vertical cylinders on (X,ω), which contradicts the fact that (X,ω) is constructed so as to prevent the moduli of vertical cylinders from satisfying such a relation. □

References

[1] **P Apisa**, *Periodic points in genus two: holomorphic sections over Hilbert modular varieties, Teichmüller dynamics, and billiards*, preprint (2017) arXiv

[2] **P Apisa**, $\mathrm{GL}_2\mathbb{R}$ *orbit closures in hyperelliptic components of strata*, Duke Math. J. 167 (2018) 679–742 MR

[3] **P Apisa**, **A Wright**, *Marked points on translation surfaces*, preprint (2017) arXiv

[4] **A Avila**, **A Eskin**, **M Möller**, *Symplectic and isometric* $\mathrm{SL}(2,\mathbb{R})$*–invariant subbundles of the Hodge bundle*, J. Reine Angew. Math. 732 (2017) 1–20 MR

[5] **F Cukierman**, *Families of Weierstrass points*, Duke Math. J. 58 (1989) 317–346 MR

[6] **C J Earle**, **I Kra**, *On sections of some holomorphic families of closed Riemann surfaces*, Acta Math. 137 (1976) 49–79 MR

[7] **A Eskin**, **S Filip**, **A Wright**, *The algebraic hull of the Kontsevich–Zorich cocycle*, Ann. of Math. 188 (2018) 281–313 MR

[8] **A Eskin**, **M Mirzakhani**, *Invariant and stationary measures for the* $\mathrm{SL}(2,\mathbb{R})$ *action on moduli space*, Publ. Math. Inst. Hautes Études Sci. 127 (2018) 95–324 MR

[9] **A Eskin**, **M Mirzakhani**, **A Mohammadi**, *Isolation, equidistribution, and orbit closures for the* $\mathrm{SL}(2,\mathbb{R})$ *action on moduli space*, Ann. of Math. 182 (2015) 673–721 MR

[10] **G Farkas**, **A Verra**, *The universal theta divisor over the moduli space of curves*, J. Math. Pures Appl. 100 (2013) 591–605 MR

[11] **S Filip**, *Splitting mixed Hodge structures over affine invariant manifolds*, Ann. of Math. 183 (2016) 681–713 MR

[12] **Q Gendron**, *The Deligne–Mumford and the incidence variety compactifications of the strata of* $\Omega\mathcal{M}_g$, Ann. Inst. Fourier (Grenoble) 68 (2018) 1169–1240 MR

[13] **S Grushevsky**, **D Zakharov**, *The double ramification cycle and the theta divisor*, Proc. Amer. Math. Soc. 142 (2014) 4053–4064 MR

[14] **J Hubbard**, *Sur la non-existence de sections analytiques à la courbe universelle de Teichmüller*, C. R. Acad. Sci. Paris Sér. A-B 274 (1972) 978–979 MR

[15] **M Kontsevich**, **A Zorich**, *Connected components of the moduli spaces of Abelian differentials with prescribed singularities*, Invent. Math. 153 (2003) 631–678 MR

[16] **A Kumar**, **R E Mukamel**, *Real multiplication through explicit correspondences*, LMS J. Comput. Math. 19 (2016) 29–42 MR

[17] **S Lelièvre**, **T Monteil**, **B Weiss**, *Everything is illuminated*, Geom. Topol. 20 (2016) 1737–1762 MR

[18] **H Masur**, *Interval exchange transformations and measured foliations*, Ann. of Math. 115 (1982) 169–200 MR

[19] **H Masur**, *Closed trajectories for quadratic differentials with an application to billiards*, Duke Math. J. 53 (1986) 307–314 MR

[20] **C T McMullen**, *Billiards and Teichmüller curves on Hilbert modular surfaces*, J. Amer. Math. Soc. 16 (2003) 857–885 MR

[21] **C T McMullen**, *Teichmüller curves in genus two: discriminant and spin*, Math. Ann. 333 (2005) 87–130 MR

[22] **C T McMullen**, *Teichmüller curves in genus two: torsion divisors and ratios of sines*, Invent. Math. 165 (2006) 651–672 MR

[23] **C T McMullen**, *Dynamics of* $\mathrm{SL}_2(\mathbb{R})$ *over moduli space in genus two*, Ann. of Math. 165 (2007) 397–456 MR

[24] **M Mirzakhani**, **A Wright**, *The boundary of an affine invariant submanifold*, Invent. Math. 209 (2017) 927–984 MR

[25] **M Mirzakhani**, **A Wright**, *Full-rank affine invariant submanifolds*, Duke Math. J. 167 (2018) 1–40 MR

[26] **S Mullane**, *On the effective cone of* $\overline{\mathcal{M}}_{g,n}$, Adv. Math. 320 (2017) 500–519 MR

[27] **F Müller**, *The pullback of a theta divisor to* $\overline{\mathcal{M}}_{g,n}$, Math. Nachr. 286 (2013) 1255–1266 MR

[28] **W A Veech**, *Gauss measures for transformations on the space of interval exchange maps*, Ann. of Math. 115 (1982) 201–242 MR

[29] **A Wright**, *Cylinder deformations in orbit closures of translation surfaces*, Geom. Topol. 19 (2015) 413–438 MR

[30] **A N Zemlyakov**, **A B Katok**, *Topological transitivity of billiards in polygons*, Mat. Zametki 18 (1975) 291–300 MR In Russian; translated in Math. Notes 18 (1975) 760–764

[31] **A Zorich**, *Flat surfaces*, from "Frontiers in number theory, physics, and geometry, I" (P Cartier, B Julia, P Moussa, P Vanhove, editors), Springer (2006) 437–583 MR

Department of Mathematics, University of Chicago
Chicago, IL, United States
Current address: *Department of Mathematics, Yale University*
New Haven, CT, United States

paul.apisa@gmail.com

Proposed: Benson Farb Received: 6 September 2018
Seconded: Anna Wienhard, Jean-Pierre Otal Revised: 29 March 2019