SYMBOLIC DYNAMICS FOR THE TEICHMÜLLER FLOW

URSULA HAMENSTÄDT

Abstract. Let S be an oriented surface of genus $g \geq 0$ with $m \geq 0$ punctures and $3g-3+m \geq 2$. The Teichmüller flow Φ^t acts on the moduli space $\mathcal{Q}(S)$ (or $\mathcal{H}(S)$) of area one holomorphic quadratic (or abelian) differentials preserving a collection of so-called strata. For each component $\mathcal{Q}$ of a stratum, we construct a subshift of finite type (Ω, σ) and Borel suspension (X, Θ^t) which admits a finite-to-one semi-conjugacy Ξ into the Teichmüller flow on $\mathcal{Q}$. This is used to show that the Φ^t-invariant Lebesgue measure λ on $\mathcal{Q}$ is the unique measure of maximal entropy. If h is the entropy of λ then for every $\epsilon > 0$ there is a compact set $K \subset \mathcal{Q}$ such that the entropy of any Φ^t-invariant probability measure μ with $\mu(\mathcal{Q}-K) = 1$ does not exceed $h-1+\epsilon$. Moreover, the growth rate of periodic orbits in $\mathcal{Q}-K$ does not exceed $h-1+\epsilon$. This implies that the number of periodic orbits for Φ^t in $\mathcal{Q}$ of period at most R is asymptotic to e^{hR}/hR. Finally we give a unified and simplified proof of exponential mixing for the Lebesgue measure on strata.

Contents

Date: December 28, 2011.
AMS subject classification: 37A20,30F60
Research partially supported by the Hausdorff Center Bonn.

1. Introduction

An oriented surface S of finite type is a closed surface of genus $g \geq 0$ from which $m \geq 0$ points, so-called *punctures*, have been deleted. We assume that $3g - 3 + m \geq 2$, i.e. that S is not a sphere with at most 4 punctures or a torus with at most 1 puncture. We then call the surface S *nonexceptional*. The Euler characteristic of S is negative.

The *Teichmüller space* $\mathcal{T}(S)$ of S is the quotient of the space of all complete finite area hyperbolic metrics on S under the action of the group of diffeomorphisms of S which are isotopic to the identity. The sphere bundle

$$\tilde{\mathcal{Q}}(S) \to \mathcal{T}(S)$$

of all *holomorphic quadratic differentials* of area one can naturally be identified with the unit cotangent bundle for the *Teichmüller metric*. If the surface S has punctures, i.e. if $m > 0$, then we define a holomorphic quadratic differential on S to be a meromorphic quadratic differential on the closed Riemann surface obtained from S by filling in the punctures, with a simple pole at each of the punctures and no other poles.

The *mapping class group* $\mathrm{Mod}(S)$ of all isotopy classes of orientation preserving self-diffeomorphisms of S naturally acts on $\tilde{\mathcal{Q}}(S)$. The quotient

$$\mathcal{Q}(S) = \tilde{\mathcal{Q}}(S)/\mathrm{Mod}(S)$$

is the *moduli space of area one quadratic differentials*. It can be partitioned into so-called *strata*. Namely, let $1 \leq m_1 \leq \cdots \leq m_\ell$ $(\ell \geq 1)$ be a sequence of positive integers with

$$\sum_i m_i = 4g - 4 + m.$$

The stratum $\mathcal{Q}(m_1, \ldots, m_\ell; -m)$ defined by the ℓ-tuple $(m_1, \ldots, m_\ell)$ is the moduli space of pairs (C, φ) where C is a closed Riemann surface of genus g and where φ is an area one meromorphic quadratic differential on C with ℓ zeros of order m_i and m simple poles and which is not the square of a holomorphic one-form.

A stratum $\mathcal{Q}(m_1, \ldots, m_\ell; -m)$ is a real hypersurface in a complex orbifold of complex dimension

$$h = 2g + \ell + m - 2.$$

Masur and Smillie [MS93] showed that the stratum $\mathcal{Q}(m_1, \ldots, m_\ell; -m)$ is non-empty unless $\ell = 2, m_1 = 1, m_2 = 3$ (and S is a closed surface of genus 2). The strata need not be connected, however they have at most two connected components [L08]. The closure in $\mathcal{Q}(S)$ of a component of a stratum $\mathcal{Q}(m_1, \ldots, m_\ell; -m)$ is a union of components of strata $\mathcal{Q}(n_1, \ldots, n_s; -m)$ where $s \leq \ell$ (note that we always fix the number of simple poles).

If the surface S is closed, i.e. if $m = 0$, then we can also consider the bundle

$$\tilde{\mathcal{H}}(S) \to \mathcal{T}(S)$$

of area one *abelian differentials*. It descends to the moduli space $\mathcal{H}(S)$ of holomorphic one-forms defining a singular euclidean metric of area one. Again this moduli space decomposes into a union of strata $\mathcal{H}(k_1, \dots, k_s)$ corresponding to the orders of the zeros of the differentials. Strata are in general not connected, but the number of their connected components is at most 3 [KZ03]. The stratum $\mathcal{H}(k_1, \dots, k_s)$ is a real hypersurface in a complex orbifold of dimension

$$h = 2g + \ell - 1.$$

The *Teichmüller flow* Φ^t acts on $\mathcal{Q}(S)$ (or $\mathcal{H}(S)$) preserving the strata. If $\mathcal{Q}$ is a component of a stratum of abelian differentials then *Rauzy induction* for interval exchange transformations can be used to construct a *symbolic coding* for the Teichmüller flow on $\mathcal{Q}$. Such a coding consists in a suspension of a subshift of finite type by a positive roof function and a semi-conjugacy of an invariant Borel subset of the suspension flow into $(\mathcal{Q}, \Phi^t)$ ([V82], see also [AGY06] for a discussion and references). Rauzy induction has been extended to strata of quadratic differentials by Boissy and Lanneau [BL09].

Our first goal is to construct a new coding for the Teichmüller flow on any component of a stratum.

Theorem 1. *Let $\mathcal{Q}$ be a component of a stratum of quadratic or abelian differentials. Then there is a subshift of finite type (Ω, σ), a σ- invariant Borel set $\mathcal{U} \subset \Omega$, a suspension (X, Θ^t) of σ over $\mathcal{U}$ given by a positive bounded continuous roof function on $\mathcal{U}$ and a finite-to-one semi-conjugacy $\Xi : (X, \Theta^t) \to (\mathcal{Q}, \Phi^t)$ which induces a continuous bijection of the space of σ-invariant Borel probability measures on $\mathcal{U}$ onto the space of Φ^t-invariant Borel probability measures on $\mathcal{Q}$.*

Our construction is valid for strata of abelian differentials, but it is different from Rauzy induction. A dictionary between these two codings has yet to be established.

Let again $\mathcal{Q}$ be a component of a stratum in $\mathcal{Q}(S)$ or $\mathcal{H}(S)$. We use Theorem 1 to investigate the space $\mathcal{M}_{\rm inv}(\mathcal{Q})$ of all Φ^t-invariant Borel probability measures on $\mathcal{Q}$ equipped with the weak*-topology. A specific example of such a measure in the Lebesgue measure class was constructed by Masur and Veech [M82, V86]. This measure λ is ergodic [M82, V86], and its *entropy* h_λ coincides with the complex dimension $2g+\ell+m-2$ (or $2g+\ell-1$) of the complex orbifold defining the stratum (note that we use a normalization for the Teichmüller flow which is different from the one used by Masur and Veech). In particular, the entropy of the Lebesgue measure on the open connected stratum $\mathcal{Q}(1, \dots, 1; -m)$ equals $6g - 6 + 2m$.

Denote by h_ν the entropy of a measure $\nu \in \mathcal{M}_{\rm inv}(\mathcal{Q})$. Define

$$h_{\rm top}(\mathcal{Q}) = \sup\{h_\nu \mid \nu \in \mathcal{M}_{\rm inv}(\mathcal{Q})\}.$$

A *measure of maximal entropy* for the component $\mathcal{Q}$ is a measure $\mu \in \mathcal{M}_{\rm inv}(\mathcal{Q})$ such that $h_\mu = h_{\rm top}(\mathcal{Q})$. A priori, such a measure need not exist. However, using Rauzy induction, Bufetov and Gurevich [BG07] showed that for components of strata of abelian differentials, the Φ^t-invariant probability measure in the Lebesgue measure class is the unique measure of maximal entropy for the component. We use Theorem 1 and the work of Buzzi and Sarig [BS03] to extend this result to all components of strata of quadratic or abelian differentials, with a simpler proof.

Theorem 2. *For every component $\mathcal{Q}$ of a stratum in $\mathcal{Q}(S)$ or $\mathcal{H}(S)$, the Φ^t-invariant Borel probability measure in the Lebesgue measure class is the unique measure of maximal entropy.*

Components $\mathcal{Q}$ of strata are non-compact. Our next result shows that there is an entropy gap for measures supported in the complement of large compact subsets of $\mathcal{Q}$.

Theorem 3. *For every $\epsilon > 0$ there is a compact subset $K \subset \mathcal{Q}$ such that the entropy of every Φ^t-invariant Borel probability measure which gives full mass to $\mathcal{Q} - K$ does not exceed $h - 1 + \epsilon$.*

We use related ideas to count periodic orbits of the Teichmüller flow in the component $\mathcal{Q}$.

Theorem 4. *For every $\epsilon > 0$ there exists a compact set $K \subset \mathcal{Q}$ and a number $c > 0$ with the following property. The number of periodic orbits of Φ^t contained in $\mathcal{Q} - K$ of period at most R does not exceed $ce^{(h-1+\epsilon)R}$.*

A *saddle connection* of a quadratic differential q is a geodesic segment for the singular euclidean metric defined by q which connects two singular points and does not contain a singular point in its interior. A closed subset K of a component $\mathcal{Q}$ of a stratum is compact if and only if there is a number $a > 0$ so that the shortest length of a saddle connection of any point in K is at least a. Thus Theorem 4 is equivalent to the following

Corollary 1. *For every $\epsilon > 0$ there exists a number $\delta > 0$ and a number $c > 0$ with the following property. The number of periodic orbits of Φ^t contained in $\mathcal{Q}$ of period at most R which consist of quadratic differentials with at least one saddle connection of length at most δ does not exceed $ce^{(h-1+\epsilon)R}$.*

For the open stratum in $\mathcal{Q}(S)$, Theorem 4 is due to Eskin and Mirzakhani [EM11]. The general case was independently shown Eskin, Mirzakhani and Rafi [EMR11]. They also obtain a more precise statement for the number of periodic orbits in the subset of $\mathcal{Q}$ of all points with $k \geq 2$ short saddle connection.

Our method provides an upper bound for the number of periodic orbits outside a large compact subset of a stratum which can explicitly be calculated. The growth rate of the number of orbits which are entirely contained in the thin part of the stratum $\mathcal{Q}$ corresponding to a Riemann surface which contains at least one simple closed curve of small extremal length can be estimated in the same way. We leave the precise calculation to the forthcoming paper [H12] where we will show that all these upper bounds are sharp. For example, we find that for every component of a stratum, the growth rate of the number of periodic orbits with one saddle connection of length at most ϵ tends to $h - 1$ as $\epsilon \to 0$ (with the only exception of some strata of small complexity). However, for every $g \geq 2$ there are components of strata on a closed surface of genus g for which the asymptotic growth rate of the number of periodic orbits in the thin part of moduli space equals $h-2$.

Theorem 4 and the main result of [H10c] immediately imply

Corollary 2. *As $R\to\infty$, the number of periodic orbits for Φ^t of length at most R which are contained in $\mathcal{Q}$ is asymptotic to e^{hR}/hR.*

Our symbolic coding can also be used to give a simpler and unified proof of the following result of Avila, Gouezel and Yoccoz [AGY06] and of Avila and Resende [AR09].

Theorem 5. *The Lebesgue measure on components of strata is exponentially mixing for the Teichmüller flow.*

The organization of the paper is as follows. In Section 2 we begin with establishing some properties of train tracks and geodesic laminations needed in the sequel. In Section 3 we associate to each component of a stratum a family of train tracks. This is used in Section 4 to construct for every connected component $\mathcal{Q}$ of a stratum a subshift of finite type (Ω,σ). In Section 5 we define a bounded roof function ρ on an invariant Borel subset of this subshift of finite type and show that there is a semi-conjugacy of the suspension of (Ω,σ) with roof function ρ into the Teichmüller flow. This completes the proof of Theorem 1. In Section 6 we use the results of the earlier sections and the work of Buzzi and Sarig [BS03] to establish Theorem 2. The proof of Theorem 5 is contained in Section 7. This is then used in Section 8 to show Theorem 3. The proof of Theorem 4 is contained in Section 9.

Acknowledement: During the work on this paper, I benefitted from numerous discussions with more collegues than I can list here. Special thanks go to Artur Avila, Yves Benoist, Alex Eskin, Francois Ledrappier, Howard Masur, Curtis McMullen, Maryam Mirzakhani, Kasra Rafi and Omri Sarig. Part of the work was carried out during two visits of the MSRI in Berkeley (in fall 2007 and in fall 2011) and at the Hausdorff Institut in Bonn in spring 2010. I am very grateful to these two institutions for their hospitality.

2. Train tracks and geodesic laminations

In this section we summarize some constructions from [T79, PH92, H10c] which will be used throughout the paper. Furthermore, we introduce a class of train tracks which will be important in the later sections, and we discuss some of their properties.

2.1. **Geodesic laminations.** Let S be an oriented surface of genus $g \geq 0$ with $m \geq 0$ punctures and where $3g - 3 + m \geq 2$. A *geodesic lamination* for a complete hyperbolic structure on S of finite volume is a *compact* subset of S which is foliated into simple geodesics. A geodesic lamination λ is called *minimal* if each of its half-leaves is dense in λ. Thus a simple closed geodesic is a minimal geodesic lamination. A minimal geodesic lamination with more than one leaf has uncountably many leaves and is called *minimal arational*. Every geodesic lamination λ consists of a disjoint union of finitely many minimal components and a finite number of isolated leaves. Each of the isolated leaves of λ either is an isolated closed geodesic and hence a minimal component, or it *spirals* about one or two minimal components [CEG87].

A geodesic lamination λ on S is said to *fill up* S if its complementary regions are all topological discs or once punctured monogons. A *maximal* geodesic lamination is a geodesic lamination whose complementary regions are all ideal triangles or once punctured monogons.

Definition 2.1. A geodesic lamination λ is called *large* if λ fills up S and if moreover λ can be approximated in the *Hausdorff topology* by simple closed geodesics. A maximal large geodesic lamination is called *complete*.

Since every minimal geodesic lamination can be approximated in the Hausdorff topology by simple closed geodesics [CEG87], a minimal geodesic lamination which fills up S is large. However, there are large geodesic laminations with finitely many leaves.

The *topological type* of a large geodesic lamination ν is a tuple

$$(m_1, \dots, m_\ell; -m) \text{ where } 1 \leq m_1 \leq \dots \leq m_\ell, \sum_i m_i = 4g - 4 + m$$

such that the complementary regions of ν which are topological discs are $m_i + 2$-gons. Let

$$\mathcal{LL}(m_1, \dots, m_\ell; -m)$$

be the space of all large geodesic laminations of type $(m_1, \dots, m_\ell; -m)$ equipped with the restriction of the Hausdorff topology for compact subsets of S.

A *measured geodesic lamination* is a geodesic lamination λ together with a translation invariant transverse measure. Such a measure assigns a positive weight to each compact arc in S with endpoints in the complementary regions of λ which intersects λ nontrivially and transversely. The geodesic lamination λ is called the *support* of the measured geodesic lamination; it consists of a disjoint union of minimal components. The space $\mathcal{ML}$ of all measured geodesic laminations on S equipped with the weak*-topology is homeomorphic to $S^{6g-7+2m} \times (0, \infty)$. Its projectivization is the space $\mathcal{PML}$ of all *projective measured geodesic laminations.*

The measured geodesic lamination $\mu \in \mathcal{ML}$ *fills up* S if its support fills up S. This support is then necessarily connected and hence minimal, and for some tuple $(m_1, \dots, m_\ell; -m)$, it defines a point in the set $\mathcal{LL}(m_1, \dots, m_\ell; -m)$. The projectivization of a measured geodesic lamination which fills up S is also said to fill up S. We call $\mu \in \mathcal{ML}$ *strongly uniquely ergodic* if the support of μ fills up S and admits a unique transverse measure up to scale.

There is a continuous symmetric pairing $\iota : \mathcal{ML} \times \mathcal{ML} \to [0, \infty)$, the so-called *intersection form*, which extends the geometric intersection number between simple closed curves.

2.2. **Train tracks.** A *train track* on S is an embedded 1-complex $\tau \subset S$ whose edges (called *branches*) are smooth arcs with well-defined tangent vectors at the endpoints. At any vertex (called a *switch*) the incident edges are mutually tangent. Through each switch there is a path of class C^1 which is embedded in τ and contains the switch in its interior. A simple closed curve component of τ contains a unique bivalent switch, and all other switches are at least trivalent. The complementary regions of the train track have negative Euler characteristic, which means that they are different from discs with $0, 1$ or 2 cusps at the boundary and different from annuli and once-punctured discs with no cusps at the boundary. We always identify train tracks which are isotopic. Throughout we use the book [PH92] as the main reference for train tracks.

A train track is called *generic* if all switches are at most trivalent. For each switch v of a generic train track τ which is not contained in a simple closed curve component, there is a unique half-branch b of τ which is incident on v and which is *large* at v. This means that every germ of an arc of class C^1 on τ which passes through v also passes through the interior of b. A half-branch which is not large is called *small.* A branch b of τ is called *large* (or *small*) if each of its two half-branches is large (or small). A branch which is neither large nor small is called *mixed.*

Remark: As in [H09a], all train tracks are assumed to be generic. Unfortunately this leads to a small inconsistency of our terminology with the terminology found in the literature.

A *trainpath* on a train track τ is a C^1-immersion $\rho : [k, \ell] \to \tau$ such that for every $i < \ell - k$ the restriction of ρ to $[k + i, k + i + 1]$ is a homeomorphism onto a branch of τ. More generally, we call a C^1-immersion $\rho : [a, b] \to \tau$ a *generalized trainpath.* A trainpath $\rho : [k, \ell] \to \tau$ is *closed* if $\rho(k) = \rho(\ell)$ and if the extension ρ' defined by $\rho'(t) = \rho(t)$ $(t \in [k, \ell])$ and $\rho'(\ell + s) = \rho(k + s)$ $(s \in [0, 1])$ is a trainpath.

A generic train track τ is *orientable* if there is a consistent orientation of the branches of τ such that at any switch s of τ, the orientation of the large half-branch incident on s extends to the orientation of the two small half-branches incident on s. If C is a complementary polygon of an oriented train track then the number of sides of C is even. In particular, a train track which contains a once punctured monogon component, i.e. a once punctured disc with one cusp at the boundary, is not orientable (see p.31 of [PH92] for a more detailed discussion).

A train track or a geodesic lamination η is *carried* by a train track τ if there is a map $F : S \to S$ of class C^1 which is homotopic to the identity and maps η into τ in such a way that the restriction of the differential of F to the tangent space of η

vanishes nowhere; note that this makes sense since a train track has a tangent line everywhere. We call the restriction of F to η a *carrying map* for η. Write $\eta \prec \tau$ if the train track η is carried by the train track τ. Then every geodesic lamination ν which is carried by η is also carried by τ.

A train track *fills up* S if its complementary components are topological discs or once punctured monogons. Note that such a train track τ is connected. Let $\ell \geq 1$ be the number of those complementary components of τ which are topological discs. Each of these discs is an $m_i + 2$-gon for some $m_i \geq 1$ $(i = 1, \dots, \ell)$. The *topological type* of τ is defined to be the ordered tuple $(m_1, \dots, m_\ell; -m)$ where $1 \leq m_1 \leq \cdots \leq m_\ell$; then $\sum_i m_i = 4g-4+m$. If τ is orientable then $m = 0$ and m_i is even for all i. A train track of topological type $(1, \dots, 1; -m)$ is called *maximal*. The complementary components of a maximal train track are all trigons, i.e. topological discs with three cusps at the boundary, or once punctured monogons.

A *transverse measure* on a generic train track τ is a nonnegative weight function μ on the branches of τ satisfying the *switch condition*: for every trivalent switch s of τ, the sum of the weights of the two small half-branches incident on s equals the weight of the large half-branch. The space $\mathcal{V}(\tau)$ of all transverse measures on τ has the structure of a cone in a finite dimensional real vector space, and it is naturally homeomorphic to the space of all measured geodesic laminations whose support is carried by τ. The train track is called *recurrent* if it admits a transverse measure which is positive on every branch. We call such a transverse measure μ *positive*, and we write $\mu > 0$ (see [PH92] for more details).

A *subtrack* σ of a train track τ is a subset of τ which is itself a train track. Then σ is obtained from τ by removing some of the branches, and we write $\sigma < \tau$. If b is a small branch of τ which is incident on two distinct switches of τ then the graph σ obtained from τ by removing b is a subtrack of τ. We then call τ a *simple extension* of σ. Note that formally to obtain the subtrack σ from $\tau - b$ we may have to delete the switches on which the branch b is incident.

Lemma 2.2. (1) *A simple extension τ of a recurrent non-orientable connected train track σ is recurrent. Moreover,*

$$\dim \mathcal{V}(\sigma) = \dim \mathcal{V}(\tau) - 1.$$

(2) *An orientable simple extension τ of a recurrent orientable connected train track σ is recurrent. Moreover,*

$$\dim \mathcal{V}(\sigma) = \dim \mathcal{V}(\tau) - 1.$$

Proof. If τ is a simple extension of a train track σ then σ can be obtained from τ by the removal of a small branch b which is incident on two distinct switches s_1, s_2. Then s_i is an interior point of a branch b_i of σ $(i = 1, 2)$.

If σ is connected, non-orientable and recurrent then there is a trainpath $\rho_0 : [0, t] \to \tau - b$ which begins at s_1, ends at s_2 and such that the half-branch $\rho_0[0, 1/2]$ is small at $s_1 = \rho_0(0)$ and that the half-branch $\rho_0[t - 1/2, t]$ is small at $s_2 = \rho_0(t)$. Extend ρ_0 to a closed trainpath ρ on $\tau - b$ which begins and ends at s_1. This is possible since σ is non-orientable, connected and recurrent. There is a closed trainpath $\rho' : [0, u] \to \tau$ which can be obtained from ρ by replacing the trainpath

ρ_0 by the branch b traveled through from s_1 to s_2. The counting measure of ρ' on τ satisfies the switch condition and hence it defines a transverse measure on τ which is positive on b. On the other hand, every transverse measure on σ defines a transverse measure on τ. Thus since σ is recurrent and since the sum of two transverse measures on τ is again a transverse measure, the train track τ is recurrent as well. Moreover, we have $\dim\mathcal{V}(\tau) \geq \dim\mathcal{V}(\sigma) + 1$.

Let p be the number of branches of τ. Label the branches of τ with the numbers $\{1, \ldots, p\}$ so that the number p is assigned to b. Let $e_1, \ldots, e_p$ be the standard basis of $\mathbb{R}^p$ and define a linear map $A : \mathbb{R}^p \to \mathbb{R}^p$ by $A(e_i) = e_i$ for $i \leq p - 1$ and $A(e_p) = \sum_i \nu(i)e_i$ where ν is the weight function on $\{1, \ldots, p - 1\}$ defined by the trainpath ρ_0. The map A is a surjection onto a linear subspace of $\mathbb{R}^p$ of codimension one, moreover A preserves the linear subspace V of $\mathbb{R}^p$ defined by the switch conditions for τ. In particular, the corank of $A(V)$ in V is at most one. However, $A(V)$ is contained in the space of solutions of the switch conditions on σ and consequently the dimension of the space of transverse measures on σ is not smaller than the dimension of the space of transverse measures on τ minus one.

Together with the first paragraph of this proof, we conclude that $\dim\mathcal{V}(\tau) = \dim\mathcal{V}(\sigma) + 1$. This completes the proof of the first part of the lemma. The second part follows in exactly the same way. □

As a consequence we obtain

Corollary 2.3. (1) $\dim\mathcal{V}(\tau) = 2g-2+m+\ell$ *for every non-orientable recurrent train track* τ *of topological type* $(m_1, \ldots, m_\ell; -m)$.
(2) $\dim\mathcal{V}(\tau) = 2g - 1 + \ell$ *for every orientable recurrent train track* τ *of topological type* $(m_1, \ldots, m_\ell; 0)$.

Proof. The disc components of a non-orientable recurrent train track τ of topological type $(m_1, \ldots, m_\ell; -m)$ can be subdivided in $4g - 4 + m - \ell$ steps into trigons by successively adding small branches. A repeated application of Lemma 2.2 shows that the resulting train track η is maximal and recurrent. Since for every maximal recurrent train track η we have $\dim\mathcal{V}(\eta) = 6g - 6 + 2m$ (see [PH92]), the first part of the corollary follows.

To show the second part, let τ be an orientable recurrent train track of type $(m_1, \ldots, m_\ell; 0)$. Then m_i is even for all i. Add a branch b_0 to τ which cuts some complementary component of τ into a trigon and a second polygon with an odd number of sides. The resulting train track η_0 is not recurrent since a trainpath on η_0 can only pass through b_0 at most once. However, we can add to η_0 another small branch b_1 which cuts some complementary component of η_0 with at least 4 sides into a trigon and a second polygon such that the resulting train track η is non-orientable and recurrent. The inward pointing tangent of b_1 is chosen in such a way that there is a trainpath traveling both through b_0 and b_1. The counting measure of any simple closed curve which is carried by η gives equal weight to the branches b_0 and b_1. But this just means that $\dim\mathcal{V}(\eta) = \dim\mathcal{V}(\tau)+1$ (see the proof of Lemma 2.2 for a detailed argument). By the first part of the corollary, we have $\dim\mathcal{V}(\eta) = 2g - 2 + \ell + 2$ which completes the proof. □

Definition 2.4. A train track τ of topological type $(m_1,\dots,m_\ell;-m)$ is *fully recurrent* if τ carries a large minimal geodesic lamination $\nu\in\mathcal{LL}(m_1,\dots,m_\ell;-m)$.

Note that by definition, a fully recurrent train track is connected and fills up S. The next lemma gives some first property of a fully recurrent train track τ. For its proof, recall that there is a natural homeomorphism of $\mathcal{V}(\tau)$ onto the subspace of $\mathcal{ML}$ of all measured geodesic laminations carried by τ.

Lemma 2.5. *A fully recurrent train track τ of topological type $(m_1,\dots,m_\ell;-m)$ is recurrent.*

Proof. A fully recurrent train track τ of type $(m_1,\dots,m_\ell;-m)$ carries a minimal large geodesic lamination $\nu\in\mathcal{LL}(m_1,\dots,m_\ell;-m)$. The carrying map $\nu\to\tau$ induces a bijection between the complementary components of τ and the complementary components of ν. In particular, a carrying map $\nu\to\tau$ is surjective. Since a minimal geodesic lamination supports a transverse measure, there is a positive transverse measure on τ. In other words, τ is recurrent. □

There are two simple ways to modify a fully recurrent train track τ to another fully recurrent train track. Namely, if b is a mixed branch of τ then we can *shift* τ along b to a new train track τ'. This new train track carries τ and hence it is fully recurrent since it carries every geodesic lamination which is carried by τ [PH92, H09a].

Similarly, if e is a large branch of τ then we can perform a right or left *split* of τ at e as shown in Figure A below. A (right or left) split τ' of a train track τ is carried by τ. If τ is of topological type $(m_1,\dots,m_\ell;-m)$, if $\nu\in\mathcal{LL}(m_1,\dots,m_\ell;-m)$ is minimal and is carried by τ and if e is a large branch of τ, then there is a unique choice of a right or left split of τ at e such that the split track η carries ν. In particular, η is fully recurrent. Note however that there may be a split of τ at e such that the split track is not fully recurrent any more (see Section 2 of [H09a] for details).

Figure A

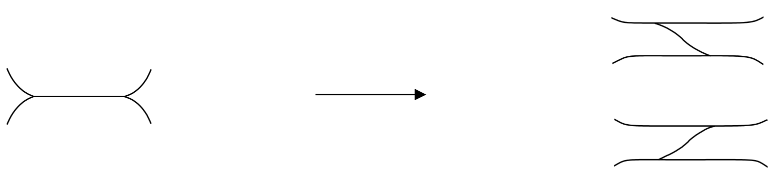

The following simple observation is used to identify fully recurrent train tracks.

Lemma 2.6. (1) *Let e be a large branch of a fully recurrent non-orientable train track τ. Then no component of the train track σ obtained from τ by splitting τ at e and removing the diagonal of the split is orientable.*

(2) *Let e be a large branch of a fully recurrent orientable train track τ. Then the train track σ obtained from τ by splitting τ at e and removing the diagonal of the split is connected.*

Proof. Let τ be a fully recurrent non-orientable train track of topological type $(m_1, \dots, m_\ell; -m)$. Let e be a large branch of τ and let v be a switch on which the branch e is incident. Let σ be the train track obtained from τ by splitting τ at e and removing the diagonal branch of the split. Then the train tracks τ_1, τ_2 obtained from τ by a right and left split at e, respectively, are simple extensions of σ.

Now assume that σ contains an orientable connected component σ_1 (not necessarily distinct from σ). Let $b_i \in \tau_i - \sigma$ be a diagonal of the split connecting τ to τ_i $(i = 1, 2)$. If $\rho_i : [0, m] \to \tau_i$ is a trainpath with $\rho_i[0, 1] = b_i$ and $\rho_i[1, 2] \in \sigma_1$ and if $\rho[j, j + 1] = b_i$ for some $j \geq 2$ then $\rho[j - 1, j]$ equals the branch $\rho_i[1, 2]$ traveled through in opposite direction. Since σ_1 is orientable, this is impossible. Therefore $\rho_i[1, m] \subset \sigma_1$ and hence once again, τ_i is not recurrent. On the other hand, since τ is fully recurrent, it can be split at e to a fully recurrent and hence recurrent train track. This is a contradiction. The first part of the corollary is proven. The second part follows from the same argument since a split of an orientable train track is orientable. $\square$

Example: 1) Figure B below shows a non-orientable recurrent train track τ of type $(4; 0)$ on a closed surface of genus two. The train track obtained from τ by a split at the large branch e and removal of the diagonal of the split track is orientable and hence τ is not fully recurrent. This corresponds to the fact established by Masur and Smillie [MS93] that every quadratic differential with a single zero and no pole on a surface of genus 2 is the square of a holomorphic one-form (see Section 3 for more information).

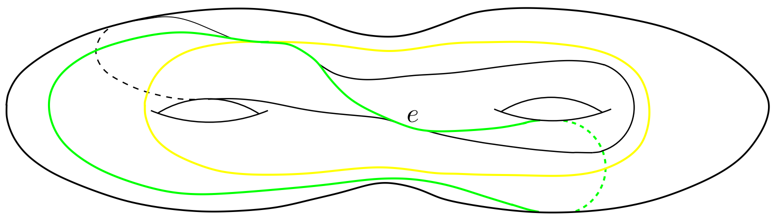

Figure B

2) To construct an orientable recurrent train track of type $(m_1, \dots, m_\ell; 0)$ which is not fully recurrent let S_1 be a surface of genus $g_1 \geq 2$ and let τ_1 be an orientable fully recurrent train track on S_1 with $\ell_1 \geq 1$ complementary components. Choose a complementary component C_1 of τ_1 in S_1, remove from C_1 a disc D_1 and glue two copies of $S_1 - D_1$ along the boundary of D_1 to a surface S of genus $2g_1$. The two copies of τ_1 define a recurrent disconnected oriented train track τ on S which has an annulus complementary component C.

Choose a branch b_1 of τ in the boundary of C. There is a corresponding branch b_2 in the second boundary component of C. Glue a compact subarc of b_1 contained in the interior of b_1 to a compact subarc of b_2 contained in the interior of b_2 so that the images of the two arcs under the glueing form a large branch e in the resulting train track η. The train track η is recurrent and orientable, and its complementary components are topological discs. However, by Lemma 2.6 it is not fully recurrent.

To each train track τ which fills up S one can associate a *dual bigon track* τ^* (Section 3.4 of [PH92]). There is a bijection between the complementary components of τ and those complementary components of τ^* which are not *bigons*, i.e. discs with two cusps at the boundary. This bijection maps a component C of τ which is an n-gon for some $n \geq 3$ to an n-gon component of τ^* contained in C, and it maps a once punctured monogon C to a once punctured monogon contained in C. If τ is orientable then the orientation of S and an orientation of τ induce an orientation on τ^*, i.e. τ^* is orientable.

There is a notion of carrying for bigon tracks which is analogous to the notion of carrying for train tracks. Measured geodesic laminations which are carried by the bigon track τ^* can be described as follows. A *tangential measure* on a train track τ of type $(m_1, \dots, m_\ell; -m)$ assigns to a branch b of τ a weight $\mu(b) \geq 0$ such that for every complementary k-gon of τ with consecutive sides $c_1, \dots, c_k$ and total mass $\mu(c_i)$ (counted with multiplicities) the following holds true.

(1) $\mu(c_i) \leq \mu(c_{i-1}) + \mu(c_{i+1})$.
(2) $\sum_{i=j}^{k+j-1} (-1)^{i-j}\mu(c_i) \geq 0$, $j = 1, \dots, k$.

(The complementary once punctured monogons define no constraint on tangential measures). The space of all tangential measures on τ has the structure of a convex cone in a finite dimensional real vector space. By the results from Section 3.4 of [PH92], every tangential measure on τ determines a simplex of measured geodesic laminations which *hit* τ *efficiently*. The supports of these measured geodesic laminations are carried by the bigon track τ^*, and every measured geodesic lamination which is carried by τ^* can be obtained in this way. The dimension of this simplex equals the number of complementary components of τ with an even number of sides. The train track τ is called *transversely recurrent* if it admits a tangential measure which is positive on every branch.

In general, there are many tangential measures which correspond to a fixed measured geodesic lamination ν which hits τ efficiently. Namely, let s be a switch of τ and let a, b, c be the half-branches of τ incident on s and such that the half-branch a is large. If β is a tangential measure on τ which determines the measured geodesic lamination ν then it may be possible to drag the switch s across some of the leaves of ν and modify the tangential measure β on τ to a tangential measure $\mu \neq \beta$. Then $\beta - \mu$ is a multiple of a vector of the form $\delta_a - \delta_b - \delta_c$ where δ_w denotes the function on the branches of τ defined by $\delta_w(w) = 1$ and $\delta_w(a) = 0$ for $a \neq w$.

Definition 2.7. A train track τ of topological type $(m_1,\dots,m_\ell;-m)$ is called *fully transversely recurrent* if its dual bigon track τ^* carries a large minimal geodesic lamination $\nu \in \mathcal{LL}(m_1,\dots,m_\ell;-m)$. A train track τ of topological type $(m_1,\dots,m_\ell;-m)$ is called *large* if τ is fully recurrent and fully transversely recurrent. A large train track of type $(1,\dots,1;-m)$ is called *complete*.

For a large train track τ let $\mathcal{V}^*(\tau) \subset \mathcal{ML}$ be the set of all measured geodesic laminations whose support is carried by τ^*. Each of these measured geodesic laminations corresponds to a tangential measure on τ. With this identification, the pairing

$$(\nu,\mu) \in \mathcal{V}(\tau) \times \mathcal{V}^*(\tau) \to \sum_b \nu(b)\mu(b) \tag{1}$$

is just the restriction of the intersection form on measured lamination space (Section 3.4 of [PH92]). Moreover, $\mathcal{V}^*(\tau)$ is naturally homeomorphic to a convex cone in a real vector space. The dimension of this cone coincides with the dimension of $\mathcal{V}(\tau)$.

Denote by $\mathcal{LT}(m_1,\dots,m_\ell;-m)$ the set of all isotopy classes of large train tracks on S of type $(m_1,\dots,m_\ell;-m)$.

Remark: In [MM99], Masur and Minsky define a large train track to be a train track τ whose complementary components are topological discs or once punctured monogons, without the requirement that τ is generic, transversely recurrent or recurrent. We hope that this inconsistency of terminology does not lead to any confusion.

3. Strata

The goal of this subsection is to relate components of strata in $\mathcal{Q}(S)$ to large train tracks.

For a closed oriented surface S of genus $g \geq 0$ with $m \geq 0$ punctures let $\tilde{\mathcal{Q}}(S)$ be the bundle of marked area one holomorphic quadratic differentials with a simple pole at each puncture over the Teichmüller space $\mathcal{T}(S)$ of marked complex structures on S. For a complete hyperbolic metric on S of finite area, an area one quadratic differential $q \in \tilde{\mathcal{Q}}(S)$ is determined by a pair (λ^+,λ^-) of measured geodesic laminations which *jointly fill up* S (i.e. we have $\iota(\lambda^+,\mu)+\iota(\lambda^-,\mu) > 0$ for every measured geodesic lamination μ) and such that $\iota(\lambda^+,\lambda^-)=1$. The *vertical* measured geodesic lamination λ^+ for q corresponds to the equivalence class of the vertical measured foliation of q. The *horizontal* measured geodesic lamination λ^- for q corresponds to the equivalence class of the horizontal measured foliation of q.

A tuple $(m_1,\dots,m_\ell)$ of positive integers $1 \leq m_1 \leq \dots \leq m_\ell$ with $\sum_i m_i = 4g-4+m$ defines a *stratum* $\tilde{\mathcal{Q}}(m_1,\dots,m_\ell;-m)$ in $\tilde{\mathcal{Q}}(S)$. This stratum consists of all marked area one quadratic differentials with m simple poles and ℓ zeros of order $m_1,\dots,m_\ell$ which are not squares of holomorphic one-forms. The stratum is a real hypersurface in a complex manifold of dimension

$$h = 2g-2+m+\ell. \tag{2}$$

The closure in $\tilde{\mathcal{Q}}(S)$ of a stratum is a union of components of strata. Strata are invariant under the action of the mapping class group $\mathrm{Mod}(S)$ of S and hence they project to strata in the moduli space $\mathcal{Q}(S) = \tilde{\mathcal{Q}}(S)/\mathrm{Mod}(S)$ of quadratic differentials on S with a simple pole at each puncture. We denote the projection of the stratum $\tilde{\mathcal{Q}}(m_1, \dots, m_\ell; -m)$ by $\mathcal{Q}(m_1, \dots, m_\ell; -m)$. The strata in moduli space need not be connected, but their connected components have been identified by Lanneau [L08]. A stratum in $\mathcal{Q}(S)$ has at most two connected components.

Similarly, if $m = 0$ then we let $\tilde{\mathcal{H}}(S)$ be the bundle of marked area one holomorphic one-forms over Teichmüller space $\mathcal{T}(S)$ of S. For a tuple $k_1 \leq \dots \leq k_\ell$ of positive integers with $\sum_i k_i = 2g-2$, the stratum $\tilde{\mathcal{H}}(k_1, \dots, k_\ell)$ of marked area one holomorphic one-forms on S with ℓ zeros of order k_i $(i = 1, \dots, \ell)$ is a real hypersurface in a complex manifold of dimension

$$h = 2g - 1 + \ell. \tag{3}$$

It projects to a stratum $\mathcal{H}(k_1, \dots, k_\ell)$ in the moduli space $\mathcal{H}(S)$ of area one holomorphic one-forms on S. Strata of holomorphic one-forms in moduli space need not be connected, but the number of connected components of a stratum is at most three [KZ03].

We continue to use the assumptions and notations from Section 2. For a large train track $\tau \in \mathcal{LT}(m_1, \dots, m_\ell; -m)$ let

$$\mathcal{V}_0(\tau) \subset \mathcal{V}(\tau)$$

be the set of all measured geodesic laminations $\nu \in \mathcal{ML}$ whose support is carried by τ and such that the total weight of the transverse measure on τ defined by ν equals one. Let

$$\mathcal{Q}(\tau) \subset \tilde{\mathcal{Q}}(S)$$

be the set of all marked area one quadratic differentials whose vertical measured geodesic lamination is contained in $\mathcal{V}_0(\tau)$ and whose horizontal measured geodesic lamination is carried by the dual bigon track τ^* of τ. By definition of a large train track, we have $\mathcal{Q}(\tau) \neq \emptyset$. The next proposition relates $\mathcal{Q}(\tau)$ to components of strata. For the purpose of its proof and for later use, define the *strong unstable manifold* $W^{su}(q)$ of a quadratic differential $q \in \tilde{\mathcal{Q}}(S)$ to consist of all quadratic differentials whose horizontal measured geodesic lamination coincides with the horizontal measured geodesic lamination of q. The *strong stable manifold* $W^{ss}(q)$ is defined to be the image of $W^{su}(-q)$ under the flip $\mathcal{F} : q \to -q$. For a component $\tilde{\mathcal{Q}}$ of a stratum in $\tilde{\mathcal{Q}}(S)$ and every $q \in \tilde{\mathcal{Q}}$, define the strong unstable (or strong stable) manifold $W^{su}_{\tilde{\mathcal{Q}}}(q)$ (or $W^{ss}_{\tilde{\mathcal{Q}}}(q)$) to be the connected component containing q of the intersection $W^{su}(q) \cap \tilde{\mathcal{Q}}$ (or $W^{ss}(q) \cap \tilde{\mathcal{Q}}$). Then $W^{i}_{\tilde{\mathcal{Q}}}(q)$ is a manifold of dimension $h-1$ $(i = ss, su)$. The manifolds $W^{su}_{\tilde{\mathcal{Q}}}(q)$ (or $W^{ss}_{\tilde{\mathcal{Q}}}(q)$) define a foliation of $\tilde{\mathcal{Q}}$ which is called the *strong unstable* (or the *strong stable*) foliation.

Proposition 3.1. (1) *For every large non-orientable train track $\tau \in \mathcal{LT}(m_1, \dots, m_\ell; -m)$ there is a component $\tilde{\mathcal{Q}}$ of the stratum $\tilde{\mathcal{Q}}(m_1, \dots, m_\ell; -m)$ such that for every $\delta > 0$ the set $\{\Phi^t q \mid q \in \mathcal{Q}(\tau), t \in [-\delta, \delta]\}$ is the closure in $\tilde{\mathcal{Q}}(S)$ of an open subset of $\tilde{\mathcal{Q}}$.*

(2) *For every large orientable train track* $\tau \in \mathcal{LT}(m_1,\dots,m_\ell;0)$ *there is a component* $\tilde{\mathcal{Q}}$ *of the stratum* $\tilde{\mathcal{H}}(m_1/2,\dots,m_\ell/2)$ *such that for every* $\delta>0$ *the set* $\{\Phi^t q \mid q\in\mathcal{Q}(\tau), t\in[-\delta,\delta]\}$ *is the closure in* $\tilde{\mathcal{H}}(S)$ *of an open subset of* $\tilde{\mathcal{Q}}$.

Proof. Let $z\in\tilde{\mathcal{Q}}(S)$ be a marked quadratic differential. A *saddle connection* for z is a geodesic segment for the singular euclidean metric defined by z which connects two singular points and does not contain a singular point in its interior. A *separatrix* is a maximal geodesic segment or ray which begins at a singular point and does not contain a singular point in its interior.

Let ξ be the support of the vertical measured geodesic lamination of z. By [L83], the geodesic lamination ξ can be obtained from the vertical foliation of z by cutting S open along each vertical separatrix and straightening the remaining leaves with respect to a complete finite area hyperbolic metric on S. In particular, up to homotopy, a vertical saddle connection s of z is contained in the interior of a complementary component C of ξ which is uniquely determined by s.

Let $\tau\in\mathcal{LT}(m_1,\dots,m_\ell;-m)$. Consider first the case that τ is non-orientable. Let $\mu\in\mathcal{V}_0(\tau)$, with support $\mathrm{supp}(\mu)$ contained in $\mathcal{LL}(m_1,\dots,m_\ell;-m)$. Then $\mathrm{supp}(\mu)$ is non-orientable since otherwise τ inherits an orientation from $\mathrm{supp}(\mu)$. If $\nu\in\mathcal{V}^*(\tau)$ then the measured geodesic laminations μ,ν jointly fill up S (since the support of ν is different from the support of μ and $\mathrm{supp}(\mu)$ fills up S) and hence if ν is normalized in such a way that $\iota(\mu,\nu)=1$ then the pair (μ,ν) defines a point $q\in\mathcal{Q}(\tau)$. Our first goal is to show that $q\in\tilde{\mathcal{Q}}(m_1,\dots,m_\ell;-m)$.

Since $\mathrm{supp}(\mu)\in\mathcal{LL}(m_1,\dots,m_\ell;-m)$, the orders of the zeros of the quadratic differential q are obtained from the orders $m_1,\dots,m_\ell$ by subdivision. There is a non-trivial subdivision, say of the form $m_i=\sum_s k_s$, if and only if q has at least one vertical saddle connection.

Choose a complete finite area hyperbolic metric on S. This choice identifies the universal covering of S with the hyperbolic plane $\mathbf{H}^2$, and it identifies the fundamental group $\pi_1(S)$ of S with a group of isometries of $\mathbf{H}^2$. Assume to the contrary that q has a vertical saddle connection s. Let $\tilde q$ be the lift of q to a quadratic differential on $\mathbf{H}^2$ and let $\tilde s\subset\mathbf{H}^2$ be a preimage of s.

The preimage $\zeta\subset\mathbf{H}^2$ of $\mathrm{supp}(\mu)$ is a closed $\pi_1(S)$-invariant set of geodesic lines in $\mathbf{H}^2$. Since μ fills up S, the complementary components of ζ are finite area ideal polygons and half-planes which are the components of the preimages of the once punctured monogons. As discussed in the second paragraph of this proof, up to homotopy the saddle connection $\tilde s$ of $\tilde q$ is contained in a complementary component $\tilde C$ of ζ. This component is an ideal polygon with finitely many sides, and it is determined by $\tilde s$.

For the singular euclidean metric defined by q, each cusp of S is a cone point with cone angle π. Let $K\subset S$ be the complement of a standard neighborhood of the cusps. A smooth geodesic arc α for the singular euclidean metric on S defined by q (i.e. an arc which does not pass through singular points) with endpoints in K can be modified near the cusps to a homotopic arc α' of roughly the same length which is entirely contained in K. Moreover, α' can be chosen to be without transverse self-intersections. In particular, the hyperbolic geodesic with the same endpoints does not enter deeply into the cusps of S, i.e. it is contained in a compact subset K' of S.

On K', the singular euclidean metric and the hyperbolic metric are uniformly quasi-isometric. Therefore a lift to $\mathbf{H}^2$ of a biinfinite vertical or horizontal geodesic is a uniform quasi-geodesic for the hyperbolic metric. Such a quasi-geodesic has well defined endpoints in the ideal boundary $\partial\mathbf{H}^2$ of $\mathbf{H}^2$ (see also [L83, PH92]).

Choose an orientation for the saddle connection $\tilde{s}$. There are two oriented vertical geodesic lines α_0, β_0 for the metric defined by $\tilde{q}$ which contain the saddle connection $\tilde{s}$ as a subarc and which are contained in a bounded neighborhood of a side α, β of $\tilde{C}$. The geodesics α_0, β_0 are determined by the requirement that their orientation coincides with the given orientation of $\tilde{s}$ and that moreover at every singular point x, the angle at x to the left of α_0 (or to the right of β_0) for the orientation of the geodesic and the orientation of $\mathbf{H}^2$ equals π (see [L83] for details of this construction).

The ideal boundary of the closed half-plane of $\mathbf{H}^2$ which is bounded by α (or β) and which is disjoint from the interior of $\tilde{C}$ is a compact subarc a (or b) of $\partial\mathbf{H}^2$ bounded by the endpoints of α (or β). The arcs a, b are disjoint (or, equivalently, the sides α, β of $\tilde{C}$ are not adjacent). A horizontal geodesic line for $\tilde{q}$ which intersects the interior of the saddle connection $\tilde{s}$ is a quasi-geodesic in $\mathbf{H}^2$ with one endpoint in the interior of the arc a and the second endpoint in the interior of the arc b. Since the horizontal length of $\tilde{s}$ is positive, this means that the support of the horizontal measured geodesic lamination of $\tilde{q}$ contains geodesics with one endpoint in the arc a and the second endpoint in b.

Since the topological types of the support of μ and of τ coincide, a carrying map $F : \mathrm{supp}(\mu) \to \tau$ is surjective and induces a bijection between the complementary components of $\mathrm{supp}(\mu)$ and the complementary components of τ. In particular, the projections to S of the geodesics α, β determine two non-adjacent sides of a complementary component C_τ of τ which is the image of the projection of $\tilde{C}$ to S.

On the other hand, by construction of the dual bigon track τ^* of τ (see [PH92]), if $\rho : (-\infty, \infty) \to \tau^*$ is any trainpath which intersects the complementary component C_τ of τ then every component of $\rho(-\infty, \infty) \cap C_\tau$ is a compact arc with endpoints on adjacent sides of C_τ. In particular, a lift to $\mathbf{H}^2$ of such a trainpath is a quasi-geodesic in $\mathbf{H}^2$ whose endpoints meet at most one of the two arcs $a, b \subset \partial\mathbf{H}^2$. Now the support of the horizontal measured geodesic lamination ν of q is carried by τ^*. Therefore every leaf of the support of ν determines a biinfinite trainpath on τ^* and hence a lift to $\mathbf{H}^2$ of such a leaf does not connect the arcs $a, b \subset \partial\mathbf{H}^2$. However, we observed above that the support of the horizontal measured geodesic lamination of $\tilde{q}$ contains geodesics connecting a to b. This is a contradiction and shows that indeed $q \in \tilde{\mathcal{Q}}(m_1, \dots, m_\ell; -m)$.

Let $\mathcal{P}(\mu) \subset \mathcal{PML}$ be the open set of all projective measured geodesic laminations whose support is distinct from the support of μ. Then the assignment ψ which associates to a projective measured geodesic lamination $[\nu] \in \mathcal{P}(\mu)$ the area one quadratic differential $q(\mu,[\nu])$ with vertical measured geodesic lamination μ and horizontal projective measured geodesic lamination $[\nu]$ is a homeomorphism of $\mathcal{P}(\mu)$ onto a strong stable manifold in $\tilde{\mathcal{Q}}(S)$.

By Corollary 2.3, the projectivization $P\mathcal{V}^*(\tau)\subset\mathcal{PML}$ of $\mathcal{V}^*(\tau)$ is homeomorphic to a closed ball in a real vector space of dimension $h-1$, and this is just the dimension of a strong stable manifold in a component of $\tilde{\mathcal{Q}}(m_1,\dots,m_\ell;-m)$. Therefore by the above discussion and invariance of domain, there is a component $\tilde{\mathcal{U}}$ of the stratum $\tilde{\mathcal{Q}}(m_1,\dots,m_\ell;-m)$ such that the restriction of the map ψ to $P\mathcal{V}^*(\tau)$ is a homeomorphism of $P\mathcal{V}^*(\tau)$ onto the closure of an open subset of a strong stable manifold $W^{ss}_{\tilde{\mathcal{U}}}(q)\subset\tilde{\mathcal{U}}$.

The above argument also shows that if $z\in\mathcal{Q}(\tau)$ is defined by $\zeta\in\mathcal{V}_0(\tau),\nu\in\mathcal{V}^*(\tau)$ and if the support of ν is contained in $\mathcal{LL}(m_1,\dots,m_\ell;-m)$ then we have $z\in\tilde{\mathcal{Q}}(m_1,\dots,m_\ell;-m)$. If $\tilde{\mathcal{P}}$ denotes the component of $\tilde{\mathcal{Q}}(m_1,\dots,m_\ell;-m)$ containing z then for every point $[\beta]$ in the projectivization $P\mathcal{V}(\tau)$ of $\mathcal{V}(\tau)$, the pair $([\beta],\nu)$ defines a quadratic differential which is contained in the strong unstable manifold $W^{su}_{\tilde{\mathcal{P}}}(z)$. The set of these quadratic differentials equals the closure of an open subset of $W^{su}_{\tilde{\mathcal{P}}}(z)$.

The set of quadratic differentials q with the property that the support of the vertical (or of the horizontal) measured geodesic lamination of q is minimal and of type $(m_1,\dots,m_\ell;-m)$ is dense and of full Lebesgue measure in $\tilde{\mathcal{Q}}(m_1,\dots,m_\ell;-m)$ [M82, V86]. Moreover, this set is saturated for the strong stable (or for the strong unstable) foliation. Thus by the above discussion, the set of all measured geodesic laminations which are carried by τ (or τ^*) and whose support is minimal of type $(m_1,\dots,m_\ell;-m)$ is dense in $\mathcal{V}(\tau)$ (or in $\mathcal{V}^*(\tau)$). As a consequence, the set of all pairs $(\mu,\nu)\in\mathcal{V}(\tau)\times\mathcal{V}^*(\tau)$ with $\iota(\mu,\nu)=1$ which correspond to a quadratic differential $q\in\tilde{\mathcal{Q}}(m_1,\dots,m_\ell;-m)$ is dense in the set of all pairs $(\mu,\nu)\in\mathcal{V}(\tau)\times\mathcal{V}^*(\tau)$ with $\iota(\mu,\nu)=1$. Thus the set $\mathcal{Q}(\tau)$ is contained in the closure of a component $\tilde{\mathcal{Q}}$ of the stratum $\tilde{\mathcal{Q}}(m_1,\dots,m_\ell;-m)$. Moreover, by reasons of dimension, $\{\Phi^t q\mid q\in\mathcal{Q}(\tau),t\in[-\delta,\delta]\}$ contains an open subset of this component. This shows the first part of the proposition.

Now if $\tau\in\mathcal{LT}(m_1,\dots,m_\ell;-m)$ is orientable and if μ is a geodesic lamination which is carried by τ, then μ inherits an orientation from an orientation of τ. The orientation of τ together with the orientation of S determines an orientation of the dual bigon track τ^* (see [PH92]). This implies that any geodesic lamination carried by τ^* admits an orientation, and if (μ,ν) jointly fill up S and if μ is carried by τ, ν is carried by τ^* then the orientations of μ,ν determine the orientation of S. As a consequence, the quadratic differential q of (μ,ν) is the square of a holomorphic one-form. The proposition follows. ☐

The next proposition is a converse to Proposition 3.1 and shows that train tracks can be used to define coordinates on strata.

Proposition 3.2. (1) *For every $q\in\tilde{\mathcal{Q}}(m_1,\dots,m_\ell;-m)$ there is a large non-orientable train track $\tau\in\mathcal{LT}(m_1,\dots,m_\ell;-m)$ and a number $t\in\mathbb{R}$ so that $\Phi^t q$ is an interior point of $\mathcal{Q}(\tau)$.*

(2) *For every* $q \in \tilde{\mathcal{H}}(k_1, \dots, k_s)$ *there is a large orientable train track* $\tau \in \mathcal{LT}(2k_1, \dots, 2k_s; 0)$ *and a number* $t \in \mathbb{R}$ *so that* $\Phi^t q$ *is an interior point of* $\mathcal{Q}(\tau)$.

Proof. Fix a complete hyperbolic metric on S of finite volume. Define the *straightening* of a train track τ to be the immersed graph in S whose vertices are the switches of τ and whose edges are the geodesic arcs which are homotopic to the branches of τ with fixed endpoints.

The hyperbolic metric induces a distance function on the projectivized tangent bundle of S. As in Section 3 of [H09a], we say that for some $\epsilon > 0$ a train track τ ϵ-*follows* a geodesic lamination μ if the tangent lines of the straightening of τ are contained in the ϵ-neighborhood of the tangent lines of μ in the projectivized tangent bundle of S and if moreover the straightening of any trainpath on τ is a piecewise geodesic whose exterior angles at the breakpoints are not bigger than ϵ. By Lemma 3.2 of [H09a], for every geodesic lamination μ and every $\epsilon > 0$ there is a transversely recurrent train track which carries μ and ϵ-follows μ.

Let $q \in \tilde{\mathcal{Q}}(m_1, \dots, m_\ell; -m)$. Assume first that the support μ of the vertical measured geodesic lamination of q is large of type $(m_1, \dots, m_\ell; -m)$. This is equivalent to stating that q does not have vertical saddle connections. For $\epsilon > 0$ let τ_ϵ be a train track which carries μ and ϵ-follows μ. If $\epsilon > 0$ is sufficiently small then a carrying map $\mu \to \tau_\epsilon$ defines a bijection of the complementary components of μ onto the complementary components of τ_ϵ. The transverse measure on τ_ϵ defined by the vertical measured geodesic lamination of q is positive.

Let $\tilde{C} \subset \mathbf{H}^2$ be a complementary component of the preimage of μ in the hyperbolic plane $\mathbf{H}^2$. Then $\tilde{C}$ is an ideal polygon whose vertices decompose the ideal boundary $\partial \mathbf{H}^2$ into finitely many arcs $a_1, \dots, a_k$ ordered counter-clockwise in consecutive order. Since q does not have vertical saddle connections, the discussion in the proof of Proposition 3.1 shows the following. Let ℓ be a leaf of the preimage in $\mathbf{H}^2$ of the support ν of the horizontal measured geodesic lamination of q. Then the two endpoints of ℓ in $\mathbf{H}^2$ either are both contained in the interior of the same arc a_i or in the interior of two adjacent arcs a_i, a_{i+1}. As a consequence, for sufficiently small ϵ the geodesic lamination ν is carried by the dual bigon track τ_ϵ^* of τ_ϵ (see the characterization of the set of measured geodesic laminations carried by τ_ϵ^* in [PH92]). Moreover, by the explicit construction of a measured geodesic lamination from a measured foliation [L83], for any two adjacent subarcs a_i, a_{i+1} of $\partial \mathbf{H}^2$ cut out by $\tilde{C}$, the transverse measure of the set of all leaves of the preimage of ν connecting these sides is positive. Therefore for sufficiently small ϵ, the horizontal measured geodesic lamination ν of q defines an interior point of $\mathcal{V}^*(\tau_\epsilon)$.

Now the set of quadratic differentials z so that the support of the horizontal measured geodesic lamination of z is large of type $(m_1, \dots, m_\ell; -m)$ is dense in the strong stable manifold $W^{ss}_{\tilde{\mathcal{Q}}}(q)$ of q. The above reasoning shows that for such a quadratic differential z and for sufficiently small ϵ, the horizontal measured geodesic lamination of z is carried by τ_ϵ^*. But this just means that $\tau_\epsilon \in \mathcal{LT}(m_1, \dots, m_\ell; -m)$.

Moreover, if $r > 0$ is the total weight which the vertical measured geodesic lamination of q puts on τ_ϵ then $\Phi^{-\log r}q$ is an interior point of $\mathcal{Q}(\tau_\epsilon)$. Thus τ_ϵ satisfies the requirement in the proposition. Note that τ_ϵ is necessarily non-orientable.

If $q \in \tilde{\mathcal{H}}(k_1, \dots, k_s)$ is such that the support of the vertical measured geodesic lamination of q is large of type $(2k_1, \dots, 2k_s; 0)$ then the above reasoning also applies and yields an oriented large train track with the required property.

Consider next the case that the support μ of the vertical measured geodesic lamination of q fills up S but is not of type $(m_1, \dots, m_\ell; -m)$. Then q has a vertical saddle connection. The set of all vertical saddle connections of q is a *forest*, i.e. a finite disjoint union T of finite trees. The number of edges of this forest is uniformly bounded. For $\epsilon > 0$ let τ_ϵ be a train track which ϵ-follows μ and carries μ. If ϵ is sufficiently small then a carrying map $\mu \to \tau_\epsilon$ defines a bijection between the complementary components of μ and the complementary components of τ_ϵ which induces a bijection between their sides as well.

Modify τ_ϵ as follows. Up to isotopy, a vertical saddle connection s of q is contained in a complementary component C_s of τ_ϵ which corresponds to the complementary component of μ determined by s (see the proof of Proposition 3.1). Since a carrying map $\mu \to \tau_\epsilon$ determines a bijection between the sides of the complementary components of μ and the sides of the complementary components of τ_ϵ, the horizontal lines crossing through s determine two non-adjacent sides c_1, c_2 of C_s (see once more the discussion in the proof of Proposition 3.1).

Choose an embedded rectangle $R_s \subset C_s$ whose boundary intersects the boundary of C_s in two opposite sides contained in the interior of the sides c_1, c_2 of C_s. We can choose these rectangles R_s where s runs through the vertical saddle connections of q to be pairwise disjoint. With a homotopy of S, collapse each of the rectangles R_s to a single segment in such a way that the two sides of R_s which are contained in τ_ϵ are identified and form a single large branch b_s as shown in Figure C. The branch b_s can be isotoped to the saddle connection s. Let η be the train track constructed in this way. Then η is of topological type $(m_1, \dots, m_\ell; -m)$.

Figure C

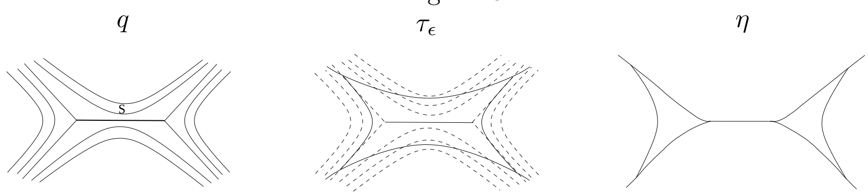

The train track τ_ϵ can be obtained from η by splitting η at each of the large branches b_s and removing the diagonal of the split. In particular, η carries τ_ϵ and hence μ. The transverse measure on η defined by the vertical measured geodesic lamination of q is positive and consequently η is recurrent. Moreover, it follows as above that for sufficiently small ϵ, the horizontal measured geodesic lamination of q is carried by η^*. In particular, if $\epsilon > 0$ is sufficiently small then η is fully transversely recurrent and in fact large. Now using once more the consideration in the first part of this proof, by possibly decreasing further the size of ϵ, we can guarantee that for some $t \in \mathbb{R}$ the quadratic differential $\Phi^t q$ is an interior point of $\mathcal{Q}(\eta)$. As a consequence, η satisfies the requirements in the proposition.

If the support μ of the vertical measured geodesic lamination of q is arbitrary then we proceed in the same way. Let $\epsilon > 0$ be sufficiently small that there is a bijection between the complementary components of the train track τ_ϵ and the complementary components of the support of μ. As before, we use the horizontal measured foliation of q to construct for every vertical saddle connection s of q an embedded rectangle R_s in S whose interior is contained in a complementary component of τ_ϵ and with two opposite sides on τ_ϵ in such a way that the rectangles R_s are pairwise disjoint. Collapse each of the rectangles to a single arc. The resulting train track has the required properties.

We discuss in detail the case that the support of μ contains a simple closed curve component α. Then τ_ϵ contains α as a simple closed curve component as well. There is a vertical flat cylinder C for q foliated by smooth circles freely homotopic to α. The boundary ∂C of C is a finite union of vertical saddle connections. Some of these saddle connections may occur twice on the boundary of C (if $\mu = \alpha$ then this holds true for each of these saddle connections). Assume without loss of generality (i.e. perform a suitable isotopy) that α is a closed vertical geodesic contained in the interior of C.

For each saddle connection s in the boundary of C choose a compact arc a_s contained in the interior of s. Choose moreover a foliation $\mathcal{F}$ of C by compact arcs with endpoints on the boundary of C which is transverse to the foliation of C by the vertical closed geodesics and such that no leaf of $\mathcal{F}$ has both endpoints in $\cup_s a_s$. We call the foliation $\mathcal{F}$ horizontal. In particular, each arc a_s which occurs twice in the boundary of the cylinder C determines an embedded rectangle R_s in S. Two opposite sides of R_s are disjoint subarcs of α; we call these sides the vertical sides. Each of the remaining two sides consists of two half-leaves of the foliation $\mathcal{F}$ which begin at a boundary point of a_s and end in a point of α. The interior of the arc a_s is contained in the interior of R_s. The rectangle R_s is foliated by horizontal arcs which are contained in leaves of $\mathcal{F}$.

The rectangles R_s are pairwise disjoint. Therefore there is a homotopy of S which collapses each of the rectangles R_s to the arc a_s by collapsing each leaf of the horizontal foliation to a single point. We may assume that the resulting graph is a train track which carries α and contains for every saddle connection s which occurs twice in the boundary of C a large branch b_s.

If $\mu = \alpha$ then $\tau_\epsilon = \alpha$ and the resulting train track η is of topological type $(m_1, \dots, m_\ell; -m)$. Moreover, the discussion in the beginning of this proof shows that it is large, and that there is some $t \in \mathbb{R}$ such that $\Phi^t q$ is an interior point of $\mathcal{Q}(\tau)$.

If $\alpha \neq \mu$ then there is a saddle connection s on the boundary of C which separates C from $S - C$. In this case the arc a_s is contained in the interior of a rectangle R_s with one side contained in α and the second side contained in the interior of a branch of a component of τ_ϵ different from α. This branch is determined by the horizontal geodesics which cross through s. As before, there is a homotopy of S which collapses the rectangle R_s to a single branch.

To summarize, the train track τ_ϵ can be modified in finitely many steps to a train track η with the required properties by collapsing for every vertical saddle connection of q a rectangle with two sides on τ_ϵ to a single large branch. This completes the construction and finishes the proof of the proposition. □

Remark: In the proof of Proposition 3.2, we constructed explicitly for every quadratic differential $q \in \mathcal{Q}(S)$ a train track τ_q belonging to the stratum of q. If q is a one-cylinder Strebel differential, i.e. if the support of the vertical measured geodesic lamination of q consists of a single simple closed curve, then the train track τ_q is uniquely determined by the glueing pattern of the vertical saddle connections on the boundary of the cylinder. This fact can be used to give an alternative proof of the classification results of Kontsevich-Zorich [KZ03] and of Lanneau [L08].

4. A symbolic system

In this section we construct a subshift of finite type which is used in the following sections to construct a symbolic coding for the Teichmüller flow on a component of a stratum in the moduli space of quadratic or abelian differentials. We continue to use the assumptions and notations from Sections 2 and 3.

Thus let $\mathcal{Q}$ be a connected component of a stratum $\mathcal{Q}(m_1,\dots,m_\ell;-m)$ of $\mathcal{Q}(S)$ (or of a stratum $\mathcal{H}(m_1/2,\dots,m_\ell/2)$ of $\mathcal{H}(S)$). Let $\tilde{\mathcal{Q}}$ be the preimage of $\mathcal{Q}$ in $\tilde{\mathcal{Q}}(S)$ (or in $\tilde{\mathcal{H}}(S)$). Let

$$\mathcal{LT}(\mathcal{Q}) \subset \mathcal{LT}(m_1,\dots,m_\ell;-m)$$

be the set of all large train tracks τ of the same topological type as $\mathcal{Q}$ such that

$$\hat{\mathcal{Q}}(\tau) = \cup_t \Phi^t \mathcal{Q}(\tau)$$

contains an open subset of $\tilde{\mathcal{Q}}$. Note that $\hat{\mathcal{Q}}(\tau)$ is just the set of all quadratic differentials q whose vertical measured geodesic lamination is carried by τ and whose horizontal measured geodesic lamination is carried by τ^*. The set $\mathcal{LT}(\mathcal{Q})$ is invariant under the action of the mapping class group.

For $\tau \in \mathcal{LT}(\mathcal{Q})$ let $\mathcal{LL}(\tau) \subset \mathcal{LL}(m_1,\dots,m_\ell;-m)$ be the set of all large geodesic laminations of topological type $(m_1,\dots,m_\ell;-m)$ which are carried by τ. It follows from [H09a] that $\mathcal{LL}(\tau)$ is compact. Define moreover

$$\mathcal{LL}(\mathcal{Q}) \subset \mathcal{LL}(m_1,\dots,m_\ell;-m)$$

to be the closure (in the restriction of the Hausdorff topology) of the set of all minimal large geodesic laminations which support the vertical measured geodesic lamination of a quadratic differential $q \in \tilde{\mathcal{Q}}$. Then $\mathcal{LL}(\mathcal{Q})$ is invariant under the action of $\mathrm{Mod}(S)$. Moreover, we have $\mathcal{LL}(\tau) \subset \mathcal{LL}(\mathcal{Q})$ for every $\tau \in \mathcal{LT}(\mathcal{Q})$.

We first look at large geodesic laminations which are models for one-cylinder Strebel differentials. Namely, define a geodesic lamination $\zeta \in \mathcal{LL}(\mathcal{Q})$ to be *special* if ζ contains a unique minimal component c which is a simple closed curve. Such a geodesic lamination has isolated leaves which spiral about c from two different sides (see [H09a]). Each such geodesic lamination ζ determines a collection of large train tracks which carry ζ and which can be described as follows.

Define a *twist connector* in a train track η to be a simple closed curve of class C^1 embedded in η which consists of a large branch and a small branch. Call a

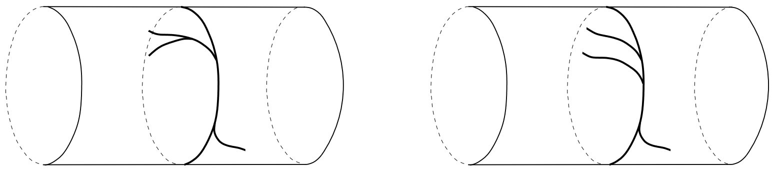

large train track $\tau \in \mathcal{LT}(m_1, \dots, m_\ell; -m)$ *special* if there is a smoothly embedded simple closed curve $c \subset \tau$ which contains every switch of τ. We also require that τ has a single large branch, that this large branch is contained in c and that τ can be modified with a sequence of shifts to a train track τ' containing c as a twist connector. For every special train track τ with embedded circle $c \subset \tau$ there is a unique special large geodesic lamination $\nu \in \mathcal{LL}(\tau)$ which contains c as its minimal component and which is carried by τ.

The next lemma implies that special geodesic laminations exist for every component of a stratum.

Lemma 4.1. (1) *Every $\tau \in \mathcal{LT}(\mathcal{Q})$ carries a special geodesic lamination $\nu \in \mathcal{LL}(\mathcal{Q})$.*
(2) *$\mathcal{LT}(\mathcal{Q})$ contains special train tracks.*

Proof. Let $\tau \in \mathcal{LT}(\mathcal{Q})$ and let c be a simple closed curve carried by τ. Then c defines a transverse measure on τ. This measure is a limit (for the topology as a space of weight functions on the branches of τ) of a sequence of transverse measures whose supports are contained in $\mathcal{LT}(\tau)$ (see the proof of Proposition 3.1 for details). Since $\mathcal{LL}(\tau)$ is compact, by passing to a subsequence we may assume that the geodesic laminations ν_i converge as $i \to \infty$ in the Hausdorff topology to a large geodesic lamination $\xi \in \mathcal{LL}(\tau)$ which contains c as a minimal component.

For some hyperbolic metric on S and for sufficiently small $\epsilon > 0$, a train track η which carries ξ and ϵ-follows ξ is carried by τ (Lemma 3.2 of [H09a]). Moreover, it contains c as an embedded curve of class C^1. The topological type of the train track η coincides with the topological type of ξ and hence with the topological type of τ. The train track η also carries every geodesic lamination β which is sufficiently close to ξ in the Hausdorff topology. Since ξ is a limit in the Hausdorff topology of the minimal large geodesic laminations ν_i, the train track η is large. Up to modifying η by a sequence of shifts we may assume that η contains c as a twist connector. In particular, there is a right or left Dehn twist φ_c about c such that $\varphi_c\eta \prec \eta$.

Choose a minimal large geodesic lamination $\zeta \in \mathcal{LL}(\eta)$. Then for each $k > 0$, $\varphi_c^k\zeta \in \mathcal{LL}(\eta)$. Since $\mathcal{LL}(\eta)$ is compact with respect to the Hausdorff topology, as $k \to \infty$ we may assume that $\varphi_c^k\zeta$ converges in $\mathcal{LL}(\eta)$ to a large geodesic lamination ν. The lamination ν contains c as a minimal component. Moreover, every leaf of

ν distinct from c spirals about c. Thus ν is special. The first part of the lemma is proven.

Now if ν is as in the previous paragraph, then for sufficiently small ϵ a train track β which carries ν and ϵ-follows ν is special. Moreover, it follows as in the second paragraph of this proof that β is large. This shows the second part of the lemma. □

Our goal is to use train tracks for a symbolic coding of the Teichmüller flow on $\mathcal{Q}$. However, the mapping class group $\mathrm{Mod}(S)$ does not act freely on large train tracks. To overcome this difficulty we extend the definition of a large train track as follows.

Definition 4.2. A *numbered large train track* is a large train track τ together with a numbering of the branches of τ.

The set

$$\mathcal{NT}(\mathcal{Q}) \tag{4}$$

of all isotopy classes of numbered large train tracks on S whose underlying unnumbered large train track is contained in $\mathcal{LT}(\mathcal{Q})$ is invariant under the natural action of the mapping class group. Since a mapping class which preserves a large train track τ as well as each of its branches is the identity (compare the proof of Lemma 3.3 of [H09a]), this action is free.

Define a *(numbered) combinatorial type* to be an orbit of a (numbered) large train track under the action of the mapping class group. The set $\mathcal{E}_0$ of numbered combinatorial types is the quotient of the set of all numbered large train tracks under the action of the mapping class group. If the numbered combinatorial type defined by a numbered large train track τ is contained in a subset $\mathcal{E}$ of $\mathcal{E}_0$, then we say that τ is *contained* in $\mathcal{E}$ and we write $[\tau] \in \mathcal{E}$. Let

$$\mathcal{E}_0(\mathcal{Q})$$

be the set of all numbered combinatorial types of numbered train tracks which are contained in $\mathcal{NT}(\mathcal{Q})$.

In the sequel we will call two large train tracks τ, τ' *shift equivalent* if τ' can be obtained from τ by a sequence of shifts. This defines an equivalence relation on the set of all large train tracks on S. If τ is numbered, then any train track which is shift equivalent to τ inherits a numbering from the numbering of τ. If the large train track τ' can be obtained from a large train track τ by a single split, then a numbering of the branches of τ naturally induces a numbering of the branches of τ' and therefore such a numbering defines a *numbered split*.

Definition 4.3. A *full split* of a (numbered) large train track τ is a (numbered) large train track τ' which can be obtained from τ by splitting τ at each large branch precisely once.

A *full (numbered) splitting sequence* is a sequence (τ_i) of (numbered) large train tracks such that for each i, the (numbered) large train track τ_{i+1} can be obtained from τ_i by a full (numbered) split.

We say that a numbered combinatorial type $x \in \mathcal{E}_0(\mathcal{Q})$ is *splittable* to a numbered combinatorial type x' if there is a numbered large train track τ contained in x which can be connected to a numbered large train track τ' contained in x' by a *full* numbered splitting sequence.

Lemma 4.5 below is the main technical tool used in this note. For the purpose of its proof and later use, we first establish a simple fact (which is well known to the experts) about the structure of quadratic differentials $q \in \mathcal{Q}$ which are recurrent under the Teichmüller flow.

Lemma 4.4. *Let $q \in \mathcal{Q}$ be a point whose orbit under the Teichmüller flow recurrs to a compact set $K \subset \mathcal{Q}$ for arbitrarily large times. Then the vertical measured geodesic lamination of a lift of q to $\tilde{\mathcal{Q}}$ is uniquely ergodic, with support in $\mathcal{LL}(\mathcal{Q})$.*

Proof. Let $q \in \mathcal{Q}$, let $\tilde q \in \tilde{\mathcal{Q}}$ be a lift of q and assume that the support of the vertical measured geodesic lamination of $\tilde q$ is not contained in $\mathcal{LL}(\mathcal{Q})$.Then q admits at least one vertical saddle connection. This saddle connection has finite length.

By continuity, for every compact set $K \subset \mathcal{Q}$ the minimal length of any saddle connection is bounded from below by a number $c > 0$. Now the length of a vertical saddle connection is exponentially decreasing under the Teichmüller flow. In particular, the orbit of q does not return to K for arbitrarily large times.

Moreover, Masur [M82] showed that the vertical measured geodesic lamination of a recurrent quadratic differential is uniquely ergodic. The lemma follows. ☐

Lemma 4.5. *For every connected component $\mathcal{Q}$ of a stratum $\mathcal{Q}(m_1, \dots, m_\ell; -m)$ (or of a stratum $\mathcal{H}(m_1/2, \dots, m_\ell/2)$) there is a set $\mathcal{E}(\mathcal{Q}) \subset \mathcal{E}_0(\mathcal{Q})$ of numbered combinatorial types with the following properties.*

1. *For all $x, x' \in \mathcal{E}(\mathcal{Q})$, x is splittable to x'.*
2. *If τ is contained in $\mathcal{E}(\mathcal{Q})$ and if (τ_i) is any full numbered splitting sequence issuing from $\tau_0 = \tau$ then τ_i is contained in $\mathcal{E}(\mathcal{Q})$ for all $i \geq 0$.*
3. *For every $q \in \mathcal{Q}$ without horizontal and vertical saddle connections and every lift $\tilde q$ of q to $\tilde{\mathcal{Q}}(S)$ (or to $\tilde{\mathcal{H}}(S)$) there is some $\tau \in \mathcal{LT}(\mathcal{Q})$ such that $[\tau] \in \mathcal{E}(\mathcal{Q})$ and $\tilde q \in \hat{\mathcal{Q}}(\tau)$.*

Proof. Let for the moment $\eta \in \mathcal{LT}(m_1, \dots, m_\ell; -m)$ be an unnumbered special large train track which contains a twist connector. This twist connector is an embedded simple closed curve c in η which contains the unique large branch e of η. Splitting η at e with a half-branch $h \subset e$ as a winner results in a large train track which is obtained from η by a 1/2-Dehn twist about the simple closed curve c.

Let σ be a large train track which is shift equivalent to η. Then for some $k \geq 0$, σ contains c as an embedded curve of class C^1 which is composed of $k+2$ branches, and k of these branches are mixed. The unique large branch e_0 of σ is contained in c. Modifying σ with a full splitting sequence, with half-branches contained in c as winners, results in a collection of train tracks which is finite up to the action of the subgroup of $\mathrm{Mod}(S)$ generated by a Dehn twist φ_c about c. Each of the large branches of any of these train tracks is contained in c. We call a train track with this property *weakly special*. Note that a special train track is weakly special. There is a train track η' which is obtained from η by a full splitting sequence, and there is a number $k \neq 0$ such that η' can be connected to $\varphi_c^k \eta'$ by a full splitting sequence.

Let $\mathcal{Q}$ be a connected component of a stratum. By the second part of Lemma 4.1, there is a special train track $\eta \in \mathcal{LT}(\mathcal{Q})$. Let $\mathcal{Z}$ be the set of all combinatorial types of weakly special train tracks in $\mathcal{LT}(\mathcal{Q})$. For $[\beta] \in \mathcal{Z}$ let $\mathcal{A}([\beta]) \subset \mathcal{Z}$ be the collection of all combinatorial types of weakly special train tracks ξ which are contained in $\mathcal{Z}$ and with the following additional property. There is an element $e \neq g \in \mathrm{Mod}(S)$ and a representative $\beta \in \mathcal{LT}(\mathcal{Q})$ of $[\beta]$ which can be connected to $g\xi$ by a nontrivial full splitting sequence. By the discussion in the first paragraph of this proof, the set $\mathcal{A}([\beta])$ is not empty.

Since $\mathcal{Z}$ and the sets $\mathcal{A}([\beta])$ are finite we can choose $[\sigma] \in \mathcal{Z}$ in such a way that $\mathcal{A}([\sigma])$ has the smallest cardinality. Let $[\xi] \in \mathcal{A}([\sigma])$. By invariance under the action of $\mathrm{Mod}(S)$, we have $\mathcal{A}([\xi]) \subset \mathcal{A}([\sigma])$. Since $[\sigma] \in \mathcal{Z}$ was chosen in such a way that the cardinality of $\mathcal{A}([\sigma])$ is minimal, we conclude that in fact $\mathcal{A}([\xi]) = \mathcal{A}([\sigma])$ and, in particular, $[\xi] \in \mathcal{A}([\xi])$. As a consequence, for all $\xi, \xi' \in \mathcal{LT}(\mathcal{Q})$ which are contained in $\mathcal{A}([\sigma])$ there is a nontrivial full splitting sequence connecting ξ to a train track $g\xi'$ for some $g \in \mathrm{Mod}(S)$.

Now let $\tau \in \mathcal{NT}(\mathcal{Q})$ be a numbered large train track obtained from a train track ξ which is contained in $\mathcal{A}([\sigma])$ by a numbering of the branches. By the above observation, there is a non-trivial full numbered splitting sequence issuing from τ and ending at a numbered train track β whose underlying unnumbered train track is contained in the $\mathrm{Mod}(S)$-orbit of ξ. In particular, up to the action of the mapping class group, β is obtained from τ by a permutation of the numbering. The permutations of the numbering of τ obtained in this way clearly form a semi-group and hence τ can be connected to a train track in its own orbit under the mapping class group by a full numbered splitting sequence.

Define $\mathcal{E}(\mathcal{Q}) \subset \mathcal{E}_0(\mathcal{Q})$ to be the set of all numbered combinatorial types of all large train tracks which can be obtained from τ by a full numbered splitting sequence. Let $\theta \in \mathcal{NT}(\mathcal{Q})$ be contained in $\mathcal{E}(\mathcal{Q})$. By Lemma 4.1, θ carries a special geodesic lamination ν and hence it carries a special train track η. Since η is special, it follows from the results of [H09b] that there is a weakly special train track η' and there is some $g \in \mathrm{Mod}(S)$ such that θ can be connected to $g\eta'$ by a splitting sequence. The discussion in the second paragraph of this proof then implies that we may assume that θ can be connected to a weakly special train track by a nontrivial full splitting sequence. By the choice of τ, we conclude that θ can be connected to an element in the $\mathrm{Mod}(S)$-orbit of τ by a full numbered splitting sequence. In other words, the first property in the lemma holds true for $\mathcal{E}(\mathcal{Q})$, and the second is true by definition.

To show the third property, let $q \in \mathcal{Q}$ be without horizontal and vertical saddle connection and let $\tilde{q} \in \tilde{\mathcal{Q}}$ be a lift of q. By Proposition 3.2, there is a train track

$\eta \in \mathcal{LT}(\mathcal{Q})$ and a number $t \in \mathbb{R}$ so that $\Phi^t\tilde{q} \in \mathcal{Q}(\eta)$ and hence $\tilde{q} \in \hat{\mathcal{Q}}(\eta)$. Since the projection of $\hat{\mathcal{Q}}(\eta)$ contains an open subset of $\mathcal{Q}$, there is a point $z \in \mathcal{Q}$ whose Φ^t-orbit is dense in $\mathcal{Q}$ in forward and backward direction and which admits a lift $\tilde{z} \in \mathcal{Q}(\eta)$.

The support of the horizontal measured geodesic lamination $\tilde{z}^h$ of $\tilde{z}$ is carried by τ^*. Since q does not have vertical saddle connections, there is a quadratic differential $\tilde{u} \in \hat{\mathcal{Q}}(\eta)$ whose horizontal measured geodesic lamination equals $\tilde{z}^h$ and whose vertical measured geodesic lamination equals the vertical measured geodesic lamination $\tilde{q}^v$ of $\tilde{q}$ up to scale. By Proposition 3.1 we have $\tilde{u} \in \tilde{\mathcal{Q}}$.

Since the backward orbit of z is dense in $\mathcal{Q}$, we have $d(\Phi^{-t}\tilde{u}, \Phi^{-t}\tilde{z}) \to 0$ $(t \to \infty)$ for any distance function d induced by a complete $\mathrm{Mod}(S)$-invariant Riemannian metric on $\tilde{\mathcal{Q}}$ [M80]. On the other hand, if τ is the train track used in the construction of $\mathcal{E}(\mathcal{Q})$ then $\hat{\mathcal{Q}}(\tau)$ contains an open subset of $\tilde{\mathcal{Q}}$. Using once more the fact that the backward orbit of z is dense in $\mathcal{Q}$, there is a number $t \in \mathbb{R}$ and there is some $g \in \mathrm{Mod}(S)$ so that $g\Phi^t\tilde{u} \in \mathcal{Q}(\tau)$. In particular, the vertical measured geodesic lamination $g\tilde{q}^v$ of $g\tilde{q}$ is carried by τ.

Now $\tilde{q}$ does not have any vertical saddle connection, and the set $\mathcal{E}(\mathcal{Q})$ has the second property stated in the lemma. From this we conclude as in the proof of Proposition 3.2 that there is a numbered train track ζ contained in $\mathcal{E}(\mathcal{Q})$ and there is some $t \in \mathbb{R}$ so that $\Phi^t g\tilde{q} \in \mathcal{Q}(\zeta)$. This shows the third property stated in the lemma. □

Remark: It is possible that $\mathcal{E}(\mathcal{Q}) = \mathcal{E}_0(\mathcal{Q})$.

Let $k > 0$ be the cardinality of a set $\mathcal{E}(\mathcal{Q}) \subset \mathcal{E}_0(\mathcal{Q})$ as in Lemma 4.5. Number the k elements of $\mathcal{E}(\mathcal{Q})$ in an arbitrary way. Identify each element of $\mathcal{E}(\mathcal{Q})$ with its number. Define $a_{ij} = 1$ if the numbered combinatorial type i can be split with a single full numbered split to the numbered combinatorial type j and define $a_{ij} = 0$ otherwise. The matrix $A = (a_{ij})$ defines a *subshift of finite type*. Its phase space is the set of biinfinite sequences

$$\Omega \subset \prod_{i=-\infty}^{\infty} \{1, \dots, k\}$$

with the property that $(x_i) \in \Omega$ if and only if $a_{x_i x_{i+1}} = 1$ for all i. Every biinfinite full numbered splitting sequence $(\tau_i) \subset \mathcal{NT}(\mathcal{Q})$ contained in $\mathcal{E}(\mathcal{Q})$ defines a point in Ω. Vice versa, since the action of $\mathrm{Mod}(S)$ on the set of numbered large train tracks is free, a point in Ω determines a $\mathrm{Mod}(S)$-orbit of biinfinite full numbered splitting sequences. We say that such a full numbered splitting sequence *realizes* (x_i).

The shift map $\sigma : \Omega \to \Omega, \sigma(x_i) = (x_{i+1})$ acts on Ω. For $n > 0$ write $A^n = (a_{ij}^{(n)})$; the shift σ is *topologically transitive* if for all i, j there is some $n > 0$ such that $a_{ij}^{(n)} > 0$. Define a finite sequence $(x_i)_{0 \leq i \leq n}$ of points $x_i \in \{1, \dots, k\}$ to be *admissible* if $a_{x_i x_{i+1}} = 1$ for all i. Then $a_{ij}^{(n)}$ equals the number of all admissible sequences of length n connecting i to j [Mn87]. The following observation is immediate from the definitions.

Lemma 4.6. *The shift (Ω, σ) is topologically transitive.*

Proof. Let $i, j \in \{1, \dots, k\}$ be arbitrary. By Lemma 4.5, there is a nontrivial finite full numbered splitting sequence $\{\tau_i\}_{0 \leq i \leq n} \subset \mathcal{NT}(\mathcal{Q})$ connecting a train track τ_0 of numbered combinatorial type i to a train track τ_n of numbered combinatorial type j. This splitting sequence then defines an admissible sequence $(x_i)_{0 \leq i \leq n}$ connecting i to j. □

Remark: Without loss of generality, we can in fact assume that the shift (Ω, σ) is topologically mixing. Namely, by the discussion on p.55 of [HK95], otherwise there are numbers $\ell, n > 0$ such that $\ell n = k$ and that the following holds true. The elements of $\mathcal{E}(\mathcal{Q})$ are divided into n disjoint sets $C_1, \dots, C_n$ of ℓ elements each so that for $x_i \in C_j$ we have $a_{x_i x_{i+1}} = 1$ only if $x_{i+1} \in C_{j+1}$ (indices are taken modulo n). Moreover, the restriction of σ^n to C_1 is topologically mixing, However, in this case we can assume that C_1 contains a weakly special train track τ. We can repeat the argument in the proof of Lemma 4.5 with a single full numbered split replaced by a full numbered splitting sequence of length n. This amounts to replacing (Ω, σ) by the topologically mixing subshift (C_1, σ^n) which has all properties stated above.

5. Symbolic dynamics for the Teichmüller flow

In this section we relate the subshift of finite type (Ω, σ) constructed in Section 4 to the Teichmüller flow. We continue to use the assumptions and notations from Section 2-4.

Let as before $\mathcal{Q}$ be a component of a stratum. Let $\mathcal{NT}(\mathcal{Q})$ be as in (4). For $\tau \in \mathcal{NT}(\mathcal{Q})$ let $\mathcal{V}(\tau)$ be the space of all measured geodesic laminations carried by τ and equipped with the topology as a set of weight functions on the branches of τ. Let moreover $\mathcal{PML}(\tau)$ be the space of all projective measured geodesic laminations whose support is carried by τ. Note that $\mathcal{PML}(\tau)$ is a *compact* subset of the compact space $\mathcal{PML}$ of all projective measured geodesic laminations on S.

Let $(\tau_i)_{0 \leq i} \subset \mathcal{NT}(\mathcal{Q})$ be any full numbered splitting sequence. Then we have $\emptyset \neq \mathcal{PML}(\tau_{i+1}) \subset \mathcal{PML}(\tau_i)$ and hence $\cap_i \mathcal{PML}(\tau_i)$ is a non-empty compact subset of $\mathcal{PML}$. If $\cap_i \mathcal{PML}(\tau_i)$ consists of a unique point then we call (τ_i) *uniquely ergodic.* Using the notations from Section 4, we call the sequence $(x_i) \in \Omega$ *uniquely ergodic* if some (and hence every) full numbered splitting sequence which realizes (x_i) is uniquely ergodic with support in $\mathcal{LL}(\mathcal{Q})$. This implies in particular that for every i a transverse measure on τ_i defined by a point $\zeta \in \cap \mathcal{V}(\tau_i)$ is positive on every branch of τ_i. Moreover, the sequence (τ_i) is uniquely determined by τ_0 and ζ [H09a].

For $\mu \in \mathcal{V}(\tau)$ let $\mu(\tau)$ be the total weight which μ puts on τ (i.e. μ is the sum of the weights over all branches of τ). Let $\mathcal{U} \subset \Omega$ be the set of all uniquely ergodic sequences. We define a function $\rho : \mathcal{U} \to \mathbb{R}$ as follows. For $(x_i) \in \mathcal{U}$ choose a full numbered splitting sequence $(\tau_i) \subset \mathcal{NT}(\mathcal{Q})$ which realizes (x_i). Let

$\mu \in \mathcal{V}_0(\tau_0)\cap\bigcap\mathcal{V}(\tau_i)$ be carried by each of the train tracks τ_i where as before, $\mathcal{V}_0(\tau)$ is the set of all transverse measures on τ of total weight one. Define

$$\rho(x_i) = -\log\mu(\tau_1).$$

By equivariance under the action of the mapping class group, the number $\rho(x_i)\in\mathbb{R}$ only depends on the sequence $(x_i)\in\mathcal{U}$. In other words, ρ is a function defined on $\mathcal{U}$. We have

Lemma 5.1. *The function $\rho:\mathcal{U}\to\mathbb{R}$ is continuous and only depends on the future.*

Proof. By construction, we have $\rho(x_i)=\rho(y_i)$ if $x_i=y_i$ for all $i\geq 0$, i.e. ρ only depends on the future.

To show continuity let $(x_i)\in\mathcal{U}$ and let $\epsilon>0$. By the definition of the topology on the shift space it suffices to show that there is some $j\geq 0$ such that

$$|\rho(y_i)-\rho(x_i)|\leq\epsilon$$

whenever $(y_i)\in\mathcal{U}$ is such that $x_i=y_i$ for $0\leq i\leq j$. For this let (τ_i) be a full numbered splitting sequence which realizes (x_i). Then (τ_i) determines a measured geodesic lamination $\lambda\in\mathcal{V}_0(\tau_1)$. By definition, $\rho(x_i)=\log\lambda(\tau_0)$.

Let $p>0$ be the number of branches of a train track in $\mathcal{LT}(m_1,\dots,m_\ell;-m)$. This number only depends on the sequence $(m_1,\dots,m_\ell;-m)$ (see [PH92]). The set $\mathcal{V}_0(\tau_1)$ of all transverse measures on τ_1 of weight one can be identified with a compact convex subset of $\mathbb{R}^p$. The natural projection $\pi:\mathcal{V}_0(\tau_1)\to\mathcal{PML}$ is a homeomorphism onto its image with respect to the weak*-topology on $\mathcal{PML}$ [PH92]. There is a neighborhood V of $\pi(\lambda)$ in $\mathcal{PML}$ with the following property. Every $\nu\in\mathcal{V}_0(\tau_1)$ with $\pi(\nu)\in V$ defines a transverse measure on τ_0 whose total weight is contained in the interval $(e^{\rho(x_i)-\epsilon},e^{\rho(x_i)+\epsilon})$.

Now for every $j>0$ the set $\mathcal{PML}(\tau_j)$ of all projective measured geodesic laminations which are carried by τ_j is a compact subset of $\mathcal{PML}$ containing $\pi(\lambda)$, and we have $\mathcal{PML}(\tau_j)\subset\mathcal{PML}(\tau_i)$ for $j\geq i$ and $\cap_j\mathcal{PML}(\tau_j)=\pi(\lambda)$. As a consequence, there is some $j_0>0$ such that $\mathcal{PML}(\tau_{j_0})\subset V$. By the definition of ρ, this implies that the value of ρ on the intersection with $\mathcal{U}$ of the cylinder $\{(y_i)\mid y_j=x_j$ for $0\leq j\leq j_0\}$ is contained in the interval $(\rho(x_i)-\epsilon,\rho(x_i)+\epsilon)$. This shows the lemma. ☐

The next lemma gives additional information on the function ρ. For its formulation, let $p>0$ be the number of branches of a train track $\tau\in\mathcal{LT}(\mathcal{Q})$.

Lemma 5.2. *The function ρ maps $\mathcal{U}$ to $(0,p\log 2]$.*

Proof. We have to show that $0<\rho(x_i)\leq p\log 2$ for every $(x_i)\in\mathcal{U}$. For this choose a full numbered splitting sequence $(\tau_i)\subset\mathcal{NT}(\mathcal{Q})$ which realizes (x_i). Using the above notations, let $\lambda\in\mathcal{V}_0(\tau_0)$ be the strongly uniquely ergodic measured geodesic lamination defined by (τ_i). Then λ is carried by each of the train tracks τ_i, and it defines a transverse measure μ on τ_0 of total weight one.

Let e be a large branch of τ_0 and let τ' be the large train track which is obtained from τ_0 by a single split at e and which is splittable to τ_1 (compare [H09a] for details). Let e' be the branch in τ' which is the *diagonal* of the split of τ_0 at e. This means that e' is the branch in τ' which is the image of e under the natural bijection Λ of the branches of τ_0 onto the branches of τ'. Let μ' be the transverse measure on τ' defined by the measured geodesic lamination λ. There are two branches b, d in τ incident on the two endpoints of e such that

$$\mu(e) = \mu'(e') + \mu'(\Lambda(b)) + \mu'(\Lambda(d)).$$

Moreover, we have $\mu(a) = \mu'(\Lambda(a))$ for every branch $a \neq e$ of τ_0 and hence $\mu'(\tau') \in [1/2, 1]$. Since τ_1 can be obtained from τ by at most p splits this immediately implies that ρ is nonnegative and bounded from above by $p \log 2$.

On the other hand, by the choice of λ, the measure μ' is positive on every branch of τ'. Then the same consideration as above shows that $\rho > 0$ on $\mathcal{U}$. Note however that ρ is not bounded from below by a positive constant. ☐

A (finite or infinite) sequence $(x_i) \subset \{1, \dots, k\}$ is called *admissible* if $a_{x_i x_{i+1}} = 1$ for all i where (a_{ij}) is the transition matrix defining the subshift of finite type (Ω, σ). Call a biinfinite sequence $(x_j) \in \Omega$ *normal* if every finite admissible sequence occurs in (x_j) infinitely often in forward and backward direction. Call a finite admissible sequence $(y_i)_{0 \leq i \leq \ell}$ *tight* if for one (and hence every) full numbered splitting sequence $(\tau_i)_{0 \leq i \leq \ell}$ realizing (y_i) the natural carrying map $\tau_\ell \to \tau_0$ maps every branch b of τ_ℓ *onto* τ_0. By the definition of $\mathcal{E}(\mathcal{Q})$ and by Lemma 4.5, tight finite admissible sequences exist.

The following observation is Lemma 5.2 of [H09b].

Lemma 5.3. *Let $(\tau_i)_{0 \leq i \leq n} \subset \mathcal{LT}(\mathcal{Q})$ be any finite numbered splitting sequence such that for some $\ell < n$ both sequences $(\tau_i)_{0 \leq i \leq \ell}$ and $(\tau_i)_{\ell \leq i \leq n}$ are tight. Then there is a compact subset K of $\mathcal{Q}$ depending on the sequence with the following property. Let (μ, ν) be a pair of measured geodesic laminations with $\iota(\mu, \nu) = 1$ and such that μ is carried by τ_n and that ν is carried by τ_0^*. Assume that μ is normalized in such a way that $\mu \in \mathcal{V}_0(\tau_0)$. Then there is a quadratic differential $q \in \tilde{\mathcal{Q}}(S)$ (or $q \in \tilde{\mathcal{H}}(S)$) with vertical and horizontal measured geodesic lamination μ, ν, respectively, and q projects into K.*

We use this fact to show

Lemma 5.4. *Let $(x_i) \in \Omega$ be normal and let $(\tau_i) \subset \mathcal{NT}(\mathcal{Q})$ be a full numbered splitting sequence which realizes (x_i). Then $\cap_i \mathcal{PML}(\tau_i)$ consists of a single uniquely ergodic projective measured geodesic lamination whose support is contained in $\mathcal{LL}(\mathcal{Q})$.*

Proof. Let $(y_i)_{0 \leq i \leq n}$ be an admissible sequence such that for some $\ell < n$, both sequences $(y_i)_{0 \leq i \leq \ell}$ and $(y_i)_{\ell \leq i \leq n}$ are tight. Let $(\tau_i)_{0 \leq i \leq n}$ be a full numbered splitting sequence which realizes $(y_i)_{0 \leq i \leq n}$. Let p be the number of branches of a large train track $\tau \in \mathcal{LT}(\mathcal{Q})$ on S. For $\lambda \in \mathcal{V}_0(\tau_n)$ the maximal weight of a branch of τ_n for the transverse measure defined by λ is not smaller than $1/p$. Then by linearity of the carrying map and by the definition of a tight sequence, the *minimal* weight put by λ on any branch of τ_ℓ is not smaller than $1/p$. Since the λ-weight of a *large* branch of τ_ℓ equals the sum of the weights of two distinct branches of τ_ℓ and since τ_ℓ has at least one large branch, the total weight $\lambda(\tau_\ell)$ of λ satisfies

$\lambda(\tau_\ell) \geq \frac{p+1}{p}$. In particular, if $(x_i) \in \mathcal{U}$ is such that $x_i = y_i$ for $0 \leq i \leq n$ then we have $\sum_{j=0}^{n-1} \rho(\sigma^j(x_i)) \geq \log \frac{p+1}{p}$.

Now let $(x_i) \in \Omega$ be normal and let (τ_i) be a biinfinite full splitting sequence which realizes (x_i). Let $\mu \in \mathcal{V}_0(\tau_0) \cap \bigcap_i \mathcal{V}(\tau_i)$. Choose a quadratic differential $q \in \mathcal{Q}(\tau_0)$ with vertical measured geodesic lamination μ. Since $(x_i) \in \Omega$ is normal, the finite admissible sequence $(y_i)_{0\leq i\leq n}$ occurs in $(x_i)_{0\leq i}$ infinitely often. By Lemma 5.3 and the discussion in the first paragraph of this proof, the projection to $\mathcal{Q}$ of the orbit $\Phi^t q$ of the Teichmüller flow intersects a fixed compact subset K of $\mathcal{Q}$ for arbitrarily large times. By Lemma 4.4, μ is uniquely ergodic with support contained in $\mathcal{LL}(\mathcal{Q})$. This completes the proof of the lemma. □

By Lemma 5.4, the set of normal points in Ω is contained in the set $\mathcal{U}$ of uniquely ergodic points. Since normal points are dense in Ω, the same is true for uniquely ergodic points.

As in Section 2, for $\tau \in \mathcal{NT}(\mathcal{Q})$ let $\mathcal{V}^*(\tau)$ be the set of all measured geodesic laminations carried by τ^* and denote by $\mathcal{PV}^*(\tau)$ the projectivization of $\mathcal{V}^*(\tau)$. If $\tau' \in \mathcal{NT}(\mathcal{Q})$ is obtained from $\tau \in \mathcal{NT}(\mathcal{Q})$ by a single split at a large branch e and if C is the matrix which describes the transformation $\mathcal{V}(\tau') \to \mathcal{V}(\tau)$ then the dual transformation $\mathcal{V}^*(\tau) \to \mathcal{V}^*(\tau')$ is given by the transposed matrix C^t [PH92].

As in the proof of Lemma 5.4 we observe

Lemma 5.5. *Let $(x_i) \in \Omega$ be normal and let $(\tau_i) \subset \mathcal{NT}(\mathcal{Q})$ be a full numbered splitting sequence which realizes (x_i); then $\cap_{i<0}\mathcal{PV}^*(\tau_i)$ consists of a single uniquely ergodic projective measured geodesic lamination whose support is contained in $\mathcal{LL}(\mathcal{Q})$.*

We call the sequence $(x_i) \in \Omega$ *doubly uniquely ergodic* if (x_i) is uniquely ergodic as defined above and if moreover for one (and hence every) full numbered splitting sequence $(\tau_i) \in \mathcal{NT}(\mathcal{Q})$ which realizes (x_i) the intersection $\cap_{i<0}\mathcal{PV}^*(\tau_i)$ consists of a unique projective tangential measure with support in $\mathcal{LL}(\mathcal{Q})$. By Lemma 5.4 and Lemma 5.5, every normal sequence is doubly uniquely ergodic and hence the Borel set $\mathcal{DU} \subset \Omega$ of all doubly uniquely ergodic sequences $(x_i) \in \Omega$ is dense.

Now let $(x_i) \in \mathcal{DU}$ and let (τ_i) be a full numbered splitting sequence which realizes (x_i). By the above discussion, (τ_i) determines a pair (μ, ν) of measured geodesic laminations by the requirement that $\mu \in \mathcal{V}_0(\tau_0) \cap \bigcap_{i\geq 0} \mathcal{V}(\tau_i)$, that $\nu \in \cap_{i\leq 0}\mathcal{V}^*(\tau_i)$ and that $\iota(\mu, \nu) = 1$. By equivariance under the action of the mapping class group, this implies that every sequence $(x_i) \in \mathcal{DU}$ determines a quadratic differential

$$\Xi(x_i) \in \mathcal{Q} \subset \mathcal{Q}(m_1, \dots, m_\ell; -m). \tag{5}$$

This quadratic differential is contained in the subset $\mathcal{UQ} \subset \mathcal{Q}$ of all area one quadratic differentials in $\mathcal{Q}$ whose vertical and horizontal measured geodesic laminations are strongly uniquely ergodic, with support in $\mathcal{LL}(\mathcal{Q})$. Note that $\mathcal{UQ}$ is a Φ^t-invariant Borel subset of $\mathcal{Q}$. Its preimage $\mathcal{U}\tilde{\mathcal{Q}}$ in $\tilde{\mathcal{Q}}(S)$ (or in $\tilde{\mathcal{H}}(S)$) is a Φ^t-invariant Borel subset of $\tilde{\mathcal{Q}}$.

Our next goal is to show that the map Ξ is finite-to-one.

Lemma 5.6. *For $q \in \mathcal{U}\tilde{\mathcal{Q}}$ there is a neighborhood V of q in $\tilde{\mathcal{Q}}$ and there are finitely many train tracks $\tau_1, \dots, \tau_n \in \mathcal{LT}(\mathcal{Q})$ (where n depends on q) with the following property. If $\eta \in \mathcal{LT}(\mathcal{Q})$ is such that $\Phi^t z \in \mathcal{Q}(\eta)$ for some $z \in V$ and some $t \in [0, p \log 2]$ then $\eta \in \{\tau_1, \dots, \tau_n\}$.*

Proof. A *marking* [MM99] consists of a pants decomposition P for S and a system of simple closed *spanning curves*. For each curve $\gamma \in P$ there is a unique spanning curve which is contained in $S - (P - \gamma)$ and which intersects γ in the minimal number of points. The spanning curves may intersect.

The mapping class group $\mathrm{Mod}(S)$ naturally acts on the set of all markings of S. By equivariance and the fact that $\mathrm{Mod}(S)$ acts on $\mathcal{LT}(\mathcal{Q})$ with finitely many orbits, there is a number $k > 0$ and for every train track $\tau \in \mathcal{LT}(\mathcal{Q})$ there is a marking F of S which consists of simple closed curves carried by τ and such that the total weight of the counting measures on τ defined by these curves does not exceed k (compare the discussion in [MM99]). We call such a marking *short for* τ. The intersection number $\iota(c, c')$ between any two closed curves c, c' which are carried by some $\tau \in \mathcal{LT}(\mathcal{Q})$ and which define counting measures on τ of total weight at most k is bounded from above by a universal constant (Corollary 2.3 of [H06]). Thus if F is a short marking for τ, with pants decomposition P, then the intersection number between any two spanning curves for P is bounded from above by a universal constant $k' > 0$. In the sequel, we require for every marking of S that this intersection bound holds true.

Let $\mathcal{T}(S)$ be the Teichmüller space of all complete finite volume hyperbolic metrics on S. By equivariance under the action of the mapping class group, there is a number $\chi_0 > 0$ and for every marking F of S (with a uniform control of intersections numbers between the spanning curves) there is a complete finite volume hyperbolic metric $h \in \mathcal{T}(S)$ such that the h-length of each marking curve is at most χ_0. We call such a hyperbolic metric *short for* F. By standard hyperbolic trigonometry, there is a number $\epsilon > 0$ such that every hyperbolic metric which is short for some marking F of S is contained in the set $\mathcal{T}(S)_\epsilon$ of all hyperbolic metrics whose *systole*, i.e. the shortest length of a closed geodesic, is at least ϵ. Moreover, the diameter in $\mathcal{T}(S)$ with respect to the *Teichmüller metric* d of the set of all hyperbolic metrics which are short for a fixed marking F is bounded from above by a universal constant.

Define a map $\Lambda : \mathcal{LT}(\mathcal{Q}) \to \mathcal{T}(S)$ by associating to a large train track τ a hyperbolic metric $\Lambda(\tau) \in \mathcal{T}(S)$ which is short for a short marking for τ. By the above discussion, there is a number $\chi_1 > 0$ only depending on the topological type of S such that if Λ' is another choice of such a map then we have $d(\Lambda(\tau), \Lambda'(\tau)) \le \chi_1$ for every $\tau \in \mathcal{LT}(\mathcal{Q})$. In particular, the map Λ is *coarsely* $\mathrm{Mod}(S)$*-equivariant*: For every $\tau \in \mathcal{LT}(\mathcal{Q})$ and every $g \in \mathrm{Mod}(S)$ we have $d(\Lambda(g\tau), g\Lambda(\tau)) \le \chi_1$. By properness of the action of $\mathrm{Mod}(S)$ on $\mathcal{T}(S)_\epsilon$, this implies that for every $x \in \mathcal{T}(S)$ and every $R > 0$ there are only finitely many large train tracks $\eta \in \mathcal{LT}(\mathcal{Q})$ with $d(\Lambda(\eta), x) \le R$.

Let
$$P : \tilde{\mathcal{Q}}(S) \to \mathcal{T}(S) \text{ (or } P : \tilde{\mathcal{H}}(S) \to \mathcal{T}(S))$$

be the canonical projection. It now suffices to show that for $q \in \mathcal{U}\tilde{\mathcal{Q}}$ there is a neighborhood V of q in $\tilde{\mathcal{Q}}$ and a number $R > 0$ with the following property. If $z \in V$, if $t \in [0, p \log 2]$ and $\eta \in \mathcal{LT}(\mathcal{Q})$ are such that $\Phi^t z \in \mathcal{Q}(\eta)$ then $d(\Lambda(\eta), Pq) \leq R$.

To show that this indeed holds true we follow the reasoning in Section 4 and Section 5 of [H09b]. Namely, up to increasing the above constant $\chi_0 > 0$, we may assume that for every quadratic differential $z \in \tilde{\mathcal{Q}}$ there is a simple closed curve on S whose *q-length*, i.e. the length with respect to the singular euclidean metric defined by q, is at most χ_0.

The *curve graph* $\mathcal{C}(S)$ of S is the metric graph whose vertices are the essential simple closed curves on S and where two such vertices are connected by an edge of length one if and only if they can be realized disjointly (see [MM99]). Define a map $\Upsilon_{\tilde{\mathcal{Q}}} : \tilde{\mathcal{Q}} \to \mathcal{C}(S)$ by associating to a quadratic differential q a simple closed curve $\Upsilon_{\tilde{\mathcal{Q}}}(q)$ of q-length at most χ_0. Then there is a number $L > 0$ such that the image under $\Upsilon_{\tilde{\mathcal{Q}}}$ of every flow line of the Teichmüller flow is an *unparametrized L-quasi-geodesic*: For every $z \in \tilde{\mathcal{Q}}$ there is an increasing homeomorphism $\varphi : \mathbb{R} \to \mathbb{R}$ such that the curve $t \to \Upsilon_{\tilde{\mathcal{Q}}}(\Phi^{\varphi(t)} z)$ is an L-quasi-geodesic in $\mathcal{C}(S)$ (see [MM99] and also [H09b, H10b]).

A *vertex cycle* for a complete train track τ is a simple closed curve carried by τ whose counting measure defines an extreme point for the space of all transverse measures on τ. The distance in $\mathcal{C}(S)$ between a vertex cycle of τ and a curve of $\Lambda(\tau)$-length at most χ_0 is bounded from above by a universal constant $b > 0$ (see the discussion in [H09b, H10b]). Moreover, for every $q \in \mathcal{Q}(\tau)$, the distance in $\mathcal{C}(S)$ between a vertex cycle for τ and the curve $\Upsilon_{\tilde{\mathcal{Q}}}(q)$ is at most b [H06].

By Lemma 4.2 of [H09b], there is a number $\ell > 0$ and for every $\epsilon > 0$ there is a number $m(\epsilon) > 0$ with the following property. Let $\sigma, \tau \in \mathcal{LT}(\mathcal{Q})$ and assume that σ is carried by τ and that the distance in $\mathcal{C}(S)$ between any vertex cycle of σ and any vertex cycle of τ is at least ℓ. Let $q \in \mathcal{Q}(\tau)$ be a quadratic differential whose vertical measured geodesic lamination q^v is carried by σ. If the total weight of the transverse measure on σ defined by q^v is not smaller than ϵ, then $d(\Lambda(\tau), Pq) \leq m(\epsilon)$.

Now let $q \in \mathcal{U}\tilde{\mathcal{Q}}$. Then $d(\Upsilon_{\tilde{\mathcal{Q}}}(\Phi^t q), \Upsilon_{\tilde{\mathcal{Q}}}(q)) \to \infty$ $(t \to \infty)$ (compare [MM99, H06, H10c]) and therefore there is a neighborhood V of q in $\tilde{\mathcal{Q}}$ and there is a number $T > 0$ such that

$$d(\Upsilon_{\tilde{\mathcal{Q}}}(\Phi^t z), \Upsilon_{\tilde{\mathcal{Q}}}(\Phi^s z)) \geq \ell + 2b + 2p \log 2$$

for all $t \geq T$, all $s \in [0, p \log 2]$ and for all $z \in V$.

Let $z \in V$, let $s \in [0, p \log 2]$ and let $\eta \in \mathcal{LT}(\mathcal{Q})$ be such that $\Phi^s z \in \mathcal{Q}(\eta)$. If $\sigma \in \mathcal{LT}(\mathcal{Q})$ is such that $\Phi^t q \in \mathcal{Q}(\sigma)$ for some $t \in [T, T + p \log 2]$ and if σ can be obtained from τ by a full splitting sequence then the distance in $\mathcal{C}(S)$ between a vertex cycle of τ and a vertex cycle of σ is at least ℓ. Thus σ, τ satisfy the hypothesis in Lemma 4.2 of [H09b] as stated above with a number $\epsilon \geq e^{-T - p \log 2}$. This implies that $d(Pz, \Lambda(\eta)) \leq m(\epsilon)$ which shows the lemma. □

Recall that the function ρ on the dense shift invariant Borel set $\mathcal{DU} \subset \Omega$ is continuous, positive and bounded from above by $p\log 2$. The *suspension* for the shift σ on the invariant subspace $\mathcal{DU}$ with roof function $\rho : \mathcal{DU} \to (0, p\log 2]$ is the space

$$X = \{(x_i) \times [0, \rho(x_i)] \mid (x_i) \in \mathcal{DU}\}/\sim$$

where the equivalence relation $\sim$ identifies the point $((x_i), \rho(x_i))$ with the point $(\sigma(x_i), 0)$. Note that $\sim$ is a closed equivalence relation on $\mathcal{DU}$ since the function ρ is continuous. There is a natural flow Θ^t on X defined by $\Theta^t(x,s) = (\sigma^j x, \tilde{s})$ (for $s > 0$) where $j \geq 0$ is such that $0 \leq \tilde{s} = t + s - \sum_{i=0}^{j-1} \rho(\sigma^i x) < \rho(\sigma^j x)$.

A *semi-conjugacy* of (X, Θ^t) into a flow space (Y, Φ^s) is a continuous map $\Upsilon : X \to Y$ such that $\Phi^t\Upsilon(x) = \Upsilon(\Theta^t x)$ for all $x \in X$ and all $t \in \mathbb{R}$. We call a semi-conjugacy Υ *finite-to-one* if the number of preimages of any point is finite. By construction, there is a natural extension of the map Ξ defined in equation (5) to the suspension flow (X, Θ^t), again denoted by Ξ, which is a semi-conjugacy of (X, Θ^t) into $(\mathcal{Q}, \Phi^t)$. The next lemma gives additional information on Ξ.

Corollary 5.7. *The semi-conjugacy $\Xi : (X, \Theta^t) \to (\mathcal{Q}, \Phi^t)$ is finite-to-one. Its image equals the Φ^t-invariant subset $\mathcal{UQ} \subset \mathcal{Q}$.*

Proof. For every $\tau \in \mathcal{NT}(\mathcal{Q})$ and every $q \in \mathcal{Q}(\tau)$ whose vertical measured geodesic lamination has support $\nu \in \mathcal{LL}(\mathcal{Q})$ there is a *unique* full numbered splitting sequence $(\tau_i)_{i>0}$ issuing from $\tau_0 = \tau$ which consists of train tracks carrying ν [H09a]. Thus Lemma 5.6 implies that for every $(x_i) \in \mathcal{DU}$ the cardinality of $\Xi^{-1}(\Xi(x_i))$ is finite. The map Ξ is clearly a semi-conjugacy. Continuity follows as in the proof of Lemma 5.1. Thus we are left with showing that the image of Ξ equals the set $\mathcal{UQ}$.

For this let $q \in \mathcal{UQ}$ and let $\tilde{q}$ be a lift of q to $\tilde{\mathcal{Q}}$. By Lemma 4.5 there is some $[\tau] \in \mathcal{E}(\mathcal{Q})$ such that $\tilde{q} \in \hat{\mathcal{Q}}(\tau)$. If μ, ν are the supports of the vertical and horizontal measured geodesic laminations of $\tilde{q}$, respectively, then there is a biinfinite full numbered splitting sequence $(\tau_i) \subset \mathcal{NT}(\mathcal{Q})$ issuing from τ such that the intersection $\cap_{i>0}\mathcal{PML}(\tau_i)$ consists of a unique point which is just the class of μ and that the intersection $\cap_{i<0}\mathcal{PML}(\tau_i^*)$ consists of a unique point which is the class of ν. Since the suspension of the orbit of a point in $\mathcal{DU}$ under the shift is mapped to a biinfinite flow line of the Teichmüller flow, this shows that the map Ξ maps $\mathcal{DU}$ onto $\mathcal{UQ}$. ☐

Remark: The Teichmüller flow is not hyperbolic, and the map Ξ is not bounded-to-one.

Since the roof function ρ on $\mathcal{DU} \subset \Omega$ is continuous, uniformly bounded and positive, every σ-invariant Borel probability measure ν on Ω which gives full mass to $\mathcal{DU}$ induces an invariant measure $\tilde{\nu}$ for the suspension flow (X, Θ^t) of total mass $\int \rho d\nu < \infty$. The measure $\tilde{\nu}$ is defined by $d\tilde{\nu} = d\nu \times dt$ where dt is the Lebesgue measure on the flow lines of the suspension flow. The image of $\tilde{\nu}$ under the semi-conjugacy Ξ is a finite Φ^t-invariant Borel measure on $\mathcal{Q}$ which we may normalize to have total mass one. Thus if $\mathcal{M}_\sigma(\mathcal{DU})$ denotes the space of all σ-invariant Borel probability measures on Ω which give full measure to $\mathcal{DU}$ then Ξ induces a map

$$\Xi_* : \mathcal{M}_\sigma(\mathcal{DU}) \to \mathcal{M}_{\rm inv}(\mathcal{Q})$$

where $\mathcal{M}_{\rm inv}(\mathcal{Q})$ is the space of of Φ^t-invariant Borel probability measures on $\mathcal{Q}$. We equip both spaces with the weak*-topology. We have

Lemma 5.8. *The map Ξ_* is continuous.*

Proof. Since Ω is a compact metrizable space, the space of all Borel probability measures on Ω equipped with the weak*-topology is compact. Thus we only have to show that whenever $\mu_i \to \mu$ in $\mathcal{M}_\sigma(\mathcal{DU})$ then $\Xi_*(\mu_i) \to \Xi_*(\mu)$.

Now by Lemma 5.2, the function ρ is continuous on $\mathcal{DU}$, bounded and positive and hence if $\mu_i \to \mu$ in $\mathcal{M}_\sigma(\mathcal{DU})$ then $\int \rho d\mu_i \to \int \rho d\mu > 0$. In particular, we have $\tilde\mu_i(X) \to \tilde\mu(X)$ where $\tilde\mu_i, \tilde\mu$ are the finite Borel measures on the suspension space (X, Θ^t) defined by the measures μ_i, μ. Therefore $\Xi_*(\mu_i) \to \Xi_*(\mu)$ if and only if for every continuous function f on $\mathcal{Q}$ with compact support we have $\int f \circ \Xi d\tilde\mu_i \to \int f \circ \Xi d\tilde\mu$. However, since Ξ is continuous this is immediate. ☐

The next result completes the proof of Theorem 1 from the introduction.

Lemma 5.9. *The map $\Xi_* : \mathcal{M}_\sigma(\mathcal{DU}) \to \mathcal{M}_{\rm inv}(\mathcal{Q})$ is surjective.*

Proof. It suffices to show that every ergodic Φ^t-invariant Borel probability measure on $\mathcal{Q}$ is contained in the image of Ξ_*.

Thus let ν be an ergodic Φ^t-invariant Borel probability measure on $\mathcal{Q}$. By the Birkhoff ergodic theorem, there is a density point $q \in \mathcal{Q}$ for ν such that the Borel probability measures

$$\nu_T = \frac{1}{T}\int_0^T \delta_{\Phi^t q} dt$$

converge weakly to ν as $T \to \infty$ where δ_x denotes the Dirac mass at x. By Lemma 4.4 and the Poincaré recurrence theorem, we have $q \in \mathcal{UQ}$. Hence up to possibly replacing q by $\Phi^t q$ for some $t \in \mathbb{R}$ there is some $(x_i) \in \mathcal{DU}$ with $\Xi(x_i) = q$.

By Lemma 5.6, there is a neighborhood V of q in $\mathcal{Q}$ such that the preimage of V under the map Ξ is a *finite* union of Borel sets $W_i \subset X$ $(i = 1, \dots, k)$. Since $\nu_T|V \to \nu|V$ weakly as $T \to \infty$ and since the map Ξ is equivariant with respect to the suspension flow and the Teichmüller flow we conclude that the restriction to $\cup_{i=1}^k W_i$ of the Borel probability measures

$$\tilde\nu_T = \frac{1}{T}\int_0^T \delta_{\Theta^t(x_i)} dt$$

converge weakly to a measure on $\cup_i W_i$ which projects to the measure ν on V. Since q was an arbitrary density point for ν we conclude that the measures $\tilde\nu_T$ converge weakly to a Θ^t-invariant Borel probability measure on X whose image under the map Ξ_* is just ν. Thus Ξ_* is surjective. ☐

6. The measure of maximal entropy

In this section we use the subshift of finite type constructed in Sections 2-5 to show that for every component $\mathcal{Q}$ of a stratum, the Φ^t-invariant probability measure λ in the Lebesgue measure class is the unique measure of maximal entropy. For strata of abelian differentials, this was earlier shown by Bufetov and Gurevich [BG07].

The strategy is as follows. Let $q \in \mathcal{Q}$ be any birecurrent point which is contained in its own α-and ω-limit set for the flow Φ^t. By the Poincaré recurrence theorem, for every Φ^t-invariant Borel probability measure μ the set of such points is of full μ-mass. We construct a topological Markov shift on a countable set $\mathcal{S}$ of symbols, given by a transition matrix $A = (a_{ij})_{\mathcal{S}\times\mathcal{S}}$. The phase space of this shift is the space

$$\Sigma = \{(y_i) \in \mathcal{S}^{\mathbb{Z}} \mid a_{y_i y_{i+1}} = 1 \text{ for all } i\}.$$

We find a positive roof function $\varphi : \Sigma \to (0,\infty)$ of bounded variation and only depending on the future such that the suspension of the shift $T : \Sigma \to \Sigma$ with roof function φ admits a bounded-to-one semi-conjugacy into $(\mathcal{Q}, \Phi^t)$. Its image $\mathcal{D}$ is Φ^t-invariant and contains all points $z \in \mathcal{Q}$ which contain the fixed quadratic differential q in their α-and ω-limit set. Since almost every orbit for the Lebesgue measure is dense in $\mathcal{Q}$, the set $\mathcal{D}$ is of full Lebesgue measure.

We use this coding and a result of Sarig [S99] to show that the supremum of the entropies of all Φ^t-invariant Borel probability measures which give full mass to $\mathcal{D}$ is the supremum of the entropies of all such measures supported in a compact invariant subset of $\mathcal{Q}$. We then apply a result of Buzzi and Sarig [BS03] to establish that the supremum is achieved by at most one measure. However, it follows from [H10a] that the entropy h of the Lebesgue measure on $\mathcal{Q}$ is the supremum of the entropies of all invariant Borel probability measures supported in compact invariant sets. This then implies Theorem 2 from the introduction.

We need the following simple observation.

Lemma 6.1. *Let $(\tau_i)_{0\le i\le k} \subset \mathcal{LT}(\mathcal{Q})$ be any full splitting sequence, let $a_0 > 0$ and let μ,ν be transverse measures on τ_k with $\mu(b) \le a_0\nu(b)$ for every branch b of τ_k. Then the transverse measures μ_0,ν_0 on τ_0 defined by μ,ν satisfy $\mu_0(e) \le a_0\nu_0(e)$ for every branch e of τ_0.*

Proof. The lemma follows immediately from the fact that the natural map $\mathcal{V}(\tau_k) \to \mathcal{V}(\tau_0)$ is the restriction of a linear map from the finite dimensional vector space of weight functions on the branches of τ_k to the vector space of weight functions on the branches of τ_0 which preserves positivity. ☐

We say that a simple closed curve c on S *fills* a large train track τ if c is carried by τ and if the transverse measure on τ defined by c is positive on every branch. A *vertex cycle* of a large train track τ is a measured geodesic lamination which is carried by τ and which spans an extreme ray in the convex cone of all measured geodesic laminations carried by τ. Masur and Minsky [MM99] observed that a vertex cycle is always spanned by a simple closed curve.

To simplify the notations, in the sequel we identify a measured geodesic lamination μ carried by a train track τ with the transverse measure it defines on τ. Using the assumptions and notations from Sections 2-5 we observe

Lemma 6.2. *Let $\tau_0 \in \mathcal{LT}(\mathcal{Q})$ and let $q \in \mathcal{Q}(\tau_0)$ be such that the vertical measured geodesic lamination ζ of q is uniquely ergodic, with support in $\mathcal{LL}(\mathcal{Q})$. Let $(\tau_i) \subset \mathcal{LT}(\mathcal{Q})$ be a full splitting sequence with $\cap_i \mathcal{V}(\tau_i) = (0,\infty)\zeta$. Then there is some $k > 0$ with the following properties.*

(1) *Every vertex cycle of τ_k fills τ_0.*
(2) *Let $\mu, \nu \in \mathcal{V}_0(\tau_0)$ be normalized transverse measures on τ_0 defined by measured geodesic laminations which are carried by τ_k. Then*
$$1/\sqrt{2} \le \mu(b)/\nu(b) \le \sqrt{2}$$
for every branch b of τ_0.
(3) *There is a number $\delta > 0$ with the following property. Let $\mu, \nu \in \mathcal{V}_0(\tau_k)$ be positive normalized measures and let $a_0 = \min\{\mu(b)/\nu(b), \nu(b)/\mu(b) \mid b\}$ where b ranges through the branches of τ_k. Then the normalized measures $\mu_0, \nu_0 \in \mathcal{V}_0(\tau_0)$ induced by μ, ν via a carrying map $\tau_k \to \tau_0$ satisfy*
$$\min\{\mu_0(b)/\nu_0(b), \nu_0(b)/\mu_0(b) \mid b\} \ge a_0 + \delta(1 - a_0).$$

Proof. Let $\tau_0 \in \mathcal{LT}(\mathcal{Q})$ and let $\zeta \in \mathcal{V}_0(\tau_0)$ be a uniquely ergodic measured geodesic lamination whose support is contained in $\mathcal{LL}(\mathcal{Q})$. Then the ζ-weight of every branch of τ_0 is positive. Let (τ_i) be the full splitting sequence such that for every i, the train track τ_i carries ζ. For each i let $A_i = \mathcal{V}_0(\tau_0) \cap \mathcal{V}(\tau_i)$ be the set of all normalized transverse measures on τ_0 defined by measured geodesic laminations which are carried by τ_i. Then $A_{i+1} \subset A_i$ and moreover $\cap_i A_i = \{\zeta\}$. Since the ζ-weight of every branch of τ_0 is positive, there is a number $\kappa > 0$ so that for sufficiently large i, say for all $i \ge i_0$, for every $\nu \in A_i$ and for every branch b of τ_0 we have $\nu(b) \ge \kappa$. In particular, every vertex cycle of τ_i fills τ_0. This shows the first part of the lemma. Moreover, by possibly increasing i_0, we can assume that for $k \ge i_0$ the conclusion in the second part of the lemma holds true.

To show the third part of the lemma, let $i_0 > 0$ and $\kappa > 0$ be as above. Let $k \ge i_0$. By Property (2), if $\mu, \nu \in \mathcal{V}(\tau_k)$ and if $\mu_0, \nu_0 \in \mathcal{V}_0(\tau_0)$ are the normalized transverse measures on τ_0 defined by μ, ν then

$$\mu_0 = \frac{1}{2}\nu_0 + \alpha \tag{6}$$

where α is a transverse measure on τ_0.

For a transverse measure $\mu \in \mathcal{V}(\tau_k)$ let $\hat{\mu} \in \mathcal{V}(\tau_0)$ be the image of μ under the carrying map $\mathcal{V}(\tau_k) \to \mathcal{V}(\tau_0)$. Let $\mu, \nu \in \mathcal{V}_0(\tau_k)$ be any two distinct *positive* normalized transverse measures on τ_k. Let $a_0 = \min\{\mu(b)/\nu(b), \nu(b)/\mu(b) \mid b\} \in (0,1)$; then $\mu - a_0\nu$ is a transverse measure on τ_k of weight $1 - a_0$. Assume that $\hat{\mu}(\tau_0) \le \hat{\nu}(\tau_0)$. Let $2\delta > 0$ be a lower bound for $\hat{\mu}(\tau_0)/\hat{\nu}(\tau_0)$ which does not depend on $\mu, \nu \in \mathcal{V}_0(\tau_k)$; such a lower bound exists since $\mathcal{V}_0(\tau_k)$ is compact. By the choice of δ and naturality under scaling, we have
$$(\hat{\mu} - a_0\hat{\nu})(\tau_0) \ge 2\delta(1 - a_0)\hat{\nu}(\tau_0).$$

By the estimate (6), applied to the normalized measures $\hat\mu - a_0\hat\nu/(\hat\mu(\tau_0) - a_0\hat\nu(\tau_0))$ and $\hat\nu/\hat\nu(\tau_0)$ on τ_0 we have

$$\hat\mu - a_0\hat\nu = \frac{1}{2}(\hat\mu(\tau_0) - a_0\hat\nu(\tau_0))\hat\nu/\hat\nu(\tau_0) + \alpha$$

for a transverse measure α on τ_0 and hence

$$\hat\mu = (a_0 + \frac{1}{2}(\hat\mu(\tau_0)/\hat\nu(\tau_0) - a_0))\hat\nu + \alpha. \tag{7}$$

However, $\hat\mu(\tau_0) = a_0\hat\nu(\tau_0) + (\hat\mu - a_0\hat\nu)(\tau_0) \geq (a_0 + 2\delta(1-a_0))\hat\nu(\tau_0)$ and consequently $\hat\mu(\tau_0)/\hat\nu(\tau_0) \geq a_0 + 2\delta(1-a_0)$. Together with inequality (7), this shows that

$$\hat\mu \geq (a_0 + \delta(1-a_0))\hat\nu.$$

But $\hat\mu(\tau_0) \leq \hat\nu(\tau_0)$ and therefore the normalized transverse measures μ_0, ν_0 on τ_0 satisfy $\mu_0 \geq (a_0 + \delta(1-a_0))\nu_0$. In particular, we have $\mu_0(b)/\nu_0(b) \geq a_0 + \delta(1-a_0)$ for every branch b of τ_0 as claimed. If $\mu(\tau_0) \geq \nu(\tau_0)$ then inequality (7) immediately yields the required inequality for $\delta = 1/2$. This shows the third part of the lemma. ☐

We call a finite full splitting sequence $(\tau_i)_{0\leq i\leq k}$ *weakly tight* if the train tracks $\tau_k \prec \tau_0$ satisfy the conclusion of Lemma 6.2.

Let $q \in \mathcal{Q}$ be any recurrent point whose α- and ω-limit set under the flow Φ^t contains itself. Let $\tilde q \in \tilde{\mathcal{Q}}$ be a lift of q. By Lemma 4.4, the vertical and horizontal measured geodesic lamination of $\tilde q$ is strongly uniquely ergodic, with support $\zeta \in \mathcal{LL}(\mathcal{Q})$. Let $p > 0$ be as in Lemma 5.2. By Lemma 4.5 and Lemma 5.6 there are finitely many large numbered train tracks $\tau_1, \dots, \tau_n \in \mathcal{NT}(\mathcal{Q})$ which are contained in $\mathcal{E}(\mathcal{Q})$ and such that $\Phi^t\tilde q \in \mathcal{Q}(\eta)$ for some $t \in [0, p\log 2]$ if and only if $\eta \in \{\tau_1, \dots, \tau_n\}$.

By Lemma 6.2 there is a number $k > 0$ such that the following holds. Let $i \leq n$ and let $(\sigma^i_j)_{0\leq j\leq k}$ be a full numbered splitting sequence of length k issuing from $\sigma^i_0 = \tau_i$ with the property that σ^i_k carries the support ζ of the vertical measured geodesic lamination of $\tilde q$. Then the sequence $(\sigma^i_j)_{0\leq j\leq k}$ is weakly tight.

Recall the definition of an admissible sequence. Define $\mathcal{S}$ to be the set of all finite admissible sequences $(x_i)_{0\leq i\leq s}$ with the following additional properties.

(1) $s \geq 2k$ and the sequences $(x_j)_{0\leq j\leq k}$ and $(x_j)_{s-k\leq j\leq s}$ are realized by one of the full splitting sequences $(\sigma^i_j)_{0\leq j\leq k}$ $(i \leq n)$.
(2) There is no number $t \in [k, s-k)$ such that the sequence $(x_j)_{t\leq j\leq t+k}$ is realized by one of the full splitting sequences $(\sigma^i_j)_{0\leq j\leq k}$.

Note that $\mathcal{S}$ is a countable set.

Define a transition matrix $A = (a_{ij})_{\mathcal{S}\times\mathcal{S}}$ by requiring that $a_{ij} = 1$ if and only if the sequence $(x_p)_{0\leq p\leq s}$ representing the symbol i and the sequence $(y_t)_{0\leq t\leq u}$ representing the symbol j satisfy $y_t = x_{s-k+t}$ for every $t \in \{0, \dots, k\}$. By construction,

$$\exists i_1, \dots, i_N \in \mathcal{S} \text{ such that } \forall \ell \in \mathcal{S}\, \exists j, k \text{ such that } a_{i_j\ell}a_{\ell i_k} = 1. \tag{8}$$

In other words, the transition matrix has the *big images and preimages (BIP) property* as defined in [S03].

Let Σ be the set of all biinfinite sequences $(y_i) \subset \mathcal{S}^{\mathbb{Z}}$ with $a_{y_i y_{i+1}} = 1$ for all i, equipped with the (biinfinite) shift $T : \Sigma \to \Sigma$. There is a natural continuous injective map

$$G : \Sigma \to \Omega$$

whose image contains the set of all normal sequences. Here Ω is as in Lemma 4.6. We claim that the image of G is contained in the set $\mathcal{DU}$ of uniquely ergodic sequences. Namely, by the definition of a weakly tight sequence and by Lemma 6.1, if (τ_i) is a full splitting sequence which realizes (y_i) the $\cap_i \mathcal{V}(\tau_i) \cap \mathcal{V}_0(\tau_0)$ consists of a unique positive transverse measure. Moreover, by the BIP-property (8) and the discussion at the end of Section 5 we may assume that the topological Markov chain (Σ, T) is topologically mixing.

Define a roof function φ on Σ by associating to an infinite sequence $(y_i) \in \Sigma$ with $y_0 = (x_i)_{0 \leq i \leq s}$ the value

$$\varphi(y_i) = \sum_{i=0}^{s-k} \rho(\sigma^i(G(y_i))).$$

By Lemma 6.2 the function φ is bounded from below by a positive constant, is unbounded and only depends on the future.

Define the n-th variation of φ by

$$\mathrm{var}_n(\varphi) = \sup\{\varphi(y) - \varphi(z) \mid y_i = z_i \text{ for } i = 0, \dots, n-1\}.$$

We have

Lemma 6.3. *There is a number $\theta \in (0,1)$ and a number $L > 0$ such that $\mathrm{var}_n(\varphi) \leq L\theta^n$ for all $n \geq 1$. In particular,*

$$\sum_{n \geq 1} \mathrm{var}_n(\varphi) < \infty.$$

Proof. Let $n \geq 1$ and let $(y_i), (z_i) \in \Sigma$ be such that $y_i = z_i$ for $i = 0, \dots, n-1$. By definition, there is a finite full numbered splitting sequence $(\tau_i)_{0 \leq i \leq u} \subset \mathcal{NT}(\mathcal{Q})$, there are numbers $k \leq \ell_0 - k < \ell_0 \leq \dots < \ell_{n-1} = u$ and there are two uniquely ergodic measured geodesic laminations $\mu, \nu \in \mathcal{V}_0(\tau_{\ell_0 - k})$ with support in $\mathcal{LL}(\mathcal{Q})$ such that the following holds.

(1) The measured geodesic lamination μ, ν projects to the vertical measured geodesic lamination of $\Xi G(y_i), \Xi G(z_i)$ (where Ξ is as in Section 5).
(2) μ, ν are carried by τ_u.
(3) The sequence $(\tau_i)_{0 \leq i \leq k}$ and each of the sequences $(\tau_i)_{\ell_j - k \leq i \leq \ell_j}$ $(j \leq n-1)$ is weakly tight.
(4) $\varphi(y_i) = \sum_{j=0}^{\ell_0 - k} \rho(\sigma^j G(y_i))$ and $\varphi(z_i) = \sum_{j=0}^{\ell_0 - k} \rho(\sigma^j G(z_i))$.

Let $\delta \in (0, \frac{1}{2})$ be as in Lemma 6.2. We show by induction on n that for $\kappa = 1-\delta \geq \frac{1}{2}$ we have

$$1 - \kappa^n \leq \min\{\mu(b)/\nu(b) \mid b \subset \tau_{\ell_0 - k}\} \leq \max\{\mu(b)/\nu(b) \mid b\} \leq \frac{1}{1-\kappa^n}.$$

For this note first that the claim holds true for $n = 1$. Namely, since μ, ν are carried by τ_{ℓ_0}, by the properties of the sequence $(\tau_i)_{\ell_0 - k \leq i \leq \ell_0}$ we obtain from the second part of Lemma 6.2 that the measures μ, ν satisfy $\mu(b)/\nu(b) \in [1/\sqrt{2}, \sqrt{2}]$ for every branch b of $\tau_{\ell_0 - k}$. This implies the claim for $n = 1$ (with the constant $1/2 < \kappa$).

By induction, assume that for some $n \geq 2$ the claim holds true for $\kappa = 1-\delta \geq \frac{1}{2}$ and all $m \in [1, n-1]$. Let $(y_i), (z_i) \in \Sigma$ be such that $y_i = z_i$ for $0 \leq i \leq n-1$. Using the above notations, let $\mu_1, \nu_1 \in \mathcal{V}_0(\tau_{\ell_1 - k})$ be normalized multiples of μ, ν. By the induction hypothesis, applied to the sequences $T(y_i), T(z_i)$, we have $1 - \kappa^{n-2} \leq \mu_1(b)/\nu_1(b) \leq (1-\kappa^{n-2})^{-1}$ for every branch b of $\tau_{\ell_1 - k}$. Lemma 6.2 now shows that the normalized transverse measures μ, ν on $\tau_{\ell_0 - k}$ which are the images of the measures μ_1, ν_1 under a carrying map $\tau_{\ell_1 - k} \to \tau_{\ell_0 - k}$ satisfy

$$\mu(b)/\nu(b) \geq 1 - \kappa^{n-2} + \delta(\kappa^{n-2}) \geq 1 - \kappa^{n-1}$$

for every branch b of $\tau_{\ell_0 - k}$ as claimed.

Now let $\mu, \nu \in \mathcal{V}_0(\tau_{\ell_0 - k})$ be normalized positive transverse measures on $\tau_{\ell_0 - k}$ so that

$$\min\{\mu(b)/\nu(b), \nu(b)/\mu(b) \mid b\} \geq a$$

for some $a > 0$. Then $\mu(\tau_0)/\nu(\tau_0) \in [a, a^{-1}]$ and therefore if $\mu_0, \nu_0 \in \mathcal{V}_0(\tau_0)$ are the normalizations of μ, ν then

$$|\varphi(\mu_0) - \varphi(\nu_0)| = |\log \mu(\tau_0) - \log \nu(\tau_0)| \leq -\log a.$$

As a consequence, the above estimate implies that $|\varphi(y_i) - \varphi(z_i)| \leq -\log(1-\kappa^{n-1})$ if $y_i = z_i$ for $0 \leq i \leq n-1$. Now $\frac{-\log(1-t)}{t} \to 1$ $(t \to 0)$ from which the lemma follows. ☐

By Lemma 6.3, the function φ can be defined on the entire space Σ. In particular, we can define the suspension (Y, Ψ^t) over Σ with roof function φ. We have

Lemma 6.4. *There is a bounded-to-one semiconjugacy $\Upsilon : (Y, \Psi^t) \to (\mathcal{Q}, \Phi^t)$. Its image contains the set of all points z whose α- and ω limit set contains q.*

Proof. There is an obvious semi-conjugacy $(Y, \Psi^t) \to (X, \Theta^t)$ which composes with the semi-conjugacy Ξ defined in Section 5 to a semi-conjugacy $\Upsilon : (Y, \Psi^t) \to (\mathcal{Q}, \Phi^t)$. Thus we only have to show that Υ is bounded-to-one. However this follows from Lemma 5.6. ☐

As before, let h be the entropy of the Lebesgue measure on $\mathcal{Q}$. Let $\mathcal{M}_T(\Sigma)$ be the space of all T-invariant Borel probability measures on Σ. For $\mu \in \mathcal{M}_T(\Sigma)$ let

$$\mathrm{pr}_\mu(-h\varphi) = h_\mu - h \int \varphi d\mu$$

where h_μ is the entropy of μ. By [S99], under the assumptions at hand, the *Gurevich pressure* of the function $-h\varphi$ is given by

$$\text{(9)} \qquad \text{pr}_G(-h\varphi) = \sup\{\text{pr}_\mu(-h\varphi) \mid \mu \in \mathcal{M}_T(\Sigma), \text{pr}_\mu(-h\varphi) \text{ is well-defined}\}.$$

The following observation relies on the results of Sarig [S99] and on [H10a]. A self-contained proof can also be derived from the discussion in Section 7.

Lemma 6.5. $\text{pr}_G(-h\varphi) \leq 0$.

Proof. By Theorem 2 of [S99], $\text{pr}_G(-h\varphi)$ equals the supremum of the quantity defined in (9) but restricted to invariant measures μ supported in compact subsets of Σ.

Now by Abramov's formula, if $A \subset \Sigma$ is any compact invariant set and if μ is a T-invariant Borel probability measure supported in A then the entropy of the induced invariant measure for the suspension flow (Y, Ψ^t) equals

$$h_\mu / \int \varphi d\mu.$$

As a consequence, the Gurevich pressure of $-h\varphi$ is nonpositive if the entropy of every Ψ^t-invariant Borel probability measure on Y which is supported in a compact set does not exceed h.

The semi-conjugacy $\Upsilon : (Y, \Psi^t) \to (\mathcal{Q}, \Phi^t)$ is bounded-to-one, and it maps the suspension of a compact invariant set A to a compact Φ^t-invariant subset of $\mathcal{Q}$. This implies that the entropy of a Ψ^t-invariant Borel probability measure on Y supported in a compact set is bounded from above by the supremum of the topological entropies of the restriction of the Teichmüller flow to compact invariant subsets of $\mathcal{Q}$. That this quantity equals h was shown in [H10a]. The lemma follows. □

Now we are ready to complete the proof of Theorem 2 from the introduction. Namely, let μ be any Φ^t-invariant ergodic Borel probability measure on $\mathcal{Q}$ and let $q \in \mathcal{Q}$ be a density point for μ which contains itself in its α- and ω-limit set. Use q to construct the countable two-sided Markov shift (Σ, T) with roof function φ and suspension flow (Y, Ψ^t).

If $\tilde{\mu}$ is a shift invariant Borel probability measure on Ω whose image under the map Ξ_* equals μ then $\tilde{\mu}$ gives full mass to $G(\Sigma)$. Since G is injective, the measure $\tilde{\mu}$ defines a shift invariant Borel probability measures $\tilde{\mu}$ on Σ which is mapped by the semi-conjugacy Υ to μ.

Let Σ^+ be the set of all one-sided infinite admissible sequences $(x_i) \subset \mathcal{S}^{\mathbb{N}}$ with $a_{x_i x_{i+1}} = 1$ for all i, equipped with the one-sided shift $T_+ : \Sigma^+ \to \Sigma^+$. Since the roof function φ only depends on the future, it defines a function on Σ^+ which will be denoted again by φ. Up to normalization, the measure μ lifts to a T^+-invariant Borel probability measure on Σ^+.

By Lemma 6.3 and Lemma 6.5 the function $-h\varphi$ on Σ^+ satisfies all the assumptions in Theorem 1.1 of [BS03]. In particular, since $\text{pr}_G(-h\varphi) \leq 0$, the entropy h_μ of the invariant Borel probability measure μ does not exceed h. As a consequence, the Lebesgue measure λ on $\mathcal{Q}$ is a measure of maximal entropy. Moreover, by Theorem 1.1 of [BS03], it is unique with this property. In other words, if $h_\mu = h$ then $\mu = \lambda$. Since μ was an arbitrary Φ^t-invariant ergodic Borel probability on $\mathcal{Q}$, this completes the proof of Theorem 2 from the introduction.

Any finite admissible sequence $(x_0, \dots, x_{n-1}) \subset \mathcal{S}$ defines the cylinder

$$[x_0, \dots, x_{n-1}] = \{(y_i) \in \Sigma^+ \mid y_i = x_i \text{ for } 0 \le i \le n-1\}.$$

A *Gibbs measure* for the function $-h\varphi$ on Σ^+ is a Borel probability measure ν with the following property. There is a number $c > 0$ so that for every cylinder set $[x_0, \dots, x_\ell]$ and every $z \in [x_0, \dots, x_\ell]$ we have

$$\nu[x_0, \dots, x_\ell] \in [c^{-1} e^{-h\sum_{0\le i\le n-1}\varphi(T^i z)}, c e^{-h\sum_{0\le i\le n-1}\varphi(T^i z)}].$$

As a consequence of [S03] and of the above discussion, we conclude

Corollary 6.6. *The Lebesgue measure is a Gibbs measure on Σ^+.*

This fact, however, can also easily be established explicitly.

7. Exponential mixing

The purpose of this section is to give a proof of exponential mixing for the Teichmüller flow on a component $\mathcal{Q}$ of a stratum of abelian or quadratic differentials. The proof follows the strategy in [AGY06] (namely, reduction to symbolic coding via Theorem 2.7 of [AGY06]), but it is more geometric than the approach in [AGY06, AR09].

To begin with, recall from Section 6 the definition of the Markov chain (Σ^+, T_+) with countable alphabet $\mathcal{S}$ and transition matrix (a_{ij}).

By construction, the set Σ^+ is a finite disjoint union

$$\Sigma^+ = \cup_{j=1}^n \Sigma_j^+$$

$(j = 1, \dots, n)$ where for each j, Σ_j^+ is determined by the large numbered train track $\tau_j \in \mathcal{NT}(\mathcal{Q})$. The Lebesgue measure λ on $\mathcal{Q}$ lifts to a T_+-invariant measure on Σ^+ which we denote by the same symbol.

Lemma 7.1. $\lambda([x, y_1] \cap [x, y_2]) = 0$ *for* $y_1 \ne y_2 \in \mathcal{S}$.

Proof. For $y_1 \ne y_2 \in \mathcal{S}$ the intersection $[x, y_1] \cap [x, y_2]$ maps to a set A of quadratic differentials whose vertical measured geodesic laminations have support which is not contained in $\mathcal{LL}(\mathcal{Q})$. Namely, if $\tau \in \mathcal{LT}(\mathcal{Q})$, if b is a large branch of τ and if τ_1, τ_2 is obtained from τ by a right and left split at b, respectively, then a geodesic lamination which is carried by both τ_1 and τ_2 is carried by the train track obtained from τ_i by removal of the diagonal of the split. This lamination can not be contained in $\mathcal{LL}(\mathcal{Q})$. By Lemma 4.4 and the Poincaré recurrence theorem, the set A has vanishing Lebesgue measure. ∎

For each $x \in \mathcal{S}$ there is some $j(x) \leq n$ such that the cylinder $[x]$ is entirely contained in $\Sigma^+_{j(x)}$. In particular, the cylinder $[x]$ can be identified with a subset of the set $\mathcal{V}_0(\tau_{j(x)})$. By Lemma 6.2 and the construction of $\mathcal{S}$, there is a number $\kappa > 0$ such that this set is contained in the set $\mathcal{P}(\tau_{j(x)}) \subset \mathcal{V}_0(\tau_{j(x)})$ of all measured geodesic laminations which give weight bigger than κ to every branch of $\tau_{j(x)}$.

The set $\mathcal{V}_0(\tau_{j(x)})$ is naturally homeomorphic to a closed ball of dimension $h-1$, and $\mathcal{P}(\tau_{j(x)})$ is an open submanifold of the interior of $\mathcal{V}_0(\tau_{j(x)})$. Define a Finsler metric $\| \, \|_{\sup}$ on $\mathcal{P}(\tau_{j(x)})$ as follows. Let $\nu \in \mathcal{P}(\tau_{j(x)})$. Then every nearby point $\mu \in \mathcal{V}_0(\tau_{j(x)})$ can be represented in the form $\mu = \nu + \alpha$ where α is a signed weight function on $\tau_{j(x)}$ satisfying the switch conditions with $\alpha(\tau_{j(x)}) = 0$. View α as a tangent vector at ν of the manifold $\mathcal{V}_0(\tau_{j(x)})$ and define

$$\|\alpha\|_{\sup} = \sup\{|\alpha(b)|/\nu(b) \mid b\}.$$

Let as before (a_{ij}) be the transition matrix of the Markov chain. For each $y \in \mathcal{S}$ with $a_{xy} = 1$, there is a full numbered splitting sequence connecting $\tau = \tau_{j(x)}$ to a train track η in the $\mathrm{Mod}(S)$-orbit of $\tau_{j(y)}$. A carrying map $\eta \to \tau$ defines a linear isomorphism β of the vector space of solutions of the switch conditions on η onto the vector space of solutions of the switch conditions on τ. This map restricts to an embedding $\mathcal{V}(\eta) \to \mathcal{V}(\tau)$ which can be projectivized to an embedding

$$B : \zeta \in \mathcal{V}_0(\eta) \to \beta(\zeta)/(\beta(\zeta)(\tau)) \in \mathcal{V}_0(\tau). \tag{10}$$

The inverse B^{-1} of B restricts to the homeomorphism $H_{xy} : [x, y] \to [y]$ determined by the shift (with the above interpretation). Since B can be identified with the restriction of the projectivization of the linear isomorphism β, the map B^{-1} is smooth. As a consequence, we can investigate its differential. To ease the notations we will write DH_{xy} to denote its differential (even though H_{xy} is only defined on a subset of $B\mathcal{V}_0(\eta)$).

Lemma 7.2. *There is a number $\kappa > 1$ and for all $x, y \in \mathcal{S}$ with $a_{xy} = 1$ there is some $c_{xy} > \kappa$ such that*

$$\kappa\|v\|_{\sup} \leq \|DH_{xy}v\|_{\sup} \leq c_{xy}\|v\|_{\sup} \text{ for all } v.$$

Proof. The lemma is a fairly immediate consequence of Part 3) of Lemma 6.2.

Namely, by construction of the alphabet $\mathcal{S}$, there is a finite full numbered splitting sequence $(\tau_i)_{0 \leq i \leq \ell}$ so that the subsequences $(\tau_i)_{0 \leq i \leq k}$ and $(\tau_i)_{\ell-k \leq i \leq \ell}$ satisfy the assumptions in Lemma 6.2 and such that the map H_{xy} is the inverse of the restriction of the map $B : \mathcal{V}_0(\tau_{\ell-k}) \to \mathcal{V}_0(\tau_0)$ to a subset of $\mathcal{P}(\tau_{\ell-k})$. Thus to show the left hand side of the inequality in the lemma it suffices to show that the restriction of B to $\mathcal{P}(\tau_{\ell-k})$ contracts the norm $\| \, \|_{\sup}$ by a fixed constant $c < 1$.

Thus let $\nu \in \mathcal{P}(\tau_{\ell-k})$ and let α be a signed weight function on $\tau_{\ell-k}$ satisfying the switch conditions with $\alpha(\tau_{\ell-k}) = 0$. We may assume that $\mu = \nu + \alpha > 0$, i.e. μ is a normalized transverse measure on $\tau_{\ell-k}$. Moreover, by perhaps replacing α by $-\alpha$ we may assume that

$$u = \|\alpha\|_{\sup} = \max\{\alpha(b)/\nu(b) \mid b\}.$$

Write

$$u+1=\max\{\frac{\alpha(b)+\nu(b)}{\nu(b)}\mid b\}=\max\{\frac{\mu(b)}{\nu(b)}\mid b\}=a_0^{-1};$$

then $a_0=\min\{\frac{\nu(b)}{\mu(b)}\mid b\}$. Let μ_0,ν_0 be the normalized transverse measures on τ_0 which are the images of μ,ν under a carrying map. Put

$$\tilde{a}_0=\min\{\frac{\nu_0(b)}{\mu_0(b)}\mid b\}.$$

The third part of Lemma 6.2 shows that $\tilde{a}_0\geq a_0+\delta(1-a_0)$ where $\delta>0$ is a universal constant.

Let $\tilde{u}+1=\tilde{a}_0^{-1}$; then we have

$$\frac{\tilde{u}}{u}=\frac{1-\tilde{a}_0}{1-a_0}\frac{a_0}{\tilde{a}_0}\leq(1-\delta)\frac{a_0}{a_0+\delta(1-a_0)}\leq 1-\delta$$

which shows the left hand side of the inequality stated in the lemma.

The map (10) can be identified with a projective linear map defined on a compact domain in a real projective space and hence the norm of its derivative is bounded away from zero by a constant depending on x,y. This shows the right hand side of the inequality in the lemma. ☐

The linear coordinates on $\mathcal{V}_0(\tau)$ also determine a volume element, unique up to scale. Let ν_τ be the normalization of this volume element of total volume one. It does not depend on any choices made. By the discussion before Lemma 7.2, the Jacobian $J[x,y]$ of an inverse branch of T_+ for these volume elements is just the Jacobian of the projectivized carrying map $(\mathcal{V}(\eta),\nu_\eta)\to(\mathcal{V}(\tau),\nu_\tau)$. Lemma 6.2 and Lemma 7.2 and its proof imply that the differential of the logarithm of $J[x,y]$ is uniformly bounded in C^0 on each inverse branch.

Now let φ be the roof function for the Markov chain (Σ^+,T_+) introduced in Section 6. We have

Lemma 7.3. *There are numbers $r_0>0,r_1>0$ with the following property.*

1. $\varphi\geq r_0$.
2. *For every cylinder $[x,y]$ of length 2, the restriction of $\varphi\circ H_{xy}^{-1}$ is differentiable, and the norm with respect to the Finsler structure $\| \|_{\sup}$ of its derivative is pointwise bounded from above by r_0.*
3. *It is not possible to write $\varphi=\alpha+\beta\circ T_+-\beta$ where α is constant on each cylinder set $[x_0,x_1]$ and β is piecewise of class C^1.*

Proof. The first part of the lemma is immediate from the construction (see Section 6 for details).

To show the second part, let $[x,y]$ be a cylinder of length two realized by a full numbered splitting sequence $(\tau_i)_{0\leq i\leq\ell}$. Let $\ell\leq s-k$ be such that the symbol $[x]$ corresponds to the sequence $(\tau_i)_{0\leq i\leq\ell}$ and that the symbol $[y]$ corresponds to $(\tau_i)_{\ell-k\leq i\leq s}$. By the definition of the Finsler metric $\| \|_{\sup}$ and the discussion in the proof of Lemma 7.2 it suffices to show the existence of a number $\kappa>0$ not depending on x, y with the following property. If $\nu \in \mathcal{P}(\tau_{\ell-k})$, if α is a weight function on $\tau_{\ell-k}$ satisfying the switch condition, with $\alpha(\tau_{\ell-k}) = 0$, and if $\mu = \nu + \alpha$ is nonnegative then

$$|\log \nu(\tau_0) - \log \mu(\tau_0)| \leq \kappa \max\{|\alpha(b)|/\nu(b) \mid b\}.$$

However, by construction, by Lemma 6.1 and the discussion in the proof of Lemma 7.2, we have $|\log \nu(\tau_0) - \log \mu(\tau_0)| \leq \log(1 + \max\{|\alpha(b)|/\nu(b) \mid b\})$ and hence the claim follows from the fact that $\lim_{t\to 0} \frac{\log(1+t)}{t} = 1$.

The third part of the lemma follows from the fact that the Lebesgue measure λ is mixing for the Teichmüller flow [V86]. ☐

Following [AGY06] we say that the roof function φ *has exponential tails* if there is a number $\beta > 0$ such that

$$\int_{\Sigma^+} e^{\beta\varphi} d\lambda < \infty.$$

The following relies on [AGY06].

Proposition 7.4. *If φ has exponential tails then the Teichmüller flow on $\mathcal{Q}$ is exponentially mixing.*

Proof. Lemma 7.3 shows that φ is a good roof function for the Markov chain (Σ^+, T_+) in the sense of Definition 2.3 of [AGY06]. Together with Lemma 7.1 and Lemma 7.2 we infer that if φ has exponential tails then all assumptions in Theorem 2.7 of [AGY06] are satisfied, and the Lebesgue measure is exponentially mixing for Φ^t. ☐

Our goal now is to show that φ has exponential tails. For this let $(\sigma_j^i)_{0\leq j\leq k}$ $(1 \leq i \leq n)$ be the full numbered splitting sequences used to define the alphabet $\mathcal{S}$. Each $x \in \mathcal{S}$ is realized by a finite full numbered splitting sequence (τ_i^x). There is a number $c > 0$ and for each x there is a number $p(x) > 0$ such that

$$\varphi(u) \in [p(x) - c, p(x) + c]$$

for all $u \in [x]$. Moreover, up to adjusting the constant c, Corollary 6.6 shows that the Lebesgue measure of each cylinder set $[x]$ is contained in the interval $[c^{-1}e^{-hp(x)}, ce^{-hp(x)}]$. Thus to show that φ has exponential tails it suffices to show that there is a number $b > 0$ with

$$\lambda(\cup\{[x] \mid p(x) \geq \ell\}) \leq e^{-b\ell} \tag{11}$$

for every integer $\ell > 0$.

Let again $\tau_1, \dots, \tau_n$ be the numbered large train tracks as in Section 6 which were used to define the set $\mathcal{S}$. Since up to the action of the mapping class group there are only finitely many large train tracks, the construction in Section 4 yields that every numbered large train track η contained in $\mathcal{E}(\mathcal{Q})$ can be connected to a train track in the Mod(S)-orbit of $\{\tau_1, \dots, \tau_n\}$ by a full numbered splitting sequence of uniformly bounded length.

For a large train track η denote as before by ν_η the normalized Lebesgue measure on $\mathcal{V}_0(\eta)$. The following is immediate from the above discussion.

Lemma 7.5. *There is a number $\chi > 0$ and for every large numbered train track η contained in $\mathcal{E}(\mathcal{Q})$ there is a closed set $V \subset \mathcal{V}_0(\eta)$ with the following properties.*

1. $\nu_\eta(V) \geq \chi$.
2. $\zeta(b) > \chi$ *for every* $\zeta \in V$ *and every branch* b *of* η.
3. $V \subset \mathcal{V}(\beta_u)$ *where* β_u *is the endpoint of a full numbered splitting sequence* $(\beta_j)_{0\leq j\leq u}$ *issuing from* $\beta_0 = \eta$ *with the property that there is some* $i \leq n$ *and some* $g \in \mathrm{Mod}(S)$ *with* $\beta_{u-k+j} = g\sigma_j^i$ *for* $0 \leq j \leq k$.

Choose some $i \in \{1, \dots, n\}$ and write $\sigma_j = \sigma_j^i$ $(0 \leq j \leq k)$ for short. Let $B_0 \subset \mathcal{Q}(\sigma_k)$ be the set of all quadratic differentials z whose horizontal measured geodesic lamination z^h is carried by σ_0^*. By the construction of the sequence σ_j, there is a number $\beta_0 > 0$ such that for every $z \in B_0$ the measured geodesic lamination z^h can be represented by a tangential measure on σ_k which associates to each branch of σ_k a weight contained in the interval $[\beta_0, \beta_0^{-1}]$. In particular, if ζ_1, ζ_2 are any two such measured geodesic laminations then $\zeta_1(b)/\zeta_2(b) \in [\beta_0^2, \beta_0^{-2}]$ for every branch b of σ_k.

For a measured geodesic lamination $\zeta \in \mathcal{V}^*(\eta)$ denote by $\zeta(\eta)$ the total weight of a tangential measure on η defined by ζ. Note that this does not depend on any choices made. With this notation, define

$$B = \{\zeta/\zeta(\sigma_k) \mid \zeta \in B_0\} \subset \mathcal{V}^*(\sigma_k).$$

For every $\zeta \in B$, the ζ-weight of any branch of σ_k is contained in the interval $[\beta, \beta^{-1}]$ for a number $\beta > 0$ only depending on β_0. As a consequence,

$$\zeta_1(b)/\zeta_2(b) \in [\beta^{-2}, \beta^2] \text{ for } \zeta_1, \zeta_2 \in B \tag{12}$$

and for every branch b of σ_k.

Using inequality (12), the next lemma is dual to Lemma 6.1. Note however that due to non-uniqueness of tangential measures representing a fixed measured geodesic lamination, the correct statement involves a suitable choice of representatives.

Lemma 7.6. *Let η be obtained from σ_k by a full splitting sequence. Then*

$$\zeta_1(b)/\zeta_2(b) \in [\beta^2, \beta^{-2}]$$

for all $\zeta_1, \zeta_2 \in B$.

Write $\eta = \sigma_k$ and let (η_j) be any full numbered splitting sequence issuing from $\eta_0 = \eta$. Let $j \geq 1$. For $\zeta \in B$ write

$$\alpha_j(\zeta) = \zeta(\eta_j)^{-1} \text{ and } \zeta(j) = \alpha_j(\zeta)\zeta$$

and note that $\alpha_j(\zeta) \leq 1$. Lemma 7.6 implies that moreover

$$\alpha_j(\zeta)/\alpha_j(\xi) \in [\beta^2, \beta^{-2}] \tag{13}$$

for all $\zeta, \xi \in B$, independent of j. For each j fix some $\alpha_j \in \{\alpha_j(\zeta) \mid \zeta \in B\}$.

For $\tau \in \mathcal{NT}(\mathcal{Q})$ call a measured geodesic lamination $\mu \in \mathcal{V}_0(\tau)$ *balanced* if it gives weight at least χ to every branch of τ where $\chi > 0$ is as in Lemma 7.5. By the choice of χ, every $\eta \in \mathcal{NT}(\mathcal{Q})$ admits a balanced measured geodesic lamination.

By the estimate (13) and the definitions we have

Lemma 7.7. *There is a number $p>0$ with the following property. Let $\mu\in\mathcal{V}_0(\eta_j)$ be balanced and define $\theta_\mu : B\to(0,\infty)$ by $\iota(\mu,\theta_\mu(\zeta)\zeta)=1$ for every $\zeta\in B$. Then*
$$\log\theta_\mu(B)\subset[\log\alpha_j-p,\log\alpha_j+p].$$

Recall from Section 3 the definition of the strong stable foliation $W^{ss}_{\tilde{\mathcal{Q}}}$ and the strong unstable foliation $W^{su}_{\tilde{\mathcal{Q}}}$ of $\tilde{\mathcal{Q}}$. The leaves of the *stable foliation* W^s are defined by $W^s_{\tilde{\mathcal{Q}}}(q)=\cup_t\Phi^tW^{ss}_{\tilde{\mathcal{Q}}}(q)$. The stable and the strong unstable foliation are transverse. The strong stable, the stable and the strong unstable foliations are invariant under the action of the mapping class group, and they descend to singular foliations $W^{ss}_{\mathcal{Q}},W^s_{\mathcal{Q}},W^{su}_{\mathcal{Q}}$ of $\mathcal{Q}$.

A closed set $W\subset\tilde{\mathcal{Q}}$ has a *product structure* if the following holds. There is some $q\in W$, there are closed neighborhoods W^s of q in $W^s_{\tilde{\mathcal{Q}}}(q)$ and W^{su} of q in $W^{su}_{\tilde{\mathcal{Q}}}(q)$, and there is a homeomorphism $H:W^s\times W^{su}\to W$ so that $H(z,w)\in W^{su}_{\tilde{\mathcal{Q}}}(z)\cap W^s_{\tilde{\mathcal{Q}}}(w)$ for all $(z,w)\in W^s\times W^{su}$.

The Lebesgue measure λ admits a family of conditional measures $\lambda^{ss},\lambda^{su}$ on strong stable and strong unstable manifolds so that
$$d\lambda=d\lambda^{ss}\times d\lambda^{su}\times dt$$
where dt is the Lebesgue measure on the flow lines of the Teichmüller flow. The measures λ^{su} transform under the Teichmüller flow via $d\lambda^{su}\circ\Phi^t=e^{ht}d\lambda^{su}$. The measures λ^{su} are invariant under the holonomy along the strong stable foliation. If $H:W^s\times W^{su}\to W$ is a homeomorphism defining a product structure on W then H is absolutely continuous with respect to the product measure $\lambda^s\times\lambda^{su}$ on $W^s\times W^{su}$ and the Lebesgue measure λ on W (see [V86]) where λ^s is the measure on W^s defined by $d\lambda^s=d\lambda^{ss}\times dt$.

Let $\mu\in\mathcal{V}_0(\eta_j)$ be a fixed balanced measured geodesic lamination. For each $\zeta\in B$ define
$$H_j(\zeta)=\{(\nu,\theta_\mu(\zeta)\zeta)\mid\nu\in\mathcal{V}(\eta_j),\iota(\nu,\theta_\mu(\zeta)\zeta)=1\}.$$
Then $H_j(\zeta)$ admits a natural identification with a subset of the strong unstable manifold $W^{su}_{\tilde{\mathcal{Q}}}(\theta_\mu(\zeta)\zeta)$ of all quadratic differentials with horizontal measured geodesic lamination $\theta_\mu(\zeta)\zeta$. The map
$$F_\mu:\nu\in\mathcal{V}_0(\eta_j)\to(\nu/\iota(\nu,\theta_\mu(\zeta)\zeta),\theta_\mu(\zeta)\zeta)\in H_j(\zeta)$$
is a homeomorphism.

Next we estimate the Jacobian of the map $F_\mu:(\mathcal{V}_0(\eta_j),\nu_{\eta_j})\to(H_j(\zeta),\lambda^{su})$ with respect to the measure ν_{η_j} on $\mathcal{V}_0(\eta_j)$ and the measure λ^{su} on $H_j(\zeta)\subset W^{su}_{\tilde{\mathcal{Q}}}(\theta_\mu(\zeta)\zeta)$.

Lemma 7.8. *Let $\mu\in\mathcal{V}_0(\eta_j)$ be a balanced measured geodesic lamination. Then for every $\zeta\in B$, the Jacobian of the homeomorphism $F_\mu:\mathcal{V}_0(\eta_j)\to H_j(\zeta)\subset W^{su}_{\tilde{\mathcal{Q}}}(\theta_\mu(\zeta)\zeta)$ is pointwise uniformly bounded, independent of j,μ and ζ. If $\nu\in\mathcal{V}_0(\eta_j)$ is balanced then it is bounded from below at ν by a universal positive constant.*

Proof. Since up to the action of the mapping class group there are only finitely many large train tracks, there is a number $a > 0$ not depending on j and there is some $\xi \in \mathcal{V}^*(\eta_j)$ with support in $\mathcal{LL}(\mathcal{Q})$ such that $\xi(b) \in [a^{-1}, a]$ for every branch b of η_j. Let

$$E = \{(\nu/\iota(\nu,\xi), \xi) \mid \nu \in \mathcal{V}_0(\eta_j)\} \subset W^{su}_{\tilde{\mathcal{Q}}}(\xi)$$

be the set of all quadratic differentials with vertical measured geodesic lamination $\nu \in \mathcal{V}(\eta_j)$ and horizontal measured geodesic lamination ξ. By formula (1), for $\nu \in \mathcal{V}_0(\eta_j)$ the intersection number $\iota(\nu,\xi)$ is bounded from above and below by a universal positive constant and hence E is compact. Thus by invariance under the action of the mapping class group, there is a number $c > 0$ not depending on j such that the Jacobian of the homeomorphism $\nu \in \mathcal{V}_0(\eta_j) \to (\nu/\iota(\nu,\xi), \xi) \in E$ with respect to the measure ν_{η_j} on $\mathcal{V}_0(\eta_j)$ and the measure λ^{su} on E is contained in $[c, c^{-1}]$.

By construction, for every $\zeta \in B$ there is a homeomorphism $\omega_\zeta : E \to H_j(\zeta)$ defined by the requirement that $\omega_\zeta(z) \in W^s_{\tilde{\mathcal{Q}}}(z)$ for all $z \in E$. There is a number $t_\zeta(z) \in \mathbb{R}$ such that

$$\omega_\zeta(z) \in \Phi^{t_\zeta(z)} W^{ss}_{\tilde{\mathcal{Q}}}(z).$$

By invariance of the measures λ^{su} under holonomy along strong stable manifolds and by the transition properties of the measures λ^{su} under the Teichmüller flow, the Jacobian of the homeomorphism ω_ζ at $z = (\nu/\iota(\nu,\xi), \xi)$ equals $e^{ht_\zeta(z)}$.

The Teichmüller flow Φ^t expands the horizontal measured foliation of a quadratic differential and hence expands the transverse measure on the vertical measured geodesic lamination. Hence by construction,

$$t_\zeta(z) = \log \iota(\nu, \theta_\mu(\zeta)\zeta) - \log \iota(\nu, \xi).$$

Now $\log \iota(\nu,\xi)$ is uniformly bounded, and since $\nu \in \mathcal{V}_0(\eta_j)$ the intersection $\iota(\nu, \theta_\mu(\zeta)\zeta)$ is uniformly bounded as well (however it may not be bounded from below by a universal positive constant). As a consequence, the holonomy shift is bounded from above by a universal constant and hence the same holds true for the Jacobian of the homeomorphism $E \to H_j(\zeta)$. This shows the first part of the lemma. If $\nu \in \mathcal{V}_0(\eta_j)$ is balanced then $\iota(\nu, \theta_\mu(\zeta)\zeta)$ is bounded from below by a universal positive constant which implies the second part. □

By invariance under the Teichmüller flow and the estimate (13), for $r > 0$ there is a function ψ_j which is uniformly bounded from above independent of j (but not bounded from below independent of j) so that the Lebesgue measure on the set

$$W_j = \cup_{t\in[-r,r]} \Phi^t(\cup_{\zeta\in B} H_j(\zeta)) \tag{14}$$

with a product structure can be represented in the form

$$d\lambda = \psi_j(e^{-h\alpha_j} d\lambda^{ss}|B) \times d\nu_{\eta_j} \times dt.$$

In particular, the Lebesgue measure of the set W_j is bounded from above by $c_0 e^{-h\alpha_j}$ where $c_0 > 0$ is a universal constant. On the other hand, Lemma 7.8 and Lemma 7.5 also show that the Lebesgue measure of W_j is bounded from below by $c_0' e^{-h\alpha_j}$ where $c_0' > 0$ is another universal constant.

Now we are ready to show

Corollary 7.9. *The roof function φ has exponential tails.*

Proof. Let $Z_0 \subset \tilde{\mathcal{Q}}$ be the set of quadratic differentials whose horizontal measured geodesic lamination is contained in the set B and whose vertical measured geodesic lamination is contained in $\mathcal{V}(\sigma_k)$. For $r > 0$ as above define $Z = \cup_{-r\leq t\leq r}\Phi^t Z_0$. Note that Z is a set with a product structure.

For $j \geq 0$ let $U(j)$ be the set of all full numbered splitting sequences of length j issuing from σ_k. By Lemma 7.6, for each j there is a decomposition

$$Z = \cup_{u\in U(j)}\Phi^{-t_u}\tilde{W}_j^u$$

with the following properties.

(1) The Lebesgue measure of the intersection of two distinct sets from the decomposition vanishes.
(2) For each $u \in U(j)$, the set $\tilde{W}_j^u$ can be represented in the form
$$\tilde{W}_j^u = \cup_{t\in[-r,r]} \cup_{\zeta\in B} \Phi^{s_j^u(\zeta)+t} H_j(\zeta)$$
where $s_j^u : B \to \mathbb{R}$ is uniformly bounded independent of j, u.
(3) There is a number $c > 0$ so that $t_u \in [h \log \lambda(\tilde{W}_j^u) - c, h \log \lambda(\tilde{W}_j^u) + c]$.

Let $V(j) \subset U(j)$ be the set of all $u \in U(j)$ with the following additional property. For $u \in V(j)$ let $(\eta_i)_{0\leq i\leq j}$ be the full numbered splitting sequence issuing from $\eta_0 = \sigma_k$ which is labeled by u. Then there is no $\ell \leq j-k$ and $g \in \mathrm{Mod}(S)$, $s \leq n$ such that $\eta_{\ell+i} = g\sigma_i^s$ for $0 \leq i \leq k$.

Let $\lambda(j) = \lambda(\cup\{\Phi^{-t_u}\tilde{W}_j^u \mid u \in V(j)\})$. It now suffices to show the existence of numbers $a < 1, m > 0$ such that $\lambda(j+m) \leq a\lambda(j)$ for all j. However, this is immediate from Lemma 7.5 and Lemma 7.8. □

Exponential mixing for the Lebesgue measure now follows from Corollary 7.9 and Proposition 7.4.

8. Entropy gaps

The goal of this section is to establish Theorem 3 from the introduction. The constructions in this section will also be used to count period orbits for the Teichmüller flow and are more involved than what is necessary for deriving entropy estimates.

Namely, to control periodic orbits we can not use the symbolic coding constructed in Section 5 directly since there are periodic orbits for the Teichmüller flow which are not images of periodic orbits for the shift space. Each of these orbits is covered by some periodic orbit for the shift, but the degree of the covering can be arbitrarily high. To overcome this difficulty we construct a new shift space which is better adapted to the counting problem.

A periodic orbit for Φ^t is defined by the conjugacy class of a pseudo-Anosov element $g \in \mathrm{Mod}(S)$. Define the *topological type* of the conjugacy class of g to be the topological type of the stratum containing the periodic orbit determined by g. Call a train track τ an *expansion* of a pseudo-Anosov element $g \in \mathrm{Mod}(S)$ if $g\tau \prec \tau$ and if moreover τ is large of the same topological type as g. In particular, the attracting fixed point for the action of g on $\mathcal{PML}$ is carried by τ, and it defines the projective class of a positive transverse measure on τ. The following result is due to Papadopoulos and Penner (Theorem 4.1 of [PP87]).

Theorem 8.1. *Every pseudo-Anosov element admits a train track expansion.*

Note that since the stabilizer in $\mathrm{Mod}(S)$ of a train track is a finite subgroup of $\mathrm{Mod}(S)$ and the cardinality of a finite subgroup of $\mathrm{Mod}(S)$ is uniformly bounded, the number of conjugacy classes of pseudo-Anosov elements with the same expansion (i.e. conjugacy classes which can be represented by pseudo-Anosov elements g, h which admit a fixed train track $\tau \in \mathcal{LT}(\mathcal{Q})$ as an expansion track, with $g\tau = h\tau$) is uniformly bounded. Moreover, any two such elements have a common power of the same degree and hence their translation lengths coincide. As a consequence, for the purpose of estimating the growth rate of periodic orbits in $\mathcal{Q}$ it suffices to estimate the growth rate of train track expansions of pseudo-Anosov elements ordered by the translation lengths of the elements.

A *splitting and shifting move* for a train track $\tau \in \mathcal{LT}(\mathcal{Q})$ is a modification of τ by a sequence of shifts composed with a single split. A *splitting and shifting sequence* is a sequence (τ_i) where for each i, τ_{i+1} can be obtained from τ_i by a splitting and shifting move. The following result is Theorem 2.4.1 of [PH92].

Proposition 8.2. *If $\sigma \prec \tau$ are large train tracks of the same topological type then τ can be connected to σ by a splitting and shifting sequence.*

In particular, for every periodic orbit for Φ^t in $\mathcal{Q}$ defined by the conjugacy class of a pseudo-Anosov element g, there is a biinfinite g-invariant splitting and shifting sequence (τ_i) whose elements are train track expansions of g.

We next construct a new topological Markov chain whose biinfinite sequences are realized by biinfinite splitting and shifting sequences in the sense described in Section 5. To avoid some technical issues arising from the existence of non-trivial stabilizers of train tracks in the mapping class group we choose once and for all a torsion free normal subgroup $\Gamma < \mathrm{Mod}(S)$ of finite index. Then Γ acts freely on $\mathcal{LT}(\mathcal{Q})$. Let E_1 be the set of all Γ-orbits of large train tracks $\tau \in \mathcal{LT}(\mathcal{Q})$. Define a transition matrix (c_{ij}) by the requirement that $c_{ij} = 1$ if and only if there are representatives τ, η of the points i, j so that η can be obtained from τ by a single splitting and shifting move. Let (Ω_1, σ_1) be the corresponding topological Markov chain. A point $(x_i) \in \Omega_1$ can be thought of as the Γ-orbit of a biinfinite splitting and shifting sequence of train tracks in $\mathcal{LT}(\mathcal{Q})$.

As in the case of the subshift Ω, we can consider the set $\mathcal{DU}_1 \subset \Omega_1$ of all biinfinite uniquely ergodic sequences. By definition, this set consists of all biinfinite admissible sequences in the alphabet E_1 which can be represented by a biinfinite splitting and shifting sequence (τ_i) with the following property. $\mathcal{V}_0(\tau_0) \cap \bigcap_i \mathcal{V}(\tau_i)$

consists of a single measured geodesic lamination with support in $\mathcal{LL}(\mathcal{Q})$, and similarly $\cap_i \mathcal{V}^*(\tau_i)$ is spanned by a single measured geodesic lamination with support in $\mathcal{LL}(\mathcal{Q})$. The set $\mathcal{DU}_1$ is invariant under the shift σ_1, and it contains all normal sequences in Ω_1. In particular, $\mathcal{DU}_1$ is dense in Ω_1.

Since Γ acts freely on $\mathcal{LT}(\mathcal{Q})$, every point $z \in \mathcal{DU}_1$ determines a quadratic differential $\Xi_1(z) \in \hat{\mathcal{Q}} = \tilde{\mathcal{Q}}/\Gamma$ (here as before, $\tilde{\mathcal{Q}}$ is the preimage of $\mathcal{Q}$ in $\tilde{\mathcal{Q}}(S)$). The image of the thus defined map

$$\Xi_1 : \mathcal{DU}_1 \to \hat{\mathcal{Q}}$$

is contained in the set of all uniquely ergodic quadratic differentials. As in Section 5, the Teichmüller flow on $\hat{\mathcal{Q}}$ induces a continuous positive roof function ρ_1 on $\mathcal{DU}_1$. Let (X_1, Θ^t) be the suspension of the shift on $\mathcal{DU}_1$ with roof fuction ρ_1. There is a natural semi-conjugacy

$$G_1 : (X_1, \Theta^t) \to (\hat{\mathcal{Q}}, \Phi^t).$$

Its image contains all quadratic differentials in $\hat{\mathcal{Q}}$ without vertical and horizontal saddle connection and with uniquely ergodic vertical and horizontal measured geodesic lamination. However, the map G_1 is countable-to-one rather than finite-to-one.

Let $\mathcal{M}_{\sigma_1}(\Omega_1)$ be the space of all σ_1-invariant Borel probability measures on Ω_1. Denote by h_ν the entropy of a measure $\nu \in \mathcal{M}_{\sigma_1}(\Omega_1)$. Define the *pressure* $\mathrm{pr}(f)$ of a continuous bounded function f on $\mathcal{DU}_1 \subset \Omega_1$ by

$$\mathrm{pr}(f) = \sup\{h_\nu + \int f d\nu \mid \nu \in \mathcal{M}_{\sigma_1}(\Omega_1),\, \nu(\mathcal{DU}_1) = 1\}.$$

The component $\mathcal{Q}$ is the quotient of $\hat{\mathcal{Q}}$ under the finite group $\mathrm{Mod}(S)/\Gamma$. Thus by Theorem 2, the entropy of every Φ^t-invariant Borel probability measure on $\hat{\mathcal{Q}}$ is at most h.

The argument in the proof of the following lemma was shown to me by Francois Ledrappier.

Lemma 8.3. $\mathrm{pr}(-h\rho_1) \le 0$.

Proof. Let μ be a σ_1-invariant ergodic Borel probability measure on Ω_1 which gives full measure to $\mathcal{DU}_1$. Let h_μ be the entropy of μ. By Abramov's formula (see [Z05] for a comprehensive account), the entropy of the induced invariant measure for the suspension flow (X_1, Θ^t) with roof function ρ_1 is $h_\mu / \int \rho_1 d\mu$. Thus we have to show that this entropy does not exceed h. For this it suffices to show that the entropy of μ as an invariant measure on (X_1, Θ^t) does not exceed the entropy of the Φ^t-invariant measure $G_1(\mu)$ on $\hat{\mathcal{Q}}$.

Since μ is ergodic and since G_1 is countable-to-one, there is a measurable set $C \subset \mathcal{DU}_1$ of full measure such that the restriction of G_1 to C is bounded-to-one. Namely, μ induces a family of conditional measures on the fibres of the map G_1. By invariance and ergodicity, the weight of a fibre is constant almost everywhere on Ω_1. Since for a typical point for μ this weight is positive, there are only finitely many fibres of non-zero weight. As a consequence, there is an invariant set $C \subset \Omega_1$ of full measure so that $G_1|C$ is bounded-to-one.

However, a bounded-to-one measure preserving semi-conjugacy is entropy preserving. This shows the lemma for ergodic measures. The general case follows from the fact that the entropy is an affine function on the space of all Borel probability measures on Ω_1 (Theorem 8.1 of [W82]). □

The roof function ρ_1 is only defined on a Borel subset of Ω_1. To overcome this difficulty we enlarge the topological Markov chain (Ω_1, σ_1) to a skew product as follows.

Recall that a point in (Ω_1, σ_1) can be realized by a biinfinite splitting and shifting sequence $(\tau_i) \subset \mathcal{LT}(\mathcal{Q})$. By Theorem 8.5.1 of [Mo03], for each such sequence (τ_i) there are finitely many minimal geodesic laminations $\lambda_1, \dots, \lambda_\ell$ $(\ell \geq 1)$ on S with pairwise disjoint support, and there is a (possibly empty) train track σ so that

$$\cap_i \mathcal{V}(\tau_i) = [\lambda_1] \times \cdots \times [\lambda_\ell] \times \mathcal{V}(\sigma).$$

Here $[\lambda_i]$ denotes the space of all (possibly trivial) transverse measures on λ_i, and $\mathcal{V}(\sigma)$ is the space of all measured geodesic laminations carried by σ. The train track σ is carried by τ. Moreover, there is some $n \geq 0$ such that σ is a *subtrack* of τ_j for all $j \geq n$.

Define $B(\tau_i) \subset \mathcal{PML}$ to be the compact set of all projective measured geodesic laminations which are projectivizations of points in $\cap_i \mathcal{V}(\tau_i)$. The mapping class group acts on biinfinite splitting and shifting sequences (τ_i) and on $\mathcal{PML}$, and the map $(\tau_i) \to B(\tau_i)$ is equivariant with respect to this action. Thus the sets $B(\tau_i)$ determine for each point $(x_i) \in \Omega_1$ a compact convex subset $B(x_i)$ of some finite dimensional projective space. Note that if $(x_i), (y_j) \in \Omega_1$ are such that $x_0 = y_0$ and if $\mu \in B(x_i), \nu \in B(y_i)$ then for some realizations $(\tau_i), (\eta_j)$ of $(x_i), (y_j)$ with $\tau_0 = \eta_0$, both μ, ν can be represented by points in $\mathcal{V}_0(\tau_0)$.

Let d_0 be a standard distance of diameter one on the subshift of finite type (Ω_1, σ_1). We require that $d_0((x_i), (y_j)) = 1$ if $x_0 \neq y_0$. Let

$$\Omega_2 = \{((x_i), \mu) \mid (x_i) \in \Omega_1, \mu \in B(x_i)\}$$

and define a distance d on Ω_2 as follows. If $((x_i), \mu), ((y_j), \nu) \in \Omega_2$ and if $x_0 \neq y_0$ then define $d(((x_i), \mu), ((y_j), \nu)) = 2$. Otherwise μ, ν can be represented by measured geodesic laminations $\tilde\mu, \tilde\nu$ of weight one on the same train track τ_0 in the class of $x_0 = y_0$. Let $\delta > 0$ be such that $d_0((x_i), (y_i)) \leq 1 - \delta$ if $x_0 = y_0$ and define

$$d(((x_i), \tilde\mu), ((y_i), \tilde\nu)) = d_0((x_i), (y_i)) + \frac{\delta}{4} \sum_b |\tilde\mu(b) - \tilde\nu(b)|$$

where the sum on the right hand side is over all branches of τ_0.

We have

Lemma 8.4. *(Ω_2, d) is a compact metric space.*

Proof. We show first that d satisfies the triangle inequality. To see this let $((x_i^j),\mu_j)$ $(j=1,2,3)$ be three points in Ω_2. If $x_0^1\neq x_0^2$ then $d(((x_i^1),\mu_1),((x_i^2),\mu_2))=2$ and also $d(((x_i^j),\mu_j),((x_i^3),\mu_j))\geq 2$ for at least one $j\in\{1,2\}$ showing the triangle inequality in this case.

Now if $x_0^1=x_0^2$ then $d(((x_i^1),\mu_1),((x_i^2),\mu_2))\leq 2$ and hence if $x_0^3\neq x_0^1$ then the triangle inequality follows from the definition. Otherwise it is just the usual triangle inequality for a product of two metric spaces.

To show that (Ω_2,d) is compact let $((x_i^j),\mu_j)\subset\Omega_2$ be any sequence. Since Ω_1 is compact, up to passing to a subsequence we may assume that $(x_i^j)\to(y_i)$ in Ω_1 $(j\to\infty)$. In particular, up to passing to another subsequence we can choose realizations of these sequences, denoted by (τ_i^j), (η_i), so that $\tau_i^j=\eta_i$ for $0\leq i\leq j$. Then for sufficiently large j we can represent $\mu_j\in B(x_i^j)$ by a measured geodesic lamination $\tilde\mu_j\in\mathcal{V}_0(\eta_0)$. Since $\mathcal{V}_0(\eta_0)$ is compact, after passing to another subsequence we may assume that $\tilde\mu_j\to\tilde\mu\in\mathcal{V}_0(\eta_j)$. But $\tilde\mu_j$ is carried by η_ℓ for each $j\geq\ell$ and hence by construction, we have $\tilde\mu\in\mathcal{V}_0(\eta_0)\cap\bigcap_\ell\mathcal{V}(\eta_\ell)$. By the definition of the metric d, this shows the claim. □

The shift σ_1 on Ω_1 extends to a shift σ_2 on Ω_2 as follows. Let $((x_i),\mu)\in\Omega_2$ and let (τ_i) be a realization of (x_i) as a biinfinite splitting and shifting sequence of large train tracks. Then μ determines a projective measured geodesic lamination which is carried by each of the train tracks τ_i. Define $\sigma_2((x_i),\mu)=((x_{i+1}),\mu)$. The first factor projection

$$P:\Omega_2\to\Omega_1$$

is a continuous semi-conjugacy of (Ω_2,σ_2) onto (Ω_1,σ_1) which is constant on the fibres of P. This is used to observe

Lemma 8.5. *The function which associates to a σ_2-invariant Borel probability measure ν on Ω_2 its entropy is upper semi-continuous.*

Proof. The first factor projection $P:\Omega_2\to\Omega_1$ commutes with the map σ_2 on Ω_2 and the shift σ_1 on Ω_1. Let ν be a σ_2-invariant ergodic Borel probability measure on Ω_2. Then $P(\nu)$ is a σ_1-invariant ergodic Borel probability measure on Ω_1. We claim that there is an invariant Borel set $A\subset\Omega_2$ of full measure such that the restriction of P to A is injective. Namely, the fibre over each point in Ω_1 is a compact subset of the space $\mathcal{PML}$ which is invariant under the map σ_2. With the interpretation of the fibre as a subset of $\mathcal{PML}$, the action of σ_2 on the fibre is constant. In particular, if for almost every point $x\in\Omega_1$ the conditional measure of ν on the fiber over x is not supported in a single point then the measure ν is not ergodic.

Now let (ν_i) be a sequence of ergodic Borel probability measures on Ω_2 converging to a measure ν. By the discussion in the previous paragraph, for each i the entropy h_{ν_i} of ν_i equals the entropy of $P(\nu_i)$, and the entropy of ν is not smaller than the entropy of $P(\nu)$. Since the subshift (Ω_1,σ_1) is expansive, the entropy of $P(\nu)$ is not smaller than $\limsup_i h_{\nu_i}$. Using once more that the entropy is an affine function on the space of all Borel probability measures the lemma follows. □

The restriction of P to $\mathcal{DU}_2 = P^{-1}\mathcal{DU}_1$ is injective and therefore the roof function ρ_1 on Ω_1 lifts to a continuous positive bounded roof function ρ_2 on $\mathcal{DU}_2 \subset \Omega_2$. Let (X_2, Θ^t) be the suspension of $\mathcal{DU}_2$ with roof function ρ_2. The semi-conjugacy $(X_1, \Theta^t) \to (\hat{\mathcal{Q}}, \Phi^t)$ induces a semi-conjugacy

$$G_2 : (X_2, \Theta^t) \to (\hat{\mathcal{Q}}, \Phi^t).$$

Extend the roof function ρ_2 on $\mathcal{DU}_2 \subset \Omega_2$ to Ω_2 by defining

$$\rho_2((x_i), \mu) = -\log \mu(\tau_1)$$

for a realization $((\tau_i), \mu)$ of $((x_i), \mu)$, i.e. $\rho_2((x_i), \mu)$ is the logarithm of the weight of μ viewed as a transverse measure on τ_1. It follows from Lemma 8.4 and Lemma 5.1 and their proofs that ρ_2 is continuous. Moreover, ρ_2 is non-negative.

We also will need to consider components $\mathcal{P}$ of strata which are distinct from $\mathcal{Q}$. To distinguish these components in our notation we sometimes write $(\Omega_1(\mathcal{P}), \sigma_1)$ for the dynamical system constructed as above for the component $\mathcal{P}$. We also allow $\mathcal{P}$ to be a component of quadratic differentials on the four punctured sphere or the one punctured torus. Define the *entropy* of the component $\mathcal{P}$ to be the entropy of the Φ^t-invariant Lebesgue measure on $\mathcal{P}$. An annulus is called an *exceptional component*, For such a component $\mathcal{P}$ we define $(\Omega_1\mathcal{P}, \sigma_1)$ to be the dynamical system whose phase space is a single point, The entropy of this space is defined to be zero. A *multi-component* $\mathcal{P}$ is a finite disjoint union of components $\mathcal{P}_i$ $(1 \leq i \leq s)$, possibly on different surfaces. The entropy of $\mathcal{P}$ is defined to be the sum of the entropies of the components.

For now and later use we record

Definition 8.6. The *direct product* $(X_1, \beta_1) \times \cdots \times (X_s, \beta_s)$ of $s \geq 1$ subshifts of finite type (X_i, β_i) with alphabet D_i and transition matrix (c^i_{uv}) is the subshift of finite type whose alphabet D is the set of all s-tuples $(d_1, \dots, d_s)$ with $d_i \in D_i$. Its transition matrix (c_{uv}) is defined by $c_{uv} = 1$ if and only if there is some $p \leq s$ such that the tuples $u = (u_1, \dots, u_s)$ and $v = (v_1, \dots, v_s)$ satisfy $u_j = v_j$ for all $j \neq p$ and $c^p_{u_p v_p} = 1$. The number p is called the *active* component of the transition which modifies u to v.

Note that in the above definition, we allow $u_p = v_p$ if $c_{u_p v_p} = 1$, but we keep track of the active component. This means that we view two constant transitions as distinct if their active components are different (we omit to formalize this idea here to keep the notations simple).

If $\mathcal{P} = \cup_{i=1}^s \mathcal{P}_i$ is a multi-component then define

$$(\Omega_1(\mathcal{P}), \sigma_1) = (\Omega_1(\mathcal{P}_1), \sigma_1) \times \cdots \times (\Omega_1(\mathcal{P}_s), \sigma_1).$$

For each $i \leq s$ there is a projection Π_i which maps each point in $(\Omega_1(\mathcal{P}), \sigma_1)$ to a finite, one-sided infinite or biinfinite admissible sequence for the shift space defined by the i-th factor $\mathcal{P}_i$. Moreover, $(\Omega_1(\mathcal{P}), \sigma_1)$ is the exact analog of the space $(\Omega_1(\mathcal{Q}), \sigma_1)$ constructed in the beginning of this section for a large train track on a disconnected surface (with the obvious interpretation). In particular, there is a natural extension $(\Omega_2(\mathcal{P}), \sigma_2)$ of $(\Omega_1(\mathcal{P}), \sigma_1)$ as before, and there is a continuous roof function ρ_2 on $(\Omega_2(\mathcal{P}), \sigma_2)$.

Define a *bounded orbit equivalence* between two dynamical systems $(X,\alpha),(Y,\beta)$ to be a measurable map $\psi : X \to Y$ so that there is a number $k > 0$ and a function $f : X \to [1,k]$ with $\beta(\psi(x)) = \psi(\alpha^{f(x)}x)$ for all x. Call the bounded orbit equivalence ψ *compactible* with two functions $\rho : X \to \mathbb{R}$, $\zeta : Y \to \mathbb{R}$ if $\zeta(\psi x) = \sum_{i=0}^{f(x)-1} \rho(\alpha^i x)$. For later reference we note

Lemma 8.7. *Let $\psi : (X,\alpha) \to (Y,\beta)$ be a bounded orbit equivalence which is compatible with functions ρ,ζ on X,Y. Let ν be an α-invariant Borel probability measure on X. Then $\psi(\nu)$ is a β-invariant Borel probability measure on Y, and*

$$\int_X \rho d\nu = \int_Y \zeta d\psi(\nu).$$

These definitions are used in the following

Lemma 8.8. *Let $\tau \in \mathcal{LT}(\mathcal{Q})$ and let η be a birecurrent (possibly disconnected) proper subtrack of τ. Then η determines a multi-component $\mathcal{P}$ of strata. There is a closed σ_2-invariant subset $B_2(\mathcal{P})$ of Ω_2, and there is a bounded orbit equivalence $(B_2(\mathcal{P}),\sigma_2) \to (\Omega_2(\mathcal{P}),\sigma_2)$ which is compatible with the roof functions ρ_2 on $B_2(\mathcal{P}) \subset \Omega_2(\mathcal{Q})$ and on $\Omega_2(\mathcal{P})$. The entropy of $\mathcal{P}$ does not exceed $h-1$.*

Proof. Let ξ be a connected train track on S. Let $N \subset S$ be a neighborhood of ξ which is sufficiently small that the inclusion $S-N \to S-\xi$ induces a bijection of components. The *subsurface filled by* ξ is the union of N with the simply connected complementary components of N.

Let $\eta < \tau$ be a birecurrent proper subtrack which fills S. Then τ can be obtained from η by sucessively subdividing some of the complementary components of η. By the discussion in Section 2, $\eta \in \mathcal{LT}(\mathcal{P})$ for a component $\mathcal{P}$ of $\mathcal{Q}(S)$ (or of $\mathcal{H}(S)$) in the closure of $\mathcal{Q}$. By Lemma 2.2, the entropy of $\mathcal{P}$ does not exceed $h-1$.

Define a splitting and shifting move of τ to be *adapted* to η if the move is at a large branch of τ contained in η and if the modified train track either contains η as a subtrack or it contains a subtrack which can be obtained from η by a splitting and shifting move.

Define $B_1(\mathcal{P}) \subset \Omega_1$ to be the set of all Γ-orbits of biinfinite splitting and shifting sequences (τ_i) with the following property. Each train track τ_i has a subtrack $\eta_i \in \mathcal{LT}(\mathcal{P})$, and the splitting and shifting move transforming τ_i to τ_{i+1} is adapted to η_i. Moreover, the train track η_{i+1} is the image of η_i in the sense described in the previous paragraph. By construction, $B_1(\mathcal{P})$ is a closed subset of Ω_1. Associating to the sequence (τ_i) the corresponding sequence $(\eta_i) \subset \mathcal{LT}(\mathcal{P})$ and quotiening by the action of Γ yields a bounded equivalence $(B_1(\mathcal{P}),\sigma_1) \to (\Omega_1(\mathcal{P}),\sigma_1)$ (see [H09c] for details why this orbit equivalence is bounded).

Let $B_2(\mathcal{P}) \subset \Omega_2(\mathcal{Q})$ be the preimage under the first factor projection $\Omega_2(\mathcal{Q}) \to \Omega_1(\mathcal{Q})$ of the set $B_1(\mathcal{P})$. Then $B_2(\mathcal{P})$ is a closed σ_2-invariant subset of Ω_2. The above discussion shows that there is a bounded orbit equivalence $(B_2(\mathcal{P}),\sigma_2) \to (\Omega_2(\mathcal{P}),\sigma_2)$ which is compatible with the roof functions ρ_2.

Next assume that $\sigma < \tau$ is connected and that the subsurface S_0 filled by σ is a proper connected subsurface of S different from an annulus. Let $c_1, \dots, c_k$ be the boundary circles of $S - S_0$. For each $i \leq k$ there is an edge-path in σ which bounds together with c_i an embedded annulus. Since σ is birecurrent, this edge path is not smooth, i.e. it contains some cusps. Note that if $S - S_0$ contains an annulus component C then there are two homotopic edge-paths in σ, one for each boundary component of C.

Cut S open along the curves c_i and let S_0 be the surface with boundary $\cup_i c_i$ which is filled by σ. There are now three cases possible for an edge path of σ corresponding to a boundary circle c_i of S_0. The first case is that this edge path contains a single cusp. In this case collapse the boundary circle c_i to a puncture.

If the edge-path contains $n \geq 2$ cusps then glue a disc to the boundary circle c_i. If $n = 2$ then the resulting graph on the modified surface contains a bigon. Collapse this bigon to a single arc.

Proceeding in this way with all edge-paths in σ homotopic to a boundary circle of $S - S_0$ produces a birecurrent train track on a surface S'. The discussion in Section 2 shows that this train track is large and hence it defines a component $\mathcal{P}$ of a stratum on S'. The entropy of this component equals $2g' - 2 + \ell + m$ where $g' \leq g$ is the genus of S_0, where ℓ is the number of cusps and where m is the number of complementary discs of σ in S'. In particular, this entropy does not exceed $h - 1$. There is closed subset $B_2(\mathcal{P})$ of Ω_2, and there is a bounded orbit equivalence $(B_2(\mathcal{P}), \sigma_2) \to (\Omega_2(\mathcal{P}), \sigma_2)$ as before.

If η is a simple closed curve then the space $\mathcal{B}_1(\mathcal{P})$ corresponds to sequences (x_i) which are realized by splitting and shifting sequences modifying a train track τ by a sequence of Dehn twists about a fixed simple closed curve. Thus in this case the lemma also holds.

If η contains more than one component then this construction can be carried out with each of the components individually. This yields the lemma. □

Definition 8.9. The multi-component $\mathcal{P}$ of the stratum in Lemma 8.8 is called a *boundary multi-component.* It is called an *ideal boundary multi-component* if there is a simple closed curve which is disjoint from the defining subtrack η.

We say that a σ_2-invariant Borel probability measure ν on Ω_2 is *supported in the boundary multi-component* $\mathcal{P}$ if it gives full mass to $B_2(\mathcal{P})$. We use Lemma 8.8 to show

Lemma 8.10. *Let μ be an ergodic σ_2-invariant Borel probability measure on Ω_2 with $\mu(\mathcal{DU}_2) = 0$. Then μ is supported in a boundary multi-component of $\mathcal{Q}$. In particular,*

$$h_\mu - (h-1)\int \rho_2 d\mu \leq 0.$$

Proof. Let $((x_i), \nu)$ be a typical point for μ. Let (τ_i) be a biinfinite splitting and shifting sequence which realizes (x_i). The image of ν under a carrying map $\nu \to \tau_0$ is a subtrack σ_0 of τ_0. By invariance and ergodicity, its topological type does not depend on (x_i).

We claim that σ_0 is a proper subtrack of τ_0. For this assume to the contrary that $\sigma_0 = \tau_0$. Since $\mu(\mathcal{DU}_2) = 0$ by assumption, either ν is not uniquely ergodic or $\cap_i \mathcal{V}^*(\tau_i)$ does not consist of a single line spanned by a uniquely ergodic point with support in $\mathcal{LL}(\mathcal{Q})$.

Consider first the case that ν is not uniquely ergodic. By the equivalent of Theorem 8.5.1 of [Mo03], there is a (possibly trivial) decomposition $\cap_i \mathcal{V}^*(\tau_i) = [\lambda_1] \times \cdots \times [\lambda_k] \times \mathcal{V}^*(\sigma)$ where $k \geq 1$ and where for each $j \leq k$, $[\lambda_j]$ is a space of measured geodesic laminations supported in a fixed minimal geodesic lamination λ_j. The geodesic laminations λ_j are pairwise disjoint. By ergodicity and the choice of (x_i), the same decomposition holds true almost everywhere.

Let $L(x_i) \subset \mathcal{PML}$ be the set of all projectivizations of all measured geodesic laminations which are contained in $[\lambda_1] \times \cdots \times [\lambda_k]$. Since the support of ν fills up S, for each $\beta \in L(x_i)$ there is a quadratic differential $q(\nu, \beta)$ with vertical measured geodesic lamination ν and whose horizontal measured geodesic lamination is contained in the class of β. The set of these quadratic differentials is compact. Associating to $((x_i), \nu)$ and $\beta \in L(x_i)$ the quadratic differential $q(\nu, \beta)$ defines a compact extension of the suspension of (Ω_2, σ_2) with roof function ρ_2 over a subset of Ω_2 of full measure. This extension admits a natural measurable semi-conjugacy into $(\hat{\mathcal{Q}}, \Phi^t)$ whose restriction to each fibre is continuous. In particular, the image of each fibre is a compact subset of $\hat{\mathcal{Q}}$. Since μ is ergodic, the Poincaré recurrence theorem holds true for μ. As a consequence, for almost every (x_i), the orbit under Φ^t of the image of the fibre over (x_i) recurs into a fixed compact subset of $\hat{\mathcal{Q}}$ for arbitrarily large times. However, since ν is not uniquely ergodic, by the results of Masur [M82] this is impossible.

If ν is uniquely ergodic with support in $\mathcal{LL}(\mathcal{Q})$ and if $\cap_i \mathcal{V}^*(\tau_i)$ does not consist of a single line spanned by a uniquely ergodic point with support in $\mathcal{LL}(\mathcal{Q})$ then the argument in the above paragraph yields a contradiction as well. Thus indeed we have $\sigma_0 \neq \tau_0$.

Lemma 8.8 now shows that there is a boundary multi-component $\mathcal{P}$ such that σ_2 is supported in $B_2(\mathcal{P})$. Thus μ defines a σ_2-invariant Borel probability measure on $\Omega_2(\mathcal{P})$. The reasoning in the proof of Lemma 8.5 shows that the entropy of μ coincides with the entropy of its image under the bounded orbit equivalence $(B_2(\mathcal{P}), \sigma_2) \to (\Omega_2(\mathcal{P}), \sigma_2)$. Moreover, the bounded orbit equivalence is compatible with the roof functions ρ_2 and consequently we have $h_\mu - (h-1)\int \rho_2 d\mu \leq 0$.

Namely, in the case that $\mathcal{P}$ consists of a single component, this is immediate from Lemma 8.3 and Lemma 8.7. If $\mathcal{P}$ is the disjoint union of $s \geq 2$ components $\mathcal{P}_i$ then the projection of μ to each of these components is an invariant Borel probability measure on the component. The roof functions on the components sum up to the roof function of the union. But this just means that the entropy of μ as an invariant measure of the suspension flow does not exceed the sum of the entropies of its projections to the suspension flows on the components. Hence by Lemma 8.8, this entropy is not bigger than $h-1$. This is what we wanted to show. $\square$

For any proper subsurface S_0 of S, the *subsurface projection* of a simple closed curve (or, more generally, of a union of simple closed curves) into S_0 is defined (Section 2 of [MM00]). Let $p_0 > 0$ be large enough so that for this number, the conclusion in Theorem 6.12 of [MM00] and in Theorem 6.1 of [R07] are satisfied. This means the following.

For a subsurface S_0 of S and two markings μ_0, μ_1 of S define $\delta(\mu_0, \mu_1; S_0)$ to be the diameter of the subsurface projection of $\mu_0 \cup \mu_1$ into S_0. By Theorem 6.12 of [MM00], up to a universal multiplicative and additive constant the distance in the *marking graph* between any two markings μ_0, μ_1 equals

$$\beta(\mu_0, \mu_1) = \sum_{\delta(\mu_0,\mu_1;S_0)\geq p_0} \delta(\mu_0, \mu_1; S_0)$$

where the sum is taken over all subsurfaces S_0 of S. We also require that the corresponding formula is valid for distances in Teichmüller space as formulated in [R07]. Namely, we require that for any two points x_0, x_1 in the thick part of Teichmüller space and for any two short markings μ_0, μ_1 for x_0, x_1 as defined in the proof of Lemma 5.6, up to a universal additive and multiplicative constant the Teichmüller distance $d_{\mathcal{T}}(x_0, x_1)$ between x_0 and x_1 equals

$$\chi(\mu_0, \mu_1) = \sum_{\delta(\mu_0,\mu_1;S_1)\geq p_0} \delta(\mu_0, \mu_1; S_1) + \sum_{\delta(\mu_0,\mu_1;A)\geq p_0} \log \delta(\mu_0, \mu_1; A)$$

where the first sum is taken over all subsurfaces S_1 of S different from annuli and where A runs through the essential annuli.

Choose once and for all a map R which associates to a large train track τ a short marking $R(\tau)$ for τ. We assume that R is equivariant with respect to the action of the mapping class group. For a number $k > 0$ define a finite splitting and shifting sequence $(\tau_i)_{0\leq i\leq n}$ to be *of k-restricted type* if

$$\chi(R(\tau_0), R(\tau_n))/\beta(R(\tau_0), R(\tau_n)) \geq 1/k.$$

Call an infinite splitting and shifting sequence (τ_i) of k-restricted type if

$$\liminf_{i\to\infty} \chi(R(\tau_0), R(\tau_t))/\beta(R(\tau_0), R(\tau_t)) \geq 1/k.$$

Call a point $((x_i), \mu) \in \Omega_2$ of k-restricted type if it can be realized by a biinfinite splitting and shifting sequence of k-restricted type.

Let $\mathcal{R}_k \subset \Omega_2$ be the set of all biinfinite sequences $((x_i), \mu)$ of k-restricted type. Note that $\mathcal{R}_k \subset \mathcal{R}_{k+1}$ for all k, moreover $\mathcal{R}_k$ is σ_2-invariant. Let $\mathcal{M}_k$ be the closure in the weak*-topology of the space of all σ_2-invariant Borel probability measures on Ω_2 which give full measure to $\mathcal{R}_k$. We have

Lemma 8.11. *For every $k > 0$ there is a number $\delta(k) > 0$ such that $\int \rho_2 d\mu \geq \delta(k)$ for every $\mu \in \mathcal{M}_k$.*

Proof. Let μ be a σ_2-invariant Borel probability measure on Ω_2 which gives full mass to $\mathcal{R}_k \cap \mathcal{DU}_2$. By invariance, each of the measures in the ergodic decomposition of μ gives full mass to $\mathcal{R}_k \cap \mathcal{DU}_2$, so it is enough to assume that μ is ergodic. Then μ defines an invariant ergodic Borel probability measure on (X_2, Θ^t) which we denote again by μ, and

$$\hat{\mu} = G_2(\mu)$$

is a Φ^t-invariant Borel probability measure on $\hat{\mathcal{Q}}$.

Identify $\mathcal{DU}_2$ with the subset $\mathcal{DU}_2 \times \{0\}$ of X_2. Let $(x_i) \in \mathcal{DU}_2$ be a typical point for μ, i.e. a point for which the Birkhoff ergodic theorem holds true. By the Poincaré recurrence theorem, there is a sequence $r_i \to \infty$ such that $\sigma_2^{r_i}(x_i) \to (x_i)$ and that $G_2(\sigma_2^{r_i}(x_i)) \to G_2(x_i)$ $(i \to \infty)$.

Let R be the map used to define sequences of k-restricted type. By construction of the subshift (Ω_1, σ_1), for each large i there is a mapping class $g_i \in \Gamma$ such that $g_i\tau_0 = \tau_{r_i}$. By Theorem 2 and Lemma 6.6 of [H09a], the image under R of a splitting and shifting sequence is a uniform quasi-geodesic in the marking graph of S. Therefore there is a number $a > 0$ not depending on (x_i) such that the distance in the marking graph between $R(\tau_0)$ and $g_iR(\tau_0)$ is contained in the interval $[a^{-1}r_i, ar_i]$. By the Birkhoff ergodic theorem, it now suffices to show the existence of a number $\delta > 0$ not depending on μ with the following property. Let (τ_i) be a splitting and shifting sequence realizing (x_i) and let $\tilde{q} \in \mathcal{Q}(\tau_0)$ be a lift of $G_2(x_i) \in \hat{\mathcal{Q}}$. Let $t_i > 0$ be such that $\Phi^{t_i}\tilde{q} \in \mathcal{Q}(\tau_{r_i})$; then $t_i/r_i \geq \delta$ for all sufficiently large i.

Let $\Pi : \tilde{\mathcal{Q}}(S) \to \mathcal{T}(S)$ be the canonical projection and let as before $d_{\mathcal{T}}$ be the Teichmüller distance on $\mathcal{T}(S)$. Since $G_2(\sigma_2^{r_i}(x_i)) \to G_2(x_i)$ and since Γ acts freely on $\mathcal{T}(S)$, by equivariance the distance between $\Pi(\Phi^{t_i}\tilde{q})$ and $g_i(\Pi\tilde{q})$ in $(\mathcal{T}(S), d_{\mathcal{T}})$ tends to zero as $i \to \infty$. In particular, for large i the distance between $\Pi\tilde{q}$ and $g_i(\Pi\tilde{q})$ is contained in the interval $[t_i - 1, t_i + 1]$. Let $x \in \mathcal{T}(S)$ be a point which is short for the short marking $R(\tau_0)$ of τ_0. Note that x is contained in the thick part of Teichmüller space as explained in the proof of Lemma 5.6. By equivariance, if $u = d_{\mathcal{T}}(x, \Pi\tilde{q})$ then

$$d_{\mathcal{T}}(x, g_ix) \in [t_i - u - 1, t_i + u + 1].$$

On the other hand, by the definition of a sequence (x_i) of k-restricted type, by Theorem 6.1 of [R07] and by the discussion in the third paragraph of this proof, there is a number $b = b(k) > 0$ not depending on (x_i) so that

$$d_{\mathcal{T}}(x, g_ix) \geq br_i - b$$

for all sufficiently large i. Together we conclude that $t_i + u + 1 \geq br_i - b$ for large i. Since the number b does not depend on μ, this is what we wanted to show.

If μ is a σ_2-invariant ergodic Borel probability measure which gives full mass to $\mathcal{R}_k - \mathcal{DU}_2$ then by Lemma 8.10, μ is supported in a boundary multi-component $\mathcal{P}$ of μ. This implies that for a suitable choice of $\mathcal{P}$, the measure μ gives full mass to the set of uniquely ergodic sequences for this boundary multi-component. The above discussion applies as well and shows the lemma. ∎

The next observation is related to a result of Masur [M93]. He showed the following. Let $\delta_{\mathcal{T}}$ be a $\mathrm{Mod}(S)$-invariant distance on $\mathcal{Q}(S)$ with the property that the canonical projection $\mathcal{Q}(S) \to \mathcal{T}(S)/\mathrm{Mod}(S)$ distorts distances only by a universal additive amount (where the moduli space $\mathcal{T}(S)/\mathrm{Mod}(S)$ is equipped with the Teichmüller metric). Then for every quadratic differential $q \in \mathcal{Q}(S)$ and almost every $\theta \in [0, 2\pi)$ we have

$$\limsup_{t\to\infty} \delta_{\mathcal{T}}(q, \Phi^t e^{i\theta} q)/\log t = \frac{1}{2}.$$

This implies in particular that with respect to the Lebesgue measure, there is some $k > 0$ such that almost every $q \in \mathcal{Q}$ is of k-restricted type.

Lemma 8.12. *For every $\epsilon > 0$ there is a number $\ell = \ell(\epsilon) > 0$ with the following property. Let μ be a σ_2-invariant Borel probability measure on Ω_2 with $\mu(\Omega_2 - \mathcal{R}_\ell) = 1$; then $h_\mu - \epsilon \int \rho_2 d\mu \leq 0$.*

Proof. We claim that for every open relative compact set $V \subset \mathcal{Q}$ and for every $\epsilon > 0$ there is some $\ell > 0$ with the following property. Let μ be a σ_2-invariant Borel probability measure on Ω_2 with $\mu(\mathcal{DU}_2 - \mathcal{R}_\ell) = 1$; then $G_2(\mu)(\overline{V}) \leq \epsilon$.

To see that this holds true assume otherwise. Then there is an open relative compact set $V \subset \mathcal{Q}$, there is a number $\epsilon > 0$ and there is a sequence μ_ℓ $(\ell \to \infty)$ of σ_2-invariant ergodic Borel probability measures with $\mu_\ell(\Omega_2 - \mathcal{R}_\ell) = 1$ and such that $G_2(\mu_\ell)(\overline{V}) \geq \epsilon$. Let μ be a weak limit of the measures μ_ℓ. Then $G_2(\mu)(\overline{V}) \geq \epsilon$ and hence the discussion in the proof of Lemma 8.10 shows that also $\mu(\mathcal{DU}_2) \geq \epsilon$. On the other hand, we have $\int \rho_2 d\mu = 0$ by continuity which contradicts the fact that the function ρ_2 is nonnegative and positive on $\mathcal{DU}_2$.

This reasoning also applies to $(\Omega_2(\mathcal{P}), \sigma_2)$ where $\mathcal{P}$ is any boundary multi-component. The lemma now follows from Lemma 8.10 and Lemma 8.5 and by induction. ☐

We use this to complete the proof of Theorem 3 from the introduction.

Proposition 8.13. *For every $\epsilon > 0$ there is a compact set $K \subset \mathcal{Q}$ such that the entropy of any Φ^t-invariant Borel probability measure μ on $\mathcal{Q}$ with $\mu(\mathcal{Q} - K) = 1$ does not exceed $h - 1 + \epsilon$.*

Proof. Let $\epsilon > 0$ and let $\ell > 0$ be sufficiently large such that $h_\mu - \epsilon \int \rho_2 d\mu \leq 0$ for every Borel probability measure μ with $\mu(\mathcal{DU}_2 - \mathcal{R}_\ell) = 1$. Such a number exists by Lemma 8.12. By Abramov's formula and the discussion in the beginning of this section, for each such measure μ the entropy of $G_2(\mu)$ does not exceed ϵ.

Let $K_1 \subset K_2 \subset \dots$ be an increasing sequence of compact subsets of $\hat{\mathcal{Q}}$ such that $\cup_i K_i = \hat{\mathcal{Q}}$. Let $A_i = G_2^{-1}(\hat{\mathcal{Q}} - K_i) \cap \mathcal{R}_\ell$. We claim that there exists some $i_0 > 0$ with the following property. If μ is a σ_2-invariant Borel probability measure with $\mu(A_{i_0}) = 1$ then

$$h_\mu - (h - 1 + \epsilon) \int \rho_2 d\mu \leq 0.$$

Namely, otherwise there is a sequence (μ_i) of σ_2-invariant ergodic Borel probability measures such that $\mu_i(A_i) = 1$ for all i and that

$$h_{\mu_i} - (h - 1 + \epsilon) \int \rho_2 d\mu_i \geq 0. \tag{15}$$

By passing to a subsequence we may assume that the measures μ_i converge weakly to a measure $\mu \in \mathcal{M}_\ell$ with $\mu(\mathcal{DU}_2) = 0$. Lemma 8.10 shows that $h_\mu - (h - 1) \int \rho_2 d\mu \leq 0$. On the other hand, we have $h_\mu \geq \limsup_i h_{\mu_i}$ by Lemma 8.5 and $\int \rho_2 d\mu \geq \delta(\ell) > 0$ by Lemma 8.11 which is a contradiction.

Now $\mathcal{R}_\ell$ is invariant under σ_2 and hence every ergodic σ_2-invariant Borel probability measure μ on Ω_2 satisfies $\mu(\mathcal{R}_\ell) = 0$ or $\mu(\mathcal{R}_\ell) = 1$. On the other hand, by Lemma 5.9 every Φ^t-invariant measure μ on $\hat{\mathcal{Q}}$ is the image under G_2 of a σ_2-invariant measure on Ω_2. Using once more Abramov's formula, the proposition follows. □

Recall that a compact subset K of a stratum is a closed subset with the additional property that for some fixed $\epsilon > 0$, the length of any saddle connection for any point $q \in K$ is at least ϵ. Thus Proposition 8.13 can be reformulated as follows.

For every $\epsilon > 0$ there is a number $\delta > 0$ with the following property. Let $C \subset \mathcal{Q}$ be the set of all differentials with at least one saddle connection of length at most δ. Then the entropy of any Φ^t-invariant Borel probability measure ν on $\mathcal{Q}$ with $\nu(C) = 1$ does not exceed $h - 1 + \epsilon$.

There are refinements of this result which follow from the same argument. We formulate explicitly one such refinement. Define h_{thin} to be the maximal entropy of an ideal boundary multi-component of $\mathcal{Q}$. Note that every disconnected boundary multi-component is an ideal boundary multi-component. Let $\overline{\mathcal{Q}}$ be the closure of $\mathcal{Q}$ in $\mathcal{Q}(S)$.

Corollary 8.14. *For every $\epsilon > 0$ there is a compact set $K \subset \overline{\mathcal{Q}}$ such that the entropy of any Φ^t-invariant Borel probability measure μ on $\mathcal{Q}$ with $\mu(\mathcal{Q} - K) = 1$ does not exceed $h_{\text{thin}} + \epsilon$.*

Proof. The proof of Proposition 8.13 directly applies. Namely, argue by contradiction and assume that the claim does not hold true. Then there is a number $\ell > 0$ and there is sequence of invariant ergodic Borel probability measures μ_i on Ω_2 which give full measure to the set of points in $\mathcal{DU}_2 \cap \mathcal{R}_\ell$ whose image under G_2 consist of quadratic differentials with at least one simple closed curve of length at most $1/i$ and whose entropy is bigger than $h_{\text{thin}} + \epsilon$. Any weak limit of this sequence gives full mass to an ideal boundary multi-component. From this observation and upper semi-continuity of the entropy the corollary follows. □

There are also extensions of Proposition 8.13 to measures on the set of quadratic differentials with $k \geq 1$ short saddle connections whose interiors are pairwise disjoint. The latter restriction is necessary because an area one quadratic differential may have an arbitrarily large number of short saddle connections which are all contained in some proper subsurface. We refrain from giving precise technical formulations here.

9. Counting periodic orbits

This final section is devoted to the proof of Theorem 4 from the introduction. It directly builds on the results in Section 8. We use all assumptions and notations from Section 8.

For a number $k > 0$ define a periodic orbit for Φ^t in $\hat{\mathcal{Q}}$ to be *of k-restricted type* if one (and hence every) point on the orbit is a quadratic differential of k-restricted type.

Proposition 9.1. *For every $\epsilon > 0$ there is an open relative compact set $U \subset \hat{\mathcal{Q}}$ and for every $k > 0$ there is a number $c = c(k) > 0$ with the following property. The number of periodic orbits of Φ^t of k-restricted type and of length at most R which are contained in $\hat{\mathcal{Q}} - U$ does not exceed $ce^{(h-1+\epsilon)R}$.*

Proof. Let $\epsilon > 0$ and let $U \subset \hat{\mathcal{Q}}$ be an open relative compact set such that

$$h_\mu - (h - 1 + \epsilon) \int \rho_2 d\mu \le 0$$

for every σ_2-invariant Borel probability measure μ on Ω_2 with $\mu(\mathcal{D}\mathcal{U}_2 - G_2^{-1}(\hat{\mathcal{Q}} - U)) = 1$. Such a set exists by Proposition 8.13.

Our goal is to show that U has the property stated in the proposition for the number 2ϵ. For this we argue by contradiction and we assume that there is a number $\kappa > 0$ and there is a sequence $R_i \to \infty$ such that the number of periodic orbits for Φ^t in $\hat{\mathcal{Q}}$ of k-restricted type and of length at most R_i which do not intersect U is not smaller than $e^{(h-1+\epsilon+2\kappa)R_i}$.

The set of all points in $\hat{\mathcal{Q}}$ whose Φ^t-orbit is contained in $\hat{\mathcal{Q}} - U$ is a closed subset of $\hat{\mathcal{Q}}$. Since G_2 is continuous, the closure $A \subset \Omega_2$ of the set of all points whose σ_2-orbit is contained in $G_2^{-1}(\hat{\mathcal{Q}} - U)$ is compact and σ_2-invariant. Define

$$f = -(h - 1 + \epsilon)\rho_2;$$

by the choice of U and by Lemma 8.10 we have $h_\mu + \int f d\mu \le 0$ for all σ_2-invariant Borel probability measures μ on Ω_2 with $\mu(A) = 1$.

Recall from Lemma 8.4 the definition of the distance d on Ω_2. For numbers $\epsilon > 0, n > 0$ call a set $E \subset A$ (ϵ, n)*-separated* if for $x \neq y \in E$ there is some $p \le n$ such that $d(\sigma_2^p x, \sigma_2^p y) \ge \epsilon$.

For a finite set $E \subset \Omega_2$ write

$$Z_n(f, E) = \sum_{x \in E} \exp \sum_{k=0}^{n-1} f(\sigma_2^k x).$$

Define

$$H(f) = \lim_{\epsilon \to 0} \limsup_{n \to \infty} \frac{1}{n} \log\{Z_n(f, E) \mid E \subset A \text{ is } (n, \epsilon)\text{-separated}\}. \tag{16}$$

By the results of Walters [W75], $H(f)$ equals the pressure of the restriction of the continuous function f to the compact invariant set A. In particular, we have $H(f) \le 0$.

Let as before $P:\Omega_2\to\Omega_1$ be the first factor projection. Let $\delta>0$ be sufficiently small that if $x,y\in\Omega_2$ and if Px,Py are contained in *distinct* standard cylinders of length n then x,y are (δ,n)-separated. Such a number exists by the construction of the metric d. Since $H(f)\leq 0$ and since the expression on the right hand side of equation (16) is increasing as $\epsilon>0$ is decreasing, we conclude the following. For $n>0$ let $E_n\subset A$ be any (δ,n)-separated set and let $\beta>0$; then

$$C_\beta=\sum_n e^{-\beta n}Z_n(f,E_n)<\infty. \tag{17}$$

Let $k>0$. By Theorem 8.1 and the construction, for each periodic orbit $\{\Phi^t q\}$ for Φ^t which is contained in $\hat{\mathcal{Q}}-U$ and is of k-restricted type there is a periodic point $x(q)\in A$ for σ_2 of some period $n(q)>0$ with $G_2(x(q))=\Phi^s q$ for some $s\in\mathbb{R}$. The point $x(q)$ is defined by a train track expansion of a pseudo-Anosov mapping class which defines the orbit. If $\tau(q)>0$ is the period of $\{\Phi^t q\}$ then

$$\sum_{i=0}^{n(q)-1} f(\sigma_2^i x(q))=-(h-1+\epsilon)\tau(q).$$

Since splitting and shifting sequences define uniform quasi-geodesics in the marking graph (Theorem 2 and Lemma 6.6 of [H09a]), by the definition of an orbit of k-restricted type and the argument in the proof of Lemma 8.11, there is a number $\ell>0$ only depending on k such that

$$\tau(q)\geq n(q)/\ell.$$

Let $\beta=\kappa/\ell$. For fixed n, the periodic points $x(q)\in A$ with $n(q)=n$ which correspond to distinct Γ-orbits of train track expansions of some pseudo-Anosov elements define a (δ,n)-separated subset of A. Since Γ is torsion free, the Γ-orbits of any two train track expansions of two distinct pseudo-Anosov elements are distinct. Thus up to choosing the (δ,n)-separated sets for the calculation of the number $C_\beta>0$ in the equation (17) appropriately, each periodic orbit $\{\Phi^t q\}$ in $\hat{\mathcal{Q}}-U$ of k-restricted type contributes to the sum with the value $e^{-(h-1+\epsilon)\tau(q)}\geq e^{-(h-1+\epsilon+\kappa)\tau(q)}e^{\beta n(q)}$.

By assumption, for each i there are at least $e^{(h-1+\epsilon+2\kappa)R_i}$ periodic orbits of Φ^t of k-restricted type and of length at most R_i which are contained in $\hat{\mathcal{Q}}-U$. The summands in (17) are all positive, and there are at least $e^{(h-1+\epsilon+2\kappa)R_i}/b$ summands with value not smaller than $e^{-(h-1+\epsilon+\kappa)R_i}$. Now the set $\{R_i\mid i\}$ is unbounded and hence this violates the fact that $C_\beta<\infty$ in (17). From this contradiction, the proposition follows. □

We are left with the (harder) task of counting the number of periodic orbits which are not of ℓ-restricted type for some large $\ell>0$. By definition, these are orbits with long subsegments twisting about a simple closed curve.

We begin with locating regions of twisting along sequences in $\mathcal{DU}_1\subset\Omega_1$. For this let c be an embedded simple closed curve in a large train track τ. Since S is oriented, a choice of an orientation of c defines a right and left side of c in S. A half-branch b in $\tau-c$ with endpoint in c is *right* or *left* according to whether it is contained in the right or left side of c in τ. It is *incoming* if any smooth oriented arc in τ which begins in b and passes through the endpoint v of b defines locally near v the orientation of c, and it is *outgoing* otherwise.

There is a number $p_1 > 0$ with the following property (see [H09c] for details). Let (τ_i) be any splitting and shifting sequence issuing from $\tau_0 = \tau$ which consists of splitting and shifting moves at large branches in c. Then for every $p \geq p_1$, either every left half-branch $b \in \tau_p - c$ with endpoint in c is incoming and every right half-branch is outgoing, or every left half-branch $b \in \tau_p - c$ with endpoint in c is incoming and every right half-branch is outgoing. In this case we call *c in standard position* in τ_p.

If c is in standard position in a large train track τ then there is a right or left Dehn twist φ_c about c so that $\varphi_c \tau \prec \tau$. The direction of the twist is determined by τ. Moreover, there is a train track τ' which contains c as a twist connector (i.e. c is the union of a single small branch and a single large branch) and such that $\varphi_c \tau \prec \tau' \prec \tau$ (we refer to [H09c] for details).

Let as before R be a Mod(S)-equivariant map which associates to a large train track $\tau \in \mathcal{LT}(\mathcal{Q})$ a short marking. By Theorem 5.3 of [MMS10] and the above discussion, after perhaps replacing the number p_0 used in the definition of the quantities $\beta(\mu_0, \mu_1)$ and $\chi(\mu_0, \mu_1)$ by a bigger constant, there are numbers $p_2 > 3(p_0 + p_1)$ and $\kappa > 0$ with the following property.

Let $(\tau_i) \subset \mathcal{LT}(\mathcal{Q})$ be a splitting and shifting sequence. Let $A \subset S$ be an annulus and assume that $\delta(R(\tau_0), R(\tau_u); A) = p \geq p_2$ for some $u > 0$. Let c be a core curve of A. Then there is some smallest number $t_0(c) \geq 0$ and a largest number $t_1(c) \leq u$ such that c is embedded in both $\tau_{t_0(c)}$ and $\tau_{t_1(c)}$ and is in standard position. Moreover, the following properties are satisfied.

(1) $\delta(R(\tau_0), R(\tau_{t_0(c)}); A) \leq p_0$, $\delta(R(\tau_{t_1(c)}), R(\tau_u); A) \leq p_0$.
(2) There is a (positive or negative) Dehn twist φ_c about c, and there is a number $k \geq 0$ such that $\tau_{t_1(c)} \prec \varphi_c^k \tau_{t_0(c)}$. Moreover, $\tau_{t_1(c)}$ can be obtained from $\varphi_c^k \tau_{t_0(c)}$ by at most κ splitting and shifting moves at large branches in c, followed by a sequence of splitting and shifting moves at branches in the complement of c.

We call the interval $[t_0(c), t_1(c)]$ the *region of twisting about* c. Note that if $j > t_1(c)$ then the curve c is not embedded in τ_j. There is an obvious extension of this definition to infinite splitting and shifting sequences where we allow $t_0(c) = -\infty$ and $t_1(c) = \infty$. The *twisting number* of c along (τ_i) is the largest number $k \geq 0$ so that property 2) above holds true for k. If (τ_i) is infinite, this can be infinite as well. We say that c is *twisted* along the sequence (τ_i) if the twisting number of c along (τ_i) is at least three. We say that c is *strongly twisted* along the sequence (τ_i) if the twisting number of c along (τ_i) is at least five.

Assume that c is twisted along (τ_i) and let $[t_0(c), t_1(c)]$ be the region of twisting. Let moreover $k \geq 3$ be the twisting number of c. By the above remark (compare [H09c]) there is a train track τ' which contains c as a twist connector and such that

$$\tau_{t_1(c)} \prec \varphi_c^{k-1} \tau' \prec \tau' \prec \tau_{t_0 c}.$$

Moreover, τ' can be obtained from $\tau_{t_0(c)}$ by a uniformly bounded number of splitting and shifting moves at large branches contained in c. If d is any simple closed curve embedded in $\varphi_c\tau'$ which intersects c then d is not embedded in τ', and it is not embedded in $\varphi_c^2\tau'$. In particular, d is not twisted along (τ_i).

Call a splitting and shifting sequence (τ_i) *in normal form* if the following holds true. Let c be any simple closed curve which is twisted along (τ_i). Let $[t_0(c), t_1(c)]$ be the region of twisting. Then there is an interval $[t_0'(c), t_1'(c)] \subset [t_0(c), t_1(c)]$ with the following properties.

(1) c is untwisted in $(\tau_i)_{i \leq t_0'(c)}$ and in $(\tau_i)_{i \geq t_1'(c)}$.
(2) c is a twist connector in each of the train tracks τ_i for $t_0'(c) \leq i \leq t_1'(c)$.

Call a train track expansion $\tau \in \mathcal{LT}(\mathcal{Q})$ of a pseudo-Anosov element $g \in \mathrm{Mod}(S)$ *in normal form* if τ can be connected to $g\tau$ by a splitting and shifting sequence in normal form. The above discussion easily implies

Lemma 9.2. *Every pseudo-Anosov element $g \in \mathrm{Mod}(S)$ admits a train track expansion in normal form.*

Proof. Let $\tau \in \mathcal{LT}(\mathcal{Q})$ be any train track expansion for g. Let (τ_i) be a g-periodic splitting and shifting sequence with $\tau_0 = \tau, \tau_k = g\tau$. Let c be a simple closed curve which is embedded in τ_0 and is in standard position. Assume that c is twisted along (τ_i). Then there is a Dehn twist φ_c about c so that either $\tau_k \prec \varphi_c\tau_0 \prec \tau_0$ or that $\tau_0 \prec \varphi_c\tau_0 \prec \tau_{-k}$.

In the first case it follows from the above discussion that we can modify τ_0 with a sequence of splitting and shifting moves at large branches in c to a train track τ' which contains c as a twist connector and such that $g\tau_0 \prec \varphi_c\tau_0 \prec \tau' \prec \tau_0$. Then $g\tau' \prec \tau'$ and hence τ' is a train track expansion for g. The second case is handled in the same way. Successively we can now construct a train track expansion for g in nomal form. □

Observe that two twist connectors c_1, c_2 in a large train track τ do not intersect. Namely, otherwise the large branch of τ contained in c_1 is contained in c_2 as well. By the definition of a twist connector, this implies that $\tau = c_1 \cup c_2$ which is impossible since S is of higher complexity by assumption. In particular, two distinct twist connectors in a large train track τ determine two disjoint not freely homotopic simple closed curves in S.

Let $b \leq 3g - 3 + m$ be the maximal number of pairwise disjoint twist connectors which are embedded in some train track $\tau \in \mathcal{LT}(\mathcal{Q})$. A *decorated train track* is a train track τ together with a labeling of all twist connectors c_i in τ with distinct numbers in $\{1, \dots, b\}$. The Γ-orbits of decorated train tracks define a finite alphabet E_3 which consists of decorated letters from the alphabet E_1. There is a decoration forgetful map $E_3 \to E_1$, and there is a Markov chain (Ω_3, σ_3) with alphabet E_3 which is determined by the requirement that the forgetful map $E_3 \to E_1$ defines a semi-conjugacy $(\Omega_3, \sigma_3) \to (\Omega_1, \sigma_1)$.

Define a finite or infinite decorated splitting and shifting sequence (η_i) to be *basic* if (η_i) is in normal form and if moreover the following properties hold true.

(1) No simple closed curve is strongly twisted along (η_i).
(2) If c is twisted along (η_i) and if c is a twist connector in both η_i and η_j then the labels of c in η_i and η_j coincide.

An *extension* of a basic decorated finite or infinite splitting and shifting sequence (η_i) is a decorated splitting and shifting sequence (τ_i) inductively defined as follows. To each i and each twist connector c in η_i with the additional property that c is twisted in η_i associate a number $b_i(c) \geq 0$. Let $\tau_0 = \eta_0$ and let $c_0, \dots, c_k$ be the twist connectors in τ_0 to which positive numbers $b_0(c_i) > 0$ have been assigned. Let φ_{c_i} be the Dehn twist about c_i so that $\varphi_{c_i}\tau_0 \prec \tau_0$. Let $u = \varphi_{c_1}^{b_0(c_1)} \cdots \varphi_{c_k}^{b_0(c_k)}$ and connect τ_0 to $u\tau_0 = \tau_s$ by a decorated splitting and shifting sequence which inherits the decoration from τ_0. Define $\tau_{s+1} = u\eta_1$ and repeat this construction with τ_{s+1} and the basic sequence $u(\eta_i)_{1 \leq i \leq p}$ for which the numbers associated to twist connectors are inherited from the numbers associated to the twist connectors in (η_i).

Informally we say that a periodic orbit in $\hat{\mathcal{Q}}$ has *small entropy* if its preimages in (Ω_2, σ_2) approximate measures with small entropy. More precisely, for a number $\alpha > 0$ we say that the entropy of a periodic orbit $\{\Phi^t q\}$ is *α-small* if there is a periodic basic decorated splitting and shifting sequence (η_i) with the following properties.

(1) There is an enlargement (τ_i) of (η_i) which is mapped by the composition of the decoration forgetful map with the map Ξ to $\{\Phi^t q\}$.
(2) The prime period of (η_i) does not exceed αR where R is the period of $\{\Phi^t q\}$.

The following proposition is not needed for the proof of Theorem 4. We include it here because its proof describes the structure of periodic orbits near the boundary of moduli space. it is possible that Theorem 4 can be proved with a similar argument.

Proposition 9.3. *For every $\epsilon > 0$ there is a number $\alpha = \alpha(\epsilon) > 0$ such that the number of periodic orbits for Φ^t of period at most R whose entropy if α-small is at most $e^{bR+\epsilon}$.*

Proof. For the proof of the proposition we evoke Minsky's product region formula. Namely, a pants decomposition P of S defines a system of Fenchel Nielsen coordinates $(\ell_1, \dots, \ell_{3g-3+m}, \tau_1, \dots, \tau_{3g-3+m})$ for Teichmüller space $\mathcal{T}(S)$. Here with respect to some numbering of the components of P, ℓ_i is the length of a pants curve γ_i, and τ_i is the *twist parameter.*

The product region theorem (Theorem 6.1 of [M96]) states the following. Let $\gamma \subset P$ be a multicurve with $u \geq 1$ components $\gamma_1, \dots, \gamma_u$. For some sufficiently small $\delta > 0$ let $V \subset \mathcal{T}(S)$ be the region of all marked hyperbolic metrics for which the length of each of the components γ_i is at most δ. Let S_0 be the (possibly disconnected) surface obtained by replacing each of the components γ_i of γ by a pair of punctures. The $(6g-6-2u+2m)$-tuple of Fenchel Nielsen coordinates at pants curves in $P-\gamma$ defines Fenchel Nielsen coordinates for S_0. The natural coordinate forgetful map determines a projection $\Pi_0 : V \to \mathcal{T}(S_0)$.

For each component γ_i of γ let $\mathbf{H}_i$ be a copy of the upper half-plane equipped with the hyperbolic metric. Define a map $\Pi_i : V \to \mathbf{H}_i$ by $\Pi_i(x) = (1/\ell_i, \tau_i)$ where (ℓ_i, τ_i) are the length and twist parameters for γ_i. In Fenchel Nielsen coordinates, the positive Dehn twist about the curve γ_i preserves the length and twist parameters of any pants curve $\zeta \in P - \gamma_i$, and it projects to the transformation $(\ell_i, \tau_i) \to (\ell_i, \tau_i + 1)$. For two points $x, y \in V$ define

$$d_P(x,y) = \max\{d_{\mathcal{T}(S_0)}(\Pi_0(x), \Pi_0(y)), d_{\mathbf{H}_i}(\Pi_i(x), \Pi_i(y))\}$$

where $d_{\mathcal{T}(S_0)}$ denotes the Teichmüller distance on $\mathcal{T}(S_0)$ and where $d_{\mathbf{H}_i}$ is the hyperbolic metric on $\mathbf{H}_i$. By Theorem 6.1 of [M96], there is a constant $a > 0$ only depending on δ such that for all $x, y \in V$,

$$|d_{\mathcal{T}(S)}(x,y) - d_P(x,y)| \le a.$$

(Ω_3, σ_3) is a topological Markov chain and hence there is a number $\kappa > 0$ so that any two sequences in Ω_3 contained in different standard cylinders of length n are (κ, n)-separated. The topological entropy of (Ω_3, σ_3) is finite and therefore for $\epsilon > 0$ there is a number $\alpha < \epsilon/a$ so that the number of distinct admissible decorated sequences of length αn is at most $e^{\epsilon n}$.

A periodic orbit $\{\Phi^t q\}$ in $\hat{\mathcal{Q}}$ of period at most R whose entropy is α-small is the image of a periodic enlargement (τ_i) of a periodic basic decorated splitting and shifting sequence (η_i). Let $g \in \mathrm{Mod}(S)$ be the pseudo-Anosov element which defines $\{\Phi^t q\}$ and which preserves $\cup_u \sigma_3^u(\tau_i)$. Then there is an element $g_0 \in \mathrm{Mod}(S)$ such that $g_0\eta_0 = \eta_s$ for some $s \le \alpha R < \epsilon R/a$, and there is a fundamental domain for the action of g on $\cup_u \sigma_3^u(\tau_i)$ which is an enlargement of the finite sequence $(\eta_i)_{0 \le i \le s}$.

Assume first that there is no twist connector in η_0 which is also a twist connector in η_s. Then we may assume that the following holds true. Let $p > s$ be such that $g\tau_0 = \tau_s$. We may assume that for every simple closed curve c, either c is not twisted in $(\tau_i)_{i \le 0}$ or c is not twisted in $(\tau_i)_{i \ge 0}$. By the product region formula, this means the following. Fix a label $j \le b$ and let $c_1, \dots, c_k$ $(k \le s)$ be the distinct twist connectors along (η_i) with label j. For each $i \le k$ let $t(c_i)$ be the twisting number of c_i in (τ_i) (with the obvious interpretation); then $\sum_i \log t(c_i) \le R + ka \le (1+\epsilon)R$.

We now count the periodic orbits in $\hat{\mathcal{Q}}$ of length at most R which are obtained as enlargements of the fixed basic decorated sequence (η_i). Our goal is to show that there are at most $ce^{(b+\epsilon)R}$ such orbits for some $c > 0$.

For this write $v = R/k$. Assign to each of the twist connectors c_i an integer $j(i) \ge 0$ so that $\sum_i j(i) \le vk(1+\epsilon)$. The integer $j(i)$ is viewed as the logarithm of the twisting number of c_i. The number of different orbits whose logarithmic length distributions are within one of the given distribution roughly equals $\sum_i e^{j(i)}$. In particular, this number does not exceed a constant multiple of $e^{vk(1+\epsilon)} = e^{R(1+\epsilon)}$.

The number of all tuples $(j(1),\dots,j(k))$ with $\sum_i j(i)\leq R(1+\epsilon)$ is the number of non-negative integral lattice points in $\mathbb{R}^k$ which are contained in a hypersimplex of side length kp where $p=v(1+\epsilon)$. The volume of such a hypersimplex is

$$(kp)^k/k!\leq p^k e^k = e^{k(1+\log p)} = e^{R(1+\epsilon)(1+\log p)/p}$$

and the number of integral lattice points in this hypersimplex is bounded from above by a constant multiple of this volume. Now as $p\to\infty$, we have $(1+\log p)/p\to 0$ and hence there is some sufficiently large $p>0$ such that this number does not exceed $e^{\epsilon R}$. As a consequence, the number of distinct orbit points counted in this way is not bigger than $e^{(1+2\epsilon)R}$.

Now if c,c' are disjoint simple closed curves then since the Teichmüller metric is a sup metric, these two curves only contribute with the max of the log of their twisting length to the length of the orbit. Since there are at most b disjoint twist connectors in a large train track $\eta\in\mathcal{LT}(\mathcal{Q})$, the number of periodic orbits of length R which are defined by the same sequence (η_i) is at most $e^{b(2\epsilon+1)R}$. This shows the claim.

The case that there is a twist connector c in η_0 which also is a twist connector in η_s follows in the same way. Namely, since g is pseudo-Anosov, we have $g_0c\neq c$ (since otherwise also $gc=c$). As a consequence, the label j of gc is different from the label i of c and hence the twisting numbers of the curves with label i determine the twisting numbers of the curves with label j.

The map g permutes the labels. If there is no twist connector in η_0 which is also a twist connector in η_{2s} then let $c_1,\dots,c_p$ be the twist connectors in η_0 which are also twist connectors in η_s. Then g maps the set of labels $\{1,\dots,p\}$ disjointly from itself. Assume that c_i has the label i and that the label of gc_i is c_{i+p}. To determine the possibilities for the twisting numbers of the labeled curves with label $j\leq p$, consider the sequence $(\eta_i)_{0\leq i\leq 2s}$. The above discussion shows that there are roughly e^{2R} possibilites, but each such choice determines the twisting numbers for the curves with label c_{j+p}. As before, we conclude that the number of possibilities does not exceed $e^{(b+2\epsilon)R}$. The case that there is a twist connector c in η_0 which also is a twist connector in η_{uk} for some $u\geq 3$ is treated in the same way. ☐

Counting periodic orbits which are neither of ℓ-restricted type nor of very small entropy is harder. Namely, such an orbit may pass through regions in the thin part of moduli space corresponding to varying boundary strata which gives rise to combinatorial difficulties for the approach in the proof of Proposition 9.3.

Instead we use the strategy employed to control orbits of ℓ-restricted type. There are two difficulties to overcome. The first is the fact that the roof function ρ_2 on Ω_2 is not bounded away from zero. Second, periodic orbits in $\hat{\mathcal{Q}}$ are not separated in regions of twisting.

To overcome these difficulties we change the topology at infinity of the flow space $\hat{\mathcal{Q}}$ in such a way that in the resulting space, periodic orbits are separated.

Let (Ω_4,σ_4) be the product of the shift space Ω_3 with b copies of the Bernoulli shift (Σ_i,ν_i) $(i=1,\dots,b)$ in the letters $0,1$. We choose the numbering of the factors in such a way that Ω_3 is the zero component. Let $\Pi_0 : \Omega_4 \to \Omega_3$ be the canonical projection, and for $1 \leq i \leq b$ let $\Pi_i : \Omega_4 \to \Sigma_i$ be the projection onto the i-th Bernoulli shift.

Define an *extended splitting and shifting sequence* to be a (finite or infinite) sequence (τ_i) so that for each i, either τ_{i+1} is obtained from τ_i by a splitting and shifting move or $\tau_{i+1} = \tau_i$. The definition of a simple closed curve c which is twisted along a (perhaps decorated) splitting and shifting sequence carries over without changes to extended splitting and shifting sequences. Similarly, there is a natural definition of a basic extended splitting and shifting sequence.

If $(x_i) \in \Omega_4$ is any sequence and if the decorated train track τ_0 is a representative of the letter $\Pi_0(x_i)_0 \in E_3$ then $\Pi_0(x_i)$ determines uniquely an extended decorated splitting and shifting sequence $\omega(x_i)$ through $\omega(x_i)_0 = \tau_0$. We say that $\omega(x_i)$ *realizes* $\Pi_0(x_i)$.

Define $\hat{\Omega}_4$ to be the largest σ_4-invariant subset of Ω_4 with the following properties.

a) An extended splitting and shifting sequence $\omega(x_i)$ which realizes $\Pi_0(x_i)$ is basic.
b) If the active component j of the transition modifying $(x_i) \in \hat{\Omega}_4$ to $\sigma_4(x_i)$ is not the zero-component, then a representative of the zero-component of x_0 contains a twist connector c labeled with j. Moreover, c is twisted along $\omega(x_i)$.
c) Let c be a twist connector in $\omega(x_i)$ labeled with $j \leq b$. Assume that the active component of the transition which modifies (x_i) to $\sigma_4(x_i)$ is the component j. Let $k > 0$ be the smallest number so that the transition which modifies $\omega(x_i)_k$ to $\omega(x_i)_{k+1}$ is a splitting and shifting move at a large branch in c. If $\ell > k$ is such that the active component of the transition which modifies σ_4^ℓ to $\sigma_4^{\ell+1}$ is the component j then c is not embedded in $\omega(x_i)_\ell$.

The following observation is immediate from the definition.

Lemma 9.4. *$\hat{\Omega}_4$ is a closed subset of Ω_4.*

Proof. Each of the conditions a),b),c) is a closed condition. □

To each sequence $(x_i) \in \hat{\Omega}_4$ and each decorated train track $\tau_0 = \omega(x_i)_0$ representing the letter $\Pi_0(x_i)_0 \in E_3$ we associate inductively a decorated splitting and shifting sequence $\Lambda(x_i) = (\tau_i)$ issuing from $\tau_0 = \Lambda(x_i)_0$ and a number $k(x_i) \geq 0$ as follows.

If the active component of the transition transforming (x_i) to $\sigma_4(x_i)$ is the zero component, then using the above notation, the decorated train track $\omega(x_i)_1$ is obtained from $\tau_0 = \omega(x_i)_0$ by a splitting and shifting move. Define

$$\Lambda(x_i)_1 = \tau_1 = \omega(x_i)_1 \text{ and } k(x_i) = 1$$

and repeat this construction with $\sigma_4(x_i)$ and $\Lambda(x_i)_1$.

If the active component is the component $p \in \{1, \dots, b\}$ then by property b) above, there is a twist connector c embedded in τ_0 with label p. Let $u \in [1, \infty]$ be the maximum of all numbers with the following properties.

(1) The active component of the transition transforming $\sigma_4^{u-1}(x_i)$ to $\sigma_4^u(x_i)$ is the component p.
(2) The extended splitting and shifting sequence $\omega(x_i)_{0 \le i \le u}$ does not contain a splitting and shifting move at a large branch in c.

Assume first that $u < \infty$. Let $v \le 0$ be the smallest number so that the curve c is a twist connector in $\omega(x_i)_v$. The string $(x_i)_{v \le i \le u}$ defines a (finite or half-infinite) string $(a_i)_{0 \le i \le s}$ ($s \in [0, \infty]$) in the numbers $0, 1$ inductively as follows.

Put $\tilde{a}_0 = \Pi_p(x_i)_u$. Let j be the largest number in the interval $[v, u-1]$ so that the active component of the transition transforming $\sigma_4^{j-1}(x_i)$ to $\sigma_4^j(x_i)$ is the component p. If no such number exists then put $s = 0$. Otherwise let $\tilde{a}_1 = \Pi_p(x_i)_j$ and proceed inductively. The last entry $\tilde{a}_s$ (here it is possible that $s = \infty$) in the string is obtained from the transition transforming $\sigma_4^{v+t}(x_i)$ to $\sigma_4^{v+t+1}(x_i)$ where $t \ge 0$ is the smallest nonnegative number such that the active component of this transition is the component p. Each entry in this (finite or half-infinite) string corresponds to exactly one transition with active component p. In particular, there is a number $\ell(x_i) \ge 0$ so that the transition transforming (x_i) to $\sigma_4(x_i)$ corresponds to the number $\tilde{a}_{\ell(x_i)}$.

For $i < 0$ and $i > s$ define $\tilde{a}_i = 0$ and let $a_i = \tilde{a}_{i+\ell(x_i)}$. Then (a_i) is an infinite string in the numbers $0, 1$. In particular, the dyadic roof function ζ as constructed in the appendix is defined for this string and any of its shifts. For $n \ge 1$ write

$$\zeta^n(a_i) = \sum_{j=0}^{n-1} \zeta(\nu^j(a_i)).$$

Let $r_0(n) \ge 0$ be the largest integer which is not larger than $2^{\zeta^n(a_i)}$. Note that if $v = 0$ then $r_0(\ell(x_i)) = 2^{\zeta^{\ell(x_i)}(a_i)}$. Put $r_0(-1) = 0$ and define

$$k(x_i) = 2(r_0(\ell(x_i)) - r_0(\ell(x_i) - 1)) \ge 0. \tag{18}$$

Let φ_c be the unique Dehn twist about c so that $\varphi_c(\tau_0) \prec \tau_0$. Connect τ_0 with a decorated splitting sequence of length $k(x_i)$ to

$$\Lambda(k(x_i)) = \varphi_c^{k(x_i)/2} \tau_0.$$

We assume that the decoration is inherited from the decoration of τ_0. Note that since c is a twist connector by assumption, the train track $\varphi_c(\tau_0)$ is obtained from τ_0 by two consecutive splits at a large branch in c. Repeat this construction with $\sigma_4(x_i)$ and $\Lambda(k(x_i))$.

If $u = \infty$ then we mark the curve c. Let $s \ge 1$ be the smallest number so that the active component of the transition from $\sigma_4^s(x_i)$ to $\sigma_4^{s+1}(x_i)$ is distinct from the component p. Define $\Lambda(x_i)_1$ with the above procedure but by replacing (x_i) by $\sigma_4^s(x_i)$. If no such s exists then define $\Lambda(x_i)_{i \ge 0}$ to be the constant sequence.

As in Section 8, for a splitting and shifting sequence (τ_i) let $B(\tau_i)$ be the set of all projectivizations of all measured geodesic laminations which are contained in $\cap_i \mathcal{V}(\tau_i)$. For $(x_i) \in \hat{\Omega}_4$ and a representative τ_0 of the zero-component of x_0 construct the sequence $\Lambda(x_i)$ through $\Lambda(x_i)_0 = \tau_0$ and define $B((x_i), \tau_0)$ to be the subset of all projective measured geodesic laminations which are contained in $\cap_i \mathcal{V}(\Lambda(x_i))$ and which moreover do not intersect any marked simple closed curve in $\cap_i \mathcal{V}(\Lambda(x_i))$. As in Section 8, for $g \in \Gamma$ we have $B((x_i), g\tau_0) = gB((x_i), \tau_0)$. In particular, the group Γ naturally acts as a group of transformations on $\{((x_i), \tau_0, \xi) \mid (x_i) \in \hat{\Omega}_4, \tau_0 \in \Pi_0(x_i)_0, \xi \in B((x_i), \tau_0)\}$ equipped with the product topology. We have

Lemma 9.5. *The quotient space*

$$\Omega_5 = \{((x_i), \tau_0, \xi) \mid (x_i) \in \hat{\Omega}_4, \tau_0 \in \Pi_0(x_i)_0, \xi \in B((x_i), \tau_0)\}/\Gamma$$

is compact.

Proof. Since $\hat{\Omega}_4$ and $\mathcal{PML}$ are compact, it suffices to show the following. Let $(x_i)_j \subset \hat{\Omega}_4$ be a sequence converging to some $(y_i) \in \hat{\Omega}_4$. Assume without loss of generality that the zero components of $(x_0)_j, y_0$ coincide. Let τ_0 be a representative of this zero component. Let $\xi_j \in B((x_i)_j, \tau_0)$; then up to passing to a subsequence, $\xi_j \to \xi \in B((y_i), \tau_0)$.

However, this is almost immediate from the construction. Namely, let $\Lambda(x_i)_j$ and $\Lambda(y_i)$, respectively, be decorated splitting and shifting sequences constructed as above from $(x_i)_j, (y_i)$ and τ_0. If there is no simple closed curve c which is a twist connector in all but finitely many of the train tracks $\Lambda(y_i)$ $(i \geq 0)$ then the sequence $\Lambda(y_i)$ is infinite, and the twisting number in $\Lambda(y_i)_{i\geq 0}$ of any simple closed curve is finite. In this case the desired conclusion follows as in the proof of Lemma 8.4.

Now assume that there is a simple closed curve c and there is a number $n \geq 0$ so that c is a twist connector in each of the train tracks $\Lambda(y_i)$ $(i \geq n)$. Let $p \leq b$ be the label of c. Then for sufficiently large j the curve c is a twist connector in $\Lambda(x_n)_j$. The case that the sequence $\Pi_p(y_i)_{i\geq 0}$ is finite follows as above. If the sequence $\Pi_p(y_i)_{i\geq 0}$ is infinite then as $j \to \infty$, the twisting number $k(j)$ of c in $\Lambda(x_i)_j$ $(i \geq n)$ tends to infinity.

If a train track η_1 is obtained from η_0 by a split at a large branch e and if μ is a transverse measure on η_1, then μ induces a transverse measure $\hat{\mu}$ on η_0. The $\hat{\mu}$-weight of a branch $b \neq e$ in η_0 coincides with the μ-weight of the corresponding branch in η_1. The $\hat{\mu}$-weight of the branch e is the sum of the μ-weight of the branch corresponding to e in η_1 and the μ-weights of the two *losing* half-branches of the split, i.e. those half-branches which are incident on the endpoints of e and which are small in the split track η_1. As a consequence, if a large train track η_0 contains a twist connector c, if φ_c is a Dehn twist about c, if $\eta_1 = \varphi_c\eta_0 \prec \eta_0$ and if λ is a measured geodesic lamination which is carried by η_0 then the λ-weight of c in η_0 is the sum of the λ-weight of c in η_1 and the λ-weights of the two half-branches which are incident on the switches of η_0 in c and which are not contained in c.

Assume without loss of generality that c is a twist connector in each of the train tracks $\Lambda(x_n)_j$ $(j > 0)$. Let $u(j)$ be the largest number so that c is a twist connector in $\Lambda(x_{u(j)})_j$. Assume for the moment that $u(j) < \infty$ for all j. Let d_j be any measured geodesic lamination which is contained $\cap_i \mathcal{V}(\Lambda(x_i)_j) \cap \mathcal{V}_0(\tau_0)$. Then the d_j-weight of c on $\Lambda(x_n)_j$ is not smaller than $k(j)$ times the d_j-weight of the two branches in $\Lambda(x_n)_j$ which are incident on switches in c but which are not contained in c. But this just means that as $j \to \infty$, every limit in $\mathcal{PML}$ of the projectivizations of d_j can be represented as a union of the weighted curve c and a measured geodesic lamination whose support is disjoint from c. Together with the above discussion, this implies that any limit of a sequence $\xi_j \in B((x_i)_j, \tau_0)$ is contained in $B((y_i), \tau_0)$.

The case that $u_j = \infty$ for infinitely many j follows from exactly the same argument. □

Let $\tilde{\Omega}_5 \subset \Omega_5$ be the set of all Γ-orbits of sequences $((x_i), \tau_0, \xi)$ so that the decorated splitting and shifting sequence $\Lambda(x_i)$ is contained in the set $\mathcal{DU}_3$ of all decorated sequences which are mapped by the decoration forgetful map into $\mathcal{DU}_1 \subset \Omega_1$. The set $\mathcal{DU}_3$ is dense in Ω_5. Define a roof function χ on $\tilde{\Omega}_5$ as follows. Let $((x_i), \tau_0, \xi)$ be a representative of a point in Ω_5. Let $(\tau_i) = \Lambda(x_i)$ be a splitting and shifting sequence beginning at τ_0 as defined above. Let $\mu \in \mathcal{V}_0(\tau_0)$ be a measured geodesic lamination which represents ξ. Let $k(x_i)$ be as in equation (18). Define

$$\tilde{\chi}((x_i), \tau_0, \xi) = \log(\mu(\tau_0)/\mu(\tau_{k(x_i)})).$$

The function $\tilde{\chi}$ is invariant under the action of Γ. In particular, it projects to a function χ on $\tilde{\Omega}_5$.

Lemma 9.6. *The function χ on $\tilde{\Omega}_5$ is continuous and bounded, and it extends to a continuous function on Ω_5.*

Proof. By construction of χ, to show that χ is bounded it suffices to show the following. Let (τ_i) be any splitting and shifting sequence. Let c be a twist connector in τ_0. Assume that $k > 0$ is such that c is a twist connector in τ_k and that c is not a twist connector in τ_{k+1}. Assume moreover that $\tau_k \prec \varphi_c^{-1}\tau_k$ where φ_c is a (positive or negative) Dehn twist about c. Let $\lambda \in \cap_i \mathcal{V}(\tau_i)$; then

$$\lambda(\varphi_c^{-s+1}\tau_k)/\lambda(\varphi_c^{-s}\tau_k) \geq 1 - 1/s. \tag{19}$$

Namely, assume that this holds true. Then for all $s > 0$ we have

$$\log \lambda(\varphi_c^{-s}\tau_k) - \log \lambda(\varphi_c^{-s+1}\tau_k) \leq c_1/s$$

for a universal constant $c_1 > 0$. Let $k(x_i)$ be as in equation (18). Then up to a uniformly bounded constant, $k(x_i)/2 = 2^{\zeta_\ell(a_i)} - 2^{\zeta_{\ell-1}(a_i)}$ for a number $\ell = \ell(x_i) > 1$. As a consequence, we conclude that

$$\log \lambda(\tau_0) - \log \lambda(\varphi_c^{k(x_i)/2}\tau_0) \leq c_2 \sum_{i=2^{\zeta^{\ell-1}(a_i)}}^{2^{\zeta^{\ell}(a_i)}} \frac{1}{i} \leq c_3$$

for numbers $c_3 > c_2 > 0$. From this boundedness of χ is immediate.

To show the estimate (19), assume that $\lambda(\tau_k)=1$. Then the sum β of the weights of the two half-branches of τ_k which are incident on switches in c and which are not contained in c is at most one. Let $\alpha\geq\beta$ be the weight of c. By the discussion in the proof of Lemma 9.5, if $s>0$ then the λ-weight of c in $\eta=\varphi_c^{-s}\tau_k$ equals $\alpha+s\beta$, in particular, we have

$$\lambda(\varphi_c^{-s}\tau_k)=1+s\beta.$$

In the case $\beta=0$ the estimate (19) is now immediate. If $\beta>0$ then

$$\lambda(\varphi_c^{-s+1}\tau_k)/\lambda(\varphi_c^{-s}\tau_k)=(1+(s-1)\beta)/(1+s\beta)=1-\beta/(1+s\beta) \tag{20}$$

which is what we wanted to show.

Continuity of χ on $\tilde{\Omega}_5$ follows as in Section 8. To show that χ extends continuously to all of Ω_5 let $[(x_i)_j,\tau_0,\xi_j]\subset\tilde{\Omega}_5$ be a sequence converging to $[(y_i),\tau_0,\xi]$. Here $[(x_i)_j,\tau_0,\xi_j]$ is the Γ-orbit of the point $((x_i)_j,\tau_0,\xi_j)$.

Let $(\tau_i)_j=\Lambda(x_i)_j$ be the splitting and shifting sequence defined by $(x_i)_j$ and τ_0. Represent ξ_j by a measured geodesic lamination $\nu_j\in\mathcal{V}_0(\eta_0)$. By the definition of the topology on Ω_5, the laminations ν_j converge to a lamination $\nu\in\mathcal{V}_0(\tau_0)$ which represents ξ.

If the active component of the transition modifying (y_i) to $\sigma_4(y_i)$ is the zero component, then the same holds true for the active component of the transition modifying $(x_i)_j$ to $\sigma_4(x_i)_j$ provided that j is sufficiently large. Then $\chi((x_i)_j,\tau_0,\xi_j)\to\chi((y_i),\tau_0,\xi)$ follows exactly as in the proof of Lemma 8.4.

The same reasoning also applies if the active component of the transition modifying (x_i) to $\sigma_4(x_i)$ is the component p for some $p\geq 1$ and if one of the following possibilites is satisfied.

(1) In the extended splitting and shifting sequence $\omega(y_i)$ constructed as above from $\Pi_0(y_i)$ and τ_0, there is some $s>0$ so that the transition modifying $\omega(y_i)_s$ to $\omega(y_i)_{s+1}$ is a splitting and shifting move at a large branch in the simple closed curve c labeled p in τ_0.
(2) The sequence $\Pi_p(x_i)_{i\geq 0}$ is finite.

Now assume that none of the above two possibilities is satisfied. Assume moreover for the moment that for the approximating sequences $(x_i)_j$, one of the two above possibilities holds. Let as before (a_i) be the sequence constructed from $\Pi_p(y_i)$ and let $(a_i)_j$ be the sequence defined by $\Pi_p(x_i)_j$ which is used to calculate the number $k(x_i)_j$.

Let $\ell=\ell(x_i)$ be as before. Using equation (20), for large j the value $\chi[(x_i)_j,\tau_0,\xi_j]$ is close to

$$\sum_{i=2^{\varsigma^{\ell-1}(a_i)_j}}^{2^{\varsigma^{\ell}(a_i)_j}}\frac{\beta}{1+i\beta}. \tag{21}$$

Now for very large ℓ the sum (21) is close to $\log(1+\beta 2^{\zeta^\ell(a_i)_j})-\log(1+\beta 2^{\zeta^{\ell-1}(a_i)_j})$. Since $\ell \to \infty$ as $j \to \infty$, we have

$$|\chi[(x_i)_j, \tau_0, \xi_j] - \zeta(a_i)\log 2| \to 0 \quad (j \to \infty). \tag{22}$$

Thus if we define $\chi[(y_i), \tau_0, \xi] = \zeta(a_i)\log 2$ then by using Lemma A.1, we conclude that this is a continuous extension of χ to Ω_5. □

As in Section 8, the shift on Ω_4 extends to a continuous transformation σ_5 on Ω_5. We have

Lemma 9.7. *There is a number $\omega > 0$ such that $\int \chi d\mu > \omega$ for every σ_5-invariant Borel probability measure on Ω_5 which gives full measure to $\tilde{\Omega}_5$.*

Proof. It suffices to show the lemma for ergodic measures.

By the Birkhoff ergodic theorem, if μ is σ_5-invariant and ergodic and gives full measure to $\tilde{\Omega}_5$ then there is some $[(x_i), \tau_0, \xi] \in \tilde{\Omega}_5$ whose σ_5-orbit is dense in the support of μ and such that $\int \chi d\mu = \lim_{n\to\infty} \frac{1}{n}\sum_{s=0}^{n-1} \chi(\sigma_5^s[(x_i), \tau_0, \xi])$.

Let (x_i) be such a point and let $(\tau_i) = \Lambda(x_i)$ be the splitting and shifting sequence constructed from (x_i) and τ_0. By assumption, $\cap_i \mathcal{V}(\tau_i) \cap \mathcal{V}_0(\tau_0)$ consists of a unique point ν. The Γ-orbit of the sequence (τ_i) defines a point $(z_i) \in \Omega_1$. Since $[(x_i), \tau_0, \xi]$ is recurrent under the shift σ_5, the point (z_i) is recurrent under the shift σ_1.

Let $k_j \to \infty$ be a sequence so that $\sigma_1^{k_j}(z_i)$ is close to (z_i) for all j. Using the above notation, let $F(\tau_i) = \tau_{i+k(x_i)}$ for all i. It now suffices to show that

$$-\frac{1}{k_j}\sum_{j=0}^{k_j-1} \log \mu(F^j(\tau_i)) \geq \delta$$

for a number $\delta > 0$ not depending on μ,

However, this follows from the construction and Rafi's distance formula together with the fact that the image under the map R of a splitting and shifting sequence is a uniform quasi-geodesic in the marking graph. □

An *extended boundary multicomponent* $\mathcal{P}$ of $\mathcal{Q}$ is the union of a boundary multicomponent $\mathcal{P}'$ of $\mathcal{Q}$ and a collection of pairwise disjoint annuli which are disjoint from $\mathcal{P}'$. The *entropy* of an extended boundary multicomponent $\mathcal{P}$ is defined to be the sum of the entropy of $\mathcal{P}'$ and the number of disjoint annuli components of $\mathcal{P}$. We have

Lemma 9.8. (1) *The entropy of an extended boundary multicomponent does not exceed $h-1$.*
(2) *Let μ be a σ_5-invariant Borel probability measure on Ω_5 with $\mu(\tilde{\Omega}_5) = 0$. Then $h_\mu - (h-1)\int \chi d\mu \leq 0$.*

Proof. The first part of the lemma follows as in Lemma 8.8.

To show the second part, note that by Lemma 9.6 and its proof, the restriction of χ to a σ_5-invariant set consisting of points $[(x_i), \tau_0, \xi]$ where each transition modifying (x_i) to $\sigma_5(x_i)$ isf a splitting and shifting move in one of the b Bernoulli shifts (Σ_i, ν_i) is just $\zeta \log 2$ where ζ is the dyadic roof function. By Lemma A.2, the pressure of $-\zeta \log 2$ equals zero or, equivalently, the topological entropy of the suspension flow for $\zeta \log 2$ equals one.

Now the entropy of any σ_5-invariant Borel probability measure which is invariant and ergodic for the suspension flow supported in an extended boundary multicomponent $\mathcal{P}$ of $\mathcal{Q}$ is not bigger than the sum of the supremums of the entropies of all σ_5-invariant Borel probability measures supported in the components of $\mathcal{P}$. The lemma now follows from the entropy calculation done in Lemma 8 together with the above remark. ☐

Now we are ready to show

Proposition 9.9. *For every $\epsilon > 0$ there is a compact set $K \subset \hat{\mathcal{Q}}$ and a number $c > 0$ such that the number of periodic orbits of period at most R which do not intersect K is at most $ce^{h-1+\epsilon}$.*

Proof. We proceed as in the proof of Proposition 9.1. Namely, by Lemma 9.8, for every $\epsilon > 0$ there is an open relative compact set $U \subset \hat{\mathcal{Q}}$ such that the pressure of the restriction of the function χ to the closure of the subset of $\tilde{\Omega}_5$ of all points which define orbits in $\hat{\mathcal{Q}}$ not intersecting U is at most $h - 1 + \epsilon$.

By Lemma 9.2, for every periodic orbit for Φ^t in $\hat{\mathcal{Q}}$ there is a corresponding periodic sequence $[(x_i), \tau_0, \xi] \in \tilde{\Omega}_5$. The closure B of the set of all such images of points which are mapped to orbits of Φ^t avoiding U is a compact σ_5-invariant subset K of Ω_5. Moreover,

$$h_\mu - (h - 1 + \epsilon) \int \chi d\mu \leq 0$$

for every σ_5-invariant Borel probability measure on K.

As in the proof of Proposition 9.1 we can use the natural metric on Ω_5 to define separation for points in K. Using the reasoning in the proof of Proposition 9.1, we are thus left with verifying the following. For every $R > 0$, there is an injection of the set of periodic orbits of Φ^t contained in $\hat{\mathcal{Q}} - U$ into a maximal separated subset of K. However, this follows as in the proof of Proposition 9.1.

The proposition now follows from the argument in the proof of Proposition 9.1.
☐

As a final conclusion we obtain

Corollary 9.10. *The number of periodic orbits for the Teichmüller flow on $\mathcal{Q}$ of period at most R is asymptotic to e^{hR}/hR.*

Proof. The corollary is an immediate consequence of Proposition 9.9 and the main result of [H10c]. □

Appendix A. The dyadic expansion flow

In this appendix we construct a suspension of the Bernoullli shift (Σ, ν) in the letters $0, 1$ which encodes dyadic expansions of positive integers.

Define a function $\zeta : \Sigma \to (0, \infty)$ as follows. If $(x_i) \in \Sigma$ is such that either $x_i = 0$ for all $i < 0$ or $x_i = 0$ for all $i \geq 0$ then define $\zeta(x_i) = 1$.

Otherwise there is a smallest number $i_0 > 0$ and a smallest number $i_1 \geq 0$ with $x_{-i_0} = x_{i_1} = 1$. Let $m > i_0$ be a large number. Consider the string $(y_i)_{0 \leq i \leq m+i_1}$ defined by $y_i = x_{i-m}$, and the string $(z_i)_{0 \leq i \leq m-i_0}$ defined by $z_i = x_{i-m}$. The strings $(y_i), (z_i)$ can be viewed as dyadic expansions of natural numbers

$$\ell = \sum_{i=0}^{m+i_1} y_i 2^i \geq k = \sum_{i=0}^{m-i_0} z_i 2^i \geq 2^{m-i_0}.$$

Define

$$\alpha_m(x_i) = (\log_2(\ell) - \log_2(k))/(i_0 + i_1)$$

and observe that α_m takes values in $[1, \log_2 3]$. Moreover, $\alpha_m(x_i)$ only depends on the finite string $(x_i)_{-m \leq i \leq i_1}$.

We claim that as $m \to \infty$, the values $\alpha_m(x_i)$ converge to a number $\zeta(x_i)$. Namely, using the notation from the previous paragraph, for $u > 0$ we have

$$\alpha_{m+u}(x_i) = (\log_2(2^u \ell + p(u)) - \log_2(2^u k + p(u)))/(i_0 + i_1)$$

for a number $p(u) \in [0, 2^u - 1]$. Now $\ell \geq k \geq 2^{m-i_0}$ and therefore the sequence $\alpha_m(x_i)$ $(m > 0)$ is a Cauchy sequence and hence it indeed converges.

The assignment $(x_i) \to \zeta(x_i)$ defines a function on Σ with values in $[1, \log_2 3]$. We call ζ the *dyadic roof function.*

Lemma A.1. *The dyadic roof function ζ is continuous.*

Proof. Let (x_i) be any sequence in the letters $0, 1$ and let $\epsilon > 0$. Assume first that there are smallest numbers $i_0 > 0$, $i_1 \geq 0$ so that $x_{-i_0} = x_{i_1} = 1$. Choose $m > i_0$ sufficiently large that $|\zeta(y_i) - \alpha_m(x_i)| \leq \epsilon/2$ for all (y_i) which satisfy $y_i = x_i$ for $-m \leq i \leq i_1$. Such a number exists by the discussion preceding this lemma. By construction, we then have $|\zeta(x_i) - \zeta(y_i)| \leq \epsilon$.

Next assume that $x_i = 0$ for all $i \geq 0$ but that there is some smallest number $i_0 > 0$ with $x_{-i_0} = 1$. Let $m \geq i_0$ be a large number. Let $n > 0$ and assume that (y_i) is such that $y_i = 0$ for $-i_0 < i \leq n$. If there is some smallest $i_1 > n$ so that $y_{i_1} = 1$ then

$$\alpha_m(y_i) = (\log_2(2^{m+i_1} + k) - \log_2 k)/(i_1 + i_0)$$

where $2^{m-i_0} \leq k \leq 2^{m-i_0+1}$ and hence $\alpha_m(y_i)$ is close to one if n is sufficiently large. The case that $x_i = 0$ for all $i \leq 0$ follows in the same way. □

Since the function ζ is continuous, uniformly bounded and bounded from below by a positive number we can consider the suspension Z of (Σ, ν) with roof function $\zeta \log 2$. This suspension is a compact space. The suspension flow β^t on Z is continuous. It will be called the *dyadic expansion flow* in the sequel.

Lemma A.2. *The topological entropy of the dyadic expansion flow equals one.*

Proof. By the variational principle and Abramov's formula, the topological entropy of the dyadic expansion flow equals one if the pressure of the function $-\zeta \log 2$ on the Bernoulli shift (Σ, ν) equals zero. To calculate this pressure, let $n > 0$ and let $E_n \subset \Sigma$ be the set of all standard cylinders of length n. Each cylinder $[x_i] \in E_n$ is defined by a sequence (x_i) of length n in the letters $0, 1$, and it is the set of all sequences (y_i) with $y_i = x_i$ for $0 \le i \le n$.

Let $\zeta_n = \sum_{0 \le i \le n-1} \zeta \circ \nu^i$; we have to estimate the sum

$$Z_n(\zeta) = \sum_{[x_i] \in E_n} \max\{e^{-\zeta_n(z) \log 2} \mid z \in [x_i]\} = \sum_{[x_i] \in E_n} \max\{2^{-\zeta_n(z)} \mid z \in [x_i]\}.$$

For this let (x_i) be a sequence of length n which defines the dyadic expansion of the number p. Let $j \le n$ be the largest number so that $x_j = 1$. If no such number exists then put $j = -\infty$. Let $(y_i) \in [x_i]$ be such that $y_i = 1$ for $i < 0$ and $y_i = 0$ for $i > j$. Then $\zeta(\nu^s y_i) = 1 \le \zeta(\nu^s z_i)$ for all $(z_i) \in [x_i]$ and all $s > j$. Moreover, since for any $n > 0$ the function $t \to \log_2(n + t) - \log_2 t$ is decreasing, for every $(z_i) \in [x_i]$ and all $m > 0$, all $s \le j$ we have $\alpha_m(\nu^s z_i) \ge \alpha_m(\nu^s y_i)$ and hence

$$\zeta_n(z_i) \ge \lim_{k \to \infty} (\log_2(2^k p + 2^k - 1) - \log_2(2^k - 1)) + n - j \ge \log_2 p + n - j.$$

Now $\log_2 p \ge j$ and hence we conclude that

$$Z_n(\zeta) \le \sum_{j=0}^{n-1} \sum_{i=2^j}^{2^{j+1}-1} \frac{1}{i} 2^{-(n-j)}.$$

Now the sum $\sum_{i=2^j}^{2^{j+1}-1} \frac{1}{i}$ is bounded independent of j. This implies that indeed we have $\lim_{n \to \infty} \frac{1}{n} \log Z_n(\zeta) = 0$ which is what we wanted to show. □

References

[AGY06] A. Avila, S. Gouëzel, J. C. Yoccoz, *Exponential mixing for the Teichmüller flow*, Publ. Math. Inst. Hautes Études Sci. No. 104 (2006), 143–211.

[AR09] A. Avila, M.J. Resende, *Exponential mixing for the Teichmüller flow in the space of quadratic differentials*, arXiv:0908.1102.

[BL09] C. Boissy, E. Lanneau, *Dynamic and geometry of the Rauzy-Veech induction for quadratic differentials*, Erg. Th. & Dyn. Syst. 29 (2009), 767–816.

[BG07] A. Bufetov, B. Gurevich, *Existence and uniqueness of the measure of maximal entropy for the Teichmüller flow on the moduli space of abelian differentials*, arXiv:math.DS/0703020.

[BS03] J. Buzzi, O. Sarig, *Uniqueness of equilibrium measures for countable Markov shifts and multidimensional piecewise expanding maps*, Erg. Th. & Dyn. Syst. 23 (2003), 1383–1400.

[CEG87] R. Canary, D. Epstein, P. Green, *Notes on notes of Thurston*, in "Analytical and geometric aspects of hyperbolic space", edited by D. Epstein, London Math. Soc. Lecture Notes 111, Cambridge University Press, Cambridge 1987.

[EM11] A. Eskin, M. Mirzakhani, *Counting closed geodesics in moduli space*, J. Mod. Dynamics 5 (2011), 71–105.

[EMR11] A. Eskin, M. Mirzakhani, K. Rafi, *Counting closed geodesics in strata*, to appear.

[H06] U. Hamenstädt, *Train tracks and the Gromov boundary of the complex of curves*, in "Spaces of Kleinian groups" (Y. Minsky, M. Sakuma, C. Series, eds.), London Math. Soc. Lec. Notes 329 (2006), 187–207.

[H09a] U. Hamenstädt, *Geometry of the mapping class groups I: Boundary amenability*, Invent. Math. 175 (2009), 545–609.

[H09b] U. Hamenstädt, *Invariant Radon measures on measured lamination space*, Invent. Math. 176 (2009), 223–273.

[H09c] U. Hamenstädt, *Geometry of the mapping class groups II: A biautomatic structure*, arXiv:0911349.

[H10a] U. Hamenstädt, *Dynamics of the Teichmüller flow on compact invariant sets*, J. Mod. Dynamics 4 (2010), 393–418.

[H10b] U. Hamenstädt, *Stability of quasi-geodesics in Teichmüller space*, Geom. Dedicata 146 (2010), 101–116.

[H10c] U. Hamenstädt, Bowen's construction for the Teichmüller flow, arXiv:1007.2289.

[H12] U. Hamenstädt, *Periodic orbits in the thin part of strata: Lower bounds*, in preparation.

[HK95] B. Hasselblatt, A. Katok, *Introduction to the modern theory of dynamical systems*, Cambridge University Press,Cambridge 1995.

[KZ03] M. Kontsevich, A. Zorich, *Connected components of the moduli spaces of Abelian differentials with prescribed singularities*, Invent. Math. 153 (2003), 631–678.

[L08] E. Lanneau, *Connected components of the strata of the moduli spaces of quadratic differentials*, Ann. Sci. École Norm. Supér. 41 (2008), 1–56.

[L83] G. Levitt, *Foliations and laminations on hyperbolic surfaces*, Topology 22 (1983), 119–135.

[Mn87] R. Mañe, *Ergodic theory and differentiable dynamics*, Springer Ergebnisse der Mathematik 8, Springer Berlin Heidelberg 1987.

[M80] H. Masur, *Uniquely ergodic quadratic differentials*, Comm. Math. Helv. 55 (1980), 255–266.

[M82] H. Masur, *Interval exchange transformations and measured foliations*, Ann. Math. 115 (1982), 169-201.

[M93] H. Masur, *Logarithmic law for geodesics in moduli space*, in "Mapping class groups and moduli spaces of Riemann surfaces", 229–245, Contemp. Math. 150, AMS, Providence 1993.

[MM99] H. Masur, Y. Minsky, *Geometry of the complex of curves I: Hyperbolicity*, Invent. Math. 138 (1999), 103-149.

[MM00] H. Masur, Y. Minsky, *Geometry of the complex of curves II: Hierarchical structure*, GAFA 10 (2000), 902-974.

[MMS10] H. Masur, L. Mosher, S. Schleimer, *On train track splitting sequences*, arXiv:1004.4564.

[M96] Y. Minsky, *Extremal length estimates and product regions in Teichmüller space*, Duke Math. J. 83 (1996), 249–286.

[MS93] H. Masur, J. Smillie, *Quadratic differentials with prescribed singularities and pseudo-Anosov diffeomorphisms*, Comment. Math. Helv. 68 (1993), 289–307.

[Mo03] L. Mosher, *Train track expansions of measured foliations*, unpublished manuscript.

[PP87] A. Papadopoulos, R. Penner, *A characterization of pseudo-Anosov foliations*, Pacific J. Math. 130 (1987), 359–377.

[PH92] R. Penner with J. Harer, *Combinatorics of train tracks*, Ann. Math. Studies 125, Princeton University Press, Princeton 1992.

[R07] K. Rafi, *A combinatorial model for the Teichmüller metric*, Geom. Funct. Anal. 17 (2007), 936–959.

[S99] O. Sarig, *Thermodynamic formalism for countable Markov shifts*, Erg. Th. & Dyn. Syst. 19 (1999), 1565–1593.

[S03] O. Sarig, *Existence of Gibbs measures for countable Markov shifts*, Proc. AMS 131 (2003), 1751–1758.

[T79] W. Thurston, *Three-dimensional geometry and topology*, unpublished manuscript.

[V82] W. Veech, *Gauss measures for transformations on the space of interval exchange maps*, Ann. Math. 115 (1982), 201–242.

[V86] W. Veech, *The Teichmüller geodesic flow*, Ann. Math. 124 (1986), 441–530.

[W75] P. Walters, *A variational principle for the pressure of continuous transformations*, Amer. J. Math. 97 (1975), 937–971.

[W82] P. Walters, *An introduction in ergodic theory*, Springer Graduate Texts in Mathematics 79, Springer New York Berlin 1982.

[Z05] R. Zweimüller, *Invariant measures for general(ized) induced transformations*, Proc. AMS 133 (2005), 2283–2295.

MATHEMATISCHES INSTITUT DER UNIVERSITÄT BONN
ENDENICHER ALLEE 60
D-53115 BONN, GERMANY
e-mail: ursula@math.uni-bonn.de