

MIRZAKHANI'S WORK ON EARTHQUAKE FLOW

ALEX WRIGHT

ABSTRACT. The Teichmüller unipotent flow can be defined concretely on certain moduli spaces of singular flat surfaces by shearing polygonal presentations of the surfaces. Thurston's earthquake flow on moduli spaces of hyperbolic surfaces is more mysterious. Both flows have deep and important connections to other areas of mathematics.

In this expository survey we give a geometric account of the main ideas behind Mirzakhani's theorem relating these two flows. Our presentation avoids some technical prerequisites that featured in the original more analytic presentation.

CONTENTS

1. Introduction	1
2. Preliminaries	5
3. Horocyclic foliations	14
4. The Fundamental Lemma on Earthquakes	22
5. Mirzakhani's isomorphism	24
6. Invariant measures	30
7. Laminations containing a pants decomposition	32
8. Hamiltonian flows	34
9. The linear structure on \mathcal{ML}_α	35
10. Other results on earthquakes	35
References	36

1. INTRODUCTION

Flat geometry. One entry point to Mirzakhani's work is the $SL(2, \mathbb{R})$ action on moduli spaces of Abelian or quadratic differentials. Due to its close connection to Teichmüller theory, applications to topics such as rational billiards, and deep analogy to Lie group actions on homogeneous spaces, this action has been extensively studied; see for example the author's short survey [Wri16] for an introduction.

Part of the appeal of this action is that it is easily defined. Abelian and quadratic differentials can be presented via polygons in \mathbb{R}^2 , with edges identified in pairs to define a surface with a singular flat metric. The $SL(2, \mathbb{R})$ action arises quite simply from the usual $SL(2, \mathbb{R})$ action on \mathbb{R}^2 , which linearly transforms one polygonal presentation into another.

Hyperbolic geometry. The other entry point is the earthquake flow defined on the bundle of measured laminations over a moduli space of hyperbolic surfaces. This flow is only easy to visualize when the measured lamination is a closed curve, in which case the result of the earthquake is obtained by cutting the geodesic representative of the curve and re-gluing with a twist. More generally, the measured lamination represents the “fault lines” along which a “sliding” or “shearing” deformation occurs. When the lamination has a fractal structure, infinitesimal shearing along each individual “fault line” can combine to a definite change to the surface.

Although the definition of earthquake flow is less elementary, its applications and connections are far reaching. Early in its history it was used prominently by Kerckhoff to resolve the longstanding Nielsen Realization Problem about mapping class groups [Ker83]. More recently, Mirzakhani and others used it to prove equidistribution results in the moduli space of hyperbolic surfaces [Mir07, AH21], which are a crucial tool for example in one of Mirzakhani’s results on counting geodesics on individual surfaces [Mir16].

Mirzakhani’s bridge between hyperbolic and flat geometry. Although defined on different spaces and featuring different notions of geometry, the two flows have similarities. Both involve some notion of shearing, and both have analogies to unipotent flows on homogeneous spaces. Both have strong non-divergence properties, established simultaneously in a single paper by Minsky and Weiss [MW02]. As discussed in Section 8, both are Hamiltonian with respect to natural functions and natural symplectic forms.

The purpose of this survey is to give an expository introduction to Mirzakhani’s discovery that, in an abstract measurable sense, both flows are the actually the same. This is the main result of [Mir08].

Theorem 1.1. *There is a measurable mapping class group equivariant conjugacy F between the earthquake flow $(\lambda, X) \mapsto (\lambda, E_{t\lambda}(X))$ on the bundle $\mathcal{ML} \times \mathcal{T}_g$ of measured laminations over Teichmüller space \mathcal{T}_g*

and the Teichmüller unipotent flow action of

$$u_t = \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}$$

on the bundle QD of nonzero quadratic differentials over Teichmüller space.

That F is a conjugacy means that the following diagram commutes.

$$\begin{array}{ccc} \mathcal{ML} \times \mathcal{T}_g & \xrightarrow{E_t} & \mathcal{ML} \times \mathcal{T}_g \\ F \downarrow & & \downarrow F \\ QD & \xrightarrow{u_t} & QD \end{array}$$

We will later discuss the natural Lebesgue class measure on $\mathcal{ML} \times \mathcal{T}_g$ and see that it is the pull back of Masur-Veech measure, but for the moment it suffices to understand that F is Borel-measurable but not continuous.

Significance and applications. Theorem 1.1 builds a surprising bridge between the more mysterious world of earthquake flow and the comparatively well understood Teichmüller unipotent flow. Perhaps the most important consequence is the following.

Corollary 1.2. *Earthquake flow is ergodic.*

Proof. It is well known that Teichmüller unipotent flow is ergodic: This follows from the Howe-Moore Theorem and the ergodicity of Teichmüller geodesic flow. (See for example the textbook [BM00, Section III] for an introduction to the Howe-Moore Theorem, which applies here because both the Teichmüller geodesic and unipotent flows are part of the $SL(2, \mathbb{R})$ action. See for example the survey [FM14, Section 4] for the ergodicity of Teichmüller geodesic flow, which was originally proved independently by Masur and Veech.) \square

This ergodicity is a key ingredient in the equidistribution applications mentioned above.

As McMullen describes in his laudation for Mirzakhani's field medal, Theorem 1.1 can be viewed not only as a bridge between flat and hyperbolic geometry, but also a bridge across the "holomorphic/symplectic" divide [McM]. Indeed, quadratic differentials live primarily in the holomorphic world, and since the only concise definition of earthquake flow is as a Hamiltonian flow it is natural to think of earthquakes as living more in the symplectic world.

Origin and intended audience. This survey is aimed at the many mathematicians working in dynamics or low dimensional topology who

may have heard the statement of Theorem 1.1 but are completely unfamiliar with its proof. We include a limited introduction to some of the necessary background material, in hopes that this will make our exposition readable to a junior graduate student who is able to get some assistance from their advisor or from other sources. The reader should have some familiarity with quadratic differentials, but we have included enough exposition on earthquake flow that the reader who is willing to take some things on faith should be able to gain a meaningful level of understanding of the proof even with no previous exposure to earthquake flow.

Our goal is to present the main proof in the most elementary and geometric way possible. Only after accomplishing that will we proceed to discuss the more sophisticated results that give additional understanding and perspective. Most of the background material we present is due to people other than Mirzakhani, especially Thurston and Bonahon. However all the material is chosen to allow us to prove and appreciate Mirzakhani's result.

Our presentation of the proof of Theorem 1.1 is not exactly the same as Mirzakhani's, in that here we do not make use of transverse cocycles or the symplectomorphism of Bonahon-Sozen [BS01]. See Remark 5.4.

These notes were originally written to accompany lectures at the 2018 summer school on Teichmüller dynamics, mapping class groups and applications at Grenoble, as well as lectures at the 2018 summer school on Teichmüller Theory and its Connections to Geometry, Topology and Dynamics at the Fields Institute in Toronto.

Updates. Since this survey was written in 2018, the author also wrote a much comprehensive survey on Mirzakhani's work aimed at a broader audience [Wri20]. Additionally, this survey has motivated further work, and the questions and conjectures proposed in Remarks 5.6 and 5.7 have now been completely resolved [AHW, CF21].

Acknowledgments. I am happy to thank Francisco Arana-Herrera, Francis Bonahon, Aaron Calderon, Steve Kerckhoff, Jeremy Kahn, and Kasra Rafi for helpful conversations. I am also grateful to thank Francisco Arana-Herrera, Dat Nguyen, Weston Ungemach, and Adva Wolf for attending and offering very helpful feedback on a test run of the lectures at Stanford the week before Grenoble, and to Yueqiao Wu for pointing out some corrections.

Some of the figures were created using [Pie]. I thank Yen Duong, Aaron Fenyes, Subhojoy Gupta, Bruno Martelli, Athanase Papadopoulos, Guillaume Théret, and Mike Wolf for permission to reproduce figures from other sources.

2. PRELIMINARIES

In this section we offer a sketchy introduction to the main objects in Theorem 1.1. It is included mainly for students coming from flat geometry, who may have seen measured foliations but may benefit from an overview of their relation to laminations and an introduction to earthquakes. Experts should certainly skip this section.

There are many sources for this material, although we have not found one that presents all the material we need in a way which is both rigorous and quickly understood. In addition to the references cited below, learners might be interested in consulting books like [Cal07, FM12, FLP12, PH92, Thu97].

Foliations and laminations. In these notes we consider only closed surfaces of genus g at least 2. A measured foliation is a foliation with finitely many prong type singularities, with a transverse measure. This measure assigns a non-negative number to each transverse arc in such a way that arcs isotopic through transverse arcs with endpoints on the same leaves have the same measure. A saddle connection of a measured foliation is an arc of the foliation joining two singularities.

Measured foliations are typically considered to be equivalent if they differ via isotopy and Whitehead moves, which are moves that collapse saddle connections to split a higher order prong singularity into lower order singularities joined by a saddle connection, as in Figure 2.1.




FIGURE 2.1. Whitehead moves. Picture from [GW17].

Remark 2.1. The typical measured foliation has only 3-pronged singularities and no saddle connections, and hence does not admit any Whitehead moves.

The space of measured foliations up to Whitehead moves and isotopy is denoted \mathcal{MF} .

Remark 2.2. A celebrated result of Thurston is that \mathcal{MF} is homeomorphic to \mathbb{R}^{6g-6} . We will not make use of this fact.

A measured geodesic lamination is a closed subset of a hyperbolic surface foliated by non-intersecting geodesics with a transverse measure

of full support. See [Kap01, Section 11.6] and [Mir08, Section 8.3] for the basic properties of measured geodesic laminations.

Remark 2.3. A closed multi-curve is an example of a measured geodesic lamination. However if you take a “typical” geodesic lamination and intersect it with a transverse arc, you will get a Cantor set with a non-atomic measure.

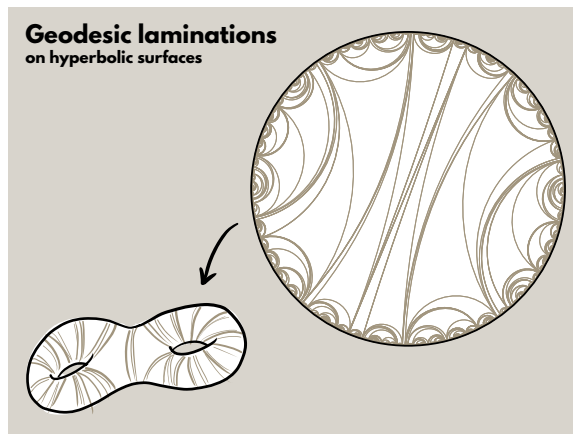


FIGURE 2.2. A geodesic lamination. Picture from [Duo], created by Aaron Fenyes. A similar figure appears in [Fena, Fenb].

If λ is a geodesic lamination on X , the connected components of $X \setminus \lambda$ are called the complementary regions. There are finitely many. Each is bounded by geodesics. The complementary regions can be ideal polygons, and can also be surfaces with genus that are bounded by closed geodesics and/or “crowns” of geodesics meeting in cusps.




FIGURE 2.3. A possibly complementary region bounded by a closed geodesic and a “crown”. Picture adapted from [Gup].

The universal cover of X can be identified with the hyperbolic plane \mathbb{H} . Geodesics in \mathbb{H} correspond to unordered pairs of distinct points on the circle S^1 at infinity. Given two different points X, Y of Teichmüller

space, one obtains an isotopy class of maps from X to Y . Lift such a map to obtain a map from $\mathbb{H} = \tilde{X}$ to $\mathbb{H} = \tilde{Y}$.

One can show that any such map from X to Y is a quasi-isometry, so the lifted map from $\mathbb{H} = \tilde{X}$ to $\mathbb{H} = \tilde{Y}$ extends to a homeomorphism between the circles at infinity. Hence geodesics on X are in correspondence with geodesics on Y , by considering the endpoints of the geodesic on S^1 . In this way a measured geodesic lamination for one hyperbolic metric uniquely determines one for any other hyperbolic metric, and we can think of measured geodesic laminations as topological rather than metric objects. Denote the set of all measured geodesic laminations by \mathcal{ML} .

We define a line of a measured foliation to be either a leaf not passing through a singularity, or any leaf that is a limit of non-singular leaves. Note that if a line passes through a singularity, it enters and exits the singularity on adjacent prongs. (Those inclined to think about quadratic differentials can think of this as having angle π at every singularity.)

Lemma 2.4. *Every line of a measured foliation also determines a pair of distinct points in S^1 .*

Cartoon of the proof. Consider a simple closed curve that the leaf passes through infinitely many times but with no unnecessary intersections that could be removed by an isotopy. The leaf gets “cornered” by lifts of the simple curve to smaller and smaller regions of \mathbb{H} , as seen from a fixed basepoint, forcing the leaf to converge to the intersection of these half-spaces, as in Figure 2.4. \square




FIGURE 2.4. The proof of Lemma 2.4. Picture from [Mar, Figure 8.12].

Remark 2.5. It is not so easy to prove the desired simple closed curve exists. One possibility is to use a “normal form” for the foliation [FLP12, Section 6.4].

Remark 2.6. If the measured foliation is known to arise from a quadratic differential, one can alternatively use the fact that the hyperbolic and flat metrics are quasi-isometric [Kap01, Section 5.3].

By replacing each line in a measured foliation by the geodesic with the same endpoints, one obtains an associated measured lamination [Kap01, p. 251]. This procedure “tightens” each leaf to a geodesic.




FIGURE 2.5. Each three pronged singularity that does not lie on any saddle connections gives an ideal triangle in the associated “tightened” lamination.

Theorem 2.7. *The tightening map $\mathcal{MF} \rightarrow \mathcal{ML}$ is a homeomorphism.*

I do not know any short proof of this result, but a good reference on the tightening map and related topics is [Lev83]. One approach is to build an inverse map using train tracks, but this involves not only showing that every lamination is carried by a train track, but also that the measured foliations constructed using different choices of train track differ by Whitehead moves. See for example [Lei12, Section 5] for an expository account.

A measured lamination is called maximal if its complementary regions are all ideal triangles. The value of the inverse map $\mathcal{ML} \rightarrow \mathcal{MF}$ on a maximal measured lamination can be visualized by a “collapsing” procedure. This procedure is perhaps initially somewhat mind bending, but nonetheless we briefly give the flavor of how it is performed. One “collapses” or “pinches” each triangle onto a “skeleton” consisting of three lines, one going towards each cusp of the triangle, meeting in a central point. This eliminates all the “empty space” not covered by the lamination, and the result is a foliation.

Remark 2.8. One should compare this collapse map to the map $x \mapsto \int_0^x d\mu$ for a measure μ on a Cantor set in \mathbb{R} . This collapses all the intervals not included in the Cantor set.

In this way we see that each complementary triangle corresponds to a three-pronged singularity. This can be easily extended to the case

when all complementary regions are ideal polygons. One could try to extend it further to the case where the complementary regions have genus by extending the lamination to a maximal lamination, but then one has to show that the measured foliations resulting from different extensions differ by Whitehead moves, and this is not obvious. The case when the measure gives positive measure to some closed curves requires a special argument, since in this case the foliation can't be obtained by collapsing.

Intersection number. There is a continuous intersection number function

$$i : \mathcal{ML} \times \mathcal{ML} \rightarrow \mathbb{R}_{\geq 0},$$

which, restricted to weighted simple multi-curves, is the linear extension of the usual geometric intersection number. One can show that weighted simple multi-curves are dense in \mathcal{ML} , so this uniquely determines i , however there are also easy direct definitions [Bon88]. Since $\mathcal{MF} \simeq \mathcal{ML}$, we also get an intersection number on \mathcal{MF} .

If $\alpha \in \mathcal{MF}$ and β is a simple curve, $i(\alpha, \beta)$ is the inf over all ways of realizing β as a sequence of arcs transverse to α of the sum of the transverse α measures of these arcs. This can be extended linearly to the case of β a simple weighted multi-curve, and by continuity to any $\beta \in \mathcal{ML}$.

Remark 2.9. The topology on \mathcal{MF} and \mathcal{ML} is the weakest topology for which the function $\lambda \mapsto i(\lambda, \gamma)$ is continuous for each simple closed curve γ .

Quadratic differentials. Define $\Delta \subset \mathcal{MF} \times \mathcal{MF}$ by

$$\Delta = \{(\alpha, \beta) : i(\alpha, \gamma) = 0 = i(\beta, \gamma) \text{ for some } \gamma \in \mathcal{MF}\}.$$

Note that Δ contains the diagonal $\{(\alpha, \alpha)\}$ (just take $\gamma = \alpha$), so we can think of Δ as a “generalized” or “fat” diagonal.

A quadratic differential q determines two measured foliations, namely the horizontal one $h(q)$ and the vertical one $v(q)$.

Lemma 2.10. *For any q , $(h(q), v(q)) \notin \Delta$.*

Proof sketch. Otherwise take a sequence of weighted simple curves γ_i converging to the γ showing that $(h(q), v(q)) \in \Delta$. Since $i(\gamma_i, h(q)) \rightarrow 0$, there is a sequence of saddle connections representing γ_i whose sum of absolute values of x -components is tiny compared to the total length. (If the holonomy of a saddle connection is $x + iy$, we call x the “ x -component” and y the “ y -component”.) Using the corresponding statement with horizontal and vertical switched, we can obtain a

contradiction. (If a curve is represented by a sum of saddle connections whose sum of x -components is less than C , then the same is true of the flat geodesic representative. As you “tighten” to get the flat geodesic representative, the x and y coordinates don’t get bigger. In fact, both components get monotonically smaller. To make this precise, you need to define an appropriate tightening procedure. Alternatively, see [Lei12, Proposition 3.9] for a proof that the flat geodesic gives intersection number with the horizontal and vertical foliations.) \square

Hence we obtain a map from the bundle QD of non-zero quadratic differentials over Teichmüller space to $\mathcal{MF} \times \mathcal{MF} \setminus \Delta$ given by $q \mapsto (h(q), v(q))$.

Remark 2.11. The intersection number $i(h(q), v(q))$ is the area of q .

Theorem 2.12. *The map $q \mapsto (h(q), v(q))$ determines a homeomorphism $QD \rightarrow \mathcal{MF} \times \mathcal{MF} \setminus \Delta$.*

Proof sketch. One can create an inverse map as follows. Given

$$(h, v) \in \mathcal{MF} \times \mathcal{MF} \setminus \Delta,$$

tighten each h, v to geodesic laminations, also denoted h, v . Since $(h, v) \notin \Delta$, we have that h and v do not share any leaves, and that each complementary region of $h \cup v$ is a compact polygon:

- Indeed, if h and v shared a leaf, then a weak star limit of the Lebesgue measure supported larger and larger segments of this leaf would give a measured foliation γ with $i(h, \gamma) = 0 = i(v, \gamma)$.
- If there is a complementary region that isn’t a polygon, one could pick a simple curve γ in that region.
- To see that the polygons are compact, i.e. that none of the vertices are at infinity, requires a bit of extra argument again using weak star limits.

Collapsing all the connected components of the complement of $h \cup v$, as well as all connected components of $h \setminus v$ and $v \setminus h$, defines a quadratic differential by picking local coordinates z for which $\operatorname{Re}(z)$ and $\operatorname{Im}(z)$ locally coincide with the two foliations. Each component gets collapsed to a single point.

If this collapsing seems too drastic, one should ponder maps from rectangles on the surface bounded by arcs of the lamination to rectangles in \mathbb{R}^2 , defined as follows: One considers arcs (or isotopy classes of arcs rel endpoints) from a designated corner to a point in the rectangle, and take the intersection numbers with the two foliations to get the two coordinates. This map accomplishes the desired collapsing. For more details, see [CB88, Proof of Lemma 6.2]. \square

Theorem 2.12 is discussed from different points of view in [GM91, Section 3] and [Pap86, Section 2].

Earthquakes on surfaces. Consider a simple closed curve α on an oriented hyperbolic surface X . The right earthquake for time t about α is the surface $E_{t\alpha}(X)$ obtained by cutting X along the geodesic representative of α and regluing with a twist to the right by hyperbolic distance t . The notion of “right” just depends on the orientation of X , and doesn't require any orientation of α : Two ants facing each other across the curve α will each see the other move to the right.

Remark 2.13. $E_{t\alpha}$ determines a flow on Teichmüller space. Since we are making a continuous change to the metric, the marking can be transported along the earthquake path. After continuously earthquaking from $t = 0$ to $t = \ell(\alpha)$ (the length of α), one arrives back at the same hyperbolic metric, but with a new marking that differs from the old marking by a Dehn twist.

Remark 2.14. In appropriate Fenchel-Nielsen coordinates, $E_{t\alpha}$ is a translation.

One can similarly define earthquakes for any simple weighted curve α , where the amount of the twist in each curve depends on the measure of a transverse arc. One then defines the earthquake in an arbitrary $\alpha \in \mathcal{ML}$ to be the limit of earthquakes in simple weighted curves α_n that converge to α ,

$$E_\alpha(X) = \lim_{n \rightarrow \infty} E_{\alpha_n}(X).$$

We will sketch a proof that this is well-defined, i.e. that the limit doesn't depend on the sequence α_n of weighted multi-curves converging to α . Our discussion will take in the universal cover.

Earthquakes on the hyperbolic plane. One can define a measured lamination on the hyperbolic plane \mathbb{H} in the same way as on a surface. However, \mathbb{H} of course doesn't have any closed curves, so instead we make the following definition: A discrete measured lamination on \mathbb{H} is a union of disjoint geodesics such that each compact set in \mathbb{H} intersects only finitely many of the geodesics, and such that each geodesic is endowed with a positive weight (the transverse measure of a small arc crossing only that geodesic). A key motivating example of a discrete measured lamination on \mathbb{H} is the preimage of a simple weighted curve on a closed surface.

Let λ be a discrete measured lamination on \mathbb{H} , and let $w_0 \in T^1\mathbb{H}$ be a fixed unit tangent vector. The choice of w_0 won't be important and so is usually suppressed from the notation. The earthquake in λ

is defined to be the unique discontinuous map $E_\lambda : \mathbb{H} \rightarrow \mathbb{H}$ which is a local isometry off λ , fixes w_0 , and “shears” along each leaf of λ by the measure of that leaf.

The last part of that definition is rather vague, so we’ll give a more explicit inductive construction of the earthquake. First, cut out λ , to get countably many regions of \mathbb{H} , each bounded by geodesics. Define the map to be the identity on the region containing w_0 . Now, having already defined the earthquake on a region R , we explain how to define it for a region R' that borders R along a geodesic γ in the support of λ . Let $m_\gamma > 0$ be the weight of γ in λ . Let I_R be the isometry of \mathbb{H} taking R to $E_\lambda(R)$. Let T be the isometry of \mathbb{H} that translates along $I_R(\gamma)$ by m_γ . Then we define

$$E_\lambda(R') = (T \circ I_R)(R').$$

Thus, as in the case of closed surfaces, we arrange for two ants staring at each other from across different sides of γ to each see the other move to the right by a distance equal to the weight m_γ .

The only reason we require the earthquake to fix a unit tangent vector is so that it is well defined, avoiding the worry that it might only be well defined up to certain isometries of \mathbb{H} .

Given any lamination in the hyperbolic plane, we define the earthquake in this lamination to be the limit of earthquakes in discrete laminations which approximate the given lamination. This should again be a discontinuous map $\mathbb{H} \rightarrow \mathbb{H}$ that is an isometry off the lamination, but as in the surface case we need to show it is well defined and doesn’t depend on the choice of discrete approximates.

The main estimate. We will now sketch a proof that earthquakes are well-defined on \mathbb{H} , following the more detailed treatment in [Ker83, Section II]. Some readers might choose to skip this.

We require two estimates, which refer to the $PSL(2, \mathbb{R})$ invariant metric on the unit tangent bundle $T^1\mathbb{H}$. We use E_{tv} to refer to the time t earthquake in the geodesic through a unit tangent vector v . (A single geodesic with weight 1 is an example of a discrete lamination, so the definition above applies in this case.)

Lemma 2.15. *For all $D, T > 0$ there exists $K = K(D, T)$ such that for all $v, v', w \in T^1\mathbb{H}$ that are pairwise distance at most D apart, and all $t \leq T$, we have*

$$d(E_{tv}(w), w) \leq Kt$$

and

$$d(E_{tv}(w), E_{tv'}(w)) \leq Ktd(v, v').$$




FIGURE 2.6. The proof that earthquakes are well-defined.

Both estimates are extremely soft, and use only the fact that a differentiable function on a compact set is Lipschitz [Ker83, Lemma 1.2].

Consider two unit tangent vectors w_0, w in $T^1\mathbb{H}$ that do not lie on the lamination. We consider two discrete measured laminations λ, λ' that both approximate the given measured lamination. *We need to show that the earthquakes corresponding to λ and λ' that fix w_0 do almost the same thing to w .* If we can do this, up to small details we will have shown that earthquakes are well defined on \mathbb{H} .

For each discrete approximation (λ or λ'), there is a finite, totally ordered set of geodesics separating the basepoints of w_0 and w . Only these geodesics and their measures are relevant to understanding the effect of the earthquake on w .

Consider the geodesic arc from the basepoint of w_0 to the basepoint of w . Consider also the finite sequence v_i (respectively v'_i) of unit tangent vectors based at the intersections of λ (respectively λ') with this arc that point along λ (respectively λ'). Let m_i (respectively m'_i) be the measure (weight) of the geodesic of λ (respectively λ') along which v_i (respectively v'_i) points.

We now qualitatively outline a quantitative argument in [Ker83]. Let us divide the arc into small chunks (subintervals). Using the first estimate above, we can reduce to the case that the two discrete measures give exactly the same mass to each chunk. This is because they must give almost the same mass to each chunk, and getting rid




FIGURE 2.7. The proof that earthquakes are well-defined. The v_i are in red, and the v'_i in blue. On each chunk of the arc from w_0 to w , they are replaced with a single unit tangent vector r_k shown in purple.

of a tiny bit of mass won't change the effect of the earthquake on w much.

On each chunk, we replace all the v_i in that chunk with a single vector close to all of them, and we earthquake in that vector with the corresponding amount of mass. The estimates above show that the collective effect of all of these changes is small, so this gives that the difference of the two earthquakes applied to w is small. Hence, the earthquakes in λ and λ' do almost the same thing to w . So earthquakes are well defined.

Equivariance (back to surfaces). If the measured lamination is invariant by a Fuchsian group, the earthquake map will be equivariant by a representation of this Fuchsian group, whose image will be a new Fuchsian group. This new Fuchsian group can be seen as the earthquake of the first Fuchsian group.

We end by being more explicit about how to get the Fuchsian group representing $E_\lambda(\mathbb{H}/\Gamma)$ from Γ . Pick a $w_0 \in T_1\mathbb{H}$ not on $\tilde{\lambda}$. For each $\gamma \in \Gamma$, we consider the earthquake that fixes w_0 , and pick $\rho(\gamma)$ such that the image of $\gamma(w_0)$ under this earthquake is $\rho(\gamma)w_0$. This is easily seen to be a homomorphism, and we define $E_\lambda(\mathbb{H}/\Gamma) = \mathbb{H}/\rho(\Gamma)$. The homomorphism ρ directly defines a marking on $\mathbb{H}/\rho(\Gamma)$ from a marking on \mathbb{H}/Γ , so we get that earthquakes are well-defined on Teichmüller space.

3. HOROCYCLIC FOLIATIONS

A very important construction, which Thurston introduced in [Thu], explains how, given a hyperbolic surface X and a certain lamination λ , we can construct a measured foliation on X . Here λ should be maximal, i.e., all the complementary regions should be triangles. (Some people

call this “complete” instead of “maximal”). But λ need not support a measure.

The construction begins by defining the foliation on each complementary triangle using horocycles based at each ideal vertex. This gives a foliation of the triangle minus a piece in the center, which can be collapsed to become a three pronged singularity without affecting the foliation along the edges of the triangle. The foliation naturally




FIGURE 3.1. A picture from [PT08] of the horocyclic foliation of a triangle.

carries a transverse measure in which the measure assigned to the set of leaves passing through a segment of an edge of the triangle is the length of that segment.




FIGURE 3.2. A picture from [Mar] of the horocyclic foliation of a triangle.

In this way we can foliate most of X , but the foliation is not yet defined on the vast majority of leaves of λ which do not bound complementary regions. However, the partial foliation defined thus far can

be checked to be Lipschitz, and hence the associated line field extends continuously to a line field that can be integrated because it is Lipschitz. (One can work with vector fields if desired instead of line fields, for example by working locally.) This gives a map

$$F_\lambda : \mathcal{T}_g \rightarrow \mathcal{MF}(\lambda),$$

where $\mathcal{MF}(\lambda)$ is the set of measured foliations μ transverse to λ , i.e. for which $(\lambda, \mu) \notin \Delta$. (Because λ, μ are literally transverse, there is an associated quadratic differential whose horizontal and vertical foliations are λ and μ , and this implies that $(\lambda, \mu) \notin \Delta$ as discussed above. But one could also just define Δ to be the set of pairs not associated to a quadratic differential.)

Our goal in this section is to sketch a proof the following result of Thurston, compare to [Thu, Proposition 4.1]. You can choose to accept this theorem as a black box and skip to the next section now.

Theorem 3.1. *F_λ is a homeomorphism*

Often this homeomorphism is followed with a certain map $\mathcal{MF}(\lambda) \hookrightarrow \mathbb{R}^{6g-6}$ and the result is called shear coordinates for Teichmüller space [Bon96], however we may refer to F_λ itself as shear coordinates.

To prove Theorem 3.1, we will explicitly build the inverse of F_λ . We will build explicitly a hyperbolic surface X whose horocyclic foliation is μ , for any $\mu \in \mathcal{MF}(\lambda)$.

Imagine we already had such an X with $\mu = F_\lambda(X)$. Then we can lift λ to $\tilde{\lambda} \subset \mathbb{H}$. If $X = \mathbb{H}/\Gamma$, then $\tilde{\lambda}$ is invariant under Γ . The idea of the proof is to construct $\tilde{\lambda}$ just from the data of μ .

To do this it helps to better understand $\tilde{\lambda}$, assuming $\mu = F_\lambda(X)$. It is this understanding that will allow us to define $\tilde{\lambda}$ in the case when μ is arbitrary. Let $\tilde{\mu}$ denote the preimage of μ in \mathbb{H} .

Consider two triangles T_1 and T_2 that are complementary regions for $\tilde{\lambda}$. Suppose there is a segment A of $\tilde{\mu}$ that goes from an edge of T_1 to an edge of T_2 , as in Figure 3.3. Consider unit vectors v_1 and v_2 that are based at the start and end points of A and are tangent to the edges of the triangles. We want to compute the Möbius transformation S , which we view as a two-by-two matrix, that maps v_1 to v_2 .

This Möbius transformation, together with the “shear”, allows us to recover the position of T_2 relative to T_1 . That is, there is a one parameter family of locations for a triangle T_2 with an edge generated by v_2 , and we call this parameter the shear. The shear can be determined by comparing the distances from the singular leaf in each of the two triangles, as in Figure 3.4.




FIGURE 3.3




FIGURE 3.4. The position of T_2 relative to T_1 : Knowing an edge of T_2 gives a one parameter family of possibilities for T_2 . To determine T_2 , we also need to use the shear, which is the signed length of the red segment. The two half rays in black are orthogeodesics.

Let I be the set of triangles in \mathbb{H} that are crossed by the segment A . Note that I is a countable totally ordered set, but the order is not a well-order. For each $i \in I$, define v_i^+ and v_i^- to be the vectors tangent to the edges of the corresponding triangle at the intersection of the edges and A , as in Figure 3.5.




FIGURE 3.5. The definition of v_i^+ and v_i^- .




FIGURE 3.6. Thurston's illustration from [Thu] of how A crosses $\tilde{\lambda}$. Its intersection with each triangle corresponds to either a stable or unstable horocycle, according to whether the third side of the triangle is to the left or to the right of A .

Let S_i be the Möbius transformation taking v_i^- to v_i^+ . We now wish to show that

$$S = \prod_{i \in I} S_i.$$

That is, the Möbius transformation moving the vector across infinitely many triangles is the product of the Möbius transformation moving the vector across each of these triangles. We need a definition to even make sense of what such an infinite product should mean.

Definition 3.2. Let I be a countable totally ordered set, and let $S_i, i \in I$ be elements of a fixed Banach algebra. Then we say that $\prod_{i \in I} S_i$ is

well-defined and equal to S if, for any increasing chain

$$I_0 \subset I_1 \subset I_2 \subset \cdots \subset I$$

of finite sets that exhausts I , we have $\lim_{k \rightarrow \infty} \prod_{i \in I_k} S_i = S$.

The only Banach algebra we will use is the algebra of 2 by 2 matrices, and the only result we will use is the following.

Lemma 3.3. *For elements s_i of any Banach algebra indexed by a countable totally ordered set, if $\sum \|s_i\| < \infty$, then $\prod(1 + s_i)$ is well-defined.*

Proof. Note that for elements s_1, \dots, s_n of a Banach algebra, and $1 \leq m \leq n$, we have

$$\begin{aligned} & \| (1 + s_1) \cdots (1 + s_n) - (1 + s_1) \cdots (1 + s_{m-1})(1 + s_{m+1}) \cdots (1 + s_n) \| \\ & \leq \|s_m\| \prod_{i=1}^n (1 + \|s_i\|). \end{aligned}$$

In the context at hand, the assumption gives that $\prod(1 + \|s_i\|)$ is bounded by some constant C , so we get that the effect of removing or adding a term s_n is at most $C\|s_n\|$. \square

To apply this lemma, we need to show the two-by-two matrices (Möbius transformations) S_i that we will use are close enough to the identity.

Lemma 3.4. *For the S_i arising as above from $\tilde{\lambda}$ and A , if we set $s_i = S_i - 1$, then $\sum \|s_i\| < \infty$. (Here 1 denotes the two-by-two identity matrix.)*

Proof. Each S_i can be realized as a time one stable or unstable horocycle flow matrix conjugated by geodesic flow, as in Figure 3.7. The basic computation

$$\begin{pmatrix} e^{-t/2} & 0 \\ 0 & e^{t/2} \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} e^{t/2} & 0 \\ 0 & e^{-t/2} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + \begin{pmatrix} 0 & e^{-t} \\ 0 & 0 \end{pmatrix}$$

shows that the s_i are small whenever the amount of geodesic flow used in the conjugation is large.

We partition all the crossings of our leaf segment A into finitely many subsets according to which “spike”, or corner of a triangle, they cross, see Figure 3.8. Then we show that the sum of the $\|s_i\|$ for each spike is bounded by a geometric progression, because the distance along the spike between neighboring crossings is always bounded below. \square

We now have the desired fact.




FIGURE 3.7. S_i can be written as geodesic flow along one orange segment, then horocycle flow for time one, and then geodesic flow backwards along the second orange segment.




FIGURE 3.8. A schematic of the intersections of an arc of the foliation with a spike.

Lemma 3.5. $S = \prod_{i \in I} S_i$.

Proof. Left to the reader as an exercise. (The hardest parts have been done above for you!) \square

Now, so far we've discussed the relative position of two triangles T_1 and T_2 which are joined by an arc A of the transverse foliation. Figure 3.9 shows that not all pairs of triangles are joined by such an arc A . The discussion may be clarified then by the following exercise.

Exercise 3.6. For any two triangles T, T' of $\tilde{\lambda}$ there is a sequence of triangles $T = T_0, T_1, \dots, T_n = T'$ of $\tilde{\lambda}$ such that T_i lies in between T_{i-1} and T_{i+1} , and there is an arc of $\tilde{\mu}$ from each triangle to the next.

Now we have reached the point where we understand $\tilde{\lambda}$ and $\tilde{\mu}$ quite well, when $\mu = F_\lambda(X)$. In fact, we understand it so well that, from the position of one triangle of $\tilde{\lambda}$, we can exactly determine the positions of all the others using the Möbius transformations S and the shears. The reader may check their understanding so far by completing the following exercise.




FIGURE 3.9. T_3 is hidden from T_1 , in that no leaf of the foliation intersects both T_3 and T_1 .

Exercise 3.7. Convince yourself that the above discussion amounts to a proof that F_λ is one-to-one.

Proof of Theorem 3.1. Now we will see that the Möbius transformations S and the shears can be defined for arbitrary $\mu \in \mathcal{MF}(\lambda)$. In this way we will define $\tilde{\lambda} \subset \mathbb{H}$, find it is invariant under a group Γ , and find $F_\lambda(\mathbb{H}/\Gamma) = \mu$, building the inverse for F_λ as desired.

To start, we use the fact that, given any $\mu \in \mathcal{MF}(\lambda)$, one can isotope μ to be actually transverse to λ , and each singularity of μ is then in a well-defined complementary triangle of λ independent of the isotopy. (The singularities of μ and the complementary triangles of λ must be in bijection to each other, because they are both in bijection to the zeros of the associated quadratic differential. Formally speaking, one should write down a more rigorous proof.)

Even without X , there is a topological version of $\tilde{\lambda}$ and $\tilde{\mu}$, defined up to isotopy on the universal cover of the topological surface. They are transverse.

First, we remark that the shears are obviously defined only in terms of topological data. Indeed, the shear is the transverse measure of the red segment in Figure 3.4. The key point is that whenever we took a hyperbolic length along an arc of a geodesic in $\tilde{\lambda}$, this was also the transverse measure assigned by $\tilde{\mu}$, because by definition the transverse measure for the horocycle foliation comes from hyperbolic length on the edges of each triangle.

Next, we recall that each S_i was defined as a conjugate of a time one horocycle flow. The amount of geodesic flow we conjugate by is

again a transverse measure assigned by $\tilde{\mu}$, so we can define the S_i from μ alone. We also check that Lemma 3.4 applies for arbitrary μ , so we can define the infinite products S .

Now, we can think of placing one triangle T of $\tilde{\lambda}$ on \mathbb{H} in an arbitrary way. (This arbitrary choice reflects the fact that everything is only defined up to Möbius transformations.) From this triangle, we can determine where we should put every triangle connected to T by a transverse arc, by using $S(v_1)$ and the shear. Continuing in this way we can determine where we should put every triangle of $\tilde{\lambda}$. We can obtain the rest of $\tilde{\lambda}$ as the closure of the set of edges.

Since the construction arises from objects on the surface, the resulting configuration of triangles in \mathbb{H} is invariant under a representation of this surface group into $PSL(2, \mathbb{R})$. More specifically, for γ in π_1 of the topological surface on which μ is defined, we may consider a pair of triangles $T_1, T_2 = \gamma(T_1)$ in the universal cover. The above discussion computes a Möbius transformation $\rho(\gamma)$ taking T_1 to T_2 . If Γ is the image of ρ , then we get that $F_\lambda(\mathbb{H}/\Gamma) = \mu$ as desired. (Note that Γ is discrete because it stabilizes a non-trivial lamination.) This concludes our proof that F_λ is a homeomorphism. \square

Remark 3.8. Because of group invariance, we can consider the shear to be defined for any two triangles on X joined by a transverse arc of the foliation.

Remark 3.9. If desired one could extend the shear by additivity to all pairs of triangles. For example, in Figure 3.9, the shear is defined for T_1 and T_2 , and also for T_2 and T_3 , and we can define the shear between T_1 and T_3 to be the sum of the shears from T_1 to T_2 and from T_2 to T_3 . This additivity makes it appropriate to refer to the shearing as a cocycle.

4. THE FUNDAMENTAL LEMMA ON EARTHQUAKES

We've discussed the shear between two triangles joined by an arc A : one follows the singular leaf from one triangle, and looks at where it lands on another triangle, and take the transverse measure, or equivalently hyperbolic length, of the arc of the boundary geodesic from that landing point to the center point. The fundamental engine of Mirzakhani's isomorphism is how this shear changes when you earthquake in λ . It is implicit that λ is maximal.

Lemma 4.1. *Denote by $\text{Shear}_X(T_1, T_2)$ the shear for two triangles joined by an arc A of the horocyclic foliation on the hyperbolic surface X . Then*

$$\text{Shear}_{E_{t\lambda}(X)}(T_1, T_2) = \text{Shear}_X(T_1, T_2) + t\lambda(A),$$

where $\lambda(A)$ denotes the transverse measure of A and t is sufficiently small.

In other words, “the change in shear is equal to the transverse measure.” Mirzakhani cites [Bon96] for this fact, but it can be seen quite easily as follows.

The restriction that t be small is absolutely not required, but it is sufficient for our purposes, and allows us to avoid thinking about, for example, the situation where T_2 and T_1 are not joined by an arc of the horocyclic foliation on $E_{t\lambda}(X)$.

Before reading the proof, the reader may first want to do the following warm up exercise.

Exercise 4.2. Let T_1 and T_2 be triangles in \mathbb{H} that share an edge γ . How does the shear change after moving one of the triangles by a hyperbolic isometry with axis γ and translation distance t ?

Proof of Lemma 4.1. Without loss of generality take $t = 1$.

T_1 and T_2 are separated by infinitely many leaves of $\tilde{\lambda}$. As discussed in the definition of earthquakes, we can understand how T_2 is moved by the earthquake E_λ (assuming T_1 is fixed, i.e. relative to T_1) by approximating the measured lamination between T_1 and T_2 by a discrete one. So we do this, picking a discrete lamination consisting of a finite subset of the leaves of λ that bound triangles. It doesn't matter to us if this is done in a group equivariant way, since we are just approximating the earthquake in \mathbb{H} . (Indeed the experts may note that it can't be done in a group equivariant way. The quotient would be a discrete lamination, and hence must consist of closed leaves, but λ has no closed leaves.)

If we earthquake along a leaf γ of $\tilde{\lambda}$ between T_1 and T_2 by an amount t , this changes the shear between T_1 and T_2 by exactly t , basically by definition. Indeed, the earthquake applies the hyperbolic isometry that translates along γ to the half plane \mathbb{H}_γ on the T_2 side of γ . This moves $\tilde{\lambda}$ by this isometry on \mathbb{H}_γ , and hence it translates the transverse horocyclic foliation on \mathbb{H}_γ . Hence, each arc of the transverse horocyclic foliation in \mathbb{H}_γ with an endpoint on γ is translated so that the new endpoint is t farther along on γ .

Similarly if we earthquake along finitely many leaves of λ with measures t_i , the shear changes by precisely $\sum t_i$. So, taking a limit, we see that the shear between T_1 and T_2 changes by an amount equal to the transverse λ measure of a transverse arc starting in T_1 and ending in T_2 . \square

Remark 4.3. It may seem strange that the “Fundamental Lemma”, as we have named it, does not apply to arbitrary earthquakes, but

only to earthquakes in maximal laminations. However any lamination can be extended, in many ways, to a maximal lamination, and the Fundamental Lemma applies to a measured maximal lamination even if the measure doesn't have full support. Recall that the horocyclic foliation, which we use to define shear, doesn't even depend on or require a measure on λ .

5. MIRZAKHANI'S ISOMORPHISM

We now turn to the proof of Theorem 1.1. We begin by specifying full measure sets on which we will build the desired conjugacy F .

Let QD_0 denote the locus of quadratic differentials over Teichmüller space that don't have any horizontal saddle connections and that have only simple zeros. Recall that a quadratic differential with simple zeros has $4g - 4$ zeros.

Remark 5.1. This condition is equivalent to the horizontal lamination being maximal. This can be checked in more than one way. For example, you can note that each simple zero without a horizontal saddle connection gives a complementary triangle, and $4g - 4$ times the area of the triangle is the area of the surface, so there is no room for any other complementary regions.

Let \mathcal{ML}_0 denote the locus of measured foliations that are maximal. We will build a mapping class group equivariant measurable isomorphism F from $\mathcal{ML}_0 \times \mathcal{T}_g$ to QD_0 that conjugates earthquake flow to unipotent flow. The map sends (λ, X) to the quadratic differential with foliations $(\lambda, F_\lambda(X))$,

$$F(\lambda, X) = q(\lambda, F_\lambda(X)).$$

Here $F_\lambda(X)$ continues to denote the horocyclic foliation of X . This map is only measurable, but its restriction to each slice $\{\lambda\} \times \mathcal{T}_g$ is a homeomorphism onto the set of quadratic differentials with horizontal lamination λ by Theorem 3.1. It follows that F is a bijection from $\mathcal{ML}_0 \times \mathcal{T}_g$ to QD_0 .

It remains only to show that the image of the earthquake flow path $(\lambda, E_{t\lambda}(X))$ is a unipotent flow path. We begin by discussing Teichmüller unipotent flow, which is of course characterized by how it changes period coordinates. But first we present a lemma that will allow us to restrict from arbitrary periods to special saddle connections.

Lemma 5.2. *Every isotopy class of path joining singularities of a quadratic differential can be realized by a sequence of paths that start at one singularity, travel in the horizontal direction, then travel in the vertical direction and end at a singularity.*

Proof. It suffices to prove this for saddle connections. This can be done by growing rectangles: Look at the rectangle from one endpoint to a point on the saddle connection nearby, and grow this rectangle until it hits a singularity. Continue in this way as in Figure 5.1. \square




FIGURE 5.1. The proof of Lemma 5.2.

Corollary 5.3. *Suppose q_t is a path of quadratic differentials. Suppose that for every t_0 and every path γ on q_{t_0} as in the lemma, the period $x_t + iy_t$ of γ satisfies*

$$\frac{d}{dt} \Big|_{t=t_0} x_t = y_{t_0}, \quad \text{and} \quad \frac{d}{dt} \Big|_{t=t_0} y_t = 0.$$

Then q_t is an orbit of Teichmüller unipotent flow.

Proof. It suffices to recall that

$$(5.0.1) \quad \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x + ty \\ y \end{pmatrix}$$

and that a function with constant derivative is linear. \square

Observe that the y component of the period of γ is given by the transverse measure of γ for the horizontal measured foliation, see Figure 5.2. The intuition of the proof of Theorem 1.1 is to think of each singularity of q as corresponding to a complementary triangle for a lamination, and to think of the x component of a period of such a γ as the shear between the two corresponding triangles. We offer the following chart to summarize this intuition, before beginning the formal proof.




FIGURE 5.2. The y component is given by the horizontal foliation.

earthquake flow	\longleftrightarrow	horocycle flow
λ	\longleftrightarrow	horizontal foliation
$F_\lambda(X)$	\longleftrightarrow	vertical foliation
$(\lambda, F_\lambda(X))$	\longleftrightarrow	quadratic differential
triangle	\longleftrightarrow	singularity
shear	\longleftrightarrow	x -component of a period
Fundamental Lemma	\longleftrightarrow	equation (5.0.1)

Proof of Theorem 1.1. We wish to show that

$$q(\lambda, F_\lambda(E_{t\lambda}(X)))$$

is a horocycle flow path using Corollary 5.3. We will consider a moment in time t_0 , which without loss of generality is $t_0 = 0$, and show that for each path γ as above, the derivative of the period of γ satisfies Corollary 5.3.

We've already done most of the work to see this. Indeed, the path γ corresponds to a path in X or $\tilde{X} = \mathbb{H}$ joining two triangles. If the period of γ is (x_t, y_t) , then we see that x_t is the transverse measure of γ given by the vertical foliation, and similarly for y_t . So $y_t = \lambda(\gamma)$ is constant. And the derivative of x_t is equal to $y_t = \lambda(\gamma)$ by the Fundamental Lemma.

Hence Corollary 5.3 gives that

$$q(\lambda, F_\lambda(E_{t\lambda}(X))) = F(\lambda, E_{t\lambda}(X))$$

is an earthquake path as desired. We already known that F is a homeomorphism from $\mathcal{ML}_0 \times \mathcal{T}_g$ to QD_0 , so this concludes the proof. \square

Remark 5.4. In [Mir08], Mirzakhani factors the map F_λ through a subset $H_+(\lambda)$ of the space $H(\lambda, \mathbb{R})$ of transverse cocycles on λ . There




FIGURE 5.3. This picture isn't geometrically accurate, but it gives an idea of how to think of γ as it lies on $\mathbb{H} = \tilde{X}$. The horizontal lines are leaves of $\tilde{\lambda}$, and the vertical lines are leaves of the horocyclic foliation.

is a map

$$I_\lambda : \mathcal{ML}(\lambda) \rightarrow H_+(\lambda)$$

that records the shearing data as a transverse cocycle. Mirzakhani considers a map

$$G_\lambda : \mathcal{T}_g \rightarrow H_+(\lambda)$$

satisfying $F_\lambda = I_\lambda^{-1} \circ G_\lambda$. The map G_λ is a symplectomorphism by [BS01], and using the fact the earthquake flow is Hamiltonian Mirzakhani concludes that G_λ conjugates $E_{t\lambda}$ with a natural linear flow on $H_+(\lambda)$. Mirzakhani also notes that if $q(\lambda, \mu_t)$ is a horocycle flow path, then $I_\lambda(\mu_t)$ is an orbit of the same linear flow on $H_+(\lambda)$. From this she concludes that F conjugates earthquake flow to horocycle flow. The proof we have presented thus differs from [Mir08] in that we do not factor through the intermediary $H_+(\lambda)$ and we use the Fundamental Lemma, rather than the Hamiltonian nature of the flows and [BS01], to establish the conjugacy.

Remark 5.5. I conjecture that the semi-conjugacy is continuous on $\mathcal{ML}_0 \times \mathcal{T}_g$. For example, if you take a sequence of maximal laminations λ_n that converge to some λ that is also maximal, then for each fixed X the horocyclic foliation on X for λ_n should converge to that of λ .

But the semi-conjugacy cannot be extended continuously to even to the locus where λ has a single quadrilateral and the rest of the complementary regions are triangles. The reason is that such λ are limits of maximal laminations in two different ways, essentially corresponding to the two different ways to turn the quadrilateral into two triangles. These two choices give two different horocyclic foliations. Note that




FIGURE 5.4. The proof that F cannot be extended to a continuous map.

this situation arises if you approximate a quadratic differential with a saddle with nearby quadratic differentials where that saddle can either slope slightly up or slightly down. (This comment is necessary because you can't just add one geodesic to the quadrilateral and get a lamination with a transverse measure of full support.)

In general, for each λ , there is possibly a finite or infinite number of ways to fill in λ to a maximal lamination (without a measure of full support), and each of these different maximal extensions gives a different horocyclic foliation that will serve as the vertical measured foliation for a quadratic differential. Perhaps one can think that Mirzakhani's map as being multivalued off of \mathcal{ML}_0 and the multiple values correspond to all these choices of maximal extension. Similarly if one wanted to compute F_λ when λ isn't maximal, one could do it by extending λ to be maximal in a number of ways, so one can also think of the inverse of F as being multi-valued.

Alternatively, one could consider earthquake flow on $\mathcal{ML}_{ext} \times \mathcal{T}_g$, where \mathcal{ML}_{ext} consists of all pairs of a measured lamination plus an extension of its support to a maximal lamination, and we use the topology that requires convergence of both the measured lamination and the

maximal unmeasured extension. This flow should map continuously onto both earthquake flow and Teichmüller unipotent flow.

Remark 5.6. I don't know how to show that there couldn't be some (totally different) continuous conjugacy between earthquake flow and Teichmüller unipotent flow. This seems like an interesting open problem.

Remark 5.7. Consider a partition κ of $4g - 4$. Consider the subset $\mathcal{I}(\kappa) \subset \mathcal{ML} \times \mathcal{T}_g$ where the complementary regions of the maximal measured lamination are all symmetric ideal hyperbolic polygons, with the number of polygons with a given number of edges given by κ . (You view κ as a partition of the area divided by π .) Minus the symmetry assumption, this would be just a condition on the measured lamination, and not the point of Teichmüller space. The symmetry condition says that there is a hyperbolic isometry that cyclically permutes the ideal vertices.

One can't even extend F to this locus. However, I conjecture that there is a different F that is a conjugacy from $\mathcal{I}(\kappa)$ to the stratum $\mathcal{Q}(\kappa)$ of quadratic differentials. To be more precise, the image of F would be the locus with no horizontal saddle connections in that stratum. This map would use the horocyclic foliation that is defined for symmetric ideal polygons.

Remark 5.8. There is a notion of hyperbolic length of geodesic laminations. The hyperbolic length of λ on $X \in \mathcal{T}_g$, is $i(\lambda, F_\lambda(X))$. If you'd like you can take this as a definition. It makes sense because the transverse measure for $F_\lambda(X)$ corresponds to arc-length along λ .

It follows that the semi-conjugacy is such that if λ has hyperbolic length ℓ on X , then the resulting quadratic differential has area ℓ . A fancy way of putting this is to say λ has extremal length ℓ on the image Riemann surface. (Don't worry if you don't know what that means.)

Remark 5.9. Mirzakhani's map F simultaneously conjugates Thurston's stretch map flow [**Thu**] to the action of

$$\begin{pmatrix} 1 & 0 \\ 0 & e^s \end{pmatrix}$$

on QD . This is because Thurston's stretch map flow is simply scalar multiplication in shear coordinates.

Remark 5.10. I conjecture that F maps co-bounded sets to co-bounded sets. Co-bounded means contained in a compact set after you quotient by the action of the mapping class group. This is related to the fact that the set of maximal unmeasured laminations on a given (unmarked)

hyperbolic surface should be compact. So given any non-maximal lamination, there should be a compact set of ways to extend it to a maximal lamination, and these should correspond to the possible limiting values of F . Similarly for F^{-1} .

6. INVARIANT MEASURES

There is a natural measure called Thurston measure μ_{Th} on \mathcal{ML} . It is basically the same thing as Masur-Veech measure. Most laminations are not orientable, but can be made so by passing to a double cover, after which they give a cohomology class. For nearby laminations, you can pass to a common (branched) double cover, so they give cohomology classes in the same vector space. Thurston measure is Lebesgue measure in this vector space. (Actually the vector space is the -1 eigenspace of the double cover.)

As discussed, QD is equal to $\mathcal{ML} \times \mathcal{ML} \setminus \Delta$. It isn't hard to show that the Masur-Veech measure (not just on the unit area locus) is equal to the restriction of $\mu_{Th} \times \mu_{Th}$ to the complement of Δ . Indeed, Masur-Veech measure also arises from taking cohomology classes on the double cover where the foliation becomes orientable.

A basic fact that we will discuss in the next section is that earthquakes are Hamiltonian flows. A corollary is the following.

Theorem 6.1. *The action of $E_{t\lambda}$ on \mathcal{T}_g leaves invariant the Weil-Petersson measure μ_{WP}*

Recall that the Weil-Petersson is nothing other than the standard Lebesgue measure in Fenchel-Nielsen coordinates.

Corollary 6.2. *For any measure ν on \mathcal{ML} , the earthquake flow leaves the measure $\nu \times \mu_{WP}$ on $\mathcal{ML} \times \mathcal{T}_g$ invariant. In particular, $\mu_{Th} \times \mu_{WP}$ is both invariant under earthquake flow and the action of the mapping class group.*

Recall from Remark 5.8 that one can take the hyperbolic length of a lamination. It isn't hard to show that this length $\ell_X(\lambda)$ is invariant under earthquake flow in λ . For example, when you earthquake in a simple closed curve, the hyperbolic length of that curve doesn't change. So earthquake flow preserves each level set \mathcal{ML}_ℓ for the hyperbolic length of λ .

If one wishes an invariant measure on the set $\mathcal{ML}_1 \times \mathcal{T}_g$ where the measured lamination has length 1, one does the same thing as for Masur-Veech measure. Namely, over a point $X \in \mathcal{T}_g$, the measure used on \mathcal{ML}_1 gives a subset of \mathcal{ML}_1 the Thurston measure of its cone in \mathcal{ML} . (The cone on a set consists of anything in the set times any

number in $[0, 1]$.) This gives a mapping class group and earthquake flow invariant measure on $\mathcal{ML}_1 \times \mathcal{T}_g$.

Similarly one gets invariant measures on any other level set $\mathcal{ML}_\ell \times \mathcal{T}_g$ for the length function. The measure of $\mathcal{ML}_\ell \times \mathcal{T}_g$ is equal to ℓ^{6g-6} times the measure of $\mathcal{ML}_1 \times \mathcal{T}_g$ (the measure is finite after quotienting by the mapping class group).

The measure on $\mathcal{ML}_\ell \times \mathcal{T}_g$ must map to a measure on QD_ℓ , the set of area ℓ quadratic differentials. The image measure is Lebesgue class and invariant under horocycle flow, so using ergodicity of horocycle flow it must be a multiple of Masur-Veech. (The first point should be true since both maps $\mathcal{ML} \times \mathcal{T}_g \rightarrow \mathcal{ML} \times \mathcal{MF}$ and $QD \rightarrow \mathcal{ML} \times \mathcal{MF}$ are pretty nice and understandable maps. For example, if you change λ a bit then $F_\lambda(X)$ only changes a bit. For a formal proof, one likely has to look at Bonahon's papers.)

The isomorphism is also a conjugacy for rescaling the λ and rescaling the horizontal foliation of the quadratic differential. The measure of QD_ℓ is also equal to ℓ^{6g-6} times the measure of QD_1 , so we get the multiple is independent of ℓ . Hence $\mu_{Th} \times \mu_{WP}$ maps to $c\mu_{Th} \times \mu_{Th}$ for some $c > 0$, and so $F_\lambda^*(\mu_{Th}) = c\mu_{WP}$. In fact, Bonahon-Sozen gave a more explicit proof of this, before Mirzakhani's isomorphism, that computes that $c = 1$ and handles the symplectic forms rather than just their associated volume forms [BS01].

Theorem 6.3. $F_\lambda^*(\mu_{Th}) = \mu_{WP}$.

Their proof was discovered using the case when λ contains a pair of pants (which doesn't fit into our setting, since such λ don't have fully supported transverse measures) and the magic formula of Wolpert.

Corollary 6.4. *The Masur-Veech volume of the principal stratum of quadratic differentials is*

$$\int_{X \in \mathcal{M}_g} \mu_{Th}(B_X(1)) d\mu_{WP},$$

where $B_X(1)$ is the unit ball in \mathcal{ML} of lamination of length at most 1 on X .

This corollary is [Mir08, Theorem 1.4].

Proof sketch. Using Fubini, the integral is the $\mu_{WP} \times \mu_{Th}$ measure of a fundamental domain for the mapping class group on

$$\{(X, \lambda) \in \mathcal{T}_g \times \mathcal{ML} : \ell_X(\lambda) \leq 1\}.$$

Using that $F_\lambda^*(\mu_{Th}) = \mu_{WP}$, this is equal to the $\mu_{Th} \times \mu_{Th}$ measure of a fundamental domain for the mapping class group on

$$\{(h, v) \in \mathcal{ML} \times \mathcal{ML} \setminus \Delta : i(h, v) \leq 1\},$$

which is the Masur-Veech volume of the principal stratum. \square

7. LAMINATIONS CONTAINING A PANTS DECOMPOSITION

We now consider the case of maximal lamination λ that contains a pants decomposition \mathcal{P} , i.e. a maximal set of disjoint curves. Such λ are seemingly irrelevant for the discussion above, because they are not in the locus where Mirzakhani's semi-conjugacy is defined. Indeed, because such λ can't have a fully supported transverse measure, they can't arise as the horizontal lamination of a quadratic differential. But the map F_λ is defined for any λ maximal, and considering this case will provide insight.

We can glue together two topological ideal triangles, i.e. triangles minus their vertices, to get a sphere minus three points, as in Figure 7.1.




FIGURE 7.1. A sphere minus three points can be obtained by gluing two triangles minus their vertices.

Let us consider gluing together two ideal hyperbolic triangles along isometries of their edges. We'll glue in the same pattern, so we know that the result will topologically be a sphere minus three points, which topologically is the same thing as a sphere minus three discs. The result will have a hyperbolic metric, but this metric might be incomplete. There are three parameters, the three shears, that we'll denote s_1, s_2, s_3 . Each shear is the distance between "center points" of edges that are glued together, in the usual way.

If you follow the horospherical foliation around a puncture, passing through both triangles, you arrive further out along the edge of the triangle by an amount equal to the sum of the shears, say $|s_1 + s_2|$. See Figure 7.2. You can then complete this horospherical path to a loop by traveling this $|s_1 + s_2|$ along the geodesic. As you slide this path




FIGURE 7.2. Here the shears are shown in red, and have opposite signs. The left and right geodesics are glued by a hyperbolic isometry that takes the basepoint of the blue arrow to its tip.

farther out along the cusps of the triangles, the distance traveled in the horospherical part of this loop goes to zero, so this loop seems to be converging to a geodesic of length $|s_1 + s_2|$.

Lemma 7.1. *The completion of the surface obtained by gluing together two triangles as above is a pair of pants with boundary geodesics of length $|s_1 + s_2|, |s_2 + s_3|, |s_3 + s_1|$. If any of these three quantities are zero, you instead get a cusp.*

For a very careful and clear proof, which proceeds using the developing map rather than the informal heuristic we have suggested, see [Mar, Section 7.4].




FIGURE 7.3. Each cusp of each triangle spirals towards one of the three boundary curves of the pants. Image from [Mar, Figure 7.20].

Remark 7.2. As we linearly interpolate between (s_1, s_2, s_3) and $(-s_1, -s_2, -s_3)$, at the halfway point $(0, 0, 0)$ each boundary component will reach zero

length and become a cusp. On one side of the interpolation the triangles will spiral towards the boundary component in one direction, and on the other side they will spiral in the other direction.

Return to the situation of a maximal lamination λ containing a pants decomposition \mathcal{P} . The s_i above are some of the shear coordinates for \mathcal{T}_g . One also requires shear coordinates for small arcs A passing through a boundary of a pants. This shear coordinate is directly seen to be similar to a Fenchel-Nielsen twist parameter, in that if we do a Fenchel-Nielsen twist by ε , the shear changes by ε .

The Fenchel-Nielsen twist is presumably not the exact same thing as the shear of the arc A , even up to a constant. This is because Fenchel-Nielsen twist parameters are usually defined by considering orthogeodesics from another boundary of the pants. The shear is related to where the central leaves of the horocyclic foliation lands on the cuff. The shear for A should be a function of the Fenchel-Nielsen twist parameter of that curve, and the 5 length parameters for the 2 pants that share this cuff. In this way one can at least see that the map from Fenchel-Nielsen twist parameters to the shear parameters preserves volume, because its derivative can be written as an upper triangular matrix with ones on the diagonal.

8. HAMILTONIAN FLOWS

In fact, both the Thurston and Weil-Petersson volume forms arise from symplectic forms. (Although it is a little tricky to talk about symplectic forms since \mathcal{ML} doesn't have a natural differential structure.) Bonahon-Sozen actually showed that the map from \mathcal{T}_g to \mathcal{ML}_λ is a symplectomorphism. The earthquake flow on \mathcal{T}_g is the Hamiltonian flow for the length of λ , and the unipotent flow on quadratic differentials with horizontal foliation λ is the Hamiltonian flow of the area function.

Consider a specific μ , and let \hat{X} denote the double cover associated to $q_{\lambda,\mu}$. We can associate λ and μ to cohomology classes $\hat{\lambda}$ and $\hat{\mu}$, and the area function A is given by

$$A(\eta) = \langle \hat{\lambda}, \eta \rangle.$$

We now claim that the Hamiltonian vector field is $\hat{\lambda}$. To show this, we compute

$$\begin{aligned} (dA)_\eta(\xi) &= \left. \frac{d}{dt} \right|_{t=0} \langle \hat{\lambda}, \eta + t\xi \rangle \\ &= \langle \hat{\lambda}, \xi \rangle. \end{aligned}$$

This exactly shows that unipotent flow is Hamiltonian.

9. THE LINEAR STRUCTURE ON \mathcal{ML}_α

If α is maximal, then all foliations $\mu \in \mathcal{ML}_\alpha$ have singularities in correspondence to the triangular regions of α . Hence one can pass consistently to a double cover where $\hat{\mu}$ gives a cohomology class $[\hat{\mu}]$. This maps \mathcal{ML}_α to a vector space.

Lemma 9.1. *The map $\mu \mapsto [\hat{\mu}]$ is injective.*

Proof. This is equivalent to the statement that if you know the horizontal foliation of an Abelian differential (up to Whitehead moves), and you know the relative cohomology class of the vertical foliation, then you can recover the Abelian differential. (Actually $\hat{\mu}$ lives in the -1 eigenspace of the cohomology of the double cover, which is isomorphic to the -1 eigenspace of the relative cohomology.) The proof is that knowing the horizontal foliation allows you to determine the IET giving the first return map to any vertical segment, and the cohomology class gives the sizes of the rectangles in an associated zippered rectangle decomposition. See [MW14]. \square

Remark 9.2. This can be interpreted as saying that, passing to the appropriate Teichmüller space, a single period coordinate chart suffices for the slice of any stratum where the horizontal foliation is held constant.

In fact one can see that the image is a convex polyhedral cone. Typically (ex if α is uniquely ergodic) this cone is a half space.

Alternatively, one can parameterize \mathcal{ML}_α by transverse distributions or transverse cocycles, and in this way see that \mathcal{ML}_α has a natural linear structure [Bon96]. Any two points in \mathcal{ML}_α can be joined by a straight line, and the resulting path in \mathcal{T}_g is called a cataclysm or shear map. It differs from an earthquake in that earthquakes always shear in one direction (right or left), and that earthquakes can be continued for all time, whereas cataclysm paths can cease to be well-defined in finite time.

10. OTHER RESULTS ON EARTHQUAKES

Thurston proved that, given any two points in \mathcal{T}_g , there is a unique earthquake path between them [Thu86]. Kerckhoff proved that hyperbolic length functions are convex along earthquake (and even cataclysm paths [The]). This was famously used by Kerckhoff to solve the Nielsen realization problem [Ker83].

REFERENCES

- [AH21] Francisco Arana-Herrera, *Equidistribution of families of expanding horospheres on moduli spaces of hyperbolic surfaces*, *Geom. Dedicata* **210** (2021), 65–102.
- [AHW] Francisco Arana-Herrera and Alex Wright, *The asymmetry of Thurston's earthquake flow*, forthcoming.
- [BM00] M. Bachir Bekka and Matthias Mayer, *Ergodic theory and topological dynamics of group actions on homogeneous spaces*, London Mathematical Society Lecture Note Series, vol. 269, Cambridge University Press, Cambridge, 2000.
- [Bon88] Francis Bonahon, *The geometry of Teichmüller space via geodesic currents*, *Invent. Math.* **92** (1988), no. 1, 139–162.
- [Bon96] ———, *Shearing hyperbolic surfaces, bending pleated surfaces and Thurston's symplectic form*, *Ann. Fac. Sci. Toulouse Math. (6)* **5** (1996), no. 2, 233–297.
- [BS01] Francis Bonahon and Yaşar Sözen, *The Weil-Petersson and Thurston symplectic forms*, *Duke Math. J.* **108** (2001), no. 3, 581–597.
- [Cal07] Danny Calegari, *Foliations and the geometry of 3-manifolds*, Oxford Mathematical Monographs, Oxford University Press, Oxford, 2007.
- [CB88] Andrew J. Casson and Steven A. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*, London Mathematical Society Student Texts, vol. 9, Cambridge University Press, Cambridge, 1988.
- [CF21] Aaron Calderon and James Farre, *Shear-shape cocycles for measured laminations and ergodic theory of the earthquake flow*, arXiv preprint arXiv:2102.13124 (2021).
- [Duo] Yen Duong, bakingandmath.com/2016/03/01/connecting-hyperbolic-and-half-translation-surfaces-part-ii-math/.
- [Fena] Aaron Fenyes, *A dynamical perspective on shear-bend coordinates*, arXiv: 1510.05757v2 (2018).
- [Fenb] ———, *Warping geometric structures and abelianizing $SL(2, \mathbb{R})$ local systems*, <https://repositories.lib.utexas.edu/handle/2152/41629>.
- [FLP12] Albert Fathi, François Laudenbach, and Valentin Poénaru, *Thurston's work on surfaces*, Mathematical Notes, vol. 48, Princeton University Press, Princeton, NJ, 2012, Translated from the 1979 French original by Djun M. Kim and Dan Margalit. MR 3053012
- [FM12] Benson Farb and Dan Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012.
- [FM14] Giovanni Forni and Carlos Matheus, *Introduction to Teichmüller theory and its applications to dynamics of interval exchange transformations, flows on surfaces and billiards*, *J. Mod. Dyn.* **8** (2014), no. 3-4, 271–436.
- [GM91] Frederick P. Gardiner and Howard Masur, *Extremal length geometry of Teichmüller space*, *Complex Variables Theory Appl.* **16** (1991), no. 2-3, 209–237.
- [Gup] Subhojoy Gupta, *Limits of harmonic maps and crowned hyperbolic surfaces*, arXiv:1805.03921 (2018).

- [GW17] Subhojoy Gupta and Michael Wolf, *Meromorphic quadratic differentials with complex residues and spiralling foliations*, In the tradition of Ahlfors-Bers. VII, Contemp. Math., vol. 696, Amer. Math. Soc., Providence, RI, 2017, pp. 153–181.
- [Kap01] Michael Kapovich, *Hyperbolic manifolds and discrete groups*, Progress in Mathematics, vol. 183, Birkhäuser Boston, Inc., Boston, MA, 2001.
- [Ker83] Steven P. Kerckhoff, *The Nielsen realization problem*, Ann. of Math. (2) **117** (1983), no. 2, 235–265.
- [Lei12] Christopher J. Leininger, *Degenerations of hyperbolic structures on surfaces*, Geometry, topology and dynamics of character varieties, Lect. Notes Ser. Inst. Math. Sci. Natl. Univ. Singap., vol. 23, World Sci. Publ., Hackensack, NJ, 2012, pp. 95–138.
- [Lev83] Gilbert Levitt, *Foliations and laminations on hyperbolic surfaces*, Topology **22** (1983), no. 2, 119–135.
- [Mar] Bruno Martelli, *An introduction to geometric topology*, arXiv:1610.02592 (2016).
- [McM] Curtis McMullen, *The work of Maryam Mirzakhani*, <https://people.math.harvard.edu/~ctm/papers/home/text/papers/icm14/icm14.pdf>.
- [Mir07] Maryam Mirzakhani, *Random hyperbolic surfaces and measured laminations*, In the tradition of Ahlfors-Bers. IV, Contemp. Math., vol. 432, Amer. Math. Soc., Providence, RI, 2007, pp. 179–198.
- [Mir08] ———, *Ergodic theory of the earthquake flow*, Int. Math. Res. Not. IMRN (2008), no. 3, Art. ID rnm116, 39.
- [Mir16] ———, *Counting mapping class group orbits on hyperbolic surfaces*, arXiv preprint arXiv:1601.03342 (2016).
- [MW02] Yair Minsky and Barak Weiss, *Nondivergence of horocyclic flows on moduli space*, J. Reine Angew. Math. **552** (2002), 131–177.
- [MW14] ———, *Cohomology classes represented by measured foliations, and Mahler's question for interval exchanges*, Ann. Sci. Éc. Norm. Supér. (4) **47** (2014), no. 2, 245–284.
- [Pap86] Athanase Papadopoulos, *Geometric intersection functions and Hamiltonian flows on the space of measured foliations on a surface*, Pacific J. Math. **124** (1986), no. 2, 375–402.
- [PH92] R. C. Penner and J. L. Harer, *Combinatorics of train tracks*, Annals of Mathematics Studies, vol. 125, Princeton University Press, Princeton, NJ, 1992.
- [Pie] Heather Pierce, www.geogebra.org/m/fUCCfAEj.
- [PT08] Athanase Papadopoulos and Guillaume Th eret, *Shift coordinates, stretch lines and polyhedral structures for Teichm uller space*, Monatsh. Math. **153** (2008), no. 4, 309–346.
- [The] Guillaume Th eret, *Convexity of length functions and Thurston's shear coordinates*, arXiv:1408.5771 (2014).
- [Thu] William Thurston, *Minimal stretch maps between hyperbolic surfaces*, arXiv: 9801039 (1998).
- [Thu86] William P. Thurston, *Earthquakes in two-dimensional hyperbolic geometry*, Low-dimensional topology and Kleinian groups (Coventry/Durham,

- 1984), London Math. Soc. Lecture Note Ser., vol. 112, Cambridge Univ. Press, Cambridge, 1986, pp. 91–112.
- [Thu97] ———, *Three-dimensional geometry and topology. Vol. 1*, Princeton Mathematical Series, vol. 35, Princeton University Press, Princeton, NJ, 1997, Edited by Silvio Levy.
- [Wri16] Alex Wright, *From rational billiards to dynamics on moduli spaces*, Bull. Amer. Math. Soc. (N.S.) **53** (2016), no. 1, 41–56.
- [Wri20] ———, *A tour through Mirzakhani’s work on moduli spaces of Riemann surfaces*, Bull. Amer. Math. Soc. (N.S.) **57** (2020), no. 3, 359–408.