

Harmonic Splittings of Surfaces

Benson Farb* Michael Wolf†

Final Version, June 8, 1999

1 Introduction

Let $\Gamma = \pi_1(S)$ be the fundamental group of a closed surface S of genus at least two. Morgan-Shalen showed [MS2], [GiSh] that every point in the Thurston compactification $\mathcal{PMF}(S)$ of the Teichmüller space $\text{Teich}(S)$ gives an isometric Γ -action on an \mathbf{R} -tree. Given a measured foliation $\mathcal{F} \in \mathcal{PMF}(S)$, the action is simply the Γ -action on the leaf space of the lift of \mathcal{F} to \mathbf{H}^2 . This action is *small* in the sense that edge stabilizers do not contain rank two free groups. It is also *minimal* in the sense that it leaves no proper subtree invariant.

Shalen [Sh] conjectured that every minimal small action of Γ on an \mathbf{R} -tree T arises in this way. This problem has several important applications in low-dimensional geometry and topology (see [Ot]). Partial results were obtained by Morgan-Shalen [MS1] and Gillet-Shalen [GiSh].

The conjecture was eventually proven in two steps: Morgan-Otal [MO] (see also [Ha]) constructed the candidate foliation, with dual \mathbf{R} -tree R , and a Γ -equivariant morphism $\phi : R \rightarrow T$ so that ϕ has no “edge folds” (see below); then Skora [Sk1, Sk2] showed that ϕ has no “vertex folds”, giving that ϕ is a Γ -equivariant isometry, completing an affirmative solution to the conjecture.

Theorem 1.1 (Morgan-Otal, Skora) *Let $\Gamma = \pi_1(S)$, S a closed surface of genus at least two. Then any small, minimal Γ -action on an \mathbf{R} -tree is dual to the lift of a measured foliation on S .*

A complete exposition of Theorem 1.1 is given in [Ot].

*Partially supported by NSF grant DMS 9704640

†Partially supported by NSF grants DMS 9300001 and DMS 9626565

The purpose of the present paper is to prove Theorem 1.1 from a different point of view, using harmonic maps. Harmonic maps were used by Gromov-Schoen [GrSc] to show that certain groups do not act nontrivially on singular spaces such as trees. Here we use harmonic maps to classify, in the special case of a surface group, all minimal, small actions on \mathbf{R} -trees, against a background where many such actions exist (namely $6g-7$ dimensions worth).

Our other interest in this proof is in the way it uses harmonic maps as a tool in combinatorial group theory. For example, combinatorial topology arguments become greatly simplified (via the maximum principle) when looking at a harmonic representative. Another example is the existence of a moduli space of harmonic maps (and harmonic map invariants) associated to a group action, allowing for an extra tool in proofs.

1.1 Outline

Here is a brief description of our approach to the proof.

Step 1 (Find a harmonic map): Given a small action of the surface group $\Gamma = \pi_1(S)$ on an \mathbf{R} -tree T , it is relatively straightforward to find a Γ -equivariant harmonic map $f : \tilde{S} \rightarrow T$. Here we have endowed S with a complex structure.

Step 2 (Associated data): The harmonic map f automatically has associated to it the following data:

- a Γ -equivariant holomorphic quadratic (Hopf) differential $\tilde{\Phi}$ on the Riemann surface \tilde{S}
- a Γ -equivariant measured foliation $\tilde{\mathcal{F}}$, the vertical foliation of $\tilde{\Phi}$
- the leaf space R of $\tilde{\mathcal{F}}$, with metric induced from the measure on $\tilde{\mathcal{F}}$, making R into an \mathbf{R} -tree

The map f is projection along the leaves of $\tilde{\mathcal{F}}$, with the possibility of several vertical leaves being sent to the same point in T . The Γ -action on S induces an isometric Γ -action on R .

Step 3 (Morphism from a geometric action to the given action):

Let $\pi : \tilde{S} \rightarrow R$ be the natural projection sending each leaf of $\tilde{\mathcal{F}}$ to a point. Here a leaf may have a countable number of k -pronged singularities. We then obtain a Γ -equivariant *morphism* $\phi : R \rightarrow T$ of \mathbf{R} -trees, where $\phi = f \circ \pi^{-1}$. We must show that ϕ is an isometry, which is the same as saying that ϕ does not *fold* at any point.

Step 4a (No edge folds): If ϕ folded at an *edge point* of R , i.e. a point whose “tangent space” has only two directions, then this would force f to locally take the form $z \mapsto |\operatorname{Re} z|$ which is not harmonic. Hence there are no edge folds, nor even *vertex folds* at trivalent vertices. The *vertex points* of R are precisely the images under π of leaves of $\tilde{\mathcal{F}}$ passing through a singularity of $\tilde{\mathcal{F}}$.

Step 4b (No vertex folds): Ruling out folds at high order vertices $v \in R$ requires a global argument (see Example 3.2.1 of a local vertex fold). The smallness hypothesis implies that, if two edges adjacent to v are folded together, then neither edge can contain a point representing the lift of a closed leaf of \mathcal{F} . This basically allows us to reduce the proof to the model case (see §5.2.3) where some leaf of \mathcal{F} is dense in S .

We now exploit the fact that we have a choice of conformal structures for S . Assuming ϕ folds at some vertex point, we can always choose a path of conformal structures S_t on $S = S_0$ so that the Hubbard-Masur differential on \tilde{S}_t (the holomorphic differential $\tilde{\Psi}_t$ whose vertical foliation projects to R) has simple zeroes for $t \neq 0$, and the edges which are folded together are represented on $\tilde{\mathcal{F}}_t$ by domains with a common one-dimensional frontier. As the harmonic map would again take the form $z \mapsto |\operatorname{Re} z|$ across this frontier, we see that $\tilde{\Psi}_t \neq \tilde{\Phi}_t$, where $\tilde{\Phi}_t$ is the Hopf differential for the harmonic map $f : \tilde{S} \rightarrow T$.

Hence there is a family of distinct \mathbf{R} -trees R_t and morphisms $\phi_t : R_t \rightarrow T$. These trees come from measured foliations on S which themselves come from interval exchange maps. But any nontrivial continuous variation in an interval exchange map forces a nontrivial variation in the tree T . As T is fixed this is impossible, so there can be no vertex folds.

1.2 Acknowledgements

We thank Mladen Bestvina for useful discussions, Howard Masur for all his help (including the idea and most of the details of the proof of Proposition 2.3), and the referee for numerous comments and corrections which greatly improved the paper. Misha Kapovich [Ka] also had the idea to use harmonic maps in the proof of Skora’s theorem.

2 Preliminaries

2.1 \mathbf{R} -trees

An \mathbf{R} -tree is a metric space T such that any two points in T are joined by a unique arc and every arc is isometric to a segment in \mathbf{R} . Let $[x, y]$ denote the unique (geodesic) segment from x to y in T .

We say that $x \in T$ is an *edge point* (resp. *vertex point*) of T if $T - \{x\}$ has precisely two (resp. more than two) components. An *edge* in T is a nontrivial, embedded segment $[x, y]$ in T .

A *morphism* of \mathbf{R} -trees is a map $\phi : T \rightarrow T'$ such that every nondegenerate segment $[x, y]$ has a nondegenerate subsegment $[x, w]$ for which $\phi|_{[x, w]}$ is an isometry.

The morphism $\phi : T \rightarrow T'$ *folds* at the point $x \in T$ if there are nondegenerate segments $[x, y_1]$ and $[x, y_2]$, with $[x, y_1] \cap [x, y_2] = \{x\}$, such that ϕ maps each segment $[x, y_i]$ isometrically onto a common segment in T' . It is easy to see that the morphism $\phi : T \rightarrow T'$ is an isometric embedding unless ϕ folds at some point $x \in T$.

By an *action* of Γ on T we mean an action by isometries. The action is *minimal* if Γ leaves no proper subtree of T invariant. For any Γ -action on T , there is a Γ -invariant proper subtree which is minimal (see, e.g., [CM]). Also, if A_γ is the isometry of T corresponding to $\gamma \in \pi_1 S$ for which $\inf_{y \in T} d(A_\gamma y, y) > 0$, then A_γ has an *axis* l_γ in T , i.e., an isometrically embedded line in T which is invariant under A_γ and which has the property that $x \in l_\gamma$ iff $d(A_\gamma x, x) = \inf_{y \in T} d(A_\gamma y, y)$. The proof is a straightforward consequence of the non-positive curvature of T .

Assumption: Henceforth we will assume, without loss of generality, that all actions are minimal.

We will need the following fact about small actions.

Lemma 2.1 *Let $\Gamma = \pi_1(S)$, S a closed surface of genus at least two. If the action $\Gamma \times T \rightarrow T$ is small then T must have a vertex point.*

Whenever speaking of vertex or edge points in a subtree of a tree T , we mean with respect to the space of directions in the subtree, not the ambient tree T .

Proof: If T has no vertex points then it is isometric to \mathbf{R} , so the action of Γ gives a representation $\psi : \Gamma \rightarrow \text{Isom}(\mathbf{R})$. As $\psi(\Gamma) < \text{Isom}(\mathbf{R})$ is virtually abelian and Γ is not solvable, it must be that the kernel of ψ contains two

noncommuting elements. But S is a closed surface of genus at least two, so sufficiently high powers of any two noncommuting elements of $\Gamma = \pi_1(S)$ generate a free group. This free group lies in the kernel of the action, in particular stabilizes any nondegenerate edge of T , a contradiction. \diamond

2.2 Holomorphic quadratic differentials

A holomorphic quadratic differential Φ on the Riemann surface S is a tensor given locally by an expression $\Phi = \varphi(z)dz^2$ where z is a conformal coordinate on S and $\varphi(z)$ is holomorphic. Such a quadratic differential Φ defines a measured foliation in the following way. The zeros $\Phi^{-1}(0)$ of Φ are well-defined and discrete. Away from these zeros, we can choose a canonical conformal coordinate $\zeta(z) = \int^z \sqrt{\Phi}$ so that $\Phi = d\zeta^2$. The local measured foliations ($\{\text{Re}\zeta = \text{const}\}$, $|d\text{Re}\zeta|$) then piece together to form a measured foliation known as the vertical measured foliation of Φ .

2.3 Actions dual to a measured foliation

Let (\mathcal{F}, μ) denote the vertical measured foliation of Φ . Lift it to a $\pi_1 S$ -equivariant measured foliation $(\tilde{\mathcal{F}}, \tilde{\mu})$ on \tilde{S} . The leaf space R of $\tilde{\mathcal{F}}$ is a Hausdorff topological space. Let $\pi : \tilde{S} \rightarrow R$ denote the projection. The leaf space R of the measured foliation $(\tilde{\mathcal{F}}, \mu)$ inherits a metric space structure from the measure μ : a geodesic segment $[x, y]$ in R is given by any path γ in \mathbf{H}^2 from a point in the leaf corresponding to x to a point in the leaf corresponding to y , such that γ is transverse to the leaves of the foliation $\tilde{\mathcal{F}}$. The distance $d_R(x, y)$ is simply $\mu(\gamma)$, and the metric space (R, d) is an \mathbf{R} -tree (see [MS2]). This tree is often not locally compact. For instance, when the leaves of the foliation on the surface S are dense, we can find sequences of arcs C_n transverse to the foliation with endpoints on singularities of $\tilde{\mathcal{F}}$ whose transverse measure $\mu(C_n)$ goes to zero, forcing the distance between the corresponding images of the (lifts of) vertices to also go to zero.

The action of Γ on \mathbf{H}^2 preserves μ , and so induces an isometric action of Γ on R . The stabilizers of this action are virtually cyclic, in particular are small.

The action of $\pi_1 S$ on \mathbf{H}^2 preserves μ , and so induces an isometric action of $\pi_1 S$ on R . The map $\pi : \tilde{S} \rightarrow R$ is equivariant with respect to this action.

2.4 The Hubbard-Masur Theorem

Holomorphic quadratic differentials on a Riemann surface S are related to classes of measured foliations via the Hubbard-Masur Theorem. To set the notation, fix a Riemann surface S and define a map $HM : \text{QD}(S) \rightarrow \mathcal{MF}(S)$ from the space $\text{QD}(S)$ of holomorphic quadratic differentials on S to the space $\mathcal{MF}(S)$ of equivalence classes of measured foliations on S that associates to $\Phi \in \text{QD}(S)$ the class of its vertical measured foliation. A fundamental result is

Theorem 2.2 (Hubbard-Masur [HM]) *HM is a surjective homeomorphism.*

Remark. A proof of Theorem 2.2 in the spirit of the current work can be found in [W2].

An alternative phrasing will be convenient for us. Let $Q(S) \rightarrow \text{Teich}(S)$ denote the bundle of holomorphic quadratic differentials over $\text{Teich}(S)$: here the fiber over $[S] \in \text{Teich}(S)$ is the space $\text{QD}(S)$ of quadratic differentials holomorphic with respect to a complex structure S in $[S]$. Let (\mathcal{F}, μ) denote a given measured foliation. Then the Hubbard-Masur Theorem shows that there is a well-defined section $\Psi_\mu : \text{Teich}(S) \rightarrow Q(S)$ which associates to $[S] \in \text{Teich}(S)$ the holomorphic quadratic differential $\Psi_\mu(S) \in \text{QD}(S)$ whose vertical measured foliation is measure equivalent to (\mathcal{F}, μ) .

2.5 Moving around in the Hubbard-Masur section

In this subsection we give a basic property of the section Ψ_μ .

Let S be a Riemann surface and let q be a holomorphic quadratic differential with vertical measured foliation (\mathcal{F}, μ) . Let p_0 be a singularity of q and let L be the maximal compact graph of singular vertical arcs through p_0 which connect p_0 to all the other singularities on the leaf through p_0 . Consider a neighborhood \mathcal{N} of L in S . We refer to the components $\{s_i\}$ of $\mathcal{N} - L$ as *sectors*, and say that two sectors *meet* along a (nondegenerate) arc if their closures intersect along that arc. We observe that there is a natural correspondence of sectors near a maximal singular arc L as above under Whitehead moves and isotopies of the foliation.

Proposition 2.3 (Sectors can be made adjacent) *Let S be a Riemann surface, let q be a holomorphic quadratic differential with vertical measured*

foliation (\mathcal{F}, μ) , and let L , p_0 and $\{s_i\}$ be as above. Choose any pair of sectors s_{i_1} and s_{i_2} from the list of sectors $\{s_i\}$. Then there is a Riemann surface S^* and a holomorphic quadratic differential q^* on S^* so that the vertical foliation of q^* is measure equivalent to (\mathcal{F}, μ) , and under the equivalence the sectors $s_{i_1}^*$ and $s_{i_2}^*$ (corresponding to s_{i_1} and s_{i_2} , respectively) meet along an arc.

Note that both q and q^* are in the image of the Hubbard-Masur section corresponding to (\mathcal{F}, μ) . A self-contained proof of Proposition 2.3 is given in the appendix.

3 Harmonic maps to trees

3.1 Definition of harmonic map

Given a Lipschitz continuous map $w : S \rightarrow (T, h)$ from a Riemann surface S to a locally finite metric tree T , we define the energy form to be the tensor

$$edz \otimes d\bar{z} = (\|w_*\partial_z\|_h^2 + \|w_*\partial_{\bar{z}}\|_h^2)dz \otimes d\bar{z}$$

Since the map w is Lipschitz, it is differentiable almost everywhere and bounded almost everywhere on closed balls; thus the form $edz \otimes d\bar{z}$ is defined almost everywhere with $edz \wedge d\bar{z}$ integrable over compacta. Note that even when T is not locally finite, the image of any closed ball in S is compact hence lies in a locally finite subtree of T , so this analysis applies.

Alternatively, for any conformal metric g on S with area form dA_g , the energy form may be expressed as follows. Choose an orthonormal frame $\{v_1, v_2\}$ at a point $z \in S$, and consider the pushforward vectors $\{w_*v_1, w_*v_2\}$. The energy form is the 2-form $\frac{1}{2}(\|w_*v_1\|_h^2 + \|w_*v_2\|_h^2)dA_g$, or alternatively $\frac{1}{2}\text{tr}_g(w^*h)dA_g$. The energy of the map w is $E = \int edz \wedge d\bar{z}$. The map w is a *harmonic map* if it is a minimum for this functional in its homotopy class of maps. We define the Hopf differential Φ for a map $w : S \rightarrow T$ by

$$\Phi = \phi dz^2 = 4 \langle w_*\partial_z, w_*\partial_z \rangle_h dz^2$$

Note that $\|\Phi\| = \|\Phi\|_{L^1} < 2E$.

3.2 Examples

In this subsection, we list some motivating examples of harmonic maps from Riemann surfaces to \mathbf{R} -trees. Each example will illustrate a principle we will later use.

1. The map $f(z) = \operatorname{Re}\{z^2\}$ as the most basic vertex fold.

The map $f(z) = \operatorname{Re}\{z^2\} : \mathbf{C} \rightarrow \mathbf{R}$ can be viewed as a harmonic map from the Riemann surface \mathbf{C} to the \mathbf{R} -tree \mathbf{R} . Observe that the preimage of the origin $O \in \mathbf{R}$ is the pair of intersecting lines $\{x = \pm y\}$ which divides \mathbf{C} into four sectors. The other level lines of a nonzero $r \in \mathbf{R}$ consist of hyperbolas $\{x^2 - y^2 = r\}$. The leaf space of the *connected* components of these level curves is the pair of coordinate axes. We conclude that the harmonic function $f(z)$ factors as a projection to the \mathbf{R} -tree of the coordinate axes followed by a vertex fold of each half-axis to its negative, which results in the image \mathbf{R} -tree \mathbf{R} .

2. Here is an example from [W3]: begin with the holomorphic differential $z^k dz^2$ on \mathbf{C} , whose vertical measured foliation is the set of curves $\{\operatorname{Re} z^{k+2} = c\}$. When we project along this foliation, we obtain a harmonic map to a tree with $k + 2$ prongs out of a single vertex.
3. Actions Dual to a Measured Foliation (\mathcal{F}, μ) , as given in §2.3.

Here the harmonic map is simply the projection along the vertical foliation of the properly normalized Hubbard-Masur differential for (\mathcal{F}, μ) . This characterization is independent of the particular Riemann surface chosen. We therefore observe the following.

Lemma 3.1 (When Φ and Ψ agree) *If the action of Γ is dual to a measured foliation (\mathcal{F}, μ) , then there is a well-defined Hopf differential section $\Phi : \operatorname{Teich}(S) \rightarrow Q(S)$ for π , and this section Φ agrees with the Hubbard-Masur differential section $\Psi : \operatorname{Teich}(S) \rightarrow Q(S)$ for \mathcal{F} .*

Proof: The lemma is effectively the content of [W2], which we now summarize; for complete details, see [W2]. (Later on, in §4, we shall give an independent proof of the existence of a harmonic map dual to a measured foliation.) As in §2.3, a measured foliation (\mathcal{F}, μ) on S lifts to an equivariant measured foliation $(\tilde{\mathcal{F}}, \tilde{\mu})$ on \tilde{S} ; we can project along the leaves to obtain an \mathbf{R} -tree (R, d) , with this construction also yielding an equivariant map $\pi_0 : \tilde{S} \rightarrow (R, d)$.

For each complex structure σ on S , we can minimize the energy in the equivariant homotopy class of π_0 obtaining [W2; proof of Prop. 3.1] a map $\pi : (S, \sigma) \rightarrow (R, d)$ whose Hopf differential $\Phi_R(\sigma)$ has vertical

foliation measure equivalent to (\mathcal{F}, μ) . (This argument is a straightforward application of Ascoli-Arzelà, with a crucial use of the axes of group elements in R to control (see [W2; Lemma 3.4]) the images of some points by elements of the minimizing sequence of maps.) This characterizes the differential uniquely [HM]; for a harmonic maps argument for this uniqueness, see [W2; §4]. Here the point is that both maps can be given as projections along minimal stretch foliations of Hopf differentials and the distance between the image points of the two maps can be equivariantly defined, and is a subharmonic function. [As the pullback of a smooth convex function off of the zeroes of the Hopf differentials, this pullback of the distance function is smooth and subharmonic (i.e. submean for balls of fixed radii) away from a discrete set of singularities and continuous across them; hence it is subharmonic everywhere. (Compare Proposition 3.2)] The maximum principle then applies, showing that the distance must be constant. Some geometry of the tree, in particular the fact that it has branches, forces that constant to vanish. Thus $\Phi_R(\sigma) = \Psi_\mu(\sigma)$. \diamond

An important part of our proof of Theorem 1.1 will be a converse (Lemma 5.3) to Lemma 3.1.

4. Actions Dual to the Measured Foliation of the Hopf Differential for an Arbitrary Harmonic Map from a Surface.

Let $f : S \rightarrow X$ be a harmonic map from the Riemann surface S to a metric space (possibly Riemannian, possibly singular). Let Φ denote the associated Hopf differential; we will see in §3.3, that this Hopf differential is a holomorphic quadratic differential on S . Lift this differential Φ to an (equivariant) differential $\tilde{\Phi}$ on the universal cover \tilde{S} , and consider its vertical (corresponding to the minimal stretch directions of \tilde{f}) measured foliation $(\tilde{\mathcal{F}}, \mu)$ and associated projection $\pi : \tilde{S} \rightarrow R$ to the leaf space (see §2.3).

Part of the content of the previous example is that the equivariant projection map $\pi : \tilde{S} \rightarrow R$ from \tilde{S} to R is harmonic.

Our proof of Skora's theorem involves a study of the relationship between the harmonic map we will construct from \tilde{S} to T and the associated harmonic map $\pi : \tilde{S} \rightarrow R$ from \tilde{S} to the leaf space R .

Remark. In [W4], we study the product harmonic map $(\tilde{f}, \pi) : \tilde{S} \rightarrow X \times \mathbf{R}$, and find that it is also conformal, after a slight homothety of

R. We also find that when X is smooth and two-dimensional, then this map is a stable minimal map.

5. A Harmonic One-Form with Integral Periods.

Project a square torus T^2 along its natural vertical foliation to the circle S^1 . This map is clearly harmonic. Now, there is a genus two surface S which is a branched cover over T^2 , and the 1-form dz lifts to yield a holomorphic one form on S . One can still project along leaves of this one form to a figure 8 which is a branched cover of the original S^1 . Hence by composing with the map to S^1 , we see that there is an associated harmonic map $f : S \rightarrow S^1$ and, via the one-form lifted from dz , an associated holomorphic one-form with integral periods on S .

As we vary S in $\text{Teich}(S)$ (say in a family S_t), the holomorphic one forms with the same A -periods (in the usual notation) varies continuously through one-forms, say, ω_t . It is interesting to consider the topology of the foliations \mathcal{F}_t that integrate $\ker \text{Re}\omega_t$.

The original surface S could be described as being constructed from a pair of cylinders C_1 and C_2 bounded by circles $\{S_{11}, S_{12}\}$ and $\{S_{21}, S_{22}\}$, respectively. Each S_{ij} is composed of two semicircles s_{ij}^t and s_{ij}^b . Now, the upshot of the notation is that S is defined by identifying s_{11}^b to s_{12}^t , s_{12}^b to s_{21}^t , s_{21}^b to s_{22}^t and s_{22}^b to s_{11}^t in the natural way, and the foliations parallel to the core curves of the cylinders become \mathcal{F}_0 .

A natural motion in Teichmüller space is to slightly rotate one of these cylinders against the other. This has the effect in our case of preserving the topology of \mathcal{F}_0 , up to a Whitehead move which splits the singularities from being locally a pair of coordinate axes as in Example 1 to a “double Y ” configuration. Of course, as \mathcal{F}_t is the foliation of a harmonic one-form, we see that this new synthetically constructed foliation inherited from the cylinders, which is actually the foliation of the Hubbard-Masur differential for (\mathcal{F}_0, μ_0) on S_t , is not \mathcal{F}_t . We conclude in this case that for the natural $\pi_1 S$ -action on \mathbf{R} defined via the one form ω , it is not the case that the Hopf differential section $\Phi : \text{Teich}(S) \rightarrow Q(S)$ agrees with the Hubbard-Masur differential section $\Psi : \text{Teich}(S) \rightarrow Q(S)$.

3.3 Local structure

R. Schoen has emphasized (see [Sc]) that a map for which the energy functional is stationary under reparametrizations of the domain has a Hopf dif-

ferential which is holomorphic: one uses suitable domain reparametrizations to show that the Hopf differential satisfies the Cauchy-Riemann equations weakly, and then Weyl's Lemma forces the Hopf differential to be (strongly) holomorphic. We observe that in this argument, the range manifold may be singular.

The vertical and horizontal foliations of the Hopf differential for $w : S \rightarrow T$ integrate the directions of minimal and maximal stretch of the gradient map dw , for smooth energy minimizing maps $w : S \rightarrow T$. As the image is one-dimensional, the harmonic map w is a projection along the minimal stretch direction. Further, if one normalizes the conformal coordinates in a domain that avoids the zeroes $\Phi^{-1}(0)$ of the Hopf differential Φ so that $\Phi = dz^2$ in that neighborhood, then one sees from the geometric definition of Φ above that the energy-minimizing map takes maximal stretch segments of measure ϵ to segments in T of length ϵ .

3.4 Effect on convex functions

A function defined on an \mathbf{R} -tree is *convex* if its restriction to every geodesic is convex in the classical sense. Recall that a function is *subharmonic* if it is submean, that is its value at any point is less than or equal to its average in a small ball around that point. Harmonic maps between Riemannian manifolds pull back convex functions to subharmonic functions (see, e.g. [I]). This important property extends to the case of \mathbf{R} -tree targets.

Proposition 3.2 *A harmonic map from a surface S to an \mathbf{R} -tree pulls back (germs of) convex functions to (germs of) subharmonic functions.*

Proof: We first argue that the map $\pi : \tilde{S} \rightarrow R$ to the leaf space R pulls back germs of convex functions on R to germs of subharmonic functions on \tilde{S} . Locally, the level set of the vertex, say $V \in R$, near a zero or pole of the Hopf differential $\tilde{\Phi}$ divides the neighborhood of the singularity into 'sectors', with the natural coordinate ζ mapping each sector conformally onto a neighborhood of zero in the upper half-plane (see [St;§7.1]). Under this mapping of a sector, the foliation of preimages of points in the tree R (in a sector) is taken to the horizontal foliation of the half-plane given by curves of the form $\{y = \text{const}\}$.

While it is not essential for the proof at this point, let us now consider a convex function F defined on the tree R near the point $V \in R$. This function pulls back to a function on a collection of sectors, which is constant on the horizontal levels $\{y = \text{const}\}$, and convex in y . Since any sector can be

taken to any other sector by an appropriate rotation, it is straightforward to see that this pullback is submean. (A more detailed argument is also given below, in the case of the tree T .)

With this in mind, let us return to the original case of the map $f : \tilde{S} \rightarrow T$. In the neighborhood of the singularity p of the Hopf differential (Φ) , we can regard our map as first projecting to a neighborhood of a vertex V in \mathbb{R} (this neighborhood of V is metrically a k -pronged star out of V , with one prong for each sector, by construction), followed by a map of R to T , in which several prongs of R map to a single prong of T , this prong of T emanating out of the image $v \in T$ of the vertex $V \in R$). Here we must have each prong taken injectively to a prong, because of the form of the map $f : \zeta \mapsto Re\zeta = \xi$ away from singularities of (Φ) .

In order to see why the map $f : \tilde{S} \rightarrow T$ pulls back germs of convex functions on T to germs of subharmonic functions on \tilde{S} , we make one crucial observation: we note that neighboring sectors on \tilde{S} must be taken to different prongs out of $f(p) \in T$; this is because a small arc transverse to the common boundary leaf of the pair of sectors is projected by f injectively into T , once again because in a neighborhood of such an arc, there are no singular points, and so the map f is of the form $\zeta \mapsto Re\zeta = \xi$. This implies that the pre-image under f of any given prong in T consists of at most half of the sectors abutting p .

Consider then a convex function F on the tree T near a point $p \in T$. This function pulls back to a function on a collection of sectors, which is constant on the horizontal levels $\{y = const\}$, and also convex in y . Suppose we have that $F(v) = 0$, so we need the mean value of f^*F to be non-negative on a disk D around p . Of course, since F is convex on $f(D) \subset T$, we know that F can be negative on at most one prong of $f(D)$, and must be non-negative on the other prongs. Moreover, since F is convex, if we average f^*F over a pair of sectors, one in which f^*F is nonpositive, and one in which f^*F is non-negative, we see that the sum of the averages must be non-negative. (To see this, conformally map each sector to a half-plane (say $\{y \geq 0\}$), and then glue the halfplanes together so that f^*F is convex in the coordinate y across the foliation.) Then we simply apply the observation of the previous paragraph to conclude that since f^*F is non-negative on at least half of the sectors, the average of f^*F on the union of sectors (i.e. the disk D) must be non-negative, as required. \diamond

4 Constructing a morphism from a geometric action to the given action

Let $\Gamma = \pi_1(S)$, S a closed surface of genus at least 2, and let $\Gamma \times T \rightarrow T$ be an action (not necessarily small) on an \mathbf{R} -tree T . In this section we construct an action of Γ on an \mathbf{R} -tree R which is dual to a measured foliation, and a Γ -equivariant morphism $\phi: R \rightarrow T$.

We will think of S as having a fixed hyperbolic structure, and so the universal cover \tilde{S} is the hyperbolic plane \mathbf{H}^2 . Since T is contractible, there is a Γ -equivariant Lipschitz continuous map $f_0: \mathbf{H}^2 \rightarrow T$. To be concrete, one may lift a triangulation of S , define f_0 by equivariance on the 0-skeleton of this triangulation, then extend (by contractibility of T) equivariantly to the 1-skeleton and 2-skeleton.

4.1 Finding the foliation using a harmonic map

Our first goal is to find a harmonic f map in the equivariant homotopy class of the Γ -equivariant continuous map $f_0: \mathbf{H}^2 \rightarrow T$ constructed at the beginning of §4. The harmonic map f will have the property that there is a measured foliation (\mathcal{F}, μ) on S so that every leaf of $\tilde{\mathcal{F}}$ gets mapped to a point under f . While it is possible to use the general theories of Korevaar-Schoen [KS] and Jost [J1, J2] on harmonic maps to nonpositively curved metric spaces, we will construct the harmonic map from elementary methods here.

To carry this out, we choose balls B_1, \dots, B_n on S so that:

- the balls are topologically trivial
- the restriction $f_0|_{\hat{B}_j}$ of f_0 to a lift \hat{B}_j of B_j is not a constant map for $j = 1, \dots, n$, and
- the set $\{B_1, \dots, B_n\}$ of balls is an open cover of S

Thus we have that each lift of B_j is disjoint from every other lift of B_j , and the union of all the lifts of all the balls $\{B_1, \dots, B_n\}$ covers \tilde{S} .

Now for each lift \hat{B}_1 of B_1 the image $f_0(\hat{B}_1)$ is a finite subtree of T . This follows from the fact that, for a basepoint $b_1 \in \hat{B}_1$, the image $f_0(\hat{B}_1)$ lies in a compact subset K of the space of directions at $f_0(b_1)$, and as this space of directions K is discrete (from the definition of \mathbf{R} -tree), it is also finite.

It is straightforward that there exists a unique harmonic map $\hat{f}_0: \hat{B}_1 \rightarrow T$ so that $\hat{f}_0|_{\partial\hat{B}_1} = f_0|_{\partial\hat{B}_1}$ (see the Appendix of [W1] for existence. To see uniqueness, apply the method of Cor. 3.2 of [W3] (see also §4 of [W1]):

the distance between any pair of solutions would be subharmonic on \widehat{B}_1 and vanishing on $\partial\widehat{B}_1$ – thus any pair of solutions coincide.). Moreover, if $h(\widehat{B}_1)$ is any other lift of B_1 , the uniqueness of the harmonic map then forces $\widehat{f}_0|_{h(\widehat{B}_1)} = h \circ \widehat{f}_0|_{\widehat{B}_1}$. Let ϕ_1 denote the map from the complete lift of B_1 to T . Then ϕ_1 , being nonconstant, also has the following properties:

- ϕ_1 is projection along the vertical measured foliation of its Hopf differential, and
- ϕ_1 is C^∞ on the interior of its domain (off of the zeroes of the Hopf differential of ϕ_1)

Set

$$f_1(z) = \begin{cases} \phi_1(z) & z \in \text{lift of } B_1 \\ f_0 & \text{otherwise} \end{cases}$$

Then f_1 is equivariantly homotopic to f_0 , and is a C^∞ equivariant projection (as above) along a measured foliation on the domain of ϕ_1 .

We repeat this procedure for lifts of the ball B_2 , using f_1 as the original map instead of f_0 . We then obtain a map f_2 . The situation is most interesting when $B_1 \cap B_2 \neq \emptyset$, as then the boundary values for ϕ_2 are defined by values of ϕ_1 , which may not agree with those of the original f_0 .

The main observation we need to make is the following: along most of a small neighborhood of $\partial\widehat{B}_2 \subset \widehat{B}_1$ we have that $\phi_1|_{\widehat{B}_1 \setminus \widehat{B}_2}$ and $\phi_2|_{\widehat{B}_2}$ extend to be a well-defined Lipschitz projection along a well-defined Lipschitz measured foliation. To see this note that $\phi_1|_{\partial\widehat{B}_1}$ is $C^{1,\alpha}$ and the measure of the vertical foliation of the Hopf differential of ϕ_1 is defined by distance between image points in T (see §2.5). As this also holds for $\phi_2|_{B_2}$, and $\partial\widehat{B}_1 \subset \widehat{B}_2$ is compact, the claim follows, except at (a discrete set of) places where the boundary values $f_1|_{\partial\widehat{B}_2}$ double back and result in small arcs in both \widehat{B}_1 and \widehat{B}_2 which close up in $\widehat{B}_1 \cup \widehat{B}_2$.

We follow the same procedures iteratively for lifts of B_3, \dots, B_n obtaining an equivariant map $f_n : \widetilde{S} \rightarrow T$ which is a projection along a Lipschitz measured foliation except for a discrete set of places where the leaves are closed and homotopically trivial.

In these places, we do an equivariant surgery to the map. For any region consisting of a union of concentric closed leaves, consider the closure of the largest such region. We then collapse the region to a segment which maps to the point defined by the boundary leaves. Call the new (collapsed) map $F : \widetilde{S} \rightarrow T$. It is evidently an equivariant map along a measured foliation with singularities that are k -pronged.

In [W2](Prop. 3.1), an elementary proof shows that the piecewise harmonic map $F : \mathbf{H}^2 \rightarrow T$ as above is equivariantly homotopic to a harmonic map $f : \mathbf{H}^2 \rightarrow T$. (This proof only requires that there are two elements of Γ whose axes in T have unbounded intersection. This property is much weaker than requiring that the action be small, but, for our purposes, follows from Lemma 2.1 above.) Moreover, attached to f is a holomorphic quadratic differential $\tilde{\Phi}_0$, the Hopf differential of f , with the following properties (see [W2; §2.2]):

- The vertical measured foliation of $\tilde{\Phi}_0$ is equivalent to $(\tilde{\mathcal{F}}, \tilde{\mu})$.
- The leaf space of the vertical foliation of $\tilde{\Phi}_0$ is R , and the vertical measure pushes down (say via $\pi : \mathbf{H}^2 \rightarrow R$) to the metric on R . This map is harmonic.
- On neighborhoods $B \subset \tilde{S}$ which are disjoint from $\tilde{\Phi}_0^{-1}(0)$, the map $f|_B$ agrees with $\pi|_B$ up to an isometry, while $\pi|_B$ is the projection $z \mapsto \operatorname{Re}z$ in a natural coordinate system.

This last property is quite important for the sequel, so we recall some the details from, for instance, [W1; p. 273] and [W2; p. 117]. By §2.2, there is a canonical coordinate $\zeta = \xi + i\eta$ so that $\Phi_0 = d\zeta^2$ on B . In its guise as a Hopf differential, of course, the definitions from §3.1 provides that $\Phi_0 = \|f_*\partial\xi\|^2 - \|f_*\partial\eta\|^2 + 2i \langle f_*\partial\xi, f_*\partial\eta \rangle$. Combining these two descriptions of Φ_0 and using that B is one-dimensional, we find that $f|_B$ is isometric to the map $\zeta \mapsto \operatorname{Re}\zeta = \xi$.

4.2 Definition of the morphism $\phi : R \rightarrow T$

Define an associated harmonic projection $\pi : \tilde{S} \rightarrow R$ via the construction in Example 3.2.4. Define also a map $\phi : R \rightarrow T$ by $\phi = f \circ \pi^{-1}$. We claim that ϕ is a morphism. To see this, let I denote a nondegenerate segment on the tree R ; we must find a non-degenerate subsegment $J \subset I$ for which $\phi|_J$ is an isometry. Well, as R is defined via projection $\pi : \tilde{S} \rightarrow R$, we can find an arc $\gamma \subset \tilde{S}$ with $\pi(\gamma) = I$. Here γ is quasi-transverse to $\tilde{\mathcal{F}}$ (in the sense of [HM], p. 231) and $\mu(\gamma) = \ell_R(I)$. On any subarc γ' of γ which avoids the zeros of $\tilde{\Phi}_0$, we may write (as we did at the end of the previous subsection) $\tilde{\Phi}_0 = dz^2$ for a choice of conformal coordinate in a neighborhood of γ' , and (again as in the previous subsection) f is an isometric submersion. Then for $J = \pi(\gamma') \subset \pi(\gamma) = I$, we have that $\phi|_J = f|_{\gamma'}$ which is an isometry by construction.

Finally, ϕ is surjective by the minimality hypothesis, and ϕ is equivariant since f is equivariant.

5 Proving that ϕ doesn't fold

5.1 No edge folds

It is a direct consequence of harmonicity of ϕ that ϕ does not fold at edge points. This is actually implicit in the proof above that ϕ is a morphism, but we give a slightly different proof in the next proposition, to which we will refer back several times in the sequel.

Proposition 5.1 (no edge-point folds) *The morphism $\phi : R \rightarrow T$ does not fold at an edge point $x \in R$.*

Proof: The pre-image of an edge point is a nonsingular leaf of the foliation $\tilde{\mathcal{F}}$. Any point z_0 on this leaf has a neighborhood \mathcal{N} foliated by non-singular arcs of leaves, and admits a conformal coordinate $z = x+iy$ with the foliation parallel to $\ker(\operatorname{Re}dz)$. If $\phi : R \rightarrow T$ were to fold at an edge point $\pi(z_0)$, then the harmonic map on the neighborhood \mathcal{N} would necessarily have the form $z \mapsto |\operatorname{Re}z|$, which is, of course, not harmonic.

Alternatively, using the same notation for the morphism ϕ folding at an edge point $\pi(z_0)$, letting p_0 denote the point $p_0 = \phi \circ \pi(z_0)$, we may apply the maximum principle to the function $h = f^*(-d_T(p_0, \cdot))$ on a neighborhood of z_0 . Here $-d_T(p_0, \cdot)$ is convex on $f(\mathcal{N})$, while $f^*(-d_T(p_0, \cdot))$ is not subharmonic on \mathcal{N} , contradicting Proposition 3.2. \diamond

Note that, at this point, we have shown that for any small Γ -action on an \mathbf{R} -tree T , there is an action on a tree R , dual to a measured foliation, and a Γ -equivariant morphism $\phi : R \rightarrow T$ which folds only at vertex points.

5.2 No vertex folds

In this section we will show that, when the action of Γ on T is small, the morphism ϕ is an isometry. A crucial feature of our argument will be a lemma that says that for actions $\Gamma \times T \rightarrow T$ which are not small, the choice of tree R is not uniquely determined.

Proposition 5.2 (no vertex folds) *If the action $\Gamma \times T \rightarrow T$ is small, then the morphism ϕ does not fold at a vertex point $v \in R$.*

The rest of this section is devoted to proving Proposition 5.2.

5.2.1 Vertex fold gives bad family

We begin with the following generalization of Lemma 3.1.

Lemma 5.3 *With notation as above, the following conditions are equivalent:*

1. *The action of Γ on T is dual to the measured foliation \mathcal{F} .*
2. *The morphism $\phi : R \rightarrow T$ is an isometry.*
3. *The Hubbard-Masur section $\psi_{\mathcal{F}} : \text{Teich}(S) \rightarrow Q(S)$ for \mathcal{F} is the same as the Hopf differential section $\Phi : \text{Teich}(S) \rightarrow Q(S)$ for T .*

Proof: As R is the dual tree of $\tilde{\mathcal{F}}$, it is clear that (2) implies (1). Lemma 3.1 states that (1) implies (3).

Now we prove that (3) implies (2). If ϕ is not an isometry, then ϕ must fold at some vertex point $v \in R$, by Proposition 5.1. Say $\phi(e_1) = \phi(e_2)$ for edges e_1, e_2 .

We may assume that R has vertices of valence at least 4: otherwise a vertex fold at a vertex $v \in R$ would be the fold of a 3-pronged star to an interval or half-interval. Thus the map f would restrict, in a neighborhood of the pre-image of $v \in R$, to a harmonic function on a disk where the corresponding Hopf differential has a 3-pronged zero. This is impossible, as harmonic functions are locally $\text{Re}(cz^k) + O(z^{k+1})$, for k an integer.

[Alternatively, the preimage of $\pi^{-1}(v)$ is a tree with discrete trivalent singularities. Near the singularities, this tree locally disconnects \tilde{S} into three sectors, with the harmonic map $f : \tilde{S} \rightarrow T$ folding the image of one sector onto the image of an adjacent sector. Yet as the sectors meet along an edge, the proof of Proposition 5.1 applies to yield a contradiction.]

Now consider the Hubbard-Masur section $\psi_{\mathcal{F}} : \text{Teich}(S) \rightarrow Q(S)$ for the foliation (\mathcal{F}, μ) . We are assuming that $\psi_{\mathcal{F}}(S)$ has zeroes of order at least two. Let s_1, s_2 be the sectors of \mathcal{F} corresponding to the edges e_1, e_2 . By Proposition 2.3 there is another quadratic differential $q' = \psi_{\mathcal{F}}(S')$ so that the sectors of the vertical foliation of q' corresponding to s_1, s_2 have closures which meet along an edge. Since by assumption $\psi_{\mathcal{F}}$ is the same as the Hopf differential section Φ , this is a contradiction: it violates the maximum principle for the map f' (defined as projection along the foliation of $\psi_{\mathcal{F}}(S') = \Phi(S')$), again as the map would locally have the form $z \mapsto |\text{Re } z|$.
 \diamond

Proof of Proposition 5.2: We now suppose, in expectation of reaching a contradiction, that the given action is not dual to a measured foliation, i.e. that ϕ is not an isometry. The equivalence of (1) and (3) in Lemma 5.3 then implies that there is a family $\{S_t\}, t \in \mathbf{R}$ of distinct Riemann surfaces for which $\psi_{\mathcal{F}}(S_t) \neq \Phi(S_t)$ for $t > 0$ and $\psi_{\mathcal{F}}(S_0) = \Phi(S_0)$ (here \mathcal{F} is defined by setting $\psi_{\mathcal{F}}(S_0) = \Phi(S_0)$ for some base point S_0), as we may as well assume for notational convenience that the two sections differ in a neighborhood of S_0 : here we get a *family* of surfaces where the sections $\psi_{\mathcal{F}}$ and Φ disagree rather than just a pair of points because the sections $\psi_{\mathcal{F}}$ and Φ are continuous.

To set notation, we rephrase this as follows: there is a family $\{S_t\}$ of distinct Riemann surfaces and corresponding Γ -equivariant harmonic maps $f_t : \tilde{S}_t \rightarrow T$, Hopf differentials $\tilde{\Phi}_t$, vertical foliations \mathcal{F}_t , and projections $\pi_t : \tilde{S}_t \rightarrow R_t$ to \mathbf{R} -trees with small Γ -actions and universal covering maps $p_t : \mathbf{H} \rightarrow S_t$ (choosing the notation so that $t = 0$ corresponds to the original action). Note that the trees R_t and morphisms ϕ_t are distinct, and that the foliations \mathcal{F}_t represent different points in $\mathcal{PMF}(S)$. If this were not true then $\psi_{\mathcal{F}}(S_{t_1}) = \Phi(S_{t_1})$ for some $t_1 > 0$, which would force the sections $\psi_{\mathcal{F}}$ and Φ to agree over S_{t_1} , contrary to the definition of the family S_t .

The heart of our argument is the case when the foliations \mathcal{F}_t are orientable and minimal. We begin with a reduction towards that case.

5.2.2 Some leaf is not closed

Let $e \in E \subset T$ denote a point of T which is not the image of a vertex in R_0 and which lies on the folded edge E of T . We consider the leaves of \mathcal{F}_0 containing $p_0 \circ f_0^{-1}(e) \subset S_0$. Since e lies on the folded edge E there are at least two of these. Each such leaf which is a closed curve represents a (conjugacy class of) element of Γ which fixes the edge $E \subset T$.

If each of these two leaves were closed, then they must represent the same element of $\pi_1(S)$: being simple closed curves, they do not represent powers of a common element of $\pi_1(S)$, hence some powers of these two elements in $\pi_1(S)$ must generate a free group since S is closed and hyperbolic; this free subgroup of $\pi_1(S)$ stabilizes E , contradicting smallness. But these two closed leaves are not even freely homotopic. If they were then they would bound an annulus A on S_0 . Since A has Euler-characteristic zero and the boundary components are leaves of \mathcal{F}_0 , no singularity of \mathcal{F}_0 lies in A . Hence the foliation \mathcal{F}_0 on this annulus would be by closed curves parallel to the boundary and the harmonic map $\pi|_A$ restricted to this annulus would map to an interval, with constant boundary values. This forces the map to be

everywhere constant, so that the Hopf differential vanishes on A , hence everywhere, an absurdity.

5.2.3 The model case

So we may assume that one of the components ℓ of $p_0 \circ f_0^{-1}(e)$ is not closed. Then consider a small arc $\alpha \subset S$ transverse to ℓ and to \mathcal{F}_0 . As the leaf ℓ is not closed, it is dense in a subsurface which we might as well take to contain α (after maybe reducing the size of α -see [St], §11). Indeed, we can find a finite number of edge points e_1, \dots, e_n so that the trajectories $p_0 \circ f_0^{-1}(e_i)$ have closure equal to all of S_0 .

Again, let α denote a small half-open arc transverse to \mathcal{F}_t on S_t with its endpoint on the singularity $q_0 \in S$; we also assume that $f_t(\alpha) \subset E$, the folded edge, and that α is chosen small enough to ensure that $\phi_t|_{\pi_t(\alpha)}$ is an isometry. By 5.2.2, we may assume that the nonsingular leaves through α are not closed on S_t . (If a non-singular leaf were closed, it would be contained in a neighborhood of non-singular closed leaves ([St], §9.3) and so there would be no leaf through α which would also be dense in a subsurface containing α . On the other hand, if every neighborhood of q_0 in α had regular closed leaves, since there are but a finite number of (maximal) ring domains (i.e. maximal neighborhoods of regular closed leaves) in \mathcal{F}_t , we see that a neighborhood of q_0 in α is contained in one of these ring domains. If this were true for all arcs α as above with $f_t(\alpha) \subset E$, we would be in the situation of 5.2.2, a contradiction.)

We begin with the model case of \mathcal{F}_t being orientable and *minimal*, i.e., every non-singular leaf is dense. The general case will follow from technical modifications to the proof in this case, but the essential ideas will be the same as in this model case.

Now, under the assumption that \mathcal{F}_t is minimal and orientable, we see that the first return map $P_t : \alpha \rightarrow \alpha$ determines an interval exchange map $\sigma_t : \alpha \rightarrow \alpha$ on α (see [St], p. 58). Moreover, one can reconstruct the measured foliation (\mathcal{F}_t, μ_t) directly from the interval exchange map $\sigma_t : \alpha \rightarrow \alpha$. We recall that this interval exchange map σ_t is determined by looking at the largest open subintervals $R_i(t)$ of α on which P_t is continuous. The endpoints $\{x_0(t) = q_0, x_1(t), \dots, x_N(t)\}$ of these subintervals are contained in singular leaves of \mathcal{F}_t , and hence (have lifts to \tilde{S} which) project to vertex points of the tree R_t .

We know that the set of vertex points in R_t is totally disconnected, as they are the image of the countable discrete set in \mathbf{H} of zeroes of Φ_t . It is also easy to see from this that the set of vertex points of the tree $\phi_t(R_t)$ in

T is also totally disconnected. We now assume, postponing the proof until the end of this subsection, that for each t there is some vertex point $v \in R_t$ such that $\phi_t(v)$ a vertex point.

Continuity argument: Our main observation is that, since the Γ -equivariant maps $f_t : \mathbf{H}^2 \rightarrow T$ are continuous in t , we see that if $f_t(\widetilde{x_i(t)})$ is a vertex in T , then as the vertices in T are a totally disconnected set, the family $f_t(\widetilde{x_i(t)})$ is constant in t . By the previous paragraph, there must be at least one endpoint $x_i(t)$ whose lift $\widetilde{x_i(t)}$ projects to a vertex in R_t . Since \mathcal{F}_t is minimal and f_t is equivariant, we have that $\Gamma f_t(\widetilde{x_i(t)}) = \Gamma f(\widetilde{x_i})$ is dense in $f(\widetilde{\alpha})$, for lifts $\widetilde{x_i(t)}$ and $\widetilde{\alpha}$ with $\widetilde{x_i(t)} \in \widetilde{\alpha}$. Letting $\Gamma_{x_i(t)} = \widetilde{\alpha} \cap \pi_t^{-1}(\Gamma \pi_t \widetilde{x_i(t)})$, we see that $f_t|_{\Gamma_{x_i(t)}}$ is constant in t , which forces $f_t(\widetilde{x_j(t)})$ to be constant in t for each j .

Since the measure of $\widetilde{\alpha}$ between consecutive vertices $\widetilde{x_i(t)}$ and $\widetilde{x_{i+1}(t)}$ (for $i = 0, \dots, N-1$) is determined by the distance $d_T(f_t(\widetilde{x_i(t)}), f_t(\widetilde{x_{i+1}(t)}))$ in the tree T , we see that these measures are also constant. Of course, after projecting from the cover \widetilde{S} to the surface S , we see that the endpoints $x_i(t) \subset \alpha$ are also constant in t .

Finally, observe that the first return maps $P_t : \alpha \rightarrow \alpha$ vary continuously in t on the interiors of the intervals $R_i(t)$ (and are affine there); hence, since the endpoints $x_i(t)$ are constant in t , we see that the maps P_t are constant in t as well. We conclude that the interval exchange maps σ_t are constant in t , so that $(\mathcal{F}_t, \mu_t) = (\mathcal{F}_0, \mu_0)$ after we reconstruct (\mathcal{F}_t, μ_t) from $\sigma_t : \alpha \rightarrow \alpha$. Hence we are done by Lemma 5.3.

Proof that some vertex point maps to a vertex point: Since this property is preserved under perturbations of the map, it is enough to prove the statement for some t .

Suppose this were not the case. Then every vertex point of every R_t maps to an edge point of T . Hence by Lemma 2.1 some edge point of each R_t maps to a vertex point of T . Since there are finitely-many Γ -orbits of vertex points, there exists $\delta_t > 0$ so that, on a Γ -fundamental domain of R_t , any such edge point of R_t has distance at least δ_t from any vertex point of R_t . For t small, we may take all $\delta_t > \delta$, for some fixed $\delta > 0$.

We first claim that by making a small perturbation in Teichmüller space from S to S_t we may assume that \mathcal{F}_t has a closed leaf λ representing an edge point $x \in R_t$ within a $\delta/6$ -neighborhood of some vertex point v_t ; necessarily, then there is a whole nondegenerate edge E containing x which is both within a distance $\delta/3$ of v_t and fixed by an element $g \in \Gamma$. This first claim follows

from essentially the same argument we used in the continuity argument above: take a small arc which abuts the vertex $v_t \in R_t$, and consider the image α on S_t of a lift of that arc. The foliation \mathcal{F}_t is determined by the interval exchange defined by the first return map on that arc α . In particular, perturbations of \mathcal{F}_0 are given by perturbations of that first return map, and we can find such a perturbation \mathcal{F}_0 so that \mathcal{F}_t has a closed leaf through α .

Now we make a few observations about our situation: since (1) all vertices are being folded away to edge points creating edges of radius at least $\delta/2$ from the image of v_t , but (2) on the surfaces S_t , no pair of adjacent sectors are having their R_t images folded together (by the argument late in the proof of Proposition 3.2), we see that for any point e' in any edge E' within $\delta/2$ of the image of v_t in T , we must have at least two distinct leaves on S_t whose lifts project to e' . But this contradicts smallness, as we showed in §5.2.2. Hence some vertex point maps to a vertex point.

Next we begin to loosen the hypotheses of the model case so as to eventually find ourselves in the general case, where \mathcal{F}_t may be non-orientable and have several minimal components.

5.2.4 Nonorientable case

Let us first drop the assumption that \mathcal{F}_t should be orientable. This is merely a matter of generalizing the correspondence between measured foliations (\mathcal{F}_t, μ_t) and interval exchange maps S_t . The idea here goes back to Strebel (see [St]). We regard one side of α as α^+ and the other side as α^- : if \mathcal{F}_t is orientable, then the rectangles $R_i(t)$ have one edge on α^+ and another on α^- , but if \mathcal{F}_t is not orientable, a rectangle may have both edges on, say, α^+ . Yet, if we now regard the first return map P_t as a map $P_t : \alpha^+ \cup \alpha^- \rightarrow \alpha^+ \cup \alpha^-$, we can consider an associated interval exchange map $S_t : \alpha^+ \cup \alpha^- \rightarrow \alpha^+ \cup \alpha^-$ from which we can reconstruct (\mathcal{F}_t, μ_t) . The endpoints $\{x_i(t)\}$ of the intervals $R_i(t)$ on $\alpha^+ \cup \alpha^-$ still (have lifts which) map continuously into the disconnected set of vertices (constant in t) of T , so then, as before the endpoints $\{x_i(t)\}$, and the map P_t , S_t are constant in t . We conclude that the measured foliations are also constant in t .

5.2.5 Breaking the model case into pieces

We come finally to the most general part, where we no longer require that \mathcal{F}_t is minimal. Then for \mathcal{F}_0 choose a collection of closed arcs $\alpha_1, \dots, \alpha_n$ which are transverse to \mathcal{F}_0 , and whose \mathcal{F}_0 -orbits both cover \mathbf{H}^2/Γ_0 and intersect at most along some compact singular leaves. At this point, we also require

the intervals α_i to have corresponding interval exchange maps for \mathcal{F}_0 which are either *irreducible*, i.e. we cannot (non-trivially) decompose $\alpha_i = \alpha'_i \cup \alpha''_i$ with the interval exchange map σ_i for α_i having a restriction $\sigma_i|_{\alpha'_i}: \alpha'_i \rightarrow \alpha'_i$ which preserves the proper subinterval α'_i , or correspond to a single cylinder in \mathcal{F}_0 , so that the interval exchange is the identity on a single cylinder.

We claim that the measured foliation \mathcal{F}_t on the whole surface $\Sigma(t)$ is constant in t . This will give a contradiction by Lemma 5.3, proving the theorem.

Let $\Sigma_i(t)$ be the subsurface of S_t obtained by taking the closure of the orbit of α_i along the leaves of \mathcal{F}_t . Our restrictions on $\{\alpha_i\}$ have the effect of forcing either $\Sigma_i(0)$ to be a cylinder or a surface on which \mathcal{F}_0 is minimal. We observe that the argument given earlier for the cases where \mathcal{F}_0 was minimal on the closed surface \mathbf{H}^2/Γ_0 continue to hold for the case where \mathcal{F}_0 is minimal on $\Sigma_i(0)$. In particular, for $\Sigma_i(0)$ a subsurface with almost every leaf dense, we see that the interval exchange maps $\sigma_i(t)$ must be constant in t . Yet, it is part of the basic construction of measured foliations from interval exchange maps that the topology of $\Sigma_i(t)$ (as well as \mathcal{F}_t) is determined from the map $\sigma_i(t)$ (see, e.g., [Ma1]). Thus, as $\sigma_i(t) = \sigma_i(0)$, we see that $\Sigma_i(t)$ is homeomorphic to $\Sigma_i(0)$.

Now each boundary circle of each $\Sigma_i(t)$ is a leaf of the foliation on that subsurface. This leaf may be taken to be singular as it would otherwise be an interior leaf of a cylinder of non-singular homotopic leaves, counter to the construction of $\{\alpha_i\}$. Thus the continuity argument also shows that the foliations on the cylindrical subsurfaces $\Sigma_j(t)$ are constant in t . Hence the measured foliation on each subsurface $\Sigma_i(t)$ is constant in t . Finally, whenever two subsurfaces $\Sigma_i(t)$ and $\Sigma_j(t)$ have a common boundary component $C(t)$, the continuity argument shows that $C(t)$ cannot become a cylinder at any time t as this would require the single vertex $f_t(C(t))$ to continuously deform into a non-trivial family of pairs of vertices, an absurdity. So we see that the identification of all the boundary components of all the $\Sigma_i(t)$ are constant over t , so that \mathcal{F}_t is constant.

6 Appendix

This appendix is dedicated to a proof of Proposition 2.3, which is partly implicit in [HM] and partly a “folklore theorem”. We provide here an elementary, geometric, and self-contained proof due almost entirely to Howard Masur (personal communication), who graciously permitted us to reproduce it here.

The proof can be reduced to the following claim: If either

1. q has a pair $\{z_1, z_2\}$ of distinct zeros connected by an arc A of a leaf of its vertical foliation, or
2. q has a k -pronged singularity at z_3 , and 2 arbitrary sectors s_1, s_2 of this k -prong are specified,

then there is a Riemann surface S^* and a holomorphic quadratic differential q^* on S^* so that the vertical foliation of q^* is measure equivalent to (\mathcal{F}, μ) and

1. (in case (1) above) the zeros of q^* corresponding to $\{z_1, z_2\}$ coincide, or
2. (in case (2) above) the images of the sectors s_1, s_2 under the equivalence of foliations meet along an arc.

The proposition follows from the claim as follows. First apply (1) above to get s_1 and s_2 as sectors abutting a common singularity z . Then apply (2) above and we are done.

We are left to prove the claim.

Single cylinder case. We first prove the claim for *Jenkins differentials*, i.e. those differentials whose vertical foliations are but one foliated open (right Euclidean) cylinder C with all singularities lying in $\partial\overline{C}$. Here S can be thought of as an identification space $\pi : \overline{C} \rightarrow S$, with identifications being made on $\partial\overline{C}$. Let C_1, C_2 denote the 2 components of $\partial\overline{C}$. Note that the graph L lies in $\partial\overline{C}$, and all the singularities on a single component of $\partial\overline{C}$ are connected by L . Moreover, there are natural correspondences between topological or geometric operations on the surface S and topological or geometric operations on \overline{C} . This means that if we continuously deform C to another right Euclidean cylinder C^* (so that there is a canonical correspondence of identifications on $\partial\overline{C^*}$), then the canonical quadratic differential q^* on C^* (defined so that the metric $|q^*|$ agrees with the Euclidean metric on C^* and all of whose vertical leaves are parallel to $\partial\overline{C^*}$) descends to a quadratic differential q^* on the identified surface S^* with the vertical foliation of q^* on S^* being Whitehead equivalent to the vertical foliation of q on S .

To prove (1) and (2) above, we will first perform the desired operation on C to obtain a new Euclidean cylinder C^* with canonically determined quadratic differential q^* as above. The important thing to check in each case is that we can do this so that the resulting Euclidean lengths $\ell(C_1^*)$ and

$\ell(C_2^*)$ of the 2 components of $\partial\overline{C}^*$ are equal. This immediately implies that the identification π determines an identification $\pi^* : C^* \rightarrow S^*$ to a Riemann surface S^* , and that the canonical quadratic differential on C^* descends to a quadratic differential q^* on S^* . By construction q^* has vertical foliation measure equivalent to that of the vertical foliation of q .

Consider case (1). Let $A_1, A_2 \subset \partial\overline{C}$ denote the 2 components of $\pi^{-1}(A)$, where we recall that A is the arc of the vertical foliation we wish to collapse. Note that $\ell(A_1) = \ell(A_2)$.

Case 1a: A_1 and A_2 lie in different components of $\partial\overline{C}$. In this case contract both A_1 and A_2 to a point to give a Euclidean cylinder C^* . Since C_1^* and C_2^* have the same Euclidean length, so we are done by the above.

Case 1b: A_1, A_2 lie on the same component of $\partial\overline{C}$. First note that the arcs of L have preimages in $\partial\overline{C}$ which come in pairs, as neighborhoods of the arcs on the identification space have full neighborhoods, while neighborhoods of arcs on $\partial\overline{C}$ have only half-neighborhoods. Hence since A_1, A_2 lie on the same component of $\partial\overline{C}$, we must be able to find some collection of pairs of arcs on the other component the sum of whose lengths is at least that of the sum of the lengths of A_1 and A_2 . (This is just a pigeon-hole principle: the arcs come in pairs whose lengths are equal and for which total lengths of all the arcs are the sum of the lengths of the boundary components of $\partial\overline{C}$, yet each of these boundary components have the same lengths, so the fact that A_1 and A_2 contribute solely to one component of $\partial\overline{C}$ forces some other family of pairs to contribute at least as much solely to the other component of $\partial\overline{C}$.) Thus we act as before, contracting A_1 and A_2 on one component of $\partial\overline{C}$ and simultaneously some other pairs of arcs the same amount on the other component of $\partial\overline{C}$. It is quite important here that the contraction of the other components has no effect on our claim or our purpose; the proof of the first part of the claim concludes as before.

Now to prove part (2) of the claim. Under the identification map $\pi : C \rightarrow S$, each sector $s_i, i = 1, 2$ has a unique pre-image on C as a neighborhood U_i of a vertex v_i .

Case 2a: v_1 and v_2 lie on different components of $\partial\overline{C}$.

We split the vertex v_1 into a pair of vertices $v_{1,1}$ and $v_{1,2}$ connected by an arc A_1 , and we split the vertex v_2 into a pair of vertices $v_{2,1}$ and $v_{2,2}$ connected by an arc A_2 of the same length as A_1 . We then re-identify the cylinder as before, with the only changes being that instead of identifying v_1 to v_2 , we send A_1 isometrically onto A_2 (there is a unique way to do this

which preserves the ordering of the sectors). The resulting surface gives S^* and q^* as required.

Case 2b: v_1 and v_2 lie on the same component, say C_1 , of $\partial\overline{C}$. Split v_1 and v_2 as in Case 2a. We now do a further deformation to make $\ell(C_1^*) = \ell(C_2^*)$.

If for some compact singular arc $B \subset L$, we have both components of $\pi^{-1}(B)$ lying on C_2 , then by lengthening B we could achieve $\ell(C_1^*) = \ell(C_2^*)$. If this isn't true, then by the pigeon-hole principle, for each such B we know $\pi^{-1}(B)$ has one component on C_1 and one on C_2 .

Now observe that any singularity on the surface S with, say, k sectors, admits a cyclic ordering of these sectors s_1, \dots, s_k (where the closure of s_2 meets the closure of s_1 on one side and the closure of s_3 on the other side, and so on). Since we are in a case where each edge incident to a singularity on S has preimages on both boundary components C_1 and C_2 of $\partial\overline{C}$, and since sectors have preimages near components of $\partial\overline{C}$ where their bounding arcs have preimages, we see that the sectors s_1, \dots, s_k also alternate between having preimages in C_1 and in C_2 . This implies that all of the singularities on the surface S have an even number of sectors.

We now claim that there are vertices w_1, w_2 in C_2 which are still identified by the identification rules, even after the splitting of v_1 and v_2 . (Here the subtlety is that by first splitting v_1 and v_2 , we have changed the identification rules, and hence the orbits of identified vertices on $\partial\overline{C}$. Our vertices w_1 and w_2 must not only then correspond to each other by the original identification rules, but they must also lie in the same new orbit of vertices on $\partial\overline{C}$, after the splitting of v_1 and v_2 .) This finishes the proof of case 2b since we then split w_1 and w_2 to make $\ell(C_1^*) = \ell(C_2^*)$.

To see that there are such vertices w_1 and w_2 , we recall that the total multiplicity of zeroes of a holomorphic quadratic differential on a Riemann surface S of genus g is equal to $4g - 4$ (Riemann-Roch). Thus, since in the case under consideration all of the singularities have an even number of sectors (and hence an even order of zero), we see that there is either one singularity z_0 with at least six sectors, or several singularities which all have at least four sectors. In the first case we see that any initial splitting of z_0 (by splitting a pair of vertices v_1 and v_2 on the same component C_1 of $\partial\overline{C}$) would leave a topological foliation with two singularities of which at least one would have four sectors, with two of those sectors having preimages on C_2 : we would then split the vertices of those sectors, say w_1 and w_2 to finish the case. In the second case, there is at the outset a singularity on S whose preimages do not include v_1 and v_2 , and which has at least a pair of sectors

with preimages on C_2 , as desired.

General case. We prove the general case by the now standard technique of approximating. By [Ma2] we may approximate q by a sequence $\{q_n\}$ of Jenkins differentials on S . In case (1), let A denote the arc of the vertical foliation of q which we wish to contract. As q_n approximates q , there is an arc $A_n \subset \partial\overline{C}_n$ which approximates A . Furthermore, there is a contractible neighborhood U on the underlying differentiable surface which is a neighborhood of the arc A and all arcs A_n for n sufficiently large.

Now, the Riemann surface S is an identification space of each cylinder \overline{C}_n , with identifications being made on $\partial\overline{C}_n$. As in the single cylinder case above, we can form new Riemann surfaces S_n^* equipped with quadratic differentials q_n^* by contracting the arcs $A_n \subset \partial\overline{C}_n$ and identifying as before; here the arc A_n on S_n bounded by a pair of low order zeros is replaced on S_n^* by a single higher order zero, say z_n^* .

The important thing to notice about this operation is that the complement $V = U^c$ of the neighborhood $U \subset S$ is approximated by the closure of an open set V_n on S_n^* which only avoids a small neighborhood of the high order zero z_n^* ; moreover, the conformal structures on V_n compare uniformly to the conformal structure on V , hence to each other. Hence, by passing to a subsequence if necessary, we have that S_n^* converges in (the interior of) Teichmüller space to a Riemann surface S^* . It also follows that q_n^* converges to a holomorphic quadratic differential, say q^* , on S^* ; as q_n approximates q and q_n^* is measure equivalent to q_n , we see that q^* is measure equivalent to q . Moreover, as the foliation of q_n^* results from a Whitehead move applied to the foliation of q_n (which contracts A_n to a point), the foliation of q^* is obtained from the foliation of q via a Whitehead move which contracts A .

Case (2) is virtually identical: we still have uniform convergence of the conformal structures outside the pair(s) of neighborhood(s) of the vertex (or arcs) we are splitting to pairs of vertices connected by an arc.

◇

References

- [CM] M. Culler and J. Morgan, *Group actions on \mathbf{R} -trees*, Proc. London Math. Soc., Vol. 3, No.55 (1987), p.571-604.
- [GiSh] H. Gillet and P. Shalen, *Dendrology of Groups in Low \mathbf{Q} -Ranks*, *J. Diff Geom.*, Vol. 32 (1990), p.605-712.

- [GrSc] M. Gromov, R. Schoen, Harmonic Maps into Singular Spaces and p -adic Superrigidity for Lattices in Groups of Rank One, *Inst. Hautes Etudes Sci. Publ. Math.* No. 76 (1992) p. 165-246.
- [Ha] A. Hatcher, Measured Lamination Spaces for Surfaces from the Topological Viewpoint, *Top. Appl.*, Vol.30 (1988), p.63-88.
- [HM] J. Hubbard and H. Masur, Quadratic Differentials and Foliations, *Acta Math.*, Vol.142 (1979), p. 221-274.
- [I] T. Ishihara, A Mapping of Riemannian Manifolds which Preserves Harmonic Functions, *J. math. Kyoto Univ.*, Vol. 19 (1979), p. 215-229.
- [J1] J. Jost, Equilibrium Harmonic Maps Between Metric Spaces, *Calc. Var. Part. Diff. Eqns.*, Vol. 2 (1994) p. 173-204.
- [J2] J. Jost, Convex Functionals and Generalized Harmonic Maps into Spaces of Non-Positive Curvature, *Comm. Math. Helv.*, Vol. 70 (1995) p. 659-673.
- [Ka] M. Kapovich, Hyperbolic Manifolds and Discrete Groups: Notes on Thurston's Hyperbolization, University of Utah Lecture Notes, 1994.
- [KS] N. Korevaar and R. Schoen, Sobolev Spaces and Harmonic Maps for Metric targets, *Comm. Anal. Geom.* Vol. 1 (1993), p. 561-659.
- [Ma1] H. Masur, Interval Exchange Transformations and Measured Foliations, *Ann. Math.* Vol. 115 (1982), p. 169-200.
- [Ma2] H. Masur, The Jenkins-Strebel differentials with one cylinder are dense, *Comment. Math. Helv.*, Vol.54, No.2 (1979), p.179-184.
- [MO] J. Morgan and J.P. Otal, Relative growth rates of closed geodesics on a surface under varying hyperbolic structures, *Comm. Math. Helv.*, No. 68 (1993), p. 171-208.
- [MS1] J. Morgan and P. Shalen, Degenerations of Hyperbolic Structures, I: Valuations, Trees, and Surfaces, *Annals of Math.*, Vol. 120 (1984), p.401-476.
- [MS2] J. Morgan and P. Shalen, Free actions of surface groups on \mathbf{R} -trees, *Topology*, Vol.30, No.2, pp.143-154, 1991.

- [Ot] J.P. Otal, Le Theoreme D'Hyperbolisation pour les Varieties Fibrees de Dimension 3, *Asterisque*, Vol. 235, 1996.
- [Sc] R. Schoen, Analytic Aspects of the Harmonic Map Problem in *Seminar on Non-linear Partial Differential Equations*, (S.S. Chern, ed.), MSRI Publ. 2, Springer-Verlag, New York 1983, p. 321-358.
- [Sh] P. Shalen, Dendrology of groups: An introduction, in *Essays in group theory*, ed. by S.M. Gersten, MSRI publications #8, Springer, 1987.
- [Sk1] R.K. Skora, Research Announcement: Splittings of Surfaces, *Bulletin A.M.S.*, Vol. 23 (1990), p.85-90.
- [Sk2] R.K. Skora, Splittings of Surfaces, *Journal A.M.S.*, Vol.9, No.2, April 1996.
- [St] K. Strebel, Quadratic Differentials, Springer-Verlag, Berlin, 1984.
- [W1] M. Wolf, On the existence of Jenkins-Strebel differentials using harmonic maps from surfaces to graphs, *Ann. Acad. Sci. Fenn. Ser. A. I. Math.* **20** (1995), 269-78.
- [W2] M. Wolf, On Realizing Measured Foliations via Quadratic Differentials of Harmonic Maps to \mathbf{R} -trees, *J. D'Analyse Math.* **68** (1996), 107-120.
- [W3] M. Wolf, Harmonic Maps from Surfaces to \mathbf{R} -trees, *Math. Zeit.*, Vol. 218 (1995), p. 577-593.
- [W4] M. Wolf, Measured Foliations and Harmonic Maps from Surfaces, *J. Differential Geom.* Vol49 (1998), p. 437-467.

Benson Farb:
 Dept. of Mathematics, University of Chicago
 5734 University Ave.
 Chicago, Il 60637
 E-mail: farb@math.uchicago.edu

Michael Wolf:
 Dept. of Mathematics, Rice University
 Houston, TX 77251
 E-mail: mwolf@math.rice.edu