ON REALIZING MEASURED FOLIATIONS VIA QUADRATIC DIFFERENTIALS OF HARMONIC MAPS TO R-TREES

Michael Wolf*

Department of Mathematics
Rice University
Houston, TX 77251
email: mwolf@math.rice.edu

Abstract. We give a brief, elementary and analytic proof of the theorem of Hubbard and Masur [HM] (see also [K], [G]) that every class of measured foliations on a compact Riemann surface $\mathcal{R}$ of genus g can be uniquely represented by the vertical measured foliation of a holomorphic quadratic differential on $\mathcal{R}$. The theorem of Thurston [Th] that the space of classes of projective measured foliations is a $6g-7$ dimensional sphere follows immediately by Riemann-Roch. Our argument involves relating each representative of a class of measured foliations to an equivariant map from $\tilde{\mathcal{R}}$ to an $\mathbf{R}$-tree, and then finding an energy minimizing such map by the direct method in the calculus of variations. The normalized Hopf differential of this harmonic map is then the desired differential.

§1. Introduction.

Measured foliations on a surface arise in two distinct ways. Complex analysts have studied the trajectory structure (see [Str]) of holomorphic quadratic differentials on a Riemann surface $\mathcal{R}$ for decades, and they occupied a prominent role in classical Teichmüller theory. However, measured foliations acquired both a fresh perspective and altogether new importance with Thurston's definition of the space $\mathcal{MF}$ of (equivalence classes of) measured foliations as a completion of the set $\mathcal{S}$ of simple closed curves on $\mathcal{R}$ and with his use of them in his pioneering works in low dimensional geometry. Our focus here is with the fundamental theorem [HM] of Hubbard and Masur, which said that every class of measured foliations can be uniquely represented by the vertical measured foliation of a holomorphic quadratic differential on $\mathcal{R}$, and which began to allow for an interplay between analytical and the synthetic approaches to Riemann surface theory.

The argument of Hubbard and Masur proved substantially more than the theorem stated above as they also described the differentiability of sections of an appropriate bundle over Teichmüller space. Because of these rather subtle complications, their proof was

*Partially supported by NSF grant DMS9300001; Alfred P. Sloan Research Fellow

somewhat lengthy and intricate. Proofs by Kerckhoff [K] and Gardiner [G] were quite brief, but relied on some of the structure of the space $\mathcal{MF}$.

Here we give a new proof of the theorem of Hubbard and Masur which is both brief and elementary, involving only the direct method in the calculus of variations, some straightforward real analysis, Weyl's lemma, the definition of an equivalence class of a measured foliation, and one slightly novel object — a real tree.

Here is a sketch of the argument, with the definitions deferred until the next section. We begin with a C^k measured foliation $(\mathcal{F}, \mu)$ on a Riemann surface $\mathcal{R}$; we lift $(\mathcal{F}, \mu)$ to a $\pi_1\mathcal{R}$-equivariant measured foliation $(\widetilde{\mathcal{F}}, \tilde{\mu})$ on the universal cover $\widetilde{\mathcal{R}}$. Then we project $\widetilde{\mathcal{R}}$ along the leaves of $\widetilde{\mathcal{F}}$ to obtain a $\mathbf{R}$-tree T; the measure $\tilde{\mu}$ projects to give a metric d on T. Any other C^k measured foliation $(\mathcal{G}, \nu)$ on $\mathcal{R}$ which is equivalent to $(\mathcal{F}, \mu)$ has a lift which projects to (T, d), and as the projection is $\pi_1\mathcal{R}$-equivariant and $(\mathcal{G}, \nu)$ is smooth, we can measure the (finite) total energy of the projection. Within the class of C^k measured foliations $(\mathcal{G}, \nu)$ equivalent to $(\mathcal{F}, \mu)$, we seek one of least energy by considering a sequence $\{(\mathcal{G}_n, \nu_n)\}$ of measured foliations whose energy tends to the infimum of energy. The bound on total energy and a length-energy argument give a uniform bound on the modulus of continuity of the projections in the sequence, and some geometry of the projection then bounds the images of the projections (associated to $(\mathcal{G}_n, \nu_n)$) on compacta in $\widetilde{\mathcal{R}}$ to compacta in (T, d): the Ascoli theorem assures uniform convergence to a limiting projection $f : \widetilde{\mathcal{R}} \to (T, d)$. The normalized Hopf differential $4\,\langle f_z, f_z\rangle\, dz^2$ is then the desired quadratic differential: it is weakly holomorphic as f is energy minimizing, thus holomorphic by Weyl's lemma.

We have several motivations for presenting what is at least the fourth proof of this result. First, one perspective on the Hubbard-Masur theorem is that it is a Hodge-like theorem for holomorphic quadratic differentials, and so we would like to present an argument similar in spirit and methods to that of the Hodge theorem for harmonic one-forms. Second, a number of authors (e.g. [W1], [J], [Tr]) have endeavored recently to give a harmonic maps treatment of Teichmüller theory; not only was the Hubbard-Masur theorem a missing part of this theory, but it was used in the harmonic maps treatment of the Thurston compactification of Teichmüller space [W1]. That treatment can now be done more or less completely within a harmonic maps framework. Third, the notion of holomorphic quadratic differentials arising from maps out of a Riemann surface rather than into a Riemann surface, and the natural role of real trees in the subject seems novel. Finally, the author found the deep and subtle investigations of [HM] into families of quadratic differentials on families of Riemann surfaces easier to appreciate once he had understood that the Hubbard-Masur theorem could be given an analytic proof involving only a single Riemann surface.

Harmonic maps into non-locally compact metric spaces have attracted considerable interest of late, greatly spurred on by the foundational work of Korevaar and Schoen [KS] (see also [W2] for an earlier relevant example). The paper [KS] provided for the existence of energy minimizing maps between a universal cover $\widetilde{M}$ and a non-positively curved metric space X, if the homotopy class of the map contained a $\pi_1 M$-equivariant map, the domain $\widetilde{M}$ was the universal cover of a compact manifold M, and either $\widetilde{M}$ had non-empty boundary or the homotopy class of maps was the lift of a class of maps between compact spaces. In order to give a self-contained and elementary approach to the Hubbard-Masur theorem, we exploit some of the specific characteristics of the $\mathbf{R}$-tree targets and Riemann surface domains, and we then do not need to cite the more general paper [KS] in our argument. Moreover, while a large part of the machinery of [KS] could easily be applied in our case, the precise situation of an $\mathbf{R}$-tree target with a $\pi_1\mathcal{R}$ action is not completely treated in their first paper [KS]; we develop some elementary methods of comparing axes of the action of group elements which may have more general applications.

We organize this article as follows. In the next section, we define our terms, give the relevant background and state the theorem. In the third section, we give our proof of the existence part of the theorem; this is the bulk of the paper. In the final section, we give a proof of uniqueness from the same perspective as the proof of existence, and note that the Thurston theorem on the topology of the space $\mathcal{PF}$ of projective measured foliations follows easily from our considerations.

§2. Background and Notation.

Let $\mathcal{R}$ be a compact Riemann surface of genus $g > 1$.

2.1 **Measured Foliations**. A C^k measured foliation on $\mathcal{R}$ with singularities $z_1, \ldots, z_l$ of order $k_1, \ldots, k_l$ (respectively) is given by an open covering $\{U_i\}$ of $\mathcal{R} - \{z_1, \ldots, z_l\}$ and open sets $V_1, \ldots, V_l$ around $z_1, \ldots, z_l$ (respectively) along with C^k real valued functions v_i defined on U_i s.t.

(i) $|dv_i| = |dv_j|$ on $U_i \cap U_j$

(ii) $|dv_i| = |\operatorname{Im}(z - z_j)^{k_j/2} dz|$ on $U_i \cap V_j$

Evidently, the kernels $\ker dv_i$ define a C^{k-1} line field on $\mathcal{R}$ which integrates to give a foliation $\mathcal{F}$ on $\mathcal{R} - \{z_1, \ldots, z_l\}$, with a $k_j + 2$ pronged singularity at z_j. Moreover, given an arc $A \subset \mathcal{R}$, we have a well-defined measure $\mu(A)$ given by

$$\mu(A) = \int_A |dv|$$

where $|dv|$ is defined by $|dv|_{U_i} = |dv_i|$. An important feature of this measure is its translation invariance: that is, if $A_0 \subset \mathcal{R}$ is an arc transverse to the foliation $\mathcal{F}$, and if we deform A_0 to A_1 via an isotopy that maintains the transversality of the image of A_0 at every time, then $\mu(A_0) = \mu(A_1)$. For $[\gamma]$ a class of curves, define $i([\gamma], \mu) = \inf_{\gamma \in [\gamma]} i(\gamma, \mu)$ where the infimum is taken over all curves γ in $[\gamma]$.

2.2 **The space** $\mathcal{MF}$. Two measured foliations $(\mathcal{F}, \mu)$ and $(\mathcal{G}, \nu)$ are said to be equivalent if after possibly some Whitehead moves on $\mathcal{F}$ and $\mathcal{G}$, there is a self-homeomorhism of $\mathcal{R}$ which takes $\mathcal{F}$ to $\mathcal{G}$, and μ to ν. Here a Whitehead move is the transformation of one foliation to another by collapsing a finite arc of a leaf between two singularities, or the inverse procedure (see [FLP]). The space of equivalence classes of measured foliations is denoted $\mathcal{MF}$.

2.3 **Holomorphic Quadratic Differentials**. A holomorphic quadratic differential Φ on the Riemann surface $\mathcal{R}$ is a tensor given locally by an expression $\Phi = \varphi(z)dz^2$ where z is a conformal coordinate on $\mathcal{R}$ 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; 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 $(\{\operatorname{Re}\zeta = \text{const}\}, |d\operatorname{Re}\zeta|)$ then piece together to form a measured foliation known as the vertical measured foliation of Φ.

2.4 **R-trees**. Let $(\mathcal{F}, \mu)$ denote the vertical measured foliation of Φ; lift it to a $\pi_1\mathcal{R}$-equivariant measured foliation $(\widetilde{\mathcal{F}}, \tilde{\mu})$ on $\widetilde{\mathcal{R}}$. The leaf space T of $\widetilde{\mathcal{F}}$ is a Hausdorff topological space; let $\pi : \widetilde{\mathcal{R}} \to T$ denote the projection. The tree T becomes a metric space once we define a metric $d = \pi_* \tilde{\mu}$ by pushing forward the measure μ by the projection π. (The real tree T is often not locally compact; for instance, when the leaves of the foliation on the surface are dense, we can find sequences of arcs C_n transverse to the foliation 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 fundamental group $\pi_1\mathcal{R}$ acts by isometries on the metric space (T,d), and the map $\pi : \widetilde{\mathcal{R}} \to (T,d)$ is equivariant with respect to this action. The only fact we will really need about such an $\mathbf{R}$-tree is that if A_γ is the isometry of (T,d) corresponding to $\gamma \in \pi_1\mathcal{R}$ for which $\inf_{y\in T} d(A_\gamma y, y) > 0$, then A_γ has an *axis* l_γ in (T,d), i.e., an isometrically embedded line in (T,d) 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,d).

2.5 The goal of this paper is to give a new proof of the following theorem

Theorem. (Hubbard-Masur [HM]) *If $(\mathcal{F},\mu)$ is a measured foliation on $\mathcal{R}$, then there is a unique holomorphic quadratic differential Φ on $\mathcal{R}$ whose vertical measured foliation is equivalent to $(\mathcal{F},\mu)$.*

Remark. Actually, the statement of the theorem in [HM] is stronger, that the map $p : \mathrm{QD}(\mathcal{R}) \to \mathcal{MF}$ from the space of holomorphic quadratic differentials on $\mathcal{R}$ to the space $\mathcal{MF}$, given by associating to $\Phi \in \mathrm{QD}(\mathcal{R})$ the class of its vertical measured foliation, is a homeomorphism. However, given the statement above that p is bijective, it is not very difficult to show that p is continuous, once we define an appropriate topology for $\mathcal{MF}$; as we have no new perspectives on this issue of continuity to offer, we omit the (straightforward) argument.

§3. Proof of Existence.

Consider a measured foliation $(\mathcal{F}, \mu)$ on a compact Riemann surface $\mathcal{R}$. Lift this measured foliation to a measured foliation $(\widetilde{\mathcal{F}}, \tilde{\mu})$ on the universal cover $\widetilde{\mathcal{R}}$. Construct the real tree (T, d) from $(\widetilde{\mathcal{F}}, \tilde{\mu})$ by projecting $\tilde{\mu}$ along the leaves of $\widetilde{\mathcal{F}}$. Let $\pi : \widetilde{\mathcal{R}} \to (T, d)$ denote the projection. The map π is equivariant with respect to the action of $\pi_1 \mathcal{R}$ on $\widetilde{\mathcal{R}}$ and (T, d) and we seek an equivariant energy minimizing map $f : \widetilde{\mathcal{R}} \to (T, d)$ which is equivariantly homotopic to π. We aim to give an elementary proof adapted to this situation; we note that large portions of our argument could be replaced by the general arguments and constructions of Korevaar and Schoen [KS].

3.1 We argue in two steps. In this subsection, we find a smooth energy minimizing equivariant map $f : \widetilde{\mathcal{R}} \to (T, d)$ which is equivariantly homotopic to π, and we describe it as the projection along the leaves of the vertical measured foliation of its normalized Hopf differential. In the next subsection, we prove the existence part of the Hubbard-Masur theorem by showing that this vertical measured foliation of the normalized Hopf differential of f is equivalent to $(\mathcal{F}, \mu)$. We begin with

Proposition 3.1. *There is an equivariant map $f : \widetilde{\mathcal{R}} \to (T, d)$ which is equivariantly homotopic to $\pi : \widetilde{\mathcal{R}} \to (T, d)$ with the properties: (i) off of a discrete set Z, f is locally a harmonic projection to a Euclidean line, (ii) at the points in Z, the map f pulls back germs of convex functions to germs of subharmonic functions.*

We consider the space $L(\mu)$ of all C^k measured foliations on $\mathcal{R}$ which are measure equivalent to $(\mathcal{F}, \mu)$. Given such a measured foliation $(\mathcal{G}, \nu) \in L(\mu)$, we can lift it to a C^k measured foliation $(\widetilde{\mathcal{G}}, \tilde{\nu})$ on $\widetilde{\mathcal{R}}$. Let $g : \widetilde{\mathcal{R}} \to (T, d)$ be the projection along the leaves of $\widetilde{\mathcal{G}}$; as $(\mathcal{G}, \nu)$ differs from $(\mathcal{F}, \mu)$ by a sequence of Whitehead moves and an isotopy, we see that g has range in the $\mathbf{R}$-tree (T, d) defined by the projection $\pi : \widetilde{\mathcal{R}} \to (T, d)$.

Moreover, as $\widetilde{\mathcal{G}}$ is C^k, we can define an energy density $e(p) = \frac{1}{2}|\nabla g|^2$ at every non-singular point of $\widetilde{\mathcal{G}}$ on $\widetilde{\mathcal{R}}$. Here we compute the gradient of p with respect to the lift $\tilde{\sigma}$ of some background (say, for concreteness, hyperbolic) conformal metric σ on $\mathcal{R}$. This energy density is equivariant with respect to the action of $\Gamma = \pi_1 \mathcal{R}$ on $\widetilde{\mathcal{R}}$, and is uniformly bounded. Thus, we can define a (finite) total energy

$$E((\mathcal{G}, \nu)) = \iint_{\widetilde{\mathcal{R}}/\Gamma} e(g) dA(\sigma).$$

Having defined this total energy for every element in $L(\nu)$, we see that we have defined an energy functional $E : L(\mu) \to \mathbf{R}$ which is finite valued. We consider the quantity

$$I = \inf_{(\mathcal{G}, \nu) \in L(\mu)} E((\mathcal{G}, \nu)) < \infty$$

and let $(\mathcal{G}_n, \nu_n)$ denote a sequence with $E((\mathcal{G}_n, \nu_n)) \to I$; we may as well take $(\mathcal{G}_0, \nu_0) = (\mathcal{F}, \mu)$ and $g_0 = \pi$.

Let $g_n : \widetilde{\mathcal{R}} \to (T, d)$ denote the associated family of maps; we may as well assume that the energies $E(g_n) = E(\mathcal{G}_n, \nu)$ satisfy $E(g_n) \le E(\pi)$.

We claim that a subsequence of $\{g_n\}$ converges uniformly to the desired map $f : \widetilde{\mathcal{R}} \to (T, d)$.

By the proof of the Courant-Lebesgue lemma (see [J], Lemma 3.1.1), the family $\{g_n\}$ is equicontinuous; for the sake of completeness, we include a proof adapted to the present situation.

Lemma 3.2. *$d(g_n(p), g_n(q)) \le (4\pi E(\pi))^{1/2} [-\log d(p,q)/2]^{-1/2}$ for $d_\sigma(p,q) < \delta_0 < (\mathrm{inj}_\sigma \mathcal{R})^2$.*

Proof. Let $\delta = d_\sigma(p,q)/2$. Connect p to q by a geodesic and let x_0 denote the midpoint. Let A denote the annulus around x_0 of radii δ and $\sqrt{\delta}$, parametrized by the polar coordinates (r, θ). We recall that the hyperbolic metric σ has the local form on the annulus A as $\sigma = dr^2 + \sinh^2 r d\theta^2$ in terms of the polar coordinates (r, θ).

We observe that since $g_n^{-1}(g_n(p))$ is a leaf of the foliation $\widetilde{\mathcal{G}}_n$ (which has no cycles), it meets every circle $C_\alpha = \{r = \alpha\}$ in A, for each $\alpha \in [\delta, \sqrt{\delta}]$. As a similar statement holds for the leaf $g_n^{-1}(g_n(q))$, we see that the image $g_n(C_\alpha)$ of C_α under g_n is connected and contains the points $g_n(p)$ and $g_n(q)$, hence the (distance achieving) geodesic between them, covered at least twice.

Thus, for any fixed $r \in [\delta, \sqrt{\delta}]$, we may integrate around C_r to obtain

$$2d(g_n(p), g_n(q)) \le \int_{C_r} \|g_n * \partial_\theta\|_d d\theta \le (2\pi)^{1/2} \left(\int_{C_r} \|g_n * \partial_\theta\|_d^2 d\theta \right)^{1/2} \tag{3.1}$$

by Cauchy-Schwarz, where we have, of course, parametrized C_r by $\theta \in [0, 2\pi]$. We next square (3.1) and multiply by $[\sinh r]^{-1}$; upon integrating the result over the domain $r \in [\delta, \sqrt{\delta}]$ for the radii in A, we obtain

$$\begin{aligned} 4\left[d(g_n(p), g_n(q))\right]^2 \int_\delta^{\sqrt{\delta}} \frac{1}{\sinh r} dr &\le 2\pi \int_\delta^{\sqrt{\delta}} \int_0^{2\pi} \frac{1}{\sinh r} \|g_n * \partial_\theta\|_d^2 d\theta dr \\ &= 2\pi \int_\delta^{\sqrt{\delta}} \int_0^{2\pi} \frac{1}{\sinh^2 r} \|g_n * \partial_\theta\|_d^2 \sinh r d\theta dr. \end{aligned} \tag{3.2}$$

Of course, the volume form for σ in these coordinates is $\sinh r dr d\theta$, and so we can rewrite (3.2) as

$$4\left[d(g_n(p), g_n(q))\right]^2 \int_\delta^{\sqrt{\delta}} \frac{1}{\sinh r} dr \le 2\pi \iint_A \frac{1}{\sinh^2 r} \|g_n * \partial_\theta\|_d dA(\sigma)$$

$$\le 2\pi \iint_A \left\{ \|g_n * \partial_r\|_d^2 + \frac{1}{\sinh^2 r} \|g_n * \partial_\theta\|_d^2 \right\} dA(\sigma) \tag{3.3}$$

$$\le 4\pi E(g_n) \tag{3.4}$$

where we have added the term $\|g_n * \partial_r\|_d^2$ in (3.3) in order to be able to identify the integral on the left hand side of (3.3) as a multiple of the energy $E(g_n)$. (We also are assuming that the annulus A embeds in $\mathcal{R}$ under the projection from $\widetilde{\mathcal{R}}$ to $\mathcal{R}$: this follows from choosing $\sqrt{\delta} < \operatorname{inj} \mathcal{R}$.) Of course,

$$E(g_n) \leq E(\pi) \tag{3.5}$$

by convention, so we combine (3.4) and (3.5) to obtain

$$d(g_n(p), g_n(q)) \leq \pi^{1/2} E(\pi)^{1/2} \left[\int_{\delta}^{\sqrt{\delta}} \frac{1}{\sinh r} \right]^{-1/2}. \tag{3.6}$$

For δ small, we use the (gross) approximation $\sinh r < 2r$ to obtain from (3.6),

$$\begin{aligned} d(g_n(p), g_n(q)) &\leq 2\pi^{1/2} E(\pi)^{1/2} (-\log \delta)^{-1/2} \\ &= 2\pi^{1/2} E(\pi)^{1/2} \left(\log \frac{2}{d(p,q)} \right)^{-1/2} \end{aligned}$$

as required. $\square$

Our goal is to apply Ascoli-Arzela to our sequence $\{g_n\}$; the last lemma 3.2 shows that the family $\{g_n\}$ is equicontinuous, so we are left to find a subsequence $\{g_{n_k}\}$ of $\{g_n\}$ for which $g_{n_k}(z)$ has compact closure for every $z \in \widetilde{\mathcal{R}}$. This relies more heavily on the geometry of the map $\pi : \widetilde{\mathcal{R}} \to (T, d)$.

We begin by considering two closed curves B and C on $\mathcal{R}$ for which the intersection numbers $i(B, \mu) \neq 0$, $i(C, \mu) \neq 0$ and $i(B, C) \neq 0$. Consider also a fundamental domain $\mathcal{D}$ for the action of Γ on $\widetilde{\mathcal{R}}$. Upon uniformizing $\widetilde{\mathcal{R}}$ as $\mathbf{H}^2$, we find axes for hyperbolic isometries γ_B and γ_C, corresponding to B and C, respectively, which meet in $\mathcal{D}$. These group elements γ_B and γ_C (to slightly abuse notation) also act on (T, d) with axes that meet in $g_n(\mathcal{D})$. So consider a point $z \in \mathcal{D}$. We know that $g_n(\gamma_B z) = \gamma_B g_n(z)$ and $g_n(\gamma_C z) = \gamma_C g_n(z)$ by equivariance; we also know that if $w \in T$ and $d(w, \text{ axis of} \gamma_C) = K_w$ then, since the isometry γ_C preserves the axis of γ_C and takes w to $\gamma_C w$, we have $d(\gamma_C w, \text{ axis of } \gamma_C) = K_w$. Moreover, the geodesic between w and $\gamma_C w$ must follow the geodesic from w to the axis of γ_C, then along that axis, and finally along the geodesic from the axis to $\gamma_C w$: we conclude $d(w, \gamma_C w) > 2K_w$.

So we consider the triple of points $z, \gamma_B z, \gamma_C z \in \mathcal{D} \cup \gamma_B \mathcal{D} \cup \gamma_C \mathcal{D}$. These three points are within a distance C_0 of each other so that our equicontinuity bound in Lemma 3.2 shows that $d(g_n z, g_n \gamma_B z) \leq C_1$ and $d(g_n(z), g_n \gamma_C z) \leq C_1$, where C_1 depends only upon C_0 and δ_0 in Lemma 3.2. Thus, by the estimates of the previous paragraph, we see that $g_n(z)$ must be within a distance of $\frac{1}{2}C_1$ from both the axes of γ_B and γ_C. Of course since the axis of γ_B and the axis of γ_C diverge from each other, the set $\{w \mid d(w, \text{ axis of } \gamma_B) < \frac{1}{2}C_1,$ $d(w, \text{ axis of } \gamma_C) < \frac{1}{2}C_1\}$ has finite diameter, and so is a bounded distance from $\pi(\mathcal{D})$, itself a compact set. We conclude

Lemma 3.3. *$d(g_n(z), \pi(z)) < K_0$ for K_0 depending only upon $E(\pi)$ and $\mathcal{R}$.* $\square$

Now, this in itself is insufficient to show that $\{g_n(z)\}$ has a convergent subsequence, as (T, d) is (generically) not locally compact anywhere. So we consider the previous construction more carefully, after a preliminary normalization.

As g_n is obtained from projecting along the leaves of the lift of a measured foliation on $\widetilde{\mathcal{R}}$, we recall that in the fundamental domain $\mathcal{D}$, the foliation $(\mathcal{G}_n, \nu_n)$ for g_n has at most $4g-4$ singularities, each with at most $4g-2$ prongs.

As the foliations associated to g_n have a uniformly bounded number of singularities and $\mathcal{D}$ has compact closure, we may pass to a subsequence (also denoted $\{g_n\}$) in which these singularities converge to a finite number of points; as the complexity of the singularities is uniformly bounded, we may assume that the complexity of each convergent sequence of singularities is constant, even if several singularities coalesce in the limit to a singularity with greater complexity.

With this new sequence $\{g_n\}$) defined, we are in a position to prove

Lemma 3.4. *There is a subsequence $\{g_{n_k}\}$ of $\{g_n\}$ so that for every $z \in \widetilde{\mathcal{R}}$, $\{g_{n_k}(z)\}$ has compact closure.*

Proof. Choose a curve $B \subset \mathcal{R}$ with $i(B, \mu) \neq 0$ as in the proof of Lemma 3.3 and such that B avoids a neighborhood of all the limits of the singularities of the foliations above. We consider the arc $\widehat{B}$ which is a lift of B in $\widetilde{\mathcal{R}}$, where we assume that $\widehat{B}$ meets $\mathcal{D}$, the fundamental domain for Γ in $\widetilde{\mathcal{R}}$. Then $g_n(\widehat{B})$ must meet the axis of γ_B on T: to see this, let u and v denote the endpoints of $\widehat{B}$, and observe that $\gamma_B u = v$. Then $g_n(v) = g_n(\gamma_B u) = \gamma_B g_n(u)$ and so to connect the endpoints of $g_n(\widehat{B})$, we proceed along the following geodesic: we first travel along the geodesic α from $g_n(u)$ to the axis of B, then along the axis of B for its translation length and then up along the geodesic $\gamma_B \alpha$ to $\gamma_B g_n(u) = g_n(v)$. Thus there is a point $z_n \in \widehat{B}$ which satisfies $g_n(z_n) \in$ axis of γ_B in T. But $\widehat{B}$ is compact, so there is a subsequence of $\{n\}$ which we'll continue to denote by $\{n\}$, so that z_n converges to, say, $z \in \widehat{B}$. Moreover, by Lemma 3.3 and the compactness of $\widehat{B}$, we have that $d(g_n(\widehat{B}), \pi(\widehat{B})) < K_1$, so that $g_n(z_n)$ must live in a portion of the axis of γ_B which is of finite length, hence compact. Thus there is a further subsequence of $\{g_n\}$, which we'll continue to denote by $\{g_n\}$, so that $g_n(z_n)$ converges, say to $w \in T$.

So we consider the sequence $\{g_n(z)\}$. We observe that

$$d\left(g_n(z), w\right) \leq d(g_n(z_n), w) + d(g_n(z), g_n(z_n)) .$$

We have just seen that the first term on the right hand side tends to zero and by Lemma 3.2, we know that $d(g_n(z), g_n(z_n)) \leq C[-\log d_\sigma(z, z_n)]^{-1/2}$ which also tends to zero, so we conclude that $g_n(z)$ converges to w.

We next show that $g_n(\tilde{z}_i)$ has compact closure where $\tilde{z}_i$ is one of limiting singularities $z_i \in \{z_1, \ldots, z_l\}$ of the foliation $\mathcal{G}_n$. Recall next that a curve C on a surface is quasitransverse to a foliation $\mathcal{G}$ if it is the union of arcs transverse to $\mathcal{G}$ with arcs contained in $\mathcal{G}$, so that upon any isotopy of the curve C along leaves of $\mathcal{G}$ which collapses the arcs contained in C while smoothing C, the resulting isotoped curve is transverse to $\mathcal{G}$; thus, a quasitransverse curve will minimize the transverse measure among curves in its free homotopy class, and will lift to a curve whose projection is an embedded arc in the tree T. Now, on the closed Riemann surface $\mathcal{R}$, we can connect the projection of z with itself by a closed curve C_n which is quasi-transverse to $\mathcal{G}_n$; here we take $\{C_n\}$ to be representatives of a single free homotopy class. The curves C_n lift to arcs $\widehat{C}_n$ in $\widetilde{\mathcal{R}}$ which connect z to some particular image $\gamma_C z$ where γ_C corresponds to the element $[C_n] \in \pi_1\mathcal{R}$. We now specialize to the case where we require C_n to pass through $z_i^{(n)}$, a singularity of $\mathcal{G}_n$ which is converging to z_i. (In the case of several singularities of $\mathcal{G}_n$ converging to z_i, the choice of singularity will not affect the argument.) Now, as C_n is quasi-transverse to $\mathcal{G}_n$, we have that $g_n(\widehat{C}_n)$ lies in an embedded interval in T connecting $g_n(z)$ and $g_n(\gamma_C z)$. Of course, as T is a tree, there is only one such embedded interval, so we see that $g_n(\widetilde{z_i^{(n)}})$ must lie in that embedded arc. As $z_i^{(n)} \to z_i$ and g_n is uniformly absolutely continuous, we see that $g_n(\tilde{z}_i)$ limits on that embedded arc, a compact set.

We pass to a subsequence, which we'll continue to denote $\{g_n\}$, so that $g_n(\tilde{z}_i)$ converges to $w_i \in T$, with $i = 1, \ldots, l$.

In fact $g_n(\widetilde{z_i^{(n)}})$ does not depend on the choice of n: to see this, use curves C_n and C_n' which are in two distinct free homotopy classes, respectively, but which agree only on an arc from z to $z_i^{(n)}$ and after which they pass into different sectors abutting the singularity $z_i^{(n)}$. (These curves C_n and C_n' can easily be constructed: by our normalizations in the proof of Lemma 3.3, all of the singularities $z_i^{(n)}$ have the same complexity, and so one can arrange for the initial portions of C_n and C_n' from z to $z_i^{(n)}$ to agree and to have local form near $z_i^{(n)}$ that is independent of n. For the latter portions of C_n and C_n', one merely connects $z_i^{(n)}$ back to z via different sectors of the foliation near $z_i^{(n)}$ along quasi-transversal curves in two distinct free homotopy classes $[C_n]$ and $[C_n']$, respectively.) The associated images $g_n(\widehat{C}_n)$ and $g_n(\widehat{C}_n')$ will then agree from $g_n(z)$ to $g_n(z_i^{(n)})$ and diverge after. Yet as the interval connecting $g_n(z)$ and $g_n(\gamma_C z) = \gamma_C g_n(z)$ and the interval connecting $g_n(z)$ and $g_n(\gamma_{C'} z) = \gamma_{C'} g_n(z)$ are independent of n, so is their common branch and its endpoint, the vertex which is the image of $g_n(\widetilde{z_i^{(n)}})$.

On the Riemann surface $\mathcal{R}$, we can connect the singularities of $\mathcal{G}_n$ with a collection $\mathcal{C}$ of arcs that are transverse to the foliation $\mathcal{G}_n$ and which cut the surface $\mathcal{R}$ into simply connected pieces with no interior singularities. The complete lift $\widetilde{\mathcal{C}}$ of this arc system connects singularities of $\widetilde{\mathcal{G}}_n$ on $\widetilde{\mathcal{R}}$ and cuts $\widetilde{\mathcal{R}}$ into precompact simply connected pieces with no interior singularities.

The fundamental domain $\mathcal{D}$ is covered by a finite number of closures, say $\overline{R}_1^{(n)}, \ldots, \overline{R}_m^{(n)}$, of these simply connected pieces. Consider the image $g_n(\overline{R}_k^{(n)})$ of one of these pieces. The vertices of such a piece have fixed image $g_n(z_i^{(n)}) = w_i$, and as the boundary edges are transverse to the foliation $\widetilde{\mathcal{G}}_n$, we see that the g_n-image of the edge connecting $z_i^{(n)}$ and

$z_j^{(n)}$ is the arc of the tree connecting w_i to w_j. As $\mathcal{G}_n$ restricts to a non-singular foliation of $\overline{R}_k^{(n)}$, we see that $g_n(\overline{R}_k^{(n)})$ lies in the finite subcomplex determined by the convex hull of the (fixed) g_n-images of the vertices of $R_k^{(n)}$. Thus $g_n(\mathcal{D}) \subset g_n\left(\bigcup_{k=1}^n \overline{R}_k^{(n)}\right)$, which itself is contained in Conv $\left(g_n\left(\bigcup_{k=1}^n \text{ vertices } (\overline{R}_k^{(n)})\right)\right)$, the convex hull of the g_n-images of the union of the vertices of $\overline{R}_k^{(n)}$. But this last set is both compact and independent of n, proving the lemma. $\square$

Conclusion of the proof of Proposition 3.1. By Lemmas 3.2 and 3.4, the sequence $\{g_n\}$ is an equicontinuous family of mappings from $\widetilde{\mathcal{R}}$ to the metric space (T, d) for which $g_n(z)$ has compact closure. Thus, the Ascoli theorem provides for a sequence which converges uniformly on compacta, hence uniformly, as each g_n is equivariant and a fundamental domain is compact. Let f be the limiting continuous map.

As all our maps g_n are equivariantly homotopic to π, so is the limit map f.

Now, locally, each map g_n maps an open precompact set Ω to a locally compact complex; indeed, for Ω chosen so that either its closure is disjoint from $\{z_i\}$ or with some z_i as interior points, we see that for n sufficiently large, the complex $g_n(\Omega)$ is contained in a single locally compact subtree $T_\Omega \subset T$. This complex T_Ω, being locally compact, embeds in $\mathbf{E}^N$ for some N. Thus there are no difficulties at this point in regarding our maps $g_n\big|_\Omega$ and $f\big|_\Omega$ as being elements of $L_1^2(\Omega, T_\Omega) \subset L_1^2(\Omega, \mathbf{E}^N)$.

So consider the form $\widetilde{\Phi} = \langle f_z, f_z \rangle \, dz^2$. As the energy functional is lower semicontinuous with respect to convergence in L_1^2, we see first that $\widetilde{\Phi} \in L_{loc}^1$ and next that, since f is also continuous, the map f is energy minimizing with respect to reparametrizations of the domain Ω. A brief and straightforward computation done explicitly in the proof of Lemma 1.1 in [S] shows that $\widetilde{\Phi}$ is then weakly holomorphic. Of course, since we have just seen that $\widetilde{\Phi} \in L_{loc}^1$, we conclude by Weyl's lemma that $\widetilde{\Phi}$ is (strongly) holomorphic.

By the equivariance of f, the form $\widetilde{\Phi}$ descends to a holomorphic quadratic differential on the compact Riemann surface $\mathcal{R}$.

We observe that $\widetilde{\Phi}$ is not identically zero, for, if it were, the map f would either be a conformal (non-constant) map from $\widetilde{\mathcal{R}}$ to (T, d) which is absurd, or constant, which would contradict its being continuous and equivariant.

Thus, the zeros of $\widetilde{\Phi}$ are discrete; as described in §2, we can introduce coordinates $\zeta = \xi + i\eta$ so that $\widetilde{\Phi} = 1 d\zeta^2$ in a neighborhood disjoint from $\widetilde{\Phi}^{-1}(0)$. Then the complex equation $4 \langle f_\zeta, f_\zeta \rangle_d \, d\zeta^2 = 1 d\zeta^2$ becomes, locally, the pair of real equations

$$\|f_\xi\|_d^2 - \|f_\eta\|_d^2 = 1 \tag{3.7a}$$

$$\langle f_\xi, f_\eta \rangle_d = 0 \tag{3.7b}$$

Now, as f is non-constant, off of the Γ-orbit $\Gamma \cdot \{z_1, \dots, z_l\}$ of $\{z_1, \dots, z_l\}$, we must have f mapping a neighborhood in $\widetilde{\mathcal{R}} - \Gamma \cdot \{z_1, \dots, z_l\}$ onto an embedded subinterval I in T. Thus, as the interval is one-dimensional, we must have $f_\xi \| f_\eta$: the equations (3.7) then combine to show that $f_\eta = 0$ and $f_\xi = \partial_t$ for unit vector ∂_t along I.

Thus, away from the Γ-orbit of $\{z_1, \ldots, z_l\}$ and the zeros of $\widetilde{\Phi}$, we see that f is a projection along the leaves of the vertical foliation of $\widetilde{\Phi}$, and that f pushes the transverse measure of the vertical foliation of $\widetilde{\Phi}$ onto the line element of the metric on embedded intervals in T. The continuity of f extends this last statement across the zeros of $\widetilde{\Phi}$ and the Γ-orbit of $\{z_1, \ldots, z_l\}$. Property (i) of the statement of the proposition follows immediately. To see property (ii), in a neighborhood Ω of a zero P of $\widetilde{\Phi}$ of order m there are $m+2$ arcs of the vertical foliation emanating from p, dividing Ω into $m+2$ sectors. For each pair of such sectors, the map f acts as a projection onto a segment I, and so the pullback of a convex function restricted to that segment I is submean. But then the pullback of a convex function, after averaging over all such pairs of sectors, is still submean, which is the content of statement (ii).

3.2 **Proof of Existence**. We claim that the holomorphic quadratic differential Φ on $\mathcal{R}$ developed in the last subsection is the desired quadratic differential. Let $(\mathcal{G}, \nu)$ be its vertical measured foliation (which integrates kerdf off of $\Phi^{-1}(0)$.

The idea is to use the uniform convergence of the family of maps $\{g_n\}$ to f to show that the measured foliations $\mathcal{G}_n$ (along which g_n projects) are equivalent (in the sense of §2) to the measured foliation of Φ (along which f projects). Our plan is the following. We partition $\mathcal{R}$ (and hence $\widetilde{\mathcal{R}}$) into a union of rectangles $\mathcal{R} = R_1^{(n)} \cup \cdots \cup R_k^{(n)}$ with two parallel sides being leaves of $\mathcal{G}_n$ (for n large) and the other two parallel sides being transverse to $\mathcal{G}_n$. We require the singularities of the foliation to lie on the boundaries of rectangles, and also to lie in the interior of the leaf segments comprising such boundary edges. If we can find a corresponding decomposition of $\mathcal{R}$ for $\mathcal{G}$, and an isotopy of $\mathcal{R}$ which takes rectangles for $\mathcal{G}_n$ to rectangles of $\mathcal{G}$ in a measure-preserving way, then we will have $[\mathcal{G}] = [\mathcal{G}_n]$, as desired.

We begin by observing that even though the map g_n has range an $\mathbf{R}$-tree, locally it is much simpler. Indeed, in the neighborhood of a non-singular point, we have already seen that we can factor g_n as $g_n = i \circ F^{(n)}$ where $F^{(n)}$ is a map to an interval I and i is the isometric inclusion of I into the tree T. Beyond that, in a neighborhood of a singular point of order $m+2$, we can factor the map g_n as $g_n = i \circ F^{(n)}_{m+2}$ where $F^{(n)}_{m+2}$ is a map to an $m+2$-pronged star S_{m+2} and i again is an isometric inclusion of S_{m+2} into T.

So to find the desired decomposition for $\mathcal{G}$, we first observe that if Ω_n is a neighborhood of a singularity of $\mathcal{G}_n$, then $F^{(n)}_{m+2}(\Omega_n)$ must contain a finite-pronged singularity where $F^{(n)}_{m+2}$ is the local factor of g_n, in the previous paragraph. The uniform convergence of g_n to f then forces singularities of $\mathcal{G}_n$ to limit on singularities of $\mathcal{G}$ and all singularities of $\mathcal{G}$ to arise as limits of singularities of $\mathcal{G}_n$. So, consider an arc of a leaf containing a singularity of $\mathcal{G}_n$ in its interior. Its lift maps under g_n to a vertex in T; the uniform convergence of g_n to f provides that near an endpoint p_n of this arc there is a point p so that $f(\widetilde{p}) = g_n(\widetilde{p}_n)$, in the obvious notation. Pick n large enough so that if p_n is such a vertex of our decomposition into rectangles $\mathcal{R} = R_1^{(n)} \cup \cdots \cup R_k^{(n)}$, then p_n is much closer to p than the width or length of any R_j.

We find other vertices similarly; for rectangles with an edge containing a singularity, an edge transverse to the foliation for $\mathcal{G}_n$ with endpoint q_n will have a lift that will project to an arc in (T, d) and there will be a point q near q_n so that $f(\widetilde{q}) = g_n(\widetilde{q}_n)$. We then connect q to the leaf containing the singularity by a transverse arc near the edge of the $\mathcal{G}_n$-rectangle connecting q_n to the leaf containing the $\mathcal{G}_n$ singularity. We then continue this way for every vertex of every rectangle, systematically building up a decomposition $\mathcal{R} = R_1 \cup \cdots \cup R_k$ relative to $\mathcal{G}$. The key points are that: (i) as the ν and ν_n-measure of an arc A is given by the lengths of $f(\widetilde{A})$ and $g_n(\widetilde{A})$ (respectively), and the endpoints of edges of rectangles for lifts of $\mathcal{G}$ and $\mathcal{G}_n$ project under f and g_n (respectively) to the same point, the corresponding edges have the same measure, and (ii) as we are eventually only interested in a measure preserving isotopy, the ambiguity in the choice of the vertices $\{p, q\}$ (coming partially from indeterminacy in the choice of $\{p_n, q_n\}$) for $R_1, R_2, \ldots, R_k$ is harmless.

Thus the ν-width of R_i agrees with the ν_n-width of $R_i^{(n)}$ and we can construct a measure preserving isotopy of $\mathcal{R} = R_1^{(n)} \cup \cdots \cup R_k^{(n)}$ onto $\mathcal{R} = R_1 \cup \cdots \cup R_k$, as desired, by piecing together such isotopies of corresponding rectangles. Therefore, $(\mathcal{G}, \nu)$ is equivalent to $(\mathcal{G}_n, \nu_n)$, hence to $(\mathcal{F}, \mu)$. $\square$

§4. Uniqueness.

Here we give a proof from the current perspective that if Φ and Ψ are holomorphic quadratic differentials with equivalent vertical measured foliations, then $\Phi = \Psi$.

Naturally, we can lift Φ and Ψ to equivariant holomorphic differentials $\widetilde{\Phi}$ and $\widetilde{\Psi}$, respectively, on $\widetilde{\mathcal{R}}$ and then produce equivariant and equivariantly homotopic projections π_Φ and π_Ψ from $\widetilde{\mathcal{R}}$ to (T, d). We consider the function $D(z) = d(\pi_\Phi(z), \pi_\Psi(z))$ on $\widetilde{\mathcal{R}}$ and compute from the Riemannian chain rule (see [J]) that

$$\Delta D(z) \;=\; \mathrm{tr}_{(\pi_\Phi, \pi_\Psi)_*(\partial_x, \partial_y)} \mathrm{Hess}_{T \times T}\, d \times d$$

off of the set $\widetilde{\Phi}^{-1}(0) \cup \widetilde{\Psi}^{-1}(0)$; here (∂_x, ∂_y) is a σ-orthonormal basis for $T_z\mathcal{R}$. We conclude that as T is negatively curved, $d : T \times T \to \mathbf{R}$ is convex, so that $D(z)$ is subharmonic. As $D(z)$ is invariant under the action of $\pi_1\mathcal{R}$, it descends to the compact surface $\mathcal{R}$ as a subharmonic function, and is therefore constant.

We claim that, in fact, $D(z) = 0$: this follows from considering a neighborhood U of a zero p of $\widetilde{\Phi}$, which π_Φ takes to a neighborhood of a vertex of finite subtree in (T, d). Then, if $\pi_\Psi(p) \neq \pi_\Phi(p)$, we can connect $\pi_\Psi(p)$ to $\pi_\Phi(p)$ by a geodesic A, which by the definition of π_Φ, will have the germ of its π_Φ-preimage at p in $\widetilde{\mathcal{R}}$ strictly within a strictly convex cone based at p. As the germ of the π_Ψ-preimage in $\widetilde{\mathcal{R}}$ of that arc will necessarily be within a half plane, we see that there is a point q near p outside of both preimages, so that the geodesic from $\pi_\Psi(q)$ to $\pi_\Phi(q)$ includes A: we conclude that $d\,(\pi_\Psi q, \pi_\Phi q) \neq d\,(\pi_\Phi(p), \pi_\Psi(p))$ unless $D(z) = 0$.

It then follows by a similar local analysis of π_Φ and π_Ψ that the foliations of Φ and Ψ coincide, so that the germs of Φ and Ψ near $\Phi^{-1}(0) = \Psi^{-1}(0)$ agree, so that $\Phi/\Psi = 1$ near those zeros. This, of course forces $\Phi \equiv \Psi$, as desired. $\square$

Remark. As noted earlier, it is easy to show that the map from the vector space $QD(\mathcal{R})$ of holomorphic quadratic differentials on the Riemann surface $\mathcal{R}$ to the space $\mathcal{MF}$ of measured foliations is continuous, once one defines the topology on $\mathcal{MF}$, proving that $QD(\mathcal{R})$ is homeomorphic to $\mathcal{MF}$. We can then conclude

Corollary. (Thurston [Th]): *The space $\mathcal{PMF}$ of projective measured foliations on a closed surface of genus g is a $6g - 7$ dimensional sphere.*

Proof. Riemann-Roch implies that $QD \approx \mathcal{MF}$ is a $6g - 6$ dimensional vector space; we check that the rays in the vector space represent projectively equivalent measured foliations. $\square$

References

[FLP] A. Fathi, F. Laudenbach, and V. Poenaru, *Traveaux de Thurston sur les Surfaces*, Asterisque (1979), 66–67.

[G] F.P. Gardiner, *Teichmüller Theory and Quadratic Differentials*, Wiley and Sons, New York, 1987.

[HM] J. Hubbard and H.A. Masur, *Quadratic Differentials and Foliations*, Acta Math. **142** (1979), 221–274.

[J] J. Jost, *Two Dimensional Geometric Variational Problems*, John Wiley and Sons, West Sussex, England, 1991.

[K] S. Kerckhoff, *The Asymptotic Geometry of Teichmüller Space*, Topology **19** (1980), 23–41.

[KS] N. Korevaar and R. Schoen, *Sobolev Spaces and Harmonic Maps for Metric Space Targets*, Comm. Anal. Geom. **1** (1993), 561–659.

[S] 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, NY, (1983), 321–358.

[Str] K. Strebel, *Quadratic Differentials*, Springer-Verlag, Berlin, 1984.

[Th] W.P. Thurston, *On the Geometry and Dynamics of Diffeomorphisms of Surfaces*, Bull. AMS **19** (1988), 417–431.

[Tr] A. Tromba, *Teichmüller Theory in Riemannian Geometry*, Birkhauser, Basel, 1992.

[W1] M. Wolf, *The Teichmüller Theory of Harmonic Maps*, J. Differential Geometry **29** (1989), 449–479.

[W2] ———, *Harmonic Maps from Surfaces to* $\mathbf{R}$*-Trees*, Math. Zeit **218** (1995), 577–593.