

HARMONIC MAPS INTO HYPERBOLIC 3-MANIFOLDS

YAIR N. MINSKY

ABSTRACT. High-energy degeneration of harmonic maps of Riemann surfaces into a hyperbolic 3-manifold target is studied, generalizing results of [M1] in which the target is two-dimensional. The Hopf foliation of a high-energy map is mapped to an approximation of its geodesic representative in the target, and the ratio of the squared length of that representative to the extremal length of the foliation in the domain gives an estimate for the energy.

The images of harmonic maps obtained when the domain degenerates along a Teichmüller ray are shown to converge generically to pleated surfaces in the geometric topology or to leave every compact set of the target when the limiting foliation is unrealizable.

1. INTRODUCTION

The geometry of hyperbolic 3-manifolds and that of Teichmüller spaces of surfaces are intricately related. The primary examples are the parametrization of the quasi-conformal deformation space of a hyperbolic manifold by the Teichmüller space of its conformal boundary at infinity, and the pleated surfaces of Thurston, which are a way of mapping hyperbolic surfaces path-isometrically to the interior of a hyperbolic 3-manifold. This paper studies a less-explored connection, provided by *harmonic maps* from Riemann surfaces to hyperbolic 3-manifolds.

Fix a homotopy class of maps $[f: S \rightarrow N]$, where S is a closed surface and N is a complete hyperbolic 3-manifold. For each choice of conformal structure σ and S there is a unique harmonic map in the homotopy class (except in certain degenerate cases—see §3). The focus of this paper is to examine the behavior of these harmonic maps as the conformal structure goes to infinity in the Teichmüller space of S . In particular, we obtain results relating the shape of images of high-energy harmonic maps to those of pleated surfaces.

The main results here are direct generalizations of [M1], and in fact that paper should be considered an introduction to this one. In most places we summarize the argument given there and supply the necessary details to make the generalization. In particular we prove Theorem 4.3 (energy is length-ratio), which characterizes the Hopf foliation of a harmonic map as the measured

Received by the editors May 1, 1990.

1991 *Mathematics Subject Classification*. Primary 53C21, 30F40, 30F60; Secondary 53A30, 53A35.

©1992 American Mathematical Society
0002-9947/92 \$1.00 + \$.25 per page

foliation maximizing (up to an additive constant) the ratio

$$\frac{1}{2} \frac{l_N^2(f(\gamma))}{E_\sigma(\gamma)}$$

over all homotopy classes of curves $\gamma \subset S$, where l_N is length in N of a geodesic representative and E_σ is extremal length in $M = (S, \sigma)$. This ratio in turn gives an estimate for the energy of the map.

The second structural result is Theorem 4.2 (map foliation near lamination), which states that the Hopf foliation is in fact mapped very near to the geodesic lamination in N that is its geodesic representative. From here it is easy to see that a geometric limit of images of harmonic maps whose energy increases is generically a pleated surface.

The last section employs these theorems to study the asymptotics of harmonic maps as the domain degenerates along a Teichmüller ray. As in [M1], the foliation determining the ray determines the limiting Hopf foliation of the maps. When this foliation corresponds to a complete realizable lamination in the manifold, the geometric limit of images is the corresponding pleated surface. Otherwise, the images escape every compact set and exit the end of N determined by the unrealizable lamination. We note here that in order to obtain this clean dichotomy we must impose the condition that N be without parabolics. Hopefully the technical difficulties posed by the presence of parabolics will be dealt with in a later paper.

The plan of the paper is as follows. Section 2 summarizes the basic notions of Thurston's theory of hyperbolic manifolds that we need, notably the uniform injectivity theorem for pleated surfaces. We prove a minor generalization of that theorem (which applies, under suitable conditions, in arbitrarily small injectivity radius) and show how it is used to control the construction of train-tracks in 3-manifolds. Section 3 generalizes to three dimensions the basic estimates on harmonic maps from [M1]. In particular, we prove an area bound on the image, and obtain exponentially small bounds on the function \mathcal{E} which measures the deviation of the map from conformality, and on the geodesic curvature of the images of leaves of the Hopf foliation. We also give a short argument showing that the harmonic maps we are studying in fact exist even when the target is not compact. Section 4 describes the generalization of the main arguments of [M1]. The idea is that, for high enough energy, a decomposition of the surface determined by the Hopf differential produces a nearly straight train-track in the image that carries "most" of the Hopf foliation. Through a careful application of the methods of §2 this train-track is extended to one that carries the entire foliation. The reason this care is needed is that although "most" of the foliation leaves are mapped to near-geodesics the lack of local control over even a small portion of the leaves can potentially destroy all information about the geodesic representative of the foliation. Section 5 gives the above-mentioned results about limits along Teichmüller rays.

The paper represents a portion of my Ph.D. thesis [M0], completed at Princeton University. I am grateful to my advisor, Bill Thurston, for his inspiration and insight.

2. HYPERBOLIC 3-MANIFOLDS AND PLEATED SURFACES

Laminations and train-tracks. As in [M1], $\mathcal{MF}(S)$ denotes the space of measured foliations (up to some topological equivalences) on a surface S . $\mathcal{ML}(M)$

denotes the space of measured geodesic laminations on $M = (S, \sigma)$, where σ denotes a hyperbolic metric on S . $QD(M)$ is the space of quadratic differentials on M holomorphic with respect to the conformal structure induced by σ . All of these spaces are canonically homeomorphic, and we will often suppress the distinction between $\mathcal{MF}(S)$ and $\mathcal{ML}(M)$. When clarity is needed we will denote by $\mathcal{L}: \mathcal{MF}(S) \rightarrow \mathcal{ML}(M)$ the canonical homeomorphism. Since this equivalence makes $\mathcal{ML}(M)$ independent of the hyperbolic metric, we will write $\mathcal{ML}(S)$ as well. $\mathcal{PMF}(S)$ and $\mathcal{PML}(S)$ are the corresponding projectivized spaces, identifying measures that are multiples of each other.

A geodesic lamination is *minimal* if it contains no proper sublamination. Every *measured* geodesic lamination is a disjoint union of finitely many minimal laminations.

A geodesic lamination is *complete* if all of its complementary regions are ideal triangles. Equivalently, a foliation is complete if all of its critical leaves are noncompact, and meet in triples at the singularities.

A lamination is carried on a *train-track*, which is a 1-complex whose branches meet at near-tangencies at vertices called “switches” (see [HP, Th2] for detailed definitions). In a hyperbolic surface we define a $(1, \epsilon)$ -*nearly-straight* train-track as one in which the branches have length at least 1 and curvature at most ϵ (where $\epsilon < 1$), and in which incoming branches meet outgoing branches at a switch in angles deviating from π by at most ϵ . It is then an easy fact of hyperbolic geometry that, for a universal c_0 , if $\epsilon < c_0$ then such a train-track is in a $C\epsilon$ -neighborhood of any lamination carried on it.

A measured lamination also deposits a weight on each branch of a train-track carrying it—the total transverse measure of the leaves passing along that branch. The weights satisfy *switch conditions*—the sum of weights on the incoming branches to any switch equals the sum on the outgoing branches. The system of train-tracks carrying a lamination and the admissible weights on them provide in a natural way the topology (in fact a piecewise linear structure) on $\mathcal{ML}(S) \cong \mathcal{MF}(S)$.

We will consider a train-tracks in 3-manifolds as well. A 1-complex τ is a $(1, \epsilon)$ -*nearly straight* train-track in N^3 if there is a map $f: S \rightarrow N$ which maps a train-track in S to τ (with the above length and curvature requirements on τ). The exact image of the rest of the surface is not important—it serves to “mark” τ and in particular gives a combinatorial meaning to the complementary regions of a train-track in a 3-manifold (see Theorem 2.4 (enlarge train-track efficiently)).

The thick-thin decomposition. If N is a complete hyperbolic n -manifold we denote by $N_{\text{thin}(\epsilon)}$ the ϵ -*thin part* of N , or that subset of N where the length of the smallest nontrivial loop is at most ϵ . For $\epsilon \leq \epsilon_0$ (where ϵ_0 is the *Margulis constant* of \mathbf{H}^n), the ϵ -thin part is of a standard form. In particular, if $n = 2$ a component of the ϵ -thin part is either a cylindrical neighborhood of a short geodesic, or a cusp (corresponding to a parabolic element in the fundamental group). For $n = 3$ a component is either a solid-torus neighborhood of a short geodesic, a (cylinder) $\times [0, \infty)$ corresponding to a one-generator parabolic subgroup (**Z**-cusp) or a (torus) $\times [0, \infty)$ corresponding to a two-generator parabolic subgroup (**Z** \oplus **Z**-cusp).

We will not dwell on the details of this except to note that ϵ_0 -thin parts corresponding to different subgroups are disjoint, and that for any $\epsilon_1 \leq \epsilon_0$ there is an $\epsilon_2 < \epsilon_1$ such that the ϵ_2 -thin part inside any component of $N_{\text{thin}(\epsilon_1)}$, if it is not empty, is arbitrarily far from the boundary of $N_{\text{thin}(\epsilon_1)}$, independently of N (this can be obtained, for example, from arguments such as in [BM]).

2.1. Pleated surfaces and the uniform injectivity theorem. A *pleated surface* (see [Th1]) is a hyperbolic surface $M = (S, \sigma)$ together with a path-isometric map $f: M \rightarrow N$ (meaning that a rectifiable path in M maps to a rectifiable path of the same length) into a hyperbolic 3-manifold N , which is totally geodesic on the complement of some geodesic lamination $\lambda \subset M$, whose leaves are mapped isometrically to geodesics in N . λ is called the *pleating locus* of f (see also [CEG]).

Given a homotopy class of maps from S to N which is *incompressible*, i.e., injective on the level of π_1 , we can ask for any lamination $\lambda \in \mathcal{ML}(S)$ whether there exists a map in the given homotopy class which maps λ geodesically. If so, we say that λ is *realizable* in the homotopy class. We say that a foliation φ is realizable in a homotopy class if its corresponding lamination $\mathcal{L}(\varphi)$ is realizable. In a geometrically finite manifold every lamination is realizable (except possibly for a finite number of simple closed curves representing parabolics). See §5 for a brief discussion of how unrealizable laminations occur.

It is an easily missed point that the choice of $\lambda \in \mathcal{ML}(S)$ (which is a topological object) determines the possible hyperbolic metrics σ of the pleated surfaces containing λ in their pleating loci. It is not possible to prescribe both σ and λ arbitrarily, and in fact the dependence of σ on λ is not well understood.

Any complete realizable lamination λ is the pleating locus of the (unique) pleated surface obtained by mapping λ geodesically to N and pasting in ideal hyperbolic triangles for each complementary region in the only possible way. If a realizable lamination λ is not complete and λ' is obtained from it by the addition of a finite number of leaves (closed or not), λ' is realizable too, except when closed curves representing parabolic elements of $\pi_1(N)$ are among the added leaves [Th2, CEG]. Thus, any realizable lamination can be completed to a pleating locus for a pleated surface.

Thurston defines a condition on a map of a surface called *double incompressibility* which we will only state for the case of a closed surface, since that is the only situation that we will encounter. A map $f: S \rightarrow N$ from a closed surface to a hyperbolic 3-manifold is doubly incompressible if:

1. f is *incompressible*; i.e., f is injective on the level of π_1 .
2. f is *acylindrical*; i.e., if $c: S^1 \times I \rightarrow N$ is a map of a cylinder into N whose boundary factors through S according to $\partial c = f \circ c_0$, where $c_0: \partial(S^1 \times I) \rightarrow S$, and c is injective on π_1 , then c_0 extends to a map of the whole cylinder into S (or roughly speaking, if two homotopy classes in S are conjugate in N then they are conjugate in S as well).
3. Each maximal abelian subgroup of $\pi_1(S)$ is mapped to a maximal abelian subgroup of $\pi_1(N)$.

In particular, if f is an isomorphism of π_1 then it easily satisfies double incompressibility. It will turn out that, for our purposes, we can always reduce to this case by taking a lift of N to the cover determined by $f_*(\pi_1(S))$, which we denote $\tilde{N}_{f_*(\pi_1(S))}$.

Let \mathbf{PN} denote the projectivized tangent bundle of N , and given a map f pleated along $\lambda \subset M$ let $g: \lambda \rightarrow \mathbf{PN}$ be the lift of the restriction of f to λ . We can now state:

Theorem 2.1 (Uniform injectivity) (Thurston). *Fix $\epsilon_0 > 0$ and $A > 0$. Suppose $f: M \rightarrow N$ is a doubly incompressible pleated surface with M a hyperbolic surface of area at most A , mapping some lamination $\lambda \subset M$ geodesically. Then for every $\epsilon > 0$ there is a $\delta > 0$, depending only on ϵ_0 and A , such that for any two points $x, y \in \lambda$ which are in the ϵ_0 -thick part of M ,*

$$d(x, y) \geq \epsilon \Rightarrow d(g(x), g(y)) \geq \delta.$$

The proof hinges on a compactness argument. Thurston shows that the set of pleated surfaces with a choice of basepoint in the thick part of a hyperbolic 3-manifold is compact in the geometric topology, and therefore that a failure of uniform injectivity would result, after taking limits, in a pleated surface for which λ is not embedded. This in turn is ruled out by the homotopy conditions of double incompressibility.

We mention a small detail of the proof, as we will need to refer to it later. The theorem is stated for x and y in the thick part of M rather than N , and though it is obvious that an incompressible map takes $M_{\text{thin}(\epsilon_0)}$ to $N_{\text{thin}(\epsilon_0)}$ the same is false for the thick part. However, it is true that there is some $\epsilon_1 > 0$, depending only on A (i.e., the genus of M in our case) and ϵ_0 , such that $f(M_{\text{thick}(\epsilon_0)}) \subset N_{\text{thick}(\epsilon_1)}$. The reason is that a point in the ϵ_0 -thick part of M is contained in two loops of length bounded by some C (depending on ϵ_0 and A) which generate a free subgroup of rank 2 in $\pi_1(M)$, and therefore the same is true of its image.

If α is an arc in M with endpoints $x, y \in \lambda$ we define $d(x, y; \alpha)$ to be the distance in the universal cover of M between endpoints \tilde{x}, \tilde{y} of a lift of α . Similarly define $d(g(x), g(y); f(\alpha))$ to be the distance in the projectivized tangent bundle of \mathbf{H}^3 between the lifts $\tilde{g}(\tilde{x})$ and $\tilde{g}(\tilde{y})$. It is then easy to see that if $\epsilon < \epsilon_1/2$ and $\delta' = \min(\delta, \epsilon_1/2)$ (where $\delta > 0$ is given by Theorem 2.1), then if $x, y \in \lambda \cap M_{\text{thick}(\epsilon_0)}$ and $d(g(x), g(y)) < \delta'$, then the arc $\alpha \subset M$ connecting x and y such that $d(x, y; \alpha) < \epsilon$ is unique up to homotopy with endpoints fixed.

2.2. Applications of uniform injectivity. We develop now some consequences of Theorem 2.1 (Uniform injectivity) that will be used as tools in §4.

If L_1 and L_2 are geodesics in \mathbf{H}^3 and $s_0 > 0$, then define the *junction* (or *s₀-junction*) of L_1 and L_2 and $J = J_1 \cup J_2$, where

$$J_i = \{x \in L_i: d(x, L_j) \leq \max(s_0, d(L_1, L_2))\}$$

for $(i, j) = (1, 2)$ or $(2, 1)$. We have

Lemma 2.2 (Closed curves separate). *Given a doubly incompressible homotopy class of maps from S to N , let γ_1 and γ_2 be the geodesic representatives in N of two simple closed curves on S (which are either homotopically the same or disjoint), and let $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ be two of their lifts to \mathbf{H}^3 (where the lifts are distinct even if the curves are identical). There is a separation constant $s_0 > 0$, independent of anything but $\chi(S)$, such that the s_0 -junction of $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ has length at most $3 \cdot \min\{l_N(\gamma_1), l_N(\gamma_2)\}$.*

Proof. In a hyperbolic surface the lemma is obvious: suppose $l(\gamma_1) \leq l(\gamma_2)$. Since the translates of $\tilde{\gamma}_2$ by the translation along γ_1 of length $l(\gamma_1)$ are disjoint, elementary hyperbolic geometry gives a constant s such that the s -juncture of the geodesics has length at most $l(\gamma_1)$.

In three dimensions, the lemma is still obvious if either γ_1 or γ_2 is the core of a thin part, because thin parts are always separated by a definite amount in the universal cover, and a closed geodesic disjoint from the core of a thin part must avoid that thin part on the surface (and thus in the 3-manifold by the previous discussion).

Assume the contrary, so that neither curve is entirely contained in the thin part. We can construct a pleated surface containing γ_1 and γ_2 in its pleating locus, find a point $\tilde{\gamma}_2$ in the universal cover of the pleated surface which is outside the s -juncture (by the two-dimensional lemma) but within a $l_N(\gamma_1)$ -neighborhood of the s -juncture, and which is in the ϵ -thick part of N . We then apply the uniform injectivity theorem to argue that it is also outside the s_0 -juncture of the geodesics in the 3-manifold (where the new constant s_0 is given by the theorem). \square

The uniform injectivity theorem, as stated, has two drawbacks for us. It does not give an explicit relationship between ϵ and δ , and it only works in the thick part. We can solve both of these problems by replacing the condition on injectivity radius with a homotopy condition which we will soon find is natural for our purposes.

As in the original statement of uniform injectivity, let $f: M \rightarrow N$ denote a doubly incompressible pleated surface which contains a lamination λ in its pleating locus, and let g denote the lift to the projective tangent bundle of the restriction of f to λ .

Lemma 2.3 (Proportional injectivity). *If x and y are points on leaves μ and ν of λ such that there is an arc $\alpha \subset M$ connecting them whose interior does not meet μ or ν , then*

$$d(x, y; \alpha) \leq Cd(g(x), g(y); f(\alpha)),$$

where $C > 0$ depends only on the genus of S .

Proof. Choose ϵ_0 as in Theorem 2.1, and choose also $\epsilon_1 < \epsilon_0/2$, and δ_1 such that for all $x, y \in M_{\text{thick}(\epsilon_0)}$, if α is an arc in M connecting x and y , then

$$d(g(x), g(y); f(\alpha)) < \delta_1 \Rightarrow d(x, y; \alpha) < \epsilon_1.$$

Suppose that $d(g(x), g(y); f(\alpha)) = \delta$. For our statement it suffices to consider small δ , so assume $\delta < \delta_1$. Let $\tilde{\mu}$ and $\tilde{\nu}$ denote the lifts of $f(\mu)$ and $f(\nu)$ to \mathbf{H}^3 connected by a lift $\tilde{\alpha}$ of $f(\alpha)$, and let \tilde{x} and \tilde{y} denote the corresponding lifts of $f(x)$ and $f(y)$. Then $\tilde{\mu}$ and $\tilde{\nu}$ are δ -nearly tangent at \tilde{x} and \tilde{y} , so their δ_1 -junctures contain intervals of length at least $r_0 = \log(c\delta_1/\delta)$ centered on \tilde{x} and \tilde{y} , where c is some constant (by standard hyperbolic trigonometry).

Choose a side of \tilde{x} in $\tilde{\mu}$ and for $r < r_0$ let \tilde{x}_r be the point of $\tilde{\mu}$ on that side such that $d(\tilde{x}_r, \tilde{x}) = r$. Let \tilde{y}_r be the point of $\tilde{\nu}$ such that $d(\tilde{y}_r, \tilde{y}) = r$ and $d(\tilde{x}_r, \tilde{y}_r) \leq \delta_1$ (see Figure 1). Let x_r and y_r be the points in M such that $f(x_r)$ and $f(y_r)$ lift to \tilde{x}_r and \tilde{y}_r . x_r and y_r are then connected by an arc $\alpha_r \subseteq M$

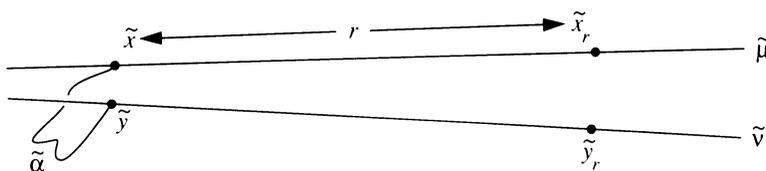


FIGURE 1

whose interior avoids μ and ν , and such that $d(g(x_r), g(y_r); f(\alpha_r)) < \delta_1(\alpha_r)$ is constructed by running from x_r to x near μ , along α to ν , and back to y_r).

If $x_r, y_r \in M_{\text{thick}(\epsilon_0)}$ then $d(x_r, y_r; \alpha_r) < \epsilon_1$. Otherwise, they are both in some component A of the ϵ'_0 -thin part of M (for ϵ'_0 slightly larger than ϵ_0), and we want to show that $d(x_r, y_r; \alpha_r)$ is nevertheless bounded (it must be the same component of $M_{\text{thin}(\epsilon'_0)}$ because it is the same component of $N_{\text{thin}(\epsilon'_0)}$ and f is doubly incompressible—or, in our case, π_1 -injective).

Observe first that α_r can be deformed into A (rel endpoints): $f(\alpha_r)$ can be deformed rel endpoints into the corresponding component B of $N_{\text{thin}(\epsilon'_0)}$ since its geodesic representative has length at most δ_1 . Since $f_*(\pi_1(A)) = \pi_1(B)$, this geodesic can be deformed into $f(A)$. Thus if α_r could not be deformed into A within M , we could compose it with this arc through A connecting x_r to y_r , so that the resulting loop is trivial in N but not in M , contradicting the π_1 -injectivity of f . The deformed arc can be assumed to have zero (geometric) intersection number with μ and ν in its interior, by making it geodesic in M .

Simple geodesics in a thin part either run all the way through, or spiral around the core. If μ or ν spiral around the core, then x_r and y_r must both be on the same side of the core, because of the existence of the arc homotopic to α_r in A with zero intersection number with μ and ν . In either case, the distance from x_r to ν is no more than ϵ'_0 , and likewise from y_r to μ . Let y' be the closest point on ν to x_r . If the distance along ν from y' to y_r is d then $f(\alpha_r)$ can be realized by a geodesic path of length d followed by a path of length at most ϵ'_0 , so $d(f(x_r), f(y_r); f(\alpha_r)) \geq d - \epsilon'_0$, implying that $d \leq \epsilon'_0 + \delta_1$.

Thus we see that the intervals of μ and ν of length r_0 centered on x and y remain a bounded distance $\epsilon_2 = \epsilon'_0 + \delta_1$ in M , as measured in the homotopy class of α . This implies that the distance from x to ν in the homotopy class of α is at most

$$c'e^{-r_0} = \left(\frac{c'\epsilon_2}{c\delta_1}\right)\delta.$$

where c' is some fixed constant, again by standard hyperbolic geometry. Repeating the last argument of the previous paragraph, we have

$$d(x, y; \alpha) \leq \left(1 + 2\frac{c'\epsilon_2}{c\delta_1}\right)\delta. \quad \square$$

This result can be applied to give the next theorem, which will be used in §4 to build train-tracks that approximate images of harmonic maps. Let $f: S \rightarrow N$ be a map taking a train-track τ' in S to a $(1, \epsilon)$ -nearly straight train-track τ in N . A complementary region P of τ' has cusps which are switches of τ'

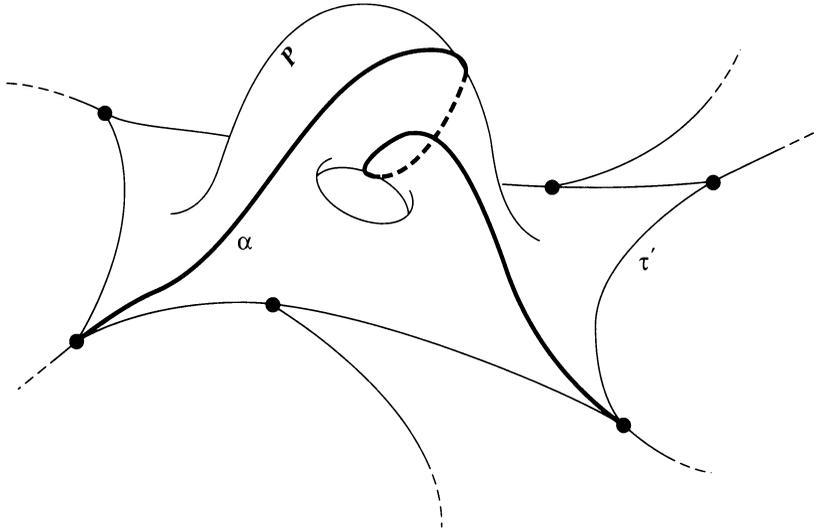


FIGURE 2. An arc in a complementary region of a train-track that joins two cusps can be added efficiently in the three-manifold.

in ∂P where the two incident branches of τ' on ∂P are incoming (see Figure 2). τ' can be *enlarged* by adding a branch α in a complementary region whose endpoints rest in two cusps (it might be the same cusp, as long as α cannot be deformed into ∂P). We say that τ can be enlarged by α to a $(1, \delta)$ -nearly straight train-track if the geodesic representative of $f(\alpha)$ in N (rel endpoints) meets the images of the cusps at angles of no more than δ . This is, in particular, obvious (with $\delta \sim \epsilon$) if N is a hyperbolic surface. In three dimensions there is the concern that τ may be folded up in some unfortunate way, but this is just what is prevented by the uniform injectivity theorem. Thus:

Theorem 2.4 (Enlarge train-track efficiently). *A $(1, \epsilon)$ -nearly straight train-track $\tau = f(\tau')$, where $f: S \rightarrow N$ is doubly incompressible, can be enlarged by any arc in its complement with endpoints in cusps of the track to give a $(1, C\epsilon)$ -nearly straight track. C depends only on the genus of S .*

Proof. Build a lamination λ that is carried on τ' and has a leaf on every branch (this can be done with a finite number of leaves, making a sequence of choices for each leaf at each switch, as long as we do not require λ to support a measure). Since τ is $(1, \epsilon)$ -nearly straight it is $c_1\epsilon$ -near to the geodesic representative of $f(\lambda)$ in N , for a universal c_1 . Extend α to a doubly infinite leaf $\tilde{\alpha}$ in the complement of λ by adding ends which follow alongside a leaf of λ at each of the cusps of τ' containing its endpoints (Figure 3). Complete $\lambda \cup \tilde{\alpha}$ to a complete lamination, and let $f_\mu: M \rightarrow N$ be the corresponding pleated surface, where M is S endowed with the appropriate hyperbolic metric, and μ is the geodesic representative of the complete lamination. Let $g_\mu: \mu \rightarrow \mathbf{P}N$ be the lift of f_μ to the projectivized tangent bundle, as before. We may assume (after an isotopy if necessary) that λ is geodesic in M .

Each cusp of τ' containing an endpoint of α corresponds to a region in $M - \mu$ bounded by leaves of μ whose images are ϵ -nearly tangent in N . In fact there is an arc β crossing this region, with endpoints, x, y on λ ,

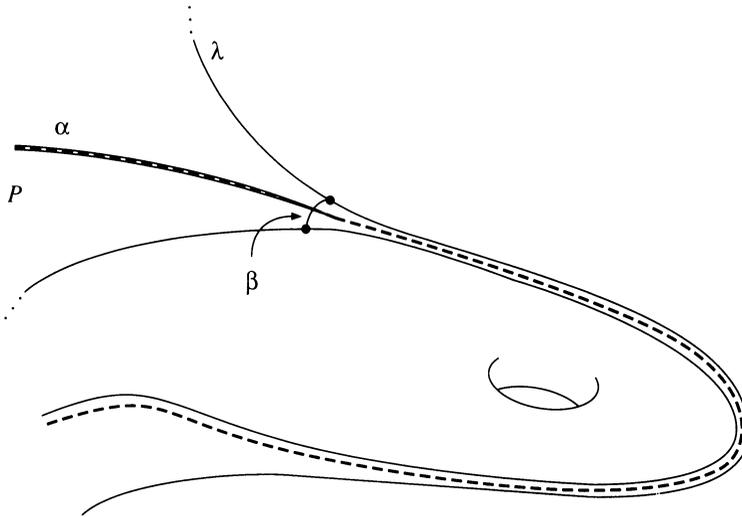


FIGURE 3. The continuation of α to an infinite leaf is shown by dotted lines. $\tilde{\alpha}$ can be defined, for example, to always follow the leaf of λ to its right.

which intersects the geodesic representing $\tilde{\alpha}$ at least once and whose interior is disjoint from λ , such that $d(g_\mu(x), g_\mu(y); f_\mu(\beta)) \leq c_1\epsilon$ (although $f_\mu(\beta)$ itself is not a priori short). An application of Lemma 2.3 then implies that in fact $d(x, y; \beta) \leq c_2\epsilon$ on M , and in particular that the geodesic representative of β in M is indeed short. Since the geodesic representative has the same geometric intersection number with $\tilde{\alpha}$ as β , it follows that the geodesic leaf representing $\tilde{\alpha}$ in $f_\mu(M)$ is $c_2\epsilon$ -nearly tangent to the image of the cusp of τ' . After a small adjustment of endpoints, we obtain the desired added branch to τ . \square

3. NOTATION AND ESTIMATES FOR HARMONIC MAPS

A harmonic map (see [ES, El, Jo]) $f: M \rightarrow N$ minimizes energy, or average squared stretching, in a homotopy class. When M is two-dimensional, harmonicity depends only on the conformal class of the metric on M . In particular, we recall from [M1] the following easy lower bound on the energy of f in terms of geometry of N and conformal invariants of M .

Proposition 3.1 (Energy lower bound). *For any simple closed curve $\gamma \subset M$ and any map $f: M \rightarrow N$, the energy of f is bounded below by*

$$\mathcal{E}(f) \geq \frac{1}{2} \frac{l_N^2 f(\gamma)}{E_M(\gamma)}.$$

Here E_M is extremal length in M , and $l_N(f(\gamma))$ is the infimum of lengths in N of representatives of $f(\gamma)$.

By continuity, this inequality extends to arbitrary measured laminations (or foliations).

Much of the geometry of f can be understood in terms of its Hopf differential, a holomorphic quadratic differential Φ defined as the $(2, 0)$ part of the pullback by f of the metric ρ of N . In particular, Φ defines a Euclidean

metric on M (with singularities at the zeros), with respect to which the pullback metric can be written

$$(3.1) \quad f^* \rho = 2(\cosh \mathcal{G} + 1)dx^2 + 2(\cosh \mathcal{G} - 1)dy^2,$$

where \mathcal{G} can be defined by

$$\sinh \mathcal{G} = \frac{\mathcal{J}}{2}$$

and \mathcal{J} is the absolute value of the Jacobian determinant of f with respect to the $|\Phi|$ -metric. We note that \mathcal{G} can blow up at the zeros of Φ , and is otherwise smooth except possibly at its zeros, where f fails to be an immersion. In this section we will show that \mathcal{G} becomes very small away from the zeros of Φ . This was done in [Wf2] and [M1] when the range is two-dimensional, and the argument here is much the same.

An area bound. If the harmonic map f is an immersion, a bound on the area of the image follows immediately from the Gauss-Bonnet theorem and the fact the the image is saddle-shaped everywhere (see [Sa]), and thus has Gaussian curvature no greater than -1 . In general, however, f may develop singularities, and a more careful argument is needed. Note that in the following we do not need to assume that f is incompressible.

Theorem 3.2 (Area bound). *If $f: M^2 \rightarrow N^n$ is a harmonic map and the sectional curvatures of N are bounded above by -1 , then*

$$\text{Area}(f(M)) \leq -2\pi\chi(M).$$

In particular, an inequality in [M1] relating pointwise energy and the Jacobian to Φ gives

$$(3.2) \quad 2\|\Phi\|_U \leq \mathcal{E}(U) \leq 2\|\Phi\|_U + 2\pi|\chi(M)|$$

for any measurable subset $U \subset M$ (where $\|\Phi\|_U$ is the area of U in the $|\Phi|$ -metric).

Proof of Theorem 3.2. We will perturb the problem slightly to remove the singularities. Let M_ϵ denote M with the hyperbolic metric scaled by ϵ , so that its curvature is $k_\epsilon = -1/\epsilon^2$. Let g_ϵ be the map of $M = M_1$ to the graph of f , with metrics as shown:

$$g_\epsilon = f \times \text{id}_\epsilon : M_1 \rightarrow N \times M_\epsilon,$$

where id_ϵ is the identity on M , considered as a map from M_1 to M_ϵ . It is clear that

- g_ϵ is an embedding,
- g_ϵ is harmonic, and
- $\lim_{\epsilon \rightarrow 0} \text{Area}(g_\epsilon(M)) = \text{Area}(f(M))$

so it would be enough to bound the area of $g_\epsilon(M)$. The complication is that $N \times M_\epsilon$ is not negatively curved in all directions—planes spanned by $\{0 \times \alpha, \beta \times 0\}$, where $\alpha \in TM_\epsilon$ and $\beta \in TN$ have zero sectional curvature. The area of $g_\epsilon(M)$ can therefore be large, but the trouble will disappear as $\epsilon \rightarrow 0$.

The sectional curvature of a plane in $T(N \times M_\epsilon)$ spanned by two vectors X and Y is

$$K(X, Y) = \frac{\langle R(X, Y)X, Y \rangle}{|X \wedge Y|^2}$$

where R is the Riemann tensor of $N \times M_\epsilon$. Any copy of M_ϵ or N is totally geodesic in the product, and this is expressed in the fact that $\langle R(X, Y)Z, W \rangle$ vanishes if some of X, Y, Z , and W are contained in TN and the rest are in TM_ϵ . On the other hand, $\langle R(X, Y)X, Y \rangle = k_\epsilon |X \wedge Y|^2$ for $X, Y \in TM_\epsilon$ and $\langle R(X, Y)X, Y \rangle \leq -|X \wedge Y|^2$ for $X, Y \in TN$.

Thus, if u and v are an orthonormal basis for TM_1 at some point, then the sectional curvature of the plane spanned by their images is

$$K(g_{\epsilon*}u, g_{\epsilon*}v) = \frac{\langle R(\text{id}_{\epsilon*}u, \text{id}_{\epsilon*}v) \text{id}_{\epsilon*}u, \text{id}_{\epsilon*}v \rangle + \langle R(f_*u, f_*v) f_*u, f_*v \rangle}{|g_{\epsilon*}u \wedge g_{\epsilon*}v|^2} \leq \frac{k_\epsilon \epsilon^4 - |f_*u \wedge f_*v|^2}{|g_{\epsilon*}u \wedge g_{\epsilon*}v|^2} = -\frac{\epsilon^2 + \mathcal{J}_f^2}{\mathcal{J}_{g_\epsilon}^2},$$

where \mathcal{J}_f and \mathcal{J}_{g_ϵ} are the absolute values of the Jacobians of f and g_ϵ , respectively, relative to the given metrics. One easily computes

$$\mathcal{J}_{g_\epsilon}^2 = \mathcal{J}_f^2 + 2e_f \epsilon^2 + \epsilon^4,$$

where e_f is the energy of f .

Now we can use the fact that $g_\epsilon(M)$ is saddle-shaped to argue that its curvature at every point is no more than the ambient sectional curvature of the tangent plane in $N \times M_\epsilon$. So, by the Gauss-Bonnet Theorem,

$$(3.3) \quad -2\pi\chi(M) \geq \int_{g_\epsilon(M)} \frac{\epsilon^2 + \mathcal{J}_f^2}{\mathcal{J}_{g_\epsilon}^2} dA(g_\epsilon(M)).$$

Fixing $\lambda > 0$, define $M_\lambda = \{p \in M : \mathcal{J}_f(p) \geq \lambda\epsilon\}$. If we let $e_M = \max_M e_f$, we see that the image area of the complement of M_λ is

$$\int_{M-M_\lambda} \mathcal{J}_{g_\epsilon} dA(M) \leq \epsilon \sqrt{\lambda^2 + 2e_M + \epsilon^2} \text{Area}(M),$$

so it will disappear in the limit as $\epsilon \rightarrow 0$. On M_λ itself we have

$$\begin{aligned} \int_{g_\epsilon(M_\lambda)} \frac{\epsilon^2 + \mathcal{J}_f^2}{\mathcal{J}_{g_\epsilon}^2} dA(g_\epsilon(M)) &= \int_{g_\epsilon(M_\lambda)} \frac{1 + \epsilon^2/\mathcal{J}_f^2}{1 + 2\epsilon^2 e_f/\mathcal{J}_f^2 + \epsilon^4/\mathcal{J}_f^2} dA(g_\epsilon(M)) \\ &\geq \int_{g_\epsilon(M_\lambda)} \frac{1}{1 + (2e_M + \epsilon^2)/\lambda^2} dA(g_\epsilon(M)) \\ &= \frac{1}{1 + (2e_M + \epsilon^2)/\lambda^2} \text{Area}(g_\epsilon(M_\lambda)). \end{aligned}$$

Combining these with (3.3), we get

$$\begin{aligned} -2\pi\chi(M) &\geq \frac{\text{Area}(g_\epsilon(M_\lambda))}{1 + (2e_M + \epsilon^2)/\lambda^2} \\ &\geq \frac{1}{1 + (2e_M + \epsilon^2)/\lambda^2} (\text{Area}(g_\epsilon(M)) - \epsilon \sqrt{\lambda^2 + 2e_M + \epsilon^2} \text{Area}(M)). \end{aligned}$$

Letting ϵ go to zero, we have

$$\text{Area}(f(M)) \leq -2\pi\chi(M)(1 + 2e_M/\lambda^2),$$

but, since λ could have been arbitrarily large,

$$\text{Area}(f(M)) \leq -2\pi\chi(M). \quad \square$$

Bounds on \mathcal{E} . In [Wf2, Sa] an equation is obtained that governs \mathcal{E} when the image is two-dimensional:

$$(3.4) \quad \Delta \mathcal{E} = -4K \sinh \mathcal{E}.$$

Here K is the sectional curvature of the image and Δ is the Laplacian relative to the $|\Phi|$ -metric. This equation holds as well when the image is three-dimensional and the map is an immersion, where now K is the intrinsic sectional curvature of the image surface. In particular, $K \leq -1$. It only fails on the zeros of Φ and on the zeros of \mathcal{E} . \mathcal{E} is therefore subharmonic away from the zeros of Φ .

As in [M1], this and the bound on $\text{Area}(f(M))$ immediately give

Lemma 3.3 (Rough bound). *Let $p \in M$ be a point with a neighborhood U such that U contains no zeros of Φ , and is a round disk of radius r centered on p in the $|\Phi|$ -metric. Then there is a bound*

$$\mathcal{E}(p) \leq \sinh^{-1}(|\chi(M)|/r^2).$$

We can also obtain, with a small adjustment in the proof, the following

Lemma 3.4 (Exponential estimate). *Let $p \in M$ be a $|\Phi|$ -distance $> d > 0$ from any zeros of Φ , and let \mathcal{E} be bounded above by B in a neighborhood of $|\Phi|$ -radius d around p . Then*

$$\mathcal{E}(p) < B / \cosh d.$$

Proof. Equation (3.4) holds away from the zero set Σ of \mathcal{E} , in a Euclidean disk of d about p , which is the lift to the universal cover of a d -neighborhood of p in the $|\Phi|$ -metric.

We can easily define a comparison function F on the disk such that $F \geq B$ on the boundary, $\Delta F = 4F$ everywhere, and $F(p) = B / \cosh d$. A simple maximum principle argument shows that $F \geq \mathcal{E}$ everywhere: on the complement of Σ ,

$$\Delta(F - \mathcal{E}) = 4F + 4K \sinh \mathcal{E} \leq 4(F - \mathcal{E}),$$

so either $F - \mathcal{E}$ achieves its minimum on Σ , where $F - \mathcal{E} = F \geq 0$, or off Σ , in which case $\Delta(F - \mathcal{E}) \geq 0$ and again $F - \mathcal{E} \geq 0$. \square

Curvature estimates. In the case of a diffeomorphic harmonic map between surfaces, Wolf [Wf2] gets derivative bounds on \mathcal{E} , which then give similar estimates for the geodesic curvature k_h of the images of the horizontal arcs of Φ . In our situation this is not enough, not just because \mathcal{E} is not everywhere smooth, but mainly because in the higher-dimensional range the geodesic curvatures depend on more than just the derivatives of \mathcal{E} .

Harmonic maps satisfy the following composition property (see [Jo]):

Theorem (Composition with convex function). *If $f: M \rightarrow N$ is harmonic and $\phi: N \rightarrow \mathbf{R}$ is a C^2 function, then the Laplacian on M of $\phi \circ f$ is*

$$\Delta(\phi \circ f) = \sum_{\alpha} \text{Hess } \phi(df(\varepsilon_{\alpha}), df(\varepsilon_{\alpha})),$$

where the sum is over an orthonormal frame $\{\varepsilon_{\alpha}\}$ in M .

In particular, if ϕ is convex then $\phi \circ f$ is subharmonic.

(Recall that a function is convex if its Hessian is nonnegative, and subharmonic if its Laplacian is nonnegative.)

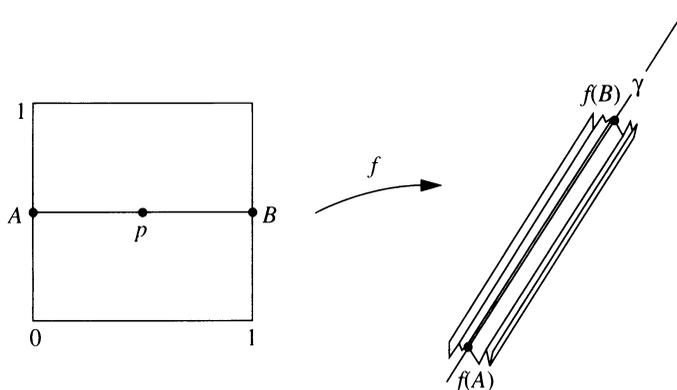


FIGURE 4. Q is parametrized as $\{(u_1, u_2): 0 \leq u_i \leq 1\}$, with the u_1 -axis in the horizontal direction of Φ , and $p = (\frac{1}{2}, \frac{1}{2})$.

Theorem 3.5 (Curvature bounds). *If $p \in M$ has a neighborhood of radius 1 in the $|\Phi|$ -metric which contains no zeros of Φ and on which $\mathcal{G} \leq \epsilon$ for some (small) $\epsilon > 0$, then the geodesic curvature of the image of the horizontal Φ -trajectory at p satisfies $k_h \leq c\epsilon$ for a fixed $c > 0$.*

We note that this estimate is undoubtedly true in a more general setting, for example when the range curvatures are bounded between two negative constants (probably using the methods of [JK1]). Our proof will be for a 3-dimensional hyperbolic target.

Proof of Theorem 3.5. Lifting to the universal covers of M and N , we have a harmonic map from a round Euclidean disk in the $|\Phi|$ -metric to \mathbf{H}^3 . We will introduce a convenient coordinate system and prove that all the relevant first and second derivatives are small in those coordinates. The constants c_1, c_2, \dots that arise are all independent of f .

Restrict the map to a square Q of sidelength 1, centered on p , whose sides are horizontal and vertical Φ -trajectories (see Figure 4), parametrized by (u_1, u_2) . In those coordinates let $A = (0, \frac{1}{2})$ and $B = (1, \frac{1}{2})$, and let $\gamma \subset \mathbf{H}^3$ denote the geodesic passing through $f(A)$ and $f(B)$. Let π_1 and π_2 be orthogonal hyperbolic planes intersecting in γ and let $x_1, x_2: \mathbf{H}^3 \rightarrow \mathbf{R}$ be signed distance functions from π_1 and π_2 , respectively (“signed” meaning that x_i takes opposite signs on opposite sides of π_i). Define a third coordinate x_3 by composing signed distance along γ (relative to some basepoint) with orthogonal projection of \mathbf{H}^3 to γ . The hyperbolic metric in the (x_1, x_2, x_3) coordinates is

$$(3.5) \quad ds_h^2 = \lambda_1 dx_1^2 + \lambda_2 dx_2^2 + \lambda_3 dx_3^2 - \mu dx_1 dx_2,$$

where λ_i and μ are analytic functions of x_1 and x_2 satisfying

$$\lambda_i(x_1, x_2) = 1 + O(|(x_1, x_2)|^2)$$

and

$$\mu(x_1, x_2) = O(|(x_1, x_2)|^2).$$

The Hessians of the x_i can be computed using standard hyperbolic geometry (with a little care to obtain an answer in terms of the x_i coordinates). One

contains an answer of the form

$$\text{Hess } x_i = \sum_{j,k=1}^3 \varphi_{jk}^i(x_1, x_2) dx_j dx_k \quad (i = 1, 2, 3),$$

where the φ_{jk}^i are analytic functions vanishing to first order on γ :

$$\varphi_{jk}^i(x_1, x_2) = O(|(x_1, x_2)|).$$

Define $f_i = x_i \circ f$. The composition theorem gives

$$(3.6) \quad \Delta f_i = \sum_{j,k=1}^3 \sum_{m=1}^2 \varphi_{jk}^i(f_1, f_2) \left(\frac{\partial f_j}{\partial u_m} \right) \left(\frac{\partial f_k}{\partial u_m} \right) \quad (i = 1, 2, 3).$$

For Δf_1 and Δf_2 there is more information: since \mathbf{H}^3 is negatively curved and π_k ($k = 1, 2$) is convex set, x_k is convex when positive and concave when negative. This together with the composition theorem implies that

$$h_k = |f_k| \quad (k = 1, 2)$$

are subharmonic. We can therefore bound h_k as follows.

Since $\mathcal{G} \leq \epsilon$ on Q , we invoke (3.1) to show that $f(Q)$ is contained in a ball of bounded radius, and that the total length of the image of a vertical arc is bounded by

$$l_{\mathbf{H}^3}(f(\{u_1\} \times [0, 1])) \leq \int_0^1 \sqrt{2(\cosh \mathcal{G}(u_1, y) - 1)} dy \leq c_1 \epsilon$$

for some constant c_1 near 1 (for small enough ϵ). Thus, using (3.5), the variation of h_k on a vertical arc of Q is bounded by $c_2 \epsilon$. In particular, $h_k \leq c_2 \epsilon$ on $\{0\} \times [0, 1]$ and $\{1\} \times [0, 1]$.

Let $H_k = \max_Q h_k$ and let $\{a_k\} \times [0, 1]$ be some vertical arc on which the maximum is achieved. Then

$$h_k(a_k, \frac{1}{2}) \geq H_k - c_2 \epsilon.$$

On the other hand, since h_k is subharmonic,

$$h_k(a_k, \frac{1}{2}) \leq \mu\text{-Avg}_{\partial Q} h_k,$$

where μ is the harmonic, or “visual,” measure, based at $(a_k, \frac{1}{2})$, on ∂Q in the Poincaré metric on Q (in other words, μ is the pullback on the standard measure of S^1 by any Riemann mapping taking Q to the unit disk and $(a_k, \frac{1}{2})$ to the origin). Since $\mu(\{0\} \times [0, 1] \cup \{1\} \times [0, 1]) \geq \frac{1}{4}$ for any a_k , it follows that

$$h_k(a_k, \frac{1}{2}) \leq \frac{3}{4} H_k + c_2 \epsilon.$$

Thus

$$H_k - c_2 \epsilon \leq \frac{3}{4} H_k + c_2 \epsilon.$$

We conclude that

$$(3.7) \quad |f_k| \leq 8c_2 \epsilon \quad (k = 1, 2).$$

Thus $f(Q)$ is contained in a thin cylinder around γ . Derivative bounds come directly from (3.1), (3.5), and the fact that $\mathcal{E} \leq \epsilon$:

$$(3.8) \quad |\nabla f_i| \leq c_3 \quad (i = 1, 2, 3).$$

Since $f(Q)$ is contained in a ball of bounded radius, there are a priori estimates (see [Jo, JK1, JK2]) on second derivatives:

$$(3.9) \quad |D^2 f_i| \leq c_4 \quad (i = 1, 2, 3).$$

We can now employ the following standard gradient estimate [GT, §3.4]:

$$|\nabla w| \leq c_5 \sup_Q |w| + c_6 \sup_Q |\nabla w|$$

in, say, a subsquare Q' of sidelength $\frac{1}{2}$. With the bounds on f_k and Δf_k given by (3.6) and (3.7) for $k = 1, 2$, we have

$$(3.10) \quad |\nabla f_k| \leq c_7 \epsilon \quad (k = 1, 2)$$

on Q' .

We observe now that the right-hand side of (3.6) is bounded by $c_8 \epsilon$ in Q' , in light of the previous estimates, and similarly for its first derivatives. This gives us an estimate on second derivatives [GT, Theorem 4.6]:

$$(3.11) \quad |D^2 f_i| \leq c_9 \epsilon \quad (i = 1, 2, 3).$$

The curvature k_h is $k_h = |\nabla_\xi^N \xi|_N$, where ξ is the unit vector

$$\xi = df(\partial/\partial u_1)/|df(\partial/\partial u_1)|_N.$$

From (3.1) we have $|df(\frac{\partial}{\partial u_1})|_N \geq 2 - c_{10} \epsilon$, and the first derivatives of this norm are bounded by $c_{11} \epsilon$ by (3.11). The Christoffel symbols for the metric (3.5) are bounded in a bounded neighborhood, and in particular $\Gamma_{33}^k = O(|(x_1, x_2)|)$ (since γ is a geodesic). A standard differential geometric computation then suffices to show that $k_h \leq c \epsilon$. \square

Existence for a noncompact target. We conclude this section with a proof that the harmonic maps in which we are interested actually exist. An argument is needed because our case (a possibly noncompact hyperbolic range) is not covered by the existence theorems generally found in the literature. The machinery for it is available, however, so the proof is just a twist on an argument of Uhlenbeck [Uh].

Theorem 3.6 (Existence). *Let M be a closed Riemann surface and N a Riemannian manifold, not necessarily compact, whose sectional curvatures are bounded above by $k < 0$. Let $f: M \rightarrow N$ be a map such that at least one simple closed curve of M is mapped to a curve whose minimal length in N is positive. Then there is a harmonic map in the homotopy class of f .*

The map is then unique, smooth, by standard methods. We are interested, of course, in the case where N is a hyperbolic 3-manifold. The condition on the simple closed curve is then easily met, for example, when N has no parabolics, or when f is injective on the level of fundamental group (and the genus of M is at least 2).

Proof of Theorem 3.6. The trouble with proving that an energy-minimizing sequence of maps converges is that a bound on energy does nothing to control

the diameter of the image of a map. Small disks on the surface can have extremely long images at a small cost in energy. Uhlenbeck solves this problem by replacing the energy functional with a perturbed functional

$$\mathcal{E}_\epsilon(f) = \mathcal{E}(f) + \epsilon \int_M |df|^p$$

(where $p > \dim M$). A bound on this functional can be shown to give a bound on the diameter of $f(M)$, and Uhlenbeck imposes a growth condition on the injectivity radius of N which then implies that a sequence of maps of bounded diameter must remain in a compact subset of N . She can then find minimizing maps for \mathcal{E}_ϵ and, letting ϵ go to zero, produce harmonic maps.

Our task, therefore, is simply to show that a bound on the diameter of a map under our conditions is enough to keep the image in a compact set. More precisely, a bound on diameter together with a bound on energy will suffice.

Suppose $\text{diam}(f(M)) = \max_{x,y \in M}(d_N(f(x), f(y))) < D$. Let $\gamma \subset M$ be the simple closed curve promised in the statement of the theorem, and let $f(\gamma)^*$ be the geodesic representative of its image in N . The curvature restriction implies that any representative of $f(\gamma)$ which is outside an R -neighborhood of $f(\gamma)^*$ has length at least $l_N(f(\gamma)^*) \cosh(\sqrt{-k}R)$. Thus if any point of $f(M)$ is outside an $(R+D)$ -neighborhood of $f(\gamma)^*$, the energy of f is at least $\frac{1}{2} \cosh^2(\sqrt{-k}R) l_N^2(f(\gamma)^*)/E(\gamma)$, by Proposition 3.1 (energy lower bound). \square

4. STRUCTURE THEOREMS FOR DEGENERATING HARMONIC MAPS

Let M denote a closed Riemann surface of genus $g > 1$ and $f: M \rightarrow N$ an incompressible harmonic map to a hyperbolic 3-manifold N . "Incompressible" means that f is injective on π_1 , and we will from now on assume (lifting if necessary to a cover of N) that it is an isomorphism of π_1 . In particular, we may assume any pleated surface in the homotopy class of f is doubly incompressible.

Let Φ be the Hopf differential of f and Φ_h its horizontal foliation. The strategy of the argument is to construct a nearly-straight train-track in N that carries the leaves of $f(\Phi_h)$.

The portion of the argument in [M1] that takes place on the domain M carries through without change, the necessary estimates having been supplied in §3. Thus we obtain, for each $R > 0$, a subsurface $\mathcal{P}_R \subset M$ containing an R -neighborhood (in the $|\Phi|$ -metric) of the zeros of Φ . Each component γ of $\partial\mathcal{P}_R$ is of one of the following two types:

1. An alternating sequence of horizontal and vertical arcs of Φ , meeting at angles of $\pi/2$ with respect to \mathcal{P}_R ; in this case we call γ a *polygonal boundary component*.
2. A geodesic core of a flat cylinder—in which case we call γ a *straight boundary component*.

The dimensions of \mathcal{P}_R (again, in the $|\Phi|$ -metric) are controlled as follows:

1. $l(\partial\mathcal{P}_R) \leq K_1 R$.
2. $\text{Area}(\mathcal{P}_R) \leq A + K_2 R^2$.
3. Each edge of a polygonal boundary component has length at least $K_3 R$.
4. The polygonal components of $\partial\mathcal{P}_R$ are s_1 -separated.

Two boundary components γ_1, γ_2 of \mathcal{P}_R are s_1 -separated if the $|\Phi|$ -distance between them in the complement of \mathcal{P}_R is at least $s_1 \max(l_{|\Phi|}(\gamma_1), l_{|\Phi|}(\gamma_2))$.

Here s_1 is an arbitrary positive parameter (chosen so as to make the subsequent arguments work) while A_1, K_1, K_2 , and K_3 depend only on $\chi(M)$ and s_1 , and on an additional parameter s_2 described below.

The straight boundary components of \mathcal{P}_R are further required to appear in pairs bounding flat (i.e., Euclidean) cylinder components of $M - \mathcal{P}_R$. Dealing with these cylinders is the major technicality in the rest of the proof, and their dimensions are regulated carefully. If $\mathcal{F}_1, \dots, \mathcal{F}_m$ are the flat-cylinder components of $M - \mathcal{P}_R$ and $\hat{\mathcal{F}}_i$ is the maximal flat cylinder containing \mathcal{F}_i for each i , then for any choice of the second parameter $s_2 > 0$ the construction is arranged so that

5. $\text{Area}(\hat{\mathcal{F}}_i \cap \mathcal{P}_R) \geq s_2 \text{Area}(\mathcal{P}_R - \bigcup \hat{\mathcal{F}}_j)$ for each i
 (where $\text{Area}(U) = \|\Phi\|_U$ is area in the $|\Phi|$ -metric).

Finally, \mathcal{P}_R has the property that bounds on \mathcal{E} hold exponentially well in its complement:

$$\mathcal{E}(p) \leq \varepsilon(R + d(p, \mathcal{P}_R)),$$

where $d(p, \mathcal{P}_R)$ is $|\Phi|$ -distance to \mathcal{P}_R , and $\varepsilon(x)$ denotes a function bounded by ae^{-bx} , with a and b depending only on $\chi(M)$.

Ignoring the cylinder components for a moment, the complement of \mathcal{P}_R has a natural train-track structure carrying the leaves of Φ_h , where a branch is obtained from a rectangle with horizontal and vertical boundaries by identifying vertical leaves to points. Switches occur at the vertical segments in the polygonal boundaries. An adjustment to the construction of \mathcal{P}_R ensures that the (horizontal) length of each branch is at least 1. The bound on \mathcal{E} and the curvature estimates in §3 then imply that the horizontal leaves in any rectangle map to paths of curvature $\varepsilon(R)$ which are no more than $\varepsilon(R)$ apart. Thus, taking one representative leaf in each rectangle and making a small adjustment at the endpoints, we obtain a train-track whose image in N is a $(1, \varepsilon(R))$ -nearly straight train-track. This train-track carries all the leaves of Φ_h except those that enter \mathcal{P}_R , and those contained in flat cylinders. The rest of the proof thus consists in showing that the train-track can be enlarged to one that carries all of Φ_h .

The main difficulty is in adjoining the cylinders. We will summarize the proof given in [M1], supplying the necessary details for its generalization.

Denote by τ the $(1, \varepsilon(R))$ -nearly straight train-track in N obtained from the image of $M - \mathcal{P}_R - \{\text{flat cylinders}\}$. The image in N of a leaf of Φ_h from any flat cylinder is an arc in N of curvature $\varepsilon(R)$ —we will inductively join each of these arcs to τ with arcs that are themselves nearly straight, that meet τ and the flat cylinder arcs at $\varepsilon(R)$ -small angles, and that are homotopic (rel endpoints) to arcs of Φ_h in \mathcal{P}_R .

Let \mathcal{F} denote the next flat cylinder to be adjoined to τ , and let b be one of its two boundary components. A generic arc α of Φ_h in \mathcal{F} enters \mathcal{P}_R through b and eventually exits again, either at (1) a polygonal boundary of \mathcal{P}_R or a flat cylinder that has already been adjoined to τ , or (2) a flat cylinder \mathcal{F}' that has not yet been adjoined (possibly $\mathcal{F} = \mathcal{F}'$). The construction of [M1] ensures that α has an initial segment α_1 which is within the maximal flat cylinder $\hat{\mathcal{F}}$ containing \mathcal{F} , and whose image $f(\alpha_1)$ is $\varepsilon(R)$ -nearly straight. In case (2) there is also a final segment $\alpha_2 \subset \hat{\mathcal{F}}'$, with the same properties.

Further, by means of an averaging argument using the area bounds on \mathcal{P}_R and (3.2), it is guaranteed that

$$l_N(f(\alpha_i)) \geq Q_1 l_N(f(\alpha_m))$$

and

$$l_N(f(\alpha_i)) \geq Q_2 R$$

for $i = 1, 2$ for any a priori choice of constants Q_1 and Q_2 , where α_m is the middle segment of α between α_1 and α_2 (to accomplish this it is also necessary to pick a particular order in which to adjoin the cylinders, based on the angle of winding of the leaves of Φ_h in each).

The task now is to show that the geodesic representative $f(\alpha)^*$ of $f(\alpha)$ (rel endpoints) meets it at angles bounded by $\varepsilon(R)$. $f(\alpha)^*$ together with any leaf of Φ_h in \mathcal{F} (and the corresponding arc on the other side of the cylinder) can be joined to τ as a new branch. We recall a lemma from [M1]:

Let β be a path in \mathbf{H}^3 which is composed of a sequence of three segments β_1, β_m , and β_2 , such that β_1 and β_2 are geodesics. Let L_i be the infinite geodesic containing β_i . Denote by b_i the endpoint common to β_i and β_m , and by p_i the other endpoint of β_i . Let R_i be the ray of L_i based at p_i and containing β_i . Let θ_i be the angle at p_i between β and its geodesic representative $\beta^* = \overline{p_i b_i}$. Recall from §2 the definition of the s_0 -juncture $J(L_1, L_2)$.

Lemma 4.1 (Straighten corners). *There is an inversely exponential function ε and a constant A (depending only on s_0) such that*

(a) *If $l(\beta_i) > l(\beta_m) + l(J_i) + A$ ($i = 1, 2$), then*

$$\theta_i \leq \varepsilon(l(\beta_i) - l(\beta_m) - l(J_i) - A).$$

(b) *If $l(\beta_2) > l(\beta_m) + A$ and $J_1 \subset R_1 - \beta_1$, then*

$$\theta_1 \leq \varepsilon(l(\beta_1)), \quad \theta_2 \leq \varepsilon(l(\beta_2) - l(\beta_m) - A).$$

(This lemma is discussed in a two-dimensional context in [M1] but it is elementary to extend it to three dimensions.)

With this lemma in mind, we proceed to break our situation down into cases. Let \mathcal{E} denote \mathcal{F} or \mathcal{F}' . A vertical arc in \mathcal{E} with both endpoints on a given leaf of Φ_h has $|\Phi|$ -length $V \leq K_1 R$ and therefore its image has N -length $v \leq \varepsilon(R)$. Denote by l the N -length of the ($\varepsilon(R)$ -nearly straight) image of the segment of Φ_h between the endpoints of the vertical arc. Choose ε_1 as in §2 to insure that the ε_0 -thick part of any pleated surface is in the ε_1 -thick part of N , and choose $\varepsilon_2 < \varepsilon_1$. We have one of:

- (a) $l \leq \varepsilon_2/2$. Then assuming R is large enough (depending only on ε_2) so that $v \leq \varepsilon_2/2$, the image of the entire cylinder lies in the ε_2 -thin part of N corresponding to its core curve.
- (b) $l > \varepsilon_2/2$. Then the horizontal segment makes an $(\varepsilon_2/2, \varepsilon(R))$ -nearly straight broken circuit. We conclude that for large enough R (again depending only on ε_2) the leaves of Φ_h in \mathcal{E} are $\varepsilon(R)$ -near to the geodesic γ in N corresponding to the core curve of \mathcal{E} .

In case (b), $l_N(\gamma) \leq K_4 R$ for a fixed constant K_4 (depending on the construction of \mathcal{P}_R). There are now several possibilities:

Case I. Both α_1 and α_2 are in case (b) above—their images spiral around closed geodesics γ_1 and γ_2 in N . Lift $f(\alpha)$ to the universal cover \mathbf{H}^3 , let $\tilde{\gamma}_i = L_i$ be the corresponding lifts of γ_i , and let β be obtained from $f(\alpha)$ by projecting $f(\alpha_i)$ to L_i . By Lemma 2.2 (Closed curves separate), with corresponding choice of s_0 , we have $l(J_i) \leq 3l_N(\gamma_1) \leq 3K_4R$. Thus by appropriate choice of Q_1 and Q_2 we can use Lemma 4.1 to conclude that the angles between β and β^* are bounded by $\varepsilon(R)$. The same estimate then holds for the $\varepsilon(R)$ -nearby $f(\alpha)$.

Case II. α_1 is in case (a), and α_2 is in case (b) (or vice versa). Lifting to \mathbf{H}^3 , let L_2, γ_2 be as before, and let $\tilde{T}_{\epsilon_i} (i = 0, 1, 2)$ be the corresponding lift of the ϵ_i -thin part in N containing $f(\alpha_1)$. Let L_1 be the infinite geodesic containing the endpoints of the lift of $f(\alpha_1)$ and define β and R_i as above. Since γ_2 and γ_1 (the core of the cylinder containing α_1) represent disjoint, homotopically distinct curves on M , the geodesic γ_2 must remain outside the ϵ_0 -thin part of γ_1 in any pleated surface containing both in its pleating locus (note that this argument works even if γ_1 is a parabolic in N). Hence, by the choice of ϵ_1 , L_2 is disjoint from \tilde{T}_{ϵ_1} . It follows that L_1 is a definite distance d , depending on the choice of ϵ_2 , from \tilde{T}_{ϵ_2} . As mentioned in §2, we can choose constants so that $d > s_0$.

If $J_1 \subset R_1 - \beta_1$, we can (given a choice of constants as in the previous paragraph) use part (a) of Lemma 4.1 to conclude that $\theta \leq \varepsilon(R)$. If J_1 is not contained in $R_1 - \beta_1$, then such $d(b_1, J_1) \leq l_N(\beta_m) + A$ and $l_N(\beta_1) \geq Q_1 l_N(\beta_m)$ we conclude that J_1 must meet β_1 (provided we choose Q_1 appropriately). But since as above $d(\beta_1, L_2) > d > s_0$ we must have $l(J_1) = 0$. In this case we can use part (b) of Lemma 4.1.

Case III. Both α_1 and α_2 are in case (a)—they are in ϵ_2 thin parts of N . Use essentially the same arguments as above. Lift to \mathbf{H}^3 and find geodesics L_1, L_2 to which the lifts are $\varepsilon(R)$ -nearly tangent. If $J_i \subset R_i - \beta_i$ for either $i = 1$ or 2 we may apply Lemma 4.1 (b). If not we can again argue that the junctures meet β_1 and β_2 , so that (using the separation of thin parts) $l(J_i) = 0$ and Lemma 4.1 (a) applies.

Case IV. $f(\alpha_1)$ spirals around γ_1 , and the other end of α meets a polygonal boundary of \mathcal{P}_R or an already adjoined flat cylinder. This must correspond to a switch of the currently existing train-track τ . Following the argument of Theorem 2.4 (Enlarge train-track efficiently), extend α to an infinite leaf μ spiraling around γ_1 at one end and following a lamination λ carried on τ at the other. As in that theorem, we can conclude that μ is realized in N as a geodesic to which $f(\alpha)$ is $\varepsilon(R)$ -nearly tangent at the far endpoint. Further, since by our choice of constants $l_N(f(\alpha_1)) - l_N(f(\alpha_m))$ is proportional to R , $f(\alpha)$ must also be $\varepsilon(R)$ -nearly tangent to μ at its α_1 endpoint.

Case V. α_1 is in case (a) and the other end of α meets a polygonal boundary of \mathcal{P}_R or an already adjoined cylinder. This time we have to look at the other end of \mathcal{F} . There is an arc α' there to which all the same considerations apply. Let $\hat{\alpha}$ be the broken path formed by α, α' , and any leaf of Φ_h in \mathcal{F} between them. If α' meets another polygonal arc or adjoined cylinder at its far endpoint, we apply Theorem 2.4 (Enlarge train-track efficiently). If it meets a case (b) cylinder we extend it to a leaf that spirals around the corresponding

geodesic and again apply the theorem. If it meets a case (a) cylinder \mathcal{F}'' we first straighten α' using Case III, and then repeat the process by looking at the opposite end of \mathcal{F}'' , until it terminates in one of the previous cases. In the end, we have a leaf terminating at both ends in switches of τ , with at most two nonstraight segments whose lengths are a definite factor shorter than any of the straight segments, so again we can apply Theorem 2.4.

Having adjoined branches representing each flat cylinder in this way, it remains to add branches to carry the leaves of Φ_h in \mathcal{P}_R that terminate on either a polygonal or a straight boundary. This is accomplished by an application of Theorem 2.4. If there is a subfoliation of Φ_h entirely contained within \mathcal{P}_R we can just ignore it, or if we like build a train-track to carry it which is entirely disjoint from τ (if it is realizable). The purpose of constructing τ is that any leaf of Φ_h which is *not* entirely contained in \mathcal{P}_R can be adjusted solely at its intersections with \mathcal{P}_R to produce a leaf carried on τ and hence near its geodesic representative. The portions outside of \mathcal{P}_R are already carried on τ and hence already near the geodesic. Thus the arguments given in §7 of [M1] produce again our main structure theorems:

Theorem 4.2 (Map foliation near lamination). *Let $f: M \rightarrow N$ be an incompressible harmonic map from a closed Riemann surface to a hyperbolic 3-manifold. There are choices of constants $s_1, s_2 > 0$ for the construction of \mathcal{P}_R and an $R_0 > 0$ (all depending only on $\chi(M)$), such that in the complement of \mathcal{P}_{R_0} there is a map π from the leaves of Φ_h to the leaves of $f(\Phi_h)^*$ that factors through f , and is a local diffeomorphism on each leaf of Φ_h , mapping it to the corresponding geodesic leaf of $f(\Phi_h)^*$. For any point p on a leaf in $(M - \mathcal{P}_{R_0})$,*

$$d_N(f(p), \pi(p)) < \varepsilon(d_{|\Phi|}(p, \mathcal{P}_{R_0})),$$

and the derivative of π (along leaves, with respect to the $|\Phi|$ -metric) satisfies

$$||d\pi| - 2| \leq \varepsilon(d_{|\Phi|}(p, \mathcal{P}_{R_0}))$$

(the factor of 2 comes from the derivative of f , which approximately 2 along the horizontal leaves). The constants determining the inversely exponential function ε depend only on $\chi(M)$.

Note that this implies in particular that for a map of high enough energy, the foliation Φ_h must be realizable in $[f]$ (except possibly for components that are entirely contained within \mathcal{P}_{R_0}). Recall also that we may use the results on doubly incompressible maps by lifting the map to a cover of N —the statement of the theorem is preserved by the projection back down to N .

The length-energy inequality carries through as well:

Theorem 4.3 (Energy is length-ratio). *There is a constant C depending only on $\chi(M)$ such that in the situation described by the previous theorem,*

$$\frac{1}{2} \frac{l_N^2(f(\Phi_h)^*)}{E_M(\Phi_h)} \leq \mathcal{E}(f) \leq \frac{1}{2} \frac{l_N^2(f(\Phi_h)^*)}{E_M(\Phi_h)} + C.$$

5. LIMITS OF TEICHMÜLLER RAYS

Fix a hyperbolic 3-manifold N , a closed Riemann surface M with underlying topological surface S , and a homotopy class of maps $[f: S \rightarrow N]$, injective

on π_1 . Fix a holomorphic quadratic differential Ψ on M and let M_K ($K \geq 1$) denote the Riemann surface whose conformal structure is given by the Beltrami differential on M , $\mu = -k\bar{\Psi}/|\Psi|$ (where $k = (K - 1)/(K + 1)$). Thus the conformal class of M_K is that of the metric obtained from $|\Psi|$ by contracting the leaves of Ψ_h by a factor of K . These same leaves then form the horizontal foliation for a differential Ψ^K , holomorphic in M_K , such that Ψ_h^K is equivalent to Ψ_h in $\mathcal{MF}(S)$.

Let $f_K: M_K \rightarrow N$ be the unique harmonic map in the homotopy class $[f]$, and let $\Phi^K \in QD(M_K)$ be its Hopf differential. In this section we will compare the asymptotic behavior of Φ_h^K to Ψ_h as $K \rightarrow \infty$, and obtain corresponding conclusions about the limiting behavior of the images of f_K .

Unlike the case treated in [M1] where N is two-dimensional, we must here distinguish between the case where Ψ_h is realizable in N , and where it is not. In the former case much of the two-dimensional analysis goes through, and we will supply here the necessary details to extend it. In the latter case the images of f_K “escape” out an end of N .

At this stage we add the technical requirement that $\pi_1(N)$ contain no parabolic elements. In the presence of parabolics the dichotomy between realizability and unrealizability is less strict—parts of the surface can remain in a compact set while others escape. This more complex case will hopefully be treated in a later paper. The assumption of no parabolics allows us to assume that any unrealizable lamination is complete [Th2].

Theorem 5.1 (Teichmüller ray limit). *If Ψ_h is realizable in f and N has no parabolics, and if λ is a limit point in $\mathcal{PML}(S)$ of the Hopf laminations $\mathcal{L}(\Phi_h^K)$ for the Teichmüller ray $\{M_{K,\psi}\}$, then $\mathcal{L}(\Psi_h)$ and λ have the same supports.*

(The support of a measured lamination is the underlying set of leaves, without regard to measure.)

Proof of Theorem 5.1. Let P be a fixed pleated surface in the homotopy class of $[f]$. For any $\alpha \in \mathcal{MF}(S)$, we have that $l_P(\alpha) \geq l_N(\alpha)$. Since Ψ_h is realizable in $[f]$ we may choose P so that it maps $\mathcal{L}(\Psi_h)$ geodesically, so in particular $l_P(\Psi_h) = l_N(\Psi_h)$.

In [M1] a weighted intersection number $I_1(\alpha, \beta) = i(\alpha, \beta)/l_N(\alpha)l_N(\beta)$ is used, which is defined on $\mathcal{PML}(S)$. Here we replace it with

$$I(\alpha, \beta) = \frac{i(\alpha, \beta)}{l_P(\alpha)l_P(\beta)}.$$

The extremal length of Ψ_h in M_K satisfies

$$E_{M_K}(\Psi_h) = \frac{1}{K}E_M(\Psi_h),$$

so by the Proposition 3.1 (energy lower bound) there is a constant c_1 such that

$$\mathcal{E}(f_K) \geq c_1K.$$

With these observations, the proof of Lemma 8.2 (Intersections vanish) of [M1] follows unchanged, and we obtain

$$I(\Psi_h, \Phi_h^K) \leq \frac{c_2}{K},$$

where $c_2 = c_2(M, P)$. Thus if $[\lambda]$ is a limit point of $[\Phi_h^K]$ in $\mathcal{PML}(S)$ we obtain

$$i(\lambda, \Psi_h) = 0.$$

If Ψ_h is complete, we are done here, since in this case $i(\lambda, \Psi_h) = 0$ implies that λ must have the same support as $\mathcal{L}(\Psi_h)$. We note that so far there has been no use of the “no parabolics” assumption. In general, $\mathcal{L}(\Psi_h)$ and λ can break up into unions of minimal components, and it remains to show that the components of each must be contained in the other.

We shall outline the proof of [M1], pointing out the necessary changes. First, we recall the fact (§2 of [M1]) that every minimal lamination ν is contained in a “supporting subsurface” S_ν , unique up to isotopy, such that (1) $\pi_1(S_\nu)$ injects in $\pi_1(S)$, and (2) every homotopically nontrivial curve $\gamma \subset S_\nu - \nu$ is deformable through $S_\nu - \nu$ into ∂S_ν .

Let ν be a minimal component of λ . The proof that the support of ν is contained in the support of $\mathcal{L}(\Psi_h)$ is virtually unchanged. Suppose first that ν is not a simple closed curve. In this case there must be a simple closed curve γ in S_ν that intersects ν , but not any of the other minimal components of either λ or $\mathcal{L}(\Psi_h)$. Thus it suffices to show that γ intersects $\mathcal{L}(\Psi_h)$, because this implies that $\mathcal{L}(\Psi_h)$ has a component in S_ν , which must then be equal to ν (up to measure).

We show this by showing that, on the one hand, if $i(\gamma, \lambda) > 0$ then $E_{M_K}(\gamma) \rightarrow \infty$, while on the other hand $E_{M_K}(\gamma)$ is bounded if $i(\gamma, \Psi_h) = 0$. The arguments of [M1], being confined almost entirely to the domain, go through unchanged, with the added observation, again, that $l_P \geq l_N$.

Next, consider the case that ν is a simple closed curve. It follows as in [M1] that ν is homotopic to the core of an annulus of modulus no less than cK , for some fixed c . Then, using the quadratic differential geometry of high-modulus annuli developed in §4 of [M1], we conclude that this annulus contains a flat annulus in the $|\Psi|$ -metric. It then follows easily that the geodesic cores of this annulus are leaves of Ψ_h , homotopic to ν .

Finally, we let ν denote a minimal component of $\mathcal{L}(\Psi_h)$, and show that it is also contained in λ . Again we consider first the case where ν is not a simple closed curve. We construct for each K a 2-complex $T_\nu^K \subset M_K$ isotopic to a deformation-retract of S_ν and bounded by $|\Phi^K|$ -geodesics. In [M1] these are shown to have $|\Phi^K|$ -area $\|\Phi^K\|_{T_\nu^K} \geq cK$ for a fixed c . There is also a sequence $R_K > R_0$ such that \mathcal{P}_{R_K} separates T_ν^K from the rest of the surface and $\lim_{K \rightarrow \infty} R_K^2/K = 0$, so that $\|\Phi^K\|_{T_\nu^K - \mathcal{P}_{R_K}}$ is proportional to K as well.

Recall the argument of §4, in which for a choice of $\epsilon_2 > 0$ the flat cylinder components of $M - \mathcal{P}_R$ are divided into type (a), which map to ϵ_2 -thin parts of N , and type (b), which spiral around closed geodesics of length bounded below by (approximately) ϵ_2 . Suppose that as $K \rightarrow \infty$ there is a choice of $\epsilon_2(K) \rightarrow 0$ so that there is always a type (a) cylinder in $T_\nu^K - \mathcal{P}_{R_K}$. Then there is a sequence of simple closed curves $\gamma_K \subset S_\nu$ such that $l_N(f_K(\gamma_K)) \rightarrow 0$. Let γ_∞ be a limit point of the γ_K in $\mathcal{PML}(S)$. Then γ_∞ is unrealizable and supported in S_ν , a contradiction to the assumption about parabolics.

We conclude that there is a choice of ϵ_2 such that $T_\nu^K - \mathcal{P}_{R_K}$ is comprised of type (b) cylinders and polygonally bounded components. These have the natural train-track structure discussed in §4 (and in greater detail in [M1]), which we

here amend slightly: each type (b) cylinder produces a long branch that spirals around a closed geodesic of length bounded by cR_K . We can replace all but a length bounded by cR_K of this branch by the closed geodesic. The result is a $(1, \varepsilon(R_K))$ -nearly straight train-track all of whose branches have length bounded by a multiple of R_K .

Further, we note that Φ_h^K deposits a measure on each rectangle branch equal to the $|\Phi^K|$ -height of the corresponding rectangle (and to the sum of the measures built up by the spiraling branch on a new closed branch). This measure fails to satisfy the switch conditions by an amount equal to the height of the vertical boundary segment of \mathcal{P}_{R_K} at each switch, which is bounded by cR_K . By choosing a large enough constant s_1 in the construction of \mathcal{P}_R we can make sure this error is small compared to the measure on any branch. It is then a short exercise in linear algebra to show that we can perturb the measure on the train-track by no more than cR_K to a positive measure that satisfies the switch conditions. This determines a lamination λ_K carried on the train-track of $T_\nu^K - \mathcal{P}_{R_K}$. λ_K must remain in a compact set of N —otherwise there would be an unrealizable lamination supported in S_ν . Thus, since the sum of the measures on all the branches is proportional to K/R_K and $R_K^2/K \rightarrow 0$, λ_K approximates in $\mathcal{PML}(S)$ the part of $\mathcal{L}(\Phi_h^K)$ carried in S_ν and $l_N(\lambda_K) \geq cK$. We conclude that λ has a component in S_ν .

The case where ν is a simple closed curve proceeds similarly. As in [M1] we show that there is a flat cylinder in the $|\Phi^K|$ -metric whose core is homotopic to ν , and whose $|\Phi^K|$ -area is proportional to K . The same train-track arguments give rise in the limit to a component of λ equal to ν . \square

One consequence of this theorem is of course a statement about the promised appearance of pleated surfaces as limits of harmonic maps. Theorem 4.2 (Map foliation near lamination) already provides most of the geometric information needed for this assertion, but without a clear notion of how the pleating lamination can be determined ahead of time. With the above, we can state

Theorem 5.2 (Pleated surface limit). *The geometric limit of the sequence of images of harmonic maps associated to a Teichmüller ray along a complete realizable foliation is the image of the pleated surface pleated along the corresponding lamination.*

As stated, this theorem is an immediate consequence of Theorems 4.2 and 5.1. The latter tells us what the limiting lamination in N will be, and the former shows that “most” of the foliation will be mapped near that lamination. It remains only to control the complementary regions, i.e., the images of the subsurfaces \mathcal{P}_R . However, since the lamination $\mathcal{L}(\Psi_h)$ is assumed complete, its complementary regions are ideal triangles, and since $\partial\mathcal{P}_R$ is mapped near their boundaries, by standard properties of harmonic maps \mathcal{P}_R itself must be mapped into small neighborhoods of ideal triangles.

We note that this version of the theorem gives no information about the convergence of derivatives, although one expects to have such control (for example on the complementary regions, by standard elliptic estimates).

If $f(\Psi_h)$ is unrealizable, we certainly do not expect the images $f_K(M_K)$ to limit to a pleated surface. We recall briefly the theory of ends of hyperbolic 3-manifolds, stated for the case without parabolics. A 3-manifold N with

finitely-generated fundamental group contains a compact submanifold $C(N)$ such that the inclusion map is an isomorphism on π_1 [Sco]. Each end of N then corresponds to a boundary component of $C(N)$. When N is hyperbolic, an incompressible boundary component of $C(N)$ gives a *tame* end (see [Th2, Bo]; see also [C] for progress on the compressible boundary case). Such an end has a neighborhood that is homeomorphic to a thickened surface. An end e is *geometrically finite* if the convex hull of N excludes a neighborhood of e . A geometrically infinite end corresponding to an incompressible component $S \subset \partial C(N)$ admits a sequence of simple closed curves in S whose geodesic representatives are eventually contained in any neighborhood of the end. The limit of these curves in $\mathcal{PML}(S)$ is a unique lamination λ_e known as the *ending lamination*. This is the only way for a lamination in N to be unrealizable. We note again that λ_e must (in the absence of parabolics) be a complete lamination.

Hence given an incompressible homotopy class of maps $[f: S \rightarrow N]$, the next theorem applies exactly if f covers a boundary component of $\partial C(N)$ and Ψ_h covers an ending lamination of a corresponding end e (note that there might be two ends, if $\pi_1 N$ is a surface group).

Theorem 5.3 (Going out the end). *If Ψ_h is unrealizable in $[f]$ and N has no parabolics, then the images $f_K(M_K)$ converge to the corresponding end e .*

(“Converging to an end” means that the images of f_K are eventually contained in any neighborhood of the end.)

Proof of Theorem 5.3. We will need the following easy lemma as a tool for bounding the images $f_K(M_K)$.

Lemma 5.4 (Always a bounded curve). *If $f: M^2 \rightarrow N^3$ is any harmonic map, then for each $x \in M$ there is a homotopically nontrivial simple closed curve γ through x such that $l_N(f(\gamma)) \leq l_0$, where l_0 is a constant depending only on $\chi(M)$.*

Proof. Recalling the notation of Theorem 3.2 (Area bound), let ρ_ϵ be the pullback metric on M via g_ϵ , the approximating diffeomorphism $g_\epsilon: M \rightarrow N \times M_\epsilon$. For small enough $\epsilon > 0$, ρ_ϵ is negatively curved and has area bounded by, say, $2\pi|\chi(M)| + 1$. And of course, as $\epsilon \rightarrow 0$, ρ_ϵ approaches the pullback metric via f_K . Since ρ_ϵ is negatively curved, an embedded disk of radius r and x has area at least πr^2 . Thus there is a constant r_0 such that the r_0 -disk around x is not embedded, which results in a simple closed curve γ through x of ρ_ϵ -length at most $2r_0$. γ cannot bound a disk, by a simple application of the Gauss-Bonnet theorem, and the lemma is proved. \square

By lifting to $\tilde{N}_{f_*(\pi_1(M))}$ we can assume that $N \cong M \times \mathbf{R}$. Note that some neighborhood of the end e lifts homeomorphically to this cover.

In order to show that $f_K(M_K)$ eventually leaves every compact set in N , we will begin by showing that otherwise the energy of f_K must grow without bound.

Suppose that there are always points x_K and y_K in M_K so that $f_K(x_K) \in \mathcal{E}$ for a fixed compact set \mathcal{E} , and $f_K(y_K)$ leaves every compact set. Lemma 5.4 (always a bounded curve) then provides nontrivial curves β_K and γ_K through x_K and y_K , respectively, whose lengths in the image are bounded by l_0 . It

follows that $f_K(\beta_K) \subset \mathcal{E}'$ for a compact set \mathcal{E}' , and thus β_K must be one of a finite number of possible homotopy classes. Let us assume (by taking a subsequence if necessary) that it is a fixed homotopy class β . γ_K on the other hand must remain disjoint from β_K , and so is supported on a fixed proper subsurface $T \subset S$. But a limit of the γ_K is unrealizable and so must be complete, a contradiction.

Therefore, if $f_K(M_K)$ does not eventually leave every compact set it must remain in a fixed compact set \mathcal{E} for all K . In this case the minimal length of Ψ_h in the image is bounded below by some constant. On the other hand, $E_K(\Psi_h) \rightarrow 0$, so by Proposition 3.1 (energy lower bound), $\mathcal{E}(f_K) \rightarrow \infty$.

As in the proof of Theorem 5.1 (Teichmüller ray limit) we can construct a sequence of approximating laminations λ_K whose (projective) limit is the same as that of Φ_h^K , and which are realized in the compact set \mathcal{E} . Since $I(\Phi_h^K, \Psi_h) \rightarrow 0$ just as before, and Ψ_h is complete, we conclude that the underlying laminations of λ_K approach that of Ψ_h , which is therefore realizable, again a contradiction.

Thus $f_K(M_K)$ leaves every compact set in N . The possibility remains that it goes out the wrong end, but this is ruled out by the following application of an argument of Bonahon. Fix a sequence $K_i \rightarrow \infty$, and let $\{\alpha_i\}$ denote a sequence of simple closed curves on S such that $\alpha_i/\sqrt{E_1(\alpha_i)} \rightarrow \Psi_h/\sqrt{E_1(\Psi_h)}$, sufficiently fast that $E_{K_i}(\alpha_i) \leq cE_1(\alpha_i)/K_i$ for a uniform c . Suppose by way of contradiction that $f_{K_i}(M_{K_i})$ exits the opposite end e' while the geodesics $f(\alpha_i)^*$ exit e . Fix a closed curve $\gamma \subset S$, observing that $i(\gamma, \Psi_h) > 0$. By considering the intersections of $f_1(\gamma)$ with a homotopy A_i between $f_{K_i}(\alpha_i)$ and $f(\alpha_i)^*$, we may conclude as in §6.4 of [Bo] that, for sufficiently large i ,

$$l_N(f_{K_i}(\alpha_i)) \geq c(\gamma)i(\gamma, \alpha_i) \geq c'(\gamma)\sqrt{E_1(\alpha_i)}.$$

Therefore, by applying Lemma 3.1 (Energy lower bound) to α_i , we conclude that $\mathcal{E}(f_K) \rightarrow \infty$. Now, by Theorem 4.2, we see that $f(\Phi_h^K)^*$ exits e' , and therefore the support of $\mathcal{L}(\Phi_h^K)$ converges to that of $\lambda_{e'}$, the ending lamination of the other end. But $i(\lambda_{e'}, \Psi_h) > 0$, so one can easily see that $E_K(\Phi_h^K)/E_1(\Phi_h^K) \rightarrow \infty$. On the other hand $l_N^2(\Phi_h^K)/E_1(\Phi_h^K)$ is bounded (by $2\mathcal{E}(f_1)$), so by Theorem 4.3 we conclude $\mathcal{E}(f_K) \rightarrow 0$, a contradiction.

REFERENCES

- [BM] R. Brooks and J. P. Matelski, *Collars in Kleinian groups*, Duke Math. J. **49** (1982), 163–182.
- [Bo] F. Bonahon, *Bouts des variétés hyperboliques de dimension 3*, Ann. of Math. (2) **124** (1986), 71–158.
- [C] R. D. Canary, *Ends of hyperbolic 3-manifolds*, J. Amer. Math. Soc. (to appear).
- [CEG] R. D. Canary, D. B. A. Epstein, and P. Green, *Notes on notes of Thurston*, Analytical and Geometric Aspects of Hyperbolic Space, London Math. Soc. Lecture Note Ser., no. 111, Cambridge Univ. Press, 1987, pp. 3–92.
- [EL] J. Eells, Jr. and L. Lemaire, *Selected topics in harmonic maps*, CBMS Regional Conf. Ser. in Math., no. 50, Amer. Math. Soc., Providence, R.I., 1980.
- [ES] J. Eells, Jr. and J. H. Sampson, *Harmonic mappings of Riemannian manifolds*, Amer. J. Math. **86** (1964), 109–160.
- [FLP] A. Fathi, F. Laudenbach, and V. Poenaru, *Travaux de Thurston sur les surfaces*, Astérisque **66–67** (1979).

- [GT] D. Gilbarg and N. S. Trudinger, *Elliptic partial differential equations of second order*, Springer-Verlag, 1983.
- [Ha] A. E. Hatcher, *Measured lamination spaces for surfaces, from the topological viewpoint*, *Topology Appl.* **30** (1988), 63–88.
- [HP] J. Harer and R. Penner, Princeton Univ. Press, 1991.
- [JK1] J. Jost and H. Karcher, *Almost linear functions and a priori estimates for harmonic maps*, *Global Riemannian Geometry (Proc. Durham Conf.)*, 1984, pp. 148–155.
- [JK2] —, *Geometrische Methoden zur Gewinnung von a-priori-Schranken für harmonische Abbildungen*, *Manuscripta Math.* **40** (1982), 27–77.
- [Jo] J. Jost, *Harmonic maps between surfaces*, *Lecture Notes in Math.*, vol. 1062, Springer-Verlag, 1984.
- [K] S. Kerckhoff, *The asymptotic geometry of Teichmüller space*, *Topology* **19** (1980), 23–41.
- [M0] Y. Minsky, *Harmonic maps and hyperbolic geometry*, Ph.D. thesis, Princeton University, 1989.
- [M1] —, *Harmonic maps, length and energy in Teichmüller space*, *J. Differential Geom.* **35** (1992), 151–217.
- [Sa] J. H. Sampson, *Some properties and applications of harmonic mappings*, *Ann. Sci. École Norm. Sup. (4)* **11** (1978), 211–228.
- [Sch] R. Schoen, *Analytic aspects of the harmonic map problem*, preprint.
- [Sco] P. Scott, *Compact submanifolds of 3-manifolds*, *J. London Math. Soc.* **7** (1973), 246–250.
- [St] K. Strebel, *Quadratic differentials*, Springer-Verlag, 1984.
- [Th1] W. Thurston, *Hyperbolic structures on 3-manifolds, I: Deformations of acylindrical manifolds*, *Ann. of Math. (2)* **124** (1986), 203–246.
- [Th2] —, *The geometry and topology of 3-manifolds*, Princeton University lecture notes, 1982.
- [Th3] —, *Minimal stretch maps between hyperbolic surfaces*, unpublished manuscript, Princeton University.
- [Uh] K. Uhlenbeck, *Harmonic maps; a direct method in the calculus of variations*, *Bull. Amer. Math. Soc.* **76** (1970), 1082–1087.
- [Wf1] M. Wolf, *The Teichmüller theory of harmonic maps*, *J. Differential Geom.* **29** (1989), 449–479.
- [Wf2] —, *High energy degeneration of harmonic maps between surfaces and rays in Teichmüller space*, *Topology* **30** (1991), 517–540.
- [Wf3] —, *Infinite energy harmonic maps and degeneration of hyperbolic surfaces in moduli space*, *J. Differential Geom.* **33** (1991), 487–539.

DEPARTMENT OF MATHEMATICS, YALE UNIVERSITY, NEW HAVEN, CONNECTICUT 06520
 Current address: Department of Mathematics, University of Michigan, Ann Arbor, Michigan 48109
 E-mail address: yair@math.lsa.umich.edu.