Harmonic maps of Riemann surfaces and applications in geometry

Thesis by

Nathaniel Sagman

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2022
Defended May 25, 2022

Nathaniel Sagman
ORCID: 0000-0002-8485-7073

ACKNOWLEDGMENTS

First I'd like to thank my advisor, Vladimir Marković. Thank you for sharing so much of your time and knowledge when I was learning the ropes, for your continued patience and guidance, and then for collaboration on top of that. Thank you for arranging my extended stay at Oxford. You were engaged in my progress, but also encouraged me to pursue my own interests and projects. Overall, I just had a lot of fun doing math and spending time together.

I'm grateful to my future postdoc mentor, Jean-Marc Schlenker, as well as to Andrea Seppi, Nicholas Tholozan, and Mike Wolf for inviting me to visit their respective institutions. I learned so much and hope to keep our conversations going.

I should certainly thank the administrators, Mika, Nam, Freddy, Meagan, and Michelle, for all their help over the years. I've leaned on you so much, and I really appreciate the baked goods. Thanks as well to the members of my thesis committee: Nets Katz, Yi Ni, Sergei Gukov, and Lei Chen.

I had the pleasure of interacting with a number of postdocs in geometry at Caltech and Oxford: Qiongling Li, Tengren Zhang, Oleg Ivrii, Lei Chen, Peter Smillie, and Anna Parlak, as well as with mathematical siblings: Dmitri Gekhtman and Ognjen Tošić. I'm especially thankful to Peter for the recent collaboration.

At both the Caltech and Oxford math departments, I made so many great friends goofing around with other grad students. A special acknowledgement goes to my entering class: Tamir Hemo, Yuhui Jin, Sunghyuk Park, and Forte Shinko.

Going back to undergrad, I benefited from the guidance of several professors in Montreal: Galia Dafni, Sam Drury, Jacques Hurtubise, and Dmitry Jakobson. I'm happy to have kept up some of the friendships that I made while doing math in Burnside. Actually, as I write this, Sid is visiting. I made him go on a hike.

I'm still not fond of driving, but I found a really lovely community of friends here to explore with. Thanks for sitting at the turtle pond, and for all the evenings out and about. It was a real treat to play in the Wind Orchestra, and even more so to hang out with the woodwinds after rehearsals. To everyone in SoC and the other student advocacy groups: thank you for fostering a space for organizing and learning together.

I'm more than lucky to have lived with such kind and thoughtful people, who always brightened my day and have made such an effort to keep a home where everyone feels comfortable and understood. A good assortment of people moved in and out of the house; I'll just name the current crop: Aeryn, Arian, Bee, Charles, Jae, and Jane. Vinicius still comes around too. And of course I acknowledge the cats: Sandwich, Tarquin, John Brown, and Artemis.

To everyone I organized with through the Pasadena Tenants Union, the Eviction Defense Network, and the wider community: I'm beyond inspired by your compassion and determination, and I can't possibly express what this has meant to me. This work has really worn me out, but you've made me a more hopeful person. In all likelihood, had I not got involved, this thesis would have been completed much more easily.

To my friends back home in Toronto, Montreal, or scattered across the world: please forgive me, but there are really too many to name. I do have to mention Alec and Shardul: thanks for the adventures, and for being there to talk, all the time. To everyone: thanks for being so supportive, and meeting up when we could. Seeing you all every so often carried me through.

Finally I'd like to thank my brothers, Gabe and Reuben, and my parents, Uri and Sandra, for the unconditional love and unwavering support. Your interest in my studies means a lot—I remember I once tried to tell you about modular forms and Monstrous Moonshine, and we all got so excited. You've been lifelong role models, working in your respective fields, and reading widely on the side. Thanks especially to my mother, for casting her editorial eye over my papers even if the math was all Greek to her.

ABSTRACT

Harmonic maps are fundamental objects in differential geometry. They play an important role in studying deformations of geometric structures and in various rigidity problems. In this thesis, we present three projects, all of which involve harmonic mappings of Riemann surfaces.

In the first project, we study infinite energy harmonic maps and spacelike maximal surfaces in pseudo-Riemannian manifolds, and give applications to domination for surface group representations and anti-de Sitter geometry. The culminating result is the existence of a new class of anti-de Sitter 3-manifolds and a parametrization of their deformation space.

The second project concerns moduli spaces of harmonic surfaces inside higher dimensional Riemannian manifolds. First, we prove a factorization theorem for harmonic maps. We then use infinite-dimensional transversality theory to prove results about the distribution of certain families of harmonic surfaces inside our moduli spaces.

The final project is motivated by the Labourie conjecture from Higher Teichmüller theory. We find unstable minimal surfaces in products of hyperbolic surfaces and products of $\mathbb{R}$-trees, and we make a connection to classical minimal surfaces.

PUBLISHED CONTENT AND CONTRIBUTIONS

[MSS22] Vladimir Marković, Nathaniel Sagman, and Peter Smillie. *Unstable minimal surfaces in $\mathbb{R}^n$ and in products of hyperbolic surfaces*. 2022. arXiv: 2206.02938 [math.DG].
All authors contributed equally. Author order is alphabetical.

[Sag21a] Nathaniel Sagman. "A factorization theorem for harmonic maps". In: *J. Geom. Anal.* 31.12 (2021), pp. 11714–11740. ISSN: 1050-6926. DOI: 10.1007/s12220-021-00699-w. URL: https://link.springer.com/article/10.1007/s12220-021-00699-w.

[Sag21b] Nathaniel Sagman. *Almost strict domination and anti-de Sitter 3-manifolds*. 2021. DOI: 10.48550/ARXIV.2105.12886. URL: https://arxiv.org/abs/2105.12886.

[Sag21c] Nathaniel Sagman. *Moduli spaces of harmonic surfaces in non-positive curvature*. 2021. DOI: 10.48550/ARXIV.2111.04142. URL: https://arxiv.org/abs/2111.04142.

[Sag19] Nathaniel Sagman. "Infinite energy equivariant harmonic maps, domination, and anti-de Sitter 3-manifolds". In: *To appear in J. Differential. Geom.* (2019). DOI: 10.48550/ARXIV.1911.06937. URL: https://arxiv.org/abs/1911.06937.

TABLE OF CONTENTS

Chapter 1

INTRODUCTION

The theory of harmonic maps is a large subject in differential geometry, with a rich history and many applications in mathematics and physics, pure and applied. A number of developments in differential geometry and geometric analysis have come about through analytic problems in harmonic maps. On the other side, questions about geometry have motivated a deeper study of harmonic maps. The subject is especially fruitful on Riemann surfaces, where harmonic maps give rise to complex analytic objects.

The utility of harmonic maps in geometry can be explained as follows.

1. When there is non-trivial geometry, harmonic maps usually exist (especially in non-positive curvature). For instance, see the resolution of the Schoen Conjecture [Mar17] and its relatives [BH21].
2. Harmonic maps are rigid: they are highly constrained by the geometric objects they live on, and hence they can reveal quite a lot of geometric information.

One striking occurence of this is Siu's superrigidity for Kähler manifolds of strongly negative curvature (see [Siu80]).

Theorem 1.0.1 (Siu)**.** *Let M be a compact Kähler manifold of complex dimension at least 2 and whose curvature tensor is strongly negative. Then any compact Kähler manifold of the same homotopy type as M must either be biholomorphic or conjugate biholomorphic to M.*

The proof is summarized in two steps.

1. Let N be another compact Kähler manifold of the same homotopy type, and $f : N \to M$ a homotopy equivalence. By the well-known existence result of Eells-Sampson [EL81], there exists a harmonic map h in the homotopy class.
2. Using the curvature assumption on M (but nothing about the curvature of N), Siu proves a rigidity result: that the harmonic map h is a biholomorphism or a conjugate biholomorphism.

Siu extended the result in [Siu82], and various authors have built on this line of thought to prove rigidity results and approach different problems. Most notably, there is the work of Corlette [Cor92a] and then Gromov and Schoen (who initiated a theory of harmonic maps to buildings) [GS92] related to Margulis superrigidity.

Harmonic maps are also useful when there is no rigidity and a manifold supports a family of geometry structures. Harmonic maps parametrize the Teichmüller space of marked Riemann surfaces [Wol89] and now play a large role in Higher Teichmüller theory via the non-abelian Hodge correspondence (see [Wie18] and [Li19]).

1.1 Questions of interest

In this thesis, we present a number of projects that all involve harmonic maps in an essential way. Below, we outline the basic questions that motivate these projects.

Infinite energy harmonic maps and AdS 3-manifolds

Let Σ_g be a closed surface of genus $g \geq 2$. A discrete and faithful representation $\rho : \pi_1(\Sigma_g) \to \mathrm{PSL}(2,\mathbb{R})$ is called Fuchsian. $\mathrm{PSL}(2,\mathbb{R})$ acts via linear fractional transformations on the two-dimensional hyperbolic space $\mathbb{H}$, and the quotient $\mathbb{H}/\rho(\pi_1(\Sigma_g))$ is a hyperbolic surface. In Higher Teichmüller theory, one studies discrete and faithful representations of surface groups into higher rank Lie groups of non-compact type. Often, such representations correspond to geometric structures on manifolds.

It is a theme in Higher Teichmüller theory that geometric structures can be parametrized and understood through analytic objects. This is where harmonic maps come into play. Given a semisimple Lie group G of non-compact type, certain families of (closed) surface group representations give rise to equivariant harmonic maps from $\mathbb{H}$ to a Riemannian or pseudo-Riemannian symmetric space G/K. The harmonic map highlights some information about the representation.

In the first part of the project, we study existence and uniqueness results for infinite energy equivariant harmonic maps. Keeping to the line of thought above, we use the harmonic maps to study representations of surface groups and anti-de Sitter geometry in dimension 3.

An anti-de Sitter (AdS) manifold is a smooth manifold with a Lorentzian metric of constant negative sectional curvature. Equivalently, it is a Lorentzian manifold locally modeled on the Anti-de Sitter space AdS^{n+1}. Anti-de Sitter manifolds originally arose as negatively curved models of spacetime in general relativity, and now there are more modern applications in mathematics and physics.

Our inquiry here is about properly discontinuous group actions on AdS^3. There are two main motivations for the subject:

1. Following Thurston's geometrization program (see Scott's article [Sco83]), it is natural to study Lorentzian structures on 3-manifolds.
2. This fits into the wider program of studying properly discontinuous actions on Clifford-Klein forms, which is itself motivated by results from affine geometry such as Bieberbach's theorems [Bie11] [Bie12], and questions like the Auslander conjecture [Aus64].

In dimension 3, AdS^3 identifies with the Lie group $\mathrm{PSL}(2,\mathbb{R})$ with (a multiple of) its Lorentzian Killing metric. The time and space orientation preserving component of the isometry group is $\mathrm{PSL}(2,\mathbb{R}) \times \mathrm{PSL}(2,\mathbb{R})$, acting by left and right multiplication

$$(x, h) \cdot g = xgh^{-1}.$$

Consequently, anti-de Sitter geometry has a lot of interesting features in dimension 3, and AdS^3 geometry is intertwined with two-dimensional hyperbolic geometry.

We give a short survey on the subject at the beginning of Chapter III. For now, we put forth some natural questions.

Question 1.1.1. *Which subgroups of $PSL(2,\mathbb{R}) \times PSL(2,\mathbb{R})$ act properly discontinuously on AdS^3? More generally, which subgroups act properly discontinuously on proper domains?*

Question 1.1.2. *Do these quotients admit good deformations? And can we describe the deformation space?*

Describing quotients of AdS^3 is intimately related to the domination problem between representations (see Chapter III for more information). Given Riemannian manifolds (X_i, g_i) and group representations $\rho_i : \Gamma \to \mathrm{Isom}(X_i, g_i)$, $i = 1, 2$, we say that ρ_1 dominates ρ_2 (strictly) if for some $\lambda \leq 1$ ($\lambda < 1$), there is a λ-Lipschitz map $f : (X_1, g_1) \to (X_2, g_2)$, with the property that for all $z \in X_1$ and $\gamma \in \Gamma$,

$$f(\rho_1(\gamma)z) = \rho_2(\gamma)f(z).$$

Deroin-Tholozan used harmonic maps in [DT16] to find strictly dominating pairs (see also Salein's thesis work [Sal00], in which he had previously used holomorphic maps to find domination). This raises another set of questions.

Question 1.1.3. *Given a representation ρ of a surface group into the isometry group of a Hadamard manifold, can we describe the Fuchsian representations that (strictly) dominate ρ? To what extent can domination be seen through the geometry of equivariant harmonic maps?*

Moduli spaces of harmonic surfaces

Harmonic maps have a robust existence and uniqueness theory in non-positive curvature. As for qualitative properties (immersedness, embeddedness, etc.), quite a lot is understood for harmonic maps between surfaces (see [SY97, Chapter 1]). For harmonic maps from surface to higher dimensional manifolds, a number of questions remain open.

Question 1.1.4. *Let $f : (\Sigma, \mu) \to (M, \nu)$ be an incompressible harmonic map from a closed hyperbolic surface to a convex cocompact hyperbolic 3-manifold. Is f an immersion?*

Even for minimal surfaces, this question is not resolved.

Question 1.1.5. *Let $f : (\Sigma, \mu) \to (M, \nu)$ be an incompressible minimal map from a closed hyperbolic surface to a convex cocompact hyperbolic 3-manifold. Is f an immersion?*

By standard theory, there is always an area minimizing minimal surface ([SY79] and independently [SU81]), and this one is an immersion. This follows from [Oss70], [Gul73], and Gabai's Simple Loop Theorem for surfaces [Gab85]. An unpublished argument of Hass and Thurston shows that there are quasi-Fuchsian 3-manifolds that contain arbitrarily many incompressible minimal surfaces (the argument is sketched in [KS07, Section 2.3]). In the paper [HW15], the authors construct quasi-Fuchsian 3-manifolds with many incompressible minimal surfaces, which are all immersed.

It is conjectured that harmonic maps for Hitchin representations are immersions [Li19, Conjecture 9.3] (see the next subsection for the definition). It is not hard to see that minimal maps for Hitchin representations are immersions [Li19, page 16]. We do not know if such minimal maps are embeddings. There is actually a number of conjectures about the geometry of harmonic maps for Hitchin representations [Li19, Part 3]. One can also ask more generally about equivariant harmonic maps for Anosov representations into symmetric spaces of non-compact type.

While it's difficult to get precise information on a single harmonic map, perhaps one can make a statement about generic harmonic maps. Marković initiated the study of moduli spaces of harmonic surfaces in 3-manifolds in an unpublished preprint [Mar18]. Here we take up the study of qualitative properties of harmonic maps of surfaces, and moduli spaces of harmonic surfaces in higher dimensional Riemannian manifolds. We try to answer the loosely posed question below.

Question 1.1.6. *Let $f : (\Sigma, \mu) \to (M, \nu)$ be a harmonic map from a closed hyperbolic surface to a Riemannian manifold of dimension at least 3 and of non-positive curvature. If f is "generic," what can we say about f? For example,*

1. *How does $f(\Sigma)$ intersect itself?*
2. *Is f is an immersion?*
3. *Is f an embedding?*

Uniqueness of minimal surfaces in products

Let Σ_g be a closed surface of genus $g \geq 2$, and let G be a simple split real Lie group of non-compact type, such as $\mathrm{SL}(n, \mathbb{R})$. Given a Fuchsian representation $\sigma : \pi_1(\Sigma_g) \to \mathrm{SL}(2, \mathbb{R})$, we compose with the unique irreducible embedding $\iota : \mathrm{SL}(2, \mathbb{R}) \to G$ to get a representation $\iota \circ \sigma : \pi_1(\Sigma_g) \to G$. A Hitchin represention into G is any representation that can be continuously deformed to one of the form $\iota \circ \sigma$.

The space of Hitchin representations (mod conjugation) forms a connected component inside the space of representations into G (mod conjugation) that contains the Teichmüller space of marked hyperbolic surfaces. Using Higgs bundles, Hitchin found that this component is contractible, just like Teichmüller space [Hit92].

The Hitchin component for $G = \mathbb{RP}^3$ coincides with the space of convex projective structures investigated by Goldman [Gol90] and then Choi-Goldman [CG93]. Hitchin components are some of the first examples of Higher Teichmüller spaces: spaces of special representations into particular Lie groups of non-compact type [Wie18, Section 3].

Given a Hitchin representation ρ, and a Riemann surface S of genus $g \geq 2$, there is a unique ρ-equivariant harmonic map from a universal cover $\tilde{S}$ to the symmetric space G/K ([Cor88] for existence, and [Lab08, Proposition 4.1.5] proves Hitchin representations are irreducible, which implies uniqueness). The Labourie conjecture asks about conformal harmonic maps (minimal maps).

The Labourie Conjecture. *Given a Hitchin representation into a rank n simple split real Lie group of non-compact type $\sigma : \pi_1(\Sigma_g) \to G$, there exists a unique equivariant minimal surface in the corresponding symmetric space.*

Labourie proves existence in general [Lab08, Theorem 1.0.1]. Loftin had independently proved uniqueness for $\mathrm{SL}(3,\mathbb{R})$ [Lof01], and Labourie gave a unified proof of uniqueness for all such G of rank $n = 2$ [Lab17] (apart from $\mathrm{SL}(3,\mathbb{R})$, this adds $G = \mathrm{Sp}(4,\mathbb{R})$ and G_2). See also [CTT19], where Collier-Tholozan-Toulisse prove the analogous statement for maximal representations into Hermitian Lie groups of rank 2.

The Labourie conjecture is already intriguing from the perspective of harmonic maps and minimal surfaces. Harmonic maps in symmetric spaces are controlled by Hitchin's self-duality equations, which are difficult and interesting in their own right.

The main interest in this conjecture stems from its connections to Higher Teichmüller theory. A positive resolution of the Labourie conjecture would be remarkable: when true for a Lie group G, it implies that there is a real analytic diffeomorphism from the Hitchin component for G onto an explicit complex manifold, equivariant with respect to actions of the mapping class group on both spaces. Moreover, the quotient by the mapping class group action is a holomorphic vector bundle over the moduli space of Riemann surfaces, with an explicit description of the fibers. See [Lab08, Section 2] and [Lab17, Section 1] for more details.

A cousin of the Labourie conjecture is the following.

Question 1.1.7. *Given a product of Fuchsian representations $\sigma : \pi_1(\Sigma_g) \to PSL(2,\mathbb{R})^n$, is there a unique minimal surface in the relevant homotopy class in the corresponding product of closed hyperbolic surfaces?*

Schoen first proved that the minimal surface is unique when $n = 2$ [Sch93]. However, Marković proved that, assuming the genus g is large enough, uniqueness fails when $n \geq 3$ [Mar22]. This casts doubt on the Labourie conjecture, and suggests that Lie groups of rank 3 should be the critical case. In the third part of the thesis, we study the following question.

Question 1.1.8. *How can we geometrically understand non-uniqueness of minimal surfaces in products of hyperbolic surfaces?*

In joint work with Vladimir Marković and Peter Smillie [MSS22], we believe that we've found a satisfactory answer. Looking ahead, this leads to another question.

Question 1.1.9. *Can the recent ideas from [Mar22] and [MSS22] be generalized to resolve the Labourie conjecture?*

1.2 Overview of thesis and main results

Most of the content in this thesis is contained in the papers [Sag19], [Sag21a], [Sag21b], and [Sag21c], and the preliminary version of our work in progress [MSS22]. In Chapter II, we provide some definitions, background constructions, and expository content. To deal with overlap in the papers, we removed quite a lot of preliminaries from Chapters III-VII, and put this in Chapter II.

Below, we give a brief overview of the thesis and results. Our main results are stated informally; we leave the precise theorems for the relevant chapters.

Infinite energy harmonic maps and AdS 3-manifolds

In Chapters III and IV, we study infinite energy harmonic maps and AdS 3-manifolds. Chapter III is a modified version of [Sag19], while Chapter IV is a modified version of [Sag21b]. We present a synthesized version of the two papers.

Beginning with Chapter III, the results progress as follows.

- We first prove our existence and classification result for tame infinite energy equivariant harmonic maps into CAT(-1) Hadamard manifolds (Theorem 3A).
- We then use this result to prove a domination result for surface group representations (Theorem 3B).
- The main application of the domination result is the existence of new complete AdS 3-manifolds (Theorem 3C).

Starting in Chapter IV, we introduce our notion of almost strict domination. In the paper [Sag19] we found representations with this property, and here we study these representations in earnest.

- The central part of the chapter is the proof that a pair of representations has the almost strict domination property if and only if one can find an equivariant spacelike maximal immersion in a certain pseudo-Riemannian product manifold (Theorem 4A)).

- We then parametrize the space of almost strictly dominating pairs in a relative representation space, or equivalently the space of maximal immersions (Theorem 4B).

- We apply our results to AdS^3: these almost strictly dominating pairs give rise to special proper domains in AdS^3 with properly discontinuous group actions, and they are parametrized by the maximal surfaces. In the end, we find the deformation space of these geometric structures (Theorem 4C).

In [AL18], Alessandrini-Li use results from [DT16] and [Tho17] to study AdS 3-manifolds in the framework of non-abelian Hodge theory. In the original arXiv version of the paper [Sag21b], we followed [AL18] to provide an alternative construction for the AdS 3-manifolds from [Sag21b, Theorem C], which highlights some different information. There was only one new computation that didn't follow from [AL18], so we removed the section from the paper. Instead, we've put this content in the thesis as Section 4.5.

Moduli spaces of harmonic surfaces

This topic occupies Chapters V and VI. In Chapter V, we present the paper [Sag21a], the factorization theorem for harmonic maps. The result doesn't require many preliminary definitions, so we state it in full here. Let $f : (\Sigma, \mu) \to (M, \nu)$ be a harmonic map from a Riemann surface to a Riemannian manifold, and assume that the image of f is not contained in a geodesic.

- Suppose that there is a conformal diffeomorphism $h : \Omega_1 \to \Omega_2$ between open subsets of Σ such that $f \circ h = f$ on Ω_1. If h is holomorphic, then there is a Riemann surface (Σ_0, μ_0), a holomorphic map $\pi : \Sigma \to \Sigma_0$, and a harmonic map $f_0 : (\Sigma_0, \mu_0) \to (M, \nu)$ such that $\pi(\Omega_1) = \pi(\Omega_2)$ and f factors as $f = f_0 \circ \pi$. If h is anti-holomorphic, Σ_0 is a Klein surface and π is dianalytic (Theorem 5A).

This result tells us that the self-intersections of the harmonic surface $f(\Sigma)$ are related to the geometry of the original Riemann surface. The factorization theorem is used as a lemma in Chapter VI, but is an interesting result on its own. As we'll explain in Chapter V, it's connected to an old and important piece of mathematics: the Plateau problem.

In Chapter VI we study the moduli spaces (the paper [Sag21c]). We fix a closed surface Σ of genus at least 2 and a Riemannian manifold (M, ν) of dimension n at least 3, and a homotopy class of maps $f : \Sigma \to M$ such that the subgroup

$$f_*(\pi_1(\Sigma, x_0)) < \pi_1(M, f(x_0))$$

is not abelian. For suitably chosen metrics on Σ and M (for example, if all metrics have negative curvature), there is a unique harmonic map in the homotopy class of f.

We consider an infinite-dimensional Banach manifold $\mathfrak{M}$ consisting of pairs of metrics on Σ and M with the desired properties (for example, the space of hyperbolic metrics on Σ cross negatively curved metrics on M). We then map $\mathfrak{M}$ into a Banach manifold of maps from Σ to M, by associating each pair to its harmonic map in the homotopy class of f. We call this the moduli space of harmonic maps. In Chapter VI, after giving the formal constructions, we prove the main results, which we informally state below.

- In all dimensions $n \geq 3$, somewhere injective harmonic maps form an open, dense, and connected subset (Theorem 6A).
- When the target manifold has dimension at least 4, harmonic maps with isolated singularities can be approximated by harmonic immersions (Theorem 6B).
- When the target manifold has dimension at least 4, harmonic maps with isolated singularities can be approximated by harmonic embeddings (Theorem 6C).

For 3-manifolds, the somewhere injective result is contained in the unpublished preprint [Mar18], although some pieces of the argument were incomplete. We've given a full proof, and we've extended the result to all dimensions.

In the introductory portion of Chapter VI, we make a conjecture that the hypothesis on singularities can be removed. At the end of the chapter, we explain our use of this hypothesis.

Uniqueness of minimal surfaces in products

The subject of Chapter VII is the paper [MSS22], joint with Vladimir Marković and Peter Smillie. Marković proved that for genus g sufficiently large, there exists a product of Fuchsian representations into $\mathrm{PSL}(2,\mathbb{R})^n$ that allows mutiple minimal surfaces in the corresponding product of closed Riemann surfaces [Mar22]. This is the starting point for our work.

- The results of Chapter VII provide a strengthening of Marković's result: we produce unstable minimal surfaces for all genus $g \geq 2$ (Theorem 7A). Moreover, the proof is simpler and more geometric.

In a way that can be made precise, harmonic maps into surfaces approximate harmonic maps into $\mathbb{R}$-trees (see [Wol95]). While the term $\mathbb{R}$-tree is not explicitly mentioned in the paper [Mar22], it should be clear to experts that harmonic maps to $\mathbb{R}$-trees come into play. The first result clarifies the role of $\mathbb{R}$-trees.

- We prove that finding unstable minimal surfaces in products of hyperbolic surfaces is equivalent to finding unstable equivariant minimal surfaces in products of $\mathbb{R}$-trees (Theorem 7B2). We also prove an auxiliary result about minimal maps to products of surfaces and $\mathbb{R}$-trees that do not minimize their energy functionals over Teichmüller space (Theorem 7B1).

The idea is that an unstable minimal surface in a product of $\mathbb{R}$-trees is approximated by unstable minimal surfaces in products of hyperbolic surfaces. We then construct minimal surfaces in products of $\mathbb{R}$-trees. We find two constructions.

- First, we show that any unstable equivariant minimal surface in $\mathbb{R}^n$ is equivalent to the data of an unstable surface in a product of $\mathbb{R}$-trees (Theorem 7C). The notion of instability in $\mathbb{R}$-trees will be discussed. An example of an unstable equivariant minimal surface in $\mathbb{R}^n$ is the lift of an unstable minimal surface in an n-torus. Minimal surfaces in the 3-torus are a classical subject! One of the basic examples is the Schwarz P-surface (see [DHS10, Section 3.5.9]).

- Secondly, we prove that any three quadratic differentials ϕ_1, ϕ_2, ϕ_2 on a Riemann surface S such that the product $\phi_1\phi_2\phi_3$ is the square of a cubic differential gives rise to an unstable minimal surface in a product of $\mathbb{R}$-trees (Theorem 7D).

The set of Riemann surfaces admitting such triples from Theorem 7D is non-empty in all genus $g \geq 2$ (see the end of Chapter VII). Thus, combining Theorem B with Theorem 7D gives the proof of Theorem 7A.

Finally, let us comment on future directions and the Labourie conjecture. It is conjectured in [Kat+15] that the asymptotics of the non-abelian Hodge correspondence are controlled by harmonic maps to buildings. See the survey [Li19, Section 8] for developments, as well as the work of Parreau on compactifications of spaces of representations [Par12]. As we discussed briefly, harmonic maps between surfaces converge in some sense to harmonic maps from surfaces to $\mathbb{R}$-trees [Wol95]. With [Wol95] and [Kat+15] in mind, there is a general idea that, in the high energy limit, harmonic maps for Hitchin representations converge in a suitable sense to harmonic maps to buildings.

Especially in view of Theorems 7B1 and 7B2 in Chapter VII, this conjectural picture suggests an approach toward uniqueness questions in higher rank symmetric spaces. At the time of writing, I believe that the ideas from [Mar22] and [MSS22] extend to this setting. I look forward to seeing how everything plays out.

Chapter 2

PRELIMINARIES

In this chapter, we give basic definitions and constructions that are used throughout the thesis. First we discuss harmonic maps between Riemannian manifolds, then we specialize to Riemann surfaces, and then we discuss connections to (Higher) Teichmüller theory. We linger on a few things that deserve explanation.

As harmonic maps are a huge subject in differential geometry, we can touch on only a few aspects here. For general references on harmonic maps, we suggest the books [EL83] and [SY97].

2.1 Harmonic maps

Let (Σ, μ) be a Riemannian m-manifold with a C^2 metric, and (M, ν) a Riemannian n-manifold with a C^2 metric. Let $f : (\Sigma, \mu) \to (M, \nu)$ be a C^2 map. If f^*TM denotes the pullback bundle, the derivative df defines a section of the endomorphism bundle $T^*\Sigma \otimes f^*TM$. We denote by ∇ the connection on the tensor bundle $T^*\Sigma \otimes f^*TM$ induced by the Levi-Civita connections $(\nabla^{\mu})^*$ and $\nabla^{f^*\nu} = f^*\nabla^{\nu}$ on $T^*\Sigma$ and f^*TM respectively.

Definition 2.1.1. The tension field of a C^2 map $f : (\Sigma, \mu) \to (M, \nu)$ is the section of f^*TM given by

$$\tau = \tau(f, \mu, \nu) = \text{trace}_{\mu} \nabla df. \tag{2.1}$$

The map f is harmonic if $\tau = 0$.

The energy density of a C^2 map $f : (\Sigma, \mu) \to (M, \nu)$ is the function

$$e(\mu, \nu, f) = \frac{1}{2}\text{trace}_{\mu} f^*\nu. \tag{2.2}$$

When Σ is compact, $\tau(f, \mu, \nu) = 0$ arises as the Euler-Lagrange equation for the Dirichlet energy functional

$$f \mapsto \mathcal{E}(\mu, \nu, f) = \int_{\Sigma} e(\mu, \nu, f) dV_{\mu}, \tag{2.3}$$

where dV_{μ} is the volume form of μ. In this thesis, we don't restrict ourselves to compact manifolds, so we take Definition 2.1.1. When there is no risk of confusion, we suppress the metrics μ, ν from our notation. For example, we write $e(f) = e(\mu, \nu, f)$.

In a local coordinate $x = (x_1, \ldots, x_m)$ on the source, and a coordinate on the target in which $f = (f_1, \ldots, f_n)$, the tension field in the i^{th} coordinate is given by

$$\Delta_\mu f^i + \nu^{\alpha\beta}(f(x))\Gamma^i_{jl}(f(x))\frac{\partial f_j}{\partial x_\alpha}\frac{\partial f_l}{\partial x_\beta} = 0. \tag{2.4}$$

Here, Δ_μ is the Riemannian Laplacian with respect to the metric μ, the Γ^i_{jl}'s are the Christoffel symbols for ν, and we are using the Einstein summation convention. Hence, $\tau(f, \mu, \nu) = 0$ defines a system of second order semilinear elliptic equations.

Due to the semilinear PDE, harmonic maps automatically have a number of good local properties: a unique continuation property [Sam78, Theorem 1] and a maximum principle [Sam78, Theorem 2], to name a few. We refer the reader to the paper [Sam78].

Basic examples

Some examples should be familiar.

1. A constant speed map $f : S^1 \to (M, \nu)$ is harmonic if and only if it parametrizes a geodesic. Likewise for intervals in the real line.

2. More generally, totally geodesic maps are harmonic.

3. A map to $\mathbb{R}^n$ with its standard flat metric is harmonic if and only if the component functions are harmonic functions. Indeed, if the target is flat, then the equation (2.4) reduces to the ordinary Laplace equation.

4. Holomorphic maps between Kähler manifolds are harmonic. In particular, holomorphic maps between Riemann surfaces are harmonic.

5. Minimal surfaces are harmonic. Given a symplectic manifold with compatible almost complex structure, pseudoholomorphic curves are minimal and hence harmonic.

6. Equipping spheres with the round metrics, the Hopf fibrations $S^3 \to S^2$, $S^7 \to S^4$, and $S^{15} \to S^8$ are harmonic.

The heat flow

Eells-Sampson developed the heat flow method to prove existence of harmonic maps in non-positive curvature [ES64]. Begin with a C^2 map $f : (\Sigma, \mu) \to (M, \nu)$. Given a function $f : (\Sigma, \mu) \times [0, t_0) \to (M, \nu)$, $f_t(z) = f(z, t)$, where $[0, t_0)$ has the flat metric and global coordinate t, the time derivative is the section of f^*TM given by

$$\frac{\partial f_t}{\partial t} = \nabla_{\frac{\partial}{\partial t}} df.$$

On some time interval, we seek a solution to the Cauchy problem

$$\frac{\partial f_t}{\partial t} = \tau(f_t), \ f_0 = f. \tag{2.5}$$

Short-time existence of a solution follows from the standard parabolic theory, without any curvature assumption. For closed manifolds, Eells-Sampson use a curvature hypothesis to prove long-time existence, i.e., to find a solution on the domain $(\Sigma, \mu) \times [0, \infty)$. Taking $t \to \infty$, the solution to the heat equation converges to a harmonic map. The end result is the following.

Theorem 2.1.2 (Eells-Sampson)**.** *Suppose* (Σ, μ) *and* (M, ν) *are closed and that* (M, ν) *has non-positive sectional curvature. Then there exists a* C^∞ *harmonic map in the homotopy class of* f_0.

The curvature hypothesis ensures that the energy density of f_t is uniformly controlled along the heat flow (see, for example, [ES64, page 135], [ES64, Lemma 8A], and the Moser-Harnack inequality [Mos64]). For details, see the original paper [ES64], or the book [LW08]. The situation is quite different in positive curvature: for sequences of maps with controlled total energy, the mass of the energy density is allowed to concentrate at isolated points. This leads to "bubbles" (see [SU81]).

Hamilton wrote a whole book in which he used the heat flow to prove the analogous result for maps into non-positively curved manifolds with convex boundary [Ham75]. The convexity condition on the boundary is very natural in the context of harmonic functions in the plane (see, for instance, the Radó-Kneser-Choquet theorem [Dur04, Chapter 2]).

Uniqueness results in non-positive curvature are typically attributed to Hartman [Har67, Theorem H] (see also [Sam78, Theorem 4]). The derivative of (2.1) in the f-direction is the Jacobi operator acting on C^2 sections in $\Gamma(f^*TM)$, which is positive semi-definite when the target space has non-positive curvature. More details are given in Chapter VI, where the Jacobi operator plays a large role.

Theorem 2.1.3 (Hartman, Sampson)**.** *Suppose that Σ is closed and that (M, ν) has non-positive curvature. Let $f : (\Sigma, \mu) \to (M, \nu)$ be a harmonic map. If $f(\Sigma)$ contains a point q at which all sectional curvatures of (M, ν) at q are negative, and if $f(\Sigma)$ is not contained in a geodesic, then f is the unique harmonic map in its homotopy class.*

The heat flow is used for the main existence theorem for equivariant harmonic maps from compact surfaces (more on this below).

2.2 Harmonic maps of Riemann surfaces

From now on, (Σ, μ) is a Riemann surface with compatible metric μ, in the sense that if z is a local holomorphic coordinate, then

$$\mu = \mu(z)|dz|^2.$$

Note that any metric μ on Σ gives rise to a unique marked Riemann surface structure in which μ is conformal.

Harmonic maps in dimension 2 are conformally invariant: if we replace the metric μ with a metric μ_φ that locally takes the form $e^{\varphi(z)}\mu(z)|dz|^2$ for some C^2 function φ, then the e^φ factor in the area form dA_{μ_φ} cancels with the $e^{-\varphi}$ factor in $e(\mu_\varphi, \nu, f)$, so that

$$\mathcal{E}(\mu_\varphi, \nu, f) = \mathcal{E}(\mu, \nu, f).$$

Consequently, critical points, or solutions to the harmonic map equation, are the same for both metrics.

Let $TM^{\mathbb{C}} = TM \otimes \mathbb{C}$ denote the complexification of the tangent bundle of M and $\mathbf{E} := f^*TM^{\mathbb{C}}$ the pullback bundle. Let $z = x + iy$ be a local complex parameter on an open subset of Σ and set

$$\frac{\partial}{\partial z} = \frac{1}{2}\Big(\frac{\partial}{\partial x} - i\frac{\partial}{\partial y}\Big),\ \frac{\partial}{\partial \overline{z}} = \frac{1}{2}\Big(\frac{\partial}{\partial x} + i\frac{\partial}{\partial y}\Big).$$

We define local sections of $\mathbf{E}$ by $df\Big(\frac{\partial}{\partial x}\Big) = f_x$, $df(\frac{\partial}{\partial y}) = f_y$, and

$$df(\frac{\partial}{\partial z}) = \frac{1}{2}df\Big(\frac{\partial}{\partial x} - i\frac{\partial}{\partial y}\Big) = \frac{1}{2}(f_x - if_y) = f_z.$$

One can check that

$$(f \circ h)_w = (f_z \circ h)h_w \tag{2.6}$$

for any holomorphic map such that $h(w) = z$. Therefore, the expression $f_z dz$ is a globally defined $\mathbf{E}$-valued $(1,0)$-form on Σ.

From a classical theorem of Koszul and Malgrange, the complex vector bundle $\mathbf{E}$ admits a unique holomorphic structure such that the $(0,1)$-component of the connection $\nabla = \nabla^{\mathbf{E}}$ is the standard $\overline{\partial}$-operator. In the complex coordinate, the harmonic map equation reduces to

$$\nabla_{\frac{\partial}{\partial \overline{z}}} f_z = 0.$$

That is, f_z is a local holomorphic section of $\mathbf{E}$. We highlight this characterization because it underscores the relation between harmonicity and complex geometry on Riemann surfaces.

For a general C^2 map f, the pullback metric $f^*\nu$ decomposes into tensors of type (1,1), (2,0), and (0,2) as

$$f^*\nu = e(f)\mu + \phi + \overline{\phi}. \tag{2.7}$$

The tensor ϕ is expressed in the local coordinate z as

$$\phi = \phi(z)dz^2 = f^*\nu(f_z, f_z)(z)dz^2.$$

From (2.6), we see that when f is harmonic, ϕ defines a holomorphic quadratic differential on Σ (see below for the definition).

Definition 2.2.1. The holomorphic quadratic differential $\phi = \phi(f)$ is called the Hopf differential.

It is clear from (2.7) that f is weakly conformal precisely when $\phi \equiv 0$.

Definition 2.2.2. f is a minimal map when it is harmonic and weakly conformal.

The weakly conformal condition is the same as being a branched immersion that is conformal away from the branch points. The image of a minimal map is a weakly minimal surface: it has zero mean curvature at immersed points.

Holomorphic quadratic differentials

Quadratic differentials will play a large role throughout the thesis. We review only the basics; the book [Str84] treats the subject carefully. Let $\mathcal{K}$ be the holomorphic cotangent bundle of the Riemann surface Σ.

Definition 2.2.3. A holomorphic quadratic differential ϕ on the Riemann surface S is a holomorphic section of the symmetric square $\mathcal{K}^{\otimes 2}$.

Informally, ϕ is a tensor on the surface, which in a local holomorphic coordinate z takes the form $\phi = \phi(z)dz^2$, with $\phi(z)$ holomorphic. If ϕ does not vanish identically, then the zeros of ϕ are independent of the parametrization and discrete. If $\phi(p) \neq 0$, then integrating a choice of square root of ϕ in a chart in which $z(p) = 0$ defines a holomorphic coordinate

$$w(z) = \int_0^z \phi^{1/2}(\zeta)d\zeta$$

in which

$$\phi(w) = dw^2.$$

If $\phi(p)$ vanishes to order n, there is a coordinate w such that

$$\phi(w) = w^n dz^2.$$

Such coordinates are called natural coordinates for ϕ.

The quadratic differential ϕ induces a singular flat metric $|\phi|$ with singularities at the zeros, which we call the ϕ-metric. Locally, the metric tensor is $|\phi| = |\phi(z)||dz|^2$.

A singular foliation $\mathcal{F}$ on Σ is a foliation that is allowed to have prong singularities (see [FLP12, Exposé Five]). A C^1 arc γ on Σ is transverse to the singular foliation if it misses the singular points and is transverse to the leaves it touches at interior points. A transverse measure μ on a foliation $\mathcal{F}$ is a function that assigns a positive real number to each C^1 arc transverse to $\mathcal{F}$, invariant under leaf-preserving isotopy and absolutely continuous with respect to the Lebesque measure on Σ. The latter condition means that away from singular points, there is a smooth chart to $\mathbb{R}^2$ with its standard coordinates (x, y) and in which μ is obtained by integration against the measure $|dy|$. A singular measured foliation $(\mathcal{F}, \mu)$ is the data of a singular foliation and a transverse measure.

The horizontal (vertical) foliation of a holomorphic quadratic differential ϕ on S is a singular foliation defined as follows.

1. The leaves are C^1 curves on S that don't go through the zeros of ϕ and whose tangent vectors evaluate under ϕ to positive (negative) numbers.

2. The singularities are standard prongs at the zeros; a zero of order n corresponds to an $(n + 2)$-pronged singularity.

Both foliations come equipped with transverse measures on arcs that avoid the zero set, defined by the densities $|\sqrt{\mathrm{Im}\phi}|$ (horizontal) and $|\sqrt{\mathrm{Re}\phi}|$ (vertical). In a natural coordinate $z = x + iy$ away from the zeros, the dx integrates the horizontal foliation, and dy the vertical. It is a theorem of Hubbard-Masur that every singular measured foliation on Σ arises from a unique holomorphic quadratic differential as the horizontal (vertical) foliation with the above transverse measure [HM79].

Returning to harmonic maps from Riemann surfaces, away from the zeros and in a natural coordinate $z = x + iy$, the pullback metric is diagonal and takes the form

$$\begin{aligned} f^*\nu(z) &= e(|\phi|, \nu, f)|dz|^2 + dz^2 + d\overline{z}^2 \\ &= (e(|\phi|, \nu, f) + 2)dx^2 + (e(|\phi|, \nu, f) - 2)dy^2. \end{aligned}$$

That is, the horizontal and vertical foliations integrate the directions of maximum and minimum stretch for the harmonic map.

Holomorphic and anti-holomorphic energies

For now, assume that (M, ν) is also a hyperbolic Riemann surface, and that ν is compatible in the sense that if w is a holomorphic coordinate on M, then $\nu = \nu(w)|dw|^2$. We consider the expressions $\frac{\partial f}{\partial z}$ and $\frac{\partial f}{\partial \overline{z}}$ in local coordinates.

Definition 2.2.4. The holomorphic and anti-holomorphic energies are the functions $H(f) = H(\mu, \nu, f)$ and $L(f) = L(\mu, \nu, f)$ on Σ defined by

$$H(f)(z) = \frac{\nu(f(z))}{\mu(z)}\Big|\frac{\partial f}{\partial z}(z)\Big|^2, \; L(f)(z) = \frac{\nu(f(z))}{\mu(z)}\Big|\frac{\partial f}{\partial \overline{z}}(z)\Big|^2. \tag{2.8}$$

Since the context is clear, we write $H = H(f)$ and $L = L(f)$, and likewise for other analytic quantities. The energy density decomposes as

$$e = H + L.$$

We also record that the Jacobian determinant $J(f)$ satisfies

$$J = H - L,$$

and the Hopf differential

$$|\phi|_\mu^2 := |\phi|^2\mu^{-2} = HL.$$

H and L play a key role in the theory through the Bochner formulae, which we now introduce. Firstly, let Δ_μ be the Laplacian with respect to the metric μ, expressed in coordinates as $\Delta_\mu = \frac{4}{\mu}\frac{\partial^2}{\partial z \partial \bar{z}}$. The Gauss curvatures K_μ and K_ν are defined in conformal coordinates by

$$K_\mu(z) = -\frac{1}{2}\Delta_\mu \log \mu(z), \ K_\nu(w) = -\frac{1}{2}\Delta_\nu \log \nu(w).$$

The Bochner formulae are as follows:

$$\frac{1}{2}\Delta_\mu \log H(f) = -K_\nu H + K_\nu H^{-1}|\phi|_\mu + K_\mu \tag{2.9}$$

$$\frac{1}{2}\Delta_\mu \log L(f) = -K_\nu L - K_\nu L^{-1}|\phi|_\mu + K_\mu \tag{2.10}$$

See [SY97, Chapter 1.7].

Now we lift our assumption that (M, ν) is a surface. The content below is not contained in the standard references; we first saw it in [DT16, Section 2]. The pullback metric $f^*\nu$ satisfies $\det f^*\nu \geq 0$, and is locally a genuine metric on the bundle when $\det f^*\nu > 0$, or equivalently at points where f is locally immersed. From the formula (2.7),

$$e(f) - 2|\phi|_\mu^2 \geq 0.$$

As a consequence, the system of equations

$$\begin{cases} x + y & = e \\ xy & = |\phi|_\mu^2 \end{cases}$$

has two non-negative solutions x and y satisfying $x \geq y$, which we suggestively denote by $H(f)$ and $L(f)$. If there is an open set on which f is a diffeomorphism onto its image, then the map itself defines a complex coordinate system. In the coordinates, H and L are equal to the functions described by (2.8) in these coordinates, provided we choose orientations correctly. The local computations from [SY97, Chapter 1.7] go through for these H and L: in the open set U on which f is an immersion (this is either empty or dense [Sam78, Theorem 3]),

$$\frac{1}{2}\Delta_\mu \log H(f) = -\kappa(f^*\nu)H + \kappa(f^*\nu)H^{-1}|\phi|_\mu + K_\mu$$

$$\frac{1}{2}\Delta_\mu \log L(f) = -\kappa(f^*\nu)L - \kappa(f^*\nu)L^{-1}|\phi|_\mu + K_\mu,$$

where $\kappa(f^*\nu)$ is the Gauss curvature of the pullback metric $f^*\nu$. The only difference with the original Bochner formulae is that K_ν becomes $\kappa(f^*\nu)$. $k(f^*\nu) \leq K_\nu$, with equality at a point if and only if the second fundamental form vanishes at that point [Sam78, Theorem 7], [DT16, Lemma 2.5]. The curvature term adds some difficulty in higher dimensions, but the equations are still very powerful in some contexts. The version here will be used in Chapter III.

2.3 Teichmüller theory

Teichmüller space

Let Σ be Riemann surface of genus g and n punctures, and $\chi(\Sigma) = 2 - 2g - n < 0$. A marking of another Riemann surface S is a homeomorphism $f : \Sigma \to S$. We demand that at every puncture, the surfaces Σ and S are conformal to a standard punctured disk $\mathbb{D}^*$. Two marked Riemann surfaces (S_1, f_1) and (S_2, f_2) are said to be equivalent if $f_2 \circ f_1^{-1} : S_1 \to S_2$ is a biholomorphism isotopic to the identity.

Definition 2.3.1. The Teichmüller space $\mathbf{T}_{g,n}$ (based at Σ) is the space of equivalence classes of marked Riemann surfaces $[(S, f)]$.

It is a consequence of the uniformization theorem that $\mathbf{T}_{g,n}$ also identifies with the space of marked finite volume hyperbolic surfaces. Teichmüller space is a cell of real dimenision $6g - 6 + 2n$, and it has a wealth of metrics and a natural complex structure. See textbooks such as [Gar87], [IT92], [Hub06] for more information.

There is an interpretation of $\mathbf{T}_{g,n}$ in terms of representations. Recall that $\mathrm{SL}(2,\mathbb{R})$ acts biholomorphically and isometrically on the upper half space

$$(\mathbb{H}, \sigma) = (\{z = x + iy \in \mathbb{C} : x > 0\}, \sigma(z) = y^{-2}|dz|^2)$$

by linear fractional transformations

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot z = \frac{az + b}{cz + d}.$$

A marking $f : \Sigma \to S$ lifts to a mapping of universal covers, and once we pick basepoints for the fundamental groups, the map intertwines the (biholomorphic) deck group actions. Identifying universal covers with the upper half space $\mathbb{H}$, the deck group action gives a discrete and faithful representation $\rho : \pi_1(\Sigma) \to \mathrm{SL}(2,\mathbb{R})$. ρ necessarily takes curves enclosing punctures to parabolic elements. Conversely, by $\mathrm{K}(G, 1)$ theory, such a representation gives a marking on the Riemann surface $\mathbb{H}/\rho(\pi_1(\Sigma_g))$. Conjugating the representation has the effect of changing the basepoint, and so $\mathbf{T}_{g,n}$ identifies with the space of conjugacy classes of discrete and faithful representations of the π_1 into $\mathrm{SL}(2,\mathbb{R})$.

A Beltrami form η on a Riemann surface S is a $(-1, 1)$-form: a section of $\mathcal{K}^* \otimes \overline{\mathcal{K}}$. Locally, $\eta = \eta(z)\frac{d\bar{z}}{dz}$. According to the measurable Riemann mapping theorem, any Beltrami form η with $||\eta||_{L^\infty} < 1$ determines a unique quasiconformal homeomorphism to another Riemann surface $f^\eta : S \to S'$ (from the transformation law, the

L^∞ norm is well-defined). In this way, a deformation of a marked Riemann surface is equivalent to a path of Beltrami forms of small L^∞ norm.

The tangent space of $\mathbf{T}_{g,n}$ at a marked Riemann surface (S, f) should be the space of variations of Beltrami forms that non-trivially distort the marked Riemann surface structure. There is a pairing between Beltrami forms and holomorphic quadratic differentials via integration over S,

$$\langle \phi, \eta \rangle = \int_S \phi\eta. \tag{2.11}$$

It turns out the tangent space in question identifies with the subspace of harmonic Beltrami forms: η such that there exists ϕ making (2.11) non-zero [IT92, Theorem 7.7]. The cotangent space identifies with the space of holomorphic quadratic differentials on S (this is an instance of Serre duality). See [IT92, Chapter 7] for more details.

Let $\mathrm{Diff}^+(\Sigma)$ be the group of C^∞ orientation preserving diffeomorphisms of Σ, with the C^∞ topology, and $\mathrm{Diff}_0^+(\Sigma)$ the normal subgroup of maps that are isotopic to the identity. The mapping class group of Σ is

$$\mathrm{MCG}(\Sigma) := \pi_0(\mathrm{Diff}^+(S)) = \mathrm{Diff}^+(S)/\mathrm{Diff}_0^+(S).$$

$\mathrm{MCG}(\Sigma)$ acts on Teichmüller space by changing the marking, and this action is isometric with respect to various metrics on $\mathbf{T}_{g,n}$ (see [FM12, Part 2]).

Minimal surfaces

Let Σ be closed. Suppose we are given (M, ν) and a homotopy class of maps $f : \Sigma \to (M, \nu)$ such that for every metric μ on Σ, there is a unique harmonic map in the homotopy class. From conformal invariance, the map is determined by the class of the source metric μ in Teichmüller space. Hence, there is a well-defined energy functional

$$\mathbf{E} : \mathbf{T}_g \to [0, \infty)$$

that records the total energy of the unique harmonic map. $\mathbf{E}$ is C^1 and has maximum regularity depending on that of ν (see [EL81, Theorem 3.1]).

Theorem 2.3.2. *Let η be a harmonic Beltrami form on (Σ, μ), and let ϕ be the Hopf differential of the harmonic map at μ. The derivative of $\mathbf{E}$ is given by*

$$d\mathbf{E}[\mu](\nu) = -4\langle \eta, \phi \rangle = -4\,\mathrm{Re} \int_\Sigma \phi\eta.$$

Consequently, at a critical point of **E**, the harmonic map is minimal. See [Wen07, page 2] for some remarks on the history of this result. We'll see many variations of **E** in the main parts of the thesis.

The usual example of this set-up is when (M, ν) is a closed manifold of negative curvature, and the induced map on the fundamental groups

$$f_* : \pi_1(\Sigma, x_0) \to \pi_1(M, f(x_0))$$

does not have abelian image. A seminal result is proved in [SY79, Theorem 3.1].

Theorem 2.3.3 (Schoen-Yau)**.** *Suppose that (M, ν) is a closed manifold of negative curvature and $f : (\Sigma, \mu) \to (M, \nu)$ is π_1-injective. Then the associated energy functional $\mathbf{E} : \mathbf{T}_g \to [0, \infty)$ is proper. Consequently, there exists a minimizer for* **E***, which yields a minimal map in the homotopy class.*

The harmonic maps parametrization

Now we assume that (Σ, μ) and (M, ν) are closed surfaces of the same genus $g \geq 2$. The following result is the starting point for applying harmonic maps to Teichmüller theory.

Theorem 2.3.4 (Schoen-Yau, Sampson)**.** *Suppose that M is negatively curved, and let $f : (\Sigma, \mu) \to (M, \nu)$ be the unique harmonic map in the homotopy class of the identity. Then f is a diffeomorphism.*

See [Sam78, Theorem 11] and [SY78, Theorem 3.1]. Let $QD(S)$ denote the Banach space of holomorphic quadratic differentials on S with the L^1 norm. Fix a source metric μ. We define a map

$$\beta : \mathbf{T}_g \to QD(S)$$

by associating to each metric ν on M the Hopf differential of the harmonic map in the homotopy class of the identity from $(\Sigma, \mu) \to (M, \nu)$. It was proved by Sampson that this map is injective: if ν_1 and ν_2 yield harmonic maps $f_i : (\Sigma, \mu) \to (M, \nu_i)$, $i = 1, 2$, with $\phi(f_1) = \phi(f_2)$, then $f_2 \circ f_1^{-1} : (M, \nu_1) \to (M, \nu_2)$ is an isometry in the homotopy class of the identity [Sam78, Theorem 12]. In his thesis [Wol89, Theorem 3.1], Wolf proved the following.

Theorem 2.3.5 (Sampson, Wolf)**.** *The map β is a homeomorphism.*

There are also independent proofs by Hitchin [Hit87, Section 11] and Wan [Wan92]. Thus, Teichmüller theory can be reformulated in terms of harmonic maps.

High energy harmonic maps

High energy harmonic maps, first studied in Wolf's thesis [Wol89] (see also Minsky's thesis work [Min92]) are at the heart of the matter in Chapter VII. The high energy behaviour is also lurking behind the scenes in Chapters III and IV.

Let Σ be a closed Riemann surface of genus $g \geq 2$, and $\phi \in QD(\Sigma)$. From the previous subsection, we know that for each $t > 0$, there is a metric σ_t such that the identity map $\text{id}_t : (\Sigma, \mu) \to (\Sigma, \sigma_t)$ is harmonic and has Hopf differential $t\phi$. Let $d(\cdot, \cdot)$ denote the distance function on S induced by the ϕ metric, and let $\mathcal{Z}$ be the zero set of ϕ. Let H_t and L_t be the holomorphic and anti-holomorphic energies for id_t, J_t the Jacobian, and ν_t the Beltrami form. There is a universal constant c such that, as we take $t \to \infty$, the following estimates hold.

1. Away from $\mathcal{Z}$, H_t and L_t satisfy the estimates

$$\frac{1}{t}H_t(z) = |\phi(z)| + o(e^{-ct^{1/2}d(z,\mathcal{Z})}), \ \frac{1}{t}L_t(z) = |\phi(z)| + o(e^{-ct^{1/2}d(z,\mathcal{Z})}). \quad (2.12)$$

2. Away from $\mathcal{Z}$, $|\nu_t|$ increases to 1, and $|\nu_t| - 1 = o(e^{-ct^{1/2}d(z,\mathcal{Z})})$.

3. It follows from (2.12) that away from $\mathcal{Z}$, $J_t = o(e^{-ct^{1/2}d(z,\mathcal{Z})})$. At the zeros, $J_t \to \infty$.

These estimates are all proved via the Bochner formulae (see [Wol89, Proposition 4.3], [Wol91a, Theorem 3.1], and [Min92, Lemma 3.2]). In terms of the geometry of the harmonic maps, these estimates show that

1. for large t, leaves of the horizontal foliation for ϕ are stretched by the harmonic map with scale $t^{1/2}$, and the images have geodesic curvature $o(e^{-ct^{1/2}d(z,\mathcal{Z})})$ (see [Wol91a, Lemma 2.2]).

2. For large t, images of leaves of the vertical foliation have length $o(e^{-ct^{1/2}d(z,\mathcal{Z})})$.

We formally define leaf spaces and harmonic maps to $\mathbb{R}$-trees in Chapter VII. Loosely speaking, the space of leaves of the vertical foliation of ϕ, when equipped with a distance function induced by the transverse measure, forms a metric space (T, d) called the leaf space of ϕ. It comes with a harmonic map $\pi : \tilde{\Sigma} \to (T, d)$. It is a consequence of the results above that as we take $t \to \infty$, the high energy harmonic maps converge in a suitable sense to the map π (see [Wol95, Corollary 5.2]).

Thurston gave a compactification of Teichmüller space on which the action of the mapping class group extends continuously to the boundary. The boundary objects are projective measured foliations on the surface S (see the book [FLP12]). Wolf compactified Teichmüller space using harmonic maps; the boundary points can be interpreted as harmonic maps to $\mathbb{R}$-trees (dual to foliations) [Wol89, Section 4]. Wolf shows [Wol89, Theorem 4.1] that this compactification agrees with the Thurston compactification.

2.4 The non-abelian Hodge correspondence

In this section, (Σ, μ) is a finite volume hyperbolic surface with compatible complex structure. We denote by Γ the Fuchsian holonomy group of (Σ, μ).

Representations and flat bundles

Let G be a Lie group, acting by isometries on a contractible manifold X. For any representation $\rho : \Gamma \to G$, there is an associated fiber bundle $X_\rho \to \Sigma$ whose total space is the quotient of $\tilde{\Sigma} \times X$ by the action of the deck group Γ via

$$\gamma \cdot (z, x) = (\gamma \cdot z, \rho(\gamma)x).$$

This is naturally endowed with a flat G-connection, and upon choosing a basepoint for the π_1, the holonomy representation is conjugate to ρ. This mapping $\rho \mapsto X_\rho$ produces a bijection between the set of conjugacy classes of representations and that of gauge equivalence classes of flat X-bundles with structure group G.

Global sections always exist because X is contractible. Under this correspondence, taking the pullback bundle with respect to the universal covering $\tilde{\Sigma} \to \Sigma$ shows that sections of X_ρ are equivalent to ρ-equivariant maps from $\tilde{\Sigma} \to X$, i.e., maps f that satisfy, for all $z \in \tilde{\Sigma}$ and $\gamma \in \Gamma$,

$$f(\gamma \cdot z) = \rho(\gamma) f(z).$$

We will pass back and forth between these two perspectives.

We now assume that (X, ν) is Hadamard, meaning that it is complete, simply connected, and non-positively curved.

Definition 2.4.1. A representation $\rho : \pi_1(\Sigma) \to (X, \nu)$ is reductive if there exists a convex set $C \subset X$, invariant under $\rho(\pi_1(\Sigma))$ and such that

1. C splits as a Riemannian product $C = C_1 \times E$, where C_1 is convex and E is a Euclidean space, and

2. $\rho(\pi_1(\Sigma))$ preserves the decomposition, and the restriction $\rho : \pi_1(\Sigma) \to \mathrm{Isom}(C_1, \nu)$ does not fix a point on the Gromov boundary $\partial_\infty C_1$.

This extends the usual notion of reductive groups.

Lemma 2.4.2. *Let G be a real algebraic group, and X a homogeneous space for G. Then $\rho : \pi_1(\Sigma) \to (X, \nu)$ is reductive if and only if the Zariski closure of $\rho(\pi_1(\Sigma))$ is a reductive subgroup of G.*

When (X, ν) has negative sectional curvature, there is a dichotomy for reductive representations. A reductive representation ρ is either

1. irreducible, meaning that there is no point ξ on the Gromov boundary $\partial_\infty X$ such that $\rho(\Gamma)$ fixes ξ; or
2. reducible, which means that the group $\rho(\Gamma)$ stabilizes a geodesic, on which it acts by translations.

Equivariant harmonic maps

Let (X, ν) be a Hadamard manifold, $G = \mathrm{Isom}(X, \nu)$, and $\rho : \Gamma \to G$ be a reductive representation. Since ρ is acting by isometries, the energy density $e(\tilde{\mu}, f)$ is invariant under the action of Γ on $\tilde{\Sigma}$, and hence descends to a function $e(\mu, f)$ on Σ. Similarly, the Hopf differential is invariant, and we call the downstairs quotient ϕ the Hopf differential of f.

Definition 2.4.3. A ρ-equivariant map $f : (\tilde{\Sigma}, \tilde{\mu}) \to (X, \nu)$ is harmonic if it satisfies the usual Euler-Lagrange equation.

If the surface is closed, we can define the usual total energy as in (2.3) by integrating $e(\mu, f)$ over Σ, and define equivariant harmonic maps as critical points. When Σ is closed and $X = \mathbb{H}^3$, Donaldson solved the heat equation (2.5) to prove existence of equivariant harmonic maps for irreducible representations [Don87]. Working in local coordinates, one can apply the usual parabolic theory to get a short-time equivariant solution f_t with energy density bounds (the methods of [ES64] apply directly). Donaldson observes that the irreducible condition guarantees that f_t takes compact sets into compact sets, independent of t. Then one can do Arzelà-Ascoli for long-time convergence.

The method is used in more general contexts by Corlette [Cor88], Labourie [Lab91], Jost-Yau [JY91], and Corlette again [Cor92b]. We compile some of the results into one:

Theorem 2.4.4. *(Corlette, Donaldson, Labourie, Jost-Yau) Suppose M is a complete Riemannian manifold possibly with boundary, $\Gamma \simeq \pi_1(M)$, X is a $CAT(-1)$ Hadamard manifold, and $\rho : \Gamma \to X$ is a reductive representation. If there exists a ρ-equivariant map with finite energy (with equivariant boundary values if $\partial M \neq \emptyset$), then there exists an equivariant harmonic map (with the same boundary values).*

The non-abelian Hodge correspondence

Here we give a very brief overview of the non-abelian Hodge correspondence for a closed Riemann surface (Σ, μ) and the Lie group $G = \mathrm{SL}(n, \mathbb{C})$. The theory holds for more general Lie groups of non-compact type, but for simplicity we restrict ourselves. We also leave out some definitions, since we discuss the non-abelian Hodge correspondence for parabolic Higgs bundles in Section 4.5. For an introduction to Higgs bundles and non-abelian Hodge theory, see [Wen16], [Gui18]. For more on the harmonic maps perspective, see [Li19].

The correspondence concerns two moduli spaces associated to the surface Σ. The Betti moduli space is the GIT quotient (we like the reference [Szé14, Section 5])

$$B(G) = \mathrm{Hom}(\pi_1(\Sigma), G)//G,$$

where the G-action is by conjugation. Alternatively, it is obtained by removing the representations that are not reductive, and then quotienting by conjugation. Hitchin introduced Higgs bundles in his seminal paper [Hit87].

Definition 2.4.5. A rank n Higgs bundle over Σ is a pair $(E, \overline{\partial}^E, \varphi)$, where $(E, \overline{\partial}^E)$ is a holomorphic vector bundle of rank n over Σ and $\phi \in H^0(\Sigma, \mathrm{End}(E) \otimes \mathcal{K})$ is a section called the Higgs field. An $\mathrm{SL}(n, \mathbb{C})$-Higgs bundle has the extra requirements that $\det E$ is the trivial line bundle, and $\mathrm{trace}\varphi = 0$.

On a closed Riemann surface S, the space of polystable Higgs bundles (see [Wen16, Definition 2.3]) mod Higgs bundle isomorphism (see [Wen16, page 7]) is called the Dolbeaut moduli space $D_S(G)$. The correspondence is below.

Theorem 2.4.6 (Corlette, Donaldson, Hitchin, Simpson)**.** *Let S be a closed Riemann surface of genus $g \geq 2$. There is a homeomorphism between the moduli spaces $B(G) \to D_S(G)$.*

The passage from $D_S(G)$ to $B(G)$ is contained in the work of Hitchin [Hit87] and Simpson [Sim88], and Donaldson [Don87] and Corlette [Cor88] proved the other direction of the correspondence. The proof of both directions involve an intermediate moduli space: the de Rham moduli space $R(G)$ of gauge equivalence classes of flat bundles with holonomy in G. As explained previously, $B(G)$ is in bijection with $R(G)$.

Harmonic maps provide the engine for one side of the correspondence. Given a representation $\rho : \pi_1(\Sigma) \to \mathrm{SL}(n, \mathbb{C})$, we form the bundle $E = \mathbb{C}^n_\rho \to \Sigma$ with flat connection ∇. Corlette provides a ρ-equivariant harmonic map to the symmetric space for $\mathrm{SL}(n, \mathbb{C})$, which is equivalent to a harmonic metric on the bundle E. From the harmonic metric and the flat connection, one can construct a complex structure $\overline{\partial}^E$ on E and a Higgs field, which, up to some identifications, can be seen as the $(1,0)$-component of the derivative of the harmonic map (see Section 4.5 in Chapter IV).

On the other side, Hitchin and Simpson show that from a polystable Higgs bundle $(E, \overline{\partial}^E, \varphi)$, one can find a metric h on the bundle such that if ∇_h is the Chern connection, and F_{∇_h} its curvature, then (∇_h, φ) solves Hitchin's self-duality equation

$$F_{\nabla_h} + [\varphi, \varphi^{*h}] = 0,$$

where the Lie bracket has been extended to $H^0(\Sigma, \mathrm{End}(E) \otimes \mathcal{K})$, and φ^{*h} is the Hermitian adjoint with respect to h. In other words, $\nabla = \nabla_h + \varphi + \varphi^{*h}$ is a flat connection on E. The holonomy of this connection defines the corresponding representation.

Unfortunately, the non-linear nature of the PDE (2.4) makes it difficult to extract information about a specific representation from its Higgs bundle. We remark here that if h is the ρ-equivariant harmonic map with Higgs bundle φ, then the Hopf differential $\phi(h)$ satisfies $\phi(h) = 2n\mathrm{trace}(\varphi^2)$ [Li19, Section 5]. For harmonic maps between surfaces of degree 1, i.e., rank 2 Higgs bundles, the Hopf differential gives the entire data of the Higgs field [Li19, Section 6].

Higher Teichmüller theory

There are many surveys on Higher Teichmüller theory, such as [Wie18]. Here we give a cursory introduction. Let Σ be closed and of genus $g \geq 2$. We mentioned previously that Teichmüller space can be seen as a space of representations. In fact, there are two components of $B(\mathrm{SL}(2, \mathbb{R}))$ that each identify with Teichmüller space.

In Section 1.2, we introduced the Hitchin components inside $B(\mathrm{SL}(n,\mathbb{R}))$, which in the case $n = 3$ parametrizes convex projective structures on surfaces. Using the non-abelian Hodge correspondence, Hitchin discovered that the Hitchin component is contractible: attaching a Riemann surface structure S, Hitchin analyzed this component by looking in the moduli space of Higgs bundles $D_S(\mathrm{SL}(n,\mathbb{R}))$. Years later, Labourie [Lab06, Theorem 1.5] and Fock-Goncharov [FG06] independently proved that Hitchin representations are discrete and faithful.

The definition below is taken from [Wie18].

Definition 2.4.7. A Higher Teichmüller space for a semisimple Lie group G of non-compact type is any connected component of $B(G)$ that consists entirely of discrete and faithful representations.

The Higher Teichmüller spaces all fit into the framework of positive representations [GLW21]. Harmonic maps and the non-abelian Hodge correspondence continue to be a valuable tool in Higher Teichmüller theory.

There is a key difference between Hitchin's parametrization and the parametrization that would follow from the Labourie Conjecture: the first one depends on a choice of marked Riemann surface structure S, while the second does not (see [Lab17, Section 1]). Also, the Labourie parametrization has mapping class group symmetry (see [Lab08, Theorem 1.0.2] and [Lab17, Section 1]). It is worth commenting here that even if Labourie's conjecture is false, there are still approaches for finding a complex structure with the desired properties (for example, see [FT21]).

In the next two chapters, we study AdS 3-manifolds. The representations giving actions on AdS^3 have a reasonable deformation space, but they do not form connected components of the $B(\mathrm{PSL}(2,\mathbb{R})^2)$. They are not Higher Teichmüller spaces in the sense above, but the work follows the same philosophy.

Chapter 3

INFINITE ENERGY HARMONIC MAPS AND ADS 3-MANIFOLDS

3.1 Introduction

Harmonic maps play a special role in the theory of geometric structures on manifolds. The existence results of Donaldson, Corlette, and Labourie link the purely algebraic data of a matrix representation of a discrete group to a geometric object—an equivariant harmonic map between manifolds—realising the prescribed transformations. In this chapter we generalize their work to a non-compact setting and apply it to the study of domination between representations.

Let Γ be a discrete group and for $k = 1, 2$ let $\rho_k : \Gamma \to \mathrm{Isom}(X_k, g_k)$ be representations into the isometry groups of Riemannian manifolds (X_k, g_k). A function $f : X_1 \to X_2$ is (ρ_1, ρ_2)-equivariant if for all $\gamma \in \Gamma$ and $x \in X_1$,

$$f(\rho_1(\gamma) \cdot x) = \rho_2(\gamma) \cdot f(x).$$

ρ_1 *dominates* ρ_2 if there exists a 1-Lipschitz (ρ_1, ρ_2)-equivariant map. The domination is *strict* if the Lipschitz constant can be made strictly smaller than 1. The translation length of an isometry γ of a metric space (X, d) is

$$\ell(\gamma) = \inf_{x \in X} d(x, \gamma \cdot x).$$

ρ_1 dominates ρ_2 *in length spectrum* if there is a $\lambda \in [0, 1]$ such that

$$\ell(\rho_2(\gamma)) \leq \lambda \ell(\rho_1(\gamma))$$

for all $\gamma \in \Gamma$. This domination is strict if $\lambda < 1$. From the definitions, (strict) domination implies (strict) domination in length spectrum.

Domination is essential to understanding complete manifolds locally modeled on $G = \mathrm{PO}(n, 1)_0$: a geometrically finite representation $\rho_1 : \Gamma \to G$ strictly dominates $\rho_2 : \Gamma \to G$ if and only if the (ρ_1, ρ_2)-action on G by left and right multiplication is properly discontinous [GK17]. For $n = 2$ these are the *anti-de Sitter* (AdS) 3-manifolds, and for $n = 3$ we have the 3-dimensional complex *holomorphic-Riemannian* 3-manifolds of constant non-zero curvature (see [DZ09] for details).

Anti-de Sitter space

The exposition here is minimal, and for more information, we suggest the recent survey [BS20]. Denote by $\mathbb{R}^{n,2}$ the real vector space $\mathbb{R}^{n+2}$ equipped with the non-degenerate bilinear form

$$q_{n,2}(x) = \sum_{i=1}^{n} x_i y_i - x_{n+1} y_{n+1} - x_{n+2} y_{n+1}.$$

We define

$$\mathbb{H}^{n,1} = \{x \in \mathbb{R}^{n,2} : q_{n,2}(x) = -1\}.$$

The quadric $\mathbb{H}^{n,1} \subset \mathbb{R}^{n,2}$ is a smooth connected submanifold of dimension $n+1$, and each tangent space $T_x\mathbb{H}^{n,1}$ identifies with the $q_{n,2}$-orthogonal complement of the linear span of x in $\mathbb{R}^{n,2}$. The restriction of $q_{n,2}$ to such a tangent space is a non-degenerate bilinear form of signature $(n,1)$, and this induces a Lorentzian metric on $\mathbb{H}^{n,1}$ of constant curvature -1 on non-degenerate 2-planes. $\mathbb{H}^{n,1}$ identifies with the Lorentzian symmetric space $O(n,2)/O(n,1)$, where $O(n,1)$ embeds into $O(2,2)$ as the stabilizer of the standard basis vector e_n.

The center of $O(n,2)$ is $\{\pm I\}$, where I is the identity matrix. The Klein model of AdS^{n+1} is the quotient

$$\mathrm{AdS}^{n+1} = \mathbb{H}^{n,1}/\{\pm I\},$$

with the Lorentzian metric induced from $\mathbb{H}^{n,1}$. This also identifies as the space of negative (timelike) directions in $\mathbb{R}^{n,2}$,

$$\mathrm{AdS}^{n+1} = \{[x] \in \mathbb{RP}^{n+1} : q_{n,2}(x) < 0\}.$$

A tangent vector $v \in T_x\mathbb{H}^{2,1}$ is timelike, lightlike, and spacelike if $q_{n,2}(v,v) < 0$, $q_{n,2}(v,v) = 0$, and $q_{n,2}(v,v) > 0$ respectively, and likewise for AdS^3. The causal character of a geodesic curve is constant, and correspondingly we call geodesics timelike, lightlike, or spacelike if every tangent vector is timelike, lightlike, or spacelike.

AdS^{n+1} space arose in physics: anti-de Sitter metrics are exact solutions of the Einstein field equations (in which the only term in the stress-energy tensor is a negative cosmological constant). Now there are more modern applications in physics. In this part of the thesis we study 3-dimensional anti-de Sitter space. The low dimensional AdS^3 appears in physics—for instance, see the work of Witten [Wit89], [Wit07].

On AdS 3-manifolds

In dimension 3, there is another model of AdS^3: the Lie group $\mathrm{PSL}(2,\mathbb{R})$. The determinant form $q = (-\det)$ defines a signature $(2,1)$ bilinear form on the lie algebra $\mathfrak{sl}(2,\mathbb{R}) = T_{[I]}\,\mathrm{PSL}(2,\mathbb{R})$ (it is a multiple of the Killing form). Translating to each tangent space via the group multiplication, we obtain a Lorentzian metric that is isometric to AdS^3. The space and time-orientation preserving component of the isometry group is $\mathrm{PSL}(2,\mathbb{R})\times\mathrm{PSL}(2,\mathbb{R})$, acting via the left and right multiplication:

$$(g,h)\cdot x = gxh^{-1}.$$

An AdS 3-manifold is a Lorentzian 3-manifold of constant curvature -1. Equivalently, such a manifold is locally isometrically modelled on $\mathrm{PSL}(2,\mathbb{R})$.

There are two main lines of research in AdS^3:

1. globally hyperbolic maximally compact AdS 3-manifolds, as studied by Mess [Mes07] and developed by many others [KS07] (see [BS20] and the references therein).

2. and the study of properly discontinuous group actions on AdS^3. Some of the main works are [KR85], [Kli96], [Sal00], [Kas10], [GKW15], [DGK16a], [DGK16b], [DT16], [Tho17], and [Tho18].

Our interest here is in the latter. More generally there is an interest in properly discontinuous group actions on Clifford-Klein forms, and study that has its roots in some famous conjectures about affine geometry.

We give a brief overview of some aspects. If an AdS 3-manifold is geodesically complete, meaning geodesics run for all time, then it comes from a proper quotient of $\widetilde{\mathrm{PSL}}(2,\mathbb{R})$ with respect to the lift of the action above. Goldman showed that the space of closed AdS 3-manifolds is larger than originally expected [Gol85], and Kulkarni and Raymond took up the problem of understanding all geodesically complete AdS 3-manifolds [KR85]. Among other things, they proved [KR85, Theorem 5.2] that any torsion-free discrete group acting properly discontinuously on AdS^3 is of the form

$$\Gamma_{\rho_1,\rho_2} = \{(\rho_1(\gamma),\rho_2(\gamma)) : \gamma\in\Gamma\},$$

where Γ is a the fundamental group of a surface, and $\rho_1,\rho_2 : \Gamma \to \mathrm{PSL}(2,\mathbb{R})$ are representations, with at least one of them Fuchsian. This is generalized for actions on rank 1 Lie groups in [Kas08].

Remark 3.1.1. Shortly after, Klingler [Kli96] proved that closed Lorentzian manifolds of constant curvature are geodesically complete. Thus, the completeness assumption can be dropped in the work of Kulkarni and Raymond on closed AdS 3-manifolds.

The natural next step is to understand which Γ_{ρ_1,ρ_2} act properly discontinuously. In the cocompact case, Salein observed it is sufficient [Sal00], and Kassel proved it is necessary [Kas10] that ρ_1 strictly dominates ρ_2 (defined below). Guéritaud and Kassel extended these results to surfaces with punctures and higher dimensional hyperbolic spaces [GK17].

When a group Γ acts on a manifold with no reference to a representation ρ_1, we may just write ρ_2-equivariant. By the Selberg lemma, we only need to consider torsion-free groups.

Theorem 3.1.2 (Guéritaud-Kassel, Theorem 1.8 in [GK17]). *A finitely generated discrete group Γ_{ρ_1,ρ_2} acts properly discontinuously and without torsion if and only if ρ_1 is Fuchsian and strictly dominates ρ_2, up to interchanging ρ_1 and ρ_2.*

The quotient is a Seifert-fibered AdS 3-manifold over the hyperbolic surface $\mathbb{H}/\rho_1(\Gamma)$ such that the circle fibers are timelike geodesics.

In the $\mathrm{PSL}(2,\mathbb{R})$ model, timelike geodesics are all of the form

$$L_{p,q} = \{X \in \mathrm{PSL}(2,\mathbb{R}) : X \cdot p = q\},$$

where (p,q) range over $\mathbb{H} \times \mathbb{H}$. These are topological circles and have Lorentzian length π.

There was an open question: is every non-Fuchsian representation $\pi_1(S_g) \to \mathrm{PSL}(2,\mathbb{R})$ strictly dominated by a Fuchsian one? This question was answered by Deroin-Tholozan in [DT16] and Guéritaud-Kassel-Wolf in [GKW15], using different methods. Deroin-Tholozan actually proved a more general result.

Theorem 3.1.3 (Deroin-Tholozan, Theorem A in [DT16]). *Let (X,ν) be a CAT(-1) Hadamard manifold with isometry group G and $\rho : \pi_1(S_g) \to G$ a representation, $g \geq 2$. Then ρ is strictly dominated by a Fuchsian representation, unless it stabilizes a totally geodesic copy of $\mathbb{H}$ on which the action is Fuchsian.*

Marché and Wolff use the domination result to answer a question of Bowditch and resolve the Goldman conjecture in genus 2 [MW16].

Tholozan completed the story for closed 3-manifolds in [Tho17].

Theorem 3.1.4 (Tholozan, Theorem 1 in [Tho17])**.** *Fix $g \geq 2$ and a $CAT(-1)$ Hadamard manifold (X, ν) with isometry group G. The space of dominating pairs within $\mathbf{T}_g \times Rep^{nf}(\pi_1(S_g), G)$ is homeomorphic to*

$$\mathbf{T}_g \times Rep^{nf}(\pi_1(S_g), G),$$

where $Rep^{nf}(\pi_1(S_g), G)$ is the space of representations that do not stabilize a totally geodesic copy of $\mathbb{H}$ on which the action is Fuchsian.

The homeomorphism is fiberwise in the sense that for each $\rho \in \mathrm{Rep}^{nf}(\pi_1(S_g), G)$, it restricts to a homeomorphism from $\mathbf{T}_g \times \{\rho\} \to U \times \{\rho\}$, where $U \subset \mathbf{T}_g$ is an open subset. The key point is that when $(X, \nu) = (\mathbb{H}, \sigma)$, this is the deformation space of closed AdS quotients of $\mathrm{PSL}(2, \mathbb{R})$. The components of the deformation space are thus organized according to Euler numbers.

Results

Henceforth a manifold that is "complete, finite volume" is implicitly understood to be non-compact. A Hadamard manifold (X, g) is $\mathrm{CAT}(-\kappa)$, $\kappa \geq 0$, if all sectional curvatures are $\leq -\kappa$. See [BH99] for information on $\mathrm{CAT}(-\kappa)$ metric spaces. When describing a fundamental group we suppress dependence on a basepoint. We often identify the fundamental group with the group of deck transformations without a change in notation. If Γ acts isometrically on a Riemannian manifold and ρ is a representation, an (id, ρ)-equivariant map is simply called ρ-equivariant. The function $\Lambda : \mathbb{R} \to \mathbb{C}$ given by

$$\Lambda(\theta) = (1 - \theta^2) - i2\theta$$

will frequently appear in the chapter. We record here that as θ increases from $-\infty$ to ∞, the complex argument of $\Lambda(\theta)$ decreases from π to $-\pi$.

All of the other relevant definitions and ambiguities will be discussed in later sections. Our first result generalizes the work of Donaldson [Don87], Corlette [Cor88], and Labourie [Lab91] and may also be regarded as an equivariant extension of [Wol91b, Theorem 3.11]. Naturally, a portion of our analysis resembles that of Wolf.

Theorem 3A. Let $\Sigma = \tilde{\Sigma}/\Gamma$ be a complete finite volume hyperbolic surface and (X, g) a Hadamard manifold. Let $\rho : \Gamma \to \mathrm{Isom}(X, g)$ be a reductive representation. There exists a ρ-equivariant harmonic map $f : \tilde{\Sigma} \to X$.

If we assume X is CAT(-1), we may construct f so that if γ is a peripheral isometry and $\theta \in \mathbb{R}$, the Hopf differential Φ has the following behaviour at the corresponding cusp:

- if $\rho(\gamma)$ is parabolic or elliptic, Φ has a pole of order at most 1 and
- if $\rho(\gamma)$ is hyperbolic, Φ has a pole of order 2 with residue

$$-\Lambda(\theta)\ell(\rho(\gamma))^2/16\pi^2.$$

Suppose that ρ does not fix a point on $\partial_\infty X$. Then all harmonic maps whose Hopf differentials have poles of order at most 2 at the cusps are of this form. If ρ stabilizes a geodesic, then any other harmonic map with the same asymptotic behaviour differs by a translation along that geodesic.

Regarding domination, the next theorem is the main result of this chapter.

Theorem 3B. Let $\Sigma = \tilde{\Sigma}/\Gamma$ be a complete finite volume hyperbolic orbifold and (X, g) a CAT(-1) Hadamard manifold. Let $\rho : \Gamma \to \text{Isom}(X, g)$ be any representation. There exists a geometrically finite representation j_Σ dominating ρ in length spectrum. If ρ is reductive, then j_Σ dominates ρ in the traditional sense. There is a family of convex cocompact Fuchsian representations strictly dominating j_Σ. Given a peripheral isometry γ,

- if $\rho(\gamma)$ is not hyperbolic, then $j_\Sigma(\gamma)$ is parabolic and
- if $\rho(\gamma)$ is hyperbolic, $j_\Sigma(\gamma)$ is hyperbolic with the same translation length.

In general j_Σ will not strictly dominate ρ. This will be discussed in detail in Section 3.6. If $X = \mathbb{H}$ and ρ is Fuchsian with no elliptic monodromy it will follow from the proof that $j_\Sigma = \rho$. For holonomy representations of closed surfaces, Thurston observed in [Thu98, Proposition 2.1] that strict domination contradicts the Gauss-Bonnet theorem and is therefore impossible.

Most of the proof of Theorem 3B is devoted to constructing j_Σ. To upgrade to a strictly dominating representation we perform a *strip deformation*, a procedure introduced by Thurston [Thu98] and further developed in [DGK16a].

Setting $X = \mathbb{H}$ in Theorem 3B, from [GK17, Theorem 1.8] we obtain:

Theorem 3C. Let $\Sigma = \tilde{\Sigma}/\Gamma$ be a complete finite volume hyperbolic orbifold and $\rho : \Gamma \to \mathrm{PSL}_2(\mathbb{R})$ any representation. Then there is a Fuchsian representation j_Σ dominating ρ and a family of convex cocompact representations (j_Σ^α) strictly dominating j_Σ such that

$$(\rho \times j_\Sigma^\alpha)(\Gamma) \subset \mathrm{PSL}_2(\mathbb{R}) \times \mathrm{PSL}_2(\mathbb{R})$$

admits a properly discontinuous action on $\mathrm{PSL}_2(\mathbb{R})$ preserving the Lorentz metric of constant curvature -1. If $\gamma \in \Gamma$ is elliptic and $\rho(\gamma)$ has smaller order than γ, then the action is torsion free as well. Consequently there exists a geometrically finite AdS 3-manifold Seifert-fibered over $\mathbb{H}/j_\Sigma^\alpha(\Gamma)$.

Note that if Σ is a manifold, the torsion condition always holds.

As an intermediate step in the proof of Theorem 3B, we obtain a result of independent interest. Let Σ be a complete finite volume hyperbolic surface with n punctures and let $T(\Sigma, p_1, \ldots, p_{d_1}, \ell_{d_1+1}, \ldots, \ell_n)$ denote the subspace of the Fricke-Teichmüller space of Σ consisting of holonomies of hyperbolic surfaces with d_1 ordered punctures and d_2 ordered geodesic boundary components of length $\ell_{d_1+1}, \ldots, \ell_{d_2} > 0$. Let $(\theta_k)_{k=d_1+1}^n \subset \mathbb{R}$ and $P := (\ell_k, \theta_k)_{k=d_1+1}^n$. Denote by $Q(\Sigma, P)$ the space of holomorphic quadratic differentials on Σ with poles of order at most one at the punctures corresponding to cusps and poles of order 2 with residue

$$-\Lambda(\theta_k)\ell_k^2/16\pi^2$$

for each puncture labelled by ℓ_k. From the results in [Wol91b], for each point in $T(\Sigma, p_1, \ldots, p_{d_1}, \ell_1, \ldots, \ell_{d_2})$, there is a unique homotopic harmonic diffeomorphism $h_f : \Sigma \to S$ whose Hopf differential lives in $Q(\Sigma, P)$.

Theorem 3D. Let Σ be a finite volume hyperbolic surface. The map

$$\Psi : T(\Sigma, p_1, \ldots, p_{d_1}, \ell_1, \ldots, \ell_{d_2}) \to Q(\Sigma, P)$$

given by $[S, f] \mapsto \mathrm{Hopf}(h_f)$ is a homeomorphism.

We expect the above result is known to experts, but could not find a proof in the literature. Hence we supply our own. The parametrization of the Teichmüller space of a closed surface by holomorphic quadratic differentials goes back to Sampson, Schoen-Yau, and Wolf (see [Wol89] for the full result). The case of Teichmüller spaces of punctured surfaces, corresponding to differentials with a pole of order at most 1, was completed by Lohkamp [Loh91]. In [Gup17], Gupta parametrized *wild Teichmüller spaces* by certain equivalence classes of holomorphic differentials with poles of order at least 3. Theorem 3D thus completes a description of the space of meromorphic quadratic differentials over a Riemann surface in terms of harmonic diffeomorphisms.

We end this subsection by presenting quick corollaries of Theorem 3B, unrelated to the rest of the chapter. When X is a CAT(-1) Hadamard manifold and $\rho : \Gamma \to \mathrm{Isom}(X, g)$ is geometrically finite, the *limit set* of $\rho(\Gamma)$ is the set of limit points of $\Gamma \cdot z$ in $\partial_\infty X$ for a fixed point z in X. It is a standard exercise to confirm that this does not depend on the point z. When $X = \mathrm{PSL}_2(\mathbb{R})$ and ρ is Fuchsian, the limit set is either the full circle $\partial_\infty \mathbb{H}$ or a Cantor set. The *critical exponent* $\delta(\rho)$ is the smallest constant s such that the Poincaré series

$$\sum_{\gamma \in \Gamma} e^{-sd(z,\rho(\gamma)\cdot z)}$$

converges, and it coincides with the Hausdorff dimension of the limit set (see [Coo93] for a proof). The analogue of the following result is known for closed surfaces and is observed in [DT16], but to the author's knowledge it is new in our context.

Corollary 3E. Let $\Sigma = \tilde{\Sigma}/\Gamma$ be a complete finite volume hyperbolic orbifold and (X, g) a CAT(-1) Hadamard manifold. Let $\rho : \Gamma \to \mathrm{Isom}(X, g)$ be a geometrically finite representation. There is a Fuchsian representation j_Σ such that the Hausdorff dimension of the limit set of $\rho(\Gamma)$ is bounded below by that of j_Σ. j_Σ has the following property around a peripheral γ:

- if $\rho(\gamma)$ is not hyperbolic, then $j_\Sigma(\gamma)$ is parabolic and
- if $\rho(\gamma)$ is hyperbolic, then $j_\Sigma(\gamma)$ is hyperbolic with $\ell(j_\Sigma(\gamma)) = \ell(\rho(\gamma))$.

The Hausdorff dimension can be estimated and sometimes fully understood from the monodromy around the punctures. For instance, if Σ is a pair of pants and ρ takes the cuffs to isometries with lengths $a, b, c > 0$, then the Hausdorff dimension of the limit set of j_Σ occurs as a zero of a certain *Selberg zeta function*

$$Z_{a,b,c}(s) = \prod_{\gamma \in \Gamma} \prod_{m=0}^{\infty} \left(1 - e^{-(s+m)\ell(\gamma)}\right).$$

The map $\gamma \mapsto \ell(\gamma)$ is determined entirely by a, b, c. These zeroes can be computed efficiently (see [PV17] for details).

Outline and strategy of proof

In the next section we introduce the relevant definitions and notations in the representation theory of discrete groups, and we prove some preparatory results about equivariant harmonic maps for surfaces with punctures. In Section 3.3 we prove the energy domination lemma, which says that the energy of an equivariant harmonic map is bounded above by that of a special harmonic diffeomorphism of the disk with the same Hopf differential. As is standard in this field, we argue via an analysis of the Bochner formula. This estimate is a central technical results of this paper, and is instrumental in proving Theorem 3B.

In Section 4.4 we prove Theorem 3D using classical techniques from the theory of harmonic maps. Section 3.5 is devoted to the proof of Theorem 3A. Infinite energy harmonic maps were constructed for some special cases in [Wol91b], [Sim90], [JZ97], and [KM08]. Consequently, there is nothing truly novel in the proof of the general existence result—it is an amalgamation of known ideas. The real work is done in studying the behaviour and uniqueness of the harmonic maps. We combine the energy estimate from Section 3/3 with Theorem 3D to control the energy locally, as well as a distance comparison to a special non-harmonic map to understand the directions in which our map should expand and contract.

In Section 3.6, we attempt to follow the approach of [DT16] to prove domination in the compact case. We take an equivariant harmonic map f from Theorem 3A and choose a harmonic diffeomorphism h from Σ to the convex core of some geometrically finite hyperbolic surface N that has the same Hopf differential as f. From our energy estimates, $f \circ h^{-1}$ is 1-Lipschitz and intertwines ρ with the holonomy of N, but it is not strictly 1-Lipschitz (this issue does not occur in [DT16]). We introduce strip deformations to strictly dominate the holonomy of N, completing the proof of Theorem 3B.

Other recent work

Shortly after a preprint of the paper was posted to the arXiv, Gupta-Su proved the same domination result for representations to $\mathrm{PSL}_2(\mathbb{C})$ [GS20]. Their proof is different: they straighten the pleated plane determined by the Fock-Goncharov coordinates associated to a framed representation, and then use strip deformations.

Acknowledgments

It is a pleasure to thank my advisor, Professor Vladimir Marković, for his support, insight, patience, and guidance. I would also like to thank Qiongling Li and Peter

Smillie for their interest and helpful conversations, as well as François Guéritaud for graciously answering some questions over email. Finally, I would like to thank my friend Arian Jadbabaie for helping create the figure on Inkscape.

3.2 Representations of discrete groups

Geometrically finite surfaces

Firstly, we review isometries for Hadamard manifolds. The translation length of an isometry $\gamma \in G$ is

$$\ell(\gamma) = \inf_{x \in X} d_\nu(\gamma \cdot x, x),$$

and γ is

- elliptic if $\ell(\gamma) = 0$ and the infimum is attained,
- parabolic if $\ell(\gamma) = 0$ and the infimum is not attained, and
- hyperbolic if $\ell(\gamma) > 0$.

Elliptic isometries fix points inside X, parabolic isometries fix points on $\partial_\infty X$, and a hyperbolic isometry stabilizes a unique geodesic on which it acts by translation of length $\ell(\gamma)$. Around a peripheral ζ, we say a representation ρ has hyperbolic, parabolic, or elliptic mondromy if $\rho(\zeta)$ is hyperbolic, parabolic, or elliptic respectively.

Let μ be any metric on Σ, with no constraint on the volume. A cusp region is a neighbourhood surrounding a puncture of Σ that identifies isometrically with

$$U(\tau) := \{z = x + iy : (x, y) \in [0, \tau] \times [a, \infty)\}/\langle z \mapsto z + \tau \rangle,$$

equipped with the hyperbolic metric $y^{-2}|dz|^2$.

Definition 3.2.1. The convex core $C(\Sigma, \mu)$ of (Σ, μ) is the quotient of the convex hull of the limit set of Γ by the action of Γ. We denote the convex hull of the limit set by $\tilde{C}(\Sigma, \mu)$. We sometimes omit the metric from the notation and write $C(\Sigma)$, $\tilde{C}(\Sigma)$.

It is the minimal convex set such that the inclusion $C(\Sigma, \mu) \to \Sigma$ is a homotopy equivalence. The convex core is a finite volume hyperbolic surface with finitely many cusps and geodesic boundary components. (Σ, μ) can be recovered from $C(\Sigma, \mu)$ by attaching infinite funnels along the boundary components. In the language of representations, the monodromy of the holonomy representation around a cusp is parabolic, and for a geodesic boundary component it is hyperbolic.

Optimal Lipschitz constants

Given Γ discrete, $\rho : \Gamma \to \mathrm{Isom}(X, g)$, and $j : \Gamma \to \mathrm{PSL}_2(\mathbb{R})$ geometrically finite, we set

$$C(j, \rho) := \inf \mathrm{Lip}(f),$$

where the infimum is taken over the family of all (j, ρ)-equivariant Lipschitz maps. The theorem below is Theorem 1.8 in [GK17].

Theorem 3.2.2. *(Guéritaud, Kassel) Let Γ be a discrete group and $\rho, j : \Gamma \to PSL_2(\mathbb{R})$ two representations with j geometrically finite. Then $C(j, \rho) < 1$ if and only if*

$$C(j, \rho)' := \sup \frac{\ell(\rho(\gamma))}{\ell(j(\gamma))} < 1,$$

unless ρ has exactly one fixed point on $\partial_\infty \mathbb{H}$ and there exists a $\gamma \in \Gamma$ such that $j(\gamma)$ is parabolic and $\rho(\gamma)$ is not elliptic.

Remark 3.2.3. As we will see later, equivariant harmonic maps only exist for reductive representations. To dominate non-reductive representations we would like to use a version of Theorem 3.2.2 that holds for variable curvature. The result and the proof of Theorem 3.2.2 do not directly transfer, and trying to extend them is outside the scope of the current work. Hence, for the non-reductive case we settle for length spectrum domination, although we expect the full domination result to be true. From the theorem above, non-reductive representations still lead to AdS 3-manifolds, which is the most important application.

Now suppose $\Sigma = \tilde{\Sigma}/\Gamma$ is a complete finite volume hyperbolic orbifold. By the Selberg lemma, Γ admits a finite index torsion free normal subgroup Γ_0. The quotient $\tilde{\Sigma}/\Gamma_0$ is a complete finite volume hyperbolic manifold. We close this section with a lemma that reduces Theorem 3B to the case of hyperbolic manifolds.

Lemma 3.2.4. *Let Γ be a discrete group and $\Gamma_0 \subset \Gamma$ a finite index normal subgroup. Let $\rho : \Gamma \to Isom(X, g)$ and $j : \Gamma \to PSL_2(\mathbb{R})$ be representations and let ρ_0 and j_0 be their restrictions to Γ_0. Then $C(j, \rho) = C(j_0, \rho_0)$.*

This is essentially done in [GK17], although the authors prove something more general and restrict to the case $X = \mathbb{H}^n$. For the convenience of the reader, we essentially repeat the proof. We use a lemma from [GK17].

Lemma 3.2.5. *Let I be any countable index set and $\alpha = (\alpha_i)_{i\in I} \subset \mathbb{R}$ a sequence summing to* 1*. Given $p \in K \subset \mathbb{H}$ and $f_i : K \to X$, $i \in I$ such that*

$$\sum_{i\in I} \alpha_i d(f_1(p), f_i(p)) < \infty,$$

the map

$$f := \sum_{i\in I} \alpha_i f_i \ , \ x \mapsto argmin\Big\{p' \in X : \sum_{i\in I} \alpha_i d(p', f_i(x)) < \infty\Big\}$$

is well-defined and satisfies

$$Lip_x(f) \leq \sum_i \alpha_i Lip_x(f_i) \ , \ Lip_Y(f) \leq \sum_i \alpha_i Lip_Y(f_i).$$

If each f_i is equivariant with respect to a pair of representations then so is f.

The authors give a proof for $X = \mathbb{H}^n$ but the proof only uses the fact that $\mathbb{H}^n$ is a CAT(0) metric space.

Proof of lemma 3.2.4. If no (j', ρ')-equivariant maps exist there is nothing to prove, so assume otherwise. The inequality $C(j', \rho') \leq C(j, \rho)$ is obvious because any (j, ρ)-equivariant map is (j', ρ')-equivariant. As for the other inequality, write

$$\Gamma = \coprod_{i=1}^{r} \gamma_i \Gamma_0$$

for some collection of coset representatives γ_i. Let f be a (j', ρ')-equivariant map. Notice that for any $\gamma \in \Gamma$, the map

$$f_\gamma := \rho(\gamma)^{-1} \circ f \circ j(\gamma)$$

depends only on the coset $\gamma\Gamma_0$. Indeed, suppose we are given $\gamma_1, \gamma_2 \in \Gamma$ such that $\gamma_1\gamma_2^{-1} \in \Gamma_0$. For $x \in \mathbb{H}$ let $y = j(\gamma_2)^{-1}x$. Then

$$f_{\gamma_1}(x) = \rho(\gamma_1)^{-1} \circ f(j(\gamma_1\gamma_2^{-1})y) = \rho(\gamma_2)^{-1} \circ f(y) = f_{\gamma_2}(x).$$

By Lemma 3.2.5 the map

$$f' := \sum_{i=1}^{r} \frac{1}{r} \cdot f_{\gamma_i}$$

satisfies

$$\rho(\gamma)^{-1} \circ f' \circ j(\gamma) = \sum_{i=1}^{r} \frac{1}{r} \cdot f_{\gamma\gamma_i} = f'$$

since the sum in the middle is just a rearrangement of the sum describing f'. By Lemma 3.2.5 again we have $\mathrm{Lip}(f') \leq \mathrm{Lip}(f)$. Taking $\mathrm{Lip}(f) \to C(j', \rho')$, the lemma follows. □

Useful lemmas

We collect some general results on harmonic maps that we'll use throughout.

Theorem 3.2.6 (Ishihara)**.** *Suppose that all sectional curvatures of a manifold* (X, ν) *are non-negative. Then* $f : (\Sigma, \mu) \to (X, \nu)$ *is harmonic if and only if it pulls back germs of convex functions to germs of subharmonic functions.*

We record a corollary.

Corollary 3.2.7. *Suppose that all sectional curvatures of a manifold* (X, ν) *are non-negative and let* $f_1, f_2 : (\Sigma, \mu) \to (X, \nu)$ *be harmonic maps. Let* d_ν *be the Riemannian distance function on* (X, ν)*. Then the function on* Σ *given by* $p \mapsto d(f_1(p), f_2(p))$ *is subharmonic.*

A harmonic function between Euclidean spaces has a representation in terms of the Poisson integral formula. Out of this formula, one can obtain local C^k bounds in terms of local C^0 bounds. For harmonic maps between manifolds, Cheng's lemma gives C^1 bounds in terms of C^0 control.

Lemma 3.2.8 (Cheng's lemma)**.** *Let X and Y be Hadamard manifolds with* $-b^2 \leq K_X \leq 0$ *and* $\dim X = k$*. Let* $z \in X$*,* $r > 0$*, and let* $h : B(x, r) \to Y$ *be a* C^∞ *harmonic map such that the image* $h(B(z, r))$ *is contained in a ball of radius* R_0*. Then*

$$||Dh(z)|| \leq 2^5 k \frac{1 + br}{r} R_0.$$

See [Che80]. Local C^k bounds are deduced from local C^1 bounds via elliptic bootstrapping.

Energy of harmonic maps

We prove some preliminary results relevant to equivariant harmonic maps from surfaces with punctures.

Proposition 3.2.9. *If* $\Sigma = \tilde{\Sigma}/\Gamma$ *is a complete finite volume hyperbolic surface,* (X, g) *is a Hadamard manifold, and* $\rho : \Gamma \to Isom(X, g)$ *is a representation, then a finite energy* ρ*-equivariant map exists if and only if* ρ *has no hyperbolic monodromy.*

Before we begin, we modify the metric to a new one that will be used throughout the chapter. Label the cusp neighbourhoods $C_1, \ldots, C_n$. Take collar neighbourhoods U_k of each ∂C_k inside $\Sigma \backslash C_k$ and consider the metric on Σ that agrees with the hyperbolic metric on $\Sigma\backslash(\cup_k C_k)$ and is flat on each $C_k \cup U_k$. Then interpolate on a neighbourhood of $\partial U_k \backslash \partial C_k$ that does not touch ∂C_k to a smooth non-positively curved metric σ', conformally equivalent to the hyperbolic metric. We will call this the flat-cylinder metric.

We also take this opportunity to introduce the *transverse horospherical flow*. With X as above, consider a horoball $B \subset X$ with horospherical boundary H centered at the fixed point ξ of a parabolic isometry ψ. The subgroup generated by ψ preserves H and B. The data (B, H, ξ) determines a flow $\varphi_t : B \times [0, \infty) \to B$ defined by

$$\varphi_t(p) = \alpha_{p,\xi}(t),$$

where $\alpha_{p,\xi} : [0, \infty) \to X$ is the unique geodesic starting from p and tending towards ξ at ∞.

Lemma 3.2.10. *The transverse horospherical flow is $\langle\psi\rangle$-equivariant.*

Proof. Notice

$$\alpha_{\psi \cdot p,\xi}(0) = \psi \cdot p = \psi \cdot \alpha_{p,\xi}(0).$$

Since $\alpha_{\psi \cdot p,\xi}(t)$ and $\psi \cdot \alpha_{p,\xi}(t)$ describe geodesics with the same starting point and end point, they are identical. □

Proof of proposition 3.2.9. By conformal invariance of energy we're permitted to do all of our computations in the flat-cylinder metric. Firstly let us assume there is a peripheral γ such that $\rho(\gamma)$ is hyperbolic. Take any equivariant map $f : \tilde{\Sigma} \to X$ and fix a cusp neighbourhood associated to the peripheral and isometric to $U(\tau)$. As $\rho(\gamma)$ is hyperbolic,

$$d_g(f(iy), f(\tau + iy)) = d_g(f(iy), \rho(\gamma) f(iy)) \geq \ell(\rho(\gamma)) > 0,$$

independent of y. For each y let γ_y be the path $x \mapsto f(x + iy)$, $x \in [0, \tau]$. The inequality above implies

$$\ell(\rho(\gamma)) \leq \int_0^\tau ||d\gamma_y||_{\sigma'} dy$$

and by Cauchy-Schwarz we obtain

$$\frac{\ell(\rho(\gamma))^2}{2\tau} \leq \frac{1}{2} \int_0^\tau ||d\gamma_y||^2_{\sigma'} dy \leq \int_0^\tau e(f)(x, y) dy.$$

Hence,

$$E(f) \geq E_V(f) = \int_a^\infty \int_a^\tau e(f)(x,y)dxdy \geq \frac{\ell(\rho(\gamma))^2}{2\tau}\int_a^\infty dy = \infty,$$

which shows all equivariant maps have infinite energy.

For the other direction, we simply produce an equivariant finite energy map. We build a finite energy map in a neighbourhood of each cusp, equivariant with respect to the subgroup generated by $\rho(\gamma_j)$ and then extend smoothly to a ρ-equivariant map on the (compact) complement of the cusps.

By induction it suffices to assume that there is only one cusp neighbourhood V. We identify it with some $U(\tau)$. Let γ be the corresponding curve. If $\rho(\gamma)$ is elliptic then we simply map all of V to a fixed point of $\rho(\gamma)$. This is clearly equivariant and has zero energy in V. Henceforward we assume $\rho(\gamma)$ is parabolic. $\langle\rho(\gamma)\rangle$ stabilizes a horoball B with horopsherical boundary H. Let g be any C^∞ $\rho|_{\langle\gamma\rangle}$-equivariant map $\mathbb{R} \to H$. Define $f : \tilde{V} \to B$ by

$$f(x+iy) = \varphi_{v\log(y+1)}(g(x)),$$

where φ is the transverse horospherical flow with respect to the fixed point and $v > 0$ will be specified later. We compute

$$|df(\partial/\partial y)|_{f(x+iy)} = |\partial/\partial y(v\log(y+1))| = \frac{v}{y+1}.$$

Next, note that

$$J_x(y) := \frac{\partial}{\partial x}f(x+iy)$$

is a Jacobi field for each x. By the curvature assumption on X, the Rauch comparison theorem shows that any Jacobi field on X along a geodesic decays exponentially in time: there is a $u > 0$ such that

$$|J_x(y)| \leq Ae^{-u\cdot v\log(y+1)}$$

for all x. Now choose v so that $uv \geq 1$. Then

$$|df(\partial/\partial x)|_{f(x+iy)} \leq \frac{A}{(y+1)^{uv}},$$

and furthermore

$$E_V(f) \leq \frac{1}{2}\int_0^\infty \int_0^\tau \frac{v^2+A^2}{(y+1)^2}dxdy = \frac{\tau}{2}(v^2+A^2) < \infty,$$

and the result follows. □

Remark 3.2.11. The total energy of a harmonic map is finite if and only if the Hopf differential is integrable. Passing to polar coordinates, we see that an integrable holomorphic quadratic differential has a pole of order at most 1 at a puncture.

Suppose a representation admits a finite energy equivariant map. If it does not fix a point on the ideal boundary, the harmonic map determined by Theorem 2.4.4 is unique. If ρ stabilizes a geodesic, there is a 1-parameter family of harmonic maps that differ by translations along that geodesic axis. The standard methods push through to give a uniqueness criterion in our setting.

Lemma 3.2.12. *Let Σ be a complete finite volume hyperbolic surface, let (X, g) be Hadamard, and let f_1 and f_2 equivariant harmonic maps for ρ such that the map $z \mapsto d(f_1, f_2)(z)$ is bounded. If ρ does not fix a point on $\partial_\infty X$ then $f_1 = f_2$. If ρ stabilizes a geodesic, then f_1 and f_2 may differ by translation along a geodesic.*

Proof. For $z \in \Sigma$ let $\{e_1, e_2\}$ be an orthonormal frame for the tangent bundle in a neighbourhood of z and let $\{v_1^0, \dots, v_n^0\}$, $\{v_1^1, \dots, v_n^1\}$ be orthonormal frames for neighbourhoods of $f_1(z)$, $f_2(z)$ respectively. In these frames we write

$$(f_k)_* e_i = \sum_{m=1}^{n} \lambda_{i,m}^k v_m^k.$$

$\{v_1^0, \dots, v_n^0, v_1^1, \dots, v_n^1\}$ is an orthonormal frame near $(f_1(z), f_2(z)) \in X \times X$. Define vector fields $X_i \in \Gamma(T(X \times X))$ so that around $(f_1(z), f_2(z))$ the projections onto the first and second factors are $f_1^* e_i$ and $f_2^* e_i$ respectively. Let $d : \tilde{\Sigma} \to \mathbb{R}$ be the function

$$d(z) = d_{g \oplus g}(f_1(z), f_2(z))$$

which is C^∞ away from the diagonal. From a computation in [SY97, Chapter 11.2], if we assume $f_1(z) \neq f_2(z)$ then from the fact that the f_k are harmonic,

$$\Delta d^2 \geq 2d \sum_{i=1}^{2} D^2 d_g(X_i, X_i)$$

around z. Above, Δ is the Laplacian on $\tilde{\Sigma}$ and $D^2 d_g$ is the Hessian of d_g, the distance function on (X, g).

By equivariance, d descends to a bounded subharmonic function on Σ. As Σ is parabolic in the potential theoretic sense, this function is constant. Therefore,

$$2d \sum_{i=1}^{2} D^2 d_g(X_i, X_i) = 0.$$

This forces $d = 0$ or $D^2 d_g(X_i, X_i) = 0$. In the first case we have $f_1 = f_2$ so let us move to the latter. From an argument in [SY97, Chapter 11.2], this implies either $f_1 = f_2$ or f_1 and f_2 have image in a geodesic and differ by a translation along that geodesic. By equivariance, this last case can only occur if ρ stabilizes a geodesic. □

3.3 Energy domination

Domination inequality

The goal of this subsection is to prove Proposition 3.3.2. Let σ denote the hyperbolic metric on $\mathbb{H}$ with constant curvature -1. Unless otherwise specified, for the rest of the chapter this is the metric on $\mathbb{H}$. We recall from [Wan92] a fundamental result in the theory of harmonic maps between surfaces.

Theorem 3.3.1. *(Wan) For any holomorphic quadratic differential Φ on $\mathbb{H}$, there is a (possibly non-surjective) harmonic diffeomorphism $h : \mathbb{H} \to \mathbb{H}$ such that*

- *$H(h) \geq 1$,*
- *The metric $H(h)\sigma(z)|dz|^2$ is complete on $\mathbb{H}$, and*
- *Hopf$(h) = \Phi$.*

The harmonic map is unique up to isometries.

We now have the machinery to state the energy estimate.

Proposition 3.3.2. *Suppose f is a harmonic map from $\mathbb{H}$ to a CAT(-1) Hadamard manifold (X, g). The energy density is always bounded above by that of any harmonic diffeomorphism $h : \mathbb{H} \to \mathbb{H}$ with the same Hopf differential and $H(h) \geq 1$. The inequality is strict unless f takes $\mathbb{H}$ into a totally geodesic plane of constant sectional curvature -1.*

This is essentially a non-compact and generalized version of [DT16, Lemma 2.1]. Proposition 3.3.2 is a consequence of the next lemma.

Lemma 3.3.3. *$H(f) \leq H(h)$ everywhere on $\mathbb{H}$, with equality for one of them if and only if and f maps $\mathbb{H}$ diffeomorphically into a totally geodesic plane $\mathbb{H} \subset X$ of constant curvature -1. If equality holds at one point, it holds everywhere.*

Proof of proposition 3.3.2. Indeed, assuming the above result, if f is not such an embedding with totally geodesic image, then $L(f) < H(h)$ and $H(f) < H(h)$ everywhere. As

$$H(f)L(f) = ||\Phi||^2 = H(h)L(h),$$

we obtain $L(f) > L(h), H(f) > L(h)$. Since $H(f) - L(f) \geq 0$, $(H(f) - L(f))^2 < (H(h) - L(h))^2$. Adding $4H(f)L(f) = 4H(h)L(h)$ to both sides, we have

$$e(f)^2 = (H(f) + L(f))^2 < (H(h) + L(h))^2 = e(h)^2,$$

which yields the desired result. □

We will prove Lemma 3.3.3 via the Bochner formula. Our main tool is the generalized maximum principle of Omori-Yau (see [Omo67] for the version we use and also [CY75, Theorem 3] for the extension to manifolds with a lower bound on the Ricci curvature).

Lemma 3.3.4. *(Omori) Let M be a Riemannian manifold such that all sectional curvatures are bounded from below. Let f be a C^2 function on M that is bounded above. There is a sequence $(x_n)_{n=1}^{\infty}$ such that $f(x_n) \to \sup f$ and*

$$|\nabla f(x_n)| \to 0\ ,\ \limsup_{n\to\infty} \Delta f(x_n) \leq 0$$

as $n \to \infty$.

Proof of lemma 3.3.3. By [Sam78, Corollary 3], the set $\mathcal{D}$ on which f^*g is nondegenerate is either empty or open and dense. From [SY97, page 10] either $H(f) = 0$ everywhere or the zeroes are isolated. We replace $\mathcal{D}$ with the open dense set

$$U := \mathcal{D} - \{z : H(f)(z) = 0\}.$$

For $H(h)$ and $H(f)$,

- $H(h) \geq 1$ and the Bochner formula is

$$\Delta \log H(h) = 2H(h) - 2\frac{||\Phi||^2}{H(h)} - 2.$$

- On U, $H(f)$ solves the Bochner formula

$$\Delta \log H(f) = -2\kappa(f^*g)H(f) + 2\kappa(f^*g)\frac{||\Phi||^2}{H(f)} - 2.$$

The next result is essentially [Sam78, Theorem 4]. It was tweaked to its present form in [DT16].

Lemma 3.3.5. *For all $x \in U$, $\kappa(f^*g) \leq -1$. We have equality iff the second fundamental form of $f(\mathbb{H})$ vanishes at x. In particullar, $\kappa(f^*g) = -1$ everywhere on U iff $f(U) \subset X$ is totally geodesic.*

When the metric f^*g is degenerate,

$$H(f) = ||\Phi|| = H(h)^{1/2}L(h)^{1/2} < H(h).$$

Hence we can dismiss the case $U = \emptyset$ and furthermore we're allowed to work only on U. As stated previously $H(h) \geq 1$ everywhere, so that $H(f)/H(h)$ never vanishes on U. Assume for the sake of contradiction that $H(f) > H(h)$ at a point x. Necessarily, $H(f)(x) \geq L(f)(x)$. From the Bochner formula we have

$$\begin{aligned}\Delta \log(H(f)/H(h)) = 2(H(h) - H(f)) + 2||\Phi||^2(H(h)^{-1} - H(f)^{-1})\\ - 2(\kappa(f^*g) + 1)(H(f) - L(f)).\end{aligned}$$

Let $w := \log(H(f)/H(h))$. Since $\kappa(f^*g) \leq -1$ and $H(h) \geq L(h)$, one can simplify the above equation to

$$\Delta w \geq 2(H(f) - H(h))(1 + ||\Phi||^2/(H(h)H(f))) = 2(H(h) + L(h))(e^w - 1),$$

and hence

$$\Delta w \geq 2(e^w - 1) > 0$$

at such an x. It now follows that this point x cannot be a local maximum for $H(f)/H(h)$ as otherwise

$$0 \geq \Delta w > 0.$$

Thus, there is a sequence contained in U and tending to the boundary of $\mathbb{H}$ along which $H(f)/H(h) > 1$ and increases to $\sup H(f)/H(h)$. We argue this supremum is finite. Let $b > a > 0$ and define $F : [b, \infty) \to \mathbb{R}$ by

$$F(s) = \int_b^s \Big(\int_a^t e^{2\tau} d\tau\Big)^{-1/2} dt + 1.$$

F is monotonically increasing, bounded above, and $F'' < 0$ everywhere. Extend F smoothly to $\mathbb{R}$ so that it is still monotonic and satisfies

$$\lim_{t\to-\infty} F(t) > 0.$$

For some large $N > 0$, let $\eta : \mathbb{H} \to [0,1]$ be a C^∞ function that is 0 on the open set $\{z : e^{w(z)} < 1/2N\}$ and is 1 on $\{z : e^{w(z)} \geq 1/N\}$. Then $F \circ (\eta w)$ is a bounded C^∞ function on $\mathbb{H}$. By the Omori-Yau maximum principle there is a sequence $(x_n)_{n=1}^\infty$ escaping to the boundary such that

$$F \circ w(x_n) = F \circ (\eta w)(x_n) \to \sup F \circ \eta w = \sup F \circ w$$

as $n \to \infty$ and

$$|\nabla F \circ w(x_n)| = |\nabla F \circ (\eta w)(x_n)| \to 0\,,\ \limsup_{n\to\infty} \Delta F \circ w(x_n) = \limsup_{n\to\infty} \Delta F \circ (\eta w)(x_n) \leq 0.$$

Above, we removed finitely many points in the sequence so we can assume $w(x_n) > 2/N$ always, and we also used the fact that $w = \eta w$ in this region.

An analogue of the computation below is contained in [CY75, Section 5], where they work with global subsolutions. We may choose a subsequence of the x_n, and abuse notation by still labelling it x_n, so that

$$0 \leq \frac{F'(w)|\nabla w|}{F(w)^2}(x_n) \leq 1/n\ (*)$$

and

$$-\frac{F''(w)|\nabla w|^2}{F^2}(x_n) - \frac{F'(w)\Delta w}{F^2}(x_n) + \frac{(F')^2|\nabla w|^2}{F^3}(x_n) \geq -1/n.$$

Multiplying the above by $(F'(w))^2/F(w)^2|F''(w)|$ we obtain

$$-\frac{F''(w)}{|F''(w)|}\frac{F'(w)^2|\nabla w|^2}{F^4} - \frac{F'(w)^3}{F^4}\frac{\Delta w}{|F''(w)|} + \frac{F'(w)^2}{F(w)|F''(w)|}\frac{F'(w)^2|\nabla w|^2}{F(w)^4}$$
$$\geq \frac{-F'(w)^2}{nF(w)^2|F''(w)|}$$

at x_n. Note F satisfies

$$\limsup_{s\to\infty} \frac{|F'(s)|^2}{F(s)|F''(s)|} < \infty,$$

and combining this with the line above yields that

$$\frac{1}{n^2} - \frac{F'(w)^3}{F^4(w)}\frac{\Delta u}{|F''(w)|} + \frac{A}{n^2} \geq -\frac{A}{n}.$$

Using $\Delta w \geq 2(e^w - 1)$ at x_n we infer

$$\frac{F'(w)^3(e^{2w}-1)}{F^4(w)|F''(w)|}(x_n) \lesssim \frac{1}{n}$$

as $n \to \infty$. However, it is straightforward to compute

$$\liminf_{s\to\infty} \frac{F'(s)^3(e^{2s}-1)}{F(s)^4|F''(s)|} > 0.$$

This means $\limsup_{n\to\infty} w(x_n) = \infty$ is impossible.

To understand this supremum we apply the generalized maximum principle to the function ηw. As with $F \circ (\eta w)$, there is a sequence $(y_n)_{n=1}^{\infty}$ leaving all compact subsets of $\mathbb{H}$ such that after refining if necessary so that $w(y_n) > 2/N$,

$$w(y_n) = \eta w(y_n) \to \sup \eta w = \sup w$$

and

$$0 \geq \limsup_{n\to\infty} \Delta \eta w(y_n) = \limsup_{n\to\infty} \Delta w(y_n) \geq 2(e^{w(y_n)} - 1) \geq 0.$$

This forces $\sup H(f)/H(h) = 1$, which contradicts our assumption that $H(f) > H(h)$ at least once. Hence, $H(f) \leq H(h)$ always. Now that we have our inequality, a special case of [Min87, Theorem 1] indicates when this inequality is strict.

Lemma 3.3.6. *Let u be a real non-positive function on a domain V in the complex plane such that $\Delta u \geq Au$ for a constant $A > 0$. Then either $u = 0$ on V or $u < 0$ on all of V.*

With this in mind, take an increasing exhaustion $(D_k)_{k=1}^{\infty}$ of $\mathbb{H}$ by pre-compact open sets. $e^x \geq x + 1$ gives

$$\Delta w \geq (H(f) - H(h))w \geq [\max_{D_k}(H(f) - H(h))]w$$

in $D_k \cap U$. It follows that either $H(f) = H(h)$ or $H(f) < H(h)$ everywhere. If $H(h) = H(f)$ then $L(h) = L(f)$ and we see that f^*g is non-degenerate everywhere. By the Bochner formula above this forces $\kappa(f^*g) = -1$, and so by Lemma 3.3.5 f maps $\mathbb{H}$ diffeomorphically into a totally geodesic plane. Identifying this plane with $\mathbb{H}$, the formulas

$$h^*\sigma = e(h)\sigma + \Phi + \overline{\Phi} \quad , \quad f^*g = e(f)\sigma + \Phi + \overline{\Phi}$$

show that f differs from h by an isometry of $\mathbb{H}$. This implies $f : (\mathbb{H}, h^*\sigma) \to (X, g)$ is an isometric embedding. □

3.4 Quadratic differentials with poles of order 2

We prove Theorem 3D.

Harmonic diffeomorphisms

Definition 3.4.1. The *Fricke-Teichmüller space* is the subset of $B(\mathrm{PSL}_2(\mathbb{R}))$ consisting of classes of geometrically finite representations.

Each representation in this space is the holonomy of a geometrically finite hyperbolic structure on Σ. Fix an n-tuple $(\ell_1, \ldots, \ell_n) \in \mathbb{R}^n_{\geq 0}$.

Definition 3.4.2. Let $T(\Sigma, \ell_1, \ldots, \ell_n)$ be the subspace of the Fricke-Teichmüller space such that the convex core of the underlying surface associated to each representation has, for each ℓ_k, either a puncture if $\ell_k = 0$ or a closed geodesic boundary component of length $\ell_k \neq 0$.

When the context is clear we just call this the Teichmüller space. We represent points as equivalence classes $[S, f]$, where S is a surface and $f : \Sigma \to C(S)$ is a diffeomorphism onto the convex core of S. Another point $[S', f']$ is equivalent if $f^{-1} \circ f'$ is an isometry and isotopic to the identity.

Recall that there is a compatible Riemann surface structure on Σ. Around any puncture p we choose a local holomorphic coordinate z such that $z(p) = 0$. A meromorphic quadratic differential Φ with a pole of order 2 at such a puncture admits a Laurent expansion

$$(a_{-2}z^{-2} + a_{-1}z^{-1} + a_0 + \ldots)dz^2.$$

The a_{-2} term is invariant under holomorphic coordinate changes that take 0 to 0, and correspondingly we call it the *residue* of Φ at p

For ease of notation we assume $\ell_1, \ldots, \ell_{d_1} = 0$, $\ell_{d_1+1}, \ldots, \ell_{d_2} > 0$, $d_1 + d_2 = n$. For any $(d_2 - d_1)$-tuple of unit norm complex numbers $\theta_{d_1+1}, \ldots, \theta_{d_2}$ let P be the vector $(\ell_k, \theta_k)^n_{k=d_1+1}$.

Definition 3.4.3. $Q(\Sigma, P)$ is the space of meromorphic quadratic differentials on Σ with poles of order at most 2 at the p_k and residues

$$-\Lambda(\theta_k)\ell_k^2/16\pi^2.$$

If $\ell_k = 0$ we have a pole of order at most 1.

The space of holomorphic quadratic differentials $Q(\Sigma)$ with pole-type singularities at the cusps of Σ is a Fréchet space with seminorms coming from the restriction of the L^1 norm to pre-compact open sets. $Q(\Sigma, P)$ inherits the subspace topology from $Q(\Sigma)$. The following result can be deduced from the work of Wolf in [Wol91b]. It links the two spaces above.

Theorem 3.4.4. *(Wolf) For any $[S, f] \in T(\Sigma, \ell_1, \ldots, \ell_n)$ there is a unique harmonic diffeomorphism*

$$h_f : \Sigma \to C(S)$$

in the isotopy class such that Hopf$(h_f) \in Q(\Sigma, P)$.

This allows us to define a map

$$\Psi : T(\Sigma, \ell_1, \ldots, \ell_n) \to Q(\Sigma, P)$$

by $[S, f] \mapsto \text{Hopf}(h_f)$.

Remark 3.4.5. In [Wol91b], Wolf only explicitly computes and writes down the residue in the event $\theta = 0$, although he outlines constructions for $\theta \neq 0$. The values listed above can be computed by following the proof of Proposition 3.5.5 in the current chapter.

The content of Theorem 3D is that the map Ψ is a diffeomorphism. The first step is a dimension count.

Lemma 3.4.6. *$T(\Sigma, \ell_1, \ldots, \ell_n)$ and $Q(\Sigma, \ell_1, \ldots, \ell_n)$ are homeomorphic to $\mathbb{R}^{6g-6+2n}$.*

Proof. For the Teichmüller space, view the punctures as nodes and double the surface across the boundary. By mapping an element to this double, $T(\Sigma, \ell_1, \ldots, \ell_n)$ then embeds into a strata of the augmented Teichmüller space consisting of surfaces of genus $2g + d_2 - 1$ with d_1 nodes. Choosing a pants decomposition that includes all of our boundary curves and nodes (the nodes correspond to pinched curves) and taking the corresponding Fenchel-Nielsen coordinates shows this strata has dimension

$$12g + 6d_2 + 4d_1 - 12.$$

Every curve in the image of $T(\Sigma, \ell_1, \ldots, \ell_n)$ has an involutive symmetry across the boundary, and so the image is determined by at most $6g - 6 + 3d_2 + 2d_1$ coordinates. Fixing the lengths of the boundary curves kills another d_2 parameters and we obtain $6g - 6 + 2n$. On the other hand, the space of holomorphic quadratic differentials with poles of order bounded above by $k_1, \dots, k_n$ at $p_1, \dots, p_n$ forms a complex vector space and by Riemann-Roch it has real dimension

$$6g - 6 + 2\sum_j k_j.$$

Specifying the Laurent expansion at the poles then removes 2 parameters for each puncture and we end up with $6g - 6 + 2n$ degrees of freedom. □

For a closed arc c on a hyperbolic surface, let $\ell(c)$ denote the hyperbolic length of the geodesic representative. Below, the surface on which the curve lives will be clear.

Proof of Theorem 3D. By Brouwer's invariance of domain, it is enough to show Ψ is continuous, injective, and proper. Continuity and injectivity follow from arguments in [Wol91b, Section 4], so we only need properness. To this end, let $K \subset Q(\Sigma, P)$ be compact. Remove cusp neighbourhoods around all punctures, each one chosen small enough so that all simple closed geodesics of Σ are contained in the resulting subsurface, which we will call Σ'. By a estimate from [Wol89, Lemma 3.2]

$$E_{\Sigma'}(h_f) \leq 2\int_{\Sigma'} |\Phi| + \mathrm{Area}(h_f(\Sigma')) \leq 2\int_{\Sigma'} |\Phi| + \mathrm{Area}(h_f(\Sigma)).$$

The Gauss-Bonnet theorem yields

$$E_{\Sigma'}(h_f) \leq 2\int_{\Sigma'} |\Phi| - 2\pi\chi(\Sigma).$$

By a minor and well-understood modification of the proof of the Courant-Lebesque lemma [Jos84, Lemma 3.1], we obtain $\ell_{Y'}(h'_f(\gamma)) \leq A_{\mathcal{F}}$ for any finite collection $\mathcal{F}$ of simple closed geodesics inside S and any choice of representative pair $(Y', h'_f) \in [(Y', h'_f)] \in \psi^{-1}(K)$. Since the boundary lengths are fixed we have an upper bound on the lengths of finite collections of simple closed geodesics in all of any Y'. We argue that we also have a uniform lower bound on such lengths. On a complete finite volume hyperbolic surface, any essential simple closed geodesic δ is contained in an embedded annulus. This annulus has a horizontal coordinate specified by δ and an orthogonal vertical coordinate. Any simple closed geodesic δ' that transversely intersects δ once must pass through the entire vertical length of the annulus. If we have a geodesic δ such that $\ell(h_{f_k}(\delta))$ shrinks to 0 along some sequence (Y_k, h_{f_k}), select a curve δ' in Σ as above. From the collar lemma we see that $\ell(h_{f_k}(\delta')) \to \infty$

as $k \to \infty$. However, we can uniformly bound $\ell(h_{f_k}(\delta'))$ from above, so this is impossible.

Now, view the punctures on Σ as nodes and double across all punctures that "are opened" to get a noded surface Σ^d. Likewise double all surfaces $(Y, h_f) \in \Psi^{-1}(K)$ across the boundaries. h_f extends by reflection and we get a pair (Y^d, h_f^d). This provides a map

$$\iota : T(\Sigma, \ell_1, \ldots, \ell_n) \to \mathbf{T}_{2g+d_1-1, 2d_2}$$

that is a diffeomorphism onto its image. By [Ham03, Lemma 3.3] on any $S_{g,n}$ there is a collection of simple closed curves $\delta_1, \ldots, \delta_{6g-5+2n}$ so that the map

$$\mathcal{L}_{g,n} : \mathbf{T}_{g,n} \to \mathbb{R}^{6g-5+2n}$$

given by

$$[X, \phi] := \chi \mapsto (\ell_\chi(\delta_1), \ldots, \ell_\chi(\delta_{6g-5+2n}))$$

is a diffeomorphism onto its image. The composition $\mathcal{L}_{2g+d_1-1,2d_2} \circ \iota$ takes $\Psi^{-1}(K)$ into a compact set, and hence Ψ is proper. As discussed above, this completes the proof. □

3.5 Existence and classification of tame harmonic maps

Existence and behaviour

Once and for all, fix a complete finite volume hyperbolic surface $\Sigma = \tilde{\Sigma}/\Gamma$, a CAT(0) Hadamard manifold (X, g), and a reductive representation $\rho : \Gamma \to \mathrm{Isom}(X, g)$. We denote both the metric on Σ and its lift to $\tilde{\Sigma}$ by σ.

There is a finite set of punctures $p_1, \ldots, p_n$ with associated peripheral isometries $\gamma_1, \ldots, \gamma_n$ such that $\rho(\gamma_j)$ is hyperbolic. If this set is empty then ρ admits a finite energy equivariant map, for which the existence is already known. Hence we declare $n \geq 1$ and by an induction argument we may reduce to $n = 1$. Let us now fix some notation: set $\gamma := \gamma_1$ and write $\Sigma = \Sigma^c \cup C$, C is a cusp neighbourhood corresponding to γ and Σ^c is the complement. C is isometric to $U(\tau)$ for some $\tau > 0$ and we equip it with the relevant coordinates $x + iy$, $0 \leq x \leq \tau$, $y > a$. For $r > a$, C_r will be $\{x + iy \in C : a < y \leq r\}$ and $\Sigma_r = \Sigma^c \cup C_r$. Let D be some fixed fundamental domain with respect to the covering $\pi : \tilde{\Sigma} \to \Sigma$ and analogously we put $D^c = D \cap \pi^{-1}(\Sigma^c)$, $D' = D \cap \pi^{-1}(C)$, $D_r = D \cap \pi^{-1}(C_r) \cup D^c$. We set $i_r : \Sigma_r \to \Sigma$ to be the inclusion map. Finally, we use $\ell(\cdot)$ to denote the length of a rectifiable curve on a surface and hope there is no confusion with isometries.

Proposition 3.5.1. *Given the data Σ, X, ρ as above, there exists a ρ-equivariant harmonic map $f : \tilde{\Sigma} \to X$.*

Proof. Let $\alpha : [0, \tau] \to X$ be a constant speed curve with image in the axis of $\rho(\gamma)$ and so that $\alpha(\tau) = \rho(\gamma)\alpha(0)$. By [Cor92b] there exists a unique harmonic section s_r of the pullback bundle $i_r^* X \to \Sigma_r$ with boundary values α. Extend s_r to Σ via $s_r(x, t) = s_r(x)$. The s_r induce equivariant maps $f_r : \tilde{\Sigma} \to X$, that are harmonic on $\pi^{-1}(\Sigma_j)$. We prove the f_r converge along a subsequence in the C^∞ topology to an equivariant harmonic map.

Let φ be any non-harmonic equivariant map corresponding to a section of $i_0^* X \to \Sigma^c$ with boundary values α. As with f_r, define φ on the rest of D by $\varphi(x, t) = \varphi(x)$ and then extend equivariantly to $\tilde{\Sigma}$. Let β be the image of α on the geodesic axis of $\rho(\gamma)$ and set

$$\beta_t^r := f_r([0, \tau] \times \{t\}).$$

Notice that $|d\varphi|_{\sigma'} = \ell(\beta)/\tau$ on C since it has constant speed. For $r > s$,

$$2E_{C_r \backslash C_s}(\varphi) = \int_{C_r \backslash C_s} |d\varphi|_{\sigma'}^2 dv_{\sigma'} = (r - s)\ell(\beta)^2/\tau.$$

As β is a geodesic arc in a negatively curved space,

$$s\ell(\beta) \leq \int_0^s \ell(\beta_t^r) dt,$$

and hence for any $r > s$,

$$E_{C_r \backslash C_s}(\varphi) \leq \frac{1}{2}\int_s^r \ell(\beta_t^r)^2/\tau dt \leq \frac{1}{2}\int_s^r \Big(\int_{S^1 \times \{t\}} |df_r| d\theta\Big)^2 \tau^{-1} dt \leq E_{C_r \backslash C_s}(f_r).$$

From the non-positive curvature hypothesis f_r minimizes energy among maps to X with the same equivariant boundary values. In particular,

$$E_{\Sigma_r}(f_r) \leq E_{\Sigma_r}(\varphi),$$

and moreover

$$E_{\Sigma_s}(f_r) = E_{\Sigma_r}(f_r) - E_{\Sigma_r \backslash \Sigma_s}(f_r) \leq E_{\Sigma_r}(\varphi) - E_{\Sigma_r \backslash \Sigma_s}(\varphi) = E_{\Sigma_s}(\varphi).$$

By a classical PDE estimate (say, from [SY97, page 171]),

$$\sup_{D_s} e(f_r) = \sup_{\Sigma_s} e(f_r) \leq A_s E_{\Sigma_{s+1}}(f_r) \leq A_s E_{\Sigma_{s+1}}(\varphi),$$

where A_s depends on the Ricci curvature of Σ_{s+1}, the injectivity radius on Σ_s, and $\text{dist}(\partial\Sigma_s, \partial\Sigma_{s+1})$. Since ρ is acting by isometries we get the same bound in all of $\pi^{-1}(\Sigma_r)$. Next, we claim there is a compact set $O_s \subset X$ such that

$$f_r(D_s) \subset O_s$$

for all r. Appealing to the energy density bound above, it is enough to show that for a fixed point $x_0 \in D_s$, $f_r(x_0)$ stays within some compact set as $r \to \infty$. We find it convenient from here to split cases. Firstly, let us assume that the image of ρ does not lie in a parabolic subgroup. Let ξ be a point in the boundary at infinity $\partial_\infty X$. There is loop $\gamma : [0, L] \to \Sigma$ parametrized by arclength such that

$$\rho(\gamma)(\xi) \neq \xi.$$

Choose ℓ so that the image of γ under π lies entirely in Σ_ℓ and let A_ℓ be a uniform bound on the derivative in $\pi^{-1}(\Sigma_\ell)$. We then have, for $r > \ell$,

$$d_g(\rho(\gamma)f_r(x_0), f_r(x_0)) = d_g(f_r(\gamma(x_0)), f_r(x_0)) \leq A_\ell L.$$

This is because lifting γ to the universal cover gives a path between x_0 and $\gamma \cdot x_0$ that remains within lifts of Σ_ℓ. Choose a neighbourhood B of ξ in $X \cup \partial_\infty X$ such that

$$d_g(B \cap X, \rho(\gamma)B \cap X) > A_\ell L.$$

Then $f_r(x_0)$ cannot enter B, no matter how large r grows. Via compactness we find a finite number of neighbourhoods $(B_k)_k$ as above that cover the boundary sphere. Choosing $O_s := X \backslash (\cup_k N_k)$ the claim follows. Notice then that f_r takes any lift of Σ_s to a compact set:

$$f_r(\gamma D_s) \subset \rho(\gamma) O_s$$

for all r, s. It now follows by a well-known argument, namely an application of the Arzelà-Ascoli theorem and a bootstrap, that a subsequence of the $(f_r)_{r>0}$ converges uniformly on compact subsets of $\tilde{\Sigma}$ to a harmonic map f_∞. By equivariance of the f_r on $\pi^{-1}(\Sigma_r)$, f_∞ is necessarily equivariant.

We next treat the case where ρ stabilizes a totally geodesic flat F. F is a symmetric space and identifies isometrically as

$$G/H := (O(n) \rtimes \mathbb{R}^n)/O(n).$$

Fix two points $x_0 \in D_s$ and $y_0 \in F$ and for each r choose $g_r \in G$ such that $g_r f_r(x_0) = y_0$. We notice that for any $y \in F$ and $\gamma \in \Gamma$, $d(g_r \rho(\gamma) g_r^{-1} y, y)$ is uniformly bounded in r. Indeed,

$$\begin{aligned} d_g(g_r\rho(\gamma)g_r^{-1}y, y) &\le d_g(g_r\rho(\gamma)g_r^{-1}y, g_r\rho(\gamma)g_r^{-1}y_0) + d_g(g_r\rho(\gamma)g_r^{-1}y_0, y_0) + d_g(y, y_0) \\ &= 2d_g(y, y_0) + d_g(g_r\rho(\gamma)g_r^{-1}y_0, y_0) \\ &= 2d_g(y, y_0) + d_g(g_r\rho(\gamma)f_r(x_0), g_r f_r(x_0)) \\ &= 2d_g(y, y_0) + d_g(g_r f_r(\gamma \cdot x_0), g_r f_r(x_0)) \\ &= 2d_g(y, y_0) + d_g(f_r(\gamma \cdot x_0), f_r(x_0)), \end{aligned}$$

and we know f_r has a uniform energy density bound on $\rho(\Gamma) \cdot D_r$. By the argument of [JY91, Lemma 2] there is a sequence $(r_n)_{n=1}^{\infty}$ increasing to ∞ and an element $g_\infty \in G$ such that for every $\gamma \in \Gamma$ and $y \in F$,

$$\lim_{n\to\infty} g_{r_n}\rho(\gamma)g_{r_n}^{-1}y = g_\infty\rho(\gamma)g_\infty^{-1}y.$$

The orbit of the point x_0 under the family of maps $g_{r_n}f_{r_n}$ is a singleton, and by our uniform energy bound we see as above that there is a compact set O_s such that

$$g_{r_n}f_{r_n}(D_s) \subset O_s.$$

Arguing as above there is a subsequence along which $g_{r_n}f_{r_n}$ converges to a harmonic map f_∞. Note that $g_{r_n}f_{r_n}$ is $g_{r_n}\rho(\Gamma)g_{r_n}^{-1}$-equivariant, so that f_∞ is $g_\infty\rho(\Gamma)g_\infty^{-1}$-equivariant. Therefore, we may take $f := g_\infty^{-1}f_\infty$ as the sought harmonic map. □

We use the ideas above to build a family of harmonic maps, indexed by a real parameter $\theta \in \mathbb{R}$. We perform a *fractional Dehn twist* on each cylinder C. This is the map given in the cusp coordinates by

$$x + iy \mapsto x + \theta y + iy$$

on C and the identity map on the rest of Σ. Lift to a map d^θ on $\tilde{\Sigma}$. The lift commutes with the relevant parabolic isometry. Define f_r^θ to be the equivariant harmonic map on $\pi^{-1}(\Sigma_r)$ with the same equivariant boundary values as $\varphi \circ d^\theta|_{\partial D_r}$. Then extend to agree with $\varphi \circ d^\theta$ on the complement. The derivative matrix of d^θ is

$$\begin{pmatrix} 1 & \theta \\ 0 & 1 \end{pmatrix},$$

so that

$$||d(\varphi \circ d^\theta)|| \le ||d\varphi||(1 + \theta).$$

Thus on Σ_s,

$$E_{\Sigma_s}(f_r^\theta) \leq E_{\Sigma_s}(f_r)(1+\theta)^2.$$

By the argument of Proposition 3.5.1 there is a subsequence along which the f_r^θ's converge to a limiting harmonic map f^θ. Of course, $f = f^0$.

We keep the same characters α and β from the proof of the above proposition. Note $\ell(\beta) = \ell(\rho(\gamma))$. Define $\varphi^\theta := \varphi \circ d^\theta$. In local Euclidean coordinates, d^θ is harmonic on C. Since $\nabla d\varphi = 0$ on C, the composition is a harmonic map there (see [EL83, Proposition 2.20]).

Lemma 3.5.2. *The function $z \mapsto d(f^\theta, \varphi^\theta)(z)$ is uniformly bounded.*

Proof. Let $\psi_r := d(f_r^\theta, \varphi^\theta)$. By equivariance, each ψ_r descends to a function on Σ. $\psi_r = 0$ on $\Sigma \backslash \Sigma_r$, and since $\psi_r > 0$ at some point we know it attains a maximum at a point in the interior of Σ_r. As ψ_r is subharmonic on C_r, $\sup_{z \in C_r} \psi_r(z)$ occurs on $\partial \Sigma^c$ and moreover ψ_r is maximized at a point in Σ^c. Meanwhile,

$$\psi_r \to d^2(f^\theta, \varphi^\theta)$$

uniformly on compacta as $r \to \infty$. By smoothness, ψ is uniformly bounded on Σ^c. This implies we have a uniform bound on the ψ_r's inside Σ^c as $r \to \infty$. Since the relevant maximum is attained inside Σ^c, this bound holds everywhere. □

Let $\Phi := \text{Hopf}(f^\theta)$. The context is clear so we do not include a θ in our notation. By equivariance we can view Φ as a holomorphic quadratic differential on any quotient of $\tilde{\Sigma}$ by a subgroup of Γ.

Lemma 3.5.3. *Φ has a pole of order 2 at the cusp.*

Proof. From the infinite energy phenomena, Φ either has a pole of order at least 2 or an essential singularity. The $(2, 0)$ component of the pullback metric by φ^θ is a section of $\mathcal{K}^2$ that is holomorphic on C. We still denote it by $\text{Hopf}(\varphi^\theta)$.

We compute this differential in C. Choose a local orthonormal basis $\partial/\partial x$, $\partial/\partial y$ of the relevant tangent spaces so that $\partial\varphi^0/\partial y = 0$ always. Starting with $\theta = 0$, we know that in local coordinates

$$\text{Hopf}(\varphi^0)(z) = \frac{1}{4}\Big(|\partial\varphi^0/\partial x|^2 - |\partial\varphi^0/\partial y|^2 - 2i\langle \partial\varphi^0/\partial x, \partial\varphi^0/\partial y\rangle\Big)dz^2.$$

Since φ^0 is constant in the vertical direction

$$\mathrm{Hopf}(\varphi^0)(z) = \frac{1}{4}|\partial\varphi^0/\partial x|^2 dz^2 = \ell(\rho(\gamma))^2/4\tau^2 dz^2.$$

From the chain rule, $d\varphi^0$ and $d\varphi^\theta$ admit matrix representations with

$$d\varphi^0 = \begin{pmatrix} v & 0 \end{pmatrix}, \ d\varphi^\theta = \begin{pmatrix} v & \theta v \end{pmatrix},$$

where v is a $1 \times \dim X$ column vector. Thus,

$$\begin{aligned}\mathrm{Hopf}(\varphi^\theta)(z) &= \frac{1}{4}(|\partial\varphi^\theta/\partial x|^2 - |\partial\varphi^\theta/\partial y|^2 - 2i\langle\partial\varphi^\theta/\partial x, \partial\varphi^\theta/\partial y\rangle)dz^2 \\ &= \frac{1}{4}(1-\theta^2 - i2\theta)|\partial\varphi^0/\partial x|^2 dz^2.\end{aligned}$$

We take the strip conformally to a punctured disk via

$$z \mapsto \zeta(z) = e^{\frac{2\pi i z}{\tau}},$$

taking the point at ∞ to 0. The transformation law multiplies by $-\zeta^{-2}\tau^2/4\pi^2$, and we see that we have a pole of order 2 with residue

$$-\Lambda(\theta)\ell(\rho(\gamma))^2/16\pi^2.$$

We now compare Φ to $\mathrm{Hopf}(\varphi^\theta)$. As φ^θ has rank 1, the formula $J = H - L$ implies

$$H(\varphi^\theta)^{1/2} = L(\varphi^\theta)^{1/2} = \frac{1}{2}e(\varphi^\theta)^{1/2},$$

so that $\mathrm{Hopf}(\varphi^\theta) = \sigma H(\varphi^\theta)^{1/2}L(\varphi^\theta)^{1/2} = \sigma e(\varphi^\theta)/4$. From Young's inequlaity,

$$||\Phi|| = \sigma H(f)^{1/2}L(f)^{1/2} \leq \frac{1}{2}\sigma e(f^\theta),$$

and hence it is enough to bound $e(f^\theta)$ by a sublinear function of $e(\varphi^\theta)$. This is not hard: for any $x_0 \in \tilde{\Sigma}$, $r_0 > 0$, and $y \in B(x_0, r_0)$,

$$\begin{aligned}d(f^\theta(x_0), f^\theta(y)) &\leq d(f^\theta(x_0), \varphi^\theta(x_0)) + d(f^\theta(y), \varphi^\theta(y)) + d(\varphi^\theta(x_0), \varphi^\theta(y)) \\ &\leq A + \sup_{B(x_0,r_0)} ||d\varphi^\theta|| d(x_0, y).\end{aligned}$$

Working in the flat cylinder metric, Cheng's lemma then gives

$$||df||(x_0) \lesssim \frac{1+r_0}{r_0}(1 + \sup ||d\varphi^\theta|| r_0).$$

In a cusp neighbourhood, the injectivity radius of the flat cylinder metric is uniformly bounded below, and hence we may choose r_0 uniformly bounded below. Squaring for the energy density gives the desired bound. □

Henceforth, we assume that X is CAT(-1). By equivariance, f^θ and φ^θ induce quotient maps

$$f_\gamma, \varphi_\gamma : \tilde{\Sigma}/\langle\gamma\rangle \to X/\langle\rho(\gamma)\rangle.$$

We suppress the θ from our notation for convenience. β projects in the quotient to a core geodesic $\overline{\beta}$. From the CAT(-1) hypothesis, this is the unique geodesic in the homotopy class. Any $D_r/\langle\gamma\rangle$ identifies isometrically with the cylinder

$$\{(x, y) = x + iy : 0 \leq x \leq \tau, a \leq y \leq r\}$$

with the usual identification.

Lemma 3.5.4. *There is a translation $\tilde{R}$ of the geodesic axis of $\rho(\gamma)$ such that the map $\Sigma \ni z \mapsto d(f^\theta, \tilde{R} \circ \varphi^\theta)(z)$ tends to 0 as we move into the puncture.*

Proof. We define C_∞ to be the infinite cylinder

$$\{(x, t) \in [0, 1] \times (-\infty, \infty) : (0, t) \sim (1, t)\}$$

with the flat metric. Let $b_s : C_\infty \to D/\langle\gamma\rangle$ be the map given by

$$\begin{cases} (x, t) \mapsto (x, s) & -\infty \leq t \leq -s \\ (x, t) \mapsto (x, 2s + t) & -s \leq t \leq s \\ (x, t) \mapsto (x, 3s) & s \leq t \leq \infty. \end{cases}$$

Then set $B_s := f_\gamma \circ b_s$ and $\varphi_s := \varphi_\gamma \circ b_s$. Both B_s and φ_s are harmonic on $-s \leq t \leq s$ because b_s is conformal there. From Lemma 3.5.2 the orbit of any point under B_s remains in a compact set as $s \to \infty$. The uniform energy bounds from Lemma 3.5.3 permit us to construct a subsequence along which both B_s and φ_s converge in the C^∞ topology to harmonic maps f_∞ and φ_∞ respectively.

Let h denote the harmonic diffeomorphism of the disk whose Hopf differential is Φ. By [Wol91b, Lemma 3.6], the Jacobian $J(h) = H(h) - L(h)$ tends to 0 as we approach the puncture. From Proposition 3.3.2, $J(f) \to 0$ as well. Therefore, $J(f_\infty) = 0$ and necessarily rank$df_\infty \leq 1$ at each point. By equivariance this is rank 1 in an open set, and by [Sam78, Theorem 3] the image is contained in a geodesic arc. Again by equivariance, the image must then be a closed geodesic arc. There is only one such arc in the quotient, and hence f_∞ maps onto the core geodesic. Lifting f_∞ and φ_∞ to maps from $\mathbb{R}^2$ to the axis of $\rho(\gamma)$, f_∞ and φ differ by a translation along $\overline{\beta}$. One can justify that last claim by observing that their distance function is a bounded subharmonic function on $\mathbb{R}^2$—hence a constant—and then following the proof of Lemma 3.2.12. Lifting back to $\tilde{\Sigma}$ this means there is a translation $\tilde{R}$ of the geodesic axis such that for any $r > 0$,

$$d(f^\theta(x, s_m + 2t), \tilde{R} \circ \varphi^\theta(x, s_m + 2t)) = d(b_{s_m}(x,t), R \circ \varphi_{s_m}(x,t)) \longrightarrow 0$$

as $m \longrightarrow \infty$ for $-r \leq t \leq r$. In particular, the quantities $d(f_\gamma(x, s_m), R \circ \varphi_\gamma(x, s_m))$ and $d(f_\gamma(x, s_{m+1}), R \circ \varphi_\gamma(x, s_{m+1}))$ are very close to 0. Since the relevant distance function is subharmonic, its maximum on

$$\{(x,t) \in C_\infty : s_m \leq t \leq s_{m+1}\}$$

is achieved on the boundary. It follows that

$$d(f_\gamma(x,t), R \circ \varphi_\gamma(x,t)) \longrightarrow 0$$

as $t \longrightarrow \infty$. Returning to the universal cover, we conclude that

$$d(f^\theta(z), \tilde{R} \circ \varphi^\theta(z)) \longrightarrow 0$$

as we move toward the puncture. □

Proposition 3.5.5. Φ *has a pole of order* 2 *at the cusp with residue* $-\Lambda(\theta)\ell(\rho(\gamma))^2/16\pi^2$.

Proof. The lemma above shows

$$\lim_{s \to \infty} B_s = R \circ \varphi_\gamma$$

in the C^0 topology, and along a subsequence in the C^∞ topology. We prove there is no need to pass to a subsequence. Indeed, if we don't have C^1 convergence we can pick a subsequence along which our maps are uniformly far from f_∞ in the C^1 norm. One can then use the argument above to pass to a subsequence that converges in the C^∞ sense to $S \circ \varphi_\gamma$ for some other rotation S. C^0 convergence to $R \circ \varphi$ forces $S = R$, which is a contradiction. Continuing inductively gives C^k convergence for any k. The Hopf differential of f then converges to Hopf(φ^θ) as we move into the puncture. The result now follows from the computation in Lemma 3.5.3. □

Uniqueness

Let f_1 and f_2 be two harmonic maps whose Hopf differentials have second order poles and such that the residues have the same complex argument $\nu \in (-\pi, \pi)$.

Lemma 3.5.6. *There exists an $A_k > 0$ such that as $y \to \infty$, the image of f_k remains in an A_k-neighbourhood of the geodesic axis of $\rho(\gamma)$.*

Proof. Let β_y^k be the curve $f_k([0,\tau] \times \{y\})$ in the usual coordinates. From Proposition 3.3.2 and [Wol91b, page 516], the energy density of f_k is uniformly bounded on Σ in the flat-cylinder metric. This implies

$$\ell(\beta_y^k) \leq A$$

for all $y > 0$. We argue each β_y^k becomes trapped close to the geodesic as $y \to \infty$. If not, there is a subsequence s_j tending to ∞ and points $f_k(z_j) \in \beta_{s_j}^k$ such that the closest-point projection onto the geodesic, say y_j, satisfies

$$d(f_k(z_j), y_j) \to \infty.$$

Then

$$\ell(\beta_{s_j}^k) \geq d(f_k(z_j), f_k(\gamma \cdot z_j)) = d(f_k(z_j), \rho(\gamma) f_k(z_j)).$$

The right most term blows up as $j \to \infty$, and this is a clear contradiction. To verify that last statement, note $f_k(z_j)$ accumulates along a subsequence to a point $\xi \in \partial_\infty X$, and since the distance from $\beta_{s_j}^k$ to the geodesic is uniformly bounded below, this is not an endpoint of the geodesic. In particular, the extension of $\rho(\gamma)$ to $\partial_\infty X$ does not fix ξ, and hence if $B_{s_j}^k$ is a neighbourhood of ξ in $X \cup \partial_\infty X$,

$$d(B_{s_j}^k \cap X, \rho(\gamma) B_{s_j}^k \cap X) \to \infty$$

as $j \to \infty$. □

Recall the cylinder C_∞. Let b_s^k be the map $b_s \circ f_\gamma^k : C_\infty \to X/\langle \rho(\gamma) \rangle$. Since the energy is controlled and it stays close to the geodesic, b_s^k converges along a subsequence to a harmonic map f_∞^k. By the same argument as in the previous subsection f_∞^k has image in a geodesic and from equivariance this must be the core geodesic $\overline{\beta}$. One can slightly modify an argument as in the previous subsection to check that a_s^k limits to f_∞^k along the whole sequence in the C^∞ topology. Moreover f_k limits onto the geodesic β as we go further into the cusp.

Lemma 3.5.7. *The residue of f_1 and f_2 is the same.*

Proof. Let $\Phi_k := \text{Hopf}(f_k)$. In the computations to follow, we use the flat-cylinder metric on Σ. Let $\gamma_y(x)$ be the curve $x \mapsto x + iy$. From the discussion above, the length of the core geodesic in $X/\langle\rho(\gamma)\rangle$ is

$$\lim_{y\to\infty} \ell_g(f_k(\gamma_y)).$$

There are differentials Φ'_k such that

$$\Phi_k = e^{i\nu}\Phi'_k.$$

That is, a differential that differs from Φ_k by a rotation and whose residue at the cusp is real. The pullback metrics can thus be written

$$f_k^* g = e(f_k)\sigma' dz d\overline{z} + e^{i\nu}\Phi'_k + e^{-i\nu}\overline{\Phi'_k} = e(f_k)\sigma' dz d\overline{z} + 2\Re e^{i\nu}\Phi'_k.$$

Writing $\Phi'_k = \phi'_k(z)dz^2$ in a local coordinate we know that in the cylinder

$$|\phi'_k| = H(f_k)^{1/2}L(f_k)^{1/2} = H(f_k)\cdot\frac{L(f_k)^{1/2}}{H(f_k)^{1/2}}.$$

From [Wol91b, Proposition 3.8], in the strip we can write

$$\Phi_k = \left(e^{i\nu}a^k_{-2} + e^{i\nu}O(e^{-Ay})\right)dz^2,$$

where $a^k_{-2} > 0$. From Proposition 3.3.2 and [Wol91b, Lemma 3.6], we also know

$$\frac{L(f_k)}{H(f_k)} \to 1$$

as we move into the puncture. The length of the core geodesic is therefore

$$\begin{aligned}
\lim_{y\to\infty} \ell_g(f_k(\gamma_y)) &= \lim_{y\to\infty}\int_0^\tau ||\dot{\gamma_y}(x)||_{f_k^* g}dx \\
&= \lim_{y\to\infty}\int_0^\tau \sqrt{e(f_k)\sigma' + 2\Re e^{i\nu}\phi'}dx \\
&= \lim_{y\to\infty}\int_0^\tau \sqrt{H(f_k)(1 + L(f_k)/H(f_k)) + 2\Re e^{i\nu}\phi'}dx \\
&= \tau\sqrt{2|a^k_{-2}|(1+\cos\nu)}
\end{aligned}$$

by the dominated convergence theorem. Meanwhile, passing to the quotient $\mathbb{H}/\langle\gamma\rangle$ we know the core geodesic has length $\ell(\rho(\gamma))$. We deduce

$$\ell(\rho(\gamma) = \tau\sqrt{2|a^k_{-2}|(1+\cos\nu)}.$$

Since ν is fixed, $|a^k_{-2}|$ does not depend on k. □

Henceforth put $a_{-2} = a^k_{-2}$ $(k = 1, 2)$.

Remark 3.5.8. From above we see that the complex argument v is related to the twist angle θ from the previous subsection by

$$\theta = \frac{-\sin v}{1+\cos v}.$$

Lemma 3.5.9. *The distance function $z \mapsto d(f_1, f_2)(z)$ is bounded.*

Proof. It suffices to bound $d(f_\infty^1, f_\infty^2)$ as then it is constant and we can lift to the universal cover. By [Wol91b, Proposition 3.8] we can express

$$\Phi_k = \Big(a_{-2}e^{iv} + e^{iv}O(e^{-Ay})\Big)dz^2$$

in the cylinder coordinates, where a_{-2} is real. Thus, upon taking $s \to \infty$, the Hopf differential of f_k^∞ is $a_{-2}e^{iv}dz^2$. That is, the Hopf differentials of f_1^∞ and f_2^∞ agree. We denote this differential by Φ_0, and highlight that the Φ_0-metric is nonsingular. Set

$$w_0(f_k) = \frac{1}{2}\log H_0(f_k)(z) - \frac{1}{2}\log|\Phi_0(z)|.$$

Here H_0 denotes the holomorphic energy in the Φ_0-metric, and analogously for the other quantities. From above it is clear that $J_0(f_k) = 0$ so $H_0(f_k) = L_0(f_k)$. From $|\Phi_0| = H_0(f_k)^{1/2}L_0(f_k)^{1/2}$ we see $w_0(f_k) = 0$. One can compute $e_0 = 2\cosh 2(w_0(f_k))$. In a coordinate $z = x+iy$ such that $\Phi_0 = dz^2$,

$$f_k^*g = (e_0+2)dx^2 + (e_0-2)dy^2 = 2dx^2.$$

Let γ_h and γ_v be horizontal and vertical curves for the Φ_0-metric. Explicitly, we mean the tangent vectors for γ_h, γ_v always evaluate under Φ_0 to positive and negative numbers respectively. Then,

$$\ell(f_k(\gamma_h)) = \int_{\gamma_h}\sqrt{e_0+2}dx\ ,\ \ell(f_k(\gamma_v)) = \int_{\gamma_h}\sqrt{e_0-2}dy$$

and we see

$$\ell(f_k(\gamma_h)) = 2\ell(\gamma_h)\ ,\ \ell(f_k(\gamma_v)) = 0.$$

Therefore, if v_a is the tangent vector to the geodesic at a point a then for all points z, $(df_k)_z(\partial_x) = 2v_{f_k(z)}$ and $(df_k)_z(\partial_y) = 0$. In particular, f_k is a constant speed map onto the geodesic in the horizontal direction and constant in the vertical direction. Any two such maps differ by a translation. This establishes the result. □

We apply Lemma 3.2.12 to obtain the uniqueness portion of Theorem 3A. If $f_1 \neq f_2$, which is only possible if ρ stabilizes a geodesic, then f_2 may be obtained from f_1 by precomposing with a lift of the translation found in Lemma 3.5.9. The results in this section constitute the proof of Theorem 3A.

3.6 Domination and AdS 3-manifolds

Non-reductive representations

When ρ is not reductive we can still produce a harmonic map that will be relevant to the domination problem. The content of the following exposition is contained in [DT16] and [GK17]. Assume ρ fixes a point ξ on $\partial_\infty X$. Given any geodesic ray $\eta : [0,\infty) \to X$ with an endpoint on $\partial_\infty X$, the *Busemann function* $\beta_\eta : X \to \mathbb{R}$ is defined by

$$\beta_\eta(x) = \lim_{t\to\infty}(d(\eta(t),x) - t).$$

The fact that this is well-defined and continuous is standard [BH99]. Now assume the endpoint is ξ. For any isometry γ with $\gamma\cdot\xi = \xi$ there is a $m(\gamma)\in\mathbb{R}$ such that

$$\beta_\eta(\gamma\cdot x) = \beta_\eta(x) + m(\gamma)$$

and $|m(\gamma)| = \ell(\gamma)$. It is easy to see the function $m\circ\rho : \Gamma\to\mathbb{R}$ is a group homomorphism. Let $\tilde{\eta}$ be any biinfinite oriented geodesic in $\mathbb{H}$ and let ρ^{red} be the representation $\Gamma\to \mathrm{PSL}_2(\mathbb{R})$ that acts by translations along $\tilde{\eta}$ with lengths $m\circ\rho$, with signs chosen according to the orientation. Since ρ^{red} stabilizes a geodesic there is a family of equivariant harmonic maps as in Theorem 3A. By construction, for all $\gamma'\in\Gamma$,

$$\ell(\rho^{red}(\gamma')) = \ell(\gamma').$$

Consequently, the problem of dominating ρ in length spectrum is equivalent to dominating ρ^{red} in length spectrum. Henceforth if ρ is not reductive we replace it with ρ^{red}.

Digression: elliptic monodromy

Looking toward domination, it is necessary to understand the behaviour of a harmonic map f when ρ has elliptic monodromy. In the event ρ has hyperbolic monodromy, the choice of parameter θ will have no effect here, so we assume $\theta = 0$. Let ξ be the point on $\partial_\infty\tilde{\Sigma}$ associated to the horocycle for γ and let F be the set of points in X fixed by $\rho(\gamma)$.

Proposition 3.6.1. *In the setting above, as $z\to\xi$ the function f limits to an element of F. Furthermore $e(f)(z)\to 0$.*

Proof. Let B be a relevant horoball for γ in the universal cover. By adapting a procedure from [GK17, Proposition 4.16], we first show that for any choice of $\delta > 0$ and ρ-equivariant map w that has a uniform Lipschitz constant in B there is a ρ-equivariant map w_δ such that

- $w_\delta = w$ on $\tilde{\Sigma}\backslash(\Gamma \cdot B)$,
- $d(w_\delta(p), w_\delta(q)) \leq d(w(p), w(q))$ for all points $p, q \in B$, and
- there is a smaller horoball $B' \subset B$ such that $f_\delta(B')$ is contained in the intersection of the convex hull of $f(B')$ and a ball of radius δ.

Towards this let $\mathcal{D}$ be a fundamental domain for the image of ∂B in the quotient and let $p \in \mathcal{D}$. Let π_t be the closest point projection from B onto the closed horoball of distance $t > 0$ from ∂B and put $p_t = \pi_t(p)$. Note the map $t \mapsto \pi_t(p)$ is nothing more than the transverse horospherical flow for $(B, \partial B, \xi)$. By hyperbolic trigonometry (see [GK17, Appendix A]),

$$d(p_t, \gamma \cdot p_t) \to 0$$

as $t \to \infty$. We next find fundamental domains $\mathcal{D}_t$ of $\pi_t(\partial B)$ containing p_t such that diam$\mathcal{D}_t \to 0$ as $t \to \infty$. By the Lipschitz condition

$$d(w(p_t), \rho(\gamma) \cdot w(p_t)) \to 0$$

and diam$f(\mathcal{D}_t) \to 0$ as $t \to \infty$.

Next, there is an $\epsilon(\delta) > 0$ such that if $d(x, \rho(\gamma) \cdot x) < \epsilon(\delta)$ then

$$d(x, F) < \delta/2.$$

In particular, for t large enough there is a $q_t \in F$ such that $d(w(p_t), q_t) < \delta/2$ and diam$w(\pi_t(\mathcal{D})) < \delta/2$. This implies the $\langle\rho(\gamma)\rangle$-invariant set $w(\pi_t(\partial B))$ is contained in $B(\delta, q_t)$. If π_{B_δ} is the closest point projection onto this ball, then take w_δ to be the ρ-eqiuivariant map that coincides with w on $\tilde{\Sigma}\backslash(\Gamma \cdot \pi_t(B))$ and with $\pi_{B_\delta} \circ w$ on $\Gamma \cdot \pi_t(B)$. This has all of the required properties.

With that cleared up recall that the total energy of f is finite in B, so that Hopf(f) has a pole of order at most 1. By Proposition 3.3.2 and [Wol91b, Proposition 3.13],

$$e(f) \leq A$$

in B. From here we make the assumption that Σ has at least two punctures, around one of which ρ has hyperbolic monodromy. This is the most complicated situation and the other cases are resolved similarly. Returning to the sequence f_r from the proof of Proposition 3.5.1, we may enlarge A if necessary to obtain

$$e(f_r) \leq A$$

for all r. This guarantees a uniform Lipschitz constant across all f_r. Next consider the maps $f_{r,\delta}$. We underline that they agree with f_r on $\partial\Sigma^c$ (here we are using the notations and conventions of Proposition 3.5.1). By definition

$$e(f_{r,\delta}) \leq e(f_r)$$

everywhere, so that

$$E_{\Sigma^c}(f_{r,\delta}) \leq E_{\Sigma^c}(f_r).$$

By the energy minimizing property of harmonic maps this forces $f_{r,\delta}$ to be harmonic. From the finite energy theory, if ρ does not fix a point on the boundary then $f_r = f_{r,\delta}$. If ρ does fix such a point then f_r and $f_{r,\delta}$ differ by a translation along a geodesic. Since they are set to be equal on $\partial\Sigma^c$ they agree everywhere. Taking $r \to \infty$ implies f has the listed properties of f_δ.

Now we put $\delta_n = 2^{-n}$ and iterate the procedure above. We obtain a sequence of horoballs tending to ξ whose image under f is contained in a closed ball of radius δ_n that intersects F non-trivially. Taking $n \to \infty$ the first result follows.

To see that the energy decays to zero, we argue by contradiction: suppose there is a $\delta_0 > 0$ and sequence z_n tending to ξ such that $e(f)(z_n) \geq \delta_0$. Consider the cylinder

$$C = \{(x,t) \in [0,1] \times [0,1] : (0,t) \sim (1,t)\}$$

and take a sequence of conformal embeddings $b_n : C \to D/\langle\gamma\rangle$ such that the projection of z_n is contained in $\mathrm{int}(b_n(C))$. The energy density of $B_n := f_\gamma \circ b_n$ is uniformly bounded, and we can choose b_n so that $e(B_n)(x,t) = e(f_\gamma)(b_n(x,t))$. By the usual argument, B_n subconverges in the C^∞ sense to a harmonic map $B_\infty : C \to X/\langle\rho(\gamma)\rangle$. From the first result, we see B_∞ is constant, which forces a contradiction in that $e(f)(z_n)$ must then tend to 0. □

Proof of domination theorem

Take any ρ-equivariant harmonic map $(\tilde{\Sigma},\sigma) \to X$ produced by Theorem 3A or a map $\tilde{\Sigma} \to \mathbb{H}$ from the previous subsection and call it f. Let Φ denote the Hopf differential. By Theorem 3D there is a surface (N,σ_0) with cusps and infinite funnels attached along closed geodesics as well as a harmonic map h with Hopf differential Φ taking Σ diffeomorphically onto the interior of the convex core N. Lift h to a map between the universal covers, that we will still denote h. We will always identify $\tilde{\Sigma}$ and the universal cover of N with $\mathbb{H}$. Let j denote the holonomy of $h^*\sigma$. Proposition 3.3.2 implies that

$$\psi := f \circ h^{-1} : (\tilde{C}(\mathbb{H}/j(\Gamma)),\sigma_0) \to (X,g)$$

is (j, ρ)-equivariant and 1-Lipschitz. Indeed, $h^*\sigma \geq f^*g$ in the sense that

$$(h^*\sigma)_z(v, v) \geq (f^*g)_z(v, v)$$

for all points z and non-zero vectors $v \in T_z\tilde{C}(\mathbb{H}/j(\Gamma))$. Hence for any two points $x, y \in \tilde{C}(\mathbb{H}/j(\Gamma))$ and path c from x to y,

$$\ell_g(\psi(c)) = \ell_{f^*g}(h^{-1}(c)) \leq \ell_{h^*\sigma}(h^{-1}(c)) = \ell_\sigma(c).$$

Now, extend ψ to the lift of the boundary of the convex core via uniform continuity, and precompose ψ with the (j, j)-equivariant 1-Lipschitz nearest point projection from $\mathbb{H}$ to $\tilde{C}(\mathbb{H}/j(\Gamma))$. This resulting map from $\mathbb{H} \to X$ is 1-Lipschitz and (j, ρ)-equivariant. This chosen j is j_Σ from the statement of Theorem 3B.

Remark 3.6.2. For this construction, it does not matter which initial harmonic map f we choose. Going forward we work with $\theta = 0$.

Lemma 3.6.3. *In general j_Σ does not strictly dominate ρ. If $X = \mathbb{H}$, a necessary and sufficient condition is that the image of any peripheral isometry under ρ is elliptic. In the general case, a sufficient condition is that*

$$\limsup_{m\to\infty} \frac{d(\psi(p), \rho(\gamma^m)\psi(p))}{2\log m} < 1.$$

This will be achieved if ρ has no hyperbolic or parabolic monodromy, but is still possible with parabolic monodromy.

Proof. If γ is a peripheral isometry such that $\rho(\gamma)$ is hyperbolic, then the translation length is fully encoded by the residue of the Hopf differential, so that $\ell(j_\Sigma(\gamma)) = \ell(\rho(\gamma))$. Hence 1 is the optimal Lipschitz constant in this setting. If $j_\Sigma(\gamma)$ is parabolic, then by elementary hyperbolic trigonometry (see [GK17, Lemma 2.7]),

$$\ell(j_\Sigma(\gamma^m)) = 2\log m + A.$$

If $\rho(\gamma)$ is parabolic then

$$\ell(\rho(\gamma^m)) \leq 2\log m + A.$$

This follows from [HI77, Theorem 1] and a minor modification of the argument in [GK17, Lemma 2.7]. We have equality above if all sectional curvatures of X are -1, which implies ρ cannot have parabolic monodromy if $X = \mathbb{H}$. When the image is elliptic we have shown $e(f) \to 0$ at the cusp. From [Wol91b, Proposition 3.13] we have $e(h) \to 1$ at such a cusp. This handles the case $X = \mathbb{H}$.

Working with arbitrary X, from the proof of Proposition 3.3.2 we know ψ either has Lipschitz constant 1 everywhere or the Lipschitz constant is strictly less than 1 on every compact set. By equivariance of ψ this implies the condition

$$\limsup_{m\to\infty} \frac{d(\psi(p), \rho(\gamma^m)\psi(p))}{2\log m} < 1$$

is sufficient. To see the last statement, simply fix $\kappa < -1$ and consider a rescaled copy of $\mathbb{H}$ with a metric of constant curvature κ. □

With j_Σ in hand, we perturb it to a convex cocompact representation that strictly dominates ρ. Let us first assume j_Σ is convex cocompact. We use the strip deformations of Thurston, which we now describe. Recall that an *arc* of a complete hyperbolic surface S is any non-trivial isotopy class of complete curves in S such that both ends exit into an infinite funnel. A *geodesic arc* is the geodesic representative of an arc.

Definition 3.6.4. A *strip deformation* of a hyperbolic surface S along a geodesic arc α is the new surface obtained by cutting along α and gluing a *strip*, the region on $\mathbb{H}$ bounded between two ultraparallel geodesics. The strip is inserted without shearing: so that the two endpoints of the most narrow cross section are identified to a single point z, which is called the *waist*. A strip deformation along a collection of pairwise disjoint and non-isotopic geodesic arcs $\alpha_1, \ldots, \alpha_n$ is the hyperbolic surface produced by performing strip deformations along α_k iteratively.

Note that strip deformations along geodesic arcs commute because the curves are disjoint. It was observed by Thurston [Thu98] and proved in full detail in [PT10] that as soon as a collection of pairwise disjoint non-isotopic arcs decomposes S into disks, the corresponding strip deformation uniformly lengthens all closed geodesics. Any geodesic arc α has two parameters associated to a strip deformation, namely the *waist* and the *width*: the thickness of the strip at its most narrow cross section. The lemma below follows from [DGK16a, Theorem 1.8].

Lemma 3.6.5. *For any choice of geodesic arcs $(\alpha_1, \ldots, \alpha_n)$ that decompose $\mathbb{H}/j_\Sigma(\Gamma)$ into disks, as well as waist and width parameters z_k, w_k, the holonomy of the corresponding strip deformation strictly dominates j_Σ.*

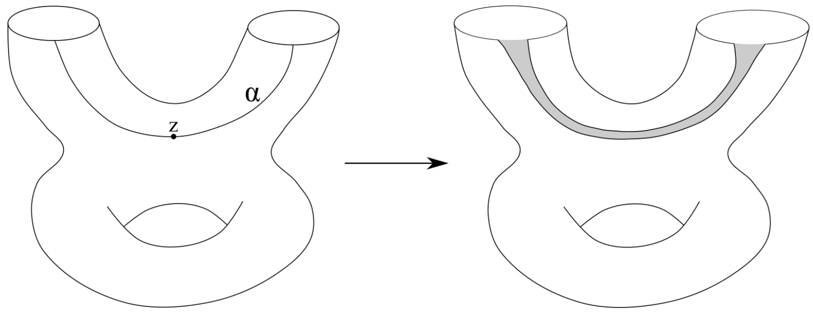

Figure 3.1: A strip deformation along a geodesic arc on a two-holed torus

Remark 3.6.6. The complex of *arc systems* $\mathcal{S}$ is the subcomplex of the arc complex obtained by removing all cells that do not divide S into disks. Danciger, Guéritaud, and Kassel established a homeomorphism between an abstract cone over $\mathcal{S}$ and the subspace of the Fricke-Teichmüller space of representations strictly dominating a convex cocompact representation. See [DGK16a] for the full description.

This solves the convex cocompact case. Without this condition we proceed as follows. For each puncture p_k in $\mathbb{H}/j_\Sigma(\Gamma)$ select disjoint biinfinite geodesic arcs α_k such that both ends of each α_k escape toward p_k. Arbitrarily choose points on α_k as a waist parameter and pick some positive width parameters. Insert a hyperbolic strip without shearing, exactly as one would do for a convex cocompact surface. The resulting surface admits a complete hyperbolic metric of infinite area, and therefore its holonomy is convex cocompact. The length spectrum of this new holonomy obviously dominates that of j_Σ. Then perform a strip deformation on the new surface to obtain a representation that strictly dominates j_Σ in length spectrum. By Theorem 3.2.2, in this context length spectrum domination implies domination in the regular sense.

This completes the proof of Theorem 3B for complete finite volume hyperbolic manifolds. The general case is now a consequence of Lemma 3.2.4.

Chapter 4

MAXIMAL SURFACES AND ADS 3-MANIFOLDS

4.1 Introduction

Near the end of the original paper [Sag19], we found something curious: pairs of representations ρ_1, ρ_2 that give rise to circle bundles with an anti-de Sitter structure that do not come from properly discontinuous actions on all of AdS^3. At the time we did not put much emphasis on the result; in fact, it's buried near the end of the paper as Proposition 7.7. The motivation for the next work is to explore representations such as (ρ_1, ρ_2) in more depth. The representations all satisfy a geometric condition, which we call almost strict domination. Below, let (X, ν) be a Hadamard manifold with isometry group G.

Definition 4.1.1. Let $\rho_1 : \pi_1(S_{g,n}) \to \mathrm{PSL}(2,\mathbb{R}), \rho_2 : \pi_1(S_{g,n}) \to G$ be two representations with ρ_1 Fuchsian. We say that ρ_1 almost strictly dominates ρ_2 if

1. for every peripheral $\zeta \in \pi_1(S_{g,n})$, $\ell(\rho_1(\zeta)) = \ell(\rho_2(\zeta))$, and
2. there exists a (ρ_1, ρ_2)-equivariant 1-Lipschitz map g defined on the convex hull of the limit set of $\rho_1(\pi_1(S_{g,n}))$ in $\mathbb{H}$ such that the local Lipschitz constants are < 1 inside the convex hull, and for peripherals ζ such that $\rho_1(\zeta)$ is hyperbolic, g takes each boundary geodesic axis for $\rho_1(\zeta)$ isometrically to a geodesic axis for $\rho_2(\zeta)$.

In the definition above, the global Lipschitz constant is

$$\mathrm{Lip}(g) = \sup_{y_1 \neq y_2} \frac{d_\nu(g(y_1), g(y_2))}{d_\sigma(y_1, y_2)},$$

where σ is the hyperbolic metric. The local one is

$$\mathrm{Lip}_x(g) = \inf_{r>0} \mathrm{Lip}(g|_{B_r(x)}) = \inf_{r>0} \sup_{y_1 \neq y_2 \in B_r(x)} \frac{d_\nu(g(y_1), g(y_2))}{d_\sigma(y_1, y_2)},$$

which by equivariance is a well-defined function on the convex core of $\mathbb{H}/\rho_1(\pi_1(S_{g,n})$. In the language of [GK17], the projection to $\mathbb{H}/\rho_1(\pi_1(S_{g,n}))$ of the stretch locus of an optimal Lipschitz map is exactly the boundary of the convex core. This property is very rare: it implies domination in the simple length spectrum (see [GS20]). In view of [GK17], we will say a Lipschitz map $g : (\mathbb{H}, \sigma) \to (X, \nu)$ is optimal if it satisfies the properties above.

Remark 4.1.2. Note that ρ_2 cannot be Fuchsian, by an application of Gauss-Bonnet. Almost strict domination is the same as strict domination when every $\rho_2(\zeta_i)$ is elliptic (see also [Sag19, Lemma 6.3]).

Before moving on, we comment that Dai-Li recently proved domination results for higher rank Hitchin representations into $\mathrm{PSL}(n, \mathbb{C})$ [DL20]. It would be interesting to see if the almost strict condition generalizes meaningfully to higher rank.

Maximal surfaces

We first fix some notations that we will keep throughout the chapter.

- Σ is surface with genus g and n punctures $p_1, \ldots, p_n$, with $\chi(\Sigma) < 0$. The deck group for the universal covering $\pi : \tilde{\Sigma} \to \Sigma$ is denoted by Γ.
- $\mathcal{T}(\Gamma)$ is the Teichmüller space of classes of complete finite volume marked hyperbolic metrics on Σ.
- $\{\zeta_1, \ldots, \zeta_n\} \subset \Gamma$ are the peripheral elements, i.e., those representing the simple closed curves enclosing p_i. If $n = 1$, write $\zeta_1 = \zeta$.
- (X, ν) is a CAT(-1) Hadamard manifold with isometry group G.
- $(\mathbb{H}, \sigma)$ denotes the hyperbolic space with constant curvature -1.

Let $\rho_1 : \Gamma \to \mathrm{PSL}(2, \mathbb{R})$, $\rho_2 : \Gamma \to G$ be reductive representations. Then $\rho_1 \times \rho_2 : \Gamma \to \mathrm{PSL}(2, \mathbb{R}) \times G$ defines its own representation.

Definition 4.1.3. A $\rho_1 \times \rho_2$-equivariant map $F : (\tilde{\Sigma}, \tilde{\mu}) \to (\mathbb{H} \times X, \sigma \oplus (-\nu))$ is maximal if the image surface has zero mean curvature. It is spacelike if the pullback metric $F^*(\sigma \oplus (-\nu))$ is non-degenerate and Riemannian.

The vanishing of the mean curvature is equivalent to the condition that F is harmonic and conformal. Using the product structure, we can write

$$F = (h, f),$$

where h, f are ρ_1, ρ_2-equivariant harmonic maps, and from (2.7),

$$\Phi(F) = \Phi(h) - \Phi(f).$$

Since F is conformal, $\Phi(h)$ and $\Phi(f)$ agree.

Definition 4.1.4. A spacelike maximal surface $F : (\tilde{\Sigma}, \tilde{\mu}) \to (\mathbb{H}\times, \sigma \oplus (-\nu))$ is called tame if the Hopf differentials of the harmonic maps have poles of order at most 2 at the cusps.

At this point, we can see a relationship to almost strict domination.

Lemma 4.1.5. *If ρ is Fuchsian, the existence of a tame spacelike maximal surface implies almost strict domination.*

Proof. Let F be such a maximal surface and split it as $F = (h, f)$. Note that $\Phi(h) = \Phi(f)$ implies $\ell(\rho_1(\zeta_i)) = \ell(\rho_2(\zeta_i))$ for all i. Indeed, $\ell(\rho_k(\zeta_i)) = 0$ if and only if the Hopf differential has a pole of order at most 1 at the cusp. And if $\ell(\rho_k(\zeta_i)) > 0$, this is because the residue at each cusp is determined entirely by the choice of twist parameter and the translation length $\ell(\rho_k(\zeta_i))$. We proved in the previous chapter that $f \circ h^{-1}$ is an optimal map in the sense of this chapter. □

Main theorems: maximal surfaces

Theorem 4A. Let $\rho_1 : \Gamma \to \mathrm{PSL}(2,\mathbb{R})$ and $\rho_2 : \Gamma \to G$ be reductive representations with ρ_1 Fuchsian. ρ_1 almost strictly dominates ρ_2 if and only if there exists a complete finite volume hyperbolic metric μ on Σ and a $\rho_1 \times \rho_2$-equivariant tame spacelike maximal immersion from

$$(\tilde{\Sigma}, \tilde{\mu}) \to (\mathbb{H} \times X, \sigma \oplus (-\nu)).$$

The maximal surfaces are not unique but are classified according to Proposition 4.2.2.

Let's give some idea of the proof. For representations ρ_1, ρ_2, we define a functional $\mathcal{E}^{\theta}_{\rho_1,\rho_2}$ on the Teichmüller space by

$$\mathcal{E}^{\theta}_{\rho_1,\rho_2}(\mu) = \int_{\Sigma} e(\mu, h^{\theta}_{\mu}) - e(\mu, f^{\theta}_{\mu}) dA_{\mu},$$

where h^{θ}_{μ}, f^{θ}_{μ} are certain harmonic maps on $(\tilde{\Sigma}, \tilde{\mu})$ that may have infinite energy, in the sense that

$$\int_{\Sigma} e(\mu, h^{\theta}_{\mu}) dA_{\mu} = \int_{\Sigma} e(\mu, f^{\theta}_{\mu}) dA_{\mu} = \infty.$$

We show that this is always finite, provided the boundary lengths for ρ_1 and ρ_2 agree (Section 4.2). We then compute the derivative (Section 4.2), showing that critical points correspond to spacelike maximal surfaces (Proposition 4.2.1). To anyone working with harmonic maps, this is expected, but with no good theory of global analysis to treat infinite energy maps on surfaces with punctures, we have to work through some thorny details directly. In the course of our analysis, we develop a new energy minimization result (Lemma 4.2.14) that may be of independent interest.

Then we show that $\mathcal{E}^{\theta}_{\rho_1,\rho_2}$ is proper if and only if ρ_1 almost strictly dominates ρ_2. Here is an indication as to why this is true. Suppose we diverge along a sequence $(\mu_n)_{n=1}^{\infty} \subset \mathcal{T}(\Gamma)$ by pinching a simple closed curve α. Then there is a collar around α whose length ℓ_n in (Σ, μ_n) is tending to ∞. Almost strict domination implies $\ell(\rho_1(\alpha)) > \ell(\rho_2(\gamma))$, and the analysis from [Sag19] shows that the total energy of the harmonic maps in the collar behaves like

$$\ell_n(\ell(\rho_1(\alpha))^2 - \ell(\rho_2(\alpha))^2) \to \infty. \tag{4.1}$$

This reasoning, however, cannot be turned into a full proof. Two problems:

1. Along a general sequence that leaves all compact subsets of $\mathcal{T}(\Gamma)$, the two harmonic maps could apriori behave quite differently in a thin collar. For instance we could have twisting in one harmonic map, which increases the energy, but no twisting in the other.
2. For a general sequence, we also have little control over the energy outside of thin collars.

We circumvent these issues as follows: if g is an optimal map, then our energy minimization Lemma 3.12 implies that

$$\mathcal{E}^{\theta}_{\rho_1,\rho_2}(\mu) \geq \int_{\Sigma} e(\mu, h^{\theta}_{\mu}) - e(\mu, g \circ h^{\theta}_{\mu}) dA_{\mu}.$$

The integrand is positive, so now we can bound below by the energy in collars. The contracting property of g then allows us to effectively study the energy in collars. In the end we make a rather technical geometric argument in order to find lower bounds similar to (4.1) along diverging sequences.

We also remark that even in the non-compact but finite energy setting, the result on the derivative of the energy functional was not previously contained in the literature. Hence we record it below.

Proposition 4.1.6. *Let $\rho : \Gamma \to G$ be a reductive representation with no hyperbolic monodromy around cusps, so that equivariant harmonic maps have finite energy. Then the energy functional $E_\rho : \mathcal{T}(\Gamma) \to [0, \infty)$ that records the total energy of a ρ-equivariant harmonic map from $(\tilde{\Sigma}, \mu) \to (X, \nu)$ is differentiable, with derivative at a hyperbolic metric μ given by*

$$dE_\rho[\mu](\psi) = -4\,\mathrm{Re}\langle \Phi, \psi \rangle.$$

Here Φ is the Hopf differential of the harmonic map at μ.

The proof can actually be extended to non-positively curved settings (see Remark 4.2.17).

Main theorems: parametrizations

The next theorem concerns the space of almost strictly dominating pairs. We denote by $\mathrm{Hom}^*(\Gamma, G) \subset \mathrm{Hom}(\Gamma, G)$ the space of reductive representations. G acts on $\mathrm{Hom}^*(\Gamma, G)$ by conjugation, and we define the representation space as

$$\mathrm{Rep}(\Gamma, G) = \mathrm{Hom}^*(\Gamma, G)/G.$$

In general this may not be a manifold, but it can have nice structure depending on G. For surfaces with punctures we would like to prescribe behaviour at the punctures.

Definition 4.1.7. Fix a collection of conjugacy classes $\mathbf{c} = (c_i)_{i=1}^n$ of elements in G. The relative representation space $\mathrm{Rep}_{\mathbf{c}}(\Gamma, G)$ is the space of reductive representations taking ζ_i into c_i, modulo conjugation.

We require one technical assumption on the group G: that if we choose a good covering of Σ and let $\chi(\Gamma, G)$ denote the space of G-local systems with respect to this covering that have reductive holonomy, then the projection from $\chi(\Gamma, G) \to \mathrm{Rep}(\Gamma, G)$ is a locally trivial principal bundle. We demand the same for the relative representation space, instead considering local systems whose holonomy representations respect $\mathbf{c}$. This assumption is satisfied under most cases of interest in Higher Teichmüller theory, for instance if G is a linear algebraic group (see [Lab13, Chapter 5]).

Within $\mathrm{Rep}_{\mathbf{c}}(\Gamma, G)$, we have the subset $\mathrm{Rep}_{\mathbf{c}}^{nf}(\Gamma, G)$ of representations that do not stabilize a plane of constant curvature -1 on which the action is Fuchsian. The almost strict domination condition is invariant under conjugation for both representations, so we can define $\mathrm{ASD}_{\mathbf{c}}(\Gamma, G)$ to be the subspace of pairs of representations $\rho_1 : \Gamma \to \mathrm{PSL}(2, \mathbb{R})$, $\rho_2 : \Gamma \to G$ such that ρ_1 is Fuchsian and almost strictly dominates ρ_2. Necessarily, ρ_1 lies in the Teichmüller space $\mathcal{T}_{\mathbf{c}}(\Gamma)$ (we use this notation when we fix the boundary monodromy according to $\mathbf{c}$).

Theorem 4B. Assume there are m peripherals such that $c_1, \dots, c_m$ are hyperbolic conjugacy classes. For each choice of parameters $\theta = (\theta_1, \dots, \theta_m) \in \mathbb{R}^m$, there exists a homeomorphism

$$\Psi^\theta : \mathcal{T}(\Gamma) \times \mathrm{Rep}^{nf}_{\mathbf{c}}(\Gamma, G) \to \mathrm{ASD}_{\mathbf{c}}(\Gamma, G).$$

Moreover, the homeomorphism is fiberwise in the sense that for each $\rho \in \mathrm{Rep}^{nf}_{\mathbf{c}}(\Gamma, G)$, it restricts to a homeomorphism

$$\Psi^\theta_\rho : \mathcal{T}(\Gamma) \times \{\rho\} \to U \times \{\rho\} \subset ASD_{\mathbf{c}}(\Gamma, G),$$

where U is a non-empty open subset of the Teichmüller space $\mathcal{T}_{\mathbf{c}}(\Gamma)$.

The mappings Ψ^θ are defined in essentially the same way as the map Ψ from [Tho17]. Theorem 4B should be compared with Theorem 3.1.4.

Main theorems: AdS 3-manifolds

Concerning AdS 3-manifolds, the following explains the relationship with spacelike immersions.

Proposition 4.1.8. *Given a ρ_1-invariant domain $V \subset \mathbb{H}$ on which ρ_1 acts properly discontinuously, there is a bijection between*

1. *(ρ_1, ρ_2)-equivariant maps $g : V \to \mathbb{H}$ that are locally strictly contracting, i.e.,*

$$d_\sigma(g(x), g(y)) < d_\sigma(x, y)$$

for $x \neq y$,

2. *and circle bundles $p : \Omega/(\rho_1 \times \rho_2(\Gamma)) \to V$, where $\Omega \subset AdS^3$ is a domain on which $\rho_1 \times \rho_2$ acts properly discontinuously and such that each circle fiber lifts to a complete timelike geodesics in AdS^3.*

Indeed, given a spacelike maximal surface (h, f) defined on $V \subset \mathbb{H}$, we will see later on that $g = h \circ f^{-1}$ is locally strictly contracting on a domain. The implication from (1) to (2) is a slight generalization of the work of Guéritaud-Kassel in [GK17], and should be known to experts.

Remark 4.1.9. The proof in [KR85] that properly discontinuous subgroups of AdS^3 are of the form Γ_{ρ_1,ρ_2} rests on their main lemma that there is no $\mathbb{Z}^2$-subgroup acting properly discontinuously. The proof is local, and one can adapt to show that any torsion-free discrete group acting properly discontinuously on a domain in AdS^3 takes this form.

Remark 4.1.10. A version of this holds more generally for geometric structures modelled on some rank 1 Lie groups. See Section 4.4.

Specializing to almost strict domination, we have the following.

Theorem 4C. Let $\rho_1, \rho_2 : \Gamma \to \mathrm{PSL}(2,\mathbb{R})$ be two reductive representations with ρ_1 Fuchsian. The following are equivalent.

1. ρ_1 almost strictly dominates ρ_2.
2. $\rho_1 \times \rho_2$ acts properly discontinuously on a domain $\Omega \subset \mathrm{AdS}^3$ and induces a fibration from $\Omega/(\rho_1 \times \rho_2(\Gamma))$ onto the interior of the convex core of $\mathbb{H}/\rho_1(\Gamma)$ such that each fiber is a timelike geodesic circle. Moreover, when there is at least one peripheral ζ with $\rho_1(\zeta)$ hyperbolic, no such domain in AdS^3 can be continued to give a fibration over a neighbourhood of the convex core.
3. There exists a complete hyperbolic metric μ on Σ and a (ρ_1, ρ_2)-equivariant embedded tame maximal spacelike immersion from $(\tilde{\Sigma}, \tilde{\mu}) \to (\mathbb{H} \times \mathbb{H}, \sigma \oplus (-\sigma))$.

Fixing a collection of conjugacy classes $\mathbf{c}$, there is a fiberwise homeomorphism

$$\Psi : \mathcal{T}_{\mathbf{c}}(\Gamma) \times \mathrm{Rep}_{\mathbf{c}}^{nf}(\Gamma, \mathrm{PSL}(2,\mathbb{R})) \to \mathrm{ASD}_{\mathbf{c}}(\Gamma, \mathrm{PSL}(2,\mathbb{R})).$$

If we restrict the domain to classes of irreducible representations, the image identifies with a continuously varying family of AdS 3-manifolds.

Components of our space of AdS 3-manifolds are classified by the relative Euler numbers (see [BIW10]). The only piece that doesn't follow quickly from Theorems 4A, 4B, and Proposition 4.1.8 is the implication from (2) to (1). To prove this part, we draw on the work of Guéritaud-Kassel on maximally stretched laminations [GK17] and show that the stretch locus (Definition 6.4) of an optimally Lipschitz map is exactly the boundary of the convex hull of the limit set.

In (2), the domain Ω is all of AdS^3 if and only if every $\rho_2(\zeta_i)$ is elliptic. This is a consequence of [GK17, Lemma 2.7] and the properness criteria, Theorem 3.1.2. Also related to (2), one can get incomplete AdS 3-manifolds fibering over larger subsurfaces, but still not extending to the whole surface, by doing strip deformations (see Section 4.4).

Outline

- In Section 4.2 we set up the proof of Theorem 4A by defining the energy functionals $\mathcal{E}^{\theta}_{\rho_1,\rho_2}$. We then show that it is well-defined and compute the derivative.
- In Section 4.3 we show that $\mathcal{E}^{\theta}_{\rho_1,\rho_2}$ is proper if and only if ρ_1 almost strictly dominates ρ_2.
- We prove Theorem 4B in Section 4.3 by studying variations of minimizers of $\mathcal{E}^{\theta}_{\rho_1,\rho_2}$ (similar to [Tho17, Section 2]).
- Section 4.4 is a change of pace. After giving an overview of the relevant aspects of AdS geometry, we prove Proposition 4.1.8 and Theorem 4C.
- We close with a section on parabolic Higgs bundles (not included in the paper, only in this thesis). This is adapted on [AL18] (closed surfaces). We explain how that the residue of the maximal surfaces (encoded as the residue of the Higgs field) is related to twisting behaviour of the associated map into the timelike Grassmanian.

Acknowledgements
I'd like to thank Qiongling Li, Vlad Marković, and Nicolas Tholozan for helpful discussions and correspondence. Thanks also to an anonymous referee for comments on a draft of the paper.

4.2 The derivative of the energy functional

In this section, we introduce the energy functional needed for the proof of Theorem 4A and compute its derivative. First, we discuss some notions that will be useful for working with tame harmonic maps.

Tame harmonic maps and twist parameters

On a hyperbolic surface (Σ, μ), recall that we have cusp regions that identify isometrically with

$$U(\tau) := \{z = x + iy : (x, y) \in [0, \tau] \times [a, \infty)\}/\langle z \mapsto z + \tau \rangle, \qquad (4.2)$$

equipped with the hyperbolic metric $y^{-2}|dz|^2$. When needed, we write τ_μ to highlight dependence on μ.

Throughout this chapter, by "hyperbolic metric on Σ" we mean a complete finite volume hyperbolic metric, unless specified otherwise. We refer to the Teichmüller space as the space of such metrics. To highlight dependence on the uniformizing Fuchsian group, we denote it by $\mathcal{T}(\Gamma)$

[Sag19, Theorem 1.1] says that for a reductive representation ρ with hyperbolic monodromy around the cusp, tame harmonic maps exist and are determined by the choice of twist parameter, with an exception if ρ is reducible. Here we introduce the twist parameters.

Take ζ_i with $\rho(\zeta_i)$ hyperbolic, and choose a constant speed parametrization α_i for the geodesic axis β_i of $\rho(\zeta_i)$. Let C be the cylinder

$$C = \{(x, y) \in [0, 1]^2\}/\langle (0, y) \sim (1, y) \rangle \qquad (4.3)$$

with the flat metric. There is a "model mapping" $\tilde{\alpha}_i^{\theta_i} : C \to \beta_i$ defined as follows. For $\theta_i = 0$, we set

$$\tilde{\alpha}_i(x, y) = \alpha_i(x),$$

a constant speed projection onto the geodesic. For $\theta_i \neq 0$, $\tilde{\alpha}_i^{\theta}$ is defined by precomposing $\tilde{\alpha}_i$ with the fractional Dehn twist of angle θ_i, the map given in coordinates by

$$(x, y) \mapsto (x + \theta_i y, y).$$

Choosing a cusp neighbourhood U of a p_i, we take conformal maps $i_r : C \to U$ that take the boundaries linearly to $\{(x, y) : y = r\}$ and $\{(x, y) : y = r + 1\}$. The mappings $f^\theta \circ i_r : C \to (X, \nu)$ are harmonic, and as $r \to \infty$ they converge in the C^∞ sense to a harmonic mapping that differs from $\tilde{\alpha}_i^{\theta_i} : C \to \beta_i$ by a constant speed translation along the geodesic (see [Sag19, Section 5]).

The energy functional

One direction of Theorem 4A is Lemma 4.1.5. The proof of the other direction will go as follows. Take a pair (ρ_1, ρ_2) such that ρ_1 almost strictly dominates ρ_2.

We want to show there is a hyperbolic metric μ on Σ such that, after choosing twist parameters in some way, the associated (ρ_1, ρ_2)-equivariant surface from $(\tilde{\Sigma}, \tilde{\mu}) \to (\mathbb{H} \times X, \sigma \oplus (-\nu))$ is spacelike and maximal. Clearly, we must have the same twist parameter θ for both ρ_1 and ρ_2, and it turns out we can choose any θ, as we will now explain. Given a metric μ, let $h_\mu^\theta, f_\mu^\theta$ be harmonic maps for ρ_1, ρ_2 respectively with twist parameter θ. Define

$$\mathcal{E}^\theta_{\rho_1,\rho_2} : \mathcal{T}(\Gamma) \to \mathbb{R}$$

by

$$\mathcal{E}^\theta_{\rho_1,\rho_2}(\mu) = \int_\Sigma e(\mu, h_\mu^\theta) - e(\mu, f_\mu^\theta) dA_\mu.$$

This does not depend on the metric $\mu \in [\mu]$, by conformal invariance. When ρ_2 is reducible, harmonic maps are not unique, but the energy density does not depend on the choice of harmonic map, so the integral is well-defined. Apriori, it is not given that the integral defining $\mathcal{E}^\theta_{\rho_1,\rho_2}(\mu)$ is finite, for the harmonic maps themselves could have infinite energy. Finiteness will be proved in Proposition 4.2.4. The main goal of this section is to prove Proposition 4.2.1.

Proposition 4.2.1. *Given (ρ_1, ρ_2) with $\ell(\rho_1(\zeta_i)) = \ell(\rho_2(\zeta_i))$ for all i, the functional $\mathcal{E}^\theta_{\rho_1,\rho_2} : \mathcal{T}(\Gamma) \to \mathbb{R}$ is differentiable with derivative at a hyperbolic metric μ given by*

$$d\,\mathcal{E}^\theta_{\rho_1,\rho_2}[\mu](\psi) = -4\,\mathrm{Re}\langle \Phi(h_\mu^\theta) - \Phi(f_\mu^\theta), \psi\rangle.$$

Thus, critical points of $\mathcal{E}^\theta_{\rho_1,\rho_2}$ correspond to maximal surfaces. And by [Sag19, Proposition 3.13], any critical point is a spacelike immersion. In Section 4.3, we prove that the almost strict domination hypothesis implies there is a metric μ that minimizes $\mathcal{E}^\theta_{\rho_1,\rho_2}$, and that this is the only critical point. Granting this, we can discuss uniqueness.

Proposition 4.2.2. *Suppose ρ_1 almost strictly dominates ρ_2 and that ρ_2 is irreducible. Assume that for every choice of twist parameter θ, each $\mathcal{E}^\theta_{\rho_1,\rho_2}$ admits a unique critical point. Then every spacelike maximal immersion is of the form $(h^\theta_{\mu_\theta}, f^\theta_{\mu_\theta})$, where μ_θ is the minimizer for $\mathcal{E}^\theta_{\rho_1,\rho_2}$. If ρ_2 is reducible, then for each θ there is a a 1-parameter family of tame maximal surfaces, and each one is found by translating f_μ^θ along a geodesic axis.*

Proof. In the irreducible case, let $F = (h', f') : (\tilde{\Sigma}, \tilde{\mu}) \to (X, \nu)$ be a spacelike maximal immersion. By the uniqueness statement in [Sag19, Theorem 1.1], h' must be of the form h_μ^θ for some parameter θ. Thus $f = f_\mu^\theta$, and by our assumption, μ is the unique critical point for $\mathcal{E}_{\rho_1,\rho_2}^\theta$.

In the reducible case, all $\rho_2(\zeta_i)$ that are hyperbolic must translate along the same geodesic. Since translating along the geodesic does not change the energy density or the Hopf differential, the one minimizer for $\mathcal{E}_{\rho_1,\rho_2}^\theta$ is the only one that yields maximal surfaces. ◻

Analytic preliminaries

The rest of this section is devoted to proving Proposition 4.2.1. We assume Σ has only one puncture p with peripheral ζ—the analysis is local so the general case is essentially the same. When $\rho_1(\zeta), \rho_2(\zeta)$ are hyperbolic, we assume the twist parameter θ is 0, and we also write ℓ in place of $\ell(\rho_k(\zeta))$. At the end of the section, we explain the adjustments for the general case. With these assumptions, it causes no harm to write $\mathcal{E}$ for $\mathcal{E}_{\rho_1,\rho_2}^\theta$. Not only in this section but for the rest of the chapter, we write h_μ, f_μ for $h_\mu^\theta, f_\mu^\theta$ when θ is zero.

We collect some notations that, in the sequel, we use without comment.

- When unspecified, C denotes a constant that may grow in the course of a proof.
- Setting $x = x_1, y = x_2$, let $e_{\alpha\beta}(f) = f^*\nu\left(\frac{\partial f}{\partial x_\alpha}, \frac{\partial f}{\partial x_\beta}\right)$, so that $e(\mu, f) = \frac{1}{2}\mu^{\alpha\beta}e_{\alpha\beta}(f)$.
- H and L are holomorphic and anti-holomorphic energies.
- For a conformal metric $\mu = \mu(z)|dz|^2$, $||\Phi||^2 := \mu^{-2}|\Phi|^2 = HL$.
- Set $|\psi| = L^{1/2}/H^{1/2}$. If $(X,\nu) = (\mathbb{H},\sigma)$ and ρ is Fuchsian, $\psi = \frac{f_{\bar{z}}}{f_z}$ is the Beltrami form.

We now prepare notation for dealing with cusps (we do things slightly differently than in the last chapter). For $r \geq 2$, we define (Σ_r, μ) to be $\Sigma \backslash \{z \in U(\tau) : y > r\}$, and we put $(\tilde{\Sigma}_r, \tilde{\mu})$ to be the preimage in $\tilde{\Sigma}$. Note that, as a set of points, this depends on μ. C is a conformal cylinder, and when we say a "conformal cylinder for μ," we mean the length is adjusted to be τ and the height is 1, so that there are conformal maps $C \to U(\tau)$. We define the maps

$$i_r : C \to \Sigma_{r+1} \backslash \Sigma_r \subset \Sigma$$

by $(x, y) \mapsto x + i(y + r)$, which are used to study asymptotics of harmonic maps.

As in the previous chapter, it is often helpful to perturb the metric in the cusp. A metric μ is expressed in the coordinate of $U(\tau)$ as $\mu(z) = y^{-2}|dz|^2$. The flat-cylinder metric μ^f is defined by μ in $\Sigma_2 = \Sigma \backslash U$, the flat metric $|dz|^2$ in the cusp coordinates on $\Sigma \backslash \Sigma_3$, and smoothly interpolated in between. This is conformally equivalent to μ, so harmonic maps for μ and μ^f are the same.

Lastly, we recall some aspects of the construction of infinite energy harmonic maps from [Sag19, Section 5]. We specialize to a reductive representation ρ such that $\rho(\zeta)$ is hyperbolic with $\theta = 0$. Denote by β the geodesic axis of $\rho(\zeta)$ and fix a constant speed parametrization $\alpha : [0, \tau] \to \beta$ such that $\rho(\zeta)\alpha(0) = \alpha(\tau)$. We define mappings f_r to be equivariant harmonic maps on $(\tilde{\Sigma}_r, \tilde{\mu})$ with equivariant boundary values specified by α on $\partial\tilde{\Sigma}_r$. In keeping with the notation of [Sag19], we set $\varphi = f_2$, extend φ vertically into the cusp by $\varphi(x, y) = \alpha(x)$ in a fundamental domain for Γ, and extend equivariantly. In this cusp, φ satisfies

$$e(\mu, \varphi) = \frac{\ell^2}{2\tau^2}.$$

In [Sag19, Proposition 5.1], we show there is a uniform bound

$$\int_{\Sigma_s} e(\mu, f_r) dA_\mu \leq \int_{\Sigma_s} e(\mu, \varphi) dA_\mu = \int_{\Sigma_2} e(\mu, \varphi) dA_\mu + \frac{(s-2)\ell^2}{2\tau}. \tag{4.4}$$

This bound will be used in the analysis below.

Next we turn to finiteness of $\mathcal{E}$. If no monodromy is hyperbolic, then harmonic maps have finite energy. So suppose $\rho(\zeta)$ is hyperbolic.

Proposition 4.2.3. *Let $f : (\tilde{\Sigma}, \tilde{\mu}) \to (X, \nu)$ be a ρ-equivariant harmonic map. Then there exists $C, c, y_0 > 0$ such that, in the cusp coordinates (4.2), for all $y \geq y_0$, the inequality*

$$|e(\mu^f, f_\mu)(x, y) - \ell^2/2\tau^2| < Ce^{-cy}$$

holds.

Proof. We implicitly work with the metric μ^f. Write

$$e = H + L = H^{1/2}L^{1/2}\Big(\frac{H^{1/2}}{L^{1/2}} + \frac{L^{1/2}}{H^{1/2}}\Big) = |\Phi|(|\psi|^{-1} + |\psi|). \tag{4.5}$$

Since Φ has a pole of order at most 2, changing coordinates to the cusp $[0, 1] \times [1, \infty)$ (see the subsection below) gives the expression

$$|\Phi| = \frac{\ell^2}{4\tau^2} + O(e^{-2\pi y}). \tag{4.6}$$

From [Wol91b, page 513], if f is Fuchsian, we have the estimate

$$1 - |\psi|^2 = O(e^{-\ell/2y}),$$

and hence

$$1 - O(e^{-cy}) \leq |\psi|^2 \leq 1.$$

If ρ is not Fuchsian, we find the Fuchsian harmonic map $h : (\tilde{\Sigma}, \tilde{\mu}) \to (\mathbb{H}, \sigma)$ with the same Hopf differential. Then from $H(f)L(f) = H(h)L(h)$ and [Sag19, Proposition 3.13] we get

$$L(h) \leq L(f), H(f) \leq H(h),$$

which implies

$$|\psi(h)| \leq |\psi(f)| \leq |\psi(h)|^{-1},$$

and furthermore

$$1 - O(e^{-cy}) \leq |\psi|^2 \leq \frac{1}{1 - O(e^{-cy})}. \tag{4.7}$$

Inserting (4.6) and (4.7) into (4.5) gives

$$e(\mu^f, f) = \frac{\ell^2}{2\tau^2} + O(e^{-c_1 y}),$$

as desired. □

The following is now evident.

Proposition 4.2.4. *For a fixed hyperbolic metric μ, the integral defining $\mathcal{E}(\mu)$ is finite. That is, $\mathcal{E} : \mathcal{T}(\Gamma) \to \mathbb{R}$ is well-defined.*

Tangent vectors of Teichmüller space.

To do analysis on Teichmüller space, we need a tractable way to study tangent vectors. Fix a metric μ_0 on Σ and let $z = x + iy$ be a conformal coordinate for the compatible holomorphic structure, so that $\mu_0 = \mu_0(z)|dz|^2$. In this subsection we describe variations of the metric

$$\mu' = \mu + \dot{\mu}.$$

Since we are working with harmonic maps, which are conformally invariant, we are permitted to work in a specified conformal class. In particular, we may restrict to variations through complete finite volume hyperbolic metrics. The hyperbolic condition is satisfied if and only if

$$\dot{\mu}_{11} + \dot{\mu}_{22} = 0$$

and

$$\phi(z)dz^2 = (\dot{\mu}_{11} - i\dot{\mu}_{12})dz^2$$

is a holomorphic quadratic differential. When Σ has a puncture, the necessary and sufficient condition on $\dot{\mu}$ to preserve the complete finite volume property is that ϕ has a pole of order at most 1 at the cusp.

Remark 4.2.5. We have described the tangent space to Teichmüller space at (Σ, μ) as the space of L^1-integrable quadratic differentials on the conjugate Riemann surface. This characterization coincides with the one we get from the Bers embedding into $\mathbb{C}^{3g-3+n}$. If $\mu(z)$ is a conformal metric, the mapping

$$\phi(z)dz^2 \mapsto \mu(z)^{-1}\phi(\bar{z})\frac{d\bar{z}}{dz}$$

yields the usual identification with the space of harmonic Beltrami forms.

We work out the growth condition on ϕ in the cusp coordinates for $U(\tau)$. For convenience put $\tau = 1$. The mapping $z \mapsto w(z) = e^{2\pi i z}$ takes a vertical strip to a punctured disk

$$\{0 \le x < 1, y > h, y^{-2}|dz|^2\} \to \{0 < |w| < e^{-h}, |w|^{-2}(\log|w|)^2|dw|^2\}$$

holomorphically and isometrically. A meromorphic quadratic differential in the disk with a pole of order at most 1 is written

$$\Phi = \phi(w) = (a_{-1}w^{-1} + \varphi(w))dw^2,$$

with φ holomorphic. Applying the above holomorphic mapping, the differential transforms according to

$$\Phi = \phi(w(z))\Big(\frac{\partial w(z)}{\partial z}\Big)^2 dz^2 = \phi(e^{iz})(ie^{iz})^2dz^2 = -(a_{-1}e^{iz} + e^{2iz}\varphi(e^{iz}))dz^2. \quad (4.8)$$

Thus, any admissible variation decays exponentially in the cusp as we take $y \to \infty$.

We also need to describe the inverse variation

$$(\mu')^{\alpha\beta} = \mu^{\alpha\beta} + \dot{\mu}^{\alpha\beta}.$$

In the Einstein notation, the relation $\mu^{\alpha\beta}\mu_{\beta\gamma} = \delta_{\alpha\gamma}$ gives $\dot{\mu}^{\alpha\beta}\mu_{\beta\gamma} + \mu^{\alpha\beta}\dot{\mu}_{\beta\gamma} = 0$, and hence

$$\dot{\mu}^{\alpha\beta} = -\mu^{\alpha\rho}\mu^{\beta\tau}\dot{\mu}_{\tau\rho}.$$

When μ is conformal, this returns

$$\dot{\mu}^{\alpha\beta} = -\mu^{-2}\dot{\mu}_{\alpha\beta}.$$

We derive that, in the cusp coordinates,

$$\dot{\mu}^{\alpha\beta} = -y^4\dot{\mu}_{\alpha\beta},$$

so the decay is still exponential. From these descriptions we deduce the following.

Proposition 4.2.6. *Suppose g is a finite energy equivariant map with respect to a finite volume hyperbolic metric μ. Then it also has finite energy for any other metric.*

Proof. By conformal invariance of energy, we are permitted to work with complete finite volume hyperbolic metrics. It is enough to prove the claim for metrics that are as close as we like to μ. For then we can connect μ to any other metric μ' via a smooth path in the Teichmüller space, cover this path with finitely many small balls, and argue inductively. That is, we can assume $\mu' = \mu + \dot{\mu}$, for some small variation $\dot{\mu}$. In a local coordinate $z = x + iy$, we write $|\mu + \dot{\mu}| = \det(\mu + \dot{\mu})$, so that the volume form is

$$dA_{\mu+\dot{\mu}} = \sqrt{|\mu + \dot{\mu}|}dz \wedge d\bar{z} = \sqrt{|\mu| - |\phi|^2}dz \wedge d\bar{z},$$

where ϕ is the holomorphic quadratic differential associated to $\dot{\mu}$. For simplicity, let's restrict to a C^1 map g. In a cusp,

$$\begin{aligned} 2e(\mu + \dot{\mu}, \nu, g)\sqrt{|\mu + \dot{\mu}|} &= \sqrt{|\mu| - |\phi|^2}(\mu + \dot{\mu}^{\alpha\beta})e_{\alpha\beta}(g) \\ &= \frac{\sqrt{|\mu| - |\phi|^2}}{\sqrt{|\mu|}} \cdot \sqrt{|\mu|}e(\mu, g) + \sqrt{|\mu| - |\phi|^2}\dot{\mu}^{\alpha\beta}e_{\alpha\beta}(g). \end{aligned}$$

By hypothesis, the first term is uniformly bounded and converges to an integrable quantity. For the second term, using the flat cylinder metric we gather

$$\int_\Sigma (e_{11} + e_{22})dA_{\mu^f} < \infty.$$

By Cauchy-Schwarz, e_{12} is integrable as well. Meanwhile, the factor $\sqrt{|\mu| - |\phi|^2}\dot{\mu}^{\alpha\beta}$ decays exponentially as we go into the cusp. □

The proof above shows that the bound depends only on the energy of g with respect to μ and the Teichmüller distance of the new metric to μ (see [Ahl06, Chapter 5] for the definition).

Variations: finite energy harmonic maps

In the proof of Proposition 4.2.1, we need to know that for a variation of hyperbolic metrics $t \mapsto \mu_t$, $f_{\mu_t} \to f_\mu$ pointwise as $t \to 0$. In this subsection and the next, we show uniform convergence of harmonic maps on compacta. To do so, we verify that analytic results for harmonic maps can be made uniform in the source metric. We start with the case of elliptic and parabolic monodromy at the cusp. Throughout, let $\rho : \Gamma \to G$ be a reductive representation.

Proposition 4.2.7. *Assume the monodromy is elliptic or parabolic and fix an admissible metric μ_0. For all μ, there exists a $C_k > 0$ depending only on the Teichmüller distance from μ_0 to μ such that for all $k > 0$,*

$$|(\nabla^{\mu^f,\nu})^{(k)} df_\mu|_{\mu^f} \leq C_k.$$

Proof. We showed in [Sag19, Proposition 3.8] that one can always find finite energy harmonic maps. Take such a map $g : (\tilde{\Sigma}, \tilde{\mu}_0) \to (X, \nu)$, which by Lemma 4.2.6 is finite energy for any other metric, with a bound depending on the Teichmüller distance to μ_0. By the energy minimizing property in negatively curved spaces,

$$\int_\Sigma e(\mu^f, f_\mu) dA_{\mu^f} = \int_\Sigma e(\mu, f_\mu) dA_\mu \leq \int_\Sigma e(\mu, g) dA_\mu \leq C.$$

For the uniform bounds on the energy density, independent of μ, one uses a Harnack-type inequality, say, from [SY97, page 171], that only depends on uniform quantities: the Ricci curvature and the injectivity radius. Since we work with the flat-cylinder metric μ^f as opposed to the hyperbolic metric μ, we do have uniform control on the injectivity radius. The estimates on higher order derivatives then come from the elliptic theory on the Sobolev space adapated to (Σ, μ_0) (see [Nic21, Chapter 10]). It is clear from the general theory that the implicit constants in these estimates can be made uniform in μ. □

Lemma 4.2.8. *Suppose ρ is irreducible. Let $K \subset \tilde{\Sigma}$ be compact and suppose $f : (K, \mu) \to (X, \nu)$ is C-Lipschitz for some $C > 0$. Then there exists a compact set $\Omega(K, C) \subset X$ that does not depend on the map f such that $f(K) \subset \Omega(K, C)$.*

This is essentially carried out in the proof of Proposition 5.1 in [Sag19] (which is a modification of the argument from the main theorem of [Don87]), but with some slightly different assumptions. We sketch a proof for the reader's convenience.

Proof. We show the claim holds for a single point ξ, and then the Lipschitz control promotes the result to general compact sets. Since ρ is irreducible, there exists $\gamma \in \Gamma$ such that $\rho(\gamma)\xi \neq \xi$. For each $x \in K$, γ may be represented by a loop $\gamma_x : [0, L_x] \to \Sigma$ based at $\pi(x)$. Since K is compact, there is an $L > 0$ such that $L_x \leq L$ for all x. Now choose a neighbourhood $B_\xi \subset X \cup \partial_\infty X$ such that

$$d_\nu(X \cap B_\xi, \rho(\gamma)X \cap B_\xi) > CL.$$

This implies that for any $x \in K$, $f(x)$ cannot lie in B_ξ. Repeating this procedure and using compactness of $\partial_\infty X$, we get a neighbourhood of $\partial_\infty X$ in $X \cup \partial_\infty X$ that $f(K)$ cannot enter. We then take $\Omega(K, C)$ to be the complement of this neighbourhood. □

Lemma 4.2.9. *Let $K \subset \Sigma$ be compact. Then there is a choice of harmonic maps f_μ that vary continuously on lifts of K inside $(\tilde{\Sigma}, \tilde{\mu})$ in the C^∞ topology.*

Proof. We argue by contradiction. First if ρ is irreducible, suppose there exists $\delta > 0$ and sequences $(k_n)_{n=1}^\infty \subset K$, $r_n \to \infty$, and $\mu_n \to \mu$ such that

$$d(f_{\mu_n}, f_\mu) \geq \delta \tag{4.9}$$

for all n. Then by Lemma 4.2.7 we have uniform derivative bounds on each f_{μ_n} and by Lemma 4.2.8 they all take a lift of K to $\tilde{\Sigma}$ into a compact subset of (X, ν). By Arzelà-Ascoli we see the f_{μ_n} C^∞-converge in K along a subsequence to a limiting harmonic map f_∞. By equivariance, we have this same convergence on the whole preimage of K.

We now show $f_\infty = f$, which contradicts (4.9). Taking a compact exhaustion of Σ and applying the same argument on each compact set, the maps f_{μ_n} subconverge on compact subsets of $\tilde{\Sigma}$ to a limiting finite energy harmonic map f'_∞ that agrees with f_∞ on lifts of K. By uniqueness for finite energy maps we get $f_\infty = f$.

If ρ is reducible, we can recenter the harmonic maps via translations along the geodesic so that they take K into a fixed compact set (see the proof of [Sag19, Proposition 5.1] for this routine procedure), and then repeat the argument above. □

Variations: infinite energy harmonic maps

Now we treat harmonic maps for representations ρ with hyperbolic monodromy at the cusp. Recall the map φ, which for a metric μ we now denote $\varphi_\mu : (\tilde{\Sigma}_2, \tilde{\mu}) \to (X, \nu)$. We emphasize that the boundary curve for $(\tilde{\Sigma_2}, \tilde{\mu})$ is varying with $\tilde{\mu}$

Lemma 4.2.10. *Near a base hyperbolic metric μ_0, φ_μ can be chosen so that the association $\mu \to \varphi_\mu$ is continuous in the C^∞ topology .*

Proof. Since the metrics are varying smoothly, the lifts of μ-horocycles vary smoothly with μ. Indeed, in conformal cusp coordinates as in (4.3), the μ-horocycles are just curves with y_μ constant. Smoothness here thus comes from the regularity theory for the Beltrami equation (if we follow the approach of Ahlfors and Bers for finding isothermal coordinates on a hyperbolic surface).

From the standard arguments (see the beginning of the section), the result amounts to choosing $\mu \to \varphi_\mu$ so that we have a uniform total energy bound on compacta, independent of μ, and such that the boundary data varies continuously. It would follow that as $\mu \to \mu_0$, the harmonic maps subconverge to a harmonic map, and continuity in the boundary values shows this is exactly φ. One can then modify the contradiction argument from Lemma 4.2.9 to see C^∞ convergence on compacta along the whole sequence.

By the energy minimizing property, it suffices to construct a family of maps ψ_μ with suitably chosen boundary values and uniformly controlled energy. To build ψ_μ, let f^μ be the unique quasiconformal diffeomorphism between (Σ_2, μ) and (Σ_2, μ_0) that, in the cusp coordinates, takes $(i2, \tau_\mu + i2, \infty) \mapsto (i2, \tau_{\mu_0} + i2, \infty)$. Then $\psi_\mu = \varphi_{\mu_0} \circ f^\mu$ has the correct boundary values. By our choice of normalizations, f^μ converges to the identity as $\mu \to \mu_0$. This gives an energy bound on f^μ, and moreover we get uniform bounds for ψ_μ. □

We can now prove the analogue of Lemma 4.2.9 for infinite energy harmonic maps.

Lemma 4.2.11. *Let $K \subset \Sigma$ be compact. Then there is a choice of harmonic maps f_μ that vary continuously on lifts of K in the C^∞ topology.*

Proof. From the estimate (4.4) and the Fatou lemma we get

$$\int_{\Sigma_s} e(\mu, f_\mu) dA_\mu \leq \int_{\Sigma_2} e(\mu, \varphi_\mu) dA_\mu + \frac{(r-2)\ell^2}{2\tau_\mu}.$$

Via the lemma above,

$$\int_{\Sigma_s} e(\mu, f_\mu) dA_\mu \leq C + \frac{2(r-2)\ell^2}{2\tau_{\mu_0}},$$

for μ close enough to μ_0. Then from the discussion in Proposition 4.2.7, we get uniform control on all derivatives on K.

If ρ is irreducible we apply Lemma 4.2.8, and if ρ is reducible we rescale by translations along the geodesic so that they take K into a fixed compact set. Well-used arguments show that for any sequence $\mu_n \to \mu$, the harmonic maps subconverge on compacta in the C^∞ sense to a limiting harmonic map f_∞. From [Sag19, Lemma 5.2], $d(f_{\mu_n}, \varphi_{\mu_n})$ is uniformly bounded, and an investigation of the proof of this lemma shows it is maximized on Σ_2. By uniform convergence on compacta, $d(f_\infty, \varphi)$ is also (non-strictly) maximized on Σ_2, and hence it is globally finite. According to the classification in [Sag19, Theorem 1.1] and the explanation of the twist parameter, f_μ is the only harmonic map with this property: if we precompose an equivariant map with a non-trivial fractional Dehn twist, the distance between the original map and the new map grows without bound as we limit toward a lift of the puncture on the boundary at infinity. □

We deduce the following.

Lemma 4.2.12. *Given a metric μ_0, there exists a uniform $C_k > 0$ such that*

$$|(\nabla^{\mu^f,\nu})^{(k)} df_\mu|_{\mu^f} \leq C_k$$

everywhere on Σ, with C_k depending on the Teichmüller distance to μ_0.

Proof. By uniform convergence on compacta, we have uniform control on the distance to φ_μ. We then couple the uniform energy bounds on φ_μ with Cheng's lemma to get uniform energy bounds, and then we appeal to the elliptic theory (as we have done many times at this point). □

Energy minimizing properties

It is well known that finite energy equivariant harmonic maps minimize the total energy among other equivariant maps. Here we show that this extends in some sense to the infinite energy setting.

Definition 4.2.13. Let $g : (\tilde{\Sigma}, \tilde{\mu}) \to (X, \nu)$ be ρ-equivariant and let α be a parametrization of the geodesic axis of the image of the peripheral, say β. We say g converges at ∞ to α if, after precomposing with the conformal mapping $i_r : C \to \tilde{\Sigma}_r \backslash \tilde{\Sigma}_{r-1}$, $g \circ i_r : C \to \beta$ converges in C^0 as $r \to \infty$ to a mapping $\tilde{\alpha} : C \to \beta$ given by

$$\tilde{\alpha}(x, y) = \alpha'(x),$$

where $\alpha'(x)$ is some translation of α along β.

Proposition 4.2.14. *Let f be the harmonic map from [Sag19, Theorem 1.1] whose Hopf differential has real reside at the cusp. Suppose that a locally Lipschitz map $g : (\tilde{\Sigma}, \tilde{\mu}) \to (X, \nu)$ converges at ∞ to α. Then $e(\mu, f) - e(\mu, g)$ is integrable with respect to μ and*

$$\int_{\Sigma} e(\mu, f) - e(\mu, g) dA_{\mu} \leq 0.$$

Remark 4.2.15. The lemma holds provided g is weakly differentiable and these weak derivatives are locally L^2. If g is not harmonic, then this inequality is strict.

Proof. Let g_r be the harmonic map with the same boundary values as g on Σ_r and set $\phi : (\tilde{\Sigma}, \tilde{\mu}) \to \mathbb{R}$ to be any Γ-invariant function that is equal to $\mu^{-1}\ell^2/2\tau_{\mu}$ in μ-conformal coordinates after $\tilde{\Sigma}_2$. For example, we can take the μ-energy of the harmonic map φ. Then

$$\int_{\Sigma_r} e(\mu, g_r) dA_{\mu} \leq \int_{\Sigma_r} e(\mu, g) dA_{\mu},$$

and hence

$$\liminf_{r \to \infty} \int_{\Sigma_r} e(g_r) - \phi dA_{\mu} \leq \int_{\Sigma} e(g) - \phi dA_{\mu}. \tag{4.10}$$

Without loss of generality, α is the limiting parametrization of the geodesic for g. Let φ_r be the unique harmonic map on $\tilde{\Sigma}_r$ with boundary values α on $\partial\tilde{\Sigma}_r$, defined by using α on one lift and then extending equivariantly. The distance function

$$p \mapsto d(\varphi_r(p), g_r(p))$$

is subharmonic and hence maximized on $\partial\Sigma_r$. By the convergence property of g, we can thus assume that for r large enough,

$$d(\varphi_r(p), g_r(p)) \leq 1$$

for all p. By Cheng's lemma, we obtain uniform bounds on the energy density of g_r in terms of that of φ_r on compact sets. Since $\varphi_r \to f$ locally uniformly on compacta, g_r also subconverges locally uniformly on compacta, and the limiting map is harmonic. As g_r has bounded distance to φ, as above the uniqueness result [Sag19, Theorem 1.1] shows the limit must be f. Returning to our integrals (4.10), Fatou's lemma then yields

$$\begin{aligned}
\int_{\Sigma} e(\mu, f) - \phi dA_{\mu} &= \int_{\Sigma_2} e(\mu, f) - \phi dA_{\mu} + \int_{\Sigma \setminus \Sigma_2} e(\mu, f) - \phi dA_{\mu} \\
&= \int_{\Sigma_2} e(\mu, f) - \phi dA_{\mu} + \int_2^{\infty} \Big(\int_0^{\tau_{\mu}} e(\mu^f, f)(x, y) - \phi(x, y) dx \Big) dy \\
&\leq \liminf_{r \to \infty} \int_{\Sigma_r} e(g_r) - \phi dA_{\mu}.
\end{aligned}$$

(4.10) then gives

$$\int_\Sigma e(\mu, f) - \phi dA_\mu \leq \int_\Sigma e(g) - \phi dA_\mu.$$

If the integral on the right is infinite then the result of our lemma is obvious, and if it is finite then we can rearrange to get the desired inequality. □

Lemma 4.2.16. *In the setting above, work with a metric μ_0. Then for any other μ, f_μ satisfies the hypothesis above.*

Proof. Let w be a complex coordinate parametrizing a cusp as a quotient of a vertical strip of length 1. Recall z is defined on a strip $y \geq a$, $0 \leq x \leq \tau_\mu$, and similar for z_μ (we take the same a). We can assume $z = f^{\lambda_0}(w)$, $z_\mu = f^\lambda(w)$, where f^{λ_0}, f^λ are smooth quasiconformal maps with complex dilatations λ_0, λ respectively. From the choice of coordinates, f^{λ_0} maps $(0, 1, \infty) \mapsto (0, \tau_{\mu_0}, \infty)$ and f^λ maps $(0, 1, \infty) \mapsto (0, \tau_\mu, \infty)$. $f^\lambda \circ (f^{\lambda_0})^{-1}$ takes $(x, y) \mapsto (x_\mu, y_\mu)$, yielding a quasiconformal mapping between the strips. Provided the metrics are close enough in Teichmüller space, λ and λ_0 are related by a small variation

$$\lambda = \lambda_0 + \dot{\lambda}.$$

Note that from a previous computation, $\dot{\lambda} \to 0$ as $y \to \infty$. It follows that $f^\lambda \circ (f^{\lambda_0})^{-1}$ has complex dilatation tending to 0 as $y \to \infty$. Mori's theorem [Ahl06, Chapter 3] implies $f^\lambda \circ (f^{\lambda_0})^{-1}$ has uniform Hölder continuity.

Upon pre-composing with cylinders that are conformal for μ, the harmonic map f_μ converges to a projection onto the geodesic axis β. We can take these μ-cylinders as large as we like. Using the Hölder continuity and our normalizations, given a family μ_0 cylinders of height 1, we can embed each cylinder in a μ-cylinder of fixed height. C^0 convergence to the projection onto the geodesic thus follows. □

Proof of Proposition 4.2.1

For the remainder of this section, fix a background metric μ_0. There is a Serre duality pairing between holomorphic quadratic differentials on Σ with at most first order poles at the cusp and harmonic Beltrami forms with appropriate decay (see Remark 4.2.5):

$$\langle \Phi, \psi \rangle = \int_\Sigma \phi \psi dz d\overline{z},$$

where $\Phi = \phi(z) dz^2$ is a coordinate expression. We recall that Proposition 4.1.6 asserts that the derivative of the total energy of a finite energy harmonic map, evaluated on a harmonic Beltrami form ψ, is

$$dE_\rho[\mu](\psi) = -4\,\mathrm{Re}\langle \Phi(f_\mu), \psi\rangle.$$

For finite energy harmonic maps, Proposition 4.2.1 is a direct corollary of Proposition 4.1.6, which we prove first. For closed surfaces, Proposition 4.1.6 is well known, although the history is unclear—the earliest computation may go back to the work of Douglas [Dou39]. The most modern version is contained in [Wen07]. The difference without compactness is that we must control variations at the cusp.

Proof of Proposition 4.1.6. Let us assume $\mu(z)$ is conformal, and let

$$\mu_t = \mu + t\dot{\mu} + t\epsilon(t)$$

be a variation through hyperbolic metrics, where $\epsilon(t) \to 0$ as $t \to 0$. We denote by ϕ_t the associated holomorphic quadratic differential, which is given in local coordinates by

$$\phi_t(z) = t\Big((\dot{\mu}_{11} + \dot{\epsilon}_{11}(t)) - i(\dot{\mu}_{12} + \dot{\epsilon}_{12}(t))\Big)dz^2,$$

and we set ψ_t to be the associated harmonic Beltrami form. We put ϕ to be the quadratic differential for the variation $\mu + \dot{\mu}$, and ψ the harmonic Beltrami form. Writing $f = f_\mu$, $f_t = f_{\mu_t}$, our objective is to show that

$$\frac{d}{dt}|_{t=0} \int_\Sigma e(\mu_t, f_t)dA_{\mu_t} = -4\,\mathrm{Re}\langle \Phi(f), \psi\rangle.$$

By energy minimization, we have the inequalities

$$\int_\Sigma e(\mu_t, f_t)dA_{\mu_t} - \int_\Sigma e(\mu, f)dA_\mu \leq \int_\Sigma e(\mu_t, f)dA_{\mu_t} - \int_\Sigma e(\mu, f)dA_\mu$$

and

$$\int_\Sigma e(\mu_t, f_t)dA_{\mu_t} - \int_\Sigma e(\mu, f)dA_\mu \geq \int_\Sigma e(\mu_t, f_t)dA_{\mu_t} - \int_\Sigma e(\mu, f_t)dA_\mu.$$

Thus, it suffices to divide by t and take the limit on the two expressions on the right. We expand

$$\frac{\sqrt{|\mu|_t}\mu_t^{\alpha\beta} - \sqrt{|\mu|}\mu^{\alpha\beta}}{t} = \Big(\frac{\sqrt{|\mu| - |\phi_t|^2} - \sqrt{|\mu|}}{t}\Big)\mu^{\alpha\beta} + \sqrt{|\mu| - |\phi_t|^2}(\dot{\mu}^{\alpha\beta} + \epsilon^{\alpha\beta}(t)). \tag{4.11}$$

Twice the first integrand is obtained by hitting the expression above with $e_{\alpha\beta}(f)$, and the second with $e_{\alpha\beta}(f_t)$, which both converge to $e_{\alpha\beta}(f)$ pointwise as $t \to 0$.

The first term in (4.11) converges to 0, while the second one to $\sqrt{|\mu|}\dot{\mu}^{\alpha\beta}$. Using the relation $\dot{\mu}^{\alpha\beta} = -\mu^{-2}\dot{\mu}_{\alpha\beta}$, the integrand converges pointwise to

$$-\frac{1}{2\mu}(\dot{\mu}_{11}e_{11}+\dot{\mu}_{22}e_{22}+2\dot{\mu}_{12}e_{12}) = -\frac{1}{2\mu}(\dot{\mu}_{11}(e_{11}-e_{22})+2\dot{\mu}_{12}e_{12}) = -4\,\mathrm{Re}\,\Phi(f)\psi.$$

Therefore, it suffices to show the integrands are always bounded above by an integrable quantity, for then we can justify an application of the dominated convergence theorem. We have the expression

$$\frac{\sqrt{|\mu|-|\phi_t|^2}-\sqrt{|\mu|}}{t} = \frac{-|\phi_t|^2}{t(\sqrt{|\mu|-|\phi_t|^2}+\sqrt{|\mu|})} = \frac{-t\Big|(\dot{\mu}_{11}+\dot{\epsilon}_{11}(t))-i(\dot{\mu}_{12}+\dot{\epsilon}_{12}(t))\Big|^2}{\sqrt{|\mu|-|\phi_t|^2}+\sqrt{|\mu|}}.$$

If a_t is the -1 Laurent coefficient in the coordinates on the punctured disk for the quadratic differential associated to $\dot{\mu}+\epsilon(t)$, then because $\epsilon(t)\to 0$, a_t converges to a constant a_0. Therefore, for t small enough, using the expression (4.8), we see that for small t, the first term in (4.11) decays at most like

$$-t(2|a_0|+1)^2y^2e^{-2y}y^2 \sim y^4e^{-2y}.$$

By similar reasoning, the second term in (4.11) decays like y^2e^{-2y}. Thus, it suffices to bound $e_{\alpha\beta}(f_t)$ by a constant. The one obstruction to applying Lemma 4.2.7 for f_t is that $e_{\alpha\beta}$ depends on the cusp coordinates for μ rather than μ_t. The argument of Lemma 4.2.16 shows that the μ and μ_0 coordinates are related by a mapping that is asymptotically Lipschitz—the Hölder exponent in Mori's theorem is the reciprocal of the quasiconformal dilatation, which is tending to 1. Thus, once high enough in the cusp, we have the same bound in the μ_0-coordinates. From the discussion above, the result follows. □

Remark 4.2.17. We have worked in negative curvature, but the proof carries through in non-positive curvature if total energies of harmonic maps to (X,ν) are unique, and if one can locally continuously associate source metrics to harmonic maps in the C^∞ topology. Our proof uses the existence of a single finite energy harmonic map and an energy minimizing property. In non-positive curvature we have energy minimization, and for existence one can modify our [Sag19, Proposition 3.8] to include NPC manifolds.

Now we turn to the main result of this section. Assume the monodromy is hyperbolic.

Proof of Proposition 4.2.1. As above, let $\mu_t = \mu + t\dot{\mu} + t\epsilon(t)$ be the variation with μ. Similarly, denote by ϕ_t and ψ the family of holomorphic quadratic differentials and the harmonic Beltrami form associated to this path respectively. Put $f = f_\mu$, $f_t = f_{\mu_t}$, $h = h_\mu$, $h_t = h_{\mu_t}$, recalling that f and h correspond to the representations $\rho_1 : \Gamma \to \mathrm{PSL}(2,\mathbb{R})$, $\rho_2 : \Gamma \to G$ respectively. From the energy minimization Lemma 4.2.14, we deduce

$$\int_\Sigma e(\mu,h) - e(\mu,f_t)dA_\mu \leq \mathcal{E}(\mu) \leq \int_\Sigma e(\mu,h_t) - e(\mu,f)dA_\mu$$

and

$$\int_\Sigma e(\mu_t,h_t) - e(\mu_t,f)dA_{\mu_t} \leq \mathcal{E}(\mu_t) \leq \int_\Sigma e(\mu_t,h) - e(\mu_t,f_t)dA_{\mu_t}.$$

We cannot apply Lemma 4.2.3 to an expression like $e(\mu,f_t)$ directly, but from the expression for the cusp coordinates (4.2), it follows that when changing $\sqrt{|\mu|}e(\mu,f)$ to $\sqrt{|\mu_t|}e(\mu_t,f_\mu)$, the energy density is asymptotically multiplied by τ_{μ_t}/τ_μ. Therefore, Lemma 4.2.3 does imply that every integral above is finite. Furthermore, it makes sense to manipulate these integrals, and so the difference $\mathcal{E}(\mu_t) - \mathcal{E}(\mu)$ is bounded above by

$$\int_\Sigma (\sqrt{|\mu_t|}e(\mu_t,h) - \sqrt{|\mu|}e(\mu,h)) - (\sqrt{|\mu_t|}e(\mu_t,f_t) - \sqrt{|\mu|}e(\mu,f_t))dzd\overline{z}$$

and bounded below by

$$\int_\Sigma (\sqrt{|\mu_t|}e(\mu_t,h_t) - \sqrt{|\mu|}e(\mu,h_t)) - (\sqrt{|\mu_t|}e(\mu_t,f) - \sqrt{|\mu|}e(\mu,f))dzd\overline{z}.$$

Dividing by t, the local computations from the proof of Proposition 4.1.6 work out almost the same. We write out the details for the upper bound and leave the lower bound to the reader. The relevant quotient may be expressed

$$\frac{1}{2}\Big(\Big(\frac{\sqrt{|\mu| - |\phi_t|^2} - \sqrt{|\mu|}}{t}\Big)\mu^{\alpha\beta} + \sqrt{|\mu| - |\phi_t|^2}(\dot{\mu}^{\alpha\beta} + \epsilon^{\alpha\beta}(t))\Big)(e_{\alpha\beta}(h) - e_{\alpha\beta}(f_t)).$$

The quotient converges pointwise to $-4\,\mathrm{Re}(\Phi(h) - \Phi(f))\cdot\overline{\psi}$, so we are left to bound it by an integrable quantity that does not depend on t. Lemma 4.2.12 and the fact that the change of coordinates is asymptotically conformal shows that the second term in the product is uniformly integrable. And it was shown in the proof of the previous proposition that the first term in the product decays exponentially. □

Different twist parameters

The proof of Proposition 4.2.1 is complete for finite energy harmonic maps, and infinite energy harmonic maps with $\theta = 0$. We sketch the necessary adjustments for the remaining harmonic maps f_μ^θ. Fix $\theta \in \mathbb{R}$. To construct the harmonic map f_μ^θ, we define the fractional Dehn twist to be the map in the cusp coordinates defined by

$$x + iy \mapsto x + \theta_i y + iy.$$

We then postcompose with the approximation maps f_r and take limits on these maps as (see [Sag19, Section 5.1]).

- The inequality in Lemma 4.2.3 becomes

  $$|e(\mu, f_\mu^\theta) - \Lambda(\theta)\ell(\rho(\gamma))^2/2\tau^2| < Ce^{-cy}.$$

  The proof is the same, except the Hopf differential has expression at the cusp according to [Sag19, Theorem 1.1].

- For continuity of harmonic maps on compacta, the main step is to find bounds on the total energy of φ_μ^θ on compacta. This is a consequence of the chain rule, since in the coordinates (4.2), the derivative matrix of the fractional Dehn twist is simply

  $$\begin{pmatrix} 1 & \theta \\ 0 & 1 \end{pmatrix}. \tag{4.12}$$

- For the energy minimizing property, we can say that a map converges at ∞ to α^θ if after pulling back to conformal cylinders approaching the cusps, the harmonic maps converge to the composition of a projection onto the geodesic with a fractional Dehn twist on the cylinder. We then have the analogue of Lemma 4.2.14, which we can show the harmonic maps f^θ satisfy.

We leave the rest of the details to the reader.

4.3 Maximal surfaces: existence, uniqueness, deformations

With Proposition 4.2.1 in hand, we prove the main theorems. For convenience we assume the twist parameter is zero—the proof has no dependence on it. The existence result is immediate from the next proposition.

Proposition 4.3.1. *The functional $\mathcal{E} = \mathcal{E}_{\rho_1,\rho_2} : \mathcal{T}(\Gamma) \to \mathbb{R}$ is proper if and only if ρ_1 almost strictly dominates ρ_2.*

Indeed, by Proposition 4.2.1, properness implies the existence of a maximal space-like immersion. Thus, the proof of the existence result reduces to showing that almost strict domination implies properness.

Mapping class groups

Preparing for the proof of Theorem 4A, we review some Teichmüller theory. Set $\mathrm{Diff}^+(\Sigma)$ to be the group of C^∞ orientation preserving diffeomorphisms of Σ, equipped with the C^∞ topology. Denote by $\mathrm{Diff}^+_0(\Sigma)$ the normal subgroup consisting of maps that are isotopic to the identity. The mapping class group of Σ is defined

$$\mathrm{MCG}(\Sigma) := \pi_0(\mathrm{Diff}^+(S)) = \mathrm{Diff}^+(S)/\mathrm{Diff}^+_0(S).$$

Given a collection of boundary lengths $\ell = (\ell_1, \dots, \ell_n) \in \mathbb{R}^n$, we can consider the Teichmüller space $\mathcal{T}_\ell(\Gamma)$ of surfaces with punctures and geodesic boundary with lengths determined by $(\ell_1, \dots, \ell_n)$. This is not a relative representation space but a union of them. The mapping class group acts on Teichmüller space by pulling back classes of hyperbolic metrics:

$$[\varphi] \cdot [\mu] \mapsto [\varphi^* \mu]. \tag{4.13}$$

The action of the mapping class group is properly discontinuous and the quotient is the moduli space of hyperbolic surfaces.

There is a wealth of metrics on Teichmüller space, and the mapping class group acts by isometries on a number of them. In the work below, we set $d_{\mathcal{T}}$ to be any metric distance function on Teichmüller space on which the mapping class group acts by isometries. One example is the Teichmüller distance.

Selecting a basepoint $z \in \Sigma$, there is an action of the mapping class group on $\pi_1(\Sigma, z)$, which identifies with Γ. Given $\varphi \in \mathrm{Diff}^+(\Sigma)$ such that $\varphi(z) = z$, the induced map $\varphi_* : \pi_1(\Sigma, z) \to \pi_1(\Sigma, z)$ is an automorphism. While a general φ may not fix z, one can find a different $\varphi_1 \in \mathrm{Diff}^+(\Sigma)$ that does fix z and is isotopic to φ. The isotopy gives rise to a path γ_1 from $\varphi(z)$ to z. If $\varphi_2 \in \mathrm{Diff}^+(\Sigma)$ also fixes z and is isotopic to φ, then we get another path γ_2 from $\varphi(z)$ to z. The automorphisms $(\varphi_1)_*$ and $(\varphi_2)_*$ of $\pi_1(\Sigma, z)$ differ by the inner automorphism corresponding to the conjugation by the class of $\gamma_1 \cdot \overline{\gamma}_2$. This association furnishes an injective homomorphism from

$$\mathrm{MCG}(\Sigma) \to \mathrm{Out}(\pi_1(\Sigma)) = \mathrm{Aut}(\pi_1(\Sigma))/\mathrm{Inn}(\pi_1(\Sigma)).$$

Through this injective mapping, the mapping class group acts on representation spaces: we can precompose a representative of a representation with a representative of an element in $\mathrm{Out}(\pi_1(\Sigma))$, and then take the corresponding equivalence class. On a Teichmüller space, this agrees with the action (4.13). For this action, we use the similar notation

$$[\varphi] \cdot [\rho] \mapsto [\varphi^* \rho] = [\rho \circ \varphi_*].$$

Note that it does not in general preserve relative representation spaces: a mapping class may permute the punctures. Finally, we record that $\mathrm{MCG}(\Sigma)$ acts equivariantly with respect to

- length spectrum: for a representation ρ and $[\varphi] \in \mathrm{MCG}(\Sigma)$, $\ell(\varphi^*\rho(\gamma)) = \ell(\rho(\varphi_*(\gamma)))$. If μ is a hyperbolic metric we set $\ell_\mu(\gamma)$ to be the μ-length of the geodesic representative of γ in (Σ, μ). A restatement of the above is $\ell_{\varphi^*\mu}(\gamma) = \ell_\mu(\varphi(\gamma))$.
- Energy of harmonic maps: if $f : (\tilde{\Sigma}, \tilde{\mu}) \to (X, \nu)$ is a ρ-equivariant map, then
$$e(\mu, f_\mu) = e(\varphi^*\mu, f_\mu \circ \varphi). \tag{4.14}$$
Furthermore, if f is harmonic, then $f \circ \varphi : (\tilde{\Sigma}, \tilde{\varphi}^*\tilde{\mu}) \to (X, \nu)$ is a $\varphi^*\rho$-equivariant harmonic map.

Existence of maximal surfaces

Suppose a simple closed curve ξ is either a geodesic boundary component or a horocycle in a hyperbolic surface. If ξ is a boundary, let $C_d(\xi)$ denote the collar around ξ consisting of points with distance to ξ at most d. If it is a horocycle, put $C_d(\xi)$ to be the union of the d-collar and the enclosed cusp (both are defined for suitably small d).

We now begin the proof of properness. Suppose ρ_1 almost strictly dominates ρ_2. We assume there is a single peripheral curve—there are no substantial changes in the case of many cusps—and we set ξ to be either the geodesic boundary component of the convex core of $\mathbb{H}/\rho_1(\Gamma)$, or a deep horocycle of $\mathbb{H}/\rho_1(\Gamma)$, depending on whether the monodromy is hyperbolic or parabolic. Let g be an optimal (ρ_1, ρ_2)-equivariant map. If the image of the peripheral curve under ρ_1 and ρ_2 is hyperbolic, then g has constant speed on the boundary components. Note that $g \circ h_\mu$ is ρ_2-equivariant, and in the hyperbolic case, converges at ∞ to a projection onto the geodesic at infinity in the sense of Definition 4.2.13. Thus, from Lemma 4.2.14 we have

$$\mathcal{E}(\mu) = \int_\Sigma e(\mu, h_\mu) - e(\mu, f_\mu) dA_\mu \geq \int_\Sigma e(\mu, h_\mu) - e(\mu, g \circ h_\mu) dA_\mu.$$

One of the main advantages of replacing f_μ with $g \circ h_\mu$ is that the integrand is now non-negative. Moving toward the proof of properness, let $([\mu_n])_{n=1}^\infty \subset \mathcal{T}(\Gamma)$ be such that

$$\mathcal{E}(\mu_n) \leq K$$

for all n, for some large K. It suffices to show $[\mu_n]$ converges in $\mathcal{T}(\Gamma)$ along a subsequence. Structurally, our proof is similar to that of the classical existence result for minimal surfaces in hyperbolic manifolds (see [SY79] and also the more modern paper [GW07]): we first prove 1) compactness in moduli space, and then 2) extend to compactness in Teichmüller space.

Toward 1), we show there is a $\delta > 0$ such that for all non-peripheral simple closed curves in Γ, we have $\ell_{\mu_n}(\gamma) \geq \delta$. We argue by contradiction: suppose this is not the case, so that there is a sequence of non-peripheral simple closed curves $(\gamma_n)_{n=1}^{\infty} \subset \Gamma$ such that

$$\ell_n := \ell_{\mu_n}(\gamma_n) \to 0$$

as $n \to \infty$. By the regular collar lemma, we know that in (Σ, μ_n), the geodesic representing γ_n is enclosed by a collar C_n of width $w_n = \operatorname{arcsinh}((\sinh \ell_n)^{-1})$. Up to a conjugation of the holonomy, C_n is conformally equivalent to $\{x + iy \in \mathbb{C} : 0 \leq x \leq w_n, 0 \leq y \leq 1\}$ with the lines $\{y = 0\}$ and $\{y = 1\}$ identified.

Let $f_n = f_{\mu_n}$, $h_n = h_{\mu_n}$. Henceforward, conformally modify the metric μ_n to be flat in C_n. Our control over the energy functional depends on our knowledge of the local Lipschitz constant of the map g. This in turns relies on the image of the harmonic map h_n, in particular how far it takes C_n into $C_d(\xi)$. For all $x, 0 < t < d$ in question, set $A_x^t = \{\theta \in \{x\} \times S^1 : h_n(x, y) \notin C_t(\xi)\}$. We also write

$$\operatorname{Lip}(g, t) = \max_{x \in C(\mathbb{H}/\rho_1(\Gamma)) \backslash C_t(\xi)} \operatorname{Lip}_x(g).$$

We have the inequalities

$$\begin{aligned}
\mathcal{E}(\mu_n)(\sigma_n) &\geq \int_{C_n} e(\mu_n, h_n) - e(\mu_n, g \circ h_n) dA_{\mu_n} \\
&= \int_0^{w_n} \int_0^1 e(\mu_n, h_n) - e(\mu_n, g \circ h_n) dy ds \\
&\geq w_n \min_x \int_{A_x^t} e(\mu_n, h_n) - e(\mu_n, g \circ h_n) dy \\
&\geq w_n \min_x (1 - \operatorname{Lip}(g, t)^2) \int_{A_x^t} e(\mu_n, h_n) dy \\
&\geq \frac{w_n}{2} \min_x (1 - \operatorname{Lip}(g, t)^2) \int_{A_x^t} \Big| \frac{\partial h_n(x, y)}{\partial y} \Big|^2 dy \\
&\geq \frac{w_n}{2} \min_x (1 - \operatorname{Lip}(g, t)^2) (\ell(\mathrm{A}_x^t))^{-1} \Big(\int_{A_x^t} \Big| \frac{\partial h_n(x, y)}{\partial y} \Big| dy \Big)^2 \\
&\geq \frac{w_n}{2} \min_x (1 - \operatorname{Lip}(g, t)^2) \Big(\int_{A_x^t} \Big| \frac{\partial h_n(x, y)}{\partial y} \Big| dy \Big)^2 .
\end{aligned}$$

Here $\ell(A_x^t)$ is the length of the (possibly broken) segment $A_{x,t}$, which is clearly bounded above by 1. Fuchsian representations have discrete length spectrum, and hence there is a $\kappa > 0$ such that $\ell(\rho_1(\gamma)) \geq \kappa$ for all $\gamma \in \Gamma$.

Lemma 4.3.2. *Set $t = d/2$. Then for any x, the inequality*

$$\int_{A_x^{d/2}} \left|\frac{\partial h_n(x,y)}{\partial y}\right| dy \geq \min\{d/2, \kappa\} =: k$$

holds.

Proof. Let α be any core curve of the form $\{x\} \times S^1$. If $h_n(\alpha)$ always remains outside $C_{d/2}(\xi)$, then

$$\int_{A_x^t} \left|\frac{\partial h_n(x,y)}{\partial y}\right| dy = \int_0^1 \left|\frac{\partial h_n(x,y)}{\partial y}\right| dy \geq \ell(\rho_1(\gamma_n)) \geq \kappa.$$

Thus, we assume $h_n(\alpha)$ intersects $C_{d/2}(\xi)$. We parametrize $h_n(x,\cdot)$ by arc length, so that the integral in question measures the hyperbolic length on $\mathbb{H}/\rho_1(\Gamma)$ of the segment of $h_n(\alpha)$ that does not enter $C_{d/2}(\xi)$. If this length is ever less than $d/2$, then the curve $h_n(\alpha)$ is contained in $C_d(\xi)$. Phrased differently, it is a simple closed curve contained in an embedded cylinder. There is only one such homotopy class of curves, namely the homotopy class of the boundary geodesic. This situation is impossible, since $h_n(\alpha)$ is homotopic to a non-peripheral simple closed curve, and hence the length is at least $d/2$. □

Returning to the inequalities above, we now have

$$\mathcal{E}(\mu_n) \geq \frac{w_n}{2}(1 - \mathrm{Lip}(g, d/2)^2)k^2.$$

We thus find $\mathcal{E}(\mu_n) \to \infty$ as $n \to \infty$, which violates the uniform upper bound.

The lower bound on the length spectrum shows the metrics $(\mu_n)_{n=1}^\infty$ satisfy Mumford's compactness criteria [Mum71] and hence project under the action of the mapping class group to a compact subset of the moduli space.

Now we promote to compactness in Teichmüller space. By compactness in moduli space, after passing to a subsequence of the μ_n's, there exists a sequence of mapping class group representatives $(\psi_n)_{n=1}^\infty$ such that $\psi_n^*\mu_n$ converges as $n \to \infty$ to a hyperbolic metric μ_∞.

Lemma 4.3.3. *Upon passing to a further subsequence, $([\psi_n])_{n=1}^\infty \subset MCG(\Sigma)$ is eventually constant.*

The main thrust of the proof is to show that for each non-peripheral $\gamma \in \Gamma$, $\ell(\rho_1 \circ \psi_n(\gamma))$ is uniformly bounded above in n. Our proof of this claim is by contradiction and similar in nature to the argument above. Suppose there is a class of curves $\gamma \in \Gamma$ such that $\ell(\rho_1 \circ \psi_n(\gamma)) \to \infty$ as $n \to \infty$. Modify the notation: set $h_n = h_{\mu_n} \circ \psi_n$ and $f_{\mu_n} \circ \psi_n$. These are harmonic for $\psi_n^*\rho_1$ and $\psi_n^*\rho_2$ respectively, and have the correct boundary behaviour according to [Sag19, Theorem 1.1]. From (4.14), we also have the equality

$$\mathcal{E}_{\rho_1,\rho_2}(\mu_n) = \mathcal{E}_{\psi_n^*\rho_1,\psi_n^*\rho_2}(\psi_n^*\mu_n).$$

On each $(\Sigma, \psi_n^*\mu_n)$, there is a collar C_n of finite width w_n around γ, and the C_n's converge to some collar C_∞ in (Σ, μ_∞) of width $w_\infty < \infty$. As before, we perturb to a flat metric and parametrize C_n by $[0, w_n] \times S^1$. We note that g is $(\psi_n^*\rho_1, \psi_n^*\rho_2)$-equivariant. Therefore, we can apply our previous reasoning to see that for n large enough, any small enough x, and $t \in [0, w_n]$,

$$\mathcal{E}_{\rho_1,\rho_2}(\mu_n) = \mathcal{E}_{\psi_n^*\rho_1,\psi_n^*\rho_2}(\psi_n^*\mu_n) \geq \frac{w_\infty}{4} \min_x (1 - \mathrm{Lip}(g,t)^2)\Big(\int_{A_x^t} \Big|\frac{\partial h_n(x,y)}{\partial y}\Big| dy\Big)^2, \tag{4.15}$$

where A_x^t is defined as above. This leads us to the analogue of Lemma 4.3.2.

Lemma 4.3.4. *Set $t = d/2$. Then, independent of the choice of x,*

$$\int_{A_x^{d/2}} \Big|\frac{\partial h_n(x,y)}{\partial y}\Big| dy \to \infty$$

as $n \to \infty$.

Proof. Choose any simple closed curve α of the form $\{x\} \times S^1$. Parametrize $h_n(x, \cdot)$ by arc length so that the integral above returns the length of the segment of $h_n(\alpha)$ that does not enter $C_{d/2}(\xi)$. If $h_n(\alpha)$ does not enter $C_{d/2}(\xi)$, then we can argue as we did in the proof of Lemma 4.3.2, bounding the length below by the $\ell(\psi_n^*\rho_1(\gamma))$, which blows up. Thus, we restrict our discussion to α such that $h_n(\alpha)$ intersects $C_{d/2}(\xi)$. Let us first assume there is a positive integer K such that every $h_n(\alpha)$ enters $C_{d/2}(\xi)$ at most K times. Of course, $h_n(\alpha)$ cannot live entirely in $C_{d/2}(\xi)$, for then it lies in the wrong homotopy class. We construct a new curve as follows: if we choose a basepoint not within $C_{d/2}(\xi)$, then every time $h_n(\alpha)$ enters $C_{d/2}(\xi)$, there is a corresponding point at which it exits. We plan to erase the segment of $h_n(\alpha)$ that connects these two points and replace it with the one of the two possible paths on $\partial C_{d/2}(\xi)$, say, β_1 and β_2. If $h_n(\alpha)$ spends some time going in a path along the boundary circle when it is entering or when it is exiting (we count just touching it once as an exit), we connect the endpoint of the entrance path to the starting point of the exiting path and concatenate with the path going along the boundary.

The question is: which path to take? We argue there is a choice so that the homotopy class of the new curve is the same as that of $h_n(\alpha)$. To this end, orient the boundary circle so that there is a "left" and a "right" path. Arbitrarily choose one of the two paths connecting the entrance and the exit point, say, β_1, and consider the loop in the cylinder obtained by concatenating this path with the piece of $h_n(\alpha) \cap C_{d/2}(\xi)$ that connects the endpoints. This is a simple closed curve in a cylinder, and hence there are three possible homotopy classes: trivial, non-trivial and left oriented, and non-trivial and right oriented. If the class is trivial, then using this path β_1 will work. If it is non-trivial and left oriented, then β_1 must be the right oriented path. Thus, if we use the opposite path β_2, it will cancel the homotopy class, and therefore the new path will indeed be homotopic to $h_n(\alpha)$. The right oriented case is similar.

Each replacement adds at most $\ell(\partial C_{d/2}(\xi))$ to the length of the path. Thus, the length of the new path is (quite crudely) bounded above by

$$\ell_{d/2}(h_n(\alpha)) + K\ell(\partial C_{d/2}(\xi)).$$

Therefore,

$$\ell_{d/2}(h_n(\alpha)) \geq \ell(\psi_n^* \rho_1(\gamma)) - K\ell(\partial C_{d/2}(\xi))$$

which tends to ∞ as $n \to \infty$, and moreover establishes the result of the lemma.

We are left to consider the case of an unbounded number of crossings into $C_{d/2}(\xi)$ as we take $n \to \infty$. The idea is that each time $h_n(\alpha)$ crosses $\partial C_{d/2}(\xi)$, there is a corresponding "down-crossing," a curve that connects the point at which it exits to the new point of entry. If each down crossing is denoted by c_j^n, then

$$\int_{A_x^{d/2}} \left| \frac{\partial h_n(x, y}{\partial y} \right| dy \geq \sum_j \ell(c_j^n).$$

Clearly, if the limit of these sums is infinite, then we are done. Hence, we assume the total length due to down-crossings is finite. This implies that there is an even integer $K > 0$ such that, as we take $n \to \infty$, there are at most K down-crossing that exit C_d. Again using our chosen basepoint, let b_1 be the first crossing into $C_{d/2}(\xi)$, and then let b_2 be the first crossing out of $C_{d/2}(\xi)$ that exits C_d. Set b_3 to be the next entry point into $C_{d/2}(\xi)$, and b_4 the next exit point from $C_d(\xi)$. We end up with $K_n \leq K$ points $b_1, \ldots b_{K_n}$. Replace the segments of $h_n(\alpha)$ between b_i and b_{i+1}

with the correct arc on $\partial C_{d/2}(\xi)$ that does not change the homotopy class, where $i = 1, 3, \ldots, K_n - 1 \leq K - 1$. As above, we end up with a curve of length at most

$$\ell_{d/2}(h_n(\alpha)) + K\ell(\partial C_{d/2}(\xi))$$

and homotopic to $\ell(\psi_n^*\rho_1(\gamma))$. Using the minimizing property of geodesics once again, we have

$$\ell_{d/2}(h_n(\alpha)) \geq \ell(\psi_n^*\rho_1(\gamma)) - K\ell(\partial C_{d/2}(\xi)) \to \infty,$$

and the resolution of this final case completes the proof. □

Returning to (4.15), this lemma shows $\mathcal{E}(\mu_n) \to \infty$, which is a contradiction. We can now conclude that each sequence $\ell(\psi_n^*\rho_1(\gamma))$ remains bounded above.

It is well-understood that the boundedness of the length spectrum implies that $(\psi_n^*\rho_1)$ converges along a subsequence to some new Fuchsian representation $\rho_\infty : \Gamma \to \mathrm{PSL}_2(\mathbb{R})$. We now restrict the $[\mu_n], [\psi_n]$ to this chosen subsequence. Since the mapping class group acts properly discontinuously on the Teichmüller space, it follows that $[\psi_n^*\rho_1] = [\rho_\infty]$ for all n large enough. In particular, ψ_n is isotopic to ψ_m for all n, m large enough. This completes the proof of Lemma 4.3.3.

To finish the proof of Proposition 4.3.1, we know there is an N such that $[\psi_n] = [\psi_N]$ for all $n \geq N$. As the mapping class group acts by isometries with respect to the metric $d_{\mathcal{T}}$,

$$d_{\mathcal{T}}([\mu_n], [(\psi_N^{-1})^*\mu_\infty]) = d_{\mathcal{T}}([\psi_N^*\mu_n], [\mu_\infty]) \to 0,$$

and hence $[\mu_n]$ converges to $[(\psi_N^{-1})^*\mu_\infty]$ as $n \to \infty$. Proposition 4.3.1 is proved.

Uniqueness of critical points

The uniqueness statement in Theorem 4A amounts to showing that critical points of $\mathcal{E}$ are unique. The main step in the uniqueness proof for closed surfaces [Tho17, Theorem 1] is a local computation ([Tho17, Lemma 2.4]) that goes through in our setting. For this reason, we omit the proof of uniqueness and invite the reader to see [Tho17]. The only real difference is that we must address convergence of various integrals, and use our Lemma 4.2.14 instead of the usual energy minimizing property. We've shown this type of calculation throughout the chapter, so we feel comfortable leaving it to the interested reader.

The maps Ψ^θ

We now begin the proof of Theorem 4B. Generalizing the map Ψ from [Tho17, subsection 2.3]), we define the map

$$\Psi^\theta : \mathcal{T}(\Gamma) \times \mathrm{Rep}_{\mathbf{c}}(\Gamma, G) \to \mathrm{ASD}_{\mathbf{c}}(\Gamma, G) \subset \mathcal{T}_{\mathbf{c}}(\Gamma) \times \mathrm{Rep}_{\mathbf{c}}(\Gamma, G)$$

as follows. When $\mathbf{c}$ has no hyperbolic classes, we implicitly assume there is no θ. Begin with a hyperbolic surface (Σ, μ) and a reductive representation $\rho_2 : \Gamma \to \mathrm{PSL}(2, \mathbb{R})$. Associated to this representation, we take a ρ_2-equivariant harmonic map with twist parameter θ, $f^\theta : (\tilde{\Sigma}, \tilde{\mu}) \to (X, \nu)$. We proved in [Sag19, Theorem 1.4] that there exists a unique Fuchsian representation $\rho_1 : \Gamma \to \mathrm{PSL}(2, \mathbb{R})$ and a unique ρ_1-equivariant harmonic diffeomorphism $h^\theta : (\tilde{\Sigma}, \mu) \to (\mathbb{H}, \sigma)$ that has Hopf differential $\Phi(h^\theta) = \Phi(f^\theta)$. From [Sag19, Proposition 3.13], the mapping $F = (h^\theta, f^\theta)$ is a spacelike maximal immersion into the pseudo-Riemannian product, and hence ρ_1 almost strictly dominates ρ_2. Ψ^θ is defined by

$$\Psi^\theta([\mu], [\rho_2]) = ([\rho_1], [\rho_2]).$$

It is clear that the map is fiberwise in the sense of the statement of Theorem 4B. Theorem 4A shows Ψ^θ is bijection, and the content of Theorem 4B is that Ψ^θ is a homeomorphism.

Remark 4.3.5. We do not know if $\Psi^\theta = \Psi^{\theta'}$ for $\theta \neq \theta'$!

Proof of Theorem 4B: continuity

From now on we refine our notation: when the representation ρ is not implicit, the ρ-equivariant harmonic map for a metric μ with zero twist parameter will be denoted f_μ^ρ, unless specified otherwise.

The proof of continuity is almost identical for each twist parameter, so we work only with $\Psi = \Psi^0$. Let us also assume for the remainder of this section that there is a single peripheral ζ. Lifting various equivalence relations, Ψ is described by a composition

$$(\mu, \rho_2) \mapsto (f_\mu^{\rho_2}, \rho_2) \mapsto (\Phi(f_\mu^{\rho_2}), \rho_2) \mapsto (\rho_1, \rho_2),$$

where ρ_1 is the holonomy of the hyperbolic metric μ' on Σ such that the associated harmonic map from $(\Sigma, \mu) \to (\Sigma, \mu')$ has Hopf differential $\Phi(f_\mu^{\rho_2})$. [Sag19, Theorem 1.4] implies the association from $\Phi(f_\mu^{\rho_2}) \mapsto \rho_1$ is continuous. Here, the topology on the space of holomorphic quadratic differentials is the Fréchet topology coming from taking L^1-norms over a compact exhaustion. To show continuity of $\mu \mapsto f_\mu^{\rho_2} \mapsto \Phi(f_\mu^{\rho_2})$ with respect to this topology, note that, by the constructions of Section 4.2, we already have continuity in the Teichmüller coordinate. Continuity will thus follow from the results below.

Lemma 4.3.6. *Suppose representations $(\rho_n)_{n=1}^\infty$ converge to an irreducible ρ in Hom(Γ, G), and all such representations project to Rep$_{\mathbf{c}}(\Gamma, G)$. Then $f_\mu^{\rho_n}$ converges to f_μ^ρ in the C^∞ sense on compacta as $n \to \infty$.*

Below, we work in fixed local sections of the principal bundles $\chi_{\mathbf{c}}(\Gamma, G) \to \mathrm{Rep}_{\mathbf{c}}(\Gamma, G)$ $\chi_{\mathbf{c}}(\Gamma, \mathrm{PSL}(2,\mathbb{R})) \to \mathrm{Rep}_{\mathbf{c}}(\Gamma, \mathrm{PSL}(2,\mathbb{R}))$ (having assumed in the introduction that such things exist). We choose a basepoint for the π_1 so that we can view these local systems through their holonomies, which are genuine representations.

Proof. We may assume each ρ_n is irreducible and that there is a single cusp. The parabolic and elliptic cases are much simpler than the hyperbolic case, so we treat only the latter. Set $f_n = f_\mu^{\rho_n}$, $f = f_\mu^\rho$ and let φ_n, φ be the approximation maps for $f_\mu^{\rho_n}$, f_μ^ρ from 3.1. These maps project a cusp neighbourhood onto a geodesic, but now the geodesic is varying with ρ_n, and converging in the Gromov-Hausdorff sense to the geodesic for ρ. If the φ_n's can be chosen to vary smoothly, then we have a uniform bound on the total energies (recall they depend on a choice of a constant speed map onto their geodesic). As discussed in Section 4.2 (in particular the proof of Lemma 4.2.11), this yields a uniform total energy bound for f_n on compacta. Via the argument described early in Section 4.2, the maps f_n C^∞-converge on compacta along a subsequence to some limiting map f^∞. The mapping is necessarily ρ-equivariant and harmonic. To see that $f^\infty = f$, recall that f is characterized by the fact that its Hopf differential has a pole of order 2 at the cusp and the residue has a specific complex argument. To check this property, from the uniqueness in [Sag19, Theorem 1.1] it suffices to check $d(f_\infty, \varphi) < \infty$, and the standard contradiction argument will also show there is no need to pass to a subsequence. As remarked previously, the proof of [Sag19, Lemma 5.2] shows $d(f_n, \varphi_n)$ is maximized on $\partial\Sigma_2$. If n_k is the subsequence, then taking $k \to \infty$, via compactness we do win

$$d(f_\infty, \varphi_\infty) \leq \limsup_{k\to\infty} \max_{\partial\Sigma_2} d(f_{n_k}, \varphi_{n_k}) \leq \max_{\partial\Sigma_2} d(f_\infty, \varphi_\infty) + 1 < \infty.$$

We are left to argue that one can choose the φ_n so that $\varphi_n \to \varphi$ in the C^∞ sense. We realize ρ, ρ_n as monodromies of flat connections ∇, ∇_n respectively on an X-bundle

E over Σ with structure group G. This is possible when all ρ_n lie in the same component of the relative representation space ([Lab13, Corollay 6.1.2] generalizes to relative representation spaces), which we are free to assume. The pullback bundle with respect to the universal covering $\tilde{\Sigma} \to \Sigma$ identifies with the trivial bundle and also pulls ∇, ∇_n back to flat connections $\tilde{\nabla}, \tilde{\nabla}_n$. That is, up to an isomorphism, we have a family of commutative diagrams

$$\begin{array}{ccc} (X \times \tilde{\Sigma}, \tilde{\nabla}_n) & \longrightarrow & (E, \nabla_n) \\ \downarrow & & \downarrow \\ \tilde{\Sigma} & \longrightarrow & \Sigma. \end{array}$$

The bundle E and connections ∇_n are constructed by choosing a good covering of Σ and depend on the local section of the principal bundle $\chi_{\mathbf{c}}(\Gamma, G) \to \mathrm{Rep}_{\mathbf{c}}(\Gamma, G)$ (see the proof of [Lab13, Lemma 6.1.1]). We can build a good covering by glueing a good covering of a relatively compact tubular neighbourhood of (Σ_2, μ) with a good covering of a tubular neighbourhood of $(\Sigma \backslash \Sigma_2, \mu)$. This ensures a lower bound on the μ-radius inside Σ_2 (note that we cannot do this on all of Σ). With this constraint, the local systems can be prescribed so that $\nabla_n \to \nabla$ in the C^∞ sense on Σ_2 in the affine space of connections.

The map φ induces a section s of the bundle (E, ∇), which can also be seen as a section s_n of (E, ∇_n). With respect to the diagram above, s_n pulls back to a ρ_n-equivariant map $\tilde{\varphi}_n : \mathbb{H} \to X$, and by our comments above, $\tilde{\varphi}_n \to \varphi$ in the C^∞ sense. The maps $\tilde{\varphi}_n$ are most likely not harmonic and may not project onto the geodesic axis of $\rho_n(\zeta)$. However, because they are converging to φ, the geodesic curvature of $\tilde{\varphi}_n(\partial\tilde{\Sigma}_2)$ is tending to 0. Moreover, $\varphi|_{\partial\tilde{\Sigma}_2}$ is arbitrarily close to some parametrization of a geodesic. So, we set φ_n to be the harmonic map with boundary data equal to this nearby geodesic projection. From energy control on $\tilde{\varphi}_n$, we get the same for φ_n, and hence convergence along a subsequence to a harmonic map φ_∞. From the boundary values, we must have $\varphi_\infty = \varphi$, and as usual we can show there is no need to pass to a subsequence. □

For a reducible representation, the harmonic maps are not unique but differ by translations along a geodesic. However, the Hopf differential is unique. We can choose a sequence $\varphi_n \to \varphi$ as in the previous proposition—boundary values are fixed so we do have uniqueness for these harmonic maps. Then we take the harmonic maps f_n built from using the approximating maps φ_n, and the proof above goes through, with the minor modification that we must rescale the approximating maps as in [Sag19, Proposition 5.1]. C^∞ convergence implies the Hopf differentials converge.

Remark 4.3.7. For different twist parameters, we simply precompose every φ_n with a fractional Dehn twist. Using (4.12), the energies are uniformly controlled.

Asymmetric metrics on Teichmüller spaces

In [GK17, Section 8], Guéritaud-Kassel define asymmetric pseudo-metrics on deformation spaces of geometrically finite hyperbolic manifolds. These metrics are natural generalizations of Thurston's asymmetric metric on Teichmüller space. We will use such a metric in the proof of bi-continuity of Ψ^θ.

For $[\rho_1], [\rho_2] \in \mathcal{T}_{\mathfrak{c}}(S_{g,n})$, we set

$$C(\rho_1, \rho_2) = \inf\{\mathrm{Lip}(f) : f : \mathbb{H} \to \mathbb{H} \text{ is } (\rho_1, \rho_2) - \text{equivariant}\}.$$

The metric $d_{Th} : \mathcal{T}_{\mathfrak{c}}(S_{g,n}) \times \mathcal{T}_{\mathfrak{c}}(S_{g,n}) \to \mathbb{R}$ is defined by

$$d_{Th}([\rho_1], [\rho_2]) = \log\Big(C(\rho_1, \rho_2)\frac{\delta(\rho_1)}{\delta(\rho_2)}\Big).$$

Here $\delta(\rho_1)$ is the critical exponent of ρ_1:

$$\delta(\rho_1) = \lim_{R\to\infty} \frac{1}{R} \log \#(j(\Gamma)\cdot p \cap B_p(R)),$$

where p is any point in $\mathbb{H}$ and $B_p(R)$ is the ball of radius r centered at p. Alternatively, $\delta(\rho_1)$ is the Hausdorff dimension of the limit set of ρ_1 in $\partial_\infty \mathbb{H} = S^1$. If all $a_j = 0$, so that all critical exponents are 1 and $\mathcal{T}_{\mathfrak{c}}(S_{g,n})$ is the ordinary Teichmüller space of a surface of finite volume, then this agrees with Thurston's asymmetric metric introduced in [Thu98]. Here asymmetric means that in general,

$$d_{Th}([\rho_1], [\rho_2]) \neq d_{Th}([\rho_2], [\rho_1]).$$

Proposition 4.3.8 (Guéritaud-Kassel, Lemma 8.1 in [GK17])**.** *The function d_{Th} : $\mathcal{T}_{\mathbf{c}}(\Gamma) \times \mathcal{T}_{\mathbf{c}}(\Gamma) \to \mathbb{R}$ is a continuous asymmetric metric.*

Remark 4.3.9. The correction factor $\delta(\rho_1)/\delta(\rho_2)$ is needed for non-negativity. For a general hyperbolic n-manifold, continuity may fail and the generalization may be just an asymmetric pseudo-metric ($d(x, y)$ need not imply $x = y$).

This metric allows us to control translation lengths nicely. Consider a left open ball around ρ_2:

$$B_{Th}(\rho_2, C) = \{[\rho_1] \in \mathcal{T}_{\mathbf{c}}(\Gamma) : d_{Th}(\rho_1, \rho_2) < C\}.$$

If $\rho_1 \in B_{Th}(\rho_2, C)$, there is a (ρ_1, ρ_2)-equivariant Lipschitz g map with Lipschitz constant $< e^C$. Therefore, for every $\gamma \in \Gamma$,

$$d(\rho_2(\gamma)g(z), g(z)) = d(g(\rho_1(\gamma)z), g(z)) < e^C d(\rho_1(\gamma)z, z),$$

and it follows that $\ell(\rho_2(\gamma)) < e^C \ell(\rho_1(\gamma))$.

Proof of Theorem 4B: bi-continuity

The inverse mapping of Ψ^θ, on the Teichmüller side, takes as input a class of an almost strictly dominating pair (ρ_1, ρ_2) and returns the unique minimizer of $\mathcal{E}^\theta_{\rho_1,\rho_2}$. We can follow the same approach as in [Tho17, subsection 2.4] here, adapted to our infinite energy setting.

Let X, Y be metric spaces and $(F_y)_{y\in Y}$ a family of continuous functions $F_y : X \to \mathbb{R}$ depending continuously on y in the compact-open topology. $(F_y)_{y\in Y}$ is said to be uniformly proper if for any $C \in \mathbb{R}$, there exists a compact subset $C \subset X$ such that for all $y \in Y$ and $x \notin K$, we have $F_y(x) > C$. We say that the family $(F_y)_{y\in Y}$ is locally uniformly proper if for all $y_0 \in Y$, there is a neighbourhood U of y_0 such that $(F_y)_{y\in U} \subset (F_y)_{y\in Y}$ is uniformly proper.

Lemma 4.3.10 (Proposition 2.6 in [Tho17])**.** *Let X and Y be two metric spaces and $(F_y)_{y\in Y}$ a locally uniformly proper family of continous functions from X to $\mathbb{R}$ depending continuously on Y (for the compact open topology). Assume that each F_y achieves its minimum at a unique point $x_m(y) \in X$. Then the function*

$$y \mapsto x_m(y)$$

is continuous.

We verify the conditions for $\mathcal{E}^\theta_{\rho_1,\rho_2}$, with (ρ_1, ρ_2) living in the space of almost strictly dominating representations. For the remainder of this subsection, we are working over local sections for our bundles of local systems over representation spaces.

Lemma 4.3.11. *The association $(\rho_1, \rho_2) \mapsto \mathcal{E}^\theta_{\rho_1,\rho_2}$ is continuous for the compact-open topology.*

In the proof, we require control on the energy of the harmonic maps at the cusp, as we vary the source metric and the representation. We defer the proof of the lemmas below to the next subsection. In these lemmas, let ρ_0 be a reductive representation and μ_0 a hyperbolic metric.

Lemma 4.3.12. *Suppose ρ_0 has parabolic monodromy. Then for every representation ρ in the same representation space that is close enough to ρ_0, and metric μ close to μ_0, there is a function $\tilde{e}$ that is integrable in the flat metric and such that*

$$\sqrt{|\mu|}e(\mu, f_\mu^\rho) \leq \tilde{e}$$

everywhere.

Lemma 4.3.13. *Suppose ρ_0 has hyperbolic monodromy. Then for every representation ρ in the same representation variety that is close enough to ρ_0, and metric μ close to μ_0, working in the cusp coordinates for μ there is a $y_0 > 0$, $C, c > 0$ such that for all $y \geq y_0$,*

$$\frac{\Lambda(\theta)\ell^2}{2\tau_\mu^2} - Ce^{-cy} \leq \sqrt{\mu}e(\mu, f_\mu^\rho) \leq \frac{\Lambda(\theta)\ell^2}{2\tau_\mu^2} + Ce^{-cy},$$

where f_μ^ρ is the harmonic map with twist parameter θ.

Proof. We want to show that if $(i_n, j_n) \to (\rho_1, \rho_2)$ in $\mathrm{Rep}_{\mathbf{c}}(\Gamma, \mathrm{PSL}(2,\mathbb{R}) \times G)$ and $K \subset \mathcal{T}(\Gamma)$ is compact, then $\mathcal{E}_n := \mathcal{E}^\theta_{i_n, j_n} \to \mathcal{E} = \mathcal{E}^\theta_{\rho_1,\rho_2}$ uniformly on K as $n \to \infty$. The finite energy case is easy: from our previous results, if $\mu_n \to \mu$ and $\rho_n \to \rho$, then the energy densities of the harmonic maps converge pointwise to $e(\mu, f_\mu^\rho)$ (recall this is independent of the harmonic map if ρ is reducible). Lemma 4.3.12 then justifies an application of the domination convergence theorem, so that the total energies converge.

Going forward, we assume the monodromoy is hyperbolic. Fixing a metric μ, set $h_n = h_\mu^{i_n}$, $f_n = f_\mu^{j_n}$, $h = h_\mu^{\rho_1}$, $f = f_\mu^{\rho_2}$. It suffices to show $\mathcal{E}_n(\mu) \to \mathcal{E}(\mu)$, and that the rate only depends on the Teichmüller distance from μ to a base metric μ_0. By Lemma 4.3.13, for $r > y_0$,

$$\begin{aligned}
|\mathcal{E}(\mu) - \mathcal{E}_n(\mu)| &= \Big| \int_{\Sigma_r} (e(\mu,h) - e(\mu,f)) - (e(\mu,h_n) - e(\mu,f_n)) dA_\mu \Big| \\
&\quad + \Big| \int_{\Sigma\setminus\Sigma_r} (e(\mu,h) - e(\mu,f)) - (e(\mu,h_n) - e(\mu,f_n)) dA_\mu \Big| \\
&\leq \Big| \int_{\Sigma_r} (e(\mu,h) - e(\mu,f)) - (e(\mu,h_n) - e(\mu,f_n)) dA_\mu \Big| + \int_{\Sigma\setminus\Sigma_r} Ce^{-cy} dA_\mu
\end{aligned}$$

holds for every μ close enough to μ_0. Fixing $\epsilon > 0$, for every $r \geq y_0$ we can find $N_r > 0$, depending only on Teichmüller distance to μ_0, such that for all $n \geq N_r$, the first integral is $< \epsilon/2$. Hence, for such n,

$$|\mathcal{E}(\mu) - \mathcal{E}_n(\mu)| < \epsilon/2 + \frac{2\pi C}{c} e^{-cr}.$$

Taking $r = c^{-1}\log(4\pi C\epsilon^{-1})$, we get $|\mathcal{E}(\mu) - \mathcal{E}_n(\mu)| < \epsilon$. $\square$

We now show that the functionals $\mathcal{E}^{\theta}_{\rho_1,\rho_2}$ are locally uniformly proper. From here we assume $\theta = 0$, because the proof is identical for every θ. We essentially show that the bounds from the proof of Proposition 4.3.1 depend continuously on $([\rho_1],[\rho_2])$. For $([\rho_1],[\rho_2]) \in \mathrm{ASD}_{\mathbf{c}}(\Gamma, G)$, we choose an open neighbourhood $U \subset \mathrm{ASD}_{\mathbf{c}}(\Gamma, G)$ containing $([\rho_1],[\rho_2])$ with compact closure. We intersect it with a product open set $U_1 \times U_2$, where U_1 is a left d_{Th}-open ball around ρ_1. We then lift via some section to the space of local sections to view these points as representations. Picking a boundary geodesic or a horocycle for $C(\mathbb{H}/\rho_1(\Gamma))$, as we perturb the representations we get a continuously varying family of such curves in the new metric. We write $C^j_d(\xi)$ for the collar neighbourhood of such a curve in $\mathbb{H}/j(\Gamma)$, $j \in U_2$. By choosing U even small enough, we can assume we have a fixed presentation for our fundamental group, and the collar neighbourhood $C^j_{d_j}(\xi)$ has uniform upper and lower bounds $\delta_1 \leq d_j \leq \delta_2$.

Set $\delta = (\delta_1 + \delta_2)/2$. We can choose a neighbourhood C_δ containing every $C^j_\delta(\xi)$ for all $j \in U_2$. For a (j,ρ)-equivariant map g, we put

$$\mathrm{Lip}_\delta(g) = \max_{x \in C(\mathbb{H}/j(\Gamma)) \setminus C_\delta} \mathrm{Lip}_x(g).$$

Lemma 4.3.14. *Shrinking U if necessary, there exists an $\epsilon > 0$ such that for every $(j,\rho) \in U$, there is a (j,ρ)-equivariant map $g_{j,\rho}$ that satisfies $Lip(g_{j,\rho}) \leq 1$ and $Lip_\delta(g_{j,\rho}) < (1-\epsilon)^{1/2}$. Moreover, we can choose it so that if $\rho(\zeta)$ is hyperbolic, then $g_{j,\rho}$ translates the geodesic axis of $j(\zeta)$ along the axis of $\rho(\zeta)$ with constant speed* 1.

Proof. We first define $g_{j,\rho}$ on the complement of C_δ. One at a time, we vary j and then ρ. By our choice of U_1, there is an $\epsilon_0 > 0$ such that for every $j \in U_1$, there is a (j,ρ_1)-equivariant $(1+\epsilon_0)$-Lipschitz map. Composing with our original optimal (ρ_1,ρ_2)-equivariant map, we get a (j,ρ_2)-equivariant map with nice control. Choosing U_1 small enough, we can shrink ϵ_0 so as to ensure the right behaviour outside of C_δ.

Now we fix a base surface $\mathbb{H}/j(\Gamma)$ and vary ρ around ρ_2. We can use flat connections as in Lemma 4.3.6. For any continuous path of classes with initial point ρ_2, the procedure detailed there gives a path of equivariant maps starting at g. From compactness of the complement of the collar and cusp neighbourhoods, the local

Lipschitz constants vary upper semicontinuously. In particular, we can achieve an upper bound $\mathrm{Lip}_\delta(\cdot) < (1-\epsilon)^{1/2}$ when close enough to $([\rho_1],[\rho_2])$.

Now we extend in C_δ and above. Note that while the local Lipschitz constants of $g_{j,\rho}$ are uniformly controlled, this can be strictly below the global Lipschitz constant, and in the hyperbolic case this global Lipschitz constant is exactly 1. To see this, we do have $d_\nu(g_{j,\rho}(x), g_{j,\rho}(y)) < d_\sigma(x,y)$ for every $x \neq y$, so $\mathrm{Lip}(g_{j,\rho}) \leq 1$ certainly. But in the case of hyperbolic monodromy, for any two points x, y that are connected by a segment that mostly fellow-travels the geodesic axis of $j(\zeta)$,

$$d_\nu(g_{j,\rho}(x), g_{j,\rho}(y)) = d_\sigma(x,y) + O(1),$$

where the implied constant depends only on the position of x and y in $\mathbb{H}/j(\Gamma)$. Since x, y can be taken as far as we like, in the limit the ratio of distances becomes 1. Using the equivariant Kirszbraun-Valentine theorem [GK17, Proposition 3.9], adapted to the stabilizer of the cusp or funnel, we extend each such equivariant map to a globally defined equivariant map with global Lipschitz constant ≤ 1 in the parabolic and elliptic cases, and exactly 1 in the hyperbolic case. The constraint $\ell(j(\zeta)) = \ell(\rho(\zeta))$ forces $g_{j,\rho}$ to translate along the geodesic. □

Remark 4.3.15. [GK17, Proposition 3.9] is only proved for equivariant maps from hyperbolic n-space to itself. However, a version still holds for maps from $(\mathbb{H}, \sigma)$ to any CAT(-1) metric space. The proof involves taking barycenters of Lipschitz maps, which can be done just the same in any CAT(0) space, and a few applications of the Toponogov theorem that go through in a CAT(-1) setting.

Remark 4.3.16. In [GK17, Proposition 3.9], only Lipschitz constant at least 1 is addressed. Using compactness of C_δ and adapting the proof using [GK17, Proposition 3.7] rather than Propositions 3.1 and Remark 3.6 from that paper, we acquire the result in this other context (with a potential loss on the Lipschitz constant).

Returning to the main proof, we write h to be a Fuchsian harmonic map, omitting dependence on the metric and representation. For any $(j,\rho) \in U$, the fact that $g_{j,\rho}$ translates along the geodesic means we can apply Lemma 4.2.14:

$$\mathcal{E}_{j,\rho}(\mu) \geq \int_\Sigma e(\mu, h) - e(\mu, g_{j,\rho} \circ h) dA_\mu.$$

And using that $\mathrm{Lip}(g) = 1$, we get

$$\mathcal{E}_{j,\rho}(\mu) \geq \int_K e(\mu, h) - e(\mu, g_{j,\rho} \circ h) dA_\mu$$

for any compact $K \subset \Sigma$.

Fix a simple closed curve $\gamma \in \Gamma$. We are positioned to repeat the initial computation in the proof of Theorem 4A, and doing so gives that for $([j],[\rho]) \in U$ and $[\mu] \in \mathcal{T}(\Gamma)$ we have

$$\mathcal{E}_{j,\rho}(\mu) \geq \frac{w_\mu}{2}\epsilon \min_x \Big(\int_{A_x^t} \Big|\frac{\partial h(x,y)}{\partial y}\Big| dy \Big)^2, \tag{4.16}$$

where A_x^t is defined as in the proof of properness, and w_μ is the μ-length of the collar associated to γ. Repeating the proof of Lemma 4.3.2, almost word for word, we can see

$$\int_{A_x^t} \Big|\frac{\partial h(x,y)}{\partial y}\Big| dy \geq \min\{\delta_1/2, \kappa_\gamma\},$$

where κ_γ is the minimum of the lengths for $j(\gamma)$. If the μ-length of any γ goes to 0, then this integral explodes. Thus, for any $([j],[\rho]) \in U$, and curve γ, there is a length ϵ_γ such that if $\ell_\sigma(\gamma) < \epsilon_\gamma$, the right-hand-side of (4.16) is greater than K. This implies there is a compact subset of the moduli space such that if we take a fundamental domain V for this subset in Teichmüller space, then we have the $\mathcal{E}_{j,\rho} > K$ on the complement of the mapping class group orbit of V.

To finish the proof, we show there are only finitely many mapping classes $[\psi]$ such that $\mathcal{E}_{j,\rho} \leq K$ on the translate $\psi^* V$. Suppose there exists a metric μ representing a point in V and a sequence of distinct mapping classes ψ_n such that $\mathcal{E}_{j,\rho}(\psi_n^*\mu) \leq K$ for all $([j],[\rho]) \in U$. Then, by proper discontinuity of the mapping class group, $(\psi_n^{-1})^* j$ diverges in Teichmüller space for every j. This implies that for each j there exists a non-trivial simple closed curve γ_j whose length under $(\psi_n^{-1})^* j$ blows up as $n \to \infty$.

Since we intersected with a left open ball for d_{Th}, we can choose all γ_j to be equal to a single curve γ. If C is the radius for our left open ball, then for all n and $\gamma \in \Gamma$,

$$\ell((\psi_n^{-1})^* j(\gamma)) \geq e^{-C} \ell(\psi_n^{-1})^* \rho_1(\gamma)). \tag{4.17}$$

Thus, if $\gamma \in \Gamma$ is such that $\ell((\psi_n^{-1})^* \rho_1(\gamma))$ blows up, then by (4.17), the same holds for every j sufficiently close by. Moreover, the rate at which $\ell((\psi_n^{-1})^* j(\gamma)) \to \infty$ is independent of j, close to that of $(\psi_n^{-1})^* \rho_1(\gamma)$. Now we have an integral estimate as in Theorem 4A:

$$K \geq \mathcal{E}_{j,\rho}(\psi_n^*\mu) = \mathcal{E}_{(\psi_n^{-1})^* j, (\psi_n^{-1})^* \rho}(\mu) \geq \frac{w}{2}\epsilon \min_x \Big(\int_{A_x^t} \Big|\frac{\partial h(x,y)}{\partial y}\Big| dy \Big)^2,$$

where w is minimum of the lengths of the collars around γ for $[\mu] \in V$. The proof of Lemma 4.3.4 can then be made uniform: by examination of the proof, the integral

$$\int_{A_x^t} \left|\frac{\partial h(x,y)}{\partial y}\right| dy$$

trails off to infinity with a rate depending on that of the translation length of the bad sequence. This is a contradiction, and thus the energy functional does have the $> K$ condition on the complement of a finite orbit. Therefore, we've satisfied Lemma 4.3.10, and modulo Lemmas 4.3.12 and 4.3.13, finished the proof of Theorem 4B.

Variations at the cusp

Here we prove Lemmas 4.3.12 and 4.3.13. Let μ be any metric close to μ_0. Uniformizing a neighbourhood of the cusp to a punctured disk, we consider the Hopf differential as a meromorphic function for the complex structure of μ with a pole of order at most 2:

$$\phi(z) = -\frac{\Lambda(\theta)\ell^2}{16\pi^2} z^{-2} + a_\mu z_\mu^{-1} + \varphi_\mu(z),$$

where z_μ is a holomorphic coordinate for μ, and ℓ is the translation length of the peripheral curve in question. If the representation does not have hyperbolic monodromy at the cusp, then it is understood that $\ell = 0$ in the expression above. We can choose a neighbourhood of the puncture containing cusp neighbourhoods for all μ that uniformize for μ to an open set containing a punctured disk of μ-radius uniformly bounded below.

It follows from the results of Section 4.2 and the proof of continuity in Theorem 4B that for any (μ_n, ρ_n) converging to (μ, ρ), the harmonic maps can be chosen, even in the reducible case, to converge in the C^∞ sense on compacta. This implies the Hopf differentials, viewed simply as smooth rather than holomorphic functions, converge to that of $f_{\mu_0}^{\rho_0}$ locally uniformly on compacta in this punctured disk. If $z = z_{\mu_0}$, then after choosing our normalizations correctly, $z_\mu \to z$ as $\mu \to \mu_0$. It follows that the Laurent coefficients converge, and hence a_μ is bounded and φ_μ is C^0 bounded.

To prove Lemma 4.3.12, we first assume ρ is Fuchsian. Then, in the coordinates on $\mathbb{D}^*$, the Beltrami form ψ satisfies

$$|\psi| = \frac{|\Phi|}{\mu H(\mu, f_\mu)} \leq \frac{|\Phi|}{\mu} \leq \frac{Cz^{-1}}{|z|^{-2}(\log|z|)^2} = C|z|(\log|z|)^{-1} \to 0$$

as $z \to 0$. Via this decay on the Beltrami form, we know that once we go high enough into the cusp, f there is a uniform bound on the quasiconformal dilatation (independent of μ and ρ). Thus from uniform convergence on compacta, we have a uniform K-quasiconformal bound everywhere. By the Schwarz lemma for quasiconformal harmonic maps [GH77], we extract the bound $H(\mu, f_\mu) \leq 2K^2$. Since $L \leq H$ in the Fuchsian case, $e(\mu, f_\mu) \leq 4K^2$. If ρ is not Fuchsian, then by Proposition 3.13 from [Sag19] we can bound the energy density above by that of the harmonic map for the Fuchsian representation with the same Hopf differential. This constant is integrable, and this proves Lemma 4.3.12.

Lemma 4.3.13 is a bit more work. Passing to the cusp coordinates, the uniform bound on the Laurent coefficients implies there is a uniform $C, c > 0$ such that ϕ satisfies

$$\frac{\Lambda(\theta)\ell^2}{4\tau_\mu^2} + Ce^{-cy} \leq |\phi| \leq \frac{\Lambda(\theta)\ell^2}{4\tau_\mu^2} + Ce^{-cy} \tag{4.18}$$

in $\Sigma \backslash \Sigma_s$. Setting $\psi = \psi_\mu$ to be $|(f_\mu)_{\bar{z}}|/|(f_\mu)_z|$, the formula

$$e(\mu^f, f_\mu) = |\Phi|(|\psi| + |\psi|^{-1}) \tag{4.19}$$

suggests we should turn to $|\psi|$. Using (4.18), we find uniform upper and lower bounds on $|\psi|$, independent of (μ, ρ). If there are no such bounds, then there is a sequence μ_n tending to μ and points z_n with $|\psi_n|(z_n) \to 0$ or $|\psi_n|(z_n) \to \infty$, where $|\psi_n| = L(f_n)^{1/2}/H(f_n)^{1/2}$, for $f_n = f_{\mu_n}$. We can assume each z_n lies in a cylinder of the form $(\Sigma_{r_n} \backslash \Sigma_{r_n - 1}, \mu)$. Taking $i_n : C \to (\Sigma_{r_n} \backslash \Sigma_{r_n-1}, \mu)$ to be a cylinder embedding, conformal for μ, uniform energy density bounds imply convergence of $F_n = f_n \circ i_n$ along a subsequence to a limiting harmonic map $F_\infty : C \to (X, \nu)$. From the inequalities (4.18), the Hopf differential is exactly $\Lambda(\theta)\ell^2/4\tau_{\mu_0}^2 dz_{\mu_0}^2$. F_∞ projects onto the geodesic and hence has rank 1. Thus, there is an $\eta \in \mathbb{R}$ such that $(F_\infty)_*(\partial/\partial x) = \eta(F_\infty)_*(\partial/\partial_y)$. Hence, writing out the Hopf differential in coordinates gives

$$\Phi(F_\infty) = \frac{1}{4}(|(F_\infty)_*(\partial/\partial x)|_\nu^2 - |(F_\infty)_*(\partial/\partial_y)|_\nu^2 - 2i\langle f_x, f_y\rangle_\nu) = \frac{1}{4}(\eta^2 - 1 - 2i\eta)dz^2.$$

Therefore, $\eta = \theta$. We thus find from the linear ODE theory that the limit is a constant speed parametrization of the geodesic, composed with a fractional Dehn twist. This implies the limiting quantity $|\psi_\infty|$ is exactly 1, which contradicts our assumption $|\psi|(z_n) \to 0$ or ∞.

From uniform bounds we upgrade to more precise control. Let us temporarily assume ρ is Fuchsian. Working in the region where we have these bounds, because the pullback metric for our harmonic map is hyperbolic, it can be deduced from the Bochner formulae [SY97, Chapter 1] that

$$\Delta_{\mu^f} \log |\psi|^{-1} = 2|\Phi_\mu| \sinh \log |\psi|^{-1}.$$

Hence,

$$\frac{\Lambda(\theta)\ell^2}{4\tau_\mu^2}\log|\psi|^{-1} \leq \Delta_{\mu^f}\log|\psi|^{-1} \leq \frac{\Lambda(\theta)\ell^2}{2\tau_\mu^2}\log|\psi|^{-1}$$

when $|\psi| < 1$, if we are high enough to get uniform control on $|\Phi_\mu|$. If $|\psi| > 1$, we have the opposite inequality

$$\frac{\Lambda(\theta)\ell^2}{2\tau_\mu^2}\log|\psi|^{-1} \leq \Delta_{\mu^f}\log|\psi|^{-1} \leq \frac{\Lambda(\theta)\ell^2}{4\tau_\mu^2}\log|\psi|^{-1}.$$

Our uniform bounds on $|\psi|$ give control on $\log|\psi|^{-1}$, which yields more bounds of the form

$$-\frac{c\Lambda(\theta)\ell^2}{2\tau_\mu^2} \leq \Delta_{\mu^f}\log|\psi|^{-1} \leq \frac{C\Lambda(\theta)\ell^2}{2\tau_\mu^2}.$$

Using the maximum principle, we can then deduce

$$-Ce^{-cy} \leq \log|\psi|^{-1} \leq Ce^{-cy}.$$

Taylor expanding $x \mapsto \log(1-x)$, we then obtain

$$1 - Ce^{-cy} \leq |\psi| \leq 1 + Ce^{-cy}.$$

If ρ is not Fuchsian, we apply an argument similar to that of Lemma 4.2.3 to get this same asymptotic. Inserting the bounds into the formula (4.19) gives

$$\frac{\Lambda(\theta)\ell^2}{2\tau_\mu^2} - Ce^{-ct} \leq e(\mu^f, f_\mu) \leq \frac{\Lambda(\theta)\ell^2}{2\tau_\mu^2} + Ce^{-ct},$$

as desired. This completes the proof of Lemma 4.3.13, and moreover the proof of Theorem 4B.

4.4 Anti-de Sitter 3-manifolds

AdS 3-manifolds with S^1-fibrations

In this subsection we prove Proposition 4.1.8, which is actually a quick consequence of the proposition immediately below. We work in the $\mathrm{PSL}(2,\mathbb{R})$ model throughout.

Proposition 4.4.1. *Let $V \subset \mathbb{H}$ be a domain. The data of a domain $\Omega \subset AdS^3$ and a fibration $\Omega \to V$ such that every fiber is a timelike geodesics is equivalent to that of a domain $V \subset \mathbb{H}$ and a locally strictly contracting map $g : V \to \mathbb{H}$.*

The proof of the first direction of the equivalence is a straightforward adaptation of the procedure from [GK17, Proposition 7.2]. There, $V = \mathbb{H}$, $\Omega = \mathrm{AdS}^3$, and h is (globally) strictly contracting. We include the proof for the readers convenience.

Proof. The key fact we use is that timelike geodesics L_{p_1,q_1} and L_{p_2,q_2} intersect if and only if

$$d_\sigma(p_1, p_2) = d_\sigma(q_1, q_2).$$

With this in mind, given a locally strictly contracting mapping $g : V \to \mathbb{H} \times \mathbb{H}$ with the properties above, timelike geodesics of the form $L_{p,g(p)}$ and $L_{q,g(q)}$ never intersect. Thus, the geodesics $L_{p,g(p)}$ sweep out a connected set $\Omega \subset \mathrm{AdS}^3$ as p ranges over V.

We argue that Ω is open. We record that $X \in L_{p,g(p)}$ if and only if

$$X^{-1} \circ g(p) = p.$$

For small $\epsilon > 0$, let $B_\epsilon(p) \subset V$ denote the ϵ-ball around p in $\mathbb{H}$. Let $B \subset \mathrm{AdS}^3$ be the open ball consisting of isometries Y such that

$$d_\sigma(p, Y^{-1}g(p)) < (1 - \mathrm{Lip}(g|_{B_\epsilon(p)}))\epsilon.$$

Then for any $q \in B_\epsilon(p)$ and $Y \in B$,

$$d_\sigma(Y^{-1} \circ g(q), p) \leq d(Y^{-1}g(q), Y^{-1}g(p)) + d(Y^{-1}g(p), p) < \epsilon.$$

Thus, $Y^{-1}g$ takes the closure of $B_\epsilon(p)$ to itself, and by the Banach fixed point theorem there is a unique $q \in B_\epsilon(p)$ such that $Y \circ g(q) = q$. So $B \subset \Omega$. This argument also shows that the fibration from $\Omega \to V$ described by $L_{f(p),h(p)} \mapsto p$ is continuous.

For the other direction, any circle fibration $\Omega \to V$ with timelike geodesic fibers determines a map $F : V \to \mathbb{H} \times \mathbb{H}$ by $F(p) = (h(p), f(p))$, where $L_{f(p),h(p)}$ is the geodesic lying over p in Ω. F preserves connectedness—using the product structure, [JM18, Theorem 2.2] guarantees it is continuous when f is non-constant. If f is a constant q, then because Ω is open, for any p and path from p to q, we can find a continuous path of isometries $r \mapsto X_r$ such that $h(r) = X_r^{-1}q$. Thus we have continuity here as well. As the timelike geodesics never intersect, $d(f(p), f(q)) \neq d(h(p), h(q))$ for $p \neq q$. As the diagonal in $\mathbb{H} \times \mathbb{H}$ has codimension 2, a connectedness argument shows $d(f(p), f(q)) < d(h(p), h(q))$ or $d(f(p), f(q)) > d(h(p), h(q))$ always for $p \neq q$. By switching coordinates, we may assume we have the former. This condition ensures that h is injective. Therefore, $g = f \circ h^{-1}$ is a well-defined locally strictly contracting map. □

Proposition 4.1.8 is just the equivariant version of this: for a pair (ρ_1, ρ_2) with ρ_1 acting properly discontinuously on V, we have

$$\rho_2(\gamma) L_{p,g(p)} \rho_1(\gamma)^{-1} = L_{\rho_1(\gamma)p, \rho_2(\gamma)g(p)},$$

so $\rho_1 \times \rho_2$ acts properly discontinuously on Ω and equivariance of the fibration is clear.

It is seen in the proof that $\Omega \subset \mathrm{AdS}^3$ consists of all isometries X such that $X^{-1} \circ g$ has a fixed point.

Remark 4.4.2. The results here generalize, almost word for word, for quotients of proper domains in the rank 1 Lie groups $G = \mathrm{O}(n,1)$, $\mathrm{SO}(n,1)$, $\mathrm{SO}_0(n,1)$, and $\mathrm{PO}(n,1)$. One can consider the action by left and right multiplication and equivariant K-fibrations $G \supset \Omega \to V \subset \mathbb{H}^n$, $n \geq 2$, where $K \subset G$ is the maximal compact subgroup. Here the fibers are copies of K, each of the form $\{X \in G : X \cdot p = q\}$ for some $p, q \in \mathbb{H}^n$.

Remark 4.4.3. Proposition 4.1.8 applies to non-reductive representations. They have been largely excluded from our discussion because harmonic maps and maximal surfaces do not exist for these representations.

Theorem 4C

Here we give the proof of Theorem 4C. We make use of results from the paper [GK17]. Fix reductive representations $\rho_1, \rho_2 : \Gamma \to \mathrm{PSL}(2, \mathbb{R})$ with ρ_1 Fuchsian.

Definition 4.4.4. Let $V \subset \mathbb{H}$ be a ρ_1-invariant domain, and $f : V \to \mathbb{H}$ a (ρ_1, ρ_2)-equivariant map realizing the minimal Lipschitz constant L among equivariant maps. The stretch locus is the set of points $x \in \mathbb{H}$ such that the restriction of f to any neighbourhood of x has Lipschitz constant exactly L and no smaller.

The result below is culled from [GK17, Theorem 1.3 and 5.1]. See the reference for more general statements and details.

Theorem 4.4.5 (Guéritaud-Kassel)**.** *Assume there exists a (ρ_1, ρ_2)-equivariant map with minimal Lipschitz constant $L = 1$, and let E be the intersection of all the stretch loci among such maps. Then there exists an "optimal" (ρ_1, ρ_2)-equivariant 1-Lipschitz map whose stretch locus is exactly E. E projects under the action of $\rho_1(\Gamma)$ to the convex core for ρ_1, and is either empty or the union of a lamination and 2-dimensional convex sets with extremal points only in the limit set $\Lambda_{\rho_1(\Gamma)} \subset \partial_\infty \mathbb{H}$.*

Proof of Theorem 4C. The equivalence between (1) and (3) is contained in Theorem 4A. Assuming (1) we prove (2). Take any optimal map g, and the map $\tilde{C}(\mathbb{H}/\rho_1(\Gamma)) \to \mathbb{H} \times \mathbb{H}$ given by $p \mapsto (p, g(p))$. In the case that there exists a peripheral on which ρ_1 is hyperbolic, suppose for the sake of contradiction that there is a choice of g so that the domain extends to give a fibration over a larger subsurface. From the other direction of Proposition 4.1.8, we obtain a (ρ_1, ρ_2)-equivariant and a locally contracting map defined on the preimage of this subsurface in the universal cover. This implies there is a peripheral γ with $\ell(\rho_1(\gamma)) > \ell(\rho_2(\gamma))$, which contradicts our original Definition 4.1.1.

Now we prove that (2) implies (1). Given such a domain and fibration, from Proposition 4.1.8 we obtain a strictly 1-Lipschitz (ρ_1, ρ_2)-equivariant map defined on $\tilde{C}(\mathbb{H}/\rho_1(\Gamma))$. If ρ_1 has no hyperbolic peripherals, then we get (1) for free and we're done. So assume there is a peripheral ζ with $\rho_1(\zeta)$ hyperbolic. Any 1-Lipschitz map g defined inside $\tilde{C}(\mathbb{H}/\rho_1(\Gamma))$ extends to a 1-Lipschitz map of the frontier inside $\mathbb{H}$, and hence

$$\ell(\rho_2(\zeta)) \leq \ell(\rho_1(\zeta)).$$

We extend g to all of $\mathbb{H}$ by precomposing with the 1-Lipschitz (ρ_1, ρ_1)-equivariant nearest point projection onto $\tilde{C}(\mathbb{H}/\rho_1(\Gamma))$, so we know that the set of globally defined Lipschitz maps is non-empty. From Lemma 4.10 in [GK17] (an application of Arzelà-Ascoli), there exists an optimal (ρ_1, ρ_2)-equivariant Lipschitz map g'. As for the optimal Lipschitz constant, g shows $L \leq 1$, and if $L < 1$ then $\rho_1 \times \rho_2$ acts properly discontinuously on the whole AdS^3, and hence $L = 1$. Applying Theorem 4.4.5, we have a stretch locus E.

E is contained in the intersection of the stretch loci of g and g'. Since g does not maximally stretch in the interior of $\tilde{C}(\mathbb{H}/\rho_1(\Gamma))$, E is contained in the boundary of $\tilde{C}(\mathbb{H}/\rho_1(\Gamma))$. If E is missing the lifts of one boundary component of $\mathbb{C}(\mathbb{H}/\rho_1(\Gamma))$, then g' is strictly contracting inside the half-spaces in $\mathbb{H}$ that project to the infinite funnel bounding this component in $\mathbb{H}/\rho_1(\Gamma)$. From Proposition 4.1.8, we can thus find a $\rho_1 \times \rho_2$-equivariant domain that yields a fibration onto the union of the convex core with this funnel, which contradicts our standing assumption. We conclude that the stretch locus is exactly these components, and hence g' is an almost strictly dominating map.

For the final statement, we use the homeomorphism $\Psi = \Psi^0$ from Theorem 4B to parametrize the space of representations. We take the domains in AdS^3 associated to the spacelike maximal immersions with 0 twist parameter (any one will do). The energy domination implies that they yield proper quotients by Proposition 4.1.8. Since the harmonic maps for irreducible representations vary continuously with the representation, so do the domains in AdS^3. Hence, when we restrict to these classes, Ψ parametrizes a deformation space of AdS 3-manifolds. □

To produce more representations that give such incomplete quotients, take an almost strictly dominating pair (ρ_1, ρ_2) (Theorem 4B shows there are many) and an optimal map g. To relax the condition that all boundary lengths agree, first choose a collection of peripherals, but not all of them. For each of the selected peripherals, there is a geodesic or a horocycle in $\mathbb{H}/\rho_1(\Gamma)$. We then specify a transversely intersecting geodesic arc that does not intersect any other peripheral geodesic or horocycle, and apply strip deformations to $\mathbb{H}/\rho_1(\Gamma)$ along these arcs (see [DGK16a], [Sag19, Section 6]). This gives a new hyperbolic surface whose holonomy is a Fuchsian representation j, and for some $\lambda < 1$, a strictly λ-Lipschitz (j, ρ_1)-equivariant map g'. We can extend g outside of the convex hull of the limit set by using the 1-Lipschitz (ρ_1, ρ_1)-equivariant closest-point projection. We then take the composition $g \circ g'$ and the corresponding circle bundle.

With the main theorems complete, we briefly digress to discuss the topology of the quotients. The quotients naturally acquire an orientation. Since the surface is not compact, the bundle is topologically trivial: $\mathrm{BU}(1) = \mathbb{CP}^\infty$ and $[\Sigma, \mathbb{CP}^\infty] = H^2(\Sigma, \mathbb{Z}) = 0$.

However, the global trivialization is by no means compatible with the AdS structure. To be precise, the 3-manifold is not "standard" in the sense of [KR85]: its casual double cover does not possess a timelike Killing field. If it did, the holonomy would normalize the isometric flow generated by the Killing field, and it follows from [KR85, pages 237-238] that this is impossible for reductive ρ_2.

4.5 Parabolic Higgs bundles

In [AL18], Alessandrini and Li use Higgs bundles to construct AdS structures on closed 3-manifolds. They build circle bundles explicitly, rather than first passing through [GK17, Theorem 1.8]. Following their work closely, we offer an alternative construction of the AdS structures from Theorem 4C for representations $\rho_1 \times \rho_2$ that lift to $\mathrm{SL}(2, \mathbb{R}) \times \mathrm{SL}(2, \mathbb{R})$. In this way, we are able to compute some geometrically meaningful quantities. And in the process, we explain how this construction is related to the previous one.

Since we are working over surfaces with punctures, we use parabolic Higgs bundles. We refer the reader to [Sim90] and [Mon16] for more background information on parabolic Higgs bundles.

More on timelike geodesics

Recall the billinear form $q_{n,2}$ from the beginning of Chapter III. In this section, we set $Q = q_{2,2}$, so that $\mathbb{H}^{2,1} = \{x \in \mathbb{R}^{2,2} : Q(x,x) = -1\}$. We use the Klein model $\mathrm{AdS}^3 = \mathbb{H}^{2,1}/\{\pm I\}$.

When restricted to a timelike geodesic, the covering $\mathbb{H}^{2,1} \to \mathrm{AdS}^3$ restricts to a covering of the circle (and the universal covering of this circle can be seen through $\widetilde{\mathrm{AdS}^3} \to \mathrm{AdS}^3$). Thus, there is a bijection between the space of timelike geodesics in $\mathbb{H}^{2,1}$ and AdS^3. In $\mathbb{H}^{2,1}$, timelike geodesics are intersections of timelike planes with $\mathbb{H}^{2,1}$. Projecting to AdS^3, timelike geodesics are projective lines contained wholly in AdS^3.

Working in $\mathbb{H}^{2,1}$, the space of timelike geodesics is the Grassmanian $\mathrm{Gr}^{\mathrm{t}}(2,4)$ of timelike planes in $\mathbb{R}^{2,2}$. To see the structure of this Grassmanian, $SO_0(2,2)$ acts transitively on the space of timelike planes, and the stabilizer of any timelike plane is conjugate to $SO(2) \times SO(2)$. Thus, $\mathrm{Gr}^{\mathrm{t}}(2,4)$ is the symmetric space

$$SO_0(2,2)/(SO(2) \times SO(2)).$$

We recall that $\mathbb{H} \times \mathbb{H}$ is the space of timelike geodesics in the $\mathrm{PSL}(2,\mathbb{R})$ model. The symmetric structure there is seen through the obvious action of $\mathrm{PSL}(2,\mathbb{R}) \times \mathrm{PSL}(2,\mathbb{R})$ on $\mathbb{H} \times \mathbb{H}$, with stabilizer of a point (x,y) being the product of the stabilizers.

This implicit identification from $\mathrm{Gr}^{\mathrm{t}}(2,4) \to \mathbb{H} \times \mathbb{H}$ actually preserves a pseudo-Riemannian structure on $\mathrm{Gr}^{\mathrm{t}}(2,4)$. Given $\mathrm{Span}(v_1, v_2) \in \mathrm{Gr}^{\mathrm{t}}(2,4)$, where v_1, v_2 are Q-orthonormal, we find spacelike vectors w_1, w_2 completing v_1, v_2 to a positively oriented Q-orthonormal basis. Lifting the classical Plücker embedding, there is a map from

$$\mathrm{Gr}^{\mathrm{t}}(2,4)(2,4) \to \wedge^2\mathbb{R}^4$$

defined by taking

$$\mathrm{Span}(v_1, v_2) \mapsto v_1 \wedge v_2.$$

The wedge product $\wedge^2\mathbb{R}^4 \times \wedge^2\mathbb{R}^4 \to \wedge^4\mathbb{R}^4$ gives rise to a signature $(3,3)$ non-degenerate bilinear form, and this restricts to a signature $(2,2)$ pseudo-Riemannian metric on the image of $\mathrm{Gr}^{\mathrm{t}}(2,4)$. Henceforth we endow $\mathrm{Gr}^{\mathrm{t}}(2,4)$ with the pullback metric m. As for the geometry of this metric, we have the result below.

Proposition 4.5.1. *$(Gr^{\mathrm{t}}(2,4), m)$ is isometric to $(\mathbb{H}\times\mathbb{H}, \sigma\oplus(-\sigma))$ via an isometry that intertwines the actions of the Lie groups $SO_0(2,2)$ and $PSL(2,\mathbb{R})\times PSL(2,\mathbb{R})$.*

This seems well-known, but we could not locate a proof or even a formal statement in the literature.

Proof. Denote by B the bilinear form on $\wedge^2\mathbb{R}^4$. For the reader's convenience we remark that if e_i is the standard basis and $v_j = \sum_i a_{ij}e_i$, then

$$B(v_1\wedge v_2, v_3\wedge v_4) = \det(a_{ij}).$$

For any Q-orthonormal basis v_1, v_2, v_3, v_4 satisfying $Q(v_1,v_1) = Q(v_2,v_2) = -1$, $Q(v_3,v_3) = Q(v_4,v_4) = 1$, the 2-vectors

$$V^1_\pm = \frac{1}{\sqrt{2}}(v_1\wedge v_2 \pm v_3\wedge v_4)\ ,\ V^2_\pm = \frac{1}{\sqrt{2}}(v_4\wedge v_2 \pm v_3\wedge v_1)\ ,\ V^3_\pm = \frac{1}{\sqrt{2}}(v_1\wedge v_4 \pm v_2\wedge v_3)$$

satisfy $B(V^j_\pm, V^j_\pm) = \pm 1$. Consider the subspaces $\wedge_\pm\mathbb{R}^4 = \mathrm{Span}(V^2_\pm, V^3_\pm, V^1_\mp) \subset \wedge^2\mathbb{R}^4$, along which $\wedge^2\mathbb{R}^4$ splits as

$$\wedge^2\mathbb{R}^4 = \wedge_+\mathbb{R}^4 \oplus \wedge_-\mathbb{R}^4.$$

The first subspace has signature $(2,1)$, while the second has signature $(1,2)$. With respect to the splitting, $SO_0(3,3)$ decomposes into an $SO_0(2,1)$ and a $SO_0(1,2)$ factor. The timelike Grassmanian embeds via

$$v_1\wedge v_2 \mapsto \frac{1}{\sqrt{2}}(v_1\wedge v_2 + v_3\wedge v_4, v_1\wedge v_2 - v_3\wedge v_4) \in \wedge_+\mathbb{R}^4 \oplus \wedge_-\mathbb{R}^4,$$

and this respects the $SO_0(2,1)$ and $SO_0(1,2)$ actions. The induced metric on the first factor is the pullback metric from the natural inclusion into Minkowski space, invariant under $SO_0(2,1)$. This is exactly the hyperboloid model of $\mathbb{H}$. The same reasoning applies for the second factor. □

Immersions in the Grassmanian

In this subsection, we explain a construction of AdS^3 from the perspective of the projective model, which is needed to approach AdS^3 quotients via Higgs bundles. Given representations ρ_1, ρ_2 to $\mathrm{SL}(2,\mathbb{R})$, we write $\rho_1\otimes\rho_2$ for the projectivization of their tensor product, which maps to $SO_0(2,2)$.

Proposition 4.5.2. *There is a bijection between $\rho_1 \otimes \rho_2$-equivariant spacelike surfaces in $Gr^t(2,4)$ and isomorphism classes of AdS circle bundles over Σ with monodromy $\rho_1 \times \rho_2$ and such that each circle fiber develops bijectively to a timelike geodesic in AdS^3.*

The work below uses the notion of a geometric structure.

Definition 4.5.3. An AdS^3 structure on a 3-manifold U is an atlas of charts $\{(\Omega_i, \varphi_i)\}_{i\in I}$ such that every φ_i is a diffeomorphism onto an open subset of ${\mathrm{AdS}^3}^3$, and every transition map $\varphi_i \circ \varphi_j^{-1}$ is the restriction of an element of $PO(2,2)$.

This is an example of a more general (G, X)-structure on a manifold U, where G is a Lie group acting transitively and faithfully on a manifold X with $\dim X = \dim U$. As in the definition above, this is an atlas of charts on a manifold U taking image in X and such that transition maps are restrictions of elements of G. We refer the reader to the survey [Ale19] for the general theory. The data of a (G, X)-structure on manifold U is equivalent to that of a representation $\rho : \pi_1(U) \to G$ and a ρ-equivariant local diffeomorphism $\tilde{D} : \tilde{U} \to X$ called the developing map. At least in one direction, from a pair $(\rho, \tilde{D})$, one can construct an atlas for (G, X) structure as follows: for a sufficiently small neighbourhood $\Omega \subset U$, choose a local section of the universal covering $s : \Omega \to \tilde{\Omega} \subset \tilde{U}$, and define $\varphi : U \to X$ by $\varphi = \tilde{D} \circ s$.

In [Bar10, Section 3.5], Baraglia finds a correspondence between real projective structures on circle bundles over surfaces and equivariant maps into Grassmanians. AdS^3 lies inside $\mathbb{RP}^3$, so Proposition 4.5.2 is found by “restricting” Baraglia’s constructions. To some extent, Proposition 4.5.2 is also observed in [AL18]. We give a proof because we need the map D below, although we omit some details.

Proof of Proposition 4.5.2. Start with a bundle $U \to S$. There is a commutative diagram

$$\begin{array}{ccccc} \tilde{U} & \longrightarrow & \overline{U} & \longrightarrow & U \\ & & \downarrow & & \downarrow \\ & & \tilde{\Sigma} & \longrightarrow & \Sigma, \end{array}$$

where $\tilde{U} \to U$ is the universal covering and $\overline{U}$ is the pullback circle bundle over $\tilde{\Sigma}$. The circle fibers of $\overline{U}$ also develop bijectively to timelike geodesics, so the developing map induces a (ρ_1, ρ_2)-equivariant map $G : \tilde{\Sigma} \to \mathrm{Gr}^t(2,4)$.

The passage above is reversible: suppose we have a (ρ_1, ρ_2)-equivariant map $G : \tilde{\Sigma} \to \mathrm{Gr}^{\mathrm{t}}(2,4)$. We define a circle bundle E over $\mathrm{Gr}^{\mathrm{t}}(2,4)$ by taking the fiber over a point to be the corresponding timelike geodesic in AdS^3. Alternatively, the Plücker embedding gives a map to a projective space

$$\mathrm{Gr}^{\mathrm{t}}(2,4) \to \mathrm{Gr}(2,r) \to \mathbb{P}(\wedge^2\mathbb{R}^4),$$

and we take the pullback to $\mathrm{Gr}^{\mathrm{t}}(2,4)$ of the tautological $\mathbb{RP}^1$-bundle. Above, $\mathrm{Gr}(2,4)$ is just the ordinary Grassmanian of 2-planes in $\mathbb{R}^4$, or lines in $\mathbb{RP}^3$. By definition, there is a natural map $i : E \to \mathrm{AdS}^3$. Set $\overline{U}$ to be the pullback bundle $\overline{U} = G^*E \to \tilde{\Sigma}$.

We obtain a map $D : \overline{U} \to \mathrm{AdS}^3$ by precomposing i with the bundle map $G^*E \to E$. Using this, we find an action of Γ on $\overline{U}$ that lifts the action on $\tilde{\Sigma}$. Let $z \in \tilde{\Sigma}$ and let $a \in \overline{U}_z$. Then $D(a) \in G(z) \subset \mathrm{AdS}^3$, and for any $\gamma \in \Gamma$,

$$(\rho_1 \otimes \rho_2)(\gamma)D(a) \in (\rho_1 \otimes \rho_2)(\gamma)F(z) = F(\gamma z).$$

Thus, there is a unique $b \in \overline{U}_{\gamma z}$ such that $D(b) = (\rho_1 \otimes \rho_2)(\gamma)D(a)$. Therefore, we define the action by $\gamma \cdot a = b$. It is easy to verify this is smooth and properly discontinuous, and moreover that $\overline{U}$ descends to a bundle $U \to \Sigma$ with respect to this action.

It follows from [Bar10, Lemma 3.5.01] that the map D is a local diffeomorphism, hence pulls back to $\tilde{U}$ to make a developing map for an AdS^3 structure, if and only if G is spacelike or timelike. After possibly passing from $\mathrm{Gr}^{\mathrm{t}}(2,4)$ to $(\mathbb{H}\times\mathbb{H}, \sigma\oplus(-\sigma))$ and applying $(x,y) \mapsto (y,x)$, we can assume the surface is spacelike. □

Relating back to Theorem 4C, if we take ρ_1 almost strictly dominating ρ_2 from Theorem 4A and an equivariant maximal spacelike immersion F, then passing through the isomorphism we get a map G to the Grassmanian. The map D here is injective, and hence a diffeomorphism onto its image. This recovers a domain Ω as $\Omega = D(\overline{U})$, and the properly discontinuous action of $\rho_1 \otimes \rho_2$ on $\overline{U}$ is mapped to an action on Ω.

Parabolic Higgs bundles

In the next few subsections, we construct the circle bundles from Theorem 4C directly. This requires us to set up the basic definitions and results for parabolic Higgs bundles.

Throughout this section, we work on a closed Riemann surface S of genus $g \geq 2$ with canonical bundle K_S. Setting $D = p_1 + \ldots p_n$ to be an effective divisor, we view $\Sigma = S \backslash D$.

Definition 4.5.4. A parabolic vector bundle E_* of rank r over S is the data of

- a holomorphic vector bundle $E \to S$ of rank r,
- at each p_i, a choice of real numbers $0 \leq \alpha_i^1 \leq \ldots \alpha_i^{n_i} < 1$ called parabolic weights,
- at each p_i, a strictly decreasing filtration $E_{p_i} = E_i^1 \supset E_i^2 \supset \cdots \supset E_i^{n_i+1} = 0$.

For bundles $E_* = (E, p_i, \alpha_i^j)$, $F_* = (F, p_i, \beta_i^k)$ over S, we can define direct sums and tensor products (see [Mon16, page 16]), as well as homomorphisms of bundles:

$$\mathrm{Hom}(E_*, F_*) = \{\varphi \in \mathrm{Hom}(E, F) : \alpha_j^i > \beta_k^i \Rightarrow \varphi(E_i^j) \subset F_i^k)\}.$$

Any holomorphic subbundle $F \subset E$ of a holomorphic subbundle acquires a parabolic structure: for each p_i, $F_{p_i} \cap E_i^j$ defines a filtration of the fiber F_{p_i}. The parabolic weight of a subspace $V \subset F_{p_i}$ is the maximum of the numbers $\{\alpha_i^j : V \subset E_i^j \cap F_{p_i}\}$. We say that a parabolic bundle is trivial if the underlying bundle is trivial and the parabolic structure is trivial ($n_i = 1$ and all weights vanish).

Over parabolic bundles, we consider logarithmic connections: first order holomorphic differential operators

$$\nabla : E \to E \otimes \Omega_S^1(\log D)$$

satisfying the Liebniz rule with respect to locally defined holomorphic sections of E and locally defined holomorphic functions. Here $\Omega_S^1(\log D) = K_S \otimes \mathcal{O}_S(D)$, $\mathcal{O}_S(D)$ is the holomorphic line bundle associated to D. Concretely, this produces a "meromorphic" connection on E, holomorphic over $E|_\Sigma$, and such that in a holomorphic chart around p_i with coordinate z, ∇ is specified by

$$\nabla = d + M(z)\frac{dz}{z}.$$

The induced endomorphism of the filtered vector space is called the residue of ∇ and denoted $\mathrm{Res}_{p_i}(\nabla)$. It just so happens that there is a choice of frame in which M is constant and equal to the matrix for the monodromy of ∇ around p_i (see [Sim90, page 725]).

There is a theory of parabolic Higgs bundles for reductive Lie groups, but here we are concerned with $G = \mathrm{SL}(2,\mathbb{R})\times\mathrm{SL}(2,\mathbb{R})$. We give definitions for $G = \mathrm{SL}(r,\mathbb{C})$, and then specialize. Given a holomorphic line bundle L and positive real numbers α_i for each p_i, we define the parabolic line bundle $L(\sum_i \alpha_i p_i)$ to be the underlying bundle $L(\sum_i [\alpha_i] p_i)$ with the weight $\alpha_i - [\alpha_i]$. Here, $[\cdot]$ is the integral part. The determinant of a parabolic vector bundle E_* is

$$\det(E_*) = \det(E) \otimes \Big(\mathcal{O}_S\Big(\sum_{i,j} \alpha_i^j p_i\Big)\Big).$$

Definition 4.5.5. A parabolic $\mathrm{SL}(r,\mathbb{C})$-Higgs bundle over (S, D) is a pair (E_*, Θ) consisting of a rank r parabolic vector bundle E_* over S with $\det(E_*)$ trivial and a holomorphic endomorphism $\Theta \in H^0(S; K_S(D) \otimes \mathrm{End}(E_*))$ called the Higgs field. The induced endomorphisms $\mathrm{Res}_{p_i}(\Theta) \in \mathrm{End}(E_*|_{p_i})$ of the filtered vector spaces $E_*|_{p_i}$ are called the residues of Θ at p_i.

In the definition above, the notation $K_S(D)$ means that we allow endomorphisms to be meromorphic on D. In this work, we only consider Higgs fields with poles of order at most one. The bundle is then called a regular parabolic Higgs bundle.

There are notions of degree, slope, and (semi, poly) stability for Higgs bundles. The parabolic degree of a parabolic bundle is

$$\mathrm{pdeg}(E_*) = \deg(E) + \sum_{i=1}^{n} \sum_{j=1}^{n_i} \alpha_i \dim(E_{p_i}^j / E_{p_i}^{j+1}).$$

The slope and (semi, poly) stability conditions are defined as in the case of Higgs bundles on closed surfaces, using this notion of degree (see [Sim90, Section 6]).

From representations to parabolic Higgs bundles

First we review the passage from representations to parabolic Higgs bundles in Simpson's non-abelian Hodge correspondence [Sim90]. Keeping the same notation as above, let $\tilde{\Sigma} = \Sigma/\Gamma$ be the universal cover and μ the hyperbolic metric on Σ that is compatible with the holomorphic structure from S.

Fixing a reductive representation $\rho : \Gamma \to \mathrm{SL}(r,\mathbb{C})$, we take the bundle $\mathbb{C}_\rho \to \Sigma$ with flat connection ∇. By a theorem of Koszul-Malgrange, the $(0,1)$-component of ∇ gives rise to a holomorphic structure $\overline{\partial}$. There is an extension to a parabolic bundle $E \to S$ with respect to which ∇ extends to a logarithmic connection such that the eigenvalues $a_i + ib_i$ of $\mathrm{Res}_{p_i}(\nabla)$ satisfy $0 \le a_1(p_i) < \cdots < a_{k_i}(p_i) < 1$ (see [Sim90, pages 718, 724]).

Given a Hermitian metric H on $V = E|_\Sigma \to \Sigma$, we can decompose the connection as

$$\nabla = \nabla_H + \Psi_H,$$

where ∇_H is an H-unitary connection and $\Psi_H \in \Omega^1(\Sigma, \mathrm{End}(V))$ is self-adjoint. A choice of ρ-equivariant harmonic map to the symmetric space $f : (\tilde{\Sigma}, \mu) \to \mathrm{SL}(r, \mathbb{C})/\mathrm{SO}(r, \mathbb{C})$ is equivalent to that of a harmonic Hermitian metric: one that satisfies $\nabla_H(*\Psi_H)$, where $*$ is the Hodge star on $\Omega^*(\Sigma, \mathrm{End}(V))$. We choose the harmonic map so that the Hopf differential has at most a pole of order 2 at each p_i. The equivalent condition for the harmonic metric is that it is tame: as a family of metrics in the symmetric space, it has at most polynomial growth (see [Sim90, Section 2]).

From these objects we can build a parabolic Higgs bundle (E_*, Θ): decompose the components of ∇ by type as

$$\nabla = \nabla_H^{1,0} + \nabla_H^{0,1} + \Psi_H^{1,0} + \Psi_H^{0,1}.$$

By Koszul-Malgrange again, there is a complex structure on V with del-bar operator $\overline{\partial}^V = \overline{\partial} - \Psi_H^{*_H}$, so that ∇_H is the Chern connection. Returning to our original bundle E, through this del-bar operator we obtain a new parabolic structure $\overline{\partial}^E$, with filtration at p_i,

$$E|_{p_i} = \mathrm{Eig}_{\geq a_1(p_i)}(\mathrm{Res}_{p_i}(\nabla_i)) \supset \cdots \supset \mathrm{Eig}_{\geq a_{n_i}(p_i)}(\mathrm{Res}_{p_i}(\nabla_i)) \supset 0,$$

and weights $\alpha_i^j = a_j(p_i)$. For this set of filtrations, $\Theta := \Psi_H^{1,0}$ is a Higgs field with a pole of order at most one at p_i [Sim90, page 723]. Here, the parabolic Higgs bundle is polystable. The flat condition means that H solves Hitchin's equation

$$F_{\nabla^H} + [\Theta, \Theta^{*_H}] = 0,$$

where F_{∇^H} is the curvature and Θ^{*_H} is the H-adjoint of Θ. Simpson's work in [Sim90, Chapters 5-7] shows that from a polystable parabolic Higgs bundle (E_*, Θ) solving the Hitchin equation with at most first order poles, one can recover the data of a parabolic bundle (and hence a representation) with a harmonic metric.

Now we take $r = 2$. Assume ρ is irreducible; we leave the general case to the reader. Reorder the punctures as $p_1, \ldots, p_m, p_{m+1}, \ldots, p_n$, so that ρ takes the monodromy around $p_1, \ldots, p_m$ to hyperbolic isometries. We assemble the parabolic vector bundle associated to $(\mathbb{C}_\rho, \nabla)$. Choose twist parameters $\theta = (\theta_1, \ldots, \theta_m) \in \mathbb{R}^m$ and form the harmonic map f^θ. In model charts around p_i, recall the connection may be expressed as $\nabla = d + \rho(\zeta_i)dz/z$. The associated harmonic metric converges to a "model metric" as we approach each p_i. That is, in the nice choice of charts for ∇, the metric has a particular form that depends

- only on the image $\rho(\zeta_i)$ if $\rho(\zeta_i)$ is elliptic or parabolic, and
- on the image $\rho(\zeta_i)$ and the choice of twist parameter.

The model metrics in the case of elliptic and parabolic mondoromy are worked out in [Mon16, Examples 5 and 7]. For hyperbolic isometries, replace μ with the flat cylinder metric μ^f and take a cusp U around p_i as in (4.2). Let

$$C_\infty = \{z = x + iy : x \in [0,1], y \in [0,\infty)\}\langle z \mapsto z+1\rangle$$

be a half-infinite cylinder and $k_r : C_\infty \to U$ a conformal mapping into U that takes $[0,1] \times \{0\}$ to $[0,1] \times \{r\}$ linearly and then extends vertically. It is shown in [Sag19, Section 5.2] that as $r \to \infty$, $f^\theta \circ k_r : C_\infty \to U$ converges in the C^∞ sense to a harmonic map with constant Hopf differential $\ell(\rho(\zeta_i))^2/4$. This shows that, in the model coordinates for ∇, the harmonic metric is determined only by the twist parameter.

Given the data of the representation, the Higgs field Θ depends only on the harmonic map. For an $\mathrm{SL}(r,\mathbb{C})$ Higgs bundle,

$$\Phi(f^\theta) = 2r\mathrm{trace}_H(\Theta^2). \tag{4.20}$$

In this way, for our $r = 2$ case, the residue of the Higgs field and the twist parameter for the harmonic map determine each other.

The circle bundle U

Assume we have representations $\rho_1, \rho_2 : \Gamma \to \mathrm{PSL}(2,\mathbb{R})$ that lift to representations to $\mathrm{SL}(2,\mathbb{R}) \subset \mathrm{SL}(2,\mathbb{C})$. Since ρ_1, ρ_2 lie in the split real form for $\mathrm{SL}(2,\mathbb{C})$, we construct real bundles $\mathbb{R}^2_{\rho_k}$ with volume forms ω_k. We can complexify to obtain bundles $(E_k^{\mathbb{C}}, \nabla_k^{\mathbb{C}}, \omega_k^{\mathbb{C}})$. As a representation of $\mathrm{SL}_2(\mathbb{R}) \times \mathrm{SL}_2(\mathbb{R})$, the tensor $\rho_1 \otimes \rho_2$ gives a real vector bundle (E, ∇, Q), where $E = E_1 \otimes E_2$, $\nabla = \nabla_1 \otimes \nabla_2$, $Q = -\omega_1 \otimes \omega_2$, and this yields a complexification $(E^{\mathbb{C}}, \nabla^{\mathbb{C}}, Q^{\mathbb{C}})$ of (E, ∇, Q) that is compatible with the tensor products. We then attach Simpson's parabolic structure.

Remark 4.5.6. In [AL18], the signature $(2,2)$ bilinear form Q is defined as $\omega_1 \otimes \omega_2$. We use a minus sign so that our conventions agree with that of Section 4.4.

We now choose harmonic maps $h^{\theta_1}, f^{\theta_2}$ for ρ_1, ρ_2 and form the associated Higgs bundles $((E_k^{\mathbb{C}})_*, \overline{\partial}_{E^k}, \Theta_k)$. From [Mon16, Lemma 3.11], the reality the representation implies that $(E_k^{\mathbb{C}})_*$ splits as a holomorphic parabolic line bundle and its inverse:

$$(E_1^{\mathbb{C}})_* = L_* \oplus L_*^{-1}\ ,\ (E_2^{\mathbb{C}})_* = N_* \oplus N_*^{-1}.$$

The parabolic degrees can be expressed in terms of the relative Euler numbers of the representations (see [BIW10, Theorem 12]),

$$\mathrm{eu}(\rho_1) = 2\mathrm{pdeg}(L_*)\ ,\ \mathrm{eu}(\rho_2) = 2\mathrm{pdeg}(N_*). \tag{4.21}$$

With respect to the splittings, the harmonic metrics H_k are diagonal and satisfy $\det(H_k) = 1$, so that $H_k = \mathrm{diag}(h_k^{-1}, h_k)$, with $h_1 \in H^0(\Sigma, L_*^{-1} \otimes L_*)$, $h_2 \in H^0(\Sigma, N_*^{-1} \otimes N_*)$. The Higgs fields take the form

$$\Theta_1 = \begin{pmatrix} 0 & \alpha \\ \beta & 0 \end{pmatrix}\ ,\ \Theta_2 = \begin{pmatrix} 0 & \gamma \\ \delta & 0 \end{pmatrix},$$

where $\alpha \in H^0(\Sigma, L_*^2 \otimes K_*), \beta \in H^0(\Sigma, L_*^{-2} \otimes K_*), \gamma \in H^0(\Sigma, N_*^2 \otimes K_*), H^0(\Sigma, N_*^{-2} \otimes K_*)$. The bundles also come equipped with real structures for the complex structure of $\overline{\partial}^E$: $\tau_k(v) = H_k^{-1}\omega_k \overline{v}$.

As for $\rho_1 \otimes \rho_2 : \Gamma \to \mathrm{SL}(2, \mathbb{R}) \times \mathrm{SL}(2, \mathbb{R}) \subset \mathrm{SL}(4, \mathbb{C})$, Simpson's correspondence is functorial with respect to the tensor product. If $\{e_k, e_k^*\}$ are local holomorphic frames for the flat $\mathrm{SL}(2, \mathbb{R})$-connections $\nabla_k^{\mathbb{C}}$, then $\{e_1 \otimes e_2, e_1 \otimes e_2^*, e_1^* \otimes e_2, e_1^* \otimes e_2^*\}$ furnishes a frame for the flat connection $\nabla^{\mathbb{C}}$. Taking the associated $\mathrm{SL}(4, \mathbb{C})$-Higgs bundle, we write all of our data down in this frame.

- The parabolic vector bundle

$$E_*^{\mathbb{C}} = (E_1^{\mathbb{C}})_* \otimes (E_2^{\mathbb{C}})_* = (L_* \otimes N_*) \oplus (L_* \otimes N_*^{-1}) \otimes (L_*^{-1} \otimes N_*) \oplus (L_*^{-1} \otimes N_*^{-1}).$$

- The harmonic metric $H = H_1 \otimes H_2 = \mathrm{diag}(h_1^{-1}h_2^{-1}, h_1 h_2^{-1}, h_1^{-1}h_2, h_1 h_2)$.

- The Higgs field

$$\Theta = \Theta_1 \otimes I + I \otimes \Theta_2 = \begin{pmatrix} 0 & \gamma & \alpha & 0 \\ \delta & 0 & 0 & \alpha \\ \beta & 0 & 0 & \gamma \\ 0 & \beta & \delta & 0 \end{pmatrix}.$$

- The bilinear form

$$Q = -\omega_1 \otimes \omega_2 = \frac{1}{2}\begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}.$$

- The real structure $\tau = \tau_1 \otimes \tau_2$

$$\begin{pmatrix} v_1 \\ v_2 \\ v_3 \\ v_4 \end{pmatrix} \mapsto \begin{pmatrix} 0 & 0 & 0 & h_1 h_2 \\ 0 & 0 & h_1^{-1} h_2 & 0 \\ 0 & h_1 h_2^{-1} & 0 & 0 \\ h_1^{-1} h_2^{-1} & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \overline{v}_1 \\ \overline{v}_2 \\ \overline{v}_3 \\ \overline{v}_4 \end{pmatrix}.$$

The key observation from [AL18], which extends to the parabolic setting, is that the real structure τ leaves invariant the two subbundles $(L_* \otimes N_*^{-1}) \otimes (L_*^{-1} \otimes N_*)$ and $(L_* \otimes N_*) \oplus (L_*^{-1} \otimes N_*^{-1})$. The degrees are communicated by (4.21), and hence depend on relative Euler numbers [BIW10]. This is related to the structure of the bundle.

Now restrict $(L_* \otimes N_*^{-1}) \otimes (L_*^{-1} \otimes N_*)$ and $(L_* \otimes N_*) \oplus (L_*^{-1} \otimes N_*^{-1})$ to Σ, effectively forgetting the parabolic structure for now, and set F_1, F_2 to be the real parts with respect τ. The proposition below is proved exactly as in [AL18, Proposition 3.2].

Proposition 4.5.7. *The real vector bundle $E \to \Sigma$ splits as a Q-orthogonal direct sum of two rank 2 real bundles $E = F_1 \otimes F_2$. With respect to Q, F_1 is timelike and F_2 is spacelike.*

We are now prepared to define the circle bundle U: it is the timelike unit circle of F_1,

$$U = \{v \in F_1 : Q(v, v) = -1\}.$$

AdS structures

We define a bundle $\mathrm{AdS}^3_\rho \to \Sigma$ to be the bundle whose fibers are the points in each fiber of $\mathbb{R}^4_{\rho_1 \otimes \rho_2} \to \Sigma$ with $Q(v, v) = -1$. This is well-defined because $\rho_1 \otimes \rho_2$ preserves Q. Each fiber is a copy of AdS^3, and there is a natural inclusion $U \to \mathrm{AdS}^3_\rho$.

$\mathrm{AdS}^3_\rho \to \Sigma$ pulls back to a bundle $p^*\mathrm{AdS}^3_\rho$ over U with respect to the projection $p : U \to \Sigma$. There is a tautological section $s : U \to p^*\mathrm{AdS}^3_\rho$ that reframes each $v \in U$ as a point in $p^*\mathrm{AdS}^3_\rho$. For the reader's convenience, we write s in local coordinates. For $v \in U$, $v \in (L_* \otimes N_*^{-1}) \otimes (L_*^{-1} \otimes N_*)$ means $v = (0, v_1, v_2, 0)^T$. The conditions $\tau v = v$ and $Q(v, v) = -1$ yield $|v_1| = (h_2 h_1^{-1})^{-1/2}$, $|v_1| = (h_2 h_1^{-1})^{1/2}$. Thus, in a local trivialization (z, θ) over an open set $\Omega \times S^1$,

$$s(v) = \begin{pmatrix} 0 \\ (h_2 h_1^{-1})^{-1/2} e^{i\theta} \\ (h_2 h_1^{-1})^{1/2} e^{-i\theta} \\ 0 \end{pmatrix}.$$

The section s pulls back via the universal covering $\tilde{U} \to U$ to a section of the pullback bundle, which yields a mapping $\tilde{D} : \tilde{U} \to \mathrm{AdS}^3$.

Remark 4.5.8. We are implicitly using the notion of a graph of a geometric structure here. See [Ale19, Section 4] for details.

For closed surfaces, a computation in local coordinates [AL18, Theorem 7.1] shows that when $\tilde{D}$ is the developing map of an AdS^3 structure, the preimage of a circle fiber develops bijectively to a timelike geodesic. The proof actually works on any surface, independent of if $\tilde{D}$ is an immersion. So we are welcome to study the associated map from $(\tilde{\Sigma}, \tilde{\mu}) \to (\mathrm{Gr}^{\mathrm{t}}(2,4), m) = (\mathbb{H} \times \mathbb{H}, \sigma \oplus (-\sigma))$ from Proposition 4.5.2 instead. As a map into the Grassmanian, this is equal to $s \wedge \nabla_\theta s$, and we can derive an explicit expression in coordinates. We write the map into the Grassmanian as G, and the equivalent map into $(\mathbb{H} \times \mathbb{H}, \sigma \oplus (-\sigma))$ as F.

Proposition 4.5.9. *The associated mapping* $F : (\tilde{\Sigma}, \tilde{\mu}) \to (\mathbb{H} \times \mathbb{H}, \sigma \oplus (-\sigma))$ *is exactly* $(h^{\theta_1}, f^{\theta_2})$.

Proof. From the computation in local coordinates on $\mathrm{Gr}^{\mathrm{t}}(2,4)$ from [AL18, Theorem 7.3], G is harmonic. Upon passing to $(\mathbb{H} \times \mathbb{H}, \sigma \oplus (-\sigma))$, F splits as a product of harmonic maps $F = (h, f)$. We will compute the Laurent expansion of the Hopf differentials of h and f in the usual uniformized punctured disk, see the poles have order 2, and appeal to the uniqueness in [Sag19, Theorem 1.1]. The following observations allow us to do so neatly.

- $\Phi_m(G) = \Phi_{\sigma \oplus (-\sigma)}(F) = \Phi(h) - \Phi(f)$.
- Consider the signature-$(4,2)$ symmetric bilinear form on $\wedge^2 \mathbb{R}^4$ defined as

$$B(v_1 \wedge v_2, v_3 \wedge v_4) = 4Q(v_1, v_4)Q(v_2, v_3) - 4Q(v_1, v_3)Q(v_2, v_4). \quad (4.22)$$

B has signature $(4,2)$, and Torralbo [Tor07] shows that with the metric m' on $\mathrm{Gr}^{\mathrm{t}}(2,4)$ induced by inclusion into $(\wedge^2\mathbb{R}^4, B)$, $(\mathrm{Gr}^{\mathrm{t}}(2,4), m')$ identifies with $(\mathbb{H}\times\mathbb{H}, \sigma\oplus\sigma)$ (he uses a slightly different form, but we appeal to Sylvester's law). With respect to this metric, G is also harmonic and $\Phi_{m'}(G) = \Phi_{\sigma\oplus\sigma}(F) = \Phi(h) + \Phi(f)$.

Above we use subscripts m, m' to designate Hopf differentials for the target metric, and the target metric for h and f is understood to be σ. From this, we see

$$\Phi(h) = \frac{1}{2}(\Phi_{m'}(G) + \Phi_m(G))\ ,\ \Phi(f) = \frac{1}{2}(\Phi_{m'}(G) - \Phi_m(G)),$$

so it suffices to compute these Hopf differentials in the coordinates on the Grassmanian, where the Higgs bundle gives us nice expressions. Write $\nabla^{\mathbb{C}}_{\partial/\partial z}s = s_z$, $\nabla^{\mathbb{C}}_{\partial/\partial\theta}s = s_\theta$, and likewise for G. In the frame $\{e_1\otimes e_2, e_1\otimes e_2^*, e_1^*\otimes e_2, e_1^*\otimes e_2^*\}$,

$$\begin{aligned}\nabla^{\mathbb{C}}_{\frac{\partial}{\partial z}} &= \partial + H^{-1}\partial H + \Phi \\ &= \partial + \begin{pmatrix} \partial\log(h_2^{-1}h_1^{-1}) & \gamma & \alpha & 0 \\ \delta & \partial\log(h_2 h_1^{-1}) & 0 & \alpha \\ \beta & 0 & \partial\log(h_2^{-1}h_1) & \gamma \\ 0 & \beta & \delta & \partial\log(h_1 h_2)\end{pmatrix}\end{aligned}$$

$\Phi_m(G) = m(G_z, G_z)$, same for m', and

$$G_z = (s\wedge s_\theta)_z = s_z\wedge s_\theta + s\wedge s_{\theta,z}.$$

In [AL18], all of these are computed in local coordinates: set $g = (h_2h_1^{-1})^{-1/2}$, $c = ig^{-1}\partial g$, $X = \gamma g e^{i\theta} + \alpha g^{-1}e^{-i\theta}$, $Y = \beta g e^{i\theta} + \delta g^{-1}e^{-i\theta}$, $X' = i\gamma g e^{i\theta} - i\alpha g^{-1}e^{-i\theta}$, $Y' = i\beta g e^{i\theta} - i\delta g^{-1}e^{-i\theta}$. Then, in a trivialization over an open set $\Omega\times S^1$ with coordinate (z,θ),

$$s = \begin{pmatrix} 0 \\ ge^{i\theta} \\ g^{-1}e^{-\theta} \\ 0\end{pmatrix},\ s_\theta = \begin{pmatrix} 0 \\ ige^{i\theta} \\ -ig^{-1}e^{-\theta} \\ 0\end{pmatrix},\ s_z = \begin{pmatrix} X \\ 0 \\ 0 \\ Y\end{pmatrix} + cs_\theta\ ,\ s_{\theta,z} = \begin{pmatrix} X' \\ 0 \\ 0 \\ Y'\end{pmatrix} - cs.$$

As in [AL18], we thus find $\Phi_m(G)$ is the determinant of the matrix specified by

$$G_z\wedge G_z = 2s_z\wedge s_\theta\wedge s\wedge s_{\theta,z},$$

which is $8(\gamma\delta - \alpha\beta)$. For $\mathrm{Hopf}_{m'}(G)$,

$$B(G_z, G_z) = B(s_z\wedge s_\theta, s_z\wedge s_\theta) + B(s_z\wedge s_\theta, s\wedge s_{\theta,z}) + B(s\wedge s_{\theta,z}, s_z\wedge s_\theta) + B(s\wedge s_{\theta,z}, s\wedge s_{\theta,z}).$$

In view of (4.22), we record $Q(s,s) = Q(s_\theta, s_\theta) = -1$, $Q(s_z, s_z) = -c^2 + XY$, $Q(s_{\theta,z}, s_{\theta,z}) = -c^2 + X'Y'$, $Q(s, s_\theta) = 0$, $Q(s, s_z) = 0$, $Q(s_\theta, s_z) = -c$. From this we calculate

$$B(s_z \wedge s_\theta, s_z \wedge s_\theta) = 4XY \ , \ B(s \wedge s_{\theta,z}, s \wedge s_{\theta,z}) = 4X'Y',$$

and

$$B(s_z \wedge s_\theta, s \wedge s_{\theta,z}) = B(s \wedge s_{\theta,z}, s_z \wedge s_\theta) = 0.$$

Thus,

$$\Phi_{m'}(G) = 4(XY + X'Y') = 8(\gamma\delta + \alpha\beta).$$

From (4.20), $\Phi(h) = 8\alpha\beta$, $\Phi(f) = 8\gamma\delta$. Meanwhile, using the Higgs fields for the harmonic metrics coming from h^θ and f^θ, $\Phi(h^{\theta_1}) = 8\alpha\beta$, $\Phi(f^{\theta_2}) = 8\gamma\delta$. Passing to the parabolic structures, we see the residues of the Hopf differentials via the residues of the Higgs fields. We conclude $h = h^{\theta_1}$, $f = f^{\theta_2}$, as desired. □

Definition 4.5.10. (ρ_1, θ_1) metric dominates (ρ_2, θ_2) if $(h^{\theta_1})^*\sigma > (f^{\theta_2})^*\sigma$, where $h^{\theta_1}, f^{\theta_2}$ are the harmonic maps for ρ_1, ρ_2 respectively.

$g_1 > g_2$ means $g_1 - g_2$ is a positive definite metric. Better yet, by Proposition 4.5.9 it is equivalent to F being a spacelike immersion. The proposition below thus follows immediately.

Proposition 4.5.11. *The map $\tilde{D}$ defines a developing map if and only if we have metric domination.*

And moreover, Theorem 4A gives the next result.

Proposition 4.5.12. *Choosing $\theta_1 = \theta_2$, there is a choice of conformal structure on Σ such that we have metric domination if and only if ρ_1 almost strictly dominates ρ_2.*

The upshot is that when we have almost strict domination, we can construct this circle bundle U so that the induced map $\tilde{U} \to \mathrm{AdS}^3$ is the developing map. The computation of the spacelike immersion shows the bundle is the same one as in Theorem 4C.

Volume

Following the local computations from [AL18], the Higgs bundles can be used to write down the AdS metric $g = (g_{ij})$ in terms of the harmonic maps data, and also to compute the volume. Explicitly, one arrives at the formula

$$\mathrm{Vol}(g_{ij}) = \pi \int_{\Sigma} (J(h) + J(f)) dx \wedge dy, \tag{4.23}$$

where h and f are the associated harmonic maps, and $J(\cdot)$ denotes the Jacobian determinant. Since infinite energy harmonic maps converge exponentially to projections onto a geodesic, both of these Jacobians can be integrated over the surface. They also satisfy the requirements from [KM08], so that the integral is the volume of the representation, depending only on the representation and computed via the relative Euler number and the total rotation (see [BIW10] for details). When the AdS^3 geometric structure is complete, this recovers Tholozan's formula for the volume of AdS^3 quotients (see [Tho18, Theorem 1]). We leave it to the interested reader to check, but it turns out that we can also arrive at the formula (4.23) above by imitating Tholozan's integration over the fibers in [Tho18].

Chapter 5

THE FACTORIZATION THEOREM

5.1 Introduction

A harmonic map $f : (\Sigma, \mu) \to (M, \nu)$ is admissible if its image is not contained in a geodesic. There is a viewpoint that while admissible harmonic maps are abundant in many contexts, they also reveal rigid geometric properties of the spaces on which they live. The result of this chapter is another instance of this phenomenon. It connects local behaviour of a harmonic map to the global complex geometry of the underlying Riemann surface.

Theorem 5A. Suppose $f : (\Sigma, \mu) \to (M, \nu)$ is an admissible harmonic map, and there is a conformal diffeomorphism $h : \Omega_1 \to \Omega_2$ between open subsets of Σ such that $f \circ h = f$ on Ω_1. If h is holomorphic, then there is a Riemann surface (Σ_0, μ_0), a holomorphic map $\pi : \Sigma \to \Sigma_0$, and a harmonic map $f_0 : (\Sigma_0, \mu_0) \to (M, \nu)$ such that $\pi(\Omega_1) = \pi(\Omega_2)$ and f factors as $f = f_0 \circ \pi$. If h is anti-holomorphic, Σ_0 is a Klein surface and π is dianalytic.

Among other solutions to geometrically flavoured PDEs, Theorem 5A has been known for minimal harmonic maps and pseudoholomorphic curves since the 1970s. Osserman in [Oss70] and Gulliver in [Gul73] studied singularites of the Douglas and Rado solutions to the Plateau problem. The only possible singularities are branch points, which are separated into so-called true branch points and false branch points. Osserman ruled out true branch points and made progress toward the non-existence of false branch points in [Oss70], and Gulliver showed there are no false branch points in [Gul73]. Alt also proved the result of Gulliver independently and in greater generality in [Alt72] and [Alt73]. This work proves that the Douglas and Rado solutions are immersed. For an exposition of the Plateau problem, see [Nit74], [DHS10, Chapter 4], and [DHT10, Chapter 4].

Curiously, very few properties specific to minimal surfaces come into play in [Gul73], but rather qualities shared by a larger class of surfaces. This prompted a deeper study of branched immersions of surfaces, which was carried out by Gulliver-Osserman-Royden in [GOR73]. A version of Theorem 5A holds for the maps considered in [GOR73]. In the next subsection we describe their theory of branched immersions of surfaces and how minimal maps fit into the framework.

Aside from connections to the Plateau problem, the result of Gulliver-Osserman-Royden has other applications. We would like to highlight the work of Moore in [Moo06] and [Moo17], where he studies moduli spaces of minimal surfaces. A map f is somewhere injective if there is a regular point p such that $f^{-1}(f(p)) = p$. Moore uses Theorem 5A for minimal maps to show that a closed minimal map in an n-manifold, $n \geq 3$, is not somewhere injective if and only if it factors through a conformal branched covering map. The same result holds for pseudoholomorphic curves [MS12, Proposition 2.5.1], whose moduli spaces are an active field of study.

If (Σ, μ) is closed with genus at least 2 and (M, ν) has negative curvature, then Σ_0 must have genus at least 2. The described results for minimal surfaces thus show the somewhere injective condition is generic, for it is very rare for a closed Riemann surface to admit a holomorphic map onto another Riemann surface with non-abelian fundamental group.

In the next chapter, we use Theorem 5A to show that somewhere injective harmonic maps are generic moduli spaces of harmonic maps. In [Moo06], [MS12], and [Moo17], as well as our own work, the somewhere injective condition plays a role in various transversality arguments.

In a different inquiry, Jost and Yau proved a version of Theorem 5A in [JY83] for harmonic maps to Kähler manifolds, using it as a tool in their study of deformations of Kodaira surfaces. Their work has played a role in the development of the theories of Kähler manifolds and Higgs bundles. See the survey of Jost [Jos08] for more information.

Minimal surfaces

Loosely following the exposition of Moore in section 4 of [Moo06], we explain how the proof of Theorem 5A for minimal maps is deduced from the results in [GOR73]. Let $f : (\Sigma, \mu) \to (M, \nu)$ be a C^1 map. A point $p \in \Sigma$ is a branch point if $df(p) = 0$. We say a branch point is a good branch point of order $m - 1$ if there is a choice of coordinates z on Σ and $(x_1, \ldots, x_n)$ on M such that f is described by the equations

$$x_1 = \mathrm{re} z^m \ , \ x_2 = \mathrm{im} z^m \ , \ x_k = \eta_k(z) \ , \ k \geq 3,$$

where $\eta_k \in o(|z|^m)$. Note that $m = 1$ implies we have a regular point.

Remark 5.1.1. These conventions could be a source of confusion. In [GOR73], Gulliver-Osserman-Royden refer to "good branch points" as simply "branch points." This causes no harm in their work, but we should distinguish here.

In [GOR73], a branched immersion is a map from a surface that is regular everywhere apart from an isolated set of good branch points. For clarity we refer to such a map here as a good branched immersion. Gulliver-Osserman-Royden use the representation formula of Hartman and Wintner [HW53] to show that a minimal map is a good branched immersion (see Propositions 2.2 and 2.4 in [GOR73]). In fact, using this same formula, Micallef and White recover finer coordinate expressions for minimal surfaces (see [MW95, Theorem 1.]).

An order $m-1$ branch point p of a good branched immersion is ramified of order $r-1$ if r is the maximal non-negative integer such that there is a disk U centered at p on which f factors through a branched covering of degree r. If $r = m$, p is called a false branch point, and true otherwise. We say f is unramified if $r = 0$. We now recast one of the key results of [GOR73].

Theorem 5.1.2 (Proposition 3.19 in [GOR73])**.** *Let Σ be a C^1 surface, M a C^1 manifold, and $f : \Sigma \to M$ a C^1 good branched immersion with the unique continuation property and no true branch points. Then there is a C^1 surface Σ_0, a C^1 good branched immersion $\pi : \Sigma \to \Sigma_0$, and an unramified C^1 good branched immersion $f_0 : \Sigma_0 \to M$ such that $f = f_0 \circ \pi$.*

We do not define the unique continue property of Gulliver-Osserman-Royden (see [GOR73, page 757]), but remark that minimal maps have this property (see [GOR73, Lemma 2.10]). The minimal case is essentially handled in [GOR73, Proposition 3.24]. If a map is conformal, one can dispense of the hypothesis that there are no true branch points, and the objects π, Σ_0, and f_0 all have the same regularity as f apart from at branch points and images of branch points.

To prove Theorem 5.1.2, Gulliver-Osserman-Royden define a relation $\sim$ on Σ as follows.

1. If p_1 and p_2 are regular points for f, $p_1 \sim p_2$ if there exists open sets Ω_i containing p_i, and an orientation preserving C^1 map $h : \Omega_1 \to \Omega_2$ such that $f \circ h = f$ on Ω_1.

2. If one of p_1 or p_2 is a branch point, then in any pair of neighbourhoods Ω_i containing p_i there exists neighbourhoods $\Omega_i' \subset \Omega_i$ of p_i consisting of only regular points such that for all $p_1' \in \Omega_1' \backslash \{p_1\}$ there exists $p_2' \in \Omega_2' \backslash \{p_2\}$ such that $p_1' \sim p_2'$, and for all $p_2' \in \Omega_2' \backslash \{p_2\}$ there exists $p_1' \in \Omega_1' \backslash \{p_1\}$ such that $p_1' \sim p_2'$.

Gulliver-Osserman-Royden show that this is an equivalence relation and define the quotient $\pi : \Sigma \to \Sigma_0$. They prove Σ_0 has the structure of a C^1 manifold and the map $f_0 : \Sigma \to M$ is defined by setting $f_0([p]) = f(p)$. Ramification leads to equivalent points, so f_0 is unramified.

When Σ is a Riemann surface and Σ and M are equipped with metrics so that f is minimal, we impose that h is holomorphic. Following the proof of [GOR73, Proposition 3.24], one can show that the transition maps on Σ_0 are holomorphic away from the branch points and extend holomorphically via the removeable singularities theorem. One checks in coordinates that the map f_0 is minimal with respect to the conformal metric on Σ_0 obtained via uniformization. The existence of a map h as in Theorem 5A amounts to saying some classes under $\sim$ are not singletons. The minimal case follows directly.

Gulliver-Osserman-Royden do not consider orientation reversing maps in the definition of $\sim$, but their construction can be modified to allow for this. In this situation, we may end up with a mapping onto a non-orientable surface. Moore notes this in [Moo06], although his context is slightly different from ours. Since we could not locate a formal proof in the literature, we explain the necessary adjustments the the end of Section 5.4.

Harmonic maps vs. minimal maps

To prove Theorem 5A in the holomorphic case, we follow the blueprint of Gulliver-Osserman-Royden. That is, we define an equivalence relation on Σ and take the quotient as our candidate for the surface Σ_0. However, it is not obvious how one should define $\sim$. The difficulty comes from the singularities of harmonic maps, in that

1. harmonic maps can have rank 1 singularities, which do not occur in the theory of Gulliver-Osserman-Royden, and

2. branch points are not good branch points. At best, we can combine the Hartman-Wintner formula with [Che76, Lemma 2.4] to see that near a branch

point p of order $m-1$ there is a C^1 coordinate z on the source and a C^∞ coordinate on the target such that $p \mapsto 0$, $f(p) \mapsto 0$, and f may be expressed $f = (f^1, \dots, f^n)$ with

$$f^1 = p^1 \ , \ f^k = p^k + r^k \ , \ k \geq 2,$$

where p^1 is a spherical harmonic of degree m, p^k is a spherical harmonic of degree at least m, and $r^k \in o(|z|^m)$.

To overcome these difficulties, we exploit the geometry of the Hopf differential. In some sense, the Hopf differential treats rank 1 and 2 points on an equal footing. Thus, if we define $\sim$ in terms of a condition on the Hopf differential, in theory we shouldn't encounter any difficulties due to rank 1 singularities. In practice this is mostly true—at some points we need to refer to the Hartman-Wintner formula. As for (ii), the Hopf differential defines a "natural coordinate" for the harmonic map near a branch point, in which the geometry can be more easily probed. At a false branch point, we see ramification behaviour similar to that displayed by minimal maps.

The only missing piece of Gulliver-Osserman-Royden's theory is the unique continuation property. In Proposition 5.2.3, we show that analytic continuation of natural coordinates for the Hopf differential induces a continuation of h. Using this proposition, we establish a "holomorphic unique continuation property" (Proposition 5.3.5).

Future directions.

It is tempting to conjecture that some version of Theorem 5A should hold without the hypothesis that h is conformal. The main motivation would be to improve our understanding of somewhere injective harmonic maps. We would like to point out that, in view of the example below, we cannot expect the map π to be holomorphic with respect to a complex structure on Σ_0.

Let (Σ_0, μ_0) be a closed hyperbolic surface and $f_0 : (\Sigma_0, \mu_0) \to (M, \nu)$ a totally geodesic map. Fix a smooth surface Σ of genus at least 2 and a homotopy class of maps $\mathbf{f} : \pi_1(\Sigma) \to \pi_1(\Sigma_0)$ with degree at least 2. Any C^2 metric μ yields a unique harmonic map $\pi : (\Sigma, \mu) \to (\Sigma_0, \mu_0)$ in the homotopy class $\mathbf{f}$. One can then find many diffeomorphisms $h : \Omega_1 \to \Omega_2$ between open subsets of Σ such that $f \circ h = f$, and by construction f factors as $f = f_0 \circ \pi$. Generically, the surface (Σ, μ) will not admit a holomorphic map onto any Riemann surface of genus at least 2.

We simplify our study of singularities using complex analytic methods. Without the conformal hypothesis, the only local information we have comes from the Hartman-Wintner representation formula. If this is the main tool, then it is also natural to ask about more general solutions to second order semiliinear elliptic systems, rather than just harmonic maps. An analysis of singularities would be related to understanding local behaviour of spherical harmonics.

A substitute for the unique continuation property seems to be a large hurdle. Implicit in the proof of the unique continuation property for minimal maps is the following result (see [GOR73, Lemma 2.10]).

Proposition 5.1.3. *Let $\mathbb{D} \subset \mathbb{R}^2$ be the unit disk. Suppose $u_1, u_2 : \mathbb{D} \to M$ are minimal maps such that, for all open sets $D_1 \subset \mathbb{D}$ containing 0, there is an open subset $D_2 \subset \mathbb{D}$ (possibly not containing 0) such that $u^2(D_2) \subset u^1(D_1)$. Then there exists an open subset $D' \subset \mathbb{D}$ containing 0 such that $u^2(D') \subset u^1(\mathbb{D})$.*

This result above fails emphatically if we replace minimal maps with harmonic maps, even if $M = \mathbb{R}^2$. Indeed, the simple example

$$u_1(x, y) = (x, xy) \ , \ u_2(x, y) = (x, y)$$

does not satisfy Proposition 5.1.3. On the other hand, our "holomorphic unique continuation property" provides a substitute for Proposition 5.1.3 (see Proposition 5.3.5). This is one of the reasons we expect a more general version of Theorem 5A to be much more delicate, and we defer this investigation to a future project.

Acknowledgements
Many thanks to Vlad Marković for encouragement and sharing helpful ideas. I would also like to thank John Wood and Jürgen Jost for comments on earlier drafts.

5.2 Local properties of harmonic maps

We'll work a lot with singular flat metrics and natural coordinates. For a holomorphic quadratic differential Φ on Σ, a disk of radius r centered at a point p in the Φ-metric shall be called a Φ-disk and written $B_r(p)$. The induced distance function is denoted $d(\cdot, \cdot)$. Although we work with different differentials in the course of our work, the use of this notation in context should be clear.

Analytic continuation

Until Section 5.4, let $f : (\Sigma, \mu) \to (M, \nu)$ be an admissible harmonic map with non-zero Hopf differential Φ and let $h : \Omega_1 \to \Omega_2$ be a holomorphic map as in the statement of Theorem 5A. We treat anti-holomorphic maps in Section 5.3. Throughout the chapter, we let $\mathcal{Z}$ denote the zero locus of Φ. By restricting, we assume Ω_1 is a Φ-disk.

We use the geometry of the Hopf differential to analytically continue h. Let $p \in \Omega_1$ be such that $\Phi(p) \neq 0$, and let $U \subset \Omega_1$ be an open subset containing p such that $\Phi \neq 0$ in U. Given a holomorphic local coordinate z in U, we define a local coordinate w on $h(U)$ by $w = z \circ h^{-1}$. In these coordinates, h is given by $w(h(z)) = z$ and

$$df_p\Big(\frac{\partial}{\partial z}\Big) = df_{h(p)}\Big(\frac{\partial}{\partial w}\Big) \in T_{f(z)}M \otimes \mathbb{C}.$$

Remark 5.2.1. Here we are viewing df as a map from $T\Sigma \to TM$ rather than as a section of the endomorphism bundle $T^*\Sigma \otimes f^*TM$.

Choosing z to be a natural coordinate with $z(p) = 0$, we obtain

$$\langle f_w, f_w \rangle (w(h(z))) = \langle f_z, f_z \rangle (z) = 1.$$

Therefore, w defines a natural coordinate on $h(U)$. We have proved the following lemma.

Lemma 5.2.2. *h is a local isometry in the Φ-metric. If Ω_1 is a Φ-disk then so is Ω_2, and h takes a natural coordinate z on Ω_1 to a natural coordinate w on Ω_2 in which $w(h(z)) = z$.*

The goal of this subsection is to prove the proposition below. In the proof we use the notion of a maximal Φ-disk. See section 5 in [Str84] for a detailed discussion on maximal Φ-disks. Let $\mathcal{Z}$ denote the zero set of Φ (which is isolated).

Proposition 5.2.3. *Suppose Ω_1, Ω_2 are Φ-disks with no zeros of Φ and that $\gamma : [0, L] \to \Sigma$ is a curve starting in Ω_1 and that γ first strikes $\partial\Omega_1$ at a point q. If there is an $\epsilon > 0$ such that*

$$min\Big\{ \inf_{s \in \gamma|_{\Omega_1}, t \in \mathcal{Z}} d(s,t), \inf_{s \in \gamma|_{\Omega_1}, t \in \mathcal{Z}} d(h(s),t) \Big\} \geq \epsilon$$

then there is a neighbourhood of q in which h can be analytically continued along γ.

Proof. We can choose an arc on $\partial\Omega_1$ centered at q on which $\Phi \neq 0$. We then connect the endpoints via an arc contained inside Ω_1 so that the enclosed region U is a topological disk. We pick these arcs in such a way that

$$\min\big\{ \inf_{s\in U, t\in \mathcal{Z}} d(s,t), \inf_{s\in U, t\in \mathcal{Z}} d(h(s),t)\big\} \geq \epsilon/2.$$

The restriction of the Φ-metric to any compact region that does not intersect $\mathcal{Z}$ is complete. As h is an isometry in the Φ-metric, we can extend it to a map $h : \overline{U} \to \overline{U}$. Therefore, we have a well-defined point $h(q)$.

For every point $p \notin \mathcal{Z}$, there is a maximal radius r_p such that we can extend any natural coordinate centered at p to a Φ-disk of radius r_p. r_p does not depend on the initial choice of natural coordinate. If $d(s,t) = \delta$, then

$$r_s - \delta \leq r_t \leq r_s + \delta.$$

Let $r_0 = \min\{r_q, r_{h(q)}\}$. Select a point $q' \in B_{r_0/4}(q) \cap \Omega_1$. This point satisfies $r_{q'} \geq 3r_0/4$ and likewise for $h(q')$. Let $\delta = d(q,q')$ and take a natural coordinate z in a Φ-disk $B_{\delta/2}(q')$. We restrict h to this Φ-disk, and as above, we use h to build a natural coordinate w on $B_{\delta/2}(h(q'))$. More precisely, we have a disk $D \subset \mathbb{C}$ of radius $\delta/2$ and two holomorphic maps

$$\varphi : D \to B_{\delta/2}(q')\ ,\ \psi : D \to B_{\delta/2}(h(q'))$$

such that $z = \varphi^{-1}$, $w = \psi^{-1}$. We can extend these maps to a larger disk $D' \subset \mathbb{C}$ with radius $3r_0/4$. The map

$$w^{-1} \circ z : B_{3r_0/4}(q') \to B_{3r_0/4}(h(q'))$$

is a holomorphic diffeomorphism that agrees with h on $B_{\delta/2}(q')$. Since $B_{r_0/2}(q) \subset B_{3r_0/4}(q')$, we see we have analytically continued h to the open set $\Omega_1 \cup B_{r_0/2}(q)$. From conformal invariance, the map $f \circ h$ is harmonic, and hence the Aronszajn theorem [Aro57] implies $f \circ h = f$ on $\Omega_1 \cup B_{r_0/2}(q)$. □

Via this result, we often find ourselves in the following situation: either h can be continued along an entire curve γ, or we have a segment $\gamma' \subset \gamma$ along which h has been continued but the endpoint of $h(\gamma')$ is a zero of Φ.

We remark that there is no guarantee that the analytic continuation is a diffeomorphism. It is at least a local diffeomorphism and a local isometry for the Φ-metric.

Harmonic singularities

Toward the proof of the main theorem, we rule out possible pathological behaviour of harmonic maps near rank 1 singularities. We need not delve too deep into the theory of singularities, but we invite the reader to see Wood's thesis [Woo74] and the paper [Woo77], in which he studies singularities of harmonic maps between surfaces in detail.

Our key tool is the Hartman-Wintner theorem [HW53], which gives a local representation formula for harmonic maps. Let z be a holomorphic coordinate centered on a disk centered at $p \in \Sigma$ with $z(p) = 0$, and let $(x_1, \dots, x_n)$ be normal (but not necessarily orthogonal) coordinates in a neighbourhood U of $f(p)$ such that $f(p) = 0$. According to the Hartman-Wintner theorem, we can write the components $(f^1, \dots, f^n)$ as

$$f^k = p^k + r^k,$$

where p^k is a spherical harmonic (a harmonic homogeneous polynomial) of some degree $m < \infty$ and $r^k \in o(|z|^m)$. We are allowing $p^k = \infty$, which means $f^k = 0$.

By permuting the coordinates, we may assume $\deg p^1 = \min_k \deg p^k$, and $\deg p^k \geq \deg p^2$ for all $k \geq 3$. Note $\deg p^1, \deg p^2 < \infty$, for otherwise Sampson's result [Sam78, Theorem 3] implies f takes its image in a geodesic.

Lemma 5.2.4. *There does not exist a sequence of points $(p_n)_{n=1}^{\infty} \subset \Sigma$ converging to p with the property that there exists a (not necessarily conformal) diffeomorphism h_n taking a neighbourhood of p_n to a neighbourhood of p that leaves f invariant.*

Proof. Arguing by contradiction, suppose there is such a sequence $(p_n)_{n=1}^{\infty}$. Since f is an embedding near regular points, p must be a singular point. Choose a coordinate z on the source and normal coordinates on the target with $p = 0$, $f(p) = 0$. We apply Hartman-Wintner to obtain the formula

$$f^k = p^k + r^k$$

with the same degree assumptions as above. It is clear that there is at least one p^k with $\deg p^k = m > 1, m \neq \infty$.

We invoke a result of Cheng [Che76, Lemma 2.4]: there is a C^1 diffeomorphism from a neighbourhood of 0 in $\mathbb{R}^2$ to a neighbourhood of p, taking 0 to 0 in coordinates, and such that

$$f^k \circ \varphi(w) = p^k(w).$$

As a spherical harmonic of degree m, the zero set of p^k consists of m distinct lines going through the origin, arranged so that the angles between two adjacent lines is constant (this is an easy consequence of homogeneity). Notice that in our neighbourhood of p,

$$\{q : f^k(q) = f^k(p)\} = \{\varphi(w) : p^k(w) = p^k(0)\}.$$

Therefore, the set $\{q : f^k(q) = f^k(p)\}$ is collection of m disjoint C^1 arcs all transversely intersecting at the origin. For n large enough, p_n lies inside the coordinate chart determined by φ, and hence it lies on one of the arcs. Fixing such a p_n, we use that h_n is a diffeomorphism to see that there should be $m-1$ more curves transversely intersecting the line containing p_n, and such that $f(q) = f(p)$ on those curves. This is a clear contradiction. □

5.3 Holomorphic factorization

Throughout this section, we continue to assume $h : \Omega_1 \to \Omega_2$ is a holomorphic diffeomorphism. Following the structure of Section 3 in [GOR73], we prove Theorem 5A holds for such h (although the technical details of our proofs are for the most part quite different).

The equivalence relation

Definition 5.3.1. Given $p_1, p_2 \in \Sigma$, we define a relation $\sim$ by

1. If $p_1, p_2 \notin \mathcal{Z}$, $p_1 \sim p_2$ if there exists open sets Ω_1, Ω_2 such that $p_i \in \Omega_i$ and a holomorphic diffeomorphism $h : \Omega_1 \to \Omega_2$ such that $f = f \circ h$ on Ω_1.
2. If one of p_1, p_2 is a zero of Φ, then for any pair of neighbourhoods Ω_i containing p_i one can find smaller neighbourhoods $\Omega_i' \subset \Omega_i$ containing p_i such that for each $q_1 \in \Omega_1' \backslash \{p_1\}$ there exists $q_2 \in \Omega_2' \backslash \{p_2\}$ such that $q_1 \sim q_2$, and for each $q_2 \in \Omega_2' \backslash \{p_2\}$ there is a $q_1 \in \Omega_1' \backslash \{p_1\}$ such that $q_2 \sim q_1$.

If $p_1 \sim p_2$ then $f(\Omega_1') = f(\Omega_2')$ and $f(p_1) = f(p_2)$ are apparent from the definition. Recall $\mathcal{Z} = \{p \in \Sigma : \Phi(p) \neq 0\}$.

Proposition 5.3.2. *$\sim$ is an equivalence relation.*

Proof. Reflexivity and symmetry are obvious. As for transitivity, this is clear if p_1, p_2, p_3 are all not zeros of Φ. If at least one is a zero, we consider two cases:

1. p_1, p_3 are zeros, or
2. p_2 is a zero while p_1, p_3 are not

The other cases are trivial. Case (i) can be seen from the definitions: take Ω_1, Ω_2 containing p_1, p_2 respectively such that for all $p_1' \in \Omega_1 \backslash \{p_1\}$ there exists $p_2' \in \Omega_2 \backslash \{p_2\}$ with $p_1' \sim p_2'$. Within Ω_2 we find an open set Ω_2', and then an open set Ω_3' containing p_3 with the same property. Set

$$\Omega_1' = \{p_1' \in \Omega_1 \backslash \{p_1\} : \text{ there exists } p_3' \in \Omega_3' \text{ such that } p_1' \sim p_3'\} \cup \{p_1\}.$$

We can find an open disk centered at p_1 inside Ω_1' by applying the definition of $\sim$ to the open sets Ω_1, Ω_2'. It is also clear that Ω_1' is open away from p_1, and hence it is open. It is now simple to check that Ω_1' and Ω_3' satisfy the definition of $\sim$.

The second case requires more work. Select Φ-disks U_1, U_3 of radius $R > 0$ around p_1 and p_3 respectively such that there are no points q_i with $q_i \sim p_i$ and no zeros of Φ. Let U_1', U_3' be Φ-disks centered at the same points with half the radius. Using $\sim$, we can find open sets $\Omega_i \subset U_i'$ containing p_i such that $f(\Omega_3) \subset f(\Omega_1)$ and every point in $q \in \Omega_3 \backslash \{p_3\}$ is equivalent to a point in $\Omega_1 \backslash \{p_1\}$. We shrink Ω_3 to turn it into an open disk in the Φ-metric centered at p_3 with radius $\delta < R/2$.

Let $p_i' \in \Omega_i$ be such that $p_3' \sim p_1'$. Viewing Ω_3 in natural coordinates, let γ be the straight line from p_3' to p_3. We have a holomorphic map h taking a neighbourhood of p_3' to one of p_1' that leaves f invariant. We analytically continue along γ as much as we can. Either $h(\gamma)$ hits a zero of Φ or we can continue up until the endpoint. The Φ-length of any segment of $h(\gamma)$ is at most δ, and we infer $h(\gamma)$ is contained in $B_{R/2+\delta}(p_2) \subset U_3$. Thus, $h(\gamma(t))$ can never be a zero for any time t, and we can continue to the endpoint. From the proof of Proposition 5.2.3, $p_3 = \gamma(1)$ is equivalent to the endpoint $h(\gamma(1))$.

To finish the proof, we need to argue $h(\gamma(1)) = p_1$. Let $q_1 = h(\gamma(1))$. We do know $p_3 \sim q_1$. We claim we could have chosen R small enough to ensure no point other than possibly p_1 is equivalent to p_3. Indeed, if this is not possible, then we get a sequence of points $(q_n)_{n=1}^{\infty}$ converging to p_1 such that $p_3 \sim q_n$ for all n. Using transitivity of $\sim$ for points in $\Sigma \backslash \mathcal{Z}$, we can then construct a sequence of points q_n' converging to p_3 that are all equivalent to p_3. This directly contradicts Lemma 5.2.4 and completes the proof. □

We use Proposition 5.3.2 to prove another useful property of $\sim$.

Lemma 5.3.3. *Suppose $p_1, p_2 \notin \mathcal{Z}$. Then there is no sequence $(q_n)_{n=1}^{\infty}$ such that $q_n \sim p_1$ for all n and $q_n \to p_2$ as $n \to \infty$.*

Proof. Again going by way of contradiction, assume such a sequence q_n exists. Firstly, by Lemma 5.2.4, we cannot have $p_1 \sim p_2$. Using the definition of $\sim$, we see that in any pair of neighbourhoods Ω_i of p_i, we can find points $p_i' \in \Omega_i$ such that $p_1' \sim p_2'$.

Let $\delta, \epsilon > 0$ and $\tau = \epsilon + 2\delta$. We choose δ, ϵ to be small enough to ensure

1. there is no point equivalent to p_1 in $B_\tau(p_1)\backslash\{p_1\}$,
2. there is no point equivalent to p_2 in $B_\delta(p_2)\backslash\{p_2\}$, and
3. there are no zeros of the Hopf differential in either ball.

Choose $p_1' \in B_\epsilon(p_1)$ that is equivalent to a point $p_2' \in B_\delta(p_2)$. In natural coordinates, let γ be the straight line path from p_2' to p_2. γ has length at most δ, and hence the image of any segment of γ along an analytic continuation of h lies in $B_\tau(p_1)$. Thus, we can continue h along γ as much as we like, and we extend to the boundary point p_2. The endpoint $h(\gamma(1))$ is then equivalent to p_2. Since $p_1 \nsim p_2$, $h(p_2) \neq p_1$.

Set $q_1' = h(p_2)$. Replace δ, ϵ, τ with smaller numbers $\delta', \epsilon', \tau'$ satisfying the same relations as above and $q_1 \notin B_{\tau'}(p_1)$. By repeating the previous procedure we secure another point $q_2' \sim p_2$ that is closer to p_1. Continuing in this way, we can build a sequence $(q_n')_{n=1}^{\infty}$ converging to p_1 such that $q_n' \sim p_2$ for all n.

We now find our contradiction. Given that both such sequences exist, f cannot be an embedding around p_1 nor p_2 and has rank 1 at both points. Choose normal coordinates on M centered at $f(p_1) = f(p_2)$, and a conformal coordinate centered at p_1 in which f takes the form

$$f^k = p^k + r^k$$

as in the previous subsection. Since f is not regular at p_1, there is at least one k such that $\deg p^k = m > 1$, $m \neq \infty$. Choosing a conformal coordinate at p_2, f takes the form

$$f^k = \tilde{p}^k + \tilde{r}^k$$

with $\tilde{p}^k$ a spherical harmonic and $\tilde{r}^k$ decaying faster. The images of p^k and $\tilde{p}^k$ in $\mathbb{R}$ intersect on open sets, so $\tilde{p}^k$ is clearly non-zero. Thus, the set of points near p_2 on which f^k is equal to $f^k(p_1)$ is some collection of arcs intersecting at that point. However, since deg $p^k > 1$, we can find the same contradiction as in Lemma 5.2.4. □

The Hausdorff condition

The main result of this subsection is Proposition 5.3.4, which implies the topological quotient of Σ by $\sim$ is Hausdorff. We say $p_1 \sim' p_2$ if for every pair of neighbourhoods U_i containing p_i, there exists $p_i' \in U_i$ with $p_1' \sim p_2'$.

Proposition 5.3.4. *Suppose $p_1 \sim' p_2$. Then $p_1 \sim p_2$.*

Proposition 5.3.4 is our "holomorphic unique continuation property." Combined with [Sam78, Theorem 3], Proposition 5.3.4 implies the following result of independent interest.

Proposition 5.3.5. *Let $\mathbb{D} \subset \mathbb{R}^2$ be the unit disk. Suppose $u_1, u_2 : \mathbb{D} \to M$ are harmonic maps maps such that, for all open sets $D_1 \subset \mathbb{D}$ containing 0, there is an open subset $D_2 \subset \mathbb{D}$ (possibly not containing 0) such that $u^2(D_2) \subset u^1(D_1)$. Moreover, assume that for any subsets $D_i' \subset D_i$ on which u_i is regular such that $u_2(D_2') \subset u_1(D_1')$, the map $u_2^{-1}|_{u_1(D_1')} \circ u_1|_{D_1'}$ is holomorphic. Then there exists an open subset $D' \subset \mathbb{D}$ containing 0 such that $u^2(D') \subset u^1(\mathbb{D})$.*

Turning toward the proof of Proposition 5.3.4, if p_1 and p_2 are both not zeros of Φ, then one can follow the argument from the proof of Proposition 5.3.2, almost word-for-word, up until the last paragraph. We just need to note that Lemma 5.3.3 shows we can choose a Φ-disk surrounding p_1 that is small enough that it contains no point equivalent to p_2. .

Going forward, we assume at least one of the two points is a zero of Φ. The main step in the proof is the next lemma.

Lemma 5.3.6. *There exists $\delta, \tau > 0$ such that every $p_1' \in B_\delta(p_1)\backslash\{p_1\}$ is equivalent to a point $p_2' \in B_\tau(p_2)\backslash\{p_2\}$.*

Proof. Let $\delta, \epsilon > 0$ and $\tau = \epsilon + 3\delta$. We choose δ, ϵ to be small enough such that $B_\delta(p_1) \cap B_\tau(p_2) = \emptyset$ and that in $B_\delta(p_1)\backslash\{p_1\}$ and $B_\tau(p_2)\backslash\{p_2\}$,

1. we have no points equivalent to the centers, and

2. there are no zeros of Φ.

We take open sets $p_1' \in B_\delta(p_1)$, $p_2' \in B_\epsilon(p_2)$ with $p_1' \sim p_2'$, and let h be the associated holomorphic diffeomorphism. Let $q \in B_\delta(p_1)$, $q \neq p_1$, and let γ be a path from a point p_1' to q. We choose γ to be either the straight line from p_1' to q, or a slight perturbation of that line to make sure the path does not touch p_1. Regardless, we can arrange so the Φ-length is bounded above by $5\delta/2$.

We analytically continue h along γ as much as we can. Since the starting point lies in $B_\epsilon(p_2)$, we see the image under h of any segment lies in $B_\tau(p_2)$. If we can continue h along γ to the endpoint, and the endpoint of $h(\gamma)$ is not p_2, then we have $q = \gamma(1) \sim h(\gamma(1))$. The only way we could not extend is if some segment of $h(\gamma)$ touches p_2. Notice that, regardless, we have a point $q \in B_\delta(p_1)$ that satisfies $q \sim' p_2$ (here we are relabelling q to be the endpoint of a bad segment if that happens). We rule this out with the lemma below.

Lemma 5.3.7. *In the setting above, we can choose our Φ-disks to be small enough so that no point $q \in B_\delta(p_1)\backslash\{p_1\}$ satisfies $q \sim' p_2$.*

Proof. We first show that given such a point q, we have $q \sim' p_1$. Let U_1, U_2, U_3 be open sets containing p_1, p_2, q respectively. Let $\delta_1, \delta_2, \delta_3 > 0$ and find $p_1' \in B_{\delta_1}(p_1)$, $p_2' \in B_{\delta_2}(p_2)$ with $p_1' \sim p_2'$, as well as $p_2'' \in B_{\delta_2}(p_2)$, $q' \in B_{\delta_3}(q)$ with $p_2'' \sim q'$. We choose the δ_j's so that $B_{\delta_3+3\delta_2}(q)$ contains no zeros of the Hopf differential, and $B_{\delta_i}(p_i)$ can only have zeros at p_i. We also choose the δ_i's so that all balls mentioned above are contained in U_1, U_2, U_3 and disjoint. Let h be the holomorphic map relating p_2'' to q'. We analytically continue h along a path γ from p_2'' to p_2' with length at most $5\delta_2/2$ that is chosen to avoid p_2. Then the image path lies in $B_{\delta_3+3\delta_2}(q)$ and so we can continue to the endpoint. The endpoint $h(\gamma(1))$ is equivalent to p_2'. If the endpoint is not q, then $h(\gamma(1)) \sim p_2' \sim p_1'$, and this proves the claim. If the endpoint $h(\gamma(1))$ is q itself, then $q \sim p_2' \sim p_1'$, and we can find q'' very close to q that is equivalent to a point very close to p_1' (in particular, contained in $B_{\delta_1}(p_1)$).

Therefore, we see that if the lemma is false, we can construct a sequence $(q_n)_{n=1}^\infty$ converging to p_1 such that $q_n \sim' p_1$ for all n. Fix a q_n, along with a $\delta' > 0$ such that $B_{4\delta'}(q_n)$ contains no zeros and no points equivalent to q_n and $B_{\delta'}(p_1)$ has no zeros other than possibly p_1. We find $q_n' \in B_{\delta'}(q_n)$ and $p_1' \in B_{\delta'}(p_1)$ such that

$q_n' \sim p_1'$. There is another point $q_N \in B_{\delta'}(p_1)$ such that $q_N \sim' p_1$. Connect p_1' to q_N via a path of length at most $5\delta'/2$ that does not touch p_1. Analytically continue the associated map h along this path. The image lies in $B_{4\delta'}(q_n)$, so we can always continue. The endpoint $h(\gamma(1)) \in B_{4\delta'}(q_n)$ is equivalent to q_N. We claim we can choose q_N with the property that $h(\gamma(1)) \neq q_n$. To this end, if $h(\gamma(1)) = q_N$, we take the straight line path σ from q_N to q_{N+1}. According to [Str84, Theorem 8.1], if p_1 is a zero of Φ of order n, then geodesics in the Φ metric are either straight lines or the concatenation of two radial lines enclosing an angle of at least $2\pi/(n+2)$. By pigeonholing, we can pass to a subsequence where every q_n lies in a closed sector of angle $\pi/(n+2)$ around the origin. This guarantees that the straight line path from any q_j to q_k is a geodesic in the Φ-metric and has length at most δ'. As $\Phi(q_n) \neq 0$, the image $h(\sigma)$ is then a straight line contained in $B_{4\delta'}(q_n)$ with initial point q_n, so it certainly cannot terminate at q_n. We prove the claim by replacing q_N with q_{N+1} and taking the concatenation of our original path with the straight line σ. We now just want to show $q_N \sim q_n$, and we will have a contradiction. Toward this, it is enough to show $q_N \sim' q_n$, since Φ does not vanish at these points.

This last step is similar to the beginning of our proof, and so we only sketch the argument. Recall that we have $p_1 \sim' q_n$ and $p_1 \sim' q_N$. Find smalls balls containing q_n, p_1, and q_N. Then within the ball containing p_1 we have two points p_1' and p_1'', with p_1' equivalent to a point near q_n and p_1'' equivalent to a point near q_N. Connect p_1' and p_1'' via a small arc that does not touch p_1. We can arrange for the arc to stay in a ball around q_n in which it can always be continued. We thus get a point near q_n that is equivalent to a point near q_N. We may need to wiggle the path so the point is not q_n. As discussed above, we are done. □

Returning to the proof of Lemma 5.3.6, we see that we can always extend our chosen segments, and moreoever each $q \in B_\delta(p_1)\backslash\{p_1\}$ has an equivalent point in $B_\tau(p_2)\backslash\{p_2\}$. □

With Lemma 5.3.6 in hand, we are now ready to complete the proof of Proposition 5.3.4. Let Ω_2' be the set of points in $B_\tau(p_2)\backslash\{p_2\}$ that have an equivalent point in $B_\delta(p_1)\backslash\{p_1\}$. Let $\Omega_2 = \Omega_2' \cup \{p_2\}$. By repeating the previous argument, we can find a very small ball $B_\alpha(p_2)$ such that every point in $B_\alpha(p_2)\backslash\{p_2\}$ is equivalent to a point in $B_\delta(p_1)\backslash\{p_1\}$. This shows that p_2 is an interior point of Ω_2. Away from p_2, Ω_2 is open by elementary considerations. It is now simple to conclude $p_1 \sim p_2$ by using the open sets $B_\delta(p_1)$ and Ω_2.

Ramification at branch points

We now investigate the local behaviour of the map f near zeros of the Hopf differential. This leads us to define a notion of ramification for branch points. Our definition is slightly different from the one given in Section 5.1.

Lemma 5.3.8. *Suppose p is a branch point of f, and hence a zero of Φ of some order n. Let $h : \Omega_1 \to \Omega_2$ be a holomorphic diffeomorphism with $f \circ h = f$, and suppose Ω_1, Ω_2 are both contained in a ball $B_\epsilon(p)$, where $\epsilon > 0$ is chosen so that there are no other zeros and no other point is equivalent to p in $B_{2\epsilon}(p)$. Then, in the natural coordinates for Φ, h is a rational rotation of angle $2\pi j/(n+2)$*

Proof. Select $p_i \in \Omega_i$ with $h(p_1) = h(p_2)$. Let $\gamma : [0,1] \to B_\epsilon$ be a straight line path starting at p_1 that terminates at the point p. We analytically continue h in a simply connected neighbourhood of γ, as far as we can. Either there is an interior point q in the straight line that is mapped via h to p, or we can continue along the whole curve and extend to the boundary point p. In the first case, Proposition 5.3.4 guarantees $q \sim h(q) = p$, which by our choice of ϵ means $q = p$, contradicting the definition of q. In the second case, Proposition 5.3.4 yields $p \sim h(p)$ and we deduce $h(p) = p$.

We now prove h is a rotation. Work in the interior of the extension of Ω_1 in which h has been continued. If we write the Hopf differential in local coordinates as $\Phi = \phi(z)dz^2$, then

$$\phi(z) = \phi(h(z))(h'(z))^2.$$

In the natural coordinate for the Hopf differential this becomes

$$z^n = (h(z))^n (h'(z))^2.$$

Since we're in a simply connected region that doesn't touch zero we can choose a branch of the square root. h then satisfies

$$z^{n/2} = (h(z))^{n/2} h'(z) = \frac{\partial}{\partial z} \frac{(h(z))^{n/2+1}}{n/2+1}.$$

Integrate to get

$$z^{n/2+1} = (h(z))^{n/2+1} + c$$

for some complex constant c. Since $h(p) = p$, taking $z \to 0$ along γ forces $c = 0$. This implies

$$z^{n+2} = (h(z))^{n+2}$$

and the result is now clear. □

Definition 5.3.9. A non-minimal harmonic map g with Hopf differential Φ is holomorphically ramified of order $r-1$ if r is the largest integer such that there exists a Φ-disk Ω centered at p and a holomorphic degree r branched cover $\psi : \Omega \to D$ with one branch point at p onto a disk D with $\psi(p) = 0$ and such that $\psi(p_1) = \psi(p_2)$ implies $f(p_1) = f(p_2)$.

A map is called unramified if $r = 1$. Clearly, a map can only ramify non-trivially at a branch point.

Lemma 5.3.10. *A non-minimal harmonic map g with Hopf differential Φ is ramified of order $r > 1$ at p if and only if for all $\epsilon > 0$, there exists $p_1, p_2 \in B_\epsilon(p)\backslash\{p\}$ such that $p_1 \sim p_2$ and $p_1 \neq p_2$, where $p_1 \sim p_2$ in the sense that there is a holomorphic map h taking a neighbourhood of p_1 to one of p_2 that leaves g invariant.*

Remark 5.3.11. A similar statement holds for minimal maps. See [GOR73, Lemma 3.12].

Proof. If g is ramified we take a Φ-disk Ω of p and a map $\psi : \Omega \to D$ as in the definition. Select two points $p_i \neq p$ such that $\psi(p_1) = \psi(p_2)$ as well as neighbourhoods Ω_i on which ψ is injective and share the same image under ψ. Setting $\psi_i = \psi|_{\Omega_i}$, the map $\psi_2^{-1} \circ \psi_1 : \Omega_1 \to \Omega_2$ is a holomorphic diffeomorphism that leaves g invariant and hence $p_1 \sim p_2$. Conversely, pick $\epsilon > 0$ such that there are no other zeros of Φ in $B_{2\epsilon}(p)$ and so we have a coordinate z such that $\Phi = z^n dz^2$. There exists $p_1, p_2 \in B_\epsilon(p)$ with $p_1 \sim p_2$ but $p_1 \neq p_2$. Lemma 5.3.8 shows there are small disks surrounding p_1, p_2 that are related by a rotation h of the form

$$z \mapsto e^{\frac{2\pi i j}{n+2}} z$$

such that $g = g \circ h$. By the Aronszajn theorem, g is invariant under this rotation in all of V. Dividing by the gcd, we see g is invariant under a rotation of the form

$$z \mapsto e^{\frac{2\pi i j_1}{r}} z,$$

where j_1 and r are coprime. It follows that $g \circ \alpha = g$ in $B_\epsilon(p)$, where α is the rotation $z \mapsto e^{2\pi i/r} z$. In these coordinates, we define a holomorphic branched cover $\psi : B_\epsilon(p) \to D$ by $\psi(z) = z^r$, and note that $\psi(p_1) = \psi(p_2)$ implies $g(p_1) = g(p_2)$. □

Lemma 5.3.12. *Let p be a branch point of f of order $m-1$ at which f is ramified of order $r-1$. Then there is a Φ-disk Ω of p such that f admits a factorization $f|_\Omega = \overline{f} \circ \psi$, where*

1. *$\psi : \Omega \to D$ is a holomorphic map onto a disk $\{|\zeta| < \delta\}$ such that $\psi|_{\Omega\setminus\{p\}}$ is an r-sheeted covering map,*

2. *$\overline{f}$ is harmonic with respect to the flat metric on D and the given metric on M, and*

3. *$\overline{f} : D \to M$ is unramified with a single branch point of order $s - 1$ at the origin, where $s = m/r$.*

Proof. Define $\overline{f}$ by $\overline{f}(\psi(z)) = f(z)$. (i) is given and we begin with (ii). Harmonicity is a local matter, and at any point away from zero we can choose a neighbourhood surrounding that point where ψ^{-1} exists and we have the factorization $\overline{f} = f \circ \psi^{-1}$ in that neighbourhood. Since ψ^{-1} is conformal, $\overline{f}$ is harmonic off 0. Near 0, we compute $\overline{f}_\zeta$ in coordinates to realise C^1 bounds. Via Schauder theory we promote to C^2 (or even C^∞) bounds. This implies that the tension field is continuous and therefore vanishes everywhere. As for (iii), we can write each component f^k in certain coordinates as

$$f^k = p^k + r^k,$$

where p^k is a spherical harmonic and r^k decays faster than p^k. In this form, it is easy to check the branching orders of f and $\overline{f}$.

It remains to show that $\overline{f}$ is unramified. Toward this, let Θ be the Hopf differential of $\overline{f}$ and note the image of a Φ-disk under ψ is a Θ-disk. Indeed, if $\Phi = \phi(z)dz^2$, $\Theta(\zeta) = \theta(\zeta)d\zeta^2$ in local coordinates, then

$$\phi(z) = \theta(z^r)\left(\frac{\partial z^r}{\partial z}\right)^2 = \theta(z^r)z^{2r-2}r^2.$$

We rearrange to see

$$\theta(\zeta) = \theta(z^r) = z^{n-2r+2}r^{-2},$$

and the fact that the image is a Θ-disk is derived from direct computation. If $\overline{f}$ is ramified, we can build another holomorphic branched covering map ψ' as in Lemma 5.3.10. Since both ψ and ψ' have finite fibers, the composition $\psi' \circ \psi$ yields a branched cover of degree greater than r, which is impossible. This finishes the proof. □

Remark 5.3.13. Our computations show that the ramification order is constrained by $r|m$, $r|(n+2)$, and $2r \leq n+2$. The last condition is superfluous, since we always have $2m \leq n+2$.

Lemma 5.3.14. *For $i = 1, 2$, let p_i be branch points of f of order $m_i - 1$ (we are allowing $m_i = 1$), ramified of order $r_i - 1$. Then $p_1 \sim p_2$ if and only if*

1. *$f(p_1) = f(p_2)$,*
2. *$m_1/r_1 = m_2/r_2$, and*
3. *if s is the common value m_i/r_i, there exist maps $\psi_i : U_i \to D$, $\overline{f}_i : D \to M$, $\psi_i(p_i) = 0$, such that $\psi_i|_{U_i \setminus \{p\}}$ is an r_i-sheeted holomorphic covering map, f factors as $f|_{U_i} = \overline{f} \circ \psi_i$, and $\overline{f}$ is a harmonic map for the flat metric on the disk with a branch point of order $s - 1$.*

Proof. If $m = 0$ this is trivial, so we assume $m > 0$. Suppose the conditions hold. Given any two open sets Ω_i containing p_i, we can radially shrink our Φ-disks to have $U_i \subset \Omega_i$ (the argument from Lemma 5.3.8 shows any two points with $\psi_i(q_1) = \psi_i(q_2)$ have the same Φ-distance to p_i). For $p_1' \in U_1 \setminus \{p_1\}$ let ψ_1' be the restriction to a neighbourhood U_1' of p_1' on which ψ_1 is injective. Let ψ_2' be the restriction onto some neighbourhood V_2' such that ψ_2 maps U_2' injectively onto $\psi_1(U_1')$. Set $p_2' = \psi_2'^{-1} \circ \psi_1'(p_1')$ and $h = \psi_2'^{-1}\psi_1$. h is holomorphic and leaves f invariant. The result follows.

Conversely, assume $p_1 \sim p_2$. (i) was already discussed. We first want to show that we can choose Φ-disks U_i that satisfy condition (2) in the definition of $\sim$. We take $\psi_i : U_i \to D_i$ and $f_i : D_i \to M$ as in Lemma 5.3.12. If $p_1' \in U_1$ is equivalent to $p_2' \in U_2$, then combining our reasoning from Proposition 5.3.2 with Proposition 5.3.4 shows $d(p_1, p_1') = d(p_2, p_2')$. We've run this type of argument a few times at this point, but we feel a duty to elaborate. Pick subdisks $U_i' \subset U_i$ that satisfy condition (2) and balls $B_\delta(p_1)$, $B_\epsilon(p_2)$ contained in the subdisks, such that in $B_{2\delta}(p_1)$ and $B_{\epsilon+2\delta}(p_2)$ there are no points equivalent to p_1, p_2 respectively and no other possible zeros of Φ. Find $p_1' \in B_\delta(p_1)\setminus\{p_1\}$ and $p_2' \in B_\epsilon(p_2)\setminus\{p_2\}$ with $p_1' \sim p_2'$. Take the straight line path γ from p_1' to p_1 and analytically continue h along γ as much as we can. The image of any segment of this path under h is also a straight line contained in $B_{\epsilon+2\delta}(p_2)$. We now have two possibilities:

1. the path $h(\gamma)$ runs into p_2 before we have finished extending, or
2. we can extend h to the boundary point $\gamma(1) = p_1$.

In the first scenario, we obtain $d(p_2, p_2') \leq d(p_1, p_1')$. In the latter, Proposition 5.3.4 ensures $h(\gamma(1)) \sim p_1 \sim p_2$, so that $h(\gamma(1)) = p_2$. Regardless of the situation, we have

$$d(p_2', p_2) \leq d(p_1', p_1).$$

To reverse the argument for the other inequality, we go via a straight line from p_2' to p_2. For any segment γ' along which we can continue, the length of $h(\gamma')$ is now bounded above by $d(p_1', p_1) < \delta$. Thus, we can continue along the whole curve so long as we don't hit p_1. In the same way as above we get the opposite inequality. This is the desired result.

Using the definition of $\sim$, we can now assume the Φ-disks U_i are such that $f(U_1) = f(U_2)$ and that for all $p_1' \in V_1, p_1' \neq p_1$, there is $p_2' \in V_2, p_2' \neq p_2$, such that $p_1' \sim p_2'$, and vice versa. We construct a holomorphic diffeomorphism $G : D_1 \to D_2$ such that

$$\overline{f}_2 \circ G = \overline{f}_1.$$

Let $w_1 \in D_1 \backslash \{0\}$. We take a small neighbourhood of w_1 and a lift to an open set via ψ_1 such that the restriction of ψ_1 is injective. Let w_1' be the given preimage under ψ_1. There is then a point $w_2' \in V_2$ related by a holomorphic map such that f agrees in neighbourhoods surrounding w_1' and w_2'. Set $w_2 = G(w_1) = \psi_2(w_2')$. We claim there can be no other point with this property. If there was such a w', then we would have $w \sim w'$ with respect to the corresponding equivalence relation for $\overline{f}_2$. However, we know the map $\overline{f}_2$ is unramified, and by Lemma 5.3.10 we can choose our disks small enough that there are no two distinct points in D_2 with this property. The association $w_1 \mapsto w_1'$ defines our map G. If we set $G(0) = 0$, then we see G is a diffeomorphism from $D_1 \backslash \{0\} \to D_2 \backslash \{0\}$, because we can invert the construction. The map G is holomorphic off $\{0\}$. Since it is bounded near 0, it extends to a holomorphic diffeomorphism on all of D_1.

From Lemma 5.3.12 the branching order of $\overline{f}_i$ is $m_i/r_i - 1$, and since G is a diffeomorphism, it is clear that these branching orders agree. Defining $D = D_1$ and $\overline{f}$ to be the common map $\overline{f}^2 \circ G = \overline{f}_1, \psi_1 = \overline{\psi}_1, \psi_2 = G^{-1} \circ \overline{\psi}_2$, (iii) can be verified easily. □

Constructing the Riemann surface

Preparations aside, we build the covering space. Our work here is drawn from Propositions 3.19 and 3.24 in [GOR73]. Let Σ_0 denote the space of equivalence classes of Σ with respect to $\sim$, equipped with the quotient topology. We denote by $\pi : \Sigma \to \Sigma_0$ the projection map.

Proposition 5.3.15. Σ_0 *is an orientable surface.*

Proof. For each $p \in M$ let U be a neighbourhood of p with no other point equivalent to p and as in Lemma 5.3.12, so that we have a map $\psi : U \to D$, a harmonic map $f : D \to M$, and a factorization $f = \overline{f} \circ \psi$. Let $\overline{U} = \{[q] : q \in U\}$. To prove such a set is open, we show any $\pi^{-1}(\overline{U}) \subset \Sigma$ is open. If $p_1 \in \pi^{-1}(\overline{U})$, then there is $p_2 \in V$ such that $p_1 \sim p_2$. Then we can find neighbourhoods Ω_i containing p_i with $\Omega_2 \subset U$ and such that for each $p_1' \in \Omega_1$ there exists $p_2' \in \Omega_2$ with $p_1' \sim p_2'$. This implies the $\overline{U}$ define an open cover of Σ_0.

On each $\overline{U}$ we have a map $\overline{\psi} : \overline{U} \to D$ given by $\overline{\psi}([q]) = \psi(q)$. We will see that these maps define charts. If $q_1, q_2 \in U$ are such that $q_1 \sim q_2$, then $\psi(q_1)$, $\psi(q_2)$ are equivalent with respect to $\overline{f}$ and hence we can choose U so that $\psi(q_1) = \psi(q_2)$, since $\overline{f}$ is unramified. This proves $\overline{\psi}$ is well-defined.

For injectivity, suppose $[p_1], [p_2] \in \overline{U}$ are such that $\overline{\psi}([p_1]) = \overline{\psi}([p_2])$. Choosing representatives p_1, p_2, either $p_1 = p_2 = p$ or neither of them is equal to p. In the second case, since ψ is a holomorphic covering map on $U \backslash \{p\}$ we can use it to build a holomorphic diffeomorphism from a neighbourhood of p_1 to a neighbourhood of p_2. Since $f = \overline{f} \circ \psi$ on U, this map leaves f invariant.

As for continuity and openness, the argument is the same as the one found in [GOR73, page 779]. The Hausdorff condition is immediate from Proposition 5.3.4. Σ_0 is orientable because π respects the orientation of Σ. □

There exists a continuous map $f_0 : \Sigma_0 \to M$ such that $f = f_0 \circ \pi$, defined by $f_0([p]) = f(p)$.

Proposition 5.3.16. *There exists a complex structure on* Σ_0 *so that* $\pi : \Sigma \to \Sigma_0$ *is holomorphic and the map* f_0 *is harmonic with respect to the conformal metric* μ_0 *obtained via uniformization.*

Proof. We use the collection of charts specified in Lemma 5.3.14. Let $(\overline{U}_1, \overline{\psi}_1)$ and $(\overline{U}_2, \overline{\psi}_2)$ be two charts for Σ_0 arising from open sets U_1, U_2 centered at points p_1, p_2. We have maps $\psi_i : U_i \to D_i$, $\overline{\psi}_i : \overline{U}_i \to D_i$, $\pi : \Sigma \to \Sigma_0$, and harmonic maps $\overline{f}_i : D_i \to M$ such that $\overline{f} = \overline{f}_i \circ \overline{\psi}_i$, $\psi_i = \overline{\psi}_i \circ \pi$. We show the map

$$\overline{\psi}_2 \circ \overline{\psi}_1^{-1} : \overline{\psi}_1(\overline{U}_1 \cap \overline{U}_2) \subset D_1 \to \overline{\psi}_2(\overline{U}_1 \cap \overline{U}_2) \subset D_2$$

is holomorphic.

By the removeable singularities theorem, it suffices to check holomorphy away from the copies of 0 in D_i. Let $[q] \in \overline{U}_1 \cap \overline{U}_2$ be so that $\overline{\psi}_i([q]) \neq 0$, and choose a neighbourhood U around $[q]$ and $U' \subset \pi^{-1}(U)$ such that

1. $0 \notin \overline{\psi}_i(U)$,
2. the map $\pi|_{U'} : U' \to U$ is injective, and so we can define an inverse $\pi^{-1} : U \to U'$, and
3. the holomorphic map ψ_i is injective in U', so that we can define a holomorphic inverse $\psi_i^{-1} : \psi_i(U') \to U'$.

Note that $\psi_i(U') = \overline{\psi}_i(U)$. Clearly, the map $\psi_2 \circ \psi_1^{-1}$ is holomorphic in $\overline{\psi}_1(U)$. Meanwhile, since we can invert π, we obtain

$$\overline{\psi}_2 \circ \overline{\psi}_1^{-1} = (\psi_2 \circ \pi_0^{-1}) \circ (\psi_1 \circ \pi^{-1})^{-1} = \psi_2 \circ \psi_1^{-1}.$$

It follows that the map in question is holomorphic near $[q]$, and hence everywhere.

In holomorphic local coordinates, the map π is of the form $z \mapsto z$ or $z \mapsto z^n$, so it is surely holomorphic. From conformal invariance of the harmonic map equation, $f_0 = \overline{f}_i \circ \overline{\psi}_i$ is harmonic away from images of branch points of π. The argument of Lemma 5.3.12 shows f_0 is globally harmonic. □

This completes the proof of Theorem 5A for holomorphic diffeomorphisms.

5.4 Klein surfaces

We explain the adjustments required to prove Theorem 5A for anti-holomorphic diffeomorphisms $h : \Omega_1 \to \Omega_2$.

Preparations.

We begin with a review of Klein surfaces. More details on the theory of Klein surfaces can be found in the book [AG71]. Set

$$\mathbb{C}_+ = \{z \in \mathbb{C} : \mathrm{Im} z \geq 0\}$$

to be the closed upper half plane.

Definition 5.4.1. Let $\Omega \subset \mathbb{C}_+$ be open. A function $f : \Omega \to \mathbb{C}$ is (anti-)holomorphic if there is an open set $U \subset \mathbb{C}$ containing Ω such that f extends to an (anti-)holomorphic function from $U \to \mathbb{C}$.

Definition 5.4.2. A map between open subsets of $\mathbb{C}$ is dianalytic if its restriction to any component is holomorphic or anti-holomorphic.

Definition 5.4.3. Let X be a topological surface, possibly with boundary. A dianalytic atlas on X is a collection of pairs $\mathcal{U} = \{(U_\alpha, \varphi_\alpha)\}$ where

1. U_α is an open subset of X, V_α is an open subset of $\mathbb{C}_+$, and $\varphi_\alpha : U_\alpha \to V_\alpha$ is a homeomorphism.
2. If $U_\alpha \cap U_\beta \neq \emptyset$, the map

$$\varphi_\alpha \circ \varphi_\beta^{-1} : \varphi_\beta(U_\alpha \cap U_\beta) \to \varphi_\alpha(U_\alpha \cap U_\beta)$$

is dianalytic.

A Klein surface is a pair $X = (X, \mathcal{U})$.

Closely related is the notion of a Real Riemann surface.

Definition 5.4.4. A Real Riemann surface is the data (X, τ) of a Riemann surface X and an anti-holomorphic involution $\tau : X \to X$.

Given a Real Riemann surface (X, τ), the quotient X/τ has the structure of a Klein surface, and as a matter of fact every Klein surface X arises in this fashion (see Chapter 1 in [AG71]). The associated Real Riemann surface is called the analytic double, and it is unique up to isomorphism in the category of Real Riemann surfaces. The boundary of the Klein surface corresponds to the fixed-point set of the involution.

Definition 5.4.5. A harmonic (minimal) map on a Klein surface is a continuous map that lifts to a harmonic (minimal) map on the analytic double with respect to the conformal metric obtained via uniformization.

To prove Theorem 5A for anti-holomorphic maps, as previously done we define an equivalence relation $\sim$ and build a dianalytic atlas on the topological quotient $\Sigma_0 = \Sigma/\sim$. Before we get into details, we make an important reduction: we apply the holomorphic case of Theorem 5A to Σ and acquire a new Riemann surface Σ', as well as maps $\pi : \Sigma \to \Sigma'$, $f' : \Sigma' \to M$. The key property of the pair (Σ', f') is that equivalences classes under Definition 5.3.1 are singletons.

We define a relation $\sim$ on Σ by taking Definition 5.3.1, but this time insisting the maps involved are merely conformal rather than holomorphic.

Lemma 5.4.6. *Given $p \in \Sigma$, there is at most one other point $q \in \Sigma'$ such that $p \sim q$.*

Proof. Suppose p, q_1, q_2 are distinct points and $p \sim q_1$ and $p \sim q_2$. If all points are not in $\mathcal{Z}$, then we have anti-holomorphic maps h_1, h_2 relating to q_1, q_2 to p. The composition $h_2 \circ h_1^{-1}$ is then a holomorphic map relating q_1 to q_2, which means they are equivalent for Definition 4.1.8, and this is impossible. If at least one of them is a zero, then we can find disjoint neighbourhoods Ω containing p and Ω_i containing q_i such that every point in $\Omega_1 \backslash \{q_1\}$ is equivalent to a point in $\Omega \backslash \{p\}$, and every point in $\Omega \backslash \{p\}$ is equivalent to a point in $\Omega_2 \backslash \{q_2\}$. This brings us to the non-zero case. □

By the previous lemma, transitivity for $\sim$ holds vacuously. Accordingly, the proof of the lemma below is trivial.

Lemma 5.4.7. *$\sim$ is an equivalence relation.*

Proof of the main theorem.

Referencing our earlier work, we prove Theorem 5A for anti-holomorphic h. Henceforth we abuse notation and set $\Sigma = \Sigma'$, $f = f'$.

The first thing to note is that h is an orientation-reversing isometry for the Φ-metric. Indeed, if Φ does not vanish on an open subset $U \subset \Omega_1$ and z is a natural coordinate for Φ, then the function

$$w = \iota \circ z \circ h^{-1}$$

defines a holomorphic coordinate on $h(U)$, where ι is the complex conjugation operator on the disk. In this coordinate, $w(h(z)) = \overline{z}$, and it can be easily checked that

$$df_p\left(\frac{\partial}{\partial z}\right) = df_{h(p)}\left(\frac{\partial}{\partial \overline{w}}\right) \in T_{f(z)}M \otimes \mathbb{C}.$$

We infer

$$\langle f_{\overline{w}}, f_{\overline{w}} \rangle = 1$$

and furthermore

$$\langle f_w, f_w \rangle = \overline{\langle f_{\overline{w}}, f_{\overline{w}} \rangle} = 1.$$

As in Lemma 5.2.2, we find that w is a natural coordinate for Φ. The result follows.

Moreover, we can analytically continue h exactly as we did in Proposition 5.2.3. Moving toward the main proof, we follow the proof of Lemma 5.3.6, word-for-word, and note that Lemma 5.3.7 is immediate from Lemma 5.4.6. The proof of the analogue of Proposition 5.3.4 follows. As for ramification, we do see new behaviour.

Lemma 5.4.8. *Suppose p is a zero of Φ of order $n \geq 0$. Let $h : \Omega_1 \to \Omega_2$ be an anti-holomorphic diffeomorphism with $f \circ h = f$, and so that Ω_1, Ω_2 are both contained in a ball $B_\epsilon(p)$, where $\epsilon > 0$ is chosen so that there are no other zeros and no other point is equivalent to p in $B_{2\epsilon}(p)$. Then, in the natural coordinates for Φ,*

$$h(z) = e^{\frac{2\pi i j}{n+2}}\overline{z}$$

on its domain.

Proof. We follow the proof of Lemma 5.3.8, except now we have a map h that satisfies

$$z^n = (\overline{h}(z))^2\left(\frac{\partial \overline{h}}{\partial z}\right)^2.$$

As in the proof of Lemma 5.3.8, h is defined in a simply connected open set whose distance to zero can be taken to be arbitrarily small. We observe that $\overline{h}$ is holomorphic, and take a branch of the square root and integrate to derive

$$\overline{h}(z) = e^{-\frac{2\pi i j}{n+2}}z$$

for some $j = 0, 1, \ldots, n+1$. We conjugate to finish the proof. □

The lemma implies that in a neighbourhood of a ramification point p, f is invariant under the map

$$\psi(z) = e^{\frac{2\pi i j}{n+2}}\overline{z}.$$

This is an anti-holomorphic involution that fixes every point on the line

$$\mathcal{L} = \{re^{\frac{\pi i j}{n+2}} : -1 < r < 1\}$$

and acts by reflection across this line on all other points.

Lemma 5.4.9. *Let p and ψ be as above. If $\psi(q) = q$, then q has no equivalent points with respect to $\sim$.*

Proof. ψ is two-to-one in a neighbourhood of q. Suppose there exists $q' \in \Sigma$ with $q \sim q'$. Then $q' \notin B_\epsilon(p)$. Using the definition of $\sim$, we can find a small disk $B_{\epsilon'}(q)$ and points $p_1, p_2 \in B_{\epsilon'}(q)$ with $p_1 \sim p_2$, but we can also find a point q'' near q' such that $p_1 \sim q''$. This contradicts Lemma 5.4.6. □

We deduce the following.

Lemma 5.4.10. *Every $q \in B_\epsilon(p) \backslash \mathcal{L}$ is equivalent to $\psi(q)$ and only $\psi(q)$.*

We say f anti-holomorphically ramifies near p if f is invariant under an anti-holomorphic involution in a neighbourhood of p. In contrast to the holomorphic definition, f can anti-holomorphically ramify near rank 1 singularities. If f does ramify at p, we form the quotient

$$\mathcal{K} = B_\epsilon(p)/\psi$$

by identifying points z and $\psi(z)$. This has the structure of a Klein surface with boundary, the boundary being identified with $\mathcal{L}$.

Lemma 5.4.11. *$p \in \Sigma$ satisfies $[p] = \{p\}$ if and only if f ramifies at p.*

Proof. We need only to show that every point at which f is unramified admits an equivalent point. Looking toward a contradiction, suppose there exists $p \in \Sigma$ with $[p] = \{p\}$ and at which f does not ramify and choose $\epsilon > 0$ so that no two points are equivalent in $B_\epsilon(p)$ and that there are no zeros of Φ in $B_{2\epsilon}(p)$.

We claim $[q] = \{q\}$ for every $q \in B_\epsilon(p)$. If not, there is a $q \in B_\epsilon(p)$ that admits an equivalent point $q' \neq q$. Let h be the anti-holomorphic diffeomorphism relating a neighbourhood of q to one of q'. In coordinates, analytically continue h along a straight line γ from q to p. It follows from our assumption $\{p\} = [p]$ that no segment $h(\gamma')$ for $\gamma' \subset \gamma$ can touch p, for otherwise we get a point equivalent to p. Thus, we can continue to the endpoint, and the endpoint of $h(\gamma)$ is p itself. This implies

$$d(p, q) = d(p, q'),$$

which contradicts our choice of $\epsilon > 0$, and therefore settles the claim.

With the claim in hand, we define a map

$$\tau : \Sigma \to \Sigma$$

as follows. If $[q] = \{q\}$, set $\tau(q) = q$. If $[q] = \{q, q'\}$, we put $\tau(q) = q'$. If f is unramified at q and $[q] = \{p, q\}$, then τ is an anti-holomorphic diffeomorphism near p. If $[q] = \{q\}$, then our claim above shows it is the identity map in a neighbourhood of q. If f ramifies at q, τ acts like the map ψ considered above. In any event, τ is real analytic. Since we know the set $\{q : |[q]| = 2\}$ is non-empty, τ is globally anti-holomorphic and moreover cannot fix the point p. This gives a contradiction. □

We now come to the main goal. Simply take the anti-holomorphic map τ defined in the proof above. Checking on a topological base for Σ, it is clear that τ is a continuous and open mapping. As $\tau^2 = 1$, it is an anti-holomorphic diffeomorphism of Σ. The quotient

$$\Sigma_0 = \Sigma/\tau = \Sigma/\sim$$

is the sought Klein surface.

Remark 5.4.12. We can read off an atlas as follows. If p is not a ramification point, $\sim$ identifies a small neighbourhood of p with no ramification points to some other neighbourhood. The coordinate chart near p then gives the chart on Σ_0. Transition maps can be holomorphic or anti-holomorphic. If p is a ramification point, the quotient gives us a space $\mathcal{K}$ as above, with two different choices for coordinates: natural coordinates for Φ, or the complex conjugation of those coordinates. Both holomorphic and anti-holomorphic transition maps exist. We omit the technical details.

With regard to Theorem 5A, we are left to discuss the projection $\pi : \Sigma \to \Sigma_0$ and the harmonic map f. The remark gives coordinate expressions for π in which we see it is dianalytic. Σ is actually the analytic double of Σ_0, and f clearly descends to a continuous map f_0 on Σ_0 that is harmonic by definition. This finishes the proof of Theorem 5A.

Minimal Klein surfaces

For completeness, we extend the work of Gulliver-Osserman-Royden on minimal maps to the anti-holomorphic case. To the author's knowledge, the result of this subsection is new.

We begin with a minimal map $f : (\Sigma, \mu) \to (M, \nu)$ and anti-holomorphic $h : \Omega_1 \to \Omega_2$ such that $f \circ h = f$. As in our approach for non-minimal maps, we first apply

[GOR73, Proposition 3.24] to assume Σ has no points that are holomorphically related. We then define $\sim$ exactly as in Section 5.1, but allow the diffeomorphisms involved to be conformal. The application of their result assures that Lemma 5.4.6 goes through for $\sim$. The proof of Proposition 3.14 in [GOR73] applies to the map f, which proves the relation $\sim$ is Hausdorff.

For ramification, the distinction is that $\Phi = 0$, so we cannot apply the usual methods. At the same time, all singular points are good branch points. Recall from Section 5.1 that near a branch point p of order m we can find a neighbourhood of p with a holomorphic coordinate z and coordinates $(x_1, \ldots, x_n)$ around $f(p)$ so that f is given by

$$x_1 = \mathrm{re} z^m \ , \ x_2 = \mathrm{im} z^m \ , \ x_k = \eta_k(z) \ , \ k \geq 3,$$

where $\eta_k(z) \in o(|z|^m)$. If we have distinct p_1, p_2 in this neighbourhood with $p_1 \sim p_2$, then the anti-holomorphic map h that relates the two must satisfy

$$(h(z))^m = \overline{z}^m.$$

Consequently, h is of the form

$$h(z) = e^{\frac{2\pi i j}{m}} \overline{z}$$

for some $j = 0, 1, \ldots, m-1$. Up until Lemma 5.4.11, almost word-for-word, one can run through the rest of the proof of the anti-holomorphic case for non-harmonic maps. The only difference is that we use coordinate disks rather than natural coordinates for a holomorphic differential. The analogue of Lemma 5.4.11 can be worked out without difficulty.

Lemma 5.4.13. *In this setting, $p \in \Sigma$ satisfies $[p] = \{p\}$ if and only if f ramifies at p.*

Proof. Even if f is minimal, analytic continuation is possible. Given a curve γ starting in Ω_1, we can analytically continue h along γ as long as γ and $h(\gamma)$ stay sufficiently far away from the set

$$\{p \in \Sigma : [p] \text{ intersects the branch set of } f\}.$$

To do so, we first can assume f is a diffeomorphism on Ω_i and injective on $\overline{\Omega_i}$. If q is the first point at which γ strikes $\partial\Omega_1$, then $h(q)$ is well-defined. We choose disks U_1 and U_2 around q and $h(q)$ respectively such that $f|_{\overline{U_i}}$ is a diffeomorphism. We then invoke the unique continuation property of Gulliver-Osserman-Royden to find a smaller disk $U_1' \subset U_1$ such that $f(U_1') \subset U_2$. Setting $U_2' = f|_{U_2}^{-1}(f(U_1'))$, the map

$$f|_{U_2'}^{-1} \circ f|_{U_1'} : U_1' \to U_2'$$

is a conformal diffeomorphism that continues h, and is therefore anti-holomorphic. This establishes the continuation result. We also note that [GOR73, Proposition 3.14] implies that if γ is a curve along which we have continued h, then $p \sim h(p)$ for all p in the image of γ.

We suppose there is a point p at which f is unramified and such that $[p] = \{p\}$. Choose a coordinate disk Ω around p in which no two points are equivalent. We show that under this assumption we must have $[q] = \{q\}$ for all $q \in \Omega$. If not, then there is a $q \in \Omega$ and a $q' \notin \Omega$ such that $q \sim q'$, and an anti-holomorphic diffeomorphism h relating a neighbourhood of q to one of q'. We analytically continue h along a simple curve from q to p that does not touch any point that is equivalent to a branch point of f. It is easy to build such a curve, since the branch set is discrete, and equivalence classes can have only two points. Using the reasoning from Lemma 5.4.11, we can continue along all of γ and $h(\gamma(1)) = p$. Now, note that by assumption there is no pair $p_1, p_2 \in B_\epsilon(p)$ with $p_1 \in \gamma([0,1])$ and $p_1 \sim p_2$. Taking γ to the endpoint gives that $h(\gamma(t))$ lies outside $B_\epsilon(p)$ for $t \in [0,1]$ sufficiently close to 1. This contradicts $h(p) = p$, and hence yields $[q] = \{q\}$ for all $q \in \Omega$. We can now conclude the proof exactly as we did in Lemma 5.4.11. □

The remainder of the content in the previous subsection goes through verbatim. The resulting map from the Klein surface to M is minimal.

Chapter 6

MODULI SPACES OF HARMONIC SURFACES

6.1 Introduction

The theory of harmonic maps from surfaces is well developed and has proved to be a useful tool in geometry and topology. There are many broadly applicable existence theorems for harmonic maps, but, compared to other objects like minimal surfaces, their geometry is neither well behaved nor easy to understand. Locally, the most we can say about a random harmonic map from a surface is that, in a good choice of coordinates, up to small perturbations it agrees with an n-tuple of harmonic homogeneous polynomials (see the Hartman-Wintner theorem [HW53]). And in contrast, minimal maps are weakly conformal and hence have much nicer local properties. In this chapter, we consider moduli spaces of harmonic surfaces and study their *generic* qualitative behaviour through transversality theory. The goal is twofold: to find nice properties shared by a wide class of harmonic maps, and to further develop the methods and analysis for future problems.

Throughout the chapter, let Σ be a closed and orientable surface of genus $g \geq 2$, and let M be an orientable n-manifold, $n \geq 3$. Fixing integers $r \geq 2$ and $k \geq 1$, as well as $\alpha, \beta \in (0,1)$ with $\alpha \geq \beta$, denote by $\mathfrak{M}(\Sigma)$ an open and connected subset of the space of $C^{r,\alpha}$ hyperbolic metrics on Σ, and by $\mathfrak{M}(M)$ an open and connected subset of the space of $C^{r+k,\beta}$ metrics on M. Set $C(\Sigma, M)$ to be the space of $C^{r+1,\alpha}$ mappings from $\Sigma \to M$. These all have C^∞ Banach manifold structures.

Definition 6.1.1. In this chapter, a homotopy class $\mathbf{f}$ of maps from Σ to M is admissible if the subgroup $\mathbf{f}_*(\pi_1(\Sigma)) \subset \pi_1(M)$ is not abelian.

In the settings in this chapter, this agrees with the definition from the previous chapter. For the whole chapter, we fix an admissible class homotopy class $\mathbf{f}$ of maps from Σ to M and assume the following.

Technical Assumption. For all $(\mu, \nu) \in \mathfrak{M}(\Sigma) \times \mathfrak{M}(M)$, there exists a unique harmonic map $f_{\mu,\nu} : (\Sigma, \mu) \to (M, \nu)$ in the class $\mathbf{f}$, and $f_{\mu,\nu}$ is a non-degenerate critical point of the Dirichlet energy functional.

The technical assumption is satisfied by a wide range of manifolds M and families of metrics. The central example is that of a closed manifold M, with $\mathfrak{M}(M)$ consisting of negatively curved metrics. In Section 6.2, we give more examples that are of interest in geometry and topology.

With this assumption, $\mathfrak{M} = \mathfrak{M}(\Sigma) \times \mathfrak{M}(M)$ may be viewed as a moduli space of harmonic surfaces inside M. More precisely, a result of Eells-Lemaire [EL81, Theorem 3.1] (a consequence of the implicit function theorem for Banach manifolds) implies that around each pair of metrics $(\mu_0, \nu_0) \in \mathfrak{M}$, there is a neighbourhood $U \subset \mathfrak{M}$ such that the mapping from $U \to C(\Sigma, M)$ given by

$$(\mu, \nu) \mapsto f_{\mu,\nu}$$

is C^k. By uniformization and conformal invariance of energy, the restriction to hyperbolic metrics on the source does not give up any information.

Our first result concerns the notion of a somewhere injective map, which is originally from symplectic topology.

Definition 6.1.2. A C^1 map $f : \Sigma \to M$ is somewhere injective if there exists a regular point $p \in \Sigma$ such that $f^{-1}(f(p)) = \{p\}$. Otherwise, we say f is nowhere injective.

Remark 6.1.3. When the somewhere injective harmonic map f has isolated singular set, or more generally the set $A(f)$ from Section 6.4 is connected, it is injective on an open and dense set of points. This follows from the Aronszajn theorem [Aro57, page 248] (see also [Sam78, Theorem 1]).

We let $\mathfrak{M}^* \subset \mathfrak{M}$ denote the space of metrics (μ, ν) such that $f_{\mu,\nu}$ is somewhere injective.

Theorem 6A. The subset $\mathfrak{M}^* \subset \mathfrak{M}$ is open, dense, and connected.

Remark 6.1.4. A minimal map on a Riemann surface is nowhere injective if and only if it factors through a holomorphic branched cover [GOR73, Section 3], or the surface admits an anti-holomorphic involution that leaves the map invariant [Sag21a, Theorem 1.1]. Pseudoholomorphic maps from a surface to a symplectic manifold have the same property (see [MS12, Chapter 2.5]). Harmonic maps, in contrast, do not have the same rigidity.

Remark 6.1.5. These results for minimal surfaces, or more general branched immersions in the sense of [GOR73], are proved using the factorization theorem of Gulliver-Osserman-Royden. The analogue of this theorem for harmonic surfaces is the subject of our paper [Sag21a]. Fittingly, the factorization theorem for harmonic maps [Sag21a, Theorem 1.1] is a crucial ingredient in the proof of Theorem 6A.

Secondly, we prove a set of results about the structure of the moduli space near somewhere injective maps. The somewhere injective condition, while not obviously significant, comes into play in transversality arguments used for moduli spaces of minimal surfaces (see the paper of Moore [Moo06] and the book that followed [Moo17, Chapter 5]) and pseudoholomorphic curves (see [MS12, Chapter 3]). In some sense, nowhere injective surfaces play the same role as reducible connections in Yang-Mills moduli spaces.

Theorem 6B. Suppose $\dim M \geq 4$, and let (μ, ν) be such that $f_{\mu,\nu}$ is somewhere injective and has isolated singularities. Then there exists a neighbourhood $U \subset \mathfrak{M}$ containing (μ, ν) such that the space of harmonic immersions in U is open and dense. If $\dim M \geq 5$, then the space of harmonic immersions in U is also connected.

Theorem 6C. Suppose $\dim M \geq 5$, and let (μ, ν) be such that $f_{\mu,\nu}$ is somewhere injective and has isolated singularities. Then there exists a neighbourhood $U \subset \mathfrak{M}$ containing (μ, ν) such that the space of harmonic embeddings in U is open and dense. If $\dim M \geq 6$, then the space of harmonic immersions in U is also connected.

We obtain the following corollary.

Corollary 6D. If $\dim M \geq 4$, then any somewhere injective harmonic map with isolated singularities can be approximated by harmonic immersions. If $\dim M \geq 5$, then any such harmonic map can be approximated by harmonic embeddings. The embeddings can be chosen to be immersed.

At the very end of the chapter, we explain our use of the hypothesis that f has isolated singularities and the possibility of removing it. We propose the following conjecture.

Conjecture 6E. The weak Whitney theorems hold for harmonic surfaces. That is,

1. if $\dim M \geq 4$, the space of harmonic immersions in $\mathfrak{M}$ is open and dense, and connected if $\dim M \geq 5$, and

2. if $\dim M \geq 5$, the space of harmonic embeddings in $\mathfrak{M}$ is open and dense, and connected if $\dim M \geq 6$.

The weak Whitney theorems [Whi36, Theorem 2] state that a regular enough map between manifolds $g : X \to Y$ can be approximated by immersions if $\dim Y \geq 2 \dim X$, and by embeddings if $\dim Y \geq 2 \dim X + 1$. Combined with the Whitney trick, they yield the Whitney immersion theorem and the Whitney embedding theorem. One can give modern proofs of the weak theorems via transversality theory. The conjecture holds for Moore's moduli spaces of minimal surfaces [Moo17, Theorem 5.1.1 and 5.1.2].

Outline of chapter and proofs

In the next section, we define harmonic maps and associated Jacobi operators, and give examples of moduli spaces of harmonic surfaces. These examples mostly require $\mathfrak{M}(M)$ to be a space of non-positively curved metrics. We prove Proposition 6.2.11 to show that some positive curvature is allowed. In Section 6.3, we compute precise expressions near singularities for reproducing kernels for Jacobi operators. We proceed by constructing parametrices for some objects that resemble Green's operators.

The proof of Theorem 6A is contained in Sections 6.4 and 6.5. Section 6.4 is the reduction to a transversality lemma and Section 6.5 is the proof of that lemma. Since the details are technical, we explain the proof here. For disjoint open disks $P, Q \subset \Sigma$ and $\delta > 0$, we set

$$\mathcal{D}(P, Q, \delta) = \{(\mu, \nu) \in \mathfrak{M} : d_\nu(f(P), f(Q)) > \delta, P, Q \subset \Sigma^{\mathcal{SR}}(f)\},$$

where $\Sigma^{\mathcal{SR}}(f)$ is the super-regular set, to be defined in Section 6.4. For Theorem 6A, it is enough to prove that somewhere injective maps are open, dense, and connected in restriction to $D(P, Q, \delta)$'s. We define a map

$$\Theta : \Sigma^2 \times (P \times Q \times \mathcal{D}) \to M^2 \times M^2,$$

$$\Theta(r, s, p, q, \mu, \nu) = (f_{\mu,\nu}(r), f_{\mu,\nu}(s), f_{\mu,\nu}(p), f_{\mu,\nu}(q)).$$

If $\Theta(r, s, p, q, \mu, \nu)$ avoids the diagonal, then $f_{\mu,\nu}$ is somewhere injective. So, if we show that Θ is transverse to the diagonal, then the preimage has codimension $2 \dim M \geq 6$. Since Σ^2 has dimension 4, the projection of $\Theta^{-1}(L)$ to $\mathcal{D}$ should be dense and connected. One complication is that this projection may not itself be a manifold, so we have to prove connectedness directly using transversality theory.

The real substance of the proof is to show that Θ is transverse to L. We argue by contradiction and suppose that Θ is not a submersion at points that map to L. Invoking an existence result for reproducing kernels, this implies that at some pair of metrics (μ, ν), there is a non-zero section $X : \Sigma \to \Gamma(\mathbf{F})$ such that for all variations through harmonic maps $V \in \Gamma(\mathbf{F})$,

$$\int \langle \mathbf{J}V, X \rangle dA_\mu = 0. \tag{6.1}$$

Above, $\mathbf{F}$ is the pullback bundle $f_{\mu,\nu}^* TM$, $\Gamma(\mathbf{F})$ is the space of sections, and $\mathbf{J}$ is the Jaocbi operator for $f_{\mu,\nu}$. X satisfies the Jacobi equation away from its singularities, and we show that these singularities can be resolved, making use of the local expressions from Section 6.3. X thus extends to a global Jacobi field, which is our contradiction.

To resolve the singularities, we vary the target metric on M to find harmonic variations V such that (6.1) gives us good information. One could also vary the source metric on Σ, but it shouldn't work too well: in some situations where the homotopy class $\mathbf{f}$ is compressible and the original harmonic surface $f_{\mu,\nu}(\Sigma) \subset M$ is totally geodesic, we will have $f_{\mu,\nu}(\Sigma) = f_{\mu+\dot{\mu},\nu}(\Sigma)$ for all admissible variations $\dot{\mu}$. Thus, we can't in general perturb away from a nowhere injective map.

The singularities of X are at intersection points of harmonic disks $f(\Omega_1), f(\Omega_2) \subset M$, and we divide into cases: either the disks are tangential at the intersection point or they are not. When not only are they tangential but also $f(\Omega_1) = f(\Omega_2)$ and $f|_{\Omega_2}^{-1} \circ f|_{\Omega_1}$ is conformal, then our approach simply cannot work. To give one example of what can go wrong, if f factors through a holomorphic branched covering map (in which case the homotopy class is compressible) and Ω_1 and Ω_2 are related by a covering transformation, then no matter how we vary the target metric, the harmonic maps will continue to factor in this way and identify the two sets. This is where we use the factorization theorem [Sag21a, Theorem 1] to say that the set of metrics giving rise to harmonic maps with this property can be removed from the moduli space without disconnecting it. The tangential case is then settled using the super-regular condition (Section 6.4). For the non-tangential case, we choose variations supported in what we call "fat cylinders" that give $\mathbf{J}V$ more support near some places than others.

In Section 6.6, we prove Theorems 6B and 6C, yet again by transversality theory. Right now, we explain only Theorem 6B, since Theorem 6C is a similar argument. We trivialize the complexified tangent bundle of M and let σ be the projection onto the $\mathbb{C}^n$ factor. Then we define a map

$$\Psi : \tilde{\Sigma} \times \mathfrak{M} \to \mathbb{C}^n, \Psi(p, \mu, \nu) = \sigma(f_z(p)),$$

where z is the uniformizing parameter for the metric μ on the universal cover $\tilde{\Sigma}$. We try to show that Ψ is transverse to $\{0\}$ and the submanifold of $\mathbb{C}^n$ consisting of rank 1 vectors. We achieve this near (μ, ν) that yield somewhere injective harmonic maps with isolated singularities. Modulo transversality details, this gives Theorem 6B. As in the proof of Theorem 6A, we suppose transversality fails, and then we find there must be a section $X : \Sigma \to \Gamma(\mathbf{F})$ that is annihilated by all $\mathbf{J}V$, where V ranges over variations through harmonic maps.

The contradiction is different from that of Theorem 6A. We attach a particular holomorphic structure to the complexification $\mathbf{E}$ of $\mathbf{F}$. Using somewhere injectivity and a lemma of Moore [Moo06], we find there is an open set Ω on which X is the real part of a holomorphic section of a special holomorphic line bundle $\mathbf{L} \subset \mathbf{E}$. Making use of the isolated singularity condition, we analytically continue the "imaginary part," so that X is the real part of a globally defined meromorphic section Z of $\mathbf{L}$. From Section 6.3 we see that Z has at most a simple pole at one point. We then check that the order of this section does not match up with the degree of $\mathbf{L}$. This final contradiction can also be seen through Riemann-Roch.

Acknowledgements

In the case of 3-manifolds, an argument for Theorem 6A is given in an unpublished manuscript of Vladimir Marković [Mar18]. The proof had a few small issues, which have been resolved here, and also more needs to be done for the argument to work in all dimensions.

I thank Vlad for allowing me to absorb content from his manuscript, and for the many discussions that we had related to this project. This work here is intended to be independent and self-contained, and the reader should not have to consult [Mar18]. I have tried to keep similar notation.

6.2 Moduli spaces

Before we begin, we review transversality theorems for Banach manifolds.

Transversality theorems

We state the transversality theorems used in the proofs. For more background on transversality theory and Banach manifolds in general, we refer the reader to

[AMR88] and [Moo17, Chapter 1].

Definition 6.2.1. Let X, Y be C^1 manifolds, $f : X \to Y$ a C^1 map, and $W \subset Y$ a submanifold. We say f is transversal to W at a point $x \in X$ if $f(x) = y \notin W$ or if $f(x) = y \in W$ and

- the inverse image $(T_x f)^{-1}(T_y W)$ splits and
- the image $T_x f(T_x X)$ contains the closure of the complement of $T_y W$ in $T_y Y$.

We say f is transversal to W if we have transversality for every $x \in X$.

The central transversality theorem is below.

Theorem 6.2.2 (Transversality Theorem for Banach Manifolds)**.** *Let X, Y be C^r manifolds ($r \geq 1$), $f : X \to Y$ a C^r map, and $W \subset Y$ a C^r submanifold. Then if f is transverse to W,*

- *$f^{-1}(W)$ is a C^r submanifold of X and*
- *if W has finite codimension in Y, then $\mathrm{codim}_X f^{-1}(W) = \mathrm{codim}_Y W$.*

Let A, X, Y be C^r manifolds and $\delta : A \to C^r(X, Y)$ a map. We say δ is a C^r representation if and only if the evaluation map $\beta : A \times X \to Y$ given by

$$\beta(a, x) = \delta(a)(x)$$

is C^r.

Theorem 6.2.3 (Parametric Transversality Theorem)**.** *Let A, X, Y be C^r manifolds and $\delta : A \to C^r(X, Y)$ a C^r representation. Let $W \subset Y$ be a C^r submanifold and let β be the associated evaluation map. Let A_W be the set of $a \in A$ such that $\delta(a)$ is transverse to W. Assume that*

- *X has finite dimension n and W has finite codimension q in Y,*
- *A and X are second countable,*
- *$r > \max(0, n - q)$, and*
- *the evaluation map is transverse to W.*

Then A_W is residual in A.

Recall that a subset of a topological space is residual if it is a countable intersection of dense open subsets. By the Baire Category theorem, residual sets are dense.

Conventions

Given two non-negative functions defined on some set X, we say

$$f \lesssim g$$

if there exists a constant $C > 0$ such that $f(x) \leq Cg(x)$ for all $x \in X$. We define $f \gtrsim g$ similarly. If $X = \mathbb{R}$, and f, g, h are functions from $X \to [0, \infty)$, we write

$$f = g + O(h)$$

to mean $|f - g| \lesssim h$. Given Banach spaces $(B_i, ||\cdot||_i)$, equipped with an inclusion map $B_1 \to B_2$, we write

$$||V||_2 \lesssim ||V||_1$$

to mean there is a uniform constant $C > 0$ such that for all $V \in B_1$, we have $||V||_2 \leq C||V||_1$.

Throughout, the space of $C^{n,\alpha}$ sections of a $C^{n,\alpha}$ vector bundle V over M is denoted $\Gamma(V)$. Here we are allowing $n = \infty$ and $n = \omega$ (real analytic). Given a map $f : \Sigma \to M$, we let $\mathbf{F} = f^*TM$ be the pullback of TM over Σ. If f is $C^{n,\alpha}$ then $\mathbf{F}$ is a $C^{n,\alpha}$ bundle. As in the preliminaries chapter, $TM^{\mathbb{C}} = TM \otimes \mathbb{C}$ is the complexification of the tangent bundle of M and $\mathbf{E} := f^*TM^{\mathbb{C}}$ is the pullback bundle. Also recall that $f_z = df(\frac{\partial}{\partial z})$ is a local holomorphic section of $\mathbf{E}$ (see the preliminaries chapter).

Under the technical assumption, $f_{\mu,\nu}$ will be the unique harmonic map from $(\Sigma, \mu) \to (M, \nu)$. When working with fixed (μ, ν) we sometimes write $f = f_{\mu,\nu}$. The connections $\nabla^{\mathbf{F}}$ and $\nabla^{\mathbf{E}}$ will be used quite often, so henceforward we condense

$$\nabla := \nabla^{\mathbf{F}},\ \nabla^{\mathbf{E}}$$

when the context is clear. A section $W \in \Gamma(\mathbf{E})$ may be uniquely written as $W = \mathrm{Re}(W) + i\mathrm{Im}(W)$, where $\mathrm{Re}(W), \mathrm{Im}(W) \in \Gamma(\mathbf{F})$.

Jacobi operators

The tension field may be seen as a map

$$\tau : \mathfrak{M} \times C(\Sigma, M) \to \Gamma(\mathbf{F}).$$

For (μ, ν) fixed, the derivative in the $C(\Sigma, M)$ direction is the Jacobi operator [EL81], which we are about to define. The Dirichlet energy is non-degenerate—or the technical assumption from the introduction is satisfied—if and only if the Jacobi operator has no kernel.

Let $f : (\Sigma, \mu) \to (M, \nu)$ be a C^2 (not necessarily harmonic) map and as before set $\mathbf{F} = f^*TM$. Let Δ denote the Laplacian induced by the connection $\nabla^{\mathbf{F}}$ and $R = R^M$ the curvature tensor of the Levi-Civita connection of ν. The Jacobi operator $\mathbf{J}_f = \mathbf{J} : \Gamma(\mathbf{F}) \to \Gamma(\mathbf{F})$ is defined

$$\mathbf{J}V = \Delta V - \text{trace}_\mu R(df, V)df \ , \ V \in \Gamma(\mathbf{F}).$$

If $z = x + iy$ is a local complex parameter and the conformal density is μ, then

$$\mathbf{J}V = -\nabla_x \nabla_x V - \nabla_y \nabla_y V - |\mu|^{-1}(R(f_x, V)f_x + R(f_y, V)f_y). \tag{6.2}$$

The Jacobi operator is a second order strongly elliptic linear operator and it is essentially self-adjoint in the sense that

$$\int_\Sigma \langle \mathbf{J}V, W\rangle dA = \int_\Sigma \langle V, \mathbf{J}W\rangle dA$$

for all $V, W \in \Gamma(\mathbf{F})$. Above, recall that $\langle \cdot, \cdot \rangle = \langle \cdot, \cdot \rangle_\nu$ is the inner product on $\mathbf{F}$ induced by the metric ν on M. The integration over Σ is with respect to the volume form $dA = dA_\mu$.

Remark 6.2.4. The assumption $r \geq 3$ guarantees the coefficients of the operator are at least C^2. This is relevant for the regularity theory, and we use this implicitly throughout the chapter.

Calculus on vector bundles

The following Banach spaces will come into play.

- For $1 \leq p < \infty$, $(L^p(\mathbf{F}), || \cdot ||_p)$ is the space of L^p-bounded measurable sections of $\mathbf{F}$.
- For $k \in \mathbb{Z}_+$, $1 \leq p < \infty$, $(W^{k,p}(\mathbf{F}), || \cdot ||_{k,p})$ is the Sobolev space of k-times weakly differentianble sections with L^p derivatives with respect to the Levi-Civita connection.
- For $k \in \mathbb{Z}_+, \alpha \in (0, 1)$, $(C^{k,\alpha}(\mathbf{F}), || \cdot ||_{k,\alpha})$ is the space of k-times differentiable sections whose k^{th} derivatives are α-Hölder.

- We can define these spaces in restriction to any open set $\Omega \subset \Sigma$. For $L^p(\mathbf{F}|_\Omega)$, we use the notation $||\cdot||_{p,\Omega}$, and likewise for the other Banach spaces.

Above, if the vector bundle is only $C^{n,\alpha}$, we restrict $k \leq n$. For precise definitions and other basic facts, see [Nic21, Chapter 10]. If we choose a different metric or connection on $\mathbf{F}$, the relevant Sobolev spaces are equal as sets of sections, and the identity map is bicontinuous. Thus, it is unambiguous to write $W^{k,p}(\mathbf{F})$ (and likewise for the other spaces), while not specifying the choices involved.

Now we recall some results relevant to the Jacobi operator. A Jacobi field is a section $V \in \Gamma(\mathbf{F})$ such that $\mathbf{J}V = 0$. We again refer the reader to [Nic21, Chapter 10]. From the basic elliptic theory, essential self-adjointness implies the following.

Proposition 6.2.5. *Suppose there are no non-zero Jacobi fields. Then for every $p > 1$ and $0 < \alpha < 1$, the operator $\mathbf{J}$ extends to a family of isomorphisms $\mathbf{J} : W^{2,p}(\mathbf{F}) \to L^p(\mathbf{F})$, $\mathbf{J} : C^{2,\alpha}(\mathbf{F}) \to C^{0,\alpha}(\mathbf{F})$. Each such isomorphism preserves the subspace of smooth sections.*

The result below is a consequence of the Weyl lemma for linear elliptic operators.

Proposition 6.2.6. *Let $\Omega \subset \Sigma$ be open and V be a measurable section over Ω such that $||V||_{p,\Omega} < \infty$ for some $1 < p < \infty$. If $\mathbf{J}V = 0$ weakly on Ω, then V is as regular as the bundle $\Gamma(\mathbf{F})$, and $\mathbf{J}V \equiv 0$ on Ω.*

Examples of moduli spaces

Here we list some examples of manifolds M and spaces of metrics $\mathfrak{M}(M)$ satisfying the technical assumption.

Example 6.2.7. M is a closed n-manifold that admits a metric of negative curvature, with $\mathfrak{M}(M)$ consisting of negatively curved metrics.

In this case, Sampson proves in [Sam78, Theorem 4] that there are no smooth Jacobi fields. In fact, he proves a more general result.

Theorem 6.2.8 (Sampson, Theorem 4 in [Sam78])**.** *Let (M, ν) be a closed Riemannian manifold with non-positive curvature. Suppose $f : (\Sigma, \mu) \to (M, \nu)$ is an admissible harmonic map and there is at least one point p at which all sectional curvatures of M at $f(p)$ are strictly negative. Then there are no non-zero Jacobi fields.*

Compactness of the target is not important.

Example 6.2.9. M is not necessarily compact, all $\mathfrak{M}(M)$ are negatively curved, and the induced mapping between the fundamental groups is irreducible.

This class of examples includes admissible classes f such that at least one simple closed curve is mapped by $\mathbf{f}_*$ to a class whose geodesic length is positive. Even more specific examples include convex cocompact manifolds of negative curvature, such as quasi-Fuchsian 3-manifolds.

To demonstrate the level of generality, we prove a slight extension of Sampson's result that allows for some positive curvature.

Definition 6.2.10. A pair $(\mu, \nu) \in \mathfrak{M}$ is $\mathbf{f}$-admissible if (M, ν) is non-positively curved and there exists a map $f \in \mathbf{f} \cap C(\Sigma, M)$ that is harmonic with respect to (μ, ν) and a point $p \in \Sigma$ such that all sectional curvatures of M are negative at $f(p)$.

As discussed, $\mathbf{f}$-admissibility implies uniqueness of the harmonic map.

Proposition 6.2.11. *Suppose (μ, ν) is $\mathbf{f}$-admissible, and let ν_n be a sequence of metrics converging to ν. Furthermore, assume $f_j : (\Sigma, \mu) \to (M, \nu_j)$ is a sequence of harmonic maps converging to a harmonic map $f : (\Sigma, \mu) \to (M, \nu)$. Then $\mathbf{J}_{f_j}$ admits no non-trivial C^2 Jacobi fields for sufficiently large j.*

This gives another example of interest.

Example 6.2.12. A sufficiently small neighbourhood of an $\mathbf{f}$-admissible pair inside the moduli space of all Riemannian metrics.

Proof. Since the harmonic maps f_j, f are homotopic through $C^{n+1,\alpha}$ maps, the bundles $f_j^*TM = \mathbf{F}_j$ and $f^*TM = \mathbf{F}$ are isomorphic in the $C^{n+1,\alpha}$ category. We identify them all with the bundle $\mathbf{F}$. Under this identification, $\mathbf{F}$ inherits a family of Riemannian metrics $\langle \cdot, \cdot \rangle_j$ with corresponding Levi-Civita connections ∇_j, as well as elliptic operators $\mathbf{J}_{f_j}$. Since $\nu_j \to \nu$ and $f_j \to f$, we have convergence of associated objects $\langle \cdot, \cdot \rangle_j \to \langle \cdot, \cdot \rangle := \langle \cdot, \cdot \rangle_\nu$, $\nabla_j \to \nabla := \nabla^\nu$, and $\mathbf{J}_{f_j} \to \mathbf{J}_f$ in the relevant topologies. Henceforth, rename $\mathbf{J}_j = \mathbf{J}_{f_j}$.

One could write our Sobolev spaces more precisely as

$$W^{k,p}(\mathbf{F}, \mu, \nu, \nabla).$$

We write $W^{1,2}(\mathbf{F})$ to denote the usual Sobolev space for $\mathbf{F}$, and $W^{1,2}(\mathbf{F}_j)$ for $W^{1,2}(\mathbf{F}_j, \mu, \nu_j, \nabla_j)$ (and likewise for the L^2 spaces). In our notation, we set $||\cdot||_2$, $||\cdot||_{1,2}$ to be the norms for $\mathbf{F}$ and $||\cdot||_{2,j}$, $||\cdot||_{1,2,j}$ to be the norms for $\mathbf{F}_j$. From bicontinuity of the identity map between these Banach spaces, there exists $C_j \geq 1$ such that for all $V \in \Gamma(\mathbf{F})$,

$$C_j^{-1}||V||_2 \leq ||V||_{2,j} \leq C_j||V||_2, \text{ and}$$
$$C_j^{-1}||V||_{1,2} \leq ||V||_{1,2,j} \leq C_j||V||_{1,2}.$$

It is an easy exercise to show that $C_j \to 1$ as $j \to \infty$.

To prove the lemma, assume for the sake of contradiction that there is a subsequence (which we still denote ν_j) and a family of non-zero sections $V_j \in C^2(\mathbf{F})$ such that $\mathbf{J}_j V_j = 0$ and $||V_j||_2 = 1$. Necessarily,

$$\int_\Sigma \langle \mathbf{J}_j V_j, V_j \rangle_j dA = 0.$$

Unravelling the definition of the Jacobi operator and integrating by parts, we obtain

$$\int_\Sigma |\nabla_j V_j|_j^2 - \int_\Sigma \langle \operatorname{trace}_\mu R^{\nu_j}(df_j, V_j)df_j, V_j \rangle_j dA = 0. \tag{6.3}$$

Remark 6.2.13. We implicitly use that ∇_j is the Levi-Civita connection for ν_j to integrate by parts. If we tried to use the metric ν, then some extra terms involving Christoffel symbols would appear.

Let σ_j denote the maximum of 0 and the largest sectional curvature of M in the image of f_j. Then

$$\langle \operatorname{trace}_\mu R^{\nu_j}(df_j, V_j)df_j, V_j \rangle_j \leq \sigma_j |\operatorname{trace}_\mu(df_j)|_j^2 |V_j|_j^2$$

pointwise. Convergence of $\nu_j \to \nu$ and $f_j \to f$ then implies

$$\langle \operatorname{trace}_\mu R^{\mu_j}(df_j, V_j)df_j, V_j \rangle_j \lesssim \sigma_j |V_j|_j^2.$$

Substituting into (6.3) we see

$$\int_\Sigma |\nabla_j V_j|_j^2 \lesssim \sigma_j \int |V_j|_j^2 dA.$$

Again using convergence of $\nu_j \to \nu$, we see $\sigma_j \to 0$ as $j \to \infty$. Choosing j large enough so that $C_j \lesssim 1$, and using $||V_j||_2 = 1$ we obtain

$$\int_\Sigma |\nabla_j V_j|_j^2 dA \lesssim \sigma_j \to 0 \tag{6.4}$$

as $j \to \infty$.

The above result gives uniform control on the $W^{1,2}(\mathbf{F}_j)$ norm of V_j, and hence we also have control on the $W^{1,2}(\mathbf{F})$ norm. Since $W^{1,2}(\mathbf{F})$ is reflexive, the Banach-Alaoglu theorem guarantees the existence of a subsequence along which V_j converges weakly in $W^{1,2}(\mathbf{F})$ to a section $V \in W^{1,2}(\mathbf{F})$. By the Rellich lemma, we may pass to a further subsequence to obtain strong convergence in L^2, so that $||V||_2 = 1$.

We now claim that $\nabla V = 0$ in the sense of distributions, i.e., it is an almost everywhere constant field. Working in a conformal parameter $z = x + iy$ for μ, we write out

$$|\nabla_j V_j|^2 = \mu^{-1}\left(|\nabla_{j,x} V_j|_\nu^2 + |\nabla_{j,y} V_j|_\nu^2\right)$$

and observe the linear maps $\nabla_{j,x}$, $\nabla_{j,y}$ converge strongly to ∇_x and ∇_y respectively in $\mathrm{Hom}(W^{1,2}(\mathbf{F}), L^2(\mathbf{F}))$ with respect to the operator norm $||\cdot||_{OP}$. Thus,

$$||\nabla_x V_j||_2 \leq ||(\nabla_x - \nabla_{x,j})V_j||_2 + ||\nabla_{x,j} V_j||_2 \leq ||\nabla_x - \nabla_{x,j}||_{OP}||V_j||_{1,2} + C_j||\nabla_{x,j} V_j||_{2,j}.$$

Our observation above shows the first term decays to 0 as $j \to \infty$. It follows from inequality (6.4) that the second term tends to 0 as well. Therefore $\nabla_x V_j \to 0$ strongly in L^2. By the same method we see $\nabla_y V_j \to 0$ strong as well. The claim follows.

We obtain a contradiction by arguing that $V = 0$ on a set of positive measure. This would force $||V||_2 = 1$ to be impossible. This is essentially Sampson's observation in [Sam78, Theorem 4]. From (6.3) it follows that

$$\int \langle \mathrm{trace}_\mu R^\mu(df, V)df, V\rangle dA = 0. \tag{6.5}$$

Since $(\mathbf{f}, \nu)$ is an admissible pair, there is a point $p_0 \subset \Sigma$ such that all sectional curvatures of M are negative at $f(p)$. We extract a neighbourhood $\Omega \subset \Sigma$ on which f is a regular embedding and there is a $c > 0$ such that all sectional curvatures of M at points in $f(\Omega)$ are bounded above by $-c$. Thus, from the non-positive curvature assumption on ν, if V is not zero almost everywhere, the left-hand side of (6.5) is strictly negative. As discussed above, this is a contradiction, and so we are done. □

Finally, the results should hold for some more examples that we don't pursue here: manifolds with boundary (see [EL81, Section 4]), non-orientable manifolds (Moore considers non-orientable minimal surfaces in [Moo06, Section 11]), and equivariant Anosov representations into Lie groups of non-compact type. For the analogue of the Eells-Lemaire result, applied to a suitable class of equivariant harmonic maps, we invite the reader to see [Sle20]. In these three cases, the only substantial missing ingredient is the factorization theorem [Sag21a]. A version of the theorem should be true in these contexts, but it would take us too far afield here.

6.3 Reproducing kernels

Let $p \in \Sigma$ and $U \in \mathbf{F}_p$. We say that $X : \Sigma \backslash \{p\} \to \mathbf{F}$ is a zeroth order reproducing kernel for the Jacobi operator if, for all $W \in \Gamma(\mathbf{F})$, we have

$$\langle W(p), U\rangle = \int_\Sigma \langle \mathbf{J}W, X\rangle dA.$$

For $V \in T_p\Sigma$, $X : \Sigma\backslash\{p\} \to \mathbf{F}$ is a first order reproducing kernel if, for all $W \in \Gamma(\mathbf{F})$,

$$\langle (\nabla_V W)(p), U\rangle = \int_\Sigma \langle \mathbf{J}W, X\rangle dA.$$

In the proof of the main theorems, we need explicit expressions for the singularities of reproducing kernels. We compute these singularities by constructing the kernels directly. Independent of the work below, one can find general existence results in [Mar18, Section 3].

Remark 6.3.1. From the self-adjoint property, kernels satisfy $\mathbf{J}X = 0$ away from the singularities.

The parametrices

Let (Ω, z) be a disk neighbourhood of p, and $\Omega' \subset \Omega$. In the local chart, extend the vector U to a C^2 section $U(z)$. Let $\phi_n : \Omega \to [0, 1]$ be a smooth function in Ω such that

- ϕ_n has support in $\{|\zeta| \le 1/n\}$,
- $\phi_n \equiv 1$ in $\{|\zeta| \le 1/2n\}$,
- ϕ_n integrates to 1 in Ω', and
- ϕ_n converges in the sense of distributions to the Dirac delta δ_p as $n \to \infty$.

Let $G(z, \zeta)$ be the ordinary Green's function on Ω', of the form

$$G(z, \zeta) = \frac{1}{2\pi} \log |z - \zeta|^{-1} + r(z, \zeta),$$

where r is smooth, and define a section S_n in Ω' by

$$S_n(z) = U(z) \int_{\Omega'} G(z, \zeta)\phi_n(\zeta) d\zeta \wedge d\overline{\zeta}.$$

Observe that

$$\int_{\Omega'} G(z,\zeta)\phi_n(\zeta)d\zeta \wedge d\overline{\zeta} \to G(z,0) = \frac{1}{2\pi}\log|z|^{-1} + r(z,0)$$

as $n \to \infty$ with maximum regularity on $\Omega' \backslash \{p\}$ and in L^p for all $1 < p < \infty$. We then extend S_n to a globally defined section of $\mathbf{F}$ with support in Ω, in a way that S_n converges as $n \to \infty$ in the C^∞ sense on $\Sigma \backslash \{p\}$ to a section S satisfying

$$S(z) = \frac{1}{2\pi}\log|z|^{-1}U(z) + r(z,0)U(z).$$

By the defining properties of $G(z,\zeta)$,

$$\frac{\partial^2}{\partial z \partial \overline{z}} \int_{\Omega'} G(z,\zeta)\phi_n(\zeta)d\zeta \wedge d\overline{\zeta} = \phi_n(z).$$

Using this, we compute that in Ω',

$$\begin{aligned}\nabla_z \nabla_{\overline{z}} S_n(z) = (\nabla_z \nabla_{\overline{z}} U(z)) \int_{\Omega'} G(z,\zeta)\phi_n(\zeta) + (\nabla_{\overline{z}} U)\frac{\partial}{\partial z}\int_{\Omega'} G(z,\zeta)\phi_n(\zeta) \\ + (\nabla_z U)\frac{\partial}{\partial \overline{z}}\int_{\Omega'} G(z,\zeta)\phi_n(\zeta) + U(z)\phi_n(z).\end{aligned}$$

Set $\Phi_n^1 = \nabla_z \nabla_{\overline{z}} S_n(z) - U(z)\phi_n(z)$.

Lemma 6.3.2. *For all $1 \leq p < 2$, Φ_n^1 converges along a subsequence in L^p as $n \to \infty$.*

Proof. It suffices to show that the three terms above all subconverge in L^p near 0. Since G splits into a log term and a regular term, we only need show L^p-subconvergence for

$$(\nabla_z \nabla_{\overline{z}} U(z)) \int_{\Omega'} \log|z-\zeta|\phi_n(\zeta),\ (\nabla_{\overline{z}} U)\frac{\partial}{\partial z}\int_{\Omega'} \log|z-\zeta|\phi_n(\zeta),\ (\nabla_z U)\frac{\partial}{\partial \overline{z}}\int_{\Omega'} \log|z-\zeta|\phi_n(\zeta).$$

By the basic properties of ϕ_n, $\int_{\Omega'} \log|z-\zeta|\phi_n(\zeta) \to \log|z|$ in L^p as $n \to \infty$, so the first term L^p-converges to $\nabla_z \nabla_{\overline{z}} U(z)\log|z|$ in Ω', and away from Ω' our regularity assumptions give L^p convergence. As for the second term, since $1/|z|$ is in $L^p(\Omega')$ for $1 \leq p < 2$, an application of dominated convergence shows it is equal to

$$\frac{1}{2}(\nabla_{\overline{z}} U) \int_{\Omega'} \frac{\phi_n(\zeta)}{z-\zeta} d\zeta \wedge d\overline{\zeta}.$$

Taking $n \to \infty$, we have convergence for such p to

$$\frac{\nabla_{\overline{z}} U}{z}$$

in Ω', and nice convergence outside of Ω' (note we can make this continuous by choosing U so that $\nabla_{\overline{z}} U = 0$, but this is not necessary). The final term is handled similarly. □

The zeroth order kernel

With the parametrices in hand, the remainder of the computation is a routine procedure. Complementary to Φ_n^1, set

$$\Phi_n^2 = \frac{1}{\sigma^2} R(S_n, f_z) f_{\overline{z}}.$$

Let $\Phi_n = \Phi_n^1 - \Phi_n^2$ and $\Psi_n = J^{-1}(\Phi_n)$. Here R is the complexified curvature tensor of M and σ^2 is the density of the conformal metric μ on Σ_μ.

Lemma 6.3.3. *For every $1 \leq p < 2$, the sequence of norms of $||\Phi_n||_p$ is uniformly bounded. Moreover, for any $\alpha \in (0,1)$, Ψ_n converges along some subsequence to a section $\Psi \in C^{0,\alpha}$.*

Proof. We showed above that Φ_n^1 converges in L^p to an L^p section. As for Φ_n^2, away from 0 it converges locally uniformly to some C^∞ section. Around 0 we have the estimate

$$\Phi_n^2 \leq C \log |z|$$

for some $C > 0$ and hence we have uniform L^p bounds for all p.

Invoking Proposition 6.2.11, Ψ_n is uniformly bounded in $W^{2,p}(\mathbf{E})$ for any $p \in [1, 2)$. The convergence result now follows from the Rellich-Kondrachov theorem, which gives a compact embedding from $W^{2,p} \to C^{0,\alpha}$ when $2 - 2/p > \alpha$. □

Proposition 6.3.4. *The reproducing kernel is of the form*

$$X(z) = -\frac{1}{2\pi} \log |z| U(p) + B(z), \tag{6.6}$$

where $B(z)$ is a $C^{0,\alpha}$ local section of $\mathbf{E}$ near p, for any $\alpha \in (0,1)$.

Proof. Let $W \in \Gamma(\mathbf{E})$. In local coordinates, the complexified Jacobi operator is given by

$$JW = \nabla_z \nabla_{\overline{z}} - \sigma^{-2} R(W, f_z) f_{\overline{z}}.$$

As $\mathbf{J}$ is essentially self-adjoint,

$$\begin{aligned}
\int_\Sigma \langle \mathbf{J}W, S_n \rangle dA &= \int_\Sigma \langle W, \nabla_z \nabla_{\overline{z}} S_n \rangle dA - \int_\Sigma \langle W, \sigma^{-2} R(S_n, f_z) f_{\overline{z}} \rangle dA \\
&= \int_\Sigma \langle W, \Phi_n \rangle dA + \int_\Sigma \langle W, \mu_n \rangle dA \\
&= \int_\Sigma \langle \mathbf{J}W, \Psi_n \rangle dA + \int_\Sigma \langle W, \mu_n \rangle.
\end{aligned}$$

We reorganize this to

$$\int_\Sigma \langle \mathbf{J}W, S_n\rangle dA - \int_\Sigma \langle \mathbf{J}W, \Psi_n\rangle dA = \int_\Sigma \langle W, \mu_n\rangle.$$

The term on the right tends to $\langle W, U(p)\rangle$ as $n \to \infty$. Meanwhile, passing to the subsequence from the previous lemma, the left-hand side converges to

$$\int_\Sigma \langle \mathbf{J}W, S - \Psi\rangle$$

as $n \to \infty$. Therefore, $X = S - \Psi$, and the expression for X is then derived from the local expression for S stated above and the fact that $\Psi \in C^{0,\alpha}$ for any $\alpha \in (0,1)$. □

Remark 6.3.5. We have made no attempt to optimize the regularity of $B(z)$.

Remark 6.3.6. If we change to a different (not holomorphic) coordinate $\varphi(z) = \varphi(x,y)$ with $\varphi(0) = 0$, the expression may not be so simple, but we know it behaves asymptotically like a constant multiple of $\log|\varphi|^{-1}$.

First order kernels

We don't need explicit information for the singularity for the first order kernel, but we do need to know the rate at which it blows up. A calculation is given in [Mar18, Appendix A], that strongly uses that $\nabla_{\overline{z}} = \overline{\partial}$ for the Koszul-Malgrange holomorphic structure. Here we give a different method that works in more generality (and applicable for higher order kernels).

We find the first order kernel with respect to the tangent vector $\frac{\partial}{\partial z}$. Taking real and imaginary parts, we can then get any kernel. Extend the vector U in a local trivialization so that $\nabla_z U(p) = 0$. For $z \in \Omega'$, $\zeta \in \Omega$, we thus have a well-defined function $X(z,\zeta)$ such that

$$\langle V(z), U(z)\rangle = \int_\Sigma \langle \mathbf{J}V(\zeta), X(z,\zeta)\rangle dA(\zeta)$$

for all $V \in \Gamma(\mathbf{F})$. Here we are changing our notation: $\tilde{X}$ is integrated in ζ rather than z. From the work above, $\tilde{X}(z,\zeta)$ takes the form

$$X(z,\zeta) = \frac{1}{2\pi} U(\zeta)\log|z-\zeta|^{-1} + B(z,\zeta),$$

where, for fixed z, $B(z,\zeta)$ is locally $C^{0,\alpha}$ in $\zeta \neq z$. This function is not regular and in fact blows up on the diagonal (unless $U(z) = 0$). Away from the diagonal, regularity in z is the maximum of regularity of U and the vector bundle: from the construction, we can choose μ_n and S_n to vary nicely with z for each n, and then we get the correct regularity in the limit.

Observe

$$\frac{\partial}{\partial z}\langle V(z), U(z)\rangle = \langle \nabla_z V(z), U(z)\rangle + \langle V(z), \nabla_z U(z)\rangle$$

in Ω'. In terms of our integrals, differentiating under the integral via dominated convergence, we get

$$\frac{\partial}{\partial z}\int_\Sigma \langle \mathbf{J}V(\zeta), X(z,\zeta)\rangle = \int_\Sigma \langle \mathbf{J}V(\zeta), \nabla_z X(z,\zeta)\rangle = \langle \nabla_z V(z), U(z)\rangle + \langle V(z), \nabla_z U(z)\rangle.$$

Setting $z = 0$, we find that the first order kernel is given by $\nabla_z X(0, \zeta)$. From this we deduce the following.

Proposition 6.3.7. *In the complex coordinate z, the reproducing kernel is of the form*

$$X(z) = \frac{1}{\pi z} U(p) + B(z) \tag{6.7}$$

where $B(z)$ is a $C^{0,\alpha}$ local section of $\mathbf{E}$ *near p, for any $\alpha \in (0, 1)$.*

6.4 Somewhere injective harmonic maps

As discussed earlier, Theorem 6A reduces to a transversality result, whose proof is given in the next section. Apart from a few things, the content of this section is adapted from [Mar18, Section 6].

Exceptional Riemann surfaces

Our proof of Theorem 6A involves a "super-regular" condition (defined below) that we would like to know is generic. The lemma below allows us to dismiss a class of metrics on which the condition fails.

Definition 6.4.1. A Riemann surface Σ is exceptional if either

- Σ is a holomorphic branched cover of another Riemann surface of genus at least 2 or
- Σ admits an anti-holomorphic involution.

The lemma below is a consequence of the factorization theorem established in [Sag21a].

Lemma 6.4.2. *Suppose there is a pair of disks $\Omega_1, \Omega_2 \subset \Sigma$ and a conformal diffeomorphism $h : \Omega_1 \to \Omega_2$ such that $f \circ h = f$ on Ω_1. Then the Riemann surface Σ is exceptional.*

Proof. According to [Sag21a, Theorem 1.1], if $h : \Omega_1 \to \Omega_2$ is a holomorphic map between open subsets of Σ such that $f \circ h = f$, then f factors through a holomorphic branched covering map onto a surface Σ_0. If Σ_0 has genus less than 2, then it is either a sphere or a torus. In both cases, the subgroup

$$f_*(\pi_1(\Sigma)) < \pi_1(M)$$

is abelian, which contradicts our assumption that the homotopy class **f** is admissible. If h is anti-holomorphic, the result follows from Theorem 1.1 and the discussion in Section 4 of [Sag21a]. □

This next result is well understood and one can find details in [Mar18, Appendix B].

Proposition 6.4.3. *We let $\mu \in \mathfrak{M}'(\Sigma)$ if Σ_μ is not exceptional. $\mathfrak{M}'(\Sigma)$ is an open, dense, and connected subset of $\mathfrak{M}(\Sigma)$.*

For ease of notation, we write $\mathfrak{M} = \mathfrak{M}'(\Sigma) \times \mathfrak{M}(M)$ instead of $\mathfrak{M}(\Sigma) \times \mathfrak{M}(M)$ throughout the rest of the chapter.

Super-regular points

Denote by $A(f)$ the set of $p \in \Sigma$ such that $f^{-1}(f(p)) \subset \Sigma^{reg}(f)$. Given metrics (μ, ν) and $p, q \in A(f)$, we say that the inner products $\mu(p)$ and $\mu(q)$ are conformal to each other via f if the tangent planes $df(T_p\Sigma)$ and $df(T_q\Sigma)$ agree in $T_{f(p)}M$, and if the push forwards $f_*\mu(p)$ and $f_*\mu(q)$ are collinear.

Definition 6.4.4. Given a map f, a point $p \in \Sigma$ is said to be super-regular if

- $p \in A(f)$ and
- if $f(p) = f(q)$, then $\mu(p)$ and $\mu(q)$ are not conformal to each other via f.

We denote the set of super-regular points for a map f by $\Sigma^{\mathcal{SR}}(f)$. We define $\mathcal{SR} \subset \Sigma \times \mathfrak{M}$ by $(p, \mu, \nu) \in \mathcal{SR}$ if $p \in \Sigma^{\mathcal{SR}}(f_{\mu,\nu})$.

Proposition 6.4.5. *Continuing to exclude the exceptional metrics from $\mathfrak{M}$, the set $\mathcal{SR}$ is open in $\Sigma \times \mathfrak{M}$ and $\Sigma^{\mathcal{SR}}(f)$ is open and dense in Σ.*

We first treat $A(f)$ on its own. It is due to Sampson [Sam78, Theorem 3] that the set of regular points of an admissible harmonic map is open and dense.

Lemma 6.4.6. *$A(f)$ is open and dense in Σ.*

Proof. Openness is obvious. As for density, suppose on the contrary that there is an open set $\Omega \subset \Sigma$ on which f is regular but no point is in $A(f)$. By shrinking Ω we may assume $f|_\Omega$ is an embedding. We then find a small tubular neighbourhood $N \subset M$ of the submanifold $f(\Omega)$, in which the nearest point projection $\pi : N \to f(\Omega)$ is well-defined. The set $S = f^{-1}(N) \subset \Sigma$ is then an open submanifold of Σ.

Let $g = \pi \circ f : S \to f(\Omega)$. If $y \in S$ is a singular point of f, then it is a singular point of g. By assumption, for each $u \in f(\Omega)$, the set $f^{-1}(u)$ contains a singular point of g. Thus, each point in $f(\Omega)$ is the image of a singular point $y \in S$ of the map g. This contradicts Sard's theorem. □

Proof of Proposition 6.4.5. It is clear that both $\Sigma^{\mathcal{SR}}(f)$ and $\mathcal{SR}$ are open. It remains to prove $\Sigma^{\mathcal{SR}}(f)$ is dense. Note that the set $f^{-1}(f(x))$ is finite provided $x \in A$. Indeed, if $|f^{-1}(f(x))| = \infty$, then the closed set $f^{-1}(f(x))$ has an accumulation point, at which the rank of df is necessarily strictly less than two (as f cannot be an embedding near that point).

From the previous lemma, we are left to show that the conformality condition holds on a dense subset of $A(f)$. Arguing by contradiction, suppose that on an open subset $\Omega \subset A$ we have that for every $p \in \Omega$ there exists $q \in f^{-1}(f(p))$ such that $\mu(p)$ and $\mu(q)$ are conformal to each other via f. Given $p \in \Omega$, we have a finite number of disks $D_1, \dots, D_n$ with centers p_i such that $f(p) = f(p_i)$ and with $\mu(p)$ and $\mu(p_i)$ conformal via f. We also assume f is a regular embedding on $\overline{D_i}$ and $\overline{\Omega}$. Let $C_i \subset D_i$ be the closed set of points $x \in \overline{D_i}$ with the property that there exists $y \in \overline{\Omega}$ with $f(x) = f(y)$ and such that $\mu(x)$ is conformal to $\mu(y)$ via f. We claim that for at least one i, C_i has non-empty interior. If not, then

$$C = \cup_i f(C_i) \cap f(\Omega)$$

has empty interior, for it is a finite union of closed nowhere dense sets. Choosing a sequence $(p_n)_{n=1}^{\infty} \subset \Omega \backslash (f^{-1}(C) \cap \Omega)$ converging to p, we can find another sequence $(q_n)_{n=1}^{\infty} \subset \Sigma \backslash (\cup_i \overline{D_i})$ with $f(p_n) = f(q_n)$ and $\mu(p_n)$ and $\mu(q_n)$ are conformal via f. Passing to a subsequence, the q_n converge to some point $q \in \Sigma \backslash (\cup_i D_i)$ such that $f(p) = f(q)$ and $\mu(p)$ and $\mu(q)$ are conformal via f. This contradicts our construction of the D_i, and so the claim is proved.

Relabelling so that $f(\Omega)$ and $f(D_1)$ intersect with non-empty interior, we can find open sets $\Omega_1 \subset \Omega$ and $\Omega_2 \subset D_1$ as well as a diffeomorphism $h : \Omega_1 \to \Omega_2$ such that $f \circ h = f$ on Ω_1. The metrics μ and $h^*\mu$ are pointwise conformally equivalent on Ω, and thus h is a conformal map. This contradicts Lemma 6.4.2. □

The map Θ

Denote by $\mathcal{J}$ the subset of nowhere injective maps.

Lemma 6.4.7. *$\mathcal{J}$ is closed.*

Proof. If a somewhere injective map f (which need not be harmonic) has an injective point at p, meaning $f^{-1}(f(p)) = \{p\}$, then there is an open set containing p that consists only of injective points. Indeed, choose a disk Ω around p on which f is regular and $f|_\Omega$ is injective. If the claim fails, we can find $p_n \to p$ and $q_n \in \Sigma\backslash\Omega$ such that $f(p_n) = f(q_n)$. By compactness, the q_n subconverge to a point q at which $f(p) = f(q)$, a contradiction.

So, suppose $((\mu_n, \nu_n))_{n=1}^{\infty} \subset \mathcal{J}$ converges to (μ, ν), and $f = f_{\mu,\nu}$ is somewhere injective with injective point p. There is a disk Ω around p such that f_{μ_n,ν_n} is injective on Ω. Thus, there exists $p_n \in \Sigma\backslash\Omega$ such that $f_{\mu_n,\nu_n}(p_n) = f_{\mu_n,\nu_n}(p)$, and again we find a contradiction by extracting an accumulation point $q \neq p$. □

For the remainder of Sections 6.4 and 6.5, let us replace $\mathfrak{M}$ with the complement of the set of pairs with exceptional metrics on Σ. We hope this harmless change of notation does not cause any confusion.

Let $P, Q \subset \Sigma$ be two disjoint open embedded disks in Σ. For $\delta > 0$ we let

$$\mathcal{D}(P, Q, \delta) = \{(\mu, \nu) \in \mathfrak{M} : d_\nu(f(P), f(Q)) > \delta, P, Q \subset \Sigma^{\mathcal{SR}}(f)\}.$$

It follows from Proposition 6.4.5 that $\mathcal{D}(P, Q, \delta)$ is an open subset of $\mathfrak{M}$. By Proposition 6.4.5 we also know that the set $\Sigma^{\mathcal{SR}}(f_{\mu,\nu})$ is dense in Σ, and therefore non-empty. Thus, each pair $(\mu, \nu) \in \mathfrak{M}$ is contained in $\mathcal{D}(P, Q, \delta)$ for some disks P, Q and $\delta > 0$.

Lemma 6.4.8. *Let $A \subset \mathfrak{M}$ be a subset. Suppose every pair $(\mu, \nu) \in \mathfrak{M}$ has a neighbourhood $\mathcal{D} \subset \mathfrak{M}$ such that $\mathcal{D}\backslash A$ is open, dense, and connected in $\mathcal{D}$. Then $\mathfrak{M} \backslash A$ is open, dense, and connected in $\mathfrak{M}$.*

The proof is trivial point-set topology and left to the reader. Thus, toward Theorem 6A it suffices to prove that every $\mathcal{D}\backslash\mathcal{J}$ is connected, where $\mathcal{D}$ ranges over connected components of $\mathcal{D}(P,Q,\delta)$. Henceforward we work on a single such component $\mathcal{D}$. Set $\Sigma^2 = \Sigma \times \Sigma$, $M^2 = M \times M$, and $\mathcal{Y} = \Sigma^2 \times (P \times Q \times \mathcal{D})$. Define the map $\Theta : \mathcal{Y} \to M^2$ by

$$\Theta(r, s, p, q, \mu, \nu) = (f(r), f(s), f(p), f(q))$$

where we abbreviate $f = f_{\mu,\nu}$. As we have noted earlier, the map $(\mu, \nu) \mapsto f_{\mu,\nu}$ is C^k and the evaluation map has the same regularity as f. We deduce Θ is C^m, where $m = \min\{k, n+1\}$.

Let L be the diagonal

$$L = \{((u, v), (u, v)) \in M^2 \times M^2\}.$$

The significance of Θ and L is contained in the fact that

$$\pi^{-1}(\mathcal{J}) \subset \Theta^{-1}(L),$$

where $\pi : \mathcal{Y} \to \mathcal{D}$ is the projection onto the last factor. Indeed, suppose $(\mu, \nu) \in \mathcal{J}$. Then for each pair of points $(p, q) \in P \times Q$ there exists $(r, s) \notin P \times Q$ such that $f(p) = f(r)$ and $f(q) = f(s)$. Thus $\Theta(r, s, p, q, \mu, \nu) \in L$.

Remark 6.4.9. $\pi^{-1}(\mathcal{J})$ also contains the set $L_{P,Q} \times \mathfrak{M}$, where $L_{P,Q}$ is the diagonal of $P \times Q$. This set has codimension 4 and will not play a role in any of our analysis.

Proof of Theorem 6A

Assuming the transversality lemma below, we prove Theorem 6A.

Lemma 6.4.10. *Let (μ, ν) be a pair of metrics with μ not exceptional. Then for all $(r, s, p, q) \in \mathcal{Y}$ such that $\Theta(r, s, p, q, \mu, \nu) \in L$, Θ is a submersion at that point. In particular, Θ is transverse to L at such points.*

Via Lemma 6.4.10, we can shrink $\mathcal{D}(P,Q,\delta)$ so that Θ is transverse to L on all of $\mathcal{Y} \times \mathcal{D}$. Beginning the proof of Theorem 6A, it is enough to show it is dense and connected. If (μ, ν) yields a somewhere injective harmonic map, then by openness there is nothing to do. According to Lemma 6.4.2, we can also dismiss pairs (μ, ν) such that μ is exceptional. Henceforth fix (r, s, p, q, μ, ν) such that μ is non-exceptional, and $f = f_{\mu,\nu}$ is nowhere injective. Define $\theta_{p,q,\mu,\nu} : \Sigma^2 \to M^2 \times M^2$ by

$$\theta_{p,q,\mu,\nu}(r, s) = \Theta(r, s, p, q, \mu, \nu).$$

The map $\theta_{p,q,\mu,\nu}$ is a direct sum of evaluations of $f_{\mu,\nu}$, and hence C^{n+1}. The evaluation map is just Θ, and hence it is C^m, with m as above. Therefore, the association

$$(p,q,\mu,\nu) \mapsto \theta_{p,q,\mu,\nu}$$

defines a C^m representation. L is of codimension $2\dim M \geq 6$, while Σ^2 has dimension 4, so the Parametric Transversality Theorem ensures that for a generic $(p,q,\mu,\nu) \in P \times Q \times \mathcal{D}$ the corresponding map $\theta_{p,q,\mu,\nu}$ is transverse to L. From dimensional considerations, $\theta_{p,q,\mu,\nu}(\Sigma^2)$ is therefore disjoint from L. It follows that (μ,ν) does not belong to $\mathcal{D} \cap \mathcal{J}$. This is true for generic (p,q,μ,ν), and therefore a generic pair (μ,ν) does not live in $\mathcal{D} \cap \mathcal{J}$. This establishes that $\mathcal{J}$ is nowhere dense.

Recalling the projection $\pi : \mathcal{Y} \to \mathcal{D}$, since $\pi(\Theta^{-1}(L))$ may not be a manifold, we cannot conclude immediately from transversality that $\mathcal{D} \cap \mathcal{J}$ is connected. We argue directly, using transversality theorems.

Let $\gamma : [0,1] \to \mathcal{D}$ be a path whose endpoints lie in $\mathcal{D} \setminus \mathcal{J}$. We show that γ can be perturbed, while keeping the endpoints fixed, to lie entirely in $\mathfrak{M} \setminus \mathcal{J}$. First we partition $[0,1]$ into sufficienltly small intervals, each of whose images under γ is contained in a sufficiently small subset of $\mathcal{D}$ that fits into a single (convex) chart in the model Banach space for $\mathfrak{M}$.

Suppose $(\mu_i,\nu_i) \in \mathcal{D}$, $i = 0,1$, are contained in such a chart. We show that one can perturb (μ_1,ν_1) so that the straight line connecting (μ_0,ν_0) and the perturbed (μ_1,ν_1) is contained in $\mathcal{D} \setminus \mathcal{J}$. Let $U \subset P \times Q \times \mathcal{D}$ be a small neighbourhood of (μ_1,ν_1) and consider the C^m map

$$\eta : [0,1] \times \Sigma^2 \times U \to M^2 \times M^2$$

given by

$$\eta(t,r,s,p,q,\mu,\nu) = \Theta(r,s,p,q,t(\mu_0,\nu_0) + (1-t)(\mu,\nu)).$$

η arises as the evaluation map for the C^m representation given by $(p,q,\mu,\nu) \mapsto \rho_{p,q,\mu,\nu}$, where

$$\rho_{p,q,\mu,\nu}(t,r,s) = \eta(t,r,s,p,q,\mu,\nu).$$

The representation does take values in C^m because it may be realized as the composition

$$(t,r,s) \mapsto (t(\mu_0,\nu_0)+(1-t)(\mu,\nu),r,s) \mapsto (f_t,r,s) \mapsto (f_t(p), f_t(q), f_t(r), f_t(s)),$$

where $f_t = f_{t(\mu_0,\nu_0)+(1-t)(\mu,\nu)}$. Using $m \geq 1$ we appeal to the Parametric Transversality Theorem to find that for a generic point $(p,q,\mu,\nu) \in U$, the map $\rho_{p,q,\mu,\nu}$ is transverse to L. Counting dimensions, we see that $\delta([0,1] \times \Sigma^2)$ does not intersect L. This implies that the path

$$\pi(\eta(t,r,s,p,q,\mu,\nu))$$

is contained in $\mathcal{D} \setminus \mathcal{J}$ and connects (μ_0,ν_0) and (μ,ν). Since $\mathcal{D} \setminus \mathcal{J}$ is open, once (μ,ν) is close enough, we can connect it to (μ_1,ν_1) via a straight line in a model chart contained in $\mathcal{D} \setminus \mathcal{J}$.

Returning to the path γ, we do the above procedure over all of the coordinate charts, which perturbs γ to a new path contained entirely in $\mathcal{D} \setminus \mathcal{J}$ and connecting the endpoints. This completes the proof.

Remark 6.4.11. A slightly simpler transversality argument is possible when M has dimension ≥ 4. We leave this for the reader to understand on their own. The proof of the analogue of Lemma 6.4.10 is essentially the same.

6.5 Proof of the transversality lemma

We give the proof of Lemma 6.4.10.

The derivative $d\Theta$

The deriative of Θ is a map $d\Theta : T\mathcal{Y} \to \mathbf{F}^2 \times \mathbf{F}^2$. The tangent space $T\mathcal{Y}$ splits as $T(\Sigma^2 \times P \times Q) \times T\mathfrak{M}$. The restriction $d\Theta : T(\Sigma^2 \times P \times Q) \times \{0\} \to \mathbf{F}^2 \times \mathbf{F}^2$ is given by

$$d\Theta(r,s,p,q,\mu,\nu,0) = df_r \times df_s \times df_p \times df_q,$$

where df_x denotes the derivative of $f = f_{\mu,\nu}$ at x. Since all four points are regular, the image of $d\Theta$ contains every quadruple of vectors $(Z_1,Z_2,Z_3,Z_4) \in \mathbf{F}^2 \times \mathbf{F}^2$ that are tangent to the surface of $f(\Sigma)$ at the corresponding points.

For derivatives in the $\mathfrak{M}$-coordinates, we leave the source metric μ fixed and vary the target metric. Let $\dot{\nu} \in T\mathfrak{M}(M)$. By [EL81, page 35], the section $V \in \Gamma(\mathbf{F})$ defined by

$$V = \frac{d}{dt}|_{t=0} f_{\mu,\nu+t\dot{\nu}}$$

satisfies

$$\mathbf{J}V = \mathcal{G}(\dot{\nu}).$$

Here, $\mathcal{G}(\dot{\nu})$ is the derivative of the tension field in the $\dot{\nu}$-direction:

$$\mathcal{G}(\dot{\nu}) = \frac{d}{dt}|_{t=0}\tau(\mu, \nu + t\dot{\nu}, f_{\mu,\nu}).$$

Accordingly, V is called a harmonic variation. It follows that

$$d\Theta(0,0,0,0,0,\dot{\nu}) = (V(r), V(s), V(p), V(q)).$$

Suppose $\Theta(r, s, p, q, \mu, \nu) \in L$. To simplify notation, we rename the points as $z_1 = p, z_2 = r, w_1 = q, w_2 = s$. The proof of Lemma 6.4.10 is another contradiction argument. Suppose the lemma is incorrect. Then, there are four vectors $Z_i \in \mathbf{F}_{z_i}$, $W_i \in \mathbf{F}_{w_i}$, not all of them zero, such that

- Z_i and W_i are either zero or normal to the surface $f(\Sigma)$ at the points $f(z_i)$ and $f(w_i)$ respectively and
- for every V such that $\mathbf{J}V = \mathcal{G}(\dot{\nu})$, the following holds:

$$\sum_{i=1}^{2}(\langle V(z_i), Z_i\rangle + \langle V(w_i), W_i\rangle) = 0. \tag{6.8}$$

We now invoke reproducing formulas for the zeroth derivative. We showed the existence of reproducing kernels in Section 6.3. Adding up the four zeroth order reproducing kernels associated to the points z_i, w_i, we find a section $X : \Sigma\backslash\{z_1, z_2, w_1, w_2\} \longrightarrow \mathbf{F}$ with maximum regularity and such that

$$\sum_{i=1}^{2}(\langle W(z_i), Z_i\rangle + \langle W(w_i), W_i\rangle) = \int_\Sigma \langle \mathbf{J}W, X\rangle dA$$

for all $W \in \Gamma(\mathbf{F})$. We also record here that $\mathbf{J}X(p) = 0$ for every $p \neq z_1, z_2, w_1, w_2$ and $X \in L^p(\mathbf{F})$ for every $p \geq 1$. X is not identically equal to zero as we can certainly find sections $W \in \Gamma(\mathbf{F})$ such that the left-hand side above is not zero. On the other hand, from (6.8) we conclude that

$$\int_\Sigma \langle \mathbf{J}V, X\rangle dA = 0$$

for every harmonic variation V.

Stepping back for a moment, if $\dot{\nu}$ has support near $f(z_1)$, then the associated $\mathbf{J}V$ is supported near all preimages of $f(z_1)$. The kernel X may have singularities at z_1 and z_2, while X is smooth at the other preimages of $f(z_1)$. The tangent planes

$df(T_{z_1}\Sigma)$ and $df(T_{z_2}\Sigma)$ are either tangential or span a k-plane for $k = 3$ or 4, and we find it convenient to treat the cases separately. In both cases, it is possible to choose $\dot{\nu}$ so that $\mathcal{G}(\dot{\nu})$ is negligible at z_2 but not so at z_1. In the tangential case, we use the argument from [Mar18, Section 7]. This is where the super-regular condition comes into play (and this is the only place it does). In this way, we can eliminate the singularity of X at z_1. Repeating the procedure, but interchanging the roles of z_1 and z_2, we're able to show that X is a global Jacobi field, which means $X \equiv 0$.

The time derivative of the tension field

We compute $\langle \mathbf{J} V, X\rangle$ in coordinates for a general variation $\dot{\nu}$.

We let (x_1, x_2) and $(u_1, \dots, u_n)$ denote local coordinates near $z_1 \in \Omega$ and $f(z_1) \in M$ such that $z_1 = (0,0)$, $f(z_1) = (0, \dots, 0)$. Near z_1, the reproducing kernel X can be expressed as a linear combination of the sections $\frac{\partial}{\partial f_j} = f^* \frac{\partial}{\partial u_j}$, $j = 1, \dots, n$. We let X^j denote the real valued functions on Ω such that

$$X_k = \sum_{j=1}^n X^j \frac{\partial}{\partial f_j}.$$

In local coordinates on Σ (not necessarily holomorphic), the tension field τ is given by

$$\tau^\gamma = \tau^\gamma(f, \mu, \nu) = \mu^{ij}\Big(\frac{\partial^2 f^\gamma}{\partial x_i \partial x_j} - {}^{\mu}\Gamma^k_{ij}\frac{\partial f^\gamma}{\partial x_k} + {}^{\nu}\Gamma^\gamma_{\alpha\beta}(f)\frac{\partial f^\alpha}{\partial x_i}\frac{\partial f^\beta}{\partial x_j}\Big),$$

where $\gamma = 1, \dots, n$, and we're using the Einstein summation convention. Here μ^{ij} are the components of the inverse of the metric tensor μ. Let $\dot{\nu}$ be a variation of ν and set $\nu_t = \nu + t\dot{\nu}$. Recall we have defined $\mathcal{G}(\dot{\nu}) = \frac{\partial}{\partial t}\tau(f_{\mu,\nu}, \mu, \nu_t)|_{t=0}$. Since $\tau(f, \mu, \nu) = 0$, we see

$$\langle \mathcal{G}(\dot{\nu}), X\rangle = \frac{d}{dt}|_{t=0}\langle \tau(f, \mu, \nu_t), X\rangle.$$

The only term that does not die upon taking the derivative is the term involving ${}^{\nu}\Gamma^\gamma_{\alpha\beta}(f)$. Thus,

$$\langle \mathcal{G}(\dot{\nu}), X\rangle = \frac{d}{dt}|_{t=0}\sum_{\alpha,\beta}\nu_{\alpha\beta}\mu^{ij}\,{}^{\nu_t}\Gamma^\alpha_{\gamma\delta}(f)\frac{\partial f^\gamma}{\partial x_i}\frac{\partial f^\delta}{\partial x_j}X^\beta. \tag{6.9}$$

Set

$${}^{\nu_t}\Gamma_{\gamma,\alpha\beta} = \frac{1}{2}(\nu^t_{\alpha\gamma,\beta} + \nu^t_{\gamma\beta,\alpha} - \nu^t_{\alpha\beta,\gamma}),$$

where $\nu^t_{\alpha\beta,\delta} = \frac{\partial \nu^t_{\alpha\beta}}{\partial u_\delta}$ and $\nu_t^{\gamma\delta}$ denote inverse components of ν^t. Under this notation, the Christoffel symbols are computed by the well-known formula

$${}^{\nu_t}\Gamma^\gamma_{\alpha\beta} = \sum_\delta \nu_t^{\gamma\delta}\cdot {}^{\nu_t}\Gamma_{\delta,\alpha\beta}.$$

Inserting back into (6.9) yields

$$\langle \mathcal{G}(\dot{\nu}), X\rangle = \frac{d}{dt}|_{t=0} \sum_{\alpha} \mu^{ij\nu_t}\Gamma_{\gamma\delta,\alpha}(f)\frac{\partial f^\gamma}{\partial x_i}\frac{\partial f^\delta}{\partial x_j}X^\alpha = \sum_{\gamma} \mu^{ij}\dot{\Gamma}_{\alpha\beta,\gamma}\frac{\partial f^\alpha}{\partial x_i}\frac{\partial f^\beta}{\partial x_j}X^\gamma. \tag{6.10}$$

Here we are using the notation

$$\dot{\Gamma}^{\gamma}_{\alpha\beta} = \lim_{t\to 0}\frac{\partial^{\nu_t}\Gamma^{\gamma}_{\alpha\beta}}{\partial t}.$$

We also record that

$$\dot{\Gamma}_{\gamma,\alpha\beta} = \frac{1}{2}(\dot{\nu}_{\alpha\gamma,\beta} + \dot{\nu}_{\gamma\beta,\alpha} - \dot{\nu}_{\alpha\beta,\gamma}).$$

Tangential harmonic disks

Let Ω be a small neighbourhood of z_1 such that $f : \Omega \to M$ is an embedding. We let (x_1, x_2) be conformal coordinates near z_1 and $(u_1, \ldots, u_n)$ normal coordinates centered at $f(z_1) \in M$ such that

- $f(z_1) = (0, \ldots, 0)$,
- the (regular) surface $f(\Omega)$ is tangent to the plane $P = \{u_3 = \cdots = u_n = 0\}$ at $f(z_1)$, and $f(x_i) = u_i$ for $i = 1, 2$, and
- $\nu_{jk} = \nu_{jk} = 0$ and $\nu_{jj} = 1$ for $k = 1, 2$ and $j = 3, \ldots, n$ when restricted to P at $f(z_1)$.

Note that, as observed in Section 6.4, the set $f^{-1}(f(z_1))$ is finite. Set

$$f^{-1}(f(z_1)) = \{z_1, z_2, \ldots, z_m\}.$$

For $\epsilon \in (0, 1)$ small enough, we let $D(\epsilon)$ denote the disk of radius ϵ in the plane P, and let D_ϵ be the ball of radius ϵ in the $(u_1, \ldots, u_n)$-coordinates centered at 0. Since $z_k \in \Sigma^{reg}(f)$, we may choose ϵ so that

$$f^{-1}(D_\epsilon) = \bigcup_{k=1}^{m} \Omega_k,$$

where $\Omega_k = \Omega_k(\epsilon)$ is the corresponding neighbourhood of z_k. If we choose a variation $\dot{\nu}$ with support in D_ϵ, then the induced variation of the pullback metric $f^*\nu$ is supported in $f^{-1}(D_\epsilon)$. If $\mathbf{J}V = \mathcal{G}(\dot{\nu})$, we will see that this implies $\mathbf{J}V$ is supported there as well, and we obtain

$$\int_\Sigma \langle \mathbf{J}V, X\rangle dA = \sum_{k=1}^{m}\int_{\Omega_k}\langle \mathbf{J}V, X\rangle dA = 0. \tag{6.11}$$

Our proof of transversality of Θ involves analyzing each integral in the sum above. We split into cases: (i) the harmonic surfaces $f(\Omega_1)$ and $f(\Omega_2)$ are tangential at $f(z_1)$ and (ii) they are not tangential. In each case, we pick a different variation of the target metric to find our contradiction.

We first treat case (i). For $\epsilon > 0$ small enough,

$$1 \lesssim |df| \lesssim 1$$

on each Ω_k. Here $|\cdot|$ is the operator norm. Since the surface $f(\Omega_k)$ is regular and proper, it follows that

$$\epsilon^2 \lesssim \int_{\Omega_k} dA \lesssim \epsilon^2 \tag{6.12}$$

for all k.

We now specify our variation. We let $\varphi_{\alpha\beta} = \varphi_{\beta\alpha}$ denote a set of real numbers such that $\varphi_{\alpha\beta} = 0$ if at least one α, β is greater than two. Let χ be a non-negative function of $(u_1, \dots, u_n)$ with support in D_2, equal to 1 on $D_{1/2}$, and such that it has total integral 1 with respect to the induced Euclidean area form on P. These conditions guarantee that, restricted to this plane, $\chi^\epsilon(u) = \epsilon^{-2}\chi(u/\epsilon)$ converges in the sense of distributions to the Dirac delta function as $\epsilon \to 0$. χ^ϵ has compact support in $D(2\epsilon)$, and the product $\chi^\epsilon \varphi_{\alpha\beta}$ is equal to $\varphi_{\alpha\beta}$ on $D(\epsilon/2)$. Define $\dot{\nu} = \dot{\nu}(\epsilon)$ by

$$\dot{\nu}_{\alpha\beta}(u) = \sum_{j=3}^{n} -2u_j \chi^\epsilon(u)\varphi_{\alpha\beta}.$$

We suppress the ϵ from our notation wherever possible. Referring back to (6.10), we are interested in the variation of $\Gamma_{\alpha\beta,\gamma}$. For $\gamma \geq 3$,

$$\dot{\Gamma}_{\alpha\beta,\gamma} = -\frac{1}{2}\dot{\nu}_{\alpha\beta,\gamma} = \varphi_{\alpha\beta}\chi^\epsilon(u) + u_\gamma \varphi_{\alpha\beta}\chi^\epsilon_\gamma(u)$$

and hence, on $(-\epsilon/2, \epsilon/2)$, this is

$$\dot{\Gamma}_{\alpha\beta,\gamma} = \dot{\nu}_{\alpha\beta,\gamma} = \varphi_{\alpha\beta}\chi^\epsilon(u).$$

For $\gamma = 1, 2$,

$$|\dot{\Gamma}_{\alpha\beta,\gamma}| \lesssim \max_{\alpha\beta,\delta} |\dot{\nu}_{\alpha\beta,\delta}| \lesssim \epsilon|\nabla\chi^\epsilon| \lesssim \epsilon^{-2}.$$

In any case, we always have a $O(\epsilon^{-2})$ bound.

The local coordinates (x_1^1, x_2^1) near z_1 satisfy

$$\frac{\partial f^\alpha}{\partial x_i^1}(z) = \delta_{i\alpha}.$$

Since $f(\Omega_2)$ is tangent to $\{u_3 = 0\}$ at z_2, we can choose coordinates (x_1^2, x_2^2) near z_2 such that

$$\frac{\partial f^\alpha}{\partial x_i^2}(z) = \delta_{i\alpha} + O(\epsilon). \tag{6.13}$$

Note that, by our restrictions on μ, μ is no longer conformal in these coordinates.

Remark 6.5.1. When the two harmonic disks are equal, the $O(\epsilon)$ term is identically zero.

Inserting these expressions into (6.9) gives, near z_1,

$$\langle \mathcal{G}(\dot{v}), X\rangle = \sum_\gamma \mu^{ij}\dot{\Gamma}_{ij,\gamma}X^\gamma \tag{6.14}$$

and near z_2,

$$\langle \mathcal{G}(\dot{v}), X\rangle = \sum_\gamma \mu^{ij}\dot{\Gamma}_{\alpha\beta,\gamma}(\delta_{i\alpha}\delta_{j\beta} + O(\epsilon))X^\gamma = \sum_\gamma \mu^{ij}\dot{\Gamma}_{ij,\gamma}X^\gamma + O(\epsilon)\sum_{\alpha,\beta,\gamma}\dot{\Gamma}_{\alpha\beta,\gamma}X^\gamma. \tag{6.15}$$

Incompatible asymptotics

The reproducing kernel X is regular near each point z_k when $k > 2$. Trivially, $|\mathbf{J}V| \lesssim 1$ near z_k. Recalling (6.12), we deduce

$$\Big|\int_{\Omega_k}\langle \mathbf{J}V, X\rangle dA\Big| \lesssim \epsilon^{-2}$$

for $k > 1, 2$. For $k = 1, 2$, it may be the case that X has a singularity near z_1 or z_2 (or both). We computed this singularity in Proposition 6.3.7, the result being that in a trivialization near z_1,

$$X(z) = \frac{1}{2\pi}\Big(\log\frac{1}{|z|}\Big)Z_k + B_1(z),$$

where $B_1(z)$ is a $C^{0,\alpha}$ local section of $\mathbf{F}$ near z_1, and $Z_k \in \mathbf{F}_{z_k}$ is the vector normal to the surface $f(\Sigma)$ at $f(z_1)$, defined above. Our coordinate is not conformal around z_2.

Here we are considering the zero vector to be normal. Recall we are assuming that for $k = 1, 2$, the patches $f(\Omega_k)$ are tangent to the plane P at $f(z_k)$. Z_k is normal to P because $v_{1\gamma} = v_{2\gamma} = 0$ for $\gamma \geq 3$. Incorporating these asymptotics into (6.14) and (6.15), we isolate that at z_1,

$$\langle \mathcal{G}(\dot{v}), X\rangle = \sum_{\gamma\geq 3}\mu^{ij}\dot{\Gamma}_{ij,\gamma}X^\gamma + \sum_{\delta=1,2}\mu^{ij}\dot{\Gamma}_{ij,\delta}X^\delta = \sum_{\gamma\geq 3}\mu^{ij}\dot{\Gamma}_{ij,\gamma}X^\gamma + O(\epsilon^{-2}) \tag{6.16}$$

and at z_2,

$$\langle \mathcal{G}(\dot{v}), X\rangle = \sum_{\gamma\geq 3} \mu^{ij}\dot{\Gamma}_{ij,\gamma} X^\gamma (1 + O(\epsilon)) + O(\epsilon^{-2}).$$

We now use the fact that the restrictions of the metric μ at the points z_1 and z_2 are not conformal to each other via f. By the choice of local coordinates, this means the matrix $\mu^{ij}(z_1)$ is not a multiple of the matrix $\mu^{ij}(z_2)$, where both matrices are found by trivializing the pullback bundle $\mathbf{F}$ over $f^{-1}(D_\epsilon)$ using a trivialization of D_ϵ. Furthermore, the two spaces of 2×2 matrices orthogonal to $\mu^{ij}(z_1)$ and $\mu^{ij}(z_2)$ respectively (with respect to the Frobenius inner product) do not coincide. Thus, we can choose φ_{ij} uniformly bounded above and such that

$$\sum_{i,j=1}^{2} \mu^{ij}(z_1)\varphi_{ij}(f(z_1)) = 1$$

and

$$\sum_{i,j=1}^{2} \mu^{ij}(z_2)\varphi_{ij}(f(z_2)) = 0.$$

Taylor expanding (6.16), we see that near z_1,

$$\begin{aligned}\langle \mathcal{G}(\dot{v}), X\rangle &= \sum_{\gamma\geq 3} \epsilon^{-2}\chi(z/\epsilon)\Big(\frac{1}{2\pi}\Big(\log\frac{1}{|x|}\Big)Z_1^\gamma(x) + B^\gamma(x)\Big) + O(\epsilon^{-2}) \\ &= \sum_{\gamma\geq 3} \epsilon^{-2}\log|x|^{-1}\frac{\chi(z/\epsilon)}{2\pi} Z_1^\gamma(0) + O(\epsilon^{-2}).\end{aligned}$$

As for z_2,

$$\sum_{i,j=1}^{2} \mu^{ij}(z_2)\varphi_{ij} \lesssim \epsilon$$

follows by Taylor expansion, and therefore

$$\begin{aligned}\langle \mathcal{G}(\dot{v}), X\rangle &= \sum_{\gamma\geq 3} \mu^{ij}(\varphi_{ij}\chi^\epsilon(u) + u_\gamma\varphi_{ij}\chi^\epsilon_\gamma(u))X^\gamma(1 + O(\epsilon)) + O(\epsilon^{-2}) \\ &\lesssim \sum_{\gamma\geq 3} \epsilon(\chi^\epsilon(u) + u_\gamma\chi^\epsilon_\gamma(u))(1 + O(\epsilon))\Big(\log\frac{1}{|x|}\Big)Z_2^\gamma(x) + B^\gamma(x)\Big) + O(\epsilon^{-2}) \\ &\lesssim \sum_{\gamma\geq 3} \epsilon^{-1}\log\epsilon^{-1}|Z_2^\gamma(0)| + O(\epsilon^{-2}).\end{aligned}$$

Taking integrals yields

$$\begin{aligned}&\int_{\Omega_1} \langle \mathbf{J}V, X\rangle dA = \frac{\epsilon^{-2}}{2\pi}\sum_{\gamma\geq 3}\int_{\Omega_1}\Big(\log\frac{1}{|x|}\Big)\chi(z/\epsilon)Z_1^\gamma(0)dA(x_1^1, x_2^1) + O(1) \\ &\int_{\Omega_2} \langle \mathbf{J}V, X\rangle dA \lesssim \epsilon\log\epsilon^{-1}|Z_2(0)| + O(1) \lesssim 1,\end{aligned}$$

and replacing back into (6.11) gives

$$\int_\Sigma \langle \mathbf{J}V, X\rangle dA = \frac{\epsilon^{-2}}{2\pi}\sum_{\gamma\geq 3}\int_{\Omega_1}\Big(\log\frac{1}{|x|}\Big)\chi(z/\epsilon)Z_1^\gamma(0)dA(x_1^1,x_2^1) + O(1).$$

Our standing assumption is that for all $\epsilon > 0$, the left-hand side is equal to 0. Therefore,

$$\begin{aligned}&\sum_{\gamma\geq 3}|Z_1^\gamma(0)|\frac{1}{2\pi}\int_{\Omega_1}\Big(\log\frac{1}{|x|}\Big)\chi(z/\epsilon)dA(x_1^1,x_2^1)\\ &=\Big|\sum_{\gamma\geq 3}\int_{\Omega_1}\frac{1}{2\pi}\int_{\Omega_1}\Big(\log\frac{1}{|x|}\Big)\chi(z/\epsilon)Z_1^\gamma(0)dA(x_1^1,x_2^1)\Big| \lesssim 1.\end{aligned}$$

In coordinates, Ω_1 is contains a ball of radius $\epsilon/2$ with respect to the Euclidean metric, and in such a ball $\chi(z/\epsilon) = 1$. Thus,

$$\frac{\epsilon^{-2}}{2\pi}\int_{\Omega_1}\Big(\chi(z/\epsilon)\log\frac{1}{|x|}\Big)dA(x_1^1,x_2^1) \gtrsim \epsilon^{-2}\int_{\Omega_1}\Big(\log\frac{1}{|x|}\Big)dx_1^1dx_2^1 \gtrsim \log(\epsilon^{-1}).$$

If there exists $\gamma \geq 3$ such that $Z_1^\gamma(0) \neq 0$, this implies

$$\log\epsilon^{-1} \lesssim 1,$$

which is nonsensical. This forces $Z_1^\gamma(0) = 0$ for all $\gamma \geq 3$. Furthermore, recalling that Z_1 is normal to $f(\Omega_1)$ at z_1, we must have that $Z_1(z_1) = 0$ identically. This proves the following lemma.

Lemma 6.5.2. *Suppose $f(\Omega_1)$ and $f(\Omega_2)$ are tangential at $f(z_1)$. Then X extends smoothly over z_1.*

Non-tangential harmonic disks

We have essentially proved transversality Θ, if we assume the images of the harmonic map are tangential at z_1, z_2, and at w_1, w_2. In this subsection, we consider other intersections. Namely, we prove the following.

Lemma 6.5.3. *Suppose $f(\Omega_1)$ and $f(\Omega_2)$ are not tangential at $f(z_1)$. Then X extends smoothly over z_1.*

Equipped with this lemma, we can prove transversality with ease.

Proof of transversality of Θ. Assume Θ is not transverse, so that we have the section X as in the work above. Applying Lemma 6.5.2 or Lemma 6.5.3, depending on the circumstance, we see X extends smoothly over the point z_1. Repeating this procedure with Z_2, W_1, and W_2 taking the role of Z_1, we can show it extends smoothly over those points as well. However, that means X extends to a global Jacobi field, which can only occur if $X \equiv 0$. This is a contradiction. □

Moving toward the proof of Lemma 6.5.3, the proof of the tangential case does not immediately adapt because (6.13) does not hold, and the super-regular condition can not be used effectively. To accommodate, we choose our variation differently. Instead of picking one supported in the ball D_ϵ, we set $C_\epsilon = D(\epsilon) \times \{|u_j| < \epsilon^2 : j = 3, \ldots, n\}$ and use

$$B_\epsilon = D_\epsilon \cap C_\epsilon.$$

The three-dimensional picture of this is the intersection of a ball with a fat cylinder. Similar to before, let Ω_k denote the connected components of $f^{-1}(B_\epsilon)$. Since B_ϵ is contained in D_ϵ, regularity gives

$$\int_{\Omega_k} dA \lesssim \int_{f^{-1}(D_\epsilon)} dA \lesssim \epsilon^2.$$

Following our previous approach, for $\alpha, \beta = 1, 2$, we continue to use functions $\varphi_{\alpha\beta} = \varphi_{\beta\alpha}$ defined on the patch $f(\Omega) \cap P$. Let χ^ϵ be exactly as before. Take ω_1 to be a smooth function with support in $P \cap \{u_1^2 + u_2^2 < 4\}$ and such that $\omega_1 = 1$ on $P \cap \{u_1^2 + u_2^2 < 1\}$, and let $\omega_2 : \mathbb{R} \to \mathbb{R}$ be a smooth function that is 1 on $(-1, 1)$ and 0 off $(-2, 2)$. Assuming $\epsilon < 1$, set

$$\omega^\epsilon(u) = \omega_1(\epsilon^{-1}u_1, \epsilon^{-1}u_2, 0, \ldots, 0)\omega_2(0, 0, \epsilon^{-2}u_3, \ldots, u_n).$$

$\chi^\epsilon\omega^\epsilon$ has compact support in $B_{2\epsilon}$, and $\chi^\epsilon\omega^\epsilon\varphi_{\alpha\beta} = \varphi_{\alpha\beta}$ in B_ϵ. Define $\dot{\nu} = \dot{\nu}(\epsilon)$ by

$$\dot{\nu}_{\alpha\beta}(u_1, u_2, u_3) = -2\sum_{j=3}^{n} u_j\chi^\epsilon(u_1, u_2, u_3)\omega^\epsilon(u_1, u_2, u_3)\varphi_{\alpha\beta}(u_1, u_2, 0)$$

and $\dot{\nu}_{\alpha\beta} = 0$ for other α, β. One can slightly adjust our previous computations to get

$$\max_{\alpha,\beta} |\dot{\Gamma}^\gamma_{\alpha\beta}| \lesssim \epsilon^{-2}.$$

Reusing our previous notation, one can now show

$$\int_\Sigma \langle \mathbf{J}V, X\rangle dA = \sum_{k=1}^{2} \int_{\Omega_k} \nu_{\alpha\beta}(f)\mathcal{G}^\alpha(\dot{\nu})X_k^\beta dA + O(\epsilon^2).$$

We aim to find asymptotics for both integrals on the right-hand side above. The key observation in bounding the integral over Ω_2 is that, since $f(\Omega_2)$ is not wholly contained in B_ϵ, we have a stronger area estimate.

Lemma 6.5.4. *The area of Ω_2 satisfies*

$$Area(\Omega_2) \lesssim \epsilon^3$$

as $\epsilon \to 0$.

Proof. Since $|df|$ is uniformly bounded above and below on Ω_2, it suffices to prove the same asymptotic for the area of $f(\Omega_2) \subset B_\epsilon$. Let $(u_1, \ldots, u_n)$, (x_1^2, x_2^2) be the coordinates from above, and let Q be the embedding of the tangent plane $df(T_{z_2}\Sigma)$ inside our coordinate patch. Our metrics are locally comparable to Euclidean metrics, and hence

$$\text{Area}(f(\Omega_2)) \lesssim \int_{\Omega_2} J(f)(x_1^2, x_2^2) dx_1^2 d_x^2 \lesssim \int_{\Omega_2} J(f)(0) dx_1^2 d_x^2 + O(\epsilon^3),$$

where $J(f)$ is the Jacobian determinant for f. Let $\tilde{\Omega}_2$ be the relevant component of $f^{-1}(C_\epsilon)$. Then

$$\int_{\Omega_2} J(f)(0) dx_1^2 dx_2^2 \leq \int_{\tilde{\Omega}_2} J(f)(0) dx_1^2 dx_2^2.$$

As f is an immersion near z_2, from multivariable calculus we have

$$\int_{\tilde{\Omega}_2} J(f)(0) dx_1^2 dx_2^2 = \int_{Q \cap C_\epsilon} dS, \tag{6.17}$$

where dS is the Euclidean area form on the parametrized surface $Q \cap C_\epsilon \subset \mathbb{R}^n$.

To compute, if P and Q span a four-dimensional subspace, then our job is very easy: the plane Q only intersects P at $f(z_2)$, and is hence contained in a ball of radius ϵ^2. So the area integral is on the order of ϵ^{2n}. The less trivial case is when P and Q intersect transversely inside a copy of $\mathbb{R}^3$. That is, Q intersects P in a line, making some acute angle $\alpha > 0$ with P. The family of planes making such an angle admits an S^1 action by rotations around the normal axis (which we now assume is the u_3-axis), which preserves the area of intersections with the cylinder. Hence, we can replace Q with any plane that makes the same angle α. A convenient choice is

$$Q = \{(u_1, u_2, u_3) : u_1 \sin\alpha + u_3 \cos\alpha = 0\}.$$

If not already the case, shrink ϵ so that $\epsilon < \tan\alpha$. We view Q as the parametrized surface specified by

$$F(u_1, u_2) = u_3 = -u_1 \tan\alpha,$$

subject to the constraints $u_1^2 + u_2^2 < \epsilon^2$, $|u_3| < \epsilon^2$. Set $F_i = F_{u_i}$. We compute

$$\int_{Q \cap C_\epsilon} dS = \int_{Q \cap C_\epsilon} \sqrt{F_1^2 + F_2^2 + 1} dS = 2 \int_{-\frac{\epsilon^2}{\tan\alpha}}^{\frac{\epsilon^2}{\tan\alpha}} \int_0^{\sqrt{\epsilon^2 - u_1^2}} \sqrt{\tan^2\alpha + 1} du_2 du_1 \lesssim \epsilon^3.$$

Thus, in both cases, inputting the estimates into (6.17) gives the desired bound. □

With this lemma in hand,

$$X_2(z) = C\Big(\log\frac{1}{|z|}\Big)Z_2 + B(z)$$

with Z_2 a vector normal to $f(\Omega_2)$ at z_2 (and thus not normal to P) and $B \in C^{0,\alpha}(\mathbf{F}|_{\Omega_2})$. Independent of the choice of function $\varphi_{\alpha\beta}$, we estimate

$$\begin{aligned}\langle \mathcal{G}(\dot{\nu}), X\rangle &= \sum_\gamma \mu^{ij}\dot{\Gamma}_{\alpha\beta,\gamma}\frac{\partial f^\alpha}{\partial x_i}\frac{\partial f^\beta}{\partial x_j}X^\gamma \lesssim \log\frac{1}{|x|}\max_{\alpha,\beta,\gamma}|\dot{\Gamma}_{\alpha\beta,\gamma}| \\ &\lesssim \Big|\int_{\Omega_2}\nu_{\alpha\beta}(f)\mathcal{G}^\alpha(\dot{\nu})X_2^\beta dA\Big| \lesssim \max_\beta\int_{\Omega_2}|X_2^\beta|dA \lesssim \int_{\Omega_2}\log\frac{1}{|x|}dA + O(\epsilon^3).\end{aligned}$$

We bound the integral on the right:

$$\begin{aligned}\int_{\Omega_2}\log\frac{1}{|x|}dA &= \int_{\Omega_2\setminus(\Omega\cap B(0,\epsilon^{3/2}))}\log\frac{1}{|x|}dA + \int_{\Omega\cap B(0,\epsilon^{3/2})}\log\frac{1}{|x|}dA \\ &\le \log\epsilon^{-3/2}\text{Area}(\Omega_2) + \int_{B(0,\epsilon^{3/2})}\log\frac{1}{|x|}dA \lesssim \epsilon^3\log\epsilon^{-1}.\end{aligned}$$

Returning to our original integral, we obtain

$$0 = \int_\Sigma\langle \mathbf{J}V, X\rangle dA = \int_{\Omega_1}\langle \mathbf{J}V, X\rangle dA + O(\epsilon^2).$$

For Ω_1, the area estimate

$$\epsilon^2 \lesssim \int_{\Omega_1} dA \lesssim \epsilon^2 \tag{6.18}$$

is obvious. We choose $\varphi_{\alpha\beta}$ exactly as in the tangential case, and if X does not extend smoothly over z_1, then using (6.18) returns

$$\int_{\Omega_1}\langle \mathbf{J}V, X\rangle dA \gtrsim \log\epsilon^{-1}$$

and produces the same contradiction as in the tangential case. This completes the proof of Lemma 6.5.3. As discussed above, this concludes our proof of transversality of Θ.

6.6 Immersions and embeddings

Preparing the arguments

Here we set up transversality arguments for Theorem 6B and Theorem 6C. We assume that M is parallelizable, the general case being a slight modification (because transversality is a local property). Accordingly, we choose an isomorphism $\sigma : TM^{\mathbb{C}} \to M \times \mathbb{C}^n$ with projection map from

$$TM^{\mathbb{C}} \to \mathbb{C}^n$$

that restricts to a family of isomorphisms $\sigma_p : TM_p^{\mathbb{C}} \to \mathbb{C}^n$, isometric with respect to the inner product induced by the metric on $TM_p^{\mathbb{C}}$ and the standard inner product on $\mathbb{C}^n$.

Let $\tilde{\Sigma}$ denote the universal cover of Σ. The metric μ on Σ lifts to a metric on the universal cover $\tilde{\Sigma}$ that we still denote by μ. Likewise, the harmonic map $f_{\mu,\nu}$ lifts to a map $f_{\mu,\nu} : (\tilde{\Sigma}, \mu) \to (M, \nu)$, and we do not distinguish our notation.

The Riemannian manifold $(\tilde{\Sigma}, \mu)$ identifies isometrically with $(\mathbb{D}, \sigma)$, the complex unit disk endowed with its hyperbolic metric. We further identify the Riemann surface Σ_μ in the conformal class of (Σ, μ) with $\mathbb{D}/\Gamma_\mu$, where Γ_μ is a smoothly varying family of Fuchsian groups acting on $\mathbb{D}$. Let $z \in \mathbb{D}$ denote the complex parameter. This provides us with a canonical complex parameter $z_\mu = z$ on $\tilde{\Sigma}$ that depends only on μ.

Unless stated otherwise, the dimension (codimension) of some object in a category (vector space, manifold, etc.) refers to the real dimension (codimension). To prove Theorem 6B, consider the subset $\mathcal{I} \subset \mathbb{C}^n$ defined by

$$\mathcal{I} = \{A \in \mathbb{C}^n : \text{rank} A < 2\}.$$

Here, $\text{Rank}(A)$ denotes the dimension of the vector space spanned by $\text{Re}(A)$ and $\text{Im}(A)$. $\mathcal{I}$ is not a submanifold, but it splits as a union of two submanifolds of $\mathbb{C}^n$: $\mathcal{I} = \mathcal{L}_0 \cup \mathcal{L}_1$, where

$$\mathcal{L}_0 = \{0\} \subset \mathbb{C}^n, \ \mathcal{L}_1 = \{A \in \mathbb{C}^n : \text{rank} A = 1\} \subset \mathbb{C}^n.$$

We define

$$\Psi : \tilde{\Sigma} \times \mathfrak{M} \to \mathbb{C}^n$$

by

$$(p, \mu, \nu) \mapsto \sigma(f_z(p)).$$

The point is that $f_{\mu,\nu}$ is an immersion as long as $\Psi(p, \mu, \nu) \notin \mathcal{I}$ for all p. $\mathcal{L}_0$ has codimension $2n$ and $\mathcal{L}_1$ has codimension $n-1$ in $\mathbb{C}^n$, so if Ψ is transverse to both submanifolds, then $\Psi^{-1}(\mathcal{I})$ is contained in a codimension $n-1$ submanifold. Let $\pi : \tilde{\Sigma} \times \mathfrak{M} \to \mathfrak{M}$ be the projection map. $\pi(\Psi^{-1}(\mathcal{I}))$ is contained in a submanifold of codimension at least $(n-1) - 2 = n - 3$, so openness and density should hold for $n \geq 4$. Heuristically speaking, we should have connectedness for $n \geq 5$, although because the image may not be a manifold, this last point does not follow immediately. Thus, once we formalize this argument, the content of Theorem 6B is that Ψ is transverse at all points in the preimage of $\mathcal{I}$.

Toward transversality of Ψ, we compute the derivative of Ψ at a point (p, μ, ν) in the direction of a variation of the target metric $(0, 0, \dot{\nu})$. As before, the infinitesimal variation of the maps $f_{\mu,\nu+t\dot{\nu}}$ is a section $V \in \Gamma(\mathbf{F})$ satisfying $\mathbf{J}V = \mathcal{G}(\dot{\nu})$. Choosing our coordinate so that $V(p) = 0$, the vector $d\Psi(0, 0, 0, \dot{\nu})$ can be identified with the vertical lift of the associated vector in $\mathbb{C}^n$ under the identification of the tangent space at 0. In such a coordinate, the derivative becomes

$$d\Psi(0, 0, \dot{\nu}) = \sigma_p(\nabla_z V(p)).$$

Thus, we have the following.

Lemma 6.6.1. *Fix a point $(p, \mu, \nu) \in \tilde{\Sigma} \times \mathfrak{M}$. Suppose that for every $W \in \mathbf{E}_p$, there exists a variation $\dot{\nu} \in T_\nu \mathfrak{M}(M)$ such that if $V \in \Gamma(\mathbf{F})$ is the section satisfying $\mathbf{J}V = \mathcal{G}(\dot{\nu})$, then $V(p) = 0$ and*

$$\sigma_p(\nabla_z V(p)) = W(p).$$

Then Ψ is transverse to $\mathcal{L}_0$ and $\mathcal{L}_1$ at (p, μ, ν).

Granting the following, we prove Theorem 6B.

Lemma 6.6.2. *Suppose $f_{\mu,\nu}$ is somewhere injective and has isolated singularities. Then the hypothesis of the lemma above is satisfied. That is, Ψ is transverse to $\mathcal{L}_0$ and $\mathcal{L}_1$ at (p, μ, ν).*

Proof of Theorem 6B. This is similar to the proof of Theorem 6A, so we don't dwell on the details. Transversality is an open property, so we can fix a neighbourhood U around (μ, ν) in which Ψ is transverse. We let U^I denote the subset of $(\mu, \nu) \in U$ corresponding to harmonic immersions. Observe $U^I = U \cap \mathfrak{M} \setminus (\pi(\Psi^{-1}(\mathcal{I})))$. The goal is to show this is open, dense, and connected.

$\mathcal{I}$ is clearly closed, from which the openness result is immediate. The Transversality Theorem for Banach Manifolds implies that $\Psi^{-1}(\mathcal{L}_0)$ has codimension $2n$ and $\Psi^{-1}(\mathcal{L}_1)$ has codimension $n - 1$, and from our comments above, we obtain density when $n = \dim M \geq 4$.

Assuming $n = \dim M \geq 5$, we prove connectedness directly. Let $\gamma : [0, 1] \to U^I$ be a path with endpoints in U^I. We show that γ can be perturbed (while keeping the endpoints fixed) to be entirely contained in U^I. By the end of the proof of Theorem 6A, we can assume that γ lives in a single chart in the model Banach space for $\mathfrak{M}$. Let the initial point and terminal point for our path be (μ_0, ν_0) and (μ_1, ν_1) respectively. We show that we can slightly perturb the endpoint so that the straight line that connects (μ_0, ν_0) to the new pair of metrics lies entirely in U^I. Then, provided our perturbation is close enough, we can use openness to connect to (μ_1, ν_1). Let $V \subset \mathfrak{M}'$ be a small convex neighbourhood of (μ_1, ν_1) and for $(\mu, \nu) \in V$ consider the map $\delta_{\mu,\nu}(p, t) : \hat{\Sigma} \times [0, 1] \to \mathbb{C}^n$ given by

$$\delta_{\mu,\nu}(p, t) = \Psi(p, t(\mu_0, \nu_0) + (1 - t)(\mu, \nu)).$$

Set $m = \min\{r - 1, k\}$ (see page 1). Using [EL81, Corollary 3.2], one can apply standard arguments to see that $\delta_{\mu,\nu}$ is C^m.

Remark 6.6.3. The reader can find an example of such an argument in Proposition 2.4.7 of [AMR88].

Moreover, the association $(\mu, \nu) \mapsto \delta_{\mu,\nu}$ gives a well-defined map $\delta : U \to C^m(\hat{\Sigma} \times [0, 1], \mathbb{C}^n)$. The evaluation map $\beta : U \times (\hat{\Sigma} \times [0, 1]) \to \mathbb{C}^n$. is given by

$$\beta((\mu, \nu), (p.t)) = \Psi(p, t(\mu_0, \nu_0) + (1 - t)(\mu, \nu)).$$

This map β has the same regularity as Ψ. As above, δ is a C^m representation. The dimension of $\tilde{\Sigma} \times [0, 1]$ is 3 and the codimension of $\mathcal{L}_0$ and $\mathcal{L}_1$ are $2n$ and $n - 1$ respectively, at least 4. Since $m \geq 1 > 0$, it is legal to apply the Parametric Transversality Theorem, and we conclude that for a generic point $(\mu, \nu) \in U$, the map $\delta_{\mu,\nu}$ is transverse to $\mathcal{L}_0, \mathcal{L}_1$. From the dimension and codimension constraints, this implies the path does not touch $\mathcal{L}_0$ or $\mathcal{L}_1$. Therefore, the path

$$\pi(p, t(\mu_0, \nu_0) + (1 - t)(\mu, \nu))$$

lies in U^I and connects (μ_0, ν_0) and (μ, ν). This shows we may connect (μ_0, ν_0) via a straight line contained in U^I to a point arbitrarily close to (μ_1, ν_1), and hence completes the proof. □

We now explain Theorem 6C. Define

$$\Phi : \Sigma^2 \times \mathfrak{M} \to M^2$$

by

$$(p, q, \mu, \nu) \mapsto (f_{\mu,\nu}(p), f_{\mu,\nu}(q)),$$

and let $\mathcal{E}$ be the diagonal

$$\mathcal{E} = \{(x,x) : x \in M\} \subset M^2.$$

If Φ is transverse to $\mathcal{E}$, then, again heuristically, the set of metrics on which $\mathfrak{M}$ fails to be an embedding has codimension $n-4$. The derivative in a $\dot{\nu}$ direction is just $(V(p), V(q))$, where V is the associated harmonic variation. The following lemma is the transversality criterion.

Lemma 6.6.4. *Fix points $(p,q,\mu,\nu) \in \Sigma^2 \times \mathfrak{M}$. Suppose that for every $W_1 \in \mathbf{F}_p$, $W_2 \in \mathbf{F}_q$, there exists a variation $\dot{\nu} \in T_\nu \mathfrak{M}(M)$ such that if $V \in \Gamma(\mathbf{F})$ is the section satisfying $\mathbf{J}V = \mathcal{G}(\dot{\nu})$, then $(V(p), V(q)) = (W_1, W_2)$. Then Ψ is transverse to $\mathcal{E}$ at (p,q,μ,ν).*

As above, Theorem 6C follows from a lemma that we leave for later.

Lemma 6.6.5. *Suppose $f_{\mu,\nu}$ is somewhere injective and has isolated singularities. Then the hypothesis of the lemma above is satisfied. That is, Ψ is transverse to $\mathcal{E}$ at (p,μ,ν).*

Assuming this lemma, the proof of Theorem 6C follows the same line as the proofs of Theorems 6A and 6B (in fact, it is simpler). Hence we omit the proof.

The holomorphic line bundle L

Working toward the lemmas, we introduce the line bundle $\mathbf{L}$. Here we follow the exposition of [Moo17, section 4.1]. Fix a pair $(\mu,\nu) \in \mathfrak{M}$ and let $f = f_{\mu,\nu}$ denote the associated harmonic map. As in Chapter II, if we take a local complex parameter $z = x + iy$, f_z is a local holomorphic section of the bundle $\mathbf{E}$, which we recall is equipped with its Koszul-Malgrange holomorphic structure.

While f_z is only locally defined, the zero set is independent of the choice of coordinate, and the projectivization $[f_z]$ is a well-defined holomorphic section of the projectivized bundle $\mathbb{P}(\mathbf{E})$. Analytically continuing to the zero set we obtain a well-defined global section

$$[f_z] : \Sigma \longrightarrow \mathbb{P}(\mathbf{E}).$$

This section defines a family of lines in $\mathbf{E}$, which patch together to form a holomorphic line bundle $\mathbf{L} \subset \mathbf{E}$, and f_z may be naturally viewed as a local holomorphic section of **L**. If p is a branch point, we can choose a coordinate z in which $z(p) = 0$ and

$$f_z = z^k g(z),$$

where g is a local section of **L** such that $g(p) \neq 0$. The integer k is called the branching order of f.

The **E**-valued $(1,0)$-form $f_z dz$ is naturally a holomorphic section of the holomorphic vector bundle $\mathbf{L} \otimes \mathbf{K}$, where **K** is the canonical bundle. If f branches at points $p_1, \ldots, p_n$ with branching orders $k_{p_1}, \ldots, k_{p_n}$, then $f_z dz$ defines a nowhere vanishing holomorphic section of the bundle

$$\mathbf{L} \otimes \mathbf{K} \otimes \zeta_{p_1}^{-k_{p_1}} \otimes \cdots \otimes \zeta_{p_n}^{-k_{p_n}},$$

where ζ_{p_j} is the holomorphic point bundle at p_j. It follows that

$$\mathbf{L} \simeq \mathbf{K}^* \otimes \zeta_{p_1}^{k_{p_1}} \otimes \cdots \otimes \zeta_{p_n}^{k_{p_n}}.$$

The degree of **L** can then be computed by the evaluation of the first Chern class against the fundamental class of Σ:

$$\deg L = \langle c_1(\mathbf{L}), [\Sigma] \rangle = 2 - 2g + \sum_p k_p,$$

where the sum is taken over the branch set.

Prescribing harmonic variations for Lemma 6.6.2

Lemma 6.6.2 is a special case of the following stronger result.

Lemma 6.6.6. *Fix a local complex coordinate $z = x + iy$ near $p \in \Sigma$. Then, for any three vectors $Z_j \in \mathbf{F}_p$, $j = 1, \ldots, 3$, we can find $\dot{v} \in T_v\mathfrak{M}^*(M)$ such that*

$$V(p) = Z_1 \ , \ \nabla_x V(p) = Z_2 \ , \ \nabla_y V(p) = Z_3,$$

where $\mathcal{G}(\dot{v}) = \mathbf{J}V$.

Suppose the lemma false, so that there are three vectors $Z_1, Z_2, Z_3 \in \mathbf{F}_p$ such that the above fails for every V of the form $\mathbf{J}V = \mathcal{G}(\dot{v})$, where $\dot{v} \in T_v\mathfrak{M}(M)$. Considering the induced inner product on $\oplus_1^3 \mathbf{F}_p$, we can find a triplet of vectors $U_1, U_2, U_3 \in \mathbf{F}_p$ (with not all of them equal to the zero vector) that is orthogonal to every of the form $V(p)$, where V is a section such that $\mathbf{J}V = \mathcal{G}(\dot{v})$. This yields the identity

$$\langle V(p), U_1 \rangle + \langle \nabla_x V(p), U_2 \rangle + \langle \nabla_y V(p), U_3 \rangle = 0$$

for every such V.

Adding together reproducing kernels, we obtain a smooth section $X : \Sigma \backslash \{p\} \to \mathbf{F}$ such that

$$\langle W(p), U_1 \rangle + \langle \nabla_x W(p), U_2 \rangle + \langle \nabla_y W(p), U_3 \rangle = \int_\Sigma \langle \mathbf{J}W, X \rangle dA \tag{6.19}$$

for every $W \in \Gamma(\mathbf{F})$. Moreover, $\mathbf{J}X(p) = 0$ for every $x \in \Sigma \backslash \{p\}$ and the growth of X is controlled by $|z|^{-1}$ at p.

Lemma 6.6.7. *X is not identically zero.*

Proof. It is an elementary exercise to show that one can construct sections of $\mathbf{F}$ with prescribed 1-jet at p (and we used this fact already in Section 6.3). So, one can choose W such that the left-hand side of equation (6.19) is positive. □

The following lemma, a very important piece of our argument, is the content of [Moo06, Lemma 3.1]. The argument can also be found in Moore's book [Moo17, page 311].

Lemma 6.6.8. *Let $\Omega \subset \Sigma^{reg}(f)$ be a small open subset of the regular set of f, and assume $f = f_{\mu,\nu}$ satisfies $f^{-1}(f(\Omega)) = \Omega$. Suppose $Y : \Omega \to \mathbf{F}$ is a smooth section. If*

$$\int_\Sigma \langle \mathbf{J}V, Y \rangle dA = 0$$

for every $\dot{\nu} \in T_\nu \mathfrak{M}(M)$ whose support is contained in $f(\Omega)$, then each point $p \in \Omega$ has a neighbourhood on which Y equals the real part of a local holomorphic section of $\mathbf{L}$.

Remark 6.6.9. The existence of such a set Ω is guaranteed by the hypothesis that f is somewhere injective.

Since the somewhere injective property is so strongly used, we give a word on the proof.

Ideas in the proof. The hypothesis that $f^{-1}(f(\Omega)) = \Omega$ implies that if we take any variation $\dot{\nu}$ with support in $f(\Omega)$, then the support of $\mathbf{J}V = \mathcal{G}(\dot{\nu})$ is contained in Ω, where V is the associated harmonic variation. Therefore,

$$\int_\Sigma \langle \mathbf{J}V, Y \rangle dA = \int_\Omega \langle \mathbf{J}V, Y \rangle dA. \tag{6.20}$$

Choosing variations $\dot{v}$ normal to $f(\Omega)$, Moore uses (6.20) to show that Y is a tangential section of $\mathbf{F}$ over Ω, i.e., it maps into the image of $df(T\Sigma|_U)$ inside the pullback bundle f^*TM. Since f is regular in Ω, one can identify $\mathbf{F}|_\Omega$ with a real subbundle of $\mathbf{L}$. This will be explained after Proposition 6.6.10. Then, choosing tangential variations, (6.20) is used to show that, under the identification, Y is the real part of a holomorphic section of $\mathbf{L}$. □

Now we return to our main argument. Choose an open set Ω as above and not containing p and apply Lemma 6.6.8 to the section X. Note that X has no singularity in Ω. Let Z be the holomorphic section of $\mathbf{L}$ defined on Ω.

We use the isolated singularity condition to analytically continue Z. Let U be any open subset of the regular set that intersects Ω with non-empty interior. We chose a conformal coordinate $z = x_1 + ix_2$ on the source as well as coordinates on the target so that we could write

$$X = X^j \frac{\partial}{\partial u_j}$$

with $\partial f^i/\partial x_j = \delta_{ij} u_i$. This identifies the first two components with the tangent bundle over Ω, and we get an orthogonal splitting into tangential and normal components as

$$f^*TU = (f^*TU)^T \oplus (f^*TU)^\perp.$$

$X|_\Omega$ is tangential, and hence the projection $\pi_{(f^*TU)^\perp}(X)$ vanishes on Ω. We see via the next proposition that this holds on all of U.

Proposition 6.6.10. $\pi_{(f^*TU)^\perp}(X) = 0$ *on all of* U. *In other words,* $X|_U$ *is tangential to the image surface* $f(U)$.

Remark 6.6.11. This is automatic when the metrics (μ, ν) are real analytic (which implies f is real analytic as well).

Proof. We prove $\pi_{(f^*TU)^\perp}(X) = 0$ in an open disk $V \subset U$ that intersects Ω. The proposition then follows from point-set considerations.

Let $\left\{\frac{\partial}{\partial f_1}, \frac{\partial}{\partial f_2}, \dots \frac{\partial}{\partial f_n}\right\}$ be a trivialization for f^*TV such that $\left\{\frac{\partial}{\partial f_1}, \frac{\partial}{\partial f_2}\right\}$ and $\left\{\frac{\partial}{\partial f_3}, \dots \frac{\partial}{\partial f_n}\right\}$ are frames for the tangential and normal subbundles respectively. In these frames, write $X = X^j \frac{\partial}{\partial f_j}$, so that

$$\pi_{(f^*TU)^\perp}(X) = \sum_{j=3}^{n} X^j \frac{\partial}{\partial f_j}.$$

Let $p \in V$ and let $z = x + iy$ be a local complex parameter for V with $z(p) = 0$. In the coordinate, $\mathbf{J}X$ is expressed as

$$\mathbf{J}X = -\mu^{-1}\Big(\sum_{j=1}^{n}(X^j_{xx} + X^j_{yy})\frac{\partial}{\partial f_j}\Big) + \mu^{-1}\sum_{j=1}^{n}\Big(2X^j_x\nabla_x\partial f_j - 2X^j_y\nabla_y\frac{\partial}{\partial f_j}\Big)$$
$$-\mu^{-1}\Big(\sum_{j=1}^{n}X^j(\nabla_x\nabla_x + \nabla_y\nabla_y)\frac{\partial}{\partial f_j} - \mu^{-1}\sum_{i,k,j,\ell}X^k({}^{\nu}R^j_{\ell ki}\circ f)(f^\ell_x f^i_x + f^\ell_y f^\alpha_y)\frac{\partial}{\partial f_j}\Big),$$

where the ${}^{\nu}R^{\gamma}_{\delta\beta\alpha}$ are the coordinate expressions for the Riemannian curvature tensor of (M, ν). Since $\mathbf{J}X = 0$ on V, we deduce that for all j,

$$\Big|\Big(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\Big)X^j\Big| \lesssim |\nabla X| + |X|,$$

where ∇ is the ordinary Euclidean gradient in the local coordinates. We can now invoke the Hartman-Wintner theorem [HW53], which asserts that, in our choice of coordinates,

$$X(z) = p(z) + r(z),$$

where $p(z)$ is a vector-valued harmonic homogeneous polynomial, and $r(z) \in O(zp(z))$. It follows immediately that $X^j = 0$ on V for $j \geq 3$. □

Next, let γ be any curve emanating from Ω that does not intersect the singular set (this includes p). In a neighbourhood U containing the first intersection point of $\gamma \cap \Omega$, we continue to choose a conformal coordinate $z = x_1 + x_2$ so that $\partial f^i/\partial x_j = \delta_{ij}u_i$. Then there is a real linear isomorphism $\tau : (f^*TU)^T \to \mathbf{L}|_U$ defined by

$$M\partial/\partial x_1 + N\partial/\partial x_2 \mapsto (M + iN)f_z.$$

In $\Omega \cap U$, the proof of Lemma 6.6.8 explicitly constructs the holomorphic section Z as

$$Z = \tau(X) = (X_1 + iX_2)\frac{\partial}{\partial z}.$$

Extending Z by this formula on all of U, it is easily checked that Z is a Jacobi field if and only if X is. Thus, arguing similarly to above, we see via Hartman-Wintner that in our coordinates,

$$Z(z) = q(z) + s(z),$$

with q a complex vector-valued harmonic homogeneous polynomial and $s(z)$ decaying faster. Differentiating, the local expression for the section $\overline{\partial}Z$ takes this form as well. Thus, since $\nabla_{\overline{z}}Z = \overline{\partial}Z$ vanishes on Ω, it vanishes everywhere. That is, Z is holomorphic on U. In this way, we continue along all of γ. The next lemma shows that the analytic continuation does not depend on the path.

Lemma 6.6.12. *Let W denote a local holomorphic section of* $\mathbf{L}$. *Then $Re(W)$ is not identically* 0.

Indeed, suppose we have two open sets $\Omega_1, \Omega_2 \supset \Omega$ and local holomorphic extensions X_1, X_2 of X. Setting $W = X_1 - X_2$, the lemma above forces $W \equiv 0$. Thus Z extends in well-defined fashion to the complement of the singular set.

Proof. This is also found in [Mar18, Proposition 4.2], but we include the proof for completeness. In a local complex parameter $z = x + iy$, W may be written $W = hf_z$ for some locally defined meromorphic function $h = h_1 + ih_2$ (with possible poles matching up with zeros of f). Then

$$\mathrm{Re}(W) = \frac{1}{2}(h_1 f_x + h_2 f_y).$$

If W is non-zero and $\mathrm{Re}(W) \equiv 0$, then f_x and f_y are linearly dependent vectors, and moreover $\mathrm{Rank}(df) < 2$ on Ω. This is impossible since, as remarked earlier, the set of regular points for f is open and dense in Σ. □

We now address singular points.

Lemma 6.6.13. *Z extends holomorphically over every singular point except possibly p.*

Proof. Let $\mathcal{S}$ be the singular set of f, so that we have a section $Y : \Sigma \backslash \mathcal{S} \to \mathbf{F}$ such that $Z = X + iY : \Sigma \backslash \mathcal{S} \to \mathbf{L}$ is holomorphic.

From the local coordinate expression for Z, the norm with respect to the natural metric on $\mathbf{F}$ blows up at singular points at worst like the inverse of the Jacobian of f. Thus, Z extends to a meromorphic section of $\mathbf{E}$ on all of Σ. Taking the projectivization gives a well-defined holomorphic section $[Z] : \Sigma \to \mathbb{P}(\mathbf{E})$, which by the identity theorem must agree with $[f_z]$. That is, Z is parallel to f_z, even at the singularities.

Let $q \neq p$ be a singularity. Choosing a local complex coordinate z with $z(q) = 0$ and any trivialization for our bundle, we can write

$$Z = z^n g(z),$$

with g holomorphic and parallel to f_z, $g(0) \neq 0$, and $n \in \mathbb{Z}$. Write $g = h(z) f_z$, with h meromorphic. Let $g = g_1 + ig_2$, $h^{-1} = h_1^{-1} + ih_2^{-1}$. Then

$$f_z = \frac{1}{2}(f_x - if_y) = \frac{1}{2}\Big((h_1^{-1}g_1 + h_2^{-1}g_2) - i(h_2^{-1}g_1 - h_1^{-1}g_2)\Big),$$

and if $X = a(z)f_x + b(z)f_y$ away from q (the coefficients may blow up at q), then we can also write

$$\begin{aligned} X &= a(z)(h_1^{-1}g_1 + h_2^{-1}g_2) + b(z)(h_2^{-1}g_1 - h_1^{-1}g_2) \\ &= (a(z)h_1^{-1} + b(z)h_2^{-1})g_1 + (a(z)h_2^{-1} - b(z)h_1^{-1})g_2. \end{aligned}$$

Since X is regular bounded at q, the coefficents on g_1 and g_2 are regular and bounded. This demonstrates that X can be expressed as a real section in the trivialization determined by g. The same can be done for Y off q. Since X is bounded at q, it follows that the singularity of Z is removeable. □

To obtain a contradiction and finish the proof of Lemma 6.6.2, we explain that no holomorphic section such as Z can exist. From the proof of Lemma 6.6.13, Z behaves at worst like the asymptotic (6.7), so it extends to a globally defined meromorphic section

$$Z : \Sigma \to \mathbf{L}$$

with a pole of order at most 1 at p. We let $\mathrm{ord}_q^{\mathbf{L}}(X)$, $\mathrm{ord}_q^{\mathbf{E}}(X)$ denote the order of vanishing of Z at q with respect to the charts on the bundles $\mathbf{L}$ and $\mathbf{E}$ respectively. In this notation,

$$\mathrm{ord}_{p_j}^{\mathbf{L}}(Z) = \mathrm{ord}_{p_j}^{\mathbf{E}}(Z) + k_j.$$

The degree of the divisor for Z with respect to $\mathbf{L}$ agrees with the degree of $\mathbf{L}$, so that

$$\deg \mathbf{L} = \sum_{q\in\Sigma} \mathrm{ord}_q^{\mathbf{L}}(Z) = \sum_{q\in\Sigma} \mathrm{ord}_q^{\mathbf{E}}(Z) + \sum_j k_j \geq -1 + \sum_j k_j.$$

Meanwhile, we showed earlier that

$$\deg \mathbf{L} = 2 - 2g + \sum_j k_j.$$

This implies $2 - 2g \geq -1$, or $g \leq 3/2$, and this contradiction establishes the result.

Harmonic embeddings: the proof of Lemma 6.6.5

This is similar to Lemma 6.6.2, so we only sketch the proof. As before, the lemma follows from a more general result.

Lemma 6.6.14. *Fix a local complex coordinate $z = x + iy$ near $p \in \Sigma$. Then for any three vectors $Z_j \in \mathbf{F}_p$, $j = 1, \ldots, 3$, we can find $\dot{v} \in T_v\mathfrak{M}^*(M)$ such that*

$$V(p) = Z_1\ ,\ \nabla_x V(p) = Z_2\ ,\ \nabla_y V(p) = Z_3,$$

where $\mathcal{G}(\dot{v}) = \mathbf{J}V$.

If the lemma fails, there are vectors $U_1 \in \mathbf{F}_p$, $U_2 \in \mathbf{F}_q$ such that for every harmonic variation V,

$$\langle U_1, V(p)\rangle + \langle U_2, V(q)\rangle = 0.$$

Taking the reproducing kernels for U_1 and U_2, we have a section $X : \Sigma\backslash\{p,q\} \to \mathbf{F}$ such that

$$\int_\Sigma \langle \mathbf{J}W, X\rangle dA = \langle U_1, W(p)\rangle + \langle U_2, W(q)\rangle$$

for all sections $W \in \Gamma(\mathbf{F})$. Invoking Moore's lemma and then repeating our argument from the previous subsection, one finds a section $Y : \Sigma\backslash\{p,q\} \to \mathbf{F}$ such that $Z = X + iY$ is holomorphic. The asymptotic (6.6) ensures that Z blows up strictly slower than any meromorphic section, and hence the singularities at p and q are removeable. Thus, Z yields a globally defined holomorphic section, a Jacobi field, and this is a contradiction.

Toward the Whitney theorems

Theorems 6B and 6C show that Conjecture 6E holds near harmonic surfaces that are somewhere injective and have isolated singularities. To conclude the chapter, we discuss our use of this hypothesis and the possibility of removing it.

Firstly, to prove Conjecture 6E, by Theorem 6A it suffices to prove it holds near surfaces that are somewhere injective. So the only extra condition that we use here is the isolated singularities. Let us assume that (μ, ν) are such that $f_{\mu,\nu}$ is somewhere injective, with no condition on singularities. Beginning the proof of Lemma 6.6.2, we find a kernel X satisfying (6.19). Then we can find an open set Ω on which f is injective, and by Moore's lemma, $X|_\Omega$ is the real part of a holomorphic section $Z : \Omega \to \mathbf{L}$. At this point, it is tempting to believe that some sort of unique continuation argument should promote this to a global result.

It is unclear if this is possible, one obstruction being that we are not aware of a coordinate-free way to express that X is the real part of a holomorphic section. The issue stems from the following remark: X is a real section of $\mathbf{E}$ with respect to the real structure induced from the splitting $\mathbf{E} = \mathbf{F} \oplus i\mathbf{F}$, but there is no reason for a transition map to the holomorphic trivialization for the Koszul-Malgrange holomorphic structure to preserve this real structure. All we can say is that in such a trivialization,

$$X(z) = K(z)X_0(z),$$

where $X_0(z)$ is a real section, and $K(z)$ is a smoothly varying family of complex matrices. Furthermore, the imaginary part Y has been defined in terms of a particular local frame for **F**. And the norm of the elements in this frame may explode as we approach singularities of f. In other words, we have no apriori uniform continuity for Z in Ω, and the imaginary part could blow up in an attempt to analytically continue.

While the holomorphic coordinates on **E** are opaque, we do have some understanding of what it means to be a holomorphic section of **L**. This is what allows for some results under stronger hypothesis. In the end we want to find our contradiction by realizing X as the real part of a global meromorphic section of **L** with constrained poles. Two steps:

1. Show that when f is regular, X is tangential to f.
2. Find mappings from the distribution $df(T\Sigma) \subset f^*TM$ to **L**, under which X corresponds to the real part of a meromorphic section (a meromorphic multiple of f_z).

When f is regular, we have well-defined splittings of f^*TM into tangential and normal components for the image of f. Thus, if the set $A(f)$ from Section 4 is connected, we can show (1) holds via a unique continuation argument (Proposition 6.6.10). Note that the isolated singularity condition is really more than what we need for this to work. Once we have (1), we can define the section Y at regular points as before, and again a connectedness assumption allows us to deduce that $Z = X + iY$ is holomorphic where defined.

The last challenge is to extend Z over singular points. If the singularity is isolated, then Z is meromorphic, and then we can argue using the boundedness of X in the right choice of coordinates. But with a more complicated singular set—say, a general fold or a meeting point of general folds—controlling Z becomes a delicate task.

At this point, we see no direct geometric reason for the argument to fail in general. It is reasonable to expect that we can relax our assumptions to include harmonic maps with particular types of (non-isolated) singularities. It is unclear how far the method goes.

Chapter 7

UNSTABLE MINIMAL SURFACES IN PRODUCTS

The content of this chapter is joint work with Vladimir Marković and Peter Smillie.

7.1 Introduction

Minimal surfaces in products of hyperbolic surfaces

Let Σ_g denote a closed surface of genus $g \geq 2$ and let $\mathbf{T}_g$ be the Teichmüller space of marked complex structures on Σ_g. Let (X, d) be the hyperbolic plane, an $\mathbb{R}$-tree, or product thereof with an action $\sigma : \pi_1(\Sigma_g) \to \mathrm{Isom}(X, d)$. For every Riemann surface structure S on Σ_g, with universal cover $\tilde{S}$, and σ-equivariant Lipschitz map $f : \tilde{S} \to (X, d)$, there is a well-defined notion of Dirichlet energy $\mathcal{E}(S, f)$ (see Section 2 for details). For admissible σ, there is an essentially unique σ-equivariant harmonic map $h : \tilde{S} \to (X, d)$, which satisfies

$$\mathcal{E}(S, h) = \inf_f \mathcal{E}(S, f).$$

This gives a function $\mathbf{E}_\sigma : \mathbf{T}_g \to \mathbb{R}$, by $\mathbf{E}_\sigma(S) = \mathcal{E}(S, h)$. When S is a critical point of $\mathbf{E}_\sigma$, we say that h is minimal; if X is a manifold and h is an immersion, this is equivalent to $h(S)$ being a minimal surface.

One case of interest is when σ is a product of Fuchsian representations into $\mathrm{PSL}(2, \mathbb{R})^n$ (also called a maximal representation), in which case each component of the harmonic map is a diffeomorphism, and critical points correspond to genuine minimal surfaces in a product of hyperbolic surfaces. The work of Schoen-Yau [SY79] implies that in this case, $\mathbf{E}_\sigma$ is proper, and therefore admits a global minimum, which is a stable critical point. For $n = 2$, Schoen proved that this is the unique critical point of $\mathbf{E}_\sigma$ [Sch93].

However, the first author proved in [Mar22] that uniqueness fails when $n \geq 3$, assuming the genus g is large enough. See also the paper [Mar21], which provides a strengthening of Schoen's result for $n = 2$. The main goal of this chapter is to show that unstable equivariant minimal surfaces in $\mathbb{R}^n$ yield unstable minimal surfaces in products of hyperbolic surfaces. In particular, this strengthens the result from [Mar22], while providing a simpler and more revealing proof. When $n \geq 3$, there are many unstable equivariant minimal surfaces in $\mathbb{R}^n$; most notably, unstable minimal surfaces in tori, which Meeks [Mee90], Hass-Pitts-Rubenstein [HPR93], and Traizet [Tra08] have shown to be abundant, lift to unstable equivariant minimal surfaces in $\mathbb{R}^n$.

We say that a critical point of $\mathbf{E}_\sigma$ is unstable if there exists a C^2 path in $\mathbf{T}_g$ starting at the point and at which the second derivative of $\mathbf{E}_\sigma$ along the path is negative.

Theorem 7A. Let $n \geq 3$. For every genus $g \geq 2$, there exists a maximal representation $\sigma : \pi_1(\Sigma_g) \to \prod_{i=1}^n \mathrm{PSL}(2, \mathbb{R})$ such that $\mathbf{E}_\sigma : \mathbf{T}_g \to (0, \infty)$ admits an unstable critical point. In particular, there are at least two minimal surfaces in the product of hyperbolic surfaces determined by σ.

Labourie conjectured that for a Hitchin representation into a simple split real Lie group G of non-compact type, there exists a unique equivariant minimal surface in the corresponding symmetric space. Labourie proves existence in general [Lab08], and that uniqueness holds when the rank of G is 2 [Lab17] (see also [CTT19], where Collier-Tholozan-Toulisse prove the analogous statement for maximal representations into Hermitian Lie groups of rank 2). The conjecture remains open in rank at least 3, and [Mar22] suggests that this is the critical case.

The key idea of the proof of Theorem 7A is to reduce it to finding unstable minimal surfaces in products of $\mathbb{R}$-trees (Theorem 7B2 below). The unstable minimal surfaces are provided by Theorems 7C and 7D. We explain in the forthcoming subsections.

Minimal surfaces in products of $\mathbb{R}$-trees

We give the definitions about harmonic maps to $\mathbb{R}$-trees in Section 7.2. Throughout the chapter, let S be a Riemann surface structure on Σ_g and $\mathrm{QD}(S)$ the space of holomorphic quadratic differentials on S. The Riemann surface structure S lifts to a Riemann surface structure on the universal cover of Σ_g, which we denote $\tilde{S}$. Given a non-zero $\phi \in \mathrm{QD}(S)$, there are two natural ways of producing an equivariant harmonic map. First, the leaf space of the vertical singular foliation of the lift $\tilde{\phi}$ to $\tilde{S}$ is an $\mathbb{R}$-tree (T_ϕ, d). The action of $\pi_1(\Sigma_g)$ on $\tilde{S}$ descends to an action $\rho : \pi_1(\Sigma_g) \to \mathrm{Isom}(T_\phi, d)$ by isometries. The quotient map $\pi : \tilde{S} \to (T_\phi, d)$ is harmonic and ρ-equivariant, with Hopf differential $\phi/4$.

On the other hand, it is proved independently by Hitchin [Hit87], Wan [Wan92], and Wolf [Wol89] that there is a unique hyperbolic structure M_ϕ on Σ_g such that the identity map from S to M_ϕ is harmonic with Hopf differential ϕ. Moreover Wolf shows that as $t \to \infty$, $M_{t\phi}$ converges in a certain sense to the rescaled tree $(T_\phi, 2d)$ (see [Wol95] for the precise statement).

Now let $\phi_1, \dots, \phi_n$ be n nonzero holomorphic quadratic differentials on the same surface S, and let X be the product of the $\mathbb{R}$-trees $(T_{\phi_i}, 2d_i)$ arising from the construction above. Let $\rho : \pi_1(\Sigma_g) \to \mathrm{Isom}(X)$ be the product of the actions ρ_i on each factor. The energy function $\mathbf{E}_\rho$ on $\mathbf{T}_g$ associated to ρ is then the sum of the energy functions E_{ρ_i} associated to each component. Also for each positive $t > 0$, let M_i^t be the hyperbolic structures associated to $t\phi_i$. We set $\mathbf{E}_\rho^t$ to be the energy functional for the product of Fuchsian representations associated to the M_i^t.

S is a critical point for $\mathbf{E}_\rho$ if and only if it is a critical point for $\mathbf{E}_\rho^t$ for every $t > 0$. In other words, minimality of the harmonic map into the product of surfaces is equivalent to the minimality of the equivariant harmonic map into the product of $\mathbb{R}$-trees. The condition occurs precisely when $\sum_{i=1}^n \phi_i = 0$.

Let $n \geq 2$. For $i = 1, \dots, n$, let ϕ_i be nonzero holomorphic quadratic differentials on the Riemann surface S such that $\sum_{i=1}^n \phi_i = 0$.

Theorem 7B1. S is not a (local) minimum for $\mathbf{E}_\rho$ if and only if there exists $t > 0$ such that S is not a (local) minimum for $\mathbf{E}_\rho^t$. In this case, for all $s > t$, S is not a (local) minimum for $\mathbf{E}_\rho^s$.

Remark 7.1.1. If $n = 2$, Schoen's result shows that the only critical point of $\mathbf{E}_\rho^t$ is a minimum, and so by (1), the same is true of $\mathbf{E}_\rho$. This was first proved by Wentworth who showed that, provided existence, the equivariant minimal surface in a product of two $\mathbb{R}$-trees is unique [Wen07, Theorem 1.6].

Remark 7.1.2. It appears to be unknown whether the energy functional on Teichmüller space for harmonic maps to $\mathbb{R}$-trees is C^2. It is always C^1, and real analytic near a Riemann surface such that the Hopf differential of the harmonic map has only simple zeros (this is a generic condition) [Mas95].

Theorem 7B1 can give critical points of $\mathbf{E}_\rho^t$ that are not minima, but this is not quite strong enough to prove Theorem 7A, which is about unstable critical points. To that end, we give a notion of instability in products of $\mathbb{R}$-trees that will be suitable for our purposes. Let S be a critical point for $\mathbf{E}_\rho$ with harmonic map $\pi = (\pi_1, \dots, \pi_n)$. Given C^∞ vector fields $V_1, \dots, V_n$ on S, let $r \mapsto f_1^r, \dots, f_n^r : S \to S$ be their flows, and construct the map $\pi_r = (\pi_1 \circ f_1^r, \dots, \pi_n \circ f_n^r)$. For any C^∞ path of Riemann surfaces $r \mapsto S_r$, there is a Beltrami form μ representing a point $T_S\mathbf{T}_g$ that is tangent to our path at $r = 0$.

Definition 7.1.3. We define a quadratic form $\mathbf{L} : T_S\mathbf{T}_g \times H^0(S, TS)^n \to \mathbb{R}$ by

$$\mathbf{L}(\mu, V_1, \dots, V_n) = \frac{d^2}{dr^2}|_{r=0}\mathcal{E}(S_r, \pi_r),$$

where $r \mapsto S_r$ is any path tangent to μ at $r = 0$.

The self-maps index of S for $\mathbf{E}_\rho$ is the maximal dimension of $T_S\mathbf{T}_g \times H^0(S, TS)^n$ on which L is negative definite. If the index is positive, we say that S is unstable.

We explain that $\mathbf{L}$ is well-defined in Section 2.2. $\mathbf{L}$ is positive semi-definite on $\{0\} \times H^0(S, TS)^n$, and hence if $\mathbf{L}$ is negative definite on a subspace $U \subset T_S\mathbf{T}_g \times H^0(S, TS)^n$, then U projects injectively to $T_S\mathbf{T}_g \times \{0\}$. Moreover, for any variations $r \mapsto S_r$ and $r \mapsto \pi_r$, we have $\mathbf{E}_\rho(S_r) \leq \mathcal{E}(S_r, \pi_r)$, and hence if $\mathbf{L}(\mu, V_1, \dots, V_n) < 0$, then $\mathbf{E}_\rho(S_r) < \mathbf{E}_\rho(S)$ for small r. See Remark 7.3.4 below for more motivation for the definition of $\mathbf{L}$.

Theorem 7B2. The index of $\mathbf{E}_\rho^t$ at S is non-decreasing with t, and converges to the self-maps index of S for $\mathbf{E}_\rho$ as $t \to \infty$. Consequently, S is unstable for $\mathbf{E}_\rho$ if and only if it is unstable for $\mathbf{E}_\rho^t$, for t sufficiently large.

Toward the proof of Theorem 7A, we only need the “only if" direction of Theorem 7B2. We include the “if" direction and Theorem 7B1 because they show that $\mathbb{R}$-trees are really at the heart of the result. A conjecture in Higher Teichmüller theory is that high energy minimal maps into symmetric spaces converge in an appropriate sense to minimal maps into buildings (see [Kat+15]). This is the higher rank generalization of [Wol95]. If our results extend to this setting, then this would suggest that any counterexample to the Labourie conjecture would have to come from an unstable minimal map into a building.

Equivariant minimal surfaces in $\mathbb{R}^n$

In order to use Theorem 7B2 to prove Theorem 7A, we construct unstable surfaces in products of $\mathbb{R}$-trees. We start by looking in a more familiar place: Euclidean space $\mathbb{R}^n$.

For $i = 1, \dots, n$, let α_i be a non-zero holomorphic 1-form on the Riemann surface S. Lifting to 1-forms $\tilde{\alpha}_i$ on a universal cover $\tilde{S}$ gives the data of a harmonic map to

$\mathbb{R}^n$ via integrating the real parts:

$$h = (h_1, \ldots, h_n),\ h_i(z) = \mathrm{Re} \int_{z_0}^{z} \tilde{\alpha}_i,$$

unique up to translation. The map h intertwines the action of $\pi_1(\Sigma_g)$ on $\tilde{S}$ with some non-trivial homomorphism $\chi : \pi_1(\Sigma_g) \to \mathbb{R}^n$. The Hopf differential of h_i is the square $\phi_i = \alpha_i^2$, which descends to a holomorphic quadratic differential on S, by the equivariance property. h is weakly conformal if and only if $\sum_{i=1}^n \phi_i = 0$, which is equivalent to h being minimal.

By the construction of Section 7.1, the Hopf differentials ϕ_i also define an action ρ of $\pi_1(\Sigma_g)$ on a product X of $\mathbb{R}$-trees and a ρ-equivariant minimal map π. The map h naturally factors through π. Let $\mathbf{E}_\chi$ and $\mathbf{E}_\rho$ be the corresponding energy functionals on Teichmüller space. In the end we prove the following near-equivalence.

Theorem 7C. For $n \geq 2$ and $i = 1, \ldots, n$, let α_i be nonzero holomorphic 1-forms on S such that $\sum_{i=1}^n \alpha_i^2 = 0$,. Let ρ, and χ be as above.

1. If S is not a (local) minimum for $\mathbf{E}_\rho$, then it is not a (local) minimum for $\mathbf{E}_\chi$.
2. The index of $\mathbf{E}_\chi$ at S is equal to the self-maps index of $\mathbf{E}_\rho$ for S. In particular, if S is unstable for $\mathbf{E}_\chi$, then S is unstable for $\mathbf{E}_\rho$.

Remark 7.1.4. As in Theorem 7B, the statement is not so interesting when $n = 2$ since every critical point is a stable minimum.

Remark 7.1.5. Instability for $\mathbf{E}_\chi$ at S is equivalent to instability for the (equivariant) area functional on the space of all equivariant maps. The second variations for both functionals have the same index (see [Eji02, Theorem 3.4]).

The final ingredient needed to prove Theorem 7A is an example of an unstable equivariant minimal surface in $\mathbb{R}^n$. Fortunately, these aren't so hard to find: when $n = 3$, every non-planar equivariant minimal surface in $\mathbb{R}^3$ is unstable, since a constant section of the normal bundle is destabilizing. Consequently, for any three 1-forms on S whose squares sum to zero, as long as they span a 2-dimensional space, S will be an unstable point of $\mathbf{E}_\chi$ (we explain the details in Section 7.5).

The most natural example is the lift of a minimal surface in a flat 3-torus; there a many classical examples, such as the Schwarz P-surface of genus 3 (see [Mee90]). In fact, for every $g \geq 3$, every flat 3-torus contains infinitely many distinct unstable minimal surfaces of genus g in the same homotopy class (see [Mee90], [HPR93], and [Tra08]).

By inclusion, this gives examples for every $n \geq 3$, as long as $g \geq 3$. We can also perturb these examples to give even more. Unfortunately, if $g = 2$, then the only triples of 1-forms whose squares sum to zero are scalar multiples of one another, so we cannot use Theorem 7C to prove Theorem 7A in the case $g = 2$.

A generalization of Theorem 7C, and the case $g = 2$

In the general setting where we start with n quadratic differentials summing to zero that are not necessarily squares of abelian differentials, we may have to lift to a branched covering of S in order to get a harmonic map to $\mathbb{R}^n$. Replacing $\pi_1(\Sigma)$ with the Deck group of this branched covering, and χ with the corresponding representation of this group, we prove that an analog of statement (2) from Theorem 7C still holds.

When $n = 3$, the minimal surface in $\mathbb{R}^3$ arising from a branched cover is no longer automatically unstable for its energy functional, because the normal bundle is not necessarily equivariantly trivial with respect to the action of the Deck group of the branched covering. Instead, we find a condition on the quadratic differentials that guarantees that the bundle is equivariantly trivial, which then yields the following theorem.

Theorem 7D. Let ϕ_1, ϕ_2, ϕ_3 be holomorphic quadratic differentials on S that are not colinear and such that $\phi_1\phi_2\phi_3$ is the square of a cubic differential. Then the corresponding equivariant minimal surface in the product of $\mathbb{R}$-trees is unstable via self-maps for its energy functional.

We show that the locus of solutions to this problem is non-empty for every genus (Proposition 7.5.8). In particular, this gives us unstable minimal surfaces in genus 2, and hence allows us to complete the proof of Theorem 7A.

Acknowledgments

We would like to thank Phillip Engel and Mike Wolf for helpful discussions. Vladimir Marković is supported by the Simons Investigator Award 409745 from the Simons Foundation.

7.2 Preliminaries

Harmonic maps to manifolds

Let ν be a smooth metric on S compatible with the complex structure. Let (M, σ) be a closed Riemannian manifold, and $h : S \to M$ a C^2 map. Assuming that (M, σ) is a surface with conformal metric σ, then in holomorphic coordinates z on S and w on M, we write $\nu = \nu(z)|dz|^2$, $\sigma = \sigma(w)|dw|^2$, and h as a complex-valued function $h(z)$. The energy density takes the form

$$e(h)(z) = \frac{\sigma(h(z))}{\nu(z)}(|h_z|^2 + |h_{\bar{z}}|^2)(z), \ \phi(h)(z) = \sigma(h(z))h_z\overline{(h_{\bar{z}})}dz^2.$$

Considering equivariant maps to the real line, again in local coordinates,

$$e(h)(z) = 2\nu(z)^{-1}|h_z|^2, \ \phi(h) = h_z^2 dz^2.$$

For a harmonic map to a product space, the definitions (2.2) and (2.7) shows that the energy density and the Hopf differential are the sum of the energy densities and the Hopf differentials respectively of the component maps. So for a mapping $h = (h_1, \ldots, h_n)$ into a product of Riemann surfaces, or an equivariant mapping into $\mathbb{R}^n$, the Hopf differential is the sum

$$\phi(h) = \sum_{i=1}^{n} \phi(h_i). \tag{7.1}$$

h is minimal if $\phi \equiv 0$.

If (M, σ) is a negatively curved surface, it is well-known that there is a unique harmonic map $h : S \to (M, \sigma)$ in the homotopy class of the identity (see [ES64] for existence, and [Har67, Theorem H] for uniqueness). If we work on a different Riemann surface structure S' on Σ_g, we get a harmonic map from $S' \to (M, \sigma)$ in the class of the identity, and the total energy depends only on the class of S' in Teichmüller space. Thus, we get a functional $\mathbf{E} : \mathbf{T}_g \to (0, \infty)$, where $\mathbf{E}(S')$ is the total energy of the harmonic map from $S' \to (M, \sigma)$. For a map into a product of surfaces, the energy functional is the sum of the energy functionals of the component mappings.

Harmonic maps to $\mathbb{R}$-trees

Definition 7.2.1. An $\mathbb{R}$-tree is a length space (T, d) such that any two points are connected by a unique arc, and every arc is a geodesic, isometric to a segment in $\mathbb{R}$.

A point $x \in T$ is a vertex if the complement $T \backslash \{x\}$ has greater than two components. Otherwise it is said to lie on an edge.

Let S be a closed Riemann surface of genus $g \geq 2$. The horizontal (resp. vertical) foliation of a holomorphic quadratic differential ϕ on S is the singular foliation whose non-singular leaves are the integral curves of the line field on $S \backslash \phi^{-1}(0)$ on which ϕ is a negative real number. The singularities are standard prongs at the zeros. Both foliations come equipped with transverse measures $|\text{Im}\sqrt{\phi}|$ and $|\text{Re}\sqrt{\phi}|$ respectively (see [FLP12, Exposé 5] for precise definitions).

In this chapter, we work with the vertical foliation, unless specified otherwise. Lifting to a singular measured foliation on a universal cover $\tilde{S}$, we define an equivalence relation under which two points $x, y \in \tilde{S}$ are equivalent if they lie on the same leaf. The quotient space is denoted T, and we can push the transverse measure down via the projection $\pi : \tilde{S} \to T$ to form a distance function d such that (T, d) is an $\mathbb{R}$-tree, with an induced action $\rho : \pi_1(S) \to \text{Isom}(T, d)$.

According to Korevaar-Schoen, for Lipschitz maps f from $\tilde{S}$ to complete and non-positively curved (NPC) length spaces such as (T, d), there is a well-defined L^1 directional energy tensor $g_{ij} = g_{ij}(f)$ that generalizes the pullback metric (see [KS07, Theorem 2.3.2]). In this way, one can define a measurable energy density function by

$$e(f) = \frac{1}{2}\text{trace}_\nu g_{ij}(f). \tag{7.2}$$

For an equivariant Lipschitz map h, the energy density $e(h)$ is invariant under the group, and we define a total energy as in the smooth setting by

$$\mathcal{E}(S, h) = \int_S e(h) dA.$$

Definition 7.2.2. We say that a ρ-equivariant map $h : \tilde{S} \to (T, d)$ is harmonic if, among other ρ-equivariant maps, it is a critical point for the energy $h \mapsto \mathcal{E}(S, h)$.

For the projection map π, we can describe the energy density explicitly. At a point on $p \in \tilde{S}$ on which $\phi(p) \neq 0$, the map locally isometrically factors through a segment in $\mathbb{R}$. In a small enough neighbourhood around that point, $e(\pi)$ is equal to the energy density of the locally defined map to $\mathbb{R}$, which is computed as usual via the formula (2.2). From this, we see that the energy density has a continuous representative that is equal to $\nu^{-1}|\phi|/2$ everywhere.

The Hopf differential is well-defined for maps f from $\tilde{S}$ to NPC spaces as above: in local coordinates, it is given by

$$\frac{1}{4}(g_{11}(f)(z) - g_{22}(f)(z) - 2ig_{12}(f)(z))dz^2. \tag{7.3}$$

The projection map $\pi : \tilde{S} \to (T, d)$ is ρ-equivariant and harmonic, with Hopf differential $\phi/4$. Instead of the equation (7.3), one can also see this by using the local isometric factoring. As in the case of maps to surfaces, a harmonic mapping into a product of trees is called minimal if the Hopf differential–which splits as a sum as in (7.1)–vanishes.

Given (T, d) as above, we always rescale the metric to $2d$, which makes it so that the Hopf differential of $\pi : \tilde{S} \to (T, 2d)$ is ϕ. For any other Riemann surface S' representing a point in $\mathbf{T}_g$, there is a unique ρ-equivariant harmonic map from $\tilde{S}' \to (T, 2d)$ (see [Wol96]). Again like the surface case, the representation ρ then defines an energy functional on Teichmüller space. The same holds for products of $\mathbb{R}$-trees with admissible actions.

Let's now address the quadratic form $\mathbf{L} : T_S\mathbf{T}_g \times H^0(S, TS)^n \to \mathbb{R}$. Given a C^∞ path of Riemann surfaces and a flow $r \mapsto f_r$, we consider

$$r \mapsto \mathcal{E}(S_r, \pi \circ \tilde{f}_r). \tag{7.4}$$

In [Moo06, Section 5], Moore computes the derivative of the two-variable energy for maps from a surface to a Riemannian manifold. Using the characterization (7.2), one can word-for-word redo that computation, but with the measurable density $e(\pi \circ \tilde{f}_r)$, to see that (7.4) is C^∞ in r. Working with a minimal map into a product of trees and n vector fields, one does the computation n times to get $\mathbf{L}$. We note that it only depends on the tangent vectors and not the specific path of Riemann surfaces and flow of maps, because a minimal map is a critical point for the two-variable energy.

7.3 Minimal surfaces in products of $\mathbb{R}$-trees

We first prove Theorem 7B2, and then Theorem 7B1.

The Reich-Strebel energy formula

Reich-Strebel computed a formula for the difference of energies of quasiconformal maps (equation 1.1 in [RS87]). Let $h : S \to M$ and $f : S \to S'$ be quasiconformal maps between Riemann surfaces, with a conformal metric on M. Let μ be the Beltrami form of f, and ϕ the Hopf differential of h, which need not be holomorphic.

Then,

$$\mathcal{E}(S', h\circ f^{-1}) - \mathcal{E}(S,h) = -4\mathrm{Re}\int_S \phi\cdot\frac{\mu}{1-|\mu|^2} + 2\int_S e(h)\cdot\frac{|\mu|^2}{1-|\mu|^2}dA. \quad (7.5)$$

The computation goes through just the same when h maps $\tilde{S}$ equivariantly into an $\mathbb{R}$-tree. For a map h to a tree (T,d) with Hopf differential ψ, the energy density satisfies $e(h) = 2\nu^{-1}|\psi|$. We obtain the Proposition below.

Proposition 7.3.1. *Let $h : \tilde{S} \to (T,d)$ be an equivariant harmonic map to an $\mathbb{R}$-tree with Hopf differential ψ, and $f : S' \to S$ a quasiconformal map. Then the following formula holds:*

$$\mathcal{E}(S', h\circ \tilde{f}^{-1}) - \mathcal{E}(S,h) = -4Re\int_S \psi\cdot\frac{\mu}{1-|\mu|^2} + 4\int_S |\psi|\cdot\frac{|\mu|^2}{1-|\mu|^2}. \quad (7.6)$$

In the formula above, $\tilde{f}$ is the lift to $\tilde{S}$ and μ is the Beltrami form.

Proof of Theorem 7B2

Let $\phi \in \mathrm{QD}(S) - \{0\}$. For $t > 0$, let M_t be the hyperbolic structure with hyperbolic metric σ_t such that the identity map $h_t : S \to M_t$ is harmonic with Hopf differential $t\phi$, let $\mathcal{E}_\rho^t$ be the two-variable energy functional for M_t, and E_ρ^t the corresponding energy functional on Teichmüller space. Similarly, let $\mathcal{E}_\rho$ and E_ρ be the energies for the $\mathbb{R}$-tree $(T, 2d)$ determined by ϕ (with a rescaled metric).

The main step in the proof of Theorem 7B2 is Lemma 7.3.2. If we rescale σ_t by t^{-1}, then for any Riemann surface structure S' and C^2 map $f : S' \to M_t$, the energy with respect to the target metric $t^{-1}\sigma_t$ is $t^{-1}\mathcal{E}_\rho^t(S', f)$. Let $r \mapsto S_r$ be a path of Riemann surfaces and $r \mapsto f_r$ a flow starting at the identity map. Lemma 7.3.2 shows that the second derivative in r of the energy of $h_t \circ f_r$ on S_r with respect to the target metric $t^{-1}\sigma_t$ converges as $t \to \infty$ to the second derivative of the energy of $\pi \circ f_r : \tilde{S}_r \to (T, 2d)$.

Lemma 7.3.2. *For $s > t$,*

$$\frac{1}{t}\frac{d^2}{dr^2}|_{r=0}\mathcal{E}_\rho^t(S_r, h_t\circ f_r) > \frac{1}{s}\frac{d^2}{dr^2}|_{r=0}\mathcal{E}_\rho^s(S_r, h_s\circ f_r) > \frac{d^2}{dr^2}|_{r=0}\mathcal{E}_\rho(S_r, \pi\circ\tilde{f}_r),$$

and

$$\lim_{t\to\infty}\frac{1}{t}\frac{d^2}{dr^2}|_{r=0}\mathcal{E}_\rho^t(S_r, h_t\circ f_r) = \frac{d^2}{dr^2}|_{r=0}\mathcal{E}_\rho(S_r, \pi\circ\tilde{f}_r).$$

Toward the proof, we first record the lemma below about the growth of the energy density.

Lemma 7.3.3. *Let $e(h_t)$ be the energy density of h_t with respect to the target metric σ_t. Then for $s \geq t$,*

$$\frac{e(h_t)}{t} \geq \frac{e(h_s)}{s}, \tag{7.7}$$

and the inequality is strict away from the zeros of ϕ. Moreover,

$$\lim_{t\to\infty} \frac{e(h_t)}{t} = 2\nu^{-1}|\phi|. \tag{7.8}$$

Proof. Let μ_t and μ_s be the Beltrami forms of h_t and h_s respectively. It is proved in [Wol89, Proposition 4.3] that away from the zeros of ϕ (at which $|\mu_t| = 0$ for every t), $|\mu_t|$ monotonically increases to 1 as $t \to \infty$. A simple computation gives

$$e(h) = \frac{2t|\phi|}{\nu} \cosh \log |\mu_t|^{-1},$$

and likewise for s. Therefore, (7.7) is equivalent to the inequality

$$\cosh \log |\mu_{t_0}|^{-1} \geq \cosh \log |\mu_t|^{-1}. \tag{7.9}$$

Since $|\mu_t| < 1$ everywhere, the inequality (7.9) follows. Using the limiting behaviour of $|\mu_t|$, we take $t \to \infty$ to obtain (7.8). □

Proof of Lemma 7.3.2. Let $(\mu, V) \in T_S\mathbf{T}_g \times H^0(S, TS)$, and let $r \mapsto f_r$ be the flow of V. Let μ_r be the Beltrami form of f_r^{-1}, and α the C^∞ $(1,-1)$-form and β the C^∞ function on S described by

$$\alpha(z) = \frac{d^2}{dr^2}\bigg|_{r=0} \frac{\mu_r(z)}{1 - |\mu_r(z)|^2}, \ \beta(z) = \frac{d^2}{dr^2}\bigg|_{r=0} \frac{|\mu_r(z)|^2}{1 - |\mu_r(z)|^2}.$$

We use the Reich-Strebel fomula (7.5). For each $t > 0$,

$$\begin{aligned} \frac{1}{t}\frac{d^2}{dr^2}\bigg|_{r=0} \mathcal{E}_\rho^t(S_r, h^t \circ f_r) &= \frac{1}{t}\frac{d^2}{dr^2}\bigg|_{r=0} \Big(\mathcal{E}_\rho^t(S_r, h^t \circ (f_r^{-1})^{-1}) - \mathcal{E}_\rho^t(S, h^t)\Big) \\ &= \frac{d^2}{dr^2}\bigg|_{r=0}\Big(- 4\mathrm{Re} \int_S \phi \cdot \frac{\mu_r}{1 - |\mu_r|^2} + 2\int_S \frac{e(h^t)}{t} \cdot \frac{|\mu_r|^2}{1 - |\mu_r|^2} dA\Big) \\ &= -4\mathrm{Re} \int_S \phi \cdot \alpha + 2\int_S \frac{e(h^t)}{t} \cdot \beta dA. \end{aligned}$$

On the other hand, by the same computation, but using (7.6),

$$\frac{d^2}{dr^2}\bigg|_{r=0} \mathcal{E}_\rho(S_r, \pi \circ \tilde{f}_r) = -4\mathrm{Re} \int_S \phi \cdot \alpha + 4\int_S |\phi(h)| \cdot \beta.$$

By Lemma 7.3.3, for $s > t$,

$$\int_S \frac{e(h^t)}{t} \cdot \beta dA > \int_S \frac{e(h^s)}{s} \cdot \beta dA,$$

so that $\frac{L_t}{t} > \frac{L_s}{s}$. By Lemma 7.3.3 again and the dominated convergence theorem,

$$\int_S \frac{e(h^t)}{t} \cdot \beta dA \to 2 \int |\phi| \cdot \beta$$

in a strictly decreasing fashion as $t \to \infty$. The result follows.

□

Moving onto the proof of Theorem 7B2, we resume the notation from the introduction: for $i = 1, \dots, n$, ϕ_i are nonzero holomorphic quadratic differentials on S summing to 0. The product of $\mathbb{R}$-trees $(T_i, 2d_i)$, which we denote by X, comes equipped with the action $\rho = (\rho_1 \times \cdots \times \rho_n)$. For each positive $t > 0$, M_i^t is the hyperbolic structure such that the identity map $h_i^t : S \to M_i^t$ is harmonic and has Hopf differential $t\phi_i$. The energy functionals are denoted $\mathcal{E}_\rho$ and $\mathbf{E}_\rho$ for the trees and $\mathcal{E}_\rho^t$ and $\mathbf{E}_\rho^t$ for the surfaces.

We define $\mathbf{L}_t : T_S\mathbf{T}_g \times H^0(S, TS)^n \to \mathbb{R}$ for the harmonic map $h_t = (h_1^t, \dots, h_n^t)$ in the same way as $\mathbf{L}$: if $r \mapsto S_r$ is a path of Riemann surfaces tangent to a Beltrami form μ at $r = 0$, and $V_1, \dots, V_n$ are vector fields giving rise to flows $\mapsto f_1^r, \dots, f_n^r$, then we set

$$\mathbf{L}_t(\mu, V_1, \dots, V_n) = \frac{d^2}{dr^2}|_{r=0}\mathcal{E}_\rho^t(S_r, h_r^t),$$

where $h_t^r = (h_1^t \circ f_1^r, \dots, h_n^t \circ f_n^r)$.

Remark 7.3.4. The two-variable energy for $\prod_{i=1}^n M_i^t$ is defined on $\mathbf{T}_g \times \prod_{i=1}^n \text{Maps}(S, M_i^t)$. Since any small perturbation of the identity map is a diffeomorphism, the space on which $\mathbf{L}_t$ acts is canonically isomorphic to the tangent space of $\mathbf{T}_g \times \prod_{i=1}^n \text{Maps}(S, M_i^t)$ at $S \times \prod_{i=1}^n \text{id}$.

Remark 7.3.5. Since S is a critical point for the two-variables energies $\mathcal{E}_\rho^t$ and $\mathcal{E}_\rho$, the second order derivatives $\mathbf{L}_t$ and $\mathbf{L}$ depend only on the first order data $\mu, V_1, \dots, V_n$ (this is not true for the second variations of each component harmonic map).

Proposition 7.3.6. *For $s > t$,*

$$\frac{\mathbf{L}_t}{t} > \frac{\mathbf{L}_s}{s} > \mathbf{L},$$

and $\lim_{t \to \infty} \frac{\mathbf{L}_t}{t} = \mathbf{L}$.

Proof. We invoke Lemma 7.3.2 n times. □

Lemma 7.3.7. *The index of* $\mathbf{L}_t$ *is equal to the index of* $\mathbf{E}^t_\rho$.

Proof. Let $r \mapsto S_r$ be a path of Riemann surfaces, tangent to the Beltrami form μ at $r = 0$, and suppose there exists $(V_1, \ldots, V_n) \in H^0(S, TS)^n$ such that $\mathbf{L}_t(\mu, V_1, \ldots, V_n) < 0$. For each fixed $t > 0$, the maps $r \mapsto \mathbf{E}^t_\rho(S_r)$ and $r \mapsto \mathcal{E}^t_\rho(S_r, h^r_t)$ have zero first variation at $r = 0$. By the minimizing property for harmonic maps, $\mathbf{E}^t_\rho(S_r) \leq \mathcal{E}^t_\rho(S_r, h^r_t)$ for every r, and it follows that

$$\frac{1}{t}\frac{d^2}{dr^2}|_{r=0}\mathbf{E}^t_\rho(S_r) \leq \frac{1}{t}\frac{d^2}{dr^2}|_{r=0}\mathcal{E}^t_\rho(S_r, h^r_t) = \frac{1}{t}\mathbf{L}_t(\mu, V_1, \ldots, V_n) < 0.$$

So, the index of $\mathbf{E}^t_\rho$ is at least that of $\mathbf{L}_t$.

For the other direction, assume $r \mapsto S_r$ lowers $\mathbf{E}^t_\rho$ to second order, and for each $r > 0$, let $k = (k^r_1, \ldots, k^r_n) : S_r \to \prod_{i=1}^n M^t_i$ be the harmonic map in the class of the identity. All h^t_i's and k^r_i's are orientation-preserving diffeomorphisms. Set $f^r_i = (h^t_i)^{-1} \circ k^r_i$ and let V_i be the infinitesimal generator of the flow $r \mapsto f^r_i$. Then

$$\frac{1}{t}\mathbf{L}_t(\mu, V_1, \ldots, V_n) = \frac{1}{t}\frac{d^2}{dr^2}|_{r=0}\mathbf{E}^t_\rho(S_r) < 0,$$

which gives the result. □

We now deduce Theorem 7B2.

Proof of Theorem 7B2. Proposition 7.3.6 implies that the index of $\mathbf{L}_t$ is non-decreasing with t, and converges to the self-maps index of S for $\mathbf{E}_\rho$. We then apply Lemma 7.3.7 to obtain the same statement for the index of $\mathbf{E}^t_\rho$ at S. □

Proof of Theorem 7B1

The proof of Theorem 7B1 is similar to that of Theorem 7B2, so we don't go through every detail. The main difference is that we replace Lemma 7.3.2 with Lemma 7.3.8 below.

As above, let M_t be the hyperbolic structure on Σ_g with hyperbolic metric σ_t such that the identity map has Hopf differential $t\phi$, with energy functional E^t_ρ, and let E_ρ be the energy functional for the $\mathbb{R}$-tree $(T, 2d)$ for ϕ.

Lemma 7.3.8. *For all Riemann surfaces* S',

$$\lim_{t\to\infty}\frac{E^t_\rho(S')}{t} = E_\rho(S').$$

In order to prove the lemma, we recall some facts about the Thurston compactification of Teichmüller space. Let $\mathcal{S}$ be the set of homotopically non-trivial simple closed curves on Σ_g and $\mathbb{R}^{\mathcal{S}}$ the product space with the weak topology. There is an embedding

$$\ell : \mathbf{T}_g \times \mathbb{R}^+ \to \mathbb{R}^{\mathcal{S}}$$

that associates the data of a hyperbolic metric σ and $s \in \mathbb{R}^+$ to the set of lengths of geodesic representatives of curves in $\mathcal{S}$ with respect to the metric $s\sigma$. Every singular measured foliation $(\mathcal{F}, \mu)$ on S also defines a point in $\mathbb{R}^{\mathcal{S}}$, by taking μ-transverse measures of simple closed curves. Furthermore, there is an injective map

$$\beta : \mathrm{QD}(S) \to \mathbb{R}^{\mathcal{S}}$$

that takes a quadratic differential to its vertical foliation, and then to $\mathbb{R}^{\mathcal{S}}$. Note that both ℓ and β are homogeneous with exponent $\frac{1}{2}$.

According to Thurston and Hubbard-Masur (see [FLP12] and [HM79]), both ℓ and β are homeomorphisms onto their images, and $\ell(\mathbf{T}_g \times \mathbb{R}^+) \sqcup \beta(QD(S))$ is homeomorphic to a cone over a closed ball, which we call C (the cone over the Thurston compactification of Teichmuller space). The following result can be gleaned from the results of [Wol89].

Theorem 7.3.9. *For any Riemann surface S, let $E_S : \mathbf{T}_g \times \mathbb{R}^+ \sqcup QD(S) \to \mathbb{R}^+$ be the function that associates to each point in $\mathbf{T}_g \times \mathbb{R}^+$ the energy of the unique harmonic map isotopic to the identity from S, and to each point of $QD(S)$ the energy of the unique equivariant harmonic map to the corresponding $\mathbb{R}$-tree. Then E_S is continuous with respect to the topology on C.*

We now explain how to deduce this theorem from the paper [Wol89]. The first ingredient is a de-projectivized version of Lemma 4.7 of that paper, whose proof is identical to the proof of the lemma in the paper.

Lemma 7.3.10. *Suppose $(\lambda_n)_{n=1}^{\infty} \subset \mathbf{T}_g$ leaves all compact subsets of the Teichmüller space, and let ϕ_n be the Hopf differential of the harmonic map from S to (S, λ_n). Suppose $(a_n)_{n=1}^{\infty} \subset (\mathbb{R}^+)^{\mathcal{S}}$ is a chosen sequence. Then $\ell(\lambda_n, a_n)$ converges in $\mathbb{R}^{\mathcal{S}}$ if and only if $\beta(a_n\phi_n)$ does, and in the case of convergence, the two sequences have the same limit.*

The second ingredient is the following computation (in which each term is linear in the scalars a_n, so the factors of a_n are superfluous).

Lemma 7.3.11 (Lemma 3.2 in [Wol89])**.** *In the notation of the previous lemma,*

$$||a_n\phi_n||_{L^1(S)} \leq E_S(a_n\lambda_n) \leq ||a_n\phi_n||_{L^1(S)} + a_n|\chi(\Sigma_g)|.$$

Proof of Theorem 7.3.9. For brevity, write $E = E_S$. First, E is continuous on $\mathbf{T}_g \times \mathbb{R}^+$ and $E(\phi) = ||\phi||_{L^1(S)}$, which is certainly continuous on $QD(S)$. To show that E is continuous on all of C, we just need to show that if $\ell(\lambda_n, a_n) \to \beta(\phi)$, then $E(a_n\lambda_n) \to E(\phi)$. By Lemma 7.3.10, $\beta(a_n\phi_n) \to \beta(\phi)$ (where as above ϕ_n is the Hopf differential of the harmonic map to λ_n), and since β is a homeomorphism onto its image, $a_n\phi_n \to \phi$ as well, so $E(a_n\phi_n) \to E(\phi)$. Finally, since a_n must tend to zero in order for the sequence $a_n\lambda_n$ to converge in $\mathbb{R}^{\mathcal{S}}$, Lemma 7.3.11 shows that $E(a_n\phi_n)$ and $E(a_n\lambda_n)$ have the same limit. □

Now the proof of Lemma 7.3.8 is easy.

Proof of Lemma 7.3.8. By definition, $E^t_\rho(S') = E_{S'}(\sigma_t)$, and $E_\rho(S') = E_{S'}(\phi)$, so by the continuity of $E_{S'}$ and its homogeneity, we just need to show that $\ell(\sigma_t/t) \to \beta(\phi)$ in C. To prove this, we use Lemma 7.3.10 applied to the surface S. Indeed, the Hopf differential of the harmonic map from S to σ_t/t is ϕ by construction, and since the constant sequence at ϕ trivially converges to ϕ, Lemma 7.3.10 implies that $\ell(\sigma_t/t)$ does as well. □

Preparations aside, we prove Theorem 7B1. We return to all of the notation from the introduction and the proof of Theorem B2. We don't recall it in full, but just record here that the energy functionals are $\mathbf{E}_\rho$ for the product of $\mathbb{R}$-trees and $\mathbf{E}^t_\rho$ for the product of surfaces. The proof is quite similar to that of Theorem 7B2, so we leave the details of the computations to the reader.

Proof of Theorem 7B1. Beginning with a Riemann surface S' such that $\mathbf{E}_\rho(S') < \mathbf{E}_\rho(S)$, applying Lemma 7.3.8 n times yields that $\mathbf{E}^t_\rho(S') < \mathbf{E}^t_\rho(S)$ for sufficiently large t.

Conversely, suppose that there exists $t > 0$ such that $\mathbf{E}^t_\rho(S') < \mathbf{E}^t_\rho(S)$, and let $k = (k^t_1, \ldots, k^t_n) : S' \to \prod_{i=1}^n M^t_i$ be the n-tuple of harmonic diffeomorphisms with lower energy. Let h^t_i be the i^{th} component of the harmonic map h_t, and set $f^t_i = (h^t_i)^{-1} \circ k^t_i$. Arguing similarly to the proof of Lemma 7.3.2, Reich-Strebel formulas (7.5) and (7.6) and the monotonicity on the level of energy densities from Lemma 7.3.3 show that for $s > t$,

$$\begin{aligned}\frac{\mathbf{E}^t_\rho(S') - \mathbf{E}^t_\rho(S)}{t} &= \frac{\sum_{i=1}^n \mathcal{E}^t_\rho(S', h^t_i \circ f^t_i) - \mathbf{E}^t_\rho(S)}{t} \\ &> \frac{\sum_{i=1}^n \mathcal{E}^s_\rho(S', h^s_i \circ f^t_i) - \mathbf{E}^s_\rho(S)}{s} \\ &> \sum_{i=1}^n \mathcal{E}_\rho(S', \pi \circ \tilde{f}^t_i) - \mathbf{E}_\rho(S).\end{aligned}$$

It follows from the minimizing property that

$$\frac{\mathbf{E}^t_\rho(S') - \mathbf{E}^t_\rho(S)}{t} > \frac{\mathbf{E}^s_\rho(S') - \mathbf{E}^s_\rho(S)}{s}$$

and

$$\frac{\mathbf{E}^t_\rho(S') - \mathbf{E}^t_\rho(S)}{t} > \mathbf{E}_\rho(S') - \mathbf{E}_\rho(S),$$

and hence the result follows. □

7.4 Unstable equivariant minimal surfaces in $\mathbb{R}^n$

We recall the setup of Theorem 7C. For $n \geq 2$ and $i = 1, \ldots, n$, let α_i be nonzero holomorphic 1-forms on S such that $\sum_{i=1}^n \alpha_i^2 = 0$. Let χ be the action of $\pi_1(S)$ on $\mathbb{R}^n$ corresponding to the 1-forms α_i, and let ρ be the action of $\pi_1(S)$ on a product $X = \prod_i (T_i, 2d_i)$ of trees corresponding to the quadratic differentials $\phi_i = \alpha_i^2$. We write $\mathcal{E}_\chi$ and $\mathcal{E}_\rho$ for the associated two-variable energies, and $\mathbf{E}_\chi$ and $\mathbf{E}_\rho$ for the energy functionals on Teichmüller space. Let $h = (h_1, \ldots, h_n)$ and $\pi = (\pi_1, \ldots, \pi_n)$ be the χ- and ρ- equivariant minimal maps respectively.

Isometric folding

We begin with the statement (1) from Theorem 7C. The result is a consequence of the proposition below.

Proposition 7.4.1. $\mathbf{E}_\rho \geq \mathbf{E}_\chi$, *with equality at S.*

The key is that there is a natural map $F : X \to \mathbb{R}^n$ intertwining ρ and χ. To see why, let's focus on a single tree T_i. Along a curve parametrizing a non-singular leaf for the vertical singular foliation of ϕ_i, α_i evaluates the tangent vectors to purely imaginary numbers. Since $dh_i = \mathrm{Re}(\tilde{\alpha}_i)$, we deduce that h_i is constant along the singular vertical foliation of $\tilde{\phi}_i$. Hence, h_i descends to a map $F_i : T_i \to \mathbb{R}$, which we call the folding map of the tree. The map $F = (F_1, \ldots, F_n)$ has the required equivariance.

Lemma 7.4.2. *If S' is any point of $\mathbf{T}_g$, and π_i' the unique ρ_i-equivariant harmonic map from $\tilde{S}'$ to $(T_i, 2d_i)$, then the energy density of π_i' is pointwise equal to the energy density of $F_i \circ \pi_i'$.*

Proof. Let ψ_i be the Hopf differential of π_i'. As discussed in Section 2, for any point p at which $\psi_i(p) \neq 0$, there exists a neighbourhood Ω of p, an open interval $I \subset \mathbb{R}$, a map $\hat{\pi}_i' : \Omega \to I$, and an isometric inclusion $\iota : I \to (T_i, 2d_i)$ such that in Ω,

$$\pi_i' = \iota \circ \hat{\pi}_i'.$$

By construction, the restriction of $F_i|_{\iota(I)} : \iota(I) \to \mathbb{R}$ is an isometric embedding. It follows by continuity that the energy densities are equal everywhere. □

Proof of Proposition 7.4.1. For any $S' \in \mathbf{T}_g$, let π' be the ρ-equivariant harmonic map to the product of trees. The map $F \circ \pi'$ is a χ-equivariant Lipschitz map to $\mathbb{R}^n$. By the minimizing property for harmonic maps,

$$\mathbf{E}_\chi(S') \leq \mathcal{E}_\chi(S', F \circ \pi').$$

By the lemma above, $\mathcal{E}(S', F \circ \pi') = \mathbf{E}_\rho(S')$, so we have

$$\mathbf{E}_\rho \geq \mathbf{E}_\chi.$$

Working on the Riemann surface S, $h_i = F_i \circ \pi_i$ for every i, so $\mathbf{E}_\chi(S) = \mathcal{E}_\chi(S', F \circ \pi)$, and we have equality. □

Remark 7.4.3. Maps of the form $F_i \circ \pi_i'$ above are subtle. They are harmonic apart from some preimages under π_i' of the vertices in $(T_i, 2d_i)$, which are typically disjoint arcs or connected sums of disjoint arcs. Even though they have finite total energy, a Weyl lemma cannot be applied because they fail to be twice weakly differentiable on these lines. The map $x \mapsto |x|$ on $\mathbb{R}$ exhibits this type of behaviour.

We see immediately from Proposition 7.4.1 that if S is not a global (resp. local) minimum of $\mathbf{E}_\rho$, then it is not a global (resp. local) minimum of $\mathbf{E}_\chi$. So (1) is proved. Furthermore, we are very close to proving one direction of (2), once we recall the definition of the self-maps index, and its basic properties. We do this after collecting some standard facts about minimal surfaces in $\mathbb{R}^n$.

Energy and area

Let f be any smooth χ-equivariant map from $\tilde{\Sigma}_g$ to $\mathbb{R}^n$. The differential of f descends to a closed $\mathbb{R}^n$-valued 1-form θ on Σ_g, and the cohomology class of θ is prescribed by the representation χ. The map f also defines a $\pi_1(\Sigma_g)$-invariant area form $dA_f = \sqrt{\det(\theta^T\theta)}$, and the area of f, which we write $A(f)$, is defined to be the integral of this form over Σ_g. If S is a Riemann surface structure on Σ_g, then $\mathcal{E}_\chi(S, f) \geq A(f)$, with equality if and only if f is minimal (in fact, the integrands are equal pointwise).

Now suppose we are in the setting of Theorem 7C, so that h is a minimal χ-equivariant map from $\tilde{S}$ to $\mathbb{R}^n$. Let B be the branch locus of h on S.

Lemma 7.4.4. *Let h_r be a smooth χ-equivariant variation of h such that $h_r = h$ in a neighborhood of B. Then for r small enough, there is a smooth variation of Riemann surface structures S_r such that h_r is minimal with respect to S_r.*

Proof. For r sufficiently small, the map h_r is still an immersion away from B, and hence uniquely defines a new conformal structure on $S - B$. Since X is compactly supported away from B, this conformal structure patches to the conformal structure of S near B, and defines a new conformal structure S_r on S, with respect to which h is minimal. □

We say that a smooth $\mathbb{R}^n$-valued vector field W on S supported on $S - B$ is a normal variation of h if it is perpendicular to the image of dh at each point of $S - B$. For any such W, let $\tilde{W}$ be the pullback to $\tilde{S}$; then the family $h_r = h + r\tilde{W}$ is a χ-equivariant deformation of h equal to h on a neighborhood of B. Taking the derivative of the corresponding S_r at $r = 0$ defines a linear map from the space of normal variations supported on $S - B$ to the tangent space of Teichmüller space at S. Let V the graph of this map, viewed as a subspace of $T_S\mathbf{T}_g \times T_h\mathrm{Map}_\chi(\tilde{\Sigma}_g, \mathbb{R}^n)$. We have shown that restricted to V, the Hessian of $\mathcal{E}_\chi$ at the critical point (S, h) is equal to the Hessian of A at the critical point h. The latter has the following formula:

Proposition 7.4.5 (Theorem 32 in [Law80])**.** *If W is a normal variation supported in $S - B$, the second derivative of the area of any equivariant variation h_r with derivative W at $r = 0$ is given by the quadratic form*

$$Q(W) = \int_S |(dW)^N|^2 - |\langle k, W\rangle|^2 \tag{7.10}$$

where $(dW)^N$ is the component of dW normal to the image of dh, k is the vector-valued second fundamental form of $h(S)$, and the second term is interpreted as the square norm of the scalar-valued 2-tensor $\langle k, W\rangle$.

Lifting to $\mathbb{R}$-trees via self-maps

In this section, we study energy and area in the context of the ρ-equivariant harmonic maps to products of $\mathbb{R}$-trees. Specifically, we relate Q to the quadratic form $\mathbf{L} : T_S\mathbf{T}_g \times H^0(S;TS)^n \to \mathbb{R}$ defined in the introduction, which defines the self-maps index for $\mathbf{E}_\rho$. Let $H^0_c(S-B, TS)$ be the subspace of $H^0(S, TS)$ of smooth vector fields supported on $S-B$.

The key to the proof of the second part of Theorem 7C is the result below.

Lemma 7.4.6. *Suppose that W is a normal variation of S with support in $S-B$, and such that $Q(W) < 0$. Then there exists a harmonic Beltrami form μ on S and vector fields $V_1, \dots, V_n \in H^0_c(S-B, TS)^n$ such that*

$$\mathbf{L}(\mu, V_1, \dots, V_n) < 0.$$

Proof. Denote the coefficients of W by W^i. For each $i = 1, \dots, n$, let V_i be the vector field which vanishes on B and is equal to $W^i\nabla x^i/|\nabla x_i|^2$ on $S-B$, where ∇x^i is the gradient of the coordinate function x^i on S, which is nonvanishing on $S-B$. We point out that V_i has compact support on $S-B$.

Let $f_i^W : \mathbb{R}\times S \to S$ be flow of V_i, so that $f_i^W(r,\cdot) = f_i^r(\cdot)$. Then the family $H : \mathbb{R}\times\tilde{S} \to \mathbb{R}^n$ defined by $H_i(r,p) = h_i \circ f_i^r(p)$ has derivative W at time zero. Moreover, the family $\Pi : \mathbb{R}\times\tilde{S} \to \prod_i(T_i, 2d_i)$ defined by $\Pi_i(r,p) = \pi_i \circ f_i^r(p)$ satisfies $F_i \circ \Pi_i = H_i$, where F_i is the folding map from T_i to $\mathbb{R}$. Let π_r be the map $(\pi_1 \circ f_1^r, \dots, \pi_n \circ f^r)$. By Lemma 7.4.4, there exists a C^∞ variation of conformal structures $r \mapsto S_r$ along which $\mathcal{E}_\rho(S_r, \pi_r) = \mathcal{E}_\chi(S_r, h_r) = A(h_r)$, and we set μ to be the Beltrami form in $T_S\mathbf{T}_g$ tangent to this path at time zero. If $Q(W) < 0$, then taking the second variation of $r \mapsto \mathcal{E}_\rho(S_r, \pi_r)$ yields $\mathbf{L}(\mu, V_1, \dots, V_n) = Q(W) < 0$. □

Log cutoff

In order to construct destabilizing variations for Q, it is helpful to do away with the condition that W is supported on $S-B$. First, we need to say what it means for W to be a normal variation over all of S. The map $S-B \to \mathbb{CP}^{n-1}$, which sends p to the (one-dimensional) image of $(\alpha_1, \dots, \alpha_n)$ at p, extends holomorphically to all of S by clearing denominators. Thus, the normal bundle also extends analytically to all of S. The quadratic form Q is still finite for normal variations that are not necessarily supported on $S - B$.

For normal variations W, which are not necessarily supported on $S - B$, we will need to show that one can replace them with variations that are supported on $S - B$ without changing the value of Q too much. This is the log cut-off trick. If r is the radial coordinate in $\mathbb{C}$ then the function $\log(r)/\log(\delta^{-1}) + 2$, defined between $r = \delta$ and $r = \delta^2$, is equal to 1 for $r = \delta$ and 0 for $r = \delta^2$ and has Dirichlet energy

$$\frac{1}{2}\int_{\delta^2}^{\delta} \frac{2\pi}{r\log(\delta^{-1})^2} = \frac{\pi}{\log(\delta^{-1})}. \tag{7.11}$$

The point is this this tends to zero as δ goes to zero. A good picture is that $\log r$ is the height coordinate on a cylinder conformal to the punctured disk, so our function is an affine function of the height of the cylinder, and its derivative is small. The extension of this function by 0 and 1 is Lipschitz. For very minor reasons, it will be convenient to use a smooth cutoff function, so we let $l_\delta(r)$ be a perturbation of $\log(r)/\log(\delta^{-1}) + 2$ which extends smoothly by 0 and 1 and has Dirichlet energy no more than $2\pi/\log(\delta^{-1})$.

We use this model to define a cut-off function as follows. For each point p_i of B, fix a holomorphic coordinate z_i with $z_i(p_i) = 0$. Then, for any value of δ small enough that each z_i is defined on the ball of radius δ around p_i and these balls do not overlap, let η_δ be the function on S defined by

- $\eta_\delta(p) = l_\delta(|z_i|)$ if $\delta^2 \leq |z_i(p)| \leq \delta$ for some i
- $\eta_\delta(p) = 0$ if $|z_i(p)| \leq \delta^2$ for some i
- $\eta_\delta(p) = 1$ otherwise.

We now use the log cut-off trick to prove the following.

Lemma 7.4.7. *Suppose that W is normal variation of h on S. Then given any $\epsilon > 0$, there is a constant $d(\epsilon, Q(W), \sup|W|)$ such that for all $\delta < d$,*

$$|Q(\eta_\delta W) - Q(W)| < \epsilon. \tag{7.12}$$

Proof. For δ to be determined, we compute $Q(\eta_\delta W)$. We first treat the normal term in the formula (7.10) applied to the variation $\eta_\delta W$:

$$\begin{aligned}\int_\Sigma |(d(\eta_\delta W))^N|^2 &= \int_\Sigma |\eta_\delta (dW)^N + (W d\eta_\delta)^N|^2 \\ &= \int_{\delta^2 \le |z| \le \delta} |\eta_\delta (dW)^N + (W d\eta_\delta)^N|^2 + \int_{|z|\ge\delta} |(dW)^N|^2,\end{aligned}$$

where $W d\eta_\delta$ is the $\mathbb{R}^n$-valued 1-form $W \otimes d\eta_\delta$. Hence,

$$\begin{aligned}|Q(\eta_\delta W) - Q(W)| &\le \int_{\delta^2 \le |z| \le \delta} |\eta_\delta (dW)^N + (W d\eta_\delta)^N|^2 + \int_{|z|\le\delta} |(dW)^N|^2 + \int_S (1-\eta_\delta^2)|W|^2|k|^2 \\ &= \int_{\delta^2 \le |z| \le \delta} |\eta_\delta (dW)^N + (W d\eta_\delta)^N|^2 + O(\delta^2),\end{aligned}$$

since $1 - \eta_\delta^2$ is supported in $|z| \le \delta$. By Cauchy-Schwarz and (7.11),

$$\begin{aligned}|Q(\eta_\delta W) - Q(W)| &\le \int_{\delta^2 \le |z| \le \delta} |\eta_\delta (dW)^N|^2 + |(W d\eta_\delta)^N|^2 + 2|\eta_\delta (dW)^N|^2 |(W d\eta_\delta)^N|^2 + O(\delta^2) \\ &= O(\delta^2) + O\Big(\frac{1}{\log \delta^{-1}}\Big) = O\Big(\frac{1}{\log \delta^{-1}}\Big).\end{aligned}$$

Thus, we can choose $\delta > 0$ so that the difference of second variations is at most ϵ. □

An immediate consequence is that we can speak without ambiguity of the index of Q.

Proposition 7.4.8. *The index of Q on the space of all normal variations is equal to the index of Q on the subspace of normal variations supported in $S - B$.*

Proof. We just need to show that if there is a k-dimensional space of normal variations on which Q is negative definite, then there is another k-dimensional space of normal variations supported in $S - B$ on which Q is still negative definite. Let V be a k-dimensional space of normal variations on which Q is negative definite. Let $S(V)$ be the unit sphere in V with respect to any metric on V. Then for δ small enough, $Q(\eta_\delta W) < 0$ for every $W \in S(V)$. Since this implies $\eta_\delta W \neq 0$, the space $\{\eta_\delta W | W \in V\}$ is a k-dimensional subspace of normal variations supported in $S - B$ on which Q is negative definite. □

We may now finish the proof of Theorem 7C.

Proof of Theorem 7C (2). Let k be the index of $\mathbf{E}_\chi$, and let $W \subset T_S\mathbf{T}_g$ be a k-dimensional subspace on which the second variation is negative definite. By the implicit function theorem, the unique harmonic 1-form in a given cohomology class varies smoothly with the conformal structure of S. We can integrate this smoothly-varying 1-form to give a smooth equivariant variation of the harmonic map h. Projecting the variation onto the normal bundle, we get from W a vector space of normal variations of h on which the second derivative of $\mathbf{E}_\chi$ is equal to Q. Since it is assumed to be positive definite, this space is still k dimensional.

By Proposition 7.4.8, we can replace this with a k-dimensional subspace of normal variations supported on $S - B$ on which Q is still negative definite. Then by Lemma 7.4.6, there is a k-dimensional subspace of $T_S\mathbf{T}_g \times H^0_c(S - B, TS)^n$ on which $\mathbf{L}$ is negative definite, and so the index of $\mathbf{E}_\rho$ by self-maps is at least k.

In the other direction, suppose W' is a k-dimensional subspace of $T_S\mathbf{T}_g \times H^0(S, TS)^n$ on which $\mathbf{L}$ is negative definite. Since $\mathbf{L}$ is positive semidefinite on $\{0\} \times H^0(S, TS)^n$, the projection of W' to $T_S\mathbf{T}_g$ is still k-dimensional. For maps to manifolds, the positive semidefinite property follows from the computation [Har67, Theorem H], and we get the same result in our setting by repeating the computation but using the measurable energy density with the characterization (7.2). Since $\mathbf{E}_\rho$ is an infimum over all maps, we get an upper bound for $\mathbf{E}_\rho$ near S by a smooth function with negative definite Hessian at S. Recall that Proposition 7.4.1 says that $\mathbf{E}_\chi \leq \mathbf{E}_\rho$, and so the index of $\mathbf{E}_\chi$ at S is at least k. □

7.5 The general case

In this section, we generalize Theorem 7C to the situation in which the quadratic differentials are not necessarily squares of abelian differentials. We then specialize to dimension 3 and give the proof of Theorem 7D.

The spectral curve

Let S_0 be a point of $\mathbf{T}_g$, and let $\phi_1, \ldots, \phi_n$ be nonzero holomorphic quadratic differentials on S_0 summing to zero. To this data, there is an associated spectral curve. This is a particular branched covering S of S_0 with abelian differentials α_i on S that square to the pullback of ϕ_i. It is always a 2^n-fold branched covering of S_0, but may be disconnected, for instance if any ϕ_i is already a square. By universality, S has n holomorphic involutions τ_i, each of which negates α_i and fixes α_j for $j \neq i$.

We let ρ be the action of $\pi_1(S_0)$ on the product X of the n $\mathbb{R}$-trees $(T_i, 2d_i)$ corresponding to the quadratic differentials ϕ_i, and π the canonical equivariant map from $\tilde{S}_0$ to X.

Since S has n abelian differentials whose squares sum to zero, the theory of the previous section applies. For instance, we can integrate $\mathrm{Re}(\tilde{\alpha}_i)$ on a simply connected covering space to get a harmonic map h to $\mathbb{R}^n$, equivariant under a representation χ of the Deck group, and well defined up to a constant on each component of S. The energy density of this map descends not only to S, but all the way to S_0, where it is equal to the energy density of π.

In the spirit of Proposition 7.4.8, we want to compare the index of $\mathbf{E}_\rho$ through self-maps at S_0 to the index of the quadratic form Q associated to h. But to get the right comparison, we need to restrict Q to a subspace of the space of normal variations. Let $G \cong (\mathbb{Z}/2\mathbb{Z})^n$ be the group generated by the τ_i. Let σ be the action of G on $\mathbb{R}^n$ such that each τ_i acts by reflection in the ith coordinate hyperplane. Let NV^σ be the space of normal variations of h that are σ-equivariant.

Proposition 7.5.1. *The index of* $\mathbf{E}_\rho$ *by self-maps is equal to the index of* Q *on* NV^σ.

Proof. Let k be the index of Q on NV^σ. The first thing we want to do is use Proposition 7.4.8 to find a k-dimensional space of σ-equivariant normal variations on $S - B$ on which Q is still negative definite. This works fine if we choose our cutoff function η_δ to be τ_i-invariant. For instance, we can define η_δ to be the pull-back to S of the similarly-defined function on S_0; then the dependence of the energy of η_δ with δ is the same up to a factor of 2 coming from the relation $\log(|\sqrt{z}|) = \log(|z|)/2$.

Next, for every W in this space, we get n tangential vector fields $V_i = W^i \nabla x^i / |\nabla x^i|^2$ on S, as in Proposition 7.4.6. Since both W^i and dx^i transform the same way under each τ_j, we have $\tau_j(V_i) = \pm W^i(\pm \nabla x^i)/|\nabla x^i|^2$, where each sign is $+$ if $i \neq j$ and $-$ if $i = j$. Hence each V_i descends to a vector field on $S_0 - B$, which we still call V_i.

For each i, let $f_i^W : \mathbb{R} \times S \to S$, $f_i^W(r, \cdot) = f_i^r(\cdot)$ be the flow of V_i. Let h_i be the component functions of h, $h_i^r = h_i \circ f_i^r$, and $h_r = (h_1^r, \ldots, h_n^r)$. The conformal structures S_r for which each h_r is conformal are still G-invariant, hence descend to conformal structures $(S_0)_r$ on Σ_g. Let $\pi_i^r = \pi \circ f_i^r$. Even though the tree T_i no longer folds to $\mathbb{R}$, the energy density of π_i^r on $(S_0)_r$ is still pointwise equal to the energy density of h_i^r on S^t; indeed, both are equal to $|(f_i^r)^* \mathrm{Re}(\alpha_i)|^2$. Therefore, the second derivative of $\mathcal{E}_\rho$ is equal to Q on this k-dimensional space so the index of $\mathbf{E}_\rho$ by self-maps is at least k.

The other inequality is easier. If the index of $\mathbf{E}_\rho$ by self-maps is k, then we can use the log-cutoff trick to find a k-dimensional space of vector fields V_i supported on $S_0 - B$ and variations μ_i of conformal structure on which L is negative definite. Lifting everything to S and differentiating the coordinate functions, we get a k-dimensional space of equivariant variations of h for which the second derivative of energy is negative definite. Taking the normal components of these variations, and using that energy dominates area, we get a k-dimensional subspace of NV^σ on which Q is negative definite. □

Unstable minimal surfaces in $\mathbb{R}^n$

In order to finish the proof of Theorem 7A, we need to construct for each $g \geq 2$ and $n \geq 3$, either an unstable equivariant minimal surface S of genus g in $\mathbb{R}^n$, or a surface S_0 of genus g whose spectral curve is a $(\mathbb{Z}/2\mathbb{Z})^n$-equivariantly unstable minimal surface in $\mathbb{R}^n$.

If $g \geq 3$, then as we discuss in the next section, there are plenty of equivariant minimal surfaces of genus g in $\mathbb{R}^n$. They are not always unstable; for instance, if the minimal map is holomorphic with respect to some complex structure on a linear subspace of $\mathbb{R}^n$, then it is calibrated by the Kähler form, and hence stable. In general, it is not straightforward to decide if a minimal surface in a flat space is unstable.

A special case is when the equivariant minimal surface is contained in a real 2-plane, and hence is stable. We call such a minimal surface flat. These at least are easy to identify.

Proposition 7.5.2. *Let $\phi_1, \ldots, \phi_n \in QD(S_0)$ sum to 0, giving a χ-equivariant map $h : \tilde{S} \to \mathbb{R}^n$ as before. The vector valued second fundamental form k of $h(\tilde{S})$ vanishes identically if and only if the quadratic differentials ϕ_i are all complex multiples of one another.*

Proof. Let $h_1, \ldots, h_n$ denote the coordinate functions of h. Since $\phi_i = ((h_i)_z)^2 dz^2$, the quadratic differentials are all complex multiples of one another if and only if the functions $(h_i)_z$ are. In one direction, assume $(h_i)_z = a_i f(z)$ for some function $f(z)$ and some complex constants a_i. Then the image of the $\mathbb{R}^n$-valued 1-form dh is contained in a two-dimensional subspace, and by integrating, we see that image of h is contained in an affine subspace of $\mathbb{R}^n$. In particular, it is totally geodesic, so the second fundamental form is zero. Conversely, if the second fundamental form is zero, then the image of dh is contained in some two-dimensional linear subspace, and so the image of h_z is contained in the complexification of that subspace, which is two-dimensional. As h is weakly conformal, $\langle h_z, h_z \rangle = 0$; since the inner product is nondegenerate on the complexification of any real two-dimensional subspace, this shows that h_z is contained in a complex line (we use analyticity to deduce this as well at the branch points), and so the functions $(h_i)_z$ are all complex multiples of one another. □

For the remaining section, we restrict to $n = 3$. For $n \geq 3$, any isometric inclusion of $\mathbb{R}^3$ into $\mathbb{R}^n$ gives examples in $\mathbb{R}^n$. Let $\mathbf{M}_g$ be the moduli space of Riemann surfaces of genus g, and let E^n be the total space of the bundle over $\mathbf{M}_g$ consisting of n-tuples of quadratic differentials that sum to 0. Instability of the corresponding equivariant minimal surfaces in $\mathbb{R}^n$ is an open condition on E^n, so by perturbing the 3-dimensional examples we get many more.

Equivariant minimal surfaces in $\mathbb{R}^3$

Every non-flat equivariant minimal surface in $\mathbb{R}^3$ is unstable. Indeed, in dimension 3, the expression $|\langle k, W \rangle|^2$ in the formula for $Q(W)$ is equal to $2|K||W|^2$, where K is the Gauss curvature of the equivariant minimal surface. The normal bundle to the minimal surface $h(\tilde{S})$ is a real line bundle on S. Since S is always orientable, the normal bundle is as well, and hence it is equivariantly trivial. If N is a unit normal section, and η is any function on S, then the second variation formula (7.10) takes the form

$$Q_0(\eta N) = \int_\Sigma |\nabla \eta|^2 - |K|^2 |\eta|^2.$$

As long as the curvature K is anywhere nonzero, a constant section of N will therefore be destabilizing: for $\eta = 1$,

$$Q_0(N) = \int_\Sigma -|K|^2 < 0.$$

When $g \geq 3$, the moduli space of $(S, \alpha_1, \alpha_2, \alpha_3)$, where S is a Riemann surface of genus g and α_i are abelian differentials on S whose squares sum to zero, but are not all mutliples of one another, is nonempty and has complex dimension $3g$ (in [FR20, Section 6] it is shown that the quotient by the natural free actions of $\mathbb{C}^*$ and $\mathrm{SO}(3, \mathbb{C})$ has dimension $3g - 4$). This proves Theorem 7A for $g \geq 3$.

Remark 7.5.3. In fact, in [Ros06, Theorem 16], Ros proves that every non-flat minimally immersed surface of genus g in a 3-torus has index at least $2g/3 - 1$.

The result easily generalizes to any non-flat equivariant minimal immersion for any representation, but we emphasize that it applies only to immersed surfaces.

Unfortunately, there are no non-flat equivariant minimal surfaces of genus 2 in $\mathbb{R}^3$, stable or not. This is because the canonical map lands in $\mathbb{P}^1$, so the canonical curve cannot be contained in a rank 3 quadric (or see the comment after Proposition 7.5.6). Hence, we are forced to study σ-equivariant deformations of the spectral curve. The key that makes this work is that the normal bundle of the spectral curve S of $(S_0, \phi_1, \phi_2, \phi_3)$ can be equivariantly trivial even if the ϕ_i are not squares (in which case S is just 8 copies of S_0).

Proposition 7.5.4. *Suppose that the sextic differential $\phi_1\phi_2\phi_3$ is the square of a cubic differential c. Then there is a σ-equivariant deformation of S of constant length 1.*

Proof. The cubic differential c distinguishes two components of S; one on which $\alpha_1\alpha_2\alpha_3 = c$, and one on which it is equal to $-c$. Each τ_i interchanges the two components of S. The subgroup $\Gamma < (\mathbb{Z}/2\mathbb{Z})^3$ preserving the components acts on $\mathbb{R}^3$ in an orientation-preserving way. Indeed, for each element $\gamma \in \Gamma$, the determinant of the matrix describing the product of hyperplane reflections is equal to the product of the monodromies of the α_i under the action of γ. We can use the orientation of $\mathbb{R}^3$, together with the orientation of the component of S, to equivariantly orient the normal bundle. Since the normal bundle is a line bundle, it therefore has an equivariant section of constant length. □

Remark 7.5.5. If each ϕ_i is the square of an abelian differential α_i, then clearly $\phi_1\phi_2\phi_3 = c^2$ with $c = \alpha_1\alpha_2\alpha_3$.

If the quadratic differentials ϕ_i are not complex multiples of one another, then neither are their lifts α_i to the spectral curve. Hence, the minimal map from the lift of the spectral curve is non-flat, so any σ-equivariant deformation of constant length will be destabilizing.

The final step is to show that there are non-flat solutions even in genus 2 to the equations $\phi_1\phi_2\phi_3 = c$ and $\phi_1 + \phi_2 + \phi_3 = 0$. For any $g \geq 2$, let $\mathcal{P}_g$ be the moduli space of genus g Riemann surfaces S together with a triple of quadratic differentials ϕ_i summing to zero whose product is a square and which are not all complex multiples of one another.

Proposition 7.5.6. *The moduli space $\mathcal{P}_2$ has dimension* 3.

Proof. Consider the three dimensional family of algebraic curves $w^2 = z(z-1)(z-a)(z-b)(z-c)$ for $(0, 1, a, b, c)$ distinct complex numbers. This is a finite covering of the moduli space of genus 2 Riemann surfaces. Every holomorphic quadratic differential on a curve in this family is of the form $p(z)(dz)^2/w^2$ for $p(z)$ a polynomial of degree at most 2. If the roots of $p(z)$ are branch points of the curve, then the quadratic differential vanishes to order two at the corresponding point of the curve. For arbitrary a and b, and c to be determined, let

$$\phi_1 = z(z-1)\frac{dz^2}{w^2}$$
$$\phi_2 = \mu(a,b)(z-a)\frac{dz^2}{w^2},$$

where $\mu(a,b) = -b(b-1)/(b-a)$ is chosen so that $\phi_1 + \phi_2$ vanishes at b (equivalently, that the corresponding quadratic polynomial vanishes at b). A short computation shows that the other root of the polynomial for $\phi_1 + \phi_2$ is at $a(b-1)/(b-a)$, so if this happens to be the value of c, then the sextic differential $\phi_1\phi_2\phi_3$ vanishes to order two at each of the six branch points of the curve (including ∞). Hence it is the square of the cubic differential dz^3/w^2, which vanishes to order one at each of these points. Including a parameter rescaling ϕ_1, ϕ_2, and ϕ_3, this shows that $\mathcal{P}_2$ has dimension 3. ◻

For example we could take $a = -1, b = i$, and $c = -i$ to get a solution on the hyperelliptic curve $w^2 = z^5 - z$. This suffices for the proof of Theorem 7A.

Remark 7.5.7. Note that the triples of quadratic differentials in genus 2 are squares of abelian differentials since the polynomials $z(z-1)$, etc., are not squares. However, they still have even order zeros.

Together with [FR20, Section 6] and the Remark 7.5.5, this shows:

Proposition 7.5.8. *For every genus $g \geq 2$, $\mathcal{P}_g$ is nonempty and every component has complex dimension at least* $3g-3$.

We give a self-contained proof of this proposition, since it is very brief in the reference.

Proof. We have already proved this for genus 2. The canonical map of a hyperelliptic curve of genus 3 is the vanishing locus of a nondegenerate quadric on $\mathbb{CP}^2$; diagonalizing this quadric gives three abelian differentials whose squares sum to zero on the curve. By Remark 7.5.5, these give points in $\mathcal{P}_3$. Since the hyperelliptic locus has dimension 5, we get a sixth dimension from rescaling the abelian differentials. This proves the result for $g = 3$.

In general, taking unramified coverings of a point in $\mathcal{P}_2$ shows that $\mathcal{P}_g$ is nonempty for every g. To get the bound on dimension, we observe that $\mathcal{P}_g$ is, up to a double cover, the intersection in the total space of the bundle $H^0(K^3)$ over $\mathbf{M}_g$ (dimension $14(g-1)$) of the sextic differentials that are squares of cubic differenitals (dimension $8(g-1)$) and those that are the product of three independent quadratic differentials summing to zero (dimension $9(g-1)$). This gives a lower bound on the dimension of $8(g-1) + 9(g-1) - 14(g-1) = 3(g-1)$. □

BIBLIOGRAPHY

[AG71] Norman L. Alling and Newcomb Greenleaf. *Foundations of the theory of Klein surfaces*. Lecture Notes in Mathematics, Vol. 219. Springer-Verlag, Berlin-New York, 1971, pp. ix+117.

[Ahl06] Lars V. Ahlfors. *Lectures on quasiconformal mappings*. Second. Vol. 38. University Lecture Series. With supplemental chapters by C. J. Earle, I. Kra, M. Shishikura and J. H. Hubbard. American Mathematical Society, Providence, RI, 2006, pp. viii+162. ISBN: 0-8218-3644-7. DOI: 10.1090/ulect/038. URL: https://doi.org/10.1090/ulect/038.

[AL18] Daniele Alessandrini and Qiongling Li. "AdS 3-manifolds and Higgs bundles". In: *Proc. Amer. Math. Soc.* 146.2 (2018), pp. 845–860. ISSN: 0002-9939. DOI: 10.1090/proc/13586. URL: https://doi.org/10.1090/proc/13586.

[Ale19] Daniele Alessandrini. "Higgs bundles and geometric structures on manifolds". In: *SIGMA Symmetry Integrability Geom. Methods Appl.* 15 (2019), Paper 039, 32. DOI: 10.3842/SIGMA.2019.039. URL: https://doi-org.clsproxy.library.caltech.edu/10.3842/SIGMA.2019.039.

[Alt72] Hans Wilhelm Alt. "Verzweigungspunkte von H-Flächen. I". German. In: *Math. Z.* 127 (1972), pp. 333–362. ISSN: 0025-5874; 1432-1823/e.

[Alt73] Hans Wilhelm Alt. "Verzweigungspunkte von H-Flächen. II". German. In: *Math. Ann.* 201 (1973), pp. 33–55. ISSN: 0025-5831; 1432-1807/e.

[AMR88] R. Abraham, J. E. Marsden, and T. Ratiu. *Manifolds, tensor analysis, and applications*. Second. Vol. 75. Applied Mathematical Sciences. Springer-Verlag, New York, 1988, pp. x+654. ISBN: 0-387-96790-7. DOI: 10.1007/978-1-4612-1029-0. URL: https://doi-org.clsproxy.library.caltech.edu/10.1007/978-1-4612-1029-0.

[Aro57] N. Aronszajn. "A unique continuation theorem for solutions of elliptic partial differential equations or inequalities of second order". In: *J. Math. Pures Appl. (9)* 36 (1957), pp. 235–249. ISSN: 0021-7824.

[Aus64] Louis Auslander. "The structure of complete locally affine manifolds". In: *Topology* 3.suppl, suppl. 1 (1964), pp. 131–139. ISSN: 0040-9383. DOI: 10.1016/0040-9383(64)90012-6. URL: https://doi.org/10.1016/0040-9383(64)90012-6.

[Bar10] David Baraglia. *G2 geometry and integrable systems*. 2010. arXiv: 1002.1767 [math.DG].

[BH21] Yves Benoist and Dominique Hulin. "Harmonic quasi-isometric maps II: negatively curved manifolds". In: *J. Eur. Math. Soc. (JEMS)* 23.9 (2021), pp. 2861–2911. ISSN: 1435-9855. DOI: 10.4171/jems/1065. URL: https://doi.org/10.4171/jems/1065.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*. Vol. 319. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, Berlin, 1999, pp. xxii+643. ISBN: 3-540-64324-9. DOI: 10.1007/978-3-662-12494-9. URL: https://doi.org/10.1007/978-3-662-12494-9.

[Bie11] Ludwig Bieberbach. "Über die Bewegungsgruppen der Euklidischen Räume". In: *Math. Ann.* 70.3 (1911), pp. 297–336. ISSN: 0025-5831. DOI: 10.1007/BF01564500. URL: https://doi.org/10.1007/BF01564500.

[Bie12] Ludwig Bieberbach. "Über die Bewegungsgruppen der Euklidischen Räume (Zweite Abhandlung.) Die Gruppen mit einem endlichen Fundamentalbereich". In: *Math. Ann.* 72.3 (1912), pp. 400–412. ISSN: 0025-5831. DOI: 10.1007/BF01456724. URL: https://doi.org/10.1007/BF01456724.

[BIW10] Marc Burger, Alessandra Iozzi, and Anna Wienhard. "Surface group representations with maximal Toledo invariant". In: *Ann. of Math. (2)* 172.1 (2010), pp. 517–566. ISSN: 0003-486X. DOI: 10.4007/annals.2010.172.517. URL: https://doi-org.clsproxy.library.caltech.edu/10.4007/annals.2010.172.517.

[BS20] Francesco Bonsante and Andrea Seppi. *Anti-de Sitter geometry and Teichmüller theory*. 2020. arXiv: 2004.14414 [math.GT].

[CG93] Suhyoung Choi and William M. Goldman. "Convex real projective structures on closed surfaces are closed". In: *Proc. Amer. Math. Soc.* 118.2 (1993), pp. 657–661. ISSN: 0002-9939. DOI: 10.2307/2160352. URL: https://doi.org/10.2307/2160352.

[Che76] Shiu Yuen Cheng. "Eigenfunctions and nodal sets". In: *Comment. Math. Helv.* 51.1 (1976), pp. 43–55. ISSN: 0010-2571. DOI: 10.1007/BF02568142. URL: https://doi-org.clsproxy.library.caltech.edu/10.1007/BF02568142.

[Che80] Shiu Yuen Cheng. "Liouville theorem for harmonic maps". In: *Geometry of the Laplace operator (Proc. Sympos. Pure Math., Univ. Hawaii, Honolulu, Hawaii, 1979)*. Proc. Sympos. Pure Math., XXXVI. Amer. Math. Soc., Providence, R.I., 1980, pp. 147–151.

[Coo93] Michel Coornaert. "Mesures de Patterson-Sullivan sur le bord d'un espace hyperbolique au sens de Gromov". In: *Pacific J. Math.* 159.2

(1993), pp. 241–270. ISSN: 0030-8730. URL: http://projecteuclid.org/euclid.pjm/1102634263.

[Cor88] Kevin Corlette. “Flat G-bundles with canonical metrics”. In: *J. Differential Geom.* 28.3 (1988), pp. 361–382. ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1214442469.

[Cor92a] Kevin Corlette. “Archimedean superrigidity and hyperbolic geometry”. In: *Ann. of Math. (2)* 135.1 (1992), pp. 165–182. ISSN: 0003-486X. DOI: 10.2307/2946567. URL: https://doi.org/10.2307/2946567.

[Cor92b] Kevin Corlette. “Archimedean superrigidity and hyperbolic geometry”. In: *Ann. of Math. (2)* 135.1 (1992), pp. 165–182. ISSN: 0003-486X. DOI: 10.2307/2946567. URL: https://doi.org/10.2307/2946567.

[CTT19] Brian Collier, Nicolas Tholozan, and Jérémy Toulisse. “The geometry of maximal representations of surface groups into $\mathrm{SO}_0(2,n)$”. In: *Duke Math. J.* 168.15 (2019), pp. 2873–2949. ISSN: 0012-7094. DOI: 10.1215/00127094-2019-0052. URL: https://doi.org/10.1215/00127094-2019-0052.

[CY75] S. Y. Cheng and S. T. Yau. “Differential equations on Riemannian manifolds and their geometric applications”. In: *Comm. Pure Appl. Math.* 28.3 (1975), pp. 333–354. ISSN: 0010-3640. DOI: 10.1002/cpa.3160280303. URL: https://doi.org/10.1002/cpa.3160280303.

[DGK16a] Jeffrey Danciger, François Guéritaud, and Fanny Kassel. In: *Invent. Math.* 204.1 (2016), pp. 133–193. ISSN: 0020-9910. DOI: 10.1007/s00222-015-0610-z. URL: https://doi.org/10.1007/s00222-015-0610-z.

[DGK16b] Jeffrey Danciger, François Guéritaud, and Fanny Kassel. “Fundamental domains for free groups acting on anti–de Sitter 3-space”. In: *Math. Res. Lett.* 23.3 (2016), pp. 735–770. ISSN: 1073-2780. DOI: 10.4310/MRL.2016.v23.n3.a10. URL: https://doi.org/10.4310/MRL.2016.v23.n3.a10.

[DHS10] Ulrich Dierkes, Stefan Hildebrandt, and Friedrich Sauvigny. *Minimal surfaces*. second. Vol. 339. Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. With assistance and contributions by A. Küster and R. Jakob. Springer, Heidelberg, 2010, pp. xvi+688. ISBN: 978-3-642-11697-1. DOI: 10.1007/978-3-642-11698-8. URL: https://doi.org/10.1007/978-3-642-11698-8.

[DHT10] Ulrich Dierkes, Stefan Hildebrandt, and Anthony J. Tromba. *Global analysis of minimal surfaces*. second. Vol. 341. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer, Heidelberg, 2010, pp. xvi+537. ISBN: 978-3-642-11705-3. DOI: 10.1007/978-3-642-11706-0. URL: https://doi-org.clsproxy.library.caltech.edu/10.1007/978-3-642-11706-0.

[DL20] Song Dai and Qiongling Li. *Domination results in n-Fuchsian fibers in the moduli space of Higgs bundles*. 2020. arXiv: 2005.13960 [math.DG].

[Don87] S. K. Donaldson. "Twisted harmonic maps and the self-duality equations". In: *Proc. London Math. Soc. (3)* 55.1 (1987), pp. 127–131. ISSN: 0024-6115. DOI: 10.1112/plms/s3-55.1.127. URL: https://doi.org/10.1112/plms/s3-55.1.127.

[Dou39] Jesse Douglas. "Minimal surfaces of higher topological structure". In: *Ann. of Math. (2)* 40.1 (1939), pp. 205–298. ISSN: 0003-486X. DOI: 10.2307/1968552. URL: https://doi-org.clsproxy.library.caltech.edu/10.2307/1968552.

[DT16] Bertrand Deroin and Nicolas Tholozan. "Dominating surface group representations by Fuchsian ones". In: *Int. Math. Res. Not. IMRN* 13 (2016), pp. 4145–4166. ISSN: 1073-7928. DOI: 10.1093/imrn/rnv275. URL: https://doi.org/10.1093/imrn/rnv275.

[Dur04] Peter Duren. *Harmonic mappings in the plane*. Vol. 156. Cambridge Tracts in Mathematics. Cambridge University Press, Cambridge, 2004, pp. xii+212. ISBN: 0-521-64121-7. DOI: 10.1017/CBO9780511546600. URL: https://doi.org/10.1017/CBO9780511546600.

[DZ09] Sorin Dumitrescu and Abdelghani Zeghib. "Global rigidity of holomorphic Riemannian metrics on compact complex 3-manifolds". In: *Math. Ann.* 345.1 (2009), pp. 53–81. ISSN: 0025-5831. DOI: 10.1007/s00208-009-0342-8. URL: https://doi.org/10.1007/s00208-009-0342-8.

[Eji02] Norio Ejiri. "A differential-geometric Schottky problem, and minimal surfaces in tori". In: *Differential geometry and integrable systems (Tokyo, 2000)*. Vol. 308. Contemp. Math. Amer. Math. Soc., Providence, RI, 2002, pp. 101–144. DOI: 10.1090/conm/308/05314. URL: https://doi.org/10.1090/conm/308/05314.

[EL81] J. Eells and L. Lemaire. "Deformations of metrics and associated harmonic maps". In: *Proc. Indian Acad. Sci. Math. Sci.* 90.1 (1981), pp. 33–45. ISSN: 0253-4142. DOI: 10.1007/BF02867016.

[EL83] James Eells and Luc Lemaire. *Selected topics in harmonic maps*. Vol. 50. CBMS Regional Conference Series in Mathematics. Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 1983, pp. v+85. ISBN: 0-8218-0700-5. DOI: 10.1090/cbms/050. URL: https://doi.org/10.1090/cbms/050.

[ES64] James Eells Jr. and J. H. Sampson. "Harmonic mappings of Riemannian manifolds". In: *Amer. J. Math.* 86 (1964), pp. 109–160. ISSN: 0002-9327. DOI: 10.2307/2373037. URL: https://doi.org/10.2307/2373037.

[FG06] Vladimir Fock and Alexander Goncharov. "Moduli spaces of local systems and higher Teichmüller theory". In: *Publ. Math. Inst. Hautes Études Sci.* 103 (2006), pp. 1–211. ISSN: 0073-8301. DOI: 10.1007/s10240-006-0039-4. URL: https://doi.org/10.1007/s10240-006-0039-4.

[FLP12] Albert Fathi, François Laudenbach, and Valentin Poénaru. *Thurston's work on surfaces*. Vol. 48. Mathematical Notes. Translated from the 1979 French original by Djun M. Kim and Dan Margalit. Princeton University Press, Princeton, NJ, 2012, pp. xvi+254. ISBN: 978-0-691-14735-2.

[FM12] Benson Farb and Dan Margalit. *A primer on mapping class groups*. Vol. 49. Princeton Mathematical Series. Princeton University Press, Princeton, NJ, 2012, pp. xiv+472. ISBN: 978-0-691-14794-9.

[FR20] Gavril Farkas and Richárd Rimányi. "Quadric rank loci on moduli of curves and $K3$ surfaces". In: *Ann. Sci. Éc. Norm. Supér. (4)* 53.4 (2020), pp. 945–992. ISSN: 0012-9593. DOI: 10.24033/asens.2437. URL: https://doi.org/10.24033/asens.2437.

[FT21] Vladimir Fock and Alexander Thomas. "Higher complex structures". In: *Int. Math. Res. Not. IMRN* 20 (2021), pp. 15873–15893. ISSN: 1073-7928. DOI: 10.1093/imrn/rnz283. URL: https://doi.org/10.1093/imrn/rnz283.

[Gab85] David Gabai. "The simple loop conjecture". In: *J. Differential Geom.* 21.1 (1985), pp. 143–149. ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1214439470.

[Gar87] Frederick P. Gardiner. *Teichmüller theory and quadratic differentials*. Pure and Applied Mathematics (New York). A Wiley-Interscience Publication. John Wiley & Sons, Inc., New York, 1987, pp. xviii+236. ISBN: 0-471-84539-6.

[GH77] Samuel I. Goldberg and Zvi Har'El. "A general Schwarz lemma for Riemannian-manifolds". In: *Bull. Soc. Math. Grèce (N.S.)* 18.1 (1977), pp. 141–148. ISSN: 0373-1391.

[GK17] François Guéritaud and Fanny Kassel. "Maximally stretched laminations on geometrically finite hyperbolic manifolds". In: *Geom. Topol.* 21.2 (2017), pp. 693–840. ISSN: 1465-3060. DOI: 10.2140/gt.2017.21.693. URL: https://doi-org.clsproxy.library.caltech.edu/10.2140/gt.2017.21.693.

[GKW15] François Guéritaud, Fanny Kassel, and Maxime Wolff. "Compact anti–de Sitter 3-manifolds and folded hyperbolic structures on surfaces". In: *Pacific J. Math.* 275.2 (2015), pp. 325–359. ISSN: 0030-8730. DOI: 10.2140/pjm.2015.275.325. URL: https://doi.org/10.2140/pjm.2015.275.325.

[GLW21] Olivier Guichard, François Labourie, and Anna Wienhard. *Positivity and representations of surface groups*. 2021. DOI: 10.48550/ARXIV.2106.14584. URL: https://arxiv.org/abs/2106.14584.

[Gol85] William M. Goldman. "Nonstandard Lorentz space forms". In: *J. Differential Geom.* 21.2 (1985), pp. 301–308. ISSN: 0022-040X. URL: http://projecteuclid.org.clsproxy.library.caltech.edu/euclid.jdg/1214439567.

[Gol90] William M. Goldman. "Convex real projective structures on compact surfaces". In: *J. Differential Geom.* 31.3 (1990), pp. 791–845. ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1214444635.

[GOR73] R. D. Gulliver II, R. Osserman, and H. L. Royden. "A theory of branched immersions of surfaces". In: *Amer. J. Math.* 95 (1973), pp. 750–812. ISSN: 0002-9327. DOI: 10.2307/2373697. URL: https://doi-org.clsproxy.library.caltech.edu/10.2307/2373697.

[GS20] Subhojoy Gupta and Weixu Su. *Dominating surface-group representations into* $\mathrm{PSL}_2(\mathbb{C})$ *in the relative representation variety*. 2020. arXiv: 2003.13572 [math.GT].

[GS92] Mikhail Gromov and Richard Schoen. "Harmonic maps into singular spaces and p-adic superrigidity for lattices in groups of rank one". In: *Inst. Hautes Études Sci. Publ. Math.* 76 (1992), pp. 165–246. ISSN: 0073-8301. URL: http://www.numdam.org/item?id=PMIHES_1992__76__165_0.

[Gui18] Olivier Guichard. "An introduction to the differential geometry of flat bundles and of Higgs bundles". In: *The geometry, topology and physics of moduli spaces of Higgs bundles*. Vol. 36. Lect. Notes Ser. Inst. Math. Sci. Natl. Univ. Singap. World Sci. Publ., Hackensack, NJ, 2018, pp. 1–63.

[Gul73] Robert D. Gulliver II. "Regularity of minimizing surfaces of prescribed mean curvature". In: *Ann. of Math. (2)* 97 (1973), pp. 275–305. ISSN: 0003-486X. DOI: 10.2307/1970848. URL: https://doi-org.clsproxy.library.caltech.edu/10.2307/1970848.

[Gup17] Subhojoy Gupta. "Harmonic maps and wild Teichmüller spaces". In: *arXiv e-prints*, arXiv:1708.04780 (2017), arXiv:1708.04780. arXiv: 1708.04780 [math.DG].

[GW07] William M. Goldman and Richard A. Wentworth. "Energy of twisted harmonic maps of Riemann surfaces". In: *In the tradition of Ahlfors-Bers. IV*. Vol. 432. Contemp. Math. Amer. Math. Soc., Providence, RI, 2007, pp. 45–61. DOI: 10.1090/conm/432/08299. URL: https://doi-org.clsproxy.library.caltech.edu/10.1090/conm/432/08299.

[Ham03] Ursula Hamenstädt. "Length functions and parameterizations of Teichmüller space for surfaces with cusps". In: *Ann. Acad. Sci. Fenn. Math.* 28.1 (2003), pp. 75–88. ISSN: 1239-629X.

[Ham75] Richard S. Hamilton. *Harmonic maps of manifolds with boundary*. Lecture Notes in Mathematics, Vol. 471. Springer-Verlag, Berlin-New York, 1975, pp. i+168.

[Har67] Philip Hartman. "On homotopic harmonic maps". In: *Canadian J. Math.* 19 (1967), pp. 673–687. ISSN: 0008-414X. DOI: 10.4153/CJM-1967-062-6. URL: https://doi.org/10.4153/CJM-1967-062-6.

[HI77] Ernst Heintze and Hans-Christoph Im Hof. "Geometry of horospheres". In: *J. Differential Geom.* 12.4 (1977), 481–491 (1978). ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1214434219.

[Hit87] N. J. Hitchin. "The self-duality equations on a Riemann surface". In: *Proc. London Math. Soc. (3)* 55.1 (1987), pp. 59–126. ISSN: 0024-6115. DOI: 10.1112/plms/s3-55.1.59. URL: https://doi.org/10.1112/plms/s3-55.1.59.

[Hit92] N. J. Hitchin. "Lie groups and Teichmüller space". In: *Topology* 31.3 (1992), pp. 449–473. ISSN: 0040-9383. DOI: 10.1016/0040-9383(92)90044-I. URL: https://doi.org/10.1016/0040-9383(92)90044-I.

[HM79] John Hubbard and Howard Masur. "Quadratic differentials and foliations". In: *Acta Math.* 142.3-4 (1979), pp. 221–274. ISSN: 0001-5962. DOI: 10.1007/BF02395062. URL: https://doi.org/10.1007/BF02395062.

[HPR93] Joel Hass, Jon T. Pitts, and J. H. Rubinstein. “Existence of unstable minimal surfaces in manifolds with homology and applications to triply periodic minimal surfaces”. In: *Differential geometry: partial differential equations on manifolds (Los Angeles, CA, 1990)*. Vol. 54. Proc. Sympos. Pure Math. Amer. Math. Soc., Providence, RI, 1993, pp. 147–162.

[Hub06] John Hamal Hubbard. *Teichmüller theory and applications to geometry, topology, and dynamics. Vol. 1*. Teichmüller theory, With contributions by Adrien Douady, William Dunbar, Roland Roeder, Sylvain Bonnot, David Brown, Allen Hatcher, Chris Hruska and Sudeb Mitra, With forewords by William Thurston and Clifford Earle. Matrix Editions, Ithaca, NY, 2006, pp. xx+459.

[HW15] Zheng Huang and Biao Wang. “Counting minimal surfaces in quasi-Fuchsian three-manifolds”. In: *Trans. Amer. Math. Soc.* 367.9 (2015), pp. 6063–6083. ISSN: 0002-9947. DOI: 10.1090/tran/6172. URL: https://doi.org/10.1090/tran/6172.

[HW53] Philip Hartman and Aurel Wintner. “On the local behavior of solutions of non-parabolic partial differential equations”. In: *Amer. J. Math.* 75 (1953), pp. 449–476. ISSN: 0002-9327. DOI: 10.2307/2372496. URL: https://doi-org.clsproxy.library.caltech.edu/10.2307/2372496.

[IT92] Y. Imayoshi and M. Taniguchi. *An introduction to Teichmüller spaces*. Translated and revised from the Japanese by the authors. Springer-Verlag, Tokyo, 1992, pp. xiv+279. ISBN: 4-431-70088-9. DOI: 10.1007/978-4-431-68174-8. URL: https://doi.org/10.1007/978-4-431-68174-8.

[JM18] István Juhász and Jan van Mill. “On maps preserving connectedness and/or compactness”. In: *Comment. Math. Univ. Carolin.* 59.4 (2018), pp. 513–521. ISSN: 0010-2628. DOI: 10.14712/1213-7243.2015.263. URL: https://doi-org.clsproxy.library.caltech.edu/10.14712/1213-7243.2015.263.

[Jos08] Jürgen Jost. “Harmonic mappings”. In: *Handbook of geometric analysis. Vol. 1*. Ed. by Lizhen Ji et al. Vol. 7. Advanced lectures in mathematics (International Press). Boston: International Press of Boston, 2008, pp. 147–194. ISBN: 978-1-57146-130-8.

[Jos84] Jürgen Jost. *Harmonic maps between surfaces*. Vol. 1062. Lecture Notes in Mathematics. Springer-Verlag, Berlin, 1984, pp. x+133. ISBN: 3-540-13339-9. DOI: 10.1007/BFb0100160. URL: https://doi.org/10.1007/BFb0100160.

[JY83] Jürgen Jost and Shing Tung Yau. "Harmonic mappings and Kähler manifolds". In: *Math. Ann.* 262.2 (1983), pp. 145–166. ISSN: 0025-5831. DOI: 10.1007/BF01455308. URL: https://doi-org.clsproxy.library.caltech.edu/10.1007/BF01455308.

[JY91] Jürgen Jost and Shing-Tung Yau. "Harmonic maps and group representations". In: *Differential geometry*. Vol. 52. Pitman Monogr. Surveys Pure Appl. Math. Longman Sci. Tech., Harlow, 1991, pp. 241–259.

[JZ97] Jürgen Jost and Kang Zuo. "Harmonic maps of infinite energy and rigidity results for representations of fundamental groups of quasiprojective varieties". In: *J. Differential Geom.* 47.3 (1997), pp. 469–503. ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1214460547.

[Kas08] Fanny Kassel. "Proper actions on corank-one reductive homogeneous spaces". In: *J. Lie Theory* 18.4 (2008), pp. 961–978. ISSN: 0949-5932.

[Kas10] Fanny Kassel. "Quotients compacts des groupes ultramétriques de rang un". In: *Ann. Inst. Fourier (Grenoble)* 60.5 (2010), pp. 1741–1786. ISSN: 0373-0956. URL: http://aif.cedram.org/item?id=AIF_2010__60_5_1741_0.

[Kat+15] Ludmil Katzarkov et al. "Harmonic maps to buildings and singular perturbation theory". In: *Comm. Math. Phys.* 336.2 (2015), pp. 853–903. ISSN: 0010-3616. DOI: 10.1007/s00220-014-2276-6. URL: https://doi.org/10.1007/s00220-014-2276-6.

[Kli96] Bruno Klingler. "Complétude des variétés lorentziennes à courbure constante". In: *Math. Ann.* 306.2 (1996), pp. 353–370. ISSN: 0025-5831. DOI: 10.1007/BF01445255. URL: https://doi.org/10.1007/BF01445255.

[KM08] Vincent Koziarz and Julien Maubon. "Harmonic maps and representations of non-uniform lattices of $PU(m,1)$". In: *Ann. Inst. Fourier (Grenoble)* 58.2 (2008), pp. 507–558. ISSN: 0373-0956. URL: http://aif.cedram.org/item?id=AIF_2008__58_2_507_0.

[KR85] Ravi S. Kulkarni and Frank Raymond. "3-dimensional Lorentz space-forms and Seifert fiber spaces". In: *J. Differential Geom.* 21.2 (1985), pp. 231–268. ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1214439564.

[KS07] Kirill Krasnov and Jean-Marc Schlenker. "Minimal surfaces and particles in 3-manifolds". In: *Geom. Dedicata* 126 (2007), pp. 187–254. ISSN: 0046-5755. DOI: 10.1007/s10711-007-9132-1. URL: https://doi.org/10.1007/s10711-007-9132-1.

[Lab06] François Labourie. "Anosov flows, surface groups and curves in projective space". In: *Invent. Math.* 165.1 (2006), pp. 51–114. ISSN: 0020-9910. DOI: 10.1007/s00222-005-0487-3. URL: https://doi.org/10.1007/s00222-005-0487-3.

[Lab08] François Labourie. "Cross ratios, Anosov representations and the energy functional on Teichmüller space". In: *Ann. Sci. Éc. Norm. Supér. (4)* 41.3 (2008), pp. 437–469. ISSN: 0012-9593. DOI: 10.24033/asens.2072. URL: https://doi.org/10.24033/asens.2072.

[Lab13] François Labourie. *Lectures on representations of surface groups*. Zurich Lectures in Advanced Mathematics. European Mathematical Society (EMS), Zürich, 2013, pp. viii+138. ISBN: 978-3-03719-127-9. DOI: 10.4171/127. URL: https://doi.org/10.4171/127.

[Lab17] François Labourie. "Cyclic surfaces and Hitchin components in rank 2". In: *Ann. of Math. (2)* 185.1 (2017), pp. 1–58. ISSN: 0003-486X. DOI: 10.4007/annals.2017.185.1.1. URL: https://doi.org/10.4007/annals.2017.185.1.1.

[Lab91] François Labourie. "Existence d'applications harmoniques tordues à valeurs dans les variétés à courbure négative". In: *Proc. Amer. Math. Soc.* 111.3 (1991), pp. 877–882. ISSN: 0002-9939. DOI: 10.2307/2048427. URL: https://doi.org/10.2307/2048427.

[Law80] H. Blaine Lawson Jr. *Lectures on minimal submanifolds. Vol. I*. Second. Vol. 9. Mathematics Lecture Series. Publish or Perish, Inc., Wilmington, Del., 1980, pp. iv+178. ISBN: 0-914098-18-7.

[Li19] Qiongling Li. "An introduction to Higgs bundles via harmonic maps". In: *SIGMA Symmetry Integrability Geom. Methods Appl.* 15 (2019), Paper No. 035, 30. DOI: 10.3842/SIGMA.2019.035. URL: https://doi.org/10.3842/SIGMA.2019.035.

[Lof01] John C. Loftin. "Affine spheres and convex $\mathbb{RP}^n$-manifolds". In: *Amer. J. Math.* 123.2 (2001), pp. 255–274. ISSN: 0002-9327. URL: http://muse.jhu.edu/journals/american_journal_of_mathematics/v123/123.2loftin.pdf.

[Loh91] Jochen Lohkamp. "Harmonic diffeomorphisms and Teichmüller theory". In: *Manuscripta Math.* 71.4 (1991), pp. 339–360. ISSN: 0025-2611. DOI: 10.1007/BF02568411. URL: https://doi.org/10.1007/BF02568411.

[LW08] Fanghua Lin and Changyou Wang. *The analysis of harmonic maps and their heat flows*. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2008, pp. xii+267. DOI: 10.1142/9789812779533. URL: https://doi.org/10.1142/9789812779533.

[Mar17] Vladimir Marković. "Harmonic maps and the Schoen conjecture". In: *J. Amer. Math. Soc.* 30.3 (2017), pp. 799–817. ISSN: 0894-0347. DOI: 10.1090/jams/881. URL: https://doi.org/10.1090/jams/881.

[Mar18] Vladimir Marković. *Harmonic surfaces in 3-manifolds and the simple loop theorem*. 2018. URL: https://doi-org.clsproxy.library.caltech.edu/10.1007/BF02867016.

[Mar21] Vladimir Marković. "Uniqueness of minimal diffeomorphisms between surfaces". In: *Bull. Lond. Math. Soc.* 53.4 (2021), pp. 1196–1204. ISSN: 0024-6093. DOI: 10.1112/blms.12493. URL: https://doi.org/10.1112/blms.12493.

[Mar22] Vladimir Marković. "Non-uniqueness of minimal surfaces in a product of closed Riemann surfaces". In: *Geom. Funct. Anal.* 32.1 (2022), pp. 31–52. ISSN: 1016-443X. DOI: 10.1007/s00039-021-00590-4. URL: https://doi.org/10.1007/s00039-021-00590-4.

[Mas95] Howard Masur. "The Teichmüller flow is Hamiltonian". In: *Proc. Amer. Math. Soc.* 123.12 (1995), pp. 3739–3747. ISSN: 0002-9939. DOI: 10.2307/2161902. URL: https://doi.org/10.2307/2161902.

[Mee90] William H. Meeks III. "The theory of triply periodic minimal surfaces". In: *Indiana Univ. Math. J.* 39.3 (1990), pp. 877–936. ISSN: 0022-2518. DOI: 10.1512/iumj.1990.39.39043. URL: https://doi.org/10.1512/iumj.1990.39.39043.

[Mes07] Geoffrey Mess. "Lorentz spacetimes of constant curvature". In: *Geom. Dedicata* 126 (2007), pp. 3–45. ISSN: 0046-5755. DOI: 10.1007/s10711-007-9155-7. URL: https://doi.org/10.1007/s10711-007-9155-7.

[Min87] David Minda. "The strong form of Ahlfors' lemma". In: *Rocky Mountain J. Math.* 17.3 (1987), pp. 457–461. ISSN: 0035-7596. DOI: 10.1216/RMJ-1987-17-3-457. URL: https://doi.org/10.1216/RMJ-1987-17-3-457.

[Min92] Yair N. Minsky. "Harmonic maps, length, and energy in Teichmüller space". In: *J. Differential Geom.* 35.1 (1992), pp. 151–217. ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1214447809.

[Mon16] Gabriele Mondello. "Topology of representation spaces of surface groups in $PSL_2(\mathbb{R})$ with assigned boundary monodromy and nonzero Euler number". In: *Pure Appl. Math. Q.* 12.3 (2016), pp. 399–462. ISSN: 1558-8599. DOI: 10.4310/PAMQ.2016.v12.n3.a3. URL: https://doi-org.clsproxy.library.caltech.edu/10.4310/PAMQ.2016.v12.n3.a3.

[Moo06] John Douglas Moore. "Bumpy metrics and closed parametrized minimal surfaces in Riemannian manifolds". In: *Trans. Amer. Math. Soc.* 358.12 (2006), pp. 5193–5256. ISSN: 0002-9947. DOI: 10.1090/S0002-9947-06-04317-0. URL: https://doi-org.clsproxy.library.caltech.edu/10.1090/S0002-9947-06-04317-0.

[Moo17] John Douglas Moore. *Introduction to global analysis*. Vol. 187. Graduate Studies in Mathematics. Minimal surfaces in Riemannian manifolds. American Mathematical Society, Providence, RI, 2017, pp. xiv+368. ISBN: 978-1-4704-2950-8.

[Mos64] Jürgen Moser. "A Harnack inequality for parabolic differential equations". In: *Comm. Pure Appl. Math.* 17 (1964), pp. 101–134. ISSN: 0010-3640. DOI: 10.1002/cpa.3160170106. URL: https://doi.org/10.1002/cpa.3160170106.

[MS12] Dusa McDuff and Dietmar Salamon. *J-holomorphic curves and symplectic topology*. Second. Vol. 52. American Mathematical Society Colloquium Publications. American Mathematical Society, Providence, RI, 2012, pp. xiv+726. ISBN: 978-0-8218-8746-2.

[MSS22] Vladimir Marković, Nathaniel Sagman, and Peter Smillie. *Unstable minimal surfaces in $\mathbb{R}^n$ and in products of hyperbolic surfaces*. 2022. arXiv: 2206.02938 [math.DG].

[Mum71] David Mumford. "A remark on Mahler's compactness theorem". In: *Proc. Amer. Math. Soc.* 28 (1971), pp. 289–294. ISSN: 0002-9939. DOI: 10.2307/2037802. URL: https://doi-org.clsproxy.library.caltech.edu/10.2307/2037802.

[MW16] Julien Marché and Maxime Wolff. "The modular action on $\mathrm{PSL}_2(\mathbb{R})$-characters in genus 2". In: *Duke Math. J.* 165.2 (2016), pp. 371–412. ISSN: 0012-7094. DOI: 10.1215/00127094-3166522. URL: https://doi-org.clsproxy.library.caltech.edu/10.1215/00127094-3166522.

[MW95] Mario J. Micallef and Brian White. "The structure of branch points in minimal surfaces and in pseudoholomorphic curves". In: *Ann. of Math. (2)* 141.1 (1995), pp. 35–85. ISSN: 0003-486X. DOI: 10.2307/2118627. URL: https://doi-org.clsproxy.library.caltech.edu/10.2307/2118627.

[Nic21] Liviu I. Nicolaescu. *Lectures on the geometry of manifolds*. Third. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2021, p. 682.

[Nit74] Johannes C. C. Nitsche. "Plateau's problems and their modern ramifications". In: *Amer. Math. Monthly* 81 (1974), pp. 945–968. ISSN: 0002-9890. DOI: 10.2307/2319297. URL: https://doi-org.clsproxy.library.caltech.edu/10.2307/2319297.

[Omo67] Hideki Omori. "Isometric immersions of Riemannian manifolds". In: *J. Math. Soc. Japan* 19 (1967), pp. 205–214. ISSN: 0025-5645. DOI: 10.2969/jmsj/01920205. URL: https://doi.org/10.2969/jmsj/01920205.

[Oss70] Robert Osserman. "A proof of the regularity everywhere of the classical solution to Plateau's problem". In: *Ann. of Math. (2)* 91 (1970), pp. 550–569. ISSN: 0003-486X. DOI: 10.2307/1970637. URL: https://doi-org.clsproxy.library.caltech.edu/10.2307/1970637.

[Par12] Anne Parreau. "Compactification d'espaces de représentations de groupes de type fini". In: *Math. Z.* 272.1-2 (2012), pp. 51–86. ISSN: 0025-5874. DOI: 10.1007/s00209-011-0921-8. URL: https://doi.org/10.1007/s00209-011-0921-8.

[PT10] Athanase Papadopoulos and Guillaume Théret. "Shortening all the simple closed geodesics on surfaces with boundary". In: *Proc. Amer. Math. Soc.* 138.5 (2010), pp. 1775–1784. ISSN: 0002-9939. DOI: 10.1090/S0002-9939-09-10195-8. URL: https://doi.org/10.1090/S0002-9939-09-10195-8.

[PV17] Mark Pollicott and Polina Vytnova. "Critical points for the Hausdorff dimension of pairs of pants". In: *Groups Geom. Dyn.* 11.4 (2017), pp. 1497–1519. ISSN: 1661-7207. DOI: 10.4171/GGD/436. URL: https://doi.org/10.4171/GGD/436.

[Ros06] Antonio Ros. "One-sided complete stable minimal surfaces". In: *J. Differential Geom.* 74.1 (2006), pp. 69–92. ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1175266182.

[RS87] Edgar Reich and Kurt Strebel. "On the Gerstenhaber-Rauch principle". In: *Israel J. Math.* 57.1 (1987), pp. 89–100. ISSN: 0021-2172. DOI: 10.1007/BF02769462. URL: https://doi.org/10.1007/BF02769462.

[Sag19] Nathaniel Sagman. "Infinite energy equivariant harmonic maps, domination, and anti-de Sitter 3-manifolds". In: *To appear in J. Differential. Geom.* (2019). DOI: 10.48550/ARXIV.1911.06937. URL: https://arxiv.org/abs/1911.06937.

[Sag21a] Nathaniel Sagman. "A factorization theorem for harmonic maps". In: *J. Geom. Anal.* 31.12 (2021), pp. 11714–11740. ISSN: 1050-6926. DOI: 10.1007/s12220-021-00699-w. URL: https://link.springer.com/article/10.1007/s12220-021-00699-w.

[Sag21b] Nathaniel Sagman. *Almost strict domination and anti-de Sitter 3-manifolds*. 2021. DOI: 10.48550/ARXIV.2105.12886. URL: https://arxiv.org/abs/2105.12886.

[Sag21c] Nathaniel Sagman. *Moduli spaces of harmonic surfaces in non-positive curvature*. 2021. DOI: 10.48550/ARXIV.2111.04142. URL: https://arxiv.org/abs/2111.04142.

[Sal00] François Salein. "Variétés anti-de Sitter de dimension 3 exotiques". In: *Ann. Inst. Fourier (Grenoble)* 50.1 (2000), pp. 257–284. ISSN: 0373-0956. URL: http://www.numdam.org/item?id=AIF_2000__50_1_257_0.

[Sam78] J. H. Sampson. "Some properties and applications of harmonic mappings". In: *Ann. Sci. École Norm. Sup. (4)* 11.2 (1978), pp. 211–228. ISSN: 0012-9593. URL: http://www.numdam.org.clsproxy.library.caltech.edu/item?id=ASENS_1978_4_11_2_211_0.

[Sch93] Richard M. Schoen. "The role of harmonic mappings in rigidity and deformation problems". In: *Complex geometry (Osaka, 1990)*. Vol. 143. Lecture Notes in Pure and Appl. Math. Dekker, New York, 1993, pp. 179–200.

[Sco83] Peter Scott. "The geometries of 3-manifolds". In: *Bull. London Math. Soc.* 15.5 (1983), pp. 401–487. ISSN: 0024-6093. DOI: 10.1112/blms/15.5.401. URL: https://doi.org/10.1112/blms/15.5.401.

[Sim88] Carlos T. Simpson. "Constructing variations of Hodge structure using Yang-Mills theory and applications to uniformization". In: *J. Amer. Math. Soc.* 1.4 (1988), pp. 867–918. ISSN: 0894-0347. DOI: 10.2307/1990994. URL: https://doi.org/10.2307/1990994.

[Sim90] Carlos T. Simpson. "Harmonic bundles on noncompact curves". In: *J. Amer. Math. Soc.* 3.3 (1990), pp. 713–770. ISSN: 0894-0347. DOI: 10.2307/1990935. URL: https://doi-org.clsproxy.library.caltech.edu/10.2307/1990935.

[Siu80] Yum Tong Siu. "The complex-analyticity of harmonic maps and the strong rigidity of compact Kähler manifolds". In: *Ann. of Math. (2)* 112.1 (1980), pp. 73–111. ISSN: 0003-486X. DOI: 10.2307/1971321. URL: https://doi.org/10.2307/1971321.

[Siu82] Yum Tong Siu. "Complex-analyticity of harmonic maps, vanishing and Lefschetz theorems". In: *J. Differential Geometry* 17.1 (1982), pp. 55–138. ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1214436700.

[Sle20] Ivo Slegers. "Equivariant harmonic maps depend real analytically on the representation". In: *arXiv preprint* (2020).

[Str84] Kurt Strebel. *Quadratic differentials*. Vol. 5. Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]. Springer-Verlag, Berlin, 1984, pp. xii+184. ISBN: 3-540-13035-7. DOI: 10.1007/978-3-662-02414-0. URL: https://doi-org.clsproxy.library.caltech.edu/10.1007/978-3-662-02414-0.

[SU81] J. Sacks and K. Uhlenbeck. “The existence of minimal immersions of 2-spheres”. In: *Ann. of Math. (2)* 113.1 (1981), pp. 1–24. ISSN: 0003-486X. DOI: 10.2307/1971131. URL: https://doi.org/10.2307/1971131.

[SY78] Richard Schoen and Shing Tung Yau. “On univalent harmonic maps between surfaces”. In: *Invent. Math.* 44.3 (1978), pp. 265–278. ISSN: 0020-9910. DOI: 10.1007/BF01403164. URL: https://doi.org/10.1007/BF01403164.

[SY79] R. Schoen and Shing Tung Yau. “Existence of incompressible minimal surfaces and the topology of three-dimensional manifolds with non-negative scalar curvature”. In: *Ann. of Math. (2)* 110.1 (1979), pp. 127–142. ISSN: 0003-486X. DOI: 10.2307/1971247. URL: https://doi-org.clsproxy.library.caltech.edu/10.2307/1971247.

[SY97] R. Schoen and S. T. Yau. *Lectures on harmonic maps*. Conference Proceedings and Lecture Notes in Geometry and Topology, II. International Press, Cambridge, MA, 1997, pp. vi+394. ISBN: 1-57146-002-0.

[Szé14] Gábor Székelyhidi. *An introduction to extremal Kähler metrics*. Vol. 152. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2014, pp. xvi+192. ISBN: 978-1-4704-1047-6. DOI: 10.1090/gsm/152. URL: https://doi.org/10.1090/gsm/152.

[Tho17] Nicolas Tholozan. “Dominating surface group representations and deforming closed anti-de Sitter 3-manifolds”. In: *Geom. Topol.* 21.1 (2017), pp. 193–214. ISSN: 1465-3060. DOI: 10.2140/gt.2017.21.193. URL: https://doi.org/10.2140/gt.2017.21.193.

[Tho18] Nicolas Tholozan. “The volume of complete anti–de Sitter 3-manifolds”. In: *J. Lie Theory* 28.3 (2018), pp. 619–642. ISSN: 0949-5932.

[Thu98] William P. Thurston. “Minimal stretch maps between hyperbolic surfaces”. In: *arXiv Mathematics e-prints*, math/9801039 (1998), math/9801039. arXiv: math/9801039 [math.GT].

[Tor07] Francisco Torralbo. “Minimal Lagrangian immersions in $\mathbb{RH}^2 \times \mathbb{RH}^2$”. In: *Symposium on the Differential Geometry of Submanifolds*. [s.n.], [s.l.], 2007, pp. 217–219.

[Tra08] Martin Traizet. “On the genus of triply periodic minimal surfaces”. In: *J. Differential Geom.* 79.2 (2008), pp. 243–275. ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1211512641.

[Wan92] Tom Yau-Heng Wan. "Constant mean curvature surface, harmonic maps, and universal Teichmüller space". In: *J. Differential Geom.* 35.3 (1992), pp. 643–657. ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1214448260.

[Wen07] Richard A. Wentworth. "Energy of harmonic maps and Gardiner's formula". In: *In the tradition of Ahlfors-Bers. IV*. Vol. 432. Contemp. Math. Amer. Math. Soc., Providence, RI, 2007, pp. 221–229. DOI: 10.1090/conm/432/08311. URL: https://doi-org.clsproxy.library.caltech.edu/10.1090/conm/432/08311.

[Wen16] Richard A. Wentworth. "Higgs bundles and local systems on Riemann surfaces". In: *Geometry and quantization of moduli spaces*. Adv. Courses Math. CRM Barcelona. Birkhäuser/Springer, Cham, 2016, pp. 165–219.

[Whi36] Hassler Whitney. "Differentiable manifolds". In: *Ann. of Math. (2)* 37.3 (1936), pp. 645–680. ISSN: 0003-486X. DOI: 10.2307/1968482. URL: https://doi-org.clsproxy.library.caltech.edu/10.2307/1968482.

[Wie18] Anna Wienhard. "An invitation to higher Teichmüller theory". In: *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. II. Invited lectures*. World Sci. Publ., Hackensack, NJ, 2018, pp. 1013–1039.

[Wit07] Edward Witten. *Three-Dimensional Gravity Revisited*. 2007. DOI: 10.48550/ARXIV.0706.3359. URL: https://arxiv.org/abs/0706.3359.

[Wit89] Edward Witten. "2 + 1-dimensional gravity as an exactly soluble system". In: *Nuclear Phys. B* 311.1 (1989), pp. 46–78. ISSN: 0550-3213. DOI: 10.1016/0550-3213(88)90143-5. URL: https://doi-org.clsproxy.library.caltech.edu/10.1016/0550-3213(88)90143-5.

[Wol89] Michael Wolf. "The Teichmüller theory of harmonic maps". In: *J. Differential Geom.* 29.2 (1989), pp. 449–479. ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1214442885.

[Wol91a] Michael Wolf. "High energy degeneration of harmonic maps between surfaces and rays in Teichmüller space". In: *Topology* 30.4 (1991), pp. 517–540. ISSN: 0040-9383. DOI: 10.1016/0040-9383(91)90037-5. URL: https://doi.org/10.1016/0040-9383(91)90037-5.

[Wol91b] Michael Wolf. "Infinite energy harmonic maps and degeneration of hyperbolic surfaces in moduli space". In: *J. Differential Geom.* 33.2 (1991), pp. 487–539. ISSN: 0022-040X. URL: http://projecteuclid.org/euclid.jdg/1214446328.

[Wol95] Michael Wolf. "Harmonic maps from surfaces to ℝ-trees". In: *Math. Z.* 218.4 (1995), pp. 577–593. ISSN: 0025-5874. DOI: 10.1007/BF02571924. URL: https://doi.org/10.1007/BF02571924.

[Wol96] Michael Wolf. "On realizing measured foliations via quadratic differentials of harmonic maps to ℝ-trees". In: *J. Anal. Math.* 68 (1996), pp. 107–120. ISSN: 0021-7670. DOI: 10.1007/BF02790206. URL: https://doi.org/10.1007/BF02790206.

[Woo74] John C. Wood. *Harmonic maps between surfaces*. Ph.D. Thesis. University of Warwick, 1974.

[Woo77] John C. Wood. "Singularities of harmonic maps and applications of the Gauss-Bonnet formula". In: *Amer. J. Math.* 99.6 (1977), pp. 1329–1344. ISSN: 0002-9327. DOI: 10.2307/2374027. URL: https://doi-org.clsproxy.library.caltech.edu/10.2307/2374027.