Geometry & Topology

Volume 28 (2024)

Higgs bundles, harmonic maps and pleated surfaces

Andreas Ott

Jan Swoboda

Richard Wentworth

Michael Wolf

msp

Geometry & Topology **28**:7 (2024) 3135–3220
DOI: 10.2140/gt.2024.28.3135
Published: 25 November 2024

Higgs bundles, harmonic maps and pleated surfaces

ANDREAS OTT
JAN SWOBODA
RICHARD WENTWORTH
MICHAEL WOLF

This paper unites the gauge-theoretic and hyperbolic-geometric perspectives on the asymptotic geometry of the character variety of $\mathrm{SL}(2,\mathbb{C})$ representations of a surface group. Specifically, we find an asymptotic correspondence between the analytically defined limiting configuration of a sequence of solutions to the $\mathrm{SU}(2)$ self-duality equations on a closed Riemann surface constructed by Mazzeo, Swoboda, Weiß and Witt, and the geometric topological shear-bend parameters of equivariant pleated surfaces in hyperbolic three-space due to Bonahon and Thurston. The geometric link comes from the nonabelian Hodge correspondence and a study of high-energy degenerations of harmonic maps. Our result has several applications. We prove: (1) the local invariance of the partial compactification of the moduli space of solutions to the self-duality equations by limiting configurations; (2) a refinement of the harmonic maps characterization of the Morgan–Shalen compactification of the character variety; and (3) a comparison between the family of complex projective structures defined by a quadratic differential and the realizations of the corresponding flat connections as Higgs bundles, as well as a determination of the asymptotic shear-bend cocycle of Thurston's pleated surface.

32G15, 53C07, 53C43; 57K32

1 Introduction

The purpose of this paper is to bring together two perspectives on the asymptotic structure of the $\mathrm{SL}(2,\mathbb{C})$ character variety of a surface group: the complex-analytic perspective from algebraic geometry and nonlinear analysis, and the synthetic perspective from hyperbolic geometry and low-dimensional topology. The former finds its incarnation in a gauge-theoretic partial compactification of the moduli space of rank two Higgs bundles on a closed Riemann surface, via decouplings of the self-duality equations. The latter is understood in terms of equivariant pleated surfaces in hyperbolic three-space.

As a means of studying the asymptotics of the character variety, the analytic and the synthetic perspectives each have advantages and disadvantages. The complex-analytic perspective presents the character variety as a fibration over a $(6g-6)$–dimensional vector space of holomorphic differentials, and so presents the compactification as a fibration over a real projective space of just one less dimension. Now, the emphasis here is holomorphic invariants, and so there is a built-in reference to a fixed Riemann surface. This dependence of the compactification on an arbitrarily chosen Riemann surface renders the compactification unnatural from the point of view of the mapping class group action on the character variety. On the other hand, the synthetic perspective relies on a choice of lamination and so avoids issues of naturality, but the most celebrated compactification from this setting, the Morgan–Shalen compactification, has the far larger codimension $6g-5$, so entails a substantial loss of information.

By uniting the perspectives, we provide a partial compactification that retains attractive features from both perspectives: it is topological, so that any dependence on an original choice of base surface has vanished by the frontier of the character variety, but it remains codimension-one and so captures some of the nuance of the fibration. Along the way, we explain the hyperbolic geometry of the gauge theory perspective, at least asymptotically. This relationship further allows us to locate the family of projective structures on a Riemann surface, again asymptotically, as a collection of nearly linear flows on the fibers at infinity.

The duality between the analytic and synthetic perspective on the $\mathrm{SL}(2,\mathbb{C})$ character variety has its roots in the identification of $\mathrm{SL}(2,\mathbb{C})$ — or rather its adjoint group $\mathrm{PSL}(2,\mathbb{C})=\mathrm{SL}(2,\mathbb{C})/\pm 1$ — as the oriented isometry group of hyperbolic three-space. The nonabelian Hodge correspondence gives a homeomorphism between the associated moduli spaces, and an important and long-standing direction of research is the study of how different geometric properties on both sides are related under this identification. In this sense, the present work describes the asymptotics of the nonabelian Hodge correspondence from a geometric point of view.

Let us now explain some of the structures of the moduli space in more detail. The first is that of an algebraically completely integrable system, where a Higgs bundle is determined by a point in the Prym variety of the spectral curve associated to a holomorphic quadratic differential on the surface. The second is in terms of solutions to the self-duality equations. Recent work proves asymptotic convergence, as the norm of the Higgs field diverges, to solutions of a decoupled equation called a limiting configuration. The data describing a limiting configuration is also a quadratic differential and a choice of Prym differential. The third aspect of the moduli space is the link between solutions to the self-duality equations and flat $\mathrm{SL}(2,\mathbb{C})$ connections that is obtained from the existence of equivariant harmonic maps to $\mathbb{H}^3$. Here, the quadratic differential appears as the Hopf differential of the map. Asymptotics of harmonic maps are well understood. While the relationship between spectral data and the Prym differential of the limiting configuration is transparent, it is perhaps not so clear how to recover this information from the asymptotic behavior of the harmonic map. This is where the fourth perspective intervenes; that of hyperbolic geometry and methods of Thurston. The new ingredient linking harmonic maps to limiting configurations is the notion of an (equivariant) pleated surface. Pleated surfaces, and their monodromy representations, can be parametrized by shear-bend cocycles with respect to a maximal geodesic lamination on the surface. This shear-bend dichotomy is parallel to the asymptotic decoupling of the self-duality equations which gives rise to limiting configurations, and a major goal of this paper is to make this similarity more concrete.

Roughly speaking, our main result proves an asymptotic correspondence between shear-bend coordinates and limiting configurations via periods of Prym differentials. More precisely, we prove the following. First, the energy of equivariant harmonic maps diverges for a sequence ρ_n of $\mathrm{PSL}(2,\mathbb{C})$–representations that leaves all compact sets in the character variety. We shall show that for sufficiently large n, there exist ρ_n–equivariant pleated surfaces in $\mathbb{H}^3$ where the bending lamination is asymptotically close to the lamination associated to the horizontal foliation of the Hopf differentials of the harmonic maps. The image of the harmonic map is itself close to this equivariant pleated surface, in an appropriate sense. Moreover, the shearing cocycle of the pleated surface is approximated projectively by the intersection number with the vertical foliation of the Hopf differential, and the asymptotic limit of the bending cocycle of the equivariant pleated surface is determined by the periods of the Prym differential associated to the limiting configuration, as described above. This analysis has several applications. First, we find that a small change in the base Riemann surface used to define the moduli space of Higgs bundles changes the data of the limiting configurations in the boundary associated to a sequence of representations by parallel transport via the Gaus–Manin connection. In other words, the analytically defined limiting configurations are topological. Next, we obtain a partial refinement of the Morgan–Shalen compactification of the character variety. Ideal points of this compactification are defined by the projective length functions of isometric actions of the surface group on $\mathbb{R}$–trees. The refinement decorates a tree in a portion of the compactification with a bending cocycle which provides more precise information on the relationship between the limiting length functions and dual trees of measured foliations. Finally, we determine the asymptotic shear-bend cocycles of the equivariant pleated surfaces that Thurston associates to complex projective structures (or *opers*). The result states that these cocycles are asymptotic to the ones defined by the "Seiberg–Witten differential" on the spectral curve, which itself is an important device that figures prominently in the WKB analysis of the differential equation defining the projective structure.

Before formulating the precise statements of these results in Section 1.2 below, we first provide some notation and important terminology.

1.1 Preliminaries

1.1.1 Notation Throughout this paper, Σ will be a fixed closed, oriented surface of genus $g \geqslant 2$ with fundamental group $\pi_1 = \pi_1(\Sigma, p_0)$, where $p_0 \in \Sigma$ is a fixed basepoint. We will typically denote a marked Riemann surface structure on Σ by X, and a marked hyperbolic structure on Σ by S. The almost complex structure on X appears as J. Universal covers are denoted by $\widetilde{\Sigma}$, etc, and $\tilde{p}_0$ will indicate a fixed lift of p_0 to $\widetilde{\Sigma}$. The two- and three-dimensional real hyperbolic spaces are written $\mathbb{H}^2$ and $\mathbb{H}^3$, respectively. Notice that $\mathbb{H}^2$ (resp. $\mathbb{H}^3$) carries a left $\mathrm{SL}(2,\mathbb{R})$ (resp. $\mathrm{SL}(2,\mathbb{C})$) action by isometries, which factors through $\mathrm{PSL}(2,\mathbb{R}) = \mathrm{SL}(2,\mathbb{R})/\{\pm 1\}$ (resp. $\mathrm{PSL}(2,\mathbb{C}) = \mathrm{SL}(2,\mathbb{C})/\{\pm 1\}$). The canonical bundle of X is indicated by K_X. We denote by $\mathcal{QD}(X)$ the space of holomorphic quadratic differentials on X, and by $\mathcal{QD}^*(X) \subset \mathcal{QD}(X)$ the cone of differentials with only simple zeroes. We let $\mathcal{SQD}(X) \subset \mathcal{QD}(X)$ denote the unit differentials with respect to some norm, and $\mathcal{SQD}^*(X) = \mathcal{SQD}(X) \cap \mathcal{QD}^*(X)$. The Teichmüller space $T(X)$ of X will sometimes be labeled by $T(\Sigma)$ when identifying it with the Fricke space of discrete $\mathrm{PSL}(2,\mathbb{R})$ representations.

1.1.2 Moduli spaces Define the *Betti moduli space*

$$M_{\mathsf{B}}(\Sigma) := \mathrm{Hom}(\pi_1, \mathrm{SL}(2,\mathbb{C})) /\!\!/ \mathrm{SL}(2,\mathbb{C}) \tag{1-1}$$

parametrizing conjugacy classes of semisimple representations of the fundamental group of Σ. The choice of $\mathrm{SL}(2,\mathbb{C})$ rather than $\mathrm{PSL}(2,\mathbb{C})$ is really a matter of convenience. For simplicity we work with $\mathrm{SL}(2,\mathbb{C})$, but we caution that it will be important at several points to track the distinction between the special and projective groups. Hence, let us also define the character variety

$$R(\Sigma) := \mathrm{Hom}(\pi_1, \mathrm{PSL}(2,\mathbb{C})) /\!\!/ \mathrm{PSL}(2,\mathbb{C}). \tag{1-2}$$

We shall be interested in the connected component of the trivial representation, $R^o(\Sigma)$, consisting of representations that lift to $\mathrm{SL}(2,\mathbb{C})$. This can be realized as a quotient $R^o(\Sigma) = M_{\mathsf{B}}(\Sigma)/J_2(\Sigma)$, where $J_2(\Sigma) := \mathrm{Hom}(\pi_1, \{\pm 1\})$. The results of this paper apply with little change to the other component of $R(\Sigma)$ consisting of representations that do not lift to $\mathrm{SL}(2,\mathbb{C})$.

Let $M_{\mathsf{DR}}(\Sigma)$ denote the *de Rham moduli space* of completely reducible flat $\mathrm{SL}(2,\mathbb{C})$ connections on Σ. Since $\dim_{\mathbb{C}} X = 1$, a holomorphic connection on a vector bundle on X is automatically flat, and so we have a canonical identification of $M_{\mathsf{DR}}(\Sigma)$ with $M_{\mathsf{DR}}(X)$, the moduli space of rank 2 holomorphic connections. We will sometimes confuse the two when the Riemann surface structure is understood. The Riemann–Hilbert correspondence gives a homeomorphism

$$\mathsf{RH}\colon M_{\mathsf{DR}}(\Sigma) \xrightarrow{\sim} M_{\mathsf{B}}(\Sigma), \tag{1-3}$$

obtained by associating to a flat connection ∇ its monodromy representation ρ. Gauge equivalent connections give conjugate representations.

Let $M_{\mathsf{H}}(X)$ denote the *Hitchin moduli space* of rank 2 Higgs bundles consisting of isomorphism classes of pairs $(\mathcal{E}, \Phi)$, where $\mathcal{E} \to X$ is a rank 2 holomorphic vector bundle on X with fixed trivial determinant, and Φ is a holomorphic Higgs field. The *nonabelian Hodge correspondence* gives a homeomorphism

$$\mathsf{NAH}\colon M_{\mathsf{DR}}(X) \xrightarrow{\sim} M_{\mathsf{H}}(X). \tag{1-4}$$

This map is described in more detail in Section 2.3.

There is a proper holomorphic map $\mathscr{H}\colon M_{\mathsf{H}}(X) \to \mathcal{QD}(X)$. The fiber $\mathscr{H}^{-1}(q)$, $q \in \mathcal{QD}^*(X)$ may be identified with the *Prym variety* $\operatorname{Prym}(\hat{X}_q, X)$, where $\hat{X}_q \to X$ is a double cover branched at the zeroes of q called the *spectral curve*. This realizes $M_{\mathsf{H}}(X)$ as a smooth torus fibration over this locus; see Section 2.2. A choice of spin structure $K_X^{1/2}$ gives a section of $\mathscr{H}$ called the *Hitchin section*, and its image is called a *Hitchin component*. This realizes $T(X) \subset M_{\mathsf{H}}(X)$. One direction in the identification (1-4) goes as follows: Given a flat connection with monodromy $[\rho] \in M_{\mathsf{B}}(\Sigma)$ there is a ρ–equivariant harmonic map $u\colon \tilde{X} \to \mathbb{H}^3$. The Hopf differential of u descends to X as a holomorphic quadratic differential. Restricted to a lift of the Fricke space to $\mathrm{SL}(2,\mathbb{R})$, this gives a diffeomorphism $\mathcal{QD}(X) \simeq T(\Sigma) \subset M_{\mathsf{B}}(\Sigma)$; see Hitchin [29] and Wolf [58]. Under this identification the harmonic maps parametrization of $T(\Sigma)$ and Hitchin section $T(X)$ agree.

1.1.3 Self-duality equations and limiting configurations The gauge-theoretical perspective on $M_{\mathsf{H}}(X)$ is in terms of solutions of Hitchin's *self-duality equations* for a pair (A,Ψ) consisting of an $\mathsf{SU}(2)$–connection A and a hermitian Higgs field Ψ. The proper setup for these will be introduced in detail in Section 2.1.1. The Kobayashi–Hitchin correspondence yields a bijection between $M_{\mathsf{H}}(X)$ and the space of unitary gauge equivalence classes of solutions of the self-duality equations. We will not distinguish the notation between $M_{\mathsf{H}}(X)$ and the latter space, and therefore we write $[(A,\Psi)] \in M_{\mathsf{H}}(X)$ if (A,Ψ) is an irreducible solution of equations (2-4). There is a partial compactification of $M_{\mathsf{H}}(X)$ in terms of *limiting configurations*, which we shall present in Section 2.1.3. Briefly, a limiting configuration $[(A_\infty,\Psi_\infty)]$ associated to a differential $q \in \mathcal{SQD}^*(X)$ is a solution of the decoupled self-duality equations (2-8) on the punctured surface $X^\times = X \setminus q^{-1}(0)$ that has a singularity of a specific type in each zero of q. Furthermore, $2q = \operatorname{tr}(\Psi_\infty \otimes \Psi_\infty)^{2,0}$. Such limiting configurations have a natural interpretation in terms of parabolic Higgs bundles. For our purposes it will be important that the set of unitary gauge equivalence classes of limiting configurations associated with any q as above is a real torus of dimension $6g-6$, and that the Hitchin map $\mathscr{H}$ extends continuously to a map $\mathscr{H}_\infty$ from the space of limiting configurations to $\mathcal{SQD}^*(X)$. The fiber $\mathscr{H}_\infty^{-1}$ may be identified with a torus of *Prym differentials* on $\hat{X}_q$; see Proposition 2.11. The Liouville form restricts to a natural Prym differential λ_{SW} called the *Seiberg–Witten differential*, and this will play an important role in the paper.

1.1.4 Pleated surfaces By a *pleated surface* we mean a 4–tuple $\mathsf{P} = (S, f, \Lambda, \rho)$, where S is a hyperbolic structure on Σ; Λ is a maximal geodesic lamination on S called the *pleating locus*; $f\colon \tilde{S} \to \mathbb{H}^3$ is a continuous map from the lift $\tilde{S}$ to $\mathbb{H}^3$ that is totally geodesic on the components of $\tilde{S} \setminus \tilde{\Lambda}$, maps leaves of $\tilde{\Lambda}$ to geodesics, and is equivariant with respect to a representation $\rho\colon \pi_1 \to \mathrm{PSL}(2,\mathbb{C})$. We sometimes abbreviate the 4–tuple P as $f\colon \tilde{S} \to \mathbb{H}^3$ when context provides the other data.

Remark What is defined here might be called an *equivariant* or *abstract* pleated surface to distinguish it from the more standard, nonequivariant situation; cf Canary, Epstein and Green [8]. Following Bonahon [5], we will simply use the term *pleated surface* in the equivariant case as well.

Let $\mathcal{H}(\Lambda,\mathbb{R})$ and $\mathcal{H}(\Lambda,S^1)$ denote the spaces of *shearing* and *bending* cocycles, respectively; see Section 4.1.1. We further set $\mathcal{H}^o(\Lambda,S^1)\subset\mathcal{H}(\Lambda,S^1)$ to be the connected component of the identity. Then $\mathcal{H}(\Lambda,\mathbb{R})$ (resp. $\mathcal{H}^o(\Lambda,S^1)$) is a $(6g-6)$–dimensional vector space (resp. torus). Bonahon proves that there is an injective map

$$B_\Lambda\colon \mathcal{C}(\Lambda)\times\sqrt{-1}\,\mathcal{H}^o(\Lambda,S^1)\to R^o(\Sigma) \tag{1-5}$$

that is a biholomorphism onto its image [5, Theorem D]. Here, $\mathcal{C}(\Lambda)\subset\mathcal{H}(\Lambda,\mathbb{R})$ is an open convex polyhedral cone that is naturally identified with $T(\Sigma)$. In fact, $[\rho]=B_\Lambda(\boldsymbol{\sigma},\boldsymbol{\beta})$ is constructed via a pleated surface $f\colon\widetilde{S}\to\mathbb{H}^3$ that is ρ–equivariant and has pleating locus Λ. The complex cocycle $\boldsymbol{\sigma}+i\boldsymbol{\beta}\in\mathcal{H}(\Lambda,\mathbb{C}/2\pi i\mathbb{Z})$ is called the *shear-bend* cocycle of the pleated surface.

In this paper we will be interested in the special case where Λ is determined by the geodesic lamination Λ_q^h associated to the horizontal measured foliation $\mathcal{F}_q^h$ of a quadratic differential q that is holomorphic for a marked Riemann surface structure X on Σ. We shall always demand that q have simple zeroes. If in addition q has no (horizontal) *saddle connections*, then $\Lambda=\Lambda_q^h$. When q does have saddle connections the situation is more complicated to formulate, but the fundamental picture described below is unchanged; see Section 2.4.3. In any case, there is a canonical transverse cocycle $\boldsymbol{\sigma}_q^{\mathrm{can}}\in\mathcal{H}(\Lambda,\mathbb{R})$ for Λ related to the transverse measure to the vertical foliation defined by q; see Example 4.10.

1.2 Results

1.2.1 Statement of the main theorem The result below gives an asymptotic comparison between Bonahon's parametrization of the character variety in terms of shear-bend cocycles (1-5), and the limiting configurations of solutions to the self-duality equations. Consider the following set-up.

Let $[\rho_n]$ be an unbounded sequence in $R^o(\Sigma)$, by which we mean it leaves every compact subset. Assume that the Hopf differentials of the ρ_n–equivariant harmonic maps $u_n\colon\widetilde{X}\to\mathbb{H}^3$ are of the form $4t_n^2q_n$, with $t_n\to+\infty$ and $q_n\to q\in\mathcal{SQD}^*(X)$. We will assume that for some fixed hyperbolic structure on Σ we have chosen maximal laminations Λ_n, Λ containing $\Lambda_{q_n}^h$ and Λ_q^h, respectively, and such that $\Lambda_n\to\Lambda$ in the Hausdorff sense. In this case, there is a notion of convergence of cocycles in $\mathcal{H}(\Lambda_n,\mathbb{R})$ and $\mathcal{H}(\Lambda_n,S^1)$; see Definition 4.1.

Lift $[\rho_n]\in R^o(\Sigma)$ to $[\widetilde{\rho}_n]\in M_{\mathsf{B}}(\Sigma)$, and let

$$[(A_n,\Psi_n)]=\mathsf{NAH}\circ\mathsf{RH}^{-1}([\widetilde{\rho}_n]) \tag{1-6}$$

be the associated solutions to the self-duality equations. Let $[(A_\infty,\Psi_\infty)]$ be any subsequential limiting configuration of the sequence $[(A_n,\Psi_n)]$.

To normalize bending cocycles, we adopt the convention that a pleated surface for a Fuchsian representation has a bending cocycle equal to zero; cf [5, Proposition 27]. This will allow us to compare the bending cocycles of pleated surfaces with the same underlying hyperbolic metric and pleating lamination. In Section 4.2 we shall describe an explicit realization of elements of $\mathcal{H}(\Lambda, \mathbb{R})$ and $\mathcal{H}^o(\Lambda, S^1)$ in terms of periods of Prym differentials. Combining this with the characterization of limiting configurations mentioned at the end of Section 1.1.3, we show that there is a 2^{2g}–sheeted covering homomorphism

$$\mathscr{H}_\infty^{-1}(q) \to \mathcal{H}^o(\Lambda, S^1). \tag{1-7}$$

Given the above we can now make the following statement.

Main Theorem *After passing to a subsequence, there is $N \geqslant 1$ such that the following hold:*

(i) *For all $n \geqslant N$, $[\rho_n] = B_{\Lambda_n}(\boldsymbol{\sigma}_n, \boldsymbol{\beta}_n)$ for some shearing and bending cocycles $\boldsymbol{\sigma}_n$ and $\boldsymbol{\beta}_n$.*

(ii) *The ρ_n–equivariant pleated surfaces $f_n \colon \widetilde{S}_n \to \mathbb{H}^3$ from* (i) *are asymptotic to the ρ_n–equivariant harmonic maps $u_n \colon \widetilde{X} \to \mathbb{H}^3$ in the sense of Definition 4.14.*

(iii) *As $n \to \infty$, the shearing cocycles satisfy $(2t_n)^{-1}\boldsymbol{\sigma}_n \to \boldsymbol{\sigma}_q^{\mathrm{can}}$.*

(iv) *As $n \to \infty$, the bending cocycles satisfy $\boldsymbol{\beta}_n \to \boldsymbol{\beta}$, where $\boldsymbol{\beta} \in \mathcal{H}^o(\Lambda, S^1)$ is the image of the limiting configuration (A_∞, Ψ_∞) under the map (1-7).*

There are a number of ways to parse this statement; we roughly describe a somewhat constructive perspective. First, fix a base Riemann surface X, a holomorphic quadratic differential $q \in \mathcal{SQD}^*(X)$ and thus a ray $\{tq\} \subset \mathcal{SQD}^*(X)$ (we restrict ourselves to rays to simplify this particular exposition). The quadratic differential has a measured horizontal foliation, and that measured foliation has an associated measured lamination; say λ. The main theorem says that a family of Higgs bundles lying over that ray has harmonic maps whose images are well-approximated by pleated surfaces bent along λ. Of course, those Higgs bundles also converge to a limiting configuration. The Prym differential associated to that limiting configuration then predicts the bending of that approximating pleated surface.

Thus, the analytic data of the X–dependent fibration of limiting configurations and Prym differentials over $\mathcal{SQD}^*(X)$ may be alternatively described in terms of the hyperbolic-geometric perspective of laminations and bending measures. The key conclusion then is that this latter description that has emerged no longer references the base Riemann surface X: we display the space of limiting configurations defined by X as actually independent of the choice of X.

1.2.2 Invariance with respect to the base Riemann surface We turn now to the motivation we framed at the outset of the paper. As we just noted, the space of limiting configurations depends for its definition on the choice of a base Riemann surface X. The space of pleated surfaces has no such dependence, so we expect some consequence of the asymptotic result in the Main Theorem that we may define a compactification of the moduli space $M_{\mathsf{B}}(\Sigma)$ or $M_{\mathsf{H}}(X)$ that is natural in the sense of having some invariance with respect to the choice of base Riemann surface X. We describe that, and the level of invariance that is presently evident, in this subsection.

Now, a key assumption above and in the work of Mazzeo, Swoboda, Weiß and Witt [42] is that quadratic differentials have simple zeroes. For this reason, limiting configurations give only a partial compactification of $M_{\mathsf{H}}(X)$, and we are unable to make a uniform statement about the topological invariance of these limit points. We therefore content ourselves here with proving the *local* invariance with respect to the Riemann surface structure X.

To make this precise, let $q_0 \in \mathcal{SQD}^*(X_0)$, and let $\mathcal{F}^v_{q_0}$ denote the associated vertical measured foliation. Let $\widetilde{U} \subset T(\Sigma)$ be the set of all equivalence classes of Riemann surfaces X such that the Hubbard–Masur differential q_X of the pair $(X, \mathcal{F}^v_{q_0})$ has simple zeroes. Then for a contractible open subset $U_0 \subset \widetilde{U}$ containing X_0, and $X \in U_0$, the Gaus–Manin connection gives an identification of the Prym varieties of X and X_0, and this in turn induces an identification of bending cocycles for the horizontal laminations associated to q_X and q_0 through (1-7). As mentioned above, a complication, described in more detail in Section 2.4.3, is that the laminations $\Lambda^h_{q_0}$ and $\Lambda^h_{q_X}$ may not be maximal.

Corollary 1.1 *The partial compactification by limiting configurations is locally independent of the base Riemann surface in the following sense. Suppose $[\rho_n]$ is a divergent sequence in $R^o(\Sigma)$, and lift $[\rho_n]$ to $[\widetilde{\rho}_n] \in M_{\mathsf{B}}(\Sigma)$. For $X \in U_0$, define $[(A_n, \Psi_n)_X]$ and $[(A_n, \Psi_n)_{X_0}]$ as in (1-6). We assume $[(A_n, \Psi_n)_{X_0}]$ has a well-defined limiting configuration, which we suppose lies in the fiber $(\mathscr{H}^{X_0}_\infty)^{-1}(q_0)$. Let $[\widehat{\eta}_{X_0}]$ be the associated Prym differential* (*as mentioned above*; *see Proposition 2.11 for the precise statement*), *and $q_X \in \mathcal{SQD}^*(X)$ chosen as above to share the projective class of vertical measured foliations with q_{X_0}.*

Then $[(A_n, \Psi_n)_X]$ has a well-defined limiting configuration in the fiber $(\mathscr{H}^X_\infty)^{-1}(q_X)$. Moreover, if $[\widehat{\eta}_X]$ is the associated Prym differential for the bending cocycle of this limiting configuration, then $[\widehat{\eta}_X]$ and $[\widehat{\eta}_{X_0}]$ are identified by parallel translation by the Gaus–Manin connection.

Here is an interpretation of this result. For $q \in \mathcal{SQD}^*(X)$ there is a natural identification of the fibers $\mathscr{H}^{-1}(tq)$ for all $t > 0$. This gives a partial compactification of $M_H(X)$, and hence via NAH and RH, also of $M_{\mathsf{B}}(\Sigma)$. A priori, this depends on the choice of base Riemann surface structure X. Corollary 1.1 states that (locally) this partial compactification is independent of X.

1.2.3 Relation to the Morgan–Shalen compactification There is a mapping class group invariant compactification of $R(\Sigma)$ due to Morgan and Shalen [49]. The ideal points of this compactification are generalized length functions on π_1, which turn out to be the translation length functions for an isometric action of π_1 on an $\mathbb{R}$–tree; see Section 2.6.2. A harmonic maps description of this compactification was partially described in Daskalopoulos, Dostoglou and Wentworth [12], which was an attempt to mirror the result of Wolf [58] for the Thurston compactification of $T(\Sigma)$. A consequence of Wolf [60] is that the $\mathbb{R}$–trees appearing in the limit of a sequence of Fuchsian representations are obtained as the leaf space of the vertical foliation $\mathcal{F}^v_q$ of the rescaled Hopf differential q on $\widetilde{X}$. This is called the dual tree T_q to q. In the case of $R(\Sigma)$, sequences of representations that are not discrete embeddings may give rise to trees that are foldings of T_q. The harmonic maps point of view gives some information about this: A folding cannot occur if q has simple zeroes and $\mathcal{F}_q^v$ has no saddle connections. It has been an open question how to describe this process completely in terms of harmonic maps. Using the Main Theorem, we can obtain a criterion ruling out folding in the case of simple zeroes, as well as a partial refinement of the Morgan–Shalen compactification by classes of limiting configurations.

Corollary 1.2 *Let $[\rho_n] \in R^o(\Sigma)$ be as in the Main Theorem. Suppose that the periods of the Prym differential associated to the limiting configuration (A_∞, Ψ_∞) are bounded away from π on every cycle defined by a saddle connection of the vertical foliation of q. Then the $\mathbb{R}$–tree defined by the Morgan–Shalen limit of any subsequence of $[\rho_n]$ is π_1–equivariantly isometric to the dual tree T_q.*

1.2.4 Limits of complex projective structures The subset of $R^o(\Sigma)$ consisting of monodromies of complex projective structures on Σ with underlying Riemann surface X is naturally an affine space modeled on $\mathcal{QD}(X)$. The corresponding local systems are called $\mathrm{SL}(2,\mathbb{C})$-*opers*. A basepoint is given by the Fuchsian projective structure Q_X. More precisely, uniformization gives rise to an isomorphism $u\colon \widetilde{X} \to \mathbb{H}^2$, equivariant with respect to π_1 and a Fuchsian representation of $\pi_1 \to \mathrm{Iso}^+(\mathbb{H}^2)$, and the Schwarzian derivative of u gives a projective connection Q_X on X. Any other projective connection is of the form $Q(q) = Q_X - 2q$, for $q \in \mathcal{QD}(X)$, and we obtain an embedding $\mathscr{P}\colon \mathcal{QD}(X) \hookrightarrow R^o(\Sigma)$ from the monodromy of the oper defined by the following differential equation on X:

$$y'' + \tfrac{1}{2}Q(q)y = 0, \tag{1-8}$$

where y is a local section of $K_X^{-1/2}$. Thurston associates to every projective structure a pleated surface $f\colon \widetilde{S}(q) \to \mathbb{H}^3$ that is equivariant with respect to $\mathscr{P}(q)$ and has pleating locus along some measured lamination $\Lambda(q)$; see [36].[1] Choosing a lift of the Fuchsian representation to $M_{\mathsf{B}}(\Sigma)$ gives a lift $\widetilde{\mathscr{P}}(q) \in M_{\mathsf{B}}(\Sigma)$. Let

$$\mathrm{Op}(q) = \mathsf{NAH} \circ \mathsf{RH}^{-1}(\widetilde{\mathscr{P}}(q)) \in M_{\mathsf{H}}(X). \tag{1-9}$$

The next application compares the limiting behavior of $\mathrm{Op}(q)$ for q large and the geometry of Thurston's pleated surface. By work of Dumas [16], the measured laminations $\Lambda(q)$ converge projectively to Λ_q^h. This allows us to compare bending cocycles. Combined with the Main Theorem, we prove the following.

Corollary 1.3 *Let $q \in \mathcal{SQD}^*(X)$.*

(i) $\lim_{t\to+\infty} t^{-2}\mathscr{H}(\mathrm{Op}(t^2q)) = \frac{1}{4}q$.

(ii) *Under the correspondence between Prym differentials and points in the Prym variety, the spectral data $[\hat{\eta}_t]$ of $\mathrm{Op}(t^2q)$ satisfies*

$$\lim_{t\to+\infty}[\hat{\eta}_t - i\,t\,\mathrm{Im}\,\lambda_{\mathrm{SW}}] = 0 \quad \textit{in } \mathrm{Prym}(\widehat{X}_q, X)/J_2(X),$$

where λ_{SW} is the Seiberg–Witten differential on $\widehat{X}_q$.

[1] Strictly speaking, $\Lambda(q)$ need not be maximal, but this possibility will not play any role in the result.

(iii) (Dumas) *If* $\Gamma_t = \boldsymbol{\sigma}_t + i\,\boldsymbol{\beta}_t$ *is the shear-bend cocycle of Thurston's pleated surface for the projective connection* $Q(t^2 q)$ *in* (1-8), *then*

$$\lim_{t\to+\infty} t^{-1}\Gamma_t = \Gamma_{\mathrm{SW}},$$

where Γ_{SW} *is the complex cocycle determined by the periods of* λ_{SW}; *see* Definition 4.13.

In part (ii) of the corollary, $J_2(X)$ denotes the 2–torsion points of the Jacobian variety of X. Its appearance in the statement of part (ii) is due to the ambiguity in the choice of the square root of K_X. Part (iii) also follows from the results of Dumas in [16; 15].

We refine Corollary 1.3(i) with an error estimate $t^{-2}\mathscr{H}(\mathrm{Op}(t^2 q)) - q/4 = O(t^{-1})$ in Proposition 6.10.

1.3 Further comments

1.3.1 Discussion of the main results The Hitchin parametrization of (an open set in) $M_{\mathrm{H}}(X)$ gives it the structure of a torus fibration over $\mathcal{QD}^*(X)$ which, via the Hopf differential of the harmonic diffeomorphism from X to a hyperbolic surface S, can be identified with Teichmüller space $T(\Sigma)$. Similarly, in the presence of a maximal lamination, Bonahon also parametrizes (an open set in) $R^o(\Sigma)$ as a torus fibration over $T(\Sigma)$. The nonabelian Hodge correspondence is a transcendental homeomorphism between these two pictures. The asymptotic decoupling of the Hitchin equations reflects the conclusion of this paper that these two torus fibrations are essentially asymptotically equivalent.

Previous work on the asymptotics of equivariant harmonic maps focused on the behavior of divergent length functions corresponding to shearing cocycles, and this is well understood. The novelty of the present work is to extract information on the *complex* length, which involves bending. Perhaps not surprisingly, through the nonabelian Hodge correspondence, bending turns out in the gauge theory picture to involve the unitary part of the flat connection.

An important subtlety happens when the quadratic differentials have saddle connections. These may occur in either the horizontal or vertical foliations, or both, and they play different roles. Saddle connections in the vertical foliation give rise to the possibility of folding in the image of harmonic maps. This will be discussed more explicitly in Section 1.2.3 below. More relevant are the saddle connections that appear in the horizontal foliation. In this case, geodesic straightening of the leaves does not produce a maximal lamination, and so a choice of *maximalization* is required; see Section 2.4. Unlike the complicated wall-crossing phenomena that emerge from this situation in other contexts, here in the asymptotic limit the choice of maximalization is a technical tool that amounts to a change of coordinates in the identification (1-7).

In terms of the consequences of the Main Theorem, let us elaborate on the discussion in the introductory paragraphs. The work of [43] analytically describes the frontier of $M_{\mathrm{H}}(X_0)$ as a torus fibration over $\mathcal{SQD}^*(X_0)$ for a chosen Riemann surface X_0; the elements of a fiber are equivalence classes of Prym differentials defined in terms of $q \in \mathcal{SQD}^*(X_0)$. It is not apparent how this parametrization of limiting configurations in terms of the Riemann surface X_0 relates to the one defined in terms of a nearby Riemann surface X. To clarify the question, imagine a pair of sequences of representations $\{\rho_n^+, \rho_n^-\} \subset R^o(\Sigma)$ where the associated solutions to the self-duality equations have distinct limiting configurations in a particular torus fiber over $q_0 \in \mathcal{SQD}^*(X_0)$. If we then change the original choice of Riemann surface from X_0 to X and consider the partial bordification of $M_{\mathsf{H}}(X)$, will the solutions of (2-8) for $\{\rho_n^+, \rho_n^-\}$ on the new surface X still have limiting configurations in a torus fiber over a single quadratic differential $q \in \mathcal{SQD}^*(X)$, or will they accumulate over different fibers? Corollary 1.1 asserts that an entire limiting torus, defined in terms of either Riemann surface, projects to a single point in the Morgan–Shalen compactification, which is defined only in terms of the topologically defined projective vertical measured foliation of q. Thus the tori of limiting configurations defined by X or X_0 are either disjoint or coincide. Moreover, and this property is more subtle, the elements of each torus fiber may be identified by periods of a differential corresponding to the limiting bending cocycle of a sequence of pleated surfaces. These periods are therefore also topologically defined. Turning this discussion around, we thus see that Corollaries 1.1 and 1.2 provide a partial refinement of the Morgan–Shalen compactification, an apparently topological result, via a construction that is geometric-analytical.

Corollary 1.3 provides an appealing picture of the space of complex projective structures. The classical Schwarzian view of the space of complex projective structures is as an affine bundle over Teichmüller space: the fibers over a point $X \in T(\Sigma)$ is the space $\mathcal{QD}(X)$ of quadratic differentials on X, and we can focus our attention on a ray $\{tq, t > 0\} \subset \mathcal{QD}(X)$ of Schwarzian derivatives on X. A basic question is to describe the image of this ray in the fibration $M_{\mathsf{H}}(X)$.

In Corollary 1.3 we see the end of such a ray, when seen as a family in $M_{\mathsf{H}}(X)$, shadows a linear flow on one of the torus boundary fibers. Different rays over a common point X in Teichmüller space shadow flows over distinct tori, depending on the vertical foliation of the common (projective) Schwarzian derivative. In short, the rays over a single Riemann surface have ends accumulating in each fiber of the partial compactification of $M_{\mathsf{H}}(X)$. On the other hand, rays over distinct Riemann surfaces, whose Schwarzians have a common projective vertical measured foliation, shadow flows over a common torus fiber in the partial compactification of $M_{\mathsf{H}}(X)$. Here the direction vectors of the flows reflect the underlying Riemann surface of the family of complex projective structures, through the horizontal foliation of the Schwarzians.

Finally, in the context of this last corollary, we provide a small bit of intuition for these claims, effectively due to Dumas in this setting. A family of complex projective structures over X with proportional Schwarzians tq for $t \gg 0$ may be seen as the "graftings" of a family of surfaces $X = \mathrm{gr}_{\lambda_t}(X_t)$.[2] Here the lamination Λ_t is the bending lamination for a pleated surface whose underlying hyperbolic surface is X_t,

[2] We provide a quick informal introduction to grafting. A complex projective structure will develop as a domain over complex projective space $\mathbb{CP}^1$. We can regard $\mathbb{CP}^1$ as $\partial\mathbb{H}^3$. Then given a pleated surface (S, f, Λ, ρ) in $\mathbb{H}^3$, we can imagine exponentiating in the normal direction from the image $f(S)$ of that pleated surface. The limiting image of a totally geodesic plaque under this flow will inject onto a domain in $\mathbb{CP}^1$ bounded by circular arcs. The image of the bending lamination Λ is more complicated, reflecting the complicated nature of a geodesic lamination, but can be imagined as (limits of) thin crescents that connect the images of plaques: for example, if Λ were only a single simple closed curve γ with bending measure θ, then

and obviously the bending measure grows with t. Of course, the result of the grafting is a fixed Riemann surface X, and so the pruned surface X_t reflects the increased bending by growing thin and long in the direction of the bending lamination. But the representation $[\rho_t] \in R^o(\Sigma)$ of the pleated surface X_t is what we focus on in this paper. We see then that translation lengths for this representation must be growing long, roughly parallel the stretched lamination Λ_t. Passing from these synthetic constructions to geometric analysis by considering the shape of the ρ_t–equivariant energy minimizing map $u_t \colon \widetilde{X} \to \mathbb{H}^3$, the Hopf differential $q_t \in \mathcal{QD}^*(X)$ of u_t will have horizontal foliation in the direction of the maximal stretch of the map, which in this case will be forced to be along the very stretched lamination Λ_t. Indeed, u_t will crowd much of its image near the maximally stretched lamination Λ_t. The regions complementary to those mapping near Λ_t will be stretched to efficiently connect components of Λ_t, and they will thus lie near portions of totally geodesic hyperbolic planes in $\mathbb{H}^3$. Taken together, these heuristics imply that the image of u_t will itself approximate the pleated surface $f \colon \widetilde{S}(q_t) \to \mathbb{H}^3$.

Now, on the one hand, the analysis of [43] tells us, as a starting point, that ρ_t will have a limiting configuration in the fiber defined by the limit of normalized Hopf differentials. The Main Theorem works by recognizing the gauge-theoretic endomorphism that represents such a limit point as an infinitesimal rotation in $\mathbb{H}^3$ about the geodesic tangent to the image of a horizontal leaf.

Ignoring for now the issue of how the harmonic map is bending near Λ_t, we note that near the preimages of Λ_t, the harmonic map is well-approximated by the very simple harmonic map $\mathbb{C} \to \mathbb{H}^3$ which takes horizontal lines in the plane to a geodesic with a parametrization proportional to arc length. As that simple model map has Hopf differential dz^2, we see that we can expect the vertical foliation of the ρ_t harmonic map to predict the length spectrum of the representation of ρ_t, at least up to its leading terms. As the length spectrum of a representation is independent of the choice X of the background Riemann surface, we find evidence for Corollaries 1.1 and 1.2. Finally, Dumas [15] and [16] makes the deep observation that Λ_t is well-approximated by the horizontal lamination of the Schwarzian, and thus the underlying geometric lamination for the bending lamination Λ_t becomes increasingly fixed as t increases, even as the amount of bending grows linearly with the measure of the horizontal foliation, ie linearly with t. That linear change in the complex translation lengths of the dominant elements of the holonomy for ρ_t, coupled with the just mentioned relationship of the gauge theory to hyperbolic geometry, suggests the linear flow in Corollary 1.3.

1.3.2 Relation with other work The literature on Hitchin systems, solutions to differential equations on Riemann surfaces, and their asymptotic properties is vast, and the Main Theorem in this paper may be viewed in that context.

Asymptotic decoupling of the self-duality equations has been studied in Taubes [57], Mazzeo, Swoboda, Weiss and Witt [43; 44], Mochizuki [48] and Fredrickson [23]. This idea is also central to the work of

each lift of γ would force the inclusion of a "lune" of width θ. Thurston observed that each complex projective structure admitted a unique description as a hyperbolic structure S as above, together with the insertion of flat lunes corresponding to the bending lamination. We might call S the "pruning" of the complex projective structure.

Gaiotto, Moore and Neitzke [25] and the conjectural structure of the hyperkähler metric on $M_{\mathsf{H}}(X)$; see Dumas and Neitzke [17] and Fredrickson [24]. The idea of "nonabelianization" also arises from this work and is related to Fock–Goncharov cluster coordinates and the Bonahon parametrization. This has been investigated by Hollands and Neitzke [30; 31] and Fenyes [22].

Corollary 1.3 is a kind of zeroth-order analog of the much more extensive results from the *exact WKB analysis* of Schrödinger equations (see Kawai and Takei [37]) where $t = 1/\hbar$. In particular, the period map

$$Z(\gamma) = \int_\gamma \lambda_{\mathrm{SW}}, \tag{1-10}$$

for γ representing an odd homology class on $\widehat{X}_q$, plays a central role in Gaiotto, Moore and Neitzke [25]. For some recent work, see Iwaki and Nakanishi [33] and Allegretti [1; 2].

1.3.3 Outline of the paper This paper is organized as follows. In Section 2 we have provided a rather large amount of background material in order to make the rest of the paper accessible to a wide readership. The main topics are the moduli space of solutions to the self-duality equations, limiting configurations and their relation to spectral data and Prym differentials. We also provide details on equivariant harmonic maps and their high-energy properties. The section concludes with background on laminations, measured foliations, train tracks and $\mathbb{R}$–trees, which will be useful in the sequel.

These preliminaries are followed in Section 3 by a discussion of "bending". We first introduce a naive geometric notion of how to measure the bending of an immersive map to hyperbolic space in terms of dihedral angles of intersecting tangent planes. In the context of the equivariant maps that appear in the nonabelian Hodge correspondence, we compare this notion to an alternative definition of bending coming from parallel translation in bundles with connections. When Higgs pairs approach a limiting configuration, the gauge-theoretic bending is shown to be determined by the periods of Prym differentials.

In Section 4 we review the notion of a transverse cocycle to a geodesic lamination, as well as Bonahon's parametrization of the character variety $R(\Sigma)$. In Lemma 4.4 we show that under certain assumptions on the pleating locus the bending cocycle can be related to the geometric notion of bending introduced in Section 3. We use this property to derive the bending cocycle of a pleated surface from the gauge-theoretic notion of bending, under the assumption that the pleated surface and the image of the equivariant harmonic map are appropriately close.

The existence of a pleated surface with the properties just mentioned is proven in Section 5. The required asymptotic results for high-energy harmonic maps are largely due to Minsky. The key idea is to compare an arrangement of geodesics in $\mathbb{H}^3$ obtained from the image of horizontal leaves of the foliation by an equivariant harmonic map to the geodesic lamination on the hyperbolic surface corresponding to a harmonic diffeomorphism with the same Hopf differential. We show that by perturbing this hyperbolic structure slightly, the geodesic configuration in $\mathbb{H}^3$ extends to a pleated surface.

Finally, in Section 6 we give the proofs of the Main Theorem and its corollaries.

Acknowledgements The authors wish to express their gratitude to a great many people for their interest in this work and for their helpful comments. Among them are Dylan Allegretti, Francis Bonahon, Marc Burger, Steve Bradlow, Aaron Fenyes, Vladimir Fock, François Labourie, Rafe Mazzeo, Yair Minsky and Andy Neitzke.

Ott was supported by the European Research Council under ERC-Consolidator Grant 614733 *Deformation spaces of geometric structures*, and by the Priority Program 2026 *Geometry at infinity* of the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under DFG grant 340014991. He further acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) grant 281869850 (RTG 2229). He was also supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC 2181/1 grant 390900948 (the Heidelberg STRUCTURES Excellence Cluster). Swoboda is supported by a Heisenberg grant of the German Research Foundation, the DFG Priority Program 2026 *Geometry at infinity*, the DFG RTG 2229 *Asymptotic invariants and limits of groups and spaces* and by DFG under Germany's Excellence Strategy EXC-2181/1 grant 390900948 (the Heidelberg STRUCTURES Excellence Cluster). Wentworth is supported in part by NSF grants DMS-1564373 and DMS-1906403. Wolf acknowledges support from NSF grants DMS-1564374 and DMS-2005551 and the Simons Foundation. Wolf and Wentworth also received funding from NSF grant DMS-1440140 administered by the Mathematical Sciences Research Institute while the authors were in residence at MSRI during the *Holomorphic differentials in mathematics and physics* program, August–December 2019. Finally, the authors acknowledge support from NSF grants DMS-1107452, DMS-1107263 and DMS-1107367 *RNMS*: *GEometric structures And Representation varieties* (the GEAR Network).

2 Background material

2.1 Higgs bundles

2.1.1 The self-duality equations We introduce the setup for Hitchin's self-duality equations for a topologically trivial rank 2 complex vector bundle E in a form that will be useful later on.

Fix a choice of spin structure $K_X^{1/2}$ as in Section 1.1.2 and consider

$$E = K_X^{-1/2} \oplus K_X^{1/2}. \tag{2-1}$$

A choice of conformal metric $ds^2 = m(z)|dz|^2$ on X induces a hermitian metric $h = (m^{1/2}, m^{-1/2})$ on E which will be fixed throughout. Notice that the determinant line bundle $\det E$ with its induced metric from h is canonically trivial. Let $\mathfrak{g}_E$ be the vector bundle of traceless skew-hermitian endomorphisms of E, and $\mathfrak{g}_E^{\mathbb{C}}$ its complexification consisting of all traceless endomorphisms. We will also often use $\sqrt{-1}\mathfrak{g}_E$, the bundle of traceless hermitian endomorphisms. The hermitian metric h on E induces a hermitian metric on the associated endomorphism bundle $\mathfrak{g}_E^{\mathbb{C}}$ which is given by

$$\langle A, B\rangle = \tfrac{1}{2}\operatorname{tr}(AB^{*_h}) \quad \text{for } A, B \in \Gamma(\mathfrak{g}_E^{\mathbb{C}}).$$

On the subbundle of traceless hermitian endomorphisms this metric reads $\langle A, B\rangle = \frac{1}{2}\operatorname{tr}(AB)$.

Denote by $\mathcal{A}(E,h)$ the space of smooth connections on E that are unitary with respect to h and which induce the trivial connection on $\det E$. This is an affine space modeled on $\Omega^1(X,\mathfrak{g}_E)$. A $\bar\partial$–operator $\bar\partial_E$ on E defines a holomorphic bundle $\mathcal{E}$ which we will often denote by $\mathcal{E} = (E,\bar\partial_E)$. There is a connection $A = (\bar\partial_E, h) \in \mathcal{A}(E,h)$ called the *Chern connection* that is uniquely determined by the requirement

$$\bar\partial_A := (d_A)^{0,1} = \bar\partial_E.$$

In this way, $\mathcal{A}(E,h)$ is identified with the space of $\bar\partial$–operators on E. Similarly, there is a real linear isomorphism

$$\Omega^1(X,\sqrt{-1}\mathfrak{g}_E) \xrightarrow{\sim} \Omega^{1,0}(X,\mathfrak{g}_E^{\mathbb{C}}), \quad \Psi \mapsto \Phi = \Psi^{1,0}, \tag{2-2}$$

with inverse

$$\Psi = \Phi + \Phi^{*_h}. \tag{2-3}$$

We shall often use this convention, $\Phi \leftrightarrow \Psi$, for the isomorphism (2-2).

A *Higgs bundle* is a pair $(\bar\partial_E,\Phi)$, where $\bar\partial_E\Phi = 0$. The *Higgs field* Φ may either be regarded as a holomorphic one-form valued in the sheaf $\operatorname{End}_0\mathcal{E}$ of traceless endomorphisms of $\mathcal{E}$, or as a holomorphic section of $\mathcal{E}nd_0\mathcal{E}\otimes K_X$. The context throughout will make clear which interpretation applies.

For a pair $(A,\Psi) \in \mathcal{A}(E,h)\times\Omega^1(X,\sqrt{-1}\mathfrak{g}_E)$, the system of PDEs

$$\begin{cases} F_A + [\Psi\wedge\Psi] = 0, \\ d_A\Psi = 0, \\ d_A(*\Psi) = 0, \end{cases} \tag{2-4}$$

is called the *self-duality equations*. A solution (A,Ψ) gives a Higgs bundle $(\bar\partial_A,\Phi)$. The holomorphicity of Φ follows from the last two equations in (2-4). Conversely, Hitchin shows that for a polystable Higgs bundle $(\bar\partial_E,\Phi)$ there is a complex gauge transformation g such that the Chern connection and Ψ associated to $g\cdot(\bar\partial_E,\Phi)$ give a solution to (2-4). Polystability will not play a role in this paper, so we omit its definition. We frequently refer to a solution (A,Ψ) of (2-4) as a *Higgs pair*.

Let $M_{\mathsf{H}}(X)$ denote the moduli space of unitary gauge equivalence classes of solutions of (2-4). Then $M_{\mathsf{H}}(X)$ is a quasiprojective variety of dimension $6g-6$. By a slight abuse of notation, when $(\bar\partial_E,\Phi)$ is polystable and (A,Ψ) the associated solution to (2-4) as in the previous paragraph, we shall write $[(\bar\partial_E,\Phi)]$ to mean the gauge equivalence class $[(A,\Psi)] \in M_{\mathsf{H}}(X)$.

A very important fact used in this paper is the following: If (A,Ψ) is a solution to (2-4) then the $\mathrm{SL}(2,\mathbb{C})$–connection

$$\nabla := d_A + \Psi \tag{2-5}$$

is flat. This follows from the first two equations of (2-4).

2.1.2 Quadratic differentials Recall the notation $\mathcal{QD}(X)$, $\mathcal{QD}^*(X)$, $\mathcal{SQD}(X)$ and $\mathcal{SQD}^*(X)$ from Section 1.1.1. We define a norm on $\mathcal{QD}(X)$ by

$$\|q\|_1 := \int_X |q(z)| \, \frac{i}{2} \, dz \wedge d\bar{z},$$

where $q = q(z)dz^2$ in local conformal coordinates, and we let $Z(q)$ denote the set of zeroes of q.

The map

$$\mathscr{H}\colon M_{\mathsf{H}}(X) \to \mathcal{QD}(X), \quad [(A,\Psi)] \mapsto \tfrac{1}{2}\operatorname{tr}(\Psi \otimes \Psi)^{2,0} = -\det \Phi, \tag{2-6}$$

is holomorphic, proper and surjective. Its restriction to $M^*_{\mathsf{H}}(X) := \mathscr{H}^{-1}(\mathcal{QD}^*(X))$ is a fibration with fibers consisting of half-dimensional complex tori; see Section 2.2 below.

As shown in [29], the Hitchin fibration $\mathscr{H}$ has a global section described as follows. The bundle E has a distinguished holomorphic structure $\bar\partial_0$ coming from the splitting (2-1) and the holomorphic structure on $K_X^{\pm 1/2}$. Let A_0 be the Chern connection associated to $(\bar\partial_0, h)$. Then the section is given by

$$\mathscr{S}_H\colon \mathcal{QD}(X) \to M_{\mathsf{H}}(X), \quad q \mapsto \left[\left(\bar\partial_{A_0}, \Phi = \begin{pmatrix} 0 & 1 \\ q & 0 \end{pmatrix}\right)\right] \tag{2-7}$$

(recall the convention concerning the notation $[(\bar\partial_E, \Phi)]$ from the previous section).

2.1.3 The partial compactification by limiting configurations Properness of the Hitchin fibration implies that every sequence (A_n, Ψ_n), with $n \in \mathbb{N}$, of solutions of equation (2-4) such that the sequence $q_n = -\det \Phi_n$ (recall $\Phi = \Psi^{1,0}$) is bounded has a subsequence that converges smoothly modulo the action of unitary gauge transformations. Conversely, a sequence (A_n, Ψ_n) diverges if the sequence q_n of holomorphic quadratic differentials diverges, ie $\|q_n\|_1 \to \infty$ as $n \to \infty$. Notice that the latter is equivalent to $\|\Psi_n\|_2 \to \infty$ as $n \to \infty$ (here the subscript refers to the L^2–norm).

By the results in [43] (see also [57]), the open and dense region $M^*_{\mathsf{H}}(X)$ of $M_{\mathsf{H}}(X)$ admits a bordification by the set $\partial M^*_{\mathsf{H}}(X)$ of so-called limiting configurations, as we explain next. To this end, we introduce the *decoupled self-duality equations*

$$\begin{cases} F_A = 0, \; [\Psi \wedge \Psi] = 0, \\ d_A \Psi = 0, \\ d_A(*\Psi) = 0, \end{cases} \tag{2-8}$$

for a Higgs field Ψ and a unitary connection A.

Definition 2.1 Let $q \in \mathcal{QD}^*(X)$. A pair (A,Ψ) is called a *limiting configuration for q* if $\det \Phi = -q$ and (A,Ψ) is a smooth solution of (2-8) on the punctured surface $X_q^\times := X \setminus Z(q)$.

This definition only applies to solutions for differentials $q \in \mathcal{QD}^*(X)$. We refer to [48] for the definition and description of limiting configurations for points $q \in \mathcal{QD}(X) \setminus \mathcal{QD}^*(X)$.

Example 2.2 Recall the connection A_0 from (2-7). For $q \in \mathcal{QD}^*(X)$, we define

$$A^0_\infty(q) = A_0 + \tfrac{1}{2}(\operatorname{Im} \bar\partial \log \|q\|) \begin{pmatrix} -i & 0 \\ 0 & i \end{pmatrix},$$
$$\Phi_\infty(q) = \begin{pmatrix} 0 & \|q\|^{1/2} \\ \|q\|^{-1/2} q & 0 \end{pmatrix}, \tag{2-9}$$
$$\Psi_\infty(q) = \Phi_\infty(q) + \Phi^{*_h}_\infty(q),$$

where $\|q\|$ means the (pointwise) norm with respect to the conformal metric ds^2. The pair $(A^0_\infty(q), \Psi_\infty(q))$ is a limiting configuration for q. It will later become important as the limiting configuration corresponding to a pleated surface with zero bending cocycle. We therefore call it the *Fuchsian limiting configuration* associated to q.

We shall often write $(A^0_\infty, \Psi_\infty)$, where the quadratic differential is understood. More generally, any other limiting configuration (A_∞, Ψ_∞) representing a point in the fiber $\mathscr{H}^{-1}(q)$ is of the form

$$A_\infty = A^0_\infty + \eta, \quad [\eta \wedge \Psi_\infty] = 0 \quad \text{and} \quad d_{A^0_\infty} \eta = 0, \tag{2-10}$$

where $\eta \in \Omega^1(X^\times_q, \mathfrak{g}_E)$. The group $\mathcal{G} = \mathcal{G}(E, h)$ of unitary gauge transformations of E acts on the space of solutions (A_∞, Ψ_∞) to equation (2-8), and we define the moduli space

$$\partial M^*_{\mathrm{H}}(X) = \{\text{all solutions to (2-8) for } q \in \mathcal{QD}^*(X)\}/\mathcal{G} \times \mathbb{R}^+.$$

Here we follow the original definition of limiting configurations in [43], where each (A_∞, Ψ_∞) is assumed to take a particular normal form in disks $\mathcal{D}_p$ around each zero of q. This normal form is given on each $\mathcal{D}_p$ by the Fuchsian limiting configuration $(A^0_\infty(q), \Phi_\infty(q))$ and identically vanishing $\eta \equiv 0$. In particular, it can be assumed that η extends over the points $q \in \mathcal{QD}(X) \setminus \mathcal{QD}^*(X)$ and therefore is an element of $\Omega^1(X_q, \mathfrak{g}_E)$. With this restriction, we divide out by unitary gauge transformations that are the identity near each $\mathcal{D}_p$; cf [44].

Second, since there is an equivalence up to positive real multiples of Ψ, it is natural to define the projection

$$\mathscr{H}_\infty \colon \partial M^*_{\mathrm{H}}(X) \to \mathcal{SQD}^*(X) \tag{2-11}$$

defined by mapping $(A_\infty, \Psi_\infty) \mapsto q/\|q\|_1$, where $q = -\det \Phi_\infty$.

We now describe the structure of the set $\partial M^*_{\mathrm{H}}(X)$ of limiting configurations more closely, summarizing the results in [43, Section 4.4]. For $(A_\infty, \Psi_\infty) \in \mathscr{H}^{-1}_\infty(q)$, define the real line bundle $L_q \to X^\times_q$ by

$$L_q = \{\eta \in \mathfrak{g}_E \mid [\Phi_\infty \wedge \eta] = 0\}. \tag{2-12}$$

Let $L^{\mathbb{C}}_q = L_q \otimes_{\mathbb{R}} \mathbb{C}$ denote the complexification. Then L_q and $L^{\mathbb{C}}_q$ are d_{A_∞}–invariant line subbundles of $\mathfrak{g}_E$ and $\mathfrak{g}^{\mathbb{C}}_E$, real and complex, respectively. Notice that the second component Φ_∞ of a limiting configuration is completely determined modulo unitary gauge by the holomorphic quadratic differential q.

Hence, the flat bundle L_q also only depends on q, which justifies the notation. The ungauged vertical deformation space at (A_∞, Φ_∞) is identified with

$$Z^1(X_q^\times; L_q) := \{\eta \in \Omega^1(X_q^\times, L_q) \mid d_{A_\infty^0}\eta = 0\}.$$

Next consider the subgroup $\mathrm{Stab}_{\Phi_\infty}$ of unitary gauge transformations which stabilize Φ_∞. If $g \in \mathrm{Stab}_{\Phi_\infty}$ lifts to a section of L_q, ie $g = \exp(\gamma)$, with $\gamma \in \Omega^0(X_q^\times, L_q)$, then g acts on $A_\infty = A_\infty^0 + \eta$, for $\eta \in \Omega^1(X_q^\times, L_q)$, by

$$g(A_\infty) = g^{-1}\eta g + g^{-1}(d_{A_\infty} g) = \eta + d_{A_\infty^0}\gamma.$$

(Recall that L_q is an A_∞–parallel line subbundle of $\mathfrak{g}_E$, so $g^{-1}\eta g = \eta$ and $d_{A_\infty}\exp(\gamma) = \exp(\gamma) d_{A_\infty}\gamma$.) Hence the infinitesimal vertical deformation space is

$$H^1(X_q^\times; L_q) = Z^1(X_q^\times; L_q)/B^1(X_q^\times; L_q), \quad \text{where } B^1(X_q^\times; L_q) := \{d_{A_\infty}\gamma \mid \gamma \in \Omega^0(X_q^\times, L_q)\}.$$

If all zeroes of q are simple, then

$$\dim_{\mathbb{R}} H^1(X_q^\times; L_q) = 6g - 6,$$

where g is the genus of Σ. To obtain the moduli space, we must also divide the infinitesimal deformation space by the residual action of the component group $\pi_0(\mathrm{Stab}_{\Phi_\infty})$. Under the correspondence above, this consists of an integral lattice $H^1_{\mathbb{Z}}(X_q^\times, L_q)$ under the exponential map.

Proposition 2.3 *The moduli space of limiting configurations with a fixed $q \in \mathcal{QD}^*(X)$ is*

$$\mathscr{H}_\infty^{-1}(q) = H^1(X_q^\times, L_q)/H^1_{\mathbb{Z}}(X_q^\times, L_q).$$

This is a torus of real dimension $6g - 6$.

2.1.4 Approximate solutions Following [43, Section 3.2], for suitable functions f, h and χ to be specified below, we define the family of *approximate solutions* $S_t^{\mathrm{app}}(q) := (A_t^{\mathrm{app}}(q) + \eta, t\,\Psi_t^{\mathrm{app}}(q))$ by

$$\begin{aligned} A_t^{\mathrm{app}}(q) &:= A_0 + \left(\tfrac{1}{2} + \chi(\|q\|)\left(4f_t(\|q\|) - \tfrac{1}{2}\right)\right) \operatorname{Im} \bar\partial \log \|q\| \begin{pmatrix} -i & 0 \\ 0 & i \end{pmatrix}, \\ \Phi_t^{\mathrm{app}}(q) &:= \begin{pmatrix} 0 & \|q\|^{1/2} e^{\chi(\|q\|)h_t(\|q\|)} \\ \|q\|^{-1/2} e^{-\chi(\|q\|)h_t(\|q\|)} q & 0 \end{pmatrix}, \\ \Psi_t^{\mathrm{app}}(q) &:= \Phi_t^{\mathrm{app}}(q) + (\Phi_t^{\mathrm{app}}(q))^{*_h}. \end{aligned} \tag{2-13}$$

Regarding the formula for Ψ_t^{app}, we follow our convention that $\Phi = \Psi^{1,0}$; cf the beginning of Section 2.1.3. We may view these approximate solutions as desingularizations of the limiting configurations introduced before. Indeed, as $t \to \infty$ there is smooth local convergence $A_t^{\mathrm{app}}(q) \to A_\infty^0(q)$ and $\Phi_t^{\mathrm{app}}(q) \to \Phi_\infty(q)$ on $X_q^\times$. Here the one-form η satisfies (2-10) and is considered as an element of $\Omega^1(X_q, \mathfrak{g}_E)$.

We now turn to a more detailed explanation of the functions used to define the approximate solution in (2-13). Here $h_t(r)$ is the unique solution to $(r\partial_r)^2 h_t = 8t^2 r^3 \sinh(2h_t)$ on $\mathbb{R}^+$ with specific asymptotic properties at 0 and ∞, and $f_t := \frac{1}{8} + \frac{1}{4} r\, \partial_r h_t$. Further, $\chi\colon \mathbb{R}^+ \to [0,1]$ is a suitable cutoff-function. The parameter t can be removed from the equation for h_t by substituting $\rho = \frac{8}{3} t r^{3/2}$; thus if we set $h_t(r) = \psi(\rho)$ and note that $r\partial_r = \frac{3}{2}\rho\partial_\rho$, then

$$(\rho\partial_\rho)^2\psi = \tfrac{1}{2}\rho^2 \sinh(2\psi).$$

This is a Painlevé III equation; there exists a unique solution which decays exponentially as $\rho \to \infty$ and with asymptotics as $\rho \to 0$ ensuring that A_t^{app} and Φ_t^{app} are regular at $r = 0$. More specifically,

- $\psi(\rho) \sim -\log\bigl(\rho^{1/3}\bigl(\sum_{j=0}^{\infty} a_j \rho^{4j/3}\bigr)$ as $\rho \downarrow 0$,
- $\psi(\rho) \sim K_0(\rho) \sim \rho^{-1/2} e^{-\rho} \sum_{j=0}^{\infty} b_j \rho^{-j}$ as $\rho \uparrow \infty$,
- $\psi(\rho)$ is monotonically decreasing (and strictly positive) for $\rho > 0$.

These are asymptotic expansions in the classical sense, ie the difference between the function and the first N terms decays like the next term in the series, and there are corresponding expansions for each derivative. The function $K_0(\rho)$ is the Bessel function of imaginary argument of order 0.

In the following result, any constant C which appears in an estimate is assumed to be independent of t.

Lemma 2.4 [43, Lemma 3.4] *The functions $f_t(r)$ and $h_t(r)$ have the following properties*:

(i) *As a function of r, f_t has a double zero at $r = 0$ and increases monotonically from $f_t(0) = 0$ to the limiting value $\frac{1}{8}$ as $r \uparrow \infty$. In particular, $0 \leqslant f_t \leqslant \frac{1}{8}$.*

(ii) *As a function of t, f_t is also monotone increasing. Further, $\lim_{t\uparrow\infty} f_t = f_\infty \equiv \frac{1}{8}$ uniformly in C^∞ on any half-line $[r_0, \infty)$, for $r_0 > 0$.*

(iii) *There are estimates*

$$\sup_{r>0} r^{-1} f_t(r) \leqslant C t^{2/3} \quad \text{and} \quad \sup_{r>0} r^{-2} f_t(r) \leqslant C t^{4/3}.$$

(iv) *When t is fixed and $r \downarrow 0$, then $h_t(r) \sim -\frac{1}{2}\log r + b_0 + \cdots$, where b_0 is an explicit constant. On the other hand, $|h_t(r)| \leqslant C \exp\bigl(-\frac{1}{8} t r^{3/2}\bigr)/(t r^{3/2})^{1/2}$ for $t \geqslant t_0 > 0$ and $r \geqslant r_0 > 0$.*

(v) *Finally,*

$$\sup_{r\in(0,1)} r^{1/2} e^{\pm h_t(r)} \leqslant C \quad \text{for } t \geqslant 1.$$

It follows from the results in [43] that the approximate solution S_t^{app} satisfies the self-duality equations up to an exponentially decaying error as $t \to \infty$ (which is uniform on the closed surface X), and there is an exact solution $(A_t, t\Phi_t)$ in its complex gauge orbit (unique up to real gauge transformations) which is no further than $Ce^{-\beta t}$ pointwise away (with respect to any C^ℓ–norm) for some $\beta > 0$.

2.1.5 Converging families of Higgs pairs For a holomorphic quadratic differential $q \in \mathcal{QD}^*(X)$, recall the fiber $\mathscr{H}^{-1}(q)$, where $\mathscr{H}$ is the Hitchin map (2-6).

Definition 2.5 Consider a family $[(A_t, t\Psi_t)] \in \mathscr{H}^{-1}(t^2 q_t)$, where $q_t \in \mathcal{SQD}^*(X)$ and $q_t \to q \in \mathcal{SQD}^*(X)$. Then $[(A_t, t\Psi_t)]$ is said to *converge* to $[(A_\infty, \Psi_\infty)] \in \mathscr{H}_\infty^{-1}(q)$ as $t \to \infty$ if, after passing to a subsequence and modifying by unitary gauge transformations (which we suppress from the notation), the family of pairs (A_t, Ψ_t) satisfies the following:

- **Convergence** The sequence (A_t, Ψ_t) converges to (A_∞, Ψ_∞) as $t \to \infty$ in $L^p(X)$ for all $1 \leqslant p < 2$, locally in $C^\ell(X_q^\times)$ for all $\ell \geqslant 0$ at an exponential rate in t.
- **Singularities** For every zero $p \in Z(q_t)$, locally on the punctured disk $\mathcal{D}_p^\times$ (equipped with polar coordinates (r, θ)) the connections A_t are in radial gauge,
$$A_t = F_t \begin{pmatrix} -i & 0 \\ 0 & i \end{pmatrix} d\theta,$$
for some uniformly C^0–bounded family of smooth functions $F_t : \mathcal{D}_p \to \mathbb{R}$ such that $F_t \to F_\infty$ pointwise for some smooth function $F_\infty : \mathcal{D}_p \to \mathbb{R}$ as $t \to \infty$.
- **Approximation** For every integer $\ell \geq 0$ there exist constants $\beta, C > 0$, not depending on t, and a one-form $\eta_t \in \Omega^1(X, \mathfrak{g}_E)$ satisfying $[\Phi_\infty(q_t), \eta_t] = 0$ and $d_{A_\infty(q_t)}\eta_t = 0$ such that
$$\|(A_t, \Psi_t) - (A_t^{\mathrm{app}}(q_t) + \eta_t, \Psi_t^{\mathrm{app}}(q_t))\|_{C^\ell(X)} \leqslant Ce^{-\beta t}$$
for all $t > 0$.

Theorem 2.6 [43] *Every family $[(A_t, t\,\Psi_t)] \in \mathscr{H}^{-1}(t^2 q_t)$ with $q_t \in K \subseteq \mathcal{SQD}^*(X)$, where K is any compact subset, subconverges to a limiting pair $[(A_\infty, \Psi_\infty)] \in \mathscr{H}_\infty^{-1}(q_\infty)$ as $t \to \infty$ in the sense of Definition 2.5. Conversely, every limiting configuration arises in this way.*

Proof By compactness of the set K, one has subconvergence $q_t \to q_\infty$ for some $q_\infty \in K$. The main part of the assertion follows from [43, Theorem 6.6] and its proof, which yields the (Approximation) axiom for any such family of Higgs pairs. There, only a polynomial bound for the C^ℓ–norm of the difference $(A_t^{\mathrm{app}}(q_t) + \eta_t, \Phi_t^{\mathrm{app}}(q_t)) - (A_t, \Phi_t)$ is stated, but the proof shows that it can be improved to an exponential bound. The other two axioms then follow, since they are satisfied by the approximating family $(A_t^{\mathrm{app}}(q_t) + \eta_t, t\,\Phi_t^{\mathrm{app}}(q_t))$ (again by the construction in [43]), and therefore also by $[(A_t, t\,\Phi_t)]$. The last statement is contained in [43, Theorem 1.2]. □

2.2 Spectral curves

2.2.1 The BNR correspondence Let $\pi : |K_X| \to X$ be the projection from the total space $|K_X|$ of K_X. There is a tautological section λ of the holomorphic line bundle $\pi^* K_X \to |K_X|$. Given $q \in \mathcal{QD}^*(X)$, the pullback $\pi^* q$ is a section of $\pi^* K_X^2 \to |K_X|$. Let

(2-14) $$\hat{X}_q = \{\hat{x} \in |K_X| \mid \lambda^2(\hat{x}) = \pi^* q(\hat{x})\} \subset |K_X|.$$

Then $\widehat{X}_q$ is a compact Riemann surface called the *spectral curve* associated to q (nonsingular, since $q \in \mathcal{QD}^*(X)$ has simple zeroes). The restriction of the projection, $\pi\colon \widehat{X}_q \to X$, realizes $\widehat{X}_q$ as a ramified double covering of X with simple branch points at the zeroes $Z(q)$ of q. By the Riemann–Hurwitz formula, $\widehat{X}_q$ has genus $4g-3$. Moreover, $\widehat{X}_q$ admits an involution $\hat{x} \mapsto -\hat{x}$, which we denote by σ.

Recall that the *Prym variety* associated to the covering is

$$\operatorname{Prym}(\widehat{X}_q, X) = \{\mathcal{L} \in \operatorname{Pic}(\widehat{X}_q) \mid \sigma^*\mathcal{L} \simeq \mathcal{L}^*\}.$$

Theorem 2.7 [3, Proposition 3.6] *There is a one-to-one correspondence between points in* $\operatorname{Prym}(\widehat{X}_q, X)$ *and isomorphism classes of Higgs bundles* $(\mathcal{E}, \Phi)$ *with* $\det \Phi = -q$.

The association in Theorem 2.7 goes as follows; see also [44, Section 2.2]. Recall that we have fixed a square root $K_X^{1/2}$. Given $\mathcal{L} \in \operatorname{Prym}(\widehat{X}_q, X)$, let $\mathcal{U} = \mathcal{L} \otimes \pi^*(K_X^{1/2})$. Then $\mathcal{E} = \pi_*(\mathcal{U})$ is a rank 2 holomorphic bundle on X with trivial determinant, and multiplication by λ gives a map

$$\Phi\colon \mathcal{E} = \pi_*(\mathcal{U}) \xrightarrow{\lambda} \pi_*(\mathcal{U} \otimes \pi^*(K_X)) = \mathcal{E} \otimes K_X,$$

with $\det \Phi = -q$. In the other direction, given a Higgs bundle $(\mathcal{E}, \Phi)$, $\mathcal{U}$ is defined (cf [3, Remark 3.7]) by the exact sequence

$$0 \to \mathcal{U} \otimes \mathcal{I}_Z \to \pi^*(\mathcal{E}) \xrightarrow{\lambda - \pi^*\Phi} \pi^*(\mathcal{E} \otimes K_X) \to \mathcal{U} \otimes \pi^*(K_X) \to 0, \tag{2-15}$$

where $\mathcal{I}_Z$ is the ideal sheaf of $Z = Z(q)$ and we regard $\pi^*\Phi$ as a holomorphic section of $\pi^*(\operatorname{End}_0 \mathcal{E} \otimes K_X)$. Since the details of this will be important in the sequel, we briefly elaborate equation (2-15). The first thing to note is that we have an exact sequence

$$0 \to \pi^*(\mathcal{E}) \to \mathcal{U} \oplus \sigma^*(\mathcal{U}) \to \mathcal{U} \otimes \mathcal{O}_Z \to 0. \tag{2-16}$$

The last map is given by mapping sections $(u, v) \in \mathcal{U} \oplus \sigma^*(\mathcal{U})$ to $u(p) - v(p)$, for $p \in Z$. To prove this statement, let $A \subset \widehat{X}_q$ be an open set. Then by definition, as $\mathcal{O}_A$–modules,

$$\pi^*(\mathcal{E})(A) = \pi_*(\mathcal{U})(\pi(A)) \otimes_{\mathcal{O}_{\pi(A)}} \mathcal{O}_A = \mathcal{U}(\pi^{-1}\pi(A)) \otimes \mathcal{O}_A = \mathcal{U}(A \cup \sigma(A)) \otimes \mathcal{O}_A.$$

Now, as an $\mathcal{O}_A$–module, $\mathcal{U}(\sigma(A)) = \sigma^*(\mathcal{U})(A)$. Hence, local sections of $\pi^*(\mathcal{E})$ are sections of $\mathcal{U} \oplus \sigma^*(\mathcal{U})$ that agree at Z; thus, (2-16).

Let $(u, v) \in \pi^*(\mathcal{E}) \subset \mathcal{U} \oplus \sigma^*(\mathcal{U})$. Now Φ acts by multiplication by λ. Since $\sigma^*\lambda = -\lambda$, we have $\pi^*\Phi(u, v) = (\lambda u, -\lambda v)$. Therefore, $(u, v) \in \ker(\lambda - \pi^*\Phi)$ if and only if $v = 0$. The condition in (2-16) forces the image to consist of sections u of $\mathcal{U}$ that vanish at Z, which is the first term in (2-15). Similarly, the image of $\lambda - \pi^*\Phi$ consists of local sections of the form $(0, 2\lambda v)$, ie sections of $\sigma^*(\mathcal{U}) \otimes \mathcal{I}_Z \otimes \pi^*(K_X)$. This is precisely the kernel of the projection $\pi^*(\mathcal{E}) \otimes \pi^*(K_X) \to \mathcal{U} \otimes \pi^*(K_X)$ given by projection onto the first factor. This proves exactness of (2-15).

Remark 2.8 The following will be important.

(i) For $t > 0$ there is a natural biholomorphism $f_t\colon \widehat{X}_q \to \widehat{X}_{t^2 q}$, and pulling back line bundles gives an isomorphism $\mathrm{Prym}(\widehat{X}_{t^2q}, X) \xrightarrow{\sim} \mathrm{Prym}(\widehat{X}_q, X)$. Under this correspondence and Theorem 2.7, $(\mathcal{E}, t\Phi) \mapsto (\mathcal{E}, \Phi)$.

(ii) Note that $d\pi$ is a holomorphic section of $K_{\widehat{X}_q} \otimes \pi^* K_X^{-1}$ with simple zeroes at $Z(q)$. Since $\lambda \in H^0(\widehat{X}_q, \pi^* K_X)$ also vanishes at $Z(q)$, it follows that $K_{\widehat{X}_q} = \pi^* K_X^2$.

2.2.2 Prym differentials Let $\pi : \widehat{X}_q \to X = \widehat{X}_q/\sigma$ be as in the previous section. Recall the exponential sequence

$$0 \to 2\pi i \mathbb{Z} \to \mathcal{O}_{\widehat{X}_q} \xrightarrow{\exp} \mathcal{O}^*_{\widehat{X}_q} \to 1,$$

and associated long exact sequence in cohomology

$$0 \to H^1(\widehat{X}_q, 2\pi i\mathbb{Z}) \to H^1(\widehat{X}_q, \mathcal{O}_{\widehat{X}_q}) \to H^1(\widehat{X}_q, \mathcal{O}^*_{\widehat{X}_q}) \xrightarrow{2\pi i c_1} H^2(\widehat{X}_q, 2\pi i\mathbb{Z}) \to 0.$$

This gives an identification

$$p\colon H^1(\widehat{X}_q, \mathcal{O}_{\widehat{X}_q})/H^1(\widehat{X}_q, 2\pi i\mathbb{Z}) \xrightarrow{\sim} \mathrm{Pic}_0(\widehat{X}_q) := \ker c_1 \subset H^1(\widehat{X}_q, \mathcal{O}^*_{\widehat{X}_q}).$$

Via the Dolbeault isomorphism, we obtain an isomorphism

$$\delta\colon H^{0,1}_{\overline{\partial}}(\widehat{X}_q)/H^1(\widehat{X}_q, 2\pi i\mathbb{Z}) \xrightarrow{\sim} H^1(\widehat{X}_q, \mathcal{O}_{\widehat{X}_q})/H^1(\widehat{X}_q, 2\pi i\mathbb{Z}). \tag{2-17}$$

Now, consider a $\overline{\partial}$–operator $\overline{\partial}_L = \overline{\partial} + \alpha$ on a trivial complex line bundle L, where $\alpha \in \Omega^{0,1}(\widehat{X}_q)$. Let $\mathcal{L}$ denote the associated holomorphic bundle. Then α defines a class $[\alpha] \in H^{0,1}_{\overline{\partial}}(\widehat{X}_q)/H^1(\widehat{X}_q, 2\pi i\mathbb{Z})$, and $\mathcal{L}$ defines a class $[\mathcal{L}] \in \mathrm{Pic}_0(\widehat{X}_q)$. We have the following well-known result.

Lemma 2.9
$$p \circ \delta(-[\alpha]) = [\mathcal{L}].$$

The map $\alpha \mapsto \alpha - \overline{\alpha}$ gives a real isomorphism $H^{0,1}_{\overline{\partial}}(\widehat{X}_q) \simeq H^1(\widehat{X}_q, i\mathbb{R})$. Combined with the Lemma 2.9 we have

$$\mathrm{Pic}_0(\widehat{X}_q) \simeq H^1(\widehat{X}_q, i\mathbb{R})/H^1(\widehat{X}_q, 2\pi i\mathbb{Z}). \tag{2-18}$$

The involution σ acts on $H^1(\widehat{X}_q, i\mathbb{R})$, giving a decomposition into even and odd cohomology

$$H^1(\widehat{X}_q, i\mathbb{R}) = H^1_{\mathsf{ev}}(\widehat{X}_q, i\mathbb{R}) \oplus H^1_{\mathsf{odd}}(\widehat{X}_q, i\mathbb{R}).$$

Clearly, $H^1_{\mathsf{ev}}(\widehat{X}_q, i\mathbb{R}) \simeq H^1(X, i\mathbb{R})$. Let

$$H^1_{\mathsf{odd}}(\widehat{X}_q, 2\pi i\mathbb{Z}) := H^1(\widehat{X}_q, 2\pi i\mathbb{Z}) \cap H^1_{\mathsf{odd}}(\widehat{X}_q, i\mathbb{R}). \tag{2-19}$$

Using (2-18), we see that there is an isomorphism

$$\mathrm{Prym}(\widehat{X}_q, X) \simeq H^1_{\mathsf{odd}}(\widehat{X}_q, i\mathbb{R})/H^1_{\mathsf{odd}}(\widehat{X}_q, 2\pi i\mathbb{Z}). \tag{2-20}$$

Canonical representatives of elements of $H^1_{\mathrm{odd}}(\widehat{X}_q, i\mathbb{R})$ are given by odd, imaginary, harmonic forms, and the space of such will be denoted by $\mathcal{H}^1_{\mathrm{odd}}(\widehat{X}_q, i\mathbb{R})$. We call $\mathcal{H}^1_{\mathrm{odd}}(\widehat{X}_q, \mathbb{C})$ the space of *harmonic Prym differentials*. Let $H^0_{\mathrm{odd}}(\widehat{X}_q, K_{\widehat{X}_q})$ denote the space of holomorphic differentials on $\widehat{X}_q$ that are odd with respect to the involution. We shall call $H^0_{\mathrm{odd}}(\widehat{X}_q, K_{\widehat{X}_q})$ the space of *holomorphic Prym differentials*.[3] We have an isomorphism

$$\mathcal{H}^1_{\mathrm{odd}}(\widehat{X}_q, i\mathbb{R}) \xrightarrow{\sim} H^0_{\mathrm{odd}}(\widehat{X}_q, K_{\widehat{X}_q}), \quad \widehat{\eta} \mapsto \widehat{\eta}^{1,0}. \tag{2-21}$$

There is a distinguished nontrivial holomorphic Prym differential associated to the Liouville form on $|K_X|$. The dual of the tangent sequence associated to π gives

$$0 \to \pi^* K_X \xrightarrow{(d\pi)^t} T^*|K_X| \to \pi^* K_X^{-1} \to 0.$$

The (holomorphic) Liouville 1–form on $|K_X|$ is by definition $(d\pi)^t \circ \lambda = \lambda \circ d\pi$ (where $\circ$ means the dual pairing). Its restriction to $\widehat{X}_q$ is a holomorphic one-form on $\widehat{X}_q$ that is odd with respect to the involution. We call this the *Seiberg–Witten differential* λ_{SW}. Viewing the restriction of $d\pi$ as a section of $T^*\widehat{X}_q \otimes \pi^* TX$, we have

$$\lambda_{\mathsf{SW}} := \lambda \otimes d\pi. \tag{2-22}$$

We shall see in Corollary 1.3 that the spectral data in $\mathrm{Prym}(\widehat{X}_q, X)$ associated to the Seiberg–Witten differential are closely related to complex projective structures.

Remark 2.10 It is customary in the literature to suppress $d\pi$ from the notation in (2-22) and denote the form λ_{SW} on $\widehat{X}_q$ simply by λ. Note that the latter is a section of $\pi^* K_X$ and not $K_{\widehat{X}_q} \simeq \pi^* K_X^2$. Since both of these differentials will figure prominently below, we prefer to keep the notational distinction.

2.2.3 Prym differentials and limiting configurations Continue with the notation of the previous section. Let us introduce

$$W_2 = \begin{pmatrix} 0 & \lambda^{-1}\|\lambda\| \\ \lambda\|\lambda\|^{-1} & 0 \end{pmatrix} \in \mathrm{End}(\pi^* E). \tag{2-23}$$

Then we may write

$$\pi^*\Phi_\infty = \lambda_{\mathsf{SW}} \otimes W_2, \tag{2-24}$$

where Φ_∞ is defined in (2-9), λ_{SW} in (2-22), and we emphasize that here we regard $\pi^*\Phi$ as the pullback of an endomorphism valued one-form. It follows that W_2 lies in $\pi^* L_q^{\mathbb{C}}$. Moreover, W_2 is hermitian, and by a direct computation (cf the proof of Proposition 3.3) we have that $d_{\widehat{A}^0_\infty} W_2 = 0$.

Let $\widehat{\eta} \in \mathcal{H}^1_{\mathrm{odd}}(\widehat{X}_q, i\mathbb{R})$. Then, because $\widehat{\eta} \otimes W_2$ commutes with Φ_∞ we see that $\eta = \widehat{\eta} \otimes W_2 \in \Omega^1(\widehat{X}_q^\times, \pi^* L_q)$ is invariant with respect to σ, and so η descends to X. By the flatness of W_2, the form η is also $d_{A^0_\infty}$–harmonic. Hence, it defines a class in $H^1(X^\times, L_q)$. Notice that $\|\eta\|$ is bounded, and therefore,

[3] The terminology we use here is somewhat nonstandard: Prym differentials as odd classes are usually defined for unramified covers; see [28; 51]. Here we follow [21, page 86].

in L^2. Conversely, suppose η is a $d_{A^0_\infty}$–harmonic form in $\Omega^1(X^\times, L_q)$ that is in L^2. Then we can write $\pi^*\eta = \hat\eta \otimes W_2$, for $\hat\eta$ a pure imaginary form on $\hat X_q^\times$ that is anti-invariant with respect to σ. Moreover, since η is harmonic and in L^2, the form $\hat\eta$ satisfies $d\hat\eta = d^*\hat\eta = 0$ weakly, and so by elliptic regularity it is a smooth harmonic Prym differential. This leads to the following identification of harmonic Prym differentials with the space of limiting configurations.

Proposition 2.11 *The maps $\hat X_q^\times \hookrightarrow \hat X_q$, and $\hat\eta \mapsto \eta$ where $\pi^*\eta = \hat\eta \otimes W_2$, induce isomorphisms*

$$\mathcal{H}^1_{\mathrm{odd}}(\hat X_q, i\mathbb{R}) \cong H^1_{\mathrm{odd}}(\hat X_q^\times; i\mathbb{R}) \cong H^1(X^\times; L_q), \tag{2-25}$$

which send the integral lattices $\mathcal{H}^1_{\mathrm{odd}}(\hat X_q, 2\pi i\mathbb{Z})$ to $H^1_{\mathbb{Z}}(X^\times, L_q)$. Hence, combined with (2-20), this gives an identification

$$\mathrm{Prym}(\hat X_q, X) \simeq \mathscr{H}_\infty^{-1}(q), \tag{2-26}$$

which is natural with respect to scaling by $t > 0$.

Proof The first isomorphism in (2-25) was shown in [44], and the second holds by the above discussion. It remains to show that under these identifications the lattices $\mathcal{H}^1_{\mathrm{odd}}(\hat X_q, 2\pi i\mathbb{Z})$ and $H^1_{\mathbb{Z}}(X^\times, L_q)$ are preserved. Indeed, suppose $[\hat\eta] \in H^1_{\mathrm{odd}}(\hat X_q, 2\pi i\mathbb{Z})$, and choose a representative $\hat\eta$ that is odd. Choose a basepoint $w_0 \in Z \subset \hat X_q$, and for $w \in \hat X_q$, let

$$g(w) = \exp\biggl(\int_{w_0}^{w} \hat\eta \otimes W_2(w)\biggr).$$

Since $\hat\eta$ has $2\pi i\mathbb{Z}$ periods, and

$$\exp(2\pi i k W_2(w)) = I \quad \text{for } k \in \mathbb{Z}, \tag{2-27}$$

one sees that $g(w)$ is well defined and independent of the path of integration from w_0 to w. Moreover, notice that

$$\int_{w_0}^{\sigma(w)} \hat\eta = \int_{w_0}^{w} \sigma^*\hat\eta = -\int_{w_0}^{w} \hat\eta \mod 2\pi i\mathbb{Z}. \tag{2-28}$$

Therefore,

$$\begin{aligned} g(\sigma(w))g(w)^{-1} &= \exp\biggl\{\int_{w_0}^{\sigma(w)} \hat\eta \otimes W_2(\sigma(w)) - \int_{w_0}^{w} \hat\eta \otimes W_2(w)\biggr\} \\ &= \exp\biggl\{-\biggl(\int_{w_0}^{\sigma(w)} \hat\eta + \int_{w_0}^{w} \hat\eta\biggr) \otimes W_2(w)\biggr\} && \text{since } W_2(\sigma(w)) = -W_2(w) \\ &= I && \text{by (2-28) and (2-27).} \end{aligned}$$

Hence, g is a well-defined $U(1)$–gauge transformation on $X^\times$, and $\hat\eta \otimes W_2 = g^{-1}\,dg$. Conversely, as mentioned in the discussion prior to Proposition 2.3, the group $H^1_{\mathbb{Z}}(X^\times, L_q)$ of components of the stabilizer of Φ_∞ is generated by global gauge transformations of this form.

The final statement holds since if $\pi_t\colon \widehat{X}_{tq}\to X$, $\pi\colon \widehat{X}_q\to X$, $f_t\colon \widehat{X}_q\to\widehat{X}_{tq}$ is given by multiplication by $t^{1/2}$, then writing

$$\pi_t^*\eta=\widehat{\eta}_t\otimes W_2 \quad\text{and}\quad \pi^*\eta=\widehat{\eta}\otimes W_2,$$

it is easy to see that $f_t^*\widehat{\eta}_t=\widehat{\eta}$. □

2.2.4 Limiting configurations and spectral data Recall the sequence (2-15). In the case where $\mathcal{U}=\pi^*(K_X^{1/2})$, we have $\mathcal{E}=K_X^{-1/2}\oplus K_X^{1/2}$, and $\sigma^*(\mathcal{U})=\mathcal{U}$. The isomorphism between the description of $\pi^*(\mathcal{E})$ in (2-16) and the pullback of this bundle is given by

$$(u,v)\mapsto\big(\tfrac{1}{2}\lambda^{-1}(u-v),\tfrac{1}{2}(u+v)\big).$$

Note that the first factor on the right-hand side above is regular because $u-v$ vanishes at the zeroes of λ.

In fact, the correspondence in (2-26) occurs at the level of spectral data as well, in a manner we now describe. The image of the map

$$\pi^*(K_X^{1/2})\otimes\mathcal{I}_Z\to\pi^*E,\quad s\mapsto(\|\lambda\|^{1/2}\lambda^{-1}s,\|\lambda\|^{-1/2}s),\tag{2-29}$$

is the kernel of $\lambda-\pi^*\Phi_\infty$. This is a *holomorphic* embedding for a limiting configuration $d_{A_\infty}=d_{A^0_\infty}+\eta$ if and only if as a holomorphic bundle, $\mathcal{U}=\mathcal{L}\otimes\pi^*(K_X^{1/2})$, and $\mathcal{L}$ is the trivial bundle on $\widehat{X}_q$ with $\bar\partial$–operator determined by the $(0,1)$ part of $\widehat{\eta}$.

Thus, the correspondence in Proposition 2.11 is between limiting configurations and "limiting spectral data". To make sense of the latter, consider the following situation. Let $q_n\to q\in\mathcal{SQD}^*(X)$, and let $B\subset\mathcal{SQD}^*(X)$ be a neighborhood of q. Then there is a smooth holomorphic fibration $p\colon\widehat{\mathcal{X}}\to B$ of complex manifolds, where for $b\in B$, $p^{-1}(b)$ is the branched covering $\widehat{X}_b\to X$. For j large, $q_n\in B$, and the Gauss–Manin connection on $\widehat{\mathcal{X}}$ gives an identification of Prym differentials on $\widehat{X}_{q_n}$ and $\widehat{X}_q$ which preserves the integral lattice; and hence also an identification of spectral data. With this understood, we have the following.

Theorem 2.12 *Suppose $\widehat{\eta}_n$ is a sequence of imaginary harmonic Prym differentials on $\widehat{X}_{q_n}$ converging to a differential $\widehat{\eta}$ on $\widehat{X}_q$, in the sense of the paragraph above. Let $t_n\to+\infty$. Let $(\mathcal{E}_n,t_n\Phi_n)$ be the Higgs bundles associated to $\widehat{\eta}_n$ via the identification (2-20) and Theorem 2.7 (see also Remark 2.8(i)), and let (A_n,Ψ_n) be the corresponding solutions to the self-duality equations. Then any accumulation point of the sequence (A_n,Ψ_n) in the space of limiting configurations is gauge equivalent to $(A^0_\infty+\eta,\Psi_\infty)$, where $\pi^*\eta=\widehat{\eta}\otimes W_2$.*

Proof Suppose without loss of generality that (A_n,Ψ_n) converges to a limiting configuration (A_∞,Ψ_∞). Then (A_∞,Ψ_∞) is in the fiber $\mathscr{H}_\infty^{-1}(q)$, and A_∞ is gauge equivalent to a connection of the form $A^0_\infty+\eta_0$, where $\pi^*\eta_0=\widehat{\eta}_0\otimes W_2$ for some $\widehat{\eta}_0\in\operatorname{Prym}(\widehat{X}_q,X)$. We must show $[\widehat{\eta}_0]=[\widehat{\eta}]$. For this, it suffices to show that $\widehat{\eta}_0$ and $\widehat{\eta}$ have the same periods on the homology $H_1^{\mathrm{odd}}(\widehat{X}_q)$, modulo integers. For any class $[\widehat{\gamma}]\in H_1^{\mathrm{odd}}(\widehat{X}_q)$, we may choose a representative $\widehat{\gamma}\subset\widehat{X}_q^\times$. The pullback connections π^*A_n converge to π^*A_∞ in C^∞_{loc} on $\widehat{X}_q^\times$ with respect to the fibration $\widehat{\mathcal{X}}$ introduced above. On the other hand, as discussed above, the class $[\hat{\eta}_n]$ of the spectral data for $(\mathcal{E}_n, t_n\Phi_n)$ is determined by the restriction of $\pi^* A_n$ to the line subbundle in the embedding (2-29), and the same is true for $[\hat{\eta}_0]$. Hence, convergence of the connections away from the branching locus implies the periods of $\hat{\eta}$ and $\hat{\eta}_0$ agree. □

Theorem 2.12 states that the partial compactification of $\mathscr{H}^{-1}(\mathcal{SQD}^*(X))$ via spectral data mentioned in the comment following Corollary 1.1 is compatible, via Proposition 2.11, with the description of ideal points in terms of limiting configurations.

2.3 Equivariant harmonic maps

The goal of this section is to relate the Riemannian geometry of the hyperbolic space $\mathbb{H}^3$ to the gauge theory of Higgs bundles. The main result is Theorem 2.14 below. All of this material is standard and is explicitly or implicitly described in Hitchin [29] and Donaldson [14], and more generally in Corlette [10], Jost and Yau [35] and Labourie [40]. Nevertheless, in order to make the exposition here self-contained and to get the correct normalizations, we wish to reformulate the general description to suit the purposes of this paper.

2.3.1 Statement of the result With a choice of lift $\tilde{p}_0 \in \tilde{X}$ of $p_0 \in X$, the fundamental group $\pi_1 = \pi_1(X, p_0)$ acts by deck transformations on $\tilde{X}$. Given $\rho\colon \pi_1 \to \mathrm{SL}(2,\mathbb{C})$, we say that a map $u\colon \tilde{X} \to \mathbb{H}^3$ is *ρ–equivariant* if $u(\gamma z) = \rho(\gamma)u(z)$ for all $z \in \tilde{X}$ and $\gamma \in \pi_1$. If u is C^2, we say that u is *harmonic* if $d^*_{\nabla^{\mathsf{LC}}} du = 0$. Here, we let $du \in \Omega^1(\tilde{X}, u^*T\mathbb{H}^3)$ denote the differential of the map u, and ∇^{LC} the Levi-Civita connection on $\mathbb{H}^3$. The key existence result is stated here.

Theorem 2.13 [10; 14; 35; 40] *Suppose that the representation $\rho\colon \pi_1 \to \mathrm{SL}(2,\mathbb{C})$ is completely reducible. Then there exists a ρ–equivariant harmonic map $u\colon \tilde{X} \to \mathbb{H}^3$. If ρ is irreducible, then u is unique.*

The *Hopf differential* of a map $u\colon \tilde{X} \to \mathbb{H}^3$ is defined as the $(2,0)$–part of the pullback of the metric tensor of $\mathbb{H}^3$,

$$\mathrm{Hopf}(u) = (u^* ds^2_{\mathbb{H}^3})^{2,0}. \tag{2-30}$$

A very important and classical fact is that $\mathrm{Hopf}(u)$ is a holomorphic quadratic differential if u is harmonic.

As before, let $E \to X$ be a hermitian rank 2 vector bundle and recall that $\mathfrak{g}_E$ and $\sqrt{-1}\mathfrak{g}_E$ denote the bundles of traceless skew-hermitian and hermitian endomorphisms of E, respectively. A central construction used in this paper is the following.

Theorem 2.14 *Let (A, Ψ) be an irreducible solution of the self-duality equations (2-4). Choose $p_0 \in X$ and $\tilde{p}_0 \in \tilde{X}$ as above and a unitary frame of the fiber E_{p_0} of E at p_0. Let $\rho\colon \pi_1(X, p_0) \to \mathrm{SL}(2,\mathbb{C})$ be the holonomy representation of the flat connection $\nabla = d_A + \Psi$. Then the unique ρ–equivariant harmonic map from Theorem 2.13 satisfies the following properties.*

(i) *The pullback $u^*T\mathbb{H}^3$ descends to a bundle on X that is isometrically isomorphic to $\sqrt{-1}\mathfrak{g}_E$. Under this identification*:

(ii) *The orthogonal connection d_A on $\sqrt{-1}\mathfrak{g}_E$ corresponds to the pullback of the Levi-Civita connection ∇^{LC} on $\mathbb{H}^3$.*

(iii) *The hermitian 1–form $-2\Psi \in \Omega^1(X, \sqrt{-1}\mathfrak{g}_E)$ corresponds to the differential $du \in \Omega^1(X, u^*T\mathbb{H}^3)$.*

(iv) *The Higgs field Ψ and the harmonic map u determine the same quadratic differential in the sense that* $\operatorname{Hopf}(u) = 2\operatorname{tr}(\Psi \otimes \Psi)^{2,0}$.

Remark 2.15 Indeed, while our focus in this paper is on harmonic maps, Theorem 2.14 holds for general ρ–equivariant maps $u \colon \widetilde{X} \to \mathbb{H}^3$, as will become clear from the discussion below.

The remainder of Section 2.3 is devoted to the proof of Theorem 2.14.

2.3.2 The matrix model of $\mathbb{H}^3$

We view the hyperbolic space $\mathbb{H}^3$ as the homogeneous space

$$\mathsf{SL}(2,\mathbb{C})/\mathsf{SU}(2).$$

The latter may in turn be identified with

$$\mathscr{D} = \{h \in \mathrm{Mat}_{2\times 2}(\mathbb{C}) \mid h = h^*, \det h = 1,\ h > 0\},$$

where the identification maps the coset $[g] \mapsto gg^*$. Note that the left action by $\mathrm{SL}(2,\mathbb{C})$ then corresponds to $g \cdot h = ghg^*$, and that $\mathscr{D}$ has a distinguished point corresponding to $h = \mathrm{id}$.

The tangent space is given by

$$T_h\mathbb{H}^3 \simeq T_h\mathscr{D} = \{H \in \mathrm{Mat}_{2\times 2}(\mathbb{C}) \mid H = H^*,\ \operatorname{tr}(Hh^{-1}) = 0\}. \tag{2-31}$$

It will be useful to have another description of the tangent space as

$$T_h\mathbb{H}^3 \simeq \{K \in \mathrm{Mat}_{2\times 2}(\mathbb{C}) \mid (Kh)^* = Kh, \operatorname{tr}(K) = 0\}. \tag{2-32}$$

The correspondence between the two descriptions is given by $H \mapsto K = Hh^{-1}$. We shall refer to (2-31) as the *hermitian model* and to (2-32) as the *traceless model*. The traceless model gives a trivialization of the complexification $T\mathbb{H}^3 \otimes \mathbb{C} \cong \mathbb{H}^3 \times \mathbb{C}^3$, and the fiber is identified with the space of traceless 2×2 complex matrices. The real bundle $T\mathbb{H}^3$ is recovered as the fixed-point set of the complex antilinear map $\tau = \tau_h$ given by

$$\tau_h(M) = hM^*h^{-1}. \tag{2-33}$$

The invariant constant curvature -1 Riemannian metric on $\mathbb{H}^3$ is defined by

$$\langle H_1, H_2\rangle_{\mathbb{H}^3,h} = \tfrac{1}{2}\operatorname{tr}(H_1h^{-1}H_2h^{-1}) \tag{2-34}$$

for $H_i \in T_h\mathbb{H}^3$ in the hermitian model. If we define the hermitian structure on $T\mathbb{H}^3 \otimes \mathbb{C}$ by

$$\langle M_1, M_2\rangle_{\mathbb{H}^3,h} = \tfrac{1}{2}\operatorname{tr}(M_1 h M_2^* h^{-1}), \tag{2-35}$$

then the map $H \mapsto K$ between models is a real isometry for the induced metric on the fixed-point set of τ.

Lemma 2.16 *In the traceless model the Levi-Civita connection of $\mathbb{H}^3$ is given by*

$$\nabla^{\mathsf{LC}} K = dK - \tfrac{1}{2}[dhh^{-1}, K]. \tag{2-36}$$

Proof On [14, page 129] it is shown that the connection in the hermitian model is given by

$$\nabla^{\mathsf{LC}} H = dH - \tfrac{1}{2}(dhh^{-1}H + Hh^{-1}dh), \tag{2-37}$$

for $H \in T_h\mathbb{H}^3$. Pulling back this connection to the traceless model means computing $\nabla^{\mathsf{LC}}(Kh)h^{-1}$, and this gives (2-36). □

2.3.3 Flat bundles on $\mathbb{H}^3$ We define a rank 2 hermitian bundle $V \to \mathbb{H}^3$ using the homogeneous space description. More precisely, endow $\mathsf{SU}(2)$ with a right action on $\mathbb{C}^2$, given by $v \cdot h = h^{-1}v$ for $v \in \mathbb{C}^2$ and $h \in \mathsf{SU}(2)$, and then define

$$V = (\mathrm{SL}(2,\mathbb{C}) \times \mathbb{C}^2)/\mathsf{SU}(2)$$

for the diagonal action. Smooth sections of V then correspond to functions $s\colon \mathrm{SL}(2,\mathbb{C}) \to \mathbb{C}^2$ satisfying $s(gh) = h^{-1}s(g)$ for $h \in \mathsf{SU}(2)$. A hermitian structure on V is then derived from the standard inner product on $\mathbb{C}^2$. We define a connection on V as follows: $(\hat{\nabla}s)(g) = ds(g) + g^{-1}dg \cdot s(g)$. One easily verifies that this connection is well-defined and *flat*.

Now consider the bundle $\operatorname{End}_0 V \to \mathbb{H}^3$ of traceless endomorphisms of V, with its flat connection $\hat{\nabla}$ induced from the connection on V described in the previous paragraph. This is a rank 3 complex vector bundle. Recall from the previous section that the trivial bundle $T\mathbb{H}^3 \otimes \mathbb{C}$ is also a rank 3 complex hermitian bundle. We endow it with the trivial connection $\nabla^{\mathbb{C}} M := dM$.

Proposition 2.17 *There is a bundle isometry $\operatorname{End}_0 V \xrightarrow{\sim} T\mathbb{H}^3 \otimes \mathbb{C}$ which intertwines the flat connections $\hat{\nabla}$ and $\nabla^{\mathbb{C}}$.*

Proof Endomorphisms T of V are given by functions $T\colon \mathrm{SL}(2,\mathbb{C}) \to \operatorname{End}_0\mathbb{C}^2$ such that $(Ts)(g) = T(g)s(g)$ on sections s. Equivariance with respect to $\mathsf{SU}(2)$ implies $(Ts)(gh) = h^{-1}(Ts)(g)$, or

$$h^{-1}T(g)s(g) = T(gh)s(gh) = T(gh)h^{-1}s(g).$$

Since the section is arbitrary, it follows that we must have $T(gh) = h^{-1}T(g)h$. In particular,

$$M(h) = gT(g)g^{-1}, \quad \text{where } h = gg^*, \tag{2-38}$$

is a well-defined traceless 2×2–matrix valued function on $\mathbb{H}^3$, and so this defines the map above. This is an isometry, since the hermitian structures are given by

$$\langle T_1, T_2\rangle = \tfrac{1}{2}\operatorname{tr}(T_1T_2^*) = \tfrac{1}{2}\operatorname{tr}(g^{-1}M_1g(g^{-1}M_2g)^*) = \tfrac{1}{2}\operatorname{tr}(M_1hM_2^*h^{-1}) = \langle M_1, M_2\rangle_h.$$

The induced connection on the endomorphism bundle $\operatorname{End}_0 V$ is

$$\hat{\nabla}T = dT + [g^{-1}dg, T].$$

On the other hand,

$$\nabla^{\mathbb{C}}M = d(gTg^{-1}) = g\,dT g^{-1} + g[g^{-1}dg, T]g^{-1} = g(\hat{\nabla}T)g^{-1},$$

which via (2-38) proves that the connections are intertwined. □

The main result of this subsection is now the following.

Proposition 2.18 *Recall that $\nabla^{\mathbb{C}}$ is the trivial connection on $T\mathbb{H}^3\otimes\mathbb{C}$. With respect to its hermitian structure, $\nabla^{\mathbb{C}} = d_A + \Psi$, where $\Psi(h) = \frac{1}{2}[dhh^{-1}, \cdot\,]$ is a hermitian endomorphism valued one-form and d_A is unitary. The real structure τ is flat with respect to d_A, and d_A induces the Levi-Civita connection ∇^{LC} on the fixed-point set of τ, which is isomorphic to $T\mathbb{H}^3$.*

Proof We calculate the hermitian part of the connection Ψ. From (2-35),

$$d\langle M_1, M_2\rangle_h = \langle dM_1, M_2\rangle_h + \langle M_1, dM_2\rangle_h + \langle [M_1, dhh^{-1}], M_2\rangle_h,$$
$$0 = 2\langle \Psi_h M_1, M_2\rangle_h + \langle [M_1, dhh^{-1}], M_2\rangle_h,$$

which implies Ψ has the form in the statement above. Hence,

$$d_A = d - \tfrac{1}{2}[dhh^{-1}, \cdot\,]. \tag{2-39}$$

Next, from (2-33),

$$\begin{aligned}(d_A\tau)(M) &:= d_A(\tau(M)) - \tau(d_AM)\\ &= d(hM^*h^{-1}) - \tfrac{1}{2}[dhh^{-1}, hM^*h^{-1}] - h(dM - \tfrac{1}{2}[dhh^{-1}, M])^*h^{-1}\\ &= h\,dM^*h^{-1} + [dhh^{-1}, hM^*h^{-1}] - \tfrac{1}{2}[dhh^{-1}, hM^*h^{-1}] - h\,dM^*h^{-1} - \tfrac{1}{2}h[h^{-1}dh, M^*]h^{-1}\\ &= 0.\end{aligned}$$

Hence, τ is flat with respect to d_A, and so d_A induces an $\mathrm{SO}(3)$ connection on $T\mathbb{H}^3$. Comparing (2-39) with (2-36), we see that this is the Levi-Civita connection. □

2.3.4 Flat connections and equivariant maps Recall from the previous paragraph the definition of the flat bundle $V\to\mathbb{H}^3$. Sections of the dual bundle V^* are functions $s^*\colon \mathrm{SL}(2,\mathbb{C})\to(\mathbb{C}^2)^*$ satisfying the condition $s^*(gh) = s^*(g)\circ h$ for all $h\in\mathrm{SU}(2)$. Moreover, the flat connection on V induces one on V^*, which we denote with the same notation $\hat{\nabla}$. In terms of this description of sections, $\hat{\nabla}s^* = ds^* - s^*\circ g^{-1}dg$. Fix a unitary frame $\{v_1, v_2\}$ for $\mathbb{C}^2$, and let $\{v_1^*, v_2^*\}$ be the dual frame. We express the matrix elements of $g\in\mathrm{SL}(2,\mathbb{C})$ as g_{ij}.

Proposition 2.19 *The functions $s_i^*(g) = \sum_{j=1}^2 g_{ij} v_j^*$ give global parallel sections of V^*. Moreover, $s_i^*(g_1 g_2) = (g_1)_{ij} s_j^*(g_2)$.*

Proof We have $ds_i^*(v_k) = dg_{ik}$. Similarly,

$$s_i^* \circ g^{-1} dg(v_k) = s_i^* \circ (g^{-1})_{jm} dg_{mk} v_j = g_{ij} (g^{-1})_{jm} dg_{mk} = dg_{ik}.$$

The second statement is clear. □

We now present the general construction. Let $E \to X$ be a hermitian vector bundle, and let ∇ be a flat $\mathrm{SL}(2,\mathbb{C})$ connection on E. The pullback of E and ∇ to the universal cover $\widetilde{X} \to X$ will be denoted with the same notation. Choose a basepoint $p_0 \in X$, and a lift $\widetilde{p}_0 \in \widetilde{X}$. Fix a unitary frame $\{e_1, e_2\}$ of the fiber E_{p_0} (and therefore also $E_{\widetilde{p}_0}$). We have a uniquely determined global frame $\{\widetilde{e}_1, \widetilde{e}_2\}$ for $E \to \widetilde{X}$ that is parallel with respect to ∇ and which agrees with $\{e_1, e_2\}$ at $\widetilde{p}_0$. Let u_{ij} be the hermitian matrix $u_{ij}(p) = \langle \widetilde{e}_i, \widetilde{e}_j \rangle (p)$. Then u_{ij} is hermitian and positive. We claim that $\det u = 1$. Indeed, write $\nabla = d_A + \Psi$, where d_A is a unitary connection on E and Ψ a 1–form with values in $\sqrt{-1}\mathfrak{g}_E$. Let $\{\hat{e}_1, \hat{e}_2\}$ be a unitary frame at p, $\widetilde{e}_i(p) = g_{ij} \hat{e}_j$, $\Psi(p)\hat{e}_i = \Psi_{ij} \hat{e}_j$. Then at the point p,

$$du_{ij} = \langle d_A \widetilde{e}_i, \widetilde{e}_j \rangle + \langle \widetilde{e}_i, d_A \widetilde{e}_j \rangle = -2 \langle \Psi(p) \widetilde{e}_i, \widetilde{e}_j \rangle = -2(g \Psi g^*)_{ij}. \tag{2-40}$$

At the same time, $u_{ij}(p) = \langle \widetilde{e}_i, \widetilde{e}_j \rangle (p) = (gg^*)_{ij}$. Hence,

$$d \log \det u = \mathrm{tr}(u^{-1} du) = -2 \,\mathrm{tr}\big((gg^*)^{-1} (g \Psi g^*)\big) = -2 \,\mathrm{tr}\, \Psi = 0,$$

since Ψ_{ij} is traceless. Therefore, $\det u(p) = \det u(\widetilde{p}_0) = 1$ for all $p \in \widetilde{X}$. Hence, $u(p) \in \mathscr{D}$, and we have therefore defined a map $u \colon \widetilde{X} \to \mathbb{H}^3$ which sends the point $\widetilde{p}_0$ to the basepoint of $\mathscr{D}$. We also note for future reference that from (2-40),

$$du\, u^{-1} = -2 g \Psi g^{-1}. \tag{2-41}$$

Let $\rho \colon \pi_1 \to \mathrm{SL}(2,\mathbb{C})$ be the holonomy representation of ∇ with respect to the frame $\{e_1, e_2\}$. Via the choice of basepoint $\widetilde{p}_0$ we may view π_1 as acting on $\widetilde{X}$ by deck transformations. By definition, if $\rho(\gamma) = (g_{ij})$, then $\widetilde{e}_i(\gamma p) = g_{ij} \widetilde{e}_j(p)$ for any $p \in \mathbb{H}^2$. Therefore,

$$u_{ij}(\gamma p) = \langle g_{ik} \widetilde{e}_k(p), g_{jm} \widetilde{e}_m(p) \rangle = g_{ik} u_{km}(p) g_{mj}^*,$$

or $u(\gamma p) = \rho(\gamma) u(p) (\rho(\gamma))^*$. Thus, u is equivariant with respect to the action of the holonomy representation ρ on $\mathbb{H}^3$.

Proposition 2.20 *There is a π_1–equivariant isometry $u^* V^* \xrightarrow{\sim} E$ that intertwines the flat connections $u^* \widehat{\nabla}$ and ∇.*

Proof Recall the sections of V^* from Proposition 2.19. Then the bundle isomorphism is defined by identifying $[g, s_i^*(g)] \mapsto \widetilde{e}_i(p)$, where $u(p) = gg^*$. By the second statement of Proposition 2.19, this identification is equivariant with respect to the action of π_1. Since the identification is between flat sections, the connections are manifestly intertwined. It remains to check that this is an isometry. But

$$\langle s_i^*, s_j^* \rangle (u(p)) = g_{ik} \overline{g_{jm}} \langle v_k^*, v_m^* \rangle = g_{ik} \overline{g_{jk}} = u_{ij}(p) = \langle \tilde{e}_i, \tilde{e}_j \rangle (p).$$

This completes the proof. □

The next proposition is the main consequence of the discussion above.

Proposition 2.21 *Let $E \to X$ be a hermitian rank 2 vector bundle with a flat* $\mathrm{SL}(2, \mathbb{C})$ *connection* ∇ *and holonomy representation $\rho \colon \pi_1 \to \mathrm{SL}(2, \mathbb{C})$. Write $\nabla = d_A + \Psi$, where d_A is a unitary connection on E and Ψ is a 1–form with values in $\sqrt{-1}\mathfrak{g}_E$. Let $u \colon \widetilde{X} \to \mathbb{H}^3$ be the ρ–equivariant map described above. Then $\sqrt{-1}\mathfrak{g}_E$ may be isometrically identified with $u^* T\mathbb{H}^3$. Under this identification the induced connection d_A corresponds to the pullback of the Levi-Civita connection on $\mathbb{H}^3$, and the 1–form -2Ψ corresponds to the differential du of the map u.*

Proof By Proposition 2.20, the connection on E pulls back from the one on V^*. The induced connection on $\mathrm{End}_0 E$ is therefore the pullback of the one on $\mathrm{End}_0 V$. Since these bundles are isometric, the subbundle $\sqrt{-1}\mathfrak{g}_E$ identifies with the bundle of traceless hermitian endomorphisms of V. By Propositions 2.17 and 2.18, the latter is isometric to $T\mathbb{H}^3$, and the induced connection is the Levi-Civita connection. The computation in (2-41) shows $du\, u^{-1} = -2g\Psi g^{-1}$. Combined with the identification (2-38) this yields the claimed relation between Ψ and the differential of the map u in the traceless model. □

Proof of Theorem 2.14 The proof of (i)–(iii) follow from Proposition 2.21 . For (iv), use (2-41) and definition of the metric (2-34) to compute

$$\mathrm{Hopf}(u) = \tfrac{1}{2} \operatorname{tr}(du u^{-1} \otimes du u^{-1})^{2,0} = 2 \operatorname{tr}(\Psi \otimes \Psi)^{2,0}.$$

This completes the proof of the theorem. □

Remark 2.22 The construction above is natural with respect to the action of unitary gauge transformations on pairs (A, Ψ). Namely, modifying (A, Ψ) to $g^*(A, \Psi)$, where $g \in \mathcal{G}$ is a unitary gauge transformation results in conjugating the representation ρ and the map u by some element in $\mathrm{SU}(2)$.

2.3.5 The self-duality equations and harmonic maps Up to this point, the choice of hermitian metric on the bundle E was arbitrary and not related to the holonomy representation ρ determined by the flat $\mathrm{SL}(2, \mathbb{C})$ connection ∇. For this reason, the pair (A, Ψ) resulting from the decomposition of ∇ into its unitary and hermitian part as in Proposition 2.21 will in general not satisfy any equation apart from the flatness of ∇, which is equivalent to the first two equations of (2-4). Likewise, the construction of the ρ–equivariant map u depends on the hermitian metric on E and hence this map will in general not enjoy any special properties. The link to the extra structure is provided by the following.

Proposition 2.23 [14] *Let $E \to X$ be a rank 2 vector bundle with hermitian metric h and a flat $\mathrm{SL}(2,\mathbb{C})$ connection ∇ and corresponding holonomy representation $\rho\colon \pi_1 \to \mathrm{SL}(2,\mathbb{C})$. Denote by $\nabla = d_A + \Psi$ the unique decomposition of ∇ into a unitary connection d_A on (E,h) and a one-form Ψ with values in $\sqrt{-1}\mathfrak{g}_E$. Let moreover $u\colon \widetilde{X} \to \mathbb{H}^3$ be a ρ–equivariant smooth map as in Proposition 2.21. Then the pair (A,Ψ) satisfies the self-duality equations (2-4) if and only if the map u is harmonic.*

Remark 2.24 A hermitian metric h on the bundle E such that the corresponding ρ–equivariant map u is harmonic is called a *harmonic metric*. If ρ is irreducible, the solution (A,Ψ) of the self-duality equation resulting from Theorem 2.13 is also irreducible. In this paper, we consider monodromies associated to pleated surfaces, and the representations are therefore automatically irreducible; cf [5, page 36].

2.4 Laminations

In this section, we briefly review some of the topological objects that will be used in our description of the images of high-energy harmonic maps.

2.4.1 Measured foliations and laminations A *measured foliation* on a surface Σ is a partial foliation $\mathcal{F}$ of the surface with a finite number of k–pronged singularities, equipped with a measure on transverse arcs. The examples we consider in this paper are the *horizontal and vertical foliations* of a holomorphic quadratic differential q with simple zeroes, $q \in \mathcal{QD}^*(X)$, which we denote by $\mathcal{F}_q^h$ and $\mathcal{F}_q^v$, respectively. In the notation of Section 2.2, these can be defined as follows. At each point of the spectral curve $\widehat{X}_q^\times$, consider a (real) unit tangent vector $\hat{u}$ with $\mathrm{Im}(\lambda_{\mathrm{SW}}(\hat{u})) = 0$. Then the flow lines of $\hat{u}$ integrate locally to give a foliation independent of the choice of sign of $\hat{u}$ and invariant under the involution σ. It therefore projects to a foliation on $X^\times$, and this is $\mathcal{F}_q^h$, the horizontal foliation of q. The vertical foliation of q, $\mathcal{F}_q^v$, is transverse to $\mathcal{F}_q^h$, and is defined similarly using the real part. We let $\widetilde{\mathcal{F}}_q^h$ and $\widetilde{\mathcal{F}}_q^v$ denote the lifts of the foliations to the universal cover $\widetilde{X}$.

A *critical leaf* of $\mathcal{F}_q^h$ is a segment of a horizontal leaf terminating at a zero of q. A *saddle connection* of the horizontal (resp. vertical) foliation is a horizontal (resp. vertical) leaf joining two zeroes. Following [41, Section 3], when we refer to a path in $\widetilde{\mathcal{F}}_q^h$ as a *horizontal leaf*, we implicitly mean that it either contains no zeroes of $\tilde{q}$, or when it meets critical points it either turns consistently to the right or to the left with respect to the cyclic ordering on the critical leaves terminating at a give zero. Saddle connections will play an important technical role in this paper, but there is a distinction between vertical and horizontal saddles, as discussed in the introduction.

The foliations $\mathcal{F}_q^h$ and $\mathcal{F}_q^v$ come equipped with transverse measures. If k is a C^1 arc transverse to $\mathcal{F}_q^h$, we can lift k to a parametrized arc $\hat{k}$ in $\widehat{X}_q$ in such a way that[4] $\mathrm{Im}(\lambda_{\mathrm{SW}}(\dot{\hat{k}})) < 0$ at all points of $\hat{k}$. The measure of k is then the integral of $-\mathrm{Im}\,\lambda_{\mathrm{SW}}$ along $\hat{k}$. We will say that a piecewise C^1 arc k is

[4]The negativity is dictated in order to agree with Bonahon's convention; see [5, 2] and Section 4.1.3 below.

quasitransverse to $\mathcal{F}_q^h$ if it is a finite union of C^1 arcs in $X \setminus Z(q)$, and if it admits a piecewise C^1 lift $\hat{k}$ in $\hat{X}_q$ in such a way that $\mathrm{Im}(\lambda_{\mathrm{SW}}(\dot{\hat{k}})) \leqslant 0$ at all points of $\hat{k}$. The definition of a path quasitransverse to $\mathcal{F}_q^v$ is defined similarly using the real part.

A *measured geodesic lamination* Λ on a hyperbolic surface S is a partial foliation of the surface by simple (not necessarily closed) geodesics, together with a measure on transverse arcs. A measured foliation may be "straightened" to a measured lamination. For example, given $\mathcal{F}$, each bi-infinite leaf of $\tilde{\mathcal{F}} \subset \tilde{S} \simeq \mathbb{H}^2$ defines a unique pair of distinct points in the circle at infinity, and hence a unique geodesic in $\mathbb{H}^2$. The collection of geodesics thus obtained are noninterlacing and form a closed set, and so define a lamination $\tilde{\Lambda}$ of $\mathbb{H}^2$. The construction is equivariant with respect to the action of the fundamental group, and so there is a well-defined quotient $\Lambda \subset S$. The transverse measure on $\mathcal{F}_q^h$ may then be transported to a measure on arcs transverse to Λ_q^h. For more details on this construction, see [41]. We will denote the measured laminations associated to $\mathcal{F}_q^h$ and $\mathcal{F}_q^v$ by Λ_q^h and Λ_q^v, respectively.

The *Hubbard–Masur theorem* [32] gives a converse to this construction. Given a measured foliation $\mathcal{F}$ (resp. measured lamination Λ) there is a unique nonzero $q \in \mathcal{QD}(X)$ such that $\mathcal{F}$ is measure equivalent to $\mathcal{F}_q^h$ (resp. Λ to Λ_q^h). We shall denote this differential by $\phi_{\mathsf{HF}}(\mathcal{F})$ (resp. $\phi_{\mathsf{HF}}(\Lambda)$). (See [61] for a proof closer to the perspective in this paper.)

For a lamination $\Lambda \subset S$, the components of $\mathbb{H}^2 \setminus \tilde{\Lambda}$ are called *plaques*, and we denote the set of such by $\mathcal{P}(\Lambda)$. When all the plaques are ideal triangles, we say that Λ is *maximal*. If $\mathcal{F}_q^h$ has saddle connections, then Λ_q^h will not be a maximal lamination, and we describe this in more detail in Section 2.4.3. For a distinct pair $P, Q \in \mathcal{P}(\Lambda)$, we say that $R \in \mathcal{P}(\Lambda)$ *separates* P and Q if any path from P to Q in $\mathbb{H}^2$ intersects R.

We end this section with two clarifying remarks. First, while a simple example of a measured lamination is a multicurve equipped with atomic transverse measures, a more typical example (obtained as a limit of multicurve examples) will meet any transverse arc in a Cantor set. Second, while geodesic laminations appear to depend on the hyperbolic structure of the surface, using the idea of straightening curves, a geodesic lamination Λ in any marked hyperbolic structure S on Σ induces a unique geodesic lamination in any other marked hyperbolic surface S' on Σ. See [5, page 7]. We will often denote these Λ without reference to the hyperbolic structure.

2.4.2 Train tracks An ingenious construction of Thurston provides for a way to organize nearby measured foliations/laminations as data on a geometric object. A *train track* on a surface Σ is an embedded finite complex τ of C^1–arcs (called *branches*) on Σ meeting at vertices (called *switches*) with a well-defined common tangency. We can and will assume the switches are always trivalent. Then one branch at a switch is *incoming* and two are *outgoing*; see [54, page 11]. Let $G = \mathbb{R}$ or $S^1 \simeq \mathbb{R}/2\pi\mathbb{Z}$. A *weight* on a train track τ is an assignment of an element of G to each branch that obeys the *switch conditions*: the weight on the incoming branch equals the sum of the weights on the outgoing branches. We denote by $\mathcal{H}(\tau, G)$ the set of G–weights on τ.

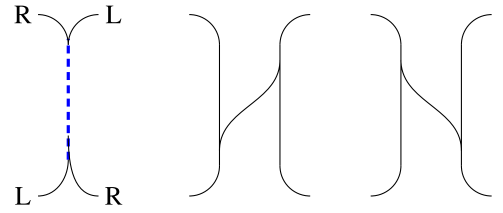

Figure 1: Splitting of train tracks.

One way to construct a train track is to consider a small ϵ neighborhood of a measured geodesic lamination, foliate that neighborhood by leaves transverse to the lamination, and then collapse the neighborhood to the leaf space of the foliation. If the resulting branches are weighted by the measure of arcs that cross the neighborhood, a measured train track that *carries* the lamination results; see [54, page 73].

A useful operation on train tracks is the *right and left splitting*; see [54, page 119]. As one chooses an increasingly small parameter ϵ in the construction above the train tracks obtained are related by splitting. Let us define splitting carefully. Recall that a branch between two switches is called *long* if it is incoming at both ends; see [54, page 118]. The orientation of Σ then orders the outgoing branches at the switches on each end of a long branch, and we label them left (L) and right (R) accordingly. A right splitting is then obtained by modifying the train track locally by replacing the long branch with two branches joining left and right at each switch, and then adding a third branch between them at whose switches the branches labeled L are incoming. The left splitting adds a branch so that the right branches are incoming. See Figure 1.

2.4.3 Maximalizations As mentioned above, in the case of horizontal saddle connections the lamination Λ_q^h is not maximal. A maximal lamination can be obtained by adding finitely many leaves to Λ_q^h [8, page 76]. Here we describe this mechanism precisely in terms of the foliation $\mathcal{F}_q^h$. Consider a connected configuration $\mathcal{S} \subset \mathcal{F}_q^h$ of saddle connections (along with their external critical leaves). We can make a train track $\tau_{\mathcal{S}}$ out of $\mathcal{S}$ by replacing each zero with a triangle, each of whose sides is outgoing.

Definition 2.25 A *maximalization of* $\mathcal{S}$ is a choice of left or right splitting $\hat{\tau}_{\mathcal{S}}$ of each branch in $\tau_{\mathcal{S}}$ corresponding to a saddle connection, in such a way that the resulting train track $\hat{\tau}_{\mathcal{S}}$ contains no long branches; see Figure 2. A *maximalization of* $\mathcal{F}_q^h$ is a choice of maximalization of every maximal connected configuration of saddle connections.

Note that maximalizations always exist: for example, one may choose right splittings for all the saddle connections. The terminology is justified by the following.

Lemma 2.26 *A maximalization of $\mathcal{F}_q^h$ uniquely determines a maximal lamination Λ containing Λ_q^h as a sublamination.*

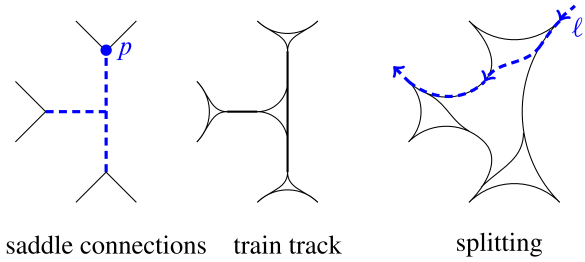

Figure 2: Maximalization.

Proof Let $\mathcal{S}$ be a maximal connected component of saddle connections, and let c be a saddle connection. There are two cases:

(1) c is part of a closed loop γ of saddle connections,

(2) there is a zero p of q at one end of c, one of whose critical leaves is in the complement of $\mathcal{S}$ (call this an *external* zero).

In case (1), Λ contains a closed geodesic $\overline{\gamma}$ homotopic to γ. The splitting now determines a train path from the critical leaf of one end point of c that is not part of γ to $\overline{\gamma}$. This corresponds to a leaf of Λ that spirals into $\overline{\gamma}$. In case (2), the splitting of the saddle connection c selects one of the other critical leaves of p; namely, the one which is incoming with respect to the switches created in the splitting. Denote this leaf by $\ell \subset \mathcal{F}_q^h$. By maximality of the component $\mathcal{S}$, the lift $\widetilde{\ell}$ of ℓ determines a geodesic half ray g in $\mathbb{H}^2$ that is asymptotic to a leaf of $\widetilde{\Lambda}_q^h$ on one end. Viewing ℓ as a train path in $\widehat{\tau}_{\mathcal{S}}$, there is a unique continuation to a path (still denoted by ℓ) that crosses the split saddle connection which ends at p, and then either exits through a branch of another external zero, or spirals around a closed branch homotopic to a closed loop of saddle connections. Uniqueness follows because all further switches the path encounters are outgoing by assumption. Thus the lift of ℓ determines a bi-infinite geodesic that is asymptotic to different leaves of $\widetilde{\Lambda}_q^h$ on either side. By the condition that there are no long branches in $\widehat{\tau}_{\mathcal{S}}$, the geodesics added in this way are disjoint, and since the interior complementary regions of $\widehat{\tau}_{\mathcal{S}}$ are triangles, the resulting lamination is maximal. □

Remark 2.27 Given a maximalization of $\mathcal{F}_q^h$ as in Definition 2.25 with lamination Λ as in Lemma 2.26:

(i) The leaves $\Lambda \setminus \Lambda_q^h$ may be represented by paths that are quasitransverse to $\mathcal{F}_q^v$, consisting of horizontal leaves coming into and exiting a neighborhood of $\mathcal{S}$, with a small vertical arc cutting one saddle connection of $\mathcal{F}_q^h$ (we shall refer to these as *additional leaves* of $\mathcal{F}_q^h$ or Λ).

(ii) There is a one-to-one correspondence between zeroes $\widetilde{Z}(q) \subset \mathbb{H}^2$ and the plaques of $\mathbb{H}^2 \setminus \widetilde{\Lambda}$.

The lift $\widehat{\Lambda}$ to $\widehat{X}_q$ can be oriented. For convenience, we always choose this so that the oriented leaves of Λ_q^h have $\operatorname{Re} \Lambda_{\mathrm{SW}} > 0$. This then gives an orientation to the homology classes in $H_1^{\mathrm{odd}}(\widehat{X}_q, \mathbb{Z})$ corresponding to the saddle connections. Suppose there is a saddle connection c from p to q. Then we can change c to an arc consisting of a vertical leaf emanating from p, followed by a horizontal leaf shadowing the saddle connection, and then another vertical leaf terminating at q. Indeed, there are two such ways of constructing such a path. However, the orientation of $\widetilde{\Lambda}$ chooses one of these: namely, the one whose lift intersects $\widetilde{\Lambda}$ positively. Notice that these paths are *not* quasitransverse with respect to $\mathcal{F}_q^v$. We shall call such arcs *modified saddle connections*.

2.5 High-energy harmonic maps

This paper focuses on asymptotics of the $\mathrm{PSL}(2,\mathbb{C})$ character variety for π_1, especially as reflected in the associated classes of Higgs bundles. The previous sections related these bundles to equivariant harmonic maps $u\colon \widetilde{X} \to \mathbb{H}^3$, and it will turn out that one leaves all compacta in the character variety (and the associated moduli spaces of Higgs bundles) exactly when the energy of the associated harmonic maps grows without bound. In this section, we collect some of the basic analytic estimates on the geometry of harmonic maps whose energies are tending to infinity. These will be used throughout the paper.

2.5.1 Minsky's results The following result due to Minsky plays a crucial role in the subsequent qualitative estimates involving high-energy harmonic maps. It will later also be needed in Section 5.1.

Let $u_n\colon \widetilde{X} \to \mathbb{H}^3$ be a sequence of ρ_n–equivariant harmonic maps with Hopf differentials $t_n^2 q_n$, with $q_n \to q$, in $\mathcal{SQD}^*(X)$. Recall that $Z(q_n)$ is the set of zeroes of q_n, which we assume to be simple. For a parameter s_n, let $\Omega_{s_n}(p)$ be a hexagonal domain for each $p \in Z(q_n)$. The s_n will be chosen so that these domains are disjoint for distinct zeroes of q_n. Set

$$\mathcal{Q}_n = \bigcup_{p \in Z(q_n)} \Omega_{s_n}(p). \tag{2-42}$$

We also assume that the boundary of each hexagon $\Omega_{s_n}(p)$ is formed from alternating horizontal and vertical edges. We let $\widetilde{Z}(q_n)$ (resp. $\widetilde{\mathcal{Q}}_n$) denote the preimage of the set $Z(q_n)$ (resp. $\mathcal{Q}_n$) under the projection map $\pi : \widetilde{X} \to X$.

Proposition 2.28 (cf [46, Theorem 4.2]) *There are constants A, c_0 and C_0, all independent of n, and an integer N such that the following hold. For $n \geqslant N$ and $s_n \leqslant c_0$, there is a ρ_n–equivariant map Π^*, from the leaves of $\widetilde{\mathcal{F}}_{q_n}^h$ in the complement of $\widetilde{\mathcal{Q}}_n$ to a collection $\widetilde{\Lambda}_n^{h,*}$ of geodesics in $\mathbb{H}^3$, which factors through u_n. Moreover, for any $p \in \widetilde{X} \setminus \widetilde{\mathcal{Q}}_n$,*

$$d_{\mathbb{H}^3}(u_n(p), \Pi^*(p)) \leqslant A \exp(-t_n C_0),$$

and the derivative along the horizontal leaf through p (in the $|q_n|$ metric) is

$$\big| |d\Pi^*| - 2 \big| \leqslant A \exp(-t_n C_0).$$

Proposition 2.28 is proven in [46] in the context of harmonic maps from surfaces to complete hyperbolic 3–manifolds, but the arguments apply equally well in the equivariant case. One important simplification in the situation here is that the domain Riemann surface X is fixed. As a consequence, the technical issues of "thin flat cylinders" that are dealt with in [46] do not play a role here. In particular, the set $\mathscr{P}_R$ in that reference may be taken to be equal to $\mathcal{Q}_n$ defined in (2-42).

The proposition is a consequence of the following construction. For s_n chosen sufficiently small and n sufficiently large, there is a train track $\tau_n \subset X \setminus \mathcal{Q}_n$ and $\varepsilon_n > 0$, $\varepsilon_n \to 0$ as $n \to +\infty$, satisfying the following.

(i) Let $\widetilde{\tau}_n \subset \widetilde{X}$ be the preimage of τ_n, and set $\widetilde{\tau}_n^* = u_n(\tau_n)$. Then the branches of $\widetilde{\tau}_n^*$ have length $O(t_n)$ and geodesic curvature $O(\varepsilon_n)$.

(ii) The images by u_n of the leaves of the horizontal foliation $\widetilde{\mathcal{F}}^h_{q_n}$ in the complement $\widetilde{X} \setminus \widetilde{\mathcal{Q}}_n$ can be straightened to give a lamination $\widetilde{\Lambda}_n^{h,*} \subset \mathbb{H}^3$.

(iii) The lamination $\widetilde{\Lambda}_n^{h,*}$ is $C^1_{\varepsilon_n}$–carried by $\widetilde{\tau}_n^*$; cf [54, page 73].

In the case where $\mathcal{F}^h_{q_n}$ has saddle connections and we have chosen a maximalization in the sense of Definition 2.25, we can enlarge the quotient $\Lambda_n^{h,*}$ of $\widetilde{\Lambda}_n^{h,*}$ to a lamination Λ_n^* as follows. By Remark 2.27, the maximalization gives rise to finitely many quasitransverse paths in X, which we may assume to lie in the complement of $\mathcal{Q}_n$. (A technical point is that Minsky creates his track by extending components of $\mathcal{Q}_n$ to "slice" through long rectangles of vertical trajectories; it is straightforward to check that this slicing can be done in a way corresponding to the maximalization discussed here.) The image by u_n of the lifts of these can be straightened to geodesics that are asymptotic on one side to leaves in $\Lambda_n^{h,*}$. The map Π^* in Proposition 2.28 can be extended to a map on these leaves satisfying the same estimates.

We next choose coordinates, which we refer to as (*canonical*) *q_n–coordinates*, that are adapted to q_n and hence to the map u_n. To this end, note that, away from the zeroes of q_n, we may choose coordinates $z_n = x_n + i y_n$ in a patch so that, in those coordinates, the quadratic differential q_n is expressed as $q_n = dz_n^2$. These are useful because the horizontal lines in this coordinates are both the leaves of the horizontal foliation of q_n, and also integrate the directions of the maximal stretch (eigendirection) of the tangent map du_n. Naturally, both the domain and the pullback metric diagonalizes with respect to these coordinates. Following Minsky [46, equation (3.1)], the pullback metric $u_n^* \, ds^2_{\mathbb{H}^3}$ with respect to the harmonic map u_n as above can be written in terms of q_n–coordinates (x_n, y_n) as

$$u_n^* \, ds^2_{\mathbb{H}^3} = 2t_n^2(\cosh \mathcal{G}_n + 1)\, dx_n^2 + 2t_n^2(\cosh \mathcal{G}_n - 1)\, dy_n^2, \tag{2-43}$$

where $\mathcal{G}_n = \sinh^{-1}(2\mathcal{J}_n)$ and $\mathcal{J}_n$ is the Jacobian determinant of the map u_n. The factor t_n^2 enters since the harmonic map u_n has Hopf differential $t_n^2 q_n$.

Proposition 2.29 *The pullback metric by u_n in terms of canonical coordinates for q_n satisfies*

$$u_s^* \, ds^2_{\mathbb{H}^3} = 4t_n^2 \, dx^2 + O(\exp(-2ct_n)) \tag{2-44}$$

in C^k for some constant $c > 0$.

Proof As shown in [46, Lemma 3.4], there is a constant B such the quantity $\mathcal{G}_n$ satisfies the pointwise estimate

$$\mathcal{G}_n(p) < \frac{B}{\cosh d}$$

for every point p at $t_n^2 q_n$–distance at least $d > 0$ to the zero set of q_n. Since we are here considering points outside some fixed neighborhood of the zero set of q_n, this distance is bounded below by ct_n for some constant $c > 0$. It follows that

$$\mathcal{G}_n(p) < 2Be^{-ct_n},$$

and consequently

$$\cosh \mathcal{G}_n(p) < 1 + 4B^2 e^{-2ct_n}.$$

Inserting this last estimate into (2-43) implies the claim. □

This last estimate implies the following properties away from the zeroes of the Hopf differential: the images of the horizontal trajectories under a high-energy map u_n are stretched by the factor t_n, up to a small and rapidly decaying error; the images by u_n of those trajectories have exponentially small geodesic curvature; and the images by u_n of the vertical trajectories have lengths exponentially decaying in t_n.

2.5.2 High-energy harmonic maps near the zeroes of q We continue with the notation of the previous section. Let $\mathbb{D}_{u(p)}$ denote the oriented totally geodesic plane in $\mathbb{H}^3$ tangent to the image of $du(p)$.

Proposition 2.30 *For every fixed $\varepsilon > 0$ there exists a constant N such that the following holds. There is an ideal hyperbolic triangle $\Delta \subset \mathbb{H}^3$ such that for every $n \geqslant N$ the distance between the tangent plane $\mathbb{D}_{u_n(p)} \subset \mathbb{H}^3$ to Δ is less than ε, for every point $p \in \widetilde{Z}(q_n)$.*

An analogous statement for two-dimensional targets is the main theorem of [59]. The present version is a reflection for harmonic maps of aspects of the approximate solutions constructions in Section 2.1.4.

Proof For each fixed n and $p \in Z(q_n)$, consider a lift $\widetilde{\Omega}_{s_n}(p)$ of the hexagon $\Omega_{s_n}(p) \subset X$ to $\widetilde{X}$. Let h_1, h_2 and h_3 denote the three horizontal edges of $\Omega_{s_n}(p)$, which we parametrize in an orientation-preserving way by a parameter $0 \leqslant s \leqslant 1$. Proposition 2.28 shows the existence of geodesics $c_i : [0,1] \to \mathbb{H}^3$ such that the distance between $u_n(h_i(s))$ and $c_i(s)$ is less than $A\exp(-t_n C_0)$ for all $0 \leqslant s \leqslant 1$. By Proposition 2.29 the length of each c_i is of order t_n. Furthermore, the distance between each consecutive pair of endpoints $u_n(h_i(1))$ and $u_n(h_{i+1}(0))$ satisfies an exponentially small bound. It follows from elementary hyperbolic geometry that there is an ideal hyperbolic triangle $\Delta \subset \mathbb{H}^3$ which is at distance at most ε to the lines $u_n(h_i)$. Since $u_n(\widetilde{\Omega}_{s_n}(p))$ is contained in the convex hull of these lines, it follows that Δ and $u_n(\widetilde{\Omega}_{s_n}(p))$ have at most distance ε, for all sufficiently large n. To see that also the tangent plane $\mathbb{D}_{u_n(p)}$ lies ε–close to Δ, we compare the harmonic map u_n with the harmonic map v_n which maps $\Omega_{s_n}(p)$ to $\mathbb{H}^3$ and has boundary values the edges of Δ. Its image is contained in Δ and, since the boundary values of u_n and v_n differ by at most ε, it follows by standard estimates on harmonic maps that both are C^1–close in the interior of $\Omega_{s_n}(p)$. This implies the assertion. □

2.6 Harmonic maps to $\mathbb{R}$–trees

2.6.1 Definitions An $\mathbb{R}$–*tree* is a complete length space T such that any two points can be joined by a unique path parametrized by arc length. This path is called the geodesic between the points, say p, q, and it is denoted by $\overline{pq}$. We shall be interested in trees admitting isometric actions of π_1, and we will always assume the action is minimal in the sense that there is no proper π_1–invariant subset of T. In such a situation, we obtain a *length function*

$$\ell_T : \pi_1 \to \mathbb{R}_{\geqslant 0}, \quad [\gamma] \mapsto \inf_{p \in T} d_T(p, \gamma p).$$

Scaling the metric (and hence ℓ_T) by positive constants defines a *projective class* of length functions; see [9].

Examples of $\mathbb{R}$–trees come from the following construction. Let $\mathcal{F}$ be a measured foliation on Σ with transverse measure μ. Define the *dual tree* $T_{\mathcal{F}}$ to the foliation as follows: if $\widetilde{\mathcal{F}}$ is the lift to the universal cover, define a pseudodistance $\tilde{d}$ by

$$\tilde{d}(p, q) = \inf\{\tilde{\nu}(c) \mid c \text{ is a rectifiable path between } p \text{ and } q\}.$$

Then the Hausdorffication of $(\widetilde{\Sigma}, \tilde{d})$ is an $\mathbb{R}$–tree with an isometric action of π_1; see [7, Corollary 2.6], and also [50; 52]. In the case of a nonzero holomorphic quadratic differential q on a Riemann surface X, we set $T_q := T_{\mathcal{F}_q^v}$.

A *morphism* of $\mathbb{R}$–trees is a continuous map $f : T \to T'$ such that given any segment $e \subset T$, either f is constant on e or e decomposes into a finite union of subsegments $e_1 \cup \cdots \cup e_k$ such that f restricted to each e_i is an isometry onto its image. It is a fact that in the latter case f is either an isometry on e or a *folding*, meaning that it identifies two or more subsegments.

Trees are examples of nonpositively curved metric spaces (NPC). Following ideas of Gromov [26], Korevaar and Schoen [38; 39] and, independently, Jost [34], developed a theory of energy-minimizing maps from Riemannian domains to NPC spaces. The fourth author [60; 61] studied the case of maps to $\mathbb{R}$–trees, which is the one relevant to this paper. We will need only very little from these results, and we package a summary statement as follows; see [13] for more details.

Theorem 2.31 *Let q be a nonzero holomorphic quadratic differential on a Riemann surface X. Then the leaf space projection map $u : \widetilde{X} \to T_q$ is an equivariant harmonic map. In general, let T be an $\mathbb{R}$–tree with an isometric action of π_1, and let $v : \widetilde{X} \to T$ be an equivariant harmonic map. Then:*

(i) *The map v is uniformly Lipschitz with constant proportional to $E(u)^{1/2}$ (the constant depends on the choice of conformal metric on X).*

(ii) *The Hopf differential* $\mathrm{Hopf}(v) = 4q$ *is well-defined, and is a holomorphic quadratic differential that is nonzero unless v is constant and the action is trivial.*

(iii) *We have $v = p \circ u$, where $u : \widetilde{X} \to T_q$ is projection as above, and $p : T_q \to T$ is a folding.*

We shall also need a version of the Korevaar–Schoen strong compactness theorem, stated here in the limited context that we require. For positive constants $t_n \to +\infty$, let $\mathbb{H}_n$ denote the hyperbolic space $\mathbb{H}^3$, but where the metric has been rescaled: $ds_{\mathbb{H}_n} = t_n^{-1}\, ds_{\mathbb{H}^3}$. For the following result, see also [11, Theorems 2.2 and 3.1].

Theorem 2.32 [39, Proposition 3.7 and Theorem 3.9] *Suppose that $u_n\colon \widetilde{X} \to \mathbb{H}_n$ is a sequence of ρ_n–equivariant continuous finite-energy maps, and assume that u_n have a uniform modulus of continuity: for each z there is a monotone function $\omega(z,R)$ such that $\lim_{R\downarrow 0}\omega(z,R) = 0$ and*

$$\max_{w\in B_R(z)} d(u_n(z), u_n(w)) \leqslant \omega(z,R).$$

Then there is an $\mathbb{R}$–tree T with an isometric action of π such that the convex hulls of the images of the u_n converge in the Gromov–Hausdorff sense to T. Moreover:

(i) *The u_n converge to a continuous finite energy map $u\colon \widetilde{X} \to T$ that is equivariant for this action.*

(ii) *If $\lim_{k\to\infty} E(u_k) \neq 0$, then u is nonconstant.*

(iii) *If the u_n are equivariant harmonic maps, then so is u; and in this case, if q_n (resp. q) is the Hopf differential of u_n (resp. u), then $t_n^{-2} q_n \to q$.*

We refer to the limiting tree T as a *Korevaar–Schoen limit*. By Theorem 2.31, T is a folding of T_q.

2.6.2 The Morgan–Shalen compactification There is a compactification of $R(\Sigma)$ that restricts on the Fricke space to Thurston's compactification of Teichmüller space. The ideal points are given by projective classes of nontrivial isometric actions of π_1 on $\mathbb{R}$–trees.

Given $[\rho] \in R(\Sigma)$, define

$$\ell_\rho\colon \pi_1 \to \mathbb{R}_{\geqslant 0}, \quad [\gamma] \mapsto \inf_{x\in\mathbb{H}^3} d_{\mathbb{H}^3}(x, \rho(\gamma)x).$$

Theorem 2.33 [49] *Consider a sequence $[\rho_n] \in R(\Sigma)$. Then, up to passing to subsequences, one of the following occurs:*

(i) *There is a $[\rho]$ such that $[\rho_n] \to [\rho] \in R^0(\Sigma)$.*

(ii) *There is a minimal nontrivial action of π_1 by isometries on an $\mathbb{R}$–tree T, and numbers $\varepsilon_n \downarrow 0$, such that for all $\gamma \in \pi_1$,*

$$\lim_{n\to\infty} \varepsilon_n \ell_{\rho_n}(\gamma) = \ell_T(\gamma).$$

For the next result we refer to [11, Theorem 3.2], and we note that in the proof of that result harmonicity is not used.

Theorem 2.34 *Suppose that there is a constant $C > 0$ such that the rescalings t_n in Theorem 2.32 satisfy*

$$C^{-1} E_n^{1/2} \leqslant t_n \leqslant C E_n^{1/2},$$

where E_n is the energy of the ρ_n–equivariant harmonic map. Then the length function of the action of π_1 on the Korevaar–Schoen limit appearing in Theorem 2.32 is in the projective class of the Morgan–Shalen limit of the sequence $[\rho_n]$.

3 Bending

In this section we introduce a geometric notion of bending along ρ–equivariant maps $u\colon \widetilde{X} \to \mathbb{H}^3$, and of pairs (A, Ψ). When (A, Ψ) is a Higgs pair, the connection $\nabla = d_A + \Psi$ has monodromy ρ, and u is the ρ–equivariant harmonic map from Theorem 2.14, then we prove that these notions coincide asymptotically at high energy; see Theorem 3.11.

3.1 Bending of maps and connections

3.1.1 Bending of maps We begin with a definition.

Definition 3.1 A *tent* T in $\mathbb{H}^3$ is a pair of totally geodesic compatibly oriented half planes meeting along a geodesic; see Figure 3. The geodesic is called the *crease* and is denoted by γ_T. By "compatibility of the orientations" we will mean the induced orientation on γ_T from the two half planes is opposite. The dihedral angle $\beta_T \in (-\pi, \pi]$ of the two planes is called the *angle of the tent*. This is the angle obtained by rotating the outward normal of one plane (call it A) to the inward normal of the other plane B, counterclockwise in the plane orthogonal to γ_T, with the orientation of this orthogonal plane being induced by the orientation of γ_T coming from plane A. Note that one obtains the same angle going from B to A. By convention, if the union of the half planes forms a totally geodesic plane, then $\beta_T = 0$; if the half planes coincide (necessarily with opposite orientations), then $\beta_T = \pi$.

We will use the following intrinsic way of measuring the angle of a tent. A *crossing* of a tent T is a continuous path $c\colon [0, L] \to T \subset \mathbb{H}^3$ satisfying the following conditions:

- (i) The points $c(0)$ and $c(L)$ lie in different components of $T \setminus \gamma_T$, say T_- and T_+, respectively.
- (ii) There is $0 < L_1 < L$ such that c restricted to the interval $[0, L_1]$ is a C^1 curve in T_- meeting γ_T at $c(L_1)$ transversely.
- (iii) There is $L_1 \leqslant L_2 < L$ such that c restricted to the interval $[L_2, L]$ is a C^1 curve in T_+ meeting γ_T at $c(L_2)$ transversely.
- (iv) The path c restricted to $[L_1, L_2]$ is a portion of γ_T.

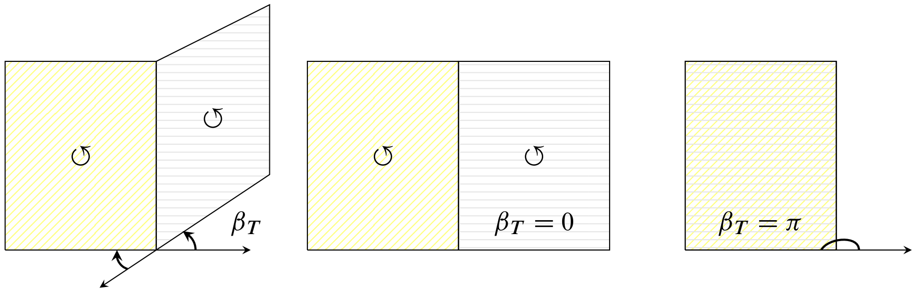

Figure 3: Tents.

The orientation of the tent gives a choice of tangent $N(L_1)$ to γ_T at the crease where a crossing intersects $c(L_1)$. More precisely, $N(L_1)$ is oriented to the left with respect to $\lim_{t\uparrow L_1} c'(t)$. Let $N(t)$ be the parallel translate of $N(L_1)$ along c. Let n_0 and n_L denote the unit normals to T_- and T_+, compatible with the orientations. Let $\tilde{n}(t)$ denote the parallel translation of n_0 along c. Then $\tilde{n}(L)$ and n_L lie in the plane orthogonal to $N(L)$. This plane inherits an orientation from $N(L)$ and the orientation on $\mathbb{H}^3$. Then β_T is the angle from n_L to $\tilde{n}(L)$ with respect to this orientation.

Let $u\colon \widetilde{X} \to \mathbb{H}^3$ be a continuous ρ–equivariant map, and fix $\tilde{p}, \tilde{q} \in \widetilde{X}$.

Assumption 1 The map u is smooth at both $\tilde{p}$ and $\tilde{q}$, and du has maximal rank there.

Definition 3.2 The *bending* $\Theta_u(\tilde{p},\tilde{q}) \in \mathbb{R}/2\pi\mathbb{Z}$ *of u from $\tilde{p}$ to $\tilde{q}$* is defined as follows. Recall that $\mathbb{D}_{u(\tilde{p})}$ and $\mathbb{D}_{u(\tilde{q})}$ denote the oriented totally geodesic planes in $\mathbb{H}^3$ tangent to the images of $du(p)$ and $du(q)$.

(i) If $\mathbb{D}_{u(\tilde{p})}$ and $\mathbb{D}_{u(\tilde{q})}$ meet along a geodesic γ_T, let β_T be the dihedral angle of the tent constructed from the two half-planes in $\mathbb{D}_{u(\tilde{p})} \setminus \gamma_T$ and $\mathbb{D}_{u(\tilde{q})} \setminus \gamma_T$ which contain $u(\tilde{p})$ and $u(\tilde{q})$, respectively, and where the orientation of the tent comes from the orientation on $\mathbb{D}_{u(\tilde{p})}$. Then set $\Theta_u(\tilde{p},\tilde{q})$ equal to β_T if the orientation of $\mathbb{D}_{u(\tilde{q})}$ is compatible with the orientation of the tent in the sense of Definition 3.1, and to $\pi + \beta_T$ if the orientation is incompatible.

(ii) If $\mathbb{D}_{u(\tilde{p})}$ and $\mathbb{D}_{u(\tilde{q})}$ do not intersect, let c be the geodesic between the planes, oriented at one endpoint to agree with the normal of $\mathbb{D}_{u(\tilde{p})}$. If this orientation of c agrees with the normal to $\mathbb{D}_{u(\tilde{q})}$ at the other end point, set $\Theta_u(\tilde{p},\tilde{q}) = 0$. If the orientation is opposite, set $\Theta_u(\tilde{p},\tilde{q}) = \pi$.

(iii) If $\mathbb{D}_{u(\tilde{p})}$ and $\mathbb{D}_{u(\tilde{q})}$ coincide, set $\Theta_u(\tilde{p},\tilde{q}) = 0$ if they have the same orientation, and $\Theta_u(\tilde{p},\tilde{q}) = \pi$ if they have opposite orientations.

By ρ–equivariance of u we clearly have

$$\Theta_u(g\tilde{p}, g\tilde{q}) = \Theta_u(\tilde{p},\tilde{q})$$

for all $\tilde{p},\tilde{q}$ satisfying Assumption 1, and all $g \in \pi_1$. It is also clear that

$$\Theta_u(\tilde{p},\tilde{q}) = \Theta_u(\tilde{q},\tilde{p}).$$

3.1.2 Bending of connections Let d_A be a unitary connection on E, inducing a connection (also denoted by d_A) on the bundle $\sqrt{-1}\mathfrak{g}_E$ of traceless hermitian endomorphisms of E. Fix a one-form $\Psi \in \Omega^1(X, \sqrt{-1}\mathfrak{g}_E)$. We will suppose that the connection $\nabla = d_A + \Psi$ is flat with monodromy ρ. Let $k\colon [0, L] \to X$ be a piecewise C^1 curve.

Assumption 2 The linear map

$$\Psi(k(\sigma))\colon T_{k(\sigma)}X \to \sqrt{-1}\mathfrak{g}_{E,k(\sigma)}$$

has maximal rank at $\sigma = 0, L$.

By analogy to the bending of maps in the previous section, we let $N(0)$ and $N(L)$ be endomorphisms that are a positive multiple of $\Psi(J(\dot{k}(0))$ and $\Psi(J(\dot{k}(L))$. We define the bending angle $\Theta_k(A, \Psi)$ of the pair (A, Ψ) along k using parallel translation with respect to A in place of the Levi-Civita connection. Namely, consider any endomorphism field $V(\sigma)$ along k that is a positive multiple of the endomorphism field $\sigma \mapsto \Psi(\dot{k}(\sigma))$. If k is C^1 on subintervals $[\sigma_{i-1}, \sigma_i]$ for $i = 1, \dots, m$, let $\Pi^{k,A}_{\sigma}$ denote parallel transport in $\sqrt{-1}\mathfrak{g}_E$ along k with respect to A. Then for $\sigma \in [\sigma_{i-1}, \sigma_i]$ let

$$\tilde{n}(\sigma) := \Pi^{k,A}_{\sigma}\Pi^{k,A}_{\sigma_{i-1}} \cdots \Pi^{k,A}_{\sigma_1} n(0) \tag{3-1}$$

be the total parallel transport, where $n(0)$ denotes the endomorphism $\sqrt{-1}[N(0), V(0)]$. Denote by $P(L) \subset \sqrt{-1}\mathfrak{g}_{E,k(\sigma)}$ the orthogonal complement to $N(L)$, and use $N(L)$ and the orientation of $\sqrt{-1}\mathfrak{g}_{E,k(L)}$ to give $P(L)$ an orientation. Now define the *bending of the pair* (A, Ψ) along the path k,

$$\Theta_k(A, \Psi) \in \mathbb{R}/2\pi\mathbb{Z},$$

to be the angle from $\sqrt{-1}[N(L), V(L)]$ to the orthogonal projection of the endomorphism $\tilde{n}(L)$ to $P(L)$ when the latter is nonzero (otherwise bending is undefined).

It is immediate from this definition that the bending $\Theta_k(A, \Psi)$ is invariant under the action of unitary gauge transformations on (A, Ψ). We may therefore write $\Theta_k([(A, \Psi)])$ for the bending of the gauge equivalence class of the pair (A, Ψ).

3.2 Asymptotic bending of Higgs pairs

In this section we relate the total bending in connections to periods of Prym differentials on the spectral curve.

3.2.1 Horizontal lifts for limiting connections Recall the definition of the spectral curve $\pi\colon \widehat{X}_q \to X$ associated to $q \in \mathcal{QD}^*(X)$ in (2-14). Our first goal here is to compute the parallel transport in the bundle $\sqrt{-1}\mathfrak{g}_E$ of hermitian endomorphisms with respect to the (singular) flat connection d_{A_∞} from (2-10). The connection A_∞ induces a unitary connection on the pullback bundle π^*E, which we denote by $\widehat{A}_\infty$.

After pulling back to the spectral curve, the calculation of the parallel transport can be carried out in terms of a suitably chosen oriented frame which we define (cf (2-23)) as

$$W_1 = \begin{pmatrix} 0 & -i\lambda^{-1}\|\lambda\| \\ i\lambda\|\lambda\|^{-1} & 0 \end{pmatrix}, \quad W_2 = \begin{pmatrix} 0 & \lambda^{-1}\|\lambda\| \\ \lambda\|\lambda\|^{-1} & 0 \end{pmatrix}, \quad W_3 = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}, \tag{3-2}$$

with commutation relations

$$[W_i, W_j] = 2i\,\mathrm{sgn}(ijk) W_k. \tag{3-3}$$

From (2-24) we have

$$\pi^*\Psi_\infty = \pi^*\Phi_\infty + \pi^*\Phi_\infty^* = 2\,\mathrm{Re}(\lambda_{\mathrm{SW}}) \otimes W_2. \tag{3-4}$$

Proposition 3.3 *The following hold*:

(i) *The hermitian endomorphism W_2 lies in $\pi^* L^{\mathbb{C}}_{\Phi_\infty}$.*

(ii) *The collection $\{W_1, W_2, W_3\}$ gives an $\widehat{A}^0_\infty$–parallel oriented orthonormal frame for the bundle $\sqrt{-1}\mathfrak{g}_{\widehat{E}}$.*

Proof The proof is a straightforward calculation. We only check that $d_{\widehat{A}^0_\infty} W_1 = 0$. For this we use that one can locally express λ as $q^{1/2}$, so that

$$d(\lambda^{-1}\|\lambda\|) = d(q^{-1/4}\bar{q}^{1/4}) = \tfrac{1}{4}q^{-1/4}\bar{q}^{-3/4}\bar\partial\bar{q} - \tfrac{1}{4}q^{-5/4}\bar{q}^{1/4}\partial q = \tfrac{1}{4}q^{-1/4}\bar{q}^{1/4}(\bar\partial \log \bar{q} - \partial \log q),$$

using that $\bar\partial q = 0$. On the other hand, recall from (2-9) that

$$A^0_\infty = A_0 + \tfrac{1}{2}(\mathrm{Im}\,\bar\partial \log \|q\|) \begin{pmatrix} -i & 0 \\ 0 & i \end{pmatrix} = A_0 + \tfrac{1}{8}(\bar\partial \log \bar{q} - \partial \log q) \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix},$$

where A_0 denotes the Chern connection. Now with $d_{\widehat{A}^0_\infty} W_1 = dW_1 + [\widehat{A}^0_\infty \wedge W_1]$, the last two calculations show that the upper-right entry of $d_{\widehat{A}^0_\infty} W_1$ vanishes, and similarly for the other entries. □

For the following, we make the same assumptions on the path k as in Section 3.2.2.

Proposition 3.4 *Let A_∞ be the unitary connection associated to a limiting configuration in $\mathscr{H}^{-1}_\infty(q)$, and write $A_\infty = A^0_\infty + \eta$, where $\eta = \hat\eta \otimes W_2$, for $\hat\eta$ a harmonic Prym differential*; *see Section 2.2.3. Define the function $\vartheta : [0, L] \to \mathbb{R}$ by*

$$\vartheta(\sigma) := -2i \cdot \int_{\hat{k}([0,\sigma])} \hat\eta. \tag{3-5}$$

Then for $j \in \{1, 2, 3\}$, the parallel transport of the hermitian endomorphisms $W_i(\hat{k}_0)$ of $\pi^* E_{\hat{k}(0)}$ along the path $\hat{k}$ with respect to the connection $d_{\widehat{A}_\infty}$ is given, for $0 \le \sigma \le L$, by

$$\begin{aligned} \Pi^{\hat{k},\widehat{A}_\infty}_\sigma W_2(\hat{k}(0)) &= W_2(\hat{k}(\sigma)), \\ \Pi^{\hat{k},\widehat{A}_\infty}_\sigma W_1(\hat{k}(0)) &= \cos(\vartheta(\sigma)) \cdot W_1(\hat{k}(\sigma)) - \sin(\vartheta(\sigma) \cdot W_3(\hat{k}(\sigma)), \\ \Pi^{\hat{k},\widehat{A}_\infty}_\sigma W_3(\hat{k}(0)) &= \sin(\vartheta(\sigma) \cdot W_1(\hat{k}(\sigma)) + \cos(\vartheta(\sigma) \cdot W_3(\hat{k}(\sigma)). \end{aligned} \tag{3-6}$$

Proof The first line of (3-6) is clear, since W_2 is parallel and commutes with η. For the rest, using (3-3),

$$d_{\widehat{A}_\infty} W_1 = \widehat{\eta} \otimes [W_2, W_1] = -2i\widehat{\eta} \otimes W_3,$$
$$d_{\widehat{A}_\infty} W_3 = \widehat{\eta} \otimes [W_2, W_3] = 2i\widehat{\eta} \otimes W_1.$$

Writing

$$\widetilde{W}_1 = \cos\vartheta \cdot W_1 - \sin\vartheta \cdot W_3 \quad \text{and} \quad \widetilde{W}_3 = \sin\vartheta \cdot W_1 + \cos\vartheta \cdot W_3,$$

we see that $d_{\widehat{A}_\infty} \widetilde{W}^i = 0$ if the derivative $\dot{\vartheta} = -2i\widehat{\eta}$. The result follows. □

3.2.2 Quasitransverse paths with vertical ends Let $q \in \mathcal{SQD}^*(X)$ be a fixed holomorphic quadratic differential, and consider a piecewise C^1 path $k\colon [0, L] \to X$ that is quasitransverse to the horizontal foliation $\mathcal{F}_q^h$ and meets the zeroes of q precisely at its endpoints. In particular, this means that the parameter interval of k admits a subdivision $0 = \sigma_0 < \sigma_1 < \cdots < \sigma_m = L$ such that k restricted to $[\sigma_{i-1}, \sigma_i]$ alternates between vertical and horizontal paths. We say that k has *vertical ends* if the following conditions are satisfied:

(i) The limits $\lim_{\sigma\downarrow 0} \dot{k}(\sigma)$ and $\lim_{\sigma\uparrow L} \dot{k}(\sigma)$ are both nonzero.

(ii) The restrictions $k|_{[0,\sigma_1]}$ and $k|_{[\sigma_{m-1},L]}$ are both vertical.

We will denote by ∂k the section of $k^*(K_X^{-1})$ induced by the derivative $\dot{k}$ of k. The quadratic differential q may be viewed as a section of $\mathrm{Sym}^2(K_X)$, and so it defines a function on $\mathrm{Sym}^2(K_X^{-1})$. We will denote this function applied to $\partial k \otimes \partial k$ by $q(\partial k, \partial k)$. In local coordinates where $q = q(z)\,dz^2$, this is simply $q(\partial k, \partial k)(\sigma) = q(z(\sigma))(\dot{z}(\sigma))^2$. With this understood, if k is parametrized by arc length locally near $\sigma = 0, L$, condition (ii) above implies that

$$\frac{q(\partial k, \partial k)(\sigma)}{\|q\|(k(\sigma))} = -1 \tag{3-7}$$

for σ in $[0, \sigma_1]$ or $[\sigma_{m-1}, L]$. Recall from Section 2.4.1 that since k is assumed to be quasitransverse, we may find a lift $\widehat{k}\colon [0, L] \to \widehat{X}_q$ of the path k to the spectral curve such that $\mathrm{Im}(\lambda_{\mathrm{SW}}(\dot{\widehat{k}})) \leqslant 0$. The path $\widehat{k}$ is piecewise C^1 and meets the zeroes of λ_{SW} precisely at its endpoints. Using condition (ii) above, it is easy to show that the endomorphisms $W_i(\widehat{k}(\sigma))$ in (3-2) extend continuously to the closed interval $[0, L]$.

In the following we suppose that $t > t_0$ is sufficiently large. Let (A_t, Ψ_t) be a solution to the self-duality equations and consider a nearby approximate solution $(A_t^{\mathrm{app}}, \Psi_t^{\mathrm{app}})$ such that the difference between these two pairs is exponentially small in t; cf Section 2.1.4. We indicate with a hat the respective pullbacks of Ψ_t and Ψ_t^{app} to $\sqrt{-1}\mathfrak{g}_{\widehat{E}}$–valued differential forms along $\widehat{k}$ on the spectral curve, ie we set

$$\widehat{\Psi}_t := \pi^*\Psi_t \in \Omega^1(\widehat{X}_q, \sqrt{-1}\mathfrak{g}_{\widehat{E}}),$$

and similarly for Ψ_t^{app}.

Proposition 3.5 *Fix a piecewise C^1 path $k\colon[0,L]\to X$ that is quasitransverse to the horizontal foliation $\mathcal{F}_q^h$ with vertical ends and meets the zeroes of q precisely at its endpoints. Then for $\sigma=0,L$, we have that*

$$\widehat{\Psi}_t^{\mathrm{app}}(\partial\hat{k}(\sigma))=W_1(\hat{k}(\sigma)) \quad \textit{and} \quad \widehat{\Psi}_t^{\mathrm{app}}(J\circ\partial\hat{k}(\sigma))=W_2(\hat{k}(\sigma)).$$

Proof Recall that near the zeroes of q,

$$\Phi_t^{\mathrm{app}}=\begin{pmatrix} 0 & e^{h_t(\|q\|)}\,\|q\|^{1/2} \\ e^{-h_t(\|q\|)}\,\|q\|^{-1/2}\,q & 0\end{pmatrix}.$$

In terms of the tautological section, $\pi^*q=\lambda^2$, the pullback of Φ_t^{app} to the spectral curve can be written in the form

$$\widehat{\Phi}_t^{\mathrm{app}}=\begin{pmatrix} 0 & e^{h_t(\|\lambda\|^2)}\|\lambda\|\lambda^{-1} \\ e^{-h_t(\|\lambda\|^2)}\|\lambda\|^{-1}\lambda & 0\end{pmatrix}\otimes\lambda_{\mathrm{SW}}. \tag{3-8}$$

Similarly,

$$(\widehat{\Phi}_t^{\mathrm{app}})^*=\begin{pmatrix} 0 & e^{-h_t(\|\lambda\|^2)}\|\lambda\|\lambda^{-1} \\ e^{h_t(\|\lambda\|^2)}\|\lambda\|^{-1}\lambda & 0\end{pmatrix}\otimes\overline{\lambda_{\mathrm{SW}}}. \tag{3-9}$$

Now we calculate

$$\begin{aligned}\lambda_{\mathrm{SW}}(\partial\hat{k})&=\lambda(\hat{k})\otimes\pi^*(\partial k),\\ (\lambda_{\mathrm{SW}}(\partial\hat{k}))^2&=(\pi^*q)(\hat{k})\otimes\pi^*(\partial k)^2=\pi^*(q(\partial k,\partial k)).\end{aligned}$$

Since k is assumed to be quasitransverse with vertical ends, by the condition (3-7) it follows that, locally near $\sigma=0,L$,

$$\left(\frac{\lambda_{\mathrm{SW}}(\partial\hat{k})}{\|\lambda\|\circ\hat{k}}\right)^2=\pi^*\left(\frac{q(\partial k,\partial k)}{\|q\|\circ k}\right)=-1,$$

and so by the choice of lift we have

$$\frac{\lambda_{\mathrm{SW}}(\partial\hat{k})}{\|\lambda\|\circ\hat{k}}=-i. \tag{3-10}$$

Similarly,

$$\frac{\lambda_{\mathrm{SW}}(J\circ\partial\hat{k})}{\|\lambda\|\circ\hat{k}}=1. \tag{3-11}$$

Inserting $\partial\hat{k}$ into (3-8) and rearranging the resulting terms slightly yields along $\hat{k}$ the endomorphism field

$$\widehat{\Phi}_t^{\mathrm{app}}(\partial\hat{k})=\begin{pmatrix} 0 & e^{h_t(\|\lambda\|^2)}\,\|\lambda\|^2\lambda^{-1} \\ e^{-h_t(\|\lambda\|^2)}\lambda & 0\end{pmatrix}\frac{\lambda_{\mathrm{SW}}(\partial\hat{k})}{\|\lambda\|},$$

and similarly

$$(\widehat{\Phi}_t^{\mathrm{app}})^*(\partial\hat{k})=\begin{pmatrix} 0 & e^{-h_t(\|\lambda\|^2)}\|\lambda\|^2\lambda^{-1} \\ e^{h_t(\|\lambda\|^2)}\lambda & 0\end{pmatrix}\frac{\overline{\lambda_{\mathrm{SW}}}(\partial\hat{k})}{\|\lambda\|}.$$

By Lemma 2.4(iv), $\exp(\pm h_t(\|\lambda\|^2)) \sim \|\lambda\|^{\mp 1}$ for $\|\lambda\|$ small. Together with (3-10) this implies the convergence

$$\widehat{\Psi}_t^{\mathrm{app}}(\partial \hat{k}) \to \begin{pmatrix} 0 & -i\,\|\lambda\|\lambda^{-1} \\ i\,\|\lambda\|^{-1}\lambda & 0 \end{pmatrix} = W_1 \quad \text{as } \|\lambda\| \to 0.$$

In a completely analogous way one obtains that along $\hat{k}$,

$$\widehat{\Psi}_t^{\mathrm{app}}(J \circ \partial \hat{k}) \to \begin{pmatrix} 0 & \|\lambda\|\lambda^{-1} \\ \|\lambda\|^{-1}\lambda & 0 \end{pmatrix} = W_2 \quad \text{as } \|\lambda\| \to 0.$$

This proves the proposition. □

3.2.3 Limit of bending for connections We relate the limit as $t \to \infty$ of the bending $\Theta_k(A_t, \Psi_t)$ defined in Section 3.1.2 to periods of Prym differentials on the spectral curve. This is the key result of this section.

Proposition 3.6 *Fix a holomorphic quadratic differential $q \in \mathcal{SQD}^*(X)$. Let $k\colon [0, L] \to X$ be a piecewise C^1 path, and fix a lift $\hat{k}\colon [0, L] \to \widehat{X}_q$ to the spectral curve such that $\pi \circ \hat{k} = k$. Assume that k is quasitransverse to the horizontal foliation $\mathcal{F}_q^h$ with vertical ends, and meets the zeroes of q precisely at its endpoints. Consider a family $[(A_t, t\,\Psi_t)] \in \mathscr{H}^{-1}(t^2 q_t)$ for $q_t \in \mathcal{SQD}^*(X)$. Letting $t \to \infty$, suppose that $q_t \to q$ and that $[(A_t, t\,\Psi_t)]$ converges to $[(A_\infty, \Psi_\infty)] \in \mathscr{H}_\infty^{-1}(q)$ in the sense of Definition 2.5. Write $A_\infty = A_\infty^0 + \eta$ with a unique one-form $\eta \in \mathcal{H}^1(X_q^\times, L_q)$ as in Proposition 3.4. Then*

$$\lim_{t \to \infty} \Theta_k([(A_t, \Psi_t)]) = -2i \int_{\hat{k}} \hat{\eta} \mod 2\pi\mathbb{Z}, \tag{3-12}$$

where $\hat{\eta} \in \mathcal{H}^1_{\mathrm{odd}}(\widehat{X}_q, i\mathbb{R})$ is the Prym differential corresponding to η from Proposition 2.11.

Proof The proof is in seven steps.

Step 1 By Definition 2.5 (Approximation), there exists a family of one-forms $\eta_t \in \Omega^1(X, \mathfrak{g}_E)$ as in equation (2-10) such that the difference $(A_t^{\mathrm{app}}(q_t) + \eta_t, \Psi_t^{\mathrm{app}}(q_t)) - (A_t, \Psi_t)$ satisfies an exponentially decaying C^ℓ bound in the parameter t. Hence the difference of the holonomies along the path k in (3-1) corresponding to the connections $A_t^{\mathrm{app}}(q_t) + \eta_t$ and A_t tends to zero as $t \to \infty$. We conclude that it suffices to prove the claim with the family (A_t, Ψ_t) replaced by the family $(A_t^{\mathrm{app}}(q_t) + \eta_t, \Psi_t^{\mathrm{app}}(q_t))$.

Step 2 Recall from Section 3.2.2 that by our assumptions on the path k, the parameter interval of k admits a subdivision $0 = \sigma_0 < \sigma_1 < \cdots < \sigma_m = L$ such that k restricted to the subintervals $[\sigma_{i-1}, \sigma_i]$ alternates between vertical and horizontal paths. Since $q_t \to q$ as $t \to \infty$, and hence also the zeroes of q_t converge to the zeroes of q, we may choose a family $k_t\colon [0, L] \to X$ of piecewise C^1 paths with the following properties:

(i) k_t meets the zeroes of q_t precisely at its endpoints,

(ii) k_t is quasitransverse to the horizontal foliation $\mathcal{F}_{q_t}^h$ with vertical ends, and

(iii) $k_t \to k$ in C^1 as $t \to \infty$ on each subinterval $[\sigma_{i-1}, \sigma_i]$ for $1 \le i \le m$.

For each t, we then fix a lift $\hat{k}_t\colon [0, L] \to \widehat{X}_{q_t}$ to the spectral curve $\pi_t\colon \widehat{X}_{q_t} \to X$ such that $\pi_t \circ \hat{k} = k$.

Step 3 For each fixed parameter t consider the family $(A_s^{\mathrm{app}}(q_t)+\eta_t, \Psi_s^{\mathrm{app}}(q_t))$ for $s>0$. We recall from Section 3.1.2 the definition of bending, and apply it to the pair $(A_s^{\mathrm{app}}(q_t)+\eta_t, \Psi_s^{\mathrm{app}}(q_t))$ and the path k_t.

We shall be working on the spectral curve $\widehat{X}_{q_t}$. After applying a unitary gauge transformation to $(A_s^{\mathrm{app}}(q_t)+\eta_t, \Psi_s^{\mathrm{app}}(q_t))$ we may assume that its pullback along the projection $\pi_t\colon \widehat{X}_{q_t}\to X$ is the pair $(\widehat{A}_s^{\mathrm{app}}(q_t)+\hat{\eta}_t, \widehat{\Psi}_s^{\mathrm{app}}(q_t))$, where $\hat{\eta}_t\in\mathcal{H}^1_{\mathrm{odd}}(\widehat{X}_{q_t}, i\mathbb{R})$ is a Prym differential as in Section 2.2.2. We also fix a lift $\hat{k}_t\colon[0,L]\to\widehat{X}_{q_t}$ of k_t such that $\pi\circ\hat{k}_t=k_t$. Keeping $\sigma=0, L$ fixed, by Proposition 3.5 we may define the endomorphisms

$$\widehat{V}_t(\sigma):=\widehat{\Psi}_s^{\mathrm{app}}(q_t)(\partial\hat{k}_t(\sigma))=W_1(\hat{k}_t(\sigma))\quad\text{and}\quad \widehat{N}_t(\sigma):=\widehat{\Psi}_s^{\mathrm{app}}(q_t)(J\circ\partial\hat{k}_t(\sigma))=W_2(\hat{k}_t(\sigma)).$$

Note that these do not depend on s. Using the commutation relations from (3-3) it follows that

$$\sqrt{-1}[\widehat{N}_t(\sigma),\widehat{V}_t(\sigma)]=\sqrt{-1}[W_2(\hat{k}_t(\sigma)),W_1(\hat{k}_t(\sigma))]=2\,W_3(\hat{k}_t(\sigma)).$$

Next we define the endomorphism

$$\hat{n}_t(0):=\sqrt{-1}[\widehat{N}_t(0),\widehat{V}_t(0)]=2\,W_3(\hat{k}_t(0))$$

and consider its parallel transport

$$\tilde{n}_{t,s}(L):=\Pi_L^{\hat{k}_t,\widehat{A}_s^{\mathrm{app}}(q_t)+\hat{\eta}_t}\,\hat{n}_t(0) \tag{3-13}$$

in $\sqrt{-1}\mathfrak{g}_E$ along the path $\hat{k}_t$ with respect to the connection $\widehat{A}_s^{\mathrm{app}}(q_t)+\hat{\eta}_t$. Let $\widehat{P}_t(L)\subset\sqrt{-1}\,\hat{\mathfrak{g}}_{E,\hat{k}_t(L)}$ be the orthogonal complement to $\widehat{N}_t(L)=W_2(\hat{k}_t(L))$. By Proposition 3.3, a frame for this complement is determined by $W_1(\hat{k}_t(L))$ and $W_3(\hat{k}_t(L))$. We use this ordering of the frame to define an orientation on the plane $\widehat{P}_t(L)$. The bending

$$\Theta_k(A_s^{\mathrm{app}}(q_t)+\eta_t, \Psi_s^{\mathrm{app}}(q_t))\in\mathbb{R}/2\pi\mathbb{Z} \tag{3-14}$$

is then given by the angle from $\sqrt{-1}[\widehat{N}_t(L),\widehat{V}_t(L)]=2\,W_3(\hat{k}_t(L))$ to the orthogonal projection of the endomorphism $\tilde{n}_{t,s}(L)$ to $\widehat{P}_t(L)$ with respect to this orientation.

Step 4 Recall from Section 2.1.4 that $A_s^{\mathrm{app}}(q_t)\to A_\infty^0(q_t)$ as $s\to\infty$ in C^∞ locally on compact subsets of $X_{q_t}^\times$, where $A_\infty^0(q_t)$ is the Fuchsian connection from (2-9). Then clearly we also have the local C^∞ convergence $A_s^{\mathrm{app}}(q_t)+\eta_t\to A_\infty^0(q_t)+\eta_t$ as $s\to\infty$. In preparation for Step 5, we now prove that there exists $t_0=t_0(q)>0$ such that the following holds: For every $\varepsilon>0$ and $\ell\geqslant 0$ there exists $s_0=s_0(\varepsilon,q,\ell)\geqslant t_0$ such that

$$\|k_t^*(A_s^{\mathrm{app}}(q_t)+\eta_t)-k_t^*(A_\infty^0(q_t)+\eta_t)\|_{C^\ell([0,L])}<\varepsilon \tag{3-15}$$

for all $s\geqslant s_0$ and every $t\geqslant t_0$. Here $k_t^*(A_s^{\mathrm{app}}(q_t)+\eta_t)$ denotes the pullback of the connection $A_s^{\mathrm{app}}(q_t)$ along the path $k_t\colon[0,L]\to X$, and likewise for $k_t^*(A_\infty^0(q_t)+\eta_t)$.

Locally on each punctured disk $\mathcal{D}_p^{\times}$ endowed with polar coordinates (r,θ), the connection $A_s^{\mathrm{app}}(q_t)$ takes the form

$$A_s^{\mathrm{app}}(q_t)(r,\theta) = f_s(r)\begin{pmatrix} -i & 0 \\ 0 & i \end{pmatrix} d\theta, \tag{3-16}$$

with a smooth function $f_s\colon [0,\infty) \to \mathbb{R}$ as in Section 2.1.4. Hence, writing the radial and angular components of the path $\sigma \mapsto k_t(\sigma)$ as $k_t(\sigma) = (r(\sigma),\theta(\sigma))$, it follows that

$$k_t^* A_s^{\mathrm{app}}(q_t)(\sigma) = f_s(r(\sigma))\begin{pmatrix} -i & 0 \\ 0 & i \end{pmatrix} \dot\theta(\sigma)\,d\sigma. \tag{3-17}$$

Since by assumption k_t has vertical ends and meets the zeroes of q_t precisely at its endpoints, we see that $\dot\theta(\sigma)$ and hence $k_t^* A_s^{\mathrm{app}}(q_t)$ vanishes identically outside some proper subinterval $[L_1,L_2] \subset [0,L]$. This subinterval may be chosen independently of t. Definition 2.5(iii) implies that (after shrinking the disk $\mathcal{D}_p$ slightly if necessary, so that $k_t([L_1,L_2])$ lies outside $\mathcal{D}_p$) the family of functions $s \mapsto f_s \circ r$ converges in $C^\ell([L_1,L_2])$ to the function $f_\infty \circ r$ as $s \to \infty$. This proves the claim.

Step 5 Keep the constant $t_0 = t_0(q) > 0$ from Step 4. We consider the bending in (3-14) for large s, and prove that for every $\varepsilon > 0$ there exists $s_0 = s_0(\varepsilon,q) \geqslant t_0$ such that

$$\left|\Theta_k(A_s^{\mathrm{app}}(q_t)+\eta_t, \Psi_s^{\mathrm{app}}(q_t)) - \left(-2i\int_{\hat k_t} \hat\eta_t \mod 2\pi\mathbb{Z}\right)\right| < \varepsilon \tag{3-18}$$

for all $s \geqslant s_0$ and every $t \geqslant t_0$.

To see this, first note that after passing to the spectral curve $\hat X_{q_t}$, by Step 3 we have the estimate

$$\|\hat k_t^* \hat A_s^{\mathrm{app}}(q_t) - \hat k_t^* \hat A_\infty^0(q_t)\|_{L^p([0,L])} < \varepsilon \tag{3-19}$$

for all $s \geqslant s_0$ and every $t \geqslant t_0$, where $\hat A_\infty^0(q_t)$ denotes the pullback of $A_\infty^0(q_t)$ along the projection $\pi\colon \hat X_{q_t} \to X$. Let us now compare the parallel transports

$$\tilde n_{t,s}(L) = -2\Pi_L^{\hat k_t, \hat A_s^{\mathrm{app}}(q_t)+\hat\eta_t} W_3(\hat k_t(0))$$

from (3-13) with the parallel transport

$$\tilde n_{t,\infty}(L) := -2\Pi_L^{\hat k_t, \hat A_\infty^0(q_t)+\hat\eta_t} W_3(\hat k_t(0)).$$

It follows from (3-19) that there is some constant $C > 0$ such that

$$|\tilde n_{t,s}(L) - \tilde n_{t,\infty}(L)| < C\varepsilon \tag{3-20}$$

for all $s \geqslant s_0$ and every $t \geqslant t_0$. Now by Proposition 3.4, we have

$$\tilde n_{t,\infty}(L) = \sin(\vartheta_t(L)) \cdot W_1(\hat k_t(L)) + \cos(\vartheta_t(L)) \cdot W_3(\hat k_t(L)), \quad \text{where } \vartheta_t(L) = -2i \cdot \int_{\hat k_t} \hat\eta_t.$$

Observe that the endomorphism $\tilde n_{t,\infty}(L)$ is contained in the plane $\hat P_t(L)$ defined in Step 2, and that the angle from $\sqrt{-1}[\hat N_t(L), \hat V_t(L)] = 2\,W_3(\hat k_t(L))$ to $\tilde n_{t,\infty}(L)$ with respect to the orientation on $P_t(L)$ is given by $\vartheta_t(L)$.

The estimate in (3-18) now follows from (3-20) and the definition of bending in Step 2.

Step 6 By assumption and Steps 1 and 2, letting $t \to \infty$ we have that $q_t \to q$, $\eta_t \to \eta$ and $k_t \to k$, which immediately implies that

$$\lim_{t\to\infty} \int_{\hat{k}_t} \hat{\eta}_t = \int_{\hat{k}} \hat{\eta}. \tag{3-21}$$

Step 7 Combining Steps 5 and 6 we infer that in the estimate

$$\begin{aligned}\Big|\Theta_k(A_t^{\mathrm{app}}(q_t)+\eta_t,\Psi_t^{\mathrm{app}}(q_t)) - \Big(-2i\int_{\hat{k}}\hat{\eta} \bmod 2\pi\mathbb{Z}\Big)\Big| \\ \leqslant \Big|\Theta_k(A_t^{\mathrm{app}}(q_t)+\eta_t,\Psi_t^{\mathrm{app}}(q_t)) - \Big(-2i\int_{\hat{k}_t}\hat{\eta}_t \bmod 2\pi\mathbb{Z}\Big)\Big| + \Big|\Big(-2i\int_{\hat{k}_t}\hat{\eta}_t\Big) - \Big(-2i\int_{\hat{k}}\hat{\eta}\Big)\Big|,\end{aligned}$$

both terms on the right-hand side tend to zero as $t \to \infty$. The proposition is proved. □

Remark 3.7 Proposition 3.6 and equation (3-12) apply equally well to the modified saddle connections (which are not quasitransverse).

3.3 Comparison of bending

In this section we show that for large energy, the bending of equivariant harmonic maps defined in Section 3.1.1 nearly coincides with the bending of the associated Higgs pair along quasitransverse paths. The main result is Theorem 3.11 below. First, we need a somewhat standard preliminary result on parallel translation, which we provide in the next subsection.

3.3.1 Parallel translation for C^1–close curves Let c and c_0 be piecewise C^1 curves $[0, L] \to \mathbb{H}^3$. Fix $\varepsilon > 0$. We say that c and c_0 are *C^0_ε–close* if

$$\max_{0\leqslant t\leqslant L} d_{\mathbb{H}^3}(c(t), c_0(t)) < \varepsilon. \tag{3-22}$$

Let us view c and c_0 as curves in the hermitian model $\mathscr{D}$ of $\mathbb{H}^3$; see Section 2.3.2. Recall the metric on $T\mathbb{H}^3 \otimes \mathbb{C}$ defined in (2-35) for the trace model. A C^0–bound on the distance in $\mathbb{H}^3$ between c and c_0 induces one on the pointwise norms of $(1 - cc_0^{-1})$ and $(1 - c_0c^{-1})$. Using this fact it is easy to prove the following.

Lemma 3.8 *There are constants $C(\varepsilon) \geqslant 1$, for which $\lim_{\varepsilon\to 0} C(\varepsilon) = 1$, with the following significance. If c and c_0 are C^0_ε–close, then for all $M \in T\mathbb{H}^3 \otimes \mathbb{C}$, and all $t \in [0, L]$,*

$$C(\varepsilon)^{-1}\|M\|_{c(t)} \leqslant \|M\|_{c_0(t)} \leqslant C(\varepsilon)\|M\|_{c(t)}.$$

Definition 3.9 Let c and c_0 be as above. Fix $\varepsilon > 0$. We say that c and c_0 are *C^1_ε–close* if they are C^0_ε–close, and

$$\max_{0\leqslant t\leqslant L} \|\dot{c}c^{-1} - \dot{c}_0c_0^{-1}\|_{c_0(t)} < \varepsilon.$$

We emphasize that here we view $\dot{c}c^{-1}$ and $\dot{c}_0c_0^{-1}$ as sections of the trivial bundle $T\mathbb{H}^3 \otimes \mathbb{C} \simeq \mathbb{H}^3 \times \mathbb{C}^3$, and using this trivialization we compare vectors at arbitrary fibers. Note that because of Lemma 3.8, the relationship of being C^1–close is symmetric (after possibly multiplying ε by a distortion that is nearly 1). The curves are not assumed to be parametrized by arc length.

Lemma 3.10 *Let c and c_0 be curves in $\mathbb{H}^3$. Suppose $v(0) \in T_{c(0)}\mathbb{H}^3$ and $v_0(0) \in T_{c_0(0)}\mathbb{H}^3$ are unit vectors, and let $v(t)$ and $v_0(t)$ denote parallel translation along $c(t)$ and $c_0(t)$, respectively. If c and c_0 are C^1_ε–close with $0 < \varepsilon \leqslant 1/4L$, then*

$$\max_{0 \leqslant t \leqslant L} \|v(t) - v_0(t)\|_{c_0(t)} \leqslant 2\|v(0) - v_0(0)\|_{c_0(0)} + 4L\varepsilon.$$

Proof By (2-36) we have

$$\dot{v}(t) = \tfrac{1}{2}[\dot{c}c^{-1}, v(t)] \quad \text{and} \quad \dot{v}_0(t) = \tfrac{1}{2}[\dot{c}_0c_0^{-1}, v_0(t)].$$

Write

$$v(t) = v_0(t) + R(t) \quad \text{and} \quad \dot{c}c^{-1} = \dot{c}_0c_0^{-1} + r(t)$$

for traceless matrix valued functions $R(t)$ and $r(t)$. Hence,

$$2\dot{R}(t) = [r(t), v_0(t) + R(t)] + [\dot{c}_0c_0^{-1}, R(t)]. \tag{3-23}$$

Now

$$\begin{aligned}\frac{d}{dt}\|R(t)\|^2_{c_0} &= \frac{d}{dt}\operatorname{tr}(Rc_0R^*c_0^{-1}) \\ &= \operatorname{tr}(\dot{R}c_0R^*c_0^{-1}) + \operatorname{tr}(Rc_0\dot{R}^*c_0^{-1}) + \operatorname{tr}(R\dot{c}_0R^*c_0^{-1}) - \operatorname{tr}(Rc_0R^*c_0^{-1}\dot{c}_0c_0^{-1}).\end{aligned}$$

One can see that the last two terms on the right-hand side of the equation above are canceled by the last term on the right-hand side of (3-23) (and the similar equation for the adjoint). Thus we are left with

$$\frac{d}{dt}\operatorname{tr}(Rc_0R^*c_0^{-1}) = \tfrac{1}{2}\operatorname{tr}([r, v_0 + R]c_0R^*c_0^{-1}) + \tfrac{1}{2}\operatorname{tr}(Rc_0[v_0^* + R^*, r^*]c_0^{-1}).$$

Since the norm of r is less than ε, and $v_0(t)$ is a unit vector, we see that

$$\frac{d}{dt}\|R(t)\|^2_{c_0} \leqslant 2\varepsilon(\|R(t)\|^2_{c_0} + \|R(t)\|_{c_0}). \tag{3-24}$$

Let $0 \leqslant t_m \leqslant L$ be the point at which $\|R(t)\|^2_{c_0}$ attains it maximum. Then from (3-24) we have

$$\begin{aligned}\|R(t_m)\|^2_{c_0} - \|R(0)\|^2_{c_0} = \int_0^{t_m} \frac{d}{dt}\|R(t)\|^2_{c_0}\,dt &\leqslant 2\varepsilon\int_0^{t_m}(\|R(t)\|^2_{c_0} + \|R(t)\|_{c_0})\,dt \\ &\leqslant 2L\varepsilon(\|R(t_m)\|^2_{c_0} + \|R(t_m)\|_{c_0}).\end{aligned} \tag{3-25}$$

Since we assume $\varepsilon \leqslant 1/4L$, it follows from (3-25) that

$$\|R(t_m)\|_{c_0} \leqslant 2(\|R(0)\|_{c_0} + 2L\varepsilon).$$

This completes the proof. □

3.3.2 Asymptotic equivalence of bending

Theorem 3.11 *Let $K \subset \mathcal{QD}^*(X)$ be a cone on a compact subset of $\mathcal{SQD}^*(X)$, and fix $\delta > 0$. Let $u\colon \widetilde{X} \to \mathbb{H}^3$ be a ρ–equivariant harmonic map, $\nabla = d_A + \Psi$ a Higgs pair. Let $\operatorname{Hopf}(u) = -q \in K$, and let $\tilde{p}^-$ and $\tilde{p}^+$ be lifts of zeroes p^- and p^+ of q. Let k be a quasitransverse path (or modified saddle connection) from p^- to p^+ that lifts to a path $\tilde{k}$ from $\tilde{p}^-$ to $\tilde{p}^+$. Assume k has small vertical ends and meets the zeroes of q only at p^- and p^+. Then if $\|q\|_1$ is sufficiently large (depending on K, ε, k), we have*

$$|\Theta_u(\tilde{p}^-, \tilde{p}^+) - \Theta_k(A, \Psi)| < \delta.$$

Proof Write $0 = \sigma_0 < \sigma_1 < \cdots < \sigma_m = L$, so that k restricted to $[\sigma_{i-1}, \sigma_i]$ alternates between C^1 vertical and horizontal paths. By assumption, $k|_{[0,\sigma_1]}$ and $k|_{[\sigma_{m-1},L]}$ are vertical. Let T be the tent with crease γ_T associated to the totally geodesic planes $\mathbb{D}_{u(\tilde{p}^-)}$ and $\mathbb{D}_{u(\tilde{p}^+)}$. By Proposition 2.30, the images by u of sufficiently small hexagonal domains $\Omega_{\tilde{p}^-}$ and $\Omega_{\tilde{p}^+}$ are C^1–close to the planes $\mathbb{D}_{u(\tilde{p}^-)}$ and $\mathbb{D}_{u(\tilde{p}^+)}$. By Proposition 2.29 it follows that the image of $k|_{[\sigma_1,\sigma_{m-1}]}$ is C^1_ε–close to γ_T. By Proposition 3.5, for sufficiently small vertical ends, the normal vector to $\mathbb{D}_{u(\tilde{p}^-)}$ is close to the vector $n(0)$ in Section 3.1.2, and similarly at $\mathbb{D}_{u(\tilde{p}^+)}$. Hence, by Lemma 3.10, parallel translation of the normal vector to $\Omega_{u(\tilde{p}^-)}$ along γ_T is close to the parallel translation $\tilde{n}$ along k. The result now follows from Theorem 2.14(i) and (ii), and the discussion in Section 3.1.1. □

4 Pleated surfaces

In this section we review the notion of a transverse cocycle for a lamination. The key results are

- Lemma 4.4, where we relate the bending cocycle of a pleated surface to its geometric bending in the sense of Section 3.1.1,
- Theorem 4.5, where we relate the group of bending cocycles to the torus of Prym differentials, and
- Theorem 4.16, where we show that the limit of a bending cocycle is determined by the periods of a Prym differential.

4.1 Transverse cocycles

4.1.1 Definitions Let Λ be a maximal geodesic lamination on a hyperbolic surface S with underlying smooth surface Σ. Recall from Section 2.4.1 that $\mathcal{P}(\Lambda)$ denotes the set of plaques in $\mathbb{H}^2 \setminus \widetilde{\Lambda}$, and note that there is a free action of π_1 on $\mathcal{P}(\Lambda)$ with finite quotient. Let G denote an abelian group that is either $\mathbb{R}$ or $S^1 \simeq \mathbb{R}/2\pi\mathbb{Z}$. A *$G$–valued transverse cocycle* for Λ is a map $\boldsymbol{\alpha}$ which sends every arc k transverse to Λ to an element $\boldsymbol{\alpha}(k) \in G$, and which satisfies the following two properties. First, for $k = k_1 \cup k_2$ a decomposition of k into two subarcs with disjoint interiors, we have $\boldsymbol{\alpha}(k) = \boldsymbol{\alpha}(k_1) + \boldsymbol{\alpha}(k_2)$. Second, $\boldsymbol{\alpha}$ is invariant under Λ–transverse homotopies, in the sense that if k can be taken to k' by a homotopy of S that preserves Λ, then $\boldsymbol{\alpha}(k) = \boldsymbol{\alpha}(k')$. In particular (see [5, page 7]), a G–valued transverse cocycle may be taken to be a function $\boldsymbol{\alpha}\colon \mathcal{P}(\Lambda) \times \mathcal{P}(\Lambda) \to G$ satisfying:

(i) **Equivariance** $\boldsymbol{\alpha}(gP, gQ) = \boldsymbol{\alpha}(P, Q)$ for all $g \in \Gamma$.

(ii) **Symmetry** $\boldsymbol{\alpha}(P, Q) = \boldsymbol{\alpha}(Q, P)$.

(iii) **Additivity** $\boldsymbol{\alpha}(P, Q) = \boldsymbol{\alpha}(P, R) + \boldsymbol{\alpha}(R, Q)$, if R separates P from Q.

We shall henceforth denote transverse cocycles by $\boldsymbol{\sigma}$ for $G = \mathbb{R}$ and $\boldsymbol{\beta}$ for $G = S^1$, and in the latter case we will continue to use additive notation in (iii). Denote by $\mathcal{H}(\Lambda, G)$ the space of transverse cocycles. For $G = \mathbb{R}$, this is a real vector space of dimension $6g-6$; see [6, page 119]. We refer to elements $\boldsymbol{\sigma} \in \mathcal{H}(\Lambda, \mathbb{R})$ as *shearing cocycles*. The space $\mathcal{H}(\Lambda, S^1)$ has two components, each of which is a $(6g-6)$–dimensional torus. We denote by $\mathcal{H}^o(\Lambda, S^1)$ the component containing the identity cocycle: $\boldsymbol{\beta}(P, Q) = 0$ for all P, Q. We refer to elements $\boldsymbol{\beta} \in \mathcal{H}^o(\Lambda, S^1)$ as *bending cocycles*. We will sometimes make reference to a norm $\|\cdot\|$ on $\mathcal{H}(\Lambda, \mathbb{R})$, which is fixed once and for all.

Convergence of transverse cocycles for families of laminations may be defined in a weak sense as functions on pairs of plaques. More precisely, recall from Remark 2.27 that there is a one-to-one correspondence between the lifts of zeroes of q and the plaques $\mathcal{P}(\Lambda)$ of any maximalization Λ of Λ_q^h. If $q_n \to q \in \mathcal{QD}^*(X)$, then $\Lambda_{q_n}^h$ converges in the Hausdorff sense to a lamination with Λ_q^h as a sublamination. For maximalizations Λ_n, we define convergence $\Lambda_n \to \Lambda$ again in the Hausdorff sense. In this case, for n sufficiently large we have bijections

$$r_{\Lambda_n}^{\Lambda} : \mathcal{P}(\Lambda) \xrightarrow{\sim} \mathcal{P}(\Lambda_n).$$

Definition 4.1 With the notation above, suppose $\Lambda_n \to \Lambda$. Let $\boldsymbol{\alpha}_n$ (resp. $\boldsymbol{\alpha}$) be either shearing or bending cocycles for Λ_n (resp. Λ). We say that $\boldsymbol{\alpha}_n$ *converges to* $\boldsymbol{\alpha}$ (and write $\boldsymbol{\alpha}_n \to \boldsymbol{\alpha}$) if

$$\lim_{n\to\infty} \boldsymbol{\alpha}_n(r_{\Lambda_n}^{\Lambda}(P), r_{\Lambda_n}^{\Lambda}(Q)) = \boldsymbol{\alpha}(P, Q) \quad \text{for all } P, Q \in \mathcal{P}(\Lambda).$$

We note for clarification that $r_{\Lambda_n}^{\Lambda}$ does not, in general, preserve the separation relations of plaques, so the pullback $\alpha_n \circ r_{\Lambda_n}^{\Lambda}$ of a cocycle α_n on Λ_n will not necessarily satisfy the additivity condition on $\mathcal{P}(\Lambda)$.

We will also need the following elementary properties of cocycles (recall the definition of $\mathcal{H}(\tau, G)$ from Section 2.4.2).

Proposition 4.2 *Let $q \in \mathcal{QD}^*(X)$ and let Λ be a maximalization of Λ_q^h.*

(i) *There is a finite set $\mathcal{P}' \subset \mathcal{P}(\Lambda)$ such that if $\boldsymbol{\alpha} \in \mathcal{H}(\Lambda, G)$ vanishes on $\mathcal{P}' \times \mathcal{P}'$, then $\boldsymbol{\alpha}$ vanishes identically.*

(ii) *There is a complete train track τ carrying Λ_q^h, and a bijection $\mathcal{H}(\tau, G) \simeq \mathcal{H}(\Lambda, G)$.*

Proof Assertion (i) follows by finite-dimensionality. More precisely, let τ' be a train track that *snugly* carries Λ in the sense of [6, page 114]. Then by [6, Theorem 11], $\mathcal{H}(\Lambda, G) \simeq \mathcal{H}(\tau', G)$. The identification assigns weights to branches of τ' that are equal to the value of the cocycle on the plaques defined by the complementary regions. Since there are only finitely many of these, the claim follows. For assertion (ii), the existence of τ follows from [54]. A train track τ'' that snugly carries Λ can be obtained by splitting τ along branches corresponding to saddle connections. Hence, $\mathcal{H}(\Lambda, G) \simeq \mathcal{H}(\tau'', G)$ as above. On the other hand, $\mathcal{H}(\tau'', G) \simeq \mathcal{H}(\tau, G)$ from properties of splittings. This completes the proof. □

4.1.2 Shearing cocycles Given a marked hyperbolic surface S with maximal geodesic lamination Λ, there is a uniquely defined transverse cocycle $\boldsymbol{\sigma} \in \mathcal{H}(\Lambda, \mathbb{R})$ called the *shearing cocycle of* S. For the precise definition see [5, page 10]. We will need the following formula for $\boldsymbol{\sigma}$.

Let $P, Q \in \mathcal{P}(\Lambda)$, and choose an arc k from P to Q in $\widetilde{S}$ that is transverse to $\widetilde{\Lambda}$. For each component d of $k \setminus \widetilde{\Lambda}$ disjoint from P and Q, let x_d^+ and x_d^- be the positive and negative endpoints, respectively, of the (oriented) segment d. We let $d_- = P \cap k$ and $d_+ = Q \cap k$. Define $x_{d_-}^+$ to be the positive endpoint of d_-, and $x_{d_+}^-$ the negative endpoint of d_+. Denote the leaves of $\widetilde{\Lambda}$ passing through $x_d^\pm$ by $g_d^\pm$, and similarly for $x_{d_\pm}$. For each component d, $d_\pm$, let $h\colon g_d^\pm \to \mathbb{R}$ denote the (signed) distance from the foot[5] determined by viewing the geodesics as boundaries of the ideal triangle corresponding to the component d. With this understood, we have the following expression for the shearing cocycle of S (see [5, Lemma 7]):

$$\boldsymbol{\sigma}(P, Q) = h(x_{d_-}^+) - h(x_{d_+}^-) + \sum_{d \neq d_+, d_-} (h(x_d^+) - h(x_d^-)). \tag{4-1}$$

We also note the following:

(i) If $\boldsymbol{\sigma}$ is the shearing cocycle of a hyperbolic surface S and $\boldsymbol{\alpha} \in \mathcal{H}(\Lambda, \mathbb{R})$ with $\|\boldsymbol{\alpha}\|$ sufficiently small, then $\boldsymbol{\sigma} + \boldsymbol{\alpha}$ is the shearing cocycle of some hyperbolic surface [5, Proposition 13]. This is the generalization of Thurston's earthquake map.

(ii) The map $T(\Sigma) \to \mathcal{H}(\Lambda, \mathbb{R})$ which associates the shearing cocycle to a hyperbolic metric is injective onto an open convex polyhedral cone $\mathcal{C}(\Lambda)$; see [5, Corollary 21].

4.1.3 Bending cocycles Recall from the introduction that a pleated surface $\mathsf{P} = (S, f, \Lambda, \rho)$ consists of a marked hyperbolic surface S, a maximal geodesic lamination $\Lambda \subset \Sigma$, and a map $f : \widetilde{S} \to \mathbb{H}^3$ that is totally geodesic on the components of $\widetilde{S} \setminus \widetilde{\Lambda}$, maps leaves of $\widetilde{\Lambda}$ isometrically to geodesics, and is ρ–equivariant for a representation $\rho\colon \pi_1 \to \mathrm{PSL}(2, \mathbb{C})$. Such a ρ, which in this paper we take to be in $R^o(\Sigma)$, is necessarily irreducible; see [5, page 36]. We sometimes denote pleated surfaces by just $f : \widetilde{S} \to \mathbb{H}^3$ when context provides the other data.

In addition to the shearing cocycle for the hyperbolic surface S, there is a uniquely defined bending cocycle $\boldsymbol{\beta} \in \mathcal{H}^0(\Lambda, S^1)$. As in the previous section, we will need a particular formula for this, which we describe below.

Choose a ρ–equivariant differentiable vector field v on $\mathbb{H}^3$ "transverse to the image" $f(\widetilde{\Lambda})$ of $\widetilde{\Lambda}$ under f. For the existence of such we refer to [6, Section 11]. As in Section 4.1.2, let k be an arc transverse to $\widetilde{\Lambda}$

[5] The *foot* of an edge of an ideal triangle is the point of intersection with the orthogonal geodesic from the opposing vertex.

from plaque P to Q. At each endpoint $x_d^{\pm}$ we have two vectors: the ambient vector field v restricted to $x_d^{\pm}$, and the vector n which is normal to the plane containing the plaque R of $\widetilde{S} \setminus \widetilde{\Lambda}$ which contains d. Here the orientation of n is such that the induced orientation of $f(R)$ by f followed by n is the orientation of $\mathbb{H}^3$. Orient these leaves of Λ (thought of here as a leaf of $\widetilde{\Lambda} \subset \mathbb{H}^3$) so that its orientation is from *right to left* with respect to k. The final geometric object we need is the normal plane N to the image $f(g_d^{\pm})$, which inherits an orientation from $f(g_d^{\pm})$ and the orientation of $\mathbb{H}^3$.

Set $a_{n,v}(x_d^{\pm})$ to be the angle from the projection of v onto the normal plane N to $n \in N$. Then we have the following expression for the bending cocycle (see [6, Lemma 36]):

$$\boldsymbol{\beta}(P,Q) = a_{n,v}(x_{d_-}^{+}) - a_{n,v}(x_{d_+}^{-}) + \sum_{d \neq d_+, d_-} [a_{n,v}(x_d^{+}) - a_{n,v}(x_d^{-})] \in \mathbb{R}/2\pi\mathbb{Z}. \tag{4-2}$$

We will use some of the details behind this expression. By [6, Lemmas 4 and 5] there are constants K, A and B, depending only on k, such that the number of components d of $k \setminus \Lambda$ with divergence radius $r(d) = r \in \mathbb{N}$ is at most K, and the length $\ell(d)$ of any such component is bounded by $Be^{-Ar(d)}$. Write the sum in (4-2) as

$$\sum_{d \neq d_+, d_-} [a_{n,v}(x_d^{+}) - a_{n,v}(x_d^{-})] = \sum_{r=0}^{\infty} \sum_{d \neq d_+, d_-;\, r(d)=r} [a_{n,v}(x_d^{+}) - a_{n,v}(x_d^{-})]. \tag{4-3}$$

Since v is Lipschitz, there is a constant $c_0 > 0$ such that

$$|a_{n,v}(x_d^{+}) - a_{n,v}(x_d^{-})| \leqslant c_0 \ell(d) \leqslant c_o B e^{-Ar(d)}.$$

Hence, the tail in the sum (4-3) is estimated by

$$\left| \sum_{r=R}^{\infty} \sum_{d \neq d_+, d_-;\, r(d)=r} [a_{n,v}(x_d^{+}) - a_{n,v}(x_d^{-})] \right| \leqslant c_0 K B \sum_{r=R}^{\infty} e^{-Ar} \leqslant \frac{c_0 K B}{A} e^{-AR}. \tag{4-4}$$

Bonahon proves that given a bending cocycle $\boldsymbol{\beta} \in \mathcal{H}^o(\Lambda, S^1)$ and a hyperbolic surface S, there is an equivariant map $f\colon \widetilde{S} \to \mathbb{H}^3$, well-defined up to isometries, totally geodesic on the plaques and pleated along the lamination Λ, and with bending cocycle $\boldsymbol{\beta}$. The map f is by construction equivariant with respect to some representation whose conjugacy class $[\rho] \in R^o(\Sigma)$ depends only on the isomorphism class of the marked hyperbolic surface S, the lamination Λ, and the bending cocycle $\boldsymbol{\beta}$. Indeed, this construction gives a parametrization of (a portion of) $R^o(\Sigma)$. Set $[\rho] = B_\Lambda(\boldsymbol{\sigma}, \boldsymbol{\beta})$, where $\boldsymbol{\sigma} \in \mathcal{C}(\Lambda)$ is the shearing cocycle of S. Then we have:

Theorem 4.3 [5, Theorem D] *The map (1-5) is a biholomorphism onto an open subset.*

4.1.4 Bending cocycles and geometric bending The following result will be crucial for the analysis later on. It provides a relationship between the bending cocycle discussed here and the geometric bending introduced in Section 3.1.1.

Lemma 4.4 *Fix $\delta > 0$ and some positive integer M. There is $\varepsilon_0 > 0$, depending only on δ, M and Λ, with the following property. Let $f : \widetilde{S} \to \mathbb{H}^3$ be a pleated surface with pleating lamination Λ. Further, given $P, Q \in \mathcal{P}(\Lambda)$ and a transverse arc k from P to Q, suppose k can be written as a union $k_1 \cup \cdots \cup k_m$, where $m \leqslant M$, and for each i the pointed geodesics bounding the plaques intersecting k_i are C^1_ε–close for $\varepsilon \leqslant \varepsilon_0$ at their intersections with k_i. Then there is a finite collection $\{P_i\}_{i=0}^N$ of plaques separating P and Q, with $P_0 = P$ and $P_N = Q$, such that for any choice of points $\tilde{p}_i \in P_i$ in the interiors of the plaques,*

$$\left|\boldsymbol{\beta}(P,Q) - \sum_{i=1}^{N} \Theta_f(\tilde{p}_{i-1}, \tilde{p}_i)\right| < \delta. \tag{4-5}$$

Proof Recall that $\beta(P,Q) = \beta(k)$. By the estimate in (4-4), we may find plaques P_i as in the statement of the lemma so that if we set $d_i = k \cap P_i$, then

$$\left|\beta(k) - \sum_{i=1}^{N} [a_{n,v}(x^+_{d_{i-1}}) - a_{n,v}(x^-_{d_i})]\right| < \tfrac{1}{2}\delta. \tag{4-6}$$

We also assume, after a possible further subdivision, the leaves g^+_{i-1} and g^-_i of Λ through $x^+_{d_{i-1}}$ and $x^-_{d_i}$ are C^1_ε–close for every $i = 1, \ldots, N$ (where ε is to be determined). This does not affect (4-6). It suffices to show that for ε sufficiently small,

$$|a_{n,v}(x^+_{d_{i-1}}) - a_{n,v}(x^-_{d_i}) - \Theta_f(x^+_{d_{i-1}}, x^-_{d_i})| < \frac{\delta}{2N}. \tag{4-7}$$

Here we have extended the definition of $\Theta_f(x^+_{d_{i-1}}, x^-_{d_i})$ from that of $\Theta_f(p_{i-1}, p_i)$ by using the tangent planes to the plaques containing $x^+_{d_{i-1}}$ and $x^-_{d_i}$ to determine the dihedral angles. To simplify notation, set

$$\Delta_i := a_{n,v}(x^+_{d_{i-1}}) - a_{n,v}(x^-_{d_i}) \quad \text{and} \quad \Theta_i := \Theta_f(x^+_{d_{i-1}}, x^-_{d_i}),$$

and let $\mathbb{D}_i$ denote the totally geodesic plane containing the plaque P_i. Equation (4-7) follows from simple estimates in $\mathbb{H}^3$. The idea is that either *both* Δ_i and Θ_i are close to 0, close to π, or the points $x^+_{d_{i-1}}$ and $x^-_{d_i}$ are close to the intersection $\gamma_T := \mathbb{D}_{i-1} \cap \mathbb{D}_i$. If the latter holds, then parallel translation of the vector v to the crease of the tent formed by $\mathbb{D}_{i-1}$ and $\mathbb{D}_i$ only changes v by a small amount, and so the difference Δ_i of angles of the parallel translates is nearly the dihedral angle of the tent.

Step 1 Let y_{i-1} be the endpoint of the geodesic A from $x^-_{d_i}$ to the plane $\mathbb{D}_{i-1}$. Define $a_{n,v}(y_{i-1})$ to be the angle from the projection of v to the normal to $\mathbb{D}_{i-1}$, where the projection is onto the parallel translation of the normal plane to the leaf g^+_{i-1} from $x^+_{d_{i-1}}$ to y_{i-1}. By the hypothesis that g^+_{i-1} and g^-_i are C^1_ε close at their basepoints $x^+_{d_{i-1}}$ and $x^-_{d_i}$, we see that $x^+_{d_{i-1}}$ and y_{i-1} are at most a distance 2ε apart; then since v is continuous, we have that $a_{n,v}(y_{i-1})$ and $a_{n,v}(x^+_{d_{i-1}})$ are equal up to an error comparable to ϵ. Let $|A|$ denote the length of A, and note that the normal to $\mathbb{D}_{i-1}$ is tangent to the segment A at the point y_{i-1}. Let us denote the normals to the planes $\mathbb{D}_{i-1}$ at y_{i-1} and $\mathbb{D}_i$ at $x^-_{d_i}$ by n_{i-1} and n_i, respectively.

Step 2 Let $z_{i-1} \in \mathbb{D}_{i-1}$ be the endpoint of the geodesic segment from $\mathbb{D}_i$ and $\mathbb{D}_{i-1}$, in the case where the planes do not intersect, and when they do intersect $z_{i-1} \in \gamma_T$ is the endpoint of the geodesic from y_{i-1} to γ_T. In either case, let B be the geodesic from y_{i-1} to z_{i-1}. The points $\{y_{i-1}, x_{d_i}^-, z_{i-1}\}$ give a geodesic triangle in $\mathbb{H}^3$ with sides A, B, and a third geodesic C from $x_{d_i}^-$ to z_{i-1} with length $|C|$. Let α, β and $\gamma = \pi/2$ be the corresponding angles of this right-angled geodesic triangle.

Step 3 Suppose that $|\Theta_i| \geqslant \delta/4N$, and $|\Theta_i - \pi| \geqslant \delta/4N$. By Definition 3.2, this means in particular that $\mathbb{D}_i$ and $\mathbb{D}_{i-1}$ intersect along a geodesic γ_T.

Now A is orthogonal to $\mathbb{D}_{i-1}$, so its parallel translate along B is orthogonal to γ_T at $B \cap \gamma_T$. As B is also orthogonal there to γ_T, we see that γ_T meets the triangle ABC orthogonally, and hence C also meets γ_T orthogonally. Thus, $\alpha = \Theta_i - \pi$, so the assumption implies $|\alpha| \geqslant \delta/4N$. By the hyperbolic law of sines,

$$\sinh|B| = \sinh|A| \cdot \frac{\sin\beta}{\sin\alpha},$$

which implies that $|B|$ and $|C|$ are of the order of $|A| = O(\varepsilon)$. Thus, γ_T is within $O(\varepsilon)$ of the points y_{i-1} and $x_{d_i}^-$. Moreover, since the dihedral angle is bounded away from 0 and π, γ_T must be C_ε^1–close to the leaves g_{i-1}^+ and g_i^-. In particular, the normal planes to all three are close. The quantity Δ_i can then be computed by parallel translation along B and C. By Lemma 3.10, it follows that Δ_i and Θ_i are close; in particular, less than $\delta/2N$ for ε sufficiently small.

Step 4 Suppose that $|\Delta_i| \geqslant \delta/4N$, and $|\Delta_i - \pi| \geqslant \delta/4N$. Then β is bounded away from $\pi/2$. For a general right-angled hyperbolic triangle one has the relation

$$\tanh|B| = \sinh|A| \cdot \tan\beta.$$

Since $\cos\beta$ is bounded away from zero, one arrives at an estimate of the form: $\tanh|B| \leqslant c_0 \sinh|A|$ for some constant c_0 depending on this bound. It again follows that γ_T is close to the points $x_{d_{i-1}}^+$ to $x_{d_i}^-$, and therefore arguing as in the previous step, Δ_i and Θ_i are close.

Step 5 Suppose that neither of the assumptions of Step 3 or 4 hold. If $|\Delta_i - \pi| < \delta/4N$, for example, then since v is continuous it follows that the orientations of $\mathbb{D}_{i-1}$ and $\mathbb{D}_i$ are compatible. Since the assumption of Step 3 fails, this forces $|\Theta_i - \pi| < \delta/4N$. A similar argument holds if $|\Delta_i| < \delta/4N$, and so in either case $|\Delta_i - \Theta_i| < \delta/2N$. □

4.2 Cocycles and Prym differentials

In this section we relate the notion of a bending cocycle to the spectral data parametrization of Higgs bundles discussed in Section 2.2. Let $q \in \mathcal{QD}^*(X)$, and choose any maximal geodesic lamination Λ containing Λ_q^h as a sublamination. Thus, if the horizontal foliation of q has no saddle connections, then $\Lambda = \Lambda_q^h$. Let $\widehat{X}_q \to X$ be the spectral curve associated to q. Recall that the Prym variety $\mathrm{Prym}(\widehat{X}_q, X)$ contains $J_2(X)$ as a subgroup. The goal is to prove the following.

Theorem 4.5 *There is a group isomorphism,*

$$\mathcal{H}^o(\Lambda, S^1) \simeq \mathrm{Prym}(\widehat{X}_q, X)/J_2(X).$$

This result is essentially contained in [54, Section 3.2] and [5, page 13]. The idea is to view a bending cocycle in terms of periods of a Prym differential. The choice of sign is fixed by a choice of orientation of the lift of the lamination on the spectral curve. Since this correspondence is so central to the present paper, we present the details below.

As mentioned in the introduction, there is a nice interpretation of Theorem 4.5, which goes as follows: the space $\mathrm{Prym}(\widehat{X}_q, X)/J_2(X)$ is the fiber over q of the Hitchin fibration for (a component of) the moduli space of $\mathrm{PSL}(2,\mathbb{C})$–Higgs bundles, whereas $\mathcal{H}^o(\Lambda, S^1)$ is a torus fiber over $\mathcal{C}(\Lambda)$ in Bonahon's parametrization of the character variety $R(\Sigma)$. Via the nonabelian Hodge correspondence, $\mathcal{QD}(X) \simeq \mathcal{C}(\Lambda)$, and the moduli space of Higgs bundles is homeomorphic to $R(\Sigma)$.

4.2.1 Homology of branched covers Here we digress to make precise the construction of a homology basis for the spectral curve. Consider the general case of a closed, oriented surface Σ. Suppose $p\colon \widehat{\Sigma} \to \Sigma$ is a connected branched double cover of Σ with branching set B and involution σ, and let p_* denote the induced map $H_1(\widehat{\Sigma}) \to H_1(\Sigma)$ on homology. Let $g, \widehat{g}$ be the genera of $\Sigma, \widehat{\Sigma}$. Recall by the Hurwitz formula that $2\widehat{g} = 2g + (2g + \#B - 2)$, where we have split the sum to indicate the dimensions of the even and odd homology of $\widehat{\Sigma}$ under the involution σ.

Proposition 4.6 *There is an exact sequence*

$$0 \to \mathbb{Z} \to H_1(\Sigma, B) \xrightarrow{\phi} H_1(\widehat{\Sigma}) \xrightarrow{p_*} H_1(\Sigma) \to 0, \tag{4-8}$$

where the map ϕ is surjective onto the odd homology of $\widehat{\Sigma}$.

Proof A topological model for the branched cover is given by decomposing B into pairs and introducing branch cuts. In this setting, generators of the homology of $\widehat{\Sigma}$ are given as follows. First, choose generators $c_1, \dots, c_{2g}$ for $H_1(\Sigma)$. Let $\widehat{c}_1, \dots, \widehat{c}_{2g}$ be lifts in $H_1(\widehat{\Sigma})$, ie $p_*(\widehat{c}_k) = c_k$. We set $\phi(c_k) = \widehat{c}_k - \sigma(\widehat{c}_k)$. Set $N = \#B/2$. Now choose generators $\{a_i, b_j\}$ for $i = 1, \dots, N$ and $j = 1, \dots, N-1$ of $H_1(\Sigma, B)$ as in the diagram below. Define closed curves $\widehat{a}_i$, $\widehat{b}_i$ on $\widehat{\Sigma}$ as follows: choose lifts $\widehat{\alpha}_i$, $\widehat{\beta}_i$ of a_i, b_i, and set $\widehat{a}_i = \widehat{\alpha}_i - \sigma(\widehat{\alpha}_i)$, $\widehat{b}_i = \widehat{\beta}_i - \sigma(\widehat{\beta}_i)$. Then $\widehat{a}_i = \phi(a_i)$, $\widehat{b}_j = \phi(b_j)$. With the orientation indicated, there is a single relation: $\sum_{i=1}^N \widehat{a}_i = 0$. The collection $\{\widehat{a}_i, \widehat{b}_j, \widehat{c}_k, \sigma(\widehat{c}_k)\}$ generate $H_1(\widehat{\Sigma})$. See Figure 4.

Note that with the appropriate choice of orientations we have the following intersection numbers:

$$\widehat{a}_i \cdot \widehat{a}_j = \widehat{b}_i \cdot \widehat{b}_j = 0, \qquad \widehat{a}_i \cdot \widehat{b}_i = +1, \qquad \widehat{a}_{i+1} \cdot \widehat{b}_i = -1, \qquad \widehat{a}_i \cdot \widehat{b}_j = 0 \quad \text{otherwise.}$$

These are compatible with the relation. The intersection numbers of the $\{\widehat{c}_k\}$ and $\{\sigma(\widehat{c}_k)\}$ are the same as those of $\{c_k\}$ on Σ, with the additional relations

$$\widehat{a}_i \cdot \widehat{c}_k = \widehat{b}_j \cdot \widehat{c}_k = \widehat{c}_k \cdot \sigma(\widehat{c}_\ell) = 0 \quad \text{for all } i, j, k, \ell. \tag{4-9}$$

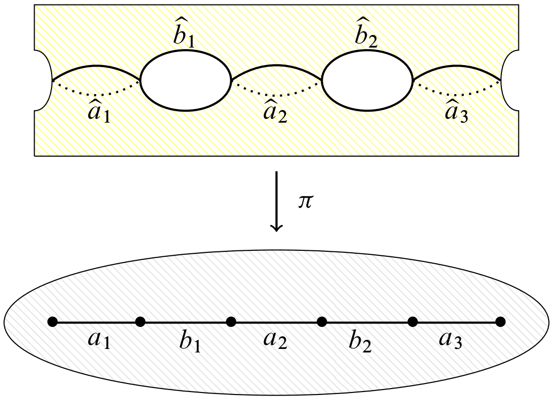

Figure 4: Branched surface.

By construction, $p_* \circ \phi = 0$. We show that ϕ is surjective onto the odd homology, which will prove exactness of the second part of the sequence. Indeed, for $\widehat{\gamma} \in H_1(\widehat{\Sigma})$, write

$$\widehat{\gamma} = \sum_{i=1}^{N} r_i \widehat{a}_i + \sum_{j=1}^{N-1} s_j \widehat{b}_j + \sum_{k=1}^{2g} m_k \widehat{c}_k + \sum_{k=1}^{2g} n_k \sigma(\widehat{c}_k).$$

If $\widehat{\gamma}$ is odd then since $\widehat{a}_i$ and $\widehat{b}_j$ are also odd, we have

$$-\widehat{\gamma} = \sigma(\widehat{\gamma}) = -\sum_{i=1}^{N} r_i \widehat{a}_i - \sum_{j=1}^{N-1} s_j \widehat{b}_j + \sum_{k=1}^{2g} m_k \sigma(\widehat{c}_k) + \sum_{k=1}^{2g} n_k \widehat{c}_k.$$

So $n_k = -m_k$, and

$$\widehat{\gamma} = \sum_{i=1}^{N} r_i \widehat{a}_i + \sum_{j=1}^{N-1} s_j \widehat{b}_j + \sum_{k=1}^{2g} m_k (\widehat{c}_k - \sigma(\widehat{c}_k)) = \phi\Big(\sum_{i=1}^{N} r_i a_i + \sum_{j=1}^{N-1} s_j b_j + \sum_{k=1}^{2g} m_k c_k\Big). \tag{4-10}$$

Let us verify that $\ker\phi \simeq \mathbb{Z}$. Suppose $\gamma \in \ker\phi$. We can write

$$\gamma = \sum_{i=1}^{N} r_i a_i + \sum_{j=1}^{N-1} s_j b_j + \sum_{k=1}^{2g} m_k c_k,$$

$$0 = \phi(\gamma) = \sum_{i=1}^{N} r_i \widehat{a}_i + \sum_{j=1}^{N-1} s_j \widehat{b}_j + \sum_{k=1}^{2g} m_k (\widehat{c}_k - \sigma(\widehat{c}_k)).$$

Taking intersections with appropriate elements $\widehat{c}_k$ and $\sigma(\widehat{c}_k)$, and using (4-9), it is easy to see that $m_k = 0$ for all k. Now

$$0 = \phi(\gamma) \cdot \widehat{a}_i = \sum_{j=1}^{N-1} s_j \widehat{b}_j \cdot \widehat{a}_i \implies 0 = s_{i-1} - s_i,$$

which implies $s_j = 0$ for all j. Similarly,

$$0 = \phi(\gamma) \cdot \widehat{b}_j = \sum_{i=1}^{N} r_i \widehat{a}_i \cdot b_j = r_j - r_{j+1},$$

which implies r_j is a fixed constant for all j. Hence, γ is a multiple of the class $\sum_{i=1}^{N} a_i$. This completes the proof of (4-8). □

Remark 4.7 The map ϕ is not canonically determined but depends rather on the choices of lifts of the cycles a_i, b_i, c_i.

4.2.2 Periods of Prym differentials Let us ignore the hyperbolic structure and consider $\Lambda \subset \Sigma$ (recall the discussion at the end of Section 2.4.1). Because Λ is maximal, if we choose a collection B of points, one in each component of $\Sigma \setminus \Lambda$, we may define a double cover $\widehat{\Sigma} \to \Sigma$ branched at B. The preimage $\widehat{\Lambda}$ is now orientable and we fix such once and for all. Recall the result of the previous section. In this case, the map ϕ in (4-8) is actually determined uniquely; see Remark 4.7. Indeed, we may assume representatives for the cycles a_i, b_i, c_i are transverse to Λ. Then choose the lifts to $\widehat{\Sigma}$ to be positively oriented with respect to the orientation of $\widehat{\Lambda}$. This determines a choice of lifts: the only possible ambiguity would be the existence of cycles not meeting Λ, but this is ruled out by maximality.

Notice that we have an identification of the π_1–orbits of plaques with the set B; let us denote this π_1–invariant map by $p\colon \mathcal{P}(\Lambda) \to B$. Let $\widetilde{\gamma}$ be a transverse C^1 curve from plaques P and Q in $\widetilde{\Sigma}$, and let γ be the projection from corresponding points p to q in B. As in the previous paragraph, there is a unique lift $\widehat{\gamma}$ that is positively oriented transverse to $\widehat{\Lambda}$, ie the lamination is oriented to the left at a point of intersection of $\widehat{\gamma} \cap \widehat{\Lambda}$. Now given a closed one-form $\widehat{\alpha}$, with $[\widehat{\alpha}] \in H^1_{\mathrm{odd}}(\widehat{\Sigma}, \mathbb{R})$, define

$$\sigma_{\widehat{\alpha}}(P, Q) := 2 \int_{\widehat{\gamma}} \widehat{\alpha}. \tag{4-11}$$

The factor of 2 is added here for convenience; see Remark 4.9 below.

We first note that $\sigma_{\widehat{\alpha}}$ is well-defined. First, it is independent of the choice of $\widetilde{\gamma}$; for a relative homotopy of $\widetilde{\gamma}$ induces one on γ, and therefore $\widehat{\gamma}$, and this does not affect the integral of the closed form $\widehat{\alpha}$. Second, it is independent of choice of representative $\widehat{\alpha}$. Any other choice can be written as $\widehat{\alpha} + df$, for an odd $\mathbb{R}$–valued function f, and since the endpoints of $\widehat{\gamma}$ lie on the fixed-point set of σ, this contributes nothing to the integral.

With this understood, we prove the following:

Proposition 4.8 *The function $\sigma_{\widehat{\alpha}}$ in (4-11) defines a transverse cocycle depending only on the class of $\widehat{\alpha}$.*

Proof Equivariance is clear, since the path $\widetilde{\gamma}$ and $g\widetilde{\gamma}$ define the same curve γ on X. If the plaque R separates P and Q, let $\widetilde{\gamma}\colon [a, b] \to \mathbb{H}$, with $\widetilde{\gamma}(a) = P$ and $\widetilde{\gamma}(b) = Q$. Then there exist $a < t_1 < t_2 < b$ such that $\widetilde{\gamma}(t_1), \widetilde{\gamma}(t_2) \in \partial R$, and $\widetilde{\gamma}[a, t_1) \cap R = \widetilde{\gamma}(t_2, b] \cap R = \varnothing$. After a homotopy, we may assume $\widetilde{\gamma}(t_1, t_2) \subset R$, and after a further homotopy we may assume there is $t_1 < c < t_2$ such that $\widetilde{\gamma}(c)$ is the point in B associated to R. It follows that $\widetilde{\gamma}$ can be written as a sum of quasitransverse paths from P to R, and R to Q. The additivity then follows from the additivity of the integral in (4-11). It remains to prove symmetry. Let $\overline{\gamma}$ denote the curve γ with the reverse orientation. Then we note that

$$\widehat{\overline{\gamma}} = \overline{\sigma(\widehat{\gamma})}. \tag{4-12}$$

Indeed, $\sigma(\hat{\gamma})$ is negatively oriented with respect to $\hat{\Lambda}$, and so both sides of (4-12) are positively oriented lifts of $\bar{\gamma}$. Using (4-12), we have

$$\tag{4-13} \tfrac{1}{2}\boldsymbol{\sigma}_{\hat{\alpha}}(Q,P) = \int_{\hat{\bar{\gamma}}} \hat{\alpha} = \int_{\overline{\sigma(\hat{\gamma})}} \hat{\alpha} = -\int_{\sigma(\hat{\gamma})} \hat{\alpha} = -\int_{\hat{\gamma}} \sigma^*\hat{\alpha} = \int_{\hat{\gamma}} \hat{\alpha} = \tfrac{1}{2}\boldsymbol{\sigma}_{\hat{\alpha}}(P,Q).$$

This completes the proof. □

Remark 4.9 Note that $\boldsymbol{\sigma}_{\hat{\alpha}}(P,Q)$ is equal to a period of the differential $\hat{\alpha}$. Indeed, from the discussion above, the curve $\hat{\gamma}_{PQ} = \hat{\gamma} \cup \hat{\bar{\gamma}}$ is a closed oriented curve on $\hat{X}_q$, and by (4-13),

$$2\boldsymbol{\sigma}_{\hat{\alpha}}(P,Q) = \boldsymbol{\sigma}_{\hat{\alpha}}(P,Q) + \boldsymbol{\sigma}_{\hat{\alpha}}(Q,P) = 2\int_{\hat{\gamma}} \hat{\alpha} + 2\int_{\hat{\bar{\gamma}}} \hat{\alpha} = 2\int_{\hat{\gamma}_{PQ}} \hat{\alpha}.$$

Conversely, by Proposition 4.6, every element of $H_1^{\mathrm{odd}}(\hat{\Sigma},\mathbb{Z})$ is represented by a linear combination of oriented curves of the form $\hat{\gamma}_{PQ}$, for some lifts P, Q of some points $p,q \in B$. It follows that the periods of $\hat{\alpha}$, and hence $[\hat{\alpha}]$ itself, are determined by $\boldsymbol{\sigma}_{\hat{\alpha}}$.

We now return to the case where Σ has a Riemann surface structure X and the lamination comes from a holomorphic quadratic differential.

Example 4.10 Let $q \in \mathcal{QD}^*(X)$, and let Λ be a maximalization of Λ_q^h in the sense of Lemma 2.26. The Seiberg–Witten differential λ_{SW} from (2-22) is a holomorphic Prym differential on $\hat{X}_q$. We can orient the lift $\hat{\Lambda}_q^h$ by the condition $\operatorname{Re}\lambda_{\mathrm{SW}} > 0$. The harmonic Prym differential $\operatorname{Re}\lambda_{\mathrm{SW}}$ defines a canonical transverse cocycle $\boldsymbol{\sigma}_q^{\mathrm{can}}$. By the previous remark, $\boldsymbol{\sigma}_q^{\mathrm{can}}$ is determined by the real parts of the periods of (1-10).

By Proposition 4.8, we have a map

$$\tag{4-14} T\colon H^1_{\mathrm{odd}}(\hat{X}_q,\mathbb{R}) \to \mathcal{H}(\Lambda,\mathbb{R}), \quad [\hat{\alpha}] \mapsto \boldsymbol{\sigma}_{\hat{\alpha}}.$$

We can do a similar construction for bending cocycles. If $[\hat{\eta}] \in H^1_{\mathrm{odd}}(\hat{X}_q, i\mathbb{R})$, set

$$\tag{4-15} \boldsymbol{\beta}_{\hat{\eta}}(P,Q) := -2i\int_{\hat{\gamma}} \hat{\eta} \mod 2\pi.$$

By Remark 4.9, $\boldsymbol{\beta}_{\hat{\eta}}$ only depends on the class of $[\hat{\eta}]$ modulo the lattice

$$\tag{4-16} L = L(\hat{X}_q) := \operatorname{Hom}(H_1^{\mathrm{odd}}(\hat{X}_q,\mathbb{Z}), 2\pi i\mathbb{Z}).$$

Hence, we have a map

$$\tag{4-17} B\colon H^1_{\mathrm{odd}}(\hat{X}_q, i\mathbb{R})/L \to \mathcal{H}^o(\Lambda,\mathbb{R}/2\pi\mathbb{Z}), \quad [\hat{\eta}] \mapsto \boldsymbol{\beta}_{\hat{\eta}}.$$

Clearly, T is linear. By Remark 4.9, it is also injective. For if $\boldsymbol{\sigma}_{\hat{\eta}} \equiv 0$, then the periods of $\hat{\eta}$ must all vanish; hence, $[\hat{\eta}] = 0$. By [5, Proposition 1], the dimensions of the two sides of (4-14) agree. In the case of the map B, notice that the lattices on either side are isomorphic under the map T. This proves the following result.

Corollary 4.11 *The maps T and B in (4-14) and (4-17) are isomorphisms.*

We observe the following:

Lemma 4.12 *The inclusion induces an exact sequence*:

$$0 \to H^1_{\mathrm{odd}}(\widehat{X}_q, 2\pi i\mathbb{Z}) \to L \to J_2(X) \to 0.$$

Proof This is most easily seen in terms of the explicit generators in Section 4.2.1. Let $\widehat{\alpha}_j$, $\widehat{\beta}_j$ and $\widehat{\delta}_j$ be Poincaré duals of $\widehat{a}_j$, $\widehat{b}_j$ and $\widehat{c}_j$, respectively. Then $\delta_j := \pi i(\widehat{\delta}_j - \sigma^*\widehat{\delta}_j) \in \Lambda$ for $j = 1, \dots, 2g$ is not $2\pi i$–integral, but $2\delta_j \in H^1_{\mathrm{odd}}(\widehat{X}_q, 2\pi i\mathbb{Z})$. It is easily seen that L is generated by $H^1_{\mathrm{odd}}(\widehat{X}_q, 2\pi i\mathbb{Z})$ and all such elements δ_j. This proves the result. $\square$

Proof of Theorem 4.5 Immediate from Corollary 4.11, Lemma 4.12 and equation (2-20). $\square$

Definition 4.13 The complex cocycle $\Gamma_{\mathrm{SW}} \in \mathcal{H}^o(\Lambda, \mathbb{R} + i\mathbb{R}/2\pi\mathbb{Z})$ is the one determined as in Corollary 4.11 by the periods of the real and imaginary parts of λ_{SW}.

4.3 Approximation by pleated surfaces

In this section, we define what it means for a family of harmonic maps and a family of pleated surfaces to be asymptotic, and we relate the corresponding notions of bending. The intuition behind the definition below may be summarized as follows: the image of a pleated surface $f : \widetilde{S} \to \mathbb{H}^3$, for a representation outside a large compact set, consists of a configuration of plaques sheared far apart from one another and related by long leaves of the lamination. At the same time, the image of a harmonic map in a neighborhood of the zeroes of a quadratic differential is nearly planar, whereas leaves of the horizontal foliation are nearly geodesic. The approximation requires these planes and approximate geodesics to be close to the plaques and leaves of the lamination of the pleated surface.

We furthermore assume that we have chosen maximal laminations Λ_n (resp. Λ) containing $\Lambda^h_{q_n}$ (resp. Λ^h_q), and that $\Lambda_n \to \Lambda$ in the Hausdorff sense.

If there exists a pleated surface $\mathsf{P}_n = (S_n, f_n, \Lambda_n, \rho_n)$, then by the discussion in Section 2.4 there is a bijective correspondence between the zeroes $\widetilde{Z}(q)$ and plaques of $\widetilde{\Lambda}_n$, and each bi-infinite leaf in $\mathcal{F}^h_{q_n}$ determines a leaf in $\Lambda^h_{q_n}$. The choice of maximalization Λ_n is determined by a finite choice of "additional" leaves; see Section 2.4.3. We also recall from Section 2.5.2 the definition of the hexagonal sets $\mathcal{Q}_n$ and $\widetilde{\mathcal{Q}}_n$. With this understood, we make the following definition.

Definition 4.14 A sequence of pleated surfaces $\mathsf{P}_n = (S_n, f_n, \Lambda_n, \rho_n)$ is *asymptotic to u_n* if for any $\varepsilon > 0$ there is N so that if $n \geqslant N$, the following holds.

(i) The images by u_n of the horizontal leaves in $\widetilde{X} \setminus \widetilde{\mathcal{Q}}_n$ are C^1_ε–close to the corresponding leaves in $\widetilde{\Lambda}_n$.

(ii) If $\widetilde{p} \in \widetilde{Z}(q_n)$, then $u_n(\widetilde{p})$ is ε–close to the image by f_n of the corresponding plaque P in $\widetilde{S}_n$. Moreover, the parallel translation of the tangent plane to the image of u_n at $u_n(\widetilde{p})$ along the geodesic to $f_n(P)$ makes an angle less than ε with the totally geodesic subspace containing $f_n(P)$.

With this definition we are in a position to compare the notion of bending for sequences of harmonic maps and of pleated surfaces that are asymptotic to each other.

Proposition 4.15 *Let $\mathsf{P}_n = (S_n, f_n, \Lambda_n, \rho_n)$ be a sequence of pleated surfaces that is asymptotic to the sequence $u_n\colon \widetilde{X} \to \mathbb{H}^3$ of ρ_n–equivariant harmonic maps in the sense of Definition 4.14. Denote by $\boldsymbol{\beta}_n \in \mathcal{H}^o(\Lambda_n, S^1)$ the bending cocycles of P_n. Fix $\delta > 0$. Then for any $P, Q \in \mathcal{P}(\Lambda)$, there are plaques $\{P_i\}_{i=0}^N$ between P and Q, with $P_0 = P$ and $P_N = Q$, with centers $\widetilde{p}_i$, such that*

$$\lim_{n\to\infty}\left(\boldsymbol{\beta}_n(r^{\Lambda}_{\Lambda_n}(P), r^{\Lambda}_{\Lambda_n}(Q)) - \sum_{i=1}^{N} \Theta_{u_n}(\widetilde{p}_{i-1}, \widetilde{p}_i)\right) \leqslant \delta.$$

Proof We shall use the setup of Lemma 4.4. By the convergence $\Lambda_n \to \Lambda$, the approximation of the bending cocycle by sums over finitely many plaques is uniform. By a further subdivision, we may assume that between any two centers $\widetilde{p}_{i-1}^{(n)} \to \widetilde{p}_{i-1}$ and $\widetilde{p}_i^{(n)} \to \widetilde{p}_i$ there are quasitransverse arcs (or modified saddle connections) $k_i^{(n)}$ with small vertical ends that meet the zeroes of the Hopf differentials q_n only at $\widetilde{p}_{i-1}^{(n)}$ and $\widetilde{p}_i^{(n)}$. Then the images by u_n of the horizontal parts of $k_i^{(n)}$ are C^1_ε–close, and therefore by the asymptotic assumption the same is true for the leaves of Λ_n along $k_i^{(n)}$. Thus, the hypotheses of Lemma 4.4 are satisfied for sufficiently large n, and we have

$$\left|\boldsymbol{\beta}_n(r^{\Lambda}_{\Lambda_n}(P), r^{\Lambda}_{\Lambda_n}(Q)) - \sum_{i=1}^{N} \Theta_{f_n}(\widetilde{p}_{i-1}^{(n)}, \widetilde{p}_i^{(n)})\right| < \tfrac{1}{2}\delta \tag{4-18}$$

for large enough n. On the other hand, an argument analogous to the one used in the proof of that lemma shows that

$$|\Theta_{f_n}(\widetilde{p}_{i-1}^{(n)}, \widetilde{p}_i^{(n)}) - \Theta_{u_n}(\widetilde{p}_{i-1}^{(n)}, \widetilde{p}_i^{(n)})| \leqslant \frac{\delta}{2N} \tag{4-19}$$

for large n. Indeed, suppose that not both $\Theta_{f_n}(\widetilde{p}_{i-1}^{(n)}, \widetilde{p}_i^{(n)})$ and $\Theta_{u_n}(\widetilde{p}_{i-1}^{(n)}, \widetilde{p}_i^{(n)})$ are within $\delta/4N$ of π, and neither are they both within $\delta/4N$ of 0. Then the image by u_n of the horizontal parts of $k_i^{(n)}$ is arbitrarily close to the crease of the tent formed by the totally geodesic planes associated to the plaques P_{i-1} and P_i. By Definition 4.14(ii), these are also close to the planes tangent to the image of u_n at p_{i-1} and p_i. The angle $\Theta_{f_n}(\widetilde{p}_{i-1}^{(n)}, \widetilde{p}_i^{(n)})$ can be computed by parallel translation of the normal vectors to the plaques, as discussed after Definition 3.1. These normal vectors are close to the normal vectors to the planes defined by u_n. By Lemma 3.10, the parallel translations along $u_n(k_i^{(n)})$ are also close to the parallel translations along the crease. Note that in the statement of that lemma, the term $L\varepsilon$ is small, since ε is exponentially small compared to the length of $u_n(k_i^{(n)})$ by Proposition 2.29.

Combining (4-19) with (4-18),

$$\left|\boldsymbol{\beta}_n(r^{\Lambda}_{\Lambda_n}(P), r^{\Lambda}_{\Lambda_n}(Q)) - \sum_{i=1}^{N} \Theta_{u_n}(\widetilde{p}_{i-1}^{(n)}, \widetilde{p}_i^{(n)})\right| < \delta.$$

Since this holds for fixed N and δ, and any sufficiently large n, this completes the proof. □

4.3.1 Bending cocycles and periods We now combine the considerations above with the results of Section 3.

Theorem 4.16 *Let $\mathsf{P}_n = (S_n, f_n, \Lambda_n, \rho_n)$ be a sequence of pleated surfaces with bending cocycles $\boldsymbol{\beta}_n$. We assume the following two conditions:*

(i) *The sequence P_n is asymptotic to the sequence $u_n : \widetilde{X} \to \mathbb{H}^3$ of ρ_n–equivariant harmonic maps in the sense of Definition 4.14.*

(ii) *The sequence (A_n, Ψ_n) of Higgs pairs for ρ_n converges to a limiting configuration with associated Prym differential $\widehat{\eta}$.*

Let $\boldsymbol{\beta}_{\widehat{\eta}}$ be defined as in (4-15). Then in $\mathcal{H}^o(\Lambda, S^1)$, $\lim_{n\to\infty} \boldsymbol{\beta}_n = \boldsymbol{\beta}_{\widehat{\eta}}$.

Proof Let $\boldsymbol{\beta}$ be any subsequential limit of $\boldsymbol{\beta}_n$. By Proposition 4.15 it suffices to estimate the geometric bending $\Theta_{u_n}(\widetilde{p}_{i-1}, \widetilde{p}_i)$. Using the assumption in the proof of that result, we have a quasitransverse path k_i from p_{i-1} to p_i that intersects the zeroes of q only at the endpoints. By Theorem 3.11, it follows that $\boldsymbol{\beta}_n(P, Q)$ is approximated by the sum of $\Theta_{k_i}(A_n, \Psi_n)$. Since the latter is additive, $\boldsymbol{\beta}_n(P, Q)$ is approximated by $\Theta_k(A_n, \Psi_n)$, where k is the image of a path from $\widetilde{p}$ to $\widetilde{q}$. By Proposition 3.6, this converges as $n \to \infty$ to the period of $\widehat{\eta}$. □

5 Realization of pleated surfaces

The goal of this section is to prove the following result, which is part (i) of the Main Theorem.

Theorem 5.1 *Let $[\rho_n] \in R^o(\Sigma)$ be a divergent sequence, $u_n : \widetilde{X} \to \mathbb{H}^3$ the ρ_n–equivariant harmonic maps, and $t_n^2 q_n$ the Hopf differentials of u_n, where $q_n \in \mathcal{SQD}^*(X)$ and $t_n \to +\infty$. We assume that $q_n \to q \in \mathcal{SQD}^*(X)$, and in some (hence any) realization of the associated geodesic laminations, choose maximal laminations Λ_n (resp. Λ) containing $\Lambda^h_{q_n}$ (resp. Λ^h_q), with $\Lambda_n \to \Lambda$. Then there is an N such that for all $n \geqslant N$, the class $[\rho_n]$ is in the image of the map B_{Λ_n} in (1-5), ie there is a pleated surface $\mathsf{P}_n = (S_n, f_n, \Lambda_n, \rho_n)$.*

5.1 Realizing laminations

Definition 5.2 (cf [8, Definition I.5.3.4]) Let $\Lambda \subset S$ be a geodesic lamination and $\rho : \pi_1(S) \to \mathrm{PSL}(2, \mathbb{C})$. Then Λ is *realizable* if there exists a continuous ρ–equivariant map $\varphi : \widetilde{S} \to \mathbb{H}^3$ that takes the leaves of $\widetilde{\Lambda}$ homeomorphically onto geodesics in $\mathbb{H}^3$.

The goal of this subsection is to prove the following.

Proposition 5.3 *Assume the hypotheses of Theorem 5.1. Then for n sufficiently large, there is a marked hyperbolic surface $\widehat{S}_n$ such that the geodesic lamination $\Lambda_n \subset \widehat{S}_n$ is realizable for ρ_n.*

The proof of Proposition 5.3 will proceed by using the result of Minsky on high-energy harmonic maps into $\mathbb{H}^3$ that we summarized in Section 2.5.1. This will lead to a suitable collection of train tracks carrying the laminations Λ_n.

5.1.1 The companion surface By [58, Theorem 3.1] and [29, Theorem 11.2] there is a marked hyperbolic surface $\widehat{S}_n$ such that the harmonic diffeomorphism $v_n\colon X \to \widehat{S}_n$ has Hopf differential $t_n^2 q_n$. Moreover, the class $[\widehat{S}_n] \in T(\Sigma)$ is uniquely determined by $t_n^2 q_n$. We let $\tilde{v}_n\colon \tilde{X} \to \mathbb{H}^2$ denote the lift to the universal cover. Via v_n, the laminations Λ_n are realized as geodesic laminations in $\widehat{S}_n$. As previously, we continue to use the notation $\Lambda_n \subset \widehat{S}_n$ to simplify the notation. Also, denote the lift by $\widetilde{\Lambda}_n \subset \mathbb{H}^2$.

Proposition 5.4 [47, Theorem 7.1] *For A, c_0 and C_0 as in Proposition 2.28, and $n \geqslant N$ and $s_n \leqslant c_0$, there is a π_1–equivariant map Π_* from the leaves of $\widetilde{\mathcal{F}}^h_{q_n}$ in the complement of $\mathcal{Q}_n$ to the leaves of $\widetilde{\Lambda}^h_n \subset \mathbb{H}^2$ which factors through v_n. Moreover, for any $p \in \tilde{X} \setminus \widetilde{\mathcal{Q}}_n$,*

$$d_{\mathbb{H}^2}(v_n(p), \Pi_*(p)) \leqslant A\exp(-t_n C_0),$$

and the derivative along the horizontal leaf through p (in the $|q_n|$ metric) is

$$\big||d\Pi_*| - 2\big| \leqslant A\exp(-t_n C_0).$$

We note that in the case of a maximalization Λ_n of Λ^h_n, the map Π_* can be extended to the additional leaves as remarked in Section 2.5.1.

The train track used in the proof of Proposition 2.28 may be chosen so that for n sufficiently large the following holds.

(i) Let $\tilde{\tau}_{n,*} = \tilde{v}_n(\tilde{\tau}_n) \subset \mathbb{H}^2$. Then the branches of $\tilde{\tau}_{n,*}$ have length comparable to t_n and geodesic curvature $O(\varepsilon_n)$ and meet tangentially.

(ii) The collection $\widetilde{\Lambda}^h_n$ is $C^1_{\varepsilon_n}$–carried by $\tilde{\tau}_{n,*}$.

Let $\hat{\boldsymbol{\sigma}}_n$ denote the shearing cocycle of $\widehat{S}_n$ with respect to the lamination Λ_n. We will need the following result from [58]. (Stronger estimates are implicit in [18].)

Lemma 5.5 *For any $\delta > 0$ there is an N such that for all $n \geqslant N$,*

$$\|\hat{\boldsymbol{\sigma}}_n - t_n \boldsymbol{\sigma}^{\mathrm{can}}_{q_n}\| < \delta t_n.$$

5.1.2 Proof of Proposition 5.3 We first choose a constant $\delta > 0$ so that there are disjoint arcs c_i, one for each branch b_i of τ_n, such that c_i intersects only b_i, and this only once. The endpoints of c_i lie in exactly two components of $X \setminus \tau_n$; see Figure 5. Viewed on $\widehat{S}_n$, we may assume that the endpoints of c_i are in $\widehat{S}_n \setminus N_{2\delta}$, where N_r is the r neighborhood of $\tau_{n,*}$. We furthermore assume $\Lambda_n \subset N_\delta$. These assumptions are made possible by Proposition 5.4.

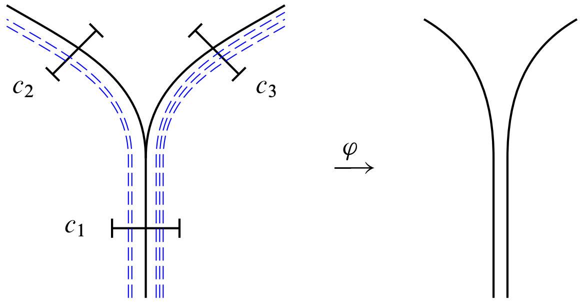

Figure 5: Realization of Λ_n.

Let g be a leaf of $\widetilde{\Lambda}_n \subset \mathbb{H}^2$. Then $u_n \circ \widetilde{v}_n^{-1}$ produces a well-defined geodesic $g^* \subset \mathbb{H}^3$. Namely, if $\ell \subset \widetilde{\mathcal{F}}^h_{q_n}$ follows a train path such that the straightening of $\widetilde{v}_n(\ell)$ is g, then g^* is the straightening of $u_n(\ell)$. For every intersection point p in $g \cap \widetilde{c}_i$, we map p to the point $\varphi(p) = p^*$ given by the nearest-point projection of $u_n \circ \widetilde{v}_n^{-1}(p)$ onto g^*. Let p_1 and p_2 be consecutive points on g, in the sense that there is no other point in $g \cap \widetilde{c}_i$ in the geodesic segment $g_{p_1 p_2}$. Extend the map along the segment

$$\varphi\colon g_{p_1 p_2} \xrightarrow{\sim} g^*_{p_1^* p_2^*}$$

as a homothety. Continuing in this way we obtain a continuous map $\varphi\colon \widetilde{\Lambda}_n \to \mathbb{H}^3$ mapping leaves homeomorphically to geodesics. Moreover, it is clearly equivariant. Since $\widetilde{\Lambda}_n$ is a closed subset, by the Tietze extension theorem we can extend φ to a continuous map $\widetilde{\widetilde{N}}_\delta$ with the same Lipschitz constant. Now we use a geodesic homotopy to join $u_n \circ \widetilde{v}_n^{-1}$ on the complement of $\widetilde{N}_{2\delta}$ to this extension. This defines the map φ, and it is equivariant.

5.2 Perturbing the companion surface

For closed 3–manifolds the existence of a realization of a lamination leads to a pleated surface; see [8, Theorem I.5.3.9]. The goal of this section is to prove the same in the equivariant case that we consider. The rough idea is that the companion surfaces $\widehat{S}_n$ obey the same asymptotics as the image of the equivariant harmonic maps u_n, so the hyperbolic structure on the putative pleated surface should be obtained from a small perturbation of that on $\widehat{S}_n$. For a similar construction, see [5, proof of Lemma 30].

5.2.1 The shearing cocycle from the realization We first describe a shearing cocycle associated to the realization of Λ_n obtained in Proposition 5.3. In order to do this, recall the notation of Section 4.1.2.

Let $\varphi_n\colon \widetilde{\widetilde{S}}_n \to \mathbb{H}^3$ be a realization of Λ_n. Let k be a transverse path to Λ_n from plaque P to plaque Q, and let d be a component of $k \setminus \Lambda_n$. Then d corresponds to a plaque $R \in \mathcal{P}(\Lambda_n)$, and therefore an ideal triangle (also denoted by R) in $\mathbb{H}^2$. (Here R depends on n, but in this passage, the index n will not vary, so we suppress the notational dependence.) Recall the lamination Λ_n^* constructed in Proposition 2.28. We first observe that under the correspondence between leaves of Λ_n and Λ_n^*, the geodesics in Λ_n^* associated to the edges of R form an ideal triangle $R^* \subset \mathbb{H}^3$. This is because first, two geodesics in Λ_n^* corresponding to a pair of edges of R must be asymptotic on one end, since the map u_n is Lipschitz.

Second, if the edges of R do not form a triangle in Λ_n^*, then two such geodesics would collapse, and this is ruled out (for sufficiently large n) by Proposition 2.30.

With this understood, set $x_d^{\pm,*} = \varphi(x_d^{\pm})$. For each d, let $h_n\colon g_d^{\pm} \to \mathbb{R}$ denote the signed distance to the foot of the geodesic, as described in Section 4.1.2. Similarly, define $h_n^*\colon g_d^{\pm,*} \to \mathbb{R}$ for the corresponding ideal triangles in $\mathbb{H}^3$. We then define

$$\boldsymbol{\sigma}_n(P,Q) := h_n^*(x_{d^-}^{+,*}) - h_n^*(x_{d^+}^{-,*}) + \sum_{d\neq d^+,d^-} (h_n^*(x_d^{+,*}) - h_n^*(x_d^{-,*})). \tag{5-1}$$

Lemma 5.6 *Equation (5-1) defines a transverse cocycle $\boldsymbol{\sigma}_n \in \mathcal{H}(\Lambda_n,\mathbb{R})$.*

Proof For each d, the quantity $|h_n^*(x_d^{+,*}) - h_n^*(x_d^{-,*})|$ may be bounded by the distance from $x_d^{+,*}$ to $x_d^{-,*}$; see [5, proof of Lemma 8]. The map φ is Lipschitz with constant M_n, say, so

$$|h_n^*(x_d^{+,*}) - h_n^*(x_d^{-,*})| \leqslant M_n \ell(d), \tag{5-2}$$

where $\ell(d)$ is the hyperbolic length of d. By the estimate in [5, Lemma 5], the sum in (5-1) converges, and $\boldsymbol{\sigma}_n$ is therefore well-defined. The symmetry and additivity conditions of Section 4.1.1 are clear. □

We shall require a more precise relationship between $\boldsymbol{\sigma}_n$ and $\hat{\boldsymbol{\sigma}}_n$.

Lemma 5.7 *Fix a finite set $\mathcal{P}' \subset \mathcal{P}(\Lambda)$ and $\delta > 0$. Then there is an N such that for all $n \geqslant N$ and all $P, Q \in \mathcal{P}'$,*

$$|\boldsymbol{\sigma}_n(P,Q) - \hat{\boldsymbol{\sigma}}_n(P,Q)| \leqslant \delta t_n.$$

Proof Let A_n be the constant defined in [5, Lemma 3] for the hyperbolic structure $\widehat{S}_n$ (this depends on P and Q), which gives a lower bound on the length of a leaf in Λ_n that intersects a transverse arc from P to Q multiple times. Since the laminations Λ_n converge, and since the $\widehat{S}_n$–length of a leaf of Λ_n is stretched by a factor of t_n (see Proposition 2.29), it follows that there is a constant $A_0 > 0$ such that $A_n \geqslant A_0 t_n$ for all n and all choices of pairs in $\mathcal{P}'$.

Appealing again to [5, Lemma 5], we have

$$\ell(d) \leqslant B \exp(-t_n A_0 r(d)), \tag{5-3}$$

where $r(d)$ is the divergence radius of d and B is independent of n. Using (5-2) and (5-3), then as in the proof of Lemma 4.4, there is a sequence of plaques separating P and Q, $P = P_0, P_1, \dots, P_N = Q$, such that

$$\boldsymbol{\sigma}_n(P,Q) = \sum_{i=1}^{N} (h_n^*(x_{d_{i-1}}^{+,*}) - h_n^*(x_{d_i}^{-,*})) + E_n^*,$$

where $E_n^* = O(\exp(-Ct_n))$ for some $C > 0$. Similarly, we have

$$\hat{\boldsymbol{\sigma}}_n(P,Q) = \sum_{i=1}^{N} (h_n(x_{d_{i-1}}^{+}) - h_n(x_{d_i}^{-})) + E_n,$$

where $E_n = O(\exp(-Ct_n))$.

Each component d_i is associated with a zero $p_i \in Z(q_n)$. By Proposition 2.30 applied to the map v_n near the zero p_i, there are curves $k_n^\pm$ in $\mathcal{F}^h_{q_n}$ such that the intersection $i(\mathcal{F}^v_{q_n}, k_n)$ is uniformly bounded, and

$$\lim_{n\to\infty} \{t_n^{-1} h_n(x^\pm_{d_i}) \pm i(\mathcal{F}^v_{q_n}, k_n^\pm)\} = 0.$$

From the construction of the map φ in Section 5.1.2, it follows that for the same curves,

$$\lim_{n\to\infty} \{t_n^{-1} h_n^*(x^{\pm,*}_{d_i}) \pm i(\mathcal{F}^v_{q_n}, k_n^\pm)\} = 0.$$

Since the sums in the expressions for $\boldsymbol{\sigma}_n$ and $\widehat{\boldsymbol{\sigma}}_n$ are finite (independent of n), and P and Q range over a finite set, we can satisfy the desired inequality for any $\delta > 0$ if n is sufficiently large. □

Recall the open convex polyhedral cone $\mathcal{C}(\Lambda_n)$ from the end of Section 4.1.2.

Proposition 5.8 *For n sufficiently large, $\boldsymbol{\sigma}_n \in \mathcal{C}(\Lambda_n)$. Hence, there is a marked hyperbolic surface S_n with shearing cocycle $\boldsymbol{\sigma}_n$.*

Proof The constant C defined in [5, Lemma 6] only depends on the combinatorics of the train track supporting the laminations (cf [5, page 26]), and therefore may be taken independent of n. By Proposition 4.2(i), one can choose a sufficiently large finite set $\mathcal{P}' \subset \mathcal{P}(\Lambda_n)$ so that transverse cocycles are determined by their values on $\mathcal{P}'$. Then using Lemma 5.7 with $\delta < A_0/2C$, we conclude that

$$\|\boldsymbol{\sigma}_n - \widehat{\boldsymbol{\sigma}}_n\| < \frac{A_n}{2C}$$

for n sufficiently large. The result then follows from [5, Proposition 13] and the proof thereof. □

5.2.2 Proof of Theorem 5.1 It remains to prove the existence of a ρ_n–equivariant pleated surface map $\widetilde{S}_n \to \mathbb{H}^3$. Here we copy a construction in [5]. Let $\mathcal{P}' \subset \mathcal{P}(\Lambda_n)$ be a finite collection of plaques. For each $P \in \mathcal{P}'$, define $f_{n,\mathcal{P}'}$ on $P \subset \widetilde{S}_n \setminus \Lambda_n$ to be the oriented isometry with the corresponding plaque $P^* \subset \mathbb{H}^3$. The complement $\widetilde{S}_n \setminus \bigcup_{P\in\mathcal{P}'} P$ consists of a union of wedges. For each wedge Σ, the boundary consists of two geodesics g and h belonging to plaques in $\mathcal{P}'$. Choose (if necessary) a diagonal γ in Σ joining opposite endpoints of g and h, and map γ to the corresponding geodesic in $\mathbb{H}^3$. The diagonal γ splits Σ into two wedges, and there is a unique way to extend $f_{n,\mathcal{P}'}$ across these to make a continuous and piecewise totally geodesic map $\widetilde{S}_n \to \mathbb{H}^3$, albeit without any equivariance property. Using Lemma 5.7, as the finite sets $\mathcal{P}'$ exhaust $\mathcal{P}(\Lambda_n)$, the $f_{n,\mathcal{P}'}$ converge locally uniformly to a map f_n. The fact that f_n is ρ_n–equivariant and has shearing cocycle σ_n follows as in [5, proof of Lemmas 14 and 16].

6 Proofs

6.1 Limiting trees

We begin by proving a general result on "factorization" of equivariant harmonic maps to $\mathbb{H}^2$ and $\mathbb{H}^3$. Let $[\rho_n] \in R^o(\Sigma)$. Suppose we are given a sequence of pleated surfaces $\mathsf{P}_n = (S_n, f_n, \Lambda_n, \rho_n)$, where Λ_n

carries a transverse measure. Let $v_n\colon \widetilde{X} \to \mathbb{H}^2$ denote the lift of the degree-one harmonic diffeomorphism $X \to S_n$, the hyperbolic surface underlying P_n. We also set $w_n = f_n \circ v_n\colon \widetilde{X} \to \mathbb{H}^3$. Note that w_n is ρ_n–equivariant, and since f_n is totally geodesic on the complement of Λ_n, which has measure zero, w_n is an L_1^2–map with the same energy as v_n. Finally, as usual, we let $u_n\colon \widetilde{X} \to \mathbb{H}^3$ denote the ρ_n–equivariant harmonic map.

Let $q_n = \mathrm{Hopf}(u_n)$ and $\psi_n = \mathrm{Hopf}(v_n)$. We may assume (after passing to a subsequence) that

$$\frac{q_n}{\|q_n\|_1} \to q, \quad \frac{\psi_n}{\|\psi_n\|_1} \to \psi.$$

Let $\phi_{\mathsf{HF}}(\Lambda_n)$ denote the Hubbard–Masur differential whose horizontal measured foliation is measure-equivalent to the one corresponding to the lamination Λ_n; see Section 2.4.1.

Proposition 6.1 *Suppose the following hold*:

(i) $\psi \in \mathcal{SQD}^*(X)$.

(ii) $\displaystyle\lim_{n\to\infty} \frac{\|4\psi_n - \phi_{\mathsf{HF}}(\Lambda_n)\|_1}{\|\psi_n\|_1} = 0.$

Then

$$\lim_{n\to\infty} \frac{\|q_n - \psi_n\|_1}{\|\psi_n\|_1} = 0.$$

The rest of this section is devoted to the proof of Proposition 6.1. We begin with the following.

Lemma 6.2 *Under assumption* (ii) *above*, *there is a constant* $0 < c \leqslant 1$ *such that*

$$c \cdot E(v_n) \leqslant E(u_n) \leqslant E(v_n).$$

Proof The inequality $E(u_n) \leqslant E(v_n)$ is automatic, since u_n is an energy minimizer among ρ_n–equivariant L_1^2–maps, and $E(v_n) = E(w_n)$. Choose any conformal metric on X, and induce a metric on $\widetilde{X}$. By the uniform Lipschitz property of harmonic maps to NPC targets (cf [55, Theorem 2.2]), there is a constant B independent of j such that for any points $p, q \in \widetilde{X}$,

$$d_{\mathbb{H}^3}(u_n(p), u_n(q)) \leqslant B\, d_{\widetilde{X}}(p,q) \cdot E^{1/2}(u_n).$$

In particular, for $\gamma \in \pi_1$,

$$\tau_{\mathbb{H}^3}(\rho_n(\gamma)) \leqslant B\, \ell_X(\gamma) \cdot E^{1/2}(u_n). \tag{6-1}$$

By (ii), the laminations Λ_n are close to $\Lambda^h_{\psi_n}$. By Poincaré recurrence, we may choose a nondegenerate leaf of $\mathcal{F}^h_{\psi_n}$ and form a closed loop γ by adding small segments of the vertical foliation; see [20, Corollary 5.3]. The image by f_n of the lift of this loop consists of nearly geodesic segments joined by tiny orthogonal segments, so the length approximates the translation length of the corresponding element in $\mathrm{Iso}(\mathbb{H}^3)$. The high-energy behavior of v_n (cf [58] and Proposition 5.4) further implies that this length is approximated by the transverse measure to $\mathcal{F}^v_{\psi_n}$. From this we deduce the existence of a constant $c_0 > 0$ such that

$$\tau_{\mathbb{H}^3}(\rho_n(\gamma)) \geqslant c_0 \cdot i(\gamma, \mathcal{F}^v_{\psi_n/\|\psi_n\|}) E(v_n)^{1/2}.$$

We then observe that since γ may be chosen to be long and nearly along the leaves of $\mathcal{F}^v_{\psi_n}$, the X– and $|\psi_n|/\|\psi_n\|$–lengths of γ are comparable. The lemma then follows from (6-1). ◻

Let $(C_n, d_{\mathbb{H}^3})$ be the closed convex hull of the image of w_n in $\mathbb{H}^3$. We now consider the rescaled metric spaces $\mathbb{H}^2_n := (\mathbb{H}^2, t_n^{-1} d_{\mathbb{H}^2})$ and $W_n := (C_n, t_n^{-1} d_{\mathbb{H}^3})$, where $t_n^2 = E(v_n)$. By the uniform Lipschitz property used in the previous proof, $v_n\colon \widetilde{X} \to \mathbb{H}^2_n$ has uniform modulus of continuity in the sense of Theorem 2.32. Since the map $f_n\colon \mathbb{H}^2 \to C_n$ is distance nonincreasing, it follows that $w_n\colon \widetilde{X} \to W_n$ has uniform modulus of continuity as well.

Lemma 6.3 *After possibly passing to a subsequence, we have the following properties:*

(i) *$\mathbb{H}^2_n$ with the isometric action of π_1 converges in the Gromov–Hausdorff sense to the $\mathbb{R}$–tree T_ψ dual to the quadratic differential ψ, up to scale, and the maps v_n converge to a surjective π_1–equivariant harmonic map $v\colon \widetilde{X} \to T_\psi$.*

(ii) *W_n with the isometric action of π_1 converges in the Gromov–Hausdorff sense to an $\mathbb{R}$–tree T with π_1–action whose projective length function is equivalent to the Morgan–Shalen limit of $\{\rho_n\}$. In particular, the action is minimal. The maps w_n converge to an equivariant map $w\colon \widetilde{X} \to T$ of finite energy.*

(iii) *The maps $f_n\colon \mathbb{H}^2_n \to W_n$ converge to a morphism of trees $f: T_\psi \to T$. There is no folding of edges in T_ψ corresponding to adjacent critical leaves of the horizontal foliation of ψ meeting at a zero. Moreover, $w = f \circ v$.*

Note that the embedding, into the surface, of the graph of critical leaves of the horizontal foliation of ψ, induces a natural notion of adjacency of critical leaves.

Proof The convergence property in item (i) follows from [4; 53], and the harmonicity, surjectivity and convergence to the map is in [60]. Convergence in (ii) follows by the construction in Theorem 2.32. Note that the result of Lemma 6.2 guarantees that the length function has a well-defined nonzero limit and that the resulting map v is nonconstant. The fact that the limiting tree is the Morgan–Shalen (minimal) tree follows from Theorem 2.34. Item (iii) can be seen by taking the images of leaves of the horizontal foliation of ψ and using the fact that the pleated maps f_n take leaves of the lamination to geodesics. The last assertion in (iii) is obvious. ◻

We shall need some further properties of the map $w\colon \widetilde{X} \to T$.

Lemma 6.4 *Fix $z_0 \in \widetilde{X}$ and $Q \in T$. On a sufficiently small disk about z_0, the function $z \mapsto d_T(w(z), Q)$ is subharmonic. Moreover, the Hopf differential of w is well-defined and equal to ψ.*

Proof We may assume z_0 is a point such that the map p folds at $v(z_0)$, since otherwise f is a local isometry, and the result follows since v is harmonic with Hopf differential ψ. Choose the disk U such that the image $v(U)$ consists of geodesic segments e_1, e_2, e_3 in T_ψ meeting at a vertex. Because f is a folding, we may assume that f maps each e_i isometrically onto corresponding geodesic segments $\overline{e}_i$ in T. Since folding cannot occur on edges, hence not on adjacent edges incident to a vertex, and the zero in U is trivalent, we see that the $\overline{e}_i$ are distinct.

Either the geodesic segment $\overline{w(z_0)Q}$ intersects each $\overline{e}_i$ in the point $w(z_0)$ (which we call Case 1); or $\overline{w(z_0)Q}$ intersects some (and hence only one) $\overline{e}_i$ in a nondegenerate segment (which we call Case 2).

Suppose we have Case 1. Then we claim that for lifts $\widetilde{Q}$ of Q to T_ψ,

$$\overline{v(z_0)\widetilde{Q}} \cap e_i = \{v(z_0)\}.$$

If this were not the case, let R be a nondegenerate segment in one of the intersections above, and set $\overline{R} = f(R) \subset T$. Set $R^- = \{v(z_0)\}$, and let R^+ denote the other endpoint of R. The geodesic γ in T_ψ from R^+ to $\widetilde{Q}$ is disjoint from the interior of R. Hence, its image $\overline{\gamma}$ in T is a path from $\overline{R}^+$ to Q that is disjoint from the interior of $\overline{R}$. On the other hand, there is another path from $w(z_0)$ to Q that is disjoint from the interior of $\overline{R}$. This contradicts the fact that T is a tree.

Since the image of v on U is $e_1 \cup e_2 \cup e_3$, it follows that

$$\begin{aligned} d_{T_\psi}(v(z), \widetilde{Q}) &= d_{T_\psi}(v(z), v(z_0)) + d_{T_\psi}(v(z_0), \widetilde{Q}) \\ &= d_T(w(z), w(z_0)) + d_{T_\psi}(v(z_0), \widetilde{Q}) \\ &= d_T(w(z), Q) - d_T(w(z_0), Q) + d_{T_\psi}(v(z_0), \widetilde{Q}). \end{aligned}$$

Since $d_{T_\psi}(v(z), \widetilde{Q})$ is subharmonic on U, so is $d_T(w(z), Q)$.

In Case 2, suppose without loss of generality that $\overline{w(z_0)Q} \cap \overline{e}_1 = \overline{R}$ is a nondegenerate segment. Then for small enough U,

$$d_T(w(z), Q) = d_T(w(z), \overline{R}^+) + d_T(\overline{R}^+, Q).$$

Now because adjacent edges do not fold,

$$d_T(w(z), \overline{R}^+) = d_{T_\psi}(v(z), R^+),$$

and the result follows as above. This proves the first part of the lemma. Taking $Q = w(z_0)$, we see that $d_T(w(z), w(z_0)) = d_{T_\psi}(v(z), v(z_0))$. Then the energy densities and Hopf differentials of w and v must coincide; see [38, Sections 1.2 and 2.3]. The lemma is proved. □

Again appealing to Theorem 2.34, the rescaled convex hulls of the images of the u_n converge to give an equivariant harmonic map $u\colon \widetilde{X} \to T$ to the Morgan–Shalen limit. Let us emphasize that since the limiting length function is not abelian, there is a unique $\mathbb{R}$–tree associated to the Morgan–Shalen limit with the given limiting length function.

Proof of Proposition 6.1 Applying assumption (ii) provides for the estimates in Lemma 6.2, and hence the existence of the nontrivial map w in Lemma 6.3. Assumption (i) is used in proof of Lemma 6.4. With this understood, consider $D(z) = d_T(u(z), w(z))$ for $z \in \widetilde{X}$. First, since u and w are equivariant for the same action, $D(z)$ descends to a function on X. Next, observe that $D(z)$ is continuous, since both u and w are Lipschitz. Let

$$\mathcal{S} = \{z \in X \mid D(z) = \max_X D\}.$$

Then $\mathcal{S}$ is closed and nonempty. Let $z_0 \in \mathcal{S}$. We may assume $D(z_0) > 0$, since otherwise $u = w$ and there is nothing to prove. Let $Q \in T$ be the midpoint of the geodesic $\overline{u(z_0)w(z_0)}$, and choose an open disk U about z_0 sufficiently small so that $u(U) \cap w(U) = \varnothing$. Since $T \setminus \{Q\}$ is disconnected and hence $u(U)$ and $w(U)$ lie in different components, any path from one image to the other must pass through Q. In particular,

$$D(z) = d_T(u(z), Q) + d_T(w(z), Q) \quad \text{for all } z \in U.$$

Since distance to a point is a convex function, and harmonic maps pull back convex functions to subharmonic functions, the first term on the right-hand side is subharmonic. For U sufficiently small, the second term is also subharmonic by Lemma 6.4. Hence, using the strong maximum principle this implies that D is constant on U, and so $\mathcal{S}$ is open. Hence, $D(z)$ is a constant function.

We claim that u and w have the same Hopf differentials. Suppose $D = D(z) > 0$. Then the claim will follow (as above using the definition in [38]) by showing that for any z_0 there is a small enough neighborhood U about z_0 such that

$$d_T(u(z), u(z_0)) = d_T(w(z), w(z_0)) \quad \text{for all } z \in U. \tag{6-2}$$

Let $R \subset T$ denote the edge from $u(z_0)$ to $w(z_0)$. Let

$$D_u^+ = \{z \in U \mid u(z) \notin R\} \quad \text{and} \quad D_u^- = \{z \in U \mid u(z) \in R\}.$$

Similarly, we define $D_w^\pm$. Notice that since $D(z) = D$ is constant, $D_w^+ = D_u^-$ and $D_w^- = D_u^+$. Hence, for $z \in D_u^+$,

$$D = d_T(u(z), w(z)) = d_T(u(z), u(z_0)) + d(u(z_0), w(z)) = d(u(z), u(z_0)) - d_T(w(z), w(z_0)) + D,$$

whereas for $z \in D_u^-$,

$$D = d_T(u(z), w(z)) = d_T(w(z), w(z_0)) + d(w(z_0), u(z)) = d(w(z), w(z_0)) - d_T(u(z), u(z_0)) + D.$$

In both cases, the equality (6-2) holds. This proves that $\psi = q$.

In fact, since the energy densities of u and w agree, w is also energy minimizing. But equivariant harmonic maps to nontrivial $\mathbb{R}$–trees are unique (see [45]), so that in fact $w = u$. Choose $p \neq p' \in \mathbb{H}^2$ to lie on a portion of a leaf $\ell \subset \widetilde{\mathcal{F}}_q^h$ away from the zeroes. We assume that the map $f : T_q \to T$ maps the image of ℓ isometrically onto a geodesic segment in T. By the definition of a folding and the dual tree, such an ℓ can always be found: one simply takes a preimage in a leaf of an arc in a tree and restricts to a small enough subleaf that is not folded. By Theorems 2.31 and 2.32, we have

$$E(w_n)^{-1/2} d_{\mathbb{H}^3}(w_n(p), w_n(p')) = E(v_n)^{-1/2} d_{\mathbb{H}^3}(v_n(p), v_n(p')) \to i(\ell, \mathcal{F}^v_\psi), \qquad E(u_n)^{-1/2} d_{\mathbb{H}^3}(u_n(p), u_n(p')) \to i(\ell, \mathcal{F}^v_q), \tag{6-3}$$

and since $q = \psi$, the right-hand sides are equal. On the other hand, from Lemma 6.3 and the fact that $w = u$,

$$\lim_{n\to\infty} E(v_n)^{-1/2}\big(d_{\mathbb{H}^3}(w_n(p), w_n(p')) - d_{\mathbb{H}^3}(u_n(p), u_n(p'))\big) = 0 \tag{6-4}$$

(recall that $t_n^2 = E(v_n)$). Equations (6-3) and (6-4) force

$$\lim_{n\to\infty} \frac{E(u_n)}{E(v_n)} = 1,$$

which implies

$$\lim_{n\to\infty} \frac{\|q_n\|_1}{\|\psi_n\|_1} = 1. \tag{6-5}$$

Indeed, this follows because the u_n and v_n are harmonic, and the energy converges [39, Theorem 3.9]. Alternatively, for v_n we have the inequality

$$\|\psi_n\|_1 - 2\pi(g-1) \leqslant E(v_n) \leqslant \|\psi_n\|_1 + 2\pi(g-1)$$

(see [19]), and so

$$\lim_{n\to\infty} \frac{E(v_n)}{\|\psi_n\|_1} = 1.$$

Similarly, using Theorem 2.14(iii) and (iv), equation (2-24), and the asymptotics in Theorem 2.6, we also have

$$\lim_{n\to\infty} \frac{E(u_n)}{\|q_n\|_1} = 1.$$

Hence, (6-5). The proposition now follows from the fact that $\psi = q$, (6-5), and the algebraic inequality

$$\frac{\|q_n - \psi_n\|_1}{\|\psi_n\|_1} \leqslant \left|1 - \frac{\|q_n\|_1}{\|\psi_n\|_1}\right| + \left\|\frac{q_n}{\|q_n\|_1} - \frac{\psi_n}{\|\psi_n\|_1}\right\|_1. \qquad \square$$

6.2 Limiting configurations and limits of representations

6.2.1 Proof of the Main Theorem Part (i) of the Main Theorem is the content of Theorem 5.1. The harmonic map estimates in Propositions 2.28 and 2.30 show that the pleated surfaces $f_n\colon \widetilde{S}_n \to \mathbb{H}^3$ are asymptotic to the images of the harmonic maps $u_n\colon \widetilde{X} \to \mathbb{H}^3$ in the sense of Definition 4.14. This is part (ii) of the theorem. Part (iii) then follows from Lemmas 5.7 and 5.5. Finally, part (iv) is a consequence of the approximation in Definition 4.14 and Theorem 4.16. This completes the proof.

6.2.2 Proof of Corollary 1.1 Let us first recast Theorem 4.5 in terms of train tracks. Let $q \in \mathcal{QD}^*(X)$. Let τ be a complete train-track (cf [54, pages 27, 175]) carrying the horizontal lamination Λ_q^h. Let $\mathcal{H}^o(\tau, S^1)$ be the connected component of the identity of the space of S^1–valued cocycles on τ. Then we have the following.

Theorem 6.5 *There is a group isomorphism*

$$\mathcal{H}^o(\tau, S^1) \simeq \mathrm{Prym}(\hat{X}_q, X)/J_2(X).$$

Proof Choose a maximalization of $\mathcal{F}_q^h$ in the sense of Definition 2.25. By Lemma 2.26, this gives a maximalization Λ of Λ_q^h. Now Λ is carried by the splitting τ' of τ corresponding to the maximalization. Since there is a natural isomorphism

$$\mathcal{H}^o(\tau, S^1) \simeq \mathcal{H}^o(\tau', S^1) \simeq \mathcal{H}^o(\Lambda, S^1),$$

the result follows from Theorem 4.5. Note that from the construction leading to Corollary 4.11, the isomorphism is independent of the choice of maximalization. □

Fix X_0 and $q_0 \in \mathcal{SQD}^*(X_0)$. The train track τ may be chosen so that for any $X \in U_0$, τ carries $\Lambda_q^h(X)$, where $q \in \mathcal{SQD}^*(X)$ is the Hubbard–Masur differential for the measured foliation $\mathcal{F}_{q_0}^v$. A maximalization $\Lambda(X)$ of $\Lambda_q^h(X)$ is carried by a splitting of τ.

We now continue with the proof of the corollary. Let T be a Morgan–Shalen limit of $[\rho_n]$. As we have noted before, by Theorem 2.34 there is an equivariant harmonic map $u\colon \tilde{X} \to T$ that factorizes through T_{q_0}. Note that since q_0 has simple zeroes, the action on T is not abelian, and so the tree T is uniquely (up to scale) associated to the projective length function of the Morgan–Shalen limit.

Consider $X \in U_0$. Then as above, we have an equivariant harmonic map $\tilde{X} \to T_{q'_X} \xrightarrow{v} T$. We claim that up to an overall scale, $T_{q'_X}$ is equivariantly isometric to T_{q_0}. From this, it follows that $q'_X = q_X$. To prove the claim, let $v\colon \tilde{X}_0 \to T_{q'_X}$ be the equivariant harmonic map, and set $w = f \circ v\colon \tilde{X}_0 \to T$. Then using exactly the same argument as in Lemma 6.4 and the proof of Proposition 6.1, we conclude that the Hopf differential of v is also q_0. Since the action on $T_{q'_X}$ is "small", it follows that the folding $T_{q_0} \to T_{q'_X}$ induced by v from Theorem 2.31(iii) is actually an isometry; see [56, Proposition 3.1]. This proves the first statement of the corollary.

To prove the statement about bending cocycles, let $\hat{S}_n(X_0)$ be the companion surfaces as in Section 5.1.1. Recall the construction of a pleated surface $\mathsf{P}_n(X_0) = (S_n(X_0), \tilde{f}_{n,X_0}, \Lambda_n(X_0), \rho_n)$, where the shearing cocycle of the hyperbolic surface $S_n(X_0)$ is obtained as a perturbation of the one for $\hat{S}_n(X_0)$. We may choose the train track $\tau_{n,*}$ used in that proof to carry both laminations $\Lambda_n^h(X)$ and $\Lambda_n^h(X_0)$. By a straightforward energy estimate the scaling factors of the quadratic differentials on X and X_0 are comparable: ie there is a constant C depending only on U_0 so that

$$C^{-1} t_n(X_0) \leqslant t_n(X) \leqslant C t_n(X_0).$$

Perturbing the shearing cocycle of $\widehat{S}_n(X_0)$ as in (5-1), but now with respect to the lamination $\Lambda_n(X)$ instead of $\Lambda_n(X_0)$, the argument in Section 5.2.1 carries over to show that there is a ρ_n–equivariant pleated surface with pleating locus $\Lambda_n(X)$. By [5, Lemma 29], this must agree (up to isotopy) with the pleated surface $\mathsf{P}_n(X) = (S_n(X), \widetilde{f}_{n,X}, \Lambda_n(X), \rho_n)$ constructed in Theorem 5.1 with the basepoint X.

Now the complementary regions of the lift $\widetilde{\tau}_{n,*}$ of the train track to $\mathbb{H}^2$ give the identification of the plaques for $\mathsf{P}_n(X)$ and $\mathsf{P}_n(X_0)$. Each plaque P is realized in two ways (say P_X and P_{X_0}) as an ideal triangle in $\mathbb{H}^3$, and where by the asymptotic estimates on the harmonic maps u_n (cf Proposition 2.30), the triples of leaves in $\Lambda_n(X)$ and in $\Lambda_n(X_0)$ bounding P are close over a large hexagonal region of P. For a pair of plaques P and Q, fix points $\widetilde{p} \in P_X$, $\widetilde{p}_0 \in P_{X_0}$, $\widetilde{q} \in Q_X$ and $\widetilde{q}_0 \in Q_{X_0}$. Then as in the proof of Lemma 4.4, the proximity of the ideal triangles for $\mathsf{P}_n(X)$ and $\mathsf{P}_n(X_0)$ when n is large gives an estimate on $|\Theta_{\widetilde{f}_{n,X}}(\widetilde{p},\widetilde{q}) - \Theta_{\widetilde{f}_{n,X_0}}(\widetilde{p}_0,\widetilde{q}_0)|$. From (4-5), we see that the bending cocycles $\boldsymbol{\beta}_n(X_0)$ and $\boldsymbol{\beta}_n(X)$ give the same limit as a cocycle in $\mathcal{H}^0(\tau, S^1)$. By the Main Theorem, this common limit determines the periods of η_{X_0} and η_X, and therefore identifies their cohomology classes under the Gaus–Manin connection. This concludes the proof of Corollary 1.1.

6.2.3 Proof of Corollary 1.2 Consider the situation of the Main Theorem. Suppose that the limiting quadratic differential q has a vertical saddle connection between $p, p' \in Z(q)$. Let p_n and p'_n be zeroes of q_n so that $p_n \to p$ and $p'_n \to p'$. If there is a folding in the Morgan–Shalen limit, then the following must happen: there is some $\delta_0 > 0$ such that for all $0 < \delta \leqslant \delta_0$, there are points z_n and z'_n with

$$
\begin{aligned}
t_n^{-1} d_{\mathbb{H}^3}(p_n^*, (p'_n)^*) &\to 0, & t_n^{-1} d_{\mathbb{H}^3}(z_n^*, (z'_n)^*) &\to 0, \\
t_n^{-1} d_{\mathbb{H}^3}(z_n^*, p_n^*) &\to \delta, & t_n^{-1} d_{\mathbb{H}^3}((z'_n)^*, (p'_n)^*) &\to \delta.
\end{aligned}
\tag{6-6}
$$

See Figure 6. The notation here means that $p_n^* = u_n(p_n)$, etc. See Theorem 2.33 and the definition of folding in Section 2.6.1. As in the proof of Lemma 4.4, the planes $\mathbb{D}_{p_n^*}$ and $\mathbb{D}_{(p'_n)^*}$ intersect, and the assumption is that the dihedral angle is bounded away from π. Let A_n denote the geodesic segment between z_n^* and p_n^*, B_n the geodesic segment between $(z'_n)^*$ and p_n^*, and C_n between z_n^* and $(z'_n)^*$, and let α_n, β_n and γ_n be the corresponding angles of the geodesic triangle thus formed. Then the assumption implies $\gamma_n \geqslant \varepsilon > 0$ for some fixed ε and n sufficiently large. From (6-6), we may assume $|B| \geqslant \delta t_n/2$. But then

$$\sinh|C| \geqslant (\sin\varepsilon)\sinh\left(\tfrac{1}{2}\delta t_n\right),$$

which contradicts the assumption that $t_n^{-1}|C| \to 0$. This completes the proof.

6.3 Complex projective structures

The goal of this section is to prove Corollary 1.3. In order to do so, insofar as the pleated surface is already given by Thurston, it is necessary to in some sense reverse the argument used in Section 5. For this, we use the result of Section 6.1, which gives criteria to identify the limiting quadratic differential for equivariant harmonic maps in terms of the lamination of the associated pleated surfaces. In Section 6.3.1,

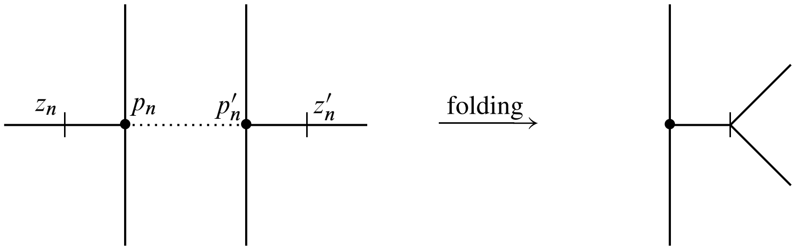

Figure 6: Folding.

we review Dumas' estimates, which show that the criteria just mentioned hold for Thurston's pleated surfaces associated to projective structures. In Section 6.3.2, we use facts about opers to derive the limiting spectral data. Finally, Corollary 1.3 is proven in the last section.

6.3.1 Dumas' estimates We recall the estimates of Dumas in [16] (see particularly Theorems 1.1 and 14.2) which relate complex projective structures and Hopf differentials. As mentioned in the Introduction, given $q \in \mathcal{QD}(X)$ the projective connection $\mathrm{Op}(q)$ produces a pleated surface $\mathsf{P}(q) = (S(q), f_q, \Lambda(q), \mathscr{P}(q))$. Moreover, the bending lamination $\Lambda(q)$ carries a transverse measure. Strictly speaking, $\Lambda(q)$ may not be maximal; it will turn out that the choice of maximalization of $\Lambda(q)$ will be immaterial, and so we suppress it from the notation.

The first result compares q with the Hubbard–Masur differential defined by the lamination.[6]

Theorem 6.6 [16, Theorem 1.1] *There is a constant $C = C(X)$ that only depends on the Riemann surface X such that*

$$\|4q - \phi_{\mathsf{HF}}(\Lambda(q))\|_1 \leqslant C(X)(1 + \|q\|_1^{1/2}). \tag{6-7}$$

The second important result is a comparison of the quadratic differentials parametrizing projective structures and those in the harmonic maps parametrization of Teichmüller space. More precisely, Dumas proves the following; see [16, Theorem 14.2 and proof], and also [15].

Theorem 6.7 *Fix $q \in \mathcal{SQD}(X)$, and let ψ denote the Hopf differential of the harmonic diffeomorphism $X \to S(q)$. Then*

$$\|4\psi(q) - \phi_{\mathsf{HF}}(\Lambda(q))\|_1 \leqslant C(X)(1 + \|q\|_1^{1/2}). \tag{6-8}$$

Recall that $\mathscr{P}(q)$ denotes the monodromy of the projective connection $Q(q)$. Combining Theorems 6.6 and 6.7 with Proposition 6.1, we have the following.

[6] Because Dumas uses the Schwarzian, the quadratic differential he uses to parametrize $S(q)$ differs from the one in (1-8) by a factor of -2.

Corollary 6.8 *Let $q \in \mathcal{SQD}^*(X)$, and let $u_t \colon \widetilde{X} \to \mathbb{H}^3$ be the $\mathscr{P}(t^2q)$–equivariant harmonic map with Hopf differential q_t. Then*

$$\lim_{t\to+\infty} \|t^{-2}q_t - q\|_1 = 0.$$

6.3.2 Spectral data for opers Here we determine the possible limiting bending cocycles of the family $\mathrm{Op}(t^2q)$. The argument we give is based on the identification of limiting configurations with limiting spectral data, and the classical fact that the underlying holomorphic bundle $\mathcal{V}$ of a (lift of a) complex projective structure is the unique nonsplit extension

$$0 \to K_X^{1/2} \to \mathcal{V} \to K_X^{-1/2} \to 0$$

(cf [27, page 201]). In terms of Higgs pairs, this means that the $\bar\partial$–operator $\bar\partial_A + \Phi^*$ must induce the holomorphic structure on $\mathcal{V}$. Moreover, since $\mathcal{V}$ has a flat connection, the holomorphic structure on $\mathcal{V}$ can be uniquely characterized by the fact that it contains $K_X^{1/2}$ as a subsheaf, and this is the criterion we shall use.

Before proceeding, it may clarify things to recall again that by Theorem 2.14 and the definition of the Hitchin map (2-6), if u is the equivariant harmonic map associated to a solution (A,Ψ) of the self-duality equations, then $\mathrm{Hopf}(u) = 4\mathscr{H}([A,\Psi])$. With this understood, we have the following.

Proposition 6.9 *Let q and q_t be as in the statement of Corollary 6.8. Let $[(A_t,\Psi_t)] = \mathrm{Op}(t^2q)$, and let η_t be the term appearing in the approximation in Definition 2.5, and $\hat\eta_t$ the Prym differential corresponding to η_t in Proposition 2.11. Then*

$$[\hat\eta_t - i\,t\,\mathrm{Im}\,\lambda_{\mathrm{SW}}] \to 0 \quad \textit{in } \mathrm{Prym}(\widehat{X}_q, X)/J_2(X).$$

Proof Let $\hat q_t = t^{-2}q_t$. Consider the spectral curves $\pi_t \colon \widehat{X}_{\hat q_t} \to X$. Denote the Seiberg–Witten differential (resp. tautological section) on $\widehat{X}_{\hat q_t}$ by $\lambda_{\mathrm{SW}}(t)$ (resp. λ_t). Since by Corollary 6.8, $\hat q_t \to q$ as $t \to +\infty$, it follows that $\lambda_{\mathrm{SW}}(t) \to \lambda_{\mathrm{SW}}$ on $\widehat{X}_q$, where the convergence is taken with respect to the Gaus–Manin connection on Prym differentials.

Let $f_t \colon \widehat{X}_{\hat q_t} \to \widehat{X}_{q_t/4}$ be as in the proof of Proposition 2.11. Then $f_t^*\lambda_{\mathrm{SW}} = (t/2)\lambda_{\mathrm{SW}}(t)$.

Now, from the discussion at the beginning of this subsection, there is an injective homomorphism of smooth bundles, $T_t \colon \pi_t^* K_X^{1/2} \to \pi_t^* E$, such that the image is preserved by the pullback $\bar\partial$–operator $\pi_t^*(\bar\partial_{A_t} + \Phi_t^*)$, and the induced $\bar\partial$–operator is isomorphic to the canonical one on $\pi_t^* K_X^{1/2}$, up to possibly twisting by a 2–torsion line bundle. As a smooth bundle, we have the splitting $\pi_t^* E \simeq \pi_t^* K_X^{-1/2} \oplus \pi_t^* K_X^{1/2}$; see (2-1). Let σ be a local trivialization of $\pi_t^* K_X^{1/2}$, and write $T_t(\sigma) = (\sigma_t^{(1)}, \sigma_t^{(2)})$, where $\sigma_t^{(i)} = \xi_t^{(i)}\sigma$ for a local smooth section $\xi_t^{(1)} \in \Gamma(\pi_t^* K_X^{-1})$ and smooth function $\xi_t^{(2)}$. By a straightforward calculation, one finds the component entries of $\pi_t^*(\bar\partial_{A_t} + \Phi_t^*)T_t$ to be (after an overall rescaling):

$$\begin{aligned} &\bar\partial(\lambda_t\xi_t^{(1)}\cdot\|\lambda_t\|^{-1/2}) + \big(\hat\eta_t'' + \tfrac12 t\bar\lambda_{\mathrm{SW}}(t)\big)\xi_t^{(2)}\|\lambda_t\|^{1/2} = R_1(t),\\ &\bar\partial(\xi_t^{(2)}\cdot\|\lambda_t\|^{1/2}) + \big(\hat\eta_t'' + \tfrac12 t\bar\lambda_{\mathrm{SW}}(t)\big)\lambda_t\xi_t^{(1)}\|\lambda_t\|^{-1/2} = R_2(t).\end{aligned} \tag{6-9}$$

Here, the remainder terms $R_i(t)$ are linear combinations of $\lambda_t\xi_t^{(1)}\cdot\|\lambda_t\|^{-1/2}$ and $\xi_t^{(2)}\|\lambda_t\|^{1/2}$ with coefficients that are exponentially small as $t\to+\infty$. To obtain (6-9), we use the expression for the limiting connection in (2-9) to calculate:

$$\pi_t^*(\bar\partial_{A^0_\infty(q_t)}+\eta_t'')(T_t(\sigma))-T_t(\bar\partial\sigma)=\begin{pmatrix}\bar\partial\xi_t^{(1)}-(\bar\partial\log\|\lambda_t\|^{1/2})\xi_t^{(1)}+\hat\eta_t''\lambda_t^{-1}\|\lambda_t\|\xi_t^{(2)}\\ \bar\partial\xi_t^{(2)}+(\bar\partial\log\|\lambda_t\|^{1/2})\xi_t^{(2)}+\hat\eta_t''\lambda_t\|\lambda_t\|^{-1}\xi_t^{(1)}\end{pmatrix}\sigma.$$

Similarly,

$$\pi_t^*(\Phi^*(q_t)(T_t(\sigma))=\tfrac12 t\bar\lambda_{\mathrm{SW}}(t)\begin{pmatrix}\lambda_t^{-1}\|\lambda_t\|\xi_t^{(2)}\\ \lambda_t\|\lambda_t\|^{-1}\xi_t^{(1)}\end{pmatrix}\sigma.$$

Multiplying the first entries by $\lambda_t\|\lambda_t\|^{-1/2}$, and the second entries by $\|\lambda_t\|^{1/2}$, we obtain the left-hand side of (6-9). The error terms come from applying the result in Theorem 2.6.

Now suppose to the contrary that there is a sequence $t_n\to+\infty$ and β_n, odd harmonic $(0,1)$ forms with periods in $2\pi i\mathbb{Z}$, such that

$$\lim_{n\to\infty}\{(\eta_{t_n}''+\tfrac12 t_n\bar\lambda_{\mathrm{SW}}(t_n))-\beta_n\}=\alpha,$$

where the class of α in the Prym variety is nonzero; see Section 2.2.2. Choose an arbitrary basepoint $z_0\in\hat X_q$, and redefine

$$\tilde\xi_n^{(i)}(z)=\exp\Bigl(-\int_{z_0}^z\beta_n\Bigr)\xi_{t_n}^{(i)}(z).$$

Then (6-9) becomes

$$\begin{aligned}&\bar\partial(\lambda_{t_n}\tilde\xi_n^{(1)}\cdot\|\lambda_{t_n}\|^{-1/2})+(\hat\eta_{t_n}''+\tfrac12 t_n\bar\lambda_{\mathrm{SW}}(t_n)-\beta_n)\tilde\xi_n^{(2)}\|\lambda_{t_n}\|^{1/2}=\tilde R_1(t_n),\\ &\bar\partial(\tilde\xi_n^{(2)}\cdot\|\lambda_{t_n}\|^{1/2})+(\hat\eta_{t_n}''+\tfrac12 t_n\bar\lambda_{\mathrm{SW}}(t_n)-\beta_n)\lambda_{t_n}\tilde\xi_n^{(1)}\|\lambda_{t_n}\|^{-1/2}=\tilde R_2(t_n),\end{aligned}\tag{6-10}$$

and where the remainder terms $\tilde R_i(t_n)$ are exponentially small as $t_n\to+\infty$, and of the order of $\lambda_{t_n}\tilde\xi_n^{(1)}\|\lambda_{t_n}\|^{-1/2}$ and $\tilde\xi_n^{(2)}\|\lambda_{t_n}\|^{1/2}$.

Fix a conformal metric on $\hat X_q$ with area form dv. Let us now normalize the sequence of homomorphisms T_{t_n} so that

$$\int_{\hat X_q}\|\lambda_{t_n}\|(\|\tilde\xi_n^{(1)}\|^2+|\tilde\xi_n^{(2)}|^2)\,dv=1.$$

Applying elliptic regularity to (6-10), we may assume that $\lambda_{t_n}\tilde\xi_n^{(1)}\|\lambda_{t_n}\|^{-1/2}\to f_1$ and $\tilde\xi_n^{(2)}\|\lambda_{t_n}\|^{1/2}\to f_2$, for functions f_i satisfying

$$\int_{\hat X_q}(|f_1|^2+|f_2|^2)\,dv=1,$$

and $\bar\partial f_1+\alpha f_2=0$, $\bar\partial f_2+\alpha f_1=0$. This, of course, implies $(\bar\partial+\alpha)(f_1+f_2)=0$. If $f_1+f_2\neq0$, then the holomorphic line bundle $\mathcal{L}$ defined by α has a nonzero holomorphic section and is therefore trivial. If $f_1+f_2=0$, then f_1 is nonzero, and $(\bar\partial-\alpha)f_1=0$; so $\mathcal{L}^*$ is trivial. In either case, this contradicts the assumption. The proposition is proved. □

6.3.3 Proof of Corollary 1.3 Fix $q \in \mathcal{SQD}^*(X)$. Let $\widetilde{S}(t^2q) \to \mathbb{H}^3$ be Thurston's pleated surface associated to the projective connection $Q(t^2q)$ with monodromy $\mathscr{P}(t^2q)$, and with bending lamination $\Lambda(t^2q)$. Part (i) of the corollary follows from Corollary 6.8 and Lemma 5.5. Part (ii) is the content of Proposition 6.9. We now move on to prove part (iii) of the corollary. Let S_t denote the companion surface defined by requiring the Hopf differential for the harmonic diffeomorphism $X \to S_t$ to be t^2q.

By Theorem 6.7, $\Lambda(t^2q)$ converges to Λ_q^h. In particular, for t sufficiently large, $\Lambda(t^2q)$ is carried by the train track $\tau_{t,*} \subset S_t$ constructed in Section 5.1.1. By (i) and arguing as in the proof of Corollary 1.1, a perturbation of the shearing cocycle of S_t as described in Section 5.2.1 results in a pleated surface for $\mathscr{P}(t^2q)$ with bending lamination $\Lambda(t^2q)$; as before by the uniqueness of pleated surfaces for a fixed lamination, the pleated surface constructed in this way must coincide with the pleated surface $S(t^2q)$. Now, by Theorem 6.6 and Corollary 6.8, the bending laminations for the harmonic map u_t and the bending lamination $\Lambda(t^2q)$ are close and hence carried by the same track $\tau_{t,*} \subset S_t$. Thus the plaques for the associated pleated surfaces are also in proximity, in the sense of the last portion of the proof of Corollary 1.1; it then follows that their bending cocycles are also close. Thus, by the Main Theorem, the bending cocycle of either is approximated by the one given by the Prym differential associated to the limiting configuration of $\mathrm{Op}(t^2q)$. The result now follows from Lemma 5.5 and Proposition 6.9. Notice that since $\operatorname{Im}\lambda_{\mathrm{SW}}$ has zero periods on saddle connections of the horizontal foliation, the choice of a possible maximalization of Λ_q^h is irrelevant.

6.3.4 Refined estimate In section, we refine the estimate Corollary 6.8 of the previous section and so prove an improvement of Corollary 1.3. We show the proposition.

Proposition 6.10 *Let $q \in \mathcal{SQD}^*(X)$, and let $u_t\colon \widetilde{X} \to \mathbb{H}^3$ be the $\mathscr{P}(t^2q)$–equivariant harmonic map with Hopf differential $q_t = \mathrm{Hopf}(u_t)$. Then*

$$\lim_{t\to+\infty} \|t^{-2}q_t - q\|_1 = O(t^{-1}).$$

Proof In outline, the argument begins as previously by using results (Theorems 6.6 and 6.7) of Dumas in [15] and [16] to show that t^2q, $\mathrm{Hopf}(v_t)$ and $\phi_{\mathrm{HF}}(\Lambda_t)$ are both of order $O(t^2)$ and within an order of $O(t)$ of one another. Then, for a properly chosen element $[\gamma] \in \pi_1(\Sigma)$, we estimate hyperbolic translation lengths $\tau_{\mathbb{H}^3}(\rho_t(\gamma))$ in two ways: first as a nearly geodesic path on the bent hyperbolic surface $S(t^2q)$, and next as a nearly geodesic path on the image in $\mathbb{H}^3$ of X by the harmonic map u_t. In both cases, for the arcs we will consider, the images of the arcs are controlled well by the intersection numbers of the arcs with the vertical measured foliations of the Hopf differentials — this is the content of Propositions 2.29 and 2.30 — and some hyperbolic geometry then asserts that the common translation length must then be predicted by the intersection numbers. This forces that those intersection numbers are close for a large family of curve classes, which are enough to in turn imply that the Hopf differentials are within a controlled error of each other.

The first step is to recall the results of Dumas on the L^1 norms of the differences of some quadratic differentials.

Theorems 6.6 and 6.7, together show that

$$\|\mathrm{Hopf}(X, S(t^2q)) - t^2q\| \leqslant O(t).$$

In particular, we may begin with the relatively weak estimate $\|\mathrm{Hopf}(X, S(t^2q))\| \asymp t^2$. On the other hand, we may use Propositions 2.29 and 2.30 to get good control on the images of a robust set of arcs $\gamma \subset \Sigma$. For example, represent $[\gamma] \in \pi_1(\Sigma)$ by a curve γ which is quasitransverse to the vertical foliation of $\mathrm{Hopf}(X, S(t^2q))$, as well as (i) piecewise vertical and horizontal with respect to that differential and also (ii) vertical near the zeroes of that differential: it is routine that this can be accomplished by simply modifying the geodesic representative of $[\gamma]$ in the flat singular metric defined by $|\mathrm{Hopf}(X, S(t^2q))|$. Then Propositions 2.29 and 2.30 assert that the image of a vertical arc through a zero is nearly a geodesic arc of some fixed positive (and finite) length, and moreover, that the image of γ is an arc on $S(t^2q)$ comprising those geodesics of uniformly bounded length meeting at images of zeroes of v_t and connected by arcs which have geodesic curvature at most $O(e^{-ct^2})$ and have length given by

$$\ell_{S(t^2q)}(v_t(\gamma_{\mathrm{hor}})) = i(\gamma, \mathcal{F}^v_{\mathrm{Hopf}(X,S(t^2q))}) + O(e^{-ct^2}).$$

Thus, by the Morse lemma in elementary hyperbolic geometry, because on the hyperbolic surface, the arc $v_t(\gamma)$ comprises long nearly geodesic arcs connected, at angles bounded away from zero, by nearly geodesic arcs of uniformly bounded length, the $S(t^2q)$–geodesic representative of $[\gamma]$ has length given by

$$\ell_{S(t^2q)}([\gamma]) = i(\gamma, \mathcal{F}^v_{\mathrm{Hopf}(X,S(t^2q))}) + O(1),$$

and moreover, outside neighborhoods of uniform size of the v_t–images of the zeroes, lies exponentially (in t^2) close to the horizontal geodesic *lamination* defined by $\mathrm{Hopf}(X, S(t^2q))$.

Of course, this is the length on a surface, so if we want to promote this estimate to an estimate of the ρ_t–translation length of $[\gamma]$, we need to consider the image $w_t(\gamma)$ after the isometry f_t. We will need to worry about curves γ which are poorly positioned with respect to a fold, so we now restrict to curve classes $[\gamma]$ which may be represented by polygonal quasitransverse arcs which also contain no vertical saddle connections; later on, we will see that these represent is a sufficiently diverse collection of elements of the fundamental group that suffice to determine the relevant Hopf differentials.

Now, by (6-8), we see that the difference $\|t^{-2}\,\mathrm{Hopf}(X, S(t^2q)) - t^{-2}\phi_{\mathsf{HF}}(\Lambda(t^2q))\|$ of normalized differentials is of order $O(t^{-1})$. Thus the corresponding laminations make an increasingly shallow angle with one another, or expressed in a way that is better for our purposes, if τ is any sufficiently split train track that carries $\Lambda(t^2q))$, then both $v_t(\gamma)$ and the horizontal geodesic *lamination* defined by $\mathrm{Hopf}(X, S(t^2q))$ meet τ at angles comparable to $O(t^{-1})$. But that train track has an image under f_t that carries the pleating lamination $\Lambda(t^2q))$ as the bending lamination for the pleated surface $S(t^2q))$.

Now, suppose we are focusing on a curve γ and γ contains a subarc $k \subset \gamma$ that connects a pair of zeroes of $\mathrm{Hopf}(X, S(t^2 q))$: it is possible that that arc k has bending with respect to the bending cocycle that is not bounded away from π. If such an arc is purely vertical, then the translation length between the w_t–image of its endpoints could be arbitrarily small, and the geodesic in $\mathbb{H}^3$ representing the ρ_t–image of $[\gamma]$ might be far away from the f_t–image of $v_t(\gamma)$. On the other hand, if all of the subarcs connecting zeroes of $\mathrm{Hopf}(X, S(t^2 q))$ have horizontal segments, then the f_t–image of those subarcs, since they have been sheared by an amount comparable to t^2 along a lamination nearly parallel to Λ_t, relative to their endpoint, will then be mapped to arcs in $\mathbb{H}^3$ that make only a shallow angle with the bending lamination Λ_t.

Thus, in that case, the v_t–image $v_t(\gamma)$ of γ will, after composition with f_t, may be seen to comprise some nearly geodesic arcs of uniformly bounded length arising from the vertical arcs of γ near the zeroes of $\mathrm{Hopf}(X, S(t^2 q))$ together with some very long, nearly geodesic arcs of length $i(\gamma, \mathcal{F}^v_{\mathrm{Hopf}(X,S(t^2 q))}) + O(e^{-ct^2})$, with only some shallow breaks of angle $O(t^{-1})$ at points far removed from their endpoints where they cross the bending lamination Λ_t.

Thus, again by hyperbolic geometry and using that the geodesic arcs make only a shallow angle with the lamination, the ρ_t–geodesic representative of such an arc $[\gamma]$ has length

$$\ell_{\mathbb{H}^3}([\gamma]) = i(\gamma, \mathcal{F}^v_{\mathrm{Hopf}(X,S(t^2 q))}) + O(1). \tag{6-11}$$

(This $\mathbb{H}^3$–geodesic representative of $[\gamma]$ also closely shadows the w_t–images of the $\mathrm{Hopf}(X, S(t^2 q))$–horizontal portions of the arc γ, but we will not need that in the sequel.)

We now turn to the map u_t. Again, we can find, for a large collection of curve classes $[\gamma]$, a representative of $[\gamma]$ that is well-positioned with respect to $\mathrm{Hopf}(u_t)$, ie it is vertical near the zeroes of $\mathrm{Hopf}(u_t)$, always quasitransverse to the vertical foliation of $\mathrm{Hopf}(u_t)$, comprising arcs that are alternately horizontal and vertical, and containing no vertical saddle connections. The image of the lift of this curve is, again by Propositions 2.29 and 2.30, an arc in $\mathbb{H}^3$ comprising images of vertical arcs that are of uniformly bounded length meeting orthogonally images of horizontal arcs that have exponentially small geodesic curvature and length given by

$$\ell_{\mathbb{H}^3}(u_t(\gamma_{\mathrm{hor}})) = i(\gamma, \mathcal{F}^v_{\mathrm{Hopf}(u_t)}) + O(e^{-ct^2}).$$

Then, again elementary hyperbolic geometry provides that the geodesic representative of $u_t([\gamma])$ lies close the images of the horizontal segments and has length

$$\ell_{\mathbb{H}^3}([\gamma]) = i(\gamma, \mathcal{F}^v_{\mathrm{Hopf}(u_t)}) + O(1). \tag{6-12}$$

We next compare equations (6-11) and (6-12) for curves γ that meet our conditions for both Hopf v_t and $\mathrm{Hopf}(u_t)$. We find that the Hopf differentials for v_t and u_t have intersection numbers, with the curve classes γ that we have considered that meet the conditions for both holomorphic differentials, that agree up to $O(1)$.

We next point out that the collection of these curve classes that meet the conditions for both Hopf v_t and Hopf$(u_t)$are sufficient for determining the vertical foliations of Hopf(v_t) and Hopf(u_t). Indeed, in the case of Hopf(v_t), we began the construction of γ by considering geodesics in the metric $|\operatorname{Hopf}(v_t)|$, and then adjusting the paths between zeroes. A subcollection of these initially chosen geodesics provides enough paths between zeroes to provide a triangulation of the surface Σ, from which the intersection numbers with the arcs suffice to determine the vertical measured foliation of, say, Hopf(v_t). When we exclude some curves that contain vertical saddle connections, we will inevitably lose the immediate means to find that those arcs have zero intersection number with the vertical foliation. To recover such information, we begin with such an arc and follow a horizontal leaf on the surface until it returns to the vertical arc, near its initial point. By either doing surgery on the original curve by adding this long horizontal segment to vertical arcs that emanate from the original pair of zeroes or, alternatively, just joining the endpoints of this horizontal arc, we find two simple curves whose intersection numbers, together with the intersection numbers obtained from the other curves in our distinguished class, determine the vertical measured foliation of Hopf(v_t). We undertake a similar process for choices of curve classes for Hopf(u_t).

It remains to compare equations (6-11) and (6-12). There is an obvious issue to address as these equations apparently refer to collections of curves that are defined independently. On the other hand, it is possible to find curve classes as in the previous paragraph that are simultaneously in the distinguished classes for both Hopf(v_t) and Hopf(u_t). The easiest cases in which to see this are when Hopf(v_t) and Hopf(u_t) are projectively equal, in which case the assertion just follows from the construction in the previous paragraph, and when they are transverse. In that latter case, we note that it is possible to realize both vertical foliations as horizontal and vertical foliations of a quadratic differential on the same surface. Then on that surface, we again triangulate the surface using arcs that are saddle connections for neither foliation, replacing any original choice of arc with a curve as in the previous paragraph, this time chosen to be at some angle with respect to both foliations. In particular, for a saddle connection, we remove a small subarc of the saddle connection and replace it with a long arc from the other foliation: the resulting arc may be replaced by an arc transverse to both foliations. The result of all these constructions is a collection of arcs which can be assembled into a collection of curves whose intersection numbers determine both Hopf(v_t) and Hopf(u_t). In the case where the vertical foliations agree on some subsurface and are transverse on another, we apply the two cases just discussed on each subsurface. In that case, we also replace any vertical subsurface boundary leaf with pairs of curves with no vertical saddle connections as described earlier.

Comparing equations (6-11) and (6-12) for the common collection of curve classes shows that

$$i(\gamma, \mathcal{F}^v_{\operatorname{Hopf}(u_t)}) = i(\gamma, \mathcal{F}^v_{\operatorname{Hopf}(X, S(t^2 q))}) + O(1).$$

Thus, since intersection numbers for a quadratic differential are computed as integrals involving $\sqrt{q}$, and our relation above holds for a class of curves whose intersection numbers may determine the quadratic differential, we see that the normalized Hopf differentials Hopf(v_t) and Hopf(u_t) agree up to O(1). Thus, applying (6-7) and (6-8), we see that

$$\|\mathrm{Hopf}(u_t) - t^2 q\| \leqslant O(t). \qquad \square$$

References

[1] **D G L Allegretti**, *Voros symbols as cluster coordinates*, J. Topol. 12 (2019) 1031–1068 MR Zbl

[2] **D G L Allegretti**, *Stability conditions, cluster varieties, and Riemann–Hilbert problems from surfaces*, Adv. Math. 380 (2021) art. id. 107610 MR Zbl

[3] **A Beauville**, **M S Narasimhan**, **S Ramanan**, *Spectral curves and the generalised theta divisor*, J. Reine Angew. Math. 398 (1989) 169–179 MR Zbl

[4] **M Bestvina**, *Degenerations of the hyperbolic space*, Duke Math. J. 56 (1988) 143–161 MR Zbl

[5] **F Bonahon**, *Shearing hyperbolic surfaces, bending pleated surfaces and Thurston's symplectic form*, Ann. Fac. Sci. Toulouse Math. 5 (1996) 233–297 MR Zbl

[6] **F Bonahon**, *Transverse Hölder distributions for geodesic laminations*, Topology 36 (1997) 103–122 MR Zbl

[7] **B H Bowditch**, *Group actions on trees and dendrons*, Topology 37 (1998) 1275–1298 MR Zbl

[8] **R D Canary**, **D B A Epstein**, **P Green**, *Notes on notes of Thurston*, from "Analytical and geometric aspects of hyperbolic space" (D B A Epstein, editor), Lond. Math. Soc. Lect. Note Ser. 111, Cambridge Univ. Press (1987) 3–92 MR Zbl

[9] **I M Chiswell**, *Abstract length functions in groups*, Math. Proc. Cambridge Philos. Soc. 80 (1976) 451–463 MR Zbl

[10] **K Corlette**, *Flat G–bundles with canonical metrics*, J. Differential Geom. 28 (1988) 361–382 MR Zbl

[11] **G Daskalopoulos**, **S Dostoglou**, **R Wentworth**, *Character varieties and harmonic maps to $\mathbb{R}$–trees*, Math. Res. Lett. 5 (1998) 523–533 MR Zbl

[12] **G Daskalopoulos**, **S Dostoglou**, **R Wentworth**, *On the Morgan–Shalen compactification of the $\mathrm{SL}(2,\mathbb{C})$ character varieties of surface groups*, Duke Math. J. 101 (2000) 189–207 MR Zbl

[13] **G D Daskalopoulos**, **R A Wentworth**, *Harmonic maps and Teichmüller theory*, from "Handbook of Teichmüller theory, I" (A Papadopoulos, editor), IRMA Lect. Math. Theor. Phys. 11, Eur. Math. Soc., Zürich (2007) 33–109 MR Zbl

[14] **S K Donaldson**, *Twisted harmonic maps and the self-duality equations*, Proc. Lond. Math. Soc. 55 (1987) 127–131 MR Zbl

[15] **D Dumas**, *Grafting, pruning, and the antipodal map on measured laminations*, J. Differential Geom. 74 (2006) 93–118 MR Zbl Correction in 77 (2007) 175–176

[16] **D Dumas**, *The Schwarzian derivative and measured laminations on Riemann surfaces*, Duke Math. J. 140 (2007) 203–243 MR Zbl

[17] **D Dumas**, **A Neitzke**, *Asymptotics of Hitchin's metric on the Hitchin section*, Comm. Math. Phys. 367 (2019) 127–150 MR Zbl

[18] **D Dumas**, **M Wolf**, *Polynomial cubic differentials and convex polygons in the projective plane*, Geom. Funct. Anal. 25 (2015) 1734–1798 MR Zbl

[19] **J Eells**, **J C Wood**, *Restrictions on harmonic maps of surfaces*, Topology 15 (1976) 263–266 MR Zbl

[20] **A Fathi**, **F Laudenbach**, **V Poénaru**, *Thurston's work on surfaces*, Math. Notes 48, Princeton Univ. Press (2012) MR Zbl

[21] **J D Fay**, *Theta functions on Riemann surfaces*, Lecture Notes in Math. 352, Springer (1973) MR Zbl

[22] **A Fenyes**, *A dynamical perspective on shear-bend coordinates*, preprint (2015) arXiv 1510.05757

[23] **L Fredrickson**, *Generic ends of the moduli space of* $\mathrm{SL}(n,\mathbb{C})$*–Higgs bundles*, preprint (2018) arXiv 1810.01556

[24] **L Fredrickson**, *Exponential decay for the asymptotic geometry of the Hitchin metric*, Comm. Math. Phys. 375 (2020) 1393–1426 MR Zbl

[25] **D Gaiotto**, **G W Moore**, **A Neitzke**, *Wall-crossing, Hitchin systems, and the WKB approximation*, Adv. Math. 234 (2013) 239–403 MR Zbl

[26] **M Gromov**, **R Schoen**, *Harmonic maps into singular spaces and* p*–adic superrigidity for lattices in groups of rank one*, Inst. Hautes Études Sci. Publ. Math. 76 (1992) 165–246 MR Zbl

[27] **R C Gunning**, *Lectures on Riemann surfaces*, Princeton Univ. Press (1966) MR Zbl

[28] **R C Gunning**, *Lectures on vector bundles over Riemann surfaces*, Univ. Tokyo Press (1967) MR Zbl

[29] **N J Hitchin**, *The self-duality equations on a Riemann surface*, Proc. Lond. Math. Soc. 55 (1987) 59–126 MR Zbl

[30] **L Hollands**, **A Neitzke**, *Spectral networks and Fenchel–Nielsen coordinates*, Lett. Math. Phys. 106 (2016) 811–877 MR Zbl

[31] **L Hollands**, **A Neitzke**, *Exact WKB and abelianization for the* T_3 *equation*, Comm. Math. Phys. 380 (2020) 131–186 MR Zbl

[32] **J Hubbard**, **H Masur**, *Quadratic differentials and foliations*, Acta Math. 142 (1979) 221–274 MR Zbl

[33] **K Iwaki**, **T Nakanishi**, *Exact WKB analysis and cluster algebras*, J. Phys. A 47 (2014) art. id. 474009 MR Zbl

[34] **J Jost**, *Equilibrium maps between metric spaces*, Calc. Var. Partial Differential Equations 2 (1994) 173–204 MR Zbl

[35] **J Jost**, **S-T Yau**, *Harmonic maps and group representations*, from "Differential geometry" (B Lawson, K Tenenblat, editors), Pitman Monogr. Surveys Pure Appl. Math. 52, Longman Sci. Tech., Harlow (1991) 241–259 MR Zbl

[36] **Y Kamishima**, **S P Tan**, *Deformation spaces on geometric structures*, from "Aspects of low-dimensional manifolds" (Y Matsumoto, S Morita, editors), Adv. Stud. Pure Math. 20, Kinokuniya, Tokyo (1992) 263–299 MR Zbl

[37] **T Kawai**, **Y Takei**, *Algebraic analysis of singular perturbation theory*, Transl. Math. Monogr. 227, Amer. Math. Soc., Providence, RI (2005) MR Zbl

[38] **N J Korevaar**, **R M Schoen**, *Sobolev spaces and harmonic maps for metric space targets*, Comm. Anal. Geom. 1 (1993) 561–659 MR Zbl

[39] **N J Korevaar**, **R M Schoen**, *Global existence theorems for harmonic maps to non-locally compact spaces*, Comm. Anal. Geom. 5 (1997) 333–387 MR Zbl

[40] **F Labourie**, *Existence d'applications harmoniques tordues à valeurs dans les variétés à courbure négative*, Proc. Amer. Math. Soc. 111 (1991) 877–882 MR Zbl

[41] **G Levitt**, *Foliations and laminations on hyperbolic surfaces*, Topology 22 (1983) 119–135 MR Zbl

[42] **R Mazzeo**, **J Swoboda**, **H Weiß**, **F Witt**, *Limiting configurations for solutions of Hitchin's equation*, Sémin. Théor. Spectr. Géom. 31 (2014) 91–116 Zbl

[43] **R Mazzeo**, **J Swoboda**, **H Weiss**, **F Witt**, *Ends of the moduli space of Higgs bundles*, Duke Math. J. 165 (2016) 2227–2271 MR Zbl

[44] **R Mazzeo**, **J Swoboda**, **H Weiss**, **F Witt**, *Asymptotic geometry of the Hitchin metric*, Comm. Math. Phys. 367 (2019) 151–191 MR Zbl

[45] **C Mese**, *Uniqueness theorems for harmonic maps into metric spaces*, Commun. Contemp. Math. 4 (2002) 725–750 MR Zbl

[46] **Y N Minsky**, *Harmonic maps into hyperbolic 3–manifolds*, Trans. Amer. Math. Soc. 332 (1992) 607–632 MR Zbl

[47] **Y N Minsky**, *Harmonic maps, length, and energy in Teichmüller space*, J. Differential Geom. 35 (1992) 151–217 MR Zbl

[48] **T Mochizuki**, *Asymptotic behaviour of certain families of harmonic bundles on Riemann surfaces*, J. Topol. 9 (2016) 1021–1073 MR Zbl

[49] **J W Morgan**, **P B Shalen**, *Valuations, trees, and degenerations of hyperbolic structures, I*, Ann. of Math. 120 (1984) 401–476 MR Zbl

[50] **J W Morgan**, **P B Shalen**, *Free actions of surface groups on $\mathbb{R}$–trees*, Topology 30 (1991) 143–154 MR Zbl

[51] **D Mumford**, *Prym varieties, I*, from "Contributions to analysis" (L V Ahlfors, I Kra, B Maskit, L Nirenberg, editors), Academic, New York (1974) 325–350 MR Zbl

[52] **J-P Otal**, *Le théorème d'hyperbolisation pour les variétés fibrées de dimension 3*, Astérisque 235, Soc. Math. France, Paris (1996) MR Zbl

[53] **F Paulin**, *Topologie de Gromov équivariante, structures hyperboliques et arbres réels*, Invent. Math. 94 (1988) 53–80 MR Zbl

[54] **R C Penner**, **J L Harer**, *Combinatorics of train tracks*, Ann. of Math. Stud. 125, Princeton Univ. Press (1992) MR Zbl

[55] **R M Schoen**, *Analytic aspects of the harmonic map problem*, from "Seminar on nonlinear partial differential equations" (S S Chern, editor), Math. Sci. Res. Inst. Publ. 2, Springer (1984) 321–358 MR Zbl

[56] **R K Skora**, *Splittings of surfaces*, J. Amer. Math. Soc. 9 (1996) 605–616 MR Zbl

[57] **C H Taubes**, *$\mathrm{PSL}(2;\mathbb{C})$ connections on 3–manifolds with L^2 bounds on curvature*, Camb. J. Math. 1 (2013) 239–397 MR Zbl

[58] **M Wolf**, *The Teichmüller theory of harmonic maps*, J. Differential Geom. 29 (1989) 449–479 MR Zbl

[59] **M Wolf**, *High energy degeneration of harmonic maps between surfaces and rays in Teichmüller space*, Topology 30 (1991) 517–540 MR Zbl

[60] **M Wolf**, *Harmonic maps from surfaces to $\mathbb{R}$–trees*, Math. Z. 218 (1995) 577–593 MR Zbl

[61] **M Wolf**, *On realizing measured foliations via quadratic differentials of harmonic maps to $\mathbb{R}$–trees*, J. Anal. Math. 68 (1996) 107–120 MR Zbl

Mathematisches Institut, Ruprecht-Karls-Universität Heidelberg
Heidelberg, Germany

Mathematisches Institut, Ruprecht-Karls-Universität Heidelberg
Heidelberg, Germany

Department of Mathematics, University of Maryland
College Park, MD, United States

School of Mathematics, Georgia Institute of Technology
Atlanta, GA, United States

aott@mathi.uni-heidelberg.de, swoboda@mathi.uni-heidelberg.de, raw@umd.edu, mwolf40@gatech.edu

Proposed: Simon Donaldson
Seconded: András I Stipsicz, Tomasz S Mrowka

Received: 21 November 2021
Revised: 13 February 2023

GEOMETRY & TOPOLOGY

msp.org/gt

See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2024 is US $805/year for the electronic version, and $1135/year (+$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFLOW® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

Geometry & Topology

Volume 28 Issue 7 (pages 3001–3510) 2024