PARABOLIC VECTOR BUNDLES AND HERMITIAN-YANG-MILLS CONNECTIONS OVER A RIEMANN SURFACE

JONATHAN A. PORITZ

3 August 1992

ABSTRACT. We study a certain moduli space of irreducible Hermitian-Yang-Mills connections on a unitary vector bundle over a punctured Riemann surface. The connections used have non-trivial holonomy around the punctures lying in fixed conjugacy classes of $U(n)$ and differ from each other by elements of a weighted Sobolev space; these connections give rise to parabolic bundles in the sense of Mehta and Seshadri. We show in fact that the moduli space of stable parabolic bundles can be identified with our moduli space of HYM connections, by proving that every stable bundle admits a unique unitary gauge orbit of Hermitian-Yang-Mills connections.

1. INTRODUCTION

In 1980 Mehta and Seshadri [16] set themselves the task of finding a holomorphic object whose moduli space could be identified with the moduli space of irreducible unitary representations of the fundamental group of a non-compact Riemann surface, in the same way that 15 years earlier Narasimhan and Seshadri [18] had found for compact Riemann surfaces Σ a bijection between $\mathrm{Hom}^{\mathrm{irr}}(\pi_1(\Sigma), U(n))/U(n)$ and the moduli space of stable holomorphic vector bundles on Σ of rank n and vanishing first Chern class. For the case of punctured surfaces Σ, the appropriate object turned out to be a holomorphic vector bundle on the compactified surface $\overline{\Sigma}$ along with the the following extra structure:

Definition 1.1. A *parabolic structure* for a holomorphic vector bundle E on $\overline{\Sigma}$ is a choice of a filtration $E\Big|_p = E_p^1 \supset \cdots \supset E_p^{k_p} \supset E_p^{k_p+1} = 0$ and *weights* $0 \leq \alpha_1^p < \cdots < \alpha_{k_p}^p < 1$ for each point $p \in \partial\Sigma = \overline{\Sigma} \setminus \Sigma$; the weight α_j^p is said to occur with

1991 *Mathematics Subject Classification.* Primary 53C07; Secondary 32G13, 14D20, 14H60.

Key words and phrases. Stable bundles, parabolic bundles, moduli spaces, Yang-Mills, singular connections, Einstein-Hermitian connections, Hermitian-Yang-Mills connections.

Appeared as Internat. J. Math. Vol. 4, No. 3 (1993) 467-501.

multiplicity $m_j^p = \dim E_p^j / E_p^{j+1}$. A *parabolic homomorphism* is a homomorphism $f : E \to F$ of holomorphic bundles which preserves the filtrations in the sense that $f\Big|_p (E_p^j) \subset F_p^{k+1}$ whenever $\alpha_j^p > \beta_k^p$; and we say (F, β) is a *parabolic subbundle* of (E, α) if $F \subset E$ as holomorphic bundles, the weights β are a subset of the weights α, and for each $p \in \partial\Sigma$ and each j we have $\beta_j^p = \alpha_k^p$, where k is the largest number such that $F_p^j \subset E_p^k$. Finally, the *parabolic degree* and *parabolic slope* are

$$\text{par-deg}\,(E, \alpha) = \deg E + \sum_{p \in \partial\Sigma} \sum_j m_j^p \, \alpha_j^p$$

and

$$\mu_{\text{par}}(E, \alpha) = \frac{\text{par-deg}(E, \alpha)}{\operatorname{rank} E},$$

and (E, α) is *stable* (resp. *semi-stable*) if for all proper parabolic subbundles (F, β) of (E, α) we have

$$\mu_{\text{par}}(F, \beta) < \mu_{\text{par}}(E, \alpha)$$

(resp. $\leq$), precisely as for bundles with no parabolic structure.

With these definitions in place, Mehta and Seshadri were able to use Mumford's geometric invariant theory [17] to construct an algebro-geometric moduli space for stable parabolic bundles of parabolic degree zero with fixed parabolic structure, and then to demonstrate that this space was in bijective correspondence with the moduli space of irreducible unitary representations of the fundamental group whose generators around the punctures lie in conjugacy classes of $U(n)$ determined by the parabolic structure.

The theory of parabolic bundles and representations of the fundamental group for non-compact Riemann surfaces remained at this point until very recently, while for compact manifolds a complex and beautiful theory has developed with connections to many diverse areas of mathematics. Much of this work (at least from the point of view needed here) derived its initial impetus from the work of Atiyah and Bott [1] which introduced the techniques of gauge theory to these problems. In particular, the crucial observation was made there that by fixing a Hermitian metric on a vector bundle E over a Riemann surface, and hence identifying the space of complex structures for E with the space $\mathcal{A}$ of unitary connections on E, one can interpret isomorphism classes of complex structures as orbits in $\mathcal{A}$ of the full group of bundle automorphisms, which then should be thought of as the complexification $\mathcal{G}^{\mathbb{C}}$ of the unitary gauge group $\mathcal{G}$. It is a relatively standard exercise in gauge theory to give the space of irreducible Hermitian-Yang-Mills (HYM) connections modulo gauge transformations $\mathcal{HYM}^{irr}/\mathcal{G}$ the structure of a smooth, finite dimensional manifold which can in turn be identified with the moduli space of stable bundles as soon as one shows that every $\mathcal{G}^{\mathbb{C}}$-orbit of stable connections contains precisely one $\mathcal{G}$-orbit of HYM connections. Settling the existence part of this last assertion amounts to solving a non-linear partial differential equation for a complex gauge transformation that transforms a given connection into one that is Hermitian-Yang-Mills; this was done by Donaldson when the base manifold is a compact Riemann surface [9] or complex algebraic surface [10] and by Uhlenbeck and Yau [24] over an arbitrary compact Kähler manifold (but see also [5], [12] and [20]). Once this picture is established a number of consequences can be derived for the topology of the moduli space of stable bundles by means of Morse theory, [1] and [6], or more direct infinite-dimensional arguments [7], and for the geometry of the moduli space, see for example [12] for applications to the geometry of the somewhat more complicated moduli space of Higgs bundles.

The goal of the present work is to bring these methods of gauge theory and non-linear partial differential equations to bear on the moduli space of stable parabolic bundles. The difficulties that we encounter are largely analytical in nature, stemming from the fact that we must either work over the punctured surface Σ and deal with the problems of a non-compact domain — such as elliptic operators not being Fredholm on the usual Sobolev spaces — or try to work on the compactified surface $\overline{\Sigma}$, but with constraints at the points of $\partial\Sigma$. Fortunately, there is a history of some success in doing gauge theory on non-compact manifolds by using the weighted Sobolev spaces $L^p_{k,\delta}$ (for the definitions of these spaces, see §1), due originally to Taubes [22] and used also for example in Floer homology and in "pulling-apart" arguments for Donaldson invariants. There are multiplication and embedding theorems for these spaces much like those for the unweighted Sobolev spaces on compact manifolds, but equally important for our work here are the fundamental results of Lockhart and McOwen [14] which tell us that an elliptic operator of order s which is asymptotically translation-invariant along the ends of Σ is Fredholm as a map from $L^p_{k,\delta}$ to $L^p_{k-s,\delta}$ for all but a countable number of δ; furthermore, the values of δ for which such an operator fails to be Fredholm are easy to determine, as is the change in its index as δ crosses such an inadmissible value.

To encode a parabolic structure with connections on the punctured surface Σ, we will take as a base connection (whose $(0,1)$-part gives a base complex structure) any unitary connection which near the puncture $p \in \partial\Sigma$ looks like the trivial connection coming from a fixed local trivialization of E plus the form $\mathrm{d}\theta \otimes X_p$, for θ the angular coordinate around p and X_p lying in a fixed conjugacy class in $\mathfrak{u}(n)$. The parabolic structure corresponding to such a connection has as weights at p the imaginary parts of the eigenvalues of X_p and the filtration is by eigenspaces of X_p for increasing weights. The full space of connections we must use is one of our base connections plus connection forms which lie in $L^2_{1,\delta}(\Sigma, T^*\Sigma \otimes \mathfrak{u}(E))$, and we denote it by $\mathcal{A}^\delta_X$. If we let $\mathfrak{u}_p$ be the Lie algebra of the centralizer $U_p = C_{U(n)}(\exp X_p)$, then the Lie algebra of the corresponding unitary gauge group $\mathcal{G}^\delta_U$ consists of sections which near p lie in $\mathfrak{u}_p \oplus L^2_{2,\delta}(\Sigma, \mathfrak{u}(n))$ while for the complex gauge group $\mathcal{G}^{\delta,\mathbb{C}}_G$ they must lie in $\mathfrak{g}_p \oplus L^2_{2,\delta}(\Sigma, \mathfrak{gl}(n,\mathbb{C}))$, where $\mathfrak{g}_p = \mathfrak{u}_p \otimes \mathbb{C}$.

It is now relatively straightforward to use the above-mentioned Fredholm theory in these weighted spaces to show that the space of irreducible Hermitian-Yang-Mills connections $\mathcal{HYM}_X^{\delta,irr} \subset \mathcal{A}_X^\delta$ modulo the unitary gauge group $\mathcal{G}_U^\delta$ is a smooth, finite-dimensional manifold, and to formalize the relationship between the moduli space $\mathcal{F}_X^{\delta,irr}/\mathcal{G}_U^\delta$ of irreducible flat connections (which is $\mathcal{HYM}_X^{\delta,irr}$ for the case of vanishing parabolic degree) and the moduli space of irreducible unitary representations of $\pi_1(\Sigma)$ with fixed conjugacy classes at the punctures. What remains then is to identify this moduli space of HYM connections with the moduli space of stable bundles $\mathcal{A}_X^{\delta,s}/\mathcal{G}_G^{\delta,\mathbb{C}}$, by showing that a stable parabolic bundle possesses a unique unitary gauge orbit of HYM connections. The uniqueness part of this is quite easy, while the existence part and its attendant machinery takes up much of this paper.

Our approach for existence is to take advantage of the fact, proven in [8], that the set $\mathcal{A}_X^{\delta,s}$ of connections representing stable parabolic bundles is connected. It is therefore up to us to demonstrate that the complex gauge orbit of the set of irreducible HYM connections $\mathcal{G}_G^{\delta,\mathbb{C}} \cdot \mathcal{HYM}_X^{\delta,irr}$ is both open and closed in $\mathcal{A}_X^{\delta,s}$. Openness follows from an application of the Implicit Function Theorem to the linearization of the complex action. To show that $\mathcal{G}_G^{\delta,\mathbb{C}} \cdot \mathcal{HYM}_X^{\delta,irr}$ is closed in $\mathcal{A}_X^{\delta,s}$, we characterize the set of HYM connections as those which represent direct sums of stable parabolic bundles and then can use the compactness of $\mathcal{HYM}_X^\delta/\mathcal{G}_U^\delta$ to reduce the problem essentially to showing that a limit of $\bar\partial_A$-operators each of which has a kernel must itself have a kernel.

We now briefly outline the contents of the sections that follow. In §2 we give local definitions near a single puncture and prove the basic theorem that every connection can be complex-gauged to be the trivial model connection on a small disk near the puncture. §3 contains global definitions of gauge groups and of spaces of connections and complex structures, as well as the identification of the moduli space of flat connections with the moduli space of representations. The category of parabolic bundles is introduced in §4 as arising from our spaces of complex structures and complex gauge transformations, and a basic lemma about this category is recalled. In §5 we prove a Coulomb gauge theorem near a single puncture and apply it to the regularity of HYM connections and the compactness of their moduli space. Next in §6 we prove that every HYM connection represents a direct sum of stable parabolic bundles, that the complex gauge orbit of the set of irreducible HYM connections is both open and closed in the space of stable connections, and that two HYM connections in the same complex orbit must be unitarily equivalent: together, these give our main result, Theorem 6.9.

As a final note, we wish to mention that there has been other work recently on these problems for non-compact base manifolds. In particular, Simpson in [20] proved the basic existence theorem (for stable Higgs bundles rather than bundles alone) over a class of non-compact Kähler manifolds that includes punctured Riemann surfaces, and in [21] he built on this and gave an elaboration of the idea of stable parabolic bundles on surfaces to include representations that are not unitary. Also, Corlette in

[4] studied non-Abelian Hodge theory, as the representation theory of the fundamental groups of Kähler manifolds has come to be called, and intersection cohomology with local coefficients over the complement of a divisor in a smooth projective variety over $\mathbb{C}$. Both of these authors used methods quite different from ours in their proofs of the existence theorem, namely a heat equation approach on a sequence of domains exhausting the non-compact manifold.

In addition, there has also been recent work on parabolic bundles using weighted Sobolev spaces, by Biquard [2] and Daskalopoulos and Wentworth [8]. Biquard's paper has a proof of the result of Mehta and Seshadri in a manner precisely analogous to that used by Donaldson [9] to prove Narasimhan and Seshadri's theorem on compact Riemann surfaces, and he is constrained by his approach to use the Banach spaces $L^p_{k,\delta}$ for very small values of $p > 1$ and relatively large values of δ. The work [8], on the other hand, uses an analytic framework in the Hilbert spaces $L^2_{k,\delta}$ with $\delta > 0$ very small, like the one in this paper, and quoting the theorem of Mehta and Seshadri in this context (*i.e.*, part of our Theorem 6.9) it goes on to study "geometric quantization", or bundles over the moduli space of parabolic bundles.

Acknowledgments. The author would like to thank his thesis advisor Karen Uhlenbeck, for suggesting this problem and for her guidance during its resolution, as well as Kevin Corlette, whose insight and generosity were vital throughout the course of this work. In addition, conversations with Carlos Simpson and Alessandra Iozzi were extremely helpful at many points.

2. Definitions and a local model theorem near a single puncture

Our local model for a Riemann surface near a puncture will be obtained by removing the origin from the unit disk $\Delta \subset \mathbb{C}$; we will denote this by Δ^* and on it mainly use the coordinate $z = re^{i\theta}$. To do analysis on Δ^*, it is more convenient however to work in the half-infinite cylinder $C^+ = \Omega^+/\{w \mapsto w + 2\pi\}$, where $\Omega^+ = \{w \in \mathbb{C} \mid 0 \leq \operatorname{Re} w \leq 2\pi, \operatorname{Im} w > 0\}$, and to use the coordinate $w = x + iy$ of Ω^+. The isomorphism we will use between these two pictures is given by the map $w \mapsto z = e^{iw}$, or, in coordinates, $(x, y) \mapsto (r, \theta) = (e^{-y}, x)$. Notice that since this map is holomorphic, the $\bar{\partial}$ operator is the same in either picture.

Using the coordinates of the cylindrical picture, we recall the definition in [14] of the weighted Sobolev space $L^p_{k,\delta}(C^+)$ of functions, for $1 \leq p \leq \infty$, $k \in \mathbb{N}$, and $\delta \in \mathbb{R}$, as the closure of $C^\infty_0(\overline{C^+}) \cap C^\infty(C^+)$ in the norm

$$\|u\|_{L^p_{k,\delta}(C^+)} = \left\|e^{\delta y} u\right\|_{L^p_k(C^+)}.$$

If E is a vector bundle over Δ^*, we must first fix a trivialization in order to define the $L^p_{k,\delta}$ norm of a section as the sum of the norms of its component functions. In [14] and [3] a number of embedding and multiplication theorems for these spaces are proven; let us summarize those that we will use in the sequel.

Proposition 2.1. For the weighted Sobolev spaces $L^p_{k,\delta}$, we have the following embeddings:

- $L^{p_1}_{k_1,\delta_1} \hookrightarrow L^{p_2}_{k_2,\delta_2}$ is continuous if $k_1 > k_2$, $k_1 - \frac{2}{p_1} \geq k_2 - \frac{2}{p_2}$, and $\delta_1 \geq \delta_2$, and is compact if these inequalities are all strict; and
- $L^p_{k_1,\delta} \hookrightarrow C^{k_2}$ is continuous if $k_1 - \frac{2}{p} > k_2$ and $\delta > 0$.

Similarly, pointwise multiplication gives a continuous map $L^{p_1}_{k_1,\delta_1} \otimes L^{p_2}_{k_2,\delta_2} \to L^p_{k,\delta}$ provided that $L^{p_1}_{k_1} \otimes L^{p_2}_{k_2} \to L^p_k$ is continuous and $\delta_1 + \delta_2 \geq \delta$.

A particular instance of the above multiplication theorem is that $L^2_{2,\delta}$ is a Banach algebra when $\delta > 0$. This immediately tells us that pointwise multiplication and bracket define a Hilbert Lie group structure on the following group:

Definition 2.2. For $\delta > 0$, the group of *local complex gauge transformations with weight* δ is

$$\mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}} = \left\{ g = \mathrm{Id} + f \;\middle|\; f \in M_n(L^2_{2,\delta}),\ g(w) \in Gl(n,\mathbb{C})\ \forall w \in C^+ \right\},$$

with Lie algebra $\mathcal{L}\mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}} = M_n(L^2_{2,\delta})$. Here, when we write $\mathrm{Id} + f$, we are thinking of Id as the constant function on C^+ with value the identity matrix.

The reason for the subscript 0 in our rather cumbersome notation for this group is that all elements $g \in \mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$ are asymptotic to the identity matrix as y goes to infinity. In fact, since such a g can be written in the form $\mathrm{Id} + e^{-\delta y} f'$ for $f' \in L^2_2(C^+)$, we see that $g(x,y) \to \mathrm{Id}$ as $y \to \infty$ uniformly in x. We can get a group of functions asymptotic instead to the elements of any closed subgroup $G \subseteq Gl(n,\mathbb{C})$ by just taking the product of G and $\mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$. More formally, noticing that G acts smoothly on $\mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$ by conjugation, as the conjugation c_a by a constant matrix $a \in G$ is smooth on $M_n(L^2_{2,\delta})$ and $c_a(\mathrm{Id} + f) = \mathrm{Id} + c_a(f)$, we make the following definition:

Definition 2.3. The group of *local complex gauge transformations asymptotic to* G *with weight* δ is $\mathcal{G}^{\delta,\mathbb{C}}_{G,\mathrm{loc}} = G \ltimes_c \mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$ with Lie algebra $\mathcal{L}\mathcal{G}^{\delta,\mathbb{C}}_{G,\mathrm{loc}} = \mathfrak{g} \oplus_{\mathrm{ad}} \mathcal{L}\mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$.

An equivalent way to define these spaces is as the plain product (*i.e.*, neither direct nor semidirect product) of their factors. That is,

$$\mathcal{G}^{\delta,\mathbb{C}}_{G,\mathrm{loc}} = G \cdot \mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}} = \left\{ g = g_\infty + f \;\middle|\; g_\infty \in G,\ f \in M_n(L^2_{2,\delta}),\ g(w) \in Gl(n,\mathbb{C})\ \forall w \right\}$$

under pointwise multiplication, and

$$\mathcal{L}\mathcal{G}^{\delta,\mathbb{C}}_{G,\mathrm{loc}} = \mathfrak{g} + \mathcal{L}\mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}} = \left\{ X = X_\infty + f \;\middle|\; X_\infty \in \mathfrak{g},\ f \in M_n(L^2_{2,\delta}) \right\}$$

under pointwise bracket. Here again, in expressions like $a + f$ for a in G or $\mathfrak{g}$ and $f \in M_n(L^2_{2,\delta})$, we mean a to be considered as a constant function on C^+ with value in G or $\mathfrak{g}$.

Our last versions of the local gauge group will be the unitary ones, defined as usual to be the subgroups of elements pointwise in $U(n)$. Since the infinite dimensional part of these groups consists of maps into a linear group, the multiplication theorem for $L^2_{2,\delta}$ is all one needs to see that these are actually Hilbert Lie groups and Lie algebras; see for example the argument in §8 of [15].

Definition 2.4. The group of *local unitary gauge transformations with weight* δ is

$$\mathcal{G}^{\delta}_{0,\mathrm{loc}} = \Big\{ g = \mathrm{Id} + f \;\Big|\; f \in M_n(L^2_{2,\delta}),\ g(w) \in U(n)\ \forall w \Big\}$$

with Lie algebra

$$\mathcal{L}\mathcal{G}^{\delta}_{0,\mathrm{loc}} = \Big\{ X \in M_n(L^2_{2,\delta}) \;\Big|\; X(w) \in \mathfrak{u}(n)\ \forall w \Big\}.$$

For a closed subgroup $U \subseteq U(n)$ with Lie algebra $\mathfrak{u}$, the group of *local unitary gauge transformations asymptotic to* U *with weight* δ is

$$\begin{aligned}\mathcal{G}^{\delta}_{U,\mathrm{loc}} &= U \ltimes_c \mathcal{G}^{\delta}_{0,\mathrm{loc}} = U \cdot \mathcal{G}^{\delta}_{0,\mathrm{loc}} \\ &= \Big\{ g = g_\infty + f \;\Big|\; g_\infty \in U,\ f \in M_n(L^2_{2,\delta}),\ g(w) \in U(n)\ \forall w \Big\}\end{aligned}$$

with Lie algebra

$$\begin{aligned}\mathcal{L}\mathcal{G}^{\delta}_{U,\mathrm{loc}} &= \mathfrak{u} \oplus_{\mathrm{ad}} \mathcal{L}\mathcal{G}^{\delta}_{0,\mathrm{loc}} = \mathfrak{u} + \mathcal{L}\mathcal{G}^{\delta}_{0,\mathrm{loc}} \\ &= \Big\{ X = X_\infty + f \;\Big|\; X_\infty \in \mathfrak{u},\ f \in M_n(L^2_{2,\delta}),\ f(w) \in \mathfrak{u}(n)\ \forall w \Big\}.\end{aligned}$$

Notice here that in the same sense that a group $G \subseteq Gl(n,\mathbb{C})$ corresponding to the complexified Lie algebra $\mathfrak{g} = \mathfrak{u} \otimes \mathbb{C}$ is the complexification of $U \subseteq U(n)$, we can say that $\mathcal{G}^{\delta,\mathbb{C}}_{G,\mathrm{loc}}$ is the complexification of $\mathcal{G}^{\delta}_{U,\mathrm{loc}}$, as $\mathcal{L}\mathcal{G}^{\delta,\mathbb{C}}_{G,\mathrm{loc}} = \mathcal{L}\mathcal{G}^{\delta}_{U,\mathrm{loc}} \otimes \mathbb{C}$.

Let us now define the weighted Sobolev spaces of connections and complex structures on which these gauge groups will act. As mentioned above, to define the space of $L^p_{k,\delta}$ sections of a vector bundle E, we must first fix a trivialization of that bundle; we were implicitly working with such a trivialization in our above definitions when we wrote $M_n(L^2_{2,\delta})$ instead of $L^2_{2,\delta}(\mathrm{End}\, E)$. Since the spaces of smooth connections and complex structures are affine spaces modeled on $\Omega^1_{\Delta^*}(\mathrm{End}\, E)$ and $\Omega^{0,1}_{\Delta^*}(\mathrm{End}\, E)$, respectively, we must now also trivialize the cotangent bundle of Δ^*. For this we will use the basis $\{\mathrm{d}x, \mathrm{d}y\}$ when working with $\Omega^1_{\Delta^*}$ and, equivalently, $\{\mathrm{d}w, \mathrm{d}\bar{w}\}$ when working with $\Omega^{1,0}_{\Delta^*} \oplus \Omega^{0,1}_{\Delta^*}$. With respect to these trivializations we can make the following definitions:

Definition 2.5. If $\mathfrak{O}$ is a subset of $\mathfrak{u}(n)$ (typically a subalgebra or conjugacy class) then the space of *unitary connections on* Δ^* *asymptotic to* $\mathfrak{O}$ *with weight* δ is

$$\mathcal{A}^{\delta}_{\mathfrak{O},\mathrm{loc}} = \Big\{ \mathrm{d}_A = \mathrm{d} + \mathrm{d}x \otimes X_\infty + f \;\Big|\; X_\infty \in \mathfrak{O},\ f \in L^2_{1,\delta}(T^*\Delta^* \otimes \mathfrak{u}(E)) \Big\},$$

endowed with the subspace topology as a subset of $\mathfrak{u}(n) \times L^2_{1,\delta}(T^*\Delta^* \otimes \operatorname{End} E)$; here $\mathfrak{u}(E)$ is the subbundle of $\operatorname{End} E$ consisting of skew-adjoint endomorphisms. Likewise, for $\mathfrak{O} \subseteq \mathfrak{gl}(n, \mathbb{C})$, the space of *complex structures asymptotic to* $\mathfrak{O}$ *with weight* δ is

$$\mathcal{C}^\delta_{\mathfrak{O},\mathrm{loc}} = \left\{ \bar\partial_A = \bar\partial + \tfrac{1}{2}\mathrm{d}\bar w \otimes X_\infty + f \,\middle|\, X_\infty \in \mathfrak{O},\ f \in L^2_{1,\delta}(T^{0,1}\Delta^* \otimes \operatorname{End} E) \right\},$$

with the subspace topology. If $\mathfrak{O}$ consists of a single element X we will write $\mathcal{A}^\delta_{X,\mathrm{loc}}$ and $\mathcal{C}^\delta_{X,\mathrm{loc}}$ instead of $\mathcal{A}^\delta_{\{X\},\mathrm{loc}}$ and $\mathcal{C}^\delta_{\{X\},\mathrm{loc}}$.

As usual, there is an identification between unitary connections and complex structures, by $\mathrm{d}_A = \mathrm{d} + A^{1,0}\mathrm{d}w + A^{0,1}\mathrm{d}\bar w \mapsto \bar\partial + A^{0,1}\mathrm{d}\bar w$ and $\bar\partial_A = \bar\partial + A \mapsto \mathrm{d} + A - A^*$, which for $X \in \mathfrak{u}(n)$ gives an isomorphism $\mathcal{A}^\delta_{X,\mathrm{loc}} \cong \mathcal{C}^\delta_{X,\mathrm{loc}}$. This identification is used in the complexified action of the following proposition:

Proposition 2.6. Let $U \subseteq U(n)$ be a closed subgroup with Lie algebra $\mathfrak{u}$ and write $\mathfrak{O}$ for the $\operatorname{Ad} U$ orbit of some fixed $X \in \mathfrak{u}$ (*i.e.*, $\mathfrak{O}$ is the conjugacy class of X under conjugation by elements of U). Then the pullback

$$(g, \mathrm{d}_A) \mapsto g^{-1} \circ \mathrm{d}_A \circ g = \mathrm{d}_A + g^{-1}\mathrm{d}_A g$$

defines a smooth action of $\mathcal{G}^\delta_{U,\mathrm{loc}}$ on $\mathcal{A}^\delta_{\mathfrak{u},\mathrm{loc}}$ and, by restriction, on $\mathcal{A}^\delta_{\mathfrak{O},\mathrm{loc}}$. Similarly, $\mathcal{G}^{\delta,\mathbb{C}}_{G,\mathrm{loc}}$ acts smoothly on $\mathcal{C}^\delta_{\mathfrak{g},\mathrm{loc}}$ and on $\mathcal{C}^\delta_{\mathfrak{O},\mathrm{loc}}$ by the pullback

$$(g, \bar\partial_A) \mapsto g^{-1} \circ \bar\partial_A \circ g = \bar\partial_A + g^{-1}\bar\partial_A g,$$

where $G \subseteq Gl(n, \mathbb{C})$ has Lie algebra $\mathfrak{g}$ and $\mathfrak{O} = \operatorname{Ad} G \cdot X$ for $X \in \mathfrak{g}$. Furthermore, if $X \in \mathfrak{u}(n)$, $\mathfrak{u} = \ker \operatorname{ad}(X)$, and $\mathfrak{g} = \mathfrak{u} \otimes \mathbb{C}$, then the action of $\mathcal{G}^\delta_{U,\mathrm{loc}}$ on $\mathcal{A}^\delta_{X,\mathrm{loc}}$ can be extended to a smooth action of the complexification $\mathcal{G}^{\delta,\mathbb{C}}_{G,\mathrm{loc}}$ by the formula

$$(g, \mathrm{d}_A) = (g, \bar\partial_A + \partial_A) \mapsto \mathrm{d}_A + g^{-1}\bar\partial_A g - \left(g^{-1}\bar\partial_A g\right)^*.$$

Proof. Since the exterior derivative gives a continuous map $L^2_{2,\delta} \to L^2_{1,\delta}$, Proposition 2.1 applied to the coordinate expressions of d_A and $\bar\partial_A$ shows immediately that these actions are smooth. □

We wish next to prove a statement analogous to that of [1, p.555]; there, in order to demonstrate that every $\bar\partial_A$ operator on a vector bundle E over a Riemann surface defines a holomorphic structure on E, Atiyah and Bott show that in a small enough disk one can always find a complex gauge transformation that takes $\bar\partial_A$ to the operator $\bar\partial$ defined with respect to a trivialization of E over the disk. In the same way, we will show that any $\bar\partial_A \in \mathcal{C}^\delta_{X,\mathrm{loc}}$ defines a holomorphic structure for E over Δ^* with a singularity of a particular form, by finding an element of $\mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$ which takes $\bar\partial_A$ to $\bar\partial + \frac{1}{2}\mathrm{d}\bar w \otimes X$. More precisely, we prove the following:

Theorem 2.7. Fix $X \in \mathfrak{u}(n)$ and write $\bar\partial_X = \bar\partial + \frac{1}{2}\mathrm{d}\bar w \otimes X$; assume that $\delta > 0$ is not an element of the set $\{l + \alpha_j - \alpha_k \mid l \in \mathbb{Z}, 1 \le j, k \le n\}$, where $i\alpha_1, \dots, i\alpha_n$ are the eigenvalues of X (such a δ will be called *compatible with* $\operatorname{ad} X$). Then there is a constant $\eta > 0$ such that for any $\bar\partial_A = \bar\partial_X + A \in \mathcal{C}^\delta_{X,\mathrm{loc}}$ with $\|A\|_{L^2_{1,\delta}} < \eta$ there exists a $g \in \mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$ such that $g(\bar\partial_X) = \bar\partial_A$.

Proof. This is a simple Implicit Function Theorem argument, with the results of Lockhart and McOwen [14] giving us control over the linearized problem. As a preliminary simplification, note that it suffices to prove our theorem for X of the form $\widehat\alpha = \begin{pmatrix} i\alpha_1 & & \\ & \ddots & \\ & & i\alpha_n \end{pmatrix}$, since conjugating by the unitary matrix which takes X to $\widehat\alpha$ is a diffeomorphism of $\mathcal{C}^\delta_{\mathfrak{u}(n),\mathrm{loc}}$ which takes $\bar\partial_X$ to $\bar\partial_{\widehat\alpha}$, $\mathcal{C}^\delta_{X,\mathrm{loc}}$ to $\mathcal{C}^\delta_{\widehat\alpha,\mathrm{loc}}$, and $\mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$-orbits to $\mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$-orbits. We will therefore assume $X = \widehat\alpha$.

Now fix an element $\bar\partial_A = \bar\partial + \frac{1}{2}\mathrm{d}\bar w \otimes \widehat\alpha + A$ of $\mathcal{C}^\delta_{\widehat\alpha,\mathrm{loc}}$, where $A \in L^2_{1,\delta}(T^{0,1}\Delta^* \otimes \operatorname{End} E)$. We need to find a $g = \mathrm{Id} + f \in \mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$ such that $g(\bar\partial_{\widehat\alpha}) = \bar\partial_A$, *i.e.*, g must satisfy

$$g^{-1}\bar\partial g + \tfrac{1}{2}\mathrm{d}\bar w \otimes (g^{-1}\widehat\alpha g - \widehat\alpha) = A. \tag{2.8}$$

Equivalently, we are looking for an $f \in M_n(L^2_{2,\delta})$ such that $P(f) = A$, where

$$P(f) = (\mathrm{Id} + f)^{-1}\,\bar\partial f + \frac{1}{2}\mathrm{d}\bar w \otimes \left[(\mathrm{Id} + f)^{-1}\,\widehat\alpha\,(\mathrm{Id} + f) - \widehat\alpha\right].$$

It is convenient to assume that f is small, so that $(\mathrm{Id} + f)^{-1} = \mathrm{Id} + \sum_{j=1}^\infty (-f)^j$; then it easily follows that $P : M_n(L^2_{2,\delta}) \to M_n(L^2_{1,\delta})$ is well-defined and smooth. Hence we may apply the Implicit Function Theorem to P and conclude that it maps onto a neighborhood of the origin in $M_n(L^2_{1,\delta})$, provided that its differential at the origin

$$Lh = DP\Big|_{f_0=0}(h) = \bar\partial h + \tfrac{1}{2}\mathrm{d}\bar w \otimes [\widehat\alpha, h],$$

is surjective. Note that the neighborhood onto which P will surject will be determined by $\widehat\alpha$, and thus there will be a single constant $\eta > 0$ such that for any $A \in L^2_{1,\delta}(T^{0,1}\Delta^* \otimes \operatorname{End} E)$ with $\|A\|_{L^2_{1,\delta}} < \eta$ we will be able to solve equation (2.8).

The theory developed in [14] applies to translation-invariant elliptic operators $A = A(x, \frac{\partial}{\partial x}, \frac{\partial}{\partial y}) : L^2_{k,\delta} \to L^2_{k-s,\delta}$ provided that δ is chosen to avoid values of the form $\delta = \operatorname{Im} d$, where $d \in \mathbb{C}$ is such that

$$A(x, \frac{\partial}{\partial x}, id)u_d(x) = 0 \tag{2.9}$$

has a nontrivial solution u_d. In our case, we can write L with respect to the trivializations we are using as a matrix of operators $L_{jk} : L^2_{2,\delta} \to L^2_{1,\delta}$, where

$$L_{jk} h_{jk} = \frac{\partial h_{jk}}{\partial \bar{w}} + \frac{i}{2}(\alpha_j - \alpha_k) h_{jk}$$

for $1 \leq j,k \leq n$; the entire operator L will then satisfy the criterion of [14] if each of the L_{jk} does separately. A brief calculation shows that for L_{jk} equation (2.9) reduces to

$$\frac{\partial u_d}{\partial x} = (d + i(\alpha_k - \alpha_j)) u_d,$$

whose solution $u_d(x) = \exp\left[(d + i(\alpha_k - \alpha_j))\, x\right]$ will be invariant with respect to $x \mapsto x + 2\pi$ only if $d + i(\alpha_k - \alpha_j) \in i\mathbb{Z}$; thus we must chose δ to avoid the bad values $l + \alpha_j - \alpha_k$ for $l \in \mathbb{Z}$.

Now it is shown in [14] that for an elliptic boundary-value problem in weighted spaces such as that for the operator L with Dirichlet boundary conditions, the image of L is closed and of finite codimension (here we mean the image under L of all of $L^2_{1,\delta}$, not of just the subspace which satisfies the boundary conditions). Moreover, this codimension is shown to be precisely the dimension of that part of the kernel of the adjoint of L consisting of sections with vanishing Cauchy data on ∂C^+. Let us therefore study the kernel of the adjoint of L_{jk}. We work here with the formal or L^2 adjoint L^*_{jk}, defined by requiring that $(L_{jk}u, v)_{L^2(C^+)} = (u, L^*_{jk}v)_{L^2(C^+)}$ for all $u, v \in C^\infty_0(C^+)$, and acting as $L^*_{jk} : (L^2_{s-1,\delta})^* = L^2_{-s+1,-\delta} \to (L^2_{s,\delta})^* = L^2_{-s,-\delta}$ when we consider $L_{jk} : L^2_{s,\delta} \to L^2_{s-1,\delta}$. However, elliptic regularity (valid here in much the same way as for unweighted spaces; see [14, §2]) tells us that $\ker L^*_{jk} \subseteq \bigcap_s L^2_{s,-\delta}$, so that in particular elements of $\ker L^*_{jk}$ are smooth on C^+. To compute the kernel explicitly, we integrate by parts and immediately find that

$$L^*_{jk} u = -\frac{\partial u}{\partial w} + \frac{i}{2}(\alpha_k - \alpha_j) u.$$

It is easy to see that the $u \in C^\infty(C^+)$ which satisfy $L^*_{jk}u = 0$ are of the form $u = e^{(\alpha_j - \alpha_k)y} v$ for v an antiholomorphic function on C^+, and hence the only element of $\ker L^*$ that has vanishing Cauchy data on ∂C^+ is identically zero. Therefore L will be surjective and so, by the Implicit Function Theorem, P surjects onto a neighborhood of the origin. In fact, by choosing $\eta > 0$ sufficiently small, we will have $\left\{ A \;\middle|\; \|A\|_{L^2_{1,\delta}} < \eta \right\} \subseteq P(O)$, for O a small enough neighborhood of the origin in $M_n(L^2_{2,\delta})$ so that every $f \in O$ has $\mathrm{Id} + f \in \mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$. $\square$

We will frequently want to use this theorem without worrying directly about the norm of the connection form, so let us state a version with weaker assumptions (and conclusion).

Corollary 2.10. Let X and δ be as above. Then for every $\bar{\partial}_A \in \mathcal{C}^{\delta}_{X,\mathrm{loc}}$ there is a $g \in \mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$ and a concentric disk $\widetilde{\Delta} \subseteq \Delta$ such that $g(\bar{\partial}_X) = \bar{\partial}_A$ on $\widetilde{\Delta}^*$ and $g = \mathrm{Id}$ near $\partial\Delta$.

Proof. If we write $\psi_c(x + iy) = \psi(y - c)$ for $\psi \in C^\infty(\mathbb{R})$ a cut-off function equal to 0 when $y \leq 1$ and equal to 1 when $y \geq 2$, then we will have $A_c = A \cdot \psi_c \to 0$ in $M_n(L^2_{1,\delta})$ as $c \to \infty$. Thus for sufficiently large c we have $\|A_c\|_{L^2_{1,\delta}} < \eta$ and hence there is a $g_0 = \mathrm{Id} + f \in \mathcal{G}^{\delta,\mathbb{C}}_{0,\mathrm{loc}}$ with $g_0(\bar{\partial}_X) = \bar{\partial}_{A_c}$. But then $g = \mathrm{Id} + \psi f$ and $\widetilde{\Delta} = \left\{ z \in \Delta \;\middle|\; |z| < e^{-c-2} \right\}$ are a gauge transformation and disk with the stated properties. □

3. Global definitions and the moduli space of representations

Let us now globalize our picture to a whole punctured Riemann surface. So choose once and for all a compact Riemann surface $\overline{\Sigma}$ and let $\Sigma = \overline{\Sigma} \smallsetminus \partial\Sigma$ be the complement in $\overline{\Sigma}$ of a set of N points. Choose also identifications of disjoint neighborhoods Δ_p of each $p \in \partial\Sigma$ with our fixed Δ and, writing C^+_p for $\Delta^*_p = \Delta_p \smallsetminus \{p\}$ identified with a copy of C^+, endow Σ with a smooth metric which is the cylindrical one on each C^+_p. Fix also smooth functions $\phi_p : \Sigma \to [0,1]$ which vanish on $\Sigma \smallsetminus \Delta_p$ and equal one in a neighborhood of p. For any C^∞ vector bundle E over Σ of rank n we must also choose trivializations of $E\big|_{\Delta^*_p}$ in order to work with weighted spaces of sections of E; these trivializations will be assumed to be unitary if E has a metric K, and to arise from the basis $\{\mathrm{d}x\big|_{\Delta^*_p}, \mathrm{d}y\big|_{\Delta^*_p}\}$ or $\{\mathrm{d}w\big|_{\Delta^*_p}, \mathrm{d}\bar{w}\big|_{\Delta^*_p}\}$ if we are working with a bundle associated to the cotangent bundle.

Having made these choices, we can define the weighted Sobolev space $L^q_{k,\delta}(\Sigma, E)$ of sections of E as the completion of $C^\infty_0(\Sigma, E)$ under the norm

$$\|s\|_{L^q_{k,\delta}(\Sigma)} = \left\| e^{\delta t} s \right\|_{L^q_k(\Sigma)},$$

where t is the function

$$t = \begin{cases} \phi_p y & \text{on } C^+_p \\ 0 & \text{elsewhere} \end{cases}$$

and L^q_k is defined with respect to the chosen trivializations over the ends of Σ.

We can now globalize the definitions of the previous section, beginning with the gauge groups:

Definition 3.1. For a closed subgroup $G = \prod_{p\in\partial\Sigma} G_p \subseteq Gl(n,\mathbb{C})^N$ with Lie algebra $\mathfrak{g} = \bigoplus_{p\in\partial\Sigma} \mathfrak{g}_p \subseteq \mathfrak{gl}(n,\mathbb{C})^N$, the group of *complex gauge transformations asymptotic* *to G with weight δ* is

$$\mathcal{G}_G^{\delta,\mathbb{C}} = \Big\{ g = f + \sum_{p\in\partial\Sigma} \phi_p g_p \,\Big|\, f \in L^2_{2,\delta}(\Sigma, \operatorname{End} E),\ g_p \in G_p\ \forall p \in \partial\Sigma, \\ g(w) \in Gl(E\big|_w)\ \forall w \in \Sigma \Big\}$$

under pointwise multiplication, with Lie algebra

$$\mathcal{L}\mathcal{G}_G^{\delta,\mathbb{C}} = \Big\{ X = f + \sum_{p\in\partial\Sigma} \phi_p X_p \,\Big|\, f \in L^2_{2,\delta}(\Sigma, \operatorname{End} E),\ X_p \in \mathfrak{g}_p\ \forall p \in \partial\Sigma \Big\}$$

under pointwise bracket. If now $U = \prod_{p\in\partial\Sigma} U_p \subseteq U(n)^N$ is a closed subgroup with Lie algebra $\mathfrak{u} = \bigoplus_{p\in\partial\Sigma} \mathfrak{u}_p \subseteq \mathfrak{u}(n)^N$, then the group of *unitary gauge transformations asymptotic to U with weight δ* is the subgroup $\mathcal{G}_U^\delta \subset \mathcal{G}_U^{\delta,\mathbb{C}}$ of automorphisms which preserve K, whose Lie algebra is the subalgebra $\mathcal{L}\mathcal{G}_U^\delta \subset \mathcal{L}\mathcal{G}_G^{\delta,\mathbb{C}}$ of K-skew-adjoint endomorphisms. As before, we will use the symbols $\mathcal{G}_0^{\delta,\mathbb{C}}$, $\mathcal{L}\mathcal{G}_0^{\delta,\mathbb{C}}$, $\mathcal{G}_0^\delta$, and $\mathcal{L}\mathcal{G}_0^\delta$ for the special case where the asymptotic values G or U are just the identity.

Using the cut-off functions ϕ_p as a way of inserting the asymptotic values for $\mathcal{G}_G^{\delta,\mathbb{C}}$ has as a consequence that this group is not the semidirect product of G with $\mathcal{G}_0^{\delta,\mathbb{C}}$, as was the case for the local gauge group. Nevertheless, $\mathcal{G}_G^{\delta,\mathbb{C}}$ is still an extension of $\mathcal{G}_0^{\delta,\mathbb{C}}$ by G, and is the appropriate group to use given that we want the asymptotic values at different punctures to be in some sense independent. Of course a similar remark holds true for the unitary groups as well.

Next let us define the spaces of connections and complex structures on E over all of Σ:

Definition 3.2. For a subset $\mathfrak{O} = \prod_{p\in\partial\Sigma} \mathfrak{O}_p \subseteq \mathfrak{u}(n)^N$ the space of *unitary connections on E asymptotic to $\mathfrak{O}$ with weight δ* is

$$\mathcal{A}_{\mathfrak{O}}^\delta = \Big\{ \mathrm{d}_A = \mathrm{d}_{A_0} + f + \sum_{p\in\partial\Sigma} \phi_p \mathrm{d}x \otimes X_p \,\Big|\, f \in L^2_{1,\delta}(T^*\Sigma \otimes \mathfrak{u}(E)),\ X_p \in \mathfrak{O}_p\ \forall p \in \partial\Sigma \Big\}$$

with the topology as identified with a subspace of $\mathfrak{u}(n)^N \times L^2_{1,\delta}(T^*\Sigma \otimes \operatorname{End} E)$. Here d_{A_0} is any smooth unitary connection on E which equals the exterior derivative on each of the trivialized pieces $E\big|_{\Delta_p^*}$, the 1-form $\mathrm{d}x$ is well-defined on each of the patches Δ_p^*, and the subbundle $\mathfrak{u}(E) \subset \operatorname{End} E$ consists of the K-skew-adjoint endomorphisms of E. Similarly, for $\mathfrak{O} = \prod_{p\in\partial\Sigma} \mathfrak{O}_p \subseteq \mathfrak{gl}(n,\mathbb{C})^N$, the space of *complex structures on E asymptotic to $\mathfrak{O}$ with weight δ* is

$$\mathcal{C}_{\mathfrak{O}}^\delta = \Big\{ \bar\partial_A = \bar\partial_{A_0} + f + \sum_{p\in\partial\Sigma} \tfrac{1}{2} \phi_p \mathrm{d}\bar w \otimes X_p \,\Big|\, f \in L^2_{1,\delta}(T^{0,1}\Sigma \otimes \operatorname{End} E), \\ X_p \in \mathfrak{O}_p\ \forall p \in \partial\Sigma \Big\}$$

with the subspace topology from $\mathfrak{gl}(n,\mathbb{C})^N \times L^2_{1,\delta}(T^{0,1}\Sigma \otimes \operatorname{End} E)$; again $\bar\partial_{A_0}$ is any smooth complex structure on E which equals $\bar\partial$ on the trivialized pieces and $\mathrm{d}\bar w$ makes sense on each Δ_p^*. As before, we will use the symbols $\mathcal{A}_X^\delta$ and $\mathcal{C}_X^\delta$ when the set of asymptotic values $\mathfrak{O}$ consists of a single element X.

An easy argument shows that these spaces are independent of the choice of basepoints d_{A_0} and $\bar\partial_{A_0}$.

Having made these definitions, it is easy to see that the statements in the previous section about the local gauge groups and spaces of connections and complex structures are valid in the global context as well, when we give asymptotic information for each puncture. In particular:

- $\mathcal{G}_G^{\delta,\mathbb{C}}$, $\mathcal{LG}_G^{\delta,\mathbb{C}}$, $\mathcal{G}_U^\delta$, and $\mathcal{LG}_U^\delta$ are Hilbert Lie groups and Lie algebras.
- If $\mathfrak{g} = \mathfrak{u} \otimes \mathbb{C}$ then $\mathcal{G}_G^{\delta,\mathbb{C}}$ is the complexification of $\mathcal{G}_U^\delta$.
- For $X \in \mathfrak{u}(n)^N$, $\mathcal{A}_X^\delta$ and $\mathcal{C}_X^\delta$ are isomorphic as affine spaces.
- There is a smooth action of $\mathcal{G}_U^\delta$ on $\mathcal{A}_\mathfrak{u}^\delta$ and of $\mathcal{G}_G^{\delta,\mathbb{C}}$ on $\mathcal{C}_\mathfrak{g}^\delta$.
- If for each $p \in \partial\Sigma$, $X_p \in \mathfrak{u}(n)$, $\mathfrak{u}_p = \ker \operatorname{ad}(X_p)$, and $\mathfrak{g}_p = \mathfrak{u}_p \otimes \mathbb{C}$, then the action of $\mathcal{G}_U^\delta$ on $\mathcal{A}_X^\delta$ can be extended to an action of the complexification $\mathcal{G}_G^{\delta,\mathbb{C}}$.

The global analogue of Corollary 2.10 is sufficiently involved that we give it a separate statement:

Corollary 3.3. Fix $X = (X_p)_{p\in\partial\Sigma} \in \mathfrak{u}(n)^N$ and, for each p, let $\bar\partial_{X_p}$ be the operator $\bar\partial + \frac{1}{2}\mathrm{d}\bar w \otimes X_p$ over Δ_p^*. Then for $\delta > 0$ compatible with $\operatorname{ad} X$ — *i.e.*, compatible with each $\operatorname{ad} X_p$ — we have: for any $\bar\partial_A \in \mathcal{C}_X^\delta$ there are $g \in \mathcal{G}_0^{\delta,\mathbb{C}}$ and concentric disks $\widetilde\Delta_p \subseteq \Delta_p$ such that $g(\bar\partial_A) = \bar\partial_{X_p}$ over each $\widetilde\Delta_p^*$.

Proof. The gauge transformations given by Corollary 2.10 are the identity near the boundary of Δ, hence in our global context we can simply extend by the identity and say that for each p we have a $g_p \in \mathcal{G}_0^{\delta,\mathbb{C}}$ with $g_p\big|_{\Sigma\smallsetminus\Delta_p} = \mathrm{Id}$ and a concentric disk $\widetilde\Delta_p \subset \Delta_p$ with $g_p(\bar\partial_A)\big|_{\widetilde\Delta_p^*} = \bar\partial_{X_p}$. But for distinct $p, p' \in \partial\Sigma$, $g_{p'}$ commutes with g_p and $g_{p'}$ is the identity on $\widetilde\Delta_p$, so defining $g = \prod_{p\in\partial\Sigma} g_p$ gives the desired global gauge transformation. $\square$

At this point we wish to use the spaces we have just defined to describe a moduli problem for unitary representations of the fundamental group of the surface Σ. Of course the fundamental group of a punctured surface is free, so unless we make some restrictions on the representations we consider we will get a very simple moduli space. It will turn out that the appropriate set of representations to work with consists of those unitary representations of $\pi_1(\Sigma)$ which map the loops around the points of $\partial\Sigma$ to specified conjugacy classes in $U(n)$. To make this notion precise, fix for each $p \in \partial\Sigma$ an element $[\ell_p] \in \pi_1(\Sigma, \sigma_0)$ which is represented by a loop ℓ_p that goes once around p (*e.g.*, let ℓ_p consist of some path from σ_0 to Δ_p, followed by going once around $\partial\Delta_p$ in a counterclockwise direction, followed by returning to σ_0 along the same path as before).

Definition 3.4. Given a conjugacy class $\mathfrak{o}_p \subset U(n)$ for each $p \in \partial\Sigma$, the *representation variety with fixed conjugacy classes around the punctures of* Σ is

$$\mathcal{R}_{\mathfrak{o}} = \Big\{\rho \in \operatorname{Hom}(\pi_1(\Sigma, \sigma_0), U(n)) \;\Big|\; \rho([\ell_p]) \in \mathfrak{o}_p \;\forall p \in \partial\Sigma\Big\},$$

where $\mathfrak{o} = \prod_{p\in\partial\Sigma} \mathfrak{o}_p \subset U(n)^N$, while the quotient $\mathcal{R}_{\mathfrak{o}}/U(n)$ of $\mathcal{R}_{\mathfrak{o}}$ by the post-conjugation action of $U(n)$ is the *moduli space of representations of* $\pi_1(\Sigma, \sigma_0)$ *with fixed conjugacy classes around the punctures.*

Note that $\mathcal{R}_{\mathfrak{o}}$ is naturally a real algebraic variety, since $\pi_1(\Sigma, \sigma_0)$ can be presented with $2g + N$ generators (where g is the genus of $\overline{\Sigma}$) and one relation.

The study of this moduli space of representations lies within the realm of the geometry of vector bundles over Σ by virtue of the classical equivalence between representations of the fundamental group and flat bundles. This correspondence gives us a flat unitary vector bundle E over Σ and an identification $E\Big|_{\sigma_0} \cong \mathbb{C}^n$ for every element $\rho \in \operatorname{Hom}(\pi_1(\Sigma, \sigma_0), U(n))$ in such a way that the holonomy representation based at σ_0 is just ρ. Since two flat unitary bundles are isomorphic precisely when they have the same holonomy representation, we can think of $\mathcal{R}_{\mathfrak{o}}/U(n)$ as the space of isomorphism classes of flat unitary bundles over Σ whose holonomies around the punctures are as specified.

We can describe this moduli space with the techniques of gauge theory by pulling-back the connections from all the above-mentioned flat unitary bundles to be connections on one fixed bundle, using the fact that all unitary vector bundles over Σ are smoothly isomorphic since Σ has the homtopy type of a wedge of circles and $U(n)$ is connected. (This step goes through even over compact surfaces since all *flat* bundles are smoothly isomorphic after forgetting their flat structure.) Then if $\mathfrak{F}_{\mathfrak{o}}$ is the set of smooth, flat, unitary connections on some bundle E whose parallel transport around ℓ_p is multiplication by an element of $\mathfrak{o}_p$ for each $p \in \partial\Sigma$, and if $\mathfrak{G}$ is the group of smooth unitary automorphisms of E, then it is a classical result that the spaces $\mathcal{R}_{\mathfrak{o}}/U(n)$ and $\mathfrak{F}_{\mathfrak{o}}/\mathfrak{G}$ are homeomorphic. It suffices in fact to consider those elements of $\mathfrak{F}_{\mathfrak{o}}$ and $\mathfrak{G}$ that lie in our weighted spaces, as we will prove after a few preliminary comments.

First of all, notice that the one relation in the fundamental group of Σ tells us that $\prod_{p\in\partial\Sigma}[\ell_p]$ is a product of commutators. This means that in order for $\mathcal{R}_{\mathfrak{o}}$ to be non-empty, we must have $\prod_{p\in\partial\Sigma} \det x_p = 1$, where the x_p are representatives of the conjugacy classes $\mathfrak{o}_p$. So that we will be making statements about a moduli space which is at least potentially non-empty, we will make this assumption for the rest of the current section.

Choose now elements $X_p \in \mathfrak{u}(n)$ with $\exp(-2\pi X_p) = x_p$ and write $\mathfrak{O}_p = \operatorname{Ad} U(n) \cdot X_p$ and $\mathfrak{O} = \prod_{p\in\partial\Sigma} \mathfrak{O}_p \subset \mathfrak{u}(n)^N$. The assumption made above on the determinants of the x_p's implies that

$$1 = \prod_{p\in\partial\Sigma} \det(\exp(-2\pi X_p)) = \exp\left(-\sum_{p\in\partial\Sigma} 2\pi \operatorname{Tr} X_p\right),$$

and hence $\sum_{p\in\partial\Sigma} \operatorname{Tr} X_p = ic$ for some $c \in \mathbb{Z}$. Let E on $\overline{\Sigma}$ be a holomorphic vector bundle of rank n and first Chern class $-c$ and choose a Hermitian metric on E and unitary trivializations of E over each Δ_p. To do weighted theory on E, we simply restrict it and these trivializations to Σ.

There are two new terms we need to state our theorem. The first is another restriction we must put on the asymptotic data.

Definition 3.5. An element $X \in \mathfrak{u}(n)$ will be called *equisingular* if the dimension of the kernel of $\operatorname{ad} X$ is the same as that of $\operatorname{Ad}\exp(2\pi X)$, *i.e.*, if no two unequal eigenvalues of X differ by an element of $i\mathbb{Z}$, and $(X_p)_{p\in\partial\Sigma} \in \mathfrak{u}(n)^N$ will be equisingular if each X_p is. (We use this terminology since although the element X may be a singular or regular element of the Lie algebra $\mathfrak{u}(n)$, it is at least no more regular than its image under $\exp$.)

The last term we need describes a stronger notion than that of compatibility of δ with X.

Definition 3.6. Given $X = (X_p)_{p\in\partial\Sigma} \in \mathfrak{u}(n)^N$, we will say that $\delta \in (0,1)$ is *adapted to* X (resp. $\operatorname{ad} X$) if it is less than the smallest positive element of the set $\{l \pm \alpha_j^p \mid p \in \partial\Sigma,\ 1 \le j \le n,\ l \in \mathbb{Z}\}$ (resp. $\{l + \alpha_j^p - \alpha_k^p \mid p \in \partial\Sigma,\ 1 \le j,k \le n,\ l \in \mathbb{Z}\}$), where $i\alpha_1^p, \ldots, i\alpha_n^p$ are the eigenvalues of X_p.

Let us also fix the notation $\mathcal{F}_{\mathfrak{D}}^{\delta}$ to indicate those elements of $\mathcal{A}_{\mathfrak{D}}^{\delta}$ whose curvature is zero almost everywhere, while $\mathfrak{F}_{\mathfrak{D}}^{\delta}$ and $\mathfrak{G}_{U(n)^N}^{\delta}$ will be the smooth elements of $\mathcal{F}_{\mathfrak{D}}^{\delta}$ and $\mathcal{G}_{U(n)^N}^{\delta}$, respectively.

Theorem 3.7. Let all asymptotic data be as above and assume that $(X_p)_{p\in\partial\Sigma}$ is equisingular and δ is adapted to X and to $\operatorname{ad} X$. Then every $\mathfrak{G}$-orbit in $\mathfrak{F}_{\mathfrak{o}}$ contains elements of $\mathcal{F}_{\mathfrak{D}}^{\delta}$, and if an element of $\mathfrak{G}$ relates two such elements of $\mathcal{F}_{\mathfrak{D}}^{\delta}$ then it actually lies in $\mathcal{G}_{U(n)^N}^{\delta}$. Indeed, there is a homeomorphism of $\mathcal{F}_{\mathfrak{D}}^{\delta}/\mathcal{G}_{U(n)^N}^{\delta}$ with $\mathfrak{F}_{\mathfrak{o}}/\mathfrak{G}$, and hence with $\mathcal{R}_{\mathfrak{o}}/U(n)$.

Proof. Beginning with the first statement of the theorem, we choose a $\mathrm{d}_A \in \mathfrak{F}_{\mathfrak{o}}$ and construct a $g \in \mathfrak{G}$ such that $g(\mathrm{d}_A) \in \mathfrak{F}_{\mathfrak{D}}^{\delta}$. Concentrate first near one puncture, in the set C_p^+. If we let s be a d_A-parallel frame for $E\big|_{C_p^+}$ which is the standard frame s_0 at i, then we can write $s(w) = h_p(w)s_0$ for $h_p : \Omega^+ \to U(n)$ a function with the property that $h_p(2\pi + iy) = \exp(-2\pi X_p)h_p(iy)$. We then define our gauge transformation g on C_p^+ by $g(x+iy) = \exp(xX_p)h_p(x+iy)$, and note that this is well-defined since $g(2\pi + iy) = g(iy)$. But now on C_p^+ we have $\mathrm{d}_A = h_p^{-1}(\mathrm{d})$, so $g(\mathrm{d}_A) = gh_p^{-1}(\mathrm{d}) = \mathrm{d} + \mathrm{d}x \otimes X_p$.

It thus remains to extend our g to all of Σ. The only obstruction to doing this is the fundamental group of $U(n)$, so we will be able to extend g as long as the concatenation of the loops in $U(n)$ formed by g restricted to each of the ℓ_p is a loop which represents c in $\mathbb{Z} \cong \pi_1(U(n))$. To check this, we work with the determinants and compute:

$$\begin{aligned}\bigodot_{p\in\partial\Sigma} \det g \circ \ell_p &\simeq \{\varphi \mapsto \prod_{p\in\partial\Sigma} \exp(\varphi \operatorname{Tr} X_p) \det g_p \circ \ell_p(\varphi)\} \\ &\simeq \{\varphi \mapsto \exp(\varphi \sum_{p\in\partial\Sigma} \operatorname{Tr} X_p)\} \odot \{\varphi \mapsto \prod_{p\in\partial\Sigma} \det g_p \circ \ell_p(\varphi)\} \\ &\simeq \{\varphi \mapsto \exp(ic\varphi)\}\end{aligned}$$

where "$\odot$" denotes concatenation of curves and "$\simeq$" homotopy. The last line follows since the curve $\{\varphi \mapsto \prod_{p\in\partial\Sigma} \det g_p \circ \ell_p(\varphi)\}$ is homotopic to a concatenation of commutators of loops (representing parallel transport around the handles of Σ) and thus is homotopically trivial as $\pi_1(U(n))$ is Abelian.

Next, say that $\mathrm{d}_{A_0}, \mathrm{d}_{A_1} \in \mathfrak{F}_{\mathfrak{O}}^\delta$ and $g(\mathrm{d}_{A_0}) = \mathrm{d}_{A_1}$ for $g \in \mathfrak{G}$. We need to show that actually $g \in \mathcal{G}_{U(n)^N}^\delta$, which we will do by looking at each puncture in turn. So, over Δ_p write our connections as $\mathrm{d}_{A_j} = \mathrm{d} + \mathrm{d}x \otimes X_p + f_j$, where f_j is a smooth element of $L_{1,\delta}^2(T^*\Sigma \otimes \operatorname{End} E)$ and, without loss of generality, we may assume X_p is of the form $X_p = \begin{pmatrix} i\alpha_1^p & & \\ & \ddots & \\ & & i\alpha_n^p \end{pmatrix}$. Then as g takes d_{A_0} to d_{A_1}, it must satisfy

$$\mathrm{d}g + \mathrm{d}x \otimes [X_p, g] = gf_1 - f_0 g.$$

Let us for a moment consider the case were the f_j are both zero in a neighborhood of p, so that in this neighborhood g solves $\mathrm{d}g + \mathrm{d}x \otimes [X_p, g] = 0$. But any such solution is of the form $g_{jk}(x+iy) = h_{jk}e^{i(\alpha_j^p - \alpha_k^p)x}$, for some fixed $h \in U(n)$. Notice this will be well-defined on Δ_p^* only if $h_{jk} = 0$ whenever $\alpha_j^p - \alpha_k^p \notin \mathbb{Z}$; however, $\alpha_j^p - \alpha_k^p$ can never be a nonzero integer, since if it were, the eigenspace of $\exp(-2\pi X_p)$ corresponding to the eigenvalue $\exp(-2\pi i\alpha_j^p) = \exp(-2\pi i\alpha_k^p)$ would be of higher dimension then that of X_p corresponding to $i\alpha_j^p$ or $i\alpha_k^p$, and this contradicts our assumption that X_p is equisingular. It therefore follows that g is constant near p and thus can be smoothly extended across p, and a quick calculation shows that such functions are elements of $\mathfrak{G}_{U(n)^N}^\delta$ for any choice of $\delta < 1$.

Notice that we did not actually need the fact that d_{A_0} and d_{A_1} were flat, all we used was the action of the asymptotic part of the connections and the smoothness of g near p. So now apply Corollary 3.3 to get gauge transformations $g_j \in \mathcal{G}_0^{\delta,\mathbb{C}}$ with the property that $g_j(\mathrm{d} + \mathrm{d}x \otimes X_p) = \mathrm{d}_{A_j}$ near p. These g_j will be smooth near p, since we can bootstrap with the equation

$$\bar{\partial} g_j + \frac{1}{2} \mathrm{d}\bar{w} \otimes [X_p, g_j] = f_j.$$

But then the argument of the preceding paragraph tells us that the gauge transformation $g' = g_1^{-1} g g_0$, which takes $g_0^{-1}(\mathrm{d}_{A_0})$ to $g_1^{-1}(\mathrm{d}_{A_1})$, must lie in the weighted complex gauge group. It follows immediately that $g = g_1 g' g_0^{-1} \in \mathcal{G}^{\delta}_{U(n)^N}$.

At this point, we have proven the first two conclusions of the theorem, but only for the smooth elements of $\mathcal{F}^{\delta}_{\mathfrak{O}}$ and $\mathcal{G}^{\delta}_{U(n)^N}$; in effect, we have seen that $\mathfrak{F}^{\delta}_{\mathfrak{O}}/\mathfrak{G}^{\delta}_{U(n)^N}$ is bijective with $\mathfrak{F}_{\mathfrak{o}}/\mathfrak{G}$. To see that this bijection also holds when we consider all flat connections and gauge transformations, and is in fact a homeomorphism, requires a bit of fairly standard regularity. One starts by applying the Coulomb gauge theorem to show that every $\mathcal{G}^{\delta}_{U(n)^N}$-orbit in $\mathcal{F}^{\delta}_{\mathfrak{O}}$ contains smooth elements; this is Theorem 5.12, below, whose hypotheses require δ to be adapted to X and $\operatorname{ad} X$. Then we must show that any element of $\mathcal{G}^{\delta}_{U(n)^N}$ that relates two elements of $\mathfrak{F}^{\delta}_{\mathfrak{O}}$ is itself smooth, which is true for the larger, complex gauge group by regularity arguments using the ellipticity of $\bar{\partial}_A$ operators (see the proof of Theorem 4.9), and hence is also true in our more restricted setting here. □

At this point we could follow the argument of [1, §6] to describe moduli spaces of weighted Yang-Mills connections with non-vanishing curvature in terms of representations of a central extension of $\pi_1(\Sigma)$, but instead we will defer such considerations until the last section of this paper, where it will be much easier to prove a stronger result.

As usual, gauge theoretic moduli spaces are smooth manifolds away from the reducible connections, so it will be useful for us to have a version of the above theorem for that smooth part. If we write $\mathcal{F}^{\delta,irr}_{\mathfrak{O}}$ for the irreducible connections in $\mathcal{F}^{\delta}_{\mathfrak{O}}$ and $\mathcal{R}^{irr}_{\mathfrak{o}}$ for the irreducible representations in $\mathcal{R}_{\mathfrak{o}}$, then essentially by definition we have:

Corollary 3.8. The above homeomorphism restricts to one of $\mathcal{F}^{\delta,irr}_{\mathfrak{O}}/\mathcal{G}^{\delta}_{U(n)^N}$ with $\mathcal{R}^{irr}_{\mathfrak{o}}/U(n)$.

There is one last rather obvious transformation of the moduli space of representations which will be convenient for us to perform. For this, noting that $\operatorname{Ad} U(n) \cdot X_p \cong U(n)/U_p$ for $U_p = C_{U(n)}(\exp X_p)$, we find immediately that $\mathcal{A}^{\delta}_{\mathfrak{O}}/\mathcal{G}^{\delta}_{U(n)^N} \cong \mathcal{A}^{\delta}_X/\mathcal{G}^{\delta}_U$. Thus we can study the moduli space of representations of $\pi_1(\Sigma)$ with fixed conjugacy classes at the punctures by examining the subspace of $\mathcal{A}^{\delta}_X/\mathcal{G}^{\delta}_U$ consisting of orbits of flat connections.

4. The category of parabolic vector bundles

In this section we will see how our space of weighted complex structures $\mathcal{C}^{\delta}_X$ gives rise to vector bundles over $\overline{\Sigma}$ with parabolic structures determined by the asymptotic data X. For this we will assume that E is some arbitrary bundle with metric K and unitary trivializations over the sets Δ_p and that the asymptotic data $X = (X_p)_{p\in\partial\Sigma} \in \mathfrak{u}(n)^N$ need not be associated with a representation of $\pi_1(\Sigma)$, so in particular we know nothing about $\sum_{p\in\partial\Sigma} \operatorname{Tr} X_p$. We will write $U_p = C_{U(n)}(\exp X_p)$, as at the end of the previous section, and now also use the complexification G_p arising from $\mathfrak{g}_p = \mathfrak{u}_p \otimes \mathbb{C}$. Whenever we like, we may conjugate our local trivializations by constant unitary matrices so as to make the X_p's diagonal with their eigenvalues $i\alpha_1^p, \ldots, i\alpha_n^p$ on the diagonal; we will usually make this simplification without explicit mention. Let us make a suggestive definition.

Definition 4.1. Given asymptotic data as above, write $\beta_1^p, \ldots, \beta_{k_p}^p$ for the distinct values of the set $\{\alpha_j^p\}_{j=1}^n$ and $\tilde\alpha_j^p$ (resp. $\tilde\beta_j^p$) for the non-integral part of α_j^p (resp. β_j^p) — *i.e.*: $\tilde\alpha_j^p = \alpha_j^p - \left\lfloor \alpha_j^p \right\rfloor$, where $\lfloor x \rfloor$ is the greatest integer less than or equal to x — and order the β_j^p's so that $\tilde\beta_1^p < \cdots < \tilde\beta_{k_p}^p$. Corresponding to each $\tilde\beta_j^p$ define also the number m_j^p to be the dimension of the $i\beta_j^p$-eigenspace of X_p. Then the collection $(\tilde\beta, m) = \left(\{\tilde\beta_j^p\}, \{m_j^p\}\right)$ will be called the *parabolic data associated to X*.

We will start now by studying smooth connections, complex structures, and gauge transformations, so in particular we denote by $\mathfrak{A}_X^\delta$, $\mathfrak{C}_X^\delta$, $\mathfrak{G}_G^{\delta,\mathbb{C}}$, and $\mathfrak{G}_U^\delta$ the smooth elements of $\mathcal{A}_X^\delta$, $\mathcal{C}_X^\delta$, $\mathcal{G}_G^{\delta,\mathbb{C}}$, and $\mathcal{G}_U^\delta$, respectively.

With these conventions in place, let us construct the global analytic objects on $\overline{\Sigma}$ arising from elements of $\mathfrak{C}_X^\delta$.

Proposition 4.2. Fix X with associated parabolic data $(\tilde\beta, m)$, δ adapted to X, and any $\bar\partial_A \in \mathfrak{C}_X^\delta$. If we set

$$\Upsilon_A^E(U) = \ker\{\bar\partial_A : L^2_{-\delta}(U\cap\Sigma, E) \to L^2_{-1,-\delta}(U\cap\Sigma, E)\}$$

for U open in $\overline{\Sigma}$, then Υ_A^E is the sheaf of sections of a holomorphic vector bundle E_A over $\overline{\Sigma}$. If δ is also adapted to $\operatorname{ad} X$, then E_A possesses a natural parabolic structure with weights and multiplicities given by $(\tilde\beta, m)$ and whose filtration step $E_{A,p}^j$ is defined as follows: Let the subsheaf $\Upsilon_{A,j}^{E,p}$ of $\Upsilon_A^E\Big|_{\Delta_p}$ be

$$\Upsilon_{A,j}^{E,p}(U) = \Upsilon_A^E(U) \cap L^2_{\tilde\beta_j^p - \delta}(U \cap \Delta_p^*, E).$$

Then $E_{A,p}^j$ is the image in $E_A\Big|_p$ of the fiber over p of $\Upsilon_{A,j}^{E,p}$ under the map induced by inclusion.

Proof. As all of the action of this proposition takes place near the punctures, we concentrate on a neighborhood of one $p \in \partial\Sigma$. Also, by Corollary 3.3, we may restrict our attention to the complex structure $\bar\partial_{X_p}$ over the disk $\widetilde{\Delta}_p^*$.

The smooth sections of $E\big|_{\widetilde{\Delta}_p^*}$ in the kernel of $\bar\partial_{X_p}$ are of the form $\sum h_j r^{\alpha_j^p} e_j$ for $h_j \in \mathcal{O}_{\widetilde{\Delta}_p^*}$ and $\{e_j\}_{j=1}^n$ the standard basis coming from our trivialization. But then checking which such sections lie in $L^2_{-\delta}$, we find that $\Upsilon_A^E(\widetilde{\Delta}_p)$ is freely generated as an $\mathcal{O}_{\widetilde{\Delta}_p}$-module by sections $\{e_j^p\}_{j=1}^n$ which over $\widetilde{\Delta}_p^*$ are of the form

$$e_j^p = r^{\alpha_j^p} z^{-\lfloor \alpha_j^p \rfloor} e_j. \tag{4.3}$$

Hence Υ_A^E is a locally free sheaf over $\widetilde{\Delta}_p$ and so represents a holomorphic vector bundle $E_A\big|_{\widetilde{\Delta}_p}$.

For the parabolic structure, notice that $\left|e_j^p\right| = r^{\tilde{\alpha}_j^p}$. Hence

$$E_{A_p}^{\,j} / E_{A_p}^{\,j+1} = \operatorname{span}\{e_k^p(p) \mid \forall k \text{ with } \alpha_k^p = \beta_j^p\},$$

since $E_{A_p}^{\,j} / E_{A_p}^{\,j+1}$ corresponds to sections with growth between $r^{\tilde{\beta}_j^p - \delta}$ and $r^{\tilde{\beta}_{j+1}^p - \delta}$ (in an L^2_δ sense) and δ is adapted to $\operatorname{ad} X$. In particular, it follows that the dimensions of these successive filtration quotients agree with the parabolic multiplicities m. $\square$

Let us notice a few things about this construction. First, it is clear from the choice of the sections in (4.3) that the following are two equivalent definitions of the sheaf Υ_A^E:

- for U open in $\overline{\Sigma}$, $\Upsilon_A^E(U)$ consists of those smooth sections of $E\big|_{U\cap\Sigma}$ which are in the kernel of $\bar\partial_A$ and are bounded near each $p \in U \cap \partial\Sigma$;
- for U open in Σ,

$$\Upsilon_A^E(U) = \ker\{\bar\partial_A : L^2_\delta(U, E) \to L^2_{-1,\delta}(U, E)\},$$

  while for $p \in \partial\Sigma$, we let $g \in \mathcal{G}_0^{\delta,\mathbb{C}}$ and the sets $\widetilde{\Delta}_p$ be as given by Corollary 3.3 and set

$$\Upsilon_A^E(\widetilde{\Delta}_p) = \ker\{\bar\partial_A : L^2_\delta(\widetilde{\Delta}_p^*, E) \to L^2_{-1,\delta}(\widetilde{\Delta}_p^*, E)\} \oplus \{g^{-1}s \mid s \in \ker X_p\},$$

  where elements of $\ker X_p$ are to be viewed here as constant functions on Δ_p^*.

In fact, the first of these alternate definitions is the one used in [16].

Next, we mention that the bundle E_A can be described without using sheaves at all by saying that it is formed by gluing the fibers of E to the fibers of a trivial bundle over each $\widetilde{\Delta}_p$ and identifying the basis $\{e_j^p\}$ (defined as in (4.3)) with the trivial basis on each $\widetilde{\Delta}_p^*$.

Finally, we should point out that the above construction behaves well with respect to homomorphisms. That is:

Proposition 4.4. Let E, X, and $\bar{\partial}_A$ be as above. Fix another set of asymptotic data X' for a vector bundle E' of rank n', a δ adapted to X, X', and $Y = -X \otimes \mathrm{Id}_{n'} + \mathrm{Id}_n \otimes X'$, and a $\bar{\partial}_{A'} \in \mathfrak{C}^\delta_{X'}$. Then $\Upsilon^{E^* \otimes E'}_{A^* \otimes A'}$ is the sheaf of parabolic homomorphisms from E_A to $E'_{A'}$.

Proof. Say that the diagonal elements of X' are ξ^p_j, with distinct values which we will denote by η^p_j. As above, we work with the complex structures $\bar{\partial}_{X_p}$ and $\bar{\partial}_{X'_p}$ over $\widetilde{\Delta}^*_p$.

We claim that a homomorphism between the parabolic bundles $(E_A, \tilde{\beta})$ and $(E'_{A'}, \tilde{\eta})$ corresponds precisely to a section of the bundle that Proposition 4.2 gives us from ${E_A}^* \otimes E'_{A'}$ and the complex structure $\bar{\partial}_{A^* \otimes A'} \in \mathfrak{C}^\delta_Y$. If s is such a section, then the component of s which maps $e^p_j = r^{\alpha^p_j} z^{-\lfloor \alpha^p_j \rfloor} e_j$ to $f^p_k = r^{\xi^p_k} z^{-\lfloor \xi^p_j \rfloor} f_k$ (where $\{f_k\}_{k=1}^{n'}$ is the local trivializing basis for E') is of the form

$$s_{jk}(z) r^{\xi^p_k - \alpha^p_j} z^{-\lfloor \xi^p_k - \alpha^p_j \rfloor} e^*_j \otimes f_k = s_{jk}(z) z^{-\lfloor \tilde{\xi}^p_k - \tilde{\alpha}^p_j \rfloor} e^{p*}_j \otimes f^p_k,$$

for $s_{jk} \in \mathcal{O}_{\widetilde{\Delta}_p}$. But this vanishes at p if $\tilde{\alpha}^p_j > \tilde{\xi}^p_k$, since $s \in \Upsilon^{E^* \otimes E'}_{A^* \otimes A'}(\widetilde{\Delta}_p)$. Hence $s\Big|_p \left({E_A}^j_p\right)$ projected to ${E'_{A'}}^k_p / {E'_{A'}}^{k+1}_p$ is zero when $\tilde{\beta}^p_j > \tilde{\eta}^p_k$, which means precisely that s is a parabolic homomorphism. □

Notice that when $E = E'$ and $X = X'$, the assumptions on δ here are that it is adapted to X and to $\mathrm{ad}\, X$.

At this point we have in fact defined a functor Υ from the category $\mathcal{U}$ of unitary vector bundles over Σ with chosen unitary trivializations near $\partial\Sigma$ together with a choice of asymptotic data X, adapted δ, and a complex structure from the corresponding space $\mathfrak{C}^\delta_X$, to the category of parabolic vector bundles. This functor is not quite an isomorphism, however, since when we try to construct its inverse we have the freedom to choose any asymptotic data X with the given associated parabolic data, and any δ adapted to X. Nevertheless, by making some choices we can restrict to a category on which Υ will be an isomorphism.

Proposition 4.5. For each rank n and set of asymptotic data $X = (X_p)_{p \in \partial\Sigma}$ of the form $X_p = \begin{pmatrix} i\alpha^p_1 & & \\ & \ddots & \\ & & i\alpha^p_n \end{pmatrix}$ with numbers α^p_j having $0 \le \alpha^p_1 \le \cdots \le \alpha^p_n < 1$ fix a δ adapted to X and to $\mathrm{ad}\, X$. Let $\mathcal{U}_{[0,1)}$ be the category whose objects are unitary vector bundles over Σ with unitary trivializations near $\partial\Sigma$ together with a choice of complex structure from $\mathfrak{C}^\delta_X$, where X and δ are as just chosen, and whose morphisms are bundle homomorphisms bounded with respect to the chosen metrics and holomorphic with respect to the chosen complex structures. Let $\mathcal{P}$ be the category of parabolic bundles over $\overline{\Sigma}$ with parabolic structure supported on $\partial\Sigma$. Then $\Upsilon : \mathcal{U}_{[0,1)} \to \mathcal{P}$ is an isomorphism of categories.

Proof. We show surjectivity on objects and leave to the reader to check the rest. So given a parabolic bundle $(E, \tilde{\alpha})$, let X be as above with $\alpha_1^p = \cdots = \alpha_{m_1^p}^p = \tilde{\alpha}_1^p$, $\alpha_{m_1^p+1}^p = \cdots = \alpha_{m_1^p+m_2^p}^p = \tilde{\alpha}_2^p$, *etc.*, for each p, where m_j^p is the multiplicity of $\tilde{\alpha}_j^p$. For each p, let $\{e_j^p\}_{j=1}^n$ be a local holomorphic frame for E near p with the property that

$$E_p^j = \operatorname{span}\{e_k^p(p) \mid m_1^p + \cdots + m_{j-1}^p < k \leq n\}. \tag{4.6}$$

We then define an element of $\mathcal{U}_{[0,1)}$ to be $E\big|_\Sigma$ with any metric that makes $\{r^{-\alpha_j^p} e_j^p\}_{j=1}^n$ our local trivializing unitary frame near $\partial\Sigma$, and with complex structure given by the restriction of $\bar{\partial}_E$ to $E\big|_\Sigma$. With respect to the chosen trivializations we have $\bar{\partial}_E = \bar{\partial}_{X_p}$ near each p, so this complex structure is certainly in $\mathfrak{C}_X^\delta$. $\square$

The main concern of our work, however, is the moduli space of bundles with a fixed parabolic structure and in this case it is not necessary to restrict our asymptotic data with the requirement that each $\alpha_j^p \in [0,1)$. We will find, in fact, that we can study the moduli space of parabolic bundles with fixed parabolic structure and degree using our weighted spaces of complex structures and gauge groups, just as in the case of representations of the fundamental group. Before stating this precisely, we need to clarify the notion of the degree of a bundle over our punctured surface.

Proposition 4.7. For $X \in \mathfrak{u}(n)^N$ and E a vector bundle with metric and unitary trivialization over the ends of Σ, we will say that the *degree of E with respect to the asymptotic structure X* is

$$\deg_X E = \int_\Sigma C_1(E, \mathrm{d}_A),$$

where $C_1(E, \mathrm{d}_A) = \frac{i}{2\pi} \operatorname{Tr} F_A$ is the first Chern form of E with respect to any connection $\mathrm{d}_A \in \mathcal{A}_X^\delta$ for some δ adapted to X and $\operatorname{ad} X$. The degree defined in this way coincides with the parabolic degree of the associated parabolic bundle, so in particular $\deg_X E$ is well-defined and independent of the choice of d_A.

Proof. We prove this statement for line bundles, the general case coming from passing to the determinant bundle of an arbitrary vector bundle. Note also that if the asymptotic data we chose were zero at each puncture and if the connection d_A were the smooth base connection, then it would continue smoothly across the punctures to be a connection on $\overline{\Sigma}$ and we would have $\deg_0 E$ equal to the parabolic degree of E_A. It thus suffices to show that

$$\frac{i}{2\pi} \int_{C_p^+} \mathrm{d}(f\phi_p + i\alpha^p \phi_p \mathrm{d}x) = \alpha^p,$$

where $f \in L^2_{1,\delta}(T^*\Sigma \otimes \mathfrak{u}(E))$, since this would make the added contribution in the degree exactly equal to that of the parabolic degree. But by Stokes' Theorem we have

$$\frac{i}{2\pi}\int_{C_p^+}\mathrm{d}(f\phi_p + i\alpha^p\phi_p\mathrm{d}x) = \alpha^p - \frac{i}{2\pi}\int_{x=0}^{2\pi} f(x+iy_0) + \frac{i}{2\pi}\int_{y>y_0}\mathrm{d}f \tag{4.8}$$

for large y_0, and the last term on the right tends to zero as $y_0 \to \infty$ since

$$\int |\mathrm{d}f| = \int e^{-\delta y}\left|e^{\delta y}\mathrm{d}f\right| \le \left\|e^{-\delta y}\right\|_{L^2}\left\|e^{\delta y}\mathrm{d}f\right\|_{L^2} < \infty.$$

The same inequality applied to f instead of $\mathrm{d}f$ tells us that f is in L^1 and hence by Fubini's Theorem the function $y \mapsto \int_{x=0}^{2\pi} f(x+iy)$ is in L^1 as well. In particular, there must be a sequence of numbers y_0 tending to infinity such that the middle term of equation (4.8) tends to zero for these values. □

We can now state the promised equivalence between the moduli of parabolic bundles of fixed parabolic structure and the moduli space of our weighted complex structures.

Theorem 4.9. *Let the asymptotic data X and the stabilizer G be as at the beginning of this section, pick δ adapted to X and $\operatorname{ad} X$, and let E be any vector bundle with metric and trivialization near $\partial\Sigma$. Let $\mathcal{P}(n,d,\tilde\beta,m)$ be the set of bundles on $\overline{\Sigma}$ of rank n and parabolic degree $d = \deg_X E$ with parabolic structure over $\partial\Sigma$ given by the data $(\tilde\beta,m)$. Then Υ induces a bijection between $\mathcal{C}_X^\delta/\mathcal{G}_G^{\delta,\mathbb{C}}$ and the set of parabolic isomorphism classes in $\mathcal{P}(n,d,\tilde\beta,m)$.*

Proof. We show first the bijection between $\mathcal{P}(n,d,\tilde\beta,m)/\mathrm{Iso}$ and $\mathfrak{C}_X^\delta/\mathfrak{G}_G^{\delta,\mathbb{C}}$, and deal with regularity issues in a moment. Since the parabolic degree and weights are fixed in $\mathcal{P}(n,d,\tilde\beta,m)$, all of its elements have the same regular degree, and hence $\mathcal{P}(n,d,\tilde\beta,m)/\mathrm{Iso}$ is the same as the set of complex structures on a fixed bundle E' of that degree modulo the set of smooth bundle automorphisms that fix the parabolic structure. Note that these automorphisms are those elements of $C^\infty(\overline{\Sigma},\mathrm{Aut}(E'))$ whose values at the points $p \in \partial\Sigma$ lie in P_p, the parabolic subgroup of $Gl(n,\mathbb{C})$ whose reductive Levi factor is G_p, with respect to a fixed basis of $E'\big|_p$ of the form of that in (4.6). We may as well choose as this bundle E' the bundle that Proposition 4.2 would attach to E with the asymptotic data $Y = (Y_p)_{p\in\partial\Sigma}$, where $Y_p = \begin{pmatrix} i\lfloor\alpha_1^p\rfloor & & \\ & \ddots & \\ & & i\lfloor\alpha_n^p\rfloor\end{pmatrix}$, as it is of the correct degree.

But then any smooth complex structure for E', when expressed in the coordinates for E, will have an expression near $p \in \partial\Sigma$ of the form $\bar\partial_{X_p} + \mathrm{d}\bar w \otimes A$, where

$$|A_{jk}| \le (\mathrm{const.})\, e^{(\tilde\alpha_k^p - \tilde\alpha_j^p)y}\left|\frac{\mathrm{d}\bar z}{\mathrm{d}\bar w}\right| = (\mathrm{const.})\, e^{(\tilde\alpha_k^p - \tilde\alpha_j^p - 1)y}.$$

Given our assumption that δ is adapted to $\operatorname{ad} X$, it follows that such a complex structure actually lies in $\mathfrak{C}_X^\delta$. Similarly, the computations in Proposition 4.4 show that if we express a smooth automorphism of E' that fixes the parabolic structure in the coordinates of E, we get a smooth map from Δ_p to $Gl(n,\mathbb{C})$ which at p takes on some value from G_p, and all such maps lie in $\mathfrak{G}_G^{\delta,\mathbb{C}}$. Hence $\mathcal{P}(n,d,\tilde{\beta},m)/\mathrm{Iso}$ is bijective with $\mathfrak{C}_X^\delta/\mathfrak{G}_G^{\delta,\mathbb{C}}$.

Finally, to pass from $\mathfrak{C}_X^\delta/\mathfrak{G}_G^{\delta,\mathbb{C}}$ to $\mathcal{C}_X^\delta/\mathcal{G}_G^{\delta,\mathbb{C}}$ is exactly as in the case of a compact Riemann surface, [1, Lemmas 14.8 and 14.9]. Briefly, the argument is as follows. Given $\bar\partial_A, \bar\partial_B \in \mathfrak{C}_X^\delta$ and $g \in \mathcal{G}_G^{\delta,\mathbb{C}}$ that takes $\bar\partial_A$ to $\bar\partial_B$, one bootstraps with the equation

$$\bar\partial_A g = g(B-A),$$

where $B-A$ is the smooth $(0,1)$-form given by $\bar\partial_B - \bar\partial_A$, to show that actually $g \in \mathfrak{G}_G^{\delta,\mathbb{C}}$. On the other hand, one shows that every $\mathcal{G}_G^{\delta,\mathbb{C}}$-orbit in $\mathcal{C}_X^\delta$ contains elements in $\mathfrak{C}_X^\delta$, by using the fact that the orbit of any $\bar\partial_A \in \mathcal{C}_X^\delta$ is locally a Banach submanifold of $\mathcal{C}_X^\delta$ of finite codimension, since the differential of the $\mathcal{G}_G^{\delta,\mathbb{C}}$-action at the identity is a Fredholm operator. □

The algebraic geometers tell us that the set of isomorphism classes of stable parabolic bundles forms a nice algebraic variety, while the set of isomorphism classes of all parabolic bundles will be quite a bad space and for semi-stable bundles the quotient is only an algebraic variety when taken with respect to a slightly coarser equivalence relation than actual isomorphism. We will thus restrict our attention to the most straightforward moduli problem, that of stable bundles. So let us write $\mathcal{P}^s(n,d,\tilde{\beta},m)$ for those elements of $\mathcal{P}(n,d,\tilde{\beta},m)$ which are stable and define the set of smooth, stable, weighted complex structures $\mathfrak{C}_X^{\delta,s} \subset \mathfrak{C}_X^\delta$ to be those whose image under Υ is a stable parabolic bundle. If the set of all stable, weighted complex structures is $\mathcal{C}_X^{\delta,s} = \mathcal{G}_G^{\delta,\mathbb{C}} \cdot \mathfrak{C}_X^{\delta,s}$, then we easily have the following:

Corollary 4.10. Υ restricts to an identification of $\mathcal{C}_X^{\delta,s}/\mathcal{G}_G^{\delta,\mathbb{C}}$ with $\mathcal{P}^s(n,d,\tilde{\beta},m)/\mathrm{Iso}$.

Since the definition of stability for parabolic bundles involves excluding bundles with certain kinds of parabolic subbundles, this Corollary and indeed the functor Υ will be useful only if we have some idea of what corresponds in our picture of weighted complex structures to the sub-objects of the category $\mathcal{P}$. Now in a fixed parabolic bundle, parabolic subbundles determine, and are determined by, holomorphic subbundles. Hence it is important for us to see that holomorphic subbundles of our fixed bundle E can be given our usual asymptotic structures.

Proposition 4.11. Let $\bar\partial_A \in \mathfrak{C}_X^\delta$ and say that E' is a holomorphic subbundle of E_A. Then there exist a trivialization near $\partial\Sigma$ of $E'\big|_\Sigma$ and asymptotic data X' such that the eigenvalues of each X'_p are subsets of the eigenvalues of X_p and $E'_{A'}$ is precisely the bundle E' with the parabolic structure it inherits from E_A, where $\mathrm{d}_{A'} = \pi_{E'} \circ \mathrm{d}_A \circ \pi_{E'}$

for $\pi_{E'}$ the unitary projection onto E'. Moreover, $\pi_{E'} \in \mathcal{LG}_G^{\delta,\mathbb{C}}$ and the degrees of these subbundles obey

$$\text{par-deg}\, E' = \deg_{X'} E' = \frac{i}{2\pi}\int_\Sigma \operatorname{Tr}(\pi_{E'} F_A) - \int_\Sigma \left|\bar\partial_A \pi_{E'}\right|_K^2. \tag{4.12}$$

Proof. It is immediate from the proof of Proposition 4.4 that $\pi_{E'} \in \mathcal{LG}_G^{\delta,\mathbb{C}}$. This then tells us that the right hand side of equation (4.12) is finite, by the same reasoning as in Proposition 4.7. Furthermore, purely local calculations tell us that what is being integrated in this equation represents the first Chern form of $\mathrm{d}_{A'}$.

Let us now construct the promised asymptotic structure near a fixed $p \in \partial\Sigma$. Assume for this that X_p is diagonalized with the diagonal elements organized in maximal blocks, so that $\alpha_1^p = \cdots = \alpha_{m_1^p}^p$, $\alpha_{m_1^p+1}^p = \cdots = \alpha_{m_1^p+m_2^p}^p$, *etc.*; it follows that $G_p = Gl(m_1^p, \mathbb{C}) \times \cdots \times Gl(m_{k_p}^p, \mathbb{C}) \subseteq Gl(n, \mathbb{C})$. Write $\pi_{E'}^p \in G_p$ for the constant part of $\pi_{E'}$ near p, and notice that $\operatorname{rank} \pi_{E'}^p = n' = \operatorname{rank} E'$. Writing $\pi_{E'}^{p,j}$ for the projection of $\pi_{E'}^p$ onto the jth factor of G_p, we can choose n' orthonormal vectors $\{f_j^p\}$ in $\mathbb{C}^n$ which span the image of $\pi_{E'}^p$, with in fact the first $m_1^{p\prime} = \operatorname{rank} \pi_{E'}^{p,1}$ spanning $\mathbb{C}^{m_1^p}$ and similarly for the remaining pieces. But then on a small enough disk around p, the vectors $\{f_j^p\}$, considered now as constant sections with respect to our fixed trivialization of E, will form the needed unitary trivialization of E' near p. The asymptotic data X_p' should then clearly have the first $m_1^{p\prime}$ elements equal to the first element of X_p, *etc.* This gives now an asymptotic structure of the correct type for E', and the degree is as it should be by Proposition 4.7. □

We conclude this section by mentioning a basic fact about the category of parabolic bundles which we will need later. The proof of this is a simple exercise with holomorphic bundles (see, *e.g.*, [19]) and is omitted.

Proposition 4.13. Let (E, α) and (F, β) be two semi-stable parabolic bundles of the same rank and parabolic degree such that at least one of them is stable. If $f : (E, \alpha) \to (F, \beta)$ is a parabolic homomorphism that is not identically zero, then f is an isomorphism.

5. A punctured Coulomb gauge theorem and its consequences

A fundamental analytic result of gauge theory in Sobolev spaces without weights is Uhlenbeck's theorem [23] on the existence of local Coulomb gauges (also called Lorentz, Hodge, or just "good" gauges) for connections with curvature that is small in some L^p norm; this has many useful applications, for example to the regularity of Yang-Mills connections and the weak compactness of the space of connections with an L^p curvature bound. As we need results of this kind, we will devote this section to proving the existence of Coulomb gauges for our weighted connections in a neighborhood of a single puncture, and then to applying this to the regularity of

Hermitian-Yang-Mills connections and the compactness of the HYM moduli space. The outline of our proof of the Coulomb gauge theorem is identical to that in the original paper [23], and we content ourselves with giving details only at those points where our spaces and estimates are significantly different from those in [23].

For this theorem we will return to a notation similar to that of §2 and work in a trivial bundle of rank n over a single punctured disk Δ^*. The connections we will consider will be those in $\mathcal{A}^{\delta}_{X,\mathrm{loc}}$ for X any fixed equisingular element of $\mathfrak{u}(n)$ and δ adapted to X and to $\operatorname{ad} X$. We will write $\mathrm{d}_X = \mathrm{d} + \mathrm{d}x \otimes X$ for the base connection.

Theorem 5.1. There exist $\kappa, c > 0$ with the property that: for every $\mathrm{d}_A \in \mathcal{A}^{\delta}_{X,\mathrm{loc}}$ with $\|F_A\|_{L^2_\delta} < \kappa$ there is a unique $g \in \mathcal{G}^{\delta}_{0,\mathrm{loc}}$ such that $A' = g(\mathrm{d}_A) - \mathrm{d}_X = A'_x\mathrm{d}x + A'_y\mathrm{d}y$ obeys

(i) $\mathrm{d}_X^* A' = 0$,

(ii) $A'_y\big|_{\partial C^+} = 0$, and

(iii) $\|A'\|_{L^2_{1,\delta}} \leq c\|F_A\|_{L^2_\delta}$.

Proof. The basic structure of the proof of this theorem is quite simple: we show that the set

$$\mathcal{A}^{\delta,\kappa}_{X,\mathrm{loc}} = \left\{ \mathrm{d}_A \in \mathcal{A}^{\delta}_{X,\mathrm{loc}} \;\middle|\; \|F_A\|_{L^2_\delta} \leq \kappa \right\}$$

of connections with small curvature is connected and then that the subset consisting of connections gauge-equivalent to ones satisfying conditions *(i)–(iii)* is both open and closed. It will be convenient to state each of these steps as a separate proposition.

Proposition 5.2. $\mathcal{A}^{\delta,\kappa}_{X,\mathrm{loc}}$ is connected.

Proof. Given $\mathrm{d}_A \in \mathcal{A}^{\delta,\kappa}_{X,\mathrm{loc}}$, let $\gamma(0) = \mathrm{d}_X$ and $\gamma(s) = T^*_{-\ln s}(\mathrm{d}_A)$ for $s \in (0,1]$, where $T_\sigma(x,y) = (x, y+\sigma)$ is translation in C^+. This is a continuous curve in $\mathcal{A}^{\delta}_{X,\mathrm{loc}}$ which connects the base connection d_X to d_A. Moreover, a simple calculation shows that

$$\left\| F_{T^*_{-\ln s}(\mathrm{d}_A)} \right\|_{L^2_\delta(C^+)} \leq s^\delta \|F_A\|_{L^2_\delta(C^+)} \leq \kappa,$$

so γ lies entirely within $\mathcal{A}^{\delta,\kappa}_{X,\mathrm{loc}}$. □

Let us write $\mathcal{A}^{\delta,\kappa}_{X,C}$ for the set of connections in $\mathcal{A}^{\delta,\kappa}_{X,\mathrm{loc}}$ satisfying the Coulomb conditions *(i)–(iii)*.

Proposition 5.3. $\mathcal{G}^{\delta}_{0,\mathrm{loc}} \cdot \mathcal{A}^{\delta,\kappa}_{X,C}$ is closed as a subspace of $\mathcal{A}^{\delta,\kappa}_{X,\mathrm{loc}}$.

Proof. The equisingularity of X implies that the kernel of d_X acting on sections of the endomorphism bundle consists only of certain of the constants, which do not lie in L^2_δ for any positive δ. Thus the fundamental elliptic inequality for the operator $\mathrm{d}_X + \mathrm{d}_X^*$ gives

$$\|g - \mathrm{Id}\|_{L^2_{1,\delta}} \leq (\mathrm{const.})\|\mathrm{d}_X\,(g - \mathrm{Id})\|_{L^2_\delta} = (\mathrm{const.})\|\mathrm{d}_X g\|_{L^2_\delta}$$

for $g \in \mathcal{G}^{\delta}_{0,\mathrm{loc}}$. Once we have this inequality, the rest of the argument in [23] applies with little change. □

In fact, conditions *(i)* and *(ii)* imply condition *(iii)* for connections that are small in L^4_δ.

Proposition 5.4. There exist $\kappa', c > 0$ such that if $\mathrm{d}_{A'} = \mathrm{d}_X + A' \in \mathcal{A}^{\delta}_{X,\mathrm{loc}}$ satisfies *(i)* and *(ii)* and also obeys $\|A'\|_{L^4_\delta} \leq \kappa'$, then in fact $\|A'\|_{L^2_{1,\delta}} \leq c\|F_{A'}\|_{L^2_\delta}$.

Proof. The boundary condition *(ii)* is elliptic for the operator $\mathrm{d}_X + \mathrm{d}^*_X$ acting on $\mathfrak{u}(n)$-valued 1-forms over C^+, and after checking the criterion of [14] (equation (2.9)) we will get the usual elliptic inequality. The condition to be checked here is that the system

$$\begin{aligned}\frac{\partial A'_{y,jk}}{\partial x} - idA'_{x,jk} + i(\alpha_j - \alpha_k)A'_{y,jk} = 0\\ \frac{\partial A'_{x,jk}}{\partial x} + idA'_{y,jk} + i(\alpha_j - \alpha_k)A'_{x,jk} = 0\end{aligned}$$

should have no solutions whenever $d \in \mathbb{C}$ has $\delta = \operatorname{Im} d$; this follows easily from our assumptions on X and δ, as we leave to the reader to verify. Then to get the cleanest possible inequality, we can do a brief computation with Fourier series to see that the only elements A' of the kernel of $\mathrm{d}_X + \mathrm{d}^*_X$ which satisfy condition *(ii)* must be of the form

$$\begin{aligned}A'_{x,jk}(x,y) &= i\sum_{l=-\infty}^{\infty} c_l\left(e^{(l+\alpha_j-\alpha_k)y} + e^{-(l+\alpha_j-\alpha_k)y}\right)e^{ilx}\\ A'_{y,jk}(x,y) &= \sum_{l=-\infty}^{\infty} c_l\left(e^{(l+\alpha_j-\alpha_k)y} - e^{-(l+\alpha_j-\alpha_k)y}\right)e^{ilx}\end{aligned}$$

for constants $c_l \in \mathbb{C}$. Again our assumptions on X and δ imply that these will lie in L^2_δ only if $A' = 0$.

The conclusion then is that all A' satisfying *(i)* and *(ii)* obey the inequality

$$\|A'\|_{L^2_{1,\delta}} \leq (\text{const.})\|\mathrm{d}_X A'\|_{L^2_\delta}.$$

Using also Sobolev embedding and the equation $F_{A'} = \mathrm{d}_X A' + A' \wedge A'$ for the curvature, we find

$$\|A'\|_{L^2_{1,\delta}} \leq (\text{const.})\|F_{A'}\|_{L^2_\delta} + (\text{const.})\|A'\|_{L^4_\delta}\|A'\|_{L^2_{1,\delta}}.$$

Rearranging terms, we see that if $\|A'\|_{L^4_\delta} < \kappa'$, we have the desired inequality $\|A'\|_{L^2_{1,\delta}} \leq c\|F_{A'}\|_{L^2_\delta}$, for constants κ' and c coming from the Sobolev theorem and the elliptic inequality above. □

An easy argument in [23] derives the openness of $\mathcal{G}^{\delta}_{0,\mathrm{loc}}\cdot\mathcal{A}^{\delta,\kappa}_{X,C}$ in $\mathcal{A}^{\delta,\kappa}_{X,\mathrm{loc}}$ for sufficiently small κ from the following lemma. (Note that Lemma 2.6 of [23] remains valid in our situation since the operator P of that lemma is defined with a cut-off function vanishing near the center of the ball, hence for us there is nothing to worry about near the puncture; the application of that lemma in Lemma 2.8 of [23] is also correct in our case just as stated, since for $\mathfrak{u}(n)$-valued functions U vanishing on ∂C^+, $\mathrm{d}_X e^U\Big|_{\partial C^+}= \mathrm{d}U\Big|_{\partial C^+}$.)

Proposition 5.5. There exists $\kappa''>0$ such that for any $\mathrm{d}_{A'}=\mathrm{d}_X+A', \mathrm{d}_{A'+\lambda}=\mathrm{d}_{A'}+\lambda\in\mathcal{A}^{\delta}_{X,\mathrm{loc}}$ satisfying $\mathrm{d}_X^* A'=0$, $\|A'\|_{L^4_\delta}<\kappa''$, $\lambda_y\Big|_{\partial C^+}=0$, and $\|\lambda\|_{L^2_{1,\delta}}<\kappa''$, the equation

$$\mathrm{d}_X^*\left(g^{-1}\mathrm{d}_X g+g^{-1}\left(A'+\lambda\right)g\right)=0$$

has a solution $g\in\mathcal{G}^{\delta}_{0,\mathrm{loc}}$. The solution in fact satisfies $\frac{\partial g}{\partial y}\Big|_{\partial C^+}=0$ and depends smoothly on λ.

Proof. We look for g of the form $g=\exp(U)$ for $U\in\mathcal{L}\mathcal{G}^{\delta}_{0,\mathrm{loc}}$. Define the spaces

$$\begin{aligned}
L^{2,\nu}_{2,\delta} &= \left\{U\in L^2_{2,\delta}\otimes\mathfrak{u}(n)\ \Big|\ \tfrac{\partial U}{\partial y}\Big|_{\partial C^+}=0\right\}\\
L^{2,\nu}_{1,\delta} &= \left\{\lambda\in L^2_{1,\delta}(T^*C^+\otimes\mathfrak{u}(n))\ \Big|\ \lambda_y\Big|_{\partial C^+}=0\right\}\\
L^{2,\perp}_{\delta} &= \left\{V\in L^2_{\delta}\otimes\mathfrak{u}(n)\ \Big|\ \textstyle\int_{C^+}V_{jk}=0\ \forall j,k \text{ s.t. } \alpha_j=\alpha_k\right\},
\end{aligned}$$

and notice that $L^{2,\nu}_{2,\delta}$ is a subspace of $\mathcal{L}\mathcal{G}^{\delta}_{0,\mathrm{loc}}$. Consider now the smooth map

$$L^{2,\nu}_{2,\delta}\times L^{2,\nu}_{1,\delta}\to L^{2,\perp}_{\delta}:(U,\lambda)\mapsto \mathrm{d}_X^*\left(e^{-U}\mathrm{d}_X e^U+e^{-U}\left(A'+\lambda\right)e^U\right)$$

whose linearization at the origin is $\mathrm{d}_X^*\mathrm{d}_{A'}$. As in [23], we need only show that

$$\mathrm{d}_X^*\mathrm{d}_X: L^{2,\nu}_{2,\delta}\to L^{2,\perp}_{\delta} \tag{5.6}$$

is an isomorphism. However, by essentially the same computation as in the proof of Proposition 5.4 above, we find that the elements of the kernel of $\mathrm{d}_X^*\mathrm{d}_X$ satisfying our Neumann boundary conditions are the constant matrices with non-zero elements only in those entries (j,k) for which $\alpha_j=\alpha_k$. As these do not lie in spaces with positive weights, (5.6) is injective. These constants do lie in $L^2_{-\delta}$, and since $\mathrm{d}_X^*\mathrm{d}_X$ is formally self-adjoint (as are our boundary conditions), they form the cokernel of $\mathrm{d}_X^*\mathrm{d}_X$ in $L^2_{-\delta}$. But $L^{2,\perp}_{\delta}$ was defined precisely as the L^2-orthogonal complement of this cokernel, so (5.6) is also surjective. □

End of the proof of Theorem 5.1 The only remaining detail to consider is the claim of uniqueness for our gauge transformation. For this, we may as well assume that we start with a $\mathrm{d}_A \in \mathcal{A}^{\delta,\kappa}_{X,C}$ and have a gauge transformation $g \in \mathcal{G}^{\delta}_{0,\mathrm{loc}}$ for which $\mathrm{d}_{A'} = g(\mathrm{d}_A) \in \mathcal{A}^{\delta,\kappa}_{X,C}$ as well. Write $g = \mathrm{Id} + f$ and notice that from the equation

$$\mathrm{d}_X f = \mathrm{d}_X g = gA' - Ag = fA' - Af + A' - A$$

it follows both that the coefficient of $\mathrm{d}y$ in $\mathrm{d}_X f$ vanishes along ∂C^+ and that

$$\mathrm{d}_X^* \mathrm{d}_X f + L_{A,A'}(\mathrm{d}_X f) = 0$$

for $L_{A,A'}$ a linear operator of norm controlled by that of A and A'. By the same sorts of arguments as above, this implies that having originally chosen κ small enough, we could always ensure that such an f would have to be zero. This concludes the proof of Theorem 5.1. $\square$

As promised above, we will now apply the Coulomb gauge theorem to the regularity of Hermitian-Yang-Mills connections and the compactness of their moduli space. First let us recall the definition of HYM connections; for this, we work over our whole punctured Riemann surface with equisingular asymptotic data X and δ adapted to X and $\mathrm{ad}\, X$.

Definition 5.7. Let $\omega \in \bigcap_k L^2_{k,\delta_\omega}(\bigwedge^2 T^*\Sigma)$ for some $\delta_\omega > 0$. Then the connection $A \in \mathcal{A}^{\delta}_X$ is said to be *Hermitian-Yang-Mills with respect to* ω (or *Einstein-Hermitian* by some authors) if there exists a constant $\mu \in i\mathbb{R}$ such that

$$F_A = \omega \otimes (\mu \mathrm{Id})\,;$$

we will write $\mathcal{HYM}^{\delta}_{X,\omega}$ for the set of such connections and $\mathcal{HYM}^{\delta,irr}_{X,\omega}$ for the subset of irreducible elements of $\mathcal{HYM}^{\delta}_{X,\omega}$. When we wish to emphasize the constant μ, we will call these connections *μ-Hermitian-Yang-Mills.*

An example of a form ω we may use in this definition is the restriction to Σ of the Kähler form of a smooth metric on $\overline{\Sigma}$, which, being of the form $fe^{-y}\mathrm{d}x \wedge \mathrm{d}y$ for $f \in C^\infty(\overline{\Sigma})$, is in L^2_{k,δ_ω} for all $k \geq 0$ and $\delta_\omega < 1$. Note also that Proposition 4.7 tells us that on a fixed bundle and with fixed asymptotic data X and Kähler form ω, all HYM connections will have the same constant

$$\mu = \frac{-2\pi i}{\int_\Sigma \omega} \mu_{\mathrm{par}}$$

for μ_{par} the slope of the associated parabolic bundle.

The key local estimate needed for both regularity and compactness is the following, for which we briefly return to the notation of Theorem 5.1.

Theorem 5.8. Assume $\delta_\omega \geq \delta$ and fix $\delta' \in [\delta, 2\delta)$ with $\delta' \leq \delta_\omega$. Then there exists a $\kappa > 0$ such that for any $\mu \in i\mathbb{R}$, $k \in \mathbb{N}$ and $y_0 > 0$, we can find $c_k > 0$ with the property that: every $\mathrm{d}_A = \mathrm{d}_X + A \in \mathcal{A}^{\delta,\kappa}_{X,C}$ which is μ-HYM with respect to ω satisfies

$$\|A\|_{L^2_{k,\delta'}(C^{y_0})} \leq c_k,$$

where $C^{y_0} = \{(x, y) \in C^+ \mid y > y_0\}$.

Proof. The crucial observation here is that we can combine the HYM equations and the Coulomb gauge condition into one elliptic equation

$$(\mathrm{d}_X + \mathrm{d}_X^*)\, A + A \wedge A = \omega \otimes (\mu \mathrm{Id})\,. \tag{5.9}$$

As usual, we will bootstrap with this equation to derive our desired estimates, and, as usual, the hardest step in this procedure is the first.

To begin, choose a $y_2 \in (0, y_0)$ and a smooth cut-off function $\psi_1 : C^+ \to [0, 1]$ which equals 1 on C^{y_2} and vanishes near $y = 0$. We then have

$$\begin{aligned}(\mathrm{d}_X + \mathrm{d}_X^*)\,(\psi_1 A) + A \wedge (\psi_1 A) &= \mathrm{d}\psi_1 \wedge A - \mathrm{d}\psi_1 \lrcorner A + \psi_1 (\mathrm{d}_X + \mathrm{d}_X^*)\, A + \psi_1 (A \wedge A) \\ &= \mathrm{d}\psi_1 \wedge A - \mathrm{d}\psi_1 \lrcorner A + \psi_1 \omega \otimes (\mu \mathrm{Id})\,. \end{aligned} \tag{5.10}$$

We will approach this equation by putting it in a global setting on a full doubly-infinite cylinder with weights that are positive on one end of the cylinder and negative on the other. That is, we work in the space $H^2_{l,\delta}$ of functions on $C = \Omega / \{w \mapsto w + 2\pi\}$, where $\Omega = \{w \in \mathbb{C} \mid 0 \le \operatorname{Re} w \le 2\pi\}$, with norm

$$\|f\|_{H^2_{l,\delta}} = \left\| e^{\delta y} f \right\|_{L^2_l}.$$

The reason to use these spaces is that a first-order translation-invariant elliptic operator on C, such as $\mathrm{d}_X + \mathrm{d}_X^*$, is actually an isomorphism from $H^2_{l,\delta} \to H^2_{l-1,\delta}$ for δ which are allowed by the criterion described in §2; this is in fact the underlying principle which makes the theory of elliptic operators in weighted spaces work, see [14].

Now the right-hand side of equation (5.10) can be extended by zero on $C^- = C \setminus C^+$ to lie in H^2_{l,δ_ω} for all l, so let us concentrate on the left-hand side. For this, define $A' = E(A)$ for $E : L^2_1(C^+) \to L^2_1(C)$ an extension operator, *e.g.*:

$$A'(x, y) = \begin{cases} A(x, y) & y \ge 0 \\ 3A(x, -y) - 2A(x, -2y) & y < 0 \end{cases}$$

and notice that our multiplication theorems in weighted spaces tell us that $A' \wedge \cdot : H^2_{2,\delta_\omega} \to H^2_{1,\delta_\omega}$ is a bounded operator of norm less than or equal to (const.)$\|A\|_{L^2_{1,\delta}(C^+)}$, since on C^- we have $A' \in L^2_{1,\delta}$, while elements of H^2_{2,δ_ω} lie in $L^2_{2,-\delta_\omega}$ when restricted to C^-.

It follows that by choosing $\kappa > 0$ sufficiently small, depending only on the operator $\mathrm{d}_X + \mathrm{d}_X^* : H^2_{2,\delta_\omega} \xrightarrow{\cong} H^2_{1,\delta_\omega}$ and the constant of the Coulomb gauge theorem, we can ensure that $\mathrm{d}_X + \mathrm{d}_X^* + A' \wedge \cdot$ is a small enough perturbation to remain an isomorphism.

Thus there is a constant c_2 such that

$$\begin{aligned}\|A\|_{L^2_{2,\delta'}(C^{y_2})} &\leq \|\psi_1 A\|_{H^2_{2,\delta_\omega}}\\ &\leq (\text{const.})\|(\mathrm{d}_X + \mathrm{d}_X^*)(\psi_1 A) + A' \wedge (\psi_1 A)\|_{H^2_{1,\delta_\omega}}\\ &\leq (\text{const.})\|A\|_{L^2_{1,\delta}(C^+)} + (\text{const.})\\ &\leq c_2.\end{aligned}$$

For $k > 2$, we will proceed by induction but using equation (5.9) in a slightly more direct fashion. So assume that we have shown that $\|A\|_{L^2_{k-1,\delta}(C^{y_{k-1}})} \leq c_{k-1}$ for some $y_{k-1} \in (0, y_0)$ and some constant $c_{k-1} > 0$. Then choosing $y_k \in (y_{k-1}, y_0)$ and a smooth function $\psi_k : C^+ \to [0,1]$ with $\psi_k\big|_{C \setminus C^{y_{k-1}}} \equiv 0$ and $\psi_k\big|_{C^{y_k}} \equiv 1$, rewrite equation (5.10) as

$$(5.11) \qquad (\mathrm{d}_X + \mathrm{d}_X^*)(\psi_k A) = \mathrm{d}\psi_k \wedge A - \mathrm{d}\psi_k \lrcorner A + \psi_k \omega \otimes (\mu \mathrm{Id}) - \psi_k (A \wedge A)\,.$$

But in this range we have the bounded multiplication map $L^2_{k,\delta} \otimes L^2_{k,\delta} \to L^2_{k,\delta'}$, so the right-hand side of equation (5.11) is an element of $L^2_{k,\delta'}$ which vanishes near $y = 0$. Hence

$$\begin{aligned}\|A\|_{L^2_{k,\delta'}(C^{y_k})} &\leq \|\psi_k A\|_{H^2_{k,\delta'}}\\ &\leq (\text{const.})\|(\mathrm{d}_X + \mathrm{d}_X^*)(\psi_k A)\|_{H^2_{k-1,\delta'}}\\ &\leq (\text{const.})\|A\|_{L^2_{k-1,\delta'}(C^{y_{k-1}})} + (\text{const.}) + (\text{const.})\|A\|^2_{L^2_{k-1,\delta'}(C^{y_{k-1}})}\\ &\leq c_k\end{aligned}$$

for a constant c_k depending on c_{k-1}, ψ_k, μ, ω, and the bound for $\mathrm{d}_X + \mathrm{d}_X^*$. □

By a careful patching argument, we can move from this local regularity to the following global regularity theorem; as this is a standard theorem in the literature on compact four-manifolds (see [11]), we will omit the proof. Here and for the rest of this section we use the notation of Definition 5.7.

Theorem 5.12. If $\delta \leq \delta_\omega$, then every $\mathcal{G}_0^\delta$-orbit in $\mathcal{HYM}^\delta_{X,\omega}$ contains smooth elements.

A very similar patching argument, applied now to a whole sequence of connections, enables us to prove the compactness of the HYM moduli space. Here it is of crucial importance that the estimate in Theorem 5.8 gives us control over our connection forms in a norm which is stronger than $L^2_{1,\delta}$ both in the number of derivatives and in the weight, so that we can take advantage of the compact inclusion $L^2_{k,\delta+\epsilon} \hookrightarrow L^2_{k-1,\delta}$. Once again, we refer the interested reader to the four-dimensional literature (*e.g.*, to [11, Theorem 8.8]) for a proof of the following theorem:

Theorem 5.13. If $\delta < \delta_\omega$, then $\mathcal{HYM}^\delta_{X,\omega}/\mathcal{G}^\delta_0$ is compact.

Letting U be the centralizer of the asymptotic data, as at the end of §3, we can see that there is still a residual action of $U = \mathcal{G}^\delta_U/\mathcal{G}^\delta_0$ on $\mathcal{HYM}^\delta_{X,\omega}/\mathcal{G}^\delta_0$ whose quotient is then $\mathcal{HYM}^\delta_{X,\omega}/\mathcal{G}^\delta_U$. Hence the latter is the continuous image of a compact set and is thus compact. That is,

Corollary 5.14. If $\delta < \delta_\omega$, then $\mathcal{HYM}^\delta_{X,\omega}/\mathcal{G}^\delta_U$ is compact.

6. Hermitian-Yang-Mills connections for stable parabolic bundles

In this section we will put together the results of the previous sections and with a little more work we will prove the existence and uniqueness of a unitary gauge orbit of HYM connections in each complex gauge orbit of stable connections. Here, as before, we work with equisingular X, δ adapted to X and to $\operatorname{ad} X$, and Kähler form ω as in Definition 5.7 for $\delta_\omega > \delta$. It will be necessary for us in this section sometimes to take formal adjoints of operators with respect to the L^2_δ norm. That is, for an operator $L : L^2_{k,\delta} \to L^2_{k-s,\delta}$ we define the operator $L^{*_\delta} : L^2_{s-k,\delta} \to L^2_{-k,\delta}$ by requiring that

$$(Lu, v)_{L^2_\delta} = (u, L^{*_\delta} v)_{L^2_\delta}$$

for all $u, v \in C^\infty_0$. This is related to the usual L^2 adjoint L^* by the simple formula $L^{*_\delta} u = e^{-2\delta t} L^* \left(e^{2\delta t} u\right)$.

Let us now study the action of the complex gauge group on the set of HYM connections, which is most conveniently examined down in the moduli space. So we first restate one of the usual slice theorems for the action of the gauge group, in a form adapted to our weighted situation; this is a direct translation of a results of [22] or [15] (which is in turn a version of the standard result for compact four-manifolds) into our two-dimensional setting, and we omit the proof. Here we write $\mathcal{A}^{\delta,irr}_X$ for the irreducible connections in $\mathcal{A}^\delta_X$.

Theorem 6.1. *(cf* [15, Theorem 3.4]*)* $\mathcal{B}^{\delta,irr}_X = \mathcal{A}^{\delta,irr}_X/\mathcal{G}^\delta_0$ is a smooth Banach manifold whose tangent space at the orbit $[\mathrm{d}_A]$ of $\mathrm{d}_A \in \mathcal{A}^{\delta,irr}_X$ can be identified with

$$\begin{aligned} T_{[\mathrm{d}_A]}\mathcal{B}^{\delta,irr}_X &= \left\{ a \in L^2_{1,\delta}(T^*\Sigma \otimes \mathfrak{u}(E)) \;\middle|\; (a, \mathrm{d}_A b)_{L^2_\delta} = 0 \;\forall b \in \mathcal{L}\mathcal{G}^\delta_0 \right\} \\ &= \ker\{\mathrm{d}_A^{*_\delta} : L^2_{1,\delta}(T^*\Sigma \otimes \mathfrak{u}(E)) \to L^2_\delta(\mathfrak{u}(E))\} \end{aligned}$$

In addition, $\mathcal{HYM}^{\delta,irr}_{X,\omega}/\mathcal{G}^\delta_0$ is a smooth, finite-dimensional submanifold of $\mathcal{B}^{\delta,irr}_X$ whose tangent space at $[\mathrm{d}_A]$ can be identified with the kernel of the map

$$\mathrm{d}_A \oplus \mathrm{d}_A^{*_\delta} : L^2_{1,\delta}(T^*\Sigma \otimes \mathfrak{u}(E)) \to L^2_\delta(\textstyle\bigwedge^2 T^*\Sigma \otimes \mathfrak{u}(E)) \oplus L^2_\delta(\mathfrak{u}(E)).$$

We should note that the gauge group used in [15] (and in [8], to which we will refer in a moment), applying the theory developed in [22], appears to be different from ours. In particular, these authors begin their definition by considering those sections of the endomorphism bundle whose covariant derivative lies in $L^2_{1,\delta}$. This space of sections is shown in [22] to include some of the constants, and the gauge group is then defined as the subset consisting of those sections which are pointwise unitary transformations and whose constant part lies in the desired subgroup of $U(n)$. Our definition with the constants added to $\mathcal{G}^\delta_0$ by hand thus gives precisely the same gauge group.

Now we can start to prove our main existence result, by a connectedness argument.

Theorem 6.2. $\mathcal{G}^{\delta,\mathbb{C}}_G \cdot \mathcal{HYM}^{\delta,irr}_{X,\omega}$ is open in $\mathcal{A}^{\delta,s}_X$.

Proof. We will show that the linearization at points $[\mathrm{d}_A] \in \mathcal{HYM}^{\delta,irr}_{X,\omega}/\mathcal{G}^\delta_0$ of the action of $\mathcal{G}^{\delta,\mathbb{C}}_G$ surjects onto $T_{[\mathrm{d}_A]}\mathcal{B}^{\delta,irr}_X$, so that, by the Implicit Function Theorem, $\mathcal{G}^{\delta,\mathbb{C}}_G \cdot \mathcal{HYM}^{\delta,irr}_{X,\omega}/\mathcal{G}^\delta_0$ will contain an open neighborhood of $\mathcal{HYM}^{\delta,irr}_{X,\omega}/\mathcal{G}^\delta_0$ in $\mathcal{B}^{\delta,irr}_X$. But it will then follow that $\mathcal{G}^{\delta,\mathbb{C}}_G \cdot \mathcal{HYM}^{\delta,irr}_{X,\omega}/\mathcal{G}^\delta_0$ is a union of open sets and is thus open in $\mathcal{B}^{\delta,irr}_X$, and hence so will be $\mathcal{G}^{\delta,\mathbb{C}}_G \cdot \mathcal{HYM}^{\delta,irr}_{X,\omega}$ up in $\mathcal{A}^{\delta,irr}_X$. As we manifestly have $\mathcal{A}^{\delta,s}_X \subset \mathcal{A}^{\delta,irr}_X$, this will finish the theorem.

But now recall [1, §8] that the infinitesimal complex action, at least up in $\mathcal{A}^\delta_X$, can be written as addition of $\mathrm{d}_A u + *\mathrm{d}_A v$, for u and v in the Lie algebra of the unitary gauge group. It would thus certainly suffice for us to show that we could write

$$L^2_{1,\delta}(T^*\Sigma \otimes \mathfrak{u}(E)) = \ker(\mathrm{d}_A \oplus \mathrm{d}_A^{*_\delta}) + \mathrm{d}_A\left(\mathcal{L}\mathcal{G}^\delta_U\right) + *\mathrm{d}_A\left(\mathcal{L}\mathcal{G}^\delta_U\right).$$

Now one thing we do know, by the Fredholm alternative, is that

$$\begin{aligned} L^2_{1,\delta}(T^*\Sigma \otimes \mathfrak{u}(E)) &= \ker(\mathrm{d}_A \oplus \mathrm{d}_A^{*_\delta}) + \mathrm{d}_A\left(L^2_{2,\delta}(\mathfrak{u}(E))\right) + \mathrm{d}_A^{*_\delta}\left(L^2_{2,\delta}(\textstyle\bigwedge^2 T^*\Sigma \otimes \mathfrak{u}(E))\right) \\ &= \ker(\mathrm{d}_A \oplus \mathrm{d}_A^{*_\delta}) + \mathrm{d}_A\left(L^2_{2,\delta}(\mathfrak{u}(E))\right) + *(\mathrm{d}_A + 2\delta\mathrm{d}t)\left(L^2_{2,\delta}(\mathfrak{u}(E))\right) \end{aligned}$$

where on the second line here we used the formula

$$\mathrm{d}_A^{*_\delta} = e^{-2\delta t}\mathrm{d}_A^* e^{2\delta t} = - * e^{-2\delta t}\mathrm{d}_A e^{2\delta t} *$$

and the fact that the Hodge star is an isomorphism from $L^2_{2,\delta}(\mathfrak{u}(E))$ to $L^2_{2,\delta}(\bigwedge^2 T^*\Sigma \otimes \mathfrak{u}(E))$. We can therefore simply show that

$$\mathrm{d}_A\left(L^2_{2,\delta}(\mathfrak{u}(E))\right) + *(\mathrm{d}_A + 2\delta\mathrm{d}t)\left(L^2_{2,\delta}(\mathfrak{u}(E))\right) \subset \mathrm{d}_A\left(\mathcal{L}\mathcal{G}^\delta_U\right) + *\mathrm{d}_A\left(\mathcal{L}\mathcal{G}^\delta_U\right),$$

i.e., that for any $u, v \in L^2_{2,\delta}(\mathfrak{u}(E))$ there exist $u', v' \in L^2_{2,\delta}(\mathfrak{u}(E))$ and $\{u'_p, v'_p\}_{p\in\partial\Sigma} \subset \mathfrak{u}$ such that

$$\mathrm{d}_A(u' + \sum_{p\in\partial\Sigma}\phi_p u'_p) + *\mathrm{d}_A(v' + \sum_{p\in\partial\Sigma}\phi_p v'_p) = \mathrm{d}_A u + *(\mathrm{d}_A + 2\delta\mathrm{d}t)\, v.$$

Let us try this even setting all the $u'_p = 0$, so rewriting the above we find that what we must solve is

$$\mathrm{d}_A(u-u') + *\mathrm{d}_A(v-v') = -2\delta v *\mathrm{d}t + \sum_{p\in\partial\Sigma} *\mathrm{d}\phi_p v'_p. \tag{6.3}$$

Notice that the right side of this expression vanishes on $\Sigma \setminus \bigcup_{p\in\partial\Sigma}\Delta^*_p$, while on each Δ^*_p it lies in $L^2_{2,\delta}(T^*\Delta^*_p \otimes \mathfrak{u}(E))$. But now calculations very similar to those in Proposition 5.4 show that we will be able to solve this problem on each Δ^*_p for $u-u', v-v' \in L^2_{2,\delta}(\mathfrak{u}(E))$ which vanish on $\partial\Delta_p$, as long as the right hand side of equation (6.3) is orthogonal to the constant elements of $\mathfrak{u}$. As we can do this by adjusting v'_p appropriately, it follows that we can solve (6.3) on $\bigcup_{p\in\partial\Sigma}\Delta^*_p$ and then extend by zero to get a solution in $L^2_{2,\delta}(\Sigma, \mathfrak{u}(E))$. This concludes the proof of the surjectivity of the linearized complex action. □

In order to prove the complementary result to this theorem, we need the following consequence of our work in §4:

Theorem 6.4. The irreducible connections in $\mathcal{HYM}^\delta_{X,\omega}$ represent stable parabolic bundles, while the reducible ones are orthogonal, holomorphic direct sums of subbundles, each of which represents a stable parabolic bundle of parabolic slope equal to $(\deg_X E)/n$. In particular, the elements of $\mathcal{HYM}^\delta_{X,\omega}$ are all at least semi-stable.

Proof. This is immediate (or see [13, Theorem V.2.5]) from Propositions 4.7 and 4.11 and the definition of HYM connections. □

This enables us to prove:

Theorem 6.5. $\mathcal{G}^{\delta,\mathbb{C}}_G \cdot \mathcal{HYM}^{\delta,irr}_{X,\omega}$ is closed in $\mathcal{A}^{\delta,s}_X$.

Proof. We consider sequences $g_j \in \mathcal{G}^{\delta,\mathbb{C}}_G$ and $\mathrm{d}_{A_j} \in \mathcal{A}^{\delta,s}_X$ such that $\mathrm{d}_{B_j} = g_j\left(\mathrm{d}_{A_j}\right) \in \mathcal{HYM}^{\delta,irr}_{X,\omega}$ and $\mathrm{d}_{A_j} \to \mathrm{d}_{A_\infty} \in \mathcal{A}^{\delta,s}_X$. Passing to a subsequence and applying Corollary 5.14, we see that there exist unitary gauge transformations $u_j \in \mathcal{G}^\delta_U$ and a limiting HYM connection $\mathrm{d}_{B'_\infty} \in \mathcal{HYM}^\delta_{X,\omega}$ with $\mathrm{d}_{B'_j} = u_j\left(\mathrm{d}_{B_j}\right) \to \mathrm{d}_{B'_\infty}$ in $\mathcal{A}^\delta_X$. If we then write $g'_j = u_j g_j$, our task will be to use the existence of this sequence g'_j to find some non-zero bundle endomorphism which takes the connection d_{A_∞} to $\mathrm{d}_{B'_\infty}$. By Theorem 6.4 and Proposition 4.13 it will then follow that this endomorphism is an isomorphism and hence $\mathrm{d}_{B'_\infty}$ will in fact be stable and the limit point d_{A_∞} will indeed lie in $\mathcal{G}^{\delta,\mathbb{C}}_G \cdot \mathcal{HYM}^{\delta,irr}_{X,\omega}$. We can do this by method similar to that used in [9, Lemma 1].

Let us switch notation now, and work in the space of weighted complex structures $\mathcal{C}^\delta_X$, so that, by Proposition 4.4, a parabolic bundle homomorphism from $\left(E, \bar\partial_{A_\infty}\right)$ to $\left(E, \bar\partial_{B'_\infty}\right)$ is an element of the kernel in $L^2_{-\delta}(\Sigma, \operatorname{End} E)$ of the complex structure

$\bar\partial_{A^*_\infty\otimes B'_\infty}$ on $\operatorname{End} E = E^*\otimes E$ induced by $\bar\partial_{A_\infty}$ and $\bar\partial_{B'_\infty}$. Now if there were no such homomorphism, we would have the elliptic inequality

$$\|g\|_{L^2_{1,-\delta}} \le (\text{const.}) \left\|\bar\partial_{A^*_\infty\otimes B'_\infty} g\right\|_{L^2_{-\delta}}$$

without kernel term for all $g\in L^2_{1,-\delta}$, hence using also the Sobolev Embedding Theorem there would be a constant $c_1>0$ with

$$c_1\|g\|_{L^4_{-\delta}} \le \left\|\bar\partial_{A^*_\infty\otimes B'_\infty} g\right\|_{L^2_{-\delta}} \qquad \forall g\in L^2_{1,-\delta}.$$

However, from the affine structure of $\mathcal{C}^\delta_X$ and the Sobolev Multiplication Theorem, we have

$$\left\|\bar\partial_{A^*_\infty\otimes B'_\infty} g\right\|_{L^2_{-\delta}} \le \left\|\bar\partial_{A^*_j\otimes B'_j} g\right\|_{L^2_{-\delta}} + c_2\|g\|_{L^4_{-\delta}}\left(\|A_\infty - A_j\|_{L^4_\delta} + \left\|B'_\infty - B'_j\right\|_{L^4_\delta}\right)$$

for all j with a fixed $c_2>0$. Putting these together we find that

$$\|g\|_{L^4_{-\delta}}\left(c_1 - c_2\left(\|A_\infty - A_j\|_{L^4_\delta} + \left\|B'_\infty - B'_j\right\|_{L^4_\delta}\right)\right) \le \left\|\bar\partial_{A^*_j\otimes B'_j} g\right\|_{L^2_{-\delta}}. \tag{6.6}$$

But $\bar\partial_{A_j}\to\bar\partial_{A_\infty}$ and $\bar\partial_{B'_j}\to\bar\partial_{B'_\infty}$ in $\mathcal{C}^\delta_X$ and therefore the connection forms converge in L^4_δ. Thus, for sufficiently large j, equation (6.6) amounts to the clean elliptic inequality showing that $\bar\partial_{A^*_j\otimes B'_j}$ has no kernel in $L^2_{1,-\delta}$, which we know to be false since g_j is an element of precisely that kernel. □

Together with Theorem 6.2 we have

Corollary 6.7. $\mathcal{G}^{\delta,\mathbb{C}}_G\cdot\mathcal{HYM}^{\delta,irr}_{X,\omega}$ is a connected component of $\mathcal{A}^{\delta,s}_X$.

This Corollary becomes useful when we combine it with a result from [8], which result is based in turn on a similar argument in [7] for connections over compact Riemann surfaces.

Theorem 6.8. *(cf.* [8, Proposition 4.3]*)* If the genus g of $\overline{\Sigma}$ is at least 2, then the set $\mathcal{A}^{\delta,s}_X$ of connections representing stable parabolic bundles is path-connected.

The necessity of avoiding the genera 0 and 1 in this result is probably based on the directness of the approach in its proof. We conjecture that the result remains true over the Riemann sphere if the number N of punctures is at least three, and over the torus if $N\ge 1$.

We are now in a position to state our main result.

Theorem 6.9. Assume $g\ge 2$. Then every $\mathcal{G}^{\delta,\mathbb{C}}_G$-orbit in $\mathcal{A}^{\delta,s}_X$ contains precisely one $\mathcal{G}^\delta_U$-orbit of Hermitian-Yang-Mills connections, so that we have an identification $\mathcal{A}^{\delta,s}_X/\mathcal{G}^{\delta,\mathbb{C}}_G \cong \mathcal{HYM}^{\delta,irr}_{X,\omega}/\mathcal{G}^\delta_U$.

Proof. Corollary 6.7 and Theorem 6.8 give the existence of at least one HYM connection in each complex gauge orbit in $\mathcal{A}_X^{\delta,s}$, so let us show that in fact if $\mathrm{d}_A, \mathrm{d}_B \in \mathcal{HYM}_{X,\omega}^{\delta,irr}$ have $\mathrm{d}_B = g\,(\mathrm{d}_A)$ for $g \in \mathcal{G}_G^{\delta,\mathbb{C}}$, then d_A and d_B are also related by an element of $\mathcal{G}_U^{\delta}$. Following [12, Theorem 2.7], we move to considering g as a $\bar{\partial}_{A^*\otimes B}$-holomorphic section of the bundle $\operatorname{End} E$, which bundle has the HYM connection $\mathrm{d}_{A^*\otimes B}$. Then, writing K also for the induced Hermitian metric on $\operatorname{End} E$, we have pointwise

$$\begin{aligned}\bar{\partial}\left(\partial_{A^*\otimes B} g, g\right)_K &= \left(\bar{\partial}_{A^*\otimes B}\partial_{A^*\otimes B} g, g\right)_K - \left(\partial_{A^*\otimes B} g, \partial_{A^*\otimes B} g\right)_K \\ &= \left(F_{A^*\otimes B} g, g\right)_K - \left(\partial_{A^*\otimes B} g, \partial_{A^*\otimes B} g\right)_K \\ &= \left(\partial_{A^*\otimes B} g, \partial_{A^*\otimes B} g\right)_K\end{aligned}$$

where we have used the fact that the curvature $F_{A^*\otimes B} = {\mathrm{d}_{A^*\otimes B}}^2 = \partial_{A^*\otimes B}\bar{\partial}_{A^*\otimes B} + \bar{\partial}_{A^*\otimes B}\partial_{A^*\otimes B}$ is central and that $\bar{\partial}_{A^*\otimes B} g = 0$. Now using an argument like that in Proposition 4.7, we can apply Stokes' Theorem and find that $\partial_{A^*\otimes B} g = 0$ as well, so g is in fact in the kernel of $\mathrm{d}_{A^*\otimes B}$. But then so will be the unitary gauge transformation

$$u = g\,(g^* g)^{-\frac{1}{2}}$$

and hence d_A and d_B will be unitarily equivalent. □

As one easy application of the existence part of this theorem, note that any parabolic line bundle possesses a HYM connection, since line bundles are always stable. We can use this fact to give an extension of the correspondence described in §3 between weighted flat connections and representations of the fundamental group, as follows. Starting with E and X as above, let F be a trivial unitary line bundle on Σ with asymptotic data $Y_p = i\beta$ for each $p \in \partial\Sigma$, where

$$\beta = \frac{-\deg_X E}{nN},$$

and choose any HYM connection d_B on F. Then tensoring with (F, d_B) gives an identification $\mathcal{HYM}_{X,\omega}^{\delta} \cong \mathcal{F}_{X'}^{\delta}$, for $X_p' = X_p + i\beta\mathrm{Id}$. If we let $\mathfrak{o}_p' \subset U(n)$ be the conjugacy class of $\exp(-2\pi X_p')$ and $\mathfrak{o}' = \prod_{p\in\partial\Sigma}\mathfrak{o}_p' \subset \mathfrak{u}(n)^N$, Theorem 3.7 then implies:

Corollary 6.10. For sufficiently small $\delta > 0$, there is a homeomorphism $\mathcal{R}_{\mathfrak{o}'}/U(n) \cong \mathcal{HYM}_{X,\omega}^{\delta}/\mathcal{G}_U^{\delta}$.

References

1. M. F. Atiyah and R. Bott, *The Yang-Mills equations over Riemann surfaces*, Phil. Trans. R. Soc. London A **308** (1982), 523–615.
2. O. Biquard, *Fibrés paraboliques stables et connexions singulières plates*, Bull. Soc. Math. France **119** (1991), 231–257.
3. Y. Choquet-Bruhat and D. Christodoulou, *Elliptic systems in $H_{s,\delta}$ spaces on manifolds which are Euclidean at infinity*, Acta Math. **146** (1981), 129–150.

4. K. Corlette, *Non-Abelian Hodge theory and intersection cohomology*, In preparation.
5. ______, *Flat G-bundles with canonical metrics*, J. Diff. Geom. **28** (1988), 361–382.
6. G. D. Daskalopoulos, *The topology of the space of stable bundles over a compact Riemann surface*, J. Diff. Geom. (to appear).
7. G. D. Daskalopoulos and K. K. Uhlenbeck, *An application of transversality to the topology of the moduli space of stable bundles*, Preprint, 1992.
8. G. D. Daskalopoulos and R. Wentworth, *Geometric quantization for the moduli space of vector bundles with parabolic structure*, Preprint, 1992.
9. S. K. Donaldson, *A new proof of a theorem of Narasimhan and Seshadri*, J. Diff. Geom. **18** (1983), 269–277.
10. ______, *Anti-self-dual Yang-Mills connections over complex algebraic surfaces and stable vector bundles*, Proc. London Math. Soc. (3) **50** (1985), 1–26.
11. D. S. Freed and K. K. Uhlenbeck, *Instantons and four-manifolds*, M.S.R.I. Publications, vol. I, Springer-Verlag, New York, 1991, 2nd ed.
12. N. J. Hitchin, *The self-duality equations on a Riemann surface*, Proc. London Math. Soc. (3) **55** (1987), 59–126.
13. S. Kobayashi, *Differential geometry of complex vector bundles*, Princeton University Press, Princeton, 1987.
14. R. B. Lockhart and R. C. McOwen, *Elliptic differential operators on noncompact manifolds*, Ann. Sc. Norm. Sup. Pisa **12** (1985), 409–447.
15. G. Matic, *$SO(3)$-connections and rational homology cobordisms*, J. Diff. Geom. **28** (1988), 277–307.
16. V. B. Mehta and C. S. Seshadri, *Moduli of vector bundles on curves with parabolic structures*, Math. Ann. **248** (1980), 205–239.
17. D. Mumford and J. Fogarty, *Geometric invariant theory*, Springer-Verlag, New York, 1982, 2nd ed.
18. M. S. Narasimhan and C. S. Seshadri, *Stable and unitary vector bundles on a compact Riemann surface*, Ann. of Math. **82** (1965), 540–567.
19. C. S. Seshadri, *Fibrés vectoriels sur les courbes algébriques*, Astérisque, vol. 96, Soc. Math. France, Paris, 1982.
20. C. T. Simpson, *Constructing variations of Hodge structure using Yang-Mills theory and applications to uniformization*, J. Amer. Math. Soc. **1** (1988), 867–918.
21. ______, *Harmonic bundles on noncompact curves*, J. Amer. Math. Soc. **3** (1990), 713–770.
22. C. H. Taubes, *Gauge theory on asymptotically periodic 4-manifolds*, J. Diff. Geom. **25** (1987), 363–430.
23. K. K. Uhlenbeck, *Connections with L^p bounds on curvature*, Comm. Math. Phys. **83** (1982), 31–42.
24. K. K. Uhlenbeck and S. T. Yau, *On the existence of Hermitian-Yang-Mills connections in stable vector bundles*, Comm. Pure Appl. Math. **39-S** (1986), 257–293.

University of Maryland, College Park, MD 20742
E-mail address: kaleb@math.umd.edu