A Geometric Existence Proof for the Vortex Equations

Baylee Verzyde

Department of Mathematics

The University of Auckland

Supervisor: Pedram Hekmati

A dissertation submitted in partial fulfillment of the requirements for the degree of BSc(Hons) in Mathematics, The University of Auckland, 2024.

Abstract

The vortex equations were originally conceived in the context of superconductivity theory, but it was shown by Mrowka, Ozsvath and Yu that the moduli space of solutions to the Seiberg-Witten equations on a Seifert fibered 3-manifold is essentially diffeomorphic to the moduli space of solutions to the vortex equations on its underlying Riemann surface. Thus, the vortex equations are of importance in the gauge-theoretic study of 3-manifolds. In this dissertation, we review a proof by Garcia-Prada which characterises the structure of the moduli space of vortices. Additionally, we provide the necessary background in complex and symplectic geometry to understand the proof.

Acknowledgements

Firstly, I would like to thank my supervisor Pedram Hekmati. It has been an aspiration of mine for years to work in this field, and I am so grateful that he has introduced me to this world of mathematics. His wisdom, patience, and guidance have been invaluable throughout this project. That he often offered this support during extremely busy times has not gone unnoticed, and is deeply appreciated.

Secondly, I would like to thank my family: my mother, and my two sisters Chanel and Bella. I am well aware that, if one is less than exhilarated by the subject, living with a mathematician can be taxing; they have not only tolerated me but remained supportive of my goals throughout.

Finally, I would like to extend my gratitude to the rest of the University of Auckland Department of Mathematics, and especially my fellow students Andrew Wood and Dean Miller. They have both given me new ways to think about mathematics, be that through guidance or informal discussion, and they have helped me feel welcome in the department.

Contents

Chapter 1

Introduction

It is relatively simple to understand and classify surfaces up to diffeomorphism, but the corresponding task for 3-manifolds and 4-manifolds is much more complicated. Though Thurston's geometrisation theorem provides a classification of the 3-dimensional model geometries, it is generally quite difficult to decompose an arbitrary 3-manifold into these model geometries; the situation is even worse for 4-manifolds, for which we have no classification. Thus, tools for understanding the structure of these manifolds are highly desirable.

One such tool was introduced by Seiberg and Witten in 1994, and is now known as Seiberg-Witten theory; it was conceived as a physical theory on a flat space, but was subsequently adapted by Witten to closed 4-manifolds. In essence, the theory allows us to extract smooth invariants from manifolds by defining a set of geometric PDEs called the Seiberg-Witten equations, and considering the *moduli space of solutions* (the space of solutions up to equivalence by a particular group action). Though the Seiberg-Witten equations are only defined on 4-manifolds a priori, one can obtain a three-dimensional version of the equations by dimensional reduction; given a 3-manifold Y, we define the equations on the 4-manifold $Y \times \mathbb{R}$ and take the $\mathbb{R}$-invariant solutions. Moreover, the three-dimensional version of the theory can be further reduced in an important special case. In 1997, it was proven by Mrowka, Ozsvath, and Yu (see [MOY97]) that, when Y is a Seifert fibered 3-manifold, the moduli space of solutions is essentially equivalent to the moduli space of another important system of equations called the *vortex equations.* The moduli space of solutions to these equations were already being studied at the time, as they had other physical interpretations.

The vortex equations originated in superconductivity theory, being first published in 1950 by Ginzburg and Landau; for a general reference on Ginzburg-Landau theory and superconductivity, see [Sch97]. In the case of a complex wavefunction-like electron field Ψ and a magnetic vector field $F^A = \nabla \times A$ for a magnetic potential A, they proposed the following free energy density for a superconductor (up to constants depending on units):

$$E = |F_A|^2 + |\mathrm{d}\Psi - iA\Psi|^2 + \frac{\mu^2}{2}(|\Psi|^2 - 1)^2, \tag{1.1}$$

where $\mu \in (0, \infty)$ determines the behaviour of the superconductor. Note that E is invariant under the gauge transformations $\Psi \mapsto e^{i\chi}\Psi$, $A \mapsto A + \nabla\chi$ for any scalar field χ, and this corresponds to the fact that gauge transformations do not change physical states. The three terms in this energy density have concrete physical interpretations: the first is the energy due to the magnetic field; the second is the energy due to the interaction between the wavefunction and the magnetic field; and the third is the energy due to the interaction of the wavefunction with itself. This self-interaction was explained microscopically by Bardeen, Cooper, and Schrieffer in [BCS57], where they found a natural interpretation for Ψ as a pair of bound electrons.

The stable configurations correspond to pairs (A, Ψ) minimising $\int E\,\mathrm{d}V$. Such minimisers are called *vortices,* because of how they tend to behave around zeros of Ψ. In particular, if Ψ is equal to zero and A varies slowly in a given direction, there will be a line along which Ψ is zero. Furthermore, the current of Ψ will circle around the line in a sufficiently small neighbourhood, meaning Ψ behaves like a vortex around a zero. In this light, it was found that the value $\mu = 1/\sqrt{2}$ was physically significant; vortices tended to attract/repel each other for smaller/larger μ. Thus, stable solutions with finitely many vortices could only be found for this critical value of μ. Indeed, it was shown by Jaffe and Taubes in [JT80] that any choice of d points on the plane $\mathbb{R}^2$ for any d corresponded uniquely (up to gauge transformations) to a vortex (A, Ψ) with zeros at each chosen point (counted with multiplicity). In more concise and suggestive language, the moduli space of vortices is isomorphic to the disjoint union of all finite symmetric products of $\mathbb{R}^2$.

There is a more general geometric perspective on the moduli space of vortices, whose specifics we will formally define in Chapters 2 and 3. Namely, instead of a complex function on $\mathbb{R}^2$ and a vector field, we can instead consider a section of a complex line bundle, called a Higgs field, and a connection on that line bundle. This means we are no longer limited to flat space, so we instead consider arbitrary Riemann surfaces; we therefore have a natural generalisation of the energy density E to any Riemann surface with a complex line bundle. The new energy density functional (with $\mu = 1/\sqrt{2}$) is referred to as the Yang-Mills-Higgs functional, so named because it generalises the Yang-Mills functional to include a Higgs field. An extra parameter τ is also introduced; this new parameter adjusts the functional slightly so that solutions are not completely eliminated by topological obstructions. (We elaborate on this in Chapter 3.)

In 1990, it was shown by Bradlow in [Bra90] that all minimisers (∇, ϕ) of the Yang-Mills-Higgs functional on a Riemann surface satisfy the following equations:

$$\begin{aligned} \nabla^{0,1}\phi &= 0; \\ *F^{\nabla} &= \frac{i}{2}(|\phi|^2 - \tau). \end{aligned} \tag{1.2}$$

This pair of geometric PDEs is called the *τ-vortex equations.* In the same paper, Bradlow published an existence proof for the τ-vortex equations in the case that τ satisfied a certain inequality, and also showed that the moduli space of τ-vortices was characterised by a choice of d vortex points. However, unlike the flat case, the number d was found to be fixed; in particular, it had to be equal to the degree of the line bundle, meaning the moduli space was equivalent to dth symmetric product of the Riemann surface. Thus, the structure of the moduli space was found to be constrained topologically. In fact, Bradlow proved an analogue of this theorem in much more generality, constructing and solving the vortex equations on arbitrary Kähler manifolds with complex vector bundles of arbitrary rank.

Bradlow's existence proof relied on the reduction of the τ-vortex equations to a specific class of PDEs called the Kazdan-Warner equations, for which solutions were already known to exist. However, in 1994, Garcia-Prada published an existence proof in [Gar94] which was much more geometric in nature. His proof was based on a method used by Atiyah and Bott in [AB83] and Donaldson in [Don83] to study the Yang-Mills equations, in which the moduli space was shown to be the zero set of a certain geometric map on the space of connections (specifically a *moment map*). Remarkably, Garcia-Prada's proof is non-constructive: he describes the entire space of solutions to the vortex equations without explicitly constructing any of them.

Thus, the main goal of this dissertation is to review Garcia-Prada's proof, as well as the necessary background in complex geometry and symplectic geometry. The outline of the dissertation is as follows:

- Firstly, we develop the prerequisite ideas in complex geometry. We start with a discussion of complex manifolds in general, and follow this with a description of complex line bundles and holomorphic structures on them. The bulk of this chapter is dedicated to the various correspondences between tools for classifying line bundles up to isomorphism, both holomorphically and topologically.

- Secondly, we introduce the Yang-Mills-Higgs functional and the vortex equations, and demonstrate that their solutions can be characterised as zeros of a moment map. We supplement this with an introduction to the theory of moment maps. Much of this chapter elaborates on claims made by Garcia-Prada in his 1994 paper.

- Thirdly, we state the structure of the moduli space of τ-vortices, and we prove that it has this structure. We present a brief introduction to the relevant functional analysis, namely Sobolev spaces and elliptic differential operators. The proof given is entirely due to Garcia-Prada, though we have spelled out the reasoning in more elementary terms, and provided clarifications and corrections wherever necessary. We conclude with some remarks on further directions.

Chapter 2

Complex Geometry

As the vortex equations are naturally defined on complex vector bundles over compact complex manifolds, we begin with some background on these objects. We start by discussing general complex manifold theory, including the Dolbeault operators, Kähler structures, and inner products on differential forms. We then discuss the basics of complex vector bundles, which includes the theory of gauge transformations, connections, and curvature. This is followed by discussions of holomorphic structures on vector bundles, and how these can be viewed in terms of complex atlases, differential operators, Hermitian connections, and divisors. Finally, we introduce the first Chern class as a tool for classifying line bundles topologically. The majority of the contents of this chapter apply to general complex manifolds, but we occasionally restrict attention to compact Riemann surfaces; these are the primary objects of interest for the vortex equations as we shall define them.

Of these topics, we only present what we will need later when we discuss vortices. For a more comprehensive overview of complex manifolds and vector bundles, see [Mor07].

2.1 Complex Manifolds and Riemann Surfaces

Complex geometry is the study of spaces which "locally resemble" complex vector spaces. This can be made precise in one of two ways. The first way is via complex atlases, which are natural extensions of smooth structures from classical geometry.

Definition 2.1. Let X be a smooth manifold. A *complex atlas on X* consists of an open cover $\mathscr{U}$ of X and a diffeomorphism $\phi_U : U \to \phi_U(U) \subseteq \mathbb{C}^n$ for each $U \in \mathscr{U}$, such that $\phi_U \circ \phi_V^{-1}$ is holomorphic as a map between subsets of $\mathbb{C}^n$ for every $U, V \in \mathscr{U}$. If X is equipped with a complex atlas, then it is called a *complex manifold;* and if $n = 1$, then it is called a *Riemann surface.* (Unless explicitly stated otherwise, we assume X is connected.)

The second way is via almost-complex structures, which derive from the observation that a complex vector space is equivalent to a real vector space with an extra automorphism corresponding to $v \mapsto iv$.

Definition 2.2. Let X be a smooth manifold. An *almost-complex structure on X* is a $(1,1)$-tensor J on TX for which $J^2 = -\mathrm{id}$. The pair (X, J) is called an *almost-complex manifold.*

Every complex vector space $\mathbb{C}^n$ possesses an almost-complex structure; on the tangent bundle $T\mathbb{C}^n \cong \mathbb{C}^n \times \mathbb{C}^n$, one defines the $(1,1)$-tensor j to correspond to scalar multiplication by i on each tangent space. Consequently, a complex atlas on a manifold generates an almost-complex structure: if $p \in X$ and (U, ϕ_U) is a holomorphic chart containing p, we define the $(1,1)$-tensor J on each coordinate chart so that

$$J|_U = (\phi_U^{-1})_* \circ j \circ (\phi_U)_*. \tag{2.1}$$

It is fairly easy to show that this definition is independent of the choice of coordinates. The conditions under which an arbitrary almost-complex structure is induced by a complex atlas are given by the Newlander-Nirenberg theorem:

Theorem 2.3 (Newlander-Nirenberg)**.** *Let (X, J) be an almost-complex manifold. Then there is a complex structure on X if and only if $T^{0,1}X \subseteq TX^{\mathbb{C}}$, the $-i$-eigenspace of J, is integrable.*

For a proof, see [Mor07].

As hinted above, there is a canonical decomposition of the complexified tangent bundle of an almost-complex manifold. In particular, we have that $TX^{\mathbb{C}} = T^{1,0}X \oplus T^{0,1}X$, where each direct summand is the $+i$- and $-i$-eigenspace of J respectively. Similarly, we have a decomposition of the complex-valued differential forms on X:

$$\Omega^k(X) = \bigoplus_{p+q=k} \Omega^{p,q}(X), \tag{2.2}$$

where $\Omega^{p,q}(X) = (\Omega^{1,0}(X))^{\wedge p} \wedge (\Omega^{0,1}(X))^{\wedge q}$, and $\Omega^{1,0}(X), \Omega^{0,1}(X) \subseteq \Omega^1(X)$ are the $\pm i$ eigenspaces of the pullback of J. This decomposition leads directly to a decomposition of the exterior derivative; where $\pi^{p,q} : \Omega^{p+q}(X) \to \Omega^{p,q}(X)$ denotes the direct sum projection, we define the *Dolbeault operators* as follows:

$$\begin{aligned} \partial_{p,q} &= \pi^{p+1,q} \circ \mathrm{d}_{p+q}, \\ \overline{\partial}_{p,q} &= \pi^{p,q+1} \circ \mathrm{d}_{p+q}. \end{aligned} \tag{2.3}$$

One can show that, whenever the almost-complex structure is integrable, $\mathrm{d}_{p+q} = \partial_{p,q} + \overline{\partial}_{p,q}$ for every p, q. It follows by separating d^2 into its direct summands that $\partial^2 = \overline{\partial}^2 = \partial\overline{\partial} + \overline{\partial}\partial = 0$. Therefore, the spaces of differential (p,q)-forms equipped with the Dolbeault operators form a double complex, meaning we can form cohomology out of them:

Definition 2.4. Given a complex manifold X, the (p,q) *Dolbeault cohomology space* is the following vector space:

$$H^{p,q}(X) = \frac{\ker(\overline{\partial}_{p,q})}{\mathrm{im}(\overline{\partial}_{p,q-1})}. \tag{2.4}$$

Unsurprisingly, there is a relationship between the Dolbeault cohomology and the other varieties of cohomology. This is made precise by Dolbeault's theorem (for a proof, see [GH94]):

Theorem 2.5 (Dolbeault)**.** *Let X be a complex manifold. Then there is an isomorphism between $H^{p,q}(X)$ and $\check{H}^q(X, \Omega^p_{\mathrm{hol}}(X))$, where $\check{H}$ denotes the Čech cohomology and $\Omega^q_{\mathrm{hol}}(X)$ is the sheaf of holomorphic q-forms.*

One case of the Dolbeault isomorphism can be spelled out explicitly, and we will use this description later. In what follows, let $\mathscr{U} = \{U_i\}_{i\in I}$ be a good open cover of X. If $p = 0$ and $q = 1$, we can describe the map $\mathscr{F} : H^{0,1}(X) \to \check{H}^1(X, \Omega^0_{\text{hol}}(X)) \cong \check{H}^1(\mathscr{U}, \mathscr{O})$ as follows. Given a class in $H^{0,1}(X)$ defined by some $(0,1)$-form α, we make the following definition for $\mathscr{F}([\alpha]) \in \check{H}^1(\mathscr{U}, \mathscr{O})$:

$$\mathscr{F}([\alpha])_{ij} = f_i|_{U_i\cap U_j} - f_j|_{U_i\cap U_j}, \tag{2.5}$$

where each $f_i : U_i \to \mathbb{C}$ is a smooth function for which $\alpha|_{U_i} = \overline{\partial} f_i$ (which we know must exist by the Poincaré lemma). There are four things to check regarding this definition:

- The output of $\mathscr{F}$ is a 1-cocycle; this is a straightforward computation.
- The output of $\mathscr{F}$ does not depend on the choice of α. If β is another $(0,1)$-form in the same class, then $\beta = \alpha + \overline{\partial} h$ for some $h : X \to \mathbb{C}$. But then $\beta|_{U_i} = \overline{\partial} f_i + \overline{\partial} h$ for each i, meaning $\mathscr{F}([\beta])_{ij} = (f_i + h)|_{U_i\cap U_j} - (f_j + h)|_{U_i\cap U_j} = \mathscr{F}([\alpha])_{ij}$.
- The output of $\mathscr{F}$ does not depend on the choice of f_i. If g_i is another function for which $\alpha = \overline{\partial} g_i$, then $g_i = f_i + c_i$ where $c_i : X \to \mathbb{C}$ is holomorphic. But then $g_i - g_j = f_i - f_j + c_i - c_j = f_i - f_j + (\delta c)_{ij}$, where c is the 0-cochain for which $c(U_i) = c_i$. It follows that the cocycle determined by the g_i differs by a coboundary to the cocycle determined by the f_i, so they are equal in $\check{H}^1(\mathscr{U}, \mathscr{O})$.
- $\mathscr{F}$ is linear; this follows from the linearity of $\overline{\partial}$.

To show that $\mathscr{F}$ is an isomorphism, we may construct its inverse. Given a cocycle $f \in \check{C}^1(\mathscr{U}, \mathscr{O})$, we can consider f as a cocycle valued in $\mathscr{E}$ instead (i.e. the sheaf of smooth functions). However, since this sheaf is acyclic, we know that $f = \delta g$ for some 0-cocycle g, i.e. $f_{ij} = g_i - g_j$. We then define a $(0,1)$-form α, whose cohomology class we define to be $\mathscr{F}^{-1}(f)$, by taking $\alpha|_{U_i} = \overline{\partial} g_i$. There are three things to check:

- α is a well-defined $(0,1)$-form. Since $g_i - g_j$ is holomorphic, we know that $\overline{\partial} g_i = \overline{\partial} g_j$ on overlaps. Therefore, the definition of α is consistent and can be glued together to form a global differential form.
- The cohomology class of α does not depend on the choice of g_i. If $f_{ij} = h_i - h_j$, then $g_i - g_j = h_i - h_j$, so $g_i - h_i = g_j - h_j$. It follows that we can glue together each $g_i - h_i$ to form a global function $g - h$ (even though neither g nor h may be glued in general), which means that $\overline{\partial} h_i = \overline{\partial} g_i - \overline{\partial}(g - h)$. It follows that the image of f under $\mathscr{F}^{-1}$ differs by a $\overline{\partial}$-exact $(0,1)$-form under different choices of g_i, meaning the cohomology class is the same.
- $\mathscr{F}^{-1}$ is indeed the inverse of $\mathscr{F}$; this is once again a straightforward computation.

It follows that $\mathscr{F}$ is an explicit isomorphism between $H^{(0,1)}(X)$ and $\check{H}^1(X, \mathscr{O})$.

Additionally, we can elegantly phrase holomorphicity in terms of the Dolbeault operators. A function $f : X \to \mathbb{C}$ is holomorphic (i.e. it is holomorphic in the complex atlas) if and only if

$$\overline{\partial} f = 0. \tag{2.6}$$

On the other hand, if X is compact, the space of holomorphic functions reduces dramatically:

Theorem 2.6. *If X is a compact complex manifold, then the space $\mathscr{O}(X)$ of holomorphic functions on X consists only of constant functions.*

For a proof of this fact, see [Wel08].

Because every complex manifold is a smooth manifold, the smooth notions of metric and symplectic structure can be applied to complex geometry. We briefly recall their definitions:

Definition 2.7. Given a smooth manifold X, a *Riemannian metric on X* is a $(0,2)$-tensor g which is pointwise a real inner product. A *symplectic structure on X* is a real-valued closed 2-form ω on X which is pointwise nondegenerate.

The analogue of a complex-valued inner product is called a Hermitian metric, and it is definable only on complex manifolds:

Definition 2.8. A *Hermitian metric* on a complex manifold X is a complex-valued 2-tensor h on TX, which is pointwise a complex inner product.

A special case occurs when all three of these structures agree with each other; a manifold with three compatible structures is called a Kähler manifold.

Definition 2.9. Let X be a smooth manifold. A *Kähler structure on X* consists of an integrable almost-complex structure J, a Riemannian metric g, and a symplectic structure ω, such that $g(u,v) = \omega(u, Jv)$ for every u, v. Equivalently, a Kähler structure on X is a complex structure on X and a Hermitian metric h on X for which $\mathrm{Im}(h)$ is a closed 2-form.

The equivalence between these two definitions is seen by taking $\mathrm{Re}(h) = g$ and $\mathrm{Im}(h) = \omega$.

Any Riemannian metric on TX has an extension to all tensor bundles $(TX)^{\otimes k} \otimes (T^*X)^{\otimes \ell}$, given by enforcing that

$$g(u_1 \otimes v_1, u_2 \otimes v_2) = g(u_1, u_2)g(v_1, v_2) \tag{2.7}$$

for any u_i, v_i in the same fibre of the tensor bundle. We also extend this to the real exterior powers $\Lambda^k(X)$ by defining

$$(\alpha, \beta) = \frac{1}{k!} g(\alpha, \beta) \tag{2.8}$$

for each $\alpha, \beta \in \Omega^k(X)$, and this is extended to the complex exterior powers $\Lambda^{p,q}(X)$ by making g conjugate-linear in its second argument. It can be shown by induction that this inner product respects the holomorphic decomposition of $\Lambda^k(X)$, i.e. that

$$\Lambda^{p,q}(X) \perp_g \Lambda^{r,s}(X) \tag{2.9}$$

whenever $(p,q) \neq (r,s)$. The inner product allows us to define a new kind of dual on the spaces ΛX:

Definition 2.10. The *Hodge dual* is the unique operator $* : \Lambda^{p,q}X \to \Lambda^{n-q,n-p}X$ such that the following holds for any $\alpha, \beta \in \Lambda^{p,q}$:

$$\alpha \wedge \overline{*\beta} = (\alpha, \beta)\mathrm{vol}_g. \tag{2.10}$$

By working in an orthonormal basis and extending linearly, it is easily shown that $*^2 = (-1)^{(p+q)(n-p-q)}$ on $\Lambda^{p,q}X$. If we further assume that X is compact, then we can integrate any smooth complex-valued function on X using the volume form vol_g, and we can therefore extend this inner product to differential forms as follows:

Definition 2.11. Given differential k-forms $\alpha, \beta \in \Omega^{p,q}(X)$ over a compact complex manifold X, we define their L^2-inner product as follows:

$$\langle \alpha, \beta \rangle = \int_X (\alpha, \beta)\mathrm{vol}_g = \int_X \alpha \wedge *\overline{\beta}, \tag{2.11}$$

where the expression (α, β) refers to the pointwise inner product of α and β in $\Lambda^{p,q}(X)$.

As previously stated, a Riemann surface is a complex manifold X of dimension 1. All of the preceding theory carries over to Riemann surfaces, but the low dimensionality results in some simplifications. Firstly, we note that $\Omega^k(X)$ vanishes for any $k > 2$, meaning there are only four nontrivial spaces of complex differential forms. We also have that every Riemannian metric automatically gives rise to a Kähler structure since every 2-form is already closed. In this case, we also have the important result that the symplectic 2-form is precisely the volume form induced by g; this is because J is orientation-preserving and orthogonal, so $\omega(u_1, u_2) = \omega(u_1, Ju_1) = 1$ for any orthonormal frame $\{u_1, u_2\}$.

Given a Riemann surface X with metric g and associated Kähler form ω, we can define a map from $\Omega^{0,0}(X)$ to $\Omega^{1,1}(X)$ given by $f \mapsto f\omega$; it is clearly an isomorphism. The inverse of this map will be denoted by $\Lambda : \Omega^{1,1}(X) \to \Omega^{0,0}(X)$; explicitly,

$$\Lambda(f\omega) = f. \tag{2.12}$$

In fact, there is an important relationship between Λ and the Dolbeault operators, that we will occasionally make use of, called the Kähler identities:

Theorem 2.12 (Kähler identities)**.** *Let X be a Riemann surface with Kähler form ω, and denote the commutator by square brackets. Then the following equations hold:*

$$\begin{aligned} [\Lambda, \partial] &= -i(* \circ \partial \circ *), \\ [\Lambda, \overline{\partial}] &= i(* \circ \overline{\partial} \circ *). \end{aligned} \tag{2.13}$$

See [Mor07] for a proof.

One final result that will be occasionally useful is the following: if X is a Riemann surface and $\alpha, \beta \in \Lambda^{1,0}X \oplus \Lambda^{0,1}X$, then

$$\alpha \wedge *\beta = - * \alpha \wedge \beta. \tag{2.14}$$

2.2 Complex Vector Bundles

Our fundamental goal is to study the space of solutions of certain PDEs on Riemann surfaces. However, since every holomorphic complex-valued function on a compact Riemann surface is necessarily constant, we will need to broaden the class of functions we consider to gain any interesting information. Thus, we will instead consider sections of vector bundles, and the PDEs will be formed using connections and Hermitian metrics. We briefly review these concepts.

Definition 2.13. Let X be a smooth manifold. A *rank k complex vector bundle over X* is a manifold E together with a projection $\pi : E \to X$, satisfying the following conditions:

- For each $p \in X$, $\pi^{-1}(p)$ has a complex vector space structure isomorphic to $\mathbb{C}^k$;
- X has an open cover $\{U_i\}_{i \in I}$ together with diffeomorphisms $\psi_i : \pi^{-1}(U_i) \to U_i \times \mathbb{C}^k$, called local trivialisations of E over U_i, such that the following diagram commutes:

$$\begin{array}{ccc} \pi^{-1}(U_i) & \xrightarrow{\psi_i} & U_i \times \mathbb{C}^k \\ & {\scriptstyle \pi}\searrow \quad \swarrow{\scriptstyle \mathrm{pr}_1} & \\ & U_i & \end{array}$$

We frequently denote the vector bundle by $E \to X$, or simply E. If $k = 1$, the vector bundle is called a *line bundle.* A *real* vector bundle of rank k is exactly the same as a complex line bundle, except that $\mathbb{C}$ is replaced with $\mathbb{R}$ wherever it appears.

Equivalently, a vector bundle $E \to X$ may be characterised as an open cover $\{U_i\}$ of X and a set of transition functions $g_{ij} : U_i \cap U_j \to \mathrm{Aut}(\mathbb{C}^k)$. Each U_i corresponds to an open set over which E is locally trivial, and each g_{ij} corresponds to the change of basis on $(U_i \cap U_j) \times \mathbb{C}^k$ when moving between local trivialisations. The transition functions must satisfy the cocycle condition in order to define a vector bundle; that is, we must have that

$$g_{ij} g_{jk} g_{ki} = 1. \tag{2.15}$$

For a detailed proof of this correspondence, see [Hus94].

We immediately define the notion of functions on manifolds with values in vector bundles. In fact, we call them sections:

Definition 2.14. Given an open set $U \subseteq X$ and a vector bundle $E \to X$, a *local section of E (over U)* is a smooth map $\sigma : U \to E$ such that $\pi \circ \sigma = \mathrm{id}$, and the set of all such sections is denoted by $\Gamma_U(E)$. If $U = X$ then σ is called a *global section,* and the set of these sections is denoted by $\Gamma(E)$.

One example of a vector bundle which always exists is the *trivial bundle,* defined to be $X \times \mathbb{C}^k \to X$. On a trivial bundle, a section is just a $\mathbb{C}^k$-valued function on X. For a more interesting class of examples, consider the various tangent bundles: the tangent bundle TX, its associated tensor bundles, and the bundles of alternating tensors $\Lambda^k X$, all constitute real vector bundles; their complexified variants, on the other hand, are complex vector bundles. Sections of these bundles correspond to vector fields, tensor fields, and differential forms.

Many linear algebraic constructions carry directly over to vector bundles, by applying them pointwise. In particular, given a vector bundle $E \to X$, there is a *dual bundle* $E^* \to X$ and a *conjugate bundle* $\overline{E} \to X$, obtained by replacing each fibre space with the dual space and the conjugate space respectively. (Note that a section of the dual bundle acts on sections of E by contracting pointwise.) Moreover, if we have another vector bundle $F \to X$, we can construct the *direct sum bundle* $E \oplus F$ by taking the direct sum of the fibres, the *tensor product bundle* $E \otimes F$ by taking the tensor product of the fibres, and the *homomorphism bundle* $\mathrm{Hom}(E, F)$ by taking the space of linear maps between fibres.

(Note that the homomorphism bundle is isomorphic to $E^* \otimes F$.) A special case of these constructions is the bundle of *E-valued differential k-forms* $\Lambda^k(E)$, which is defined by taking an arbitrary vector bundle E and taking its tensor product with $\Lambda^k(X)$. The space of sections of this bundle is denoted by $\Omega^k(E)$. For a more careful construction of a topology on these bundles, see [Hus94].

With these notions, we can define automorphisms of vector bundles. These are referred to as gauge transformations:

Definition 2.15. Let $\pi : E \to X$ be a complex vector bundle. A *gauge transformation on* L is a section $g \in \Gamma(\mathrm{Aut}(E))$ of the automorphism bundle. We denote by $\mathscr{G}$ (or $\mathscr{G}(E)$) the group of all gauge transformations of a vector bundle E.

We define a group action of $\mathscr{G}(E)$ on $\Gamma(E)$ as follows: an element $g \in \mathscr{G}(E)$ acts on a section $\sigma \in \Gamma(E)$ by taking

$$\sigma \mapsto g^{-1}\sigma. \tag{2.16}$$

This ostensibly unnatural choice ensures that the action of $\mathscr{G}(E)$ on $\Gamma(E)$ is a group action, i.e. that $(\sigma^g)^h = \sigma^{gh}$.

In order to differentiate sections of a vector bundle, we need an extra structure called a connection. A connection is essentially interpreted as a generalisation of the directional derivative in flat space, except that the input vector is a tangent vector.

Definition 2.16. Let $\pi : E \to X$ be a complex vector bundle. A *connection on E* is a linear map $\nabla : \Gamma(E) \to \Omega^1(E)$ satisfying the Leibniz rule: for every $f \in C^\infty(X)$ and every $\sigma \in \Gamma(E)$, we have

$$\nabla(f\sigma) = \mathrm{d}f \otimes \sigma + f\nabla\sigma. \tag{2.17}$$

We denote by $\mathscr{A}(E)$ the set of all connections on E, or simply $\mathscr{A}$ is E is obvious from context.

Note that there is a natural action of $\mathscr{G}(E)$ on $\mathscr{A}(E)$, defined so that $g \in \mathscr{G}(E)$ acts on $\nabla \in \mathscr{A}(E)$ by taking

$$\nabla \mapsto g^{-1}\nabla g. \tag{2.18}$$

This choice ensures that $(\nabla\sigma)^g = \nabla^g\sigma^g$.

Given connections ∇^E and ∇^F on vector bundles $E, F \to X$, we can define a connection $\nabla^{E\otimes F}$ on $E \otimes F$ by defining

$$\nabla^{E\otimes F}(\sigma \otimes \eta) = (\nabla^E\sigma) \otimes \eta + \sigma \otimes (\nabla^F\eta) \tag{2.19}$$

for every $\sigma \in \Gamma(E)$ and $\eta \in \Gamma(F)$. Likewise, we can define a connection ∇^{E^*} on E^* by defining

$$(\nabla^{E^*}\omega)(\sigma) = \mathrm{d}(\omega(\sigma)) - \omega(\nabla^E\sigma) \tag{2.20}$$

for every $\omega \in \Gamma(E^*)$ and every $\sigma \in \Gamma(E)$. This leads immediately to the following useful result:

Theorem 2.17. *Let $\pi : E \to X$ be a complex vector bundle with connection ∇. Then there is a unique system of connections on each tensor product bundle $E^{\otimes m} \otimes (E^*)^{\otimes n}$ satisfying the following conditions:*

- *On the trivial bundle,* $\nabla = \mathrm{d}$.
- *If* $\delta : E^{\otimes m} \otimes (E^*)^{\otimes n} \to E^{\otimes(m-1)} \otimes (E^*)^{\otimes(n-1)}$ *denotes a contraction operator (i.e. a trace), then* $\delta \circ \nabla = \nabla \circ \delta$.
- *For any sections* A *and* B *of any two tensor product bundles,* $\nabla(A \otimes B) = (\nabla A) \otimes B + A \otimes (\nabla B)$.

Similarly, there is a unique system of connections on each bundle of E*-valued differential forms* $\Lambda^k(E)$ *satisfying the Leibniz rule: for any* $\omega \in \Omega^k(X)$ *and any* $\sigma \in \Gamma(E)$*, we have that*

$$\nabla(\omega \wedge \sigma) = (\nabla\omega) \otimes \sigma + (-1)^p \omega \wedge \nabla\sigma. \tag{2.21}$$

Thus, once we define a connection on sections of a bundle, we immediately have sections of every "derived" bundle. We will often refer to these as the same connection. For more details on this, see [Kob87].

It is an idea from classical differential geometry that derivatives give information regarding intrinsic curvature; specifically, one can compute derivatives in different directions, and if they commute at a given point then the space is flat there. This idea admits a natural generalisation in the vector bundle formalism:

Definition 2.18. Let $\pi : E \to X$ be a complex vector bundle with connection ∇. The *curvature 2-form* $F^\nabla \in \Omega^2(\mathrm{End}(E))$ is defined such that

$$F^\nabla(\sigma) = \nabla(\nabla\sigma). \tag{2.22}$$

It is linear over $C^\infty(X)$ and alternating in its input tangent vectors, meaning it is indeed a 2-form. Note also that, under a gauge transformation $g \in \mathscr{G}(E)$, the curvature 2-form transforms as follows:

$$(F^\nabla)^g = g^{-1} F^\nabla g. \tag{2.23}$$

It will be occasionally useful to conceptualise our global definitions in terms of local trivialisations. To this end, the following proposition is illuminating:

Proposition 2.19. *Let* $\pi : E \to X$ *be a complex vector bundle of rank* k*. A local trivialisation of* E *over* U *may be equivalently described as a collection of nowhere-zero local sections* $\sigma_1, \ldots, \sigma_k \in \Gamma_U(E)$ *which is pointwise linearly independent (that is, a smoothly varying basis for* $E|_U$*).*

Proof. Given a local trivialisation $\psi : \pi^{-1}(U) \to U \times \mathbb{C}^k$, we define $e_i : U \to U \times \mathbb{C}^k$ to be the constant vector in $\mathbb{C}^k$ in the ith direction, and we then take $\sigma_i \in \Gamma_U(E)$ to be $\psi^{-1}(e_i)$. Conversely, given a smoothly varying basis $\sigma_1, \ldots, \sigma_k \in \Gamma_U(E)$, we define $\psi : \pi^{-1}(U) \to U \times \mathbb{C}^k$ to take $e \in \pi^{-1}(U)$ to $(\pi(e), v)$, where $v \in \mathbb{C}^k$ consists of the components of e in the pointwise basis $\sigma_1, \ldots, \sigma_k$. The verification that all maps are smooth, and that ψ is a diffeomorphism, is left to the reader. □

We will often refer to a smoothly varying basis as a local trivialisation, and vice versa. If $\{\sigma_i\}_{i=1}^k \in \Gamma_U(E)$ is a local trivialisation of E over an open set U, the connection can be written locally as follows:

$$\nabla\sigma = A\sigma, \tag{2.24}$$

where $A \in \Omega^1(\mathrm{End}(E)|_U)$ is an endomorphism-valued 1-form on U called the *connection form.* In terms of the connection form, the curvature 2-form is locally expressed as follows:

$$F^{\nabla} = \mathrm{d}A + A \wedge A, \tag{2.25}$$

where the wedge product is taken component-wise and therefore does not vanish in general.

In the case of a complex line bundle $L \to X$, several reductions make the theory slightly simpler. Firstly, since the fibres are 1-dimensional, a trivialisation of L corresponds to a nowhere-vanishing section of L. Secondly, $\mathrm{End}(\mathbb{C}) = \mathbb{C}$ as a vector space over $\mathbb{C}$, and the identity map constitutes a nowhere-vanishing section of $\mathrm{End}(\mathbb{C})$, making the bundle globally trivial. Thus, a gauge transformation of L is really just a function $g : X \to \mathbb{C}^*$, and the action of $\mathscr{G}(L)$ on F^{∇} is trivial. Moreover, the connection form becomes an ordinary 1-form, i.e. $A \in \Omega^1(U)$, and this means that the curvature reduces to $F^{\nabla} = \mathrm{d}A$. We also have that $F^{\nabla^g} = F^{\nabla}$. It is worth noting that, for a line bundle, we recover the classical theory of electromagnetism with no sources by taking A to be the vector potential and F^{∇} to be the field strength; then Maxwell's equations are simply $\mathrm{d}F^{\nabla} = \mathrm{d} * F^{\nabla} = 0$.

2.3 Holomorphic Vector Bundles

All that we have outlined so far is in the realm of smooth geometry, and (aside from the dimension of the fibres) there is no complex structure in the objects we have defined. We now introduce this aspect by specifying what it means for a function to be holomorphic.

Definition 2.20. Let X be a complex manifold, and let $\pi : E \to X$ be a rank k vector bundle. A *holomorphic atlas on E* consists of an open cover $\{U_i\}$ of X and local trivialisations $\psi_i : \pi^{-1}(U_i) \to U_i \times \mathbb{C}^k$ such that, for each intersecting U_i, U_j, we have $(\psi_i \circ \psi_j^{-1})(p, z) = (p, g_{ij}(p)z)$ for holomorphic functions $g_{ij} : U_i \cap U_j \to \mathbb{C}^k$. We call E a *holomorphic vector bundle* if it has a holomorphic atlas.

A section is then called *holomorphic* if its local representation in the holomorphic atlas is holomorphic. A related analytic notion is that of a holomorphic structure:

Definition 2.21. Let $\pi : E \to X$ be a complex vector bundle over a complex manifold. A *holomorphic structure on E* is a collection of operators $\overline{\partial}_E : \Omega^{p,q}(E) \to \Omega^{p,q+1}(E)$ for each (p, q) such that $\overline{\partial}_E^2 = 0$, satisfying the Leibniz rule: for any $\omega \in \Omega^{p,q}(X)$ and any $\sigma \in \Omega^{r,s}(E)$, we have that

$$\overline{\partial}_E(\omega \wedge \sigma) = (\overline{\partial}\omega) \wedge \sigma + (-1)^{p+q}\omega \wedge (\overline{\partial}_E\sigma). \tag{2.26}$$

A section $\sigma \in \Gamma(E)$ is called *holomorphic* if $\overline{\partial}_E\sigma = 0$.

It turns out that these two descriptions of holomorphicity are equivalent. We will prove that the two are equivalent in the case of line bundles, and we will do so with the following lemma:

Lemma 2.22. *Let $\pi : L \to X$ be a complex line bundle over a complex manifold X, and let $\overline{\partial}_L$ be a holomorphic structure. Around every point $p \in X$, there is a neighbourhood U and a nowhere-zero local section $\sigma \in \Gamma_U(L)$ such that $\overline{\partial}\sigma = 0$.*

Proof. We give only an outline of the proof; for the complete proof of this lemma, see [Mor07]. The argument essentially consists of three steps:

- Work in a local trivialisation of the line bundle, i.e. some local section σ in a neighbourhood $U \subset X$ of the given point (with holomorphic coordinates z_α). We can then represent the holomorphic structure as the $(0,1)$-form τ for which $\overline{\partial}_L \sigma = \tau\sigma$, and the target trivialisation as a nowhere-zero complex-valued function $f : U \to \mathbb{C}$ satisfying the equation $\overline{\partial} f + f\tau = 0$.
- Use the holomorphic structure τ to put a complex structure on $\pi^{-1}(U) \cong U \times \mathbb{C}$. In particular, if w represents the coordinate on $\mathbb{C}$, there is a unique complex structure on $U \times \mathbb{C}$ for which $\{\mathrm{d}z_\alpha, \mathrm{d}w - \tau w\} \subseteq \Lambda^1(U \times \mathbb{C})$ generates the $(1,0)$-forms on $U \times \mathbb{C}$.
- Use the Newlander-Nirenberg theorem to complete the holomorphic coordinates z_α to a local coordinate system (z_α, u) on $U \times \mathbb{C}$. The desired function f is given by a component of $\mathrm{d}u$ in the basis of $(1,0)$-forms.

□

Theorem 2.23. *Let $\pi : L \to X$ be a complex line bundle over a complex manifold. Then it is holomorphic if and only if it has a holomorphic structure.*

Proof. It is fairly easy to construct a holomorphic structure for a holomorphic bundle: in a local trivialisation $\sigma \in \Gamma_U(L)$, we define $\overline{\partial}_L(\alpha\sigma) = (\overline{\partial}\alpha)\sigma$ for each $\alpha \in \Omega^{p,q}(L)$, where $\overline{\partial}$ is a Dolbeault operator. It can be easily shown that this definition is coordinate-independent, squares to 0, and satisfies the Leibniz rule.

We construct a system of holomorphic transition functions from a holomorphic structure $\overline{\partial}_L$ by using the above lemma. In particular, by the above lemma, there is an open cover $\{U_i\}$ of X and local trivialisations (i.e. nowhere-zero local sections) $\sigma_i \in \Gamma_{U_i}(L)$ such that $\overline{\partial}_L \sigma_i = 0$. The transition functions between the trivialisations σ_i and σ_j is the function $g_{ij} : U_i \cap U_j \to \mathbb{C}$ for which $\sigma_j = g_{ij}\sigma_i$. It follows from the Leibniz rule that $\overline{\partial} g_{ij} = 0$ for every i, j, so each g_{ij} is holomorphic. □

Henceforth, we will identify systems of holomorphic transition functions with holomorphic structures. We will also use meromorphic sections when discussing divisors:

Definition 2.24. A *meromorphic section* of a holomorphic line bundle is a section which is meromorphic in the holomorphic atlas. We denote by $\mathscr{M}(L)$ the set of all meromorphic sections of L.

Note that the complexified tangent bundle and bundles of complex differential $(p,0)$-forms possess natural holomorphic structures, so we have natural notions of holomorphic vector fields and holomorphic differential forms. On a compact Riemann surface, the existence of nonvanishing holomorphic differential forms is entirely determined by the topology of the surface:

Proposition 2.25. *The only compact Riemann surface with a nonvanishing holomorphic 1-form is the torus.*

For a proof, see [Don11].

2.4 The Chern Correspondence

Given a vector bundle $\pi : E \to X$, there is another way to think of holomorphic structures in the special case that E possesses a Hermitian structure; we begin with a description of this structure:

Definition 2.26. Let $\pi : E \to X$ be a vector bundle over a complex manifold. A *Hermitian structure* h *on* E is a smooth section of $(E \otimes \overline{E})^*$ which is a pointwise Hermitian inner product (i.e. it is conjugate-symmetric and nondegenerate). A vector bundle with a Hermitian structure is called a *Hermitian vector bundle,* and is denoted by E_h. A connection ∇ on E_h is called a *Hermitian connection* (or a unitary connection) if $\nabla h = 0$. We will denote by $\mathscr{A}^h$ the space of all Hermitian connections, or simply $\mathscr{A}$ when the Hermitian structure is understood from context.

If a vector bundle possesses a Hermitian structure h, the gauge transformations which preserve h are of special importance. Specifically, we call a gauge transformation $g \in \mathscr{G}(E)$ *unitary* if

$$(g^* h)(u, v) = h(u^g, v^g) = h(u, v) \tag{2.27}$$

for any $u, v \in E$ in the same fibre. We will sometimes refer to more general gauge transformations as *complex.*

Given a section $\sigma \in \Gamma(E)$ and a connection ∇, it is true by definition that $\nabla\sigma$ is in $\Omega^1(E)$. Furthermore, if we complexify these differential forms, this splits into $(1,0)$- and $(0,1)$-forms $\Omega^{1,0}(E)$ and $\Omega^{0,1}(E)$. By projecting $\nabla\sigma$ into each of these spaces, we obtain new connections $\nabla^{1,0}$ and $\nabla^{0,1}$. Using these connections, we can demonstrate a correspondence between Hermitian connections and holomorphic structures on vector bundles:

Theorem 2.27 (Chern correspondence)**.** *Let* $\pi : E_h \to X$ *be a Hermitian vector bundle over a complex manifold. Given a holomorphic structure* $\overline{\partial}_E$*, there is a unique Hermitian connection* ∇ *(called the Chern connection) for which* $\nabla^{0,1} = \overline{\partial}_E$*. Additionally, if* ∇ *is any Hermitian connection on* E_h *for which* $F^\nabla \in \Omega^{2,0}(X) \oplus \Omega^{1,1}(X)$ *(meaning* F^∇ *has no* $(0,2)$*-component), then* $\nabla^{0,1}$ *is a holomorphic structure on* E_h*.*

Proof. The following proof is from [Mor07]. First, let $\overline{\partial}_E$ be a holomorphic structure. Since h is nondegenerate, we can think of h as a map from $\Gamma(E)$ to $\Gamma(\overline{E}^*)$ by tensoring h with a given section, and applying a contraction. If ∇ is a Hermitian connection, then by using the Leibniz rule (and the action of connections on tensor products and contractions), we observe that

$$\nabla(h(\sigma)) = \nabla(\delta(h \otimes \sigma)) = \delta((\nabla h) \otimes \sigma + h \otimes \nabla\sigma) = h(\nabla\sigma), \tag{2.28}$$

for every $\sigma \in \Gamma(E)$. Given an arbitrary complex tangent vector $v \in TX^{\mathbb{C}}$, it follows from the conjugate linearity of h that $\nabla_v(h(\sigma)) = h(\nabla_{\overline{v}}\sigma)$. Since conjugation exchanges $T^{1,0}X$ and $T^{0,1}X$, and since $(1,0)$-forms annihilate $(0,1)$-vectors (and vice versa), it follows that $h(\nabla^{1,0}\sigma) = \nabla^{0,1}(h(\sigma))$. Applying the inverse of h, we see that $\nabla^{1,0} = h^{-1} \circ \nabla^{0,1} \circ h$.

Now, define the Chern connection to be

$$\nabla = h^{-1} \circ \overline{\partial}_E \circ h + \overline{\partial}_E. \tag{2.29}$$

It is easy to verify that ∇ is linear and satisfies the Leibniz rule, meaning it is a connection. Furthermore, its $(0,1)$-component is simply $\overline{\partial}_E$. Any other connection with this $(0,1)$-component must have $(1,0)$-component $h^{-1} \circ \overline{\partial}_E \circ h$ by the above calculation, meaning it must be the same as the Chern connection.

Conversely, let ∇ be an arbitrary connection on E_h. Since $R^\nabla = \nabla \circ \nabla$, the linearity of F^∇ and the splitting of ∇ shows that $F^\nabla = (\nabla^{1,0})^2 + \{\nabla^{1,0}, \nabla^{0,1}\} + (\nabla^{0,1})^2$. The only $(0,2)$-component of R^∇ vanishes precisely when $(\nabla^{0,1})^2 = 0$; since $\nabla^{0,1}$ already satisfies the Leibniz rule, it follows that $\nabla^{0,1}$ is a holomorphic structure. □

Observe that if X is a Riemann surface then its real dimension is 2, meaning there are no $(0,2)$-forms. It follows that $\nabla^{0,1}$ is always a holomorphic structure, so Hermitian connections correspond exactly to holomorphic structures on vector bundles over Riemann surfaces.

If we endow our underlying complex manifold X with a Kähler form ω, we can naturally extend the operator $\Lambda : \Omega^{p,q}(X) \to \Omega^{p-1,q-1}(X)$ to $\Omega^{p,q}(E)$. If we do this, we get an analogue of the Kähler identities, called the Nakano identities:

Proposition 2.28 (Nakano)**.** *Let X be a Kähler manifold, let E be a Hermitian vector bundle over X, and let ∇ be a unitary connection on E. Then the following identities holds:*

$$\begin{aligned} [\Lambda, \nabla^{0,1}] &= i(* \circ \nabla^{0,1} \circ *), \\ [\Lambda, \nabla^{1,0}] &= -i(* \circ \nabla^{1,0} \circ *). \end{aligned} \tag{2.30}$$

Proof. See [Huy05]. □

2.5 Divisors on Riemann Surfaces

It is also possible to completely classify holomorphic line bundles using only the structure of the underlying complex manifold. The fundamental objects for this classification are called divisors. (We henceforth restrict discussion to compact Riemann surfaces, as the theory is considerably simpler. The natural setting for the study of arbitrary complex manifolds is algebraic geometry; see [Vak24] for an overview of divisors in this light.)

Definition 2.29. Let X be a compact Riemann surface. A *divisor on X* is an element of the free abelian group generated by the points on X; that is, a divisor D is a finite integral linear combination of points:

$$D = \sum_{i=1}^{m} n_i p_i, \quad n_i \in \mathbb{Z}, \quad p_i \in X, \quad m < \infty. \tag{2.31}$$

A divisor D is called *effective* if each of its integers is nonnegative; we write $D \geq 0$ if D is effective. The *degree of a divisor* is the sum of its integer entries:

$$\deg(D) = \sum_{i=1}^{m} n_i. \tag{2.32}$$

It is worth noting that the set of all divisors on X form a group $\mathrm{Div}(X)$ under addition, and that $\deg : \mathrm{Div}(X) \to \mathbb{Z}$ is a group homomorphism.

Intuitively, divisors are an abstraction of the zeros and poles of a meromorphic function/section. Indeed, given a meromorphic section $\sigma \in \mathscr{M}(L)$, we can define an associated divisor as follows:

$$(\sigma) = \sum_{p \in X} \deg(\sigma(p))p, \tag{2.33}$$

where $\deg(\sigma(p))$ is the degree of σ at p; observe that $(\sigma) \geq 0$ if and only if σ is holomorphic. In the special case of a meromorphic function, the divisor is called a *principal divisor.*

Definition 2.30. Given a compact Riemann surface X, two divisors D and D' on X are *linearly equivalent* if their difference is a principal divisor, that is, if there exists some meromorphic function $f \in \mathscr{M}(X)$ for which $(f) = D - D'$.

The principal divisors form a subgroup $\mathrm{PrDiv}(X)$ of $\mathrm{Div}(X)$; the sum of two principal divisors corresponds to the product of their respective functions. Thus, an alternative definition of a linear equivalence class is an element of the quotient group $\mathrm{Div}(X)/\mathrm{PrDiv}(X)$. On the other hand, the class of holomorphic line bundles over a given Riemann surface also constitutes a group:

Proposition 2.31. *Let* $\mathrm{Pic}(X)$ *denote the set of holomorphic line bundles over* X *up to isomorphism. Then, with the (complex) tensor product as a binary operation,* $\mathrm{Pic}(X)$ *constitutes a group (called the* Picard group*).*

Proof. The dimension of a tensor product of m- and n-dimensional spaces is mn, so the tensor product of two line bundles is another line bundle. Furthermore, the tensor product is associative. The identity element is given by the trivial bundle $X \times \mathbb{C}$, and the inverse of a line bundle L is given by L^*; this is consistent because $L^* \otimes L \cong \mathrm{End}(L) \cong X \times \mathbb{C}$. □

With this observation, there is a precise correspondence between line bundles and divisors:

Theorem 2.32. *For any* $D \in \mathrm{Div}(X)$*, there is some holomorphic line bundle* $L_D \to X$ *with a section* $\sigma \in \Gamma(L_D)$ *for which* $D = (\sigma)$*. This defines a map taking* $\mathrm{Div}(X) \mapsto \mathrm{Pic}(X)$*, which is a group epimorphism with kernel* $\mathrm{PrDiv}(X)$*. (It follows from the isomorphism theorems that* $\mathrm{Pic}(X) \cong \mathrm{Div}(X)/\mathrm{PrDiv}(X)$*.)*

Proof. The following proof is from [Don11]. First, let $p \in X$ be an arbitrary point; we shall construct a line bundle over X which admits a holomorphic section vanishing only at p. To do this, take an open cover $\{U_k\}$ of X and holomorphic functions $f_k : U_k \to \mathbb{C}$ with degree 1 at p and degree 0 elsewhere. We then define L_p to be locally trivial on each U_k, with transition functions given as follows:

$$g_{k\ell} = f_k/f_\ell. \tag{2.34}$$

Observe that each such $g_{k\ell}$ is holomorphic and nowhere-zero, and they jointly satisfy the cocycle condition (Equation 2.15). It follows that $g_{k\ell}$ may be used to define a system of holomorphic transition functions, giving rise to a holomorphic line bundle L_p. Furthermore, since the transition function $g_{k\ell}$ transforms f_ℓ into f_k, the collection of all such f_k can be considered to be local representations of a global holomorphic section σ_p with a single zero at p. If we instead consider the dual bundle L_p^* then, since the transition functions are inverted for the dual bundle, there is instead a meromorphic section on L_p^* with a single pole at p; it is given in a local trivialisation over U_k by f_k^{-1}.

Now, let $D = \sum_i n_i p_i$ be a divisor on X. We define the corresponding line bundle to be the following:

$$L_D = \bigotimes_i (L_{p_i})^{\otimes n_i}, \tag{2.35}$$

where a negative tensor power is the positive tensor power of the dual; there is a corresponding meromorphic section given by

$$\sigma_D = \bigotimes_i (\sigma_{p_i})^{\otimes n_i}. \tag{2.36}$$

The zeros and poles of σ_D are given by the zeros and poles of each $\sigma_{p_i}^{\otimes n_i}$, which is clearly n_i at p_i and 0 elsewhere. It follows that $(\sigma_D) = D$.

The map $D \mapsto L_D$ is clearly a group homomorphism. We verify that it is an isomorphism:

- We first show that its kernel consists of principal divisors. Assume that L_D is in the kernel for some divisor D, i.e. L_D is trivial. Then L_D has a global section 1, so $\sigma_D = f \cdot 1$ for some meromorphic function $f : X \to \mathbb{C}$. It follows that $D = (\sigma_D)$ is a principal divisor. Moreover, for every principal divisor (f), we can construct a nowhere-zero holomorphic section $f^{-1} \cdot \sigma_{(f)}$ of the bundle $L_{(f)}$, making $L_{(f)}$ a trivial bundle.

- Secondly, we show that the homomorphism is surjective. But this follows from the fact that every line bundle admits a meromorphic section. (This result is nontrivial but we will not prove it; see [Don11] for this part of the proof.)

□

Given these three correspondences, we can describe a holomorphic structure on a given Hermitian line bundle L_h over a Riemann surface X in four different ways:

- As a restriction on the allowed maps on the bundle (through the given holomorphic trivialisations and transition functions in Definition 2.20);

- As a differential operator $\overline{\partial}_L$, which essentially serves as an extension of the Dolbeault operator $\overline{\partial}$ to L_h (as in Definition 2.21);

- As a special connection ∇ on the line bundle, under which h is parallel (as in Theorem 2.27);

- As a finite linear combination of points in X (up to linear equivalence), representing the zeros and poles of a meromorphic section on the line bundle (as in Theorem 2.32).

2.6 The First Chern Class

Finally, we summarise some basic theory regarding the topology of line bundles. Naturally, such a topic will rely heavily on algebraic topology, and in our case we will be using Čech cohomology theory extensively; for a summary of the notation and the main theorems we will be using, see the Appendix.

It turns out that the topology of line bundles is completely classified by a single invariant, a singular cohomology class called the first Chern class. There are several equivalent definitions of the first Chern class, and we will present the definition which most clearly relates to line bundles. We need the following lemma:

Lemma 2.33. *Let X be a complex manifold. Every line bundle over X corresponds uniquely to an element of $H^1(X, \mathscr{E}^*)$, the first Čech cohomology space of X over the sheaf $\mathscr{E}^*$ of nowhere-vanishing smooth $\mathbb{C}$-valued functions on X.*

Proof. Throughout this proof, we take the group operation in $\check{C}^k(\mathscr{U}, \mathscr{F})$ to be multiplicative. Choose a good cover $\mathscr{U} = \{U_\alpha\}_{\alpha \in I}$ of X, and let $(U_\alpha, U_\beta, U_\gamma)$ be a set of intersecting open sets in $\mathscr{U}$. Let $g \in C^1(\mathscr{U}, \mathscr{E}^*)$ be a cochain for which $(\delta_1 g)(U_{\alpha\beta\gamma}) = 0$; then

$$g(U_{\alpha\beta})g(U_{\alpha\gamma})^{-1}g(U_{\beta\gamma}) = 1 \tag{2.37}$$

on $U_{\alpha\beta\gamma}$, so g satisfies the cocycle condition for line bundles. We can therefore use each $g(U_{\alpha\beta}) = g_{\alpha\beta}$ as transition functions for a line bundle, so we have a homomorphism $\ker(\delta_1) \to \mathrm{Pic}(X)$ which is clearly surjective.

The kernel of this homomorphism consists of transition functions which result in a globally trivial bundle; this happens precisely when a global section f of $\mathscr{E}^*$ can be found. Now, if the bundle made by $g \in H^1(\mathscr{U}, \mathscr{E}^*)$ is trivial, the requirement that g is a set of transition functions for f translates precisely to the statement that $g = \delta_0 f$. It follows that the kernel consists of coboundaries, meaning $\mathrm{Pic}(X) \cong \ker(\delta_1)/\mathrm{im}(\delta_0) = H^1(X, \mathscr{E}^*)$. □

Now, consider the exponential sheaf exact sequence $0 \to \underline{\mathbb{Z}} \to \mathscr{E} \to \mathscr{E}^* \to 1$, where $\underline{\mathbb{Z}}$ is the constant sheaf on X. From this short exact sequence we get a long exact sequence of cohomology spaces, including the following exact subsequence:

$$H^1(X, \mathscr{E}) \to H^1(X, \mathscr{E}^*) \xrightarrow{c_1} H^2(X, \mathbb{Z}) \to H^2(X, \mathscr{E}).$$

Because $H^1(X, \mathscr{E}) = H^2(X, \mathscr{E}) = 0$, exactness implies that the map c_1 is an isomorphism. Thus, we make the following definition:

Definition 2.34. Let $\pi : L \to X$ be a line bundle over a compact Riemann surface. L is uniquely represented by an element of $H^1(X, \mathscr{E}^*)$; the image of L under the connecting morphism $c_1 : H^1(X, \mathscr{E}^*) \to H^2(X, \mathbb{Z})$ is called the *first Chern class of L,* and is denoted by $c_1(L)$.

As we have defined it, the first Chern class is a complete invariant for smooth line bundles over Riemann surfaces. It turns out there is an important correspondence between the first Chern class and the degree of a divisor on a line bundle; namely, they are Poincaré duals.

Theorem 2.35. *Let X be a compact Riemann surface with a divisor D, and let $c_1(L_D) \in H^2(X,\mathbb{Z})$ be the first Chern class of the line bundle associated to D. Specify an orientation on X by choosing a fundamental class $[X] \in H_2(X,\mathbb{Z})$. Then $\langle c_1(L_D), [X]\rangle = \deg(D)$.*

Proof. Refer to [Huy05]. □

We can also compute the first Chern class analytically, using connections on our line bundles:

Theorem 2.36. *Let ∇ be any connection on a line bundle L over X. Then the de Rham cohomology class of $(i/2\pi)F^\nabla \in \Omega^2(X)$ corresponds to a class in $H^2(X,\mathbb{Z}) \hookrightarrow H^2(X,\mathbb{R})$, and this class is precisely $c_1(L)$.*

Proof. We use the fact that the sheaves of differential k-forms Ω^k are all fine, meaning their nontrivial cohomology vanishes. We also use the snake lemma to produce an explicit formula for the connecting homomorphisms in homology sequences. For details, see the Appendix.

We choose a good open cover $\{U_\alpha\}_{\alpha\in I}$ of X, and we refine it so that L is locally trivial over each U_α. Then L corresponds to a set of transition functions $g_{\alpha\beta} : U_{\alpha\beta} \to \mathbb{C}^*$, and by Lemma 2.33, these transition functions correspond to an element of $H^1(X,\mathscr{O}^*)$. By the snake lemma, the corresponding element of $H^2(X,\mathbb{Z})$ is as follows:

$$c_1(L) = \left\{\frac{1}{2\pi i}(\log(g_{\alpha\beta}) + \log(g_{\beta\gamma}) + \log(g_{\gamma\alpha}))\right\}_{\alpha,\beta,\gamma\in I}. \tag{2.38}$$

Here, by $\log(f)$, we are referring to any function for which $e^{\log(f)} = f$. This choice is not unique, but it is guaranteed not to have an effect on $c_1(L)$ by the snake lemma. We wish to show that $[\frac{1}{2\pi i}F^\nabla]$ is equal to this class.

To do this, observe that we have the following short exact sequences of sheaves:

$$0 \to \mathbb{R} \to \Omega^0 \xrightarrow{d} \mathscr{K}^1 \to 0; \tag{2.39}$$

$$0 \to \mathscr{K}^1 \to \Omega^1 \xrightarrow{d} \mathscr{K}^2 \to 0, \tag{2.40}$$

where $\mathscr{K}^p$ denotes the sheaf of closed differential p-forms. By the long exact cohomology sequence and the fineness of Ω^p, we obtain the following exact sequences:

$$0 \to H^1(X,\mathscr{K}^1) \xrightarrow{a} H^2(X,\mathbb{R}) \to 0, \tag{2.41}$$

$$\Omega^1 \xrightarrow{d} \mathscr{K}^2 \xrightarrow{b} H^1(X,\mathscr{K}^1) \to 0, \tag{2.42}$$

where $a : H^1(X,\mathscr{K}^1) \to H^2(X,\mathbb{R})$ and $b : \mathscr{K}^2 \to H^1(X,\mathscr{K}^1)$ are connecting homomorphisms. The first exact sequence shows that a is an isomorphism, and the second shows that b induces an isomorphism from $H^2_{\mathrm{dR}}(X) \to H^1(X,\mathscr{K}^1)$. We claim that $a \circ b : H^2_{\mathrm{dR}}(X) \to H^2(X,\mathbb{R})$ takes the cohomology class of $\frac{i}{2\pi}F^\nabla$ to an integral cohomology class which corresponds to $c_1(L)$.

Given U_α, let $\sigma_\alpha \in \Gamma_{U_\alpha}(L)$ be nonvanishing. We write F^∇ in terms of connection forms; that is, $F^\nabla|_{U_\alpha} = \mathrm{d}A_\alpha$ where $A_\alpha \in \Omega^1(U_\alpha)$ satisfies $A_\alpha\sigma_\alpha = \nabla\sigma_\alpha$. By the snake lemma, we find that

$$b([F^\nabla]) = \{A_\alpha - A_\beta\}_{\alpha,\beta\in I}. \tag{2.43}$$

This is clearly a well-defined cocycle in $H^1(X, \mathscr{K}^1)$. However, we have the following computation for any $\alpha, \beta \in I$:

$$(A_\alpha - A_\beta)(\sigma_\alpha) = \nabla\sigma_\alpha - g_{\beta\alpha}\nabla\sigma_\beta = \nabla(g_{\beta\alpha}e_\beta) - g_{\beta\alpha}\nabla\sigma_\beta = (\mathrm{d}g_{\beta\alpha})\sigma_\beta = -(g_{\alpha\beta}^{-1}\,\mathrm{d}g_{\alpha\beta})\sigma_\alpha. \tag{2.44}$$

Since σ_α spans $L|_{U_\alpha}$, it follows that $A_\alpha - A_\beta = -\mathrm{d}\log g_{\alpha\beta}$, so $b([F^\nabla]) = \{-\mathrm{d}\log g_{\alpha\beta}\}_{\alpha,\beta\in I}$. We can find the explicit form of the isomorphism a using the snake lemma, and we conclude that

$$(a \circ b)([F^\nabla]) = \{-\log(g_{\alpha\beta}) - \log(g_{\beta\gamma}) - \log(g_{\gamma\alpha})\}_{\alpha,\beta,\gamma\in I}. \tag{2.45}$$

Dividing by $2\pi i$ concludes the proof. □

Thus, we have several different ways of thinking about the topology of line bundles on Riemann surfaces:

- As a cohomology class in $H^2(X, \mathbb{Z})$, corresponding to the class in $H^1(X, \mathscr{E}^*)$ representing the line bundle.
- As an integer $d \in \mathbb{Z} \cong H^0(X, \mathbb{Z})$, corresponding to the net number of zeros and poles a meromorphic section is allowed to have.
- As the integral cohomology class associated to the curvature 2-form of any connection.

We have only needed the first Chern class to classify line bundles, but the more general notion of k-th Chern classes is important for classifying higher-rank vector bundles. For a more comprehensive discussion of Chern classes and their relationship to vector bundles, see [Hat03].

Chapter 3

Vortex Equations

Now that we have defined the relevant background, we are able to discuss the vortex equations themselves; it is the purpose of this chapter to define and understand these equations. We start by defining the Yang-Mills-Higgs functional, and showing that its minima can be characterised as solutions to the vortex equations. Then, we provide a brief discussion of moment maps, before showing that the vortex equations can be naturally interpreted in terms of a moment map. We will find a moment map which represents the 0-vortex equations, and then we will slightly alter the moment map to represent the general vortex equations.

Throughout the rest of this chapter, we will be working with a fixed compact Riemann surface X with metric g and associated volume/Kähler form vol_g, and a Hermitian line bundle $L_h \to X$ of degree $d \in \mathbb{Z}_+$ with Hermitian inner product h. (We choose the degree to be nonnegative so that holomorphic solutions exist.) The gauge group $\mathscr{G} = \mathrm{Map}(X, \mathrm{U}(1))$ acts on $\Gamma(L)$ by inverted multiplication and on $\mathscr{A}^h$ by conjugation, and its Lie algebra is $\mathfrak{g} = i\mathscr{E}(X)$. The complexified gauge group $\mathscr{G}^{\mathbb{C}} = \mathrm{Map}(X, \mathbb{C}^*)$ acts similarly, though its Lie algebra is given by $\mathfrak{g}^{\mathbb{C}} = \mathrm{Map}(X, \mathbb{C})$.

3.1 The Yang-Mills-Higgs Functional

With this structure, we are able to define the functional:

Definition 3.1. The *Yang-Mills-Higgs functional* is a functional depending on a real number τ (called the *vortex parameter*), which takes a section $\phi \in \Gamma(L)$ and a connection $\nabla : \Gamma(L) \to \Omega^1(L)$, and returns a nonnegative real number according to the following formula:

$$\mathrm{YMH}_\tau(\nabla, \phi) = \left\|F^\nabla\right\|_{L^2}^2 + \|\nabla\phi\|_{L^2}^2 + \tfrac{1}{4}\left\||\phi|^2 - \tau\right\|_{L^2}^2. \tag{3.1}$$

(Where the context is clear, we will omit the L^2-subscript on norms.)

We are interested in the minima of this functional up to gauge equivalence; these will be called *τ-vortices*. In order to find these minima, it will be convenient to rewrite the functional.

Proposition 3.2. *The Yang-Mills-Higgs functional can be rewritten as follows:*

$$\mathrm{YMH}_\tau(\nabla,\phi) = 2\left\|\nabla^{0,1}\phi\right\|_{L^2}^2 + \left\|i\Lambda F^\nabla + \tfrac{1}{2}(|\phi|_h^2 - \tau)\right\|_{L^2}^2 + 2\pi\tau d. \tag{3.2}$$

Proof. We make use of the Nakano identities (Proposition 2.28). In particular, since X has real dimension 2, they reduce to the following: $\Lambda\nabla^{1,0} = -i(*\circ\nabla^{1,0}\circ *)$ and $\Lambda\nabla^{0,1} = i(*\circ\nabla^{0,1}\circ *)$. Note also that $-(*\circ\nabla^{0,1}\circ *)$ is precisely the formal adjoint of $\nabla^{1,0}$ with respect to the inner product on k-forms, and likewise $-(*\circ\nabla^{1,0}\circ *)$ is the formal adjoint of $\nabla^{0,1}$.

We expand the second term in Equation 3.1:

$$\left\|i\Lambda F^\nabla + \tfrac{1}{2}(|\phi|_h^2 - \tau)\right\|^2 = \left\|\Lambda F^\nabla\right\|^2 + \left\|\tfrac{i}{2}(|\phi|_h^2 - \tau)\right\|^2 + 2\,\mathrm{Re}(\langle i\Lambda F^\nabla, \tfrac{1}{2}|\phi|_h^2\rangle - \langle i\Lambda F^\nabla, \tfrac{\tau}{2}\rangle). \tag{3.3}$$

The first term is clearly just $\|F^\nabla\|^2$ and the second is clearly $\frac{1}{4}\left\||\phi|_h^2 - \tau\right\|^2$. Note also that $[iF^\nabla] = 2\pi c_1(L) \in H^2(X,\mathbb{R})$, meaning every remaining term is real (so they are equal to their real parts, and we can omit the projection onto the real part in Equation 3.3). We simplify the remaining two terms separately:

- $\langle i\Lambda F^\nabla, |\phi|_h^2\rangle$: Expanding this using our definitions, this is equal to

$$\int_X i\Lambda F^\nabla h(\phi,\phi)\mathrm{vol}_g. \tag{3.4}$$

  On the other hand, since $i\Lambda F^\nabla \in \Omega^0(X)$ is simply a real-valued function, it can be moved inside of h; thus, the term reduces to

$$\int_X h(i\Lambda F^\nabla\phi,\phi)\mathrm{vol}_g = \left\langle i\Lambda F^\nabla\phi,\phi\right\rangle. \tag{3.5}$$

  On the other hand, $\Lambda F^\nabla\phi = \Lambda\nabla(\nabla\phi) = \Lambda(\nabla^{1,0}\nabla^{0,1} + \nabla^{0,1}\nabla^{1,0})\phi$ on a Riemann surface. We now use the Nakano identities to commute Λ through these operators, and then we apply the formal adjoint property; we eventually find that

$$\begin{aligned}\left\langle i\Lambda F^\nabla\phi,\phi\right\rangle &= i\left\langle(-i(*\circ\nabla^{1,0}\circ *)\nabla^{0,1} + i(*\circ\nabla^{1,0}\circ *)\nabla^{1,0})\phi,\phi\right\rangle \\ &= \left\langle\nabla^{1,0}\phi,\nabla^{1,0}\phi\right\rangle - \langle\nabla^{0,1}\phi,\nabla^{0,1}\phi\rangle \\ &= \|\nabla^{1,0}\phi\|^2 - \|\nabla^{0,1}\phi\|^2.\end{aligned} \tag{3.6}$$

  We rewrite this slightly by observing that $\|\nabla\phi\|^2 = \|\nabla^{1,0}\phi + \nabla^{0,1}\phi\|^2 = \|\nabla^{1,0}\phi\|^2 + \|\nabla^{0,1}\phi\|^2 + 2\,\mathrm{Re}(\langle\nabla^{1,0}\phi,\nabla^{0,1}\phi\rangle)$, but this last term is 0 (since $\Omega^{0,1}$ and $\Omega^{1,0}$ are orthogonal with respect to the induced inner product). As such, we find that

$$\langle i\Lambda F^\nabla, |\phi|_h^2\rangle = \|\nabla\phi\|^2 - 2\|\nabla^{0,1}\phi\|^2. \tag{3.7}$$

- $\langle i\Lambda F^\nabla,\tau\rangle$: Using the fact that $[iF^\nabla/2\pi] = c_1(L)$, we compute this term explicitly:

$$\begin{aligned}\langle i\Lambda F^\nabla,\tau\rangle &= \int_X i\tau\Lambda F^\nabla\mathrm{vol}_g \\ &= 2\pi\tau\int_X \frac{i}{2\pi}F^\nabla = 2\pi\tau d.\end{aligned} \tag{3.8}$$

Substituting each term into Equation 3.3, and the resulting term into Equation 3.2, the equivalence follows. □

In this form, it is clear that the Yang-Mills-Higgs functional is bounded below by $2\pi\tau d$. If this lower bound is achieved by some (∇, ϕ), then they necessarily satisfy the following equations:

$$\begin{aligned} \nabla^{0,1}\phi &= 0, \\ \Lambda F^{\nabla} - \frac{i}{2}(|\phi|_h^2 - \tau) &= 0. \end{aligned} \tag{3.9}$$

These equations are called the *τ-vortex equations.* Observe that the equations are also invariant under the action of $\mathscr{G}$, meaning a solution (∇, ϕ) and a unitary gauge transformation $g \in \mathscr{G}$ give rise to another solution $(g^{-1}\nabla g, g^{-1}\phi)$. Quotienting the space of solutions by $\mathscr{G}$, we get the *moduli space of τ-vortices.*

Note that, by the Chern correspondence, the first equation simply states that ϕ is holomorphic with respect to the holomorphic structure induced by ∇. It is therefore the second equation which is of the most interest, and the most difficult to solve. Nevertheless, we can obtain a constraint on the existence of solutions fairly easily:

Proposition 3.3. *If there is a solution to the τ-vortex equations, then $d < \frac{\tau \mathrm{Vol}(X)}{4\pi}$ (where $\mathrm{Vol}(X) = \int_X \mathrm{vol}_g$).*

Proof. Let (∇, ϕ) be a solution to the second equation. Multiplying through by $(i/2\pi)\mathrm{vol}_g$ and integrating, we see that

$$\int_X \frac{i}{2\pi} F^{\nabla} + \int_X \frac{1}{4\pi}(|\phi|_h^2 - \tau)\mathrm{vol}_g = 0. \tag{3.10}$$

But since $|\phi|_h^2 \geq 0$, this means that

$$\int_X \frac{i}{2\pi} F^{\nabla} < \frac{\tau}{4\pi}\mathrm{Vol}(X). \tag{3.11}$$

The result follows from the observation that $[(i/2\pi)F^{\nabla}] = c_1(L)$. □

This proposition sheds light on the parameter τ: if τ is sufficiently large, we are able to avoid an obstruction which depends on the volume of X. We therefore naturally interpret τ as a scaling factor for $\mathrm{Vol}(X)$, and we will eventually find that the moduli space does not depend on the value of τ we choose (insofar as τ satisfies the constraint we just derived).

3.2 Moment Maps

It turns out that the vortex equations can be naturally interpreted in terms of a map on the space of sections and connections, called a *moment map.* Thus, in this section, we introduce the concept of a moment map. The basic idea comes from physics, in which the state of a system evolves based on an energy function H. Geometrically, this idea manifests as a symplectic manifold, wherein the space of all states is a manifold M with a symplectic form ω, and the dynamics of the state space is given by the flow of the Hamiltonian vector field

X_H (which is ω-dual to the covector field $\mathrm{d}H$, i.e. $X_H \lrcorner\, \omega = \mathrm{d}H$). A well-known theorem in physics is Noether's theorem, which states that every differentiable symmetry of the laws of physics (represented by a Lie group G acting on the state space) corresponds to a quantity which is conserved over time. The corresponding symplectic notion is the moment map.

Definition 3.4. Let (M, ω) be a (possibly infinite-dimensional) symplectic manifold, and let G be a Lie group which acts symplectically on M (i.e. for any $g \in G$ we have $g^*\omega = \omega$). Let $\mathfrak{g}$ be the Lie group of G, and given $\xi \in \mathfrak{g}$, denote by $\tilde{\xi} \in \Gamma(TM)$ the *fundamental vector field of* ξ (i.e. the vector field generated by the flow $(t, p) \mapsto \exp(t\xi) \cdot p$; explicitly, $\tilde{\xi}_p = \frac{\mathrm{d}}{\mathrm{d}t}|_{t=0}(e^{t\xi} \cdot p)$). Then a *moment map* is a smooth map $\mu : M \to \mathfrak{g}^*$ satisfying the following conditions:

- *Equivariance:* for any $g \in G$ and any $p \in M$, the map μ must satisfy

$$\mu(g \cdot p) = \mathrm{Ad}_g^*(\mu(p)). \tag{3.12}$$

- *Hamiltonian property:* for any $\xi \in \mathfrak{g}$, we must have

$$\mathrm{d}\langle \mu, \xi \rangle = \tilde{\xi} \lrcorner\, \omega. \tag{3.13}$$

We also have the dual notion of a co-moment map:

Definition 3.5. Let (M, ω) be a symplectic manifold on which a Lie group G (with Lie algebra $\mathfrak{g}$) acts symplectically. A *co-moment map* is a smooth map $\mu^* : \mathfrak{g} \to C^\infty(M)$ satisfying the following conditions:

- If $C^\infty(M)$ is equipped with the Poisson bracket, then μ^* is a Lie algebra homomorphism.
- For any $\xi \in \mathfrak{g}$, we must have that $\mathrm{d}\mu^*(\xi) = \tilde{\xi} \lrcorner\, \omega$.

A moment map gives rise to a co-moment map by taking $\mu^*(\xi) = \langle \mu, \xi \rangle$.

There are some circumstances in which a moment map can necessarily be found. For example, let ξ be an element of the Lie algebra with fundamental vector field $\tilde{\xi}$. The fact that G acts symplectically implies that $\mathscr{L}_{\tilde{\xi}}(\omega) = 0$, where $\mathscr{L}$ is the Lie derivative. But by Cartan's homotopy formula (and the fact that ω is closed), it follows that $\mathrm{d}(\tilde{\xi} \lrcorner\, \omega) = 0$. Now, if $H^1(M, \mathbb{R}) = 0$ (which happens whenever M is an affine space, for instance), it follows that there is some function $\mu_X \in C^\infty(M)$ such that $\mathrm{d}\mu_\xi = \tilde{\xi} \lrcorner\, \omega$. We can patch together these maps to obtain a moment map which may not be equivariant, and then we can adjust the redundant parameters to make the map equivariant. Moreover, if we can find a moment map then it is essentially unique:

Proposition 3.6. *Let $\mu, \nu : M \to \mathfrak{g}^*$ be moment maps for a connected symplectic manifold (M, ω) with a symplectic action by G. Then $\mu - \nu : M \to \mathfrak{g}^*$ is constant over M.*

Proof. Consider the co-moment map difference $\mu^* - \nu^* : \mathfrak{g} \to C^\infty(M)$. For any $\xi \in \mathfrak{g}$, we know that $\mathrm{d}(\mu^* - \nu^*)(\xi) = 0$ since they both have the Hamiltonian property. It follows that $(\mu^* - \nu^*)(\xi)$ is locally constant, and since M is connected, we have a well-defined function $c^* : \mathfrak{g} \to \mathbb{R}$ for which $\mu^*(\xi) - \nu^*(\xi) = c^*(\xi)$. It can be easily verified that c^* is linear, meaning it is an element of $\mathfrak{g}^*$, and we therefore find that $\mu - \nu = c^*$. □

We will also be interested in combining and reducing moment maps on different spaces. The following proposition (for which the proof is simply a matter of working through definitions) gives us a canonical method for combining them:

Proposition 3.7. *Let (M_1, ω_1) and (M_2, ω_2) be symplectic manifolds, and let G be a Lie group acting symplectically on both M and N. Let $\mathfrak{g}$ be the Lie algebra of G, and let $\mu_1 : M_1 \to \mathfrak{g}^*$ and $\mu_2 : M_2 \to \mathfrak{g}^*$ be corresponding moment maps. Then $(M_1 \times M_2, \omega_1 \oplus \omega_2)$ is a symplectic manifold, and $\mu_1 + \mu_2$ is a moment map.*

Moreover, the following proposition (whose proof is found in [Mar+07]) allows us to combine moment maps of different actions:

Proposition 3.8. *Let (M, ω) be a symplectic manifold, and let G, H be Lie groups acting symplectically on M with Lie algebras $\mathfrak{g}, \mathfrak{h}$, in such a way that the actions commute with each other (which guarantees that there is a canonical action of $G \times H$ on M). Let $\mu_G : M \to \mathfrak{g}^*$ and $\mu_H : M \to \mathfrak{h}^*$ be corresponding moment maps, and assume that μ_G is H-invariant and μ_H is G-invariant. Then $\mu_G \times \mu_H : M \to \mathfrak{g}^* \times \mathfrak{h}^*$, defined so that $(\mu_G \times \mu_H)(p) = \mu_G(p) + \mu_H(p)$, is a moment map for $G \times H$.*

Proof. The Hamiltonian property follows simply by the observation that $\widetilde{(\xi, \eta)} = \tilde{\xi} + \tilde{\eta}$ for $\xi \in \mathfrak{g}$ and $\eta \in \mathfrak{h}$, which means that $\widetilde{(\xi, \eta)} \lrcorner\, \omega = \tilde{\xi} \lrcorner\, \omega + \tilde{\eta} \lrcorner\, \omega = \mathrm{d}\langle \mu_G, \xi \rangle + \mathrm{d}\langle \mu_H, \eta \rangle = \mathrm{d}\langle \mu_G \times \mu_H, (\xi, \eta) \rangle$. Equivariance follows from the equivariance of each moment map with respect to their own group, and the invariance of each moment map with respect to the complementary group. □

It follows from these results that the moment map for a quotient of Lie groups corresponds to the difference between the two moment maps, if it exists.

In fact, the main practical reason we are interested in moment maps is that they allow us to reduce a problem of finding zeros to finding minima:

Theorem 3.9. *Let (M, g, ω, J) be a Kähler manifold on which a Lie group G acts symplectically and isometrically, and endow the Lie algebra $\mathfrak{g}$ with a G-invariant inner product $(\cdot, \cdot)$. Let $\mu : M \to \mathfrak{g}^*$ be a moment map for G. If $\|\mu\|^2$ has a minimum at p when restricted to an orbit of $G^{\mathbb{C}}$ (the complexification of G), and if the isotropy subgroup of $G^{\mathbb{C}}$ at p is trivial, then μ has a zero on this orbit.*

Proof. Throughout this proof, we identify $\mathfrak{g}$ and $\mathfrak{g}^*$ using the inner product. Define $f : M \to \mathbb{R}$ by taking $f(p) = \|\mu(p)\|^2 = (\mu(p), \mu(p))$. For any vector field $v \in \Gamma(TM)$, we have the following:

$$g(\mathrm{grad}(f)_p, v_p) = \mathrm{d}f_p(v_p) = \mathrm{d}_{v_p}(\mu(p), \mu(p)) = 2(\mathrm{d}_{v_p}\mu(p), \mu(p)), \tag{3.14}$$

where we have used the Leibniz rule and the symmetry of $(\cdot, \cdot)$ in the last equality. On the other hand, if $\xi \in \mathfrak{g}$ does not depend on the point in M, we know that $\mathrm{d}_{v_p}\langle \mu(p), \xi \rangle = \langle \mathrm{d}_{v_p}\mu(p), \xi \rangle$. It follows from the Hamiltonian property that

$$g(\mathrm{grad}(f)_p, v_p) = 2(\widetilde{\mu(p)} \lrcorner\, \omega)(v_p), \tag{3.15}$$

where $\widetilde{\mu(p)}$ is the fundamental vector field of $\mu(p) \in \mathfrak{g}$ evaluated at p. One can use the compatibility of g and ω to show that $\omega(\widetilde{\mu(p)}, v_p) = g(J\widetilde{\mu(p)}, v_p)$. Since this holds for any v_p and g is nondegenerate, we have that $\mathrm{grad}(f) = 2J\widetilde{\mu(p)}$. Thus, the flow lines of $\mathrm{grad}(f)$ must be contained in the orbits of $G^{\mathbb{C}}$.

Next, suppose f has a minimum at p inside an orbit Γ of $G^{\mathbb{C}}$. Then $\mathrm{grad}(f)(p)$ must be orthogonal to the tangent vectors along which one remains in Γ; but by the above formula, this implies that the gradient is 0. It follows that $\widetilde{\mu(p)} = 0$. Now, suppose $\mu(p)$ is nonzero; we will relabel $\mu(p)$ by ξ, and $\widetilde{\mu(p)}$ by $\tilde{\xi}_p$. We then have that $\frac{\mathrm{d}}{\mathrm{d}t}(e^{t\xi}p) = 0$, which implies that $\frac{\mathrm{d}}{\mathrm{d}t}(e^{t\xi}(e^{s\xi}p)) = 0$ for any $s \in \mathbb{R}$. Observe that the set of all such $e^{s\xi}p$ is a 1-dimensional submanifold N of M on which $\tilde{\xi} = 0$ (they must all be distinct points since p has trivial isotropy), and this clearly forces the action of $e^{t\xi}$ on N to be trivial for all t. But we claimed that the isotropy subgroup of $G^{\mathbb{C}}$ at $p \in N$ was finite, which is a contradiction. It follows that $\mu(p)$ must be zero. □

We will eventually show that the solutions to the vortex equations correspond precisely to the zeros of a certain moment map, and we will therefore find solutions by minimising $\|\mu\|^2$ on each orbit of $\mathscr{G}^{\mathbb{C}}$ and applying Proposition 3.9. Furthermore, we will find that there is *exactly one* zero on each orbit of $\mathscr{G}^{\mathbb{C}}$, up to unitary gauge equivalence. Recall that we quotient the space of solutions by unitary gauge transformations to obtain the moduli space; to this end, the following theorem will be useful:

Theorem 3.10 (Marsden-Weinstein). *Let (M, ω) be a symplectic manifold with a symplectic action by a Lie group G, and let $\mu : M \to \mathfrak{g}^*$ be a moment map for this action. Suppose the action of G is free and properly discontinuous. Define $c \in \mathfrak{g}^*$ to be* regular *if $\mu^{-1}(c)$ is a manifold and μ_{*p} is surjective for every $p \in M$. Then, for any regular value $c \in \mathfrak{g}^*$, the quotient space $\mu^{-1}(c)/G$ is a manifold. Further, the manifold inherits a symplectic structure from M.*

Proof. See [MW74]. □

3.3 The Yang-Mills-Higgs Functional as a Moment Map

The space of solutions to the vortex equations is a subspace of $\mathscr{A}^h \times \Gamma(L)$, the space of pairs of unitary connections and sections of L. Moreover, by taking $\nabla \mapsto g^{-1}\nabla g$ and $\phi \mapsto g^{-1}\phi$ for each $g \in \mathscr{G}$, we get a natural action of $\mathscr{G}$ on this space. In this section, we will construct a moment map for this space with respect to $\mathscr{G}$. We will do so by considering both $\mathscr{A}^h$ and $\Gamma(L)$ independently, and demonstrating that they are infinite-dimensional (Hilbert) manifolds. Then, we will put a symplectic structure on them which is compatible with the action of $\mathscr{G}$, and we will then explicitly construct a moment map for each action. Finally, we will use Proposition 3.7 to combine them.

We start by analysing $\mathscr{A}^h$. By the Chern correspondence (Theorem 2.27), we can characterise its structure using the following observation.

Proposition 3.11. *Let L_h be a Hermitian line bundle over the Riemann surface X, and let $\mathscr{C}$ be the space of holomorphic structures on L. Then $\mathscr{C}$ is an affine space modelled on $\Omega^{0,1}(X)$.*

Proof. We must prove three things:

- $\mathscr{C}$ is nonempty. This follows from the fact that every complex manifold has a unitary connection (we can always define a unitary connection locally, and then stitch them together using a partition of unity), of which we can take the $(0,1)$-part to get a holomorphic structure.
- The difference of any two elements $\overline{\partial}_L, \overline{\partial}_L' \in \mathscr{C}$ is a $(0,1)$-form. It follows quickly from the Leibniz rule that $\overline{\partial}_L - \overline{\partial}_L'$ is homogeneous of degree 1 over $C^\infty(X)$, and it is clearly additive, so it is tensorial. Moreover, since it takes as input a tangent vector and a section of L and returns a section of L, it is a section of $\Omega^1(\mathrm{End}(L)) \cong \Omega^1(X)$ (since $\mathrm{End}(L) \cong X \times \mathbb{C}$, being a trivial complex line bundle). It is also easy to check that $\overline{\partial}_L - \overline{\partial}_L'$ annihilates $(1,0)$-vectors, meaning it is a $(0,1)$-form.
- Given $\overline{\partial}_L \in \mathscr{C}$ and $\alpha \in \Omega^{0,1}(X)$, the sum $\overline{\partial}_L + \alpha$ is another holomorphic structure. But this follows simply from the linearity of all of the operators involved, and from the fact that $\overline{\partial}_L$ already satisfies the Leibniz rule. □

By choosing a reference point in $\mathscr{C}$, we can identify $\mathscr{C}$ with $\Omega^{0,1}(X)$; moreover, the latter space is a vector space with a natural inner product (given by Equation 2.11). It follows that $\mathscr{C}$ has the natural topology of an infinite-dimensional manifold. We can therefore also think of $\mathscr{A}^h$ as an affine space over $\Omega^{0,1}(X)$; in fact, we have the following:

Proposition 3.12. *Let $\overline{\partial}_L \in \mathscr{C}$ be a holomorphic structure, and let ∇ be the Chern connection for $\overline{\partial}_L$. If $\overline{\partial}_L \mapsto \overline{\partial}_L + \alpha$ for $\alpha \in \Omega^{0,1}(X)$, then $\nabla \mapsto \nabla + \alpha - \overline{\alpha}$.*

Proof. The Chern connection for $\overline{\partial}_L + \alpha$ is the unique h-unitary connection with $(0,1)$-part given by $\overline{\partial}_L + \alpha$. Thus, we only need to check that $\nabla + \alpha - \overline{\alpha}$ satisfies these conditions. We can check that h is parallel under the new connection by a computation: if $\sigma, \eta \in \Gamma(L)$, then

$$\begin{aligned}((\nabla + \alpha - \overline{\alpha})h)(\sigma,\eta) &= \mathrm{d}(h(\sigma,\eta)) - h(\nabla\sigma + \alpha\sigma - \overline{\alpha}\sigma, \eta) - h(\sigma, \nabla\eta + \alpha\eta - \overline{\alpha}\eta) \\ &= (\nabla h)(\sigma,\eta) - \alpha h(\sigma,\eta) + \overline{\alpha} h(\sigma,\eta) - \overline{\alpha} h(\sigma,\eta) + \alpha h(\sigma,\eta) \\ &= 0 + 0 = 0,\end{aligned} \tag{3.16}$$

where in the second equality we have used the sesquilinearity of h to conjugate α and $\overline{\alpha}$. Moreover, since α is a $(0,1)$-form and $\overline{\alpha}$ is a $(1,0)$-form, the $(0,1)$-part of $\nabla + \alpha - \overline{\alpha}$ is indeed $\overline{\partial}_L + \alpha$. □

Note that, because of these relationships, the tangent bundles of both $\mathscr{A}^h$ and $\mathscr{C}$ are trivial with tangent space isomorphic to $\Omega^{0,1}(X)$. As such, we will often refer to a $(0,1)$-form-valued function on $\mathscr{A}^h$ as a vector field.

As we have remarked previously (Equation 2.11), the Hodge star induces a natural inner product on the spaces of (p,q)-forms. If we think of $\mathscr{C}$ as an infinite-dimensional manifold, this amounts to the existence of a natural Hermitian form on $\mathscr{C}$: for $\alpha, \beta \in T_{\overline{\partial}_L}\mathscr{C} \cong \Omega^{0,1}(X)$, we define a Hermitian form as follows:

$$\langle \alpha, \beta \rangle = \int_X \alpha \wedge *\overline{\beta}. \tag{3.17}$$

We also get a natural 2-form by essentially taking the imaginary part:

$$\omega(\alpha,\beta) = i(\langle\alpha,\beta\rangle - \langle\beta,\alpha\rangle) = i\int_X \alpha\wedge *\overline{\beta} - \beta\wedge *\overline{\alpha}. \tag{3.18}$$

Since the 2-form ω does not depend on the point $\overline{\partial}_L$, it is clear that $\mathrm{d}\omega = 0$ and hence ω is a symplectic form. We therefore find that $\mathscr{A}^h$ is a symplectic manifold (and indeed a Kähler manifold).

As we have mentioned previously, there is a natural action of $\mathscr{G}$ on $\mathscr{A}^h$. There is also a natural action on $\mathscr{C}$ given by $\overline{\partial}_L \mapsto g^{-1}\overline{\partial}_L g$, and it is easy to show that these two actions are compatible under the Chern correspondence. In fact, one can show that the action of $g \in \mathscr{G}$ on $\mathscr{C}$ is simply a translation by $g^{-1}\overline{\partial}g \in \Omega^{0,1}(X)$, meaning the tangent spaces $\Omega^{0,1}(X)$ are invariant under the action of $\mathscr{G}$, so $\langle\cdot,\cdot\rangle$ and ω are also invariant under $\mathscr{G}$. In short, the gauge group $\mathscr{G}$ acts symplectically on $\mathscr{A}^h$. This indicates the possibility of a moment map on $\mathscr{A}^h$ with respect to the Lie group $\mathscr{G}$. Indeed, the following moment map was found in [AB83]:

Proposition 3.13. *Define a map $\mu : \mathscr{A}^h \times \mathfrak{g} \to \mathbb{R}$ as follows: if $i\xi \in \mathfrak{g}$ and $\nabla \in \mathscr{A}^h$, then*

$$\mu(\nabla, i\xi) = \int_X i\xi F^\nabla. \tag{3.19}$$

Then μ induces a moment map with respect to $\mathscr{G}$.

Proof. First, we show that μ is $\mathscr{G}$-equivariant, which is to say that $\mu(\nabla^g, \mathrm{Ad}_g(i\xi)) = \mu(\nabla, i\xi)$. Since $\mathscr{G}$ is abelian, we know that the adjoint action on $\mathfrak{g}$ is trivial. Moreover, we know that $F^{\nabla^g} = g^{-1}F^\nabla g$, and once again the commutativity of $\mathscr{G}$ shows that this is simply F^∇. Equivariance follows at once.

The main difficulty in the proof is showing that $\mu_{i\xi}$ has the Hamiltonian property for each $\xi \in \mathscr{E}(X)$. Before we can start this, we need to have a formula for the fundamental $\mathscr{A}^h$-vector field (i.e. the $(0,1)$-form-valued function on $\mathscr{A}^h$) induced by $i\xi$, which we denote by $\widetilde{i\xi}$. Let $\nabla \in \mathscr{A}^h$ be arbitrary; by definition, the vector field is given by

$$\begin{aligned}\widetilde{i\xi}_\nabla &= \frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0}\Big(e^{i\xi t}\cdot\nabla\Big)\\ &= \lim_{t\to 0}\tfrac{1}{t}\Big(e^{-i\xi t}\circ\nabla\circ e^{i\xi t} - \nabla\Big).\end{aligned} \tag{3.20}$$

This limit is most easily evaluated by applying it to a section $\sigma \in \Gamma(L)$ and using the Leibniz rule:

$$\tfrac{1}{t}\Big(e^{-i\xi t}\circ\nabla\circ e^{i\xi t} - \nabla\Big)\sigma = \tfrac{1}{t}\Big(\nabla\sigma + (it\overline{\partial}\xi)\sigma - \nabla\sigma\Big) = i\overline{\partial}\xi\sigma, \tag{3.21}$$

from which it follows that $\widetilde{i\xi}_\nabla = i\overline{\partial}\xi$.

Now that we have this, let $\alpha \in T_\nabla\mathscr{A}^h \cong \Omega^{0,1}(X)$; we wish to show that $\mathrm{d}_{\alpha_\nabla}\langle\mu, i\xi\rangle =$

$(\widetilde{i\xi} \lrcorner \omega)(\alpha)$. We first compute $\mathrm{d}_{\alpha_\nabla}\langle \mu, i\xi\rangle$ using a limit:

$$\begin{aligned}
\mathrm{d}_{\alpha_\nabla}\langle \mu, i\xi\rangle &= \lim_{t\to 0} \tfrac{1}{t} \left(\mu(\nabla + (\alpha - \overline{\alpha})t, i\xi) - \mu(\nabla, i\xi)\right) \\
&= \lim_{t\to 0} \frac{1}{t} \int_X i\xi \left(F^{\nabla + (\alpha - \overline{\alpha})t} - F^{\nabla}\right) \\
&= \lim_{t\to 0} \frac{1}{t} \int_X i\xi \left(F^{\nabla} + \mathrm{d}(\alpha - \overline{\alpha})t - F^{\nabla}\right) \\
&= \int_X i\xi \mathrm{d}(\alpha - \overline{\alpha}),
\end{aligned} \tag{3.22}$$

where we have used the fact that $F^{\nabla + a} = F^{\nabla} + \mathrm{d}a$ for any 1-form a. Our goal is to show that this is equal to $(\widetilde{i\xi} \lrcorner \omega)(\alpha)$, or equivalently $\omega(-i\overline{\partial}\xi, \alpha)$.

Expanding this expression in terms of integrals, we see that

$$(\widetilde{i\xi} \lrcorner \omega)(\alpha) = -\int_X \overline{\partial}\xi \wedge *\overline{\alpha} + \alpha \wedge *\partial\xi. \tag{3.23}$$

However, the term $\alpha \wedge *\partial\xi$ can be rewritten as $\partial\xi \wedge *\alpha$ (in fact $v \wedge *w = w \wedge *v$ for any v, w on a 1-dimensional complex vector space). Moreover, because the real dimension of X is 2 and each derivative of ξ is wedged with a 1-form, we can replace $\overline{\partial}\xi$ and $\partial\xi$ with $\mathrm{d}\xi$ wherever they appear. It follows that

$$(\widetilde{i\xi} \lrcorner \omega)(\alpha) = -\int_X \mathrm{d}\xi \wedge *\overline{\alpha} + \mathrm{d}\xi \wedge *\alpha = -\int_X \mathrm{d}\xi \wedge *(\alpha + \overline{\alpha}). \tag{3.24}$$

Integrating by parts and using Stokes' theorem, we can transfer the d over to the α terms:

$$-(\widetilde{i\xi} \lrcorner \omega)(\alpha) = \int_X \xi \mathrm{d}(*(\alpha + \overline{\alpha})). \tag{3.25}$$

Distributing, identifying d with ∂ and $\overline{\partial}$ where possible, and using the Hodge star to reintroduce the volume element, we see that

$$(\widetilde{i\xi} \lrcorner \omega)(\alpha) = \int_X \xi \left(*\partial(*\alpha) + *\overline{\partial}(*\overline{\alpha})\right)\mathrm{vol}_g. \tag{3.26}$$

On the other hand, we may use the Kähler identities $- * \overline{\partial} * = i[\Lambda, \overline{\partial}]$, and $- * \partial * = -i[\Lambda, \partial]$ to simplify this expression. Moreover, since X has real dimension 2, $\overline{\partial}\Lambda = \partial\Lambda = 0$. We can thus rewrite the expression accordingly:

$$\begin{aligned}
(\widetilde{i\xi} \lrcorner \omega)(\alpha) &= \int_X i\xi \left(\Lambda\partial\alpha - \Lambda\overline{\partial}\overline{\alpha}\right)\mathrm{vol}_g \\
&= \int_X i\xi \left(\partial\alpha - \overline{\partial}\overline{\alpha}\right) \\
&= \int_X i\xi \mathrm{d}(\alpha - \overline{\alpha}) = \mathrm{d}_{\alpha_\nabla}\langle \mu, i\xi\rangle.
\end{aligned} \tag{3.27}$$

So μ does have the Hamiltionian property, meaning it is indeed a moment map. □

We now turn attention to $\Gamma(L)$. Since $\Gamma(L)$ is a vector space with an inner product (induced by h and the volume element), we can consider it as an infinite-dimensional manifold with tangent space naturally isomorphic to $\Gamma(L)$. A symplectic structure naturally arises on $\Gamma(L)$ from this inner product: for $\sigma, \eta \in T_\phi\Gamma(L) \cong \Gamma(L)$, we define

$$\omega(\sigma, \eta) = \frac{i}{2}(\langle\sigma, \eta\rangle - \langle\eta, \sigma\rangle) = \frac{i}{2}\int_X (h(\sigma, \eta) - h(\eta, \sigma))\mathrm{vol}_g. \tag{3.28}$$

Again, since ω does not depend on the point ϕ, we have $\mathrm{d}\omega = 0$ and hence ω is a symplectic form. It follows immediately that ω is preserved by $\mathscr{G}$, since each $g \in \mathscr{G}$ is pointwise unit-length. Again, we are able to find a moment map for this action:

Proposition 3.14. *Define a map $\nu : \Gamma(L) \times \mathfrak{g} \to \mathbb{R}$ as follows: if $i\xi \in \mathfrak{g}$ and $\phi \in \Gamma(L)$, then*

$$\nu(\phi, i\xi) = -\frac{i}{2}\int_X |\phi|^2(i\xi)\mathrm{vol}_g. \tag{3.29}$$

Then ν induces a moment map with respect to $\mathscr{G}$.

Proof. It is clear that ν is well-defined and $\mathscr{G}$-equivariant. To show that it is a moment map, we must show that it satisfies the Hamiltonian property, that is, $\mathrm{d}_{\sigma_\phi}\langle\nu, i\xi\rangle = \omega(\widetilde{i\xi}_\phi, \sigma_\phi)$ for any $i\xi \in \mathfrak{g}$ and $\sigma_\phi \in T_\phi\Gamma(L) \cong \Gamma(L)$. The term $\mathrm{d}_{\sigma_\phi}\langle\nu, i\xi\rangle$ can be evaluated explicitly:

$$\begin{aligned} \mathrm{d}_{\sigma_\phi}\langle\nu, i\xi\rangle &= \lim_{t\to 0}\tfrac{1}{t}(\nu(\phi + t\sigma, i\xi) - \nu(\phi, i\xi)) \\ &= \lim_{t\to 0}\frac{1}{2t}\int_X \xi\left(|\phi + t\sigma|^2 - |\phi|^2\right)\mathrm{vol}_g \\ &= \lim_{t\to 0}\frac{1}{2t}\int_X \xi\left(th(\sigma, \phi) + th(\phi, \sigma) + t^2|\sigma|^2\right)\mathrm{vol}_g \\ &= \frac{1}{2}\int_X \xi(h(\sigma, \phi) + h(\phi, \sigma))\mathrm{vol}_g. \end{aligned} \tag{3.30}$$

To evaluate the term $\omega(\widetilde{i\xi}_\phi, \sigma_\phi)$, observe that

$$\widetilde{i\xi}_\phi = \frac{\mathrm{d}}{\mathrm{d}t}\bigg|_{t=0}(e^{-it\xi}\phi) = -i\xi\phi; \tag{3.31}$$

we then simply use the definition of ω:

$$\begin{aligned} \omega(\widetilde{i\xi}_\phi, \sigma_\phi) &= \frac{i}{2}\int_X (h(-i\xi\phi, \sigma) - h(\sigma, -i\xi\phi))\mathrm{vol}_g \\ &= \frac{1}{2}\int_X \xi(h(\phi, \sigma) + h(\sigma, \phi))\mathrm{vol}_g = \mathrm{d}_{\sigma_\phi}\langle\nu, i\xi\rangle. \end{aligned} \tag{3.32}$$

Thus, the map ν satisfies the Hamiltionian property, making it a moment map. □

We have now defined a moment map on both $\mathscr{A}^h$ and $\Gamma(L)$. By Proposition 3.7, each moment map can be combined into a single moment map on $\mathscr{A}^h \times \Gamma(L)$ by simply adding the two moment maps together:

$$\mu_{i\xi}(\nabla, \phi) = \int_X i\xi\left(\Lambda F^\nabla - \frac{i}{2}|\phi|^2\right)\mathrm{vol}_g. \tag{3.33}$$

We can write this as a map $\mu : \mathscr{A}^h \times \Gamma(L) \to \mathfrak{g}$ by defining a natural inner product on $\mathfrak{g}$ and identifying $\mathfrak{g}^*$ with $\mathfrak{g}$; for $i\xi, i\eta \in \mathfrak{g} = i\mathscr{E}(X)$, we define $\langle i\xi, i\eta\rangle = -\int_X (i\xi)(i\eta)\mathrm{vol}_g$.

Using this inner product, it is clear that the moment map is induced by the following map $\mu : \mathscr{A}^h \times \Gamma(L) \to \mathfrak{g}$:

$$\mu(\nabla, \phi) = \Lambda F^\nabla - \frac{i}{2}|\phi|^2. \tag{3.34}$$

Moreover, we restrict our attention to the following submanifold of $\mathscr{A}^h \times \Gamma(L)$:

$$\mathscr{N} = \{(\nabla, \phi) \in \mathscr{A}^h \times \Gamma(L) \,:\, \nabla^{0,1}\phi = 0,\ \phi \text{ is not identically } 0\}. \tag{3.35}$$

This subspace is clearly closed under the action of $\mathscr{G}$, and it is also a Kähler submanifold of the entire space, so the same moment map μ can be used on $\mathscr{N}$. The elements of this space automatically satisfy the first vortex equation, meaning we can focus our attention on the second.

3.3.1 Introduction of the Vortex Parameter

Observe that the integrand in μ is identical to the τ-vortex equation for $\tau = 0$, meaning the 0-vortex equations are satisfied precisely when μ is equal to 0. We would like to redefine the setup so that the τ-vortex equations may be represented for arbitrary τ. In fact, we will do this by highlighting a dimension-1 subgroup of $\mathscr{G}$ whose moment map will be a constant, and using Proposition 3.8 to subtract the resulting moment map out.

We consider the subgroup $\mathrm{U}(1) \subseteq \mathscr{G}$ consisting of constant maps. Restricting to a subgroup clearly does not affect equivariance or the Hamiltonian property, so we can get the moment map for the action of $\mathrm{U}(1)$ simply by restricting the moment map for $\mathscr{G}$. Identifying $\mathfrak{u}(1)$ with $i\mathbb{R}$ and $\mathfrak{u}(1)^*$ with $\mathfrak{u}(1)$ by using the inner product, we see that the moment map is given by

$$\tilde{\mu}(\nabla, \phi) = -\frac{i}{2\mathrm{Vol}(X)} \int_X |\phi|^2 \mathrm{vol}_g. \tag{3.36}$$

Given any $c \in i\mathbb{R}^-$, we will be able to solve the equation $\tilde{\mu}(\nabla, \phi) = c$, and it is easy to check that every such c is a regular value for $\tilde{\mu}$. Thus, by Theorem 3.10, we can reduce to $\tilde{\mathscr{N}} = \tilde{\mu}^{-1}(c)/\mathrm{U}(1)$. Furthermore, Proposition 3.8 allows us to simply subtract these moment maps to get the moment map on $\tilde{\mathscr{N}}$ for $\mathscr{G}/\mathrm{U}(1)$, which we will define to be $\tilde{\mathscr{G}}$. We subtract an extra constant term in this moment map, and conclude that a moment map for $\tilde{\mathscr{G}}$ on $\tilde{\mathscr{N}}$ is the following:

$$\mu(\nabla, \phi) = \Lambda F^\nabla - \frac{i}{2}|\phi|^2 + \frac{2\pi i d}{\mathrm{Vol}(X)} - c. \tag{3.37}$$

We define $i\tau/2$ to be this entire constant term; when this is done, the requirement that c is strictly negative is precisely the requirement that $\tau > 4\pi d/\mathrm{Vol}(X)$. It follows that, on the manifold $\tilde{\mathscr{N}}$, the moment map with respect to $\tilde{\mathscr{G}}$ is given by

$$\mu(\nabla, \phi) = \Lambda F^\nabla - \frac{i}{2}|\phi|^2 + \frac{i}{2}\tau. \tag{3.38}$$

Thus, the τ-vortex equations are satisfied precisely when $\mu(\nabla, \phi) = 0$, so the problem of finding solutions reduces to finding zeros of μ. As indicated previously, Theorem 3.9 allows us to reduce the problem further to finding minima of $|\mu|^2$ on orbits of $\tilde{\mathscr{G}}^{\mathbb{C}}$.

We have been conceptualising $\tilde{\mathscr{G}}^{\mathbb{C}}$ as the complexification of $\tilde{\mathscr{G}}$, which has a natural action on $\tilde{\mathscr{N}}$. However, we can also realise it explicitly as the group $\mathscr{G}^{\mathbb{C}}$ modulo the constant gauge transformations $\mathbb{C}^*$. Under this interpretation, the action of $g \in \mathscr{G}^{\mathbb{C}}$ on an element $(\nabla, \phi) \in \mathscr{N}$ gives rise to the corresponding action of $[g] \in \tilde{\mathscr{G}}^{\mathbb{C}}$ on $[(\nabla, \phi)] \in \tilde{\mathscr{N}}$: we choose a representative g' of $[g]$ in $\tilde{\mathscr{G}}^{\mathbb{C}}$ for which $\|g'\phi\|_{L^2} = \|\phi\|_{L^2}$. Henceforth, we will essentially identify elements of $\tilde{\mathscr{G}}^{\mathbb{C}}$ acting on $\tilde{\mathscr{N}}$ with elements of $\mathscr{G}^{\mathbb{C}}$ acting on $\mathscr{N}$ in this way.

In the coming chapter, we will find a zero for μ on every orbit of $\tilde{\mathscr{G}}^{\mathbb{C}}$, so long as τ satisfies the constraint in Proposition 3.3; furthermore, it will be unique up to the action of $\mathscr{G}$. This will establish a correspondence between the moduli space of solutions and $\tilde{\mathscr{N}}/\tilde{\mathscr{G}}^{\mathbb{C}}$. However, observe that elements of $\tilde{\mathscr{G}}^{\mathbb{C}}$ correspond to nonzero complex-valued maps on X up to constant multiples, and elements of $\tilde{\mathscr{N}}$ correspond to elements of $\mathscr{N}$ with a certain average value of $|\phi|^2$ up to constant unitary transformations. It follows with thought that each element of $\tilde{\mathscr{N}}/\tilde{\mathscr{G}}^{\mathbb{C}}$ corresponds to a family of holomorphic maps ϕ whose zeros agree (with multiplicity), meaning we can identify the quotient with the set of divisors on X. In more concise terms, we have the following:

Proposition 3.15. *The quotient space $\tilde{\mathscr{N}}/\tilde{\mathscr{G}}^{\mathbb{C}}$ is diffeomorphic to S^dX, the dth symmetric power of X.*

Thus, once we show that every $\tilde{\mathscr{G}}^{\mathbb{C}}$-orbit has a unique zero up to the action of $\mathscr{G}$, we will have that the moduli space is precisely S^dX.

Chapter 4

Existence Proof

We have now established that the moduli space of vortices can be naturally interpreted in terms of a moment map, and we have further demonstrated that a zero can be found if we have a minimum on an orbit of $\mathscr{G}^{\mathbb{C}}$. This insight was used by Garcia-Prada in [Gar94] to characterise the structure of the moduli space of vortices. In this chapter, therefore, we reproduce Garcia-Prada's existence proof in clear and self-contained language. We emphasise the big picture of the proof throughout the chapter, and provide background on the relevant functional analysis.

Before we continue, we briefly restate the main constructions of Chapter 3, and the overarching strategy of the existence proof. We have a space $\tilde{\mathscr{N}} = \tilde{\mu}^{-1}(c)/\mathrm{U}(1)$ with a natural action of $\tilde{\mathscr{G}}^{\mathbb{C}} = \mathrm{Map}(X, \mathbb{C}^*)/\mathbb{C}^*$, and a moment map $\mu : \tilde{\mathscr{N}} \to C^\infty(X)$ given by the Yang-Mills-Higgs functional (up to a constant). The approach is to find a minimum of $\|\mu\|^2$ on a $\tilde{\mathscr{G}}^{\mathbb{C}}$-orbit of $\tilde{\mathscr{N}}$. We shall do this by constructing an infinite sequence in a $\tilde{\mathscr{G}}^{\mathbb{C}}$-orbit $\Gamma \subseteq \tilde{\mathscr{N}}$ which minimises $\|\mu\|^2$, and constructing a subsequence which converges to an element of Γ (up to gauge equivalence); this limit will therefore attain the desired minimum.

4.1 Generalised Sections and Connections

The goal of attaining a minimum for $\|\mu\|^2$ by finding a convergent sequence of sections and connections is hindered by the fact that spaces of smooth functions tend to be incomplete. We will therefore be interested in the completions of these spaces; in particular, we take the completion with respect to the Sobolev norm to get spaces of Sobolev sections and connections. Before discussing the spaces of sections and connections, we give a brief overview of Sobolev spaces and elliptic operator theory.

4.1.1 Sobolev Spaces

Let (X, g) be a compact n-dimensional Riemannian manifold, and let $\pi : E \to X$ be a vector bundle. Though the space $\Gamma(E)$ of C^∞ sections is well-behaved pointwise, its natural topologies are quite poorly behaved. More concretely, $\Gamma(E)$ is incomplete under every possible p-norm. The natural solution is to take the metric completion under a chosen p-norm, which results in the spaces $L^p(E)$; however, these spaces do not carry any derivative information, so they are not suitable for analysis on PDEs. The corresponding notion to solve this problem is that of a Sobolev space.

In order to define the space $L^p_k(E)$, we need some extra structure on our vector bundle. First, we put a bundle metric h on E; we will take it to be a Hermitian metric in our context. Then we take a Hermitian connection ∇ on E, as well as the Levi-Civita connection ∇ on X. This induces connections on every tensor power of E and X, and in particular, the tensor power $(T^*X)^{\otimes k} \otimes E$ for every $k \in \mathbb{N}$. Now, by the definition of a connection, we have for any $\sigma \in \Gamma(E)$ a section $\nabla\sigma \in \Gamma(T^*X \otimes E)$; but by the above extension, we also have sections $\nabla^j\sigma \in \Gamma((T^*X)^{\otimes j} \otimes E)$ for every $j \in \mathbb{N}$. Additionally, the metric on X and the Hermitian metric on E extends to a metric on all tensor powers; in particular, we can define $|\nabla^j\sigma|$ for all j (and it will be a smooth function on X).

We now define a series of norms on $\Gamma(E)$:

Definition 4.1. Let $(E, h, \nabla) \to (X, g, \nabla)$ be as above. Given $p \in [1, \infty]$ and $k \in \mathbb{N}$, the (k,p)*-Sobolev norm* on $\sigma \in \Gamma(E)$ is defined as follows:

$$\|\sigma\|_{L^p_k} = \sum_{j=0}^{k} \|\nabla^j\sigma\|_{L^p} = \sum_{j=0}^{k} \left(\int_X |\nabla^j\sigma|^p \operatorname{vol}_g \right)^{1/p}. \tag{4.1}$$

Note that we recover the usual L^p-norm if we take $k = 0$.

Unsurprisingly, the space $\Gamma(E)$ is incomplete under this norm; we therefore complete it in just the same way as the L^p spaces:

Definition 4.2. The (k,p)*-Sobolev space of sections of* E*,* or the L^p_k space of sections, is the completion of $\Gamma(E)$ with the (k,p)-Sobolev norm. It is denoted by $\Gamma(E)_{L^p_k}$. Moreover, if E is the trivial $\mathbb{R}$-bundle or $\mathbb{C}$-bundle, the space is simply denoted by L^p_k.

Observe that a section of E is L^p_k if and only if its components are real-valued L^p_k functions in a local trivialisation.

Though we need to make several choices throughout the construction, the choice of connection leads to no topological difference.

Proposition 4.3. *Fix* $k \in \mathbb{N}$ *and* $p \in [1, \infty]$*. Any two* (k,p)*-Sobolev norms corresponding to different metric connections on* E *are equivalent.*

Proof. For a proof in the $k = 1$ case, see [Hay24]. The general case is similar. □

The notion of a Sobolev space readily applies to three new kinds of spaces:

- Since differential ℓ-forms are sections of the bundle $\Lambda^\ell(T^*X)$; we therefore have the notion of differential ℓ-forms of class L^p_k, denoted by $\Omega^\ell(X)_{L^p_k}$.
- Recalling that the space $\mathscr{A}$ of smooth connections forms an affine space over $\Omega^{0,1}(X)$, we can generalise the notion of a connection by fixing some $\nabla \in \mathscr{A}$ and adding elements of $\Omega^{0,1}(X)_{L^p_k}$ instead; this gives the notion of an L^p_k *connection.* If $p = 2$, we denote the space of such connections by $\mathscr{A}^k$.

- Conceptualising the gauge groups $\mathscr{G} = \mathrm{Map}(X, \mathrm{U}(1))$ and $\mathscr{G}^{\mathbb{C}} = \mathrm{Map}(X, \mathbb{C}^*)$ as spaces of sections of trivial U(1)-bundles or $\mathbb{C}^*$-bundles, we therefore have the notion of an L^p_k *gauge transformation.* If $p = 2$, we denote the corresponding groups by $(\mathscr{G})^k$ and $(\mathscr{G}^{\mathbb{C}})^k$ respectively.

We will need three important results from the theory of Sobolev spaces.

Theorem 4.4 (Sobolev Embedding)**.** *Let $p, q \in [1, \infty)$ and let $k, \ell \in \mathbb{N}$. Consider the Sobolev spaces $L^p_k(E)$ and $L^q_\ell(E)$ on an n-dimensional manifold M with a vector bundle $E \to M$. If $k - n/p \geq \ell - n/q$ and $k \geq \ell$, then $L^p_k(E) \subset L^q_\ell(E)$. Furthermore, if the inequalities are strict, then the inclusions are compact operators. Additionally, if $k - n/p > \ell$, then there is a compact inclusion $L^p_k(E) \subset C^\ell$.*

Proof. See [Heb96]. □

Theorem 4.5 (Weak compactness)**.** *Any bounded subset of the space $\Gamma(E)_{L^p_k}$ is weakly compact for $p \in (1, \infty)$ and $k \geq 0$. That is, every bounded sequence of L^p_k sections of E has a subsequence which converges weakly, in that their images under any bounded linear functional converge in $\mathbb{R}$.*

Proof. The space $\Gamma(E)_{L^p_k}$ can be embedded in a finite product of reflexive spaces, namely the L^p spaces over the possible multi-indices up to order k. A finite product of reflexive spaces is reflexive, and L^p_k is therefore a closed subspace of a reflexive space, making it reflexive. The result then follows from a corollary of Alaoglu's theorem, namely that the unit ball in every reflexive space is weakly compact (see [Con90]). □

Theorem 4.6 (Sobolev Multiplication)**.** *Let n be the dimension of the manifold, and let k_1, k_2, k be natural numbers and $p_1, p_2, p \in [1, \infty)$ be real numbers satisfying the following constraints:*

- $k \leq k_1, k_2$;
- $k_i - k \geq n\left(\frac{1}{p_i} - \frac{1}{p}\right)$ *for each i;*
- $k_1 + k_2 - k > n\left(\frac{1}{p_1} + \frac{1}{p_2} - \frac{1}{p}\right) \geq 0.$

Then the multiplication map $L^{p_1}_{k_1} \times L^{p_2}_{k_2} \to L^p_k$ is continuous and well-defined. In particular, the multiplication map $L^2_k \times L^2_k \to L^{2-\varepsilon}_k$ is continuous and well-defined for every $k \in \mathbb{N}$ and every $\varepsilon > 0$.

Proof. The case for which the underlying manifold is $\mathbb{R}^n$ and the underlying vector bundle is trivial is done in [FWX25]. The general case follows from the fact that the statement is purely local, and every vector bundle over a manifold can be locally reduced to this case. □

4.1.2 Differential Operators and Ellipticity

The vortex equations utilise several operators which belong to a certain class: namely, they are elliptic linear differential operators. We will define this class of operators, demonstrate that $\nabla^{0,1}$ and $\overline{\partial}$ are of this class, and state some important results from elliptic operator theory.

A linear differential operator between vector bundles is, roughly speaking, a sum of partial derivatives in any coordinate system. More precisely:

Definition 4.7. Let $E, F \to X$ be vector bundles of rank a and b over a smooth manifold X, and let $L : \Gamma(E) \to \Gamma(F)$ be a linear map. Let $U \subseteq X$ be open and homeomorphic to $\mathbb{R}^n$ through a chart $x : U \to \mathbb{R}^n$, and further suppose E and F are both locally trivial over U with local diffeomorphisms $\phi_E : \pi_E^{-1}(U) \to U \times \mathbb{R}^a$ and $\phi_F : \pi_F^{-1}(U) \to U \times \mathbb{R}^b$. Then L is a *linear differential operator of order* k if, for every such U, there is some operator $\tilde{L} : \mathrm{Map}(U, \mathbb{R}^a) \to \mathrm{Map}(U, \mathbb{R}^b)$ of the form

$$\tilde{L}(f) = \sum_{|\alpha| \leq k} A_\alpha(x) \frac{\partial^{|\alpha|} f}{\partial x^\alpha} \tag{4.2}$$

for every $f : U \to \mathbb{R}^a$ (where $A_\alpha : U \to \mathbb{R}^{a \times b}$ for each multi-index α), such that $\phi_F \circ L = \tilde{L} \circ \phi_E$.

It is clear that this definition is independent of the coordinate chart $x : U \to \mathbb{R}^n$, and also independent of the local trivialisations.

Elliptic operators are those differential operators which are locally "well-behaved." The notion of local behaviour is captured by the *principal symbol:*

Definition 4.8. Let the k-th order differential operator $L : \Gamma(E) \to \Gamma(F)$, the open set U, and the local representation $\tilde{L} : \mathrm{Map}(U, \mathbb{R}^a) \to \mathrm{Map}(U, \mathbb{R}^b)$ be as above. The *principal symbol of* L *at* $x \in X$ is the map $\sigma(L) : \mathbb{R}^n \to \mathbb{R}^{a \times b}$ defined as follows:

$$[\sigma(L)](\omega) = \sum_{|\alpha| = k} A_\alpha(x) \omega^\alpha, \tag{4.3}$$

where we define $\omega^\alpha = \omega_1^{\alpha_1} \cdots \omega_n^{\alpha_n}$ for the multi-index $\alpha = (\alpha_1, \ldots, \alpha_n)$.

Once again, this can be shown to be independent of coordinates. We then have the notion of an elliptic operator:

Definition 4.9. A linear differential operator $L : \Gamma(E) \to \Gamma(F)$ is *elliptic* if, for any $\omega \in \mathbb{R}^n \backslash \{0\}$, its principal symbol $[\sigma(L)](\omega)$ is invertible.

Note that all of our definitions only apply to smooth sections. Nevertheless, they can be extended to Sobolev spaces (refer to [Hay24]):

Proposition 4.10. *Any k-th order differential operator $L : \Gamma(E) \to \Gamma(F)$ admits a unique bounded extension to a map $L : \Gamma(E)_{L^2_\ell} \to \Gamma(E)_{L^2_{\ell-k}}$.*

We will essentially regard these two maps as the same, without further comment.

The standard example of an elliptic linear operator is the Laplacian $\Delta : C^\infty(\mathbb{R}^n) \to C^\infty(\mathbb{R}^n)$, defined by taking $\Delta(f) = \frac{\partial^2 f}{\partial x_1^2} + \cdots + \frac{\partial^2 f}{\partial x_n^2}$. By definition, it is a second-order differential operator on the trivial $\mathbb{R}^n$-bundle on $\mathbb{R}^n$. The principal symbol of Δ is simply $[\sigma(\Delta)](\omega) = (\omega_1^2 + \cdots + \omega_n^2)\mathrm{id}_{\mathbb{R}^n}$, and this is clearly invertible whenever each $\omega_i \neq 0$. Two more relevant examples are as follows:

Proposition 4.11. *If X is a Riemann surface, the Dolbeault operator $\overline{\partial} : \Omega^{0,0}(X) \to \Omega^{0,1}(X)$ is an elliptic differential operator.*

Proof. In a holomorphic coordinate system, the Dolbeault operator is represented as $(\frac{\partial}{\partial x} + i\frac{\partial}{\partial y})\mathrm{d}\overline{z}$; this is clearly a first-order differential operator. Its principal symbol is given by

$$(\omega_x, \omega_y) \mapsto (\omega_x + i\omega_y)\mathrm{d}\overline{z}, \tag{4.4}$$

which is clearly invertible whenever ω_x or ω_y are nonzero, making $\overline{\partial}$ elliptic. □

Corollary 4.12. *If $L \to X$ is a line bundle over a Riemann surface, the holomorphic structure $\overline{\partial}_L : \Omega^{0,0}(L) \to \Omega^{0,1}(L)$ is an elliptic differential operator.*

Proof. In a holomorphic coordinate system on X, the holomorphic structure is represented as $\overline{\partial} + \tau$ where $\tau \in \Omega^{0,1}(X)$, which is clearly a first-order differential operator. Moreover, its highest-order part coincides with $\overline{\partial}$, which we just proved is elliptic. □

Our motivation for introducing this theory is the powerful results that can be obtained for elliptic operators. We will be using two in particular:

Theorem 4.13 (Elliptic regularity). *Let $D : \Gamma(E) \to \Gamma(F)$ be any elliptic ℓ-th order differential operator. If ϕ is a section of E for which $D\phi$ is an L^2_k section of F, then ϕ is $L^2_{k+\ell}$. In particular, $\ker(D)$ consists of smooth sections of E.*

Theorem 4.14 (Elliptic estimate). *Let $\sigma \in \Gamma(E)$ be an L^2_2 section of a vector bundle E over a compact manifold X, and let $D : \Gamma(E) \to \Gamma(F)$ be a k-th order elliptic operator. Then there is some $C \in (0, \infty)$ which does not depend on σ, such that the following inequality holds:*

$$\|\sigma\|_{L^2_2} \leq C(\|D\sigma\|_{L^2_1} + \|\sigma\|_{L^2}). \tag{4.5}$$

Furthermore, if σ is L^2-orthogonal to $\ker D$, we may omit the $\|\sigma\|_{L^2}$ term.

For proofs of each of these theorems, see [Hay24].

4.1.3 Application to Gauge Theory

When we come to prove the existence of vortices, most of the work will be done within the space of L^2_1 connections (which we denote by $\mathscr{A}^1$), the space of L^2_1 sections (which we denote by $\Gamma(L)_{L^2_1}$), and the spaces of L^2_2 gauge transformations (which we denote by $\mathscr{G}^2$ and $(\mathscr{G}^{\mathbb{C}})^2$). We first check that the action of L^2_2 on L^2_1 is well-defined.

Lemma 4.15. *L^p_j is a topological L^p_k-module whenever $j \leq k$ and $k > n/p$.*

Proof. The statement to be proved is that the multiplication map on $L^p_j \times L^p_k$ maps continuously into L^p_j. This clearly follows if it is true on every open subset of X, so it can be reduced to the local statement (i.e. multiplication of two Sobolev functions over $\mathbb{R}^n$). But this is a special case of Theorem 6.1 in [BH21]. □

Setting $j = 1$, $k = 2$, $p = 2$ and $n = 2$ gives that L^2_1 is a topological L^2_2-module, and hence that $(\mathscr{G}^{\mathbb{C}})^2$ acts continuously on $\mathscr{A}^1$. Additionally, we need to check that the Yang-Mills-Higgs functional still makes sense on these extended spaces:

Proposition 4.16. *The Yang-Mills-Higgs functional can be continuously extended to L^2_1 sections and connections.*

Proof. It is first shown that the output of a map taking an L^2_1 connection to its curvature is L^2; our method here is based on the method in [Uhl82]. Let $\nabla' \in \mathscr{A}^1$ be an arbitrary element; then $\nabla' = \nabla + \alpha$, where ∇ is smooth and $\alpha \in \Omega^1(X)_{L^2_1}$. Then $F^{\nabla'} = F^{\nabla} + \mathrm{d}\alpha$; the first term is smooth, and the second is L^2 (since α has sufficiently regular zeroth and first derivatives, so $\mathrm{d}\alpha$ will only have sufficiently regular zeroth derivative). This makes $F^{\nabla'}$ an L^2 object.

Next, the Sobolev inclusion $L^2_1(E) \subseteq L^4(E)$ shows that any $\phi \in L^2_1(E)$ is also in $L^4(E)$, and this implies that $|\phi|^2_h$ is in L^2. But since the Yang-Mills-Higgs functional is essentially the L^2-norm of $F^{\nabla'}$ and $|\phi|^2_h$, this means that it is well-defined in the extended sense. □

In the existence proof, we shall eventually produce a zero of μ on an orbit of $(\mathscr{G}^{\mathbb{C}})^2$ consisting of an L^2_1 section and connection, thus producing an L^2_1 solution to the vortex equations. On the other hand, we are interested in producing smooth solutions up to unitary gauge equivalence. Contrary to what one might expect, there is no way to transform an arbitrary L^2_1 pair into a smooth pair using an L^2_2 gauge transformation; if ∇ has nonsmooth curvature, any gauge transformation will preserve this nonsmooth curvature and hence cannot map ∇ to a smooth connection. However, if an L^2_1 pair is a *solution to the vortex equations,* such a gauge transformation does exist. Instrumental to finding such a gauge transformation is the following result, which establishes the existence of the Coulomb gauge in a special case (this was initially proved in [Uhl82]):

Theorem 4.17 (Uhlenbeck). *Let B^2 be the unit disk in $\mathbb{C}$, let $L \to B^2$ be the trivial line bundle, and let ∇ be an L^2_1 connection on L. Then there is some $\kappa > 0$ and some $c < \infty$ such that, whenever $\|F^{\nabla}\|_{L^2} < \kappa$, the connection ∇ is L^2_2-gauge equivalent to an L^2_1 connection $\mathrm{d} + A$ for which the following conditions hold:*

$$\begin{aligned} \mathrm{d}^* A &= 0; \\ \|A\|_{L^2_1} &\leq c\|F^{\nabla}\|_{L^2}. \end{aligned} \tag{4.6}$$

Once we show that L^2_1 solutions can be gauge-transformed into smooth solutions, we can proceed with L^2_1 sections and connections without further issues.

With this preparation, we can state the two major convergence theorems that we will use to construct our solution. The first is the weak compactness of L^p_k, which we have already stated and proved. The second is Uhlenbeck's weak compactness theorem:

Theorem 4.18 (Uhlenbeck Weak Compactness). *Let X be a compact Riemann surface with a vector bundle $E \to X$, and let $(A^k)_{k\in\mathbb{N}} \in \mathscr{A}^1$ be a sequence of connections on E for which $\|F^{A_k}\|$ is uniformly bounded. Then there is a subsequence $(A^{k_i})_{i\in\mathbb{N}}$ and a sequence of unitary gauge transformations $g_i \in \mathscr{G}^2$ such that $g_i(A^{k_i})$ converges weakly in $\mathscr{A}^1$, where weak convergence refers to convergence in the weak topology.*

More concisely, this theorem states that any closed and bounded subset of $\mathscr{A}^1$ is compact in the weak topology, *up to gauge equivalence.* For a proof, see [Weh04].

There is one more result that we will need, and it relates to the Dolbeault cohomology of the Riemann surface. To illustrate the issue, we consider the following question: if we enlarge the space of differential $(0,1)$-forms to the Sobolev space $\Omega^{0,1}(X)_{L_1^2}$, should we compute the Dolbeault cohomology in terms of these k-forms or the smooth k-forms? It turns out that the distinction is immaterial for a Riemann surface.

Proposition 4.19. *On a Riemann surface X, the Dolbeault cohomology spaces $H^{0,1}(X)_{L_1^2}$ and $H^{0,1}(X)_{C^\infty}$ are isomorphic.*

Proof. Define a map $\mathscr{F} : \Omega^{0,1}(X)_{C^\infty} \to H^{0,1}(X)_{L_1^2}$ by taking $\mathscr{F}(\alpha) = [\alpha]_{L_1^2}$, i.e. the map $\mathscr{F}$ takes each smooth $(0,1)$-form to its L_1^2 cohomology class. Observe that $\ker \mathscr{F}$ consists of smooth $(0,1)$-forms of the form $\overline{\partial} f$, where $f : X \to \mathbb{C}$ is of class L_2^2. However, by elliptic regularity, the smoothness of $\overline{\partial} f$ implies that f itself is smooth, meaning $\ker \mathscr{F}$ consists entirely of C^∞-exact $(0,1)$-forms. By the first isomorphism theorem, $\mathscr{F}$ induces the desired isomorphism. □

4.2 The Proof

We are now ready to begin the proof. Recall that, on each $(\tilde{\mathscr{G}^{\mathbb{C}}})^2$-orbit of $\tilde{\mathscr{N}}$, we wish to attain a minimum of the squared norm of the moment map

$$\mu(\nabla, \phi) = \Lambda F^\nabla - \frac{i}{2}(|\phi|_h^2 - \tau), \tag{3.38}$$

which is equal to $\mathrm{YMH}(\nabla, \phi) = \|F^\nabla\|_{L^2}^2 + \||\phi|_h^2 - \tau\|^2$ up to a constant difference of $2\pi\tau d$.

We begin with the following lemma:

Lemma 4.20. *On any $(\tilde{\mathscr{G}^{\mathbb{C}}})^2$ orbit in $\tilde{\mathscr{N}}$, there is an L_1^2-weakly convergent sequence of sections and connections $(\nabla_n, \phi_n) \rightharpoonup (\nabla, \phi)$ for which $\|\mu(\nabla, \phi)\|^2$ is the minimum possible value of $\|\mu\|^2$ on the orbit.*

Proof. We begin by choosing a representative (∇_0, ϕ_0) of an orbit of $(\tilde{\mathscr{G}^{\mathbb{C}}})^2$. Since $\|\mu\|^2$ is a nonnegative real-valued function on this orbit, there is some L_1^2 sequence (∇_n, ϕ_n) such that $\|\mu(\nabla_n, \phi_n)\|^2$ converges to the minimum value of $\|\mu\|^2$ on the orbit. It follows that $\|\mu(\nabla_n, \phi_n)\|^2$ is uniformly bounded, and since $\|F^{\nabla_n}\|^2 \le \mathrm{YMH}(\nabla_n, \phi_n)$, we get a uniform bound on the curvature of ∇_n. By Uhlenbeck's weak compactness theorem, there is a subsequence of ∇_n and a sequence of L_2^2 unitary gauge transformations which transform the subsequence to a weakly convergent subsequence in $\mathscr{A}^1$. Henceforth, we relabel the sequence (∇_n, ϕ_n) so that ∇_n converges weakly to $\nabla \in \mathscr{A}^1$.

We can similarly find a uniform upper bound on $\|\phi_n\|_{L^2_1}$; to do this, we show that a uniform upper bound on $\|\phi_n\|_{L^2}$ is sufficient. Note that $\nabla_n^{0,1}$ is elliptic and ϕ_n satisfies the equation $\nabla_n^{0,1}\phi_n = 0$ (since it is in $\mathscr{N}$). By the elliptic estimate (Theorem 4.14), we find that there are constants $C_n > 0$ for each ∇_n (independent of $\phi \in \Gamma(L)$) for which

$$\|\phi\|_{L^2_1} \leq C_n(\|\nabla_n^{0,1}\phi\|_{L^2} + \|\phi\|_{L^2}) = C_n\|\phi\|_{L^2}, \tag{4.7}$$

where the final equality holds if ϕ is holomorphic (which is the case for each ϕ_n). Now, choose some $\phi \in \Gamma(L)_{L^2_1}$. Since $\nabla_n^{0,1}$ converges weakly to $\nabla^{0,1}$, we have that $\langle\nabla_n^{0,1}\phi, \alpha\rangle$ converges to $\langle\nabla^{0,1}\phi, \alpha\rangle$ for any $\alpha \in \Omega^{0,1}(X)_{L^2}$; in particular, if we choose an orthonormal basis $\{e_k\}$ for $\Omega^{0,1}(X)_{L^2}$, it follows that

$$\lim_{n\to\infty}\|\nabla_n^{0,1}\phi\|^2_{L^2} = \lim_{n\to\infty}\sum_k|\langle\nabla_n^{0,1}\phi, e_k\rangle|^2 = \sum_k|\langle\nabla^{0,1}\phi, e_k\rangle|^2 = \|\nabla^{0,1}\phi\|^2_{L^2}. \tag{4.8}$$

On the other hand, we can take the constants C_n to be the following:

$$C_n = \sup_{\phi\neq 0}\frac{\|\phi\|_{L^2_1}}{\|\nabla_n^{0,1}\phi\|_{L^2} + \|\phi\|_{L^2}}. \tag{4.9}$$

However, by the computation above, the limit of the denominator as $n \to \infty$ is $\|\nabla^{0,1}\phi\|_{L^2} + \|\phi\|_{L^2}$, and it follows from the ellipticity of $\nabla^{0,1}$ that C_n also converges as $n \to \infty$. But then the constants C_n must be uniformly bounded over n, meaning we can replace all the C_n by some universal constant $C > 0$ which is independent of ∇_n. It follows that $\|\phi_n\|_{L^2_1} \leq C\|\phi_n\|_{L^2}$ for all n, so we need only find uniform bounds on $\|\phi_n\|_{L^2}$.

To do this, observe that Hölder's inequality gives us the estimate

$$\|1\cdot\phi_n\|_{L^2} \leq \|1\|_{L^4}\|\phi_n\|_{L^4} = \mathrm{Vol}(X)^{1/4}\|\phi_n\|_{L^4}. \tag{4.10}$$

Moreover, since $\mathrm{YMH}(\nabla_n, \phi_n)$ is bounded, so too is $\||\phi_n|^2_h - \tau\|^2_{L^2}$ (see Equation 3.1). We equivalently write this expression as follows:

$$\||\phi_n|^2_h - \tau\|^2_{L^2} = \|\phi_n\|^4_{L^4} + \tau^2\mathrm{Vol}(X) - 2\tau\|\phi_n\|^2_{L^2}. \tag{4.11}$$

We then substitute $\frac{1}{\mathrm{Vol}(X)^{1/4}}\|\phi_n\|_{L^2}$ into this expression and multiply through by $\mathrm{Vol}(X)$; the following expression is therefore also bounded:

$$\begin{aligned}&\|\phi_n\|^4_{L^2} - 2\tau\mathrm{Vol}(X)\|\phi\|^2_{L^2} + (\tau\mathrm{Vol}(X))^2\\ &= (\|\phi_n\|^2_{L^2} - 2\tau\mathrm{Vol}(X))^2.\end{aligned} \tag{4.12}$$

But this clearly implies an upper bound on $\|\phi_n\|_{L^2}$, and hence on $\|\phi_n\|_{L^2_1}$. By the weak compactness of L^2_1, we conclude that ϕ_n has a weakly convergent subsequence to some L^2_1 section ϕ. If we take the corresponding subsequence of ∇_n, it will still converge to ∇. □

Thus, we obtain an L^2_1 connection and section which attains the minimum possible value of $\|\mu\|^2$. However, it is not yet clear that they belong to the *same orbit* as (∇_0, ϕ_0); without this detail, we cannot conclude that μ has a zero at all. The rest of the proof is devoted to this fact.

Lemma 4.21. *Let $(\nabla_n, \phi_n) \rightharpoonup (\nabla, \phi)$ be a convergent sequence of L^2_1 sections and connections for which each element of the sequence is in the same $(\mathscr{G}^{\mathbb{C}})^2$ orbit. Then (∇, ϕ) is also in the same $(\mathscr{G}^{\mathbb{C}})^2$ orbit.*

Proof. To be in the same orbit means that there is a sequence $g_n \in (\mathscr{G}^{\mathbb{C}})^2$ such that $(\nabla_n, \phi_n) = g_n \cdot (\nabla_0, \phi_0)$, and therefore that the following hold for each $n \in \mathbb{N}$:

$$\begin{aligned} \nabla_n^{0,1} - \nabla_0^{0,1} &= g_n^{-1} \overline{\partial} g_n, \\ \|g_n \phi_0\|_{L^2} &= \|\phi_0\|_{L^2}. \end{aligned} \tag{4.13}$$

We wish to show that g_n can be chosen to converge to a holomorphic isomorphism g between $(L, \nabla_0^{0,1})$ and $(L, \nabla^{0,1})$.

Define $\alpha_n = g_n^{-1} \overline{\partial} g_n$. By definition, this is a $(0,1)$-form on X of class L^2_1, and it defines a cohomology class in $H^{0,1}(X)_{L^2_1}$ since $\overline{\partial} = 0$ identically on $(0,1)$-forms. Moreover, we showed above (Proposition 4.19) that the cohomology of the space of L^2_1 differential forms is the same as the ordinary C^∞ differential forms. This is isomorphic to $H^1(X, \mathscr{O})$, the first sheaf cohomology of the sheaf of holomorphic functions, by the Dolbeault isomorphism. Thus, we can identify the cohomology class of α_n with a class in $H^1(X, \mathscr{O})$.

We will show that α_n can be taken to be exact, meaning it is of the form $\overline{\partial} f_n$ for some $f_n \in L^2_2$, and we will do this by analysing the exponential exact sequence of sheaves. Note that it induces the following subsequence of the long exact cohomology sequence:

$$H^1(X, \underline{\mathbb{Z}}) \to H^1(X, \mathscr{O}) \xrightarrow{\exp(2\pi i \cdot)} H^1(X, \mathscr{O}^*), \tag{4.14}$$

and we can consider each α_n to be an element of $H^1(X, \mathscr{O})$ via the Dolbeault isomorphism. In fact, by Equation 2.5, we can write an explicit representation of α_n in $H^1(X, \mathscr{O})$ in a good open cover $\{U_i\}_{i \in I}$ by taking $\overline{\partial}$-primitives of α_n over each U_i and subtracting them on overlaps. On the other hand, $\alpha_n = \overline{\partial}(\ln(g_n))$ on small enough open sets, meaning the differences between the primitives on overlaps can be chosen to be constant integer multiples of $2\pi i$. Each one evaluates to 1 under the exponential map, meaning the class defined by α_n is trivial in $H^1(X, \mathscr{O}^*)$.

Now, by exactness, each α_n is an image of a class in $H^1(X, \underline{\mathbb{Z}})$ which is isomorphic to the first singular cohomology of X. But X is compact, so this is finitely generated and hence discrete. If we identify each α_n with its corresponding harmonic representative, we see that the Cauchy nature of α_n implies that its sequence of classes is also Cauchy. (We choose the harmonic representative because it has the smallest norm in the cohomology class of α_n; for more details on harmonic forms, see [Voi02].) But since $H^1(X, \underline{\mathbb{Z}})$ is discrete, it follows that the classes (and hence α_n) must be eventually constant. We can even perform a complex gauge transformation to make $[\alpha_n]$ eventually equal to 0: if one of the representatives of the eventual class is of the form $g^{-1} \overline{\partial} g$ for $g \in \mathscr{G}^{\mathbb{C}}$, we can simply redefine α_n as follows:

$$\alpha_n \mapsto \alpha_n^{g^{-1}} = g g_n^{-1} \overline{\partial}(g^{-1} g_n). \tag{4.15}$$

This necessarily converges to the zero class. Replacing each (∇_n, ϕ_n) with $g^{-1} \cdot (\nabla_n, \phi_n)$, it follows that $[\alpha_n]$ is eventually 0. Thus, by applying g and shifting n, we can assume that *every* α_n is exact, and we can therefore take $\alpha_n = \overline{\partial} f_n$ for some sequence $f_n \in L^2_2$. We will additionally assume that $\int_X f_n \operatorname{vol}_g = 0$ for all n, since α_n only determines f_n up to a constant.

Now that we have shown that α_n can be taken to be exact for every n, we will show that g_n can be taken to converge to a holomorphic isomorphism. By the weak convergence of the α_n, the sequence $\|\alpha_n\|_{L^2_1} \in \mathbb{R}$ is uniformly bounded. We can use this fact and the elliptic estimate (Theorem 4.14) to uniformly bound $\|f_n\|_{L^2_2}$. Observe that any complex function h for which $\overline{\partial} h = 0$ must be holomorphic on X and hence constant, but this means that

$$\langle f_n, h \rangle = \int_X f_n \overline{h} \operatorname{vol}_g = \overline{h} \int_X f_n \operatorname{vol}_g = 0. \tag{4.16}$$

It follows that each f_n is orthogonal to the kernel of $\overline{\partial}$. Since $\overline{\partial}$ is elliptic, we can use the elliptic estimate to uniformly bound the norms of f_n:

$$\|f_n\|_{L^2_2} \leq C\|\alpha_n\|_{L^2_1}. \tag{4.17}$$

Thus, $\|f_n\|_{L^2_2}$ must also be uniformly bounded. By the weak compactness of L^2_2, the sequence f_n has a weakly convergent subsequence; we relabel f_n to be weakly convergent in L^2_2. But by the Sobolev embedding theorems, there is a compact inclusion $i : L^2_2 \hookrightarrow C^0$, so the f_n are also weakly convergent in $(C^0, \|\cdot\|_\infty)$ and hence must be uniformly bounded in the sup-norm. Thus, we can find some $M > 0$ such that $f_n(x) \leq M$ for all $x \in X$ and all $n \in \mathbb{N}$.

We now have that $g_n^{-1}\overline{\partial} g_n = \overline{\partial} f_n$ where f_n are uniformly bounded functions. The log function exists locally in the complex plane, so this implies that $\overline{\partial}\log(g_n) = \overline{\partial} f_n$ and hence $g_n = K_n e^{f_n}$ for some nonzero constants $K_n \in \mathbb{C}^*$. It follows that

$$|K_n|e^{-M} \leq |g_n| \leq |K_n|e^M. \tag{4.18}$$

However, recall that the $\tilde{\mathscr{G}}^{\mathbb{C}}$ acts on ϕ by preserving its L^2 norm; in particular, we must have that $\|g_n\phi_0\|_{L^2} = \|\phi_0\|_{L^2}$. Consequently, we have the following inequalities:

$$|K_n|e^{-M}\|\phi_0\|_{L^2} \leq \|g_n\phi_0\|_{L^2} = \|\phi_0\|_{L^2} \leq |K_n|e^M\|\phi_0\|_{L^2}. \tag{4.19}$$

Since $\phi_0 \neq 0$, we see that $|K_n|e^{-M} \leq 1 \leq |K_n|e^M$, which implies that $|K_n| \in [e^{-M}, e^M]$. Therefore, $|g_n|$ is bounded for $n \in \mathbb{N}$, meaning $\|g_n\|_\infty$ is uniformly bounded. From the Sobolev inclusion $L^2_2 \hookrightarrow C^0$, it follows that $\|g_n\|_{L^2_2}$ is uniformly bounded, so by the weak compactness of L^2_2 it has an L^2_2-weakly convergent subsequence. But then its weak limit is a complex gauge transformation relating (∇_0, ϕ_0) to (∇, ϕ), demonstrating that (∇, ϕ) is in the same orbit. □

We now have an L^2_1-solution on each orbit of $\tilde{\mathscr{G}}^{\mathbb{C}}$, and we need to be able to transform these into smooth solutions using unitary gauge transformations.

Lemma 4.22. *Let (∇, ϕ) be an L^2_1 solution to the vortex equations. Then there exists an L^2_2 gauge transformation g for which (∇^g, ϕ^g) is a smooth solution.*

Proof. The proof for this lemma is an adaptation of the proof in [FWX25], which constructs a similar gauge transformation for L^2_1 solutions to the Hitchin equations.

Since (∇, ϕ) is a solution to the vortex equations, we know that $\|F^\nabla\|_{L^2} = \frac{1}{2}\||\phi|^2 - \tau\|_{L^2} < \infty$, which means that we can make $F^\nabla|_U$ have arbitrarily small norm by shrinking an open set $U \subseteq X$. As such, we can choose a finite open cover $\{U_i\}_{i\leq n}$ of X for which each U_i is diffeomorphic to the unit (open) disk in $\mathbb{C}$, the line bundle $L|_{U_i}$ is trivial, and for which $\|F^\nabla|_{U_i}\|_{L^2}$ is arbitrarily small. On each open set U_i, we can write $\nabla = \mathrm{d}+A$ where A is an L^2_1 complex-valued 1-form, and we can interpret ϕ as an L^2_1 complex-valued function on U_i as long as we interpret the hermitian metric h as a smooth complex-valued function on U_i as well. The local version of the vortex equations then becomes the following:

$$\begin{aligned} \overline{\partial}\phi + A^{0,1}\phi &= 0; \\ \mathrm{d}A &= \frac{i}{2}(h|\phi|^2 - \tau)\mathrm{vol}_g. \end{aligned} \tag{4.20}$$

Note that $\overline{\partial}$ is an elliptic operator and $A^{0,1}\phi \in L^{2-\varepsilon}_1$ for every $\varepsilon > 0$ by Theorem 4.6, so by elliptic regularity we must have that $\phi \in L^{2-\varepsilon}_2$. Moreover, taking d^* of both sides of the second equation, we have that $\mathrm{d}^*\mathrm{d}A = \frac{i}{2}\mathrm{d}^*(h|\phi|^2\mathrm{vol}_g)$. But $\mathrm{d}^*\mathrm{d}A = \Delta A$ where $\Delta = \mathrm{d}^*\mathrm{d} + \mathrm{d}\mathrm{d}^*$ is the Laplacian, as long as A is in the Coulomb gauge (since $\mathrm{d}^*A = 0$ in the Coulomb gauge), and Δ is known to be elliptic. By Theorem 4.17, the Coulomb gauge always exists. Moreover, we have that $\mathrm{d}(h|\phi|^2)$ is $L^{2-\varepsilon}$, and elliptic regularity ensures that A is $L^{2-\varepsilon}_2$. By applying this method inductively, we see that A and ϕ are $L^{2-\varepsilon}_k$ for every k and every $\varepsilon > 0$, which means they must be smooth (the intersection of L^p_k over k is C^∞).

We now have that (∇, ϕ) are locally gauge equivalent to a smooth pair, and we need to glue these local gauge transformations into a global gauge transformation. To do this, we first ensure that they can be written in terms of exp by finding arbitrarily small corresponding gauge transformations. Let $\{g_i : U_i \to \mathrm{U}(1)\}$ be the local L^2_2-gauge transformations for which $(\nabla|_{U_i})^{g_i}$ and $(\phi|_{U_i})^{g_i}$ are smooth. Since C^∞ is dense in L^2_2, there exist smooth $h_i : U_i \to \mathrm{U}(1)$ such that $\|h_i - g_i^{-1}\|_{L^2_2}$ is arbitrarily small. It follows that $\|1 - h_i g_i\|_{L^2_2}$ can also be made arbitrarily small, so by replacing each g_i with $h_i g_i$, we can choose the g_i to be arbitrarily close to the constant function 1 on U_i. It follows also that the "transition" gauge transformation $g_{ij} := g_i g_j^{-1}$ can be made arbitrarily close to the identity. Observe that $g_{ij}[g_j(\nabla, \phi)|_{U_i \cap U_j}]$ is smooth, meaning g_{ij} takes a smooth pair to a smooth pair. It is a theorem of Atiyah that such a gauge transformation must itself be smooth, so each g_{ij} is smooth. Now, since $\exp : i\mathbb{R} \to \mathrm{U}(1)$ is a diffeomorphism close to the identity, we can write each g_{ij} as $\exp(if_{ij})$ where $f_{ij} : \mathbb{R} \to \mathbb{R}$ is smooth. In other words, $\ln(g_{ij})$ is well-defined.

We now begin the process of constructing a global smoothing gauge transformation. Let $d : X \times X \to [0, \infty)$ be the metric on X, and define

$$U_i^\delta = \{p \in U_i \ : \ d(p, \partial U_i) > \delta\} \tag{4.21}$$

for each $\delta > 0$. Note that U_i^δ is obtained from U_i by shrinking away from the boundary, and if δ is sufficiently small, the collection $\{U_i^{n\delta}\}_{i \le n}$ is also an open cover of X (where n is the number of open sets in the collection). We will find a global L^2_2 gauge transformation $\tilde{g} : X \to \mathrm{U}(1)$ for which $g_i^{-1}\tilde{g}$ is smooth on each $U_i^{n\delta}$; once this is done, we will have that $\tilde{g}(\nabla, \phi)$ is smooth, since $(g_i^{-1}\tilde{g})(g_i(\nabla, \phi))$ is smooth on each $U_i^{n\delta}$. We construct a sequence of local gauge transformations $\tilde{g}^{(k)}$ on $\cup_{i \le k} U_i^{k\delta}$ inductively as follows:

- We define the local gauge transformation $\tilde{g}^{(1)}$ on U_1^δ by taking $\tilde{g}^{(1)} = g_1|_{U_1^\delta}$. Clearly $g_1^{-1}\tilde{g}^{(1)}$ is smooth, as it is constant.
- Suppose there is a local gauge transformation $\tilde{g}^{(k)}$ on $\cup_{i \le k} U_i^{k\delta}$ such that $g_i^{-1}\tilde{g}^{(k)}$ is smooth for every $i \le k$. We now define new local gauge transformations $\tilde{g}_i^{(k+1)}$ on each $U_i^{(k+1)\delta}$ in two cases:

– If $U_{k+1}^{(k+1)\delta} \cap (\cup_{i\leq k} U_i^{(k+1)\delta}) = \varnothing$, we define $\tilde{g}_i^{(k+1)} = \tilde{g}^{(k)}|_{U_i^{(k+1)\delta}}$ for $i \leq k$ and $\tilde{g}_{k+1}^{(k+1)} = g_{k+1}|_{U_{k+1}^{(k+1)\delta}}$. We clearly have that $g_i^{-1}\tilde{g}^{(k+1)}$ is smooth for every $i \leq k+1$.

– Otherwise, we know that there must be some smooth function $\phi_{k+1} : U_{k+1}^{(k+1)\delta} \to [0,1]$ which is equal to 0 on $U_{k+1}^{(k+1)\delta} \cap (\cup_{i\leq k} U_i^{(k+1)\delta})$, and is 1 on $U_{k+1}^{(k+1)\delta} \setminus (\cup_{i\leq k} U_i^{k\delta})$ (this follows from the fact that X is normal, being a Riemann surface). We now define the local gauge transformations $\tilde{g}_i^{(k+1)} = \tilde{g}^{(k)}|_{U_i^{(k+1)\delta}}$ for $i \leq k$, and we define $\tilde{g}_{k+1}^{(k+1)}$ on $U_{k+1}^{(k+1)\delta}$ as follows:

$$\tilde{g}_{k+1}^{(k+1)} = \begin{cases} \tilde{g}^{(k)} \exp\Big(\phi_{k+1} \cdot \ln((\tilde{g}^{(k)})^{-1} g_{k+1})\Big) & \text{on } U_{k+1}^{(k+1)\delta} \cap (\cup_{i\leq k} U_i^{(k+1)\delta}) \\ g_{k+1} & \text{otherwise.} \end{cases} \tag{4.22}$$

Note that ϕ_{k+1} is being used to interpolate between $\tilde{g}^{(k)}$ on $\cup_{i\leq k} U_i^{(k+1)\delta}$ and g_{k+1} on the rest of U_{k+1}. In fact, we have that $\tilde{g}_{k+1}^{(k+1)} = \tilde{g}_i^{(k+1)}$ for $i \leq k$, since $\phi_{k+1} = 0$ on the overlap. Thus, we can stitch together the local gauge transformations into one gauge transformation $g^{(k+1)}$ on all of $\cup_{i\leq k+1} U_i^{(k+1)\delta}$. Moreover, $g_i^{-1}\tilde{g}^{(k+1)}$ is smooth for every $i \leq k+1$; for $i \leq k$ it is constant and for $i = k+1$ we have that

$$g_{k+1}^{-1}\tilde{g}_{k+1}^{(k+1)} = (g_{k+1}^{-1} g_i)(g_i^{-1}\tilde{g}^{(k)} \exp\Big(\phi_{k+1} \cdot \ln((\tilde{g}^{(k)})^{-1} g_i g_i^{-1} g_{k+1})\Big) \tag{4.23}$$

on $U_i^{(k+1)\delta}$. But this is smooth, as all of the components are smooth by construction.

As such, we proceed inductively up to $k = n$ to find a global gauge transformation $\tilde{g} = \tilde{g}^{(n)}$ for which $g_i^{-1}\tilde{g}$ is smooth for every i. But then $\tilde{g}(\nabla, \phi)$ is a smooth solution to the vortex equations. □

Finally, we must verify that solutions are unique on an orbit of $\tilde{\mathscr{G}}^{\mathbb{C}}$, up to unitary gauge equivalence.

Lemma 4.23. *Let $(\nabla, \phi) \in \tilde{\mathscr{N}}$, and let Γ be the orbit of (∇, ϕ) under $\tilde{\mathscr{G}}^{\mathbb{C}}$. Suppose $\mu(\nabla, \phi) = 0$, and there is some $g \in \tilde{\mathscr{G}}^{\mathbb{C}}$ for which $\mu(\nabla^g, \phi^g) = 0$. Then g is a unitary gauge transformation, meaning the element of Γ for which $\mu = 0$ is unique up to unitary gauge equivalence.*

Proof. If $\mu(\nabla, \phi) = \mu(\nabla^g, \phi^g) = 0$ then, by the definition of μ and the gauge invariance of F^∇, we must have the following:

$$|g\phi|_h^2 = \tau - 2i\Lambda F^\nabla = |\phi|_h^2. \tag{4.24}$$

Since g and ϕ are both smooth, this implies that $|g|^2 = 1$ everywhere, making g a unitary gauge transformation. □

Theorem 4.24 (Bradlow, Garcia-Prada)**.** *There is a smooth solution to the vortex equations, and up to unitary gauge equivalence, it is uniquely determined by the positions of d zeros on X. In other words, the moduli space of τ-vortices is diffeomorphic to $S^d(X)$.*

Proof. In Lemma 4.20 we have shown that, on each $(\tilde{\mathscr{G}}^{\mathbb{C}})^2$ orbit in $\tilde{\mathscr{N}}$, there is an L_1^2-weakly convergent sequence of sections and connections for which $\|\mu\|^2$ converges to its minimum value on the orbit. We have also shown in Lemma 4.21 that the limit of this sequence is in the same $(\tilde{\mathscr{G}}^{\mathbb{C}})^2$ orbit. Thus, we have an L_1^2 pair (∇, ϕ) on each orbit which minimises $\|\mu\|^2$, and by Lemma 4.22 we can apply a unitary gauge transformation to make it into a C^∞ pair. By Lemma 3.9 and the invariance of μ under $\mathscr{G}^2$, this implies that $\mu(\nabla, \phi) = 0$, meaning (∇, ϕ) satisfy the vortex equations.

This demonstrates that there is a solution on each orbit. Each one is unique up to unitary gauge equivalence by Lemma 4.23, meaning each point of the quotient space $\tilde{\mathscr{N}}/\tilde{\mathscr{G}}^{\mathbb{C}}$ corresponds to a solution. But $\tilde{\mathscr{N}}/\tilde{\mathscr{G}}^{\mathbb{C}}$ is diffeomorphic to $S^d(X)$ by Proposition 3.15, which completes the proof. □

4.3 Further Directions

We have demonstrated that the moduli space of τ-vortices is isomorphic to S^dX whenever τ is sufficiently large. An interesting observation due to Garcia-Prada is that S^dX inherits a Kähler structure from this theorem. This is because $\mathscr{N}$ already has a Kähler structure according to Equations 3.18 and 3.28, and symplectic reduction theory demonstrates that this is preserved under the quotients by U(1) and $\tilde{\mathscr{G}}^{\mathbb{C}}$ (for details see [Mar+07]).

As we mentioned in the Introduction, it was shown in [MOY97] that the moduli space of the vortex equations on a Riemann surface X is essentially equivalent to the moduli space of the Seiberg-Witten equations on a Seifert fibered 3-manifold over X. This result, in conjunction with the theorem we have just proved, has interesting implications on the theory of 3-manifolds. In a sense, however, the result is incomplete: even if we restrict to Seifert fibered 3-manifolds which are orientable, there are some which fibre over non-orientable surfaces such as $\mathbb{RP}^2$. On the other hand, every Riemann surface is necessarily orientable; a complex atlas on any manifold constitutes an orientation, since the transition maps are biholomorphisms and hence positively oriented. This leads to the following question: can we define the vortex equations on non-orientable surfaces?

If the question is taken literally, the answer is no: the complex structure is necessary to define holomorphic sections of a line bundle, and even if we drop this condition, the Hodge star cannot be defined on non-orientable surfaces either. However, one can work around these obstructions by forming the *orientable double covering* of the surface (an orientable 2-fold covering space, which always exists), performing the analysis on a holomorphic line bundle over this space, and then projecting back down to the non-orientable surface. More generally, one may consider a Riemann surface X with a *Real structure* on it, that is, an anti-holomorphic involution σ, and form a *Real holomorphic line bundle* L over X, that is, a holomorphic line bundle with an anti-holomorphic involution τ compatible with σ. (One obtains a non-orientable surface by taking the quotient of X by σ.) A Riemann surface with a Real structure is called a *Klein surface.*

In fact, this strategy is not new. It was shown in [Sch17] that the Yang-Mills equations extend to Klein surfaces in this way, and that an important theorem (the Narasimhan-Seshadri theorem) carried over as well. Thus, the moduli space of vortices on Klein surfaces could conceivably be found using the same method we used for Riemann surfaces, and subsequently be used to complete the work in [MOY97] to all Seifert fibered 3-manifolds.

Appendix A

Cohomology Theory

We use cohomology theory extensively in Chapter 2 to characterise line bundles topologically, and again in Chapter 4 to demonstrate the exactness of certain differential forms. In this appendix, we give a brief primer in cohomology.

Note that this introduction is very brief, and very few proofs are done. For a more complete introduction to singular cohomology and de Rham cohomology, we refer the reader to [BT82] and [Hat02]. For a more comprehensive discussion of Čech cohomology and homological algebra, we refer the reader to [Bre97] and [Osb00].

A.1 Abstract Theory

We begin by defining cochain complexes.

Definition A.1. Fix a unital ring R (which we usually take to be $\mathbb{Z}$ or $\mathbb{R}$). A *cochain complex with coefficients in* R is a sequence of R-modules $\{A^i\}_{i\in\mathbb{N}}$ together with module homomorphisms $\mathrm{d}^i : A^i \to A^{i+1}$ called *coboundary maps* for which $\mathrm{d}^2 = 0$. Given two cochain complexes $(A^\bullet, \mathrm{d}^\bullet)$ and $(A^\bullet, \delta^\bullet)$, a *cochain map* is a collection of module homomorphisms $f^i : A^i \to B^i$ commuting with the coboundary maps. The collection of cochain complexes forms a category, with cochain maps as morphisms.

That $\mathrm{d}^2 = 0$ implies that $\mathrm{im}(\mathrm{d}^{i-1}) \subseteq \ker(\mathrm{d}^i)$ for every i. If the reverse inclusion holds then the sequence is exact; if not, the cohomology measures the extent to which exactness fails.

Definition A.2. Let $(A^\bullet, \mathrm{d}^\bullet)$ be a cochain complex. The *nth cohomology group* of the chain complex is the following quotient group:

$$H^n(A;R) = \frac{\ker(\mathrm{d}^n)}{\mathrm{im}(\mathrm{d}^{n-1})}. \tag{A.1}$$

It is worth noting that there is a dual notion of a *chain complex,* obtained by reversing all arrows and removing the prefix "co" wherever it appears. Given a cochain complex $(A^\bullet, \mathrm{d}^\bullet)$, one can produce a chain complex $(A_\bullet, \partial_\bullet)$ by taking each A^i to $A_i := \mathrm{Hom}(A^i, R)$ and each coboundary map $\mathrm{d}^i : A^i \to A^{i+1}$ to the boundary map $\overline{\partial}_{i+1} = (\mathrm{d}^i)^* : A_{i+1} \to A_i$.

By taking $\ker(\partial_n)/\mathrm{im}(\partial_{n+1})$ for an arbitrary chain complex, we get the *homology groups* of the chain complex.

We will also be interested in cohomology with coefficients in a sheaf over a topological space. We therefore briefly define sheaves:

Definition A.3. Let X be a topological space, and denote by $\mathscr{U}(X)$ the category whose objects are open sets and whose morphisms are inclusions. Then a *presheaf of abelian groups on* X is a contravariant functor $\mathscr{F} : \mathscr{U}(X) \to \mathrm{Ab}$, the category of abelian groups. Given an open set U, elements of the group $\mathscr{F}(U)$ are called *sections;* given an inclusion $U \hookrightarrow V$, the induced homomorphism $\mathscr{F}(V) \to \mathscr{F}(U)$ is called a *restriction morphism* and is denoted by $\sigma \mapsto \sigma|_U$.

A *sheaf* is a presheaf satisfying the following two conditions:

- If $U \subseteq X$ is open, $\{U_i\}_{i\in I}$ is an open cover of U, and $\sigma, \eta \in \mathscr{F}(U)$ are sections for which $\sigma|_{U_i} = \eta|_{U_i}$ for all i, then $\sigma = \eta$.
- If $U \subseteq X$ is open, $\{U_i\}_{i\in I}$ is an open cover of U, and $\{\sigma_i \in \mathscr{F}(U_i)\}_{i\in I}$ is a family of sections for which $\sigma_i|_{U_i\cap U_j} = \sigma_j|_{U_i\cap U_j}$ for all $i, j \in I$, then there is a section $\sigma \in \mathscr{F}(U)$ for which $\sigma|_{U_i} = \sigma_i$ for all $i \in I$.

We will make extensive use of several sheaves. Firstly, given any abelian group A and any connected topological space X, we can form the constant sheaf $\underline{A}$ over X by defining $\underline{A}(U) = A$ for every $U \subseteq X$, and making restriction morphisms trivial. We will be especially interested in the constant sheaves $\underline{\mathbb{Z}}$ and $\underline{\mathbb{R}}$. Less trivially, for many reasonable function spaces, one can form a sheaf by taking U to the abelian group of sections over U, with the corresponding restriction morphisms being simply restrictions of sections. For instance:

- If X is a connected smooth manifold, we can form the sheaf of smooth $\mathbb{C}$-valued functions $\mathscr{E}$, as well as the subsheaf of nonvanishing smooth $\mathbb{C}$-valued functions $\mathscr{E}^*$, and the sheaves of differential k-forms Ω^k.
- If X possesses a complex structure, we can analogously define the sheaf of holomorphic functions $\mathscr{O}$ and the sheaf of nonvanishing holomorphic functions $\mathscr{O}^*$.

We can also define a *morphism of sheaves:* given two sheaves $\mathscr{F}$ and $\mathscr{G}$ on a topological space X, a morphism α from $\mathscr{F}$ to $\mathscr{G}$ is just a natural transformation (i.e. a collection of homomorphisms $\alpha_U : \mathscr{F}(U) \to \mathscr{G}(U)$ for each $U \in \mathscr{U}(X)$ which commutes with the restriction morphisms). We can define a sheaf $\ker(\alpha)$ on X by taking $(\ker(\alpha))(U) = \ker(\alpha_U)$, and it inherits restriction morphisms from $\mathscr{F}$. We can also define a *presheaf* $\mathrm{im}(\alpha)$ on X by taking $(\mathrm{im}(\alpha))(U) = \mathrm{im}(\alpha_U)$, and it inherits restriction morphisms from $\mathscr{G}$. (This is not a sheaf in general, but there is a process called *sheafification* which produces a unique sheaf for each presheaf, so we can take this to be the image sheaf of α.) We define this machinery because we are now able to speak of exact sequences of sheaves, and one important example is as follows:

$$0 \to \underline{\mathbb{Z}} \to \mathscr{O} \xrightarrow{\exp(2\pi i\cdot)} \mathscr{O}^* \to 1. \tag{A.2}$$

This is called the *exponential sheaf sequence.* The inclusion $\underline{\mathbb{Z}} \to \mathscr{O}$ is defined by mapping any section $n \in \mathbb{Z}(U)$ to the constant holomorphic function $p \mapsto n$ for $p \in X$, and the surjection $\mathscr{O} \to \mathscr{O}^*$ is defined by mapping a holomorphic function $f : X \to \mathbb{C}$ to the nonvanishing holomorphic function $e^{2\pi i f}$. It is clear that this sequence is exact. Moreover, we can get an analogous sequence by replacing holomorphic functions with smooth functions.

A.2 Čech Cohomology

There is one cohomology theory we will use frequently for computations, namely Čech cohomology. It is built out of a limit of Čech complexes, which in turn are defined based on the combinatorics of open covers of topological spaces, so we will define these concepts first.

Let X be a topological space with a sheaf $\mathscr{F}$, and let $\mathscr{U} = \{U_i\}_{i \in A}$ be an open cover of X. An *n-simplex* is an ordered collection of $n+1$ open sets with nonempty joint intersection. We denote by $\alpha = (\alpha_0, \alpha_1, \ldots, \alpha_n)$ an arbitrary n-simplex (with $\alpha_i \neq \alpha_j$ for $i \neq j$), corresponding to the collection $(U_{\alpha_0}, \ldots, U_{\alpha_n})$, and we denote by $U_\alpha \neq \varnothing$ the joint intersection. A *face* of an n-simplex α is an $(n-1)$-subsimplex with the same ordering. We denote by $\alpha \backslash \alpha_i$ the face of α with U_{α_i} missing.

We now define the relevant cochain complex. An *n-cochain* is a function f on n-simplices $\alpha \subseteq A$, for which $f(\alpha) \in \mathscr{F}(U_\alpha)$. We denote by $\check{C}^n(\mathscr{U}, \mathscr{F})$ the abelian group of all n-cochains. We then define the *coboundary map* $\delta^n : \check{C}^n(\mathscr{U}, \mathscr{F}) \to \check{C}^{n+1}(\mathscr{U}, \mathscr{F})$ as follows:

$$\delta^n(f) = \sum_{k=0}^{n+1} (-1)^k f(\alpha \backslash \alpha_k)|_{U_{\alpha_k}}, \tag{A.3}$$

for all $f \in \check{C}^n(\mathscr{U}, \mathscr{F})$. It is easy to verify that $\delta^2 = 0$, meaning the collection of cochains with the coboundary maps forms a cochain complex. We can therefore form a $\mathscr{U}$-dependent cohomology on X:

$$\check{H}^n(\mathscr{U}, \mathscr{F}) = \frac{\ker(\delta^n)}{\operatorname{im}(\delta^{n-1})}. \tag{A.4}$$

The dependence on the open cover $\mathscr{U}$ is inconvenient. However, if $\mathscr{V}$ is another open cover which refines $\mathscr{U}$, one can show that there is a map $\check{H}^\bullet(\mathscr{U}, \mathscr{F}) \to \check{H}^\bullet(\mathscr{V}, \mathscr{F})$. We are now finally ready to define the Čech cohomology:

Definition A.4. Let X be a topological space, and let $\mathscr{F}$ be a sheaf over X. The *Čech cohomology of X with coefficients in $\mathscr{F}$* is the following direct limit:

$$\check{H}^n(X, \mathscr{F}) = \varinjlim \check{H}^n(\mathscr{U}, \mathscr{F}), \tag{A.5}$$

where the direct limit is taken over the directed system of open covers, with refinement as the relation.

This cover-free definition is superficially much more cumbersome. It turns out that we can work with open covers after all, so long as we choose sufficiently nice ones:

Theorem A.5. *Let X be a topological space with a sheaf $\mathscr{F}$, and suppose $\mathscr{U}$ is a good open cover; that is, it is locally finite with contractible intersections. Then $\check{H}^n(X, \mathscr{F}) \cong \check{H}^n(\mathscr{U}, \mathscr{F})$.*

Supplemented with the following theorem, this alleviates the issue:

Theorem A.6. *Every paracompact manifold admits a good open cover.*

Proof. See, for instance, [BT82]. □

There are two other cohomology theories we will make use of. The first cohomology theory, singular cohomology, arises from the singular chain complex whose R-modules are defined as follows: where Δ^k is the standard k-simplex, we define

$$C_k^{\text{sing}}(X, R) = \left\{ \sum_i r_i \sigma_i \;:\; r_i \in R,\, \sigma_i : \Delta^k \to X \right\}. \tag{A.6}$$

The boundary maps are defined as follows:

$$\partial\sigma = \sum_j (-1)^j \sigma|_{F^j(\Delta^k)}, \tag{A.7}$$

where $F^j(\Delta^k)$ is the j-th face of Δ^k. Since $\partial^2 = 0$, this naturally generates homology groups, and we can dualise the chain complexes and boundary maps to obtain cohomology groups. The second cohomology theory, de Rham cohomology, arises from the chain complex of differential forms with each coboundary map given by the exterior derivative. It follows that the cohomology of the de Rham complex consists of the space of closed forms modulo exact forms.

One might wonder if the three cohomology theories are related. It turns out that they are isomorphic in important cases:

Theorem A.7 (de Rham). *If X is a smooth manifold, then $H^\bullet_{\text{sing}}(X, \mathbb{R}) \cong H^\bullet_{\text{dR}}(X)$.*

Theorem A.8. *If X is a triangulable topological space, then $\check{H}^\bullet(X, \mathbb{Z}) \cong H^\bullet_{\text{sing}}(X, \mathbb{Z})$.*

Proof. For proofs of both of these results, see [GH94]. □

Since we are primarily interested in Riemann surfaces, which are smooth and triangulable, we can essentially identify these cohomology groups. When there is no ambiguity, we will write all cohomology groups as simply $H^\bullet(X, \mathscr{F})$ or $H^\bullet(X, R)$.

A.3 Technical Results

We will be using some technical theorems in our proofs, which we will state but not fully prove.

Theorem A.9 (Poincaré Duality). *Let X be a connected, compact, orientable n-manifold. Then $H^n(X, \mathbb{Z}) \cong \mathbb{Z}$. Furthermore, if $[X]$ is a choice of generator of $H^n(X, \mathbb{Z})$ (called a* fundamental class*), then there is an isomorphism $\alpha : H^k(X, \mathbb{Z}) \to H_{n-k}(X, \mathbb{Z})$ for each $k \in \mathbb{N}$ defined by taking $\alpha(\omega) = \omega \frown [X]$ (where $\frown$ is the cap product).*

Proof. See [Hat02]. □

Theorem A.10. *Fine sheaves are acyclic. That is, if a sheaf $\mathscr{F}$ over a topological space X admits partitions of unity on any open set U (meaning there is a collection of local sections on any open cover of U which are locally finite and sum to the identity on U), then $H^k(X,\mathscr{F}) = 0$ for $k \geq 1$.*

Proof. See [GH94]. □

Theorem A.11. *The sheaves $\mathscr{E}$ and Ω^k are fine, so by the above theorem, their nontrivial cohomology vanishes.*

Proof. See [GH94]. □

Theorem A.12 (Snake Lemma). *Suppose the following diagram commutes in the category of abelian groups, with exact rows:*

$$\begin{array}{ccccccc} & & A & \xrightarrow{f} & B & \xrightarrow{g} & C & \longrightarrow & 0 \\ & & {\scriptstyle a}\uparrow & & {\scriptstyle b}\uparrow & & {\scriptstyle c}\uparrow & & \\ 0 & \longrightarrow & A' & \xrightarrow{f'} & B' & \xrightarrow{g'} & C' & & \end{array}$$

Then there is a map $\delta : \ker(c) \to \operatorname{coker}(a)$ making the following sequence exact:

$$\ker(a) \longrightarrow \ker(b) \longrightarrow \ker(c) \xrightarrow{\delta} \operatorname{coker}(a) \longrightarrow \operatorname{coker}(b) \longrightarrow \operatorname{coker}(c).$$

In fact, $\delta = f^{-1} \circ b \circ (g')^{-1}$, and this map is well-defined on $\ker(c)$ under the quotient $A \to \operatorname{coker}(a)$.

Proof. See [Osb00]. □

Corollary A.13. *Let $0 \to \mathscr{F} \xrightarrow{a} \mathscr{G} \xrightarrow{b} \mathscr{H} \to 0$ be an exact sequence of sheaves over a fixed topological space X. Then there is an induced short exact sequence of chain complexes $0 \to \check{C}^\bullet(\mathscr{U},\mathscr{F}) \xrightarrow{a} \check{C}^\bullet(\mathscr{U},\mathscr{G}) \xrightarrow{b} \check{C}^\bullet(\mathscr{U},\mathscr{H}) \to 0$ for any open cover $\mathscr{U}$ of X, and consequently there is a long exact sequence of abelian groups given as follows:*

$$\begin{array}{l} 0 \longrightarrow H^0(X,\mathscr{F}) \longrightarrow H^0(X,\mathscr{G}) \longrightarrow H^0(X,\mathscr{H}) \\ \quad \hookrightarrow H^1(X,\mathscr{F}) \longrightarrow H^1(X,\mathscr{G}) \longrightarrow H^1(X,\mathscr{H}) \\ \quad \hookrightarrow H^2(X,\mathscr{F}) \longrightarrow H^2(X,\mathscr{G}) \longrightarrow H^2(X,\mathscr{H}) \\ \quad \cdots \\ \quad \dashrightarrow H^k(X,\mathscr{F}) \longrightarrow H^k(X,\mathscr{G}) \longrightarrow H^k(X,\mathscr{H}) \longrightarrow \cdots \end{array}$$

The connecting morphisms $\delta : H^k(X,\mathscr{H}) \to H^{k+1}(X,\mathscr{F})$ are given by $\mathrm{d}^k \circ b^{-1}$ on cocycles.

References

[AB83] Michael Atiyah and Raoul Bott. “The Yang-Mills Equations over Riemann Surfaces”. In: *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* 308.1505 (1983), pp. 523–615.

[BCS57] John Bardeen, Leon Cooper, and John Schrieffer. “Theory of Superconductivity”. In: *Physics Review* 108.5 (1957), pp. 1175–1204.

[BH21] A. Behzadan and M. Holst. “Multiplication in Sobolev spaces, revisited”. In: *Arkiv för Matematik* 59.2 (2021), pp. 275–306. DOI: `10.4310/ARKIV.2021.v59.n2.a2`. URL: `https://doi.org/10.4310/ARKIV.2021.v59.n2.a2`.

[Bra90] Steven Bradlow. “Vortices in Holomorphic Line Bundles over Closed Kähler Manifolds”. In: *Communications in Mathematical Physics* 135 (1990), pp. 1–17.

[Bre97] Glen Bredon. *Sheaf Theory*. Berlin: Springer-Verlag, 1997.

[BT82] Raoul Bott and Loring Tu. *Differential Forms in Algebraic Topology*. New York: Springer-Verlag, 1982.

[Cav22] Gil Cavalcanti. *Complex geometry. Holomorphic vector bundles, elliptic operators and Hodge theory*. 2022.

[Con90] John Conway. *A Course in Functional Analysis*. 2nd ed. New York: Springer-Verlag, 1990.

[Don11] Simon Donaldson. *Riemann Surfaces*. New York: Oxford University Press, 2011.

[Don83] Simon Donaldson. “A New Proof of a Theorem of Narasimhan and Seshadri”. In: *Journal of Differential Geometry* 18 (1983), pp. 269–277.

[FWX25] Yu Feng, Shuo Wang, and Bin Xu. *A note On the existence of solutions to Hitchin’s self-duality equations*. 2025. arXiv: `2501.10976 [math.DG]`. URL: `https://arxiv.org/abs/2501.10976`.

[Gar94] Oscar García-Prada. “A Direct Existence Proof for the Vortex Equations over a Compact Riemann Surface”. In: *Bulletin of the London Mathematical Society* 26 (1994), pp. 88–96.

[GH94] Phillip Griffiths and Joseph Harris. *Principles of Algebraic Geometry*. John Wiley & Sons, Ltd, 1994.

[Hat02] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.

[Hat03] Allen Hatcher. *Vector Bundles and K-Theory*. Cambridge University Press, 2003.

[Hay19] Andriy Haydys. *Introduction to gauge theory*. 2019. URL: `https://arxiv.org/abs/1910.10436`.

[Hay24] Andriy Haydys. *Global analysis*. 2024. URL: https://haydys.net/misc/GlobAnalysis.pdf.

[Heb96] Emmanuel Hebey. *Sobolev Spaces on Manifolds*. Lecture Notes in Mathematics 1635. Berlin: Springer-Verlag, 1996.

[Hus94] Dale Husemoller. *Fibre Bundles*. New York: Springer New York, 1994.

[Huy05] Daniel Huybrechts. *Complex Geometry. An Introduction*. Heidelberg: Springer Berlin, 2005.

[JT80] Arthur Jaffe and Clifford Taubes. *Vortices and Monopoles. Structure of Static Gauge Theories*. Progress in Physics 2. Switzerland: Birkhauser, 1980.

[Kob87] Shoshichi Kobayashi. *Differential Geometry of Complex Vector Bundles*. New Jersey: Princeton University Press, 1987.

[Mar+07] Jerrold Marsden et al. *Hamiltonian Reduction by Stages*. Lecture Notes in Mathematics 1913. Heidelberg: Springer Berlin, 2007.

[Mor07] Andrei Moroianu. *Lectures on Kähler geometry*. New York: Cambridge University Press, 2007.

[MOY97] Tomasz Mrowka, Peter Ozsvath, and Baozhen Yu. "Seiberg-Witten Monopoles on Seifert Fibered Spaces". In: *Communications in Analysis and Geometry* 5.4 (1997), pp. 685–793.

[MW74] Jerrold Marsden and Alan Weinstein. "Reduction of symplectic manifolds with symmetry". In: *Reports on Mathematical Physics* 5.1 (1974), pp. 121–130.

[Osb00] M. Scott Osborne. *Basic Homological Algebra*. New York: Springer New York, 2000.

[Ram04] S. Ramanan. *Global Calculus*. Vol. 65. Graduate Studies in Mathematics. Rhode Island: American Mathematical Society, 2004.

[Sch17] Florent Schaffhauser. "On the Narasimhan–Seshadri correspondence for real and quaternionic vector bundles". In: *Journal of Differential Geometry* 105.1 (2017), pp. 119–162. DOI: 10.4310/jdg/1483655861.

[Sch97] Vadim Vasil'evich Schmidt. *The Physics of Superconductors*. Ed. by Paul Müller and Alexey Ustinov. Trans. by Irina Grigorieva. Berlin: Springer-Verlag, 1997.

[Uhl82] Karen Uhlenbeck. "Connections with L^p Bounds on Curvature". In: *Communications in Mathematical Physics* 83 (1982), pp. 31–42.

[Vak24] Ravi Vakil. *The Rising Sea. Foundations of Algebraic Geometry*. New Jersey: Princeton University Press, 2024.

[Voi02] Claire Voisin. *Hodge Theory and Complex Algebraic Geometry*. Vol. 1. United Kingdom: Cambridge University Press, 2002.

[Weh04] Katrin Wehrheim. *Uhlenbeck Compactness*. Zürich: European Mathematical Society, 2004.

[Wel08] Raymond Wells. *Differential Analysis on Complex Manifolds*. New York: Springer New York, 2008.