Moduli Spaces of Seiberg-Witten Monopoles on Seifert 3-manifolds

Baylee Verzyde

Department of Mathematics

The University of Auckland

A thesis submitted in fulfilment of the requirements for the degree of MSc in Mathematics, The University of Auckland, 2026.

Abstract

The Seiberg-Witten equations are a system of geometric PDEs whose solutions are of interest in defining 3-manifold invariants. In 1997, it was shown by Mrowka, Ozsvath and Yu that the Seiberg-Witten equations on Seifert 3-manifolds were related to another system of geometric PDEs on their U(1)-orbit space called the vortex equations. The solutions to these equations can, in turn, be characterised geometrically using a moment map argument; this was accomplished by Garcia-Prada in 1994.

The purpose of this thesis is to explore the structure of the moduli spaces of vortices and Seiberg-Witten monopoles, and to explain the correspondence between the two in detail. These concepts are also generalised to the equivariant category: in the case that a Seifert 3-manifold is equipped with the action of a finite group Λ commuting with the circle action, it is shown that the moduli space of Λ-equivariant Seiberg-Witten monopoles on the 3-manifold has the same relationship with the moduli space of Λ-equivariant vortices on the orbit space. Equivariant generalisations of the relevant existence theorems are given; this includes an orbifold generalisation of Uhlenbeck's weak compactness theorem and Uhlenbeck's gauge theorem. In the case that $\Lambda = \mathbb{Z}_2$, the Klein surface formalism is used to generalise the characterisation of the moduli space to Seifert 3-manifolds fibring over non-orientable surfaces.

Acknowledgements

First, I wish to thank my supervisor Pedram Hekmati. I am extremely grateful for the huge amount of time and effort he has taken to support my ambitions. His seemingly endless wisdom and patience have made me the mathematician I am today.

Second, I wish to thank my mother and my two sisters. I could not have gotten close to this far in my mathematical journey without your consistent support over my life. I hope that I am eventually able to return the favour in some way.

Third, I wish to thank my friends Halle, Trav, Tony, and Leander. You have all helped me in different ways so much throughout the years, be that through directly working through a difficult problem of mine or just being there. I suspect you have done much more than you know.

Finally, I wish to express my gratitude for my fellow postgraduate students at Auckland, including (but definitely not limited to) Andrew, Niv, Thomas, Marc, and Jake. I sincerely hope our paths all cross again some time, for both the maths and the babbling in the postgraduate room.

Contents

Introduction

The study of closed orientable 3-manifolds is deep and diverse, tracing back to the work done by Poincaré and Heegaard in the late 1800s [Gor99]. Manifolds in dimension 3 (and 4) lend naturally to interesting topology; closed orientable surfaces are simple enough to be classified by an integer, and in dimension greater than 4 there are techniques such as the Whitney trick [MSS65] that allow for more maps to be constructed between manifolds. Thus, the introduction of specialised machinery is necessary to gain information about the topology of 3-manifolds.

Inequivalent manifolds are generally distinguished through the assignment of topological invariants: an algebraic object is assigned to each diffeomorphism class of manifolds, and one shows that different objects are assigned to two manifolds of interest. One programme for developing such invariants is *Floer theory:* to each manifold with a given structure (spin, symplectic, etc.), one assigns a higher-dimensional manifold and computes its homology using Morse theory. The associated manifold is often the configuration space for some physics-inspired system, and the chain complex for the homology is generated by the critical points of an energy functional. Floer's original application in [Flo88b] was restricted to symplectic manifolds, but it was later adapted in various forms to 3-manifolds.

One of these adaptations of Floer theory was built from a specific system of geometric PDEs as follows. To a 3-manifold Y equipped with a vector bundle called a spinor bundle, one associates the space of sections and connections on that bundle; the configuration space of the theory is the quotient of this space by the gauge group. The critical points of the relevant energy functional are the solutions to the *Seiberg-Witten equations:*

$$*F_A = -\rho^{-1}(\psi \otimes \psi^*), \tag{1}$$

$$D_A\psi = 0. \tag{2}$$

The quotient of the solution space by the gauge group is called the *moduli space of Seiberg-Witten monopoles.* The original context for these equations was supersymmetric field theory on flat 4-space, but they were subsequently adapted to 4-manifolds by Witten in [Wit94], and then to 3-manifolds in [KM97] and [MST96]. That these equations could be used to define Floer homology groups was first pointed out by Donaldson in [Don96]; this idea was realised soon after by Mrowka, Ozsvath and Yu in [MOY96]. Seiberg-Witten Floer homology groups would later be used by Manolescu in [Man15] to disprove the triangulation conjecture in dimension greater than 4.

Thus, the moduli space of solutions to the Seiberg-Witten equations in dimension 3 is of interest for low-dimensional topology. However, as a system of nonlinear geometric PDEs on manifolds, it is difficult even to determine if these moduli spaces are nonempty, let alone their structure. Nevertheless, in 1996, the structure of the moduli space was understood for a fairly large class of 3-manifolds by Mrowka, Ozsvath and Yu in [MOY96]. In this paper, the 3-manifolds of interest were *Seifert 3-manifolds,* i.e., 3-manifolds equipped with a U(1)-action with finite stabilisers. The authors related the Seiberg-Witten equations on the 3-manifold to another system of equations on the U(1)-orbit space called the *vortex equations,* another system of geometric PDEs defined on line bundles over Riemann surfaces as follows:

$$*F_A = \tfrac{i}{2}(|\phi|_h^2 - \tau), \tag{3}$$

$$\overline{\partial}_A \phi = 0. \tag{4}$$

Once the moduli space of solutions to the vortex equations was determined, therefore, they were able to characterise the moduli space of Seiberg-Witten monopoles.

Indeed, solutions to the vortex equations had already been analysed in detail several years prior by Bradlow in [Bra90], where he was able to show that each gauge equivalence class of solutions could be identified with an effective divisor on the underlying surface. Another proof of this result was published in 1994 by Garcia-Prada in [Gar94]; this proof was based on the approach used by Donaldson in [Don83] to study the moduli space of flat vector bundles, which was later adapted to solve Hitchin's self-duality equations in [Hit87]. In particular, Garcia-Prada showed that the vortex equations could be interpreted as the zeros of a *moment map,* and then used a result in the theory of symplectic group actions to construct a solution on each complex gauge orbit. The functional-analytic flavour of this proof strategy is primarily informed by the use of Uhlenbeck's weak compactness theorem, which asserts that the space of connections with bounded curvature is weakly compact up to gauge equivalence; this was first proved by Uhlenbeck in [Uhl82] to prove strong compactness for connections satisfying the Yang-Mills equations.

The results obtained in [MOY96] apply to all Seifert 3-manifolds whose orbit space can be equipped with a Riemann surface structure; the complex structure on the orbit space is necessary to define holomorphicity for the vortex equations. However, there are some orientable Seifert 3-manifolds whose orbit spaces are non-orientable, and therefore do not admit complex structures; examples are the lens spaces $L(4,1)$ and $L(4,3)$ fibring over $\mathbb{RP}^2$. A potential work-around uses the theory of *Klein surfaces.* Instead of attempting to define a complex structure on a non-orientable surface, one may instead define it on the orientable double covering and ensure that the action of $\mathbb{Z}_2$ on the covering flips its sign. Once this is done, one may proceed as usual on the double covering, defining everything to be $\mathbb{Z}_2$-invariant, and project back down onto the base space at the end. The theory of Klein surfaces was initially developed by Klein, Harnack, and Weichold in the late 1800s [Nat90], and its application to gauge-theoretic moduli spaces has been pioneered by Schaffhauser, Ho, and Liu in [Sch17], [HLR08], [LS13].

Broadly speaking, the objective of this thesis is to present the characterisation of the moduli space in [MOY96], including all major prerequisite topics, and to generalise their work to include a finite group action. More precisely, this generalisation aims to characterise the moduli space of Λ-equivariant Seiberg-Witten monopoles, where Λ is a finite group acting on the Seifert 3-manifold in a way that commutes with the circle action. In the light of the above paragraph, this characterisation would allow for an extension of the result in [MOY96] to those Seifert 3-manifolds fibring over non-orientable base spaces, by simply taking $\Lambda = \mathbb{Z}_2$. From a more practical perspective, however, the extra requirement of Λ-equivariance could potentially lead to richer Floer invariants. Equivariant varieties of

Seiberg-Witten Floer homology have been developed for certain classes of 3-manifolds in [Man01], [Lin16], [BH24], and the resulting homology groups were found to have useful extra algebraic structure; the computation of equivariant moduli spaces could be the first step in the development of a new equivariant Floer homology for Seifert 3-manifolds.

The structure of this thesis is summarised as follows.

- Chapter 1 covers the theory of orbifolds; though the Seiberg-Witten equations will always be defined on a 3-manifold Y, the vortex equations will be defined on the orbit space $Y/\mathrm{U}(1)$ which is an orbifold. Particular emphasis is given on orbifold fibre bundles, and how they may be used to define gauge theories on orbifolds. The notion of a Seifert 3-manifold is also discussed at the end of the chapter, primarily through the lens of orbifold circle bundles. Proofs are often omitted unless there is a perceived gap in the literature, as the content is introductory.

- Chapter 2 is devoted to the vortex equations on orbifold Riemann surfaces, including the characterisation of the moduli space of vortices in terms of effective divisors. The theory is presented in the form of its equivariant generalisation, which carries through smoothly when the finite group action preserves orientation. The existence proof is adapted from Garcia-Prada's existence proof in [Gar94], which lends more naturally to equivariant generalisation than the original existence proof by Bradlow. Its adaptation to orbifolds requires a generalisation of Uhlenbeck's results to the orbifold category, so this generalisation is also presented; even though we only ever take the structure group to be $\mathrm{U}(1)$, the extension proved in this chapter applies for any compact Lie group.

- Chapter 3 is an introduction to the Seiberg-Witten equations on 3- and 4-manifolds. The author has attempted to separate the pointwise linear algebra from the bundle-theoretic differential constructions as cleanly as possible. This chapter also serves to set up the notation used in the remainder of the thesis.

- Chapter 4 explores the correspondence between the vortex equations and the Seiberg-Witten equations proved by Mrowka et al., and the correspondence is generalised to include a finite group action. The characterisation of the irreducible component of the moduli space is essentially a direct adaptation from [MOY96], but the topology of the reducible component differs significantly in the equivariant case. Once the moduli spaces have been characterised in general, the result is applied to Seifert 3-manifolds over non-orientable base spaces, as well as lens spaces. Further directions are briefly discussed at the end of the chapter.

- Two appendices accompany the main text. The first appendix outlines the functional-analytic setup for geometric analysis on manifolds, with interspersed remarks regarding generalisations to orbifolds. This is essentially only used in Chapter 2. The second appendix details the convention differences that exist in defining the Seiberg-Witten equations across the literature, and what we have chosen throughout the thesis. This is used throughout Chapters 3 and 4.

Conventions and notation: Orbifolds are locally quotients by *faithful* group actions. All groups act on the left unless otherwise stated. All cohomology groups are singular $\mathbb{Z}$-valued cohomology or de Rham cohomology unless otherwise indicated. All manifolds are compact and connected unless otherwise stated. All line bundles are complex Hermitian, and all connections preserve the Hermitian structure. From Chapter 2 onwards, every chosen object is assumed equivariant under the action of the finite group Λ unless otherwise stated. All complex inner products are conjugate linear in the second entry.

Chapter 1

Orbifolds

This chapter is dedicated to the geometry and topology of orbifolds and orbifold fibre bundles. We begin with a brief review of the definitions of orbifolds and smooth maps, and establish some basic examples. We then discuss orbifold fibre bundles, from both a local and a global perspective, and we briefly comment on the algebraic invariants involved in the study of bundles. This theory is applied to the specific case of complex line bundles over orbifold Riemann surfaces; in particular, we define the invariants that allow such bundles to be completely classified, and we touch on their relationship to unitary connections. Finally, we define Seifert fibred spaces in terms of bundles over orbifold Riemann surfaces, and we describe the relationship between bundles on Seifert fibred spaces and their underlying orbifold Riemann surfaces.

In our exposition, we frequently use results without proof or gloss over subtleties. For a more comprehensive discussion of orbifolds, the reader is referred to [Jr22], [CHK00], and [ALR07]; for a more comprehensive discussion of Seifert fibred spaces, the reader is referred to [Hat01] and [Sco83].

1.1 Basic Definitions

Orbifolds are essentially equivariant generalisations of manifolds; whereas manifolds are built by gluing Euclidean patches together, orbifolds are instead built from quotients of Euclidean spaces by finite groups. Described this way, it may be tempting to simply define an orbifold to be locally homeomorphic to a quotient of $\mathbb{R}^n$ by a finite group. However, this approach loses valuable information about the group structure; for instance, note that $\mathbb{R}^2/\mathbb{Z}_\alpha$ (with the natural action by rotations) would be equivalent to $\mathbb{R}^2$ for every α. The definition of an orbifold is built to incorporate the group structure of the quotient directly.

We begin by defining the notion of coordinate charts on orbifolds.

Definition 1.1. Let X be a topological space. An *n-dimensional orbifold chart on X* consists of an open subset $\widetilde{U} \subseteq \mathbb{R}^n$, a finite group Γ acting faithfully via diffeomorphisms on $\widetilde{U}$, and a continuous map $\varphi : \widetilde{U} \to X$ which is invariant under the action of Γ and induces a homeomorphism between $\widetilde{U}/\Gamma$ and $\varphi(\widetilde{U})$. We will often refer to the tuple $(\widetilde{U}, \Gamma, \varphi)$ as an orbifold chart, and we will denote $\varphi(\widetilde{U})$ by U.

Remark. Though Γ may act by diffeomorphisms in general, we can always reduce an orbifold chart such that Γ acts by linear orthogonal transformations on $\widetilde{U}$. The idea is that we can put a Γ-invariant Riemannian metric on $\widetilde{U}$, and then identify an open subset of $T_p\widetilde{U}$ (on which Γ acts linearly) with an open subset of $\widetilde{U}$ using the exponential map. For details, see [Jr22].

Definition 1.2. Let $(\widetilde{U}, \Gamma, \varphi)$ and $(\widetilde{V}, \Lambda, \psi)$ be two orbifold charts on X. An *embedding of* $(\widetilde{U}, \Gamma, \varphi)$ *into* $(\widetilde{V}, \Lambda, \psi)$ is a smooth embedding $i : \widetilde{U} \to \widetilde{V}$ for which $\psi \circ i = \varphi$.

At this point, the reader may wonder why the notion of an orbifold chart embedding makes no reference to the group structure of each chart. It turns out that this definition is sufficient to tightly constrain the behaviour of i under Γ and Λ, so much so that Γ must be a subgroup of Λ and i must be Γ-equivariant. For a proof of this fact, refer to [MP97] and [Jr22]. We will henceforth assume that an embedding induces a subgroup inclusion $\Gamma \leq \Lambda$.

We are now ready to define orbifolds.

Definition 1.3. Let X be a topological space. An *n-dimensional orbifold atlas on X* is a collection of n-dimensional orbifold charts $\{(\widetilde{U}_i, \Gamma_i, \varphi_i)\}_{i\in I}$ for which any $x \in U_i \cap U_j$ has a corresponding chart $(\widetilde{U}_k, \Gamma_k, \varphi_k)$ containing x that admits embeddings into U_i and U_j. An orbifold atlas is a *refinement* of another atlas if every chart in the first embeds into a chart in the second, and two orbifold atlases are *equivalent* if they admit a common refinement.

Definition 1.4. An *n-dimensional orbifold* is a paracompact Hausdorff space X equipped with an equivalence class of n-dimensional orbifold atlases (or equivalently one maximal orbifold atlas). If the embeddings are all orientation-preserving then we call the atlas an *oriented atlas,* and X is an *oriented orbifold.* The underlying space will be denoted by $|X|$ if we wish to distinguish between the orbifold and its underlying topological space.

Remark. If the embeddings are constrained to be differentiable, smooth, or complex, then we call the resulting orbifold differentiable, smooth, or complex. Unless otherwise stated, every orbifold we define is smooth.

Example 1.5. If $\mathbb{R}^n$ is equipped with a faithful linear action by a finite group Γ, the quotient space $\mathbb{R}^n/\Gamma$ has a natural orbifold atlas consisting of the chart $(\mathbb{R}^n, \Gamma, \pi)$ (where $\pi : \mathbb{R}^n \to \mathbb{R}^n/\Gamma$ is the quotient map). Note that if $\Gamma = \{e\}$, the orbifold $\mathbb{R}^n/\Gamma$ is Euclidean; if $\Gamma = \mathbb{Z}_2$ acts via reflection in the hyperplane $[x_n = 0]$, the orbifold $\mathbb{R}^n/\mathbb{Z}_2$ is naturally identified with the upper half-plane; and if $\Gamma = (\mathbb{Z}_2)^{\oplus n}$ acts by reflection in n perpendicular hyperplanes, the orbifold $\mathbb{R}^n/(\mathbb{Z}_2)^{\oplus n}$ is naturally identified with the n-dimensional orthant. It follows that manifolds-with-boundary and manifolds-with-corners can be endowed with natural orbifold structures.

Example 1.6. For our purposes, 2-dimensional orbifolds are of particular interest. As $\mathrm{O}(2)$ can be described as $\mathrm{U}(1) \rtimes \mathbb{Z}_2$, the only finite subgroups of $\mathrm{O}(2)$ are of the form $\mathbb{Z}_\alpha$ or $\mathbb{Z}_\alpha \rtimes \mathbb{Z}_2$ for $\alpha \in \mathbb{N}$, the latter being the dihedral group of order α. These correspond to the finite groups which can be faithfully represented on $\mathbb{R}^2$, meaning a 2-orbifold must look locally like $\mathbb{R}^2/\mathbb{Z}_\alpha$ or $\mathbb{R}^2/D_\alpha$ for some $\alpha \in \mathbb{N}$ (around points where Γ is nontrivial). $\mathbb{R}^2/\mathbb{Z}_\alpha$ can be interpreted as the boundary of a 3-dimensional cone, whereas $\mathbb{R}^2/D_\alpha$ can be naturally interpreted as a flat wedge.

One might notice that certain points are distinguished from others in the orbifolds we have defined thus far: the boundary/corners are distinguished in manifolds with corners, and the cone points are distinguished in $\mathbb{R}^2/\mathbb{Z}_\alpha$. The distinguishing feature is that these points have nontrivial isotropy under the finite group when they are pulled back to $\mathbb{R}^n$. Thus, we make the following definition:

Definition 1.7. Let X be an orbifold, and let $x \in X$. Given a chart $(\widetilde{U}, \Gamma, \varphi)$ containing x, the *local group at* x is defined to be the isotropy subgroup of Γ at any point in $\varphi^{-1}(x)$. It is denoted by Γ_x, and its isomorphism class is independent of the choice of chart and the element of $\varphi^{-1}(x)$.

A point at which the local group is trivial is called a *regular point,* and the collection of all such points is denoted by X_{reg}; this defines an open submanifold of X. Its complement is said to be *singular,* and it defines a subspace of codimension at least 1.

The local group dictates the structure of the orbifold around a given point: if the local group at x is Γ_x, then the orbifold looks like $\mathbb{R}^n/\Gamma_x$ around x. More precisely, we have the following:

Proposition 1.8. *Every orbifold has an atlas of the form $\{(\widetilde{U}_p, \Gamma_p, \varphi_p)\}_{p \in X}$, where each U_p is a neighbourhood of p and the preimage $\varphi_p^{-1}(p)$ consists of a single point.*

Proof. Let $(\widetilde{U}, \Gamma, \varphi)$ be a chart around $p \in X$, and choose some $x \in \varphi^{-1}(p)$. Observe that $\varphi^{-1}(p)$ is a discrete subset of $\widetilde{U}$, so there is a connected neighbourhood $\widetilde{V} \ni x$ such that $g\widetilde{V} \cap \widetilde{V} \neq \varnothing$ if and only if $g \in \Gamma_p$. By replacing $\widetilde{V}$ with $\cap_{g \in \Gamma_p} g\widetilde{V}$, we can assume that $\widetilde{V}$ is a Γ_p-invariant subset of $\widetilde{U}$. We then define $\lambda : \widetilde{V} \to \widetilde{U}$ to be the inclusion, and $\psi : \widetilde{V} \to V$ to be $\varphi|_{\widetilde{V}}$ (where $V = \psi(\widetilde{V})$ by definition). It is clear that $(\widetilde{V}, \Gamma_p, \psi)$ is a chart for which $\psi^{-1}(p) = \{x\}$, and moreover that λ is an embedding of charts. Repeating this construction for every p, we obtain an equivalent atlas with the desired property. □

We may now discuss the concept of a smooth map between orbifolds. Intuitively, a smooth map should be one which descends to a smooth map between charts; however, the finite quotient complicates matters, as we need some notion of "local equivariance" for smooth functions. This has led to several inequivalent notions of a smooth map between orbifolds. We will use the following definition in general, but we will introduce a specialisation when we come to discuss fibre bundles and pullbacks.

Definition 1.9. Let X and Y be orbifolds, and let $f : X \to Y$ be a continuous map. Given a point $x \in X$, a *smooth local lift for* f *at* x consists of the following:

- a chart $(\widetilde{U}, \Gamma_x, \varphi)$ about x and a chart $(\widetilde{V}, \Gamma_{f(x)}, \psi)$ about $f(x)$, for which $f(U) \subseteq V$;
- a group homomorphism $\alpha_x : \Gamma_x \to \Gamma_{f(x)}$;
- and a smooth map $\widetilde{f}_x : \widetilde{U} \to \widetilde{V}$;

for which $\widetilde{f}_x$ is Γ_x-equivariant (i.e., $\widetilde{f}_x(g \cdot y) = \alpha_x(g) \cdot \widetilde{f}_x(y)$ for every $y \in \widetilde{U}$ and every $g \in \Gamma_x$) and for which the following diagram commutes:

$$\begin{array}{ccc} U & \xrightarrow{f} & V \\ {\scriptstyle\varphi}\uparrow & & \uparrow{\scriptstyle\psi} \\ \widetilde{U} & \xrightarrow{\widetilde{f}_x} & \widetilde{V} \end{array}$$

The map f is called *smooth* if it admits a smooth local lift at every point. Analogously, f is called *analytic* or *holomorphic* if the map $\widetilde{f}_x$ shares these properties at each point.

The idea behind this definition is that the map f can be locally thought of as a Γ_x-equivariant smooth map $\widetilde{f}_x$ between subsets of $\mathbb{R}^n$, and the homomorphism α_x specifies exactly how elements of Γ_x should be thought of as acting on the image.

Example 1.10. Let $f : [0,\infty) \to \mathbb{R}$ be continuous, and equip $[0,\infty) \cong \mathbb{R}/\mathbb{Z}_2$ with its natural orbifold structure. The smoothness of f around any nonzero input is equivalent to the usual definition. Around $x = 0$, however, the local lift condition implies that f lifts to an even function from $\mathbb{R}$ to $\mathbb{R}$, meaning $f'(0)$ must be zero for f to be smooth. More generally, if a manifold with boundary M is considered to be an orbifold, a local lift of f at $x \in \partial M$ gives local coordinates around x in which the normal derivative of f along the boundary is zero. Thus, smooth functions on manifolds with boundary naturally satisfy Neumann boundary conditions in the orbifold formalism.

1.2 Orbifold Fibre Bundles

To develop differential geometry on orbifolds, we need to generalise the notion of a fibre bundle. In order for a fibre bundle to respect the local group structure of its base space, the local triviality condition needs to be altered. We are led to the following definition.

Definition 1.11. Let X and F be orbifolds. An *orbifold fibre bundle over X with fibre F* consists of another orbifold E and a surjective smooth map $p : E \to X$ with the following property: for every $x \in X$, there is an orbifold chart $(\widetilde{U}, \Gamma_x, \varphi)$ about x and a smooth extension of the action of Γ_x on $\widetilde{U}$ to $\widetilde{U} \times F$ (which respects the fibre-wise structure of F), with a diffeomorphism $(\widetilde{U} \times F)/\Gamma_x \leftrightarrow p^{-1}(U)$ making the following diagram commute:

$$\begin{array}{ccc} (\widetilde{U} \times F)/\Gamma_x & \longleftrightarrow & p^{-1}(U) \\ {\scriptstyle \mathrm{pr}_1/\Gamma_x}\downarrow & & \downarrow{\scriptstyle p} \\ \widetilde{U}/\Gamma_x & \longleftrightarrow & U \end{array}$$

A *section of* E is a smooth map $\sigma : X \to E$ which is a right-inverse for p, and the space of such sections is denoted by $\Gamma(E)$. In the special case that F is an n-dimensional vector space and the action of Γ_x on $\widetilde{U} \times F$ is fibrewise linear, E is called an *orbifold vector bundle with rank n*.

Remark. In the case that $\pi : E \to X$ is an orbifold vector bundle, it is no longer necessarily true that the fibres of E possess a vector space structure; indeed, a fibre over a singular point may not even be homeomorphic to a Euclidean space. However, since the action of Γ_x is taken to be linear in the above definition, there is still a well-defined vector space structure on the lifting $\widetilde{U} \times F$.

Intuitively, while a fibre bundle is locally a product with F, an orbifold fibre bundle is locally a quotient of a product with F by a finite group action. In general, if $\widetilde{U} \times F$ is equipped with a Γ-action extending the action on $\widetilde{U}$, we can write this action as follows:

$$\gamma \cdot (x, v) = (\gamma \cdot x, \rho_x(\gamma) \cdot v). \tag{1.1}$$

Here $\gamma \in \Gamma$, $x \in \widetilde{U}$ and $v \in F$ are arbitrary, while $\rho_x : \Gamma \to G$ denotes the action of each element of Γ on the fibre F over x, and $G \leq \mathrm{Aut}(F)$ is the *structure group of* F. The requirement that $\gamma \cdot (\gamma' \cdot (x, v)) = (\gamma\gamma') \cdot (x, v)$ manifests in this language as follows:

$$\rho_x(\gamma\gamma') = \rho_{\gamma' \cdot x}(\gamma) \circ \rho_x(\gamma'). \tag{1.2}$$

Thus, the map ρ_x behaves somewhat like an $\mathrm{Aut}(F)$-representation of Γ. In fact, if x^* is a fixed point of Γ, we see that ρ_{x^*} is a genuine $\mathrm{Aut}(F)$-representation of Γ.

It turns out that this representation ρ_{x^*} characterises the group action up to equivalence, where, as in the classical case, an equivalence between two fibre bundles $E, E' \to X$ is a bijection $\psi : E \to E'$ commuting with the projections onto X and preserving the fibre structure. In the coming proposition, we take x^* to be the origin and Γ to act linearly (recall that this can always be done). We also assume that the structure group G is a Lie group. The idea of the proof is adapted from similar results in [Las82].

Proposition 1.12. *Let Γ be a finite group acting linearly and isometrically on $\mathbb{R}^n$, and let F be an orbifold with structure group $G \leq \mathrm{Aut}(F)$. Assume G is a finite-dimensional Lie group. Let $\rho, \eta : \mathbb{R}^n \times \Gamma \to G$ be maps for which the action*

$$\gamma \cdot (x, v) = (\gamma \cdot x, \rho_x(\gamma)v) \tag{1.3}$$

on $\mathbb{R}^n \times F$ is a Γ-action (and likewise for η). Then the following are equivalent:

- *The maps ρ and η define equivalent Γ-equivariant fibre bundles, in the sense that there is a map $h : \mathbb{R}^n \to G$ such that*

$$h_{\gamma \cdot x} = \eta_x(\gamma) h_x \rho_x(\gamma)^{-1}. \tag{1.4}$$

  for every $\gamma \in \Gamma$.

- *The maps ρ and η define equivalent Γ-representations in G at the point $0 \in \mathbb{R}^n$, in the sense that there is some $h_0 \in G$ such that*

$$\eta_0(\gamma) = h_0 \rho_0(\gamma) h_0^{-1}. \tag{1.5}$$

  for every $\gamma \in \Gamma$.

Proof. If $P := \mathbb{R}^n \times G$ is equipped with the Γ-action

$$\gamma \cdot (x, g) = (\gamma \cdot x, \eta_x(\gamma) g \rho_x(\gamma)^{-1}), \tag{1.6}$$

then a smooth map $h : \mathbb{R}^n \to G$ defining a Γ-equivariant equivalence of representations is the same as a smooth Γ-invariant section $\psi \in \Gamma(P)$ (the section is given by $\psi(x) = (x, h_x)$). Thus, we wish to show that the existence of such a section is equivalent to the existence of a conjugation between the two maps at the identity.

It is easy to show that a Γ-invariant section of P gives a conjugation between the two representations ρ_0 and η_0. Since Γ acts linearly, the point $0 \in \mathbb{R}^n$ is a fixed point for the Γ-action on $\mathbb{R}^n$; it follows that the condition for $h : \mathbb{R}^n \to G$ to be Γ-invariant at 0 is exactly that h_0 conjugates between the two representations.

Conversely, assume that $h_0 \in G$ satisfies $h_0\rho_0(\gamma) = \eta_0(\gamma)h_0$. We will define a smooth Γ-invariant section of P around the point $0 \in \mathbb{R}^n$ using the Riemannian exponential map. Given a smooth Riemannian metric δ on P, the following is a Γ-invariant metric:

$$g^P := \sum_{\gamma \in \Gamma} \gamma^* \delta. \tag{1.7}$$

This allows us to define an exponential map $\exp_p : T_pP \to P$ for any $p \in P$, which takes the zero tangent vector to p, and takes a neighbourhood U_p of zero diffeomorphically to a neighbourhood of p. Moreover, the Riemannian exponential is Γ-equivariant in the following sense: where $L_\gamma : P \to P$ denotes the action of $\gamma \in \Gamma$, we have that

$$\gamma \cdot \exp_p = \exp_{\gamma \cdot p} \circ \mathrm{d}(L_\gamma)_p. \tag{1.8}$$

(This follows from the fact that g^P is Γ-invariant, meaning Γ acts isometrically on P.)

We choose the point p to be $(0, h_0)$, and we shrink the neighbourhood U_p so that it is Γ-invariant. Note that p is fixed under the Γ-action since $\eta_0(\gamma)h_0\rho_0(\gamma)^{-1} = h_0$ and $\gamma \cdot 0 = 0$; thus, the above equation for the equivariance of $\exp_p$ shows that $\exp_p(U_p)$ is also Γ-invariant as a neighbourhood of $p \in P$. Furthermore, consider the subspace $V_p \subseteq U_p$ consisting of tangent vectors to $\mathbb{R}^n \times G$ pointing entirely in the $\mathbb{R}^n$-direction. This is a Γ-invariant n-submanifold of U_p, and it is therefore taken to a Γ-invariant n-submanifold under $\exp_p$.

Since $\exp_p$ is smooth, we can shrink U_p (and by restriction V_p) so that V_p always intersects the fibres of $\mathbb{R}^n \times G$ transversally. It follows that the projection map $\pi|_{V_p} : V_p \to \mathbb{R}^n$ is injective, and its image is open by the constant rank theorem. Thus, we define a section $h : \pi(V_p) \to \pi(V_p) \times G$ as follows: $h(x)$ is the unique element of V_p which is also contained in the fibre $\{x\} \times G$. We have just shown that it is well-defined, and it is smooth and Γ-invariant by the smoothness and Γ-equivariance of $\exp_p$ and the Γ-invariance of V_p. □

Corollary 1.13. *Let x be a point in an orbifold X with local group Γ_x, and let F be a fibre for an orbifold fibre bundle $E \to X$ with finite-dimensional structure group G. The local behaviour of E around X is completely determined up to equivalence by a single representation $\rho_x : \Gamma_x \to G$. In particular, every $x \in X$ admits an orbifold chart $(\widetilde{U}, \Gamma_x, \varphi)$ over which $E|_U$ is equivalent to $(\widetilde{U} \times F)/\Gamma_x$, where Γ_x acts on F according to the representation ρ_x.*

Henceforth, we call an equivalence of the kind in the above corollary a *local trivialisation,* even though the bundle may be topologically nontrivial. Note that we recover the statement that all fibre bundles are locally product bundles if every point in X is regular.

1.2.1 Transition functions and local constructions

It is desirable to have an analogue of the fibre bundle construction lemma for orbifolds (see, for instance, [Hus94]). However, the above theorem demonstrates an added complication: while classical fibre bundles always have the same local behaviour, orbifold fibre bundles come with the additional structure of a representation of Γ_x at each point. In order to define a transition functions between two charts, therefore, we will need to enforce some kind of Γ_x-equivariance; this is only possible if one chart embeds into another, since this induces an inclusion of local groups which are otherwise incomparable. We therefore have the following analogous result.

Proposition 1.14 (Fibre bundle construction lemma, existence). *Let X be an orbifold with orbifold atlas $\{(U_i, \Gamma_i, \varphi_i)\}_{i \in I}$, and let F be an orbifold with the Lie group G as its structure group. Then an orbifold vector bundle with fibre F is specified by the following data:*

- *For each $i \in I$ there is a representation $\rho_i : \Gamma_i \to G$, which induces an action on $\widetilde{U}_i \times F$ by taking $\gamma \cdot (x, v) = (\gamma \cdot v, \rho_i(\gamma) \cdot v)$.*
- *For each embedding $\lambda_{ji} : \widetilde{U}_i \to \widetilde{U}_j$ there is a Γ_i-equivariant fibrewise linear map $\alpha_{ji} : \widetilde{U}_i \times F \to \widetilde{U}_j \times F$ for which $\mathrm{pr}_1 \circ \alpha_{ji} = \lambda_{ji} \circ \mathrm{pr}_1$ (i.e., α_{ji} is a lift of the embedding λ_{ji}).*

Additionally, the embedding lifts are required to satisfy the orbifold cocycle condition: for a sequence of embeddings making the following diagram commute:

$$\widetilde{U}_i \xrightarrow{\lambda_{ji}} \widetilde{U}_j \xrightarrow{\lambda_{kj}} \widetilde{U}_k, \qquad \lambda_{ki} : \widetilde{U}_i \to \widetilde{U}_k$$

the induced bundle maps make the following diagram commute:

$$\widetilde{U}_i \times F \xrightarrow{\alpha_{ji}} \widetilde{U}_j \times F \xrightarrow{\alpha_{kj}} \widetilde{U}_k \times F. \qquad \alpha_{ki} : \widetilde{U}_i \times F \to \widetilde{U}_k \times F$$

Remark. The lifts $\alpha_{ji} : \widetilde{U}_i \times F \to \widetilde{U}_j \times F$ can also be thought of as maps $\psi_{ji} : \widetilde{U}_i \to G$; the relationship between the two is as follows:

$$\alpha_{ji}(x, v) = (\lambda_{ji}(x), \psi_{ji}(x) \cdot v). \tag{1.9}$$

If α_{ji} is required to be Γ_i-equivariant, then ψ_{ji} is required to be equivariant under the action of Γ_i by conjugation:

$$\psi_{ji}(\gamma \cdot x) = \rho_j(\gamma)\psi_{ji}(x)\rho_i(\gamma)^{-1}. \tag{1.10}$$

We will call a collection of such ψ_{ji} a set of *transition functions.* In particular, if $x \in \widetilde{U}_i$ is a fixed point of Γ_i, the representations ρ_i and $\rho_j|_{\Gamma_i}$ are conjugate under the element $\psi_{ji}(x)$. We can therefore assume without loss of generality that $\rho_j|_{\Gamma_i} = \rho_i$, and that $\psi_{ji}(x)$ is the identity at a fixed point.

Proof. We start by showing that this data can be retrieved from an existing orbifold vector bundle. As proved in Corollary 1.13, an orbifold chart on the base space of an orbifold vector bundle induces a representation of the local group $\rho_i : \Gamma_i \to G$. Furthermore, an embedding of orbifold charts $\lambda_{ji} : \widetilde{U}_i \to \widetilde{U}_j$ induces the Γ_i-equivariant map α_{ji} if we take $\alpha_{ji}(x, v) = (\lambda_{ji}(x), v)$. (Note that each fibre of $\widetilde{U}_i \times F$ has an induced action by Γ_j, and therefore also by Γ_i since $\Gamma_i \leq \Gamma_j$.) These induced maps clearly satisfy the cocycle condition.

Conversely, suppose all of the data above has been given. We proceed as in the nonsingular case: we use the data to define an equivalence relation on the union of charts, and quotient by this equivalence relation. If $[(x_i, v_i)] \in (\widetilde{U}_i \times F)/\Gamma_i$ and $[(x_j, v_j)] \in (\widetilde{U}_j \times F)/\Gamma_j$, and there is an embedding $\lambda_{ji} : \widetilde{U}_i \to \widetilde{U}_j$, we say that $[(x_i, v_i)] \sim [(x_j, v_j)]$ if $[\alpha_{ji}(x_i, v_i)] = [(x_j, v_j)]$ (where square brackets denote equivalence classes under the group action). The Γ_i-equivariance of α_{ji} ensures that this relation does not depend on the representative chosen. Note that this relation implies that $\lambda_{ji}(x_i)$ is in the Γ_j-orbit of x_j. We have a reflexive binary relation on the space

$$\bigsqcup_i (\widetilde{U}_i \times F)/\Gamma_i, \tag{1.11}$$

and we take the symmetric and transitive closure of this relation to get an equivalence relation. We denote by E the quotient of this space by the equivalence relation, and we define the map $p : E \to X$ as follows:

$$p([(x_i, v_i)]) = \varphi_i(x_i). \tag{1.12}$$

If an equivalent representative (x_j, v_j) is chosen, we know that $\varphi_j(x_j) = \varphi_i(x_i)$ by the equivariance of each φ, as well as the fact that $\lambda_{ji}(x_i)$ is in the same Γ_j orbit as x_j. Hence, p is well-defined. The map p also clearly admits a smooth local lift, since it is built to respect the action of Γ_i on each fibre. The local triviality condition is essentially satisfied by definition. □

Note that the notion of an orbifold vector bundle is recovered by taking F to be a vector space and $G = \text{Aut}(F)$ to be the collection of linear automorphisms of F. The resulting Γ_i-actions are linear by definition.

Proposition 1.15 (Fibre bundle construction lemma, equivalence)**.** *Let X, $\{(U_i, \Gamma_i, \varphi_i)\}_{i \in I}$, F, G be as above. Let $\rho_i, \eta_i \in \text{Hom}(\Gamma_i, G)$ be a pair of representations for each Γ_i, and let $\phi_{ji} : \widetilde{U}_i \to G$ be a set of transition functions with respect to ρ_i and $\psi_{ji} : \widetilde{U}_i \to G$ a set of transition functions with respect to η_i. The fibre bundle defined by the two sets of transition functions are equivalent if and only if there is a collection of smooth maps $h_i : \widetilde{U}_i \to G$ satisfying the following two properties:*

- *The h_i conjugates between the transition functions ϕ_{ji} and ψ_{ji} in the following sense:*

$$\psi_{ji}(x) = h_j(x)\phi_{ji}(x)h_i(x)^{-1}. \tag{1.13}$$

- *The h_i are Γ_i-equivariant under the conjugation action:*

$$h_i(x) = \eta_i(\gamma)h_i(x)\rho_i(\gamma)^{-1}. \tag{1.14}$$

Remark. If $x \in \widetilde{U}_i$ is a Γ_i-fixed point, then the Γ_i-equivariance condition on h_i implies that $h_i(x)$ conjugates between the representations η_i and ρ_i. Thus, two fibre bundles can only be equivalent if the representations of their local groups in G are equivalent.

Proof. If the two fibre bundles defined by the transition functions are equivalent, let $\mu_i, \nu_i : \pi^{-1}(U_i) \to (\widetilde{U}_i \times F)/\Gamma$ be the trivialisations associated with ϕ_{ji} and ψ_{ji} respectively. Then we have that $\nu_i \circ \mu_i^{-1}$ lifts to an equivariant map taking $\widetilde{U}_i \times F$ to itself, and it can be written in the form $(x, v) \mapsto (x, h_i(x)v)$ for Γ_i-equivariant functions $h_i : \widetilde{U}_i \to G$. These are the desired functions for the proposition; using the approach in [Hus94] (Proposition 2.5), we see that they conjugate between the two transition functions.

Conversely, suppose the collection of functions h_i exists. Define a collection of smooth Γ_i-equivariant maps $f_i : \widetilde{U}_i \times F \to \widetilde{U}_i \times F$ so that $f_i(x,v) = (x, h_i(x)v)$. We claim that the f_i glue together to form an equivalence f between the fibre bundles induced by ϕ_{ji} and ψ_{ji}. It is clear that the f_i are fibrewise isomorphisms of F and, if they do glue together, the equivalence locally looks like a Γ_i-equivariant isomorphism from $\widetilde{U}_i \times F$ with the ρ_i-action to $\widetilde{U}_i \times F$ with the η_i-action. Thus, we only need to show that f is well-defined, meaning that if $(x,v) \in \widetilde{U}_i \times F$ and $(y,w) \in \widetilde{U}_i \times F$ are equivalent, then their image under f_i and f_j is also equivalent (in the sense of Proposition 1.14).

If (x,v) and (y,w) are equivalent, we assume without loss of generality that $U_i \hookrightarrow U_j$; this means that $(y,w) = (\lambda_{ji}(x), \phi_{ji}(x)v)$. But then

$$f_i(x,v) = (x, h_i(x)v),$$

and

$$f_j(y,w) = (y, h_j(y)w) = (\lambda_{ji}(x), h_j(y)\phi_{ji}(x)v) = (\lambda_{ji}(x), \psi_{ji}(x)h_i(x)v).$$

But this is precisely the action of $\alpha_{ji} = (\lambda_{ji}, \psi_{ji})$ on $(x, h_i(x)v) = f_i(x,v)$, meaning $f_j(y,w)$ and $f_i(x,v)$ are equivalent. It follows that f_i and f_j glue across the quotient to form a well-defined equivalence f between the fibre bundles. □

We recover the classical theorem by taking each representation to be trivial and each α_{ji} to be an inclusion; this clearly satisfies the conditions of the theorem. This also demonstrates that an orbifold fibre bundle over an orbifold X reduces to a regular fibre bundle over X_{reg}.

The sections of a fibre bundle can be interpreted in this light as follows:

Proposition 1.16. *Using the above notation, let $E \to X$ be an orbifold vector bundle with fibre F, and let $\psi_{ji} : \widetilde{U}_i \to G$ and $\rho_i : \Gamma_i \to G$ constitute a set of transition functions compatible with E. The smooth sections of E are in one-to-one correspondence with the collections of Γ_i-equivariant smooth maps $\sigma_i : \widetilde{U}_i \to F$ satisfying the following condition:*

$$\sigma_j(\lambda_{ji}(x)) = \psi_{ji}(x)\sigma_i(x). \tag{1.15}$$

Proof. The result follows immediately from the observation that, in order for $(x, \sigma_i(x))$ and $(\lambda_{ji}(x), \sigma_j(\lambda_{ji}(x)))$ to be equivalent, the second component must be related by a factor of $\psi_{ji}(x)$. □

We therefore obtain the following local picture of orbifold fibre bundles with finite-dimensional structure group:

- Each orbifold fibre bundle is locally isomorphic to $(\mathbb{R}^n \times F)/\Gamma$, where Γ is a finite group acting linearly on $\mathbb{R}^n$ and as a representation on F. These local charts are isomorphic to one another if and only if the representations of Γ are conjugate.

- Given an orbifold atlas $\{(U_i, \Gamma_i, \varphi_i)\}$, an orbifold fibre bundle can be specified with representations $\rho_i : \Gamma_i \to G$ for each i, and transition functions $\psi_{ji} : \widetilde{U}_i \to G$ for every chart embedding $\widetilde{U}_i \hookrightarrow \widetilde{U}_j$. The representations can be chosen so that $\rho_j|_{\Gamma_i} = \rho_i$ for any embedding, and the transition functions are required to be Γ_i-invariant under the conjugation action and satisfy the cocycle condition.

- Two orbifold bundles can only be isomorphic if the representation of the local group at each point is conjugate; we can therefore always choose trivialisations where the local group representations $\rho_i : \Gamma_i \to G$ are the same. Given two systems of transition functions ϕ_{ji} and ψ_{ji}, they define equivalent bundles if and only if there is a smooth Γ_i-equivariant function conjugating between them for every i.

Henceforth, we will freely move between the global description of an orbifold (as an orbifold projection $\pi : E \to X$) and the local description in terms of transition functions $\psi_{ji} : U_i \to G$ and representations ρ_i.

1.2.2 Examples of orbifold bundles

We now present some key examples of orbifold bundles, as well as some of their associated structures. All of the examples we shall present have natural analogues in the non-singular category.

Firstly, from orbifold vector bundles $E, F \to X$, we can always construct new vector bundles corresponding to linear algebraic operations on each fibre: there is the *dual bundle* E^*, the *direct sum* $E \oplus F$, the *tensor product* $E \otimes F$, and the k*th exterior power* $\Lambda^k E$. Note that elements of the dual bundle can be naturally interpreted in a local trivialisation over $(\widetilde{U}_i, \varphi_i, \Gamma_i)$ as Γ_i-equivariant linear functionals on each fibre.

We will be dealing extensively with complex line bundles, i.e., rank 1 complex vector bundles. The space of such bundles is closed under the tensor product operation; furthermore, the trivial bundle serves as an identity for the operation, and the dual bundle corresponds to the inverse. We have the following definition:

Definition 1.17. Let X be an orbifold. The collection of isomorphism classes of complex line bundles over X, together with group operation given by the tensor product, is called the *topological Picard group.* It is denoted by $\mathrm{Pic}^t(X)$.

We may now also define the tangent bundle on an orbifold X; given an orbifold chart $(\widetilde{U}, \Gamma, \varphi)$ on X, there is a natural extension of the action of Γ on $\widetilde{U}$ to $T\widetilde{U} \cong \widetilde{U} \times \mathbb{R}^n$ given by differentiating the action. An embedding of charts naturally lifts to an equivariant embedding of tangent bundles, since an inclusion is linear and therefore differentiates to the inclusion itself. We therefore have a system of transition functions, and the associated bundle TX is called the *orbifold tangent bundle.*

Remark. Depending on whether we take the Γ_x-quotient, the tangent space is referred to by different names: the fibre $T_y\widetilde{U}$ for $y \in \varphi^{-1}(x)$ is called the *tangent space at x,* whereas $p^{-1}(x) \cong T_y\widetilde{U}/\Gamma_x$ is called the *tangent cone at x.* An element of the tangent cone is what we refer to as a tangent vector. Note that a tangent vector field on X will be naturally valued in the tangent cone at each point.

Just as we have constructed the tangent bundle on a manifold, we may also construct tensor bundles and exterior bundles, resulting in the notion of tensors and differential forms on orbifolds. A rank k covariant tensor takes k tangent vectors and returns a number, and it follows that, over a chart $(\widetilde{U}, \Gamma, \varphi)$, we can regard such a tensor as a Γ-equivariant map from $(T\widetilde{U})^k$ to $\mathbb{R}$. Since tangent vectors and smooth functions both have locally equivariant lifts, we can also define the exterior derivative of a smooth function as $\mathrm{d}f(v) = v(f)$ in local charts. We therefore get extensions of several classical geometric notions:

- A complex orbifold comes with a natural almost-complex structure $J : TX \to TX$, since a complex chart $(\widetilde{U}, \Gamma, \varphi)$ requires that the action of Γ on $\widetilde{U}$ is complex-linear, meaning multiplication by i on $\widetilde{U}$ can be transported to TX by φ. It is easy to verify that this definition does not depend on the coordinate chart. As in the manifold case, we obtain splittings of the complexified k-forms $\Omega^k(X)$ into (p,q)-forms $\Omega^{p,q}(X)$ with $p+q=k$.

- A Riemannian metric g still induces a unique connection ∇^g on TX, the Levi-Civita connection, as well as a unique parallel volume form vol_g given by the wedge product of an orthonormal coframe at any point, and an operator $* : \Omega^k(X) \to \Omega^{n-k}(X)$ for which $\alpha \wedge *\beta = \langle \alpha, \beta \rangle \mathrm{vol}_g$. If we additionally have a complex structure and the 2-form $(v,w) \mapsto g(v, Jw)$ is closed, we have a Kähler structure on g; the Hodge star can be extended so that $* : \Omega^{p,q}(X) \to \Omega^{n-p,n-q}(X)$. We assume that the Hodge star is complex antilinear throughout the thesis.

- We also get the various correspondences for complex vector bundles over complex orbifolds: a holomorphic structure on a vector bundle $E \to X$ is still uniquely specified by a map $\overline{\partial}_E : \Gamma(E) \to \Omega^1(E)$ satisfying the Leibniz rule, and over a Hermitian vector bundle $(E,h) \to X$, we can naturally identify holomorphic structures $\overline{\partial}_E$ with Hermitian connections A by taking $\overline{\partial}_E = (\nabla^A)^{0,1}$ and $\nabla^A = \overline{\partial}_E + h^{-1} \circ \overline{\partial}_E \circ h$.

The notion of integration calls for a slight modification. To define integration of a compactly supported n-form ω over an n-orbifold X, we assume (as in the nonsingular case) that X is orientable, meaning it can be given an oriented atlas $\{\widetilde{U}_i, \Gamma_i, \varphi_i\}$. Given a smooth partition of unity $\{\psi_i : U_i \to [0,1]\}$ subordinate to the open cover $\{U_i\}$, we define the *integral of ω over X* as follows:

$$\int_X \omega = \sum_i \frac{1}{|\Gamma_i|} \int_{\widetilde{U}_i} \varphi_i^*(\psi_i\, \omega). \tag{1.16}$$

It is routine to verify that this integral does not depend on the choice of charts or partition of unity.

Tools from functional analysis also transfer to the orbifold setting; in particular, the definitions of Sobolev spaces transfer directly to the orbifold setting. Given a rank n orbifold vector bundle $E \to X$, we define the Sobolev p-norm on X to be the following:

$$\|\sigma\|_{L^p_k} = \sum_{j=0}^{k} \|\nabla^j \sigma\|_{L^p} = \sum_{j=0}^{k} \left(\int_X |\nabla^j \sigma|^p \,\mathrm{vol}_g \right)^{1/p}. \tag{1.17}$$

Taking the completion with respect to this norm, we obtain the space of L^p_k sections of E, denoted by $\Gamma(E)_{L^p_k}$. By specialising E to various bundles, we obtain Sobolev vector-valued functions, vector fields, k-forms, and tensors. Additionally, since the space of connections on an orbifold vector bundle E is an affine space modelled on $\Omega^1(E)$, we also get the notion of a Sobolev connection. Finally, if Y is a manifold, we can embed Y into a Euclidean space $\mathbb{R}^N$ for some N by Whitney's embedding theorem; we can therefore define the space of Sobolev functions $L^p_k(X,Y)$ to be the subspace of $L^p_k(X,\mathbb{R}^N)$ for which the image of the function is contained in Y (this definition only works when $kp > n$). This also allows for a notion of Sobolev sections of any fibre bundle whose fibre is a manifold, by reducing such sections to local Sobolev functions.

In order to make these constructions precise, several technical details need to be resolved. We refer the reader to the Appendix for a more precise treatment of this topic.

1.2.3 Gauge theory on orbifolds

With the notion of orbifold fibre bundles, we may begin the study of gauge theory on orbifolds. Once again, most of the structures from gauge theory on manifolds transfer to the orbifold setting, with only minor technical ramifications.

We begin with the notion of a principal bundle.

Definition 1.18. Let G be a compact Lie group, and let X be an orbifold. A *principal G-bundle over X* consists of an orbifold G-bundle $P \to X$ equipped with a smooth *right* action by G preserving the fibres of P, satisfying the following condition: for every $x \in X$ and every corresponding fibre P_x, there is a homomorphism $\alpha : \Gamma_x \to G$ such that the G-action on P_x is transitive with isotropy group (conjugate to) $\alpha(\Gamma_x)$.

The essence of this definition is that G acts freely and transitively on fibres over regular points, but it is only G/Γ that acts freely and transitively over a point with nontrivial local group Γ. Note that this quotient may still be all of G if Γ is chosen to act trivially on G. In general, the local picture of a principal G-bundle is a Γ_x-equivariant G-bundle over an open subset of $\mathbb{R}^n$, and we will make a point of phrasing global definitions in this language when possible.

Example 1.19. Let TX be the tangent bundle of an n-orbifold X. Just as in the nonsingular case, we can define the *frame bundle of* X to be the principal $\mathrm{GL}(n)$-bundle $\pi : FX \to X$ defined as follows: given a local orbifold chart $(\widetilde{U}, \Gamma_x, \varphi)$ around x, the fibre of FX over x is defined to be the collection of all equivalence classes of frames $\{e_1, \ldots, e_n\}$ for $T_x\widetilde{U}$, where two frames are considered equivalent if they are in the same Γ_x-orbit (this does not depend on the chart). Because of the equivalence relation, the fibre over x is acted on transitively but not in general freely by $\mathrm{GL}(n)$; the action of Γ_x on $T_x\widetilde{U}$ is faithful, so there is an inclusion $\Gamma_x \hookrightarrow \mathrm{GL}(n)$. Essentially by definition, the isotropy subgroup of a given frame is Γ_x.

Additionally, if the orbifold X has extra structure, one may define reduced frame bundles. For instance, if X has a Riemannian metric, the frames in the above construction can be chosen to be orthonormal.

Definition 1.20. Given a principal G-bundle $\pi : P \to X$, a *gauge transformation* is a smooth G-equivariant map $\psi : P \to P$ preserving the fibres of P.

Remark. An equivalent definition for a gauge transformation is a smooth map $h : P \to G$ which is equivariant when G is equipped with the conjugation action on itself; that is, for every $p \in P$ and $g \in G$,

$$h(p \cdot g) = g^{-1} h(p) g. \tag{1.18}$$

The associated automorphism of P is $p \mapsto p \cdot h_p$. Conversely, given an automorphism $\psi : P \to P$ and a nonsingular point p, there is a unique $h(p) \in G$ for which $\psi(p) = p \cdot h(p)$ (since G acts freely and transitively on $\pi^{-1}(p)$); by using smooth local lifts, the smoothness of ψ allows $h(p)$ to be smoothly defined even when p is singular.

The equivalent definition in the remark is much more useful for defining gauge transformations locally:

Definition 1.21. Given a principal G-bundle $P \to X$ and a trivialising orbifold chart $(\widetilde{U}, \Gamma, \varphi)$ with local Γ-action specified by the representation $\rho : \Gamma \to G$, a *local gauge transformation over* U is a map $h : \widetilde{U} \to G$ which is Γ-equivariant in the following sense:

$$h(\gamma \cdot x) = \rho(\gamma) h(x) \rho(\gamma)^{-1}. \tag{1.19}$$

There is an associated automorphism of $P|_U \cong (\widetilde{U} \times G)/\Gamma$ given by taking $(x, g) \mapsto (x, h(x)g)$.

One of the reasons it is convenient to use the language of principal bundles is the associated bundle construction. This construction makes precise the relationship between similar bundles, such as the various tensor bundles and the bundles of differential forms.

Definition 1.22. Let $\pi : P \to X$ be a principal G-bundle, let F be an orbifold, and let $\alpha : G \to \mathrm{Aut}(F)$ denote a left G-action on F. Equip F with the *right* G-action given by $v \mapsto \alpha(g^{-1})v$ for $v \in F$. The *associated bundle* $P \times_\alpha F$ is defined to be the following quotient:

$$P \times_\alpha F = (P \times F)/G, \tag{1.20}$$

where G acts diagonally on $P \times F$. It is straightforward to verify that this is an orbifold vector bundle over X with fibre F.

Example 1.23. Given the frame bundle $FX \to X$ with structure group $\mathrm{GL}(n)$, one may take the tautological representation of $\mathrm{GL}(n)$ on $\mathbb{R}^n$, the dual representation, direct sum representations, tensor product representations, and exterior power representations. Under the associated bundle construction, this leads respectively to the tangent bundle, the cotangent bundle, direct sums of bundles, tensor products of bundles, and bundles of differential forms.

Example 1.24. Particularly important for the vortex equations is the interplay between complex line bundles and $\mathrm{U}(1)$-bundles. Given a complex line bundle $L \to X$ with Hermitian metric h, one can form a principal $\mathrm{U}(1)$-bundle $F_{\mathrm{U}} \to X$ by taking the frame bundle construction on L but with single unit-length vectors instead. Conversely, given a $\mathrm{U}(1)$-bundle $P \to X$, one can get an associated complex line bundle $P \times_{\mathrm{id}} \mathbb{C}$ by taking the tautological representation of $\mathrm{U}(1)$ in $\mathbb{C}$.

Example 1.25. Denoting by $c : G \to \mathrm{Aut}(G)$ the map $c(g)h = ghg^{-1}$, the space of all gauge transformations $h : P \to G$ can be naturally identified with the sections of the associated bundle $P \times_c G$. A gauge transformation h defines a section given by $x \mapsto [p, h(p)]$ for any $p \in \pi^{-1}(x)$, which is independent of the p chosen, and this correspondence is bijective. Note that, when c acts trivially on G, the associated bundle reduces to $X \times G$, meaning the space of gauge transformations is just $C^\infty(X, G)$. This occurs whenever G is abelian.

Recall the following from Proposition 1.15: given two sets of transition functions $\psi_{ji}, \phi_{ji} : \widetilde{U}_i \to G$ compatible with the same representations $\rho_i : \Gamma_i \to G$, their respective fibre bundles are equivalent if and only if there exist Γ_i-equivariant functions $h_i : \widetilde{U}_i \to G$ which conjugates between ψ_{ji} and ϕ_{ji}. By Definition 1.21, it follows that the h_i can be naturally interpreted as local gauge transformations. In Section 2.4.2, we will need to modify a series of local gauge transformations so that they "glue together" to form a global gauge transformation;

by this we mean that they can all be interpreted as the local representations of a map $h : P \to G$ defined on all of P. It turns out that this can be interpreted in terms of transition functions.

Proposition 1.26. *Let $P \to X$ be a principal G-bundle with transition functions $\{\phi_{ji}\}$, and let $g_i : \widetilde{U}_i \to G$ be a series of local gauge transformations. There is a corresponding global gauge transformation $g : P \to G$ which matches g_i on the orbifold charts U_i if and only if the following holds:*

$$g_{ji} := g_j \phi_{ji} g_i^{-1} = \phi_{ji}. \tag{1.21}$$

Proof. As noted above, gauge transformations correspond bijectively with sections of $P \times_c G$. If the transition functions for P are given by multiplying by ϕ_{ji}, then the transition functions for $P \times_c G$ are given by conjugating by ϕ_{ji}. Applying Proposition 1.16 finishes the proof. □

This leads to the following result. If a system of local gauge transformations g_i is modified by multiplication by another system h_i (so $g_i \mapsto h_i g_i$), then the modified g_i glue to form a global gauge transformation if and only if h_i constitutes an equivalence between ϕ_{ji} and g_{ji}, in the sense of the fibre bundle construction lemma (Proposition 1.15).

1.2.4 Pullback bundles

Though we have described many of the structures one would wish to transfer from manifolds to orbifolds, we have not yet defined pullback of bundles. In order to do this, we need to give our maps slightly more structure. The smooth maps along which pullback is allowed will be called good maps. For a more comprehensive overview of pullback orbifold bundles, refer to [CR01].

Definition 1.27. Let X be an orbifold. A *compatible collection of charts on X* is an orbifold atlas $\{(\widetilde{U}_i, \Gamma_i, \varphi_i)\}_{i\in I}$ on a *subset* of X, such that any inclusion $U_i \hookrightarrow U_j$ has a corresponding embedding of orbifold charts. If the collection of charts covers X, then it is called a *compatible cover.*

Definition 1.28. Let $f : X \to Y$ be a smooth map between orbifolds. A *compatible system for f* consists of the following data:

- A compatible cover $\{(\widetilde{U}_i, \Gamma_i, \varphi_i)\}_{i\in I}$ on X, and a corresponding compatible collection of charts $\{(\widetilde{V}_i, \Lambda_i, \psi_i)\}_{i\in I}$ on Y, such that $f(U_i) \subseteq V_i$ for all i and $U_i \subseteq U_j$ implies $V_i \subseteq V_j$ for all i, j;
- A collection of smooth local lifts $\widetilde{f}_i : \widetilde{U}_i \to \widetilde{V}_i$ over each chart $\widetilde{U}_i$ (cf. Definition 1.9);
- For each chart embedding $\lambda_{ji} : \widetilde{U}_i \to \widetilde{U}_j$ a corresponding chart embedding $\mu(\lambda_{ji}) : \widetilde{V}_i \to \widetilde{V}_j$, such that
$$\mu(\lambda_{ji}) \circ \widetilde{f}_i = \widetilde{f}_j \circ \lambda_{ji}, \tag{1.22}$$
and
$$\mu(\lambda_{kj} \circ \lambda_{ji}) = \mu(\lambda_{kj}) \circ \mu(\lambda_{ji}). \tag{1.23}$$

We denote the compatible system by $\mathscr{S} = \{U_i, \Gamma_i, V_i, \Lambda_i, \widetilde{f}_i, \mu_{ji}\}$. If a smooth map admits a compatible system, then it is called a *good map.*

It turns out that there is a well-defined notion a pullback bundle along a good map.

Theorem 1.29. *Let X and Y be orbifolds, and let $f : X \to Y$ be a smooth map admitting a compatible system $\mathscr{S} = \{U_i, \Gamma_i, V_i, \Lambda_i, \widetilde{f}_i, \mu_{ji}\}$. If there is an orbifold F-bundle $\pi : E \to Y$, then there is a canonically constructed bundle $\pi' : f^*E_{\mathscr{S}} \to X$ together with a smooth map $f^* : f^*E_{\mathscr{S}} \to E$ lifting f. We call this bundle a* pullback of E along f.

Proof. We assume that f is surjective; if not, we may simply redefine Y to be the union of the subsets V_i.

Let $(\widetilde{V}_i, \Lambda_i, \psi_i)$ be a chart on Y in the compatible system $\mathscr{S}$. We can take the pullback of the bundle $E|_{V_i} \cong (\widetilde{V}_i \times F)/\Lambda_i$ along $\widetilde{f}_i$ to get a Γ_i-equivariant bundle $\widetilde{f}_i^*(\widetilde{V}_i \times F)$ over $\widetilde{U}_i$. Additionally, since each $\widetilde{f}_i$ is a smooth local lift, there is a natural Γ_i-action on the pullback bundle given by the induced homomorphism $\Gamma_i \to \Lambda_i$.

Now, suppose we also have a system of transition maps $\alpha_{ji} : \widetilde{V}_i \times F \to \widetilde{V}_j \times F$ characterising $E \to Y$ (as in Lemma 1.14). We obtain a natural system of transition maps $\beta_{ji} : \widetilde{f}_i^*(\widetilde{V}_i \times F) \to \widetilde{f}_j^*(\widetilde{V}_j \times F)$ by simply precomposing with $\widetilde{f}_i$, i.e., $\beta_{ji} = \alpha_{ji} \circ \widetilde{f}_i$. Since the α_{ji} are lifts of the μ_{ji}, which in turn correspond to the λ_{ji}, we know that the β_{ji} also satisfy the cocycle condition. As such, we obtain a bundle $f^*E_{\mathscr{S}}$ over X which locally looks like the pullback of E, and there is clearly a map from this bundle to E given by the usual induced pullback map. □

Note that a good map may have several inequivalent compatible structures, and there is nothing compelling their pullbacks to be isomorphic. However, there is an important class of smooth maps for which we can expect exactly one pullback bundle.

Definition 1.30. A smooth map $f : X \to Y$ is called *regular* if $f^{-1}(Y_{\text{reg}})$ is open, dense, and connected in X.

Theorem 1.31. *Every regular map between orbifolds is a good map, and any two compatible systems for a regular map induce naturally isomorphic pullback bundles.*

As such, for a regular map $f : X \to Y$ and a bundle $E \to Y$, we can essentially speak of "the" pullback bundle along f. All of the maps we define pullbacks along will naturally be regular; in particular, the projection map $\pi : E \to X$ and any section $\sigma : X \to E$ are both always regular.

1.2.5 Algebraic topology on orbifolds

The results we have thus far on the structure of orbifold fibre bundles are primarily local. In order to understand and classify the global structure of these bundles, we will need orbifold analogues of several algebraic topological invariants. Thus, we make a brief note on algebraic invariants for orbifolds.

We begin with the notion of the fundamental group. A naive definition for the fundamental group of an orbifold X is $\pi_1(|X|)$, the space of homotopy classes of continuous loops on the underlying topological space. However, this definition fails to capture the relationship between the fundamental group and covering spaces: the space $\mathbb{R}^2/\mathbb{Z}_\alpha$ would have trivial fundamental group, but we would like to say that $\mathbb{R}^2$ is its α-fold "universal cover". The following modification is due to Thurston [Thu02]:

Definition 1.32. Let X be an orbifold. An *orbifold loop in* X is a loop in $|X|$ equipped with a partition into finitely many components, and a lift of each component into an orbifold chart. Two orbifold loops are equivalent if their partitioned components can be subdivided and mapped diffeomorphically onto each other by changes of charts, and a homotopy of orbifold loops is a homotopy between the components of one loop into another such that each slice of the homotopy defines an orbifold loop. The *orbifold fundamental group of* X is defined to be the space of equivalence classes of orbifold loops on X, up to subdivision, equivalence, and homotopy. It is denoted by $\pi_1(X)^{\mathrm{orb}}$.

Additionally, there is the following result:

Theorem 1.33. *Every orbifold* X *has a universal covering* $\widetilde{X} \to X$ *for which* $\pi_1(\widetilde{X})^{\mathrm{orb}} = 1$.

The definition of the fundamental is essentially designed so that $\pi_1(X)^{\mathrm{orb}}$ acts on the universal cover by deck transformations. As an example, the orbifold fundamental group of $\mathbb{R}^n/\Gamma$ is always isomorphic to Γ for any finite group Γ acting effectively on $\mathbb{R}^n$. In the case that $X = Y/\Lambda$ is a global quotient of a smooth space by a finite group, there is a relationship between the orbifold fundamental group of X and the fundamental group of Y [ALR07]:

Proposition 1.34. *If* Y *is smooth then the orbifold fundamental group of* $X = Y/\Lambda$ *is a* $\pi_1(Y)$*-extension of* Λ*, i.e., the following sequence is exact:*

$$0 \to \pi_1(Y) \to \pi_1(X)^{\mathrm{orb}} \to \Lambda \to 0. \tag{1.24}$$

For more details on the definition and covering space theory, refer to [Jr22].

Next, we consider the notion of cohomology on orbifolds. The construction of characteristic classes for vector bundles is one of the main reasons why cohomology is so useful in gauge theory; this is especially true when studying complex line bundles, for which the first Chern class $c_1 : \mathrm{Pic}^t(X) \to H^2(X;\mathbb{Z})$ is a complete algebraic invariant. However, the naive definition of cohomology on an orbifold X in terms of $|X|$ again fails: orbifold bundles generally have nontrivial behaviour around singularities, which clearly cannot be detected by $H^\bullet(|X|;\mathbb{Z})$.

The natural definition of orbifold cohomology makes use of the description of orbifolds as groupoids, a tangent we will not discuss further. We refer the reader to [ALR07] for more details on the definition of the orbifold cohomology $H^\bullet(X)^{\mathrm{orb}}$. In the case that an orbifold is a global quotient Y/Λ where Y is smooth and Λ is finite, the definition reduces to the following related cohomology theory:

Definition 1.35. Let Y be a smooth manifold acted upon by a finite group Λ. The Λ-*equivariant cohomology of* Y is defined to be the cohomology of the stable homotopy quotient of Y by Λ:

$$H^\bullet_\Lambda(Y) := H^\bullet((Y \times E\Lambda)/\Lambda), \tag{1.25}$$

where $E\Lambda$ is the universal principal Λ-bundle over the classifying space $B\Lambda$, and the Λ-action is diagonal.

Finally, we make a brief note on de Rham cohomology theory. In Satake's first paper on orbifolds [Sat56], the following theorem is proved:

Theorem 1.36. *For any orbifold X, the de Rham cohomology of X is isomorphic to the real-valued singular cohomology of $|X|$.*

Thus, the global behaviour of differential forms on orbifolds is unaffected by their singularities.

1.2.6 Flat complex line bundles on orbifolds

We can use the algebraic invariants we have thus far defined to determine the structure of the space of flat Hermitian line bundles on a given orbifold, that is, complex line bundles equipped with a fibrewise Hermitian inner product and a unitary connection. In the nonsingular case, each flat Hermitian line bundle corresponds uniquely to a U(1)-representation of the fundamental group; this representation comes from the holonomy of the connection, which is homotopy-invariant by flatness. The representation may be extended to the orbifold case as follows:

Proposition 1.37. *The gauge equivalence classes of flat Hermitian line bundles over an orbifold X are in one-to-one correspondence with $\mathrm{Hom}(\pi_1(X)^{\mathrm{orb}}, \mathrm{U}(1))$ under the holonomy representation.*

Proof. First, we need to verify that a flat Hermitian line bundle $(L, A) \to X$ defines a holonomy representation. Let $\gamma : [a, b] \to X$ be an orbifold path; we can cover γ by finitely many orbifold charts $\{\widetilde{U}_i, \Gamma_i, \varphi_i\}$, each of which has a specified smooth lifting of $\gamma|_{U_i}$ to $\widetilde{U}_i$ (since γ is an orbifold path). By subdividing, we may also choose these orbifold charts so that L is trivialised over them and A admits a smooth lift as well. We define the holonomy in the usual way on each chart, and patch them together to obtain the holonomy of the entire curve. It is straightforward to verify that two equivalent orbifold paths define equivalent holonomy. One can also verify (as in the manifold case) that the holonomy of an orbifold loop homotopic to the trivial orbifold loop is the integral of the curvature over a bounding surface, and therefore that a flat connection gives a well-defined holonomy representation on homotopy classes of loops. It follows that a flat bundle over X defines a holonomy representation of the orbifold fundamental group.

Conversely, given a representation $\rho : \pi_1(X)^{\mathrm{orb}} \to \mathrm{U}(1)$, we take the universal covering $\widetilde{X} \to X$, and define a Hermitian line bundle $\widetilde{X} \times_\rho \mathbb{C}$ via the associated bundle construction. The flat connection on this line bundle is defined so that its pullback to $\widetilde{X}$ is equal to the exterior derivative, and the fact that ρ is valued in U(1) gives the associated bundle a Hermitian structure.

Unitary gauge transformations do not affect the holonomy of a connection, and conversely, a connection is uniquely determined by its holonomy up to gauge equivalence (this has been proven for manifolds in [Lew93], and extends to orbifolds since it has been shown in [Lan20] that X_{reg} is connected, open, and dense in X). It follows that the two constructions are mutually inverse. □

Additionally, the following result allows us to restrict to trivial flat bundles if necessary.

Proposition 1.38. *The space of trivial flat vector bundles on an orbifold X is given by $H^1(|X|; \mathbb{R})/H^1(|X|; \mathbb{Z})$.*

Proof. The space of unitary connections on a Hermitian line bundle over X is affinely modelled on $i\Omega^1(X)$, and $\mathrm{d} : C^\infty(X) \to i\Omega^1(X)$ defines a flat connection on any trivial bundle. It follows that any flat connection on a trivial line bundle is given by a closed imaginary 1-form on X, and a 1-form $i\alpha$ corresponds to the following holonomy representation:

$$\gamma \mapsto \exp\left(\int_\gamma i\alpha\right). \tag{1.26}$$

By the previous proposition, two flat connections are gauge equivalent if and only if they induce the same holonomy representation. It is easy to see that any exact 1-form must then be gauge-equivalent to d, meaning we have a map $H^1_{\mathrm{dR}}(X) \to \mathrm{Hom}(\pi_1(|X|), \mathrm{U}(1))$; by Satake's theorem, this cohomology group is isomorphic to $H^1(|X|, \mathbb{R})$. Additionally, it follows by definition that α defines a cohomology class valued in $2\pi\mathbb{Z}$ if and only if its holonomy representation is trivial. Applying the first isomorphism theorem completes the proof. □

The quotient space $H^1(|X|;\mathbb{R})/H^1(|X|;\mathbb{Z})$ is called the *Jacobian torus* of X, and it is sometimes denoted by $\mathfrak{J}(X)$.

Example 1.39. To illustrate the behaviour of flat line bundles over orbifolds, we briefly outline their restrictions to a two-dimensional coordinate chart. In particular, we take the orbifold X to be $\mathbb{R}^2/\mathbb{Z}_\alpha$ for $\alpha \in \mathbb{N}$, so that all complex line bundles on X may be represented by $D^2 \times \mathbb{C}$ a representation $\rho : \mathbb{Z}_\alpha \to \mathrm{U}(1)$ (cf. Proposition 1.12). The exterior derivative always constitutes a $\mathbb{Z}_\alpha$-equivariant flat connection on $D^2 \times \mathbb{C}$, and all other flat connections may be obtained by adding a closed, $\mathbb{Z}_\alpha$-equivariant 1-form $\alpha \in i\Omega^1(\mathbb{R}^2)$.

Note that $\exp(\int_\gamma i\alpha)$ will be zero for any closed loop γ in $\mathbb{R}^2$, since every 1-form on $\mathbb{R}^2$ is exact. However, this does not mean the holonomy representation is trivial; if γ is an *orbifold* loop on $\mathbb{R}^2/\mathbb{Z}_\alpha$ then a vector field along the loop γ may suffer jump discontinuities across liftings, according to the specified partition of γ and the representation ρ. Since $\pi_1(X)^{\mathrm{orb}} = \mathbb{Z}_\alpha$, each representation ρ gives rise to a different holonomy representation for a flat connection. In the case that ρ is trivial, the exterior derivative has trivial holonomy and it is the only flat connection up to unitary gauge equivalence.

1.3 Orbifold Riemann Surfaces

We will be particularly interested in compact orbifold Riemann surfaces, and complex line bundles over them. Throughout this section, we outline the topology and geometry of these spaces.

Let Σ be a compact complex orbifold of dimension 1; then the underlying topological space $|\Sigma|$ is a surface of genus g. As Σ is compact and orientable, its singular locus must consist of finitely many isolated points, all with cyclic local group. We can therefore conceptualise Σ as a complex curve with a collection of *marked points* $x_1, \ldots, x_n \in \Sigma$ and corresponding *local invariants* $\alpha_1, \ldots, \alpha_n \in \mathbb{N}_{\geq 2}$, with the understanding that Σ is biholomorphic to $\mathbb{C}/\mathbb{Z}_{\alpha_i}$ in a neighbourhood U_i of x_i, and is locally biholomorphic to $\mathbb{C}$ everywhere else. Thus, we have an open cover $\{U_x\}_{x \neq x_1,\ldots,x_n}$ of Σ_1 and biholomorphisms $\phi_x : D^2 \to U_x$, as well as holomorphic $\mathbb{Z}_{\alpha_i}$-equivariant maps $\phi_i : D^2 \to U_i$ which induce homeomorphisms. We will continue to use this notation throughout the document.

Note that the complex structure on Σ restricts its orbifold atlas to be orientable. We will eventually consider non-orientable 2-orbifolds in Section 2.1, but via the Klein orbifold formalism.

1.3.1 Geometry and topology of orbifold Riemann surfaces

A Riemannian metric can be specified over each U_i as a $\mathbb{Z}_{\alpha_i}$-equivariant positive symmetric bilinear form on TD^2. Just as in the non-singular case, we can construct the Levi-Civita tensor and hence the Riemann curvature tensor on an orbifold Riemann surface. One may wonder whether the uniformisation theorem for surfaces can be generalised to orbifolds; the following proposition details the exceptions:

Proposition 1.40. *A compact orbifold Riemann surface Σ admits a metric of constant curvature if and only if it is a very good orbifold, i.e., it is diffeomorphic to M/G for some smooth surface M and some finite group G.*

Proof. As shown in [Thu02], a 2-orbifold admits a covering by a smooth 2-manifold if and only if it admits an elliptic, parabolic, or hyperbolic structure; this is to say it admits a representation of the form S/Γ, where S is one of S^2, $\mathbb{E}^2$, and $\mathbb{H}^2$, and Γ is a discrete group acting isometrically on S. By projecting the homogeneous metric on S down to the 2-orbifold, we get a constant-curvature metric. On the other hand, it is known that all 2-orbifolds covered by smooth manifolds admit a finite smooth covering (see [Sco83]), so we can take the group G to be the deck transformations of this covering. □

Remark. The proof can be easily adapted to show the following: if Σ is a compact orbifold Riemann surface acted on by a finite group Λ, then there is a constant-curvature metric on Σ which is invariant under Λ. More precisely, the proposition gives us a constant-curvature metric on Σ/Λ (which is also an orbifold), and this can be pulled up to a Λ-invariant metric on Σ through the quotient $\Sigma \to \Sigma/\Lambda$.

It follows that, whenever we claim that Σ has a metric with constant curvature, we are restricting to the very good orbifolds. However, not all 2-orbifolds enjoy this property. Bad 2-orbifolds all have topological type S^2; these are the *teardrop surfaces* with one marked point, and the *spindle surfaces* with two marked points of distinct multiplicity [Sco83]. Though the restriction is minor, it does have some consequences; we will eventually consider Seifert fibred spaces over orbifolds with constant curvature, and spaces fibring only over bad orbifolds will be excluded.

We will also use results regarding the algebraic topology of orbifold Riemann surfaces, specifically the following two theorems:

Proposition 1.41. *The orbifold fundamental group of Σ has the following presentation [Sco83]:*

$$\pi_1(\Sigma)^{\mathrm{orb}} = \left\langle a_1, b_1, \ldots, a_g, b_g, g_1, \ldots, g_n \middle| g_i^{\alpha_i} = 1, \prod_{j\leq g}[a_j, b_j]\prod_{i\leq n} g_i = 1 \right\rangle. \tag{1.27}$$

Proposition 1.42. *If Σ is an orbifold Riemann surface, then the following holds [FS92]:*

$$H^1(\Sigma;\mathbb{Z})^{\mathrm{orb}} \cong H^1(|\Sigma|;\mathbb{Z}), \tag{1.28}$$

$$H^2(\Sigma;\mathbb{Z})^{\mathrm{orb}} \cong \mathrm{Pic}^t(\Sigma). \tag{1.29}$$

Remark. The presentation of the orbifold fundamental group above is closely related to the standard presentation of $\pi_1(|\Sigma|)$; the only difference is the inclusion of another generator of order α_i for each marked point x_i.

In fact, there is another topological invariant which we make use of, the *orbifold Euler characteristic.* Its generalisation to orbifolds can be understood best in the context of the generalisation of the Gauss-Bonnet theorem, proved by Satake in [Sat57]:

Theorem 1.43 (Satake-Gauss-Bonnet)**.** *Given an orbifold Riemann surface Σ with a Riemannian metric and associated Gaussian curvature $K : \Sigma \to \mathbb{R}$, we have the following formula:*

$$\int_\Sigma K \mathrm{vol}_\Sigma = 2\pi\chi(\Sigma), \tag{1.30}$$

where $\chi(\Sigma)$ is the orbifold Euler characteristic:

$$\chi(\Sigma) = 2 - 2g - \sum_{i \le n} \left(1 - \frac{1}{\alpha_i}\right). \tag{1.31}$$

1.3.2 Complex line bundles

Recall that a complex line bundle over an orbifold is an orbifold fibre bundle with fibre $\mathbb{C}$. Here, we describe the structure of complex line bundles over orbifold Riemann surfaces in some detail. From this point forth, we will always assume that a complex line bundle is equipped with a Hermitian metric, and that connections preserve this metric. We will also assume that all line bundles are complex.

Using the fibre bundle construction lemma (Proposition 1.14), we may characterise complex line bundles over Σ explicitly. According to this lemma, we may think of a bundle over Σ as an ordinary bundle everywhere except over marked points; thus, we briefly restrict attention to the complex orbifold $D^2/\mathbb{Z}_\alpha$. Any complex line bundle over this orbifold is uniquely specified by a U(1)-representation of $\mathbb{Z}_\alpha$. Because $\mathbb{Z}_\alpha$ is cyclic, each representation is described by taking the generator of $\mathbb{Z}_\alpha$ to some α-th root of unity in $\mathbb{C}$. Moreover, if the generator is taken to a primitive root of unity, say $\zeta_\alpha \in \mathrm{U}(1)$, the induced line bundle clearly generates $\mathrm{Pic}^t(D^2/\mathbb{Z}_\alpha)$. The holomorphic sections of a line bundle over $D^2/\mathbb{Z}_\alpha$ are naturally represented by maps $f : D^2 \to \mathbb{C}$ which are equivariant under the action by $\mathbb{Z}_\alpha$; if the line bundle is induced by taking the generator of $\mathbb{Z}_\alpha$ to $(\zeta_\alpha)^\beta$ for $0 \le \beta < \alpha$, then $f(\zeta_\alpha \cdot z) = (\zeta_\alpha)^\beta f(z)$. By taking the power series expansion of f, it follows that $f(z) = z^\beta g(z^\alpha)$ for some holomorphic map $g : D^2 \to \mathbb{C}$.

These examples naturally extend to orbifold line bundles over Σ. We denote by H_{x_i} the bundle which is trivial everywhere except the neighbourhood U_i, where it is induced by the representation taking the generator of $\mathbb{Z}_{\alpha_i}$ to $\zeta_\alpha \in \mathrm{U}(1)$. A section looks locally like a holomorphic function *except around x_i*, where it looks like $z^{\beta_i} g_i(z^{\alpha_i})$ for some holomorphic map $g_i : D^2 \to \mathbb{C}$.

In general, the total space of a complex line bundle $L \to \Sigma$ will be singular around each marked point; more precisely, its restriction to each neighbourhood U_i is isomorphic to $H_{x_i}^{\otimes \beta_i}$ for some $\beta_i \in [0, \alpha_i)$, which is singular if $\beta_i \neq 0$. We can recover a nonsingular total space by tensoring with $H_{x_i}^{\otimes -\beta_i}$ for each i.

Definition 1.44. Let $L \to \Sigma$ be an orbifold line bundle over an orbifold Riemann surface. If $E \otimes H_{x_1}^{\otimes -\beta_1} \otimes \cdots \otimes H_{x_n}^{\otimes -\beta_n}$ is locally trivial over each marked point, we call it the *desingularisation* of L and denote it by $|L|$. From this, we define the following invariants:

- The Poincaré dual of the first Chern class of $|L|$ (thought of as a line bundle over $|\Sigma|$) is called the *background degree of* L.
- The integers β_i mod α_i are called the *local invariants of* L *at* x_i.

The tuple $(\deg(|L|), \beta_1, \ldots, \beta_n)$ is called the *Seifert invariant* of L.

Note that the bundle $|L|$ has a section with $\deg(|L|)$ zeros (counted with multiplicity), all of which can be chosen to lie away from the marked points x_i. By tensoring with all of the bundles $H_{x_i}^{\beta_i}$, we obtain a section of L with an additional zero at each x_i (if $\beta_i \neq 0$). It is natural to define the multiplicity of these zeros to be β_i/α_i, which leads to the following:

Definition 1.45. Let Σ be a compact orbifold Riemann surface, and let $L \to \Sigma$ be a complex orbifold line bundle. The *degree of* L is defined to be the following sum:

$$\deg(L) = \deg(|L|) + \sum_i \frac{\beta_i}{\alpha_i}. \tag{1.32}$$

The degree defines a group homomorphism $\deg : \mathrm{Pic}^t(\Sigma) \to \mathbb{Q}$.

Though the degree is a complete invariant for non-singular line bundles up to isomorphism, it does not completely classify orbifold line bundles; if Σ has two marked points x and y of multiplicity 2, for instance, the line bundle $H_x \otimes H_y$ has the same degree as the non-singular line bundle with degree 1, even though they are clearly not isomorphic. We can recover a complete invariant by considering in addition the local invariants:

Theorem 1.46. *Let Σ be a compact orbifold Riemann surface. The map*

$$\begin{aligned} \mathrm{Pic}^t(\Sigma) &\to \mathbb{Q} \oplus \mathbb{Z}_{\alpha_1} \oplus \cdots \oplus \mathbb{Z}_{\alpha_n} \\ L &\mapsto (\deg(L), \beta_1, \ldots, \beta_n) \end{aligned} \tag{1.33}$$

is injective (though it is not *a group homomorphism). Its image consists of the tuples $(b + \sum_i \gamma_i/\alpha_i, \gamma_1, \ldots, \gamma_n)$ for $\gamma_i \in \mathbb{Z}_{\alpha_i}$ and $b \in \mathbb{Z}$.*

Proof. The proof is due to [FS92]. Denote by $\mathrm{Pic}^t(\Sigma)_{\beta_1,\ldots,\beta_n}$ the subset of the Picard group consisting of line bundles with local invariants given by β_i, and define a map $f : \mathrm{Pic}^t(\Sigma) \to \mathbb{Z}$ by taking $f(L) = \langle c_1(|L|), [|\Sigma|] \rangle$, i.e., the background degree of L. We claim that f is bijective on $\mathrm{Pic}^t(\Sigma)_{\beta_1,\ldots,\beta_n}$.

First, suppose $f(L) = f(L')$ for line bundles $L, L' \in \mathrm{Pic}^t(\Sigma)_{\beta_1,\ldots,\beta_n}$. We then see that $c_1(|L^{-1}| \otimes |L'|) = 0$; but since L and L' have the same local invariants, we know that $|L^{-1}| \otimes |L'| = |L^{-1} \otimes L'|$. It follows that the background degree of $L^{-1} \otimes L'$ is zero, and its local invariants are all zero as well, so it is trivial. Hence, f is injective. Conversely, given a bundle $L \in \mathrm{Pic}^t(\Sigma)_{\beta_1,\ldots,\beta_n}$ and some regular point x, we know from the theory of line bundles on Riemann surfaces (see [Don11]) that there is a bundle L_x admitting a section vanishing only at x (with multiplicity 1). It follows that L_x has trivial local invariants, but

$f(L_x) = c_1(|L_x|) = c_1(L_x) = 1$. As such, we see that $f(L \otimes L_x^{\otimes k}) = f(L) + k$ for any $k \in \mathbb{Z}$, meaning f is surjective.

Now, since f is bijective over each collection of local invariants, we can make this into an injective map $L \mapsto (c_1(|L|), \beta_1, \ldots, \beta_n)$ from $\mathrm{Pic}^t(\Sigma)$ to $\mathbb{Z} \oplus \mathbb{Z}^{\oplus n}$. The rest follows easily. □

Corollary 1.47. $\mathrm{Pic}^t(\Sigma)$ *has the following abelian group presentation:*

$$\mathrm{Pic}^t(\Sigma) = \langle g, h_1, \ldots, h_n \,|\, \alpha_i h_i = g \text{ for all } i \rangle. \tag{1.34}$$

In particular, if the local invariants α_i are pairwise coprime, then $\mathrm{Pic}^t(\Sigma)$ *is generated by a single line bundle with degree* $(\alpha_1 \cdots \alpha_n)^{-1}$.

We complete this section by discussing the canonical bundle on Σ, which coincides with the holomorphic cotangent bundle $\Lambda^{1,0}\Sigma$, and is denoted by K_Σ.

Proposition 1.48. *The Seifert invariant of the canonical bundle is* $(2g-2, \alpha_1 - 1, \ldots, \alpha_n - 1)$. *The degree of the canonical bundle is* $-\chi(\Sigma)$.

Proof. The desingularisation of the canonical bundle is just the canonical bundle of $|\Sigma|$, and it is a basic result from complex geometry that the degree of this bundle is $-\chi(|\Sigma|) = 2g-2$. To compute the local invariants around the marked points, we simply note that $\mathbb{Z}_{\alpha_i}$ acts locally on the tangent bundle by taking $(z, w) \mapsto (\gamma \cdot z, \mathrm{d}\gamma \cdot w) = (\gamma \cdot z, \gamma \cdot w)$ (since the group action is linear). It follows that the tangent bundle representation at x_i takes the generator of $\mathbb{Z}_{\alpha_i}$ to $e^{2\pi i/\alpha_i}$; taking the conjugate transpose, we see that the cotangent bundle representation at x_i takes each generator to $e^{-2\pi i/\alpha_i}$. It follows that the local invariant at x_i for K_Σ is $-1 \cong \alpha_i - 1$.

To calculate the degree of K_Σ, one needs only use the formula for the degree above. □

1.3.3 Connections on line bundles

The notion of a connection is fundamental to mathematical gauge theory. In particular, connections on complex line bundles are convenient in several ways; the following result illustrates the close link between the curvature of a connection and the topology of its underlying line bundle.

Theorem 1.49. *Let $L \to \Sigma$ be a complex orbifold line bundle over a compact orbifold Riemann surface, and let $A \in \mathscr{A}(L)$ be a connection on L. Then we have the following equation:*

$$\frac{i}{2\pi} \int_\Sigma F_A = \deg(L). \tag{1.35}$$

Proof. The proof is an adaptation of the approach in [Huy05]. We begin by choosing a series of objects on Σ and L as follows. Let $\{U_x\}_{x \in \Sigma_{\mathrm{reg}}} \cap \{U_i\}_i$ be an open cover of Σ respecting the orbifold structure over which L is trivial, and let $\varphi_i : D^2 \to U_i$ be a $\mathbb{Z}_{\alpha_i}$-equivariant map reducing to a biholomorphism under the quotient by $\mathbb{Z}_{\alpha_i}$. Let $\sigma \in \Gamma(L)$ be a section of L with zeros of multiplicity β_i/α_i at x_i, and $\deg(|L|)$ zeros elsewhere. Importantly, we ensure that all zeros of $|\sigma|$ are away from marked points. Denote by h the Hermitian structure on L, and let ∇ be the corresponding Chern connection.

All of these objects can be written locally over each U. If L is trivial over each U, then there are biholomorphic trivialisations $\psi_x : U_x \times \mathbb{C} \to L|_{U_x}$ for each $x \in \Sigma_{\mathrm{reg}}$, and holomorphic maps $\psi_i : U_i \times \mathbb{C} \to L|_{U_i}$ which reduce to biholomorphisms under the quotient of $U_i \times \mathbb{C}$ by the action $(z, w) \mapsto (\zeta_{\alpha_i} \cdot z, (\zeta_{\alpha_i})^{\beta_i} \cdot z)$ of $\mathbb{Z}_{\alpha_i}$. We can describe σ as a holomorphic function $\sigma_x : U_x \to \mathbb{C}$ over each U_x, and as a holomorphic function $\sigma_i : D^2 \to \mathbb{C}$ of the form $\sigma_i(z) = z^{\beta_i} g_i(z^{\alpha_i})$ for $g_i : D^2 \to \mathbb{C}$ holomorphic over each U_i. We can describe the Hermitian structure as a map $h_x : U_x \to [0, \infty)$ over each U_x, and as a map $h_i : D^2 \to [0, \infty)$ such that $h_i(\zeta_{\alpha_i} \cdot z) = h_i(z)$ over each U_i; the explicit correspondence is that we take $h(\sigma(y)) = h_x(y)|\sigma_x(y)|^2$ for $y \in U_x$, and[1] $h(\sigma(y)) = h_i(\varphi_i^{-1} n(y))|\sigma_i(\varphi_i^{-1}(y))|^2$ for $y \in U_i$. Finally, the connection matrix for ∇ can be locally expressed as $h^{-1}\partial h = \partial \log h$ by the Chern correspondence (see Proposition 4.2.14 from [Huy05]), meaning the curvature 2-form is given by $\overline{\partial}\partial \log h$.

We now begin the proof. First, observe that $\overline{\partial}\partial \log(h(\sigma, \sigma)) = \overline{\partial}\partial \log(h_i)$ over U_i, since σ is a holomorphic section and $h(\sigma, \sigma) = h_i \cdot |\sigma_i|^2$ locally. It follows that $F^\nabla = \overline{\partial}\partial \log(h \circ \sigma)$ wherever $\sigma \neq 0$. Around each zero of σ, we can find a neighbourhood on which $|h \circ \sigma| < \varepsilon$ for any $\varepsilon > 0$; calling the union of these disks D_ε, we can rewrite the integral of F^∇ as follows:

$$\begin{aligned} \frac{i}{2\pi} \int_\Sigma F^\nabla &= \lim_{\varepsilon \to 0} \int_{\Sigma \setminus D_\varepsilon} F^\nabla \\ &= \frac{i}{2\pi} \lim_{\varepsilon \to 0} \int_{\Sigma \setminus D_\varepsilon} \overline{\partial}\partial \log(h \circ \sigma) \\ &= \frac{i}{2\pi} \lim_{\varepsilon \to 0} \int_{\Sigma \setminus D_\varepsilon} \frac{1}{2} \mathrm{d}(\partial - \overline{\partial}) \log(h \circ \sigma) && (d = \partial + \overline{\partial} \text{ and } \partial^2 = \overline{\partial}^2 = 0) \\ &= \frac{i}{4\pi} \lim_{\varepsilon \to 0} \int_{\partial D_\varepsilon} (\partial - \overline{\partial}) \log(h \circ \sigma). && \text{(Stokes' theorem)} \end{aligned}$$

Locally, we have that $h \circ \sigma \to h_i \cdot |\sigma_i|^2 = h_i \cdot \sigma_i \cdot \overline{\sigma_i}$. It follows that $\log(h \circ \sigma)$ locally becomes $\log(h_i) + \log(\sigma_i) + \overline{\log(\sigma_i)}$, and noting that the second term is holomorphic and the third is antiholomorphic, we observe the following:

$$\begin{aligned} (\partial - \overline{\partial}) \log(h \cdot |\sigma_i|^2) &= (\partial - \overline{\partial}) \log(h_i) + \partial \log(\sigma_i) - \overline{\partial \log(\sigma_i)} \\ &= (\partial - \overline{\partial}) \log(h_i) + 2i \operatorname{Im}(\partial \log(\sigma_i)). \end{aligned}$$

Integrating this over ∂D_ε and taking $\varepsilon \to 0$, we note that h_i is positive-definite and smooth, meaning $(\partial - \overline{\partial}) \log(h_i)$ is bounded; its integral over D_ε will therefore vanish as $\varepsilon \to 0$. If we write D_ε as a union of disks V_y^ε around y, where y is a zero of σ, we have the following expression for the integral of the curvature:

$$\frac{i}{2\pi} \int_\Sigma F^\nabla = -\frac{1}{2\pi} \lim_{\varepsilon \to 0} \left(\operatorname{Im} \sum_{y \text{ zero of } \sigma} \int_{\partial V_y^\varepsilon} \partial \log(\sigma_i) \right). \tag{1.36}$$

If y is not a marked point of Σ and σ vanishes at y with multiplicity k, it is shown in [Huy05] that the integral is simply $-2\pi i k$. If y is one of the marked points x_i, we note that $\sigma_i(z) = z^{\beta_i} g_i(z^{\alpha_i})$ where g_i is holomorphic, meaning $\partial \log(\sigma_i) = \partial(\log z^{\beta_i} + \log g(z)) = \beta_i / z + \partial \log g(z)$. Moreover, since we chose $|\sigma|$ to have no zeros at marked points, we know

[1] Note that this expression is well-defined since both h_i and $|\sigma_i|^2$ is unaffected by the action of $\mathbb{Z}_{\alpha_i}$.

that g is nonzero in a neighbourhood of zero; it follows that $\partial \log g(z)$ is holomorphic in a small enough neighbourhood, and its line integral will vanish. Since $\partial V_{x_i}^{\varepsilon}$ is a closed simple curve containing the point x_i, we can represent it as a closed curve C_ε around 0 in a local trivialisation of L. Finally, since the definition of the integral over an orbifold point is slightly modified, we need to introduce an extra factor of $1/\alpha_i$ in our result. Thus, we get the following:

$$\begin{aligned}\lim_{\varepsilon \to 0} \int_{\partial V_{x_i}^{\varepsilon}} \partial \log(\sigma) &= \lim_{\varepsilon \to 0} \frac{1}{\alpha_i} \int_{C_\varepsilon} \partial \log \sigma_i(z) \mathrm{d}z \\ &= \frac{1}{\alpha_i} \lim_{\varepsilon \to 0} \int_{C_\varepsilon} (\beta_i/z + \partial \log g(z)) \mathrm{d}z \\ &= -\frac{2\pi i \beta_i}{\alpha_i}.\end{aligned}$$

(As stated in [Huy05], the negative sign appears because the initial domain of integration was the complement of the interior.) Using our formula for the integral of the curvature, we conclude that

$$\frac{i}{2\pi} \int_\Sigma F^\nabla = \deg(|L|) + \sum_i \frac{\beta_i}{\alpha_i},$$

which is how we defined the degree of an orbifold line bundle. □

The following result is related to the previous theorem; we will use it extensively when discussing reducible solutions to gauge-theoretic equations.

Proposition 1.50. *Given an orbifold Riemann surface Σ, any complex line bundle $L \to \Sigma$ can be given a connection whose curvature is a constant scalar multiple of* vol_Σ*, the volume element of Σ. The constant of proportionality is $-2\pi i \deg(L)/\mathrm{Vol}(\Sigma)$. In particular, if* $\deg(L) = 0$*, then L admits a flat connection.*

Proof. Consider $[\mathrm{vol}_\Sigma]$ as a class in $H^2_{\mathrm{dR}}(\Sigma) \cong H^2(|\Sigma|; \mathbb{R}) \cong \mathbb{R}$. We claim that vol_Σ defines a nonzero class; indeed, if vol_Σ defined the zero class, then it would be exact and $\mathrm{Vol}(\Sigma)$ would be zero. Now, let $A \in \mathscr{A}(L)$ be some connection (which always exists by a partition of unity argument); then $F_A = i f \mathrm{vol}_\Sigma$ for some $f \in C^\infty(\Sigma)$. Since vol_Σ spans the de Rham cohomology space, there is some constant $\zeta \in i\mathbb{R}$ for which $\zeta \mathrm{vol}_\Sigma - i f \mathrm{vol}_\Sigma$ is exact.

Let $\alpha \in \Omega^1(\Sigma)$ be the 1-form for which

$$\mathrm{d}\alpha = \zeta \mathrm{vol}_\Sigma - i f \mathrm{vol}_\Sigma. \tag{1.37}$$

Then $A + \alpha$ is the desired connection, since $F_{A+\alpha} = F_A + \mathrm{d}\alpha = \zeta \mathrm{vol}_\Sigma$. To determine the value of ζ, we need only use the above result for the degree to determine that

$$\deg(L) = \frac{i}{2\pi} \int_\Sigma F_{A+\alpha} = \frac{i}{2\pi} \zeta \int_\Sigma \mathrm{vol}_\Sigma = \frac{i\zeta \mathrm{Vol}(\Sigma)}{2\pi}; \tag{1.38}$$

the result follows from rearranging. □

1.4 Seifert Fibred Spaces

From the perspective of Seiberg-Witten theory, the interest in orbifold Riemann surfaces is that they give a window into 3-manifold topology. The object through which this is expressed is the Seifert fibred space.

Definition 1.51. Let Σ be a compact orbifold Riemann surface. A *Seifert fibred space over* Σ is a 3-manifold Y together with a projection $\pi : Y \to \Sigma$ with the property that $Y \cong S(N)$, the unit circle bundle of an orbifold line bundle $N \to \Sigma$. Equivalently, a *Seifert 3-manifold* is a 3-manifold Y equipped with a smooth action by $\mathrm{U}(1)$ whose stabiliser subgroups are everywhere finite.

The equivalence between the two definitions is seen by taking Σ to be the space of $\mathrm{U}(1)$-orbits of Y; this orbit space is an orbifold since all stabilisers are finite, and one may recover the line bundle N as the associated bundle $Y \times_{\mathrm{id}} \mathbb{C}$. It is fairly straightforward to characterise the circle bundles which define smooth manifolds:

Proposition 1.52. *A circle bundle $S(N)$ over a Riemann surface Σ is a smooth manifold (every point of $S(N)$ is regular) if and only if $\gcd(\beta_i, \alpha_i) = 1$ for all i.*

Proof. It clearly suffices to consider the local behaviour around a marked point; that is, we assume the underlying Riemann surface is D^2 with the usual action of $\mathbb{Z}_\alpha$, and the line bundle N is represented by $(D^2 \times \mathbb{C})/\mathbb{Z}_\alpha$ where the generator of $\mathbb{Z}_\alpha$ acts on $D^2 \times \mathbb{C}$ by taking

$$(p, v) \mapsto (\zeta_\alpha \cdot p, (\zeta_\alpha)^\beta \cdot v) \tag{1.39}$$

for some $\beta \in \{0, 1, \ldots, \alpha - 1\}$. We now restrict to Y, meaning we take v to be unit-length; a fixed point of this action must then be of the form $(0, v)$ for $v \in \mathbb{C}$, and the elements of $\mathbb{Z}_\alpha$ fixing v correspond to integers m satisfying

$$(\zeta_\alpha)^{\beta m} = 1, \tag{1.40}$$

or equivalently $\beta m = 0 \bmod \alpha$. It follows from Bezout's identity that the subgroup of $\mathbb{Z}_\alpha$ fixing $(0, v)$ is the cyclic group $\mathbb{Z}_{\gcd(\alpha,\beta)}$, meaning the local group is trivial if and only if $\gcd(\alpha, \beta) = 1$. □

A circle bundle over an orbifold Riemann surface with singularities will be called a *Seifert 3-orbifold.*

Note that the algebraic topology of a Seifert fibred space reflects its underlying orbifold structure, as in the following theorem.

Theorem 1.53. *Let $N \to \Sigma$ be a line bundle with Seifert invariant $(b, \beta_1, \ldots, \beta_n)$ over the closed orbifold Riemann surface Σ, and let $Y = S(N)$ be the associated Seifert fibred space. Then the fundamental group of Y has the following presentation:*

$$\pi_1(Y) = \Big\langle a_1, b_1, \ldots, a_g, b_g, g_1, \ldots, g_n, h \,\Big|\, [a_j, h] = [b_j, h] = [g_i, h] = g_i^{\alpha_i} h^{\beta_i} = 1, \prod_{j \le g} [a_j, b_j] \prod_{i \le n} g_i = h^b \Big\rangle. \tag{1.41}$$

Additionally, the first and second cohomology of Y can be written in terms of the cohomology of Σ as follows:

$$H^1(Y) = \begin{cases} H^1(|\Sigma|) & \deg(N) \neq 0 \\ H^1(|\Sigma|) \oplus \mathbb{Z} & \deg(N) = 0, \end{cases} \tag{1.42}$$

$$H^2(Y) = (\mathrm{Pic}^t(\Sigma)/\langle N \rangle) \oplus \mathbb{Z}^{2g}, \tag{1.43}$$

where $\langle N \rangle$ denotes the cyclic subgroup of $\mathrm{Pic}^t(\Sigma)$ generated by N.

Proof. Refer to [FS92]. □

Remark. If Y is the trivial circle bundle over Σ, then the background degree b and the local invariants β_i are all trivial. The presentation of $\pi_1(Y)$ then reduces to the direct product of $\pi_1(\Sigma)^{\mathrm{orb}}$ with $\mathbb{Z} = \langle h \rangle$, as expected.

1.4.1 Line bundles over Seifert 3-manifolds

A line bundle $L \to \Sigma$ induces a line bundle over Y in a natural way via pullback along π, which manifests in cohomology as a map $\pi^* : H^2(\Sigma;\mathbb{Z}) \to H^2(Y;\mathbb{Z})$. The topological type of the pullback bundle is then determined by the following theorem, again proven in [FS92]:

Proposition 1.54. *The subgroup $\mathrm{Pic}^t(\Sigma)/\langle N \rangle \leq H^2(Y) \cong \mathrm{Pic}^t(Y)$ is naturally the image of the pullback map $\pi^* : \mathrm{Pic}^t(\Sigma) \to \mathrm{Pic}^t(Y)$. In particular, given two line bundles $L, L' \to \Sigma$, the pullbacks π^*L and π^*L' are isomorphic if and only if $L' \cong L \otimes N^{\otimes k}$ for some $k \in \mathbb{Z}$ (here a negative tensor power is a tensor power of the dual bundle).*

A natural question is whether there are any line bundles over Y which cannot be obtained via pullback along π^*. The fact that $H^2(Y) = \pi^*\mathrm{Pic}^t(\Sigma) \oplus \mathbb{Z}^{2g}$ gives us an immediate answer: as long as $g \neq 0$, the nonzero elements of $\mathbb{Z}^{2g}$ correspond to such bundles. However, if we restrict attention to bundles which can be equipped with special connections, we can ensure that every bundle may be obtained via pullback.

Proposition 1.55. *Let $\pi : Y \to \Sigma$ be a Seifert fibred space. There is a natural one-to-one correspondence between line bundles with connection over Σ and line bundles with connection over Y whose fibrewise holonomy is trivial over regular points in Σ (that is, the holonomy of the connection around any $\pi^{-1}(x) \subseteq Y$ for $x \in \Sigma_{\mathrm{reg}}$ is trivial). This correspondence is induced by pullback, where the pullback of a connection $\nabla \in \mathscr{A}(L)$ over a line bundle $L \to \Sigma$ is defined so that*

$$(\pi^*\nabla)_v(\pi^*\sigma) = \nabla_{\pi_*(v)}\sigma, \tag{1.44}$$

for any $\sigma \in \Gamma(L)$ and any $v \in TY$. (Note that this condition fully determines $\pi^*\nabla$ by the Leibniz rule.)

Proof. The proof is due to [MOY96]. If $L \to \Sigma$ is a line bundle and $\nabla \in \mathscr{A}(L)$, the connection $\pi^*\nabla \in \mathscr{A}(\pi^*L)$ will have trivial fibrewise holonomy over the nonsingular points of Σ; this follows from the fact that $\pi^*(\frac{\mathrm{d}}{\mathrm{d}\varphi}) = 0$. This is one direction in the correspondence.

Conversely, suppose $E \to Y$ is a line bundle and $A \in \mathscr{A}(E)$ is a connection with trivial fibrewise holonomy. It clearly suffices to reduce to the trivial case, where $\Sigma = D^2/\mathbb{Z}_\alpha$ and

$Y = (D^2 \times S^1)/\mathbb{Z}_\alpha$, with $\mathbb{Z}_\alpha$ acting on S^1 by the representation $z \mapsto z^\beta$ for $\beta \in \mathbb{Z}_\alpha$. The idea is to find a section of π, pull the bundle E back to Σ along this section, and then show that this does not depend on the section using holonomy.

Let $q : D^2 \to \Sigma$ and $q' : D^2 \times S^1 \to Y$ be the quotient maps induced by $\mathbb{Z}_\alpha$. Instead of working directly with sections $\sigma : \Sigma \to Y$, it will be easier to work with $\mathbb{Z}_\alpha$-invariant maps $\tau : D^2 \to Y$ for which $\pi \circ \tau = q'$; the two kinds of map are in one-to-one correspondence. An example of such a map is given by $q' \circ \overline{\tau}$, where $\overline{\tau} : D^2 \to D^2 \times S^1$ is the map $\overline{\tau}(p) = (p, 1)$. We then define the line bundle over D^2 to be $\tau^* E$, with the connection induced by pullback, and we then quotient by $\mathbb{Z}_\alpha$ to obtain a line bundle over Σ. The relevant maps are summarised in the following commutative diagram:

$$\begin{array}{ccc} D^2 \times S^1 & \xrightarrow{q'} & Y \\ \overline{\tau}\uparrow & \nearrow^{\tau} & \downarrow\pi \\ D^2 & \xrightarrow{q} & \Sigma \end{array}$$

Now we verify that this bundle is well-defined. Suppose $\tau, \tau' : D^2 \to (D^2 \times S^1)/\mathbb{Z}_\alpha$ are two $\mathbb{Z}_\alpha$-invariant maps for which $\tau \circ \pi = \tau' \circ \pi = q'$; we wish to show that $\tau^* E$ and $(\tau')^* E$ are canonically isomorphic. Note that $\tau(p)$ and $\tau'(p)$ are in the same fibre of π for any $p \in D^2$, so there is some $u(p) \in S^1$ for which $\tau(p) = u(p)\tau'(p)$. This defines a $\mathbb{Z}_\alpha$-invariant smooth map $u : D^2 \to S^1$, but since D^2 is contractible, u is homotopic to a constant map $p \mapsto 1$ via a homotopy $F : D^2 \times [0,1] \to S^1$. We therefore define a map from $(\tau')^* E$ to $\tau^* E$ as follows: the fibre $(\tau^* E)_p$ is mapped to the fibre $((\tau')^* E)_p$ by first identifying them with subbundles in $E \to Y$, and then taking parallel transport along the path $t \mapsto F(p,t)\tau'(p) \in Y$ according to the connection A. Though the two bundles are always locally isomorphic, they are only canonically isomorphic because A has fibrewise trivial holonomy: this fact ensures that parallel transport along the specified path does not depend on the specific choice of homotopy $F : u \simeq 1$. These bundles pass down to the quotient Σ so long as we ensure that F is also $\mathbb{Z}_\alpha$-invariant. □

There is a special case in which the fibrewise holonomy condition on the connection is satisfied:

Proposition 1.56. *Let A be a connection on a line bundle $E \to Y$. If $F_A = \pi^*\omega$ for some $\omega \in i\Omega^2(\Sigma)$, and there is a point $x \in \Sigma_{\text{reg}}$ for which $\pi^{-1}(x)$ has trivial holonomy, then A has trivial fibrewise holonomy.*

Proof. Let $y \in \Sigma_{\text{reg}}$ be another regular point. Since Σ is assumed to be connected and there are only finitely many singular points, there is a non-self-intersecting path $\gamma : [0,1] \to \Sigma_{\text{reg}}$ connecting x and y. Note that $\pi^{-1}(\gamma)$ can be naturally interpreted as a cylinder, with $\pi^{-1}(x)$ and $\pi^{-1}(y)$ constituting the boundary of the cylinder. First, we show that the holonomy around $\pi^{-1}(x)$ and $\pi^{-1}(y)$ are related by a factor of $\exp(\int_{\pi^{-1}(\gamma)} F_A)$.

Let $\Gamma_1, \Gamma_2 : [0,1] \to Y$ be non-intersecting lifts of γ, with sources p_1, p_2 and targets q_1, q_2 respectively. Let $\alpha : [0,1] \to \pi^{-1}(x)$ and $\beta : [0,1] \to \pi^{-1}(y)$ be loops at p_1 and q_1 respectively, define $\alpha_1 = \alpha|_{[0,1/2]}$ and $\alpha_2 = \alpha|_{[1/2,1]}$, and define β_1 and β_2 likewise. Observe that $\alpha_1 * \Gamma_2 * \beta_1^{-1} * \Gamma_1^{-1}$ and $\Gamma_1 * \beta_2^{-1} * \Gamma_2^{-1} * \alpha_2$ are contractible loops which collectively bound the entire area of the cylinder $\pi^{-1}(\gamma)$; it follows that the product of their holonomies is equal to $\int_{\pi^{-1}(\gamma)} F_A$. On the other hand, the product of the loops is given by the

$$(\alpha_1 * \Gamma_2 * \beta_1^{-1} * \Gamma_1^{-1}) * (\Gamma_1 * \beta_2^{-1} * \Gamma_2^{-1} * \alpha_2) = \alpha_1 * \Gamma_2 * \beta^{-1} * \Gamma_2^{-1} * \alpha_2. \tag{1.45}$$

This can be rewritten as $(\alpha^{\alpha_2}(\beta^{-1})^{\Gamma_2})^{\alpha_2^{-1}}$. Conjugation does not affect the holonomy, since the structure group is abelian; thus, we have the following:

$$\int_{\pi^{-1}(\gamma)} F_A = \mathrm{Hol}_A\left((\alpha^{\alpha_2}(\beta^{-1})^{\Gamma_2})^{\alpha_2^{-1}}\right) = \mathrm{Hol}_A(\alpha)\mathrm{Hol}_A(\beta)^{-1}, \tag{1.46}$$

where we have used the fact that the holonomy is multiplicative.

However, if $F_A = \pi^*\omega$, it is clear that $\int_{\pi^{-1}(\gamma)} F_A = 0$; observe that $\Lambda^2(\pi^{-1}(\gamma))$ is spanned by a 2-form which contracts nontrivially with $\frac{\mathrm{d}}{\mathrm{d}\varphi}$, but $\frac{\mathrm{d}}{\mathrm{d}\varphi} \lrcorner \pi^*\omega = 0$. It follows that the holonomies around $\pi^{-1}(x)$ and $\pi^{-1}(y)$ are the same. □

1.4.2 Flat line bundles

Finally, we make a brief note on flat line bundles over Y. Recall that Propositions 1.37 and 1.38 allowed us to characterise the flat vector bundles on an orbifold in terms of the orbifold fundamental group, and furthermore realise the trivial flat vector bundles as a subspace. When the base space is a manifold, we may characterise the bundles which admit flat connections more explicitly:

Proposition 1.57. *A line bundle over an orientable manifold admits a flat connection if and only if its first Chern class is torsion.*

Proof. Observe that the first Chern class of a line bundle $L \to M$ is torsion if and only if its image in $H^2(M;\mathbb{R})$ is trivial. Let $A \in \mathscr{A}(L)$ be a connection; using Chern-Weil theory, we see that $c_1(L)$ is torsion if and only if $[F_A] \in H^2(M;\mathbb{R})$ defines the trivial class. By definition, this means that $F_A = \mathrm{d}\omega$ for some 1-form ω, or equivalently that $A - \omega$ is a flat connection on L. □

By writing the homology of Y in terms of Σ, we get a characterisation of the space of flat bundles on Y:

Proposition 1.58. *Let $Y \to \Sigma$ be a Seifert 3-manifold for which $Y = S(N)$, and denote by $\mathfrak{J}(\Sigma)$ the Jacobian torus of Σ. Then the space of trivial flat vector bundles on Y is given by $\mathfrak{J}(\Sigma)$ when $\deg(N) \neq 0$, and $\mathfrak{J}(\Sigma) \times \mathrm{U}(1)$ when $\deg(N) = 0$. If $\deg(N) \neq 0$, then the $\mathrm{U}(1)$-representation variety of Y fits into the following short exact sequence:*

$$0 \to \frac{H^1(|\Sigma|;\mathbb{R})}{H^1(|\Sigma|;\mathbb{Z})} \to \mathrm{Hom}(\pi_1(Y), \mathrm{U}(1)) \to \mathrm{Pic}^t(\Sigma)/\langle N\rangle \to 0. \tag{1.47}$$

Proof. The space of trivial flat vector bundles on Y is given by its Jacobian torus, so the first part of the theorem follows immediately from Theorem 1.53. In the case that $\deg(N) \neq 0$, the torsion subgroup of $H^2(Y;\mathbb{Z})$ can also be directly read off from Theorem 1.53 as the third nonzero term in the above sequence. The exact sequence then follows at once. □

Chapter 2

Equivariant Vortices on Orbifold Riemann Surfaces

Now that we have developed the theory of orbifolds and fibre bundles over them, it is possible to define gauge theories on orbifolds. We are in particular interested in the vortex equations on orbifolds, which are necessary for understanding the Seiberg-Witten equations on Seifert 3-manifolds (as per the correspondence in [MOY96]). We will also eventually define an equivariant moduli space for the Seiberg-Witten equations; this means it is important to consider how vortices behave under finite group action as well.

As such, this chapter is devoted to the study of the equivariant vortex equations. We begin by defining Λ-equivariant orbifolds and fibre bundles over them, where Λ is a finite group; in doing so, we generalise the notion of a Klein surface to a *Klein orbifold* as the special case where $\Lambda = \mathbb{Z}_2$. We then move on to discuss the Λ-equivariant vortex equations, and we present an existence proof based on the moment map approach in [Gar94] on the assumption that Uhlenbeck's gauge theorem and Uhlenbeck's weak compactness theorem generalise to orbifolds. We also touch on the Kähler vortex equations, a related set of gauge-theoretic equations of relevance to Seiberg-Witten theory. Finally, we give a proof of Uhlenbeck's results for orbifolds based on the approach in [Weh04].

Throughout this chapter, Λ always denotes a finite group acting on an orbifold Riemann surface Σ. We further assume that Σ admits a constant-curvature metric; by Proposition 1.40, this is equivalent to the requirement that Σ is very good. We also assume that the metric is invariant under the action of Λ, so that the orientation-preserving elements act holomorphically on Σ and the orientation-reversing elements act antiholomorphically. Unless otherwise stated, *every collection of objects is equivariant under the action of* Λ. The Λ-equivariant subset of some collection X is traditionally denoted by X^Λ, but we will generally omit these superscripts to avoid cumbersome notation.

2.1 Λ-equivariant Orbifolds

We begin with the equivariant notion of an orbifold.

Definition 2.1. Let Λ be a finite group. A Λ-*equivariant orbifold,* or simply a Λ-*orbifold,* is an orbifold X equipped with a smooth action by Λ. Note that X/Λ inherits orbifold charts from X, so is itself an orbifold.

Example 2.2. Given an orbifold Y, the trivial example of a Λ-orbifold is $Y \times \Lambda$ with Λ-action given by multiplication in the second factor. This example serves as a useful sanity check for the constructions that follow; any conditions we specify for data defined on a Λ-orbifold X should be met by the trivial example $Y \times \Lambda$, so long as the data is defined on $Y \times \{1\}$ and extended to $Y \times \Lambda$ in the obvious way.

Definition 2.3. Let X be a Λ-orbifold. A Λ-*equivariant bundle* over X is a bundle $\pi : E \to X$ equipped with an action of Λ by automorphisms which commutes with π. That is to say the following: given any $e \in E$ and any $h \in \Lambda$, we must have that $\pi(h \cdot e) = h \cdot \pi(e)$. In particular, if E is a complex vector bundle, we further require that the action of $h \in \Lambda$ is complex linear if h preserves orientation, and complex antilinear if h reverses orientation.

Definition 2.4. Let $E \to X$ be a Λ-equivariant bundle. The *gauge group of* E consists of all Λ-equivariant automorphisms of E, i.e., the collection of all maps $g : E \to E$ commuting with π and respecting the fibre-wise structure of E, which satisfy the additional constraint that $g(h \cdot e) = h \cdot g(e)$ for any $e \in E$, $h \in \Lambda$. It is denoted by $\mathscr{G}$.

The notion of an equivariant bundle behaves nicely under bundle operations:

- If E and E' are Λ-equivariant fibre bundles, then their fibred product $E \times_X E'$ has a natural Λ-action given by $h \cdot (e, e') = (h \cdot e, h \cdot e')$. In particular, this means the direct sum of equivariant vector bundles over X is also equivariant.

- If E and E' are Λ-equivariant vector bundles, then their tensor produt $E \otimes E'$ has a natural Λ-action given on simple tensors by $h \cdot (e \otimes e') = (h \cdot e) \otimes (h \cdot e')$.

- If $P \to X$ is a Λ-equivariant principal G-bundle and F is an orbifold on which G acts by a representation $\rho : G \to \mathrm{Aut}(F)$, then $P \times_\rho F$ is naturally a Λ-equivariant bundle under the action $h \cdot [(p, v)] = [(h \cdot p, v)]$. In particular, any vector bundle associated to an equivariant principal bundle is equivariant; as a special case, this implies that the dual and tensor powers of an equivariant vector bundle are equivariant.

Example 2.5. The simplest example of a Λ-equivariant bundle on a Λ-orbifold X is its tangent bundle; the action of Λ on the tangent bundle is given by its derivative. By the above discussion, all of the tensor bundles and the bundles of k-forms also admit a natural Λ-action.

In order to make a gauge theory on a Λ-equivariant space X descend to the quotient X/Λ, we will need to ensure that every object we define is Λ-equivariant. This includes the following examples:

- Every section σ of a Λ-equivariant bundle E is required to be equivariant under Λ, meaning $\sigma(h \cdot x) = h \cdot \sigma(x)$ for every $h \in \Lambda$ and every $x \in X$. Equivalently, if $\Gamma(E)$ is equipped with the left Λ-action $\sigma \mapsto h\sigma h^{-1}$, we require that σ is invariant under the action.

- Every connection on a Λ-equivariant principal G-bundle P can be identified with a G-equivariant $\mathfrak{g}$-valued 1-form on P, i.e., $A \in \Omega^1(P; \mathfrak{g})$, or equivalently a section of the bundle $\Lambda^1(P) \times \mathfrak{g}$. Since the action of Λ on $\mathfrak{g}$ is trivial, we require that A is a Λ-equivariant section of this bundle in the exact same way as above.

- As mentioned in the definition, a gauge transformation will be presumed invariant under Λ, meaning $g \circ h = g$ for every $h \in \Lambda$ and every $g \in \mathscr{G}$. Equivalently, if $\mathscr{G}$ is equipped with the left Λ-action $g \mapsto h^* g = g \circ h$, we require that g is invariant under the action.

- Finally, every map between Λ-equivariant objects will be required to be Λ-equivariant: if $f : X \to Y$ is a map between two Λ-orbifolds, then $h \circ f = f \circ h$ for every $h \in \Lambda$. Equivalently, we require that f is invariant under the conjugation action $f \mapsto h \circ f \circ h^{-1}$. In fact, the three previous examples are special cases of this phenomenon.

Because of the equivalence between equivariance and invariance stated above, we will essentially use the terms interchangeably when discussing well-defined data on a Λ-orbifold.

Given a Λ-invariant n-form ω on an n-dimensional Λ-orbifold, we define its integral as follows:

$$\int_{X/\Lambda} \omega = \frac{1}{|\Lambda|} \int_X \omega. \tag{2.1}$$

We also redefine the degree of a line bundle (background or total) to be $\deg(L)/|\Lambda|$. Note that this preserves the Chern relation: given a unitary connection A on L, we have that

$$\frac{i}{2\pi} \int_{X/\Lambda} F_A = \deg(L). \tag{2.2}$$

2.1.1 Klein orbifolds

As an application of the prior theory, we consider the following problem: how do we extend problems initially defined on orientable manifolds to the non-orientable setting? One answer is to observe that every unoriented manifold X has a double covering, called the *orientation double covering,* which is itself orientable; the $\mathbb{Z}_2$-fibre over each point $x \in X$ is given by a choice of an orientation of $T_x X$ (i.e., a connected component of $\mathrm{GL}(T_x X)$). Thus, orientable theory may be extended to the non-orientable setting by lifting to the double covering and ensuring that everything defined is $\mathbb{Z}_2$-invariant.

If we restrict to the case where X is an unoriented surface, and the data we wish to define is a complex structure, then the above idea specialises to the theory of Klein surfaces and Real structures on vector bundles. We therefore briefly discuss these objects here, and in doing so we shall generalise them to the orbifold category. (For a deeper overview of this topic, refer to [Sch16b] and [Nat90].)

Definition 2.6. A *Klein orbifold* is an orbifold Riemann surface Σ (without boundary) together with an antiholomorphic involution $\tau : \Sigma \to \Sigma$. Explicitly, τ must be a smooth map such that the complex conjugate of its smooth local lift around each point is holomorphic, and τ^2 must be the identity.

It is worth noting that (Σ, τ) is sometimes called a Real orbifold Riemann surface, as there is another common definition of Klein structures:

Definition 2.7. Given a smooth 2-orbifold Σ_0, a *dianalytic structure* on Σ_0 is a compatible orbifold atlas such that every embedding between charts is either holomorphic or antiholomorphic. A 2-orbifold with a dianalytic structure is called a *Klein orbifold.*

The equivalence of the two definitions is seen by taking the quotient Σ/τ to be the 2-orbifold Σ_0, with the inverse operation given by the orientation double covering of Σ_0. We will default to the former definition from this point forth.

Given a Klein orbifold (Σ, τ), the quotient $|\Sigma|/\tau$ is naturally a 2-orbifold; the fixed point set Σ^τ naturally corresponds to topological boundary points. Each orbifold Riemann surface Σ_0 can be naturally considered a Klein orbifold if we take Σ to be $\Sigma_0 \sqcup \Sigma_0$ and τ to permute Σ_0 antiholomorphically. However, even though $|\Sigma|$ must be orientable, $|\Sigma|/\tau$ need not be; for instance, we can obtain $\mathbb{RP}^2$ as the quotient by taking $\Sigma = \mathbb{CP}^1$ and τ to be the antipodal map.

In the case that Σ has no marked points, the pair (Σ, τ) is called simply a Klein surface. The topological types of these objects are well-understood, to the point that there is a relatively simple classification:

Theorem 2.8. *Let (Σ, τ) be a compact Klein surface. Then the topology of Σ/τ is completely determined by the genus g of Σ, the number k of boundary components of $\partial(\Sigma/\tau)$, and the orientability of Σ/τ. If Σ/τ is orientable then k is a number between 1 and $g+1$ with opposite parity to g, and if Σ/τ is non-orientable then $0 \leq k \leq g$. Conversely, a pair (g, k) satisfying these conditions gives rise to a Klein surface.*

We may extend this classification to the case of Klein orbifolds, by simply adding marked points to Σ of various multiplicities. Note that, in order for $\tau : \Sigma \to \Sigma$ to be a smooth involution, the marked points should come in τ-related pairs; thus, every marked point $x \in \Sigma$ has a corresponding marked point $\tau(x) \in \Sigma$ with the same multiplicity.

Additionally, we have the notion of a vector bundle over a Klein surface:

Definition 2.9. Let (Σ, τ) be a Klein orbifold, and let $\pi : E \to \Sigma$ be a holomorphic orbifold vector bundle. A *Real structure on E* is an antiholomorphic involution $\widetilde{\tau} : E \to E$ which is fibrewise antilinear and lifts the action of τ, i.e., $\pi \circ \widetilde{\tau} = \tau \circ \pi$. The pair $(E, \widetilde{\tau})$ is called a *Real vector bundle over* (Σ, τ).

2.2 The Λ-equivariant Vortex Equations

With the theory of Λ-equivariant bundles established, we begin our discussion of the Λ-equivariant moduli space of vortices on Λ-orbifolds. Our exposition of this theory will closely follow the non-equivariant non-singular treatment of the equations; the reader is referred to [Gar91] and [Bra90] for more details.

The natural setting for the vortex equations is a Riemann surface with a Hermitian line bundle. Thus, in the Λ-equivariant setting, we take Σ to be a Λ-orbifold Riemann surface and choose a Λ-equivariant line bundle $L \to \Sigma$.

Definition 2.10. The *(unitary) gauge group* is the group of Λ-equivariant unitary automorphisms of L, which is simply[1] $\mathscr{G} = C^\infty(X, \mathrm{U}(1))^\Lambda$. Its complexification $\mathscr{G}^\mathbb{C} = C^\infty(X, \mathbb{C}^*)^\Lambda$ is called the *complex gauge group*. These both have a natural infinite-dimensional Lie group structure, with Lie algebra $\mathfrak{g} = C^\infty(\Sigma, i\mathbb{R})^\Lambda$ and $\mathfrak{g}^\mathbb{C} = C^\infty(\Sigma, \mathbb{C})^\Lambda$ respectively. We assume that $\mathfrak{g}$ is equipped with the natural L^2 inner product, which allows us to identify $\mathfrak{g}$ with $\mathfrak{g}^*$.

[1] Recall from Example 1.25 that the gauge group reduces to maps into the structure group G when G is abelian.

While there is a natural action of $\mathscr{G}$ on $\mathscr{A}(L)$ given by conjugation, the action of $\mathscr{G}^{\mathbb{C}}$ is more complicated: given $A \in \mathscr{A}(L)$ and $g \in \mathscr{G}^{\mathbb{C}}$, we define $A^g - A$ to be $-g^{-1}\overline{\partial}g + (g^{-1}\overline{\partial}g)^*$.

We are now ready to define the space of objects on which the vortex equations will be defined.

Definition 2.11. The *pre-configuration space of the vortex equations* is the affine space $\mathscr{C}(L) = \Gamma(L) \times \mathscr{A}(L)$, the space of pairs of Λ-invariant sections and unitary connections on L. This space is acted upon smoothly by $\mathscr{G}$, and the quotient $\mathscr{B}(L) = \mathscr{C}(L)/\mathscr{G}$ is called the *configuration space.* The subset of $\mathscr{C}(L)$ on which $\mathscr{G}$ acts freely (i.e., with trivial isotropy) is called the *irreducible locus* $\mathscr{C}^*(L)$, and its complement is called the *reducible locus.* The quotient $\mathscr{C}^*(L)/\mathscr{G}$ is denoted by $\mathscr{B}^*(L)$.

Note that a configuration may be written either as (ϕ, A) or (A, ϕ), where ϕ is a section and A is a connection.

Proposition 2.12. *A gauge transformation $g \in \mathscr{G}$ fixes an element of the pre-configuration space if and only if $g \in \mathrm{U}(1)$, the subgroup of constant gauge transformations of $\mathscr{G}$. An element $(\phi, A) \in \Gamma(L) \times \mathscr{A}(L)$ is reducible if and only if ϕ is identically zero.*

Proof. A is fixed by g if and only if $g^{-1}\mathrm{d}g = 0$ everywhere, which is equivalent to the constancy of g. The pair (ϕ, A) is fixed by a nontrivial g if and only if g is constant and ϕ is fixed by g; since $g \in \mathrm{U}(1)$, ϕ must be identically zero. □

Definition 2.13. Let $\tau \in \mathbb{R}$ be a fixed real number. The *Yang-Mills-Higgs functional associated to* τ is the following real-valued functional on the pre-configuration space:

$$\mathrm{YMH}_\tau(\phi, A) = \int_\Sigma (|F_A|^2_{i\Lambda^2\Sigma} + |\nabla^A\phi|^2_{\Lambda^1\otimes L} + ||\phi|^2 - \tau|^2)\mathrm{vol}_\Sigma. \tag{2.3}$$

Here, each subscript refers to the space in which the norm is being taken. Also, $|\phi|^2 = h(\phi, \phi)$, where h is the Hermitian structure on L; we generally suppress h in the notation. Note that YMH_τ is invariant under the action of $\mathscr{G}$, meaning it descends to a well-defined functional on $\mathscr{B}(L)$.

By using the twisted versions of the Kähler identities and expanding inner products, we can rewrite the Yang-Mills-Higgs functional in the following form:

$$\mathrm{YMH}_\tau(\phi, A) = 2\left\|\overline{\partial}_A\phi\right\|^2_{L^2} + \left\|*F_A - \tfrac{i}{2}(|\phi|^2 - \tau)\right\|^2_{L^2} + 2\pi\tau\deg(L). \tag{2.4}$$

Here $\overline{\partial}_A : \Gamma(L) \to \Omega^{0,1}(L)$ denotes the holomorphic structure associated to the connection A. We therefore find that a sufficient condition for the Yang-Mills-Higgs functional to be minimised is for the *vortex equations* to be satisfied:

Definition 2.14. A pair $(\phi, A) \in \mathscr{C}(L)$ *satisfies the vortex equations* if the following hold:

$$*F_A = \frac{i}{2}(|\phi|^2 - \tau), \tag{2.5a}$$
$$\overline{\partial}_A\phi = 0. \tag{2.5b}$$

These equations are also invariant under the action of $\mathscr{G}$, so their solution set descends to a subset $\mathscr{M}_{\mathrm{vtx}}(L) \subseteq \mathscr{B}(L)$. This subset is called the *moduli space of vortices.*

Remark. It is important to note that even if the connection A is invariant under the action of Λ, the invariance of its associated holomorphic structure $\overline{\partial}_A$ is only guaranteed if Λ acts holomorphically on Σ. In particular, if the action of any element of Λ is antiholomorphic, the only Λ-equivariant holomorphic section of a line bundle is the zero section.

By integrating the first equation over Σ, one obtains an obstruction to the existence of vortices on L:

Proposition 2.15. *If $L \to \Sigma$ admits an irreducible vortex, then $\tau > 4\pi \deg(L)/\mathrm{Vol}(\Sigma)$. If $L \to \Sigma$ admits a reducible vortex, then $\tau = 4\pi \deg(L)/\mathrm{Vol}(\Sigma)$.*

Proof. Let (ϕ, A) be a vortex on L. Noting that $\int_\Sigma F_A = -2\pi i \deg(L)$, we find that the following holds:

$$-2\pi i \deg(L) = \frac{i}{2}\left(\|\phi\|^2_{L^2} - \tau \mathrm{Vol}(\Sigma)\right), \tag{2.6}$$

or equivalently

$$\tau = \frac{1}{\mathrm{Vol}(\Sigma)}\left(\|\phi\|^2_{L^2} + 4\pi \deg(L)\right). \tag{2.7}$$

Note that, since ϕ is holomorphic, $\phi^{-1}(0)$ is either finite if the vortex is irreducible, or all of Σ if the vortex is reducible. It follows that $\|\phi\|_{L^2} > 0$ for an irreducible vortex and $\|\phi\|_{L^2} = 0$ for a reducible vortex, which gives the desired obstructions. ☐

The vortex equations have a natural interpretation in terms of symplectic geometry. In particular, observe that the pre-configuration space is an affine space locally modelled on $\Gamma(L) \times \Omega^{0,1}(\Sigma)$, which can be naturally equipped with the following Kähler structure:

$$\begin{aligned} ((\phi,\alpha),(\psi,\beta)) &= \int_\Sigma \mathrm{Re}\, h(\phi,\psi) + \mathrm{Re}\, g^{\mathbb{C}}(\alpha,\beta), \\ \omega((\phi,\alpha),(\psi,\beta)) &= \int_\Sigma \mathrm{Im}\, h(\phi,\psi) + 2\,\mathrm{Im}\, g^{\mathbb{C}}(\alpha,\beta). \end{aligned} \tag{2.8}$$

The action of $\mathscr{G}$ on the pre-configuration space respects this Kähler structure, so $\mathbb{C}(L)$ has the structure of a Kähler manifold equipped with a symplectic and isometric action by a Lie group $\mathscr{G}$. The first vortex equation takes on special meaning in this interpretation:

Proposition 2.16. *The map $\mu_0 : \Gamma(L) \times \mathscr{A}(L) \to \mathfrak{g} \cong \mathfrak{g}^*$ defined by taking*

$$\mu_0(\phi, A) = *F_A - \tfrac{i}{2}|\phi|^2 \tag{2.9}$$

is a moment map[2] for the $\mathscr{G}$-action on $\Gamma(L) \times \mathscr{A}(L)$.

Proof. The map $\mu_{\mathscr{A}} : \mathscr{A}(L) \to \mathfrak{g}$ defined by taking $\mu_{\mathscr{A}}(A) = *F_A$ was shown to be a moment map for the $\mathscr{G}$-action on $\mathscr{A}(L)$ by [AB83] in the case that Λ acts trivially. The equivariant connections constitute a subspace of the space of all connections, and the equivariant gauge group is a Lie subgroup of the entire gauge group. A restriction of moment maps to a subspace and a subgroup is still a moment map, so we find that $\mu_{\mathscr{A}}$ is a moment map in general. The map $\mu_\Gamma : \Gamma(L) \to \mathfrak{g}$ defined by taking $\mu_\Gamma(\phi) = -\frac{i}{2}|\phi|^2$ is easily shown to be a moment map for the $\mathscr{G}$-action on $\Gamma(L)$. The moment map of the product of $\mathscr{G}$-spaces is simply the sum of the moment maps on each respective space [Mar+07]. ☐

[2] By a moment map for a symplectic G-action on a symplectic Hilbert manifold (M, ω), we mean here a G-equivariant map $\mu : M \to \mathfrak{g}^*$ satisfying $\mathrm{d}\langle \mu, \xi \rangle = \widetilde{\xi} \lrcorner \omega$ for every $\xi \in \mathfrak{g}$, where $\widetilde{\xi} \in \mathfrak{X}(M)$ denotes the infinitesimal group action of $\xi \in \mathfrak{g}$ on M, and $\langle \mu, \xi \rangle : M \to \mathbb{R}$ is the pairing $\langle \mu, \xi \rangle(p) = \mu(p)(\xi)$.

Thus, consider the submanifold[3] $\mathscr{N} \subseteq \mathscr{C}^*(L)$ given by pairs (ϕ, A) for which $\overline{\partial}_A\phi = 0$. This submanifold inherits the symplectic $\mathscr{G}$-action, and it therefore has the moment map $\mu_0|_{\mathscr{N}}$. This is not quite the first vortex equation for arbitrary τ, but it turns out that we can introduce the parameter τ by quotienting a particular subspace by the U(1)-action.

Proposition 2.17. *Given an arbitrary $\tau > 4\pi\deg(L)/\mathrm{Vol}(\Sigma)$, define $\mathscr{N}_\tau \subseteq \mathscr{N}$ to be the subspace for which $\|\phi\|^2_{L^2} = \tau\mathrm{Vol}(\Sigma) - 4\pi\deg(L)$. Then the moment map for $\mathscr{G}/\mathrm{U}(1)$ on $\mathscr{N}_\tau/\mathrm{U}(1)$ is the following:*

$$\mu_\tau(\phi, A) = *F_A - \tfrac{i}{2}(|\phi|^2 - \tau), \tag{2.10}$$

where $\mathrm{U}(1) \hookrightarrow \mathscr{G}$ *is the subgroup of constant gauge transformations.*

Proof. The direct product of two Lie group actions has a moment map given by the sum of the two corresponding moment maps [Mar+07], so it follows that the moment map for the quotient of a Lie group action can be obtained by subtracting the moment maps. On the space $\mathscr{N}_\tau$, the restriction of μ_0 to U(1) is given as follows:

$$\mu_0(\phi, A)|_{\mathrm{U}(1)} = \int_\Sigma (*F_A - \tfrac{i}{2}|\phi|^2)\mathrm{vol}_\Sigma = -2\pi i\deg(L) - \frac{i}{2}\|\phi\|^2_{L^2} = -\frac{i\tau\mathrm{Vol}(\Sigma)}{2} = \int_\Sigma -\frac{i\tau}{2}\mathrm{vol}_\Sigma.$$

This leads to the desired formula when subtracted from the moment map μ_0 on all of $\mathscr{G}$. □

Henceforth, we denote by $\widetilde{\mathscr{N}}$ the quotient $\mathscr{N}_\tau/\mathrm{U}(1)$, and by $\widetilde{\mathscr{G}}$ the quotient $\mathscr{G}/\mathrm{U}(1)$. Accordingly, we denote by $\widetilde{\mathscr{G}^{\mathbb{C}}}$ the quotient $\mathscr{G}^{\mathbb{C}}/\mathbb{C}^*$. We also drop the subscript on μ_τ.

We see, therefore, that the irreducible vortices on $L \to \Sigma$ are in one-to-one correspondence with zeros of the moment map μ on $\widetilde{\mathscr{N}}$. By the following two theorems, we can use this fact to reduce the work in finding solutions significantly.

Proposition 2.18. *Let $\mu : M \to \mathfrak{g}^*$ be a moment map for a free symplectic and isometric G-action on a Kähler manifold M, let $G^{\mathbb{C}}$ be the complexification of G, and let $S \subseteq M$ be a $G^{\mathbb{C}}$ orbit. Then μ has a zero on S if and only if there is some $p \in S$ for which $\inf_{q\in S}\|\mu(q)\|^2 = \|\mu(p)\|^2$.*

Proof. The original proof is due to Hitchin; see [Hit87]. Define $f : M \to [0, \infty)$ so that $f(p) = \langle\mu(p), \mu(p)\rangle$, and denote by $\nabla f \in \mathfrak{X}(M)$ the gradient vector field of f. Fix some $p \in M$. Given any $v_p \in T_pM$, observe the following computation:

$$g(\nabla f_p, v_p) = \mathrm{d}_{v_p} f = 2\langle \mathrm{d}_{v_p}\mu, \mu(p)\rangle = 2\mathrm{d}_{v_p}\langle\mu, \mu(p)\rangle = 2\omega(\widetilde{\mu(p)}, v_p) = 2g(J\widetilde{\mu(p)}, v_p). \tag{2.11}$$

By the nondegeneracy of g, it follows that $\nabla f_p = 2J\widetilde{\mu(p)}$. This means the gradient flow lines of f are confined to the orbits of $G^{\mathbb{C}}$.

Now, suppose $f|_S$ is minimised at $p \in S$. Then ∇f_p must be orthogonal to TS, which means $\nabla f_p = 0$ by the above computation. It follows that $\widetilde{\mu(p)} = 0$, meaning $\mu(p) \in \mathfrak{g}$ induces the trivial action on M. But since the action of G on M is free, this means $\mu(p) = 0$. □

Proposition 2.19. *If two irreducible vortices are in the same $\mathscr{G}^{\mathbb{C}}$-orbit then they are unitarily equivalent.*

[3] The regular value theorem allows us to prove that the holomorphic sections define a further submanifold.

Proof. Let (ϕ, A) and $g \cdot (\phi, A)$ be two pairs satisfying the vortex equations, where $g \in \mathscr{G}^{\mathbb{C}}$ is some complex gauge transformation. Observe that $F_{g \cdot A} = F_A$, meaning the following holds:

$$\frac{i}{2}(|\phi|^2 - \tau) = *F_A = *F_{g \cdot A} = \frac{i}{2}(|g \cdot \phi|^2 - \tau). \tag{2.12}$$

This implies that $|g||\phi| = |\phi|$; since ϕ is a holomorphic section of a line bundle, it is nonzero on a dense open subset of Σ. It follows that $|g| = 1$ almost everywhere, meaning g is unitary everywhere by continuity. □

The latter proposition implies that there is at most one vortex on each $\mathscr{G}^{\mathbb{C}}$-orbit, and the former proposition implies that a vortex can be found on each $\mathscr{G}^{\mathbb{C}}$-orbit by showing that μ achieves its infimum on each orbit.

2.3 Vortex Existence Proof

We now come to the proof of existence for Λ-equivariant vortices on orbifolds. We begin with the characterisation for irreducible vortices, and then briefly discuss reducible vortices. These results will subsequently be applied to specific cases.

2.3.1 Irreducible vortices

We now state precisely the existence theorem for irreducible Λ-equivariant vortices.

Theorem 2.20. *Let Λ be a finite group, and let L be a Λ-equivariant unitary line bundle with nonnegative degree over a Λ-orbifold Riemann surface Σ, on which Λ acts holomorphically. If $\tau > 4\pi\deg(L)/\mathrm{Vol}(\Sigma)$, then the moduli space of irreducible vortices is diffeomorphic to the space of $(\mathscr{G}^{\mathbb{C}})$ orbits in $\mathscr{N}$. This, in turn, is diffeomorphic to $S^b(\Sigma/\Lambda)$, where b is the background degree of L. If $\tau \leq 4\pi\deg(L)/\mathrm{Vol}(\Sigma)$, then the moduli space of irreducible vortices is empty.*

Before proving the theorem, we first outline the proof strategy in broad terms.

- First, we enlarge the spaces of smooth sections and connections to include sections and connections with Sobolev class L^2_1; the details for the constructions of these spaces are in the Appendix. We denote by $\mathscr{A}^1$ and $(\mathscr{G}^{\mathbb{C}})^2$ the space of L^2_1 connections and the space of L^2_2 gauge transformations. The essential motivation for this step is that weak limits are generally only well-behaved for Banach spaces, so the space of smooth sections and connections needs to be completed before the next step is possible.

- On each $(\widetilde{\mathscr{G}^{\mathbb{C}}})^2$-orbit of $\widetilde{\mathscr{N}}$, we choose a sequence of (representatives of) pairs converging to the infimum of μ over the orbit. We show that this sequence can be chosen to be weakly convergent to some (ϕ, A) not necessarily contained in the orbit.

- Since all elements of the sequence are in the same $(\mathscr{G}^{\mathbb{C}})^2$-orbit, we can write them in the form $g_n \cdot (\phi_0, A_0)$ for $g_n \in (\mathscr{G}^{\mathbb{C}})^2$. We show that the complex gauge transformations g_n can be chosen to be of the form $K_n e^{f_n}$, where $f_n \in L^2_2(\Sigma, \mathbb{C})^{\Lambda}$ are L^2_2-bounded and $K_n > 0$ are bounded above and below by strictly positive numbers. This allows us to choose a subsequence of g_n which converges weakly to some complex gauge transformation g, and therefore write $(\phi, A) = g \cdot (\phi_0, A_0)$.

- We then revert the enlargement of the pre-configuration space, by showing that every L^2_1 solution to the vortex equations is L^2_2-unitarily equivalent to a smooth one. This is done by first finding a local gauge transformation to a smooth solution, and then patching them together to find a global smoothing transformation.
- Finally, we show that the space of orbits $\widetilde{\mathscr{N}}/\widetilde{\mathscr{G}}^{\mathbb{C}}$ is diffeomorphic to the space of effective divisors with specified degree.

The proof of this theorem will require several technical results in analysis, generally pertaining to weak compactness. The first is an elementary result in the theory of Sobolev spaces.

Theorem 2.21. *The closed unit ball in L^p_k is weakly compact for every $k \in \mathbb{N}$ and every $p \in (1, \infty)$.*

Proof. By the Sobolev embedding theorems, we have a compact inclusion $L^p_k \hookrightarrow L^p$ for every k, p, and moreover L^p is reflexive for $p \in (1, \infty)$. By the Banach-Alaoglu theorem (which states that the unit ball of the dual of a Banach space is weakly compact; see [Lan93]), the space L^p_k is weakly compact. □

The second is a technical theorem demonstrating that the space of Sobolev connections with bounded curvature is weakly compact, but only up to unitary gauge equivalence.

Theorem 2.22 (Weak Uhlenbeck compactness for orbifolds)**.** *Let P be a principal* U(1)*-bundle over a compact 2-orbifold X. Let $(A^\nu)_{\nu\in\mathbb{N}} \subseteq \mathscr{A}(P)$ be a sequence of L^2_1 connections for which $\|F_{A^\nu}\|_{L^p}$ is uniformly bounded. Then there exists some subsequence of connections, also denoted by A^ν, and a sequence of L^2_2 gauge transformations $(u^\nu)_{\nu\in\mathbb{N}} \subseteq \mathscr{G}(P)$ for which the sequence $u^\nu(A^\nu)$ converges weakly in L^2_1.*

The third is another technical theorem demonstrating the existence of a special local trivialisation for a bundle-with-connection.

Theorem 2.23 (Uhlenbeck gauge for orbifolds)**.** *Let P be a principal* U(1)*-bundle over a compact 2-orbifold X, and let $A \in \mathscr{A}(P)$ be an L^2_1 connection on P. There exists a finite open covering $\{U_i\}_{i\leq n}$ of X trivialising P and a collection of local gauge transformations $u^i \in \mathscr{G}^2|_{U_i}$ such that $\mathrm{d}^*(u^i(A|_{U_i})) = 0$ for every i.*

The final result is elementary, but fundamental in demonstrating that the classical results generalise to the setting of Λ-equivariant orbifolds.

Proposition 2.24. *Weak and strong limits preserve invariance under a linear group action. That is, if V is a normed space equipped with a linear action by a group G, and $(v_n)_{n\in\mathbb{N}}$ is a sequence of G-invariant vectors which converge weakly or strongly to $v \in V$, then v is also G-invariant.*

Proof. The weak case implies the strong case, since strong convergence implies weak convergence. Let $g \in G$, and let $\phi \in V^*$ be an arbitrary continuous linear functional. Denote the group action by $\rho : G \to \mathrm{Aut}(V)$. Then, by the weak convergence of ϕ, we have the following:

$$\phi(v_n) = \phi(\rho_g(v_n)) = (\rho_g^*\phi)(v_n) \to (\rho_g^*\phi)(v) = \phi(g \cdot v). \tag{2.13}$$

Since weak limits are unique, it follows that $g \cdot v = v$. □

The proofs of Uhlenbeck's results for orbifolds are deferred to the end of the chapter.

We now begin the proof with the following lemma. In what follows, we relax our assumption that all objects are smooth; we instead assume that sections and connections have at least L^2_1 regularity, and that gauge transformations have at least L^2_2 regularity.

Lemma 2.25. *Every* $(\widetilde{\mathscr{G}}^{\mathbb{C}})^2$*-orbit* $S = (\widetilde{\mathscr{G}}^{\mathbb{C}})^2 \cdot [(\phi_0, A_0)] \subseteq \widetilde{\mathscr{N}}^1$ *admits a weakly convergent sequence of elements* $(\phi_n, A_n) = g_n \cdot (\phi_0, A_0)$ *with weak limit* (ϕ, A)*, such that* $\|\mu(\phi, A)\|^2 = \inf_{p \in S} \|\mu(p)\|^2$.

Proof. Let (ϕ_0, A_0) be a representative for a $(\widetilde{\mathscr{G}}^{\mathbb{C}})^2$-orbit in $\widetilde{\mathscr{N}}$; there is some sequence (ϕ_n, A_n) in the orbit for which $\|\mu(\phi_n, A_n)\|^2$ converges to the infimum of $\|\mu\|^2$ over the orbit. Since $\|\mu\|^2 = \mathrm{YMH}_\tau$ up to a constant term, and since $\|\mu(\phi_n, A_n)\|^2$ is bounded (by convergence), it follows that $\mathrm{YMH}_\tau(\phi_n, A_n)$ and hence $\|F_{A_n}\|^2_{L^2}$ are also uniformly bounded sequences. By Uhlenbeck's weak compactness theorem, we can therefore assume that A_n converges weakly in L^2_1 to some $A \in \mathscr{A}(L)$.

Since the ϕ_n are all in $\widetilde{\mathscr{N}}$, they all have the same L^2 norm. By the elliptic estimate, as well as the holomorphicity of each ϕ_n with respect to A_n, we have that $\|\phi_n\|_{L^2_1} \leq C_n \|\phi_n\|_{L^2}$ for constants C_n. Furthermore, by using the weak convergence of A_n, it can be shown that the C_n are uniformly bounded, meaning there is some $C > 0$ for which $\|\phi_n\|_{L^2_1} \leq C\|\phi_n\|_{L^2}$. Thus, we can assume that the sequence ϕ_n weakly converges in L^2_1 as well.

Finally, note that the weak limit is also Λ-invariant since weak limits commute with the action of Λ. □

Lemma 2.26. *Let* $g_n \cdot (\phi_0, A_0) \rightharpoonup (\phi, A)$ *be as above. The gauge transformations* g_n *can be chosen to be of the form* $g_n = K_n e^{f_n}$*, where the sequence* $K_n \in (0, \infty)$ *is strictly bounded below and above and the sequence* $f_n \in L^2_2(\Sigma, \mathbb{C})^\Lambda$ *is* L^2_2*-bounded.*

Proof. The holomorphic structures defined by A_0 and A_n are related by the addition of $g_n^{-1}\overline{\partial} g_n$, which is of class L^2_1. We first show that these $(0,1)$-forms can be assumed exact; each f_n will be defined to be the $\overline{\partial}$-primitive of these $(0,1)$-forms.

Since $g_n^{-1}\overline{\partial} g_n$ is $\overline{\partial}$-closed, it defines an element of $H^{0,1}(\Sigma)$; this is in turn isomorphic to $H^1(|\Sigma|; \mathscr{O})$ by Dolbeault's theorem, where $\mathscr{O}$ is the sheaf of holomorphic functions on Σ. The exponential sheaf sequence $\underline{\mathbb{Z}} \to \mathscr{O} \to \mathscr{O}^*$ induces the following exact sequence in sheaf cohomology:

$$H^1(|\Sigma|; \mathbb{Z}) \to H^1(|\Sigma|; \mathscr{O}) \to H^1(|\Sigma|; \mathscr{O}^*). \tag{2.14}$$

Now, if we choose a finite open cover $\mathscr{U}$ for Σ with contractible intersections, then we can identify each element of $H^1(|\Sigma|; \mathscr{O})$ with a map taking each nonempty intersection $U_i \cap U_j$ of open sets in $\mathscr{U}$ to a section of $\mathscr{O}$ over $U_i \cap U_j$. In particular, by direct computation, we can see that $g_n^{-1}\overline{\partial} g_n$ takes the intersection $U_i \cap U_j$ to the section[4] $\ln(g_n)$ (which exists uniquely up to addition of $2\pi i$ since $U_i \cap U_j$ is contractible); under the exponential map, this is taken to the trivial cocycle in $H^1(|\Sigma|; \mathscr{O}^*)$. By exactness, each $g_n^{-1}\overline{\partial} g_n$ corresponds to an element of $H^1(|\Sigma|; \mathbb{Z})$ under the inclusion.

[4]The fact that g_n is L^2_2 ensures that $\ln(g_n)$ is also L^2_2, and therefore continuous; this is part of Lemma A.4 in the appendix.

Using the open cover $\mathscr{U}$, we can give $H^1(|\Sigma|;\mathscr{O})$ a natural topology by taking the compact open topology on each $\Gamma(U_i \cap U_j;\mathscr{O})$. By choosing the harmonic representative of $g_n^{-1}\overline{\partial}g_n$, we see that the sequence of $(0,1)$-forms defines a Cauchy sequence in $H^1(|\Sigma|;\mathbb{Z})$. However, since Σ is closed and orientable, this is finitely generated and hence discrete, meaning the sequence is eventually constant. By applying a fixed unitary gauge transformation, we can ensure that every associated class is the zero class. Thus, we can assume from the outset that each g_n is $\overline{\partial}$-exact, and therefore that we can find constants $K_n > 0$ and L^2_2-functions $f_n : \Sigma \to \mathbb{C}$ for which $g_n = K_n e^{f_n}$ and $\int_\Sigma f_n = 0$. (The latter condition may be fulfilled because f_n is determined only up to an additive constant.)

Since $g_n^{-1}\overline{\partial}g_n = \overline{\partial}f_n$, we can use the elliptic estimate to ensure that

$$\|f_n\|_{L^2_2} \leq C(\|g_n^{-1}\overline{\partial}g_n\|_{L^2_1} + \|f_n\|_{L^2})$$

for some $C > 0$, and we can ignore the term $\|f_n\|_{L^2}$ if f_n is orthogonal to the kernel of $\overline{\partial}$. However, the kernel of $\overline{\partial}$ is the space of holomorphic functions, which are all constant since Σ is compact, and the orthogonality condition follows from our choice that $\int_\Sigma f_n = 0$. Now, since $g_n^{-1}\overline{\partial}g_n$ is uniformly bounded in L^2_1, this implies that the sequence f_n is uniformly bounded in L^2_2, and by the Sobolev embedding theorems it is also uniformly bounded in C^0. We therefore have the following inequalities for some $M > 0$:

Inequalities are out of order

$$|K_n|e^{-M} \leq |g_n| \leq |K_n|e^M. \tag{2.15}$$

Multiplying through by $|\phi_0|$ and taking the L^2-norm of each term in the inequalities, we see that

$$|K_n|e^{-M}\|\phi_0\|_{L^2} \leq \|\phi_n\|_{L^2} \leq |K_n|e^M\|\phi_0\|_{L^2}. \tag{2.16}$$

But since ϕ_0 and ϕ_n have the same L^2-norm, it follows that $|K_n|$ is itself uniformly bounded above and below. Additionally, by Lemma A.4, an L^2_2-bound on f_n gives an L^2_2-bound on $\exp(f_n)$.

We therefore find that the L^2_2-norms of the g_n are uniformly bounded, meaning we can find a weakly convergent subsequence of the g_n. This weak limit is a gauge transformation relating (ϕ_0, A_0) to the weak limit (ϕ, A), which achieves the infimum of $\|\mu\|^2$, and therefore satisfies the vortex equations. Additionally, as a weak limit of Λ-invariant gauge transformations, it is also Λ-invariant. □

Lemma 2.27. *Let (ϕ, A) be an L^2_1 solution to the vortex equations on the compact Λ-equivariant orbifold line bundle $L \to \Sigma$. Around each point, there exists a local Λ-equivariant L^2_2 unitary gauge transformation taking (ϕ, A) to a smooth solution of the vortex equations.*

Proof. The proofs of this lemma and the next are based on the approach in [FWX25] to establish regularity for the Hitchin equations. The idea behind the proof is as follows. First, we find a series of local smoothing transformations, by putting the connections into Uhlenbeck gauge and using elliptic bootstrapping techniques. Then, we patch together the local smoothing transformations to form a global one, by successively adding local charts and gluing them together with cutoff functions.

Since (ϕ, A) satisfies the vortex equations we know that $\|F_A\|_{L^2} = \frac{1}{2}\||\phi|^2 - \tau\|_{L^2}$, but the latter is bounded above by the Yang-Mills-Higgs functional. Given an open neighbourhood U of a given point, therefore, we can control the size of $F_A|_U$ by shrinking U. We therefore choose U sufficiently small that A may be put into Uhlenbeck gauge.

On a given U, we can write locally $A = \mathrm{d} + A_0$ where A_0 is an L^2_1 complex-valued Γ_i-equivariant 1-form on D^2, and we can interpret ϕ as an L^2_1 complex-valued Γ-equivariant function ϕ_0 on D^2 as long as we interpret the Hermitian metric h as a similar smooth function h_i. The vortex equations can then be stated locally as follows:

$$\overline{\partial}\phi_0 + A_0^{0,1}\phi_0 = 0; \tag{2.17}$$

$$\mathrm{d}A_0 = \frac{i}{2}(h_0|\phi_0|^2 - \tau)\mathrm{vol}_g. \tag{2.18}$$

Since $\overline{\partial}$ is elliptic and $A_0^{0,1}\phi_0$ is $L_1^{2-\delta}$ for every $\delta > 0$ (by the Sobolev multiplication theorem), we see by elliptic regularity that ϕ_0 is $L_2^{2-\delta}$. Moreover, taking the codifferential of both sides of the second equation, we see that $\mathrm{d}^*\mathrm{d}A_0 = \frac{i}{2}\mathrm{d}^*(h_0|\phi_0|^2\mathrm{vol}_g)$. By the Uhlenbeck gauge theorem for orbifolds, we see that there is a local gauge transformation g_0 under which $\mathrm{d}^*g_0(A_0) = 0$ on U; it follows that $\mathrm{d}^*\mathrm{d}g_0(A_0) = \Delta g_0(A_0)$, where Δ is the Laplacian (which is elliptic). We know that $\mathrm{d}(h_0|g_0(\phi_0)|^2)$ is of class $L^{2-\delta}$ for every δ, so elliptic regularity shows that $g_0(A_0)$ is $L_2^{2-\delta}$. By applying this method inductively, we see that $g_0(A_0)$ and $g_0(\phi_0)$ are $L_k^{2-\delta}$ for every k and every δ, meaning they must in fact be smooth. As such, we have found a local gauge transformation g_0 over U which takes (ϕ, A) to a smooth solution to the vortex equations. □

Lemma 2.28. *Let (ϕ, A) be an L^2_1 solution to the vortex equations on the compact Λ-equivariant orbifold line bundle $L \to \Sigma$. The local gauge transformations in Lemma 2.27 can be used to make a global Λ-equivariant L^2_2 unitary gauge transformation taking (A, ϕ) to a smooth solution of the vortex equations.*

Proof. First, we choose appropriate open sets. Given some $\varepsilon > 0$, we can choose a finite atlas $\{(U_i, \Gamma_i, \varphi_i)\}_{i=1,\dots,n}$ for Σ satisfying the following conditions:

- each U_i is an open set with closure diffeomorphic to D^2/Γ_i;
- the line bundle $L|_{U_i}$ is equivalent to a standard line bundle of the form $(D^2 \times \mathbb{C})/\Gamma_i$;
- each $\|F_A|_{U_i}\|_{L^2}$ is bounded above by ε;
- the charts U_i and the line bundles $L|_{U_i}$ are permuted equivariantly by the Λ-action;[5] we will say that U_i and U_j are Λ-*related* if there is some $g \in \Lambda$ for which $g(U_i) = U_j$. We furthermore assume that the atlas has been ordered so that Λ-related charts are indexed consecutively, meaning that if U_i and U_{i+m} are Λ-related for $m > 0$ then U_i and U_{i+k} are also Λ-related for all $0 < k < m$.

From the above lemma, we get local gauge transformations g_i on each U_i making (ϕ, A) smooth locally. Additionally, since (ϕ, A) is Λ-equivariant, we can assume that the gauge transformations g_i and g_j for two Λ-related charts U_i and U_j are also Λ-related.

Next, we modify the gauge transformations $g_i : \widetilde{U}_i \to \mathrm{U}(1)$ so that they can be glued together to form a global smoothing transformation. (Recall that we are conceptualising local gauge transformations over U_i as Γ_i-equivariant maps from $\widetilde{U}_i$ to $\mathrm{U}(1)$.) First, we will ensure that the g_i are sufficiently small that they can be written in terms of $\exp : i\mathbb{R} \to \mathrm{U}(1)$.

[5] It is always possible to find an atlas satisfying this condition; we can simply choose an open cover of Σ/Λ and lift each open set to Σ.

Since C^∞ is dense as a subset of L^2_2, we can always find Γ_i-equivariant smooth functions $h_i : \widetilde{U}_i \to \mathrm{U}(1)$ that make $\|h_i - g_i^{-1}\|_{L^2_2}$ as small as we like. It follows that, by replacing g_i with $g_i h_i$, we can also choose the g_i to be arbitrarily close to the constant function $1 : \widetilde{U}_i \to \mathrm{U}(1)$ in L^2_2.

We now glue together the g_i to form a global smoothing transformation. The strategy is to find a global gauge transformation $\widetilde{g}$ for which $\widetilde{g} \cdot g_i^{-1}$ is smooth for all i; this ensures that $\widetilde{g} \cdot g_i^{-1} \cdot (g_i(\phi, A))$ is smooth. We construct $\widetilde{g}$ inductively on the orbifold charts, but we will find that the charts must successively shrink. Thus, we make the following definition: where $d : \Sigma \times \Sigma \to [0, \infty)$ is the metric on Σ inherited from the orbifold metric g, and for $\delta > 0$, the open subsets $U_i^\delta \subseteq \Sigma$ are defined as follows:

$$U_i^\delta = \{p \in U_i \,:\, d(p, \partial U_i) > \delta\}\,. \tag{2.19}$$

By the Λ-invariance of the metric, these open subsets are still Λ-related to each other. Furthermore, if δ is small enough to ensure that the open sets U_i^δ still cover Σ, these open sets still correspond to an orbifold atlas; the orbifold chart mappings $\varphi_i : \widetilde{U}_i/\Gamma_i \to U_i$ can simply be restricted to cover U_i^δ, and it is easy to verify that a chart embedding $U_i \hookrightarrow U_j$ restricts to an embedding $U_i^\delta \hookrightarrow U_j^\delta$. We choose $\delta > 0$ small enough that the $U_i^{n\delta}$ still cover Σ, where n is the number of orbifold charts. At this point, we also reorder the sets U_i so that the following condition is met:

- If $U_{k+1}^{(k+1)\delta} \cap (\cup_{i\le k} U_i^{(k+1)\delta})$ is nonempty, and if U_j is another orbifold chart which embeds into U_{k+1}, then $j \le k$.

This condition ensures that when $U_{k+1}^{(k+1)\delta}$ is added to the domain of $\widetilde{g}$, any points already in the domain of $\widetilde{g}$ are in charts that can be embedded into $U_{k+1}^{(k+1)\delta}$.

For each $k = 1, \ldots, n$, a series of local gauge transformations $\widetilde{g}_i^{(k)}$ on each $U_i^{k\delta}$ for $i = 1, \ldots, k$ is defined inductively in terms of the g_i as follows:

- The gauge transformation $\widetilde{g}^{(1)}$ is defined on U_1^δ to be $g_1|_{U_1^\delta}$; note that $\widetilde{g}^{(1)} \cdot g_1^{-1}$ is clearly smooth. We also need not adjust for Λ-equivariance; the charts are permuted by Λ, and if $\Lambda \cdot U_1 = U_1$ then g_1 is already Λ-equivariant.
- Assume the local gauge transformation $\widetilde{g}^{(k)}$ has been defined on $\cup_{i\le k} U_i^{k\delta}$ and $\widetilde{g}^{(k)} g_i^{-1}$ is smooth for every $i \le k$. We now define the new local gauge transformations $\widetilde{g}_i^{(k+1)}$ on each $U_i^{(k+1)\delta}$ for every $i \le k+1$, and the definition depends on two things: which charts U_{k+1} is Λ-related to, and whether the new open set $U_{k+1}^{(k+1)\delta}$ is disjoint from the remaining open sets.
  - If U_{k+1} is Λ-related to U_j via $\gamma \in \Lambda$ for some $j \le k$, we define $\widetilde{g}_i^{(k+1)} = \widetilde{g}_i^{(k)}$ for $i \le k$ and $\widetilde{g}_{k+1}^{(k+1)} = (\widetilde{g}_k^{(k)})^\gamma$. Note that this definition does not depend on the choice of j by the ordering of the U_i.
  - Assume U_{k+1} is not Λ-related to any U_j for $j \le k$. If $U_{k+1}^{(k+1)\delta}$ and $\cup_{i\le k} U_i^{(k+1)\delta}$ are disjoint, we define $\widetilde{g}_i^{(k+1)}$ as follows:

$$\widetilde{g}_i^{(k+1)} = \begin{cases} \widetilde{g}^{(k)}|_{U_i^{(k+1)\delta}} & \text{for } i \le k \\ g_{k+1}|_{U_{k+1}^{(k+1)\delta}} & \text{for } i = k+1. \end{cases} \tag{2.20}$$

These local gauge transformations clearly glue together to a gauge transformation $\widetilde{g}^{(k+1)}$ on all of $\cup_{i\leq k+1}U_i^{(k+1)\delta}$ (since we have only added a single disjoint chart), and also $\widetilde{g}^{(k+1)}g_i^{-1}$ is smooth for every $i \leq k+1$.
Again, Λ-equivariance is not an issue: either there is some $\gamma \in \Lambda$ for which $\gamma \cdot U_{k+1} = U_{k+1}$, in which case g_{k+1} is also invariant under γ; or there is no such γ, in which case the subsequent steps of the induction will provide Λ-equivariance.

- Assume that $U_{k+1}^{(k+1)\delta}$ is not disjoint from $\cup_{i\leq k}U_i^{(k+1)\delta}$, and additionally that U_{k+1} is not Λ-related to itself. The latter assumption allows us to ignore Λ-equivariance for $\widetilde{g}_{k+1}^{(k+1)}$, since it will be accounted for in the next inductive step. We now use a cutoff function to interpolate between $\widetilde{g}^{(k)}$ on $\cup_{i\leq k}U_i^{k\delta}$ and g_{k+1} on the remainder of $U_{k+1}^{(k+1)\delta}$. That is, we use a function $\phi_{k+1} : U_{k+1}^{(k+1)\delta} \to [0,1]$ which is 0 on $U_{k+1}^{(k+1)\delta} \cap (\cup_{i\leq k}U_i^{(k+1)\delta})$ and 1 on $U_{k+1}^{(k+1)\delta} \backslash (\cup_{i\leq k}U_i^{k\delta})$.
Thus, we define $\widetilde{g}_i^{(k+1)} = \widetilde{g}_i^{(k)}$ for $i \leq k$, and we define $\widetilde{g}_{k+1}^{(k+1)}$ as follows:[6]

$$\widetilde{g}_{k+1}^{(k+1)} = \begin{cases} \exp\left(\phi_{k+1}\cdot \ln(g_{k+1}(\widetilde{g}_i^{(k)})^{-1})\right)\widetilde{g}_i^{(k)} & \text{on } U_{k+1}^{(k+1)\delta}\cap U_i^{k\delta} \text{ for } i\leq k \\ g_{k+1} & \text{otherwise.}\end{cases} \tag{2.21}$$

If U_i and U_j overlap then, by the inductive hypothesis, we know that $\widetilde{g}_i^{(k)}$ and $\widetilde{g}_j^{(k)}$ coincide. Moreover, the ln operator is well-defined since we chose the g_i to be sufficiently close to the identity. Thus, $\widetilde{g}_{k+1}^{(k+1)}$ is well-defined. Additionally, since $\phi_{k+1} = 0$ on the domain of $\widetilde{g}^{(k)}$, the local gauge transformations $\widetilde{g}_i^{(k+1)}$ glue together into a global gauge transformation.
We now show that $\widetilde{g}_i^{(k+1)}g_j^{-1}$ is smooth for any $i, j \leq k+1$: for $i \leq k$ it follows directly from the inductive hypothesis, and for $i = k+1$ we observe the following on $U_{k+1}^{(k+1)\delta} \cap U_j^{k\delta}$:

$$\widetilde{g}_{k+1}^{(k+1)}g_j^{-1} = \exp\left(\phi_{k+1}\cdot \ln\left((g_{k+1}g_j^{-1})(g_j(\widetilde{g}_i^{(k)})^{-1})\right)\right)(\widetilde{g}_i^{(k)}g_j^{-1}). \tag{2.22}$$

But by the inductive hypothesis, the terms $g_j(\widetilde{g}_i^{(k)})^{-1}$ and $\widetilde{g}_i^{(k)}g_j^{-1}$ are both smooth. Additionally, $g_{k+1}g_j^{-1}$ transforms the smooth connection $g_j(A_j)$ into the smooth connection $g_{k+1}(A_{k+1})$, and it is shown in [AB83] (Lemma 14.9) that such a gauge transformation must be smooth. It follows that $\widetilde{g}_{k+1}^{(k+1)}$ is smooth on $U_{k+1}^{(k+1)\delta} \cap U_j^{k\delta}$ and therefore on all of $U_{k+1}^{(k+1)\delta}$.

- Finally, if $U_{k+1}^{(k+1)\delta}$ is not disjoint from $\cup_{i\leq k}U_i^{(k+1)\delta}$ and U_{k+1} is nontrivially Λ-related to itself, then we use the same definition of $g_{k+1}^{(k+1)}$ given in Equation 2.21. We just need to prove that it is equivariant under $\Lambda_{k+1} \leq \Lambda$, the subgroup of Λ stabilising U_{k+1}. Before we do this, we need to make the smooth cutoff function ϕ_{k+1} satisfy Λ_{k+1}-equivariance; this is easily done by redefining ϕ_{k+1} to

[6] Here each $\widetilde{g}_i^{(k)} : \widetilde{U}_i^{k\delta} \to \mathrm{U}(1)$ is interpreted as map taking a Γ_{k+1}-equivariant subset of $U_{k+1}^{(k+1)\delta}$ to Λ; this is possible because we assumed that the first k charts are either disjoint from or embed into the $(k+1)$-th chart. Note that this interpretation makes $\widetilde{g}_i^{(k)}$ a Γ_{k+1}-equivariant map. We also interpret ϕ_{k+1} as a Γ_{k+1}-invariant smooth map from $\widetilde{U}_{k+1}^{(k+1)\delta}$ to $[0,1]$; together, these ensure that the definition of $\widetilde{g}_{k+1}^{(k+1)}$ is Γ_{k+1}-equivariant.

be $\frac{1}{|\Lambda_{k+1}|}\sum_{\gamma\in\Lambda_{k+1}}\gamma^*\phi_{k+1}$. From here, however, the Λ-equivariance follows from the induction: if U_i and U_j are Λ-related for $i, j \leq k$ then we have inductively chosen $\widetilde{g}_i^{(k)}$ and $\widetilde{g}_j^{(k)}$ to be Λ-related, meaning we have the following on the overlap for every $\gamma \in \Lambda$:

$$\begin{aligned}\gamma^*\widetilde{g}_{k+1}^{(k+1)} &= \exp\left(\gamma^*\phi_{k+1}\cdot\ln(\gamma^* g_{k+1}(\gamma^*\widetilde{g}_i^{(k)})^{-1})\right)\gamma^*\widetilde{g}_i^{(k)}\\ &= \exp\left(\phi_{k+1}\cdot\ln(g_{k+1}(\widetilde{g}_j^{(k)})^{-1})\right)\widetilde{g}_j^{(k)} = \widetilde{g}_{k+1}^{(k+1)},\end{aligned}$$

where the last equality follows from the fact that $\widetilde{g}_{k+1}^{(k+1)}$ is defined independently of which local representative of $\widetilde{g}^{(k)}$ is used. It follows that the definition is Λ_{k+1}-equivariant, and the induction extends this equivariance to all of Λ.

By induction, we obtain a global Λ-equivariant gauge transformation $\widetilde{g}^{(n)}$ for which $\widetilde{g}^{(n)}g_i^{-1}$ is smooth for all i. But then, since $g_i(\phi, A)$ is smooth, $\widetilde{g}^{(n)}(\phi, A)$ is smooth everywhere. □

Lemma 2.29. *The space of $\widetilde{\mathscr{G}}^{\mathbb{C}}$-orbits of $\widetilde{\mathscr{N}}$ is diffeomorphic to $S^b(\Sigma/\Lambda)$.*

Proof. Firstly, it is clear that the space of $\widetilde{\mathscr{G}}^{\mathbb{C}}$-orbits of $\widetilde{\mathscr{N}}$ is the same as the space of $\mathscr{G}^{\mathbb{C}}$-orbits of $\mathscr{N}$. Given a representative (ϕ, A) of such an orbit, the corresponding symmetric b-tuple of points is given by the zeros of ϕ; note that there will be precisely $b|\Lambda|$ zeros of ϕ on Σ (up to multiplicity), since ϕ is a holomorphic section of a Λ-equivariant line bundle with background degree b over Σ, and they will all come in sets of size $|\Lambda|$ which descend to Σ/Λ. Since $\mathscr{G}^{\mathbb{C}}$ acts smoothly and equivariantly, this tuple is preserved by the action of $\mathscr{G}^{\mathbb{C}}$. Given two sections ϕ and ϕ' which define the same tuple of points, we define a corresponding complex gauge transformation relating them to be the unique number g at each point for which $\phi' = g\phi$; this defines g almost everywhere since ϕ' and ϕ are holomorphic, and the definition can be extended continuously since the zeros of ϕ' and ϕ have the same order at each point. □

Proof of Theorem 2.20. There is an L^2_1 solution to the vortex equations on every orbit of $(\widetilde{\mathscr{G}})^{\mathbb{C}}$ by the first and second lemma, and this is L^2_2 gauge equivalent to a smooth solution by the third and fourth. Thus, the moduli space of vortices is simply given by the quotient of the configuration space by the complexified gauge group, which was noted in the preceding lemma to be $S^b(\Sigma/|\Lambda|)$. □

2.3.2 Reducible vortices

The discussion of reducible vortices is much simpler:

Theorem 2.30. *The moduli space of reducible Λ-equivariant vortices on a line bundle L over a compact orbifold Riemann surface Σ is nonempty if and only if $\tau = 4\pi\deg(L)/\mathrm{Vol}(\Sigma)$. In this case, it is homeomorphic to the Jacobian torus of Σ/Λ.*

Proof. We first assume the Λ-action is trivial. When $\phi = 0$, the vortex equations reduce to $*F_A = -i\tau/2$, and as shown in Proposition 1.50, any line bundle over an orbifold Riemann surface admits a connection whose curvature is a constant multiple ξ of the volume form. The constraint on ξ is equivalent to the requirement on τ stated in the theorem. Any two connections with the same curvature differ by tensoring with a flat line bundle, and conversely tensoring with a flat line bundle does not change the curvature, so the space of reducible vortices is a torsor on the space of trivial flat bundles. This was shown to be homeomorphic to the Jacobian torus of Σ in Proposition 1.38.

The analogous statement for a nontrivial Λ-action can be recovered by repeating the above argument for Σ/Λ, which is an orbifold. □

2.3.3 Vortices on Klein orbifolds

The above results provide insight for non-orientable orbifolds.

Theorem 2.31. *Let Σ be a non-orientable 2-orbifold. The moduli space of irreducible vortices on Σ is always empty, and the moduli space of reducible vortices is nonempty if and only if $\tau = 4\pi\deg(L)/\mathrm{Vol}(\Sigma)$, in which case it is homeomorphic to the Jacobian torus of Σ.*

Proof. If Σ is non-orientable, then there is some connected Klein orbifold (Σ_0, τ) for which $\Sigma_0/\tau \cong \Sigma$; this is equivalent to a $\mathbb{Z}_2$-orbifold Riemann surface for which the $\mathbb{Z}_2$-action is antiholomorphic. However, a vortex (ϕ, A) on the Real line bundle L must satisfy $\overline{\partial}_A\phi = 0$ and also must be τ-invariant. It follows that if ϕ is holomorphic at one point then it is antiholomorphic at the corresponding point under τ; since Σ_0 is connected, this implies that $\phi = 0$. Thus, a non-orientable orbifold admits no irreducible vortices.

The claim regarding reducible vortices follows immediately from Theorem 2.30. □

2.3.4 Kähler vortices

When we come to discuss the Seiberg-Witten equations on Seifert fibred spaces, the vortex equations will appear as a dimensional reduction. However, they will be in a slightly different form.

Definition 2.32. Let $L \to \Sigma$ be a Λ-equivariant line bundle over an orbifold Riemann surface with a constant-curvature metric, let $(\alpha, \beta) \in \Gamma(L) \times \Gamma(K_\Sigma^{-1} \otimes L)$ be a pair of Λ-equivariant sections, and let $A \in \mathscr{A}(L)$ be a Λ-equivariant unitary connection. Then (α, β, A) satisfy the Kähler vortex equations if the following equations hold:

$$2F_A - F_{K_\Sigma} = i(|\alpha|^2 - |\beta|^2)\mathrm{vol}_\Sigma, \tag{2.23a}$$

$$\overline{\partial}_A\alpha = 0 \text{ and } \overline{\partial}_A^*\beta = 0, \tag{2.23b}$$

$$\alpha = 0 \text{ or } \beta = 0. \tag{2.23c}$$

(Here F_{K_Σ} refers to the Levi-Civita connection on the canonical bundle.) We call the vortex *positive* if α is nonzero, *negative* if β is nonzero, and *reducible* if both α and β are zero. The *moduli space of positive/negative vortices* is the space of solutions up to unitary gauge equivalence, and is denoted by $\mathscr{M}_{\mathrm{vtx}}^{\pm}(L)$.

If $\beta = 0$ (i.e., the vortex is positive or reducible), then F_{K_Σ} is a constant multiple of the volume form; defining this constant multiple so that $F_{K_\Sigma} = -i\tau\mathrm{vol}_\Sigma$, we recover the usual vortex equations. Additionally, we can understand the negative vortices as positive vortices over different bundles:

Proposition 2.33. *If (β, A) defines a negative solution to the Kähler vortex equations on L, then $(*\beta, A^* \otimes \omega^{\mathrm{SO}})$ defines a positive solution to the Kähler vortex equations on $L^* \otimes K_\Sigma$.*

Proof. Observe that the Hodge star defines an isomorphism $* : \Gamma(L \otimes K_\Sigma^{-1}) \to \Gamma(\overline{L \otimes K_\Sigma^{-1}} \otimes \Lambda^{1,1})$; the conjugation arises from the fact that it is antilinear. However, the Hermitian structure on L and K_Σ induce isomorphisms between the conjugate and the dual, meaning we have the following sequence of isomorphisms:

$$L \otimes K_\Sigma^{-1} \overset{*}{\leftrightarrow} \overline{L \otimes K_\Sigma^{-1}} \otimes \Lambda^{1,1} \overset{h\otimes\Lambda}{\longleftrightarrow} L^* \otimes K_\Sigma. \tag{2.24}$$

(Here $\Lambda : \Lambda^{1,1} \to \Lambda^0$ is the contraction with the volume form, not the finite group.) It follows that $*\beta$ is a section of $L^* \otimes K_\Sigma$.

Furthermore, we know that $\overline{\partial}_A^* \beta = 0$, meaning $\partial_A \overline{*\beta} = 0$. Since the holomorphic 1-forms are precisely the 1-forms in the kernel of $\overline{\partial}$, we see that $\overline{\partial}$ is the holomorphic structure of K_Σ, and therefore that ∂ is the holomorphic structure of K_Σ^{-1}. It follows from the tensor product that ∂_A is the holomorphic structure of $L \otimes K_\Sigma^{-1}$, so $\overline{*\beta} \in \Gamma(L \otimes K_\Sigma^{-1})$ is holomorphic if and only if $\partial_A \overline{*\beta} = 0$. This means $*\beta$ is holomorphic, which makes it correspond to a positive vortex on $L^* \otimes K_\Sigma$. □

Corollary 2.34. *The moduli space of Λ-equivariant positive vortices on L is equivalent to the moduli space of Λ-equivariant negative vortices on $L^* \otimes K_\Sigma$. This is nonempty if and only if $0 \leq \deg(L) < \deg(K_\Sigma)/2 = -\chi(\Sigma)/2$, in which case it is equivalent to $S^b(\Sigma/\Lambda)$.*

Corollary 2.35. *The moduli space of Λ-equivariant reducible Kähler vortices on L is equivalent to the moduli space of Λ-equivariant reducible vortices on L with $\tau = 2\pi \deg(K_\Sigma)/\mathrm{Vol}(\Sigma)$. It is nonempty if and only if $\deg(L) = \deg(K_\Sigma)/2$, in which case it is equivalent to the Jacobian torus of Σ/Λ.*

2.4 Uhlenbeck's Theorems for Orbifolds

We now prove the extensions of Uhlenbeck's results to the orbifold category. In characterising the moduli space of vortices, the only relevant case will be for L_1^2 connections and L_2^2 gauge transformations on a U(1)-bundle over a 2-orbifold whose local groups are all cyclic; nevertheless, we will prove the theorems in full generality. We start with Uhlenbeck's gauge theorem and use it to prove Uhlenbeck's weak compactness theorem.

2.4.1 Uhlenbeck's gauge theorem

Before stating the theorem explicitly, we make some definitions. In what follows, X will be an n-orbifold, G will be a compact Lie group with Lie algebra $\mathfrak{g}$, $P \to X$ will be a principal G-bundle over X, and U will be a subset of X diffeomorphic to B/Γ, where $B \subseteq \mathbb{R}^n$ is the closed unit n-ball and Γ is a finite group acting linearly on B. The connection A that we refer to will always be L_1^p for some $p \in (1, \infty)$. We denote the collection of L_1^p-connections by $\mathscr{A}^{1,p}$, and we denote by $\mathscr{G}^{2,p}$ the collection of L_2^p-gauge transformations.

First, given some $q \in [1, \infty]$ and some connection A on a bundle over a Riemannian manifold X, we define the *q-energy of A* to be the (qth power of) the L^q-norm of its curvature:

$$\mathscr{E}(A) = \|F_A\|_{L^q}^q = \int_X |F_A|^q \mathrm{vol}_X. \tag{2.25}$$

Over the subset $U \cong B/\Gamma$, on which we can represent a connection by a Γ-equivariant $\mathfrak{g}$-valued 1-form on B, we say that a connection $A \in \Omega^1(U, \mathfrak{g})$ is in *Coulomb gauge* if it satisfies the following differential equation and boundary condition:

$$\mathrm{d}^* A = 0, \tag{2.26}$$

$$*A|_{\partial U} = 0. \tag{2.27}$$

Additionally, given some $q \in (1, p]$ and some $C \geq 0$, we say that A is in *q-Uhlenbeck gauge with constant C* if it is in Coulomb gauge and satisfies the following inequalities:

$$\|A\|_{L_1^q} \leq C\|F_A\|_{L^q}; \tag{2.28}$$

$$\|A\|_{L_1^p} \leq C\|F_A\|_{L^p}. \tag{2.29}$$

The constant C and the regularity q are omitted if they are irrelevant or clear from context, so the connection A is simply said to be in Uhlenbeck gauge.

With these definitions, the Uhlenbeck gauge theorem for orbifolds is the following:

Theorem 2.36 (Uhlenbeck Gauge Theorem)**.** *Let X be a Riemannian n-orbifold and G a compact Lie group, and equip X with a principal G-bundle P. Let $p, q \in (1, \infty)$ be real numbers satisfying the constraints that $q \leq p$, $q \geq n/2$, and $p > n/2$, as well as the constraint that $p \leq \frac{nq}{n-q}$ if $q < n$. There exists some $\varepsilon > 0$ and some $C > 0$ for which the following holds:*

Every point in X has a neighbourhood U such that any L_1^p-connection A on P satisfying $\mathscr{E}(A|_U) \leq \varepsilon$ admits an L_2^p-gauge transformation u for which $u(A)$ is in q-Uhlenbeck gauge with constant C.

In essence, the theorem states that we can always locally transform a sufficiently regular connection into Uhlenbeck gauge over a sufficiently small neighbourhood. To this end, we will need to prove the special case of this theorem for the model case:

Theorem 2.37 (Uhlenbeck Gauge Theorem, model case)**.** *Let $B \subseteq \mathbb{R}^n$ be the closed unit n-ball equipped with a linear action by a finite group Γ, let $\mathfrak{g}$ be a Lie algebra, and let $B \times \mathfrak{g}$ be a vector bundle with a lift of the Γ-action given by a representation $\rho : \Gamma \to \mathfrak{g}$. Let $p, q \in (1, \infty)$ be real numbers satisfying the constraints that $q \leq p$, $q \geq n/2$, and $p > n/2$, as well as the constraint that $p \leq \frac{nq}{n-q}$ if $q < n$. There exists some $\varepsilon > 0$, some $\delta > 0$, and some $C > 0$ for which the following holds:*

If B is equipped with a smooth metric g for which $\|g - 1\|_{L_2^\infty} \leq \delta$, then any Γ-equivariant L_1^p-connection A on $B \times \mathfrak{g}$ satisfying $\mathscr{E}(A) \leq \varepsilon$ admits a Γ-equivariant L_2^p-gauge transformation u for which $u(A)$ is in q-Uhlenbeck gauge with constant C.

Our proofs of each of these results are essentially direct adaptations of Wehrheim's proofs to the case of orbifolds. We will state several results without proof, and notably much of the proof that the relevant connections are in Uhlenbeck gauge will be omitted.[7] Note also that Wehrheim's statement is slightly more general and accounts for manifolds-with-boundary.

[7]From a broad perspective, this is because Wehrheim has already constructed the relevant connections and shown that they are in Uhlenbeck gauge; our contribution is merely to show that every construction is also Γ-equivariant, to ensure that they are well-defined under the orbifold quotient.

In order to prove the model case, we define a *modified q-energy* $\mathscr{E}' : \mathscr{A}^{1,p} \to [0,\infty)$ as follows:[8]

$$\mathscr{E}'(A) = \frac{1}{|\Gamma|}\int_B |F_A|_g^q \,\mathrm{d}^n x \tag{2.30}$$

where $|\cdot|_g$ denotes the norm of a tensor with respect to the metric g and a suitable inner product on $\mathfrak{g}$. Given $\varepsilon > 0$ and $C > 0$, we define two topological spaces of connections:

$$\mathscr{A}_\varepsilon = \{A \in \mathscr{A}^{1,p}(B) \,:\, \mathscr{E}'(A) \le \varepsilon\}; \tag{2.31}$$

$$\mathscr{S}_{\varepsilon,C} = \{A \in \mathscr{A}_\varepsilon \,:\, \text{there is some } u \in \mathscr{G}^{2,p} \text{ for which } u(A) \text{ is in } q\text{-Uhlenbeck gauge with constant } C\}. \tag{2.32}$$

The topology on $\mathscr{A}_\varepsilon$ and $\mathscr{S}_{\varepsilon,C}$ is inherited from the L^p_1-norm on $\mathscr{A}^{1,p}$. Note that $\mathscr{S}_{\varepsilon,C}$ is a nonempty subset of $\mathscr{A}_\varepsilon$, since the trivial connection is in Uhlenbeck gauge no matter what q and C are and has zero energy (modified or not). The strategy of the proof is therefore to show that $\mathscr{A}_\varepsilon$ is connected and that $\mathscr{S}_{\varepsilon,C}$ is a closed and open subset for some C.

We now break the connectedness argument into three key lemmas.

Lemma 2.38. *$\mathscr{A}_\varepsilon$ is connected. That is, any $A \in \mathscr{A}_\varepsilon$ can be connected to the trivial connection by a path which is continuous in the topology inherited from the L^p_1-norm.*

Proof. Define a path of connections A_t for $t \in [0,1]$ as follows: for every $x \in D^2$,

$$A_t(x) = tA(tx). \tag{2.33}$$

It is clear that $A_0 = 0$ and $A_1 = A$, and moreover, A_t is clearly of class L^p_1 for every t. By the linearity of the Γ-action on D^2, the connections are also all Γ-equivariant. Moreover, the modified q-energy of A_t is always bounded above by ε when $q \ge 1$: it is fairly easy to show that $F_{A_t}(x) = t^2 F_A(tx)$, meaning the energy of A_t is bounded above by $t^{2q-n}\mathscr{E}'(A)$, which in turn is bounded above by $\mathscr{E}'(A) < \varepsilon$ since we assumed $q \ge n/2$. (It is this point at which the modification to $\mathscr{E}$ is necessary.) This path of connections is continuous; for the proof of this fact, see [Weh04]. ∎

Lemma 2.39. *The subset $\mathscr{S}_{\varepsilon,C}$ is always closed. In more detail, let g be a smooth Γ-equivariant metric on B and let $\varepsilon > 0$. Let $(A_i)_{i\in\mathbb{N}} \in \mathscr{A}_\varepsilon$ be a sequence of Γ-equivariant connections converging to the Γ-equivariant connection $A \in \mathscr{A}_\varepsilon$ in the L^p_1-topology, and suppose there exists a sequence of Γ-equivariant L^p_2-gauge transformations $(u_i)_{i\in\mathbb{N}} \in \mathscr{G}^{2,p}$ for which $u_i(A_i)$ is in q-Uhlenbeck gauge with constant C. Then there is a Γ-equivariant L^p_2 gauge transformation u for which $u(A)$ is also in q-Uhlenbeck gauge with the same constant C.*

Proof. Since each $u_i(A_i)$ is in Uhlenbeck gauge, we know that $\|u_i(A_i)\|_{L^p_1} \le C\|F_{A_i}\|_{L^p}$. Moreover, there is a uniform upper bound on $\|F_{A_i}\|_{L^p}$ by the L^p_1-convergence of the A_i (see A.11 in [Weh04]), meaning the $u_i(A_i)$ are uniformly bounded as well. By the Banach-Alaoglu theorem, there is a subsequence of the $u_i(A_i)$ which converges weakly to some L^p_1

[8] Note that this is not the same as the energy $\mathscr{E}$ defined above; the metric g is being used to take the pointwise norm of F_A, but the Euclidean metric is being used to integrate this norm. Consequently, the energy $\mathscr{E}$ is modulated by an extra factor of $\sqrt{\det(g)}$ in the integrand. We can nevertheless control the deviation of g from the Euclidean metric 1 with the parameter δ, so $\mathscr{E}'$ can be made as close to $\mathscr{E}$ as we like by shrinking δ. As such, we will treat bounds on $\mathscr{E}$ and bounds on $\mathscr{E}'$ as equivalent.

connection $\widetilde{A}$, and by the Sobolev inequalities, this sequence can be chosen to converge in the L^{2p}-norm. Since weak limits commute with the group action, $\widetilde{A}$ is Γ-equivariant.

Additionally, there is a subsequence of the gauge transformations u_i which converges in the C^0-topology to some L^p_2 gauge transformation u, which is also invariant under Γ since strong limits commute with the group action, with the property that $u_i^{-1}\mathrm{d}u_i$ converges in the L^{2p}-topology to $u^{-1}\mathrm{d}u$. On the other hand, since $u_i(A_i) = u_i^{-1}A_iu_i + u_i^{-1}\mathrm{d}u_i$, it follows that $u_i(A_i)$ converges in L^{2p} to $u^{-1}Au + u^{-1}\mathrm{d}u = u(A)$; since this definitionally converges to $\widetilde{A}$, it follows that $\widetilde{A} = u(A)$. We have already noted that $\widetilde{A}$, u, and A are all Γ-invariant, and the proof that it is in Uhlenbeck gauge is clearly independent of this fact; we thus refer to [Weh04] for the remainder of the proof. □

Lemma 2.40. *The constants ε, δ and C can be chosen such that the subset $\mathscr{S}_{\varepsilon,C} \subseteq \mathscr{A}_\varepsilon$ is open. In more detail, there exists some $\varepsilon > 0$, some $\delta > 0$, and some $C > 0$ satisfying the following: if $\|g-1\|_{L^\infty_2} < \delta$ then, given any $A_0 \in \mathscr{A}_\varepsilon$ which is L^p_2-gauge equivalent to a connection in q-Uhlenbeck gauge with constant C, there is an open neighbourhood of A_0 in $\mathscr{A}_\varepsilon$ consisting of connections which permit L^p_2 gauge transformations under which they are also in Uhlenbeck gauge.*

Proof. The idea behind the proof is to use the implicit function theorem to find a neighbourhood of A_0 which can be transformed into the Coulomb gauge, and then refine this neighbourhood so that its elements also satisfies the curvature bound for the Uhlenbeck gauge. We henceforth redefine A_0 so that it is already in q-Uhlenbeck gauge with constant C.

In order to use the implicit function theorem, we will introduce three new Banach spaces. Firstly, $L^p_2(B,\mathfrak{g})_m$ is the collection of Γ-equivariant L^p_2 functions $v : B \to \mathfrak{g}$ subject to the condition that $\int_B v = 0$. Secondly, $L^p_1(B,\mathfrak{g})_\partial$ is the quotient space of $L^p_1(B,\mathfrak{g})$, the space of Γ-equivariant functions $\phi : B \to \mathfrak{g}$, by the closed subspace consisting of functions vanishing on ∂B. Thirdly, $\mathscr{Z}$ is defined as follows:

$$\mathscr{Z} = \{(f,\phi) \in L^p(B,\mathfrak{g}) \times L^p_1(B,\mathfrak{g})_\partial \,:\, \int_B f + \int_{\partial B} \phi = 0\}. \tag{2.34}$$

We now define the map on which the implicit function theorem will be used: it is the map $D : \mathscr{A}^{1,p}(B) \times L^p_2(B,\mathfrak{g})_m \to \mathscr{Z}$ defined by taking the pair (A,v) to the pair $(\mathrm{d}^*(\exp(v)\cdot A), *(\exp(v)\cdot A)|_{\partial B})$. Note that $D^{-1}(0)$ corresponds to connections and gauge transformations satisfying the Coulomb gauge, and $(A_0,0)$ is necessarily in this set.

In order to apply the implicit function theorem around this point we need to linearise D with respect to the second coordinate v, and then show that this linearisation is an isomorphism. The linearisation of the gauge group action $A \mapsto e^v \cdot A$ with respect to v is given as follows:

$$\mathscr{G}_(A,v)(\xi) = \mathrm{d}\xi + \mathrm{d}_{e^{-v}}\mathrm{Ad}(\mathrm{d}_{-v}e^{-\xi})A, \tag{2.35}$$

where $\xi : B \to \mathfrak{g}$ is an arbitrary L^p_2 map and d_p denotes the derivative at p. By the linearity of the Hodge star and the codifferential, we see that the linearisation of D is simply given by

$$\partial_2 D_{(A,v)}(\xi) = (\mathrm{d}^*\mathscr{G}_{(A,v)}(\xi), *\mathscr{G}_{(A,v)}(\xi)|_{\partial B}). \tag{2.36}$$

Specifically, at the point $(A_0,0)$, the linearisation of the gauge group action simplifies to $\mathrm{d}\xi - \mathrm{ad}_{-\xi}(A_0)$. Through algebraic manipulations, as well as the defining properties of the

Uhlenbeck gauge for A_0, we find that the linearisation of D simplifies to the following:

$$\partial_2 D_{(A_0,0)}(\xi) = \left(\Delta\xi, \frac{\partial\xi}{\partial\nu}\right) + (*[d\xi \wedge *A_0], 0), \tag{2.37}$$

where Δ is the Laplacian induced by the metric g, $\frac{\partial}{\partial\nu}$ denotes the normal derivative on ∂B, and $[A \wedge B]$ denotes the wedge product combined with the Lie bracket. We therefore set out to prove that this sum of maps is an isomorphism.

As noted in [Weh04], the first map is simply the operator of the inhomogeneous Neumann problem, and it is shown that this problem always has solutions which are uniquely determined when an additive constant is chosen (this was done in effect when we restricted the elements of $L^p_1(B, \mathfrak{g})_m$ to average to zero over B). Additionally, it was shown that the inverse of this operator is continuous as well. Though all of these results are obtained without considering Γ-equivariance, it is a simple matter to add this element to the picture: since the Neumann operator is linear and Γ acts linearly, a solution to the Neumann problem with Γ-equivariant sources must also be Γ-equivariant. (If u is a solution to the system $\Delta u = f$ where $f^\Gamma = f$, for instance, then $f = f^\Gamma = (\Delta u)^\Gamma = \Delta(u^\Gamma)$, meaning u^γ is another solution for any $\gamma \in \Gamma$; but solutions are uniquely determined up to additive constants, meaning $u^\gamma = u$.)

In order to prove that the added term does not change the fact that the map is an isomorphism, the following lemma from functional analysis is used: if $T, S \in B(X, Y)$ are bounded linear operators between Banach spaces for which T is bijective and $\|S\| < 1/\|T^{-1}\|$, then $T + S$ is also bijective. Thus, the remaining operator is bounded above by $1/2\|T^{-1}\|$ by shrinking the deviation of g from the usual metric on B. With this, we conclude that the linearisation $\partial_2 D_{(A_0,0)}$ is an isomorphism; it follows that there is some $\Delta > 0$ for which every connection $A \in \mathscr{A}_\varepsilon$ within Δ of A_0 is gauge equivalent to a connection in the Coulomb gauge, with gauge transformations given by corresponding elements of an open neighbourhood of $0 \in L^p_2(B, \mathfrak{g})_m$.

To complete the proof, it is shown via functional-analytic estimates that all of these connections immediately satisfy the curvature bound for some C when ε and Δ are chosen small enough. For details, see [Weh04]. (Note that it is this step at which C and ε must be chosen; everything preceding this point holds for any value of C and ε.) □

Proof of Theorem 2.37. As we have shown above, $\mathscr{S}_{\varepsilon,C}$ is a nonempty closed and open subset of $\mathscr{A}_\varepsilon$ for some choice of ε and C. It follows that $\mathscr{S}_{\varepsilon,C} = \mathscr{A}_\varepsilon$, so by the definition of $\mathscr{S}_{\varepsilon,C}$ and $\mathscr{A}_\varepsilon$, every connection with low enough energy is L^p_2-gauge equivalent to a connection in q-Uhlenbeck gauge with constant C. □

From the model case, it can be shown that the general case holds via a rescaling argument.

Proof of Theorem 2.36. Any given point in a Riemannian orbifold admits normal coordinates, which is to say a chart $\psi : tB \to X$ for some $t \in (0, 1]$ which satisfies $\psi^* g(0) = 1$ (the diagonal Euclidean metric) and $\nabla(\psi^* g)(0) = 0$. Since g is assumed smooth, it follows that by making t small enough, we can make $\|\psi^* g - 1\|_{L^\infty_1}$ arbitrarily small. In order to apply the theorem above to ψ, its domain would need to be all of B; modifying ψ to the map $\psi_t : B \to X$ defined so that $\psi_t(z) = \psi(tz)$ allows for this, but $\psi_t^* g$ is no longer close to 1 around 0. However, the rescaled metric $t^{-2}\psi_t^* g$ is close to 1 around 0. Additionally, its derivatives are all zero, and its second derivatives can be computed to be $t^2\nabla^2(\psi^* g)$, which can be made arbitrarily small by shrinking t and noting that g is smooth. Thus, we conclude that $\|t^{-2}\psi_t^* g - 1\|_{L_2^\infty}$ can also be made arbitrarily small.

The conclusion of the preceding paragraph is that we may apply Theorem 2.37 to B equipped with the rescaled metric $g_t := t^{-2}\psi_t^* g$ for small enough t. In more detail, this means the following: given $p, q \in (1,\infty)$ satisfying the relevant constraints, there is some $\varepsilon > 0$ and some $C > 0$ such that any L_1^p-connection A satisfying $\mathscr{E}_{g_t}(A) < \varepsilon$ admits an L_2^p-gauge transformation u taking A to a connection in q-Uhlenbeck gauge with constant C, *with respect to the rescaled metric* g_t. The remainder of the proof is in showing that we can take g_t to be simply g in this statement. This consists of two parts:

- In the general case, we start with a connection A on P for which $\mathscr{E}_g(A|_U) \le \varepsilon$; we need to show that this implies that $\mathscr{E}_{g_t}(A|_U) < \varepsilon$. We do this by rescaling a general metric h by t^2 and observing the effect on the energy:

$$\mathscr{E}_{t^2 h}(A) = \frac{1}{|\Gamma|}\int_B (t^{-2}h^{ik}t^{-2}h^{j\ell}(F_A)_{ij}(F_A)_{k\ell})^{q/2}\sqrt{t^{2n}\det(h)}\mathrm{d}^n x = t^{n-2q}\mathscr{E}_h(A). \quad (2.38)$$

  In our case, the metric h plays the role of g_t. If we assume that $\mathscr{E}_{t^2 h}(A) < \varepsilon$, therefore, this implies that $\mathscr{E}_h(A) < t^{2q-n}\varepsilon$; however, since we assumed $q \ge n/2$ and $t \le 1$, it follows that $\mathscr{E}_h(A) < \varepsilon$ as well. This completes the first part of the proof.

- Now that we have shown that A permits a gauge transformation u for which $u(A)$ is in Uhlenbeck gauge with respect to g_t, we need to show that $u(A)$ is also in Uhlenbeck gauge with respect to g. The Coulomb gauge conditions both depend only on the metric in the form of the Hodge star, which is conformally invariant, and the rescaling of g is a conformal transformation; thus, $u(A)$ also satisfies the Coulomb gauge condition with respect to g. For the extra bounding conditions, it can be shown by expanding out the relevant norms into integrals and rescaling the metric wherever it appears that the following equations hold:

$$(\|A\|_{t^2 h, L^p})^p = t^{n-p}(\|A\|_{h,L^p})^p; \quad (2.39)$$
$$(\|F_A\|_{t^2 h, L^p})^p = t^{n-2p}(\|F_A\|_{h,L^p})^p. \quad (2.40)$$

  Using the assumptions that $p \ge n/2$ and $t < 1$, it follows that the inequalities are preserved. Thus, $u(A)$ is also in Uhlenbeck gauge with respect to g, which completes the proof. □

2.4.2 Uhlenbeck compactness

We now turn attention to Uhlenbeck's weak compactness theorem, and we will use the gauge theorem to assist in the proof. We restate the theorem as it will be proved:

Theorem 2.41. *Let G be a compact Lie group, and let $P \to X$ be a principal G-bundle over an n-orbifold X. Let $(A^\nu)_{\nu\in\mathbb{N}} \in \mathscr{A}^{1,p}(P)$ be a sequence of connections for which $\|F_{A^\nu}\|_{L^p}$ is uniformly bounded, and for which $p > n/2$. Then there exists some subsequence of* *connections, also denoted by* A^ν, *and a sequence of* L^p_2 *gauge transformations* $(u^\nu)_{\nu\in\mathbb{N}} \in \mathscr{G}^{2,p}$ *for which* $u^\nu(A^\nu)$ *converge weakly*[9] *in* $A^{1,p}$.

In proving this theorem, the extent to which the Uhlenbeck gauge theorem is used is rather limited; we only need the part of the theorem pertaining to the curvature bound. We will also need a technical lemma which allows us to patch together gauge transformations on open covers of the orbifold in a way that retains the curvature bound. Before this lemma is proved, we need a result from Lie theory which is also used in [Weh04], and which is proved in [GHL90].

Lemma 2.42. *Let* G *be a compact Lie group with Lie algebra* $\mathfrak{g}$, *and equip* G *with a left-invariant metric* d. *Then there is some convex geodesic ball around the identity in* G. *That is, there is some real number* $\Delta_{\exp} > 0$ *satisfying the following two properties:*

- *The map* $\exp : \mathfrak{g} \to G$ *induces a bijection between the balls of radius* $\Delta_{\exp}$ *around the origin in* $\mathfrak{g}$ *and the identity in* G.
- *The ball of radius* $\Delta_{\exp}$ *around the identity in* G *contains unique minimal geodesics, meaning any two points in this ball can be joined by a unique geodesic whose image lies entirely within the ball and whose length is minimal among paths joining the two points.*

Lemma 2.43. *Let* X *be a compact* n*-orbifold and let* $p > n/2$. *Let* G *be a compact Lie group with Lie algebra* $\mathfrak{g}$, *fix a left-invariant metric* d *on* G, *and let* $\Delta_{\exp}$ *be the radius of a convex geodesic ball in* G. *Fix a finite atlas* $\{(\widetilde{U}_\alpha, \Gamma_\alpha, \varphi_\alpha)\}_{1\le\alpha\le N}$ *on* X *for which each* U_α *has compact closure in* X. *For each* α *define a homomorphism* $\rho_\alpha : \Gamma_\alpha \to G$, *and extend the action of* Γ_α *from* $\widetilde{U}_\alpha$ *to* $\widetilde{U}_\alpha \times G$ *by taking* $\gamma\cdot(x,g) = (\gamma\cdot x, \rho_\alpha(\gamma)g)$.

Then there exists another atlas $\{(\widetilde{V}_\alpha, \Gamma_\alpha, \varphi_\alpha)\}_{1\le\alpha\le N}$ *for which* $V_\alpha \subseteq U_\alpha$, *such that the following holds:*

(i) Let $k \in \mathbb{N}$, *and for every inclusion* $U_\alpha \hookrightarrow U_\beta$, *let* $g_{\alpha\beta}, h_{\alpha\beta} : \widetilde{U}_\alpha \to G$ *be systems of* L^p_{k+1} *transition functions for a principal* G*-bundle on* X. *If*

$$\sup_{x\in U_\alpha} d(g_{\alpha\beta}(x), h_{\alpha\beta}(x)) \le \Delta_{\exp} \tag{2.41}$$

for every chart inclusion $U_\alpha \hookrightarrow U_\beta$, *then there exist* L^p_{k+1} *local gauge transformations* $h_\alpha : \widetilde{V}_\alpha \to G$ *satisfying the following condition on* V_α*:*

$$h_\alpha^{-1} h_{\alpha\beta} h_\beta = g_{\alpha\beta}. \tag{2.42}$$

(ii) Let the $h_{\alpha\beta}$ *in (i) run through a sequence* $h^\nu_{\alpha\beta}$ *of sets of transition functions such that* $g_{\alpha\beta}, h^\nu_{\alpha\beta}$ *are of class* L^p_{k+1} *for all* $k < K$, *where* $K \ge 2$ *is an integer or* ∞. *Assume that for every* α, β *and every* $k < K$ *the sequence of 1-forms* $(h^\nu_{\alpha\beta})^{-1}\mathrm{d}h^\nu_{\alpha\beta}$ *is uniformly bounded in* L^p_k. *Then the gauge transformations* h^ν_α *in (i) can be chosen so that, for every* α *and* $k < K$, *the sequence of 1-forms* $(h^\nu_\alpha)^{-1}\mathrm{d}h^\nu_\alpha$ *is uniformly bounded in* L^p_k.

[9] By weak convergence we mean here that if a reference connection A^0 is chosen, the sequence of $\mathfrak{g}$-valued 1-forms $u^\nu(A^\nu) - A^0 \in \Omega^1(X,\mathfrak{g})_{L^p_1}$ converges weakly; this in turn means that there is some $\mathfrak{g}$-valued 1-form α for which $\varphi(u^\nu(A^\nu) - A^0)$ converges to $\varphi(\alpha)$ for every continuous linear functional φ on $\Omega^1(X,\mathfrak{g})_{L^p_1}$.

Remark. In light of the fibre bundle construction lemma, this lemma can be interpreted as follows: if a sequence of systems of gauge transformations of a principal G-bundle are sufficiently close together then they are all equivalent, and if they are uniformly bounded then the equivalences are also uniformly bounded. We will use this lemma to glue local gauge transformations together uniformly: over a bundle with transition functions $\phi_{\alpha\beta}$, a gluing of the local gauge transformations g_α is the same thing as an equivalence $\phi_{\alpha\beta} \sim g_\alpha^{-1}\phi_{\alpha\beta}g_\beta$.

Proof. The idea of the proof is to construct the h_α inductively as follows. If the chart U_α does not overlap with any previous charts U_β on which h_β has already been defined, then h_α is taken to be the identity gauge transformation. However, if any U_β does overlap, h_α is defined on $U_\alpha \cap U_\beta$ to be whatever makes the equation $h_\beta^{-1}h_{\beta\alpha}h_\alpha = g_{\beta\alpha}$ true (namely the gauge transformation $h_{\beta\alpha}^{-1}h_\beta g_{\beta\alpha}$). On the rest of U_α, it is smoothly transformed into the identity using a cutoff function; in order for this to be possible, the h_α need to be sufficiently close to the identity that they can be written in terms of the exponential map.

We begin the proof by reusing the notation and some of the restrictions for the bundle atlas that were used in Lemmas 2.27 and 2.28. In particular, given some $\delta > 0$, we define the set $U_\alpha^\delta = \{p \in U_\alpha : d_X(p, \partial U_\alpha) > \delta\}$, where d_X is a metric corresponding to a Riemannian metric on X. We then reorder the atlas so that the following holds: if $U_{j+1}^{(j+1)\delta}$ and $\cup_{\alpha\leq j}U_\alpha^{(j+1)\delta}$ overlap, and if U_β is another orbifold chart which embeds into U_{j+1}, then $\beta \leq j$. Since transition functions on orbifolds only make sense when one chart is embedded into the other, this effectively means $\alpha \leq \beta$ must hold in order for $g_{\alpha\beta}$ to be defined. We will eventually define the V_α in the statement of the theorem to be $U_\alpha^{n\delta}$, where δ is sufficiently small that this collection of open sets still covers X.

We now begin the inductive construction of the h_α. At the jth stage, we will obtain a collection of local gauge transformations $h_\alpha \in \mathscr{G}^{k+1,p}(U_\alpha^{j\delta})$ which give an equivalence between the transition functions $g_{\alpha\beta}$ and $h_{\alpha\beta}$, and which additionally satisfy the following constraint for every $\alpha \leq j$ and every $\beta \geq j+1$:

$$d(h_{\alpha\beta}^{-1}h_\alpha g_{\alpha\beta}, e_G) \leq \Delta_{\exp}. \tag{2.43}$$

- The local gauge transformation $h_1 : \widetilde{U}_1^\delta \to G$ is defined simply to be the identity gauge transformation (which is clearly Γ_1-equivariant). Since $g_{11} = h_{11} = e_G$, any choice for h_1 would have constituted an equivalence between the two. Additionally $d(h_{1\beta}^{-1}g_{1\beta}, e_G) \leq \Delta_{\exp}$ for all β, by the G-invariance of the metric and by the assumptions on $g_{\alpha\beta}$ and $h_{\alpha\beta}$.

- Suppose there exist L_{k+1}^p gauge transformations $h_\alpha : \widetilde{U}_\alpha^{(j-1)\delta} \to G$ for every $\alpha \leq j-1$ such that $h_\alpha^{-1}h_{\alpha\beta}h_\beta = g_{\alpha\beta}$ whenever $U_\alpha \hookrightarrow U_\beta$, and additionally Equation 2.43 holds. We wish to construct an L_{k+1}^p gauge transformation $h_j : \widetilde{U}_j^{j\delta} \to G$ which maintains these properties. To do this, first note that we have ordered the U_α so that U_j cannot embed into any of the U_α for $\alpha \leq j-1$, so it suffices to ensure that $h_\alpha^{-1}h_{\alpha j}h_j = g_{\alpha j}$ on $U_\alpha^{j\delta}$ for the case that $\alpha \leq j-1$ (we do not need to consider $\alpha = j$), or equivalently that $h_j = h_{\alpha j}^{-1}h_\alpha g_{\alpha j}$.

  First, we must check that this condition is consistent for any embedding $U_\alpha \hookrightarrow U_j$. If $U_\beta \hookrightarrow U_j$ is another embedding which overlaps U_α, then we may assume without loss of generality that $U_\beta \hookrightarrow U_\alpha$; it follows from the cocycle condition and the inductive hypothesis that

$$h_{\beta j}^{-1} h_\beta g_{\beta j} = (h_{\beta\alpha}^{-1} h_{\beta j})^{-1} h_{\beta\alpha}^{-1} h_\alpha g_{\beta\alpha}(g_{\beta\alpha}^{-1} g_{\beta j}) = h_{\alpha j}^{-1} h_\beta g_{\alpha j}, \tag{2.44}$$

so the two definitions are consistent. Now, in the induction, the requirement that $h_{\alpha\beta}^{-1} h_\alpha g_{\alpha\beta}$ is sufficiently close to the identity of G for $\alpha \leq j < \beta$ implies that we can write it in terms of the exponential map: there is some Γ_α-equivariant map $(\xi_j)_\alpha : \widetilde{U}_\alpha^{j\delta} \to \mathfrak{g}$ for which $h_{\alpha j}^{-1} h_\alpha g_{\alpha\beta} = \exp(\xi_j)$. Now that we have this local representation of h_j on $U_\alpha \cap U_j$, we smoothly connect it to the identity on the rest of U_j as follows. There is a smooth Γ_j-invariant map $\psi_j : \widetilde{U}_j^{j\delta} \to [0,1]$ which is 1 on $\widetilde{U}_j^{j\delta} \cap (\cup_{\alpha \leq j-1} \widetilde{U}_\alpha^{j\delta})$ and 0 on $\widetilde{U}_j^{j\delta} \backslash (\cup_{\alpha \leq j-1} \widetilde{U}_\alpha^{(j-1)\delta})$. We define $h_j : \widetilde{U}_j^{j\delta} \to G$ to be the following:

$$h_j = \begin{cases} \exp(\psi_j(\xi_j)_\alpha) & \text{on } \widetilde{U}_j^{j\delta} \cap \widetilde{U}_\alpha^{j\delta} \\ e_G & \text{otherwise.} \end{cases} \tag{2.45}$$

Since the $(\xi_j)_\alpha$ match each other on overlaps and ψ_j tends to zero away from overlaps, h_j is well-defined and continuous. In fact, we use the same definition for part (ii) of the theorem: given transition functions $h_{\alpha\beta}^\nu$ we define h_α^ν to be $\exp(\psi_j(\xi_j^\nu)_\alpha)$, where $(\xi_j^\nu)_\alpha$ is defined to be $\exp^{-1}((h_{\alpha j}^\nu)^{-1} h_\alpha^\nu g_{\alpha\beta})$.

To complete the induction, we need to verify that the functions $h_{j\beta}^{-1} h_j g_{j\beta}$ are close to the identity for any $U_j \hookrightarrow U_\beta$. Around any point not in $U_\alpha^{(j-1)\delta}$ for $\alpha \leq j-1$, the function h_j is the identity; thus, $d(h_{j\beta}^{-1} h_j g_{j\beta}, e_G) = d(h_{j\beta}, g_{j\beta}) \leq \Delta_{\text{exp}}$. Otherwise, there is a chart $U_\alpha \hookrightarrow U_j$ containing any given point, so we may use the cocycle condition and the inductive hypothesis:

$$\begin{aligned} d(h_{j\beta}^{-1} h_j g_{j\beta}, e_G) &= d((h_{\alpha j} h_{j\beta})^{-1} h_{\alpha j} h_j g_{\alpha j}^{-1} g_{\alpha j} g_{j\beta}, e_G) \\ &= d(h_{\alpha\beta}^{-1} h_\alpha g_{\alpha\beta}, e_G) \leq \Delta_{\text{exp}}. \end{aligned}$$

This process is carried out for every $h_{\alpha\beta}^\nu$, meaning we complete the induction with a series of equivalences h_α^ν.

Now that the equivalences h_α^ν have been defined, it remains to verify the regularity and uniform boundedness conditions for the relevant maps. For these purposes, we will be using several technical lemmas in the theory of Sobolev connections and gauge transformations; refer to the Appendix for discussion of these lemmas.

We begin by proving that each h_α is L_{k+1}^p by induction on α. The base case is obvious as h_1 is constant, so assume that h_α is L_{k+1}^p for $\alpha < j$. Then, since gauge group multiplication is well-defined for L_{k+1}^p, we know that each $h_{\alpha\beta}^{-1} h_\alpha g_{\alpha\beta}$ is also L_{k+1}^p; the exponential map gives a local bijective correspondence between L_{k+1}^p functions, meaning $(\xi_j)_\alpha$ is also L_{k+1}^p. By the smoothness of ψ_j the map $\psi_j(\xi_j)_\alpha$ is L_{k+1}^p, and therefore $h_j := \exp(\psi_j(\xi_j)_\alpha)$ is L_{k+1}^p. This proves the inductive hypothesis, so h_α is L_{k+1}^p.

We now find uniform bounds for $\|(h_\alpha^\nu)^{-1} \mathrm{d} h_\alpha^\nu\|_{L_k^p}$ over ν by induction on α. Once again, the base case is obvious, so assume that $\|(h_\alpha^\nu)^{-1} \mathrm{d} h_\alpha^\nu\|_{L_k^p}$ is uniformly bounded over ν for all $\alpha < j$. We will find the required uniform bounds by first finding bounds on $(\xi_\alpha)_j^\nu$, and in turn on $\exp(-(\xi_\alpha)_j^\nu) \mathrm{d} \exp((\xi_\alpha)_j^\nu)$.

To find bounds on the logarithmic derivative of $\exp((\xi_\alpha)_j^\nu)$, observe that it is equal to $(h_{\alpha j}^\nu)^{-1} h_\alpha^\nu g_{\alpha j}$; by Lemma A.14, it suffices to find uniform bounds on $\|(h_{\alpha j}^\nu)^{-1} \mathrm{d} h_{\alpha j}^\nu\|_{L_k^p}$, $\|(h_\alpha^\nu)^{-1} \mathrm{d} h_\alpha^\nu\|_{L_k^p}$, and $\|(g_{\alpha j}^\nu)^{-1} \mathrm{d} g_{\alpha j}\|_{L_k^p}$. But the first of these has uniform bounds by assumption, the second has uniform bounds by the inductive hypothesis, and the third does not depend on ν. It follows that $\exp(-(\xi_\alpha)_j^\nu) \mathrm{d} \exp((\xi_\alpha)_j^\nu)$ is uniformly bounded in L_k^p, and by Lemma A.4, this is equivalent to a uniform bound on $(\xi_\alpha)_j^\nu$ in L_{k+1}^p. Multiplying by the smooth function ψ_j and using Lemma A.4, we get a uniform bound on $(h_j^\nu)^{-1} \mathrm{d} h_j^\nu$ in L_k^p, proving the inductive hypothesis. □

Proof of Theorem 2.41. Choose $q \in (1, p)$ such that $q \geq n/2$ and $q \geq pn/(p+n)$, so that Theorem 2.36 holds with the L^q-energy $\mathscr{E}$. Let $C > 0$ and $\varepsilon > 0$ be the constants from the theorem (C being the bounding constant for the L_1^p-norm of the connection and ε being the maximum energy of a suitable connection). Over a small trivialising chart U for the principal bundle $P \to X$, the q-energy of the connection A^ν restricted to U is given as follows:

$$\mathscr{E}(A^\nu|_U) = \int_U |F_{A^\nu}|^q \mathrm{vol}_U \leq (\mathrm{Vol}(U))^{1-q/p} \|F_{A^\nu}\|_{L^p}^q, \tag{2.46}$$

where we have applied Hölder's inequality in the last step. Since the L^p-norm of the curvature is uniformly bounded, the q-energy of the connections over U can be made smaller than ε as long as U is made small enough.

By Uhlenbeck's gauge theorem for orbifolds and the compactness of X, there is a finite open covering of X over which P is trivial and each connection can be put into Uhlenbeck gauge; we may assume without loss of generality that this covering can be refined to form an atlas $\{(\widetilde{U}_\alpha, \Gamma_\alpha, \varphi_\alpha)\}_{1 \leq \alpha \leq N}$ over X. Note that, over a trivialising open set U_α, the connection A^ν can be represented by an Γ_α-equivariant $\mathfrak{g}$-valued 1-form A_α^ν of class L_1^p on $\widetilde{U}_\alpha \subseteq \mathbb{R}^n$. We denote by $u_\alpha^\nu : \widetilde{U}_\alpha \to G$ the local gauge transformations given by the Uhlenbeck gauge theorem, so that

$$\|u_\alpha^\nu(A_\alpha^\nu)\|_{L_1^p} \leq C \|F_{A_\alpha^\nu}\|_{L_2^p}; \tag{2.47}$$

note that the u_α^ν are L_2^p and are Γ_α-equivariant. It follows that $\|u_\alpha^\nu(A_\alpha^\nu)\|_{L_1^p}$ is uniformly bounded over ν for every α. If the gauge transformations u_α^ν matched on overlaps then we could use the Banach-Alaoglu theorem to construct a convergent subsequence; thus, the remainder of the proof is devoted to patching together the gauge transformations while maintaining the uniform L_1^p bound.

For each embedding $U_\alpha \hookrightarrow U_\beta$, let $\phi_{\alpha\beta} : \widetilde{U}_\alpha \times G \to \widetilde{U}_\beta \times G$ be the Γ_α-equivariant transition function for P, and denote by $u_{\alpha\beta}^\nu : \widetilde{U}_\alpha \to G$ the following map:

$$u_{\alpha\beta}^\nu = (u_\alpha^\nu)^{-1} \phi_{\alpha\beta} u_\beta^\nu. \tag{2.48}$$

Note that this sequence of systems of functions satisfy the cocycle condition, meaning they constitute transition functions, and additionally note that an equivalence $u_{\alpha\beta}^\nu \sim \phi_{\alpha\beta}$ gives a patching of the u_α^ν together.

We will use the Lemma above to construct a series of equivalences, but in order to do so, we need to ensure that the $u_{\alpha\beta}^\nu$ are sufficiently close together in G. Observe that $u_{\alpha\beta}^\nu(u_\alpha^\nu(A_\alpha^\nu)) = u_\beta^\nu(A_\beta^\nu)$, and both $u_\alpha^\nu(A_\alpha^\nu)$ and $u_\beta^\nu(A_\beta^\nu)$ are uniformly bounded in L_1^p by definition. By Lemma A.13, the sequence $(u_{\alpha\beta}^\nu)^{-1} \mathrm{d} u_{\alpha\beta}^\nu$ is uniformly bounded in L_1^p and it has a C^0-convergent subsequence. In particular, it follows that every gauge transformation in this subsequence is eventually within Δ_{exp} of every other gauge transformation. We therefore relabel the sequence in the following way: we define $g_\alpha = u^1_\alpha$ and $g_{\alpha\beta} = u^1_{\alpha\beta}$, and we ensure that $\mathrm{d}(u^\nu_{\alpha\beta}, g_{\alpha\beta}) \leq \Delta_{\text{exp}}$ for every α, β and every $\nu \in \mathbb{N}$.

As a result, by taking $u^\nu_{\alpha\beta}$ to be the $h^\nu_{\alpha\beta}$ in the above theorem (with $K = 2$), we get a series of equivalences h^ν_α between $u^\nu_{\alpha\beta}$ and $g_{\alpha\beta}$ whose logarithmic derivatives are uniformly bounded in L^p_1. By multiplying each h^ν_α by g^{-1}_α, we get the following equivalence:

$$(h^\nu_\alpha g^{-1}_\alpha)^{-1} u^\nu_{\alpha\beta} h^\nu_\alpha g^{-1}_\alpha = \phi_{\alpha\beta}. \tag{2.49}$$

It follows that the local gauge transformations defined by $\widetilde{u}^\nu_\alpha := u^\nu_\alpha h^\nu_\alpha g^{-1}_\alpha$ glue together to give a global gauge transformation.

We now prove that $\widetilde{u}^\nu_\alpha(A^\nu_\alpha)$ is bounded in L^p_1 over every U_α. The key observation is that Lemma A.11 gives uniform L^p_1 bounds on $u(A)$ so long as $u^{-1}\mathrm{d}u$ and A are both uniformly bounded in L^p_1, meaning we need only find uniform L^p_1 bounds on the following quantities:

- $u^\nu_\alpha(A^\nu_\alpha)$ is uniformly bounded by assumption, as u^ν_α is the Uhlenbeck gauge.
- $(h^\nu_\alpha)^{-1}\mathrm{d}h^\nu_\alpha$ is uniformly bounded by the Lemma above.
- $g^{-1}_\alpha \mathrm{d}g_\alpha$ is uniformly bounded as it does not depend on ν.

We conclude that $\widetilde{u}^\nu_\alpha(A^\nu_\alpha)$ is bounded in L^p_1. By the Banach-Alaoglu theorem, it has an L^p_1 weakly convergent subsequence over every α; by induction, this can be made into a weakly convergent subsequence for all α. Once these are patched together, we obtain a single L^p_1 convergent subsequence $\widetilde{u}^\nu(A^\nu)$. □

Chapter 3

Seiberg-Witten Theory

In this chapter, we at last turn attention to Seiberg-Witten theory. The Seiberg-Witten equations of interest in this thesis are the 3-manifold versions; however, their original form was defined on 4-manifolds, and the solutions to the 4-manifold equations on cylinders are of great importance in assigning invariants to 3-manifolds. As such, we devote time to both dimensions 3 and 4 throughout the chapter.

In contrast to the vortex equations, whose parameters consisted simply of a section and unitary connection on a complex line bundle, the definition of the configuration space of the Seiberg-Witten equations requires considerable algebraic machinery. Thus, the first two sections provide an exposition of the requisite algebra and geometry. In particular, the first section covers the complexified spin group and Clifford algebras, and the next section explores their parametrised versions as Spin^c-structures and spinor bundles. After these sections, the Seiberg-Witten equations are defined on 3-manifolds and 4-manifolds, and the interpretation of the 3-manifold equations in terms of the Chern-Simons-Dirac functional is discussed. The 4-manifold equations are related to the Chern-Simons-Dirac functional in the final section of the chapter, in the special case that the 4-manifold is a cylinder $Y \times \mathbb{R}$ over a 3-manifold.

We do not discuss equivariant generalisations in this chapter. However, from Section 3.2 onwards, it is possible to define a Λ-equivariant analogue for every geometric object in the usual way. An equivariant Spin^c-structure over a Λ-equivariant 3-manifold, for instance, is a Λ-equivariant principal $\mathrm{Spin}^c(3)$-bundle together with a Λ-equivariant projection onto the frame bundle. These equivariant generalisations will be used in Chapter 4, but we will generally assume that the equivariant version of a geometric object can be intuited based on the approach in Section 2.1.

Due to the wealth of sign and scaling choices involved in the development of Spin^c-geometry and the Seiberg-Witten equations, Appendix B compiles convention differences across the literature and the choices we have made in this exposition.

3.1 Algebraic Preliminaries

Contrary to the vortex equations with structure group $\mathrm{U}(1)$, the structure group of Seiberg-Witten theory is the complexified spin group. We therefore devote a section to the fundamental properties of this group. For a more detailed study of Clifford algebras and spin groups, refer to [Gar11], [Wer19], and [LM89].

3.1.1 Spin groups

The spin groups are usually defined via Clifford algebras:

Definition 3.1. Let (V, g) be a quadratic space over a field $\mathbb{K}$. The *Clifford algebra of* (V, g) is the associative unital algebra generated by V subject to the relations $v^2 = -g(v, v)$ for all $v \in V$. We denote this algebra by $\mathrm{Cl}(V, g)$. If V is an $\mathbb{R}$-vector space, the Clifford algebra over the complexification of V is denoted by $\mathrm{Cl}^c(V, g)$, and is called the *complexified Clifford algebra;* note that $\mathrm{Cl}^c(V, g) \cong \mathrm{Cl}(V, g) \otimes \mathbb{C}$.

Remark. Another common convention is to enforce the relation $v^2 = +g(v, v)$ instead, which results in sign differences between sources. However, the literature on Seiberg-Witten theory generally uses the negative sign convention.

In our setting, g will always be an inner product, and where there is no ambiguity we will simply denote the Clifford algebra by $\mathrm{Cl}(V)$. A g-orthonormal basis for V will generically be denoted by $\{e_i\}$; note that $e_ie_i = -1$ for all i and $e_ie_j = -e_je_i$ for $i \neq j$.

Definition 3.2. Let V be a real n-dimensional vector space with an inner product g. The *spin group* $\mathrm{Spin}(n)$ is defined to be the subset of $\mathrm{Cl}(V, g)$ which can be written in the form $\pm v_1 \cdots v_{2k}$, where each v_i has length 1 (and k is allowed to be zero). This forms a group under the Clifford product, with inverses given by taking $v_1 \cdots v_{2k} \mapsto v_{2k} \cdots v_1$.

There is a natural representation of the spin group on V given by conjugation: an element $\gamma \in \mathrm{Spin}(n)$ acts on a vector $x \in V$ by taking $x \mapsto \gamma x \gamma^{-1}$ (it is a straightforward exercise to show that $\gamma x \gamma^{-1} \in V$). Moreover, for any unit vector $v \in V$, the map $x \mapsto -vxv$ is a reflection across the ray $\mathbb{R}v$. Since the action of any $\gamma \in \mathrm{Spin}(n)$ is a composition of these transformations, and since these reflections generate $\mathrm{O}(n)$ under composition (see [Gar11]), it follows that the representation of $\mathrm{Spin}(n)$ in $\mathrm{GL}(n)$ covers $\mathrm{SO}(n)$. It happens that this is always a double covering:

Proposition 3.3. *The map* $\mathrm{Spin}(n) \to \mathrm{SO}(n)$ *defined by the conjugation action of* $\mathrm{Spin}(n)$ *on* $\mathbb{R}^n$ *is a smooth double covering.*

Proof. Refer to [Wer19]. □

A standard theorem of algebraic topology is that $\pi_1(\mathrm{SO}(n)) = \mathbb{Z}_2$ for $n \geq 3$ (see [Ark11]), meaning $\mathrm{Spin}(n)$ is even the universal cover of $\mathrm{SO}(n)$ for $n \geq 3$ (note that $\mathrm{Spin}(n)$ is connected for $n \geq 2$).

It also follows from this proposition that $\mathrm{Spin}(n)$ is a Lie group, whose Lie algebra is isomorphic to that of $\mathrm{SO}(n)$. We will occasionally need to make use of the isomorphism between these two Lie algebras, so we state the explicit correspondence here. We identify $\mathfrak{so}_n$ with the Lie algebra of skew-symmetric matrices (with the commutator bracket), and in turn we identify this with $\Lambda^2\mathbb{R}^n$ by the linear isomorphism $e_i \wedge e_j \mapsto E_{ji} - E_{ij}$; here E_{ij} is the elementary matrix with nonzero entry at position (i, j).

Proposition 3.4. *The Lie algebra* $\mathfrak{spin}_n$*, interpreted as the tangent space at* $1 \in \mathrm{Spin}(n)$*, can be identified with the subspace of* $\mathrm{Cl}(\mathbb{R}^n)$ *spanned by* $\{e_i e_j\}_{1 \le i < j \le n}$ *with the commutator bracket. This space, in turn, can be identified with* $\Lambda^2\mathbb{R}^n$ *via the Lie algebra isomorphism* $e_i e_j \mapsto 2(e_i \wedge e_j)$*.*

Proof. The following proof is adapted from [Wer19]. Observe that the map $\gamma : \mathbb{R} \to \mathrm{Cl}(\mathbb{R}^n)$ defined by taking $\gamma(t) = \cos(t) + \sin(t)e_i e_j$ is valued in $\mathrm{Spin}(n)$, since it can be written as $-(\cos(t/2)e_i + \sin(t/2)e_j)(\cos(t/2)e_i - \sin(t/2)e_j)$. Moreover, we have that $\gamma(0) = 1$ and $\gamma'(0) = e_i e_j$, meaning $e_i e_j \in \mathfrak{spin}_n$ for each i, j. Moreover, since there are $n(n-1)/2$ choices for $1 \le i < j \le n$, and since $\dim(\mathrm{Spin}(n)) = \dim(\mathrm{SO}(n)) = n(n-1)/2$, these elements span the Lie algebra. Also, since $\mathrm{Spin}(n)$ is a Lie subgroup of the group of units of $\mathrm{Cl}(\mathbb{R}^n)$, their Lie brackets coincide. This proves the first correspondence.

Let $\rho : \mathrm{Spin}(n) \to \mathrm{SO}(n)$ denote the conjugation action; given some $e_i e_j \in \mathfrak{spin}_n$, we compute the corresponding element of $\mathfrak{so}_n$ by computing ρ_{*0} applied to $\gamma'(0)$. We do this using the map $\rho \circ \gamma : \mathbb{R} \to \mathrm{SO}(n)$; computing this explicitly on some $x \in \mathbb{R}^n$, we find that

$$\begin{aligned}(\rho \circ \gamma)(t) \cdot x &= \gamma(t) x \gamma(t)^{-1} \\ &= (\cos(t) + \sin(t)e_i e_j)x(\cos(t) - \sin(t)e_i e_j) \\ &= \cos^2(t)x + \cos(t)\sin(t)(e_i e_j x - x e_i e_j) - \sin^2(t) e_i e_j x e_i e_j.\end{aligned}$$

Differentiating this at $t = 0$, we obtain a tangent vector in $\mathfrak{so}_n$ (applied to x):

$$(\rho \circ \gamma)'(0) \cdot x = e_i e_j x - x e_i e_j. \tag{3.1}$$

The matrix representation of this endomorphism in the $\{e_i\}$ basis is given by $2(E_{ji} - E_{ij})$, which corresponds to the bivector $2(e_i \wedge e_j)$ under the standard identification with $\Lambda^2\mathbb{R}^n$. □

Remark. Given two vectors $v, w \in \mathbb{R}^n$ and a metric on $\mathbb{R}^n$ inducing an isomorphism with $(\mathbb{R}^n)^*$, we can define an endomorphism $v \wedge w \in \mathrm{End}(\mathbb{R}^n)$ given by $v \wedge w = vw^* - wv^*$. (Note that the matrix representation of $e_i e^j - e_j e^i$ is $E_{ji} - E_{ij}$, so the two interpretations of $e_i \wedge e_j$ are consistent.) This is another way to describe the correspondence between $\mathfrak{so}_n$ and $\Lambda^2\mathbb{R}^n$. We also find that $[v, w] \in \mathfrak{spin}_n$, since all grade-zero components of vw are cancelled by those of wv. Under ρ_{*0}, we find that $[v, w]$ is mapped to the following:

$$\begin{aligned}\rho_{*0}([v, w]) &= \rho_{*0}\left(\sum_{i,j} v^i w^j [e_i, e_j]\right) \\ &= \sum_{i,j} v^i w^j \rho_{*0}(2e_i e_j) \\ &= \sum_{i,j} 4v^i w^j (e_i \wedge e_j) = 4(v \wedge w).\end{aligned}$$

In other words, we find that the Lie algebra element $v \wedge w \in \Lambda^2\mathbb{R}^n \cong \mathfrak{so}_n$ corresponds to $\frac{1}{4}[v, w] \in \mathfrak{spin}_n$.

3.1.2 Spin^c groups

The Seiberg-Witten equations make use of the complex analogue of the spin groups.

Definition 3.5. Let V be a real n-dimensional vector space with an inner product g. The complexified Spin group $\mathrm{Spin}^c(n)$ is defined to be the subset of $\mathrm{Cl}^c(V,g)$ generated by $\mathrm{Spin}(n) \otimes 1$ and $\mathrm{U}(1)$; that is, the group consists of all elements of the form $e^{i\theta} v_1 \cdots v_{2k}$ where each $v_i \in V$ has unit length.

Though every element of $\mathrm{Spin}^c(n)$ is a product of elements of $\mathrm{Spin}(n)$ and $\mathrm{U}(1)$, the group $\mathrm{Spin}^c(n)$ itself is not the direct product $\mathrm{Spin}(n) \times \mathrm{U}(1)$; observe that $e^{i\theta}\gamma = (-e^{i\theta})(-\gamma) \in \mathrm{Spin}^c(n)$. Instead, the complexified group has a *fibred product structure:* it is given by

$$\mathrm{Spin}^c(n) \cong \mathrm{Spin}(n) \times_{\mathbb{Z}_2} \mathrm{U}(1) := (\mathrm{Spin}(n) \times \mathrm{U}(1))/\mathbb{Z}_2, \tag{3.2}$$

where $\mathbb{Z}_2$ acts on each factor by negation. We will often denote elements of $\mathrm{Spin}^c(n)$ by pairs $(\gamma, e^{i\theta})$, with the understanding that they are equivalent to $(-\gamma, -e^{i\theta})$.

Nevertheless, we do get natural inclusions $\mathrm{Spin}(n) \hookrightarrow \mathrm{Spin}^c(n)$ and $\mathrm{U}(1) \hookrightarrow \mathrm{Spin}^c(n)$ by taking $\gamma \mapsto (\gamma, 1)$ and $e^{i\theta} \mapsto (1, e^{i\theta})$ respectively. The image of each inclusion is normal in $\mathrm{Spin}^c(n)$, and the quotient groups are naturally modulated by the $\mathbb{Z}_2$ action:

$$\mathrm{Spin}^c(n)/\mathrm{U}(1) \cong \mathrm{Spin}(n)/\mathbb{Z}_2 \cong \mathrm{SO}(n), \tag{3.3}$$

$$\mathrm{Spin}^c(n)/\mathrm{Spin}(n) \cong \mathrm{U}(1)/\mathbb{Z}_2 \cong \mathrm{U}(1). \tag{3.4}$$

We therefore have the following exact sequences:

$$1 \longrightarrow \mathrm{Spin}(n) \longrightarrow \mathrm{Spin}^c(n) \xrightarrow{q_{\mathrm{U}}} \mathrm{U}(1) \longrightarrow 1$$

$$1 \longrightarrow \mathrm{U}(1) \longrightarrow \mathrm{Spin}^c(n) \xrightarrow{q_{\mathrm{SO}}} \mathrm{SO}(n) \longrightarrow 1$$

Remark. The quotient map $q_{\mathrm{U}} : \mathrm{Spin}^c(n) \to \mathrm{U}(1)$ can be described explicitly as follows:

$$q_{\mathrm{U}}(\gamma, e^{i\theta}) = e^{i(2\theta)}. \tag{3.5}$$

This factor of 2 has subtle implications for Seiberg-Witten theory. In particular, when we "break up" $\mathrm{Spin}^c(n)$ into a spin piece and a $\mathrm{U}(1)$ piece, we will often find that the gauge theory of $\mathrm{Spin}^c(n)$ differs by a factor of 2 from $\mathrm{U}(1)$ gauge theory.

Combining the quotient maps, we obtain the following exact sequence:

$$0 \longrightarrow \mathbb{Z}_2 \longrightarrow \mathrm{Spin}^c(n) \xrightarrow{(q_{\mathrm{SO}}, q_{\mathrm{U}})} \mathrm{SO}(n) \times \mathrm{U}(1) \longrightarrow 1$$

It follows that $\mathrm{Spin}^c(n)$ is a double covering of $\mathrm{SO}(n) \times \mathrm{U}(1)$. We once again obtain an isomorphism of Lie algebras, but in this instance it is between $\mathfrak{spin}^c(n)$ and $\mathfrak{so}_n \oplus \mathfrak{u}_1$; the exact form of this isomorphism is the content of the following proposition:

Proposition 3.6. *The Lie algebra $\mathfrak{spin}^c(n)$, interpreted as the tangent space at the identity in* $\mathrm{Spin}^c(n)$, *can be identified with the (real) subspace of* $\mathrm{Cl}^c(\mathbb{R}^n)$ *spanned by* $\{e_i e_j\}_{1 \le i < j \le n}$ *and the grade-zero element i. This space, in turn, can be identified with $\mathfrak{so}_n \oplus i\mathbb{R}$ via the Lie algebra isomorphism $\mathfrak{a} : \mathfrak{spin}^c(n) \to \mathfrak{so}_n \oplus i\mathbb{R}$ defined on generators as follows:*

$$\begin{aligned} \mathfrak{a}(e_i e_j) &= 2(e_i \wedge e_j), \\ \mathfrak{a}(i) &= 2i. \end{aligned} \tag{3.6}$$

Proof. Since $\mathrm{Spin}(n)$ is a subgroup of $\mathrm{Spin}^c(n)$, the non-complex generators can be found via the exact same argument as in Proposition 3.4. To find the generator of the $\mathrm{U}(1)$ action, we consider the path $\gamma : \mathbb{R} \to \mathrm{Spin}^c(n)$ given by $\gamma(t) = e^{it}$; this clearly has derivative i. Since these generators are all linearly independent and exhaust the dimension of $\mathrm{SO}(n) \times \mathrm{U}(1)$, we have the first correspondence.

For the second, we clearly still have that the $\mathfrak{spin}_n$ Lie subalgebra of $\mathfrak{spin}^c(n)$ corresponds to $\Lambda^2\mathbb{R}^n$. Thus, we need only consider the induced Lie algebra isomorphism from its complement to $\mathfrak{u}_1$ by the map q_{U}. On the other hand, differentiation of q_{U} shows that this map is simply multiplication by 2. □

3.1.3 Representations of Clifford algebras

We now turn our attention to the representations of Clifford algebras. We will only consider Clifford algebras over $\mathbb{R}^3$ and $\mathbb{R}^4$ with a positive-definite inner product from this point forth, since the Seiberg-Witten equations are generally only considered on 3- and 4-manifolds. Additionally, we restrict attention to irreducible representations. (In algebra, these are often called simple Clifford modules.)

The Clifford algebras $\mathrm{Cl}^c(\mathbb{R}^3)$ and $\mathrm{Cl}^c(\mathbb{R}^4)$ have concrete characterisations in terms of the Pauli matrices, which are defined as follows:

$$\sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad \sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \qquad \sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{3.7}$$

Proposition 3.7. *The complexified Clifford algebra of $\mathbb{R}^3$ can be identified (as an ungraded algebra) with $M_2(\mathbb{C}) \oplus M_2(\mathbb{C})$ via the map $e_j \mapsto (i\sigma_j, -i\sigma_j)$, where σ_j is the j-th Pauli matrix. Additionally, the complexified Clifford algebra of $\mathbb{R}^4$ can be identified (as an ungraded algebra) with $M_4(\mathbb{C})$ via the following map:*

$$\begin{aligned} e_0 &\mapsto \begin{pmatrix} 0 & -I_2 \\ I_2 & 0 \end{pmatrix}, \\ e_j &\mapsto \begin{pmatrix} 0 & i\sigma_j \\ i\sigma_j & 0 \end{pmatrix}. \end{aligned} \tag{3.8}$$

These representations will be referred to as the chiral representations.

To identify the irreducible representations, we use the following theorem (see [Lan12]):

Theorem 3.8. *Let $\mathbb{F}$ be a field, and let $n \in \mathbb{N}$. The matrix algebra $M_n(\mathbb{F})$ of $n \times n$ matrices with coefficients in $\mathbb{F}$ has a unique irreducible representation (up to isomorphism), and it is given by the action on $\mathbb{F}^n$ by matrix multiplication. Additionally, a direct sum of matrix algebras $A = M_{n_1}(\mathbb{F}) \oplus \cdots \oplus M_{n_k}(\mathbb{F})$ has precisely k irreducible representations (up to isomorphism), with the j-th irreducible representation given by taking $(B_1, \ldots, B_k) \in A$ to $B_j \in M_{n_j}(\mathbb{F})$ and acting on $\mathbb{F}^{n_j}$.*

As a corollary, we find that (up to isomorphism) there are two irreducible representations of $\mathrm{Cl}^c(\mathbb{R}^3)$, and one of $\mathrm{Cl}^c(\mathbb{R}^4)$. We denote these representations by

$$\rho_3^{\pm} : \mathrm{Cl}^c(\mathbb{R}^3) \to \mathrm{End}(\Delta_3^{\pm}); \tag{3.9}$$

$$\rho_4 : \mathrm{Cl}^c(\mathbb{R}^4) \to \mathrm{End}(\Delta_4), \tag{3.10}$$

where $\Delta_3 \cong \mathbb{C}^2$ and $\Delta_4 \cong \mathbb{C}^4$ as complex vector spaces. In terms of the chiral representations, the representations of $\mathrm{Cl}^c(\mathbb{R}^3)$ correspond to the maps $e_j \mapsto \pm i\sigma_j$, and the representation of $\mathrm{Cl}^c(\mathbb{R}^4)$ is simply given by Equation 3.8. Note also that each of these representations are skew-Hermitian and trace-free when restricted to $\mathbb{R}^3$ or $\mathbb{R}^4$, so we will always assume that $\rho_3^{\pm}|_{\mathbb{R}^3}$ and $\rho_4|_{\mathbb{R}^4}$ are valued in $\mathfrak{su}(\Delta_3^{\pm})$ and $\mathfrak{su}(\Delta_4)$ respectively.

Remark. For now, we will write every representation with subscripts and superscripts to distinguish between them. However, we will eventually drop this convention, and the default meaning of ρ will be ρ_3^+, i.e., the representation of $\mathrm{Cl}^c(\mathbb{R}^3)$ taking $e_j \mapsto ie_j$. Additionally, because $\mathbb{R}^n \subseteq \mathrm{Cl}^c(\mathbb{R}^n)$ is a generating set, we will often abuse notation and denote this representation by $\rho : \mathbb{R}^n \to \mathfrak{su}(\Delta)$ (and simply assume the Clifford relation is satisfied).

In the three-dimensional case, there is a coordinate-free way to distinguish the two representations. Specifically, consider the volume element $\mathrm{vol}_3 = e_1e_2e_3$ of $\mathbb{R}^3$: if $e_j \mapsto \pm i\sigma_j$, then $\mathrm{vol}_3 := e_1e_2e_3 \mapsto \mp i\sigma_1\sigma_2\sigma_3 = \pm 1$. Thus, we say that the two representations have opposite *chirality;* if $\mathrm{vol}_3 \mapsto 1$ then we say the representation has *positive chirality,* and if $\mathrm{vol}_3 \mapsto -1$ then it has *negative chirality.*

Remark. This difference presents another sign ambiguity, as there are multiple conventions for the definition of the volume element in the literature. We have chosen the convention consistent with [MOY96].

In the four-dimensional case, the unique Clifford algebra representation Δ_4 splits into two complementary two-dimensional spaces, which is again related to the complex volume element. If we define $\omega_c := -e_0e_1e_2e_3$ then we see that $\omega_c^2 = 1$ and $\rho(\omega_c)$ is self-adjoint, meaning we can split Δ_4 into two complementary eigenspaces of $\rho_4(\omega_c)$ with eigenvalues ± 1. Furthermore, since $\rho_4(\omega_c)\rho_4(v) = -\rho_4(v)\rho_4(\omega_c)$ for any $v \in \mathbb{R}^4$, the two eigenspaces of $\rho_4(\omega_c)$ are isomorphically exchanged by the action of any v (which means both eigenspaces are two-dimensional). We denote each eigenspace by Δ_4^+ and Δ_4^- respectively. In the chiral representation, ω_c acts on $\mathbb{C}^4$ as the following matrix:

$$\rho_4(\omega_c) = \begin{pmatrix} -I_2 & 0 \\ 0 & I_2 \end{pmatrix}. \tag{3.11}$$

It follows that the eigenspaces are given by $\Delta_4^- = \mathrm{span}\{e_0, e_1\}$ and $\Delta_4^+ = \mathrm{span}\{e_2, e_3\}$. (It is also worth noting that the eigenspaces are preserved by the action of $\mathrm{Spin}^c(4)$, so Δ splits into two inequivalent irreducible representations $\Delta^{\pm}$ of $\mathrm{Spin}^c(4)$.)

There is a correspondence between the three- and four-dimensional Clifford algebra representations, and we will use this later to show that there is a correspondence between the Seiberg-Witten equations on 3- and 4-manifolds. Specifically, we have the following:

Proposition 3.9. *Denote by $i : \mathbb{R}^3 \to \mathbb{R}^4$ the standard inclusion $x \mapsto (0, x)$, and denote by $p : \mathbb{R}^4 \to \mathbb{R}^3$ the standard projection $(t, x) \mapsto x$; these induce inclusions and projections of the corresponding Clifford algebras and spin groups, which we also denote by i and p. Define a representation $\alpha : \mathbb{R}^3 \to \mathfrak{su}(\Delta_4^+)$ as follows:*

$$\alpha(v) = -\rho_4(e_0) \circ \rho_4(i(v)), \tag{3.12}$$

and define a representation $\beta : \mathbb{R}^4 \to \mathfrak{su}(\Delta_3^+ \oplus \Delta_3^+)$ *as follows:*

$$\beta(e_0) = \begin{pmatrix} 0 & I_2 \\ -I_2 & 0 \end{pmatrix}, \tag{3.13}$$

$$\beta(e_j) = \begin{pmatrix} 0 & \rho_3^+(p(e_j)) \\ \rho_3^+(p(e_j)) & 0 \end{pmatrix} \tag{3.14}$$

Then there are maps $f : \Delta_3^+ \to \Delta_4^+$ *and* $g : \Delta_4 \to \Delta_3^+ \oplus \Delta_3^+$ *which constitute equivalences* $\rho_3^+ \sim \alpha$ *and* $\rho_4 \sim \beta$*. Furthermore,* f *and* $g|_{\Delta_4^+}$ *are mutually inverse, and the first and second direct summands of* $\Delta_3^+ \oplus \Delta_3^+$ *can be identified with* Δ_4^- *and* Δ_4^+ *respectively.*

Proof. Here we will write the definitions of f and g; the rest of the proof is straightforward. Without loss of generality, we may write $\Delta_3^{\pm}$ as $\mathbb{C}^2$ and Δ_4 as $\mathbb{C}^4$ with the action given in terms of the Pauli matrices; as noted above, the eigenspaces of ω_c on Δ_4 are then given by $\Delta_4^- = \operatorname{span}\{e_0, e_1\}$ and $\Delta_4^+ = \operatorname{span}\{e_2, e_3\}$. The map $f : \Delta_3^+ \to \Delta_4^+$ is defined to take $e_j \mapsto e_{j+1}$, and the map $g : \Delta_4 \to \Delta_3^+ \oplus \Delta_3^+$ is defined such that $g(z_0, z_1, z_2, z_3) = ((z_0, z_1), (z_2, z_3))$. □

This theorem allows us to essentially identify $\Delta_3^{\pm}$ with $\Delta_4^{\pm}$, and we will omit the subscript 3 and 4 henceforth. Elements of Δ will be called *spinors,* and elements of Δ^+ will be called *positive spinors* in dimension 4.

Remark. The sign ambiguity for positive/negative chirality mentioned above results in differing conventions for the map β in the literature; in particular, $\beta(e_0)$ is sometimes defined to be the negative of the block matrix we have defined.

Though everything we have discussed in this section relates exclusively to Clifford algebras, the group $\mathrm{Spin}^c(n)$ appears naturally as the automorphism group of a Clifford module.

Definition 3.10. Let $\rho_i : V_i \to \mathfrak{su}(\Delta_i)$ be equivalent complex irreducible representations of $\mathrm{Cl}(V_i)$ for $i = 1, 2$, where the V_i are real oriented inner product spaces and the Δ_i are complex inner product spaces. An *isomorphism from* ρ_1 *to* ρ_2 is a pair consisting of an oriented isometric isomorphism $T : V_1 \to V_2$ and a unitary isomorphism $\alpha : \Delta_1 \to \Delta_2$ intertwining ρ_1 and $T^*\rho_2$, meaning $\rho_2(Tv) \circ \alpha = \alpha \circ \rho_1(v)$ for every $v \in V_1$.

Proposition 3.11. *The group* $\mathrm{Aut}(\rho)$ *of automorphisms of* $\rho : V \to \mathfrak{su}(\Delta)$ *is canonically isomorphic to* $\mathrm{Spin}^c(V)$*. The space* $\mathrm{Iso}(\rho_1, \rho_2)$ *of isomorphisms from* $\rho_1 : V_1 \to \mathfrak{su}(\Delta_1)$ *to* $\rho_2 : V_2 \to \mathfrak{su}(\Delta_2)$ *is a torsor on* $\mathrm{Spin}^c(V_1) \cong \mathrm{Spin}^c(V_2)$*.*

Proof. Recall the canonical projection $q_{\mathrm{SO}} : \mathrm{Spin}^c(V) \to \mathrm{SO}(V)$. It follows from the structure of $\mathfrak{spin}^c(n)$ that $\rho|_{\mathrm{Spin}^c(V)}$ is valued in $\mathrm{U}(\Delta)$, so we define the map $F : \mathrm{Spin}^c(V) \to \mathrm{Aut}(\rho)$ as follows:

$$F(g) = (q_{\mathrm{SO}}(g), \rho(g)). \tag{3.15}$$

Note that $\rho(g)$ is indeed an intertwining map, since $\rho(q_{\mathrm{SO}}(g)v) = \rho(gvg^{-1}) = \rho(g)\rho(v)\rho(g)^{-1}$ for any $v \in \mathbb{R}^n$. It is clear that F is a group homomorphism, so we prove that it is an isomorphism:

- Suppose $F(g) = (\mathrm{id}, \mathrm{id})$. If conjugation by g acts as the identity on V then g must be a scalar. But the only scalar for which $\rho(g) = 1$ is 1; this proves that F is injective.

- Suppose $T \in \mathrm{SO}(V)$ and $\alpha \in \mathrm{U}(\Delta)$ constitute an automorphism of ρ, meaning α intertwines ρ and $T^*\rho$. Since $\mathrm{Spin}^c(V)$ covers $\mathrm{SO}(V)$, we know there is some $g \in \mathrm{Spin}^c(n)$ such that $q_{\mathrm{SO}}(g) = T$. Then $F(g) = (T, \rho(g))$, so the action of $\rho(g)$ intertwines ρ and $T^*\rho$ as well. By Schur's lemma, $\rho(g)$ and α must differ by some $e^{i\theta} \in \mathrm{U}(1)$, so $\rho(e^{-i\theta}g) = \alpha$ and $q_{\mathrm{SO}}(e^{-i\theta}g) = T$ (the latter since q_{SO} is U(1)-invariant).

The statement about isomorphisms follows immediately. □

3.1.4 The quadratic map

One of the key components of the Seiberg-Witten equations is a map which takes as input a spinor field ψ and outputs a differential form, and which varies quadratically with respect to the input. In both the 3- and 4-dimensional cases, this can be interpreted as a differential form naturally associated to the trace-free part of the self-adjoint endomorphism $\psi \otimes \psi^*$, where ψ^* denotes the Hermitian conjugate of ψ according to the inner product on Δ. To write down this quadratic map, we will use the following two lemmas:

Lemma 3.12. *The representation $\rho_3 : \Lambda^1(\mathbb{R}^3) \to \mathfrak{su}(\Delta^+)$ constitutes a linear isomorphism between $\Lambda^1(\mathbb{R}^3)$ and $\mathfrak{su}(\Delta^+)$.*

Proof. Both $\Lambda^1(\mathbb{R}^3)$ and $\mathfrak{su}(\Delta^+)$ are 3-dimensional, so we need only prove that ρ_3 is injective. But $\rho_3(e^j) = i\sigma^j$ in the chiral representation, and the Pauli matrices are linearly independent. □

In dimension 4, the Hodge star $* : \Lambda^*(\mathbb{R}^4) \to \Lambda^*(\mathbb{R}^4)$ acts as an involution on the subspace of bivectors. We denote by $\Lambda^2_\pm(\mathbb{R}^4)$ the subspace of bivectors α satisfying $*\alpha = \pm\alpha$; these two subspaces are complementary in $\Lambda^2(\mathbb{R}^4)$, and we call them self-dual and anti-self-dual respectively.

Lemma 3.13. *The representation $\rho_4 : \mathrm{Cl}(\mathbb{R}^4) \to \mathrm{End}(\Delta)$ restricts to an isomorphism*[1] $\Lambda^2_\pm(\mathbb{R}^4) \cong \mathfrak{su}(\Delta^\pm)$.

Proof. Firstly, note that $\rho_4(\alpha)$ does indeed preserve the splitting $\Delta = \Delta^- \oplus \Delta^+$ whenever $\alpha \in \Lambda^2(\mathbb{R}^4)$; this is clear when ρ_4 is taken to be the chiral representation. Moreover, the spaces $\Lambda^2_\pm(\mathbb{R}^4)$ and $\mathfrak{su}(\Delta^\pm)$ are both three-dimensional, so we need only verify that ρ_4 is injective on these subspaces. But this can be easily verified in the chiral representation. □

It follows that any trace-free self-adjoint endomorphism of Δ^+ can be identified with an element of $i\Lambda^1(\mathbb{R}^3)$ in three dimensions, or with an element of $i\Lambda^2_+(\mathbb{R}^4)$ in four dimensions. This leads to the following definition:

Definition 3.14. The *three-dimensional quadratic map* is the map $q : \Delta^+ \to i\Lambda^1(\mathbb{R}^3)$ defined as follows:

$$q(\psi) = \rho_3^{-1}(\psi \otimes \psi^*)_0, \tag{3.16}$$

[1] We are regarding bivectors as elements of $\mathrm{Cl}(\mathbb{R}^4)$ on the basis that $\mathrm{Cl}(\mathbb{R}^4)$ and $\Lambda^*(\mathbb{R}^4)$ are linearly isomorphic.

where the subscript zero denotes the trace-free part of the endomorphism. Similarly, the *four-dimensional quadratic map* is the map $q' : \Delta^+ \to i\Lambda^2_+(\mathbb{R}^4)$ defined as follows:

$$q'(\phi) = \rho_4^{-1}(\phi \otimes \phi^*)_0. \tag{3.17}$$

The three-dimensional map has a useful algebraic characterisation as a quadratic "adjoint" to ρ_3.

Proposition 3.15. *For any $\psi \in \Delta^+$ and any $\alpha \in \Lambda^1(\mathbb{R}^3)$, we have the following:*

$$\langle q(\psi), \alpha \rangle_{\Lambda^1} = \tfrac{1}{2}\langle \rho_3(\alpha) \cdot \psi, \psi \rangle_{\Delta^+}. \tag{3.18}$$

Proof. It suffices to show that the following equation holds in the chiral representation:

$$(\psi \otimes \psi^*)_0 = \tfrac{1}{2}\rho_3 \left(\sum_j \langle i\sigma_j \cdot \psi, \psi \rangle e_j \right). \tag{3.19}$$

Using the linearity of ρ_3, and writing $\psi = (\alpha, \beta)^T$ for $\alpha, \beta \in \mathbb{C}$, this is equivalent to the following matrix equation:

$$\frac{1}{2}\begin{pmatrix} |\alpha|^2 - |\beta|^2 & 2\alpha\overline{\beta} \\ 2\overline{\alpha}\beta & |\beta|^2 - |\alpha|^2 \end{pmatrix} = \sum_j \langle \sigma_j \psi, \psi \rangle \sigma_j. \tag{3.20}$$

But this equation follows from the definitions of the Pauli matrices. □

Remark. This characterisation is often taken to be the definition of q in the literature, but some authors include a negative sign in Equation 3.18. This is a manifestation of the differing conventions for positive and negative chirality.

3.2 Spin^c Geometry

In formulating the Seiberg-Witten equations, we essentially parameterise all of the preceding constructions over 3- and 4-manifolds. The key geometric structure that allows for this parametrisation is a Spin^c-structure.

3.2.1 Spin^c-structures and spinor bundles

Let (M, g) be a connected Riemannian n-manifold, and let $\pi_{\mathrm{SO}} : F_{\mathrm{SO}} \to M$ denote the bundle of oriented orthonormal frames on M. As above, denote by $q_{\mathrm{SO}} : \mathrm{Spin}^c(n) \to \mathrm{SO}(n)$ the canonical surjection. There are two key structures on M in the spin geometry of the Seiberg-Witten equations.

Definition 3.16. A Spin^c*-structure on M* is a principal $\mathrm{Spin}^c(n)$-bundle $\pi : P \to Y$ with a $\mathrm{Spin}^c(n)$-invariant bundle surjection $\widehat{\pi} : P \twoheadrightarrow F_{\mathrm{SO}}$. That is, $\widehat{\pi}$ must satisfy $\widehat{\pi}(g \cdot p) = q_{\mathrm{SO}}(g) \cdot \widehat{\pi}(p)$ for all $g \in \mathrm{Spin}^c(n)$ and all $p \in P$, as well as $\pi_{\mathrm{SO}} \circ \widehat{\pi} = \pi$.

Definition 3.17. A *spinor bundle over M* is a Hermitian vector bundle $W \to M$ together with a *Clifford module structure,* i.e., a smoothly varying irreducible representation $\rho : T^*M \to \mathfrak{su}(W)$ of the Clifford algebra of T^*M. Explicitly, we have a map $\rho : T^*_xM \to \mathfrak{su}(W_x)$ for each $x \in Y$ such that $\rho(\alpha) \circ \rho(\alpha) = -g(\alpha, \alpha)\mathrm{id}_{W_x}$ for every $\alpha \in T^*_xM$. A section of W is called a *spinor.* We will always assume that ρ is isomorphic to the *positive* spinor representation in odd dimensions; for a 3-manifold Y, this means that $\rho(\mathrm{vol}_Y) = 1$.

In the light of Proposition 3.11, there is a natural correspondence between the two notions. In what follows, we regard each element $e_x \in F_{\mathrm{SO}}$ as an isomorphism $e_x : T_xM \to \mathbb{R}^n$ via fibrewise basis decomposition.

- Given a Spin^c-structure $P \to M$, we can form a spinor bundle over M by taking the associated vector bundle $P \times_\rho \Delta$, where $\rho : \mathbb{R}^n \to \mathfrak{su}(\Delta)$ is an irreducible complex Clifford representation. The Clifford module structure is defined on a fibre over $x \in M$ in the natural way:
$$\rho(\alpha) \cdot [(p, \psi)] = [(p, \rho(\widehat{\pi}(p) \cdot \alpha) \cdot \psi)]. \tag{3.21}$$
Note that $\rho|_{\mathrm{Spin}^c(n)}$ is valued in $\mathrm{U}(\Delta)$, which gives the bundle a natural Hermitian structure.

- Conversely, given a spinor bundle $W \to M$, we get a Spin^c-structure by taking the space of all fibrewise isomorphisms from the representation $\rho : T^*_xM \to \mathfrak{su}(W_x)$ to the standard representation $\rho : \mathbb{R}^n \to \Delta$; by Proposition 3.11, this is a Spin$^c(n)$-torsor over each point. The projection $\widehat{\pi}$ onto F_{SO} is given by taking the corresponding isomorphism $T^*_xM \to \mathbb{R}^n$.

As such, we can always translate between a Spin^c-structure and a spinor bundle, so we will essentially identify them henceforth. We will think of them as the “background” geometry for the Seiberg-Witten equations.

Upon encountering the definition of a Spin^c-structure, a reasonable concern is that any given manifold might not even admit such a structure. Indeed, there are some manifolds that do not admit Spin^c-structures; these spaces can be identified with the Stiefel-Whitney classes (see [LM89]):

Proposition 3.18. *The Riemannian n-manifold M admits a* Spin^c *structure if and only if its second Stiefel-Whitney class $w_2(TM)$ is the* mod 2 *reduction of an integral class (meaning the map $H^2(M; \mathbb{Z}) \to H^2(M; \mathbb{Z}_2)$ induced by the projection $\mathbb{Z} \to \mathbb{Z}_2$ on cochains contains $w_2(TM)$ in its image).*

In particular, since the tangent bundle of a closed oriented 3-manifold is trivial [KM07], the Stiefel-Whitney classes of such a manifold all vanish, meaning every closed 3-manifold has a Spin^c-structure. Once a Spin^c-structure is found, the following result allows us to obtain all remaining Spin^c-structures:

Proposition 3.19. *If an n-manifold M has a* Spin^c*-structure, then the space of all* Spin^c*-structures on M is a* Pic$^t(M)$*-torsor. The action of $L \in$* Pic$^t(M)$ *on a spinor bundle $S \to M$ is given by the spinor bundle $L \otimes S \to M$, with Clifford module structure given by $u \otimes \psi \mapsto u \otimes (\rho(\omega) \cdot \psi)$ for every $\omega \in T^*_xM$, $u \in L_x$, $\psi \in S_x$.*

Proof. The proof is adapted from [KM07]. First, we verify that the given action is well-defined.

- If L and L' are equivalent as line bundles by $\alpha : L \to L'$, the following map is an equivalence between $L \otimes S$ and $L' \otimes S$:

$$u \otimes \psi \mapsto \alpha(u) \otimes \psi. \tag{3.22}$$

- Since the Clifford action on $L \otimes S$ commutes with the tensor product, it clearly satisfies the Clifford relation. Additionally, by counting dimensions, the resulting Clifford module is still irreducible under the tensor product, and it clearly has the same chirality.

Thus, the action of $\mathrm{Pic}^t(M)$ on the space of Spin^c-structures is well-defined. We now prove that it is free and transitive:

- *Transitivity:* Let S_1 and S_2 be spinor bundles with Clifford module structures $\rho_i : T^*M \to \mathfrak{su}(S_i)$. We define the complex line bundle E_{12} to be the bundle consisting of fibrewise linear maps $f_x : (S_1)_x \to (S_2)_x$ which intertwine ρ_1 and ρ_2. (Note that E_{12} is a complex line bundle by Schur's lemma.) Now, observe that $E_{12} \otimes S_1$ is equivalent as a spinor bundle to S_2 under the following map:

$$(f_x, \psi_x) \mapsto f_x(\psi_x) \text{ for } f_x \in (E_{12})_x \text{ and } \psi_x \in (S_1)_x. \tag{3.23}$$

- *Freeness:* Suppose S is a spinor bundle and E is a line bundle for which $E \otimes S$ and S are equivalent spinor bundles, with equivalence given by $\alpha : E \otimes S \to S$. Now, fix some $u_x \in E_x$. By Schur's lemma, the map $S_x \to S_x$ given by $\psi_x \mapsto \alpha(u_x \otimes \psi_x)$ for each $\psi_x \in S_x$ can be represented by a complex number $f(u_x)$ (equal to zero if and only if $u_x = 0$); thus, we must have that $\alpha(u_x \otimes \psi_x) = f(u_x)\psi_x$. The map f constitutes a smoothly varying fibrewise linear isomorphism from E_x to $\mathbb{C}$, meaning it extends to a bundle equivalence from E to the trivial bundle, which implies that $c_1(E) = 0$. □

As usual, automorphisms of a given Spin^c-structure are called gauge transformations.

Definition 3.20. A *gauge transformation of* W is a section $u \in \Gamma(\mathrm{End}(W))$ which is a fibrewise unitary automorphism respecting the Clifford module structure (i.e., $u(\rho(\alpha)\cdot\psi) = \rho(\alpha)\cdot u(\psi)$ for all $\alpha \in T^*M$ and $\psi \in W$). Equivalently, a *gauge transformation of a* Spin^c-*structure* P is an automorphism of P under which F_{SO} is invariant. The collection of all such maps is denoted by $\mathscr{G}$.

As we noted above, the irreducibility of Δ implies that a map from Δ to itself respecting the action of a complexified Clifford algebra must be scalar multiplication. It follows that u can be considered to be an element of $\mathrm{U}(1)$ at each point, meaning we can identify $\mathscr{G}$ with $\mathrm{Map}(M, \mathrm{U}(1))$. Thus, even though Spin^c geometry is based on a principal Spin$^c(n)$-bundle, the structure group is essentially $\mathrm{U}(1)$.

3.2.2 Connections on spinor bundles

The notion of a Spin^c-compatible connection comes in two forms, depending on whether we conceptualise a Spin^c-structure in terms of spinor bundles or principal bundles. We begin with the spinor description.

Definition 3.21. Let ∇^{SO} be an SO(n)-connection on TY, that is, a connection derived from a connection 1-form on F_{SO}. A Hermitian connection A on the spinor bundle W is called *spinorial with respect to* ∇^{SO} if the following holds for every $\alpha \in \Omega^1(M)$ and every $\psi \in \Gamma(W)$:

$$\nabla^A(\rho(\alpha)\cdot\psi) = \rho(\alpha)\cdot(\nabla^A\psi) + \rho(\nabla^{\mathrm{SO}}\alpha)\cdot\psi. \tag{3.24}$$

More concisely, we require that $[\nabla^A, \rho(\alpha)] = \rho(\nabla^{\mathrm{SO}}\alpha)$. We denote by $\mathscr{A}(W)$ the space of all such connections. There is a natural action of $\mathscr{G}$ on $\mathscr{A}(W)$ given by conjugation by u, and this action is clearly compatible with the action on $\Gamma(W)$ (in that $\nabla^{u(A)}u(\psi) = u(\nabla^A\psi)$ for any connection A and any section ψ).

It is fairly easy to see that $\mathscr{A}(W)$ is an affine space modelled on $i\Omega^1(M)$. For any $\eta \in \Omega^1(M)$ we know that $A + i\eta$ will be Hermitian and spinorial; the fact that $i\eta$ is purely imaginary ensures that $A + i\eta$ will be Hermitian, and $[\nabla^A + i\eta, \rho(\alpha)] = [\nabla^A, \rho(\alpha)] + [i\eta, \rho(\alpha)] = \rho(\nabla^{\mathrm{SO}}\alpha)$. Moreover, if A and A' are two connections, their difference will be an End(W)-valued 1-form that commutes with any α and therefore must be scalar multiplication by an element of U(1). Finally, we know that a spinorial connection always exists by a partition of unity argument and the fact that a spinorial connection always exists locally (see Equation 3.26).

Definition 3.22. Let $\omega_{\mathrm{SO}} \in \Omega^1(F_{\mathrm{SO}}; \mathfrak{so}(n))$ denote the connection form of some SO(n)-connection. A *spinorial connection form on* P is a connection form $\alpha \in \Omega^1(P; \mathfrak{spin}^c(n))$ which can be written in the form

$$\alpha = \mathfrak{a}^{-1}((q_{\mathrm{SO}})^*\omega_{\mathrm{SO}}) + i\widetilde{A}, \tag{3.25}$$

where $\widetilde{A} \in \Omega^1(P;\mathbb{R})$ is a Spin$^c(n)$-invariant 1-form satisfying $i\widetilde{A}(\widetilde{\xi}_p) = \frac{1}{2}\mathrm{tr}(\mathfrak{a}(\xi))$ for any $\xi \in \mathfrak{spin}^c(n)$, and where $\widetilde{\xi}$ is the fundamental vector field on P generated by the action of ξ. The gauge group acts on such connections by taking $iA \mapsto iA + 2g^{-1}dg$ for any $g \in \mathscr{G}$.

This definition has the benefit that connection matrices may be computed in local trivialisations with relative ease. By applying the correspondence between connection forms and connection matrices, and writing ω_{SO} in an orthonormal basis $\{e_i\}$ with the summation convention, we find that the connection matrix is as follows:

$$\begin{aligned}
[\alpha] &= \rho(\mathfrak{a}^{-1}((\omega_{\mathrm{SO}})_i^j e^i \otimes e_j)) + \mathfrak{a}^{-1}(iA) && (3.26)\\
&= \rho(\mathfrak{a}^{-1}(\tfrac{1}{2}(\omega_{\mathrm{SO}})_{ij} e_i \wedge e_j)) + \tfrac{i}{2}A && (\text{where } e_i \wedge e_j = e_i \otimes e^j - e_j \otimes e^i)\\
&= \tfrac{1}{4}(\omega_{\mathrm{SO}})_{ij}\rho(e_i e_j) + \tfrac{i}{2}A && (\text{definition of } \mathfrak{a})\\
&= \tfrac{1}{8}(\omega_{\mathrm{SO}})_{ij}\rho(e_i \wedge e_j) + \tfrac{i}{2}A. && (\text{where } e_i \wedge e_j = \tfrac{1}{2}(e_i e_j - e_j e_i))
\end{aligned}$$

This argument also clearly shows that spinorial connection forms always exist locally, and that the collection of all such forms is an affine space over $i\Omega^1(Y)$: we just take $A = 0$ to get a suitable connection form, and all others differ from this connection form by a purely imaginary differential form.

It remains to show that these two notions of connection are indeed equivalent; this is the content of the following proposition.

Proposition 3.23. *A spinorial connection form α on P locally induces a spinorial connection* $\mathrm{d} + [\alpha]$ *on W.*

Proof. The connection matrix for the Levi-Civita connection acting on a vector v can be written as follows:

$$\begin{aligned}\omega(v) &= \omega_{ij} v^i e_j = \omega_{ij}(e_j \otimes e^i)(v) \\ &= \tfrac{1}{2}\omega_{ij}(e_j \otimes e^i - e_i \otimes e^j)(v). \qquad\qquad (\omega_{ij} = -\omega_{ji} \text{ since } \omega \in \mathfrak{so}_n)\end{aligned}$$

We also have that $e_i e_j v - v e_i e_j = 2(e_j \otimes e^i - e_i \otimes e^j)(v)$. The remainder of the proof is a computation:

$$\begin{aligned}&\nabla^\alpha(\rho(v)\cdot\psi) - \rho(v)\cdot\nabla^\alpha\psi \\ &= \left(\mathrm{d}(\rho(v)\cdot\psi) + \tfrac{1}{4}(\omega_{\mathrm{SO}})_{ij}\rho(e_i e_j) + \tfrac{i}{2}A\rho(v)\cdot\psi\right) - \rho(v)\cdot\left(\mathrm{d}(\psi) + \tfrac{1}{4}(\omega_{\mathrm{SO}})_{ij}\rho(e_i e_j) + \tfrac{i}{2}A\psi\right) \\ &= \rho(\mathrm{d}v)\cdot\psi + \tfrac{1}{4}(\omega_{\mathrm{SO}})_{ij}\rho(e_i e_j v - v e_i e_j)\cdot\psi \\ &= (\rho(\mathrm{d}v) + \tfrac{1}{2}(\omega_{\mathrm{SO}})_{ij}\rho(e_j\otimes e^i - e_i \otimes e^j))\cdot\psi = \rho(\nabla^{\mathrm{SO}}v)\cdot\psi,\end{aligned}$$

where the last two lines were direct consequences of the two identities we mentioned. □

Since both spaces of connections are affinely modelled on $i\Omega^1(M)$, we conclude that the two types of connection coincide.

A new perspective on the space of spinorial connections is offered by the notion of a determinant bundle.

Definition 3.24. Let $P \to M$ be a Spin^c-structure. The *determinant principal bundle* P_{det} is the U(1) bundle defined by taking $P/\mathrm{Spin}(n)$ (via the inclusion $i : \mathrm{Spin}(n) \to \mathrm{Spin}^c(n)$), with the natural projection induced by the quotient. By taking the canonical representation of U(1) in $\mathbb{C}$ and taking the associated bundle, we obtain the *determinant line bundle* $L_{\mathrm{det}} := P_{\mathrm{det}} \times_{\mathrm{U}(1)} \mathbb{C}$. Note that the gauge group $\mathscr{G}$ acts on P_{det} and L_{det} as follows: for $g \in \mathscr{G}$, we take $p \in P_{\mathrm{det}}$ to $g^2 \cdot p$, and $z \in L_{\mathrm{det}}$ to $g^2 \cdot z$ (this is a direct consequence of the definition of q_{U} in Equation 3.5).

Remark. An equivalent definition for the determinant line bundle is the top exterior power of the spinor bundle.

Remark. It was shown in Proposition 3.19 that any two spinor bundles on a manifold differ by tensoring with a line bundle. Under tensor product with the line bundle L, the determinant line bundle is transformed by tensoring with $L \otimes L$; see [Sal14] for details on this fact.

Let $\alpha = \mathfrak{a}^{-1}((q_{\mathrm{SO}})^*\omega_{\mathrm{SO}}) + i\widetilde{A}$ be a spinorial connection form. Since the 1-form $i\widetilde{A}$ is Spin$^c(n)$-invariant, it descends to a 1-form A_{det} on P_{det}. Note that this does not define a connection on P_{det}: the U(1)-factor is effectively squared under the quotient by Spin(n), meaning

$$iA_{\mathrm{det}}(\widehat{it}_p) = \frac{1}{2}it,$$

for any $it \in \mathfrak{u}(1)$. In any case, a spinorial connection can essentially be reduced to a U(1)-invariant imaginary 1-form on P_{det}. Given such a 1-form $\widetilde{A} \in \Omega^1(P_{\text{det}}; i\mathbb{R})$, we can recover a connection on P simply by taking the pullback $(q_{\text{U}})^*\widetilde{A}$ along the quotient and adding $(q_{\text{SO}})^*\omega_{\text{SO}}$ back in. Conversely, given a spinorial connection, we can take the trace of the connection matrix in a local trivialisation to get the invariant 1-form on P_{det}; by the computation above, the component corresponding to the Levi-Civita connection is proportional to the trace-free endomorphism $\rho(e_i \wedge e_j)$. We therefore make the following definition:

Definition 3.25. The *curvature 2-form* of a spinorial connection A is given by the trace of $\frac{1}{2}F_{\nabla^A} \in \Omega^2(\text{End}(W))$ when considered as a matrix-valued 2-form. Equivalently, it is the curvature 2-form[2] of the corresponding invariant 1-form $\widetilde{A}$ on P_{det}.

From this point forth, F_A will always denote the traced version, i.e., F_A will be an element of $i\Omega^2(M)$. If we wish to refer to the curvature of the spinorial connection ∇^A as an endomorphism of W, we will use the notation $F_{\nabla^A} \in \Omega^2(\text{End}(W))$.

3.2.3 The Dirac operator

A connection on a spinor bundle naturally induces a differential operator on sections called the Dirac operator. It was first introduced by Dirac on flat spacetime as a "square root" of the Laplacian, but was later adjusted and generalised to manifolds.

Definition 3.26. Denote by $\delta : \Gamma(T^*M \otimes W) \to \Gamma(W)$ the map given by the Clifford module structure, and let $A \in \mathscr{A}(W)$ be a spinorial connection. The *Dirac operator induced by A* is the differential operator $D_A : \Gamma(W) \to \Gamma(W)$ defined as follows:

$$D_A = \delta \circ \nabla^A.$$

In an orthonormal oriented frame $\{e_i\}$, we can equivalently express this operator as $D_A\psi = \sum_i \rho(e_i) \cdot \nabla_{e_i}\psi$.

Physically, the Dirac operator informs the dynamics of spin-1/2 fermions; the classical field corresponding to a massless fermion under the influence of an electromagnetic field represented by A naturally lies in the kernel of D_A. Elements of $\ker(D_A)$ are called *harmonic spinors,* and the PDE $D_A\psi = 0$ is called the *harmonic equation.*

The Dirac operator has several convenient analytic properties. We state three of these which will be useful later.

Proposition 3.27. *If ∇^{SO} is the Levi-Civita connection, then the Dirac operator for a ∇^{SO}-spinorial connection is formally self-adjoint. That is, for a given connection $A \in \mathscr{A}(W)$, the following holds for every $\phi, \psi \in \Gamma(W)$:*

$$\langle D_A\phi, \psi\rangle = \langle \phi, D_A\psi\rangle. \tag{3.27}$$

Proof. Around any point in M, there exists a coordinate frame for which the Christoffel symbols for the Levi-Civita connection are all zero; let $\{e_i\}$ be such a frame. The formal

[2] Note that $\widetilde{A}$ does not define a connection on P_{det} in general, but its curvature can still be defined as $d\widetilde{A}$. This still descends to a well-defined 2-form on M.

self-adjointness of D_A then simply follows from the fact that ∇^A is a Hermitian connection and $\rho(e_i)$ is anti-self-adjoint:

$$\begin{aligned}
\langle D_A\phi, \psi\rangle &= \sum_j \langle \rho(e_j)\cdot \nabla^A_{e_j}\phi, \psi\rangle \\
&= -\sum_j \langle \nabla^A_{e_j}\phi, \rho(e_j)\cdot\psi\rangle \\
&= \sum_j \langle \phi, \nabla^A_{e_j}(\rho(e_j)\cdot\psi)\rangle \\
&= \sum_j \langle \phi, \rho(e_j)\cdot\nabla^A_{e_j}\psi\rangle + \langle \phi, \rho(\nabla^{\mathrm{SO}}_{e_j}e_j)\cdot\psi\rangle \\
&= \langle \phi, D_A\psi\rangle.
\end{aligned}$$

□

Proposition 3.28. *The symbol of the Dirac operator is given by $\xi \mapsto \rho(\xi)$ for every $\xi \in \Omega^1(M)$. As a consequence, the Dirac operator is elliptic: its principal symbol on nonzero 1-forms is always an isomorphism.*

Proof. The principal symbol is given by $\xi \mapsto \sum_i \xi_i\rho(e_i) = \rho(\sum_i \xi_i e_i) = \rho(\xi) \in \mathrm{End}(W)$. When ξ is nonzero, the inverse of this endomorphism is given by $-\frac{1}{|\xi|^2}\rho(\xi) \in \mathrm{End}(W)$. □

Proposition 3.29 (Unique continuation)**.** *The Dirac operator has the unique continuation property. That is, any harmonic spinor ψ vanishing on a nonempty open set is the zero spinor.*

For a proof, refer to [BW93]. The analogue of this property for the Dolbeault operator $\overline{\partial}$ on a complex manifold is the well-known identity theorem.

In the special case that the $\mathrm{SO}(n)$-connection is the Levi-Civita connection on a 3-manifold, we can characterise some of its behaviour with respect to the Clifford product:

Proposition 3.30. *If ∇^{SO} is the Levi-Civita connection on a 3-manifold Y and $\alpha \in \Omega^1(Y)$ is a 1-form with dual vector field $\alpha^\flat \in \mathfrak{X}(Y)$, then the following anticommutator equation holds:*

$$\{D_A, \rho(\alpha)\} = \rho((*\mathrm{d} + \mathrm{d}*)\alpha) - 2\nabla_{\alpha^\flat}. \tag{3.28}$$

Proof. Once again, we choose an orthonormal coframe $\{e^0, e^1, e^2\}$ for which all Christoffel symbols vanish at a given point. Writing α locally as $\sum_i a_i e^i$ where a_i are real-valued functions, we see that

$$\begin{aligned}
\{D_A, \rho(\alpha)\}(\psi) &= D_A(\rho(\alpha)\psi) + \rho(\alpha)D_A\psi \\
&= \sum_j \left[\rho(e^j)\nabla_j\left(\sum_i a_i\rho(e^i)\psi\right) + \rho(\alpha)\rho(e^j)\nabla_j\psi\right] \\
&= \sum_j \left[\rho(e^j)\sum_i\left(\rho(e^i)\frac{\partial a_i}{\partial x^j}\right) + \rho(e^j)\rho(\alpha)\nabla_j\psi + \rho(\alpha)\rho(e^j)\nabla_j\psi\right] \\
&\qquad\qquad (\text{since } \Gamma^i_{jk} = 0) \\
&= \sum_{i,j}\left[\rho(e^je^i)\frac{\partial a_i}{\partial x^j}\right] + \sum_j\left[\rho(\{e^j, \alpha\})\nabla_j\psi\right].
\end{aligned}$$

We claim that the first and second term of this equation match the first and second term in Equation 3.28, beginning with the second term. The anticommutator of two grade-1 elements of the Clifford algebra is known to be given by $\{e^i, e^j\} = -2g^{ij}$, where g^{ij} are the matrix elements of the inverse metric; thus, we see that $\{e^j, \alpha\} = -2\sum_i g^{ij} a_i = -2e^j(\alpha^\flat)$, meaning the second term is simply $-2\sum_j e^j(\alpha^\flat)\nabla_j \psi = -2\nabla_{\alpha^\flat}\psi$.

For the first term, we note that $*\mathrm{d}\alpha = (\nabla \times \alpha^\sharp)^\flat$ and $\mathrm{d} * \alpha = *\nabla \cdot \alpha^\sharp$, where $\nabla \times \alpha^\sharp = \sum_{i,j,k} e_i \varepsilon^{ijk} \partial_j \alpha_k$ and $\nabla \cdot \alpha^\sharp = \sum_{i,j} \delta^{ij} \partial_i a_j$, and where ε^{ijk} is the Levi-Civita symbol. On the other hand, we have that $e^j e^i = -\delta^{ij} - e^i \wedge e^j = -\delta^{ij} - * \sum_k \varepsilon^{ijk} e^k$. From these formulas, the correspondence between the two terms follows easily. □

3.3 The Seiberg-Witten Equations

Now that we have introduced the necessary machinery, we are ready to discuss the Seiberg-Witten equations. We begin by discussing the three-dimensional equations.

3.3.1 Three-dimensional equations

In the three-dimensional theory, we denote our underlying 3-manifold by Y, with Spin^c-structure $P_Y \to M$ and spinor bundle $W \to Y$. If we wish to avoid ambiguity then we will denote the Clifford module structure by $\rho_Y : T^*Y \to \mathfrak{su}(W)$; however, we will usually drop the subscript. With this data, we define the *pre-configuration space* to consist of pairs of sections and (spinorial) connections:

$$\mathscr{C}(W) := \Gamma(W) \times \mathscr{A}(W). \tag{3.29}$$

There is a natural action of the gauge group $\mathscr{G}$ on the pre-configuration space: the gauge transformation $g \in \mathscr{G}$ acts on a pair (ψ, A) as follows:

$$(\psi, A) \mapsto (g\psi, A + 2g^{-1}\mathrm{d}g). \tag{3.30}$$

A pair is called *irreducible* if it is not fixed by any element of $\mathscr{G}$; otherwise, the pair is called *reducible*. We denote by $\mathscr{C}^*(W)$ the collection of all irreducible elements. Just as for vortices, the section ψ is what dictates reducibility:

Proposition 3.31. *A gauge transformation $g \in \mathscr{G}$ fixes some pair if and only if g is constant. A pair $(\psi, A) \in \mathscr{C}(W)$ is fixed by some gauge transformation if and only if ψ is identically zero.*

We define the *configuration space* to be the quotient of $\mathscr{C}(W)$ by the action of $\mathscr{G}$, and we denote it by $\mathscr{B}(W)$. Similarly, the irreducible configuration space is denoted by $\mathscr{B}^*(W)$.

With this preparation, we write the Seiberg-Witten equations on 3-manifolds.

Definition 3.32. A connection $A \in \mathscr{A}(W)$ and a spinor $\psi \in \Gamma(W)$ are said to satisfy the *3-dimensional Seiberg-Witten equations* if the following two equations hold:

$$*F_A + q(\psi) = 0; \tag{3.31a}$$

$$D_A\psi = 0. \tag{3.31b}$$

The pair (ψ, A) will be referred to as a Seiberg-Witten monopole, or simply a monopole. Note that the equations are $\mathscr{G}$-invariant, i.e., each solution (ψ, A) to the equations gives rise to another solution $(u\psi, u(A))$ for any $u \in \mathscr{G}$.

We will be interested in the space of solutions to the Seiberg-Witten equations modulo gauge equivalence.

Definition 3.33. The *moduli space of Seiberg-Witten monopoles on Y* is the space of monopoles up to the action of the gauge group $\mathscr{G}$. It is denoted by $\mathscr{M}_{\text{sw}}$; the 3-manifold and spinor bundle are generally clear from context. The irreducible component of the moduli space is denoted by $\mathscr{M}^*_{\text{sw}}$.

3.3.2 The Chern-Simons-Dirac functional

It turns out that there is a natural interpretation of the three-dimensional Seiberg-Witten equations in terms of an action functional.

Definition 3.34. Fix a reference connection $A_0 \in \mathscr{A}(W)$. The *Chern-Simons-Dirac functional*[3] is the functional $\text{csd} : \mathscr{C}(W) \to \mathbb{R}$ defined as follows:

$$\text{csd}(\psi, A) = \frac{1}{2}\int_Y \Big(\text{Re}\langle\psi, D_A\psi\rangle \text{vol}_g - (A - A_0) \wedge (F_A + F_{A_0})\Big). \tag{3.32}$$

Remark. The constant factors on each term in the Chern-Simons-Dirac functional vary across the literature. In [MOY96], the functional is twice our definition, while in [KM07], the term with no ψ dependence is modulated by a factor of 1/4. Differing conventions will result in different forms of the equation $*F_A + q(\psi) = 0$

The reason we speak unambiguously of "the" Chern-Simons-Dirac functional, despite its seeming dependence on a reference connection, is the following proposition:

Proposition 3.35. *Given two reference connections $A_0, A_1 \in \mathscr{A}(W)$, the corresponding Chern-Simons-Dirac functionals* csd_0 *and* csd_1 *differ by a global constant which depends only on A_0 and A_1, and not the input connection and section.*

Proof. Clearly the term involving ψ has no dependence on the reference section. Let $A_{\text{in}} \in \mathscr{A}(W)$ be the input connection, and define $\alpha = A_1 - A_0$ and $\beta = A_{\text{in}} - A_0$, so that $F_{A_1} = F_{A_0} + \mathrm{d}\alpha$ and $F_{A_{\text{in}}} = F_{A_0} + \mathrm{d}\beta$. We see that the difference between the two functionals is (half of) the following:

$$\begin{aligned} 2\text{csd}_1(\psi, A) - 2\text{csd}_0(\psi, A) &= \int_Y -(\beta - \alpha) \wedge (2F_{A_0} + \mathrm{d}\alpha + \mathrm{d}\beta) + \beta \wedge (2F_{A_0} + \mathrm{d}\beta) \\ &= \int_Y \alpha \wedge 2F_{A_0} - \beta \wedge \mathrm{d}\alpha + \alpha \wedge \mathrm{d}\alpha + \alpha \wedge \mathrm{d}\beta. \end{aligned}$$

But $\mathrm{d}(\alpha \wedge \beta) = \mathrm{d}\alpha \wedge \beta - \alpha \wedge \mathrm{d}\beta = \beta \wedge \mathrm{d}\alpha - \alpha \wedge \mathrm{d}\beta$ by the Leibniz rule; using Stokes' theorem, we conclude that

$$\text{csd}_1(\psi, A) - \text{csd}_0(\psi, A) = \frac{1}{2}\int_Y \alpha \wedge (2F_{A_0} + \mathrm{d}\alpha) = \frac{1}{2}\int_Y (A_1 - A_0) \wedge (F_{A_1} + F_{A_0}),$$

which clearly has no A_{in}-dependence. □

[3] The name derives from each term in the integrand: the first term is essentially the Lagrangian density for a massless Dirac spinor in an electromagnetic field, and the second term is familiar from Chern-Simons theory.

Theorem 3.36 (Kronheimer, Mrowka)**.** *The pairs in $\mathscr{C}(W)$ satisfying the Seiberg-Witten equations are precisely the critical points of the Chern-Simons-Dirac functional.*

Proof. Let $(\psi, A) \in \mathscr{C}(W)$ be arbitrary. A tangent vector in $\mathscr{C}(W)$ is described by a 1-form $i\alpha \in i\Omega^1(Y)$ and a section $\phi \in \Gamma(W)$. Given any $\varepsilon > 0$, we have the following:

$$\begin{aligned}\mathrm{csd}(\psi + \varepsilon\phi, A + i\varepsilon\alpha) &= \mathrm{csd}(\psi, A) + \frac{\varepsilon}{2}\int_Y \Big((\mathrm{Re}\langle\phi, D_A\psi\rangle + \mathrm{Re}\langle\psi, D_A\phi + \rho(i\alpha)\cdot\psi\rangle)\mathrm{vol}_g \\ &\qquad - i\alpha \wedge (F_A + F_{A_0}) - (A - A_0)\wedge i\mathrm{d}\alpha\Big) + o(\varepsilon^2) \\ &= \mathrm{csd}(\psi, A) + \frac{\varepsilon}{2}\int_Y \Big((2\,\mathrm{Re}\langle\phi, D_A\psi\rangle + \mathrm{Re}\langle\psi, i\rho(\alpha)\cdot\psi\rangle)\mathrm{vol}_g \\ &\qquad + \langle i\alpha, *(F_A + F_{A_0})\rangle + \langle A - A_0, *i\mathrm{d}\alpha\rangle\Big) + o(\varepsilon^2).\end{aligned}$$

(In the second line, we used the self-adjointness of D_A as well as the fact that $\beta\wedge\gamma = -\langle\beta, *\gamma\rangle$ whenever β is a 1-form and γ is a 2-form.) Differentiating with respect to ε at $\varepsilon = 0$, we find that the gradient of csd in the $(\phi, i\alpha)$-direction is the following:

$$\begin{aligned}\nabla\mathrm{csd}_{(\psi,A)}(\phi, i\alpha) &= \int_Y \mathrm{Re}\langle\phi, D_A\psi\rangle\mathrm{vol}_g + \frac{1}{2}\int_Y \langle\psi, i\rho(\alpha)\cdot\psi\rangle\mathrm{vol}_g \\ &\qquad + \frac{1}{2}\int_Y \langle *(F_A + F_{A_0}) + *\mathrm{d}(A - A_0), i\alpha\rangle\,\mathrm{vol}_g \\ &= \langle D_A\psi, \phi\rangle_{L^2} + \langle q(\psi), i\alpha\rangle_{L^2} + \langle *F_A, i\alpha\rangle_{L^2} \\ &= \Big((D_A\psi, q(\psi) + *F_A), (\phi, i\alpha)\Big)_{L^2}.\end{aligned}$$

(In the first line, we used the fact that $*\mathrm{d}$ is self-adjoint; in the last line, we denoted the L^2-inner product on $T_{(\phi,A)}\mathscr{C}(W)$ by round brackets with a subscript L^2. We have also used Proposition 3.15 to introduce the quadratic map.) The pair (ψ, A) is extremal if and only if this inner product is always zero, which is equivalent to satisfying the Seiberg-Witten equations. □

It turns out that the Chern-Simons-Dirac functional is not gauge invariant in general.

Proposition 3.37. *The Chern-Simons-Dirac functional is well-defined on gauge orbits modulo $8\pi^2\mathbb{Z}$.*

Proof. Let $g : Y \to \mathrm{U}(1)$ be a unitary gauge transformation. We assign to g the de Rham cohomology class $[g]$ associated to $\frac{1}{2\pi i}g^{-1}\mathrm{d}g$. Note that this is a class valued in $\mathbb{Z}$, as it may be locally written as $\frac{1}{2\pi i}\mathrm{d}(\ln(g))$. (As a side note, the gauge transformations exhaust $H^1(Y;\mathbb{Z})$; this is because the connected components of $\mathscr{G}$ are the homotopy classes of maps from Y to $\mathrm{U}(1) = B\mathbb{Z}$, which is isomorphic as a group to H^1.) Now, let $(\psi, A) \in \mathscr{C}(W)$ be arbitrary; the following can be verified by a short computation:

$$\mathrm{csd}(g\cdot(\psi, A)) - \mathrm{csd}(\psi, A) = \int_Y (g^{-1}\mathrm{d}g)\wedge(F_A + F_{A_0}). \tag{3.33}$$

On the other hand, both F_A and F_{A_0} define the integral cohomology class $-2\pi i c_1(W)$ by Chern-Weil theory. We can therefore rewrite the functional as

$$\mathrm{csd}(g\cdot(\psi, A)) - \mathrm{csd}(\psi, A) = 8\pi^2\langle[g]\smile c_1(W), [Y]\rangle, \tag{3.34}$$

which is an integer multiple of $8\pi^2$. □

Remark. The constant $8\pi^2$ changes depending on the definition of the Chern-Simons-Dirac functional. In [KM07], for instance, the constant changes to $2\pi^2$ because the Chern-Simons term in their functional differs from ours by a factor of 4.

Because of this ambiguity, the Chern-Simons functional is often taken to be valued in $\mathbb{R}/8\pi^2\mathbb{Z}$. Note that the unmodified functional is still gauge invariant if one restricts to the identity component of $\mathscr{G}$, since the $8\pi^2$-shift comes from different homotopy classes of gauge transformations. This is consistent with the fact that the critical points are gauge invariant. Alternatively, if one restricts to spaces with trivial H^1 (such as homology spheres), csd is also gauge invariant.

3.3.3 Manifold structure

In general, it is difficult to characterise the moduli space of Seiberg-Witten monopoles; most of our conclusive results will focus on the case where Y is a Seifert fibred space. However, it is possible to characterise the local topology of the irreducible component via linearisations.

Since Y is compact with a Riemannian metric g, the space $\Gamma(W) \times i\Omega^1(Y)$ comes equipped with the following natural Hilbert space structure (used in the proof of Theorem 3.36):

$$\langle(\psi, i\alpha), (\phi, i\beta)\rangle = \int_Y (\mathrm{Re}\langle\psi, \phi\rangle_W + \langle\alpha, \beta\rangle_{\Lambda^1})\mathrm{vol}_Y. \tag{3.35}$$

It follows that $\Gamma(W)\times\mathscr{A}(W)$ has the structure of a Hilbert manifold (being affinely modelled over a Hilbert space). We can also realise the quotient $\Gamma(W) \times \mathscr{A}(W)/\mathscr{G}$ as a manifold; in order to do this, we need the following elementary concept:

Definition 3.38. In what follows, we do not assume any space is finite-dimensional. Let X be a manifold, and let G be a Lie group acting smoothly on X. Given some $x \in X$, denote by G_x the isotropy subgroup of G at X. A *slice of the G-action through x* is a smooth submanifold $S \subseteq X$ containing x, which satisfies the following:

- S is invariant under G_x;
- The orbit $G \cdot S$ is open in X;
- The map $(G \times S)/G_x \to G \cdot S$ given by taking $(g, s) \mapsto gs$ is a diffeomorphism (here the G_x-action on $G \times S$ is defined by taking $(g, s) \mapsto (gh^{-1}, hs)$).

We are particularly interested in the case where the G-action is free. In this case, all isotropy groups are trivial; this means that the first condition is always fulfilled, and the third condition is equivalent to requiring that S intersects the orbits of G transversely. It follows that a slice of a free G-action is essentially a smooth local section of the G-bundle $X \to X/G$. From this reasoning, a result follows immediately:

Proposition 3.39. *If a manifold with a free action by a Lie group G has a slice through every point, then its G-orbit space is also a manifold.*

Since the subspace $\mathscr{C}^*(W) := (\Gamma(W)\backslash\{0\}) \times \mathscr{A}(W)$ is acted upon freely by the gauge group $\mathscr{G}$, we can realise the $\mathscr{G}$-orbit space $\mathscr{B}^*(W)$ as a manifold simply by constructing a slice of the $\mathscr{G}$-action through each point.

Proposition 3.40. *Given $\gamma = (\psi, A) \in \mathscr{C}^*(W)$, define the space $S_\gamma \subseteq \Gamma(W) \times i\Omega^1(Y)$ as follows:*

$$S_\gamma = \{(\phi, a) \in \Gamma(W) \times i\Omega^1(Y) \, : \, -\mathrm{d}^* a + \mathrm{Re}\langle i\psi, \phi\rangle_W = 0\}. \tag{3.36}$$

Then the subspace $\gamma + S_\gamma$ is a slice for the $\mathscr{G}$-action on $\mathscr{C}^*(W)$.

The proof for this proposition can be found in [KM07][4]. The essence of the proof is that the infinitesimal gauge group action defines a closed subspace of tangent vectors at each point $(\psi, A) \in \Gamma(W) \times \mathscr{A}(W)$ given by elements of the form $(\xi\psi, -\mathrm{d}\xi)$ for $\xi : Y \to i\mathbb{R}$, and its orthogonal complement is given by pairs (ϕ, a) satisfying $\mathrm{d}^* a = \mathrm{Re}\langle i\psi, \phi\rangle$. This orthogonal complement therefore defines a space of directions in the pre-configuration space which are transverse to the group action.

The moduli space of irreducible Seiberg-Witten monopoles is also a smooth manifold. For a proof of this fact, see [KM97]; the idea is essentially to take the function $\nabla\mathrm{csd} : \Gamma(W)^* \times \mathscr{A}(W) \to i\Omega^1(W) \times \Gamma(W)$, show that its linearisation is surjective over 0, and use the implicit function theorem to put a smooth manifold structure on $(\nabla\mathrm{csd})^{-1}(0)$. Once we have this fact, we can describe the tangent space of the solution set explicitly as the kernel of the linearisation. More precisely, we can describe $\nabla\mathrm{csd}$ as the following nonlinear function:

$$(\psi, A) \mapsto (*F_A + q(\psi), D_A\psi). \tag{3.37}$$

The linearisation of this map at (ψ, A) is computed by taking some $(\phi, a) \in \Gamma(W) \times i\Omega^1(Y)$, computing $\frac{1}{\varepsilon}(\nabla\mathrm{csd}(\psi + \varepsilon\phi, A + \varepsilon a) - \nabla\mathrm{csd}(\psi, A))$ for $\varepsilon \neq 0$, and taking the limit as $\varepsilon \to 0$; this yields the map $D_{(\psi,A)}\nabla\mathrm{csd} : \Gamma(W) \times i\Omega^1(Y) \to \Gamma(W) \times i\Omega^1(Y)$ defined as follows:

$$D_{(\psi,A)}\nabla\mathrm{csd}(\phi, a) = (D_A\phi + \rho(a) \cdot \psi, *\mathrm{d}a + \rho^{-1}(\psi \otimes \phi^* + \phi \otimes \psi^*)_0). \tag{3.38}$$

The kernel of this map defines the tangent space of the solution set at (ψ, A). We have therefore proved the following:

Proposition 3.41. *The tangent space of $\mathscr{M}^*_{\mathrm{sw}}$ at (ψ, A) is given by the linear subspace of $\Gamma(W) \times i\Omega^1(Y)$ satisfying the following system of equations:*

$$*\mathrm{d}a + \rho^{-1}(\psi \otimes \phi^* + \phi \otimes \psi^*)_0 = 0, \tag{3.39}$$
$$D_A\phi + \rho(a) \cdot \psi, \tag{3.40}$$
$$\mathrm{d}^* a = \mathrm{Re}\langle i\psi, \phi\rangle. \tag{3.41}$$

The reducible locus does not necessarily have a manifold structure. However, we may still understand its local structure in terms of the linearisations.

Definition 3.42. The *tangent space* at $(0, A) \in \mathscr{M}_{\mathrm{sw}}$ is defined to be the linear subspace of $\Gamma(W) \times i\Omega^1(Y)$ consisting of pairs (ϕ, a) for which ϕ is harmonic and a is both closed and co-closed.

[4]The statement of the proposition in this reference is much more general than our statement here. First, their statement gives extra conditions on S corresponding to boundary conditions on ∂Y. Second, they allow ψ and A to have Sobolev regularity L^2_j, where we have assumed they must be smooth. Third, while we have simply taken the subspace of $\mathscr{C}(W)$ on which $\mathscr{G}$ acts freely, they analyse the local structure of the blow-up of the pre-configuration space along the fixed points of $\mathscr{G}$.

By Hodge theory, it follows immediately that whenever (ϕ, a) in the tangent space, a is also harmonic. Thus, we may identify a with its corresponding class in $H^1(Y;\mathbb{R})$ and obtain the following result:

Proposition 3.43. *The tangent space to a reducible monopole* $(0, A)$ *is isomorphic to* $H^1(Y;\mathbb{R}) \oplus \ker(D_A)$.

Remark. In the equivariant case, the de Rham theorem for orbifolds allows us to characterise the Λ-equivariant tangent space to an equivariant reducible monopole $(0, A)$ as $H^1(|Y/\Lambda|;\mathbb{R}) \oplus \ker(D_A)$.

3.3.4 Four dimensions

In the four-dimensional theory, the underlying Riemannian manifold is denoted by (M, g_M) and the Spin^c-structure is denoted by $P_M \to M$. We denote the spinor bundle over M by $S \to M$, and we denote the Clifford module structure by $\rho_M : T^*M \to \mathfrak{su}(S)$. Unlike the 3-manifold case, the spinor bundle admits a nontrivial splitting according to the eigenspaces of $\rho_M(\mathrm{vol}_M)$; we denote the summand bundles of this splitting by $S^\pm$ and call them the *positive and negative spinor bundles* respectively, so that $S = S^+ \oplus S^-$. As in the unparametrised case, multiplication by odd-grade elements of $\mathrm{Cl}^c(T^*M)$ (such as 1-forms) exchanges the two summands. In particular, given a spinorial connection $A \in \mathscr{A}(S)$, the Dirac operator $D_A : \Gamma(S) \to \Gamma(S)$ takes sections of S^+ to sections of S^- and vice versa; we therefore define the *chiral Dirac operators* to be the restricted maps $D_A^\pm : \Gamma(S^\pm) \to \Gamma(S^\mp)$. (Note that the connection A respects the splitting, i.e., positive spinors are mapped to positive spinors.)

With this data defined, we define the *pre-configuration space* $\mathscr{C}(S)$ to be the collection of pairs of *positive* spinors and spinorial connections $(\phi, A) \in \Gamma(S^+) \times \mathscr{A}(S)$. Once again, an element of the pre-configuration space is irreducible if ϕ is not identically zero, and reducible otherwise; the gauge group still acts freely on the irreducible part of the pre-configuration space. This leads to the Seiberg-Witten equations on 4-manifolds:

Definition 3.44. A connection $A \in \mathscr{A}(S)$ and a positive spinor $\phi \in \Gamma(S^+)$ are said to satisfy the *4-dimensional Seiberg-Witten equations* if the following two equations hold:

$$\tfrac{1}{2}F_A^+ + q'(\psi) = 0; \tag{3.42}$$

$$D_A^+\phi = 0. \tag{3.43}$$

Here, F_A^+ denotes the self-dual part of F_A (i.e., $F_A^+ = \frac{1}{2}(F_A + *F_A)$), and $q' : \Gamma(S^+) \to \Omega_+^2(M)$ denotes the quadratic map from Definition 3.14.

Unlike the three-dimensional case, the four-dimensional Seiberg-Witten equations do not admit a natural interpretation in terms of an action functional. In the following section, however, we will see that the equations arise as the *flow lines* of the Chern-Simons-Dirac functional.

3.4 Seiberg-Witten Theory on Cylinders

Though there are some key differences, the 3- and 4-dimensional equations are similar in many ways. These similarities can be made precise if M is taken to be a cylinder over the

3-manifold Y. The results in this section will not be used to characterise the moduli space in Chapter 4, but we include them to give hints toward applications.

As before, we take Y to be a Riemannian 3-manifold and M is a 4-manifold. However, we now restrict to the case for which $M = \mathbb{R} \times Y$. There is a natural projection $p : M \to Y$, and we also choose a series of right inverses $i_t : Y \to M$ of p for each $t \in \mathbb{R}$ by taking $i_t(x) = (t, x)$. Note that the projection p induces a splitting of TM via pullback as $\mathbb{R}\frac{\mathrm{d}}{\mathrm{d}t} \oplus p^*TY$, where $\frac{\mathrm{d}}{\mathrm{d}t}$ is the vector field induced by the flow lines $t \mapsto (t, x)$ for each $x \in Y$. Similarly, there is a splitting of T^*M into $\mathbb{R}\,\mathrm{d}t \oplus p^*T^*Y$, where $\mathrm{d}t \in \Omega^1(\mathbb{R})$ is the dual of $\frac{\mathrm{d}}{\mathrm{d}t}$ (according to the usual Riemannian metric on $\mathbb{R}$). There is a natural Riemannian metric on M given by $p^*g + \mathrm{d}t^2$, which induces the volume form $\mathrm{d}t \wedge p^*\mathrm{vol}_Y$ on M.

3.4.1 Fibre bundles

We now turn attention to general fibre bundles over Y and M. A bundle over M corresponds to a bundle over Y by restricting to $i_t(Y)$, while a bundle over Y corresponds to a bundle over M by pullback along p. Given a bundle $\pi : E \to Y$, we denote by $\widehat{\pi} : \widehat{E} \to M$ the corresponding pullback by p. It turns out that this induces a correspondence between sections of E and sections of $\widehat{E}$:

Proposition 3.45. *Let $\pi : E \to Y$ and $\widehat{\pi} : \widehat{E} \to M$ be corresponding bundles, and denote by $p_* : \widehat{E} \to E$ and $(i_t)_* : E \to \widehat{E}$ the induced bundle maps under pullback by the maps $p : M \to Y$ and $i_t : Y \to M$. A section $\widehat{\sigma} \in \Gamma(\widehat{E})$ has a corresponding path of sections $\sigma_\bullet : \mathbb{R} \to \Gamma(E)$ given by the following formula:*

$$\sigma_t(x) = p_*(\widehat{\sigma}(t, x)); \tag{3.44}$$
$$\widehat{\sigma}(t, x) = (i_t)_*(\sigma_t(x)). \tag{3.45}$$

Proof. If we write $\widehat{\sigma} : M \to \widehat{E}$ in terms of the usual explicit realisation of the pullback bundle $\widehat{E} \subseteq M \times E$, we can write $\widehat{\sigma}(t, x) = ((t, x), f(t, x))$ where $f(t, x) \in E_x$; we then have that $p_*\widehat{\sigma}(t, x) = f(t, x)$. Conversely, if we write each $\sigma_t : Y \to E$ explicitly, we see that $\sigma_t(x) = (x, g_t(x))$ where $g_t(x) \in \widehat{E}_{(t,x)} \cong E_x$; we then have that $(i_t)_*\sigma_t(x) = g_t(x)$. □

For the purposes of Seiberg-Witten theory, the primary bundles of interest are the spinor bundles $S \to M$ and $W \to Y$. The above proposition allows us to construct a correspondence between the two types of spinor bundle:

Proposition 3.46. *The maps $p : M \to Y$ and $i_t : Y \to M$ induce a one-to-one correspondence between spinor bundles over Y and M.*

Proof. Let $S \to M$ be a spinor bundle. Since M is four-dimensional, we can split $S \cong S^+ \oplus S^-$ into chiral components. Since we have an inclusion $i_t : Y \to M$ for each t and an inclusion $p^* : T^*Y \to T^*M$, we can restrict S^+ to a bundle over $i_0(Y) \cong Y$. We obtain a Clifford module structure on this bundle by defining $\rho : T^*Y \to \mathfrak{su}(S^+)$ as follows (cf. Proposition 3.9):

$$\rho(\alpha) = -\rho_M(\mathrm{d}t) \circ \rho_M(p^*\alpha). \tag{3.46}$$

Conversely, let $W \to Y$ be a spinor bundle with positive chirality. Since we have a projection $p : M \to Y$ and a splitting $T^*M \cong \mathbb{R}\mathrm{d}t \oplus p^*T^*Y$, we can again use the above proposition to make $\widehat{W} \oplus \widehat{W}$ into a spinor bundle over M if we take the following Clifford module structure:

$$\begin{aligned} \rho_M(\mathrm{d}t) &= \begin{pmatrix} 0 & \mathrm{id}_{\widehat{W}} \\ -\mathrm{id}_{\widehat{W}} & 0 \end{pmatrix}, \\ \rho_M(p^*\alpha) &= \begin{pmatrix} 0 & \rho(\alpha) \\ \rho(\alpha) & 0 \end{pmatrix}. \end{aligned} \tag{3.47}$$

□

Thus, we can think of spinors on M as paths of spinors on Y. This is the essential idea underlying the correspondence between the 3-manifold and 4-manifold Seiberg-Witten equations, and we will carry out this analysis for each term in the equations.

3.4.2 Differential forms and the quadratic map

Before we discuss the terms in the Seiberg-Witten equations, we need to explicitly define the correspondence between differential forms on M and on Y. Paths of differential forms on Y do give rise to differential forms on M, but since there is an extra direction in which to move on M, the relationship is slightly more complicated than that in Proposition 3.45.

First, consider the k-covectors $\omega \in \Lambda^k(T^*M)$ for which $\frac{\mathrm{d}}{\mathrm{d}t} \lrcorner\, \omega = 0$. This condition is equivalent to the statement that $\omega = p^*\alpha$ for some $\alpha \in \Lambda^k(T^*Y)$, which can be easily shown in coordinates. On the other hand, it can also be shown that an arbitrary k-form $\omega \in \Omega^k(M)$ can be decomposed into pieces of this form: we can always find a k-form $\alpha \in \Omega^k(M)$ and a $(k-1)$-form $\beta \in \Omega^{k-1}(M)$ for which $\omega = \alpha + \mathrm{d}t \wedge \beta$, and which contract trivially with $\frac{\mathrm{d}}{\mathrm{d}t}$. Together with Proposition 3.45, we get the following correspondence:

Proposition 3.47. *Every pair of paths $\alpha : \mathbb{R} \to \Omega^k(Y)$ and $\beta : \mathbb{R} \to \Omega^{k-1}(Y)$ corresponds uniquely to a k-form on M given by $\omega = \widehat{\alpha} + \mathrm{d}t \wedge \widehat{\beta}$.*

The next two propositions simply express the usual operations on differential forms in this new language.

Proposition 3.48. *Let $\alpha : \mathbb{R} \to \Omega^k(Y)$ and $\beta : \mathbb{R} \to \Omega^{k-1}(Y)$ be paths of forms on Y, so that $\omega = \widehat{\alpha} + \mathrm{d}t \wedge \widehat{\beta}$ is an arbitrary k-form on M. Then, we have the following formula for the Hodge star:*

$$*_M\omega = (-1)^k \mathrm{d}t \wedge \widehat{*_Y\alpha} + \widehat{*_Y\beta}. \tag{3.48}$$

Proof. It is easy to verify the formula in coordinates on simple k-forms in $\Omega^k(M)$, which allows us to extend by linearity to the whole space. □

Proposition 3.49. *As above, let $\omega = \widehat{\alpha} + \mathrm{d}t \wedge \widehat{\beta}$ be an arbitrary k-form on M. Then, we have the following formula for the exterior derivative on M:*

$$\mathrm{d}_M\omega = \widehat{\mathrm{d}_Y\alpha} + \mathrm{d}t \wedge (\dot{\alpha} - \mathrm{d}_Y\beta)\widehat{\ }. \tag{3.49}$$

Proof. By working in a coordinate system x, expanding α and β into linear combinations of $\mathrm{d}x^I$ for multi-indices I, and using the formula $\mathrm{d}_M(f\mathrm{d}x^I) = \mathrm{d}_Y f \wedge \mathrm{d}x^I + \dot{f}\mathrm{d}t \wedge \mathrm{d}x^I$, the result follows. □

Recall from Lemmas 3.12 and 3.13 that $\mathbb{R}^3$ is linearly isomorphic to $\mathfrak{su}(\Delta^+)$ and $\Lambda^2_\pm\mathbb{R}^4$ is linearly isomorphic to $\mathbb{R}^4$. These suggest the following:

Proposition 3.50. *There is a one-to-one correspondence between self-dual 2-forms on M and paths of 1-forms on Y.*

Proof. By the above formula for the Hodge star, a self-dual 2-form $\omega = \widehat{\alpha} + \mathrm{d}t \wedge \widehat{\beta}$ satisfies $\mathrm{d}t\wedge\widehat{*_Y\alpha}+\widehat{*_Y\beta} = \widehat{\alpha}+\mathrm{d}t\wedge\widehat{\beta}$. Matching components with and without $\mathrm{d}t$, we see that $*_Y\beta = \alpha$ and $*_Y\alpha = \beta$, meaning α and β are Hodge duals of one another. We therefore take the self-dual 2-form ω to the path of 1-forms given by β. This process is clearly invertible. □

Given a path of 1-forms $\alpha : \mathbb{R} \to \Omega^1(Y)$, we denote by $\mathrm{sd}(\alpha) \in \Omega^2_+(M)$ the corresponding self-dual 2-form on M. The correspondence between quadratic terms on 3-manifolds and 4-manifolds may now be phrased as follows:

Lemma 3.51. *Let $S \to M$ be a spinor bundle over M with Clifford module structure $\rho_M : T^*M \to \mathfrak{su}(S)$, and let $W \to Y$ be the corresponding bundle over Y with Clifford module structure $\rho_Y : T^*Y \to \mathfrak{su}(Y)$. Let $\widehat{\phi} \in \Gamma(S^+)$ be a positive spinor, corresponding to a path of spinors $\phi : \mathbb{R} \to \Gamma(W)$ over Y. Then, we have the following relation:*

$$\rho_M^{-1}(\widehat{\phi} \otimes \widehat{\phi}^*)_0 = \tfrac{1}{2}\mathrm{sd}(\rho_Y^{-1}(\phi_t \otimes \phi_t^*)_0). \tag{3.50}$$

Proof. We can explicitly write the right-hand expression as follows:

$$\mathrm{sd}(\rho_Y^{-1}(\phi_t \otimes \phi_t^*)_0) = \left(*_Y\rho_Y^{-1}(\phi \otimes \phi^*)_0\right)^{\widehat{}} + \mathrm{d}t \wedge \left(\rho_Y^{-1}(\phi \otimes \phi^*)_0\right)^{\widehat{}}. \tag{3.51}$$

Thus, it suffices to apply ρ_M to this expression, and show that it reduces to $2\widehat{\phi} \otimes \widehat{\phi}^*$. To evaluate ρ_M on the first term, we note that $\rho_Y(*_Y\alpha) = \rho_Y(\alpha)$ for every $\alpha \in \Omega^1(Y)$; this is because $\alpha \wedge *\alpha = |\alpha|^2\mathrm{vol}_Y$, which means that $\rho_Y(\alpha)\rho(*\alpha) = -|\alpha|^2\mathrm{id}_W = \rho_Y(\alpha)^2$. It follows from this and Equation 3.47 that

$$\rho_M\left(\left(*_Y\rho_Y^{-1}(\phi \otimes \phi^*)_0\right)^{\widehat{}}\right) = \widehat{\phi} \otimes \widehat{\phi}^*. \tag{3.52}$$

For the second term, we note that $\rho_M(\mathrm{d}t \wedge \widehat{\beta}) = \rho_M(\mathrm{d}t) \circ \rho_M(\widehat{\beta})$ for every $\beta : \mathbb{R} \to \Omega^1(Y)$. The action of $\rho_M(\widehat{\beta})$ on S^+ is the same as the action of $\rho_Y(\beta)$ on W, except that S^+ is mapped to S^- by β. Since $\rho_M(\mathrm{d}t)$ acts as the identity, it follows that the second term also evaluates to $\widehat{\phi} \otimes \widehat{\phi}^*$. □

3.4.3 Spinorial connections

We now turn our attention to connections, curvature, and the Dirac operator on each spinor bundle. For a general vector bundle $E \to Y$ and an arbitrary connection $\nabla \in \mathscr{A}(E)$, we obtain a connection on $\widehat{E} \to M$ as follows: for $v \in \Gamma(TY)$ and $\sigma_\bullet : \mathbb{R} \to \Gamma(E)$, we define

$$\widehat{\nabla}_{\alpha\frac{\mathrm{d}}{\mathrm{d}t}+\widehat{v}}\widehat{\sigma} = \alpha\widehat{\frac{\mathrm{d}\sigma}{\mathrm{d}t}} + \widehat{\nabla_v\sigma}. \tag{3.53}$$

Note that the induced connection $\widehat{\nabla}$ is $\mathbb{R}$-invariant, i.e., $\widehat{\nabla}^{\tau_s} = \widehat{\nabla}$, where $\tau_s : \mathbb{R} \times Y \to \mathbb{R} \times Y$ translates by $s \in \mathbb{R}$ along the factor of $\mathbb{R}$. Conversely, a connection on $\widehat{E}$ induces a connection on E if the connection is $\mathbb{R}$-invariant.

This construction carries over to spinorial connections. Given a spinor bundle $W \to Y$ and a spinorial connection $B \in \mathscr{A}(W)$, we denote by $\widehat{B}$ the induced connection on $\widehat{W} \cong S^+$. Since we take $S \cong \widehat{W} \oplus \widehat{W}$, this gives rise to a spinorial connection $A := \widehat{B} \oplus \widehat{B} \in \mathscr{A}(S)$. All of this information behaves nicely when converted to the language of principal bundles; if W is associated to the principal bundle $P \to Y$ and α is the connection 1-form of B, the connection 1-form of $\widehat{B}$ is simply $p^*\alpha$ and the connection 1-form of A becomes $p^*\alpha \oplus p^*\alpha$.

Since the space of all spinorial connections is an affine space modelled on the space of imaginary 1-forms, we can write every spinorial connection on S in terms of a reference connection A in the form $A + \widehat{\eta} + i\lambda \mathrm{d}t$, where η is a path in $i\Omega^1(Y)$ and $\lambda \in C^\infty(M, \mathbb{R})$. Due to the presence of λ in this decomposition, there is not a one-to-one correspondence between connections on S and paths of connections on W; however, up to gauge equivalence, this correspondence can be restored.

Lemma 3.52. *For any spinorial connection $A \in \mathscr{A}(S)$, there is a gauge transformation $g \in \mathscr{G}_M$ which takes A to a connection of the form $A_0 + \widehat{\eta}$, where $A_0 = \widehat{B}_0 \oplus \widehat{B}_0$ for a connection $B_0 \in \mathscr{A}(W)$ and η is a path in $\Omega^1(Y)$. Moreover, this gauge transformation is unique up to gauge transformations on Y, i.e., two such gauge transformations will differ by a gauge transformation g' satisfying $\frac{\mathrm{d}g'}{\mathrm{d}t} = 0$.*

Proof. If we choose some $B_0 \in \mathscr{A}(W)$, we know that we can write $A = A_0 + \widehat{\eta} + i\lambda \mathrm{d}t$ for some path of 1-forms η and some smooth function λ on M. Define a gauge transformation $g : \mathbb{R} \times Y \to \mathrm{U}(1)$ as follows:

$$g(t, x) = \exp\left(-i \int_0^t \lambda(t', x) \mathrm{d}t'\right). \tag{3.54}$$

Note that $g^{-1}\mathrm{d}_M g = g^{-1}\mathrm{d}_Y g + g^{-1}\dot{g}\mathrm{d}t$, if we think of g as a path of gauge transformations of Y. But then $\dot{g} = -i\lambda g$, meaning $g^{-1}\dot{g}\mathrm{d}t = -i\lambda \mathrm{d}t$. Applying the gauge transformation g to A corresponds to shifting A by $g^{-1}\mathrm{d}_M g$; observe that the λ term is cancelled, and the remaining 1-forms can all be considered paths in $i\Omega^1(Y)$. Thus, A^g is of the desired form.

To prove the second claim, suppose $A = A_0 + \widehat{\eta}$ and A^g is also of the desired form. Then $g^{-1}\mathrm{d}_M g$ has no $\mathrm{d}t$ component; however, since this can essentially be identified with the t-derivative of g, the result follows. □

It follows that paths of imaginary-valued 1-forms can be identified with connections on S up to gauge equivalence, and likewise paths of spinors in W can be identified with positive spinors in S up to gauge equivalence. Thus, we find that there is an equivalence between paths in the pre-configuration space of Y and elements of the pre-configuration space of Z, up to gauge equivalence.

We may now relate the curvature 2-forms and the Dirac operators on Y and M; we begin with the Dirac operator. Fix a reference connection $B_0 \in \mathscr{A}(W)$, and let $A \in \mathscr{A}(S)$ be an arbitrary connection of the form $A_0 + \widehat{\eta} + i\lambda \mathrm{d}t$. By definition, $\nabla^{\widehat{B}_0} = \frac{\mathrm{d}}{\mathrm{d}t} + \nabla^{B_0}$. Additionally, we know that ϕ corresponds to a positive spinor when we pass to S, and $\mathrm{d}t$ acts as the identity on positive spinors; it follows that the positive chiral Dirac operator for A is given by the following formula:

$$D_A^+ \widehat{\phi} = \widehat{\dot{\phi}} + \widehat{D_{B_0}\phi} + \widehat{\rho(\eta)\phi} + i\widehat{\lambda\phi}. \tag{3.55}$$

To relate the curvature 2-forms of each connection, we use the language of principal bundles. In this case, we can write the connection form of B_0 as $\omega^Y_{\mathrm{SO}} + C$, where C is an imaginary 1-form and ω^Y_{SO} is the Levi-Civita connection form of Y. In this form, we can write the connection form of A_0 as $\omega^Y_{\mathrm{SO}} \oplus \omega^Y_{\mathrm{SO}} + C \oplus C$; but since $C \oplus C$ is simply $C \cdot \mathrm{id}_S$, we conclude that the induced connection on the determinant line bundle for W and S are the same (i.e., they are related under pullback). Moreover, the curvature 2-form of each connection is also the same under pullback:

$$F_{A_0} = p^* F_{B_0} = \widehat{F_{B_0}}. \tag{3.56}$$

As we showed above, an arbitrary connection on S is gauge equivalent to a connection of the form $A_0 + \widehat{\mu}$ for a path of imaginary 1-forms μ on Y. It follows that the curvature of an arbitrary connection is simply

$$F_A = \widehat{F_{B_0}} + \mathrm{d}_M \widehat{\mu} = \widehat{F_{B_0} + \mathrm{d}_Y \mu} + \mathrm{d}t \wedge \widehat{\dot{\mu}}. \tag{3.57}$$

From here, we can use the correspondence between self-dual 2-forms on M and paths of 1-forms on Y to conclude that F_A^+ corresponds to the path of 1-forms given by the following:

$$F_A^+ \leftrightarrow *_Y(F_{B_0} + \mathrm{d}_Y \mu) + \dot{\eta}. \tag{3.58}$$

3.4.4 The Seiberg-Witten equations in terms of flow lines

With this, we can relate the two equations. Let $\phi \in \Gamma(S^+) \cong \Gamma(\widehat{W})$ be a positive spinor, and let $A \in \mathscr{A}(S)$ be a spinorial connection. By the above decompositions, we can write positive spinors as paths of spinors on Y, i.e., $\phi = \widehat{\psi}$, where $\psi : \mathbb{R} \to \Gamma(W)$ is a path of spinors. If we choose a reference connection $B_0 \in \mathscr{A}(W)$ on Y, we can also write $A = \widehat{B_0} \oplus \widehat{B_0} + \widehat{\mu} + i\lambda \mathrm{d}t$ where $\mu : \mathbb{R} \to i\Omega^1(Y)$ is a path of 1-forms and $\lambda : M \to \mathbb{R}$ is a real-valued function; up to gauge equivalence, we can assume $\lambda = 0$. In this decomposition, we find that the 4-dimensional Seiberg-Witten equations take the following form:

$$\tfrac{1}{2}(*_Y F_{B_0} + *_Y \mathrm{d}_Y \mu + \dot{\mu}) + \tfrac{1}{2}\rho^{-1}(\psi \otimes \psi^*)_0 = 0, \tag{3.59}$$

$$\dot{\psi} + D_{B_0}\psi + \rho(\mu)\psi = 0. \tag{3.60}$$

(We have already identified the self-dual 2-forms with paths of imaginary 1-forms in the first equation, and we have removed all hats in the second equation.) By defining $B := B_0 + \mu$ and isolating $\dot{\mu}$ and $\dot{\psi}$, we can rewrite this as follows:

$$\dot{\mu} = -\left(q(\psi) + *_Y F_B\right), \tag{3.61}$$

$$\dot{\psi} = -D_B \psi. \tag{3.62}$$

On the other hand, it has been shown in the proof of Theorem 3.36 that the right-hand expressions constitute the components of $\nabla \mathrm{csd}$, the gradient of the Chern-Simons-Dirac functional. We now state the correspondence between the two equations explicitly.

Theorem 3.53. *Let Y be a 3-manifold, and let $W \to Y$ be a spinor bundle over Y. Let $M = \mathbb{R} \times Y$ be a corresponding 4-manifold, denote by $\widehat{W}$ the pullback bundle of W by the projection $\mathbb{R} \times Y \to Y$, and let $S := \widehat{W} \oplus \widehat{W} \to M$ be a spinor bundle over M. Then:*

- *The three-dimensional Seiberg-Witten equations on $W \to Y$ are defined by the extrema of the Chern-Simons-Dirac functional.*

- *The four-dimensional Seiberg-Witten equations on $S \to M$ are defined (up to gauge equivalence) by the upwards flow lines of the Chern-Simons-Dirac functional.*

This theorem was first stated by Kronheimer and Mrowka in [KM07] and by Morgan, Szabo, and Taubes in [MST96]. It was later observed by Donaldson in [Don96] that, in the light of this theorem, one might be able to construct *Floer homology groups* by using Morse theory on the configuration space $\mathscr{B}(W)$.

Recall that the finite-dimensional form of Morse theory runs as follows. Given a Riemannian manifold $\mathscr{M}$ and a generic function $f : M \to \mathbb{R}$, one can define a flow under the associated vector field ∇f; the lines of this flow will naturally start and end at the fixed points of f, and two fixed points with index differing by 1 will have finitely many flow lines connecting them. One can use these flow lines to form a chain complex called the Morse complex, and the homology of this complex turns out to be isomorphic to the singular homology of the manifold. This strategy for computing homology groups was adapted to the study of 3-manifolds by Floer in [Flo88a], in which the manifold $\mathscr{M}$ was instead taken to be the space of SU(2)-connections, and the Morse function was taken to be the Chern-Simons functional.

The natural generalisation of this idea to the Seiberg-Witten equations was first carried out by Mrowka, Ozsvath and Yu in [MOY96]. In this case, the manifold $\mathscr{M}$ was taken to be the pre-configuration space $\mathscr{C}^*(W)$, and the function was taken to be the Chern-Simons-Dirac functional. The lines of the flow were then identified with monopoles on $\mathbb{R} \times Y$, and the fixed points of the flow were identified with monopoles on Y. Using this framework, the authors were able to generate new invariants for 3-manifolds in the form of *Seiberg-Witten Floer homology.* We will not discuss these invariants in further detail, but we refer the reader to [KM07] for a refined development of the invariants.

Chapter 4

Seiberg-Witten Monopoles on Seifert 3-manifolds

In this chapter, we develop results characterising the moduli space of Λ-equivariant Seiberg-Witten monopoles on Seifert 3-manifolds. Our proof is adapted from the non-equivariant proof in [MOY96]. The essence of the proof is that solutions to the vortex equations on the U(1)-orbit space of a Seifert 3-manifold pull up to Seiberg-Witten monopoles, and conversely all Seiberg-Witten monopoles on a Seifert 3-manifold are U(1)-invariant so that they descend to vortices on the orbit space.

The structure of the chapter is as follows. We begin with a brief discussion of Riemannian structures on a Seifert 3-manifold. We then discuss spinors on Riemann surfaces, and the reduction of spinors on a Seifert 3-manifold to its U(1)-orbit space; this includes an overview of the equivariant generalisation, which involves group extensions of the finite group. With these preliminaries, we extend the correspondence of Mrowka et al. between Seifert 3-manifolds and vortices to the equivariant category. The irreducible component is treated first; we relate the Kähler vortex equations on line bundles over Riemann surfaces to a equivalent system on a spinor bundle, and we then prove that the pullback of these spinors along the Seifert 3-manifold projection induces a diffeomorphism of moduli spaces. We then treat the reducible locus, relating the reducible solutions to flat connections on line bundles. Finally, we present some examples and discuss further directions for research.

Throughout this chapter, Λ will always be a finite group, $\widehat{\Lambda}$ will be a cyclic group extension, Σ will be a Λ-equivariant orbifold Riemann surface, and $\pi : Y \to \Sigma$ will denote a Seifert 3-manifold equipped with a Λ-action extending the action on Σ. The vector field on Y generated by the U(1) action will be denoted by $\frac{\partial}{\partial \varphi}$. A spinor bundle over Y will be denoted by $W \to Y$, whereas a spinor bundle over Σ will be denoted by $S \to \Sigma$; in both cases, the Clifford module structure is denoted by ρ. A spinorial connection on W will be denoted by A until we discuss reducible monopoles; a Hermitian connection on a line bundle over Σ will be denoted by B; and a spinorial connection on a spinor bundle over Σ will be denoted by $\widetilde{B}$. We use the letter ψ to denote spinors over both Y and Σ.

4.1 Metrics and Connections on Seifert 3-manifolds

Let $\pi : Y \to \Sigma$ be a Seifert fibred space, and equip the fibration with a compatible Λ-action. We assume that Σ has been given a constant-curvature Λ-invariant Riemannian metric g_Σ, with volume form vol_Σ and Levi-Civita connection ∇^Σ. As in Chapter 2, we assume Λ *preserves the orientation of* Σ unless otherwise stated. To construct a metric on Y, we use Proposition 1.50 to find a constant-curvature U(1)-connection $i\eta \in i\Omega^1(Y)$ for the U(1)-bundle $Y \to \Sigma$; explicitly, we require that η satisfies the following three conditions:

- η is U(1)-invariant.
- η satisfies $\eta(\frac{\partial}{\partial\varphi}) = 1$.
- $\mathrm{d}\eta \in i\Omega^2(Y)$ is a constant multiple of $\pi^*\mathrm{vol}_\Sigma \in \Omega^2(\Sigma)$. By Proposition 1.50, this constant multiple is given by $2\xi \in \mathbb{R}$, where

$$\xi = -\frac{\pi \deg(Y)}{\mathrm{Vol}(\Sigma)}. \tag{4.1}$$

With this prescription, there is the following natural metric on Y:

$$g_Y = \eta^2 + \pi^* g_\Sigma. \tag{4.2}$$

The metric g_Y is sometimes called the *adiabatic metric;* see [Nic96]. The tangent and cotangent bundle then have the following orthogonal splittings:

$$\begin{aligned} TY &\cong \mathbb{R}\frac{\partial}{\partial\varphi} \oplus \pi^* T\Sigma; \\ T^*Y &\cong \mathbb{R}\eta \oplus \pi^* T^*\Sigma. \end{aligned} \tag{4.3}$$

We denote by $\Pi_{\eta^\perp} : T^*Y \to \pi^*T^*\Sigma$ the natural orthogonal projection; note that the other projection is given simply by $\langle \cdot, \eta\rangle\eta : T^*Y \to \mathbb{R}\eta$.

Already, however, we have reached a point of ambiguity: which SO(3)-connection should be chosen for Y? There are two desirable properties one might expect of a natural choice, namely that the connection preserves the metric and that it preserves the orthogonal splittings above. The natural connection with the first property is the Levi-Civita connection with respect to g_Y, which we denote by $\widehat{\nabla}^{\mathrm{SO}}$, and the natural connection with the second property is the connection[1] $\mathrm{d} \oplus \pi^*\nabla^\Sigma$ which we denote by ∇^{SO}. These two connections are different in general.

For the purposes of relating the Seiberg-Witten monopoles on Y to vortices on Σ, the property of preserving the orthogonal splittings in Equation 4.3 is more important than preserving the metric. The induced spinorial connections for each connection are related as follows:

Proposition 4.1. *Let $W \to Y$ be a spinor bundle with Clifford module structure ρ, and let ∇ and $\widehat{\nabla}$ be connections which are spinorial with respect to ∇^{SO} and $\widehat{\nabla}^{\mathrm{SO}}$ respectively. If the traces of ∇ and $\widehat{\nabla}$ are the same, then they are related as follows for every $v \in TY$:*

$$\widehat{\nabla}_v = \nabla_v + \xi(\tfrac{1}{2}\rho(v) - \eta(v)\rho(\eta)). \tag{4.4}$$

[1] By this, we mean the connection ∇^{SO} satisfying $\nabla^{\mathrm{SO}}\eta = 0$ and $\nabla^{\mathrm{SO}}_v \pi^* X = \nabla^\Sigma_{\pi_*(v)} X$ for every $X \in \mathfrak{X}(\Sigma)$.

Additionally, if D and $\widehat{D}$ denote the Dirac operators for each connection, the two are equal up to a constant shift:

$$\widehat{D} = D - \frac{\xi}{2}. \tag{4.5}$$

Proof. It suffices to verify each formula in local coordinates. Let $\{e^0, e^1, e^2\}$ be an orthonormal coframe on $U \subseteq Y$ compatible with the splitting in Equation 4.3, meaning $e^0 = \eta$ and $e^1 = \pi^* f^1$, $e^2 = \pi^* f^2$ for a given orthonormal coframe $\{f^1, f^2\}$ on $V \subseteq \Sigma$. Denote by $(\omega^{\mathrm{SO}})^i_j \in \Omega^1(V)$ the connection matrix for the Levi-Civita connection on $T^*\Sigma$; since the connection matrix is skew-symmetric, the only independent component is $(\omega^{\mathrm{SO}})^1_2$. We may identify the connection coefficients of ∇^{SO} with those of ∇^Σ, since all others are clearly zero; thus, ∇^{SO} has the following connection matrix:

$$[\nabla^{\mathrm{SO}}] = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & (\omega^{\mathrm{SO}})^1_2 \\ 0 & -(\omega^{\mathrm{SO}})^1_2 & 0 \end{pmatrix}. \tag{4.6}$$

On the other hand, since $\widehat{\nabla}^{\mathrm{SO}}$ is the Levi-Civita connection, we may use Cartan's structural equation

$$\mathrm{d}e^i = \sum_j \omega^i_j \wedge e^j \tag{4.7}$$

to find its connection coefficients ω^i_j. First, observe that $\mathrm{d}\eta = 2\xi\mathrm{vol}_\Sigma = 2\xi e_1 \wedge e_2 = \xi e^1 \wedge e^2 - \xi e^2 \wedge e^1$. Moreover, using the structural equations on Σ, we find that $\mathrm{d}e^1 = (\omega^{\mathrm{SO}})^1_2 \wedge e^2$ and $\mathrm{d}e^2 = -(\omega^{\mathrm{SO}})^1_2 e^1$. Adding various terms to ensure that ω^i_j is skew-symmetric, we find the following connection matrix for $\widehat{\nabla}^{\mathrm{SO}}$:

$$[\nabla^{\mathrm{SO}}] = \begin{pmatrix} 0 & -\xi e^2 & \xi e^1 \\ \xi e^2 & 0 & \xi e^0 + (\omega^{\mathrm{SO}})^1_2 \\ -\xi e^1 & -\xi e^0 - (\omega^{\mathrm{SO}})^1_2 & 0 \end{pmatrix}. \tag{4.8}$$

Thus, the two connections differ by the following $\mathfrak{so}_3$-valued 1-form:

$$\alpha^i_j := [\nabla^{\mathrm{SO}}] - [\widehat{\nabla}^{\mathrm{SO}}] = \begin{pmatrix} 0 & -\xi e^2 & \xi e^1 \\ \xi e^2 & 0 & \xi e^0 \\ -\xi e^1 & -\xi e^0 & 0 \end{pmatrix}. \tag{4.9}$$

We may now use Equation 3.26 to find the difference between two spinorial connections locally. As the two connections reduce to the same connection on $P_{\det}$, the imaginary 1-forms in the equation cancel, and we are left with the following local expression:

$$[\nabla] - [\widehat{\nabla}] = \frac{1}{4}\sum_{i,j} \alpha^i_j \otimes \rho(e^i \wedge e^j) = \tfrac{\xi}{2}(-e^2 \otimes \rho(e^0 \wedge e^1) + e^1 \otimes \rho(e^0 \wedge e^2) + e^0 \otimes \rho(e^1 \wedge e^2)). \tag{4.10}$$

But $e^0 \wedge e^1 = -e^2\mathrm{vol}_Y$, and likewise $e^0 \wedge e^2 = e^1\mathrm{vol}_Y$ and $e^1 \wedge e^2 = -e^0\mathrm{vol}_Y$. It follows that $\rho(e^0 \wedge e^1) = -\rho(e^2)\rho(\mathrm{vol}_Y) = \rho(e^2)$, since we chose the positive chiral representation. We therefore rewrite this expression as follows:

$$[\nabla] - [\widehat{\nabla}] = \tfrac{\xi}{2}(e^0 \otimes \rho(e^0) - e^1 \otimes \rho(e^1) - e^2 \otimes \rho(e^2)) \tag{4.11}$$

$$= \xi \left(\eta \otimes \rho(\eta) - \tfrac{1}{2} \sum_j e^j \otimes \rho(e_j) \right) \tag{4.12}$$

Applying this to an element $v \in TU$, and noting that the result is independent of coordinates, we find that

$$\nabla_v - \widehat{\nabla}_v = \xi \left(\eta(v)\rho(\eta) - \tfrac{1}{2}\rho(v) \right). \tag{4.13}$$

This proves Equation 4.4. Using this formula, the proof of Equation 4.5 is merely a computation:

$$\begin{aligned} D - \widehat{D} &= \sum_j \rho(e_j)(\nabla_j - \widehat{\nabla}_j) \\ &= \xi \sum_j \rho(e_j) \left(\eta(e_j)\rho(\eta) - \tfrac{1}{2}\rho(e_j) \right) \\ &= \xi \sum_j \left(\eta(e_j)\rho(e_j e_0) - \tfrac{1}{2}\rho(e_j)^2 \right) \\ &= \xi \left(-1 + \frac{3}{2} \right) = \frac{\xi}{2}, \end{aligned}$$

where in the last line we have used the Clifford relation. □

This allows us to extend Equation 3.28 to the non-metric-preseving connection ∇^{SO}:

Corollary 4.2. *The anticommutator of D with $\rho(\alpha)$, where $\alpha \in \Omega^1(Y)$, is the following:*

$$\{D, \rho(\alpha)\} = \rho((*\mathrm{d} + \mathrm{d}*)\alpha) - 2\nabla_{\alpha^\sharp} + 2\xi\langle \alpha, \eta \rangle. \tag{4.14}$$

Proof. This follows immediately from Equation 3.28 for the Levi-Civita connection. □

The derivation of this formula is the sole use of the Levi-Civita connection on Y in the proof, and we will henceforth assume the connection on TY is ∇^{SO}.

Under the orthogonal splitting of the tangent bundle, the Dirac operator also admits a splitting. It will be useful to understand how this splitting behaves when the Dirac operator is squared.

Proposition 4.3. *Let D_B be the Dirac operator associated to a spinorial connection B on a spinor bundle $W \to Y$ (with respect to ∇^{SO}). Define D_2 to be the part of the Dirac operator orthogonal to η in the following sense:*

$$D_B = \rho(\eta)\nabla^B_{\frac{\partial}{\partial\varphi}} + D_2. \tag{4.15}$$

Then the square of the Dirac operator is given by the following formula:

$$(D_B)^2 = (\rho(\eta)\nabla^B_{\frac{\mathrm{d}}{\mathrm{d}\varphi}})^2 + (D_2)^2 - \rho(\Pi_{\eta^\perp} * F_B). \tag{4.16}$$

Proof. It suffices to prove the following anticommutator relation:

$$\left\{ \rho(\eta)\nabla^B_{\frac{\mathrm{d}}{\mathrm{d}\varphi}}, D_2 \right\} = -\rho(\Pi_{\eta^\perp} * F_B). \tag{4.17}$$

Thus, let $\{\eta, e^1, e^2\}$ denote an orthonormal coframe on a neighbourhood of Y, lifted from one on Σ, and let $\{\frac{\partial}{\partial\varphi}, e_1, e_2\}$ be the corresponding orthonormal frame. By definition, then, $D_2 = \rho(e^1)\nabla^B_{e_1} + \rho(e^2)\nabla^B_{e_2}$. We now wish to evaluate the expression of the form

$$\left\{ \rho(\eta)\nabla^B_{\frac{\partial}{\partial\varphi}}, \rho(e^1)\nabla^B_{e_1} \right\}. \tag{4.18}$$

Since the connection is spinorial, we know that $\nabla^B \circ \rho(v) = \rho(v) \circ \nabla^B + \rho(\nabla^{\mathrm{SO}} v)$. However, ∇^{SO} respects the splitting in Equation 4.3, meaning $\nabla^{\mathrm{SO}}\eta = 0$; and since e^1 pulls up from Σ, we also know that it is parallel in the $\frac{\partial}{\partial\varphi}$-direction. We can therefore rewrite the term as follows:

$$\left\{\rho(\eta)\nabla^B_{\frac{\partial}{\partial\varphi}}, \rho(e^1)\nabla^B_{e_1}\right\} = \rho(\eta)\rho(e^1)\left[\nabla^B_{\frac{\partial}{\partial\varphi}}, \nabla^B_{e_1}\right]. \tag{4.19}$$

On the other hand, the curvature of ∇^B is defined to be the $\mathrm{End}(W)$-valued tensor F_{∇^B} for which $F_{\nabla^B}(u,v) = [\nabla_u, \nabla_v] - \nabla_{[u,v]}$. Since e^1 pulls up from Σ, its Lie derivative in the fibre direction is necessarily 0; the commutator term is therefore simply $F_{\nabla^B}(\frac{\partial}{\partial\varphi}, e_1)$. However, a short computation shows that the following holds for every $\alpha \in \Omega^1(Y)$:

$$F_{\nabla^B}(e_0,e_1)\rho(\alpha) = \rho(\alpha)F_{\nabla^B}(e_0,e_1) + \rho(F_{\nabla^{\mathrm{SO}}}(e_0,e_1)\alpha). \tag{4.20}$$

Since ∇^{SO} pulls up from Σ, the second term is zero; thus, the curvature tensor commutes with the Clifford action. By Schur's lemma, the curvature must be equal to its trace part, which is simply $F_B(e_0,e_1)$. We therefore find that

$$\left\{\rho(\eta)\nabla^B_{\frac{\partial}{\partial\varphi}}, \rho(e^1)\nabla^B_{e_1}\right\} = \rho(\eta)\rho(e^1)F_B(e_0,e_1). \tag{4.21}$$

The same computation applies to the e_2-component of D_2, which means that

$$\left\{\rho(\eta)\nabla^B_{\frac{\mathrm{d}}{\mathrm{d}\varphi}}, D_2\right\} = \rho(\eta)\rho(e^1)F_B(e_0,e_1) + \rho(\eta)\rho(e^2)F_B(e_0,e_2) = \rho(\eta \wedge \iota_{e_0}F_B). \tag{4.22}$$

By breaking F_B into its components in e^ie^j, one can verify that $*\eta \wedge \iota_{e_0}F_B = \Pi_{\eta^\perp} * F_B$. But then $\rho(*\alpha) = -\rho(\alpha)$ since $\rho(\mathrm{vol}_Y) = 1$, which completes the proof. □

4.2 Spinors on Riemann Surfaces

To relate Seiberg-Witten monopoles to vortices, we need a notion of spinors on Riemann surfaces. This section is devoted to the spin geometry of complex manifolds.

4.2.1 Spin^c-structures on complex manifolds

A central fact in the correspondence between spinors on Y and on Σ is that every complex manifold can be given a canonical $\mathrm{Spin}^c(2n)$-structure. We prove this via the following lemma:

Lemma 4.4. *For any $n \in \mathbb{N}$, let $j : \mathrm{U}(n) \to \mathrm{SO}(2n)$ be the canonical map induced by the $\mathbb{R}$-linear isomorphism $(x_1+iy_1,\ldots,x_n+iy_n) \mapsto (x_1,y_1,\ldots,x_n,y_n)$, and let $\det : \mathrm{U}(n) \to \mathrm{U}(1)$ be the determinant. Then the map $(j,\det) : \mathrm{U}(n) \to \mathrm{SO}(2n) \times \mathrm{U}(1)$ admits a unique lift $\sigma : \mathrm{U}(n) \to \mathrm{Spin}^c(2n)$ making the following diagram commute:*

$$\begin{array}{ccc} & & \mathrm{Spin}^c(2n) \\ & \overset{\sigma}{\nearrow} & \big\downarrow{\scriptstyle (q_{\mathrm{SO}}, q_{\mathrm{U}})} \\ \mathrm{U}(n) & \xrightarrow[(j,\det)]{} & \mathrm{SO}(2n) \times \mathrm{U}(1) \end{array}$$

where $q_{\mathrm{SO}} : \mathrm{Spin}^c(2n) \to \mathrm{SO}(2n)$ is the $\mathrm{U}(1)$ quotient map and $q_{\mathrm{U}} : \mathrm{Spin}^c(2n) \to \mathrm{U}(1)$ is the $\mathrm{Spin}(2n)$ quotient map (as defined in Section 3.1).

Proof. The following is an elaboration on a proof in [LM89]. Since $\mathrm{Spin}^c(2n)$ is a double covering of $\mathrm{SO}(2n) \times \mathrm{U}(1)$, we may use the lifting criterion to construct σ; this states that the map σ exists if and only if the image of $\pi_1(\mathrm{U}(n))$ under $(j, \det)_*$ is contained in the image of $\pi_1(\mathrm{Spin}^c(2n))$ under $(q_{\mathrm{SO}}, q_{\mathrm{U}})_*$, which is equivalent to the requirement that $\mathrm{im}\, j_* \subseteq \mathrm{im}\,(q_{\mathrm{SO}})_*$ and $\mathrm{im}\,\det_* \subseteq \mathrm{im}\,(q_{\mathrm{U}})_*$.

To prove this, we will use the fact that $\mathrm{U}(n) \cong \mathrm{SU}(n) \rtimes \mathrm{U}(1)$, where $\mathrm{U}(1)$ is embedded in $\mathrm{U}(n)$ via the map $k : \mathrm{U}(1) \to \mathrm{U}(n)$ given by $e^{i\theta} \mapsto \mathrm{diag}(e^{i\theta}, 1, \ldots, 1)$, which has the property that $\det \circ k = \mathrm{id}_{\mathrm{U}(1)}$. Blowing up the product $\mathrm{SO}(2n) \times \mathrm{U}(1)$ into its constituent factors, we summarise this setup in the following commutative diagram of Lie groups:

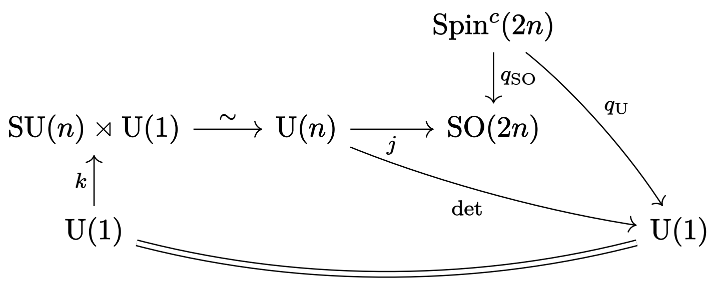

Once we apply π_1 to this diagram, we obtain another commutative diagram of the fundamental groups. To construct this diagram, we note the following:

- By applying the long homotopy exact sequence to the fibration $\mathrm{U}(n-1) \hookrightarrow \mathrm{U}(n) \twoheadrightarrow S^{2n-1}$, as well as the fact that $\pi_1(S^{2n-1}) = \pi_2(S^{2n-1}) = 1$ for all $n \geq 2$, we find that the fundamental group of $\mathrm{U}(n)$ is $\mathbb{Z}$ for all n. Similarly, we find that $\mathrm{SU}(n)$ is simply connected for all n, and $\pi_1(\mathrm{SO}(2n)) = \mathbb{Z}_2$ for $n \neq 1$ while $\pi_1(\mathrm{SO}(2)) = \mathbb{Z}$.

- Let $\gamma : [0,1] \to \mathrm{SO}(2n)$ be a loop corresponding to the generator of $\pi_1(\mathrm{SO}(2n))$ which starts at the identity. Since $\mathrm{Spin}(2n)$ doubly covers $\mathrm{SO}(2n)$, the loop γ can be lifted to a path $\widetilde{\gamma} : [0,1] \to \mathrm{Spin}(2n)$ for which $\widetilde{\gamma}(0) = 1$ and $\widetilde{\gamma}(1) = -1$. Now, consider the loop $\eta : [0,1] \to \mathrm{Spin}^c(2n)$ defined such that $\eta(t) = e^{i\pi t}\widetilde{\gamma}(t)$; it is clear that $q_{\mathrm{SO}} \circ \eta$ generates $\pi_1(\mathrm{SO}(2n))$ and $q_{\mathrm{U}} \circ \eta$ generates $\pi_1(\mathrm{U}(1))$. It follows that the subgroup of $\pi_1(\mathrm{SO}(2n) \times \mathrm{U}(1))$ generated by $(1,1)$ is contained in the image of $(q_{\mathrm{SO}}, q_{\mathrm{U}})_*$.

  In fact, it turns out that this subgroup covers $\mathrm{im}(q_{\mathrm{SO}}, q_{\mathrm{U}})_*$ for $n > 1$; using the long homotopy exact sequence on the $\mathbb{Z}_2$-fibration of $\mathrm{Spin}^c(2n)$ over $\mathrm{SO}(2n) \times \mathrm{U}(1)$ shows that $\pi_1(\mathrm{Spin}^c(2n))$ constitutes an index-2 subgroup of $\pi_1(\mathrm{SO}(2n) \times \mathrm{U}(1))$. When $n = 1$, the image is generated by $(1,1)$ and $(1,-1)$.

- Since semi-direct products of topological groups correspond to topological products, and since $\mathrm{SU}(n)$ is simply connected, we find that k and $\det$ constitute mutually inverse isomorphisms between the fundamental groups of $\mathrm{U}(n)$ and $\mathrm{U}(1)$. This gives a direct computation of $\det_*$, but it also allows us to compute j_*. In particular, by identifying $\mathfrak{u}(1)$ with $\mathrm{SO}(2)$, we find that $j \circ k$ corresponds to the natural inclusion of $\mathrm{SO}(2)$ into $\mathrm{SO}(2n)$; this inclusion induces a surjection from $\pi_1(\mathrm{SO}(2)) \cong \mathbb{Z}$ to $\pi_1(\mathrm{SO}(2n))$, which is either $\mathrm{id}_{\mathbb{Z}}$ if $n = 1$ or $\bmod 2$ if $n > 1$. Thus, we conclude that $\mathrm{im}(j, \det)_*$ is generated by $(1,1)$. But we just proved that this generator is contained in the image of $(q_{\mathrm{SO}}, q_{\mathrm{U}})_*$, which completes the proof. The corresponding commutative diagram of fundamental groups is shown below.

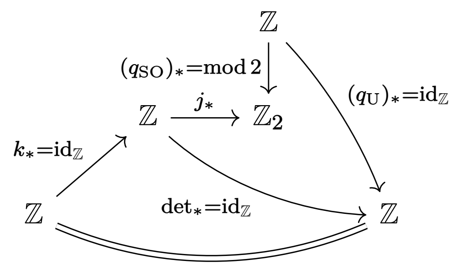

□

Proposition 4.5. *Every Kähler manifold M with real dimension $2n$ has a canonically defined* $\mathrm{Spin}^c(2n)$*-structure, with a spinor bundle isomorphic to $\Lambda^{0,*}TM$ with the following Clifford product: given $v \in TM$ and $\alpha \in \Lambda^{0,*}TM$,*

$$\rho(v) \cdot \alpha = v \wedge \alpha - v^* \lrcorner\, \alpha, \tag{4.23}$$

where $\lrcorner$ denotes contraction and v^* is the covector dual to v.

Proof. Let M be a $2n$-dimensional Kähler manifold. The almost-complex structure J on M naturally induces a reduction of the special orthogonal frame bundle $F_{\mathrm{SO}} \to M$ to the unitary frame bundle $F_{\mathrm{U}} \to M$ (and M needs to be Kähler so that $F_{\mathrm{U}} \subseteq F_{\mathrm{SO}}$). In the above Lemma, we demonstrated the existence of a canonical map $\alpha : \mathrm{U}(n) \to \mathrm{Spin}^c(2n)$, and we therefore define the $\mathrm{Spin}^c(2n)$-structure to be the associated principal bundle $P = F_{\mathrm{U}} \times_\alpha \mathrm{Spin}^c(2n)$, together with the natural projection induced by q_{SO} onto F_{SO}. The spinor bundle is then given by the associated bundle $P \times_\iota \Lambda^{0,*}\mathbb{C}^{2n}$, where $\iota : \mathrm{Spin}^c(2n) \to \mathrm{Aut}(\Lambda^{0,*}\mathbb{C}^{2n})$ is the conjugation map. □

In the case that M is an orbifold Riemann surface Σ, this bundle can be described explicitly in terms of antiholomorphic 1-forms. Recall that K_Σ denotes the canonical bundle.

Definition 4.6. The *canonical spinor bundle over* Σ is the 2-plane bundle $S_0 = \mathbb{C} \oplus K_\Sigma^{-1}$, with Clifford module structure given for any $\xi \in \Lambda^1\Sigma$, any $a \in \mathbb{C}$, and any $\beta \in K_\Sigma^{-1}$ by the following:

$$\rho(\xi) \cdot (a + \beta) = \sqrt{2}(-h(\xi^{1,0}, \beta) + a\xi^{0,1}). \tag{4.24}$$

The trivial bundle $\mathbb{C}$ inherits a fibrewise trivial action from Σ, and K_Σ inherits an action from Σ since Λ acts holomorphically. The Clifford module structure is clearly equivariant, so S_0 defines a Λ-equivariant spinor bundle.

The Levi-Civita connection can be naturally extended to an equivariant spinorial connection on the bundle; if ∇^Σ denotes the Levi-Civita connection on Σ, we define the spinorial connection $A_0 \in \mathscr{A}(S_0)$ as follows:

$$\nabla^{A_0}(a + \beta) = \mathrm{d}a + \nabla^\Sigma \beta. \tag{4.25}$$

The Dirac operator associated to this connection can be calculated to be the following:

$$D_{A_0}(a + \beta) = \sqrt{2}(\overline{\partial} a + \overline{\partial}^* \beta); \tag{4.26}$$

here $\overline{\partial}^* := - * \overline{\partial} *$ is the formal adjoint to $\overline{\partial}$. We can then use Proposition 3.19 to get the following classification:

Proposition 4.7. *The Λ-equivariant spinor bundles over Σ can be identified with equivariant line bundles $L \to \Sigma$. Given such a line bundle, we construct the spinor bundle $L \otimes S_0$ with the Clifford module structure $1 \otimes \rho$, and conversely, a spinor bundle $S \to \Sigma$ gives rise to a line bundle $L = S^+$ by restricting to the $(+1)$-eigenspace of $\rho(i\mathrm{vol}_\Sigma)$.*

The unitary connections on L are in bijective correspondence with the spinorial connections on $L \otimes S_0$: a unitary connection $A \in \mathscr{A}(L)$ induces the spinorial connection $A \otimes A_0$ on L:

$$\nabla^{A\otimes A_0}(u \otimes (a+\beta)) = \nabla^A u \otimes (a+\beta) + u \otimes (\mathrm{d}a + \nabla^{\mathrm{SO}}\beta). \tag{4.27}$$

Conversely, a spinorial connection on S gives rise to a unitary connection on S^+ by restriction. The Dirac operator for the unitary connection A is given by

$$D_A(u \otimes (a+\beta)) = \sqrt{2}\Big(\overline{\partial}_A(u \otimes a) + \overline{\partial}_A^*(u \otimes \beta)\Big), \tag{4.28}$$

where $\overline{\partial}_A : \Gamma(L) \to \Omega^{0,1}(L)$ is the holomorphic structure associated to A, and $\overline{\partial}_A^* := - * \overline{\partial}_A *$ is its formal adjoint.

As in four dimensions, the complex volume element $i\mathrm{vol}_\Sigma$ splits the bundle $L \otimes S_0$ into ± 1 eigenspaces.

Definition 4.8. Elements of the (± 1)-eigenspaces of $i\mathrm{vol}_\Sigma$ on a spinor bundle are called *positive and negative spinors,* and the subbundles they define are denoted by $S^\pm$. Note that, if $S \cong L \otimes S_0$, we have that $S^+ \cong L$ and $S^- \cong L \otimes K_\Sigma^{-1}$.

4.2.2 Pullback to the Seifert 3-manifold

Now that we have parametrised the Spin^c-structures on the underlying orbifold Riemann surface, we can essentially lift the constructions to the total 3-manifold Y by pullback. The pullback bundle $\pi^*(L \otimes S_0)$ becomes a spinor bundle for Y when we extend the Clifford action to the connection 1-form η as follows:

$$\rho(\eta) \cdot u \otimes (a+\beta) = iu \otimes (a-\beta). \tag{4.29}$$

Expressed another way, $\rho(\eta) = -\rho(\mathrm{vol}_\Sigma)$; this comes from the chirality condition $\rho(\mathrm{vol}_Y) = 1$. Additionally, the spinorial connection associated to a connection $A \in \mathscr{A}(L)$ lifts to a $\mathrm{U}(1)$-invariant connection on $\pi^* Y$, and its Dirac operator is easily seen to be the following:

$$D_A(u \otimes (a+\beta)) = \sqrt{2}\Big(\overline{\partial}_A(u \otimes a) + \overline{\partial}_A^*(u \otimes \beta)\Big) + i\frac{\mathrm{d}}{\mathrm{d}\varphi}(u \otimes (a-\beta)). \tag{4.30}$$

Note that, by Proposition 1.55, any line bundle over Y with a connection that has trivial fibrewise holonomy can be pulled up from Σ.

In general, because the odd-dimensional volume element acts homogeneously, spinors over odd-dimensional manifolds tend not to admit splittings according to chirality. However, for a Seifert 3-manifold, we can define chirality in terms of the action of vol_Σ instead:

Definition 4.9. Elements of the (± 1)-eigenspaces of the action of $i\pi^*\mathrm{vol}_\Sigma$ on the Λ-equivariant spinor bundle $W \to Y$ are called *positive and negative spinors,* and the orthogonal subbundles they define are denoted by $W^\pm$. Note that each of these bundles inherit Λ-equivariance from the equivariance of vol_Σ.

This splitting is useful in understanding the quadratic map that appears in the Seiberg-Witten equations:

Proposition 4.10. *Let α be a positive spinor, and let β be a negative spinor. The quadratic map $q : \Gamma(W) \to i\Omega^1(Y)$ decomposes into η-components as follows:*

$$\langle q(\alpha+\beta), \eta\rangle = -\frac{i}{2}(|\alpha|^2 - |\beta|^2), \tag{4.31a}$$

$$\Pi_{\eta^\perp} q(\alpha+\beta) = \rho^{-1}(\alpha\otimes\beta^* + \beta\otimes\alpha^*). \tag{4.31b}$$

Proof. Recall from the definition of q that

$$\rho(q(\alpha+\beta)) = ((\alpha+\beta)\otimes(\alpha+\beta)^*)_0. \tag{4.32}$$

Observe that $\alpha\otimes\alpha^*$ has trace $|\alpha|^2$ (likewise for $\beta\otimes\beta^*$), while $\alpha\otimes\beta^*$ and $\beta\otimes\alpha^*$ are trace-free by the orthogonality of $W^\pm$. Thus, an equivalent formulation of the proposition is the following:

$$-\tfrac{i}{2}(|\alpha|^2-|\beta|^2)\rho(\eta)+\alpha\otimes\beta^*+\beta\otimes\alpha^* = \alpha\otimes\alpha^*+\beta\otimes\beta^*+\alpha\otimes\beta^*+\beta\otimes\alpha^*-\tfrac{1}{2}(|\alpha|^2+|\beta|^2); \tag{4.33}$$

or after cancelling the trace-free terms,

$$-\tfrac{i}{2}(|\alpha|^2-|\beta|^2)\rho(\eta) = \alpha\otimes\alpha^* + \beta\otimes\beta^* - \tfrac{1}{2}(|\alpha|^2+|\beta|^2). \tag{4.34}$$

But this equation can simply be checked by operating on α and β individually, and observing that $\rho(\eta)|_{W^\pm} = \mp i$. Assuming without loss of generality that α and β are both nonzero, and extending by linearity, the equation holds for all elements of $W^+\oplus W^- = W$. □

4.3 Group Actions on Seifert 3-Manifolds

Let Y be a Seifert 3-manifold fibring over Σ, and equip Σ with an isometric action by a finite group Λ. We have already characterised the Λ-equivariant vortices on Σ, and the correspondence in [MOY96] relates vortices on orbifold Riemann surfaces to monopoles on Seifert 3-manifolds. We wish to define a compatible action on Y; the precise notion of compatibility is the following:

Definition 4.11. Let $Y \to \Sigma$ be a Seifert 3-manifold over an orbifold Riemann surface Σ equipped with an isometric action by a finite group Λ. A *lifting of the Λ-action* is a Λ-action on Y with the following two properties:

- The U(1)-action on Y commutes with the orientation-preserving elements of Λ, and anticommutes with the orientation-reversing elements of Λ.
- The Λ-action preserves the fibres of Y: for all $\gamma\in\Lambda$ and all $y\in Y$, we have that $\gamma\cdot\pi(y) = \pi(\gamma\cdot y)$.

The first condition is necessary to account for the non-orientable apparatus we have set up.

Aiming to generalise the correspondence of moduli spaces, a natural question arises: does the group action on Σ always lift to Y? In general, the answer is no:

Proposition 4.12. *If the action of Λ on Σ lifts to Y, then the induced action of Λ on bundles via pullback either preserves or reverses the isomorphism class of Y, depending on orientation.*

Proof. Let $\gamma \in \Lambda$ be arbitrary. If Λ acts compatibly on Y, observe that the following diagram commutes:

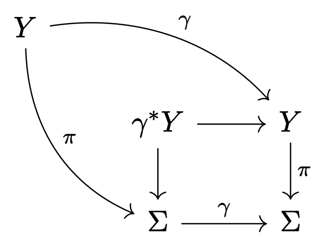

By the universal property, there is a unique induced map from Y to γ^*Y compatible with π; its inverse is clearly given by the corresponding map for γ^{-1} applied to γ^*Y, and the fact that it either commutes or anticommutes with the U(1)-action means it defines an isomorphism with Y or Y^*. □

Remark. If Σ is a smooth Riemann surface, then this condition is always met because the action of Λ induces isomorphisms in cohomology. For a Riemann surface with singularities, however, the Λ-action must also preserve the local invariants of Y around each point.

In particular, the pullback along $\gamma \in \Lambda$ must preserve the local invariants of Y. Even when this condition is imposed, however, there is an extra obstruction given by group cohomology as follows. (For more on group cohomology, refer to [Bro82].)

Proposition 4.13. *Let Λ be a finite group acting on an orbifold Riemann surface Σ preserving the Seifert fibred space $Y \to \Sigma$ under pullback. Given a function $\sigma : \Lambda \to \mathrm{U}(1)$, define the map $\omega : \Lambda \times \Lambda \to \mathrm{U}(1)$ as follows:*

$$\omega(\gamma, \gamma') = \sigma(\gamma)\sigma(\gamma')\sigma(\gamma\gamma')^{-1}. \tag{4.35}$$

Then the group action admits a lifting to Y if and only if there is some choice of σ for which ω defines the trivial group cohomology class in $H^2(\Lambda; \mathrm{U}(1))$. If there is no such choice, then there is an extension $\widehat{\Lambda}$ of Λ by a cyclic group which acts compatibly on Y.

Proof. For the moment, we pick a single point $p \in \Sigma$ and assume it is fixed by Λ. In this case, a choice of lifting of the Λ-action over p is just a map $\overline{\sigma} : \Lambda \to \mathrm{U}(1)$ satisfying the group action condition $\overline{\sigma}(\gamma)\overline{\sigma}(\gamma') = \overline{\sigma}(\gamma\gamma')$. Thus, let $\sigma : \Lambda \to \mathrm{U}(1)$ be an *arbitrary* map; there is then a map $\omega : \Lambda \times \Lambda \to \mathrm{U}(1)$ satisfying the following:

$$\sigma(\gamma)\sigma(\gamma') = \omega(\gamma, \gamma')\sigma(\gamma\gamma'). \tag{4.36}$$

By using the associativity of U(1), we obtain the following relation for all $\gamma, \gamma', \gamma'' \in \Lambda$:

$$\omega(\gamma, \gamma')\omega(\gamma\gamma', \gamma'') = \omega(\gamma, \gamma'\gamma'')\omega(\gamma', \gamma''). \tag{4.37}$$

This is the condition for ω to define a class in $H^2(\Lambda; \mathrm{U}(1))$. This class is trivial if and only if $\omega = \delta\mu$ for some $\mu : \Lambda \to \mathrm{U}(1)$, where $\delta\mu(\gamma, \gamma') = \mu(\gamma)\mu(\gamma')\mu(\gamma\gamma')^{-1}$; defining $\overline{\sigma} = \sigma \cdot \mu^{-1}$ gives us the desired lifting.

If the class is nontrivial, we define an extension $\widehat{\Lambda}^{\mathrm{U}(1)}$ of Λ by taking the set $\Lambda \times \mathrm{U}(1)$ and equipping it with the following product:

$$(\gamma, e^{i\theta}) \cdot (\gamma', e^{i\theta'}) = (\gamma\gamma', \omega(\gamma, \gamma')e^{i(\theta+\theta')}). \tag{4.38}$$

This group naturally acts on Y by multiplication by $e^{i\theta}\sigma(\gamma)$, and the map $\widehat{\Lambda}^{\mathrm{U}(1)} \to \Lambda$ given by projection onto the first factor is clearly an epimorphism with kernel $\mathrm{U}(1)$. Furthermore, since every group cohomology class for a finite group is torsion, we can take ω to be valued in some finite subgroup of $\mathrm{U}(1)$; these must all be cyclic. This allows us to reduce to a cyclic extension by simply taking the product of Λ with that cyclic subgroup.

Since this argument is independent of p, we can extend to all of Σ by noting that the local behaviour of Y over p is preserved under Λ. □

We shall deviate slightly from our established paradigm for equivariant structures from this point forth. In all that preceded this proposition, there was only ever *one* finite group Λ acting and everything was required to be equivariant with respect to that single group. In the light of this proposition, however, we see that the natural notion of equivariance should be with respect to a finite group Λ, and some *cyclic extension* $\widehat{\Lambda} \to \Lambda$ on Y, whose corresponding cyclic group acts along the fibres of Y. This will be assumed implicitly henceforth.

Note that we get a correspondence of bundles as expected:

Proposition 4.14. *A Λ-equivariant bundle over Σ induces a $\widehat{\Lambda}$-equivariant bundle over Y under pullback, and the pullback of Λ-equivariant sections of a bundle are $\widehat{\Lambda}$-equivariant.*

Proof. Let $E \to \Sigma$ be a Λ-equivariant bundle. Conceptualising the pullback π^*E as a subset of $Y \times E$, and denoting by $q : \widehat{\Lambda} \to \Lambda$ the canonical epimorphism, we define a $\widehat{\Lambda}$-action on π^*E as follows:

$$\widehat{\gamma} \cdot (y, v) = (\widehat{\gamma} \cdot y, q(\widehat{\gamma}) \cdot v). \tag{4.39}$$

The statement regarding sections is easily verified. □

By choosing the constant-curvature metric on the orbifold $\mathrm{U}(1)$-bundle $Y/\widehat{\Lambda} \to \Sigma/\Lambda$ and lifting it back up to Y, we may assume that the connection $i\eta \in i\Omega^1(Y)$, and hence the adiabatic metric 4.2 on Y, are $\widehat{\Lambda}$-invariant. On the other hand, let $W \to Y$ be a spinor bundle; then, in the same way that a Λ-action on Σ does not lift to one on Y, the $\widehat{\Lambda}$-action on Y will generally only lift to W in the form of a group extension $\widehat{\widehat{\Lambda}} \to \widehat{\Lambda}$. For the sake of simplicity, we will assume henceforth that this group extension splits, meaning we can lift the $\widehat{\Lambda}$-action on Y to the Spin^c-structure without extending $\widehat{\Lambda}$. In other words, we will assume that the Spin^c-structure we define on Y is $\widehat{\Lambda}$-equivariant.

4.4 Irreducible Moduli Correspondence

We now state a rough version of the correspondence between the irreducible moduli space of Kähler vortices and the irreducible moduli space of Seiberg-Witten monopoles. (A more precise version is given in Theorem 2.20.)

Theorem 4.15. *Let $\mathscr{M}^*_{\text{vtx}}(\Sigma)$ denote the moduli space of irreducible Λ-equivariant Kähler vortices on some line bundle over Σ, and let $\mathscr{M}^*_{\text{sw}}(Y, \pi^*(L \otimes S_0))$ denote the moduli space of irreducible $\widehat{\Lambda}$-equivariant Seiberg-Witten monopoles. The pullback map π^* induces a diffeomorphism between the two moduli spaces.*

This theorem will be proved in three parts. First, we will verify that the map π^* is well-defined, in that the pullback of a Kähler vortex is indeed a Seiberg-Witten monopole and the map respects the gauge action. Then we will demonstrate that π^* induces a homeomorphism between the two spaces, which will involve a decomposition of the Dirac operator on Y according to the Seifert fibration. Finally, we will prove that the map π^* is a diffeomorphism by analysing the linearisations of the Yang-Mills-Higgs functional and the Chern-Simons-Dirac functional, and showing that their kernels are identified under π^*.

4.4.1 Rewriting the Kähler vortex equations

We recall the Kähler vortex equations here for convenience. Given a Hermitian line bundle $L \to \Sigma$ with a unitary connection $B \in \mathscr{A}(L)$ and sections $\alpha \in \Gamma(L)$, $\beta \in \Gamma(L \otimes K_\Sigma^{-1})$, the triple (B, α, β) is a Kähler vortex if and only if the following is satisfied:

$$2F_A - F_{K_\Sigma} = i(|\alpha|^2 - |\beta|^2)\text{vol}_\Sigma, \tag{2.23a}$$

$$\overline{\partial}_A \alpha = 0 \text{ and } \overline{\partial}_A^* \beta = 0, \tag{2.23b}$$

$$\alpha = 0 \text{ or } \beta = 0. \tag{2.23c}$$

The idea of the proof is to rewrite the Kähler vortex equations in terms of a connection and section on a spinor bundle $S \to \Sigma$. Once this is done, we essentially prove that every solution to the Seiberg-Witten equations on $W \to Y$ lifts from a solution of the rewritten equations on $S \to \Sigma$.

Proposition 4.16. *Let $L \to \Sigma$ be a Hermitian line bundle over a Riemann surface, let $\alpha \in \Gamma(L)$ and $\beta \in \Gamma(L \otimes K_\Sigma^{-1})$ be sections, and let $B \in \mathscr{A}(L)$ be a unitary connection. Let $\psi\Gamma(L \otimes S_0)$ be the spinor with chiral decomposition $\alpha + \beta$, and let $\widetilde{B} \in \mathscr{A}(L \otimes S_0)$ be the spinorial connection corresponding to B (i.e., $\widetilde{B} = B \oplus (B \otimes \omega_g)$, where ω_g is the Levi-Civita connection on K_Σ^{-1}). Then (α, β, B) is a solution to the Kähler vortex equations if and only if $(\psi, \widetilde{B})$ is a solution to the following system of equations:*

$$\rho(F_{\widetilde{B}}) = (\psi \otimes \psi^*)_0; \tag{4.40a}$$

$$D_{\widetilde{B}}\psi = 0; \tag{4.40b}$$

$$\psi \in \Gamma(S^\pm). \tag{4.40c}$$

Here, as in the three-dimensional case, $F_{\widetilde{B}}$ refers to half of the trace of the curvature 2-form of $\nabla^{\widetilde{B}}$.

Proof. Equation 4.40c is equivalent to the vanishing of α or β (depending on whether $\psi \in S^-$ or $\psi \in S^+$), and Equation 4.40b is equivalent to the holomorphicity of α and the coholomorphicity of β by the definition of the Dirac operator on S (Equation 4.28). Thus, we focus attention on the first equation (4.40a).

It is easy to show that the curvature of a direct sum or tensor product of connections is the direct sum or tensor product of the curvatures. Thus, it follows that the *trace* of the curvature of $\widetilde{B}$ is simply $F_B + (F_B + F_{K_\Sigma^*}) = 2F_B - F_{K_\Sigma}$. Thus, it suffices to prove that

$$i(|\alpha|^2 - |\beta|^2)\rho(\mathrm{vol}_\Sigma) = 2(\psi \otimes \psi^*)_0. \tag{4.41}$$

To rewrite the right-hand side, first observe that either α or β vanishes identically, so

$$(\psi \otimes \psi^*)_0 = \alpha \otimes \alpha^* + \beta \otimes \beta^* - \tfrac{1}{2}(|\alpha|^2 + |\beta|^2). \tag{4.42}$$

By computing the matrix representation in a basis of the form $\{s^+, s^-\}$ with $s^\pm \in S^\pm$, one can show that

$$(\psi \otimes \psi^*)_0 = \rho\left(\frac{i}{2}(|\alpha|^2 - |\beta|^2)\mathrm{vol}_\Sigma\right). \tag{4.43}$$

Additionally, ρ takes imaginary 2-forms on Σ of the form $it\mathrm{vol}_\Sigma$ isomorphically to linear endomorphisms of $L \otimes S_0$ which act by scaling $S^\pm$ by $\pm t$. This completes the proof in the case that the Λ-action is trivial.

If Λ acts nontrivially, the equivariance of α and β implies the equivariance of $\psi = \alpha + \beta$ by linearity; moreover, since ω_g is the Levi-Civita connection of an equivariant metric, the equivariance of B implies the equivariance of $\widetilde{B}$. By splitting into the Λ-invariant subbundles $S^\pm$, the converse follows immediately. □

Henceforth, we will also refer to the pair $(\psi, \widetilde{B})$ as a Kähler vortex, and to Equations 4.40a–4.40c as the Kähler vortex equations.

4.4.2 The pullback map

With this, we can define the pullback map taking Kähler vortices to monopoles more precisely.

Definition 4.17. Let (α, β, B) be a Kähler vortex on $L \to \Sigma$, and let $\pi : Y \to \Sigma$ be a Seifert fibred space. Then the *pullback of* (α, β, B) *to* Y is defined to be

$$\pi^*(\alpha, \beta, B) = (\pi^*\psi, \pi^*\widetilde{B}) \in \Gamma(\pi^*(L \otimes S_0)) \times \mathscr{A}(\pi^*(L \otimes S_0)), \tag{4.44}$$

where $\psi = \alpha + \beta \in \Gamma(L \otimes S_0)$ and $\widetilde{B} \in \mathscr{A}(L \otimes S_0)$ are the associated spinor and spinorial connection as in Proposition 4.16.

In the remainder of this section, we will show via a series of lemmas that the pullback map induces a diffeomorphism between the moduli spaces of vortices and monopoles. We begin with the following:

Lemma 4.18. *The map* $\pi^* : \mathscr{M}_v^*(\Sigma, L) \to \mathscr{M}_{sw}^*(Y, \pi^*(L \otimes S_0))$ *is well-defined.*

Proof. We need to check two conditions: that the pullback of an equivariant Kähler vortex is an equivariant monopole, and that equivariantly gauge-equivalent vortices define equivariantly gauge-equivalent monopoles. The latter condition follows simply from the fact that a gauge transformation $g : \Sigma \to \mathrm{U}(1)$ pulls back to a gauge transformation $\pi^*g : Y \to \mathrm{U}(1)$

and commutes with π^* (since U(1) is abelian), so we focus our attention on the former condition for the remainder of the proof.

Let (α, β, B) be a Kähler vortex, and let $(\pi^*\psi, \pi^*\widetilde{B})$ be its pullback. By Proposition 4.16, we know that $(\psi, \widetilde{B})$ satisfy the equations $\rho(F_{\widetilde{B}}) = (\psi \otimes \psi^*)_0$ and $D_{\widetilde{B}}\psi = 0$. We show that each of the Seiberg-Witten equations hold individually:

- Observe that ρ commutes with pullback, and the pullback of the curvature of a connection is the curvature of its pullback. Thus, we have the following:

$$\rho(*F_{\pi^*\widetilde{B}}) = -\pi^*\rho(F_{\widetilde{B}}) = -\pi^*(\psi \otimes \psi^*)_0. \tag{4.45}$$

  In the second equality, we have used the fact that $(\psi, \widetilde{B})$ corresponds to a Kähler vortex. However, it is easy to verify that $\pi^*(\psi \otimes \psi^*)_0 = ((\pi^*\psi) \otimes (\pi^*\psi)^*)_0$; using the fact that ρ is an isomorphism when restricted to 1-forms, we find that

$$*F_{\pi^*\widetilde{B}} = -q(\pi^*\psi). \tag{4.46}$$

- Using Equation 4.28, as well as the fact that the pullback along π is U(1)-invariant, we see that $D_{\pi^*\widetilde{B}}(\pi^*\psi) = D_{\widetilde{B}}\psi = 0$.

Thus, we find that the pullback of the Kähler vortex does indeed satisfy the Seiberg-Witten equations. Equivariance follows from the fact that the pullback of Λ-equivariant sections of line bundles over Σ is $\widehat{\Lambda}$-equivariant over Y (see Proposition 4.14). □

We see that the pullback of a Kähler vortex on a spinor bundle S is always a Seiberg-Witten monopole on the pullback spinor bundle $W = \pi^*S$.

The next part of the proof is to show that all Seiberg-Witten monopoles are pullbacks of Kähler vortices. In order to do this, we must contend with the following fact: even if the bundle $W \to Y$ admits some bundle $S \to \Sigma$ for which $\pi^*S = W$, there may be some other bundle $S' \to \Sigma$ which W also pulls up from. Indeed, if Y is the circle bundle corresponding to the line bundle $N \to \Sigma$, one may recall from Proposition 1.54 that $\pi^*N \to Y$ is a trivial bundle; this means the spinor bundles S and $N \otimes S$ pull up to isomorphic spinor bundles over Y. Thus, any sensible correspondence must incorporate all tensor powers of N.

Lemma 4.19. *Let $W \to Y$ be a $\widehat{\Lambda}$-equivariant spinor bundle.*

- *If $W \cong \pi^*S$ for some Λ-equivariant spinor bundle $S \to \Sigma$, then the map*

$$\pi^* : \bigsqcup_{k \in \mathbb{Z}} \mathscr{M}^*_{\mathrm{vtx}}(N^{\otimes k} \otimes S) \to \mathscr{M}^*_{\mathrm{sw}}(W) \tag{4.47}$$

  is a bijection. (Note that the disjoint union ranges over all spinor bundles which pull up to W.)

- *If $W \not\cong \pi^*S$ for any such S, then $\mathscr{M}^*_{\mathrm{sw}}(W)$ is empty.*

Proof. We begin by showing the following: if (ψ, A) is a solution to the Seiberg-Witten equations on $W \to Y$, then there is some spinor bundle $S \to \Sigma$ for which $\pi^*S = W$. Given such a solution, decompose $\psi = \alpha + \beta$ into its positive and negative parts. Using the decomposition of $(D_A)^2$ (Proposition 4.3), the fact that $D_A\psi = 0$ implies that

$$\langle \psi, D_A^2\psi\rangle = \left\langle \psi, \left(\rho(\eta)\cdot\nabla_{\frac{\partial}{\partial\varphi}}\right)^2 \psi\right\rangle + \left\langle \psi, (D_2)^2\psi\right\rangle - \left\langle \psi, \rho(\Pi_{\eta^\perp} * F_A)\psi\right\rangle = 0. \tag{4.48}$$

Observe that $\rho(\eta)\cdot\nabla_{\frac{\partial}{\partial\varphi}}$ and D_2 are both self-adjoint, meaning the first two terms are nonnegative. The third term is also nonnegative: we have that $*F_A = -q(\psi)$ since (ψ, A) is a Seiberg-Witten monopole, meaning

$$\begin{aligned}
-\langle \psi, \rho(\Pi_{\eta^\perp} * F_A)\psi\rangle &= \langle \psi, \rho(\Pi_{\eta^\perp} q(\psi))\psi\rangle \\
&= \langle \psi, (\alpha\otimes\beta^* + \beta\otimes\alpha^*)\psi\rangle \\
&= |\alpha|^2|\beta|^2.
\end{aligned}$$

It follows that, if $\langle\psi, (D_A)^2\psi\rangle = 0$, the following three equations must hold:

$$\nabla_{\frac{\partial}{\partial\varphi}}\alpha = \nabla_{\frac{\partial}{\partial\varphi}}\beta = 0, \tag{4.49a}$$

$$\Pi_{\eta^\perp} * F_A = 0, \tag{4.49b}$$

$$|\alpha||\beta| = 0. \tag{4.49c}$$

By Equation 4.49a, the fibrewise holonomy acts trivially on α and β; by Equation 4.49c, at least one of α and β is nonzero on some regular fibre of Y (we are working in the irreducible locus). It is straightforward to show that the fibrewise holonomy of any spinorial connection form (with respect to ∇^{SO}) must be in $\mathrm{U}(1)$ by Schur's lemma; thus, the fibrewise holonomy must be trivial over some regular fibre of Y. Additionally, the second equation shows that F_A pulls up from Σ; it follows from the correspondence detailed in Propositions 1.55 and 1.56 that there is some bundle-with-connection $(S, B) \to \Sigma$ for which $\pi^*S = W$ and $\pi^*B = A$. We defer the construction of a Λ-equivariant structure on S until ψ_0 has been defined.

Using Equations 4.49a–4.49c, we can also prove that the pullback map is surjective. Since α and β are both fibrewise constant, it follows that ψ is fibrewise constant; thus, there is a corresponding section $\psi_0 \in \Gamma(S)$ for which $\pi^*\psi_0 = \psi$. Given the Seiberg-Witten monopole (ψ, A), we therefore define a corresponding vortex (ψ_0, B). Once it is verified that this pair satisfies the Kähler vortex equations, the surjectivity of π^* is immediate:

- As per the argument in Lemma 4.18, the fact that (ψ_0, B) satisfies $\rho(*F_B) = -(\psi_0 \otimes \psi_0^*)_0$ on $S \to \Sigma$ essentially follows from the fact that all relevant objects commute with π^*.
- Similar logic shows that $D_B\psi_0 = 0$.
- To prove that ψ_0 is either strictly positive or strictly negative, we use the fact that $|\alpha||\beta| = 0$ and both α and β are smooth. This means that one of the two must be zero on an open set. However, both α and β are in the kernel of an elliptic operator (either the Dolbeault operator or its formal adjoint); by the unique continuation theorem for elliptic operators, one of them must vanish everywhere. Thus, ψ is either positive or negative, and the pullback preserves the chiral splitting, meaning ψ_0 is as well.

We now construct a Λ-equivariant structure on $S \to \Sigma$ as follows: given $\widehat{\gamma} \in \Lambda$ and the canonical projection $q : \widehat{\Lambda} \to \Lambda$, we define $q(\widehat{\gamma}) \cdot \psi_0$ so that the following holds:

$$\pi^*(q(\widehat{\gamma}) \cdot \psi_0) = \widehat{\gamma} \cdot \pi^* \psi_0. \tag{4.50}$$

Since ψ_0 is nonzero on a dense open set, this defines the action everywhere by smoothness. Moreover, the action is well-defined; if $\widehat{\gamma}$ projects to $1 \in \Lambda$, then the fact that $\pi^*\psi_0$ has trivial fibrewise holonomy means it is unaffected by the action of $\widehat{\gamma}$. The action is compatible with the action on $\widehat{\Lambda}$ essentially by definition, so π^* is equivariant. This proves that (ψ, B) is an equivariant Kähler vortex.

We now prove that π^* is injective. First, we will show that Kähler vortices on non-isomorphic spinor bundles over Σ pull up to distinct Seiberg-Witten monopoles on Y; we do this by showing that, if two spinors on Σ pull up to the same spinor on Y, then they must be the same spinor. Let S and S' be two spinor bundles satisfying $\pi^*S \cong \pi^*S'$, and let $\psi \in \Gamma(S)$ and $\psi' \in \Gamma(S')$ be spinors satisfying

$$\pi^*\psi' = \pi^*\psi. \tag{4.51}$$

We assume that both spinors are positive in order to view them as holomorphic sections of line bundles S^+ and $(S')^+$; the negative case is analogous. The strategy is to use the behaviour of ψ and ψ' around their zeros to prove that the two bundles are isomorphic.

Recall from Theorem 1.46 that a line bundle L is classified by its Seifert invariant $(\deg(L), b_1, \ldots, b_n)$. This invariant may be recovered by looking at a global meromorphic section of L and examining the behaviour of its zeros and poles (in particular all that is required is the order of a given zero or pole). Note that $\psi \in \Gamma(S)$ encodes this data for $S^{\pm}$, and it can be recovered from $\pi^*\psi \in \Gamma(\pi^*S)$ as follows: given some $p \in \Sigma$ for which $\psi(p) = 0$, take a smooth local section τ of π around an open neighbourhood $U \ni p$ and compute the order of the zero on $(\pi^*\psi)|_{\tau(U)}$ (considered as a section of $(\pi^*S)^{\pm}$). This does not depend on the section or the open neighbourhood, meaning all of the zeros of ψ and ψ' have the same order. It follows that the Seifert invariants of S^+ and $(S')^+$ match, meaning the two spinor bundles are isomorphic. The fact that $\pi^*(\psi' - \psi) = 0$ immediately implies that $\psi = \psi'$.

Thus, Kähler vortices on non-isomorphic spinor bundles over Σ pull up to distinct monopoles on Y. To finish the proof of injectivity, we need to show that any two Kähler vortices whose pullbacks are gauge equivalent must themselves be gauge equivalent. Thus, let (ψ, B) and (ψ', B') be Kähler vortices, and let $g : Y \to \mathrm{U}(1)$ be a gauge transformation for which

$$\pi^*(\psi', B') = g \cdot \pi^*(\psi, B). \tag{4.52}$$

In particular, this means that $\pi^*\psi' = g \cdot \pi^*\psi$. Applying $\nabla_{\frac{\partial}{\partial\varphi}}$ to this equation, we see that $(\nabla_{\frac{\partial}{\partial\varphi}} g)\psi = 0$; since ψ is nonzero on a dense open set, we conclude that g is fibrewise constant. Thus, there is some $h : \Sigma \to \mathrm{U}(1)$ for which $g = \pi^*h$, meaning

$$\pi^*(\psi', B') = \pi^*(h \cdot (\psi, B)). \tag{4.53}$$

By the above result, this implies that each Kähler vortex comes from the same spinor bundle, and furthermore that $(\psi', B') = h \cdot (\psi, B)$. Since h relates two Λ-equivariant spinors whose nonzero set is open and dense, it must be Λ-equivariant. We conclude that the two Kähler vortices are themselves gauge equivalent. □

Finally, in order to prove that the moduli spaces are diffeomorphic, we must show that π^* identifies not only the elements of $\mathscr{M}_{\text{vtx}}$ and $\mathscr{M}_{\text{sw}}$, but also their respective tangent spaces under linearisation.

Lemma 4.20. *The map π^* is a diffeomorphism, in the sense that the tangent spaces of $\mathscr{M}^*_{\text{sw}}$ and $\mathscr{M}^*_{\text{vtx}}$ are identified by the action of π^*.*

Proof. First, we construct the tangent spaces. As shown in Proposition 3.41, the tangent space at (ψ, A) of the moduli space of Seiberg-Witten monopoles can be characterised as the linear subspace of $\Gamma(W) \times i\Omega^1(Y)$ consisting of pairs (ϕ, a) satisfying the following system of equations:

$$\rho(*\mathrm{d}a) = -(\psi \otimes \phi^* + \phi \otimes \psi^*)_0, \tag{4.54a}$$
$$D_A\phi + \rho(a) \cdot \psi = 0, \tag{4.54b}$$
$$\mathrm{d}^*a = i\mathrm{Re}\langle i\psi, \phi\rangle. \tag{4.54c}$$

The tangent space of the moduli space of $\widehat{\Lambda}$-equivariant Seiberg-Witten monopoles is easily seen to be the subspace for which (ϕ, a) are $\widehat{\Lambda}$-equivariant.

The construction of the tangent spaces for $\mathscr{M}_{\text{vtx}}$ parallels the construction for $\mathscr{M}^*_{\text{sw}}$: we take the kernel of the linearisation of the vortex equations to get the tangent space of the solution set, and we restrict to a linearised slice of the group action to get the tangent space of the quotient. The modified vortex equations can be described as the zero set of the map $\mathscr{V} : \Gamma(S^\pm) \times \mathscr{A}(S) \to \mathrm{End}(S) \times \Gamma(S^\mp)$ defined as follows:

$$\mathscr{V}(\psi, B) = (\rho(F_B) - (\psi \otimes \psi^*)_0, D_B\psi). \tag{4.55}$$

The linearisation $D_{(\psi,B)}\mathscr{V} : \Gamma(S^\pm) \times i\Omega^1(\Sigma) \to \mathrm{End}(S) \times \Gamma(S^\mp)$ is computed by perturbing the basepoint (ψ, B) by some $\varepsilon(\phi, b) \in \Gamma(S^\pm) \times i\Omega^1(\Sigma)$ for $\varepsilon > 0$ and observing the first-order change. This yields the following:

$$D_{(\psi,B)}\mathscr{V}(\phi, b) = (\rho(\mathrm{d}b) - (\psi \otimes \phi^* + \phi \otimes \psi^*)_0, D_B\phi + \rho(b) \cdot \psi). \tag{4.56}$$

The kernel of this linearisation describes the tangent space at (ψ, B) of the space of vortices, when taken as a subset of $\Gamma(S) \times \mathscr{A}(S)$. Next, according to the inner product

$$\langle(\phi, b), (\phi', b')\rangle_{L^2} = \int_\Sigma (\mathrm{Re}\langle\phi, \phi'\rangle_S + \langle b, b'\rangle_{\Lambda^1}) \,\mathrm{vol}_\Sigma \tag{4.57}$$

on $\Gamma(S) \times i\Omega^1(\Sigma)$, the linearisation of the gauge group action given by

$$\widehat{\xi}_{(\psi,B)} = (\xi \cdot \psi, -\mathrm{d}\xi) \text{ for } \xi : \Sigma \to i\mathbb{R} \tag{4.58}$$

defines a closed subspace of $\Gamma(S) \times i\Omega^1(\Sigma)$ whose orthogonal complement defines a slice of the gauge group action:

$$S_{(\psi,B)} = \{(\phi, b) \in \Gamma(S) \times i\Omega^1(\Sigma) \,:\, -\mathrm{d}^*b + i\mathrm{Re}\langle i\psi, \phi\rangle = 0\}. \tag{4.59}$$

As such, the tangent space at (ψ, B) of the moduli space of Λ-equivariant vortices can be characterised as the linear subspace of $\Gamma(S) \times i\Omega^1(\Sigma)$ consisting of Λ-equivariant pairs (ϕ, b) satisfying the following system of equations:

$$\rho(\mathrm{d}b) = (\psi \otimes \phi^* + \phi \otimes \psi^*)_0, \tag{4.60a}$$
$$D_B\phi + \rho(b) \cdot \psi = 0, \tag{4.60b}$$
$$\mathrm{d}^*b = i\mathrm{Re}\langle i\psi, \phi\rangle. \tag{4.60c}$$

From here, we may begin the proof that π^* induces identifications between the tangent spaces. First, we will show that π^* induces an injective map from the tangent space of $\mathscr{M}^*_{\mathrm{vtx}}$ at (ψ, B) to the tangent space of $\mathscr{M}^*_{\mathrm{sw}}$ at $(\pi^*\psi, \pi^*B)$. Let (ϕ, b) be a tangent vector to $\mathscr{M}^*_{\mathrm{vtx}}$ at (ψ, B); this means it satisfies Equations 4.60a–4.60c. The Clifford module structure, the Dirac operator, and the Hermitian product on S commute with π^* by definition. Since π is an isometry, π^* also commutes with the Hodge star, the exterior derivative, and the codifferential. Moreover, it is fairly easy to show the following:

$$(\pi^*\psi \otimes (\pi^*\phi)^*)_0 = \pi^*(\psi \otimes \phi^*)_0. \tag{4.61}$$

Lastly, our conventions ensure that $\rho(*\alpha) = -\rho(\alpha)$ for all 2-forms α. It follows immediately that, if (ϕ, b) satisfies Equations 4.60a–4.60c, then $(\pi^*\phi, \pi^*b)$ satisfy Equations 4.54a–4.54c. Injectivity is immediate from the fact that π^* is injective and Λ-equivariant.

We now begin the difficult part of the proof, namely that π^* induces a surjection onto the tangent spaces of $\mathscr{M}^*_{\mathrm{sw}}$. Let $(\phi, a) \in \Gamma(W) \times i\Omega^1(Y)$ define a tangent vector to $\mathscr{M}^*_{\mathrm{sw}}$ at the solution (ψ, A); we need to show three key facts:

- If $\psi \in \Gamma(S^\pm)$, then $\phi \in \Gamma(S^\pm)$ as well.
- The spinor ϕ lifts from Σ, i.e., $\nabla_{\frac{\partial}{\partial\varphi}} = 0$.
- The 1-form a lifts from Σ, i.e., $a \in \pi^*(i\Omega^1(\Sigma))$.

Note that we can always assume that $\psi \in \Gamma(S^\pm)$ by Lemma 4.19; every Seiberg-Witten monopole is gauge equivalent to the pullback of a Kähler vortex, which is naturally either positive or negative.

We suppose for the remainder of the proof that $\psi \in \Gamma(S^+)$; the case where $\psi \in \Gamma(S^-)$ is analogous. We first work towards proving that ϕ is positive; to this end, decompose $\phi = \alpha + \beta$ into its positive and negative parts. Using the fact that $D_A\phi + \rho(a)\cdot\psi = 0$, we have the following:

$$\begin{aligned}
0 &= \langle D_A(D_A\phi + \rho(a)\cdot\psi), \beta\rangle \\
&= \langle (D_A)^2\phi + D_A(\rho(a)\cdot\psi), \beta\rangle \\
&= \langle (D_A)^2\phi, \beta\rangle + \langle -\rho((*\mathrm{d} + \mathrm{d}*)a)\cdot\psi, \beta\rangle + 2\langle -\nabla_{a^\flat}\psi + \xi\langle a, \eta\rangle\rho(\eta)\cdot\psi, \beta\rangle.
\end{aligned}$$

(by Equation 4.14)

Each of these terms can be individually evaluated:

- $\langle (D_A)^2\phi, \beta\rangle$: We have a decomposition of $(D_A)^2$ into terms by Proposition 4.3, and we have shown in Lemma 4.19 (Equation 4.49b) that the zeroth-order term vanishes. Thus, we can write $(D_A)^2$ as follows:

  $$(D_A)^2 = (\rho(\eta)\nabla^B_{\frac{\partial}{\partial\varphi}})^2 + (D_2)^2. \tag{4.62}$$

  These remaining terms preserve the splitting $W = W^+ \oplus W^-$, meaning we can replace ϕ with β in $\langle (D_A)^2\phi, \beta\rangle$; moreover, the hermiticity of the operators allows us to conclude that

  $$\langle (D_A)^2\phi, \beta\rangle = \left|\rho(\eta)\nabla_{\frac{\partial}{\partial\varphi}}\beta\right|^2 + |D_2\beta|^2. \tag{4.63}$$

  These terms are both nonnegative real numbers.

- $\langle -\rho((*\mathrm{d}+\mathrm{d}*)a)\cdot\psi,\beta\rangle$: We further break this term down into two more terms:

$$\langle -\rho((*\mathrm{d}+\mathrm{d}*)a)\cdot\psi,\beta\rangle = \langle \rho(\mathrm{d}a)\cdot\psi,\beta\rangle + \langle -\rho(\mathrm{d}*a)\psi,\beta\rangle. \tag{4.64}$$

  For the first term, we have by the definition of the Seiberg-Witten tangent space (Equation 4.54a) that $\rho(\mathrm{d}a)\cdot\psi = (\psi\otimes\phi^*+\phi\otimes\psi^*)_0\cdot\psi$; by expanding ϕ into positive and negative components, using the fact that ψ is positive, and discarding all positive terms under the inner product with β, this term reduces down to simply $\langle(\beta\otimes\psi^*)\psi,\beta\rangle = |\psi|^2|\beta|^2$.

  For the second term, observe that d^*a is a 3-form and therefore a scalar multiple of the volume form. However, Clifford multiplication by the volume form preserves the splitting $W = W^+\oplus W^-$, meaning $\rho(\mathrm{d}^*a)\psi$ and β are orthogonal and the second term vanishes. Thus, we observe that

$$\langle -\rho((*\mathrm{d}+\mathrm{d}*)a)\cdot\psi,\beta\rangle = |\psi|^2|\beta|^2. \tag{4.65}$$

- $2\langle -\nabla_{a^\flat}\psi + \xi\langle a,\eta\rangle\rho(\eta)\cdot\psi,\beta\rangle$: Observe that the covariant derivative and $\rho(\eta)$ both preserve the splitting $W^+\oplus W^-$. (The volume element of Σ is parallel under the Levi-Civita connection, meaning its eigenspaces are preserved by the spinorial connection; η is a product of factors which each preserve the splitting, namely $\eta = -(\mathrm{vol}_Y)(\mathrm{vol}_\Sigma)$.) It follows that $-\nabla_{a^\flat}\psi + \xi\langle a,\eta\rangle\rho(\eta)\cdot\psi$ is a positive spinor, meaning it is orthogonal to β:

$$2\langle -\nabla_{a^\flat}\psi + \xi\langle a,\eta\rangle\rho(\eta)\cdot\psi,\beta\rangle = 0. \tag{4.66}$$

Thus, we have the following:

$$\left|(\rho(\eta)\nabla_{\frac{\partial}{\partial\varphi}}\beta\right|^2 + |D_2\beta|^2 + |\psi|^2|\beta|^2 = 0. \tag{4.67}$$

All of these terms are nonnegative, so $|\psi||\beta| = 0$. But ψ is nonzero on a dense open set, meaning $\beta = 0$ everywhere. This proves that ϕ is positive.

We now set out to prove that $\nabla_{\frac{\partial}{\partial\varphi}}\phi = 0$. A consequence of the fact that $\psi,\phi\in\Gamma(S^+)$ is that the operator $-\rho(*\mathrm{d}a) = (\psi\otimes\phi^*+\phi\otimes\psi^*)_0$ preserves the splitting $W^+\oplus W^-$, meaning $*\mathrm{d}a$ must be a multiple of η, which in turn means that $e_0\lrcorner\,\mathrm{d}a = 0$. Since the connection on T^*Y is designed to pull up from Σ, it follows that $\nabla_{\frac{\partial}{\partial\varphi}}a = 0$, and using the slice condition, we therefore have the following:

$$\frac{\partial}{\partial\varphi}\mathrm{Im}\langle\psi,\phi\rangle = i\frac{\partial}{\partial\varphi}\mathrm{d}^*a = 0. \tag{4.68}$$

Additionally, since a is parallel in the $\frac{\partial}{\partial\varphi}$-direction, so too is $\mathrm{d}a$. It can be shown that $\mathrm{d}a = i\mathrm{Re}\langle\psi,\phi\rangle\mathrm{vol}_\Sigma$, which implies that

$$\frac{\partial}{\partial\varphi}\mathrm{Re}\langle\psi,\phi\rangle = 0 \tag{4.69}$$

as well.

With this in mind, write $a = a_i\mathrm{d}x^i$, where a_i are imaginary-valued functions on Y. The positive spinor part of the equation $D_A\phi + \rho(a)\cdot\psi = 0$ is equivalent to

$$i\nabla_{\frac{\partial}{\partial\varphi}}\phi + ia_0\psi = 0. \tag{4.70}$$

We know that $i\nabla_{\frac{\partial}{\partial\varphi}}|_{W^+}$ is self-adjoint; applying this operator to the above equation and taking the inner product with ϕ, we observe that

$$0 = \|i\nabla_{\frac{\partial}{\partial\varphi}}\phi\|_{L^2}^2 - \int_Y \frac{\partial a_0}{\partial\varphi}\langle\psi,\phi\rangle. \tag{4.71}$$

However, the latter term is zero; integrating by parts, one finds a total derivative and a φ-derivative of the inner product, both of which are zero. Iv follows that $\nabla_{\frac{\partial}{\partial\varphi}}\phi = 0$, meaning ϕ pulls up from Σ.

Finally, we use the equation $i\nabla_{\frac{\partial}{\partial\varphi}}\phi + ia_0\psi = 0$ to show that a_0 must vanish on the nonzero set of ψ, which is dense and open, meaning a_0 vanishes everywhere. It follows that $a = a_1\mathrm{d}x^1 + a_2\mathrm{d}x^2$, which manifestly pulls up from Σ.

Thus, we may write $\phi = \pi^*\phi_0$ and $a = \pi^*b$, where $(\phi_0, b) \in \Gamma(S^+) \times i\Omega^1(\Sigma)$. Using the exact same logic as in the injectivity argument, it is easy to show that (ϕ_0, b) must be a tangent vector to $\mathscr{M}^*_{\mathrm{vtx}}$ at the point (ψ_0, B), where $\psi = \pi^*\psi_0$ and $A = \pi^*B$. We conclude that the two tangent spaces are isomorphically identified by π^*. □

4.4.3 Characterisation of the irreducible moduli space

Combining these results together, we have the following characterisation of the irreducible component of the moduli space.

Theorem 4.21. *Let $Y = S(N) \to \Sigma$ be a Seifert fibred space equipped with a compatible action by a finite group Λ, and given some Λ-equivariant line bundle $L \to \Sigma$, denote by $\mathscr{C}^{\pm}(L)$ the moduli spaces of Λ-equivariant positive and negative Kähler vortices on L, respectively. Let $W = E \otimes (\mathbb{C} \oplus \pi^*K_\Sigma^{-1})$ be a $\widehat{\Lambda}$-equivariant $\mathrm{Spin}^c(3)$-structure on Y, where $E \to Y$ is some equivariant line bundle. The moduli space of irreducible equivariant Seiberg-Witten monopoles on W is nonempty if and only if $E \cong \pi^*(L_0)$ for some equivariant line bundle $L_0 \to \Sigma$. In this case, there is the following relationship between the two moduli spaces:*

$$\mathscr{M}^*_{\mathrm{sw}}(W) = \bigsqcup_k \mathscr{C}^+(L_0 \otimes N^{\otimes k}) \sqcup \mathscr{C}^-(L_0 \otimes N^{\otimes k}), \tag{4.72}$$

where $k \in \mathbb{Z}$ runs over the integers for which $0 \leq \deg(L_0 \otimes N^{\otimes k}) < -\frac{\chi(\Sigma)}{2}$. The spaces $\mathscr{C}^{\pm}(L)$ are each diffeomorphic to the space of effective divisors on Σ/Λ with degree $\deg(L)/|\Lambda|$ and $\deg(L^* \otimes K_\Sigma)/|\Lambda|$ respectively.

Proof. We have shown above that the moduli space of Seiberg-Witten monopoles is diffeomorphic to the disjoint union of $\mathscr{M}^*_{\mathrm{vtx}}(S)$ over all spinor bundles S related to W by pullback. Since the line bundle of positive spinors characterises S up to isomorphism, and since the map $\pi^* : \mathrm{Pic}^t(\Sigma) \to \mathrm{Pic}^t(Y)$ has kernel generated by the bundle N, every such spinor bundle can be written in the form $S \otimes N^{\otimes k}$ for a fixed bundle S and an integer k. Furthermore, we have shown in Chapter 2 that the moduli space of irreducible positive (resp. negative) Kähler vortices on $L \to \Sigma$ is diffeomorphic to the space of effective divisors of Σ corresponding to L (resp. $K_\Sigma \otimes L^*$) when $0 \leq \deg(L) < \deg(K_\Sigma)/2$, and it is empty otherwise. □

Due to the restriction on the degree of the line bundle L, we obtain the following topological restriction on the existence of irreducible monopoles:

Corollary 4.22. *If $\chi(\Sigma) \geq 0$, then the moduli space of irreducible monopoles is empty. In particular, if Σ is a sphere with less than three marked points, or if Σ is a torus with no marked points, then there are no irreducible monopoles.*

Proof. Apply Equation 1.31 for the orbifold Euler characteristic, and set $g = 0$ or $g = 1$. □

4.5 Reducible Moduli Correspondence

Once again, let $Y \to \Sigma$ be an equivariant Seifert fibration with respect to the Λ-action on Σ and the $\widehat{\Lambda}$-action on Y. Recall that the equivariant $\mathrm{Spin}^c(3)$-structures on Y are in one-to-one correspondence with equivariant line bundles over Y, and a spinorial connection on a spinor bundle $W \to Y$ can be identified with the corresponding unitary connection on W^+, which we denote by E. Under the assumption of reducibility, the Seiberg-Witten equations for $(0, A)$ then simplify to the following:

$$F_A = \tfrac{1}{2} F_{K_\Sigma} = -\frac{i\pi \deg(K_\Sigma)}{\mathrm{Vol}(\Sigma)} \pi^* \mathrm{vol}_\Sigma. \tag{4.73}$$

In this section, we give a characterisation of the structure of the space of solutions to this equation.

First, we have the following existence theorem for constant curvature connections over Y. (Information about reducible monopoles is recovered by taking $\gamma = -\pi \deg(K_\Sigma)/\mathrm{Vol}(\Sigma)$.)

Proposition 4.23. *Let $Y \to \Sigma$ and $E \to Y$ be as above, and let $\gamma \in \mathbb{R}$ be arbitrary. Then there exists a connection $A \in \mathscr{A}(E)$ satisfying $F_A = i\gamma\pi^* \mathrm{vol}_\Sigma$ if and only if one of the following conditions is satisfied:*

- $\deg(Y) \neq 0$ *as a* $\mathrm{U}(1)$*-bundle over* Σ*, and* E *pulls up from a* Λ*-equivariant line bundle* L *over* Σ*.*
- $\deg(Y) = 0$ *as a* $\mathrm{U}(1)$*-bundle over* Σ*,* E *pulls up from a* Λ*-equivariant line bundle* L *over* Σ*, and* $\gamma = -2\pi \deg(L)/\mathrm{Vol}(\Sigma)$*.*

Proof. If $F_A = i\gamma\pi^* \mathrm{vol}_\Sigma$ for a $\widehat{\Lambda}$-equivariant connection A, then the fact that E pulls up from Σ is guaranteed by Proposition 1.55, and it inherits a Λ-equivariant structure by the same argument as in the proof of Lemma 4.19. We assume in the remainder of the proof that $E = \pi^* L$, where $L \to \Sigma$ is a Λ-equivariant line bundle, and we use Proposition 1.50 to find a connection $B \in \mathscr{A}(L)$ with curvature $(-2\pi i \deg(L)/\mathrm{Vol}(\Sigma))\mathrm{vol}_\Sigma$ on L. (We may assume that B is Λ-equivariant by assuming that it lifts from a constant-curvature connection on Σ/Λ.)

The statement to be proved now is as follows: the existence of a connection $A \in \mathscr{A}(\pi^* L)$ with curvature $i\gamma\pi^* \mathrm{vol}_\Sigma$ is guaranteed if $\deg(Y) \neq 0$, and equivalent to the condition that $\gamma = -2\pi \deg(L)/\mathrm{Vol}(\Sigma)$ if $\deg(Y) = 0$. In the case that $\deg(Y) \neq 0$, the following choice of $A \in \mathscr{A}(E)$ is readily computed to have the desired curvature:

$$A = \pi^* B + \left(\frac{-\frac{1}{2\pi}\gamma \mathrm{Vol}(\Sigma) - \deg(L)}{\deg(N)} \right) i\eta. \tag{4.74}$$

Additionally, if $\deg(Y) = 0$ and $\gamma = -2\pi\deg(L)/\mathrm{Vol}(\Sigma)$, the choice $A = \pi^*B$ clearly has the desired curvature. The remainder of the proof is devoted to the following claim: if $\gamma \neq -2\pi\deg(L)/\mathrm{Vol}(\Sigma)$ and a connection exists with curvature $i\gamma\pi^*\mathrm{vol}_\Sigma$, then $\deg(Y) \neq 0$.

Let $A \in \mathscr{A}(\pi^*L)$ be a connection with curvature $i\gamma\pi^*\mathrm{vol}_\Sigma$ for $\gamma \neq -2\pi\deg(L)/\mathrm{Vol}(\Sigma)$. Defining a 1-form $\alpha = A - \pi^*B$, we see that $d\alpha = F_A - \pi^*F_B$ is a nonzero constant multiple of $\pi^*\mathrm{vol}_\Sigma$; it follows by rearranging that $\pi^*\mathrm{vol}_\Sigma \in \Omega^2(Y)$ is exact, or equivalently its corresponding de Rham cohomology class in $H^2(Y;\mathbb{R})$ is zero. Thus, the map $\pi^* : H^2(|\Sigma|;\mathbb{R}) \to H^2(Y;\mathbb{R})$ takes the class defined by vol_Σ to the zero class; without loss of generality, we may assume that vol_Σ corresponds to an integral cohomology class by dividing by $\mathrm{Vol}(\Sigma)$.

Recall that Theorem 1.53 gives the following formula for the second integral cohomology of Y:

$$H^2(Y;\mathbb{Z}) = (H^2(\Sigma;\mathbb{Z})/\langle c_1(N)\rangle) \oplus \mathbb{Z}^{2g}, \tag{4.75}$$

where the first direct summand is the image of $H^2(\Sigma;\mathbb{Z})$ under π^*. By tensoring with $\mathbb{R}$, we see that $\pi^*\mathrm{vol}_\Sigma$ is trivial in $H^2(Y;\mathbb{R})$ if and only if there is some nonzero rational $q \in \mathbb{Q}^*$ for which $c_1(N) = q\mathrm{vol}_\Sigma$. However, by pairing vol_Σ and $c_1(N)$ with the fundamental class of Σ, we see that

$$\deg(N) = \langle c_1(N), [\Sigma]\rangle = \langle q\mathrm{vol}_\Sigma, [\Sigma]\rangle = q\mathrm{Vol}(\Sigma), \tag{4.76}$$

which proves that $\deg(N)$ must be nonzero for such a connection to exist.[2] □

We use this existence result to prove the following general structure theorem for the moduli space of reducible monopoles.

Theorem 4.24. *Let $Y \to \Sigma$ be a Seifert fibred space equipped with a compatible $\widehat{\Lambda}$-action, and let $W \to Y$ be a $\widehat{\Lambda}$-equivariant spinor bundle over Y. The moduli space of reducible monopoles is nonempty if and only if one of the following conditions is met:*

- $\deg(Y) \neq 0$ *and* $W^+ \cong \pi^*L$ *for some Λ-equivariant line bundle $L \to \Sigma$;*
- $\deg(Y) = 0$, $W^+ \cong \pi^*L$ *for some Λ-equivariant line bundle $L \to \Sigma$, and* $2\deg(L) = \deg(K_\Sigma)$.

In this case, the moduli space is homeomorphic to the space of flat $\widehat{\Lambda}$-equivariant trivial bundles on Y, which is in turn equivalent to $H^1(|Y/\widehat{\Lambda}|,\mathbb{R})/H^1(|Y/\widehat{\Lambda}|,\mathbb{Z})$, the Jacobian torus of $Y/\widehat{\Lambda}$.

Proof. The existence result is obtained by taking $\gamma = -\pi\deg(K_\Sigma)/\mathrm{Vol}(\Sigma)$ in Proposition 4.23. To find all other possible monopoles, note that any two bundles with the same curvature differ by tensoring with a flat line bundle, and conversely tensoring with a flat line bundle does not change the curvature. It follows that the space of equivariant reducible monopoles on W is a torsor on the space of equivariant trivial flat bundles. Such a bundle descends to the quotient $Y/\widehat{\Lambda}$ as a flat trivial orbifold line bundle by definition, and by Proposition 1.38, the space of trivial flat bundles on $Y/\widehat{\Lambda}$ is given by its Jacobian torus. □

[2] We have glossed over the introduction of a fundamental class for orbifolds, but such a cohomology class does exist with rational coefficients so long as Σ is orientable and has singularities with codimension at least 2. The reader is referred to [GM80] for details.

Corollary 4.25. *In the case that Λ is trivial and* $\deg(N) \neq 0$*, the moduli space (if nonempty) reduces to the space of flat connections on Σ, i.e., the Jacobian torus of Σ.*

Proof. This follows from the fact that $H^1(Y) \cong H^1(\Sigma)$ whenever $\deg(N) \neq 0$ (cf. Theorem 1.53). □

Finally, we briefly discuss the tangent space structure of the reducible locus. Recall that the reducible locus does not necessarily have a manifold structure, but we defined its tangent space at $(0, A)$ to consist of harmonic 1-forms and A-harmonic spinors. Since the Jacobian torus is a quotient of $H^1(|Y/\widehat{\Lambda}|, \mathbb{R})$ by a discrete lattice, its tangent space can be identified with $H^1(|Y/\widehat{\Lambda}|, \mathbb{R})$ as well. This leads to the following result.

Proposition 4.26. *The tangent space of $\mathscr{M}_{\mathrm{sw}}$ at $(0, A)$ is identified with the tangent space of $H^1(|Y/\widehat{\Lambda}|, \mathbb{R})/H^1(|Y/\widehat{\Lambda}|, \mathbb{Z})$ if and only if* $\ker(D_A)$ *is trivial.*

On the other hand, the kernel of the Dirac operator for a reducible monopole is not generally trivial. Partial results on the kernel of the Dirac operator are given in [MOY96].

4.6 Examples

Now that the structure of the moduli space has been established, we reserve the remainder of this chapter for examples of specific Seifert fibred spaces. Before applying the results to specific spaces, however, we make the following general remark on the structure of Seifert fibred spaces:

Proposition 4.27. *Every Seifert 3-orbifold fibring over a very good orbifold Riemann surface may be written in the form $\widetilde{Y}/\widehat{\Lambda} \to \widetilde{\Sigma}/\Lambda$, where $\widetilde{Y} \to \widetilde{\Sigma}$ is a Seifert 3-manifold, Λ is a finite group acting compatibly on $\widetilde{Y}$ and $\widetilde{\Sigma}$, and $\widehat{\Lambda}$ is a cyclic group extension of Λ.*

Proof. Let $Y \to \Sigma$ be a Seifert orbifold, with Σ very good. Then there is a finite covering $\widetilde{\Sigma} \to \Sigma$ by a (nonsingular) Riemann surface, and we denote its group of deck transformations by Λ. By pulling back along this covering, we obtain another circle bundle $\widetilde{Y} \to \widetilde{\Sigma}$; since $\widetilde{\Sigma}$ is nonsingular, so too is $\widetilde{Y}$. We then lift the action of Λ to $\widetilde{Y}$ via Proposition 4.13, which is compatible with the projections by definition. □

This demonstrates that, theoretically, there is a wealth of examples for which the group Λ does not act freely. Recall that the U(1)-bundle with Seifert invariant $(b, \beta_1, \dots, \beta_n)$ over the orbifold Riemann surface $\Sigma_{\alpha_1,\dots,\alpha_n}$ defines a Seifert 3-manifold precisely when α_i and β_i are coprime for all i. The case for which α_i and β_i share factors therefore defines a Seifert 3-orbifold. Of course, if we wish to define equivariant Seiberg-Witten theory on a *given* Seifert 3-manifold $Y' \to \Sigma'$, these examples are only useful insofar as we can demonstrate that the Seifert 3-manifold agrees with $\widetilde{Y} \to \widetilde{\Sigma}$ for some Seifert orbifold $Y \to \Sigma$. Further research into the isomorphism types of these liftings may therefore be of interest.

4.6.1 Seifert 3-manifolds fibring over non-orientable 2-orbifolds

Recall that a Seifert 3-manifold over a non-orientable 2-orbifold may be conceptualised as a Seifert 3-manifold equipped with a compatible $\mathbb{Z}_2$-action, which acts antiholomorphically on the base. By using the decomposition of the irreducible solutions into positive and negative spinors, we see that the irreducible moduli space reduces drastically:

Theorem 4.28. *If Y is a Seifert 3-manifold fibring over a non-orientable 2-orbifold Σ with underlying space $|\Sigma| \cong (\mathbb{RP}^2)^{\# k}$, then the moduli space of irreducible Seiberg-Witten monopoles is empty. The moduli space of reducible Seiberg-Witten monopoles is nonempty if and only if $\chi(\Sigma) = 0$, in which case it is homeomorphic to $\mathrm{U}(1)^k$.*

Proof. We first realise Σ as a Klein orbifold. Let $q : \widetilde{\Sigma} \to \Sigma$ be the (connected) orientation doubling of Σ, and equip $\widetilde{\Sigma}$ with a holomorphic structure which flips sign under the natural involution. By pulling back along q, we can realise the Seifert fibration $Y \to \Sigma$ as a $\mathbb{Z}_2$-equivariant Seifert fibration $\widetilde{Y} := q^*Y$ over $\widetilde{\Sigma}$; data on Y and Σ corresponds to invariant data on $\widetilde{Y}$ and $\widetilde{\Sigma}$. Note that the volume element on $\widetilde{\Sigma}$ naturally flips sign under $\mathbb{Z}_2$.

Given a spinor bundle $W \to Y$, we may once again use pullback to get a $\mathbb{Z}_2$-equivariant spinor bundle $\widetilde{W} := q^*W \to \widetilde{Y}$. We get a natural Clifford module structure $\widetilde{\rho} : T\widetilde{Y} \to \mathfrak{su}(\widetilde{W})$ as follows: given $\widetilde{v} \in T\widetilde{Y}$ lifting $v \in TY$ and $\widetilde{\psi} \in \widetilde{W}$ lifting $\psi \in W$, we define $\widetilde{\rho}(\widetilde{v}) \cdot \widetilde{\psi}$ to be the unique element of $\widetilde{W}$ corresponding to $\rho(v) \cdot \psi$ under the quotient $\widetilde{W} \to W$. This is well-defined because the projection $q^*W \to W$ is surjective and $\mathbb{Z}_2$-invariant, implying that such an element exists, and the projection is also a fibrewise linear isomorphism so it is unique.

Thus, the moduli space of irreducible Seiberg-Witten monopoles on Y is equivalent to the moduli space of $\mathbb{Z}_2$-equivariant monopoles on $\widetilde{Y}$, which in turn is equivalent to the moduli space of $\mathbb{Z}_2$-equivariant Kähler vortices on $\widetilde{\Sigma}$. Additionally, a Kähler vortex is contained entirely in either the positive or the negative spinor bundle, meaning the volume element acts on the spinor by multiplication by ± 1; thus, the action of vol_Σ must exchange positive and negative spinors. If a Kähler vortex is required to be $\mathbb{Z}_2$-equivariant, therefore, it must be both positive and negative at once. We conclude that the spinor must be zero, meaning all $\mathbb{Z}_2$-equivariant monopoles must be reducible.

Now, if $A \in \mathscr{A}(W^+)$ is a $\mathbb{Z}_2$-invariant connection representing a monopole, the equation $F_A = \frac{1}{2} K_\Sigma$ reduces to $F_A = 0$, since the $\mathbb{Z}_2$-action takes the canonical bundle to its dual and $F_{K_\Sigma^*} = -F_{K_\Sigma}$. It follows that $\deg(K_\Sigma) = 0$, meaning $\chi(\Sigma) = 0$. The moduli space of reducible solutions is therefore nonempty if and only if $\chi(\Sigma) = 0$, in which case it is given by the Jacobian torus of Y. Since Σ is non-orientable, its second cohomology vanishes; it follows that all line bundles over Σ are trivial, so $\deg(Y) = 0$. It follows that $H^1(Y) \cong H^1(|\Sigma|) \oplus \mathbb{Z}$, and $H^1(|\Sigma|) \cong \mathbb{Z}^{k-1}$. Thus, the Jacobian torus of Y is given by $\mathrm{U}(1)^k$. □

4.6.2 Lens spaces

The next set of examples is generated by the lens space construction.

Definition 4.29. Let p, q be coprime integers with $p > 0$ and $q \neq 0$, and consider $S^3 \subseteq \mathbb{C}^2$. The *lens space* $L(p, q)$ is defined to be $S^3/\mathbb{Z}_p$, where the generator of $\mathbb{Z}_p$ acts on $(z, w) \in S^3$

as follows:

$$(z, w) \mapsto (e^{2\pi i/p} z, e^{2\pi i q/p} w). \tag{4.77}$$

The definition is extended to arbitrary coprime integers p, q by defining $L(0,1) = S^2 \times S^1$ and $L(-p,-q) = L(p,q)$.

Note that the coprimality of p and q ensures that the action is free, making all lens spaces nonsingular.

Proposition 4.30. *Every lens space can be equipped with a Seifert fibration, the base of which is either S^2 (in which case there are at most two singular fibres) or $\mathbb{RP}^2$ (in which case there is at most one singular fibre).*

Proof. Defining a Seifert fibration for $L(p,0)$ is trivial, so we assume $q \neq 0$. Given coprime nonzero integers k, ℓ, define the following U(1)-action on $L(p,q)$:

$$e^{i\theta} \cdot (z, w) = (e^{ik\theta} z, e^{i\ell\theta} w), \tag{4.78}$$

for any $e^{i\theta} \in \mathrm{U}(1)$. Note that the action commutes with the $\mathbb{Z}_p$-action, so it is well-defined on $L(p,q)$. Since the U(1)-action is free except at $z = 0$ or $w = 0$, where the stabiliser subgroup is either $\mathbb{Z}_\ell$ or $\mathbb{Z}_k$, we see that $L(p,q)/\mathrm{U}(1)$ is a 2-orbifold and the quotient map $L(p,q) \to L(p,q)/\mathrm{U}(1)$ defines a Seifert fibration.

To characterise the base, note that every loop in the base of a Seifert fibration may be lifted to a loop in the total space, meaning the map $\pi_1(L(p,q)) \to \pi_1(L(p,q)/\mathrm{U}(1))$ induced by the quotient must be surjective. But $\pi_1(L(p,q)) \cong \mathbb{Z}/p\mathbb{Z}$, so the fundamental group of the base is also cyclic; the only closed surfaces with cyclic fundamental group are S^2 and $\mathbb{RP}^2$.

Finally, to understand the singular fibres of the base, we may analyse the fundamental group of a Seifert fibred space (cf. Theorem 1.53) and verify that the conditions on the number of fibres are sufficient to ensure that the fundamental group is finite cyclic. For details, refer to [GL18]. □

In order to use the equivariant machinery we have developed thus far on a lens space $L(p,q)$, we first need to equip such a space with an isometric action of a finite group preserving the fibres of its Seifert fibration. Such group actions have been classified in [KM]; compiling results from this paper, we have the following:

Proposition 4.31. *Let $L(p,q)$ be a lens space with $p > 2$. The diffeomorphism group of $L(p,q)$ consists only of orientation-preserving diffeomorphisms if and only if $q^2 \neq -1$ mod p. A finite isometric group action on $L(p,q)$ preserves the fibres of a Seifert fibration of $L(p,q)$ if and only if it preserves orientation. If $q^2 = 1$ mod p, then the group $D_m \times \mathbb{Z}_\ell$ with $m, \ell \in \mathbb{N}$ and ℓ odd acts isometrically and non-freely on $L(p,q)$, so $L(p,q)$ admits a finite isometric group action compatible with a Seifert fibration.*

By applying Theorems 4.21 and 4.24 to $L(p,q) \to L(p,q)/\mathrm{U}(1)$, we state the structure of the moduli space for $L(p,q)$.

Proposition 4.32. *Let $L(p,q)$ be a lens space with $p > 2$ and $q^2 = 1$ mod p, and let $\Lambda = D_m \times \mathbb{Z}_\ell$. Given any Λ-equivariant spinor bundle over $L(p,q)$, the moduli space of irreducible monopoles is empty, and the moduli space of reducible monopoles is homeomorphic to the Jacobian torus of $|L(p,q)/\Lambda|$.*

4.7 Further Directions

Up to this point, we have established a framework for Λ-equivariant Seiberg-Witten theory on Seifert 3-manifolds. The next natural step is to follow the programme developed by Kronheimer and Mrowka in [KM07], and use these moduli spaces to define and compute Λ-equivariant Floer homology groups.

A related approach has been carried out by Baraglia and Hekmati in [BH24], but as an equivariant generalisation of Manolescu's construction of a Seiberg-Witten-Floer stable homotopy type [Man01] rather than Mrowka et al.'s construction of a moduli space on a Seifert 3-manifold. While the Seiberg-Witten-Floer homology groups are known to have a module structure over the ring $H^*_{\mathrm{U}(1)}(\mathrm{pt}) \cong \mathbb{Z}[U]$, it was found that the G-equivariant homology groups formed modules over $H^*_{G\times \mathrm{U}(1)}(\mathrm{pt}) \cong H^*(BG)[U]$; this richer module structure leads to more sophisticated invariants. However, a notable obstruction to constructing these invariants for a given 3-manifold Y is that the first Betti number of Y must be zero. Beyond this class of 3-manifolds, results are limited; a version of equivariant Seiberg-Witten-Floer homology was touched upon in [LM18], but only for free group actions and instead within the framework of Heegaard-Floer homology.

The work we have done in this thesis could potentially lead to extensions beyond this work. In general, there is no restriction on $b_1(Y)$ for Seifert 3-manifolds, so a generalisation of the equivariant Seiberg-Witten-Floer homology based on the moduli spaces we have constructed could be applied to a larger class of manifolds. A potential obstruction to this programme is in the transversality of flow lines: the gradient flow lines of csd must be transverse to the gauge orbits in $\Gamma(W) \times \mathscr{A}(W)$ in order to have finitely many flow lines between critical points, but the introduction of a finite group action adds extra conditions to this transversality which may not be met in general.

If the equivariant generalisation is successful, the next step would be to understand the ways in which a finite group may act compatibly on a Seifert 3-manifold, and the possible cyclic extensions that may be obtained. Indeed, as was noted earlier, every Seifert 3-orbifold can be constructed as the quotient of a Seifert 3-manifold by a compatible group action. In order to apply the results from this thesis to such a space, however, the group must act isometrically on Y and Σ. It may be interesting to investigate the compatible isometries of Seifert 3-manifolds in more detail.

Appendix A

Sobolev Spaces on Orbifolds

In this appendix, we briefly summarise the results from global analysis on manifolds that will be used. At the end of each section, we briefly comment on their generalisations to the orbifold category.

A.1 Definitions of Sobolev Spaces

First, we define the notion of Sobolev spaces of sections of vector bundles. All of the following definitions generalise readily to the orbifold category.

Definition A.1. Let (X, g) be a smooth Riemannian manifold (not necessarily compact and possibly with boundary), and let $E \to X$ be a rank-r vector bundle. Equip E with a bundle metric h and covariant derivative $\nabla^E : \Gamma(E) \to \Omega^1(E)$. Extend the bundle metric using the (inverse) Riemannian metric on X to a metric on $(T^*M)^\ell \otimes E$ for all ℓ, and extend the covariant derivative using the Levi-Civita connection to a covariant derivative $\nabla : \Gamma((T^*X)^\ell \otimes E) \to \Gamma((T^*X)^{\ell+1} \otimes E)$.

Denote by $\Gamma(E)_c$ the space of smooth compactly supported sections of E. For every $k \in \mathbb{N}$ and every $p \in [1, \infty)$, define the L^p_k *Sobolev norm* to be the following for every $\alpha \in \Gamma(E)_c$:

$$\|\alpha\|_{L^p_k} = \left(\sum_{j=0}^{k} \int_X |\nabla^j \alpha|^p \right)^{1/p}. \tag{A.1}$$

In the case that $k = 0$, the L^p_k norm reduces to the L^p norm which we simply denote by $\|\alpha\|_p$.

The completion of $\Gamma(E)_c$ with respect to the L^p_k norm is called the *Sobolev space of L^p_k sections of E,* and it is denoted by $\Gamma(E)_{L^p_k}$ or $L^p_k(E)$. In the special case that E is the trivial rank-n bundle over X, we denote the space by $L^p_k(X, \mathbb{R}^n)$ and call its elements L^p_k functions from X to $\mathbb{R}^n$; if $n = 1$, the space is denoted by $L^p_k(X)$.

In the case that X is compact, the resulting Sobolev spaces do not depend on the choice of g, h, or ∇^E; however, the Sobolev norm itself does change with these structures.

This definition allows us to make sense of L^p_k tensors and ℓ-forms as L^p_k sections of various powers of the (co)tangent bundle. We may use this definition of Sobolev spaces to make sense of L^p_k connections as well, but to do this we must use the following fact: if $P \to X$ is a principal G-bundle with Lie algebra $\mathfrak{g}$, and $\mathrm{ad}(P)$ is the associated vector bundle $P \times_{\mathrm{Ad}} \mathfrak{g}$ with $\mathrm{Ad} : G \to \mathrm{Aut}(\mathfrak{g})$ being the adjoint action, then the space of smooth connections on G is an affine space modelled on the space of $\mathrm{ad}(P)$-valued 1-forms $\Omega^1(\mathrm{ad}(P))$ on X.

Definition A.2. Let X be a Riemannian manifold, let $P \to X$ be a principal G-bundle for a Lie group G, and choose a smooth connection $\widetilde{A} \in \mathscr{A}(P)$. Given $p \geq 1$ and $k \in \mathbb{N}$, the *Sobolev space of L^p_k connections* $\mathscr{A}^{k,p}(P)$ is defined as follows:

$$\mathscr{A}^{k,p}(P) = \widetilde{A} + \Omega^1(\mathrm{ad}(P))_{L^p_k}. \tag{A.2}$$

For compact base manifolds X, this definition does not depend on the choice of reference connection $\widetilde{A}$.

To define Sobolev spaces of mappings from X to another manifold Y, we essentially treat Y as Euclidean space using a coordinate chart and use Definition A.1. Whereas we previously allowed p and k to be any numbers and X to be any manifold, here we will require that $kp > \dim(X)$ and that X is compact; this requirement implies that all L^p_k maps are continuous (in the light of the Sobolev embedding theorems which will be stated later), which in turn ensures that the following definition is independent of the atlas.

Definition A.3. Let X and Y be Riemannian manifolds with X compact, and let $p \geq 1$ and $k \in \mathbb{N}$ be numbers for which $kp > \dim(X)$. The *Sobolev space of L^p_k maps from X to Y* is the collection of continuous maps $f : X \to Y$ such that, in any coordinate chart $\phi : U \to \mathbb{R}^{\dim(Y)}$ on Y, the map $\phi \circ f|_{f^{-1}(U)} : f^{-1}(U) \to \mathbb{R}^{\dim(Y)}$ is an L^p_k function.

We denote this space by $L^p_k(X, Y)$, and we define a topology on $L^p_k(X, Y)$ as follows: a sequence $f_j : X \to Y$ converges to $f : X \to Y$ if and only if f_i converges to f in the $\mathscr{C}^0$-topology and $\phi \circ f_j$ converges to $\phi \circ f$ in the L^p_k norm for all coordinate charts $\phi : U \to \mathbb{R}^{\dim(Y)}$ on Y.

It is worth noting that there are two other equivalent definitions for L^p_k maps. If $\Phi : Y \to \mathbb{R}^N$ is a smooth embedding, then $f : X \to Y$ is L^p_k if and only if $\Phi \circ f : X \to \mathbb{R}^N$ is L^p_k, which in turn holds if and only if there is some smooth map $s : X \to Y$ and some smooth section v of the pullback bundle $s^*TY \to X$ such that $f(x) = \exp_{s(x)}(v_x)$, where exp is the Riemannian exponential on Y.

Lemma A.4. *Let X and Y be orbifolds, let $\Phi : Y \to \mathbb{R}^N$ be a smooth embedding, and let $kp > n$. Then the following are equivalent for a continuous map $f : X \to Y$:*

- *The map f is L^p_k, in the sense of Definition A.3;*
- *The map $\Phi \circ f : X \to \mathbb{R}^N$ is L^p_k, in the sense of Definition A.1;*
- *There is a smooth map $s : X \to Y$ and an L^p_k vector field $v \in \Gamma(s^*TY)$ for which $f(p) = \exp_{s(p)}(v_p)$, where* exp *denotes the Riemannian exponential on Y.*
- *In the case that Y is a compact Lie group G with Lie algebra $\mathfrak{g}$: $f^{-1}\mathrm{d}f$ is an L^p_k $\mathfrak{g}$-valued 1-form, in the sense of A.1.*

Furthermore, given a sequence of L^p_k maps $f_j : X \to Y$, the following are equivalent:

- *$f_j \to f$ in the topology defined by Definition A.3;*
- *$\Phi \circ f_j \to \Phi \circ f$ in the Sobolev norm on the trivial bundle $X \times \mathbb{R}^N$;*
- *There exists a smooth map $s : X \to Y$, an L^p_k-section $v \in \Gamma(s^*TY)$, and for sufficiently large j there exist smooth maps $s_j : X \to Y$ and L^p_k-sections $v_j \in \Gamma(s_j^*TY)$ such that $f = \exp_s(v)$, $f_j = \exp_{s_j}(v_j)$, and $v_j \to v$ in the L^p_k-norm.*
- *In the case that Y is a compact Lie group G with Lie algebra $\mathfrak{g}$: $f_j^{-1}\mathrm{d}f_j$ is a sequence of L^p_{k-1} $\mathfrak{g}$-valued 1-forms converging to $f^{-1}\mathrm{d}f$ in the L^p_{k-1}-norm.*

Similar results apply for uniform L^p_k-bounds.

This definition allows us to consider the notion of an L^p_k gauge transformation of a principal G-bundle $P \to X$, since these can be regarded either as G-equivariant maps $\psi : P \to P$ or as G-equivariant maps $g : P \to G$. We denote the space of L^p_k gauge transformations of P by $\mathscr{G}^{k,p}(P)$.

Definition A.5. Let X and F be manifolds, and let $E \to X$ be an F-fibre bundle. The *Sobolev space of L^p_k sections of E* is the space of $\mathscr{C}^0$ sections $\sigma \in \Gamma(E)_{\mathscr{C}^0}$ such that, for any smooth local trivialisation $\psi : E|_U \to U \times F$, the map $\psi \circ \sigma|_U : U \to F$ is an L^p_k function. We denote this space of sections by $\Gamma(E)_{L^p_k}$ or $L^p_k(E)$, and we require that a sequence of sections converges in L^p_k if and only if their local representations all converge in L^p_k.

These definitions generalise readily to orbifolds in the case that Y is a manifold (this is required for a smooth embedding to exist). The only meaningful diversion is in the definition of a local trivialisation, but this can be fixed by thinking of the section σ over an orbifold chart $(\widetilde{U}, \Gamma, \varphi)$ as a Γ-invariant map $\widetilde{\sigma} : \widetilde{U} \to E|_U$, lifting this map to a Γ-equivariant map $\widetilde{\sigma} : \widetilde{U} \to \widetilde{U} \times F$, and proceeding in the same way.

A.2 Embedding and Multiplication Theorems

The Sobolev embedding theorems give a series of relationships between the different Sobolev spaces; though they were originally stated for bounded domains in $\mathbb{R}^n$, they generalise to compact manifolds and even compact orbifolds. Henceforth, we will denote by n the dimension of the compact manifold X.

Theorem A.6. *Let X be a compact n-manifold, and let $\pi : E \to X$ be a rank-r vector bundle. Let $j, k \in \mathbb{N}$ be natural numbers for which $j < k$, and let $p, q \in [1, \infty)$.*

- *If $k - n/p \geq j - n/q$ then the inclusion $L^p_k(E) \hookrightarrow L^q_j(E)$ is continuous, meaning there is some constant $C > 0$ such that $\|\sigma\|_{L^q_j} \leq \|\sigma\|_{L^p_k}$ for all $\sigma \in \Gamma(E)_{L^p_k}$. If the inequality is strict then the inclusion is a compact map, meaning a bounded sequence of L^p_k sections has a convergent subsequence in L^q_j.*
- *If $k - n/p > j$ then the inclusion $L^p_k(E) \hookrightarrow C^j(E)$ is continuous and compact.*

The generalisation of this theorem from domains in $\mathbb{R}^n$ to sections of vector bundles over compact manifolds is straightforward. To do so, choose a bundle atlas $\phi_i : \pi^{-1}(U_i) \to U_i \times \mathbb{R}^r$ for E; since X is compact, only finitely many charts (say N) are needed and each U_i can be assumed Euclidean. Then, one can represent an arbitrary section of E as a collection of maps $\sigma_i : U_i \to \mathbb{R}^r$, and furthermore break this into components $(\sigma_i)_j : U_i \to \mathbb{R}^r$ for j ranging from 1 to r. The following norm is equivalent to the usual Sobolev norm:

$$\|\sigma\|'_{L^p_k} = \sum_{i \leq N} \sum_{j \leq r} \|(\sigma_i)_j\|_{L^p_k(U_i)}, \tag{A.3}$$

where $L^p_k(U_i)$ is the Sobolev space of real-valued functions on the Euclidean domain U_i. The Sobolev embedding theorems for bounded domains allow us to bound the L^p_k norms in terms of the L^q_j norms, which can be made into a bound for the norm $\|\cdot\|'_{L^p_k}$ in terms of $\|\cdot\|'_{L^q_j}$ since there are only finitely many charts. Given an L^p_k-bounded sequence of sections $(\sigma^\nu)_{\nu \in \mathbb{N}}$, the compact inclusions on bounded domains allows us to find L^q_j-convergent subsequences of the components of σ^ν over each chart; since there are only finitely many components and charts, we can make this into an L^q_j-convergent subsequence of the σ^ν over all of X. Note that this proof generalises easily to the case of orbifolds, since orbifold vector bundles afford similar bundle atlases and therefore allow for the same strategy.

The second major theorem that we will use extensively is the Sobolev multiplication theorem, which gives sufficient conditions for multiplying two Sobolev functions on a compact manifold.

Theorem A.7. *Let X be a compact manifold, let $k \in \mathbb{N}$, and let $p, r, s \in [1, \infty)$ be real numbers satisfying $r, s \geq p$ and $\frac{1}{r} + \frac{1}{s} \geq \frac{k}{n} + \frac{1}{p}$, with the additional constraint that at least one of these inequalities is strict. Then the multiplication function $L^r_k(X) \times L^s_k(X) \to L^p_k(X)$ is continuous, which is to say that there is a constant $C > 0$ such that the following holds for every $f \in L^r_k(X)$ and $g \in L^s_k(X)$:*

$$\|fg\|_{L^p_k} \leq C\|f\|_{L^r_k}\|g\|_{L^s_k}. \tag{A.4}$$

The proof of this theorem is given in [Weh04]; it is entirely functional-analytic, meaning it only relies on facts about the Sobolev norms (including the embedding theorems above), and therefore the logic applies equally to orbifolds. One special case of this theorem are important for gauge theory:

Corollary A.8. *If $kp > n$, $q \leq p$, and $j \leq k$, then $L^p_k(X)$ acts continuously on $L^q_j(X)$, and $L^p_k(X)$ has a natural topological ring structure.*

This gives sufficient conditions for the regularity of L^p_k gauge transformations so that they act continuously on spaces of sections and connections. The geometric analysis we do will essentially be limited to vortices on Klein orbifolds for which $n = 2$; it follows that a sufficient regularity for gauge transformations is L^2_2.

The third major theorem we will use is the weak compactness properties of the Sobolev spaces.

Theorem A.9 (Banach-Alaoglu). *If $E \to X$ is a vector bundle over a compact manifold, then $L^p_k(E)$ is weakly compact. That is to say, every bounded sequence σ^ν in $L^p_k(E)$ has a subsequence also denoted σ^ν and a corresponding section σ such that, for any continuous linear functional $\phi : L^p_k(E) \to \mathbb{R}$, the sequence $\phi(\sigma^\nu)$ converges to $\phi(\sigma)$.*

This theorem also holds true for orbifolds.

There is also a result on the regularity of gauge transformations that we will need.

Proposition A.10. *Let $P \to X$ be a principal bundle over a compact n-manifold X with compact structure group G and Lie algebra $\mathfrak{g}$, and let $k \in \mathbb{N}$ and $p \in [1, \infty)$ be numbers satisfying $kp > n$. Then a continuous gauge transformation $u : P \to G$ is L^p_k if and only if $u^{-1}\mathrm{d}u \in \Omega^1(\mathfrak{g})$ is L^p_{k-1}. Furthermore, a sequence of gauge transformations u^ν converges in L^p_k to a gauge transformation u^∞ if and only if $(u^\nu)^{-1}\mathrm{d}u^\nu$ converges in L^p_{k-1} to $(u^\infty)^{-1}\mathrm{d}u^\infty$.*

The proof is given in [Weh04]. Moreover, the result generalises directly to orbifolds as follows. Since the definition of L^p_k gauge transformations is purely local, it suffices to prove the following: given an open Euclidean domain $\widetilde{U}$ with a smooth action by a finite group Γ, and an extension of this action to $\widetilde{U} \times G$, a continuous Γ-equivariant gauge transformation $u : \widetilde{U} \times G \to G$ is L^p_k if and only if $u^{-1}\mathrm{d}u$ is L^p_{k-1}. This is a special case of the above proposition.

Finally, there are several estimates on the Sobolev norms of connections and gauge transformations on manifolds that we will need to prove Uhlenbeck's theorems. Once again, since the definition of L^p_k functions is entirely local and an object defined around a singular point in an orbifold is the same as a Γ-equivariant object defined in a neighbourhood of $\mathbb{R}^n$, the results transfer immediately to orbifolds.

Lemma A.11. *Let $k \in \mathbb{N}$ and $p \in [1, \infty)$ be such that $kp > n$, let U be a domain in $\mathbb{R}^n$, and let G be a Lie group with Lie algebra $\mathfrak{g}$. There is a constant $C > 0$ such that, for every $u \in \mathscr{G}^{k,p}(U)$ and every $A \in \mathscr{A}^{k-1,p}(U)$, the following inequality holds:*

$$\|u(A)\|_{L^p_{k-1}} \leq \|u^{-1}\mathrm{d}u\|_{L^p_{k-1}} + C\|A\|_{L^p_{k-1}}(1 + \|u^{-1}\mathrm{d}u\|_{L^{2p}_{k-2}})^{k-1}. \tag{A.5}$$

Lemma A.12. *Let U, G and $\mathfrak{g}$ be as above, let $k, \ell \in \mathbb{N}$ be integers such that $k > \ell$, and let $p \in [1, \infty)$ be such that $kp > n$ and $p \geq n/2$. There is a constant $C > 0$ such that for any $u \in \mathscr{G}^{k,p}$ and any $\tau \in \Omega^\bullet(U, \mathfrak{g})_{L^p_\ell}$, the following inequality holds:*

$$\|u^{-1}\tau u\|_{L^p_\ell} \leq C\|\tau\|_{L^p_\ell}(1 + \|u^{-1}\mathrm{d}u\|_{L^{2p}_{\ell-1}})^\ell. \tag{A.6}$$

Lemma A.13. *Let $k \in \mathbb{N}$ and $p \in [1, \infty)$ be such that $kp > n$ and $p > n/2$. Let $P \to X$ be a principal G-bundle over a compact manifold X, and let $\{U_\alpha\}_{\alpha \in I}$ be a bundle atlas. Let $A^\nu \in \mathscr{A}^{k-1,p}$ be a sequence of connections and $u^\nu \in \mathscr{G}^{k,p}$ a sequence of gauge transformations for which both $\|A^\nu\|_{L^p_{k-1}}$ and $\|u^\nu(A^\nu)\|_{L^p_{k-1}}$ are uniformly bounded, and denote by A^ν_α and u^ν_α the local representations of the connections and gauge transformations over U_α. Then $\|(u^\nu_\alpha)^{-1}\mathrm{d}u^\nu_\alpha\|_{L^p_{k-1}}$ is uniformly bounded over ν for all α, and there is a C^0-convergent subsequence of the u^ν.*

We will use these results in the following lemma:

Lemma A.14. *There exist constants $C, D > 0$ such that the following inequality holds for all $u, v \in \mathscr{G}^{k+1,p}$:*

$$\|(uv)^{-1}\mathrm{d}(uv)\|_{L^p_k} \leq C\|u^{-1}\mathrm{d}u\|_{L^p_k}(1 + D\|v^{-1}\mathrm{d}v\|_{L^p_k})^k + \|v^{-1}\mathrm{d}v\|_{L^p_k}. \tag{A.7}$$

Proof. First, observe that $(uv)^{-1}\mathrm{d}(uv) = v^{-1}(u^{-1}\mathrm{d}u)v + v^{-1}\mathrm{d}v$ by the product rule; thus, the L^p_k norm of the first term is bounded above by the sum of the norms of the other two terms. The second of these terms is already found in Equation A.7. For the first, we use Lemma A.12 to find a constant $C > 0$ for which

$$\|v^{-1}(u^{-1}\mathrm{d}u)v\|_{L^p_k} \leq C\|u^{-1}\mathrm{d}u\|_{L^p_k}(1 + \|v^{-1}\mathrm{d}v\|_{L^{2p}_{k-1}})^k, \tag{A.8}$$

and we then use the Sobolev embedding $L^p_k \hookrightarrow L^{2p}_{k-1}$ to find a constant $D > 0$ for which $\|v^{-1}\mathrm{d}v\|_{L^{2p}_{k-1}} \leq D\|v^{-1}\mathrm{d}v\|_{L^p_k}$. This completes the proof. □

A.3 Elliptic Operators

Finally, we will make use of the ellipticity of the Dirac operator, and of the Dolbeault operator on a Riemann surface. We briefly review elliptic operators here; for more details, refer to [Ram04].

Let $E, F \to X$ be vector bundles over a compact orbifold X with rank r and s respectively, and let $D : \Gamma(E) \to \Gamma(F)$ be a linear map. Then D is called a *differential operator of order k* if every open set U which trivialises E and F admits a series of maps $A_\alpha : U \to \mathrm{Hom}(\mathbb{R}^r, \mathbb{R}^s)$ for each multi-index α on $\mathbb{R}^r$, such that D is locally represented as follows:

$$D|_U = \sum_{|\alpha| \leq k} A_\alpha \frac{\partial^{|\alpha|}}{\partial x^\alpha}. \tag{A.9}$$

The *principal symbol of D* is the k-homogeneous map $\sigma(D) : T^*X \to \mathrm{Hom}(E, F)$ defined as follows in local coordinates:

$$[\sigma(D)](\omega) = \sum_{|\alpha| \leq k} A_\alpha \omega^\alpha, \tag{A.10}$$

where $\omega^\alpha = \omega_1^{\alpha_1} \cdots \omega_n^{\alpha_n}$. This definition is independent of coordinates. The operator D is called *elliptic* if the principal symbol is an isomorphism whenever $\omega \neq 0$. Note that the Dirac operator associated to a Clifford module structure $\rho : T^*X \to \mathfrak{su}(S)$ is first-order elliptic with principal symbol given by $\omega \mapsto \rho(\omega)$, and the Dolbeault operator on an orbifold Riemann surface Σ is first-order elliptic with principal symbol given by $\omega \mapsto (\omega_x + i\omega_y)\mathrm{d}\bar{z}$.

We will use the following two theorems on elliptic operators.

Theorem A.15 (Elliptic estimate)**.** *Let $D : \Gamma(E) \to \Gamma(F)$ be an elliptic operator of order k on bundles over a compact orbifold. Then there is some constant $C > 0$ such that the following holds for every $\sigma \in \Gamma(E)_{L^2_\ell}$:*

$$\|\sigma\|_{L^2_\ell} \leq C(\|D\sigma\|_{L^2_{\ell-k}} + \|\sigma\|_{L^2}). \tag{A.11}$$

Additionally, if $\sigma \perp \ker(D)$, then the L^2-term may be omitted.

Theorem A.16 (Elliptic regularity)**.** *If D is an elliptic operator of order k and $D\sigma$ is an L^2_ℓ section, then σ is an $L^2_{\ell+k}$ section. In particular, $\ker(D)$ consists of smooth sections.*

Appendix B

Spin Geometry Conventions

In setting up the Seiberg-Witten equations, many of the relevant spin-geometric objects can be chosen in different equivalent ways. As a result, different resources on Seiberg-Witten theory will give the equations in slightly different forms. In this Appendix, we briefly review the convention differences that exist in the literature, the effect that they have on the Seiberg-Witten equations, and what we have chosen in our exposition. For detailed definitions of the objects we describe in this Appendix, refer to Chapters 3 and 4.

B.1 Inner products

A complex inner product is defined to be a nondegenerate sesquilinear form; this means either the first or the second entry is conjugate linear. It is essentially consistent throughout the mathematical literature that the second component is conjugate linear (though not in the physics literature), so we have also adopted this convention.

B.2 The Clifford relation

The fundamental relation for the Clifford algebra over a vector space is either $v^2 = -g(v,v)$ or $v^2 = +g(v,v)$; accordingly, a Clifford module structure $\rho : T^*Y \to \mathfrak{su}(W)$ could satisfy either $\rho(v)^2 = -g(v,v)$ or $\rho(v)^2 = +g(v,v)$. We have chosen the former convention, which is consistent with the literature on mathematical gauge theory (but not the literature on algebra). One of the benefits of the former convention is that the associated Dirac operator is formally self-adjoint, whereas it is formally anti-self-adjoint for the latter convention.

B.3 The double covering of SO(n)

The group $\mathrm{Spin}(n) \leq \mathrm{Cl}(n)$ has a natural orthogonal action on $\mathbb{R}^n$ by conjugation. This action could be interpreted as $x \mapsto \gamma x \gamma^{-1}$ or $x \mapsto \gamma^{-1} x \gamma$. We have chosen the former convention, which is consistent with the literature; this convention has the benefit that the action of $\mathrm{Spin}(n)$ on $\mathbb{R}^n$ is a left action, just as the action of $\mathrm{SO}(n)$ by matrix multiplication is a left action.

B.4 Lie algebra isomorphisms

The Lie algebra $\mathfrak{spin}_n$ is canonically identified with the subspace of $\mathrm{Cl}(\mathbb{R}^n)$ consisting of bivectors, in the sense that the Lie group exponential (in terms of the tangent space at the identity) and the matrix exponential (in terms of the infinite series) coincide. In the same way, the Lie algebra $\mathfrak{so}_n$ is canonically identified with the subspace of $\mathrm{End}(\mathbb{R}^n)$ consisting of skew-symmetric endomorphisms. However, this can in turn be identified with $\Lambda^2(\mathbb{R}^n)$ when $v \wedge w$ is interpreted as an endomorphism $vw^* - wv^*$; this identification is canonical with respect to the standard metric on $\mathbb{R}^n$. The map $\mathrm{Spin}(n) \to \mathrm{SO}(n)$ induces an isomorphism of Lie algebras given by $e_i e_j \mapsto 2(e_i \wedge e_j)$. Additionally, the commutator of two vectors $[v, w]$ is in $\mathfrak{spin}_n$, and it naturally corresponds to the wedge product $4(v \wedge w)$.

Since $\mathrm{Spin}^c(n)$ doubly covers $\mathrm{SO}(n) \times \mathrm{U}(1)$, its Lie algebra is isomorphic to $\mathfrak{so}_n \oplus i\mathbb{R}$. This covering induces the isomorphism given by taking $e_i e_j \mapsto 2(e_i \wedge e_j)$ and $i \mapsto 2i$.

The above is consistent across the literature, but it has been reproduced here to emphasise the points at which choices are being made.

B.5 Clifford algebra representations and the volume element

When n is odd, there are exactly two irreducible representations of $\mathrm{Cl}^c(\mathbb{R}^n)$; when n is even, there is exactly one irreducible representation of $\mathrm{Cl}^c(\mathbb{R}^n)$, but it splits into two irreducible representations of $\mathrm{Spin}^c(n)$. In either case, the two representations are distinguished by the action of $\Lambda^n\mathbb{R}^n$, according to *chirality.* Specifically, a distinguished spanning element $\omega \in \Lambda^n\mathbb{R}^n$ called the volume element is chosen, and the chirality is positive (or negative) if ω acts by positive (or negative) multiplication.

The ambiguity in this definition is in the choice of the element ω. On an oriented real inner product space with some oriented basis $e_1, \dots, e_n$, there is a canonical choice of ω given by $e_1 \cdots e_n$. However, this element only squares to the identity when n is equal to 3 or 4 mod 8; for other values of n its square is negative, and this means that ω has eigenvalues $\pm i$ on a given representation. To fix this, ω is replaced by a *complex volume element* ω_c for which $\omega_c^2 = 1$, and which is conventionally defined as follows:

$$\omega_c := i^{\lfloor (n+1)/2 \rfloor} e_1 \cdots e_n. \tag{B.1}$$

In particular, for $n = 3, 4$, we have that $\omega_c = -\omega$; this still squares to 1, but the $+1$ and -1 eigenspaces are flipped.

In the literature, and in particular in [MOY96], the convention chosen for dimension 3 and 4 is to use the real volume element ω, not ω_c, to define positive and negative spinors. In particular, the positive representation is chosen to be that for which $e_0 e_1 e_2$ and $e_0 e_1 e_2 e_3$ are mapped to the identity element. However, in Section 5.4 of the paper, the volume element of the Riemann surface Σ is used to split the spinor bundle over Y; in this context, the complex volume element of Σ, namely $ie_1 e_2$, is used implicitly instead.

The choice to map $e_1 e_2 e_3$ to the identity results in the following identity for 1-forms $\alpha \in \Omega^1(Y)$:

$$\rho(*\alpha) = -\rho(\alpha). \tag{B.2}$$

Note that the Hodge star is defined to be complex-antilinear throughout this document where relevant.

Finally, we note that in odd dimensions, the negative-chirality representation can be achieved simply by negating the positive-chirality representation. This can serve as a useful sanity check for understanding which quantities do and do not depend on chirality.

B.6 The quadratic map

In both three and four dimensions, there is a quadratic map relating spinors to k-forms; it is valued in 1-forms for 3-manifolds and self-dual 2-forms for 4-manifolds. We will review the differences in conventions for the three-dimensional map; the conventions for four dimensions are very similar, but they are largely irrelevant for the development of the three-dimensional theory.

Depending on the application, the three-dimensional quadratic map $q : \Gamma(W) \to i\Omega^1(Y)$ can be defined in three ways; each of these definitions differ by a scalar multiple. They are as follows.

- As a pointwise "quadratic adjoint" to the action of Clifford multiplication by a spinor: for arbitrary α, we define q to satisfy

$$\langle \alpha, q(\psi) \rangle_{\Lambda^1(Y)} = C \langle \rho(\alpha) \cdot \psi, \psi \rangle_W. \tag{B.3}$$

  Here C is some complex constant, chosen by each respective author. Equivalently, it is defined in a basis $\{e_1, e_2, e_3\}$ as follows:

$$q(\psi) = C^* \sum_{j=1}^{3} \langle \psi, \rho(e_j) \psi \rangle e^j. \tag{B.4}$$

- As a pointwise Clifford inverse of the endomorphism associated to a spinor:

$$q(\psi) = \rho^{-1}(\psi \otimes \psi^*)_0. \tag{B.5}$$

- As a pointwise Clifford adjoint of the endomorphism associated to a spinor:

$$q(\psi) = \rho^\dagger(\psi \otimes \psi^*). \tag{B.6}$$

  This expression makes sense, since $\mathrm{End}(W)^\dagger \cong \mathrm{End}(W)$ and $\rho^\dagger : \mathrm{End}(W)^\dagger \to T^*Y$. Also, since $\mathrm{id}_W^* = \mathrm{tr}$ and $\rho(v)$ is always traceless, it is equivalent to take $(\psi \otimes \psi^*)_0$ as the argument of $\rho^\dagger$.

Using the first definition as a reference point, we state how the other two definitions are related:

- The second definition corresponds to the first by the choice $C = +1/2$.
- The third definition corresponds to the first by the choice $C = -1$. This also means that $\rho^\dagger = -2\rho^{-1}$.

The verifications of these facts reduce to computations in the chiral representation. Additionally, if the chirality of the representation is flipped, all three definitions of q flip sign; this simply follows from the fact that only one factor of ρ appears in each definition.

The most common definition in the modern literature appears to be the second; in fact, the issue of defining the quadratic map is often sidestepped by taking $\rho(*F_A) = -(\psi \otimes \psi^*)_0$ instead of $*F_A = q(\psi)$. In [MOY96], the quadratic map is defined in the first way with $C = -i/2$; this offers the advantage that the codomain of q can be taken to consist of purely real 1-forms, allowing the imaginary unit to be shown explicitly. In this instance we divert from the convention chosen in [MOY96], instead opting for the second convention.

B.7 Spinorial connections

A spinorial connection is a connection on a spinor bundle commuting with the Clifford action. While its curvature is an $\mathrm{End}(W)$-valued 2-form, the curvature in the Seiberg-Witten equations is an associated imaginary 2-form. There are two ways to understand this 2-form:

- It is the trace of the curvature 2-form with values in $\mathrm{End}(W)$. It can be shown using Schur's lemma that this trace must always be imaginary.
- It is the derivative of the associated U(1)-invariant 1-form on the determinant line bundle $L_{\det} = (P/\mathrm{Spin}(n)) \times_{\mathrm{U}(1)} \mathbb{C}$. Since this connection is unitary, its curvature 2-form is manifestly imaginary.

These two approaches differ by a factor of 2; specifically, the curvature of the connection on $L_{\det}$ is 1/2 the trace of the curvature 2-form on W. In [MOY96], the trace viewpoint is exclusively adopted; we will make use of the determinant line bundle interpretation, as it is the more modern approach in the literature. To be explicit, the symbol F_A will always refer to the imaginary-valued 2-form given by half the trace of F_{∇^A}, and F_{∇^A} will refer to the $\mathrm{End}(W)$-valued curvature 2-form of the connection ∇^A.

B.8 The Chern-Simons-Dirac functional

The Chern-Simons-Dirac functional has essentially one definition across the literature: it is a sum of a Chern-Simons functional and a Dirac Lagrangian functional. However, these two summands will often be modulated by constant factors. In general, the Chern-Simons-Dirac functional is always of the following form:

$$\mathrm{csd}(A, \varphi) = \int_Y -\frac{K_1}{2}(A - A_0)(F_A + F_{A_0}) + \frac{K_2}{2}\langle \psi, D_A \psi \rangle, \tag{B.7}$$

where $K_1, K_2 > 0$ are constants. In [KM07], the constants are taken to be $K_1 = 1/4$ and $K_2 = 1$. In our exposition, we take both K_1 and K_2 to be 1.

B.9 The Seiberg-Witten equations

The first five sections in this appendix describe conventions which are relatively consistent across the mathematical literature. The remaining three, however, vary considerably across resources. Each of these choices will result in a slightly different version of the Seiberg-Witten equation relating the curvature and the quadratic map (the harmonic equation is unaffected).

Nevertheless, it should be possible to write every version of the equation into the following form for some $M > 0$:

$$*F_A + M\rho^{-1}(\psi \otimes \psi^*)_0 = 0. \tag{B.8}$$

The constant M depends on which definition is used for the quadratic map (i.e., the value of C), the interpretation of F_A, and the definition of csd (in particular the value of K_1/K_2). In fact, when written in this form, one can prove the moduli space of irreducible Seiberg-Witten monopoles does not depend on the above choices; the solution (A, ψ) to the $M = 1$ version corresponds to the solution $(A, \sqrt{M_0}\psi)$ to the $M = M_0$ version of the equation. However, it is important to note that M must have no imaginary component and must be strictly positive in order for the solutions to be unaffected.

References

[AB83] M. Atiyah and R. Bott. "The Yang-Mills equations over Riemann surfaces". In: *Philos. Trans. R. Soc. Lond. A* 308.1505 (1983), pp. 523–615. DOI: 10.1098/rsta.1983.0017.

[AG71] N. L. Alling and N. Greenleaf. *Foundations of the Theory of Klein Surfaces.* 1st ed. Springer Berlin, 1971. DOI: 10.1007/BFb0060987.

[ALR07] A. Adem, J. Leida, and Y. Ruan. *Orbifolds and Stringy Topology.* Cambridge Tracts in Mathematics. Cambridge University Press, 2007. DOI: 10.1017/CBO9780511543081.

[Ark11] M. Arkowitz. *Introduction to homotopy theory.* Springer Science & Business Media, 2011. DOI: 10.1007/978-1-4419-7329-0.

[BH24] D. Baraglia and P. Hekmati. "Equivariant Seiberg–Witten–Floer cohomology". In: *Algebr. Geom. Topol.* 24.1 (2024), pp. 493–554. DOI: 10.2140/agt.2024.24.493.

[Bra90] S. B. Bradlow. "Vortices in holomorphic line bundles over closed Kähler manifolds". In: *Comm. Math. Phys.* 135.1 (1990), pp. 1–17. DOI: 10.1007/BF02097654.

[Bro82] K. S. Brown. *Cohomology of Groups.* Springer New York, 1982. DOI: 10.1007/978-1-4684-9327-6.

[BS85] F. Bonahon and L. Siebenmann. "The classification of Seifert fibred 3-orbifolds". In: *Low Dimensional Topology.* Ed. by R. Fenn. London Mathematical Society Lecture Note Series. Cambridge University Press, 1985, pp. 19–85.

[BW93] B. Booß-Bavnbek and K. P. Wojciechowski. *Elliptic Boundary Problems for Dirac Operators.* Boston, MA: Birkhäuser Boston, 1993. DOI: 10.1007/978-1-4612-0337-7_8.

[CHK00] D. Cooper, C. D. Hodgson, and S. P. Kerckhoff. *Three-dimensional Orbifolds and Cone-Manifolds.* The Mathematical Society of Japan, 2000.

[CR01] W. Chen and Y. Ruan. "Orbifold gromov-witten theory". In: *arXiv preprint math/0103156* (2001). URL: 10.48550/arXiv.math/0103156.

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

[Don83] S. K. Donaldson. "A new proof of a theorem of Narasimhan and Seshadri". In: *J. Differ. Geom.* 18.2 (1983), pp. 269–277. DOI: 10.4310/jdg/1214437664.

[Don96] S. K. Donaldson. "The Seiberg-Witten equations and 4-manifold topology". In: *Bull. Am. Math. Soc.* 33.1 (1996), pp. 45–70. DOI: 10.1090/S0273-0979-96-00625-8.

[Flo88a] A. Floer. "An instanton-invariant for 3-manifolds". In: *Comm. Math. Phys.* 118.2 (1988), pp. 215–240. DOI: 10.1007/BF01218578.

[Flo88b] A. Floer. "Morse theory for Lagrangian intersections". In: *J. Differ. Geom.* 28.3 (Jan. 1988), pp. 513–547. DOI: 10.4310/jdg/1214442477.

[FS92] M. Furuta and B. Steer. "Seifert fibred homology 3-spheres and the Yang-Mills equations on Riemann surfaces with marked points". In: *Adv. Math.* 96.1 (1992), pp. 38–102. DOI: 10.1016/0001-8708(92)90051-L.

[FWX25] Yu Feng, Shuo Wang, and Bin Xu. *A note On the existence of solutions to Hitchin's self-duality equations.* 2025. DOI: 10.48550/arXiv.2501.10976. arXiv: 2501.10976 [math.DG].

[Gar11] D. J. H. Garling. *Clifford Algebras: An Introduction.* Cambridge University Press, 2011. DOI: 10.1017/CBO9780511972997.

[Gar91] O. Garcia-Prada. "The Geometry of the Vortex Equation". PhD thesis. Wolfson College, Oxford, 1991.

[Gar94] O. García-Prada. "A direct existence proof for the vortex equations over a compact Riemann surface". In: *Bull. Lond. Math. Soc.* 26.1 (1994), pp. 88–96. DOI: 10.1112/blms/26.1.88.

[GHL90] S. Gallot, D. Hulin, and J. Lafontaine. *Riemannian geometry.* Vol. 2. Springer Berlin, 1990. DOI: 10.1007/978-3-642-18855-8.

[GL18] H. Geiges and C. Lange. "Seifert fibrations of lens spaces". In: *Abh. Math. Semin. Univ. Hamb.* 88.1 (2018), pp. 1–22. DOI: 10.1007/s12188-017-0188-z.

[GM80] M. Goresky and R. MacPherson. "Intersection homology theory". In: *Topology* 19.2 (1980), pp. 135–162. DOI: 10.1016/0040-9383(80)90003-8.

[Gor99] C. M. A. Gordon. "3-Dimensional Topology up to 1960". In: *History of Topology.* Ed. by I. M. James. North Holland, 1999, pp. 449–489. DOI: 10.1016/B978-044482375-5/50016-X.

[Hat01] A. E. Hatcher. *Notes on Basic 3-Manifold Topology.* 2001. URL: https://api.semanticscholar.org/CorpusID:9792594.

[Hit87] N. J. Hitchin. "The Self-Duality Equations on a Riemann Surface". In: *Proc. London Math. Soc.* 55.3 (1987), pp. 59–126. DOI: https://doi.org/10.1112/plms/s3-55.1.59.

[HLR08] N. Ho, C. M. Liu, and D. A. Ramras. "Orientability in Yang–Mills theory over nonorientable surfaces". In: *Comm. Anal. Geom.* 17.5 (2008), pp. 903–953. DOI: 10.4310/CAG.2009.v17.n5.a3.

[HT97] M. Hutchings and C. H. Taubes. "An introduction to the Seiberg-Witten equations on symplectic manifolds". In: 1997. URL: https://api.semanticscholar.org/CorpusID:39656956.

[Hus94] D. Husemöller. *Fibre Bundles.* Graduate Texts in Mathematics. Springer New York, 1994. DOI: 10.1007/978-1-4757-2261-1.

[Huy05] D. Huybrechts. *Complex geometry: an introduction.* Springer Berlin, 2005. DOI: 10.1017/S0025557200181379.

[Jr22] F. C. Caramello Jr. *Introduction to orbifolds.* 2022. DOI: 10.48550/arXiv.1909.08699. arXiv: 1909.08699 [math.DG].

[KM] J. Kalliongis and A. Miller. "Geometric Group Actions on Lens Spaces". In: *Kyungpook Mathematical Journal* 42.2 ().

[KM07] P. Kronheimer and T. Mrowka. *Monopoles and Three-Manifolds*. New Mathematical Monographs. Cambridge University Press, 2007. DOI: 10.1017/CBO9780511543111.

[KM97] P. Kronheimer and T. Mrowka. "The Genus Of Embedded Surfaces In The Projective Plane". In: *Math. Res. Lett.* 1.6 (1997), pp. 797–808. DOI: 10.4310/MRL.1994.v1.n6.a14.

[Lan12] S. Lang. *Algebra*. Vol. 211. Springer Science & Business Media, 2012. DOI: 10.1007/978-1-4613-0041-0.

[Lan20] C. Lange. "Orbifolds from a metric viewpoint". In: *Geom. Dedicata* 209.1 (2020), pp. 43–57. DOI: 10.1007/s10711-020-00521-x.

[Lan93] S. Lang. *Real and Functional Analysis*. 3rd ed. Springer New York, 1993. DOI: 10.1007/978-1-4612-0897-6.

[Las82] R. K. Lashof. "Equivariant bundles". In: *Illinois J. Math.* 26.2 (1982), pp. 257–271. DOI: 10.1215/ijm/1256046796.

[Lew93] J. Lewandowski. "Group of loops, holonomy maps, path bundle and path connection". In: *Class. Quant. Grav.* 10 (1993), pp. 879–904. DOI: 10.1088/0264-9381/10/5/008.

[Lin16] F. Lin. "Monopoles and Pin(2)-symmetry". PhD thesis. University of Cambridge, Jan. 2016.

[LM18] T. Lidman and C. Manolescu. "Floer homology and covering spaces". In: *Geom. & Topol.* 22.5 (2018), pp. 2817–2838. DOI: 10.2140/gt.2018.22.2817.

[LM89] H. B. Lawson and M. Michelson. *Spin Geometry (PMS-38)*. Princeton University Press, 1989. DOI: 10.1515/9781400883912.

[LS13] C. M. Liu and F. Schaffhauser. "The Yang–Mills equations over Klein surfaces". In: *J. Topol.* 6.3 (2013), pp. 569–643. DOI: 10.1112/jtopol/jtt001.

[Man01] C. Manolescu. "Seiberg-Witten-Floer stable homotopy type of three-manifolds with b_1=0". In: *Geom. Topol.* 7.2 (2001), pp. 889–932. DOI: 10.2140/gt.2003.7.889.

[Man15] C. Manolescu. "Pin(2)-Equivariant Seiberg-Witten Floer homology and the Triangulation Conjecture". In: *J. Amer. Math. Soc.* 29 (2015), pp. 147–176. DOI: doi.org/10.1090/jams829.

[Mar+07] J. Marsden et al. *Hamiltonian Reduction by Stages*. Lecture Notes in Mathematics 1913. Springer Berlin, 2007. DOI: 10.1007/978-3-540-72470-4.

[Mor07] A. Moroianu. *Lectures on Kähler geometry*. Vol. 69. Cambridge University Press, 2007. DOI: 10.1017/CBO9780511618666.

[MOY96] T. Mrowka, P. Ozsvath, and B. Yu. "Seiberg-Witten Monopoles on Seifert Fibered Spaces". In: *Commun. Anal. Geom.* 5.4 (1996), pp. 685–791. DOI: 10.4310/CAG.1997.v5.n4.a3.

[MP97] I. Moerdijk and D. A. Pronk. "Orbifolds, Sheaves and Groupoids". In: *Indag. Mathem.* 12 (1997), pp. 3–21. DOI: 10.1023/A:1007767628271.

[MSS65] J. Milnor, L. Siebenmann, and J. Sondow. *Lectures on the h-Cobordism Theorem.* Princeton University Press, 1965. DOI: 10.1017/S0013091500012050.

[MST96] J. W. Morgan, Z. Szabó, and C. H. Taubes. "A product formula for the Seiberg-Witten invariants and the generalized Thom conjecture". In: *J. Differ. Geom.* 44.4 (1996), pp. 706–788. DOI: 10.4310/jdg/1214459408.

[Nat90] S. M. Natanzon. "Klein surfaces". In: *Russian Math. Surveys* 45.6 (1990), pp. 53–108. DOI: 10.1070/RM1990v045n06ABEH002713.

[Nic96] L. I. Nicolaescu. "Adiabatic limits of Seiberg-Witten equations on Seifert manifolds". In: *Comm. Anal. Geom.* 6.2 (1996), pp. 331–392.

[Ram04] S. Ramanan. *Global Calculus.* Vol. 65. Graduate Studies in Mathematics. Rhode Island: American Mathematical Society, 2004. DOI: 10.1090/gsm/065.

[Sal14] D. A. Salamon. "SPIN GEOMETRY AND SEIBERG-WITTEN INVARIANTS". In: 2014. URL: https://api.semanticscholar.org/CorpusID:116235731.

[Sat56] I. Satake. "On a Generalization of the Notion of Manifold". In: *Proc. Natl. Acad. Sci. U.S.A.* 42.6 (1956), pp. 359–363. DOI: 10.1073/pnas.42.6.359.

[Sat57] I. Satake. "The Gauss-Bonnet Theorem for V-manifolds". In: *J. Math. Soc. Japan* 9.4 (1957), pp. 464–492. DOI: 10.2969/jmsj/00940464.

[Sch16a] F. Schaffhauser. "Lectures on Klein Surfaces and Their Fundamental Group". In: *Geometry and Quantization of Moduli Spaces.* Ed. by L. A. Consul, J. E. Andersen, and I. M. i Riera. Springer International Publishing, 2016, pp. 67–108. DOI: 10.1007/978-3-319-33578-0_2.

[Sch16b] F. Schaffhauser. "Lectures on Klein surfaces and their fundamental group". In: *Geometry and quantization of moduli spaces.* Birkhäuser, Cham, 2016, pp. 67–108. DOI: 10.1007/978-3-319-33578-0_2.

[Sch17] F. Schaffhauser. "On the Narasimhan–Seshadri correspondence for real and quaternionic vector bundles". In: *J. Differ. Geom.* 105.1 (2017), pp. 119–162. DOI: 10.4310/jdg/1483655861.

[Sco83] P. Scott. "The Geometries of 3-Manifolds". In: *Bull. Lond. Math. Soc.* 15.5 (1983), pp. 401–487. DOI: 10.1112/blms/15.5.401.

[Thu02] W. P. Thurston. *The Geometry and Topology of Three-Manifolds.* 2002. URL: https://library.slmath.org/books/gt3m/.

[Uhl82] K. Uhlenbeck. "Connections with Lp bounds on curvature". In: *Comm. Math. Phys.* 83.1 (1982), p. 42. DOI: 10.1007/BF01947069.

[Weh04] K. Wehrheim. *Uhlenbeck Compactness.* European Mathematical Society, 2004. DOI: 10.4171/004.

[Wer19] K. Wernli. *Lecture notes on spin geometry.* 2019. DOI: 10.48550/arXiv.1911.09766.

[Wit94] E. Witten. "Monopoles and four manifolds". In: *Math. Res. Lett.* 1 (1994), pp. 769–796. DOI: 10.4310/MRL.1994.v1.n6.a13.