The Kapustin-Witten Equations on Cylindrical Manifolds

Alfred Holmes

University College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

July 2025

Abstract

This thesis contains two parts.

In Part 1 we introduce a perturbed complex Chern-Simons functional, $\mathfrak{C}_\eta$, on the space of $SL(2;\mathbb{C})$ connections on a three manifold and study the critical points. We explain how picking a particular perturbations lead to a discrete set of critical points. This discrete set of points admits a compactification coming from Culler-Shalen theory. The perturbation can also be used to study links, and we give some example critical points of $\mathfrak{C}$ that relate to certain knots and links. We then study the gradient flow of $\mathrm{Re}(e^{i\theta}\mathfrak{C}_\eta)$ and consider compactness and exponential decay properties of solutions under finite energy assumptions. Solutions to this gradient flow are solutions to the Kapustin-Witten equations on the cylinders $I \times Y^3$. We explain how the functionals $\mathfrak{C}_\eta$ can be used in proposals of new three manifold invariants by Haydys and Witten.

In Part 2 we study solutions to the unperturbed Kapustin-Witten equations that admit Nahm pole singularities. We prove that if a solution to the gradient flow becomes singular then it admits a Nahm pole singularity and satisfies the conditions outlined in a proposal by Witten.

Contents

Notation Conventions

Throughout this thesis we will, unless stated otherwise, stick to the following notation conventions.

- The letters X and Y will always denote smooth manifolds. Most of the time the manifold Y will be closed (compact without boundary) but we will sometimes consider Y with boundary. If we are considering a four manifold, we will denote it by X. Y will be used to denote three manifolds. The four manifolds X will typically, although not always, be of the form $X = I \times Y$, for an interval $I \subseteq \mathbb{R}$. If the primary application of a result will be in a certain dimensional setting we will use the corresponding letter in the statement of the result.

- Principal G–bundles, for G compact, will be denoted by $P \to X$, where X is the base manifold. For a compact lie group, G, with complexification $G_{\mathbb{C}}$, we will use $P_{\mathbb{C}}$ to denote a principal $G_{\mathbb{C}}$ bundle. We will often consider $P_{\mathbb{C}}$ as the complexification of a bundle P, where $P \subseteq P_{\mathbb{C}}$ is a G reduction.

- For a principal bundle $P \to X$, the bundle associated to the adjoint action on the Lie algebra will be denoted by $\mathfrak{g}_P$. Hence the bundle $\mathfrak{g}_{P_{\mathbb{C}}} \cong \mathfrak{g}_P \otimes \mathbb{C}$. The bundle $AdP \to X$ denotes the bundle associated to the adjoint action of G on G, where G is the structure group of P.

- Let $E \to X$ be a smooth fibre bundle over X. We will denote the space of smooth sections of E by $\Gamma(E)$. This is the space of maps $X \to E$ such that they are the identity when composed with the bundle projection. We will denote by $\Omega^p(E)$ the space of E–valued p forms. That is

$$\Omega^p(E) \cong \Gamma(E \otimes \Lambda^p T^* V)$$

- We will use capital fraktur letters, for example $\mathfrak{A}$, $\mathfrak{B}$, to denote connections on bundles with complex structure group and Latin letters, for example A, B to denote connections on bundles with a compact structure group. We will use Greek letters, typically ϕ, to denote imaginary parts of complex connections (see Definition 3.7). Hence, when considering connections on a principal $G_{\mathbb{C}}$–bundle $P_{\mathbb{C}} \to X$, we will often write $\mathfrak{A} = A + i\phi$ where A is a connection on a G reduction of $P_{\mathbb{C}}$ and $\phi \in \Omega^1(\mathfrak{g}_P)$

- We will want to consider completions of the spaces of smooth sections with respect to certain norms. These norms will be tend to be standard norms, for example L^2 or $C^{0,\alpha}$. Given a bundle with an appropriate structure, for example, a normed vector bundle, we will use the following convention to denote spaces of completions. For example $L^2(\Gamma(E))$ will denote the L^2 completion of $\Gamma(E)$, and $C^{0,\alpha}(\Gamma(E))$ will denote the Hölder space with exponent α taking values in E.

- If we have a section $\phi \in \Omega^1(\mathfrak{g}_P)$ of the adjoint bundle, then we define

$$\phi \wedge \phi = \frac{1}{2}[\phi, \phi] = \frac{1}{2}\sum_{i,j}[\phi_i, \phi_j] dx^i \wedge dx^j$$

where $\phi = \phi_i dx^i$ in local coordinates.

- When considering principal $G_{\mathbb{C}}$–bundles we will assume that the bundle comes with a given choice of faithful representation

$$G_{\mathbb{C}} \to GL(n; \mathbb{C}).$$

  In particular, it will make sense to talk of the trace and matrix products of elements of the adjoint bundle.

- When considering the Hodge star, $*$, of forms valued in complex vector bundles, we will always mean the anti-complex linear extension of $*$.

- When considering $S^1 = \{e^{i2\pi t}\}$ we will use $t \in [0, 1]$ to denote coordinates on S^1.

Chapter 1

Introduction and Background

Let $P_{\mathbb{C}} \to Y$ be a principal $G_{\mathbb{C}}$ bundle on a compact manifold Y, where $G_{\mathbb{C}}$ is the complexification of a compact group G. Denote by $\mathcal{A}_{\mathbb{C}}$ the space of connections on $P_{\mathbb{C}}$ and pick a flat background connection $\mathfrak{B} \in \mathcal{A}_{\mathbb{C}}$. The main object of study in this thesis is a complexified version of the Chern-Simons functional,

$$\mathfrak{C} : \mathcal{A}_{\mathbb{C}} \to \mathbb{C}, \quad \mathfrak{A} \mapsto \int_Y \operatorname{tr}\left(a \wedge d_{\mathfrak{B}} a + \frac{2}{3} a \wedge a \wedge a\right), \tag{1.1}$$

where $a = \mathfrak{A} - \mathfrak{B} \in \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}})$. The Chern-Simons functional, when working with a compact structure group, is a key tool in the gauge theoretic study of three-manifolds. In particular, its critical points correspond to flat G connections, which can be counted to recover the Casson invariant of the three-manifold, [Tau90], and its gradient flow (with respect to the L^2 metric on $T\mathcal{A} = \Omega^1(\mathfrak{g}_P)$) is used to define Floer's instanton homology groups, [Flo88].

The motivation for studying the complex version of the Chern-Simons functional is to use it to construct new invariants of three-manifolds, and to give gauge theoretic interpretations of known knot invariants.

Formally it is quite straightforward to replace a compact group G with its complexification $G_{\mathbb{C}}$ when considering the setup of various gauge theories, but one quickly runs into issues. An obvious first problem is the non-compactness of the groups $G_{\mathbb{C}}$ (see Proposition 3.3). This non-compactness implies that for any relevant vector bundle in the theory, there is no choice of inner product invariant under $G_{\mathbb{C}}$ gauge transformations. This means that one cannot define norms on the relevant spaces of sections required for the linear analysis. As we discuss in Chapter 3, to solve these issues one can restrict the bundle automorphisms to those coming from G valued gauge transformations, where $G \subseteq G_{\mathbb{C}}$ is a maximal compact subgroup.

When the structure group has been reduced to G, it makes sense to talk of the real and imaginary part of a connection $\mathfrak{A} \in \mathcal{A}_{\mathbb{C}}$ – see Definition 3.7. That is, we can write

$$\mathfrak{A} = A + i\phi$$

for $\phi \in \Omega^1(\mathfrak{g}_P)$, where $P \subseteq P_{\mathbb{C}}$ is the relevant G reduction, and A is a connection on P. If we consider the $\mathcal{G}_{\mathbb{C}}$ orbit of a connection $\mathfrak{A} \in \mathcal{A}_{\mathbb{C}}$ then this is the disjoint union of $\mathcal{G}$ orbits, as $\mathcal{G} \subseteq \mathcal{G}_{\mathbb{C}}$. Since we have removed gauge freedom by reducing the structure group, in order to study the space $\mathcal{A}_{\mathbb{C}}/\mathcal{G}_{\mathbb{C}}$ of connections modulo the complex gauge group, we should pick a representative $\mathcal{G}$ orbit of each $\mathcal{G}_{\mathbb{C}}$ orbit. The L^2 norm of ϕ, with respect to a $\mathcal{G}$ invariant inner product, is fixed on $\mathcal{G}$ orbits. A natural choice of $\mathcal{G}$ gauge equivalence classes to choose are the ones which minimise the norm

$$\|\phi\|_{L^2}.$$

From an analytic point of view, this is the same as imposing the extra condition $d_A^*\phi = 0$, see Lemma 3.19. We call a complex connection $\mathfrak{A}$ stable if $d_A^*\phi$ is satisfied by some connection in its $\mathcal{G}_{\mathbb{C}}$ orbit. Not all elements of $\mathcal{A}_{\mathbb{C}}$ are gauge equivalent to a connection with $d_A^*\phi = 0$. For the case of flat $G_{\mathbb{C}}$ connections, whether a connection admits such a representative was solved by Corlette [Cor88], building on the work of Donaldson, [Don85], see Section 3.3.

With the Chern-Simons functional in particular, another issue that arises from replacing G with $G_{\mathbb{C}}$ comes from the properties of flat $G_{\mathbb{C}}$ connections. Flat connections are the critical points of the Chern-Simons functional, see Proposition 5.6. It is a standard result that flat connections are in one to one correspondence with the space of representations of the fundamental group modulo conjugation,

$$\mathrm{Hom}(\pi_1(Y); G_{\mathbb{C}})/G_{\mathbb{C}},$$

see Section 3.5.2. Taken as read, the space of flat $G_{\mathbb{C}}$ connections will typically be non-Hausdorff. This is due to non-reductive representations, which can be conjugated to a representation arbitrarily close to being reductive (not stable in the sense of Corlette), see Lemma 8.2. This issue can be solved by restricting to stable flat connections. Even when restricting to stable connections, however, many integer homology three-spheres have character varieties with positive dimensional components. These positive dimensional components mean that care needs to be taken when studying the critical points of the Chern-Simons functional. The critical points will not be isolated in general and instead be described by certain non-compact algebraic varieties coming from the representations of the fundamental group of the manifold.

A flat connection can be in a positive dimensional family if and only if it is degenerate, see Lemma 3.63. In three dimensions, due to Poincaré duality. Our technique for handling degenerate connections is to introduce a holonomy perturbation

$$\eta : \mathcal{A}_{\mathbb{C}} \to \mathbb{C}.$$

This perturbation is a generalisation of the perturbation used in Floer theory to remove non-degenerate flat connections. Floer was able to give a topological description of the critical points of the perturbed Chern-Simons functional, see [BD95]. These critical points can be described in terms of flat connections on a manifold with boundary, and the perturbation imposes constraints on the holonomy of the connection when restricted to the boundary. We generalise this description of the critical points to the complex case and it to prove that the critical points are isolated.

The structure of this thesis is as follows. In this introductory chapter, we give an overview instanton Floer homology and Casson invariants, Picard-Lefschetz theory and the Kapustin-Haydys-Witten equations, and then discuss applications of gauge theory to knot theory.

In the first part of the thesis we study perturbing the complexified Chern-Simons functional. This is done by considering so called holonomy perturbations. These involve integrating connection one forms around closed curves in the manifold. When working with connections defined in Sobolev completions we need a general codimension trace theorem. We prove such a result, Corollary 2.6, in Chapter 2 and also give some relevant results from Sobolev theory that we will need in later chapters to prove that certain operators are bounded. In Chapter 3 we study principal $G_{\mathbb{C}}$–bundles with a choice of G–reduction, which we call complexified principal bundles. In Chapter 4 we study harmonic maps on manifolds with boundary. This allows us to generalise the Corlette-Donaldson theorem to this setting. This is an important result to have in Chapter 5 when we consider critical points of the perturbed Chern-Simons functional that satisfy $d_A^* \phi = 0$. To study these critical points, the idea is to split the manifold into two components. When restricted to one of the components, away from the support of the perturbation, the critical points are flat $G_{\mathbb{C}}$ connections satisfying $d_A^* \phi = 0$, but this component has boundary. On the other component the connection reduces to an

Abelian connection which are easier to handle analytically. Generalising the result of Floer to the $G_{\mathbb{C}}$ case, we are able to give a geometric interpretation of solutions to the perturbed equations in terms of flat connections satisfying certain holonomy conditions on the complement of the perturbation. We finish the chapter on the complexified Chern-Simons functional by deriving perturbed versions of the Kapustin-Witten equations on cylindrical manifolds.

Once we have defined the perturbations we can then apply them to three-manifold invariants. In Chapter 6 we study flat $G_{\mathbb{C}}$ connections on Heegaard decompositions. In Chapter 8 we consider example critical points when the perturbations are supported on neighbourhoods of certain knots and links in S^3.

The last three chapters of part one consider the analytic setup for the Kapustin-Witten equations on cylindrical four-manifolds. The ideas here are based off those given in [Flo88] and [Don02] for compact structure groups, generalised to the complexified case. The linear analysis in Chapter 9 relies heavily on the ideas discussed in Chapter 3. The main results that we prove are compactness and exponential decay results for solutions to the perturbed Kapustin-Witten equations. These are key results in the study of the moduli spaces of solutions to the Kapustin-Witten equations on cylindrical manifolds.

Part two takes on a rather different character. Here we study the alternative approach to three-manifold invariants using the Kapustin-Witten equations proposed by Witten in [Wit11]. These invariants are defined by solutions admitting a certain type of singularity known as a Nahm pole. We prove that under certain assumptions on the real part of the curvature, the only singularities that can occur are Nahm pole singularities. This is Theorem 12.5. Such solutions occur in Witten's proposal for computing Khavonov Homology using gauge theory. The motivation for studying these boundary conditions is to try and prove existence of solutions satisfying the Nahm pole boundary condition, and to gain an understanding of the boundary conditions suggested by Witten.

1.1 The Casson Invariant and Floer Homology

In 1985 Casson introduced an invariant for homology three-spheres, see, for example, [AM90]. Let Y be an integer homology three-sphere. The Casson invariant is a signed count of representations

$$\rho : \pi_1(Y) \to SU(2).$$

modulo conjugation, defined by picking a Heegaard splitting of Y and counting the intersections of the representation varieties defined by each of the elements of the splitting. In 1990 Taubes proved that one could recover the Casson invariant as a signed count of flat $SU(2)$ connections modulo gauge transformations on principal $SU(2)$ bundles over Y, [Tau90]. It is a standard result that flat connections correspond to representations of $\pi_1(Y)$. Hence the content of [Tau90] is to show that one can orient the moduli space of flat connections in a way that recovers the orientation used to define the Casson invariant, as well as showing that there are no issues coming from singular points in the space of connections. To do this, Taubes uses a perturbed version of the Chern-Simons functional, orienting the moduli space using the spectral flow of the Hessian of the perturbed Chern-Simons functional on a path of connections from the trivial connection to the critical point.

The perturbation used by Taubes is essentially the same as the perturbation used by Floer in the construction of instanton Floer homology, [Flo88]. An outline of the perturbation is as follows. One first picks a finite collection of disjoint, embedded, loops $\gamma_i : S^1 \to Y$ along with disjoint neighbourhoods, $N(\gamma_i)$ with torus boundary for each loop. Since these neighbourhoods are solid tori, one can parametrise their points by picking a point $z \in D^2$ and a point on S^1. Hence $N(\gamma_i)$ defines an embedding

$$\Gamma_i : S^1 \times D^2 \to Y.$$

For a given connection, one can then consider

$$\eta_i : A \mapsto \int_{D^2} f_i(\mathrm{tr}(H_A))d\mu$$

where μ is a compactly supported two form on D^2, with integral 1, $f_i : \mathbb{R} \to \mathbb{R}$ is a smooth function, and $H_A : D^2 \to SU(2)$ is the holonomy around the curve given by the image of $S^1 \times \{z\}$ under Γ_i. The full perturbation to the Chern-Simons functional is then given by

$$\eta(A) = \sum_i \eta_i(A).$$

Floer Homology categorifies the Casson invariant, in the sense that it defines a cohomology theory whose Euler characteristic is the Casson invariant.

1.1.1 $SL(2;\mathbb{C})$ Casson invariants

The Casson invariant has been generalised (in different ways) to $SL(2;\mathbb{C})$, see [Cur01; AM19]. The construction of this invariant in [Cur01] is very similar to $SU(2)$ invariant except that one has to restrict the count to connections that lie in a suitable subset of the representations. These representations are the ones that lie in a 0 dimensional irreducible component of the moduli space of flat connections. Since the $SL(2;\mathbb{C})$ character variety is indeed an algebraic variety, it decomposes into a finite collection of irreducible subvarieties. Some of these will be 0 dimensional, and so there will be a compact set of suitable representations to count.

The construction in [AM19] is more involved. A Heegaard decomposition of gives two Lagrangians in the $SL(2;\mathbb{C})$ character variety of the boundary of the handlebodies. They then consider the Lagrangian Floer homology in these two Lagrangians to define a homology groups associated to the Heegaard decomposition. They prove that for a given three-manifold these homology groups are independent of the choice of Heegaard decomposition. The Euler characteristic of these homology groups can be used to define the Casson invariant. The Casson invariants from [Cur01] and [AM19] are different, see [AM19, discussion after Proposition 8.4].

In [Muñ25], for the special case of Seifert-Fibred integer homology three-spheres, it is shown that one can recover the Casson invariant defined in [AM19] from a count of perturbed flat $SL(2;\mathbb{C})$ connections. The count is carried out by considering the critical points of a perturbation to the Chern-Simons functional given by

$$S^{\varepsilon}_{\theta} : A + i\phi \mapsto \mathrm{Re}(e^{i\theta}\mathfrak{C}(A + i\phi)) + \varepsilon\|\phi\|^2_{L^2} + \varepsilon^2\eta,$$

where η is a holonomy perturbation.

The procedure for carrying out these counts is quite involved. This is because there can be infinitely many critical points for any $\varepsilon > 0$. To select a subset to count, Mũnoz-Enchániz proves that as $\varepsilon \to 0$ critical points either localise around a compact set or are such that $\|\phi\|_{L^\infty}$ is large. The precise formulation of this is that there exists a compact, zero dimensional, subset of irreducible flat connections, $\mathcal{L}$, such that for sufficiently large constant $C > 0$ and open set $\mathcal{U} \supseteq \mathcal{L}$, where $\mathcal{U} \subseteq \{[A + i\phi] | \|\phi\|_{L^\infty} < C\}$, it can be arranged (by making ε sufficiently small) that any critical point of S^{θ}_{ε} is either contained in $\mathcal{U}$ or $\|\phi\|_{L^\infty} > C$.

With our perturbation, in general we get infinitely many discrete critical points. One could hope to carry out some sort of localisation procedure as described but we have been unable to do so.

1.2 Picard-Lefschetz Theory

Picard-Lefschetz theory, see, for example, [Nic11, Chapter 5], is an analogue of Morse theory for symplectic manifolds. With many caveats, instanton Floer homology is the study of Morse theory for the Chern-Simons functional on the space of connections (modulo gauge) on a three-manifold. The resulting instanton Floer homology groups are essentially the Morse homology groups defined using the Chern-Simons functional. For the complexified Chern-Simons theory, it has been proposed, see [Hay15] and [Wit10], that one can use Picard-Lefschetz theory to study the topology of the space of complex connections on a three-manifolds to construct invariants.

Definition 1.1. A Lefschetz fibration is a J–holomorphic map

$$f : Z \to \mathbb{C}P^1$$

where Z and is an exact symplectic manifold, such that f is a submersion away from finitely many points critical points $\{p_1, \dots, p_n\}$, and at these critical points the Hessian of f is non-degenerate. For a given Lefschetz fibration, it is useful to cut out the singular fibres, so we define $\mathbb{C}P^1_0 = \mathbb{C}P^1 \setminus \{p_1, \dots, p_n\}$, and $Z_0 = f^{-1}(\mathbb{C}P^1_0)$. In this section we describe a diffeogeometric approach to the Fukaya-Siedel Category of a Lefschetz fibration due to Haydys, andonsider the application of the theory to the complex Chern-Simons functional. This results in the derivation of the Haydys-Witten equations and the Kapustin-Witten equations.

In [Wit10] Witten uses ideas from Picard-Lefschetz theory to define the evaluation of the complexified Chern-Simons path integral. Witten's idea is that the integration of certain holomorphic functions on a dimension n submanifold of a $2n$ dimensional complex manifold can be evaluated in terms of the homology of the domain of integration – the prototypical example being contour integration on $\mathbb{C}$. Witten uses these ideas to suggest ways of calculating the partition of function of the physical theory defined with the Chern-Simons action in an analytically continued setting. This has lead to his conjectures on the Jones polynomial and to his definition of the Haydys-Witten equations.

Haydys' approach to the Haydys-Witten equations is more explicitly related to Picard-Lefschetz theory. It turns out that the homology theory defined by a

Lefschetz fibration can be categorified, see [Sei08]. We call the resulting category the Fukaya-Seidel category of the Lefschetz fibration. In [Hay15], Haydys gives a differential geometric description of the construction of the Fukaya-Seidel category using the gradient flow of the real functions

$$\Phi_\theta : Z \to \mathbb{C}, \quad \Phi_\theta = \mathrm{Re}(e^{i\theta} f).$$

This proposal was studied further in [Wan22]. Haydys suggests that these same ideas applied in the infinite dimensional setting of the complex Chern-Simons functional on the space of $SL(2;\mathbb{C})$ connections on a three-manifold should allow one to construct a categorial invariant of the three-manifold – to the three-manifold, one associates the Fukaya-Seidel category defined using the three-manifold. We now briefly describe the ingredients that go into Haydys' construction of the Fukaya-Seidel category.

The Fukaya-Seidel category is an A_∞ category, and so as well as having morphisms between objects there are also n–morphisms for each $n \in \mathbb{Z}_{>0}$ which are morphisms between morphisms. The objects in the Fukaya-Seidel category are chain complexes that generate certain Lagrangian Floer groups. The idea is as follows. One picks a point $z \in \mathbb{C}$ that is not the image of a critical point of f. For each critical point, p, of f, one can then consider the manifolds

$$S_p \text{ and } W_p$$

of, respectively, stable and unstable critical points of $\mathrm{Re}(e^{i\theta} f)$. For the stable points we consider $\theta = \mathrm{Arg}(p - z)$ and S_p consists of the solutions to the equation

$$\gamma'(t) = \nabla\, \mathrm{Re}(e^{i\theta} f), \quad \lim_{t\to\infty} \gamma(t) = p, \quad f(\gamma(0)) = z$$

The solutions W_p are much the same but these are valued on $(-\infty, 0)$ and approach w as $t \to -\infty$. We also have to replace θ by $\mathrm{Im}\, z - p$. To illustrate what is going on, the following lemma is useful.

Lemma 1.2. Under the gradient flow of $\mathrm{Re}(f)$, $\mathrm{Im}(f)$ is fixed.

Proof. We have

$$\frac{d}{dt}(\mathrm{Im}\, f(\gamma(t))) = \langle \nabla\, \mathrm{Im}\, f, \nabla\, \mathrm{Re}\, f)\rangle.$$

Since f is J holomorphic $J\nabla\, \mathrm{Re}\, f = \nabla\, \mathrm{Im}\, f$ and so we have

$$\frac{d}{dt}(\mathrm{Im}\, f(\gamma(t))) = -\langle \nabla\, \mathrm{Im}\, f, J\nabla\, \mathrm{Re}\, f)\rangle = 0.$$

since the right hand side is the symplectic form $\square$

As a corollary to the previous lemma we have that $\mathrm{Im}(e^{i\theta} f)$ is fixed under the gradient flow of $\mathrm{Re}(e^{i\theta} f)$. The picture to have in mind is Figure 1.1.

Hence the projection of the curves of S_p and W_p into $\mathbb{C}$ is just the line connecting $f(p)$ and z. We also have that

Lemma 1.3. The map $i : S_p \to f^{-1}(z)$ given by $\gamma \mapsto \gamma(z)$ is injective.

Proof. This follows from the fact that a solution to the gradient flow is uniquely determined once we know its value at some time t. $\square$

Lemma 1.4. We can associate the intersection

$$i(S_{p_+}) \cap i(W_{p_-})$$

with the continuous solutions to the problem

$$\frac{d\gamma}{dt} = \nabla \operatorname{Re} e^{i\theta_v}, f(\gamma(0)) = z, \lim_{t\to\pm\infty} = p_\pm.$$

where

$$\theta_v = \begin{cases} \mathrm{Arg}(p_+ - z) & t > 0 \\ \mathrm{Arg}(z - p_i) & t < 0. \end{cases}$$

It turns out that this intersection is a Lagrangian submanifold of $f^{-1}(z)$.

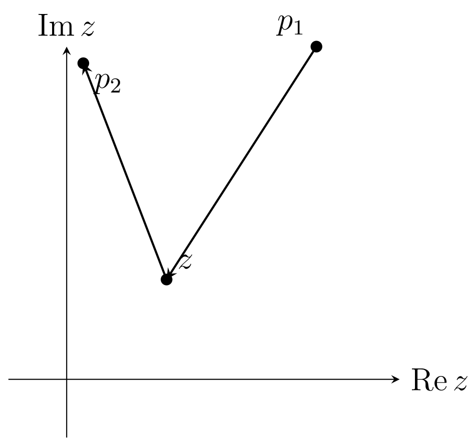

Figure 1.1: The image under f of the curves $\gamma \in S_{p_2} \cap W_{p_1}$

Proof. If $p_0 \in i(S_{p_+}) \cap i(W_{p_-})$ then we can construct a continuous curve by concatenating the corresponding curves. $\square$

We can now outline the Fukaya-Seidel category conjectured by Haydys, see [Hay15; Wan22] for details.

Definition 1.5. Given a Lefschetz fibration $f : Z \to \mathbb{C}$ we can define a Fukaya-Seidel category in the following way. Let $p_1, \ldots, p_n$ be the critical points of the functional f, and suppose that $\mathrm{Arg}(f(p_j) - f(p_i)) > 0$ for $i > j$. Then we define the Lagrangians

$$L_j = i(S_p) \cap i(W_p) \subseteq f^{-1}(z).$$

These will be the objects of the category are the Lagrangians L_j. For the morphisms in the category, we define

$$\hom(L_j, L_k) = \begin{cases} CF(L_j, L_k) & j < k \\ \mathbb{Z}_2.\,\mathrm{id} & j = k \\ 0 & j > k. \end{cases}$$

We give an indication on how to define the higher order morphisms of the A_∞ category. These are maps

$$\mu^d : \hom(L_{j_d}, L_{j_{d+1}}) \otimes \cdots \otimes \hom(L_{j_1}, L_{j+2}) \to \hom(L_{j_1}, L_{j+d})$$

For μ^1, this map is the differential of the chain complex which corresponds to a count of perturbed holomorphic strips with boundary conditions given by L_{j_1} and L_{j_2}. Written in terms of the broken flow lines γ this is a solution to the perturbed Cauchy-Riemann problem

$$\frac{\partial}{\partial \bar{z}} u = -i \nabla \,\mathrm{Re}(e^{i\theta_v} f)$$

that interpolates between two solutions γ_i. The differential takes the starting curve as input and returns the formal sum of all the other curves for which there are solutions to the holomorphic strip equation.

In the next three sections we study Haydys' construction of the Fukaya-Category when applied to the Chern-Simons functional.

1.2.1 Flat $SL(2;\mathbb{C})$ Connections

As in the case of Instanton Floer Theory, flat connections are the critical points of the complexified Chern-Simons theory. We prove this in Proposition 5.6. It follows that to consider an analogue of Picard-Lefschetz theory with the complexified

Chern-Simons functional, one needs to have a good understanding of the moduli space of flat $G_{\mathbb{C}}$ connections. In [Hay15], when constructing the Fukaya-Category for a three-manifold it is assumed that there are finitely many critical points of the Chern-Simons functional. This is not true in general, even for the case of integer homology three-spheres.

Theorem 1.6. There exist integer homology three-spheres with infinitely many irreducible critical points of the $SL(2;\mathbb{C})$ Chern-Simons functional.

Proof. See [GSZ23, Theorem 1.5]. □

On the other hand, there do exist some integer homology three-spheres that have finitely many critical points of the $SL(2;\mathbb{C})$ Chern-Simons functional. This is implied by the manifold having a finite fundamental group. This is because, up to conjugation, $SL(2;\mathbb{C})$ has finitely many finite subgroups of a given order.

In the setup above, we want to consider the complex Chern-Simons functional on the space of $G_{\mathbb{C}}$ connections modulo $G_{\mathbb{C}}$ gauge transformations. As discussed, a naive consideration of the space of $G_{\mathbb{C}}$ connections modulo gauge would be non-Hausdorff, and so it makes sense to impose additional restrictions when considering this quotient. We can solve these issues by restricting the gauge transformations to those coming from a G reduction and also restricting to $G_{\mathbb{C}}$ gauge equivalence classes of connections for which there exists a connection on which $\|\phi\|_{L^2}$ takes a minimum value. This idea is similar to what is done in Donaldson theory where one restricts to irreducible connections. All in all, the above considerations, result in considering the critical points of the Chern-Simons functional as solutions to the system

$$F_{\mathfrak{A}} = 0, d_{\mathfrak{A}}^{*}\phi = 0,$$

where $\mathfrak{A} = A + i\phi$ with respect to the G reduction, see Chapter 3.

1.2.2 The Kapustin-Witten Equations

In the analogy with Picard-Lefschetz theory, one wants to study the gradient flow of the functionals

$$\Phi_{\theta_v} : \mathcal{A}_{\mathbb{C}} \to \mathbb{R}, \quad \mathfrak{A} \mapsto \mathrm{Re}(e^{i\theta_v}\mathfrak{C}(\mathfrak{A})).$$

The resulting equations are known as the Kapustin-Witten equations, first derived in [KW07].

Lemma 1.7. The negative gradient flow of the functional Φ_θ is the equation

$$\frac{d\mathfrak{A}}{dt} = - * e^{i\theta} F_{\mathfrak{A}},$$

where $*$ is the extension of the hodge star operator to $\Omega^2(\mathfrak{g}_{P_\mathbb{C}})$ that is anti-complex linear.

Proof. From (5.1) we have that

$$\begin{aligned} d_a \operatorname{Re}(e^{i\theta}\mathfrak{C}(A)) &= \operatorname{Re}(e^{i\theta} d_a \mathfrak{C}(A)) \\ &= \operatorname{Re}\left(e^{i\theta}\int_Y \operatorname{tr}(a \wedge F_{\mathfrak{A}})\right) \\ &= \int_Y \operatorname{Re}\left(\operatorname{tr}(a \wedge e^{2i\theta} F_{\mathfrak{A}})\right). \end{aligned}$$

Hence we have that $*\nabla\Phi_\theta = e^{i\theta} F_{\mathfrak{A}}$. Here we are using formulas for the L^2 inner product with respect to our choice of inner product from Chapter 3. □

Remark 1.8. When studying the Picard-Lefschetz theory of the Chern-Simons functional, one should replace the constant θ by a function $t \mapsto \theta_v(t)$, and study the related flow equation

$$\frac{d\mathfrak{A}}{dt} = - * e^{i\theta_v} F_{\mathfrak{A}}.$$

Note that imposing such a solution removes translation invariance of the equations. Since the index of the gradient flow equations is 0, see Corollary 9.12, for a suitably generic setup it would be expected that there are no solutions to the equations with θ_v constant, and instead there would be a codimension 1 subset of parameters for which the constant θ equations admit solutions. This phenomena is discussed in [Wit10].

As with the case for flat connections on Y, we impose the condition

$$d^*_{\mathfrak{A}}\phi = 0$$

at each time t to consider a specific set of $G_\mathbb{C}$ orbits. Imposing this condition makes sense because the quantity $d^*_A\phi$ is conserved under the flow, see Corollary 5.13, and also that our critical points will satisfy $d^*_A\phi$. Note however that, due to the use of the L^2 gradient and the resulting appearance of the hodge star, the equations for the gradient flow are not invariant under $G_\mathbb{C}$ gauge transformations, unless one also changes the metric with the gauge transformation. By this we mean that, since we define the equations by fixing a choice of G reduction, with respect to this fixed G reduction the gradient flow is not $G_{\mathbb{C}}$ gauge invariant. However, if we consider $G_{\mathbb{C}}$ gauge transformations as also changing the G reduction, then of course the gradient flow equation is conserved, as is the condition $d_A^*\phi = 0$.

Just as for the anti-self-dual equations on $I_t \times Y$, we can in fact find an equation on all four manifolds that recovers the gradient flow equation in the special case that $\mathfrak{A}(\partial_t) = 0$.

Definition 1.9. The θ–Kapustin-Witten equation on a four manifold X, where $P_{\mathbb{C}} \to X$ is a principal $G_{\mathbb{C}}$ with a choice of G reduction is given by

$$(e^{i\theta} F_{\mathfrak{A}})^+ = 0, \quad d_A^*\phi = 0.$$

Here $+$ refers to the self-dual part

$$(e^{i\theta} F_{\mathfrak{A}})^+ = \frac{1}{2}\left(e^{i\theta} F_{\mathfrak{A}} + *e^{i\theta} F_{\mathfrak{A}}\right),$$

where $*$ is again the anti-complex linear extension to the Hodge star operator. In the special case of manifolds of the form $I_t \times \mathbb{R}$ when $(A)(\partial_t) = 0$ these equations are the negative gradient flow of $\mathrm{Re}(e^{i2\theta}\mathfrak{C})$.

Remark 1.10. Given a smooth function $\theta_v : X \to \mathbb{R}$, we can define the $\frac{\theta_v}{2}$–Kapustin-Witten equation as solutions to

$$(e^{i\frac{\theta_v}{2}} F_{\mathfrak{A}})^+ = 0, d_A^*\phi = 0.$$

When going from the four manifold picture to the cylindrical picture, there is an added complication to justify that we should look for solutions with $\phi_0 = 0$. This condition is best viewed as a boundary condition for the operator $d_A^*\phi = 0$. We want to consider the critical points at infinity up to $G_{\mathbb{C}}$ gauge equivalence and so we want to allow for gauge transformations on the cylinder that tend to a value different from the identity at the ends. Hence the condition $d_A^*\phi = 0$ needs to include the boundary condition $\phi_0 = 0$ at infinity. See the discussion after Definition 3.14 for a discussion. Once this condition is included, it then follows that $\phi_0 = 0$ for all t, either by a maximum principle argument or just from looking at the gradient flow equations.

1.2.3 The Haydys-Witten Equations

The final element of the Picard-Lefschetz construction with the Chern-Simons functional is to consider are the holomorphic disk equations. This leads to the

Haydys-Witten equations. The Haydys-Witten equations are gauge theoretic equations on five dimensional manifolds. The equations have been studied independently by Haydys and Witten, [Hay13; Hay15; Wit11].

As discussed above, for the holomorphic disk equations with respect to the Lefschetz firbation $f : Z \to \mathbb{C}$, we consider maps

$$u : \mathbb{C} \to Z$$

which satisfy

$$\frac{du}{d\bar{z}} = -i\nabla \operatorname{Re}(e^{i\theta} f)$$

and appropriate boundary conditions relating to flow lines. In the case where f is the Chern-Simons functional these equations become

$$\frac{d\mathfrak{A}}{d\bar{z}} = *ie^{i\theta} F_{\mathfrak{A}}, \quad d^*_{\mathfrak{A}}\phi = 0$$

Again, as in the four dimensional case these equations may be written for general five-manifolds that admit a nowhere vanishing vector field. In the following we pick a vector field v on a five manifold (note that all closed five manifolds admit such a field) and we define the orthogonal tangent space as

$$W = \langle v \rangle^{\perp} \subseteq TX.$$

We can then define ΛW^*, $\Lambda^2_+ W$ and define projections of differential forms onto these spaces. When $X = \mathbb{R} \times Y$ for a four-manifold Y, these bundles correspond to restrictions to the and self-dual exterior algebra. Note that if $Y = \mathbb{R} \times Z$ for a three manifold Z, then we can identify $\Lambda^2_+ W$ with the tangent bundle of Z.

Definition 1.11. The Haydys-Witten equations on a five-manifold with respect to a nowhere vanishing vector field v and connection $\mathfrak{A}$ can be described as

$$\begin{aligned} i_v F_A &= (d_A^+)^* \phi \\ F_A^+ &= -\nabla_v \phi + \phi \wedge \phi. \end{aligned}$$

Here A is a connection on a principal bundle $P \to X$, $\phi \in \Gamma(\Lambda^2_+ W \otimes \mathfrak{g}_P)$ and the product $\phi \mapsto \phi \wedge \phi$ is such that it corresponds with the one form map $\phi \mapsto *(\phi \wedge \phi)$ when $\Lambda^2_+ W$ is identified with the tangent bundle of a three manifold. See [Hay15, Section 3] for details.

1.3 Gauge Theory and Knot Theory

The key motivation for Witten's proposals in [Wit10; Wit11] is to use the Kapustin-Witten equations to study the Jones polynomial for a given knot. Since the beginning of the study of mathematical gauge theory and gauge theories in physics, it has been clear that there are deep relations between gauge theory and knot theory. For example, in [Wit89], Witten proved that the Partition function of a singular version of Chern-Simons theory (where the solutions in the theory are singular along a given knot) is related to the Jones polynomial of the knot. Moreover, Floer started work extending his instanton homology groups to a homology theory of knots, see [Flo91; BD95]. Floers idea, much like the definition of the knot Casson invariant, was to compare his instanton homology groups for manifolds obtained by Dehn surgery on a knot, using an exact triangle. More recently, Kronheimer and Mrowka have defined singular instanton Floer homology [KM11; KM10]. The singular instanton Floer homology framework, as the name suggests, involves studying Floer homology groups on knot complements with specified singular boundary conditions.

These approaches all seem different but they can be related through perturbations to the Chern-Simons functional. Witten's starting point in [Wit89] is to consider a physical theory defined by an action functional which is the Chern-Simons action modified by Wilson loop operators.

Definition 1.12. A Wilson loop operator is defined as the trace of the holonomy of a connection around a given loop. That is, if we have a principal G bundle $P \to Y$ and a representation $R : G \to GL(n;\mathbb{C})$, then the Wilson loop operator is the map

$$W_R^\gamma : \mathcal{A} \to \mathbb{C}, \quad W_R^\gamma(A) = \mathrm{tr}(R(H_A)), \tag{1.2}$$

where H_A is the holonomy of a connection around the given loop γ.

Witten then studies the physical theory defined by the action

$$\int dA \exp(iCS(A)) \prod_{i=1}^{r} W_{R_i}^{\gamma_i}(A), \tag{1.3}$$

and proves that the partition function of this theory relates to the Jones polynomial. The ground states of this theory become flat connections on the complement of the curves γ_i, and these flat connections must satisfy certain conditions coming from the Wilson loop operators.

This modification to the theory bears many similarities to the perturbations used in [Tau90] and [Flo88]. As described above, in these papers perturbations to the Chern-Simons functional of the form

$$\eta : A \mapsto \int_{D^2} f(\mathrm{tr}(H_A))\mu \tag{1.4}$$

are used, where $S^1 \times D^2 \hookrightarrow Y$ is a knot neighbourhood and the integral is taken over a slice $\{t\} \times D^2$ and μ is a two form on Y. See, Section 5.3 for details. The resulting perturbed flat connection equation is

$$F_A = -f'(\mathrm{tr}(H_A)) \left(H_A - \frac{1}{2}\,\mathrm{tr}(H_A) \right) \mu. \tag{1.5}$$

We can relate these two points ideas by sending the radius of the knot neighbourhood to 0. Hence, it should be possible to study the singular ground states using these perturbed functions.

Another relation to the perturbed Chern-Simons functional is Kronheimer and Mrowka's singular instanton Floer homology. In this theory they impose boundary conditions in terms of the holonomy of the connection around $\mathbb{R} \times \gamma \subseteq \mathbb{R} \times Y$ and study the resulting Floer homology theory. In the perturbed picture, this corresponds to studying the gradient flow of the perturbed Chern-Simons functional with a specific choice of function f to impose the holonomy constraints, see Theorem 5.27. In some sense the solutions of the perturbed instanton picture are a smoothed out version of the singular instanton moduli space.

Witten's proposal for knots in [Wit10] is to study the Kapustin-Witten equations with singularities along $\mathbb{R} \times K$, in much the same way as for singular instanton Floer homology. As far as the author is aware, since this work was published there hasn't been much work on this incarnation of Witten's proposal. In many ways it was succeeded by Witten's study of solutions to the Kapustin-Witten equations with Nahm pole solutions.

The perturbations we study in Chapter 5 could be used to study Witten's proposal in the same way that Floer's approach to Knot Homology has strong links to singular instanton Floer homology. A potential advantage of this is that the analysis in Witten's singular boundary conditions could be difficult, especially with non-compactness issues. In the approach given here, one of the main motivations for introducing the perturbation is to solve compactness issues. Another benefit of using these perturbations is that we get a topological understanding of the critical points, and a relationship between the solutions and three manifold topology. In particular Dehn surgery relations as in Floer's original work.

1.3.1 The Nahm Pole Boundary Condition

We end this introduction by considering the Nahm Pole solutions to the Kapustin-Witten equations. This relates to the content of the last few chapters of the thesis. In [Wit11], Witten conjectures that a count of certain singular solutions to the Kapustin-Witten equations to give a gauge theoretic interpretation of the Jones polynomial. The solutions that Witten proposes to count are solutions to the $\frac{\pi}{4}$-Kapustin-Witten equations on $[0, \infty) \times Y$ with a so called Nahm pole singularity at 0. A solution to the Kapustin-Witten equations, $\mathfrak{A} = A + i\phi$, on a $SL(2;\mathbb{C})$ bundle over $(0, \infty) \times Y$ is a Nahm pole solution if

$$|\phi| \sim \frac{1}{t} \text{ as } t \to 0, \tag{1.6}$$

and in the limit $\phi - t * (\phi \wedge \phi)$ is bounded. A slightly different characterisation of the Nahm pole boundary condition is that the one-form ϕ, when suitably rescaled, gives an isometry

$$\phi : TM \to \mathfrak{g}_P \tag{1.7}$$

such that the pullback of the connection A is the Levi-Civita connection of Y. The Nahm pole boundary condition can also be extended to include knots. In this case, one looks for solutions that are Nahm pole solutions away from the knot and more singular along the knot embedded in $\{0\} \times Y$. See [MW13; MW18] for detailed analysis of the Nahm pole boundary condition.

The motivation for defining this boundry condition comes from physical dualities. The idea is that due to certain dualities in physical theories, two different physical theories will share the same partition function (or there will be a bijective correspondence between the data defining them). In this sense, the Nahm pole theory is dual to the Floer style Kapustin-Witten theory.

The difficulty when considering these physical theories is that the physical dualities do not have a solid mathematical underpinning. It might be reasonable to suggest that two different theories should be dual to each other without a rigorous mathematical proof. Witten proposes instead that one should take the dual theory as the starting point and prove directly that the invariants give the Jones polynomial. Again with this theory there are difficult issues with compactness of the space of solutions. We take a first step in this direction in the final two chapters of the thesis by studying solutions to the $\frac{\pi}{4}$–Kapustin-Witten equations for which $|\phi| \to \infty$ in finite time. We show that if this occurs then $\mathfrak{A}$ admits a Nahm pole singularity.

Part I

Analytic Foundations for Kapustin-Witten Theory

Chapter 2

Sobolev Spaces and Embeddings

We will be perturbing the Chern-Simons functional using holonomy perturbations. These perturbations result in studying connections that are solutions to the equation

$$F_{\mathfrak{A}} + f'(U_{\mathfrak{A}})V_{\mathfrak{A}} = 0,$$

see Chapter 5. When we look at the deformation theory of solutions we will end up considering integrals of the form

$$\int_{S^1} U(1;\tau)a_\tau(\tau)U(\tau;0)d\tau,$$

where $U(b,a)$ denotes parallel transport with respect to $\mathfrak{A}$ along a given path $S^1 \hookrightarrow Y$. This is the derivative of the holonomy of a connection as one varies the connection in the direction a_τ. Here we will impose $a \in L^p_k(\Omega^1(\mathfrak{g}_{P_{\mathbb{C}}}))$, where L^p_k denotes an appropriate Sobolev completion of the space of one forms. Hence we will require Sobolev embedding theorems onto one dimensional submanifolds of a three manifold. In this chapter, for lack of a good reference, we prove the required embedding for Bessel potential spaces. For completeness we also define the spaces that we are working with.

2.1 Banach Spaces of Sections and Connections

To generalise function spaces on subsets of $\mathbb{R}^n$ to spaces on compact manifolds, a standard procedure is to define the function spaces as those functions which restrict to the relevant spaces on $\mathbb{R}^n$ with respect to a partition of unity. To start we define function spaces on $\mathbb{R}^n$. Many extra details on these spaces may be found in [JW84, Chapter 1]. In this thesis, we will mainly use function spaces defined using the Fourier transform. The main reason for this is that we can easily prove Sobolev trace theorems for arbitrary codimension using these spaces.

Definition 2.1 (Potential Spaces)**.** The Bessel potential space, $L^p_\alpha(\mathbb{R}^n)$, for $1 \leq p \leq \infty$ and $\alpha \geq 0$ is the image of $L^p(\mathbb{R}^n)$ under the map

$$C_\alpha : L^p(\mathbb{R}^n) \to L^p(\mathbb{R}^n), f \mapsto G_\alpha * f$$

where $*$ denotes convolution, and $G_\alpha \in L^1$ is the function with Fourier transform

$$\hat{G}_\alpha(\xi) = (1 + |\xi|^2)^{-\frac{\alpha}{2}}.$$

The norm is chosen such that C_α defines an isometry – that is

$$\|f\|_{L^p_\alpha(\mathbb{R}^n)} = \|g\|_{L^p(\mathbb{R}^n)}$$

where $f = G_\alpha * g$

Definition 2.2 (Fractional Sobolev Spaces on Manifolds)**.** Let Y be a compact manifold and $E \to Y$ a vector bundle. We define $L^2_\alpha(Y; E)$ to be completion of $\Gamma(E)$ with respect to the norm

$$\|f\|^2_{L^2_\alpha} = \sum_i \|(\varphi_i f)|_{U_i}\|^2_{L^2_\alpha(\mathbb{R}^n;B)},$$

where we pick a finite cover $\{U_i\}$ such that the bundle E may be trivialised so that

$$\psi : E|_{U_i} \cong V_i \times B$$

and $f|_{U_i} : V_i \to B$ is the projection of $\psi(f(\pi\psi^{-1}(v, 0)))$ onto B. φ_i is a partition of unity subordinate to the cover. For a compact manifold the space $L^2_\alpha(Y, B)$ does not depend on the choices made to define it.

Now let X be a non-compact, complete Riemannian manifold with positive injectivity radius and bounded geometry (curvature in C^∞). We can define the Sobolev spaces, for vector bundles $E \to X$, $L^2_\alpha(X, E)$ in much the same way as the compact case, but in this case we can pick a countable cover U_i of X where for each point, $p \in X$, there are finitely many elements of the cover containing p.

The potential spaces recover traditional Sobolev spaces for integer k.

Definition 2.3. Let X be a, possibly non-compact, Riemannian manifold, and let $E \to X$ be a Hermitian vector bundle, with a connection ∇. Using the Levi-Civita connection we can extend ∇ to a connection

$$\nabla : (T^*X)^{k-1} \otimes E \to (T^*X)^k \otimes E.$$

We say that a section $f \in \Gamma(E)$ has an L^2 weak derivative if there exists a section $\eta \in T^*X \otimes E$ such that

$$\langle f, \nabla^* g \rangle = \langle \eta, g \rangle,$$

for all $g \in (T^*X) \otimes E$. We can iterate taking weak derivatives.

Theorem 2.4. Let $W^{k,2}(X;E)$ for $k \in \mathbb{N}$ be the space of functions with k weak L^2 derivatives. Then

$$L^2_k(X;E) \cong W^{k,2}(X;E).$$

Proof. See, [AF12, Example 7.62]. □

2.2 Generalised Sobolev Trace Theorems

Proposition 2.5. Let $f \in C^\infty(\mathbb{R}^n)$. Then, for a subspace $\mathbb{R}^d \subseteq \mathbb{R}^n$ we have

$$\|f\|_{L^2(\mathbb{R}^d)} \leq C\|f\|_{L^2_s(\mathbb{R}^n)}$$

provided

$$s > \frac{n-d}{2}.$$

Proof. We first note that

$$\hat{f}|_{\mathbb{R}^d}(\xi) = B \int_{\mathbb{R}^{n-d}} \hat{f}(\xi, \xi'),$$

for some constant B, where the left hand side is the Fourier transform of the restriction, and the integrand on the right hand side is the Fourier transform on $\mathbb{R}^n$. To see this, we write

$$f(x,0) = B \int_{\mathbb{R}^d} e^{ix.\xi} \left(\int_{\mathbb{R}^{n-d}} e^{i0.\xi'} \hat{f}(\xi, \xi') d\xi' \right) d\xi$$

and so $f(x,0)$ is the Fourier transform of

$$\xi \mapsto \int_{\mathbb{R}^{n-d}} \hat{f}(\xi, \xi') d\xi',$$

up to a constant. Using the Cauchy-Schwartz inequality, we can then write

$$|\hat{f}|_{\mathbb{R}^d}(\xi)|^2 \leq \int_{\mathbb{R}^{n-d}} (1+|\xi'|^2)^{-s} d\xi' \int_{\mathbb{R}^{n-d}} (1+|\xi'|^2)^s |\hat{f}(\xi,\xi')|^2 d\xi'.$$

This integral on the right converges due to the assumptions on f. The left hand integral will converge provided the function decays sufficiently quickly, which happens if

$$s > \frac{n-d}{2}.$$

If we then integrate the inequality over $\mathbb{R}^d$, and use the fact that $(1+|\xi|^2) < (1+|(\xi,\xi')|^2)$, we get

$$\|f\|_{L^2(\mathbb{R}^d)} \leq C\|f\|_{L^2_s(\mathbb{R}^n)},$$

if $s > \frac{n-d}{2}$. □

The argument in the above proof generalised to higher derivatives and we have the following generalised trace theorem

Corollary 2.6. Let $f \in C^\infty_c(\mathbb{R}^n)$. Then

$$\|f\|_{L^2_k(\mathbb{R}^d)} \leq \|f\|_{L^2_s(\mathbb{R}^n)},$$

provided $k < s + \frac{n-d}{2}$.

Proof. Note that, for $\xi' \in \mathbb{R}^d$ and $\xi \in \mathbb{R}^{n-d}$ we have

$$(1+|\xi|^2)(1+|\xi'|^2) = (1+|\xi|^2+|\xi|^2+2|\xi|^2|\xi'|^2) \leq (1+|\xi|^2+|\xi'|^2)^2$$

Hence we can write

$$(1+|\xi|^2)^k(1+|\xi'|^2)^l \leq (1+|\xi|^2+|\xi'|^2)^{k+l}$$

and we then have that

$$\begin{aligned}\int_{\mathbb{R}^d}(1+|\xi|^2)^k|\hat{f}|_{\mathbb{R}^d}|^2 d\xi &\leq \left(\int_{\mathbb{R}^{n-d}}(1+|\xi'|^2)^l d\xi'\right)\int_{\mathbb{R}^n}(1+|(\xi,\xi')|^2)^{k+l}|\hat{f}(\xi,\xi')|^2 \\ &\leq C\|f\|_{L^2_{k+l}\mathbb{R}^n},\end{aligned}$$

provided $l > \frac{n-d}{2}$. This proves the result □

We now show that in the case for $\mathbb{R} \subseteq \mathbb{R}^3$ the bound on k is sharp. This is unfortunate since it means that for an L^2_1 $SL(2;\mathbb{C})$ connection it is impossible to bound the norm the holonomy of a connection by its L^2_1 norm on a particular path – for a generic path we could do this via Tonelli's theorem. We are, however, able to bound it by the $L^2_{1+\delta}(Y^3)$ norm, or if the connection is on a four dimensional space, by its $L^2_{\frac{3}{2}+\delta}(X^4)$ norm.

Proposition 2.7. There is no embedding $L^2_1(\mathbb{R}^3) \to L^q(\mathbb{R})$ for $q \in [1, \infty)$.

Proof. We can define a counter example to such an embedding via the function on $\mathbb{R}^3$ given by

$$f(x, y, z) = f(r, z) = \sin(\log(-\log r))\,\phi(r, z)$$

where $\phi : \mathbb{R}^2 \to [0, 1]$ is a non-negative smooth function with support in $B(0, 1)$, and $r > 0$ satisfies $r^2 = x^2 + y^2$. With this function we have

$$df = \cos(\log(-\log r))\,\frac{1}{r\log r}dr + \sin(\log - (\log r))\,d\phi$$

Hence $f \in L^2_1$ since sin is bounded, $d\phi \in L^2$, and over $\mathbb{R}^2$ we have

$$\int_{r<1} \frac{1}{r^2\log^2 r} = \int_0^1 \frac{\pi}{r\log^2 r}dr = \pi\int_{-\log 1}^{\infty} \frac{1}{t^2}dt = \frac{\pi}{\log 1}.$$

From this we have that $f \in L^2_1(\mathbb{R}^3)$. We will conclude by saying that f has no trace on the line $r = 0$. To define the restriction of f to $r = 0$ we would define

$$f_{|r=0}(z) = \lim_{r\to 0} \frac{1}{\mathrm{vol}(B(0, r))} \int_{B(0,r)} \sin(\log(-\log r))\phi(r, z)d\mu$$

where we integrate over the disk $B(0, 1)$ in the affine plane perpendicular to $r = 0$ intersecting $(0, 0, z)$. However this integral, after a change of the variables $t = \log r$ is given by

$$\lim_{t\to -\infty} e^{-2t} \int_{-\infty}^{e^t} e^{2t}\sin(t)dt$$

$\square$

We finish this section with a result on restricting the directions of the derivatives. These results are obvious when the Sobolev spaces are defined by weak partial derivatives, but less clear when defined using the Fourier transforms.

Definition 2.8. Let X be a complex Banach space. Then the space

$$L^2_k(\mathbb{R}^n; B)$$

is defined to be the space of smooth functions such that the norm

$$\int_{\mathbb{R}^k} (1+|\xi|^2)^k |\hat{f}(\xi)|^2$$

Here the Fourier transform of the function taking values in the Banach space is defined in terms of the Bochner integral. See, for example, [DU77].

Lemma 2.9. Let $f \in L^2_k(\mathbb{R}^n)$. Then we can view f as an element

$$f \in L^2_j(\mathbb{R}^k; L^2_{k-j}(\mathbb{R}^{n-k})).$$

Proof. We can define a map of functions

$$\mathrm{ev} : (\mathbb{R}^n \to \mathbb{R}) \to (\mathbb{R}^d \to (\mathbb{R}^{n-d} \to \mathbb{R}))$$

via

$$\mathrm{ev} : f \mapsto (x \mapsto f(x, \cdot)).$$

The statement of the lemma is that ev defines a continuous map of Banach spaces

$$\mathrm{ev} : L^2_k(\mathbb{R}^n) \to L^2_j(\mathbb{R}^d; L^2_{k-j}(\mathbb{R}^{n-d})).$$

To see this, note first that $\hat{\mathrm{ev}(f)}(\xi) = y \mapsto \int_{\mathbb{R}^d} e^{-ix.\xi} f(x,y)dx$, and so

$$\|\hat{\mathrm{ev}} f(\xi)\|^2_{L^2_{k-j}(\mathbb{R}^{n-d})} = \int_{\mathbb{R}^{n-d}} (1+|\eta|^2)^{k-j} \int_{\mathbb{R}^{n-d}} e^{iy.\eta} \left(\int_{\mathbb{R}^d} e^{-ix.\xi} f(x,y) dx dy \right) d\eta$$

Hence, we have that for $f \in L^2_k(\mathbb{R})$

$$\begin{aligned}
&\int_{\mathbb{R}^d} (1+|\xi|^2)^{k-j} \|\hat{\mathrm{ev}(f)}(\xi)\|^2 = \\
&\quad \int_{\mathbb{R}^d} (1+|\xi|^2)^j \int_{\mathbb{R}^{n-d}} (1+|\eta|^2)^{k-j} \int_{\mathbb{R}^{n-d}} e^{iy.\eta} \left(\int_{\mathbb{R}^d} e^{-ix.\xi} f(x,y) dx dy \right) d\eta \\
&\quad \leq \|f\|_{L^2_k(\mathbb{R}^n)}.
\end{aligned}$$

□

Corollary 2.10. In the special case that $j = 0$, we have that $f \in L^2_k(\mathbb{R}^n)$ defines a function

$$f \in L^2(\mathbb{R}^d; L^2_k(\mathbb{R}^{n-d})).$$

Remark 2.11. Note that the norms on the spaces $L^2_j(\mathbb{R}^d; L^2_{k-j}(\mathbb{R}^{n-d}))$ are weaker than that of $L^2_k(\mathbb{R}^n)$ since, for example, we miss information about the j–th derivative in the $\mathbb{R}^d$ subspace. The space $L^2_k(\mathbb{R}^n)$ is the intersection of all the spaces $L^2_j(\mathbb{R}^d; L^2_{k-j}\mathbb{R}^{n-d})$ for all j.

2.3 Derivative Estimates for Integral Functions

In this section we consider derivatives of functions of the form

$$f(x) = \int_{\mathbb{R}} h(y, x) dy. \tag{2.1}$$

for functions $h \in C^\infty(\mathbb{R} \times \mathbb{R}^n)$. The application we have in mind is considering the derivative of the holonomy of a connection around a loop with respect to changing the connection. One can write this in the form

$$\frac{dH}{da} = \int_0^1 U(1;\tau) a_\tau U(\tau; 0) d\tau.$$

where $U(a, b)$ denotes a parallel transport operator which we assume is bounded and smooth, and the derivative should be interpreted as the functional derivative in the direction a.

Lemma 2.12. Let $f(x)$ be as in (2.1). Then

$$\hat{f}(\xi) = \hat{h}(0, \xi).$$

Proof. We have

$$\begin{aligned} \hat{h}(0, \xi) &= \int_{\mathbb{R}^n} h(x, y) e^{-ix.\xi} dx dy \\ &= \int_{\mathbb{R}^n} e^{-ix.\xi} \left(\int_{\mathbb{R}} h(x, y) dy \right) dx. \end{aligned}$$

□

Lemma 2.13. Suppose $h \in L^2_k(\mathbb{R}^n \times S^1)$. Then the function

$$f(x) = \int_{S^1} h(x, y) dy$$

is in the space $L^2_k(\mathbb{R}^n)$.

Proof. We have, from Lemma 2.9, that

$$h \in L^2_k(\mathbb{R}^n; L^2(S^1)).$$

By Tonelli's theorem, we have that the map

$$f : x \mapsto \int_{S^1} h(x,y)dy$$

is almost everywhere defined, and we can estimate the L_k^2 norm by

$$\begin{aligned}
\|f\|_{L_k^2} &= \int_{\mathbb{R}^n} (1+|\xi|^2)^k \|\hat{f}\|^2_{L^2(S^1)} \\
&= \int_{S^1} \int_{\mathbb{R}^n} (1+|\xi|^2)^k \left(\int_{\mathbb{R}^n} e^{ix.\xi} f(x,y)dx \right)^2 d\xi dy \\
&= \int_{S^1} \|f(\cdot,y)\|^2_{L_k^2} dy \leq \|f\|^2_{L_k^2(\mathbb{R}^k \times S^1)}.
\end{aligned}$$

□

Lemma 2.14. If we restrict to compact domains $\Omega \subseteq \mathbb{R}^n$, the map

$$h \mapsto f : L_s^2(\Omega \times S^1) \to L^2(\Omega)$$

as defined in (2.1) is compact for all $s > 0$

Proof. This follows from the compact embedding $L_s^2(\Omega) \to L^2(\Omega)$ for bounded domains. □

Chapter 3

Geometric Setup

In this chapter we study principal bundles with structure group $G_{\mathbb{C}}$, where $G_{\mathbb{C}}$ is the complexification of a compact group G. This is a principal bundle with a complex, non-compact structure group whose adjoint bundle is given by $\mathfrak{g}_P \otimes \mathbb{C}$, where P is an appropriate principal G bundle. We call such a bundle a complexified bundle, or the complexification of the bundle P. Studying the Kapustin-Witten equations on a complexified bundle ensures that there is a distinguished group of gauge transformations, those induced by automorphisms of P, for which there is an invariant norm on the adjoint bundle.

For a section of a complex vector bundle bundle $E \to X$ with a real structure, one can consider the real and imaginary parts of that section. This is what is often done in the study of the Kapustin-Witten equations. For example, in [Wit11] the $\frac{\pi}{4}$–Kapustin-Witten equations are presented as

$$F_A - \phi \wedge \phi + *d_A\phi = 0, d_A^*\phi = 0. \tag{3.1}$$

We take a slightly different approach and view the equations as equations for a complex connection on a $G_{\mathbb{C}}$ bundle $P_{\mathbb{C}} \to X$, with respect to a fixed G reduction $P \subseteq P_{\mathbb{C}}$. As explained in Chapter 1, the $\frac{\pi}{4}$ Kapustin-Witten equations may also be written as

$$\left(e^{i\frac{\pi}{4}} F_{\mathfrak{A}}\right)^+ = 0, \quad d_A^*\phi = 0$$

where $\mathfrak{A} = A + i\phi$. These two points of view are largely equivalent, and mostly a notational preference.

We chose the latter approach to keep the relation of the pairs (A, ϕ) to complex connections explicit. As we describe in this chapter, there is also a natural analogue of the Coulomb gauge for complex connections, which is different from imposing only a gauge condition on A and then using $d_A^*\phi = 0$ to remove the gauge equivalence on ϕ. This makes the analysis of the gradient flow lines of $\mathfrak{C}$ much more analogous to the case of Floer homology.

In this chapter, after recalling the notions of complexification of Lie groups and defining the complexification of a principal bundle, we consider the geometry of the space of complex connections, $\mathcal{A}_{\mathbb{C}}$, with respect to the action the gauge group $\mathcal{G}$ of automorphisms of the underlying real bundle. We then introduce the complexified Coulomb gauge condition and prove a gauge fixing result for connections that are L_1^2 close, and then consider notions of gauge fixing on manifolds with boundary. We then take a step back and consider the action of the full $\mathcal{G}_{\mathbb{C}}$. We consider the corresponding equivalence classes of connections as well the space of G reductions of a given $G_{\mathbb{C}}$ bundle. By doing this, we lay the groundwork for the results in Chapter 4, which generalise the results of [Cor88] to manifolds with boundary. This generalisation is a key result that we require for studying the critical points of the perturbed complex Chern-Simons functional

We then consider the complex Yang-Mills equations and prove estimates for solutions to these equations. Solutions to the Kapustin-Witten equations and flat connections are automatically solutions to the Yang-Mills equations, so the estimates proved here apply in these cases. Our main result of this section is a proof of interior C^∞ regularity for solutions to the complex Yang Mills equations in the complex Coulomb gauge.

We then consider flat $G_{\mathbb{C}}$ connections. As described in Chapter 1, stable flat $G_{\mathbb{C}}$ connections are the critical points of the gradient flow of $\mathfrak{C}$. We recall the notions of the representation variety and twisted de Rham cohomology, which are key ideas to describe the moduli spaces of flat connections.

We end the chapter by considering the metric induced by the L^2 metric on $\Omega^1(\mathfrak{g}_{P_{\mathbb{C}}})$ on the moduli space of flat connections modulo real gauge transformations. We prove some lower bounds on the geodesic distance between connections which we use in when studying the Kapustin-Witten equations.

3.1 Complexified Principal Bundles

Definition 3.1. Let G be a compact Lie group. The complexification of G, $G_{\mathbb{C}}$, is the unique complex Lie group satisfying the universal property that there exists a map $\gamma : G \to G_{\mathbb{C}}$ of real Lie groups such that for any complex Lie group, K, and real Lie group homomorphism, $u : G \to K$, there exists a complex lie group homomorphism $U : G_{\mathbb{C}} \to K$ such that $U \circ \gamma = u$.

Example 3.2. The complexification of S^1 is $\mathbb{C}^*$, the complexification of $SU(2)$ is $SL(2;\mathbb{C})$ and the complexification of $SO(3)$ is $PSL(2;\mathbb{C})$.

The following results on complex Lie groups are standard.

Proposition 3.3. The complexification of a non-trivial compact Lie group is never compact.

Proof. By the Cartan Decomposition, $G_{\mathbb{C}}$ is diffeomorphic to $G \times \mathfrak{g}$. This is not compact if $\mathfrak{g}$ is non-trivial. □

Proposition 3.4. Let $\mathfrak{g}$ be the Lie algebra of a compact group G. Then the Lie algebra of $G_{\mathbb{C}}$ is isomorphic to $\mathfrak{g} \otimes \mathbb{C}$.

Definition 3.5. Let G be a compact Lie group, and suppose that $P \to X$ is a principal G–bundle. Then the complexification of P, denoted $P_{\mathbb{C}}$, is the associated bundle

$$P \times_\gamma G_{\mathbb{C}}$$

where γ is the left multiplication action.

Since it is an associated bundle, points in $P_{\mathbb{C}}$ are given by the equivalence classes $[(p,g)], p \in P, g \in G$ under the action of G. As such it admits a natural $G_{\mathbb{C}}$ action where, for h in $G_{\mathbb{C}}$, the action is given by

$$\rho_h : [p,g] \mapsto [p,gh].$$

Lemma 3.6. For a compact Lie group G and principal G–bundle $P \to X$, the complexified bundle $P_{\mathbb{C}}$ is a principal $G_{\mathbb{C}}$ bundle.

Proof. A principal $G_{\mathbb{C}}$ is bundle is a locally trivialisable bundle over the base, admitting a right $G_{\mathbb{C}}$ action that is free and acts transitively on the fibres. This property follows from the construction. □

Definition 3.7 (Real and Imaginary Part of a Complex Connection)**.** Let $P_{\mathbb{C}}$ be the complexification of a bundle P, and let $\mathfrak{A}$ be a connection on $P_{\mathbb{C}}$. Then there is a canonical decomposition of the connection, such that

$$\mathfrak{A} \cong A + i\phi.$$

We call A the real part of the connection and ϕ the imaginary part.

Proof. Any connection B on P induces a connection on $P_{\mathbb{C}}$. An easy way to see this is view the connection as an Ehresmann distribution of horizontal subspaces – G invariant complements to the kernel of the projection $\pi : P \to X$. We can use the embedding $i : P \hookrightarrow P_{\mathbb{C}}$ to push forward the horizontal subspace to define the horizontal subspace on the subset of $P_{\mathbb{C}}$ isomorphic to P, and then use $G_{\mathbb{C}}$ invariance to define the horizontal subspace on the whole of P.

Now consider the connection B as a connection on $P_{\mathbb{C}}$. We have that

$$\mathfrak{A} = B + \mathfrak{a}$$

for $\mathfrak{a} = a + i\phi \in \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}}) = \Omega^1(\mathfrak{g}_P) \oplus i\Omega^1(\mathfrak{g}_P)$. Hence we can write

$$\mathcal{A} = B + a + i\phi = A + i\phi$$

for a connection A on P. □

Lemma 3.8. Let $\mathfrak{A} = A + i\phi$ be a connection on the bundle $P_{\mathbb{C}}$ with respect to the reduction P. Then

$$F_{A+i\phi} = F_A + id_A\phi - \phi \wedge \phi$$

Proof. This usual formula for the curvature of a connection carries over to complexified bundles. □

3.2 Gauge Choices

In this section we consider gauge choices when working with connections on complexified bundles. As is common in gauge theory, when working locally in the space of connections, we can remove gauge freedom coming from automorphisms of the bundle by imposing certain gauge conditions on the one forms that describe nearby connections. This gives us a description of the tangent space to the space of connections modulo gauge.

When studying complex connections, a natural choice of gauge condition is to consider how the Coulomb gauge generalises to the complex setting. In this section we review the derivation of the Coulomb gauge condition in the compact structure group case and derive an appropriate analogue in the complexified bundle setting. In the literature such a gauge choice has been called the Kapustin-Witten gauge, [He19], but we prefer to use the term complexified Coulomb gauge to refer to this gauge condition.

Let $P \to X$ be a principal G–bundle. Consider the action of the gauge group $\mathcal{G}$ on $\mathcal{A}$. This is the map

$$\mathcal{G} \times \mathcal{A} \to \mathcal{A}, \quad g(A) = gAg^{-1} - (\nabla_A g)g^{-1}$$

Fixing a connection, A, we can consider the map given by the image of A under $\mathcal{G}$, a map $\mathcal{G} \to \mathcal{A}$. The derivative of this at $(A, 0)$ is the map

$$\Omega^0(\mathfrak{g}) \to \Omega^1(\mathfrak{g}_P)$$

given by

$$\xi \mapsto -d_A\xi.$$

In an appropriate topologies, for example the L^2_2 completion of the one-forms, we have a natural decomposition

$$T_A\mathcal{A} = L^2_2(\Omega^1(\mathfrak{g}_P)) = \mathrm{Im} d_A \oplus \mathrm{Ker} d_A^*$$

where $d_A : L^2_3(\Omega^0(\mathfrak{g}_P)) \to L^2_2(\Omega^1(\mathfrak{g}_P))$. This leads naturally to the Coulomb gauge condition due to the following proposition. A proof of this proposition can be found in, for example, [Don83].

Proposition 3.9. The map $\mathcal{G} \times \ker d_A^* \to A$

$$(g, a) \mapsto g.(A + a)$$

is locally surjective.

We now consider the same construction for the complexified connections with the real gauge group action.

Lemma 3.10. The derivative of the action $\mathcal{G} \times \mathcal{A}_{\mathbb{C}} \to \mathcal{A}_{\mathbb{C}}$ at the connection $A + i\phi$ is given by

$$-d_{A+i\phi} : \Omega^0(\mathfrak{g}_P) \to \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}}), \quad \xi \mapsto -d_{A+i\phi}\xi \tag{3.2}$$

Proof. In a local trivialisation, the action of an element g is given by

$$g.(A + i\phi) = gAg^{-1} - dgg^{-1} + ig\phi g^{-1}$$

Hence for $g = \exp(t\xi)$ we have

$$\frac{d}{dt}\bigg|_{t=0} \exp(t\xi).(A + i\phi) = [\xi, A + i\phi] - d\xi = -d_{A+i\phi}\xi.$$

□

Lemma 3.11. The adjoint of the map (3.2) is

$$\mathrm{Re}(-d^*_{A+i\phi}\cdot) : \Omega^1(\mathfrak{g}_{P_\mathbb{C}}) \to \Omega^1(\mathfrak{g}_P)$$

Proof. Let $\xi \in \Omega^0(\mathfrak{g}_P)$, and $\varphi \in \Omega^1(\mathfrak{g}_{P_\mathbb{C}})$. Then we have

$$\begin{aligned}\langle d_{A+i\phi}\xi, \varphi\rangle &= \langle \xi, d^*_{A+i\phi}\varphi\rangle \\ &= \langle \xi, \mathrm{Re}(d^*_{A+i\phi}\varphi)\rangle\end{aligned}$$

since ξ is real. □

From this we are lead to the following definition.

Definition 3.12. The Complexified Coulomb Gauge condition, for $\varphi \in \Omega^1(\mathfrak{g}_{P_\mathbb{C}})$, is the condition

$$\mathrm{Re}\left(d^*_{A+i\phi}\varphi\right) = 0.$$

We now prove a gauge fixing theorem.

Proposition 3.13. Let $\mathfrak{A}, \mathfrak{B}$ be connections on a complexified principal bundle $P_\mathbb{C} \to X$, where X is a manifolds of dimension 3 or 4. Suppose that $\|\mathfrak{A} - \mathfrak{B}\|_{L^2_1} \leq K(\mathfrak{A})$. Then there exists an L^2_2 gauge transformation $g \in \mathcal{G}$ such that if

$$a + i\varphi = \mathfrak{A} - g.\mathfrak{B}$$

then

$$\mathrm{Re}\left(d^*_{\mathfrak{A}}(a + i\varphi)\right) = 0.$$

Proof. To prove this result we use the implicit function theorem. First note that if $a + i\varphi \in \Omega^1\left(\mathfrak{g}_{P_\mathbb{C}}\right)$ and $u \in \mathcal{G}$ then

$$u(A + a) = \mathfrak{A} + u(a + i\varphi)u^{-1} - (\nabla_{\mathfrak{A}}u)u^{-1}$$

and we want to solve the equation

$$\mathrm{Re}\left(d^*_{\mathfrak{A}}(u(A + a + i\varphi) - A)\right) = 0$$

for u, where $a + i\varphi = B - A$. To do this consider the operator

$$F : L^2_2(\mathcal{G}) \times L^2_1(\ker \mathrm{Re}\,(d_A^*)) \to L^2(\Omega^1(\mathfrak{g}_{P_\mathbb{C}})),$$

given by

$$F(u,\eta) = u(A+\eta) - A = u\eta u^{-1} - (\nabla_A u)\, u^{-1}.$$

Given the restriction on the dimension of the manifold, we have an embedding $L^2_2 \to L^p$ for all $p < \infty$ and $L^2_1 \to L^4$. At a point (u,η) the derivative in the direction (ξ,ν) is given by

$$F'_{u,\eta}(\xi,\nu) = u[\xi,\eta]u^{-1} + u\nu u^{-1} - u\,(\nabla_{\mathfrak{A}}\xi)\,u^{-1} + (\nabla_{\mathfrak{A}} u)\,[\xi, u^{-1}].$$

Note that this map $F'_{(u,\eta)} : L^2_2(\Omega^0(\mathfrak{g}_P)) \times L^2_1(\Omega^1(\mathfrak{g}_{P_\mathbb{C}})) \to L^2(\Omega^1(\mathfrak{g}_{P_\mathbb{C}}))$ is bounded since $L^2_2 \hookrightarrow L^8$ and $L^2_1 \hookrightarrow L^4$. For example, we can bound

$$\begin{aligned}\|u[\xi,\eta]u^{-1}\|^2_{L^2} &\le \||u|^4|\xi|^2\|_{L^2}\||\eta|^2\|_{L^2}\\ &\le \||u|^8\|_{L^2}\||\xi|^4\|_{L^2}\|\eta\|^2_{L^4}\\ &\le C\|u\|^4_{L^2_2}\|\xi\|^2_{L^2_2}\|\eta\|^2_{L^2_1}\end{aligned}$$

If we consider the derivative at $(I,0)$ we get the map

$$F' : \Omega^0(\mathfrak{g}_P) \oplus \ker \mathrm{Re} d_A^* \to \Omega^1(\mathfrak{g}_{P_\mathbb{C}}), \quad F' : (\xi,\eta) \mapsto -d_A\xi + \eta.$$

This map is surjective because of the decomposition

$$\begin{aligned}L^2(\Omega^1(\mathfrak{g}_{P_\mathbb{C}})) &= \mathrm{Im}\,(d_A) \oplus \mathrm{Im}\,(d_A)^\perp\\ &= \mathrm{Im}\,(d_A) \oplus \ker \mathrm{Re}\; d_A^*\end{aligned}$$

where we view $d_A : L^2_1 \to L^2$. Note that d_A is bounded and hence has closed image and so it makes sense to consider its perpendicular subspace. This shows that F is locally surjective at $(I,0)$ and so for $\|\mathfrak{A} - \mathfrak{B}\|_{L^2_1}$ sufficiently small constant depending on $\mathfrak{A}$ we can find a u and α solving $F(u,\alpha) = 0$ using the implicit function theorem. Moreover this gauge transformation is unique. □

We end this section by considering gauge conditions on manifolds with boundary. There is an added complication, as compared to the closed manifold case because the gauge conditions above are defined using the adjoint of the derivative operator. As with the condition $d_A^*\phi = 0$, one needs to be careful when defining the boundary conditions. This relates to whether or not one allows gauge transformations that are non-trivial on the boundary. For example, one may impose certain boundary as part of the equation that one is working with, and impose less gauge freedom, or allow more gauge freedom.

Definition 3.14. Let A be a connection on a principal G–bundle $P_{\mathbb{C}} \to X$, where X is a compact manifold with boundary. We define $\mathcal{G}_0$ to be the space of gauge transformations that are trivial on the boundary.

As before, the infinitesimal action of the gauge group $\mathcal{G}_0$ at a connection A is given by

$$\xi.(A) = -d_A\xi.$$

Note that if $\xi \in \mathrm{Lie}(\mathcal{G}_0)$ then we have that $\xi_{\partial X} = 0$. When considering the adjoint of d_A we can consider the map

$$\eta_\xi : \Omega^1(\mathfrak{g}_P) \to \Omega^{n-1}(\mathfrak{g}_P), \quad \eta_\xi(v) = \langle \xi, *v \rangle,$$

where $\langle \cdot, \cdot \rangle$ denotes taking the value of the inner product on $\mathfrak{g}_P$ of each component of the form. By Stokes' theorem,

$$0 = \int_{\partial X} \eta = \int_X d\eta = \int_X \langle d_A\xi, *v \rangle + \langle \xi, d_A^* v \rangle.$$

Hence we do not need to impose any boundary conditions for the Coulomb gauge condition.

If we instead allow gauge transformations that are non-trivial on the boundary, then the integral over the boundary will not vanish in general. In this case, we can impose the boundary condition

$$*v|_{\partial X} = 0,$$

in order to consider a subspace that is L^2 orthogonal to the infinitesimal image of the gauge transformations. That is, $v(\partial_\nu) = 0$, where ∂_ν is the normal to the boundary.

Remark 3.15. In the non-compact case, this is particularly relevant in Floer theory, where one studies connections on manifolds of the form $\mathbb{R} \times Y$ that converge to flat connections as $t \to \pm\infty$. The natural notion of gauge equivalence of flat connections means that in the setup of the theory means that one should allow gauge transformations that are non-trivial at infinity. Hence the Coulomb gauge condition in this case should also include

$$\lim_{t\to\infty}\int_{\{t\}\times Y}\langle\xi, *\eta\rangle = 0$$

for any $\xi \in L^2_1(\mathbb{R}\times Y)$.

3.3 The Gauge Group $\mathcal{G}_{\mathbb{C}}$

For a given connection $\mathfrak{A} = A + i\phi$ on a complexified bundle $P_{\mathbb{C}}$, we can consider its orbit under the action of automorphisms of $P_{\mathbb{C}}$ as well as the automorphisms of P. If we act by an element $g \in \mathcal{G}$ then we have

$$g^*(A + i\phi) = g^*A + i\mathrm{Ad}_g\phi.$$

In particular, the $\mathcal{G}$ action preserves $\|\phi\|_{L^2}$. This is not true for a general element of $\mathcal{G}_{\mathbb{C}}$. In fact, we have the following theorem due to Corlette, [Cor88].

Proposition 3.16. Let $A+i\phi$ be a connection on a complexified bundle $P_{\mathbb{C}} \to X$, such that there are no covariantly constant sections of $\mathfrak{g}_{P_{\mathbb{C}}}$. Then there exists at most one $G_{\mathbb{C}}$ orbit for which $\|\phi\|_{L^2}$ is minimised.

Proof. See [Cor88, Proposition 2.3]. □

Remark 3.17. The previous proposition does not make any assumptions on the compactness or boundary properties of the manifold X.

Lemma 3.18. Let X be a manifold with boundary and consider the group $\mathcal{G}_{\mathbb{C}0}$ of $P_{\mathbb{C}}$ automorphisms that are the identity on the boundary. Fix a connection $A+i\phi$ and consider the functional on $\mathcal{G}_{\mathbb{C}}$ given by

$$g \mapsto \|\phi_g\|^2_{L^2}$$

then the derivative of this map at the identity in the direction of $\xi + i\eta$ is given by

$$\langle\eta, d_A^*\phi\rangle,$$

and we impose no boundary conditions on ϕ. If we consider all gauge transformations in $\mathcal{G}_{\mathbb{C}}$, then, in line with the discussion at the end of the previous section, $A + i\phi$ is a critical point of the functional provided

$$d_A^*\phi = 0, \text{ and } \phi(\partial_n) = 0,$$

where ∂_n is the normal at the boundary.

Proof. The infinitesimal action by $\xi + i\eta$ on the connection $A + i\phi$ is given by

$$(\xi + i\eta).(A + i\phi) = -d_{A+i\phi}(\xi + i\eta).$$

Hence we have that, along the path of gauge transformations $g_t = \exp(t\,(\xi + i\eta))$

$$\begin{aligned}\frac{d}{dt}\|\phi_t\|_{L^2} &= 2\int_X \langle \mathrm{Im}(d_{A+i\phi}\xi + i\eta), \phi\rangle \\ &= 2\langle d_A\eta, \phi\rangle - \langle [\phi, \xi], \phi\rangle\end{aligned}$$

Note that $\langle [\phi, \eta], \phi\rangle = 0$, and hence

$$\frac{d}{dt}\|\phi\|_{L^2} = 2\int_X \langle \eta, d_A^*\phi\rangle,$$

provided the integration by parts boundary conditions vanish. See the end of the previous section for details. □

An alternative way of phrasing this is the following.

Lemma 3.19. The equation $d_A^*\phi = 0$ is the Euler-Lagrange equation for the functional

$$\phi \mapsto \|\phi\|_{L^2}^2.$$

as $A + i\phi$ varies in the orbit given by the $SL(2;\mathbb{C})$ gauge transformations. In other words, $d_A^*\phi = 0$ if and only the representation of the $SL(2;\mathbb{C})$ connection is such that connection.

3.3.1 $G_{\mathbb{C}}/G$ and the space of G reductions

In this section we consider how one can describe the space of choices of G–reductions of a $G_{\mathbb{C}}$–bundle in terms of sections of associated bundles. This will result, in Chapter 4, in being able to phrase the condition of finding a $\mathcal{G}_{\mathbb{C}}$ equivalence class for a given flat connection in terms of finding a harmonic map into the space $G_{\mathbb{C}}/G$.

Definition 3.20. A representation $\rho : G \to \mathrm{GL}(n;\mathbb{C})$ is reductive if any ρ invariant subspace of $\mathbb{C}^n$ has a ρ invariant complement.

Given a $G_{\mathbb{C}}$-bundle, we will need to consider the space of possible G reductions of that bundle.

Definition 3.21. The space $G_{\mathbb{C}}/G$ is the space of left cosets of G in $G_{\mathbb{C}}$. We can give it a metric by picking an invariant metric on G, h, and defining

$$h_g(X,Y) = h_e((g^{-1})_*X, (g^{-1})_*Y)$$

This metric is G invariant and hence defines a metric on the quotient.

Lemma 3.22. The space $G_{\mathbb{C}}/G$ is diffeomorphic to

$$TG/G \cong \mathfrak{g}.$$

Proof. This follows from $G_{\mathbb{C}} \cong TG$. $\square$

Proposition 3.23. There is a one to one correspondence between G reductions of $P_{\mathbb{C}}$ and sections of the associated bundle

$$E = P_{\mathbb{C}} \times_\rho G_{\mathbb{C}}/G.$$

where ρ is left multiplication.

Proof. A section of E is a choice of equivalence class $[p, xG]$ at each point, for $p \in P_{\mathbb{C}}, xG \in G_{\mathbb{C}}/G$. Given such a section, u, we can select a subbundle of $P_{\mathbb{C}}$ by picking all the elements $p \in P_{\mathbb{C}}$ such that

$$u(x) = [p, e.G]\,.$$

To go the other way around, if we have $P \subseteq P_{\mathbb{C}}$, then we can write $u(x) = [p, e.G]$ where $p \in P_x$. $\square$

We now consider the geometry of the symmetric spaces $G_{\mathbb{C}}/G$. In particular, their curvature properties and boundary structure at infinity under certain compactifications.

Lemma 3.24. The space $(G_{\mathbb{C}}/G, h_g(X,Y))$ has negative sectional curvature.

Proof. This follows from the formula for the sectional curvature of a Lie group. We have that at the origin in $G_{\mathbb{C}}$,

$$K(X,Y) = -|[X,Y]|^2$$

for orthogonal $X, Y \in i\mathfrak{g} \cong \mathfrak{g}$. It follows that the sectional curvature at the identity is negative, since G is compact and semi-simple. However, $G_{\mathbb{C}}$ acts transitively by isometries. Hence there is an isometry sending each point of $G_{\mathbb{C}}/G$ to the origin and so $K(X, Y)$ has negative sectional curvature. $\square$

We will be studying harmonic maps into the spaces $G_{\mathbb{C}}/G$. As such, when studying singular behaviour of such maps, it is important to have an understanding of the asymptotic geometry of these spaces. We now describe the boundary at infinity. For more details on these constructions, see [Ebe97] and [BJ06].

Definition 3.25 (Geodesic Boundary at Infinity of a Symmetric Space of Non–Compact Type)**.** Denote by $S \subseteq TM$ the sphere bundle of unit vectors in the tangent bundle. Let $p, q \in M = G_{\mathbb{C}}/G$ and let $v \in S_pM, w \in S_qM$ be unit tangent vectors. Then we say that $v \sim w$ if and only if the geodesics $\exp(tv)$ and $\exp(tw)$ satisfy

$$\sup_t d(\exp(tv), \exp(tw)) \leq C$$

for some $C \in \mathbb{R}$. We define the geodesic boundary at infinity to be the space

$$M_\infty = SM/\sim .$$

Definition 3.26. Let M be a symmetric space of non-compact type. Its compactification by the geodesics at infinity is the space

$$\bar{M} = M \cup M_\infty$$

Pick a base point $p \in M$. We give $\bar{M}$ the topology such that if x_n is a sequence in M with $x_n \to \infty$, $x_n \to [\exp_p(tv)]$ if and only if

$$v_n \to v$$

where $\exp_p(tv_n)$ is the geodesic from p to x_n.

As an example it is useful to understand the above construction for $SL(n; \mathbb{C})/SU(n)$.

Example 3.27. The boundary at infinity of $SL(n; \mathbb{C})/SU(n)$ can be identified with the space of flags in $\mathbb{C}^n$.

A key result is the following.

Proposition 3.28 ([BJ06, Proposition Proposition I.2.6])**.** For any point in M_∞, its stabilizer in $G_{\mathbb{C}}$ is a parabolic subgroup. Conversely, every proper parabolic subgroup of $G_{\mathbb{C}}$ is the stabilizer of some point in M_∞.

We finish this section by considering connections, and in particular connections on the bundle $P_{\mathbb{C}}$ and how they relate to connections on P that are defined by a section of the bundle with fibre $G_{\mathbb{C}}/G$. These definitions and results are fairly standard, but they have been included so that it is clear what mathematical objects we are talking about. In the literature there is sometimes a bit of imprecision when it comes to describing connections. For a comprehensive treatment of the ideas given here, see [Hay19] and also [DK97, Chapter 2].

Proposition 3.29. Let $\mathfrak{A}$ be a connection on the bundle $P_{\mathbb{C}}$, and let $u \in \Gamma(E)$ be a section that defines a G reduction $P \subseteq P_{\mathbb{C}}$. Then we can decompose $\mathfrak{A}$ as

$$\mathfrak{A} = A + i\phi$$

where $\phi = d_{\mathfrak{A}}u$, and A is a connection on P.

Proof. The exterior derivative of the section u with respect to the connection, may be seen as a section

$$d_{\mathfrak{A}}u \in \Omega^1(\mathfrak{g}_P).$$

To prove the proposition we just need to show that

$$A = \mathfrak{A} - d_{\mathfrak{A}}u$$

defines a connection on P. This is tautological. $\square$

Remark 3.30. The condition $d_A^*\phi = 0$ can be reformulated as the twisted harmonic condition

$$d_{\mathfrak{A}}^* d_{\mathfrak{A}} u = 0.$$

In the flat connection case, we can lift to the universal cover and gauge transform to the trivial bundle using parallel transport. Then the condition $d_{\mathfrak{A}}^* d_{\mathfrak{A}} u = 0$ becomes an equivariant harmonic map condition, where the fundamental group acts on u via the holonomy representation.

3.4 The Complex Yang-Mills equations

In this section we consider Yang-Mills connections on complexified bundles. We first recall the theory in the real case.

Definition 3.31. A connection, A, on a principal G–bundle $P \to X$ is Yang-Mills if

$$d_A^* F_A = 0.$$

This is the Euler-Lagrange equation for the energy functional $\|F_A\|^2_{L^2(X)}$, and the equation is elliptic modulo gauge transformations.

Definition 3.32. Let $P_{\mathbb{C}} \to Y$ be a complexified bundle. We define the complex Yang-Mills energy functional for connections on $P_{\mathbb{C}}$ as

$$\mathrm{YM}(A + i\phi) = \|F_A - \phi \wedge \phi\|^2_{L^2} + \|d_A\phi\|^2_{L^2}.$$

Note that the complex Yang-Mills functional is not invariant under complex gauge transformations, as the inner product is not invariant, but it is invariant under real gauge transformations.

Lemma 3.33. The equation $d^*_{A+i\phi}F_{A+i\phi}$ is the Euler-Lagrange equation for the complex Yang-Mills functional.

Proof.
$$\frac{d}{dt}\mathrm{YM}(A + i\phi + t(a + i\varphi)) = 2\langle F_{A+i\phi}, d_{A+i\phi}(a + i\varphi)\rangle.$$
Hence $A + i\phi$ is Yang-Mills if and only if $d^*_{A+i\phi}F_{A+i\phi} = 0$. □

The linearisation of the map

$$F : \mathcal{A} \to \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}}), F : A + i\phi \mapsto d^*_{A+i\phi}F_{A+i\phi}$$

is $\mathbb{C}$ linear, and hence if $F^{-1}(0)$ is smooth then it is a complex submanifold of $\mathcal{A}_{\mathbb{C}}$. However, F is not elliptic modulo real gauge transformations or invariant under complex gauge transformations. To remedy this situation we can instead consider the space

$$F^{-1}(0) \cap \mu^{-1}(0)/\mathcal{G},$$

where $\mu(\mathfrak{A}) = d^*_A\phi$. That is, we consider the system equations

$$d^*_{A+i\phi}F_{A+i\phi} = 0, d^*_{A+i\phi}\phi = 0.$$

Definition 3.34. A connection $A + i\phi$ on a complexified bundle is Yang-Mills if it satisfies the complex Yang-Mills equations

$$d^*_{A+i\phi}F_{A+i\phi} = 0, d^*_{A+i\phi}\phi = 0.$$

Example 3.35. Two important examples of solutions to the complex Yang-Mills equation are the following.

- Flat connections $A + i\phi$ with $d_A^*\phi = 0$. If $F_{A+i\phi} = 0$ then $d_{A+i\phi}^* F_{A+i\phi} = 0$.
- Solutions to the Kapustin-Witten equations. If $\left(e^{i\theta} F_{A+i\phi}\right)^+ = 0$ then

  $$d_{A+i\phi}^* F_{A+i\phi} = *d_{A+i\phi} * F_{A+i\phi}$$

  The Kapustin-Witten equations imply that $*F_{A+i\phi} = e^{2i\theta} F_{A+i\phi}$, and hence

  $$d_{A+i\phi}^* F_{A+i\phi} = e^{2i\theta} * d_{A+i\phi} F_{A+i\phi}$$

  This is equal to zero by the Bianchi identity. The other condition, $d_A^*\phi = 0$, is imposed by the Kapustin-Witten equations.

Proposition 3.36. The Complex Yang-Mills equation is elliptic modulo gauge transformations.

Proof. Picking a background connection $A + i\phi$, which solves $d_A^*\phi = 0$, we can write a general connection as

$$\mathfrak{A} = A + i\phi + a + i\varphi.$$

We impose the gauge condition

$$\mathrm{Re}(d_{A+i\phi}^*(a + i\varphi)) = 0.$$

Note that

$$\begin{aligned}\mathrm{Re}\left(d_{A+i\phi}^*(a + i\varphi)\right) + i d_{A+a}^*(\phi + \varphi) &= d_A^* a + *[\phi, *\varphi] + d_A^* i\varphi - i * [a, *\phi + *\varphi] \\ &= d_{A+i\phi}^*(a + i\varphi) - i * [a, *\varphi].\end{aligned}$$

Hence, up to a zero order quadratic term, the gauge condition and moment map condition are together equivalent to the condition $d_{A+i\phi}^*(a + i\varphi) = 0$. Hence we have that the leading order terms in the Complex Yang Mills equations, working in a trivialisation, are then

$$d^* d(a + i\varphi) = 0, d^* a = 0, d^*\varphi = 0.$$

Hence the principal symbol of the Complex Yang-Mills is the same as two copies of the Yang-Mills equation. It follows that the system is elliptic. $\square$

3.4.1 Regularity Results for Complex Yang-Mills Connections

In this section we review some regularity results for solutions of the complex Yang-Mills equations. We will apply these results to solutions of the Kapustin-Witten equations. This section combines results and methods of proof from [GU12] and [Tau13]. We start with a Weitzenböck formula.

Proposition 3.37. Let B be a connection on $P \to X$. Then for any smooth section $\varphi \in \Omega^1(\mathfrak{g}_P)$ we have that

$$\nabla_A^* \nabla_A \varphi + \varphi \circ \mathrm{Ricci} + F_A \lrcorner \varphi = (d_A^* d_A + d_A d_A^*)\, \varphi, \tag{3.3}$$

where at a point $p \in X$, $\varphi \circ \mathrm{Ricci}(Z) = \varphi(\sum_i R(Z, e_i) e_i)$, and $F_A \lrcorner \varphi(X) = \sum_i [F_A(X, e_i), \varphi(e_i)]$ for $Z \in \Gamma(TX)$, R the Riemann curvature operator, and e_i an orthonormal frame.

Proof. See, for example, [BL81, Theorem 3.3]. □

Corollary 3.38. For a solution $A + i\phi \in \mathcal{A}_{\mathbb{C}}$ complex Yang-Mills equations, we have

$$\nabla_A^* \nabla_A \phi + \phi \circ \mathrm{Ricci}(X) + (\phi \wedge \phi) \lrcorner \phi = 0.$$

Proof. For this we take the imaginary part of the Yang-Mills equation

$$d_{A+i\phi}^* F_{A+i\phi} = 0,$$

to get that

$$d_A^* d_A \phi - *[\phi, *F_A - *\phi \wedge \phi] = 0.$$

Working in an orthonormal frame, note that for a two form $\omega \in \Omega^2(\mathfrak{g}_P)$ and one form $\varphi \in \Omega^1(\mathfrak{g}_P)$, we have that

$$\begin{aligned} *[\varphi, *\omega] &= *\frac{1}{2} \varepsilon_{jklm} [\varphi_i, \omega_{jk}] dx^i \wedge dx^l \wedge dx^m \\ &= \frac{1}{2} \varepsilon_{ilmn} \varepsilon_{jklm} [\varphi_i, \omega_{jk}] dx^n \\ &= \sum_i [\varphi_i, \omega_{in}] dx^n, \end{aligned}$$

where we have used the Einstein summation convention and ε_{ijkl} denotes the Levi-Civita symbol. Hence $*[\varphi, *\omega](X) = \sum_i [\omega(X, e_i), \varphi(e_i)]$, and so

$$*[\phi, *F_A - *\phi \wedge \phi] = (F_A - \phi \wedge \phi) \lrcorner \phi.$$

Recall that a solution to the complex Yang-Mills equations also satisfies $d_A^*\phi = 0$. Putting all this together gives the result. $\square$

Theorem 3.39 ([GU12], Corollary 4.7)**.** Let $A + i\phi$ be a solution to the complex Yang-Mills equations on a four manifold X. Then there exists a constant C, independent of A, ϕ, such that

$$\|\nabla_A \phi\|^2_{L^2(X)} + \frac{1}{2}\|[\phi, \phi]\|^2_{L^2(X)} \leq C\|\phi\|^2_{L^2(X)}$$

Proof. We have

$$\begin{aligned}\langle \nabla_A \phi, \nabla_A \phi\rangle &= \langle \nabla_A^* \nabla_A \phi, \phi\rangle \\ &= -\langle \phi \circ \mathrm{Ricci} + \phi \wedge \phi \lrcorner \phi, \phi\rangle.\end{aligned}$$

From the proof of the previous corollary we have that

$$\phi \wedge \phi \lrcorner \phi = *[*\phi \wedge \phi, \phi].$$

Hence

$$\langle \phi \wedge \phi \lrcorner \phi, \phi\rangle = \langle *[*\phi \wedge \phi, \phi], \phi\rangle = -\frac{1}{2}|[\phi, \phi]|^2.$$

$\square$

Corollary 3.40. Suppose that (X, g) is Ricci positive. Then any solution to the complex Yang-Mills equation is real, that is, $\phi = 0$.

Lemma 3.41. Let $A + i\phi$ be a solution to the complex Yang-Mills equations. Then the norm of ϕ, $|\phi|$ satisfies

$$\frac{1}{2}\Delta|\phi|^2 + |\nabla_A \phi|^2 + \frac{1}{2}|[\phi, \phi]|^2 - \langle \phi \circ \mathrm{Ricci}, \phi\rangle = 0. \tag{3.4}$$

Proof. This is a direct calculation. We have, in local orthonormal coordinates

$$\Delta\langle \phi, \phi\rangle = -2\sum_i \langle \nabla_i^A \nabla_i^A \phi, \phi\rangle - 2\sum_i \langle \nabla_i^A \phi, \nabla_i^A \phi\rangle$$

and so

$$\frac{1}{2}\Delta|\phi|^2 + |\nabla_A \phi|^2 + \langle \nabla_A^* \nabla_A \phi, \phi\rangle = 0.$$

The result then follows from the Weitzenböck formula Equation (3.3). $\square$

We now prove a local version of Theorem 3.39.

Theorem 3.42. Let $p \in X$ and $B_p(r)$ be a ball around p of radius r, and let $B_p(R) \supset B_p(r)$ be a ball of larger radius. Then there exists a constant C such that

$$\|\nabla_A \phi\|^2_{L^2(B_p(r))} + \frac{1}{2}\|[\phi, \phi]\|^2_{L^2(B_p(r))} \leq C\|\phi\|^2_{L^2(B_p(R))}$$

Proof. Let $f : B_p(R) \to \mathbb{R}$ be a smooth cutoff function, such that $f = 0$ on an open neighbourhood of $\partial B_p(R)$ and $f = 1$ on $B_p(r)$. Then we have that

$$\begin{aligned}\int_{B_p(r)} |\nabla_A \phi|^2 + \frac{1}{2}|[\phi, \phi]|^2 &\leq \int_{B_p(R)} -f\Delta|\phi|^2 + \langle \phi \circ \text{Ricci}, \phi\rangle \\ &\leq \int_{B_p(R)} -(\Delta f)|\phi|^2 + \langle \phi \circ \text{Ricci}, \phi\rangle \leq C\|\phi\|^2_{L^2(B_p(R))}.\end{aligned}$$

□

Proposition 3.43. We have bounds of the form

$$|\phi|_{L^\infty(B_p(r))} < C\|\phi\|_{L^2(U)}$$

where U is an open set containing the closure of $B_p(r)$. The constant depends on the set U and $B_p(r)$.

Proof. See [Tau13, Equation 2.10]. □

Theorem 3.44. Let $P_{\mathbb{C}} \to X$ be a principal $G_{\mathbb{C}}$ bundle on a compact manifold, X of dimension 3 or 4. Let $P \subseteq P_{\mathbb{C}}$ be a choice of G reduction. Let $\mathfrak{A}_0 = A + i\phi$ be a smooth background connection on $P_{\mathbb{C}}$ with respect to the reduction P, and suppose that the connection $\mathfrak{A} = A + i\phi + a + i\varphi$ satisfies the Complex Yang Mills equation with $a + i\varphi$ small in L^2_1. Then there exists an L^2_2 gauge transformation, ψ such that $\psi^*(\mathfrak{A}) - \mathfrak{A}_0$ is C^∞ and in the complexified Coulomb gauge relative to $\mathfrak{A}_0$.

Proof. This is a standard gauge theory bootstrapping argument, but we r since we are using the complexified Coulomb gauge condition. By Proposition 3.13 we can find a gauge transformation such that $a + i\varphi = \mathfrak{A} - \mathfrak{A}_0$ satisfies the complexified Coulomb gauge condition. Hence we can assume that

$$\begin{aligned} d^*_{\mathfrak{A}_0}(a + i\varphi) &= \text{Re}(d^*_{\mathfrak{A}_0}(a + i\varphi)) + id_{\mathfrak{A}}(\phi + \varphi) - d^*_A\phi \pm i * [a, *\varphi] \\ &= d^*_A\phi \pm i * [a, *\varphi]\end{aligned}$$

with the sign depending on the dimension. This implies that, for example,

$$\|d^*_{\mathfrak{A}}(a + i\varphi)\|_{L^2_1} \leq C + \|\{a, \varphi\}\|_{L^2_1}.$$

Once we've fixed the gauge then we can consider elliptic estimates using $d_{\mathfrak{A}} \oplus d^*_{\mathfrak{A}}$. We have

$$\begin{aligned}
\|F_{\mathfrak{A}}\|_{L^2_1} &\leq C\left(\|d_{\mathfrak{A}_0} F_{\mathfrak{A}}\|_{L^2} + \|d^*_{\mathfrak{A}_0} F_{\mathfrak{A}}\|_{L^2} + \|F_{\mathfrak{A}}\|_{L^2}\right) \\
&\leq C\left(\|d_{\mathfrak{A}} F_{\mathfrak{A}}\|_{L^2} + \|d^*_{\mathfrak{A}} F_{\mathfrak{A}}\| + \|\{a + i\varphi, F_{A+i\varphi}\}\|_{L^2} + \|F_{\mathfrak{A}}\|_{L^2}\right) \\
&\leq C\left(\|a + i\varphi\|_{L^2_1}\|F_{A+i\varphi}\|_{L^2_1} + \|F_{\mathfrak{A}}\|_{L^2}\right)
\end{aligned}$$

Hence, for $\|a + i\varphi\|_{L^2_1} \leq \frac{1}{C}$, we have

$$\|F_{\mathfrak{A}}\|_{L^2_1} \leq C'\|F_{\mathfrak{A}}\|_{L^2}$$

We can then bound $a + i\varphi$ in L^2_2 using the estimate

$$\begin{aligned}
\|a + i\varphi\|_{L^2_2} &\leq C\left(\|d_{\mathfrak{A}_0}(a + i\varphi)\|_{L^2_1} + \|d^*_{\mathfrak{A}_0}(a + i\varphi)\|_{L^2_1} + \|a + i\varphi\|_{L^2}\right) \\
&\leq C\left(\|F_{\mathfrak{A}}\|_{L^2_1} + \|F_{\mathfrak{A}_0}\|_{L^2_1} + \|(a + i\varphi) \wedge (a + i\varphi)\|_{L^2_1} + \|d^*_A\phi\|_{L^2_1} + \|\{a, \varphi\}\|_{L^2_1}\right).
\end{aligned}$$

For the terms of the form $\|a + i\varphi, a + i\varphi\|_{L^2_1}$, we can write

$$\begin{aligned}
\|\{a + i\varphi, a + i\varphi\}\|_{L^2_1} &\leq C\left(\|\{\nabla_{\mathfrak{A}_0}(a + i\varphi), (a + i\varphi)\}\|_{L^2} + \|\{a + i\varphi, a + i\varphi\}\|_{L^2}\right) \\
&\leq C\left(\|a + i\varphi\|_{L^2_2}\|a + i\varphi\|_{L^2_1} + \|a + i\varphi\|^2_{L^2_1}\right).
\end{aligned}$$

Hence, for $\|a + i\varphi\|_{L^2_2} \leq \frac{1}{C}$ we can bound the L^2_2 norm of $a + i\varphi$. Now we can repeat this procedure to get an L^2_2 bound on the curvature which will lead to an L^2_3 bound on $a + i\varphi$. Note that we have

$$\begin{aligned}
\|\{a + i\varphi, F_{\mathfrak{A}}\}\|_{L^2_1} &\leq C\left(\|\{\nabla_{\mathfrak{A}_0}(a + i\varphi), F_{\mathfrak{A}}\}\|_{L^2} + \|\{a + i\varphi, \nabla_{\mathfrak{A}_0} F_{\mathfrak{A}}\}\|_{L^2} + \|\{a + i\varphi, F_{\mathfrak{A}}\}\|_{L^2}\right) \\
&\leq C\left(\|a + i\varphi\|_{L^2_2}\|F_{\mathfrak{A}}\|_{L^2_1} + \|a + i\varphi\|_{L^2_1}\|F_A\|L^2_2 + \|a + i\varphi\|_{L^2_1}\|F_{\mathfrak{A}}\|_{L^2_1}\right)
\end{aligned}$$

and so we have

$$\|F_{\mathfrak{A}}\|_{L^2_2} \leq C'\|F_{\mathfrak{A}}\|_{L^2_1}$$

provided $\|a + i\varphi\| \leq \frac{1}{C}$. Now we can estimate

$$\|a + i\varphi\|_{L^2_3} \leq C\left(\|F_{\mathfrak{A}}\|_{L^2_2} + \|\{a + i\varphi, a + i\varphi\}\|_{L^2_2} + \|d^*_A\phi\|_{L^2_2}\right).$$

and as before we can bound

$$\|\{a+i\varphi, a+i\varphi\}\|_{L^2_2} \leq C\sum_{j=0}^{2}\|\{\nabla^j_{\mathfrak{A}}(a+i\varphi), \nabla^{2-j}_{\mathfrak{A}}(a+i\varphi)\}\|_{L^2} + \|\{a+i\varphi, a+i\varphi\}\|_{L^2_1}.$$

The only terms in the sum that cause any trouble are the ones proving second derivatives. For these we can again estimate

$$\|\{\nabla^2_{\mathfrak{A}}(a+i\varphi), a+i\varphi\}\|_{L^2} \leq C\|a+i\varphi\|_{L^2_3}\|a+i\varphi\|_{L^2_1}.$$

And so, for $\|a+i\varphi\|_{L^2_1}$ sufficiently small we have

$$\|a+i\varphi\|_{L^2_3} \leq C'\left(\|F_{\mathfrak{A}}\|_{L^2_2} + \|a+i\varphi\|^2_{L^2_2}\right).$$

From the Sobolev embedding theorem, we can now deduce that $\|a+i\varphi\|_{L^\infty}$ is bounded. Now we have the following inductive step. Suppose that $a+i\varphi$ is bounded in L^2_k, for $k \geq 3$ then it is bounded in L^2_{k+1}. To prove this, note that

$$\|F_{\mathfrak{A}}\|_{L^2_k} \leq C\left(\|a+i\varphi, F_{\mathfrak{A}}\|_{L^2_k} + \|F_{\mathfrak{A}}\|_{L^2}\right)$$

using the elliptic estimates for the operator $d_{\mathfrak{A}} \oplus d^*_{\mathfrak{A}}$. Expanding the derivative of the product out, the only term we need to consider

$$\|\{a+i\phi, \nabla^{k-1}_{\mathfrak{A}}F_{\mathfrak{A}}\}\|_{L^2}$$

as all the other terms can be bounded by the using the embedding $L^2_j \hookrightarrow L^4_{j-1}$. But to bound this term we can use the L^∞ embedding. Hence we have $F_{\mathfrak{A}}$ is bounded in L^2_k. We can then estimate

$$\|a+i\varphi\|_{L^2_{k+1}} \leq C\left(\|F_{\mathfrak{A}}\|_{L^2_k} + \|\{a+i\varphi, a+i\varphi\}\|_{L^2_k} + \|d^*_A\phi\|_{L^2_k}\right).$$

Again we need can bound all the terms of the form

$$\|\nabla^l(a+i\varphi), \nabla^{k-l}(a+i\varphi)\|_{L^2}$$

either using the L^∞ estimates when $l \geq 3$ or $k-l \leq 3$. The only situation where both of these do not hold is when $k=3$ and $l=1,2$ or when $k=4$, $l=2$. In both of these cases we can then use the $L^2_j \hookrightarrow L^4_{j-1}$ Sobolev embedding. Hence we can find constants C_k, D_k such that for $\|a+i\varphi\|_{L^2_1}$ sufficiently small and a solution of the complex Yang-Mills equations we have

$$\|a+i\varphi\|_{L^2_k} \leq C_k, \|F_{\mathfrak{A}}\| \leq D_k$$

after applying a L^2_2 gauge transformation such that $\mathrm{Re}(d^*_{\mathfrak{A}_0}a+i\varphi)=0$.

$\square$

3.5 Flat $G_{\mathbb{C}}$ connections

In later sections, it will be important to have a good understanding of flat $G_{\mathbb{C}}$ connections on three manifolds. In the same way that flat connections are critical points of the Chern-Simons functional for compact groups, flat $G_{\mathbb{C}}$ connections satisfying $d_A^*\phi = 0$ are critical points of the complexified Chern-Simons functional. This is discussed in Chapter 5. In this section we recall some on parallel transport and the holonomy representation of flat connections, and the deformation theory for flat connections. We include proofs for standard results since we will need variations on the formulae used later.

3.5.1 Parallel Transport and Holonomy

In this section we record some results on the holonomy of connections on complexified bundles that we will use when studying perturbations of the Chern-Simons functional.

Definition 3.45 (Parallel Transport)**.** Let A be a connection on a principal bundle $P \to X$, let $\gamma : [0, 1] \to X$ be a path and let v be a point in an associated bundle

$$\pi : P \times_\rho V = E \to X.$$

such that $\pi(v) = \gamma(0)$. The parallel transport of v along γ is the solution $u(t)$ of the ODE

$$\frac{du}{dt} + A(\dot\gamma(t))u = \nabla_{\dot\gamma} u = 0, \quad u(0) = v. \tag{3.5}$$

Definition 3.46 (Parallel Transport Map)**.** Given a curve γ, connection A, and a bundle E, as in Definition 3.45, we define the parallel transport map to be

$$U : E_{\gamma(0)} \to E_{\gamma(1)}, \quad v \mapsto u(1),$$

where we map each element $v \in E_{\gamma(0)}$ to the value of the solution of the parallel transport at the end of the curve. We also define the partial parallel transport maps to be the values

$$U(t) : E_{\gamma(0)} \to E_{\gamma(t)}, \quad v \mapsto u(t).$$

Proposition 3.47. Let A be a flat connection on a principal H bundle $Q \to Y$. Let $H : [0,1] \times S^1 \to Y$ be a homotopy of loops based at $p \in Y$. Then the holonomy of A around the loop $H(t, \cdot)$ is independent of s.

Proof. Let $U(t,s) : E_{H(0,s)} \to E_{H(t;s)}$ be the parallel transport map for the connection A on an associated vector bundle E. We will differentiate this map with respect to s. Note that $U(t,s)$ satisfies

$$\frac{\partial}{\partial t}U(t,s) + a_t.U(t,s) = 0.$$

where $A = \theta + a$, is the pullback of the connection A to the (t,s) plane, θ is the trivial connection the pullback bundle and $a_t = a(\partial_t)$. Hence we have

$$\frac{\partial}{\partial s}\left(\frac{\partial}{\partial t}U(t,s) + a_t(\partial_t H).U(t,s)\right) = \frac{\partial^2}{\partial s \partial t}U(t,s) + \left(\frac{\partial}{\partial s}a_t\right)U(t,s) + a_t\frac{\partial}{\partial s}U(s,t)$$

Changing the order of differentiation, we have that the derivative $D = \frac{d}{ds}U(t,s)$ satisfies the differential equation

$$\frac{\partial}{\partial t}D + a_t D = -\frac{\partial a_t}{\partial s}U(t,s).$$

This has solution

$$D(t,s) = \frac{\partial}{\partial s}U(t,s) = -\int_0^t U_s(\tau;t)\frac{\partial a_t}{\partial s}U_s(\tau;0)d\tau,$$

where the notation $U_s(b;a) : E_a \to E_b$ is the parallel transport along the curve

$$t \mapsto H(t,s), t \in [a,b].$$

Note that the connection A is flat, and so, in particular

$$0 = \frac{\partial a_t}{\partial s} - \frac{\partial a_s}{\partial t} + [a_t, a_s].$$

Hence,

$$U_s(\tau;t)\frac{\partial a_t}{\partial s}U_s(\partial;0) = U_s(1;\tau)\left(d_A a_s(\partial_t)\right)U_s(\tau;0).$$

Note that this is the derivative

$$\frac{d}{d\tau}(U(1;\tau)a_s U(\tau;0)) = U(1;\tau)d_A(a_s)(\partial_t)U(\tau;0)$$

We then have that

$$D(t,s) = -a_s(0,s)U_s(1;0) + U_s(1;0)a_s(0,s).$$

However, $a_s(0,s) = a(0,s)(\partial_s) = 0$ since the homotopy is of closed loops with a fixed base point, so $\partial_s(0,s) = 0$. Hence we have that $D = 0$. $\square$

Corollary 3.48. If $E \to X$ is a bundle with a flat connection and X is simply connected, then the parallel transport operator $E_x \to E_y$ along a path from x to y is independent of the choice of path. In particular, E is trivial.

Proof. If X is simply connected, then by the lemma, the holonomy around any loop is trivial. Hence if γ_1, γ_2 are two paths from x to y, concatenating one with the inverse of the other gives a closed loop, over which the holonomy is trivial. Hence the parallel transport along γ_1 is equal to the parallel transport along γ_2. We can trivialise the bundle by picking a point $p \in X$ and then defining a map

$$f : X \times E_p \to E$$

where f maps each point (x, v) to the parallel transport of v along a curve from p to x. $\square$

Definition 3.49. From now on, when considering parallel transport, we will use the notation

$$U(b;a) : E_{\gamma(b)} \to E_{\gamma(a)}$$

to denote the parallel transport along the subset of the curve $\gamma : I \to Y$ starting at $t = a$ and ending at $t = b$.

Remark 3.50. Note that $U(b;a) = U(a,b)^{-1}$ and for $a, b, c \in I$ $U(c;a) = U(c;b)U(b;a)$.

Lemma 3.51. Let $\mathfrak{A} = A + i\phi$ be an $G_{\mathbb{C}}$ connection on a complexified bundle $P_{\mathbb{C}} \to X$. Let $\gamma : S^1 \to X$ be a curve and let $U : E_{\gamma(1)} \to E_{\gamma(1)}$ be the holonomy of the connection around γ. Then the holonomy of the connection

$$\bar{\mathfrak{A}} = A - i\phi.$$

is equal to $(U^{-1})^{\dagger}$.

Proof. The parallel transport of the connection $\mathfrak{A}$ solves the equation

$$\frac{d}{dt}U(t;0) = -\mathfrak{A}_{\gamma'}U(t;0)$$

Differentiating the equation

$$I = U(t;0)U(0;t)$$

gives

$$\frac{d}{dt}\left(U(t;0)^{-1}\right) = U^{-1}\mathfrak{A}.$$

Taking the conjugate transpose of both sides gives

$$\frac{d}{dt}U(0;t)^{\dagger} = -\bar{\mathfrak{A}}_{\gamma'}U(0;t)^{\dagger},$$

but this is the equation for the parallel transport of $\bar{\mathfrak{A}}$. □

We now give a useful calculation on an integral operator that appears in the linearisation of the perturbation operator. This calculation is used to prove that the perturbed Hessian is self adjoint.

Lemma 3.52. Let $P_{\mathbb{C}} \to S^1$ be a complexified bundle over S^1, and pick a connection $\mathfrak{A}$ on $P_{\mathbb{C}}$. Let $\hat{U}$ denote the holonomy with respect to the connection $\bar{\mathfrak{A}}$. Define the maps

$$I_{\pm} : \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}}) \to (\mathfrak{g}_{P_{\mathbb{C}}})_1$$

as

$$I_+(a) = \int_0^{2\pi} U(2\pi;\theta)a_{\theta}(\theta)U(\theta;0)d\theta$$

and

$$I_-(a) = \int_0^{2\pi} \hat{U}(-2\pi;-\theta)a_{-\theta}(-\theta)\hat{U}(-\theta;0)d\theta.$$

Then

$$I_+(a)^{\dagger} = I_-(\bar{a}).$$

Proof. This is a direct calculation. We have

$$I_+(a)^{\dagger} = \int_0^{2\pi} U(\theta;0)^{\dagger}a_{\theta}^{\dagger}(\theta)U(2\pi;\theta)^{\dagger}d\theta$$

Using Lemma 3.51 we can write this as

$$\begin{aligned}
I_+(a)^\dagger &= \int_0^{2\pi} \hat{U}(0;\theta)a_\theta^\dagger(\theta)\hat{U}(\theta;2\pi)d\theta \\
&= -\int_0^{2\pi} \hat{U}(0;2\pi-\theta)a_\theta^\dagger(2\pi-\theta)\hat{U}(2\pi-\theta;2\pi)d\theta \\
&= \int_0^{2\pi} \hat{U}(-2\pi;-\theta)a_{-\theta}^\dagger(-\theta)\hat{U}(-\theta;0)d\theta
\end{aligned}$$

□

Corollary 3.53. Let $P_{\mathbb{C}} \to X$ be a complexified bundle over a manifold X and let $\mathfrak{A}$ be a connection on $P_{\mathbb{C}}$. Define I_γ, $I_{\bar{\gamma}}$. Then

$$I_\gamma(a)^\dagger = I_{\bar{\gamma}}(a^\dagger).$$

We now consider how the norm of the holonomy depends on the imaginary part of the connection. If we restricted to the real connections, that is, $\phi = 0$, then we would have that the holonomy is in G and thus has bounded norm. For $\|\phi\|_{L^\infty}$ bounded we can generalise this result.

Lemma 3.54. Let $\mathfrak{A}$ be a complex connection with $\|\phi\|_{C^0}$ bounded. Then for any path $\gamma : [0,1] \to Y$, we have

$$|U(1;0)| \leq \exp \int_\gamma |\phi|.$$

Proof. Let $v \in E_{\gamma(0)}$ and consider the parallel transport $v(t)$ defined by $\mathfrak{A}$ along γ. Then, on a trivialisation of $P_{\mathbb{C}}$ over γ, $n(t) = |v(t)|^2$ satisfies the differential equation

$$\frac{dn}{dt} = 2\langle v, \nabla_A v\rangle = -2\langle v, i\phi.v\rangle.$$

Hence we have

$$\left|\frac{dn}{dt}\right| \leq |\phi||n|,$$

and so

$$n \leq |n(0)| \exp \int_0^1 |\phi|.$$

The bound on the operator norm of the parallel transport operator follows. □

Corollary 3.55. Suppose $\dim Y = 3$ and $\phi \in L^2_{1+\delta}$ for some $\delta > 0$. Then we have

$$|U(1;0)| \leq C \exp \|\phi\|_{L^2_{1+\delta}}.$$

Proof. This is just the previous result and an application of Corollary 2.6. $\square$

3.5.2 The Representation Variety

We now describe the space of flat connections on a manifold. These results are well known but will be used extensively so we recall them here.

Definition 3.56. The representation variety of a manifold Y with respect to the group G is the space

$$\mathrm{Hom}(\pi_1(Y), G)/G,$$

where G acts by conjugation.

Proposition 3.57. The space of flat connections on principal G bundles modulo gauge transformation on a manifold Y is in one to one correspondence with elements of the representation variety with respect to the group G.

Proof. Consider the universal cover $\tilde{Y}$ of Y, equipped with a trivial G bundle $\tilde{P}$. The group $\pi_1(Y)$ acts on the total space of $\tilde{P}$ space by diffeomorphisms via deck transformations and the given representation of $\pi_1(Y)$. We then take the quotient by $\pi_1(Y)$, giving a principal bundle over Y. The trivial connection on the covering induces a flat connection on the base, and by construction, the holonomy representation of the flat connection is given by the action of π_1 on G.

If we start with a flat connection on a bundle over Y, we can construct a group homomorphism $\pi_1(Y) \to G$ given by taking the holonomy of the connection around representatives of the homotopy classes of loops. Acting by gauge transformations will conjugate the parallel transport maps. $\square$

Remark 3.58. If we have a flat connection on a bundle $P \to Y$ then we can consider the pullback to the universal cover $\tilde{Y}$. By Corollary 3.48, we can trivialise this bundle. This gives a bundle isomorphism $\pi^*P \to \tilde{Y} \times G$. Moreover, if we have a section of an associated bundle $E \to Y$, then we can consider this as a $\pi_1(Y)$ invariant section on π^*E. When we trivialise the bundle, this section will be mapped to a $\pi_1(Y)$ equivariant section, where $\pi_1(Y)$ acts by the element of the representation variety that defines the connection.

3.5.3 Twisted De Rham Cohomology

We now consider the associated de Rham complex associated to a complex connection. This is a key idea used in the deformation theory of flat connections. In this section we restrict to three-manifolds.

Definition 3.59 (Twisted De Rham Complex)**.** Let $\mathfrak{A}$ be a flat $G_{\mathbb{C}}$ connection. Then we can consider the complex

$$0 \to \Omega^0(\mathfrak{g}_{P_{\mathbb{C}}}) \xrightarrow{d_{\mathfrak{A}}} \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}}) \xrightarrow{d_{\mathfrak{A}}} \Omega^2(\mathfrak{g}_{P_{\mathbb{C}}}) \xrightarrow{d_{\mathfrak{A}}} \Omega^3(\mathfrak{g}_{P_{\mathbb{C}}}) \to 0.$$

Since $\mathfrak{A}$ is flat, $d_{\mathfrak{A}}^2 = 0$ and so this is a genuine complex. Its cohomology is called the twisted de Rham cohomology with respect to the connection $\mathfrak{A}$. We denote the cohomology groups by

$$H^*_{\mathfrak{A}}(Y; \mathfrak{g}_{P_{\mathbb{C}}})$$

We can also make the following definitions describing flat $G_{\mathbb{C}}$ connections. These definitions are standard and can be found in, for example, [Don02] and [He19].

Definition 3.60. Let $\mathfrak{A}$ be a flat $G_{\mathbb{C}}$ connection. It is

- degenerate if $H^1_{\mathfrak{A}}(Y; \mathfrak{g}_{P_{\mathbb{C}}}) = 0$, otherwise non-degenerate,
- obstructed if $H^2_{\mathfrak{A}}(Y; \mathfrak{g}_{P_{\mathbb{C}}}) \neq 0$, otherwise unobstructed,
- acyclic if $H^*_{\mathfrak{A}}(Y; \mathfrak{g}_{P_{\mathbb{C}}}) = 0$.

Lemma 3.61. A connection $\mathfrak{A}$ is obstructed if and only if it is non-degenerate.

Proof. We will prove that $(H^2_{\mathfrak{A}})^* \cong H^1_{\mathfrak{A}}$. To see this, consider map

$$H^k_{\mathfrak{A}} \times H^{3-k}_{\mathfrak{A}} \to \mathbb{C}, \quad ([\alpha], [\beta]) \mapsto \int_Y \operatorname{tr}(\alpha \wedge \beta).$$

This defines an isomorphism

$$H^k_{\mathfrak{A}} \cong \left(H^{3-k}_{\mathfrak{A}}\right)^*.$$

and so we have in particular that $\dim H^1_{\mathfrak{A}} = \dim H^2_{\mathfrak{A}}$. □

The acyclic condition is related to reducibility in the following way. The kernel of $d_{\mathfrak{A}} : \Omega^0(\mathfrak{g}_{P_{\mathbb{C}}}) \to \Omega^1(\mathfrak{g}_{P\mathbb{C}})$. If the connection $\mathfrak{A}$ is reducible, then the holonomy is contained in a strict subgroup.

Corollary 3.62. A connection is acyclic if and only if it is non-degenerate and $H^1_{\mathfrak{A}} = 0$.

Lemma 3.63. A flat connection, $\mathfrak{A} = A + i\phi$, satisfying $d_A^*\phi = 0$, is isolated if and only if $H^1_{\mathfrak{A}} = 0$.

Proof. We use a trick to simplify the analysis. Note that a connection $F_{\mathfrak{A}}$ is flat if and only if there exists an element $\xi \in \Omega^0(\mathfrak{g}_{P_{\mathbb{C}}})$ such that

$$*F_{\mathfrak{A}} = d_{\mathfrak{A}}\xi.$$

Note that if $F_{\mathfrak{A}}$ is indeed flat then we can take ξ to be zero. For the other direction, note that see this note that the Bianchi identity implies that $d_A^* d_A \xi = 0$. Hence we have that

$$0 = \int_Y \langle d_{\mathfrak{A}}^* d_{\mathfrak{A}}\xi, \xi\rangle = \int_Y |d_{\mathfrak{A}}\xi|^2$$

and so the connection is flat. Hence we can describe flat connections near $\mathfrak{A}$ as solutions to the equation

$$f : \ker d_{\mathfrak{A}}^* \oplus \Omega^0(\mathfrak{g}_{P_{\mathbb{C}}}) \to \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}}), \quad (a, \xi) \mapsto *(d_{\mathfrak{A}}a + a \wedge a) + d_{\mathfrak{A}}\xi.$$

The derivative of this map is

$$df(a, \xi) = *d_{\mathfrak{A}}a + d_{\mathfrak{A}}\xi$$

which is an isomorphism if and only if $H^1_{\mathfrak{A}} = 0$ (when restricted to appropriate Sobolev completions), since we have

$$\Omega^1(\mathfrak{g}_{P_{\mathbb{C}}}) \cong H^1_{\mathfrak{A}} \oplus \ker d_{\mathfrak{A}}^* \oplus \operatorname{Im} d_{\mathfrak{A}}.$$

Hence flat connections are isolated with respect to $\mathcal{G}_{\mathbb{C}}$ gauge transformations. □

Remark 3.64. The above result also proves that flat connections with $d_A^*\phi = 0$ are isolated in $\mathcal{A}_{\mathbb{C}}/\mathcal{G}$.

3.6 Moduli Spaces of Connections

A key object of study in Gauge theory is the moduli space of connections. That is the space of connections modulo gauge equivalence. In the complex setting, as we have described, we are interested in the space of connections such that $d_A^*\phi = 0$ at some point in the $\mathcal{G}_{\mathbb{C}}$. This is the complexified analogue of considering $\mathcal{B}^*$ – reducible real connections modulo gauge. In the following we denote by μ the map

$$\mu(\mathfrak{A}) = d_A^*\phi,$$

where $\mathfrak{A} = A + i\phi$ with respect to a given G–reduction of $P_{\mathbb{C}}$.

Definition 3.65. We define the space $\mathcal{B}_{\mathbb{C}} = \mu^{-1}(0)/\mathcal{G}$, and we define $\mathcal{B}^*_{\mathbb{C}}$ to be the subset of those connections that are irreducible, that is, restricting to gauge equivalence classes $[\mathfrak{A}]$ with $H^0_{\mathfrak{A}}(\mathfrak{g}_{P_{\mathbb{C}}}) = 0$.

We can use a G reduction $P \subseteq P_{\mathbb{C}}$ to define a metric on the space of complex connections modulo gauge.

Definition 3.66. We define the metric on the space of connections to be the L^2 norm on the tangent bundle, restricted to kernel of the complex Coulomb gauge. That is, for a tangent vector $\xi \in T_{\mathfrak{A}}\mathcal{A}_{\mathbb{C}} = \Omega^1(\mathfrak{g}_P)$, we define the norm of it's projection under $\pi : \mathfrak{A}_{\mathbb{C}} \to \mathfrak{B}_{\mathbb{C}}$ to be

$$\|\pi_*\xi\|_{L^2} = \|\xi_0\|,$$

where $\xi_0 + \xi_1 = \xi$ is the decomposition given by $\Omega^1(\mathfrak{g}_{P_{\mathbb{C}}}) \cong \ker \operatorname{Re}(d^*_{\mathfrak{A}}) \oplus \operatorname{Im} d_{\mathfrak{A}}$.

Using the metric, we can define the geodesic distance between connections. This will be important in Chapter 5 where we deduce a finite energy result for connections with bounded perturbed Chern-Simons functional.

Definition 3.67. Let $\mathfrak{A}, \mathfrak{B}$ be two complex connections. Let S be the set of paths connecting A and B. Then the geodesic distance between them is

$$d(\mathfrak{A}, \mathfrak{B}) = \inf_{\gamma \in S} l(\gamma),$$

where $l(\gamma)$ is the length of the curve.

Lemma 3.68. Let $\mathfrak{A}$ be a background connection and suppose that $a + i\varphi$ is in the complex Coulomb gauge relative to $\mathfrak{A}$. Then if $\mathfrak{B} = A + a + i\varphi$,

$$d(\mathfrak{A}, \mathfrak{B}) = \|a + i\varphi\|_{L^2}$$

Proposition 3.69. This is just the reverse triangle inequality on the space of connections. Let $\mathfrak{A} = A_1 + i\phi_1$, $\mathfrak{B} = A_2 + i\phi_2$ be complex L^2_1 connections. Then with respect to the metric induced by a G–reduction,

$$d(\mathfrak{A}, \mathfrak{B}) > \left| \|\phi_1\|_{L^2} - \|\phi_2\|_{L^2} \right|.$$

Proof. For a connection $\mathfrak{D} = D + i\varphi$, $\|\varphi\|_{L^2}$ is gauge invariant. Hence we have that taking the L^2 norm of the imaginary part of the connection gives a $\mathcal{G}$–invariant map $\mathcal{A} \to [0, \infty)$. From Lemma 3.68 we have that the result is true if $\mathfrak{B}$ is in

Coulomb gauge relative to $\mathfrak{A}$. If not then we can split up the domain of a path of connections into small intervals and bound the geodesic distance of each part of the path from below by the difference in the norms, and then apply the reverse triangle inequality. $\square$

Corollary 3.70. Let $\gamma : [0,1] \to \mathcal{A}$ be a path of connections and let

$$E = \int_a^b \|\gamma'\|_{L^2}^2 .$$

for $a, b \in \mathbb{R}$. Then

$$\frac{1}{2} d(\gamma(0), \gamma(1))^2 \leq (b-a)^2 E$$

Proof. Apply the Cauchy-Schwartz inequality to the integral defining the length of γ. $\square$

Remark 3.71. There is not a version of the above inequality for paths $\gamma : \mathbb{R} \to \mathcal{A}$. If this were the case then one could bound $\|\phi\|_{L^2}$ by the change in the Chern-Simons functional for a solution of the Kapustin-Witten equations.

Chapter 4

Harmonic Maps and Flat Connections

The work of Corlette [Cor88], shows that flat connections on principal $G_{\mathbb{C}}$ bundles $P_{\mathbb{C}} \to Y$ that satisfy

$$d_A^* \phi = 0,$$

with respect to some choice of G reduction, are in one to one correspondence with $\pi_1(Y)$ equivariant harmonic maps

$$u : \hat{Y} \to G_{\mathbb{C}}/G.$$

Here $\hat{Y}$ is the universal cover of Y and $\pi_1(Y)$ acts on the codomain by the holonomy representation

$$\rho : \pi_1(Y) \to G_{\mathbb{C}},$$

defined by the connection. Hence, the existence of a $G_{\mathbb{C}}$ gauge transformation mapping a given flat connection to one with $d_A^* \phi = 0$ is equivalent to showing the existence of a suitable harmonic map.

For the case when Y is closed, that is compact with boundary, [Cor88, Theorem 3.3] gives that such a harmonic map exists if and only if the representation ρ is reductive. In this section we prove the following theorem.

Theorem 4.1. Let $\mathfrak{A}$ be a flat connection on a principal $P_{\mathbb{C}}$ bundle on a manifold with boundary with a reductive associated representation. Then, for any choice of G reduction on the boundary, there exists a G reduction of $P_{\mathbb{C}}$ that coincides with the given G reduction on the boundary, with respect to which

$$d_A^* \phi = 0.$$

Here $\mathfrak{A} = A + i\phi$ with respect to the G reduction, in the sense of Definition 3.7.

Remark 4.2. The implication of theorem could be restated by saying that for a given choice of G reduction there exists a $G_{\mathbb{C}}$ gauge transformation such that the image of the connection under the gauge transformation satisfies $d_A^*\phi = 0$.

An overview of a proof of this theorem is as follows. First we note that any choice of G reduction corresponds to picking a metric on an appropriate bundle. By Proposition 3.23, G reductions on Y are in one to one correspondence with sections of an associated bundle with fibre $G_{\mathbb{C}}/G$. When we pull this back to the universal cover and gauge transform to trivialise the bundle, we get a corresponding equivariant section

$$h : \hat{Y} \to G_{\mathbb{C}}/G.$$

By Proposition 3.29, we have that $dh = \phi$. It follows that h is harmonic if and only if $d_A^*\phi = 0$. Hence the proof of the theorem comes down to proving the existence of a harmonic map from $\hat{Y}$ into $G_{\mathbb{C}}/G$. The reductivity assumption is really the statement that the monodromy representation splits as a direct sum of irreducible representations with no invariant subspaces. This is the key condition that allows us to rule out bubbling in the harmonic maps we consider. In this chapter we first review the relevant harmonic map theory and then give the proof of Theorem 4.1.

4.1 Basic Definitions

Definition 4.3 (Harmonic map from a compact manifold)**.** Let Y and M be two Riemannian manifolds, with M complete and Y compact. A map $f : Y \to M$ is said to be harmonic if it is the critical point of the energy functional

$$E(f) = \int_Y |df|^2$$

where $df \in T^*Y \otimes f^*TM$ is the derivative of f. If Y has non-empty boundary we will assume that f satisfies Dirichlet boundary conditions, that is, has prescribed boundary values.

The second type of harmonic maps that we consider are $\pi_1(Y)$ equivariant harmonic maps from the universal cover of a compact manifold, Y. Note that the universal cover of a compact manifold will be non-compact if $|\pi_1(Y)|$ is infinite.

Definition 4.4. Let Y be a compact manifold, possibly with boundary, and let $\pi_1(Y)$ act by isometries on a complete Riemannian manifold M. A $\pi_1(Y)$ equivariant map

$$f : \hat{Y} \to M$$

satisfying Dirichlet boundary conditions, is said to be harmonic if it is a critical point of the energy functional

$$E(f) = \int_Y |df|^2,$$

where f varies over L^2_1 maps satisfying the boundary conditions. Note that, since f is equivariant and $\pi_1(Y)$ acts by isometries, the real valued function

$$|df|^2 : Y \to \mathbb{R}, \quad x \mapsto |df([\gamma]x)|^2$$

is well defined, where $[\gamma] : \hat{Y} \to \hat{Y}$ is the deck transformation corresponding $[\gamma] \in \pi_1(Y)$.

The application that we have in mind for these maps is the case where Y is a compact manifold and we have a flat connection A on a bundle $P \to Y$. In this case, as explained in Section 3.5.2, the bundle P is the quotient of $\hat{Y} \times G$ by $\pi_1(Y)$, where $\pi_1(Y)$ acts via the holonomy representation. If we have a section u of an associated bundle E with fibre V, then we can pull this section back to a section of the pullback bundle over $\hat{Y}$, $\hat{E} = \pi^* E$. This section will be $\pi_1(Y)$ invariant. If we then trivialise the bundle via parallel transport then we get a map bundle map $\Phi : \hat{E} \to \hat{Y} \times E$. The image $\hat{u} = \Phi(\pi^* u)$ is a $\pi_1(Y)$ equivariant section of the trivial bundle. Moreover, given an equivariant section of a trivial bundle, $\hat{u}$, we can consider the section of the associated bundle by applying the inverse transform to recover the pullback bundle.

There has been extensive work on harmonic maps. For details, a good references is the paper [GS92] on Harmonic maps into singular spaces, which contains details on earlier work on harmonic maps into smooth spaces. A key theorem we will need is a weak existence theorem for the Dirichlet problem of harmonic maps on manifolds with boundary.

Lemma 4.5 (Existence of Weak Energy Minimizers)**.** Let Y be a compact manifold with smooth boundary and M a Riemannian manifold. Then, for any choice of $f \in C^1(\partial Y, M)$, there exists a map

$$u : Y \to M \in L^2_1(Y; M)$$

such that $u|_{\partial Y} = f$ with the property that for any $v \in L^2_1$ equal to f on the boundary we have

$$E(u) \leq E(v).$$

Proof. This is a standard result but we recall the proof for completeness. Consider the set

$$\mathcal{S} = \left\{ u \in L^2_1(Y; M) \mid u_{\partial Y} = f \right\}.$$

Define the energy

$$E : S \to \mathbb{R}, u \mapsto \int_Y |\nabla u|^2$$

and let u_i be an energy minimizing sequence. From the Poincaré inequality, we have that u is bounded in L^2, and hence the sequence u_i is bounded and so contains a weak L^2_1 convergent subsequence $u_i \rightharpoonup u$. The set $\mathcal{S}$ is closed in L^2_1 and so is weakly closed, and hence $u \in S$ since the L^2 norm is weakly lower semicontinuous. □

A key result is the regularity result on compact subsets for harmonic maps. This was originally shown in [ES64], and is generalised with fewer regularity assumptions in [GS92].

Theorem 4.6. Let M be a Riemannian manifold with non-positive sectional curvature. Suppose $f : B(0; 1) \to M \in L^2_1$ is an energy minimising map for some metric on the ball. Then there exists a constant C such that

$$\sup_{B\left(0;\frac{1}{2}\right)} |\nabla f|^2 \leq C \int_{B(0;1)} |\nabla f|^2$$

Proof. See [GS92, Theorem 2.4]. □

Remark 4.7. The choice of $B\left(0; \frac{1}{2}\right)$ in the theorem is fairly arbitrary, and could instead be stated for any choice of compact subset in the interior of $B(0, 1)$. The theorem will be applied in this form.

4.2 Existence Theory for Equivariant Harmonic Maps

We now prove the existence of $\pi_1(Y)$ equivariant maps under the assumption that the action on the target is Abelian or contains no fixed points on the boundary. For the application to flat $G_{\mathbb{C}}$ connections, this corresponds to the holonomy representation being either Abelian or irreducible (and non-Abelian). We treat the Abelian case first. This case is simpler and is a result of the existence of harmonic functions on manifolds with boundary, which we now state.

Proposition 4.8. Let Y be a compact Riemannian manifold with non-trivial boundary, and let $f_0 : \partial Y \to \mathbb{R}$ be continuous. Then there exists a function

$$f \in C^2(\mathrm{int}(Y)) \cap C^0(\bar{Y})$$

such that $f_{\partial Y} = f_0$ and $d^*df = \eta$, for any C^0 function η.

Proof. The bilinear form on $L^2_1(Y)_0$ given by

$$B(u,v) = \int_Y \langle du, dv\rangle$$

is coercive. This follows from the Poincaré inequality. Hence it follows from the Lax-Milgram theorem that there exists an element $v \in L^2_1(Y)_0$ such that

$$B(u,v) = \int_Y \langle u, \eta\rangle = \langle u, d^*dv\rangle.$$

for all $u \in L^2_1(Y)_0$. It follows that v is our desired solution. □

We now consider the existence of $\pi_1(Y)$ equivariant maps from the universal cover $\hat{Y}$.

Theorem 4.9 (Existence of equivariant harmonic maps)**.** Suppose there exists an equivariant map $f_0 : \hat{Y} \to M$. Then either there exists a harmonic map $f : Y \to M$ such that $f|_{\partial\hat{Y}} = f_0|_{\partial\hat{Y}}$ or there exists a point at infinity in M that is fixed by $\pi_1(Y)$.

Proof. We follow the strategy used in [GS92, Theorem 7.1]. Pick a finite cover of Y, $\mathscr{U}$, such that each element $\hat{U}$ in the connected components of preimages under π of the cover is, such that

$$g(\hat{U}) \cap \hat{U} = \emptyset, \forall g \in \pi_1(Y) \setminus \{\mathrm{Id}\}\,.$$

We may assume that each U has smooth boundary. Let f_i be a minimising sequence for the equivariant energy and satisfying the boundary condition. Pick an element, U, of the cover and pick a connected component $\hat{U} \subseteq \pi^{-1}(U)$. We create a new minimising sequence as follows. Let $\check{f}_i$ be defined by

$$\check{f}_i(x) = \begin{cases} g(v_i(g^{-1}x)) & x \in g(U), \\ f_i(x) & x \notin \pi_1(Y)(U) \end{cases}$$

where we pick $v_i : \hat{U} \to M$ to be the minimizer to the Dirichlet harmonic map problem on $\hat{U}$, satisfying $v_i|_{\partial\hat{U}} = f_i|_{\partial\hat{U}}$, coming from Lemma 4.5. Now from Theorem 4.6 we have that $\check{f}_i$ has bounded first derivative on any compact subset of $\hat{U}$, and hence it is Lipschitz there. Moreover the Lipschitz constant is bounded by the energy of f_i, and the energy is uniformly bounded on the sequence. Hence $\check{f}_i$ has a subsequence that is uniformly convergent on compact subsets of $\hat{U}$ (and on any given image of $\hat{U}$ under deck transformations). It follows that the limit of the sequence either maps compact subsets of $\hat{U}$ to a bounded subset of M, or it maps the whole compact set to a single point at infinity in M. We will conclude the proof of this theorem by showing that in this latter case, the translates of the compact set by deck transformations must also map to the same point. It then follows that the action of $\pi_1(Y)$ on M fixes a point at infinity.

To do this, suppose that γ is a smooth embedded curve from $x \in \hat{U}$ to $\tau(x)$ for some $\tau \in \pi_1(Y)$. We have that

$$\int_\gamma |\nabla_{\gamma'} \check{f}_i|^2 dt \leq \int_\gamma |\nabla \check{f}_i|^2 dt$$

and that

$$\left(\int_{N(\gamma)} |\nabla \check{f}_i|^2 dc \right) dz \leq E(f_i),$$

where $N(\gamma)$ is a smooth neighbourhood of the interior points of γ. In particular, $N(\gamma)$ can be parametrised as a smooth family of curves from x to $\tau(x)$. Suppose we parametrise $N(\gamma)$ by the disk D^{n-1}. Then, by Tonelli's theorem there exists a point $z \in D^{n-1}$ such that

$$\int_{\gamma_z} |\nabla \check{f}_i|^2 \leq E(f_i).$$

From this it follows that the length of $\check{f}_i(\gamma_z)$ is bounded, and hence if $\check{f}_i$ is converging to an element at infinity in M, $\tau(\check{f}_i(x)) = \check{f}_i(x)$.

Hence, in the case where there does not exist a point in the boundary fixed by $\pi_1(Y)$, we may assume that the minimizing sequence converges uniformly on compact subsets of $\hat{Y}$. □

4.3 Flat $G_{\mathbb{C}}$ Bundles on Manifolds with Boundary

Suppose that we have a given flat connection $\mathfrak{A}$ on Y which has an associated holonomy representation

$$\rho : \pi_1(Y) \to G_{\mathbb{C}} \subseteq GL(n; \mathbb{C}).$$

We assume that ρ is reductive, meaning that we can write

$$\mathbb{C}^n = E_1 \oplus \cdots \oplus E_k,$$

as a direct sum of $G_{\mathbb{C}}$ irreducible representations. By considering each E_k separately, we can reduce to the case where $k = 1$, and we consider separately the case where $\dim E_1 = 1$, and where $\dim E_1 > 1$.

We first consider the case where $E_1 = 1$. In this case we are essentially considering a $\mathbb{C}^*$ connection.

Proposition 4.10. Let $P_{\mathbb{C}} \to Y$ be a $\mathbb{C}^*$ bundle, and let $\mathfrak{A}$ be a flat connection on P. Let $E = P_{\mathbb{C}} \times_{\mathbb{C}^*} \mathbb{C}$ and pick a metric on $E_{\partial Y}$. Then there exists a metric on E, equal to the given metric on the boundary, such that with respect to this metric

$$\mathfrak{A} = A + i\phi$$

and $d^*\phi = 0$

Proof. In the Abelian case we can work directly on Y and find a gauge transformation to solve the boundary value problem. $\mathbb{C}^*$ gauge transformations are given by maps $f : Y \to \mathbb{C}^*$. If we write such a map as

$$f(x) = e^z$$

for $z = w + iy \in \mathbb{C}$, we then have that

$$f(A + i\phi) = A + i\phi + ie^{-z}de^z = A + i\phi + idz$$

Hence $\phi \mapsto \phi + dw$. The condition $d_A^*\phi = 0$ under the gauge transformation can be written as

$$d^*\phi + d^*dw = 0,$$

which we can solve using Proposition 4.8 (and imposing that $w = 0$ on the boundary). ☐

We now consider the case of non-abelian representations.

Lemma 4.11. Let $\rho : \pi_1(Y) \to GL(n;\mathbb{C})$ be an irreducible representation. Then ρ induces a action on $GL(n;\mathbb{C})/U(k)$. Moreover, this representation has no fixed points at infinity.

Proof. We can define the action of $\pi_1(Y)$ on $GL(n;\mathbb{C})/U(k)$ simply by defining

$$g[M] = [\rho(g)M].$$

Suppose this fixes a point at infinity in $GL(n;\mathbb{C})/U(k)$. Then by Proposition 3.28, each isometry $\rho(g)$ for $g \in \pi_1(Y)$ is contained in the stabilizer of this point at infinity, and so the image of ρ is contained in a parabolic subgroup of $GL(n;\mathbb{C})$. This contradicts the irreducibility assumption. ☐

We can now state out existence result for canonical metrics for flat connections on manifolds with boundary.

Proposition 4.12. Let $\rho : \pi_1(Y) \to G_{\mathbb{C}} \subseteq GL(n;\mathbb{C})$ be an irreducible representation defining a flat $G_{\mathbb{C}}$ connection, $\mathfrak{A}$, on a bundle $P_{\mathbb{C}} \to Y$, where Y is a compact manifold with boundary. Pick a metric on the associated bundle with fibre $\mathbb{C}^n$. Then there exists a metric on Y, agreeing with the metric on the boundary, such that with respect to the induced G reduction of P, the associated connection has the form

$$\mathcal{A} = A + i\phi,$$

and $d_A^*\phi = 0$.

Proof. Let E denote the associated bundle

$$E = P_{\mathbb{C}} \times_{GL(n;\mathbb{C})} \mathbb{C}^n$$

and pull this back to the bundle $\pi^*E \to \hat{Y}$ over the universal cover. Pick a background metric on E. This lifts to a π_1 invariant metric on $\pi^*E \to \hat{Y}$. If we pick a point $p \in \hat{Y}$ then we can define a trivialisation of π^*E via parallel transport (see Corollary 3.48). As explained after Definition 4.4, under the bundle isomorphism trivialising π^*E, the pullback metric becomes $\pi_1(Y)$ equivariant and the flat connection is the product connection. Let h_0 denote this equivariant metric. Then, from Proposition 3.23 the choice of metric gives an equivariant section, which we also call h_0, of the trivial $G_{\mathbb{C}}/G$ bundle

$$h_0 : \hat{Y} \to G_{\mathbb{C}}/G.$$

From Proposition 3.29 we have that the pullback of ϕ in the trivialisation is given by dh_0.

Now, using Theorem 4.9, we can solve the harmonic map boundary value problem to get an equivariant harmonic map

$$h : \hat{Y} \to G_{\mathbb{C}}/G$$

that is equal to h_0 on the boundary. It now follows from Lemma 3.19 and Proposition 3.29 that the metric on the pullback bundle π^*E defines a G reduction with respect to which

$$\mathcal{A} = A + i\phi$$

and $d_A^*\phi = 0$. □

Chapter 5

The Complexified Chern-Simons Functional

The Chern-Simons functional plays an important role in quantum field theory in three dimensions, and has strong links to four dimensional gauge theory and knot theory.

In four dimensional Chern-Weil theory an important fact is that, for any connection A on a principal bundle P, the form

$$\operatorname{tr}(F_A \wedge F_A)$$

is closed and its de-Rham homology depends only on the bundle. If P is a $SU(n)$ bundle on a closed, oriented, four manifold then this integral is a multiple of the second Chern class of the bundle.

One can verify the following by a direct calculation.

Lemma 5.1. Let X be four dimensional an let A be a connection on a principal bundle $P \to X$ and $a \in \Omega^1(\mathfrak{g}_P)$, then

$$\operatorname{tr}(F_{A+a} \wedge F_{A+a}) - \operatorname{tr}(F_A \wedge F_A) = d \operatorname{tr}\left(a \wedge d_A a + \frac{2}{3} a \wedge a \wedge a\right).$$

Now, if we now consider a four manifold with boundary, it will no longer be true that the integral of $\operatorname{tr}(F_A \wedge F_A)$ is independent of the choice of connection. In general we will have

$$\int_X \operatorname{tr}(F_{A+a} \wedge F_{A+a}) - \operatorname{tr}(F_A \wedge F_A) = \int_{\partial X} \operatorname{tr}\left(a \wedge d_A a + \frac{2}{3} a \wedge a \wedge a\right)$$

We make the following definition.

Definition 5.2. With respect to a background connection B on a bundle $P \to Y$, where Y is a three-manifold, the Chern-Simons functional on the space of connections $\mathcal{A}$ is the map

$$\mathfrak{C} : \mathcal{A} \to \mathbb{R}, \quad \mathfrak{C}(A) = \int_Y \operatorname{tr}\left(a \wedge d_B a + \frac{2}{3} a \wedge a \wedge a\right),$$

where $a = A - B \in \Omega^1(Y)$.

Remark 5.3. When one considers connections on a trivial bundle, it is common to take B to be the trivial (product) connection.

Lemma 5.4. *The Chern-Simons functional is invariant under the action of elements of $\mathcal{G}_0$, the gauge transformations in the connected component of the identity. Moreover, the Chern-Simons functional defines a map*

$$\mathcal{G} \to \mathbb{Z}$$

that detects the connected components of the gauge group.

Proof. We consider the case of $SU(2)$ bundles. Φ be a gauge transformation. Pick an oriented four manifold with boundary Y. All $SU(2)$ bundles on X are trivial, so we can pick a trivialisation. Gluing the bundle along the gauge transformation gives a bundle $P_\Phi \to X$. Let $8\pi^2 c_2(P)$ denote the integral of $\operatorname{tr}(F_A \wedge F_A)$ for a connection A on Y. When restricted to $Y \hookrightarrow X$. We have

$$8\pi C_2(P_\Phi) = \int_{X \cup X'} \operatorname{tr}(F_A \wedge F_A) \in \mathbb{Z} = \mathfrak{C}(A) - \mathfrak{C}(\Phi^*(A)),$$

where to get the right hand side we have split the integral along the boundary Y. Note that this value is locally constant and hence constant on the connected components of $\mathcal{G}$. □

In this chapter we focus on the Chern-Simons functional on complexified bundles.

Definition 5.5. Let $P \to Y$ be a principal bundle over a three-manifold Y. The complexified Chern-Simons functional is the Chern-Simons functional for a connection on the principal bundle $P_{\mathbb{C}}$. It is a map

$$\mathcal{G}_{\mathbb{C}} \to \mathbb{C}.$$

Since the complex gauge group has the same number of connected components as the real gauge group, the complex Chern-Simons functional is still gauge invariant modulo $\mathbb{Z}$.

Proposition 5.6. The critical points of the Chern-Simons functional are flat connections.

Proof. We can calculate the derivative of the Chern-Simons functional directly. We have, writing out the terms in (1.1),

$$\mathfrak{C}(\mathfrak{A} + \alpha) = \mathfrak{C}(A) + \mathfrak{C}(\alpha + \mathfrak{B}) + \int_Y \mathrm{tr}(a \wedge d_{\mathfrak{A}}\alpha + \alpha \wedge d_{\mathfrak{B}+\alpha} a)$$

where $\mathfrak{B}$ is our choice of flat background connection and $a = \mathfrak{A} - \mathfrak{B}$. If we consider the above formula with α replaced by $t\alpha$ then we have

$$\begin{aligned} \mathfrak{C}(\mathfrak{A} + t\alpha) &= \mathfrak{C}(A) + t \int_Y \mathrm{tr}\left(a \wedge d_{\mathfrak{A}}\alpha + \alpha \wedge d_{\mathfrak{B}} a\right) + O(t^2) \\ &= \mathfrak{C}(A) + 2t \int_Y \mathrm{tr}\left(\alpha \wedge (d_{\mathfrak{B}} a + a \wedge a)\right) + o(t^2). \end{aligned} \tag{5.1}$$

Here we've used an integration by parts formula for $\mathrm{tr}(\alpha \wedge d_{\mathfrak{B}} a)$. Hence $\mathfrak{A}$ is a critical point if and only if

$$F_{\mathfrak{A}} = d_{\mathfrak{B}} a + a \wedge a = 0.$$

Note that this formula holds because $\mathfrak{B}$ is flat. If we were to chose another connection then the critical points would be connections with curvature equal to that of the background connection.

□

Proposition 5.7. The complexified Chern-Simons functional is a holomorphic map

$$\mathfrak{A}_{\mathbb{C}} \to \mathbb{C}.$$

Proof. This follows from the fact that all of the operations in the definition of the Chern-Simons functional (integration over Y, taking the trace, wedge products and the taking the derivative with respect to a complex connection are holomorphic).

□

5.1 The $SL(2;\mathbb{C})$ Representation Variety

In this section we consider the values of the Chern-Simons functional on the $SL(2;\mathbb{C})$ representation variety.

Lemma 5.8. The Chern-Simons functional takes finitely many values at critical points.

Proof. The critical points of the Chern-Simons functional are precisely those coming from $F_A = 0$. If we also require that $d_A^* \phi = 0$ then critical loci of the Chern-Simons functional correspond to the irreducible components of the $SL(2;\mathbb{C})$ character variety. One each irreducible component the Chern-Simons functional is constant. This follows from the fact that the derivative of the Chern-Simons functional at the connection $\mathfrak{A}$ is

$$d_a \mathfrak{C}(\mathfrak{A}) = \int_Y \operatorname{tr}(a \wedge F_{\mathfrak{A}}) = 0.$$

and so as we move along a path in an irreducible component then the Chern-Simons functional is constant. The result follows from the fact that there are finitely many irreducible components. □

In general, it may be the case that two different irreducible components of the representation variety share the same critical points, but we will always have that there are finitely many values. This gives a nice visualisation of the moduli space of flat connections in terms of Picard-Lefschetz theory, as described in Chapter 1 which we will come back to when considering the perturbed Chern-Simons functional.

5.2 The Gradient Flow of Invariant Functionals

We now start to develop the perturbation theory for the complexified Chern-Simons functional. We will study the gradient flow of holomorphic functionals

$$f : \mathcal{A}_{\mathbb{C}} \to \mathbb{C}$$

that are invariant under gauge transformations in the connected component $(\mathcal{G}_{\mathbb{C}})_0$.

Definition 5.9. We call such a functional f an invariant functional. The (negative) gradient flow of such a functional, for a parameter θ, is the ordinary differ-

ential equation
$$\frac{d(A+i\phi)}{dt} = -\nabla \operatorname{Re}\left(e^{i\theta} f(A+i\phi)\right)$$
where the gradient is with respect to the L^2 inner product defined by the choice of G reduction of the $G_{\mathbb{C}}$ bundle.

Remark 5.10. Note that the definition of the L^2 gradient depends on the choice of G reduction as one needs to pick a metric to define the L^2 norm.

An obvious example of an invariant functional is the Chern-Simons functional. The perturbations that we construct will also be invariant functionals. We now prove that studying the gradient flow of invariant functionals preserves stability of connections. This can be explained intuitively by the fact that the gradient flow is finding the path of steepest descent for the functional. If some component was in a direction of a $\mathcal{G}_{\mathbb{C}}$ orbit then the component of the tangent vector of the path in this direction would not be changing the value of the functional. Hence there is no component of the velocity of the path in these directions.

Theorem 5.11. Suppose that $A+i\phi : \mathbb{R} \to \mathcal{A}_{\mathbb{C}}$ is a solution to the gradient flow of an invariant functional, with respect to some choice of G reduction. Then the function
$$t \mapsto d^*_{A(t)}\phi(t)$$
is constant under the flow.

Remark 5.12. This is proven for the special case of the gradient flow of the Chern-Simons functional (when starting at $d^*_A\phi = 0$) in [GU12, Proposition 6.2] and stated in [Muñ25, Remark following Proposition 2.11].

Proof. If we have a path $A+i\phi : \mathbb{R} \to \mathcal{A}_{\mathbb{C}}$ then we have that
$$\frac{d}{dt} d^*_A\phi = d^*_A \frac{d\phi}{dt} + * \left[\frac{dA}{dt}, *\phi\right] = \operatorname{Im}\left(d^*_{A+i\phi}\frac{d(A+i\phi)}{dt}\right)$$
Now if $A+i\phi$ is the gradient flow of an invariant functional, then we have that the image of the infinitesimal action of the gauge group is contained in the kernel of the derivative of $e^{i\theta}f$ is in other words
$$\operatorname{Im} d_{A+i\phi} \subseteq \ker de^{i\theta} f.$$
Hence we have that, for all $\xi \in \Omega^0(\mathcal{G}_{\mathbb{C}})$,
$$\langle d^*_{A+i\phi}\nabla e^{i\theta} f, \xi\rangle = \langle \nabla e^{i\theta} f, d_{A+i\phi}\xi\rangle = d_{\left(d_{A+i\phi}\xi\right)} e^{i\theta} f = 0.$$

□

Corollary 5.13. The value $d_A^*\phi$ is conserved under the gradient flow of the complexified Chern-Simons functional.

5.3 Holonomy Perturbations

Our method for producing perturbations will be to consider the holonomy of a connection. In this section we adapt the ideas outlined in [BD95; Flo88; Tau90] to complexified bundles. The goal is to construct a gauge invariant function η with geometric significance so that we can define a perturbed functional.

Definition 5.14 (Perturbed Chern-Simons Functional)**.** Let $\eta : \mathcal{A}_{\mathbb{C}} \to \mathbb{C}$ be $\mathcal{G}_{\mathbb{C}}$ invariant function on the space of connections. Then the η–perturbed Chern-Simons functional is

$$\mathfrak{C}_\eta : \mathcal{A}_{\mathbb{C}} \to \mathbb{C}, \quad \mathfrak{C}_\eta : A + i\phi \mapsto (\mathfrak{C} + \eta)(A + i\phi).$$

Lemma 5.15. The section $d_A^*\phi \in \Omega^0(\mathfrak{g}_{P_{\mathbb{C}}})$ is preserved under the gradient flow of the perturbed Chern-Simons functional.

Proof. The perturbed Chern-Simons functional is invariant under gauge transformations in the connected component of the identity. The result then follows from Theorem 5.11. □

In what follows, we will need to understand how the parallel transport map varies when one varies the connection. The proof of the following result is very similar to Proposition 3.47, and we use the same notation here.

Proposition 5.16. Let $a \in \Omega^1(\mathfrak{g})$, and consider

$$U_s : [0,1] \to \mathrm{Isom}(E_{\gamma(0)}, E_{\gamma(t)})$$

which is the parallel transport defined with respect to the connection $A_s = A + sa$. Then

$$\frac{d}{ds}U_s(t) = U(t)^{-1}\int_0^t U(\tau)^{-1}(a + i\varphi)(\dot{\gamma}(\tau))U(\tau)d\tau$$

Proof. Differentiating (3.5) with respect to s, at $s = 0$, gives

$$\frac{d}{dt}\left(\frac{dU_s}{ds}\right) + A\frac{dU_s}{ds} + aU_0 = 0.$$

We can look for solutions of this equation (as a first order ODE in t) of the form

$$V(t) = U_0 f,$$

where $f(t) \in \mathfrak{isom}(E_{\gamma(0)}, E_{\gamma(t)})$. Substituting in gives

$$U_0 \frac{df}{dt} + aU_0 = 0,$$

and so

$$f = \int_0^t U(\tau)^{-1} a(\dot{\gamma}(\tau)) U(\tau) d\tau.$$

Hence we have that

$$\frac{dU_s}{ds} = (U(t))^{-1} \int_0^t U(\tau)^{-1} a(\dot{\gamma}(\tau)) U(\tau) d\tau.$$

□

Remark 5.17. Using the notation

$$U(b; a) : E_{\gamma(a)} \to E_{\gamma(b)}$$

of Definition 3.49, we have

$$\begin{aligned} \frac{dU_s}{ds} &= \int_0^t U(1;0) U(0,\tau) a(\dot{\gamma}(\tau)) U(\tau, 0) d\tau \\ &= \int_0^t U(1,\tau) a(\dot{\gamma}(\tau)) U(\tau, 0) d\tau. \end{aligned}$$

We now can start constructing the perturbations. Fix a loop $\gamma : S^1 \to Y$ and thicken it to get an embedding of the solid torus $\iota : D^2 \times S^1 \to Y$. For $z \in D^2$, denote by γ_z the path

$$e^{i\theta} \mapsto \iota(z, e^{i\theta}).$$

At each point x in the image of ι, where $x = \gamma_z(b)$ for some b, we can consider the parallel transport of elements in the bundle associated to the standard representation around γ_z starting and ending at x. This gives an element $U^z(1 + b; b) \in G_{\mathbb{C}}$ at each point, where $x = \iota(z, e^{ib})$.

To define η, let $f : \mathbb{C} \to \mathbb{C}$ be an entire holomorphic function and consider the map

$$\eta(A + i\phi) = \int_{D^2} f\left(\operatorname{tr}\left(U^z(1; 0)\right)\right) \mu \tag{5.2}$$

where μ is a two form on D^2 with support contained in the interior of D^2.

Proposition 5.18. The derivative of η with respect to $A + i\phi$ is given by

$$d_{A+i\phi}\eta(a+i\varphi) = \int_{D^2} f'(\mathrm{tr}(U^z(0))\,\mathrm{tr}\left(\int_{S^1} U^z(1;\tau)(a+i\varphi)(\dot\gamma_z(\tau))U^z(\tau;0)\right) d\tau d\mu.$$

Proof. Consider $F(s) = \eta(A+i\phi+s(a+i\varphi))$, for $a+i\varphi \in \Omega^1(\mathfrak{g}_{\mathbb{C}})$. Differentiating at $s = 0$ we have

$$\left.\frac{dF}{ds}\right|_{s=0} = \int_{D^2} f'(\mathrm{tr}(U^z(1;0))\,\mathrm{tr}\left(\frac{d}{ds}U^z(1;0)\right).$$

The result then follows from Proposition 5.16 □

We now introduce some notation to help with the formulae so that we can efficiently describe the gradient of the perturbations.

Definition 5.19. We define the sections $U_{\mathfrak{A}} \in \Omega^0(N(\gamma); AdG)$ and $V_{\mathfrak{A}} \in \Omega^0(N(\gamma); \mathfrak{g}_{P_{\mathbb{C}}})$, for $x = \iota(z,b)$, to be, respectively, the holonomy and the trace free part of the holonomy of the connection $\mathfrak{A}$ around γ_z starting at b, that is

$$U_{\mathfrak{A}}(x) = U^z(1+b;b)$$
$$V_{\mathfrak{A}}(x) = U^z(1+b;b) - \frac{1}{n}\mathrm{tr}\left(U^z(1+b;b)\right)\mathrm{Id},$$

here n is the dimension of the group $GL(n;\mathbb{C})$ that we are embedding $G_{\mathbb{C}}$ in, to define the trace and matrix products.

Corollary 5.20. The critical points of the gradient flow of $\mathfrak{C}_\eta = \mathfrak{C} + \eta$ are given by

$$F_{\mathfrak{A}} = -f'\left(\mathrm{tr}(U_{\mathfrak{A}})\right)V_{\mathfrak{A}}\mu. \tag{5.3}$$

Here $U(x)$ is defined to be 0 on the exterior of the image of $S^1 \times D^2$ and equal to the holonomy of the connection $A + i\phi$ starting and ending at x going around the S^1 when x is in the image. Similarly V is the trace free part of U.

5.4 Sobolev Estimates for the Perturbed Hessian

In this section we consider the Hessian of our perturbation

$$\eta : \mathfrak{A} \mapsto \int_{D^2} f(\mathrm{tr}(\mathfrak{A}))\mu.$$

Lemma 5.21. The Hessian of the perturbation η, viewed as a map $T_{\mathfrak{A}}\mathcal{A}_{\mathbb{C}} \to T_{\mathfrak{A}}\mathcal{A}_{\mathbb{C}}$

$$H_\eta : a \mapsto \frac{d}{dt}(\nabla\mathfrak{C})_{\mathfrak{A}+ta}\Big|_{t=0}$$

is given by

$$H_\eta(a) = *\alpha V_{\mathfrak{A}} \int_{S^1} tr(Ub_\tau)\mu + *\beta W_{\mathfrak{A}},$$

where α and β are constants depending on $\mathfrak{A}$ and f, and W is the trace free part of the operator

$$b \mapsto \int_{S^1} U(1;\tau)b_\tau U(\tau;0)d\tau \in \mathrm{Ad}G_x,$$

where the integral is taken around the loop γ_z starting at $x \in N(\gamma)$.

Proof. Let $a, b \in \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}})$. Then

$$H_\eta = *f''(\mathrm{tr}(U_{\mathfrak{A}}))\,\mathrm{tr}\left(\frac{dU_{\mathfrak{A}}}{db}\right) V_{\mathfrak{A}}\mu + *f'(U)\frac{dV_{\mathfrak{A}}}{db}\mu,$$

here

$$\frac{dU_{\mathfrak{A}}}{db} = \frac{dU_{\mathfrak{A}+tb}}{dt}\Big|_{t=0}$$

and the same for V. We have that

$$\frac{dU(x)}{db} = \int_{S^1} U(1;\tau)b_\tau U(\tau;0)d\tau$$

and so the trace of this is

$$\int_{S^1} \mathrm{tr}(Ub_\tau).$$

The result follows. ☐

Proposition 5.22. Let $\mathfrak{A}$ be a smooth connection on $P_{\mathbb{C}} \to Y$. The operator $H_\eta : L^2_1(\mathfrak{g}_{P_{\mathbb{C}}}) \to L^2(\mathfrak{g}_{P_{\mathbb{C}}})$ at Y is a compact operator.

Proof. The proof of compactness comes down to the fact that for a fixed connection, $\mathfrak{A}$, defining the parallel transport maps U, the map

$$a \mapsto \mu \int_0^1 U(1;\tau)aU(\tau, 0) \tag{5.4}$$

is a bounded endomorphism of $L^2_k(Y)$. To see this, first restrict to the support of μ. We can consider the projection $\pi : D^2 \times S^1 \times S^1$ given by

$$\pi : (z, e^{i\theta}, e^{i\psi}) \mapsto (z, e^{i(\theta+\psi)})$$

and pull back the section $a \in L^2_k(\Omega^1(\mathfrak{g}_{P_\mathbb{C}}))$. Since S^1 is compact, we have

$$\pi^*(a) \in L^2_k(D^2 \times S^1 \times S^1).$$

Moreover, we can consider the integral in (5.4) as

$$a(x) \mapsto I = \int_{S^1} \pi^* a(\tau(x), e^{i\psi}) d\psi.$$

where $\tau : D^2 \times S^1 \hookrightarrow D^2 \times S^1 \times S^1$ given by

$$(z, e^{i\theta}) \mapsto (z, e^{i\theta}, 1).$$

By Lemma 2.14 it follows that

$$\|I\|_{L^2_k} \leq \|a\|_{L^2_k},$$

and hence the compactness of $a \mapsto I$ follows from the compactness of the embedding $L^2_k \to L^2$. $\square$

Remark 5.23. The above theorem shows that the deformation theory of the perturbed complexified Chern-Simons functional is much the same as the normal Chern-Simons functional, since it differers by a compact operator. For example, the index theory is the same.

5.5 The Moduli Space of Perturbed Flat Connections

In this section we consider the moduli space of perturbed flat connections.

Definition 5.24. An element of a real or complex Lie group, $g \in G$, is semisimple if the adjoint representation

$$\mathrm{Ad}_g : \mathfrak{g} \to \mathfrak{g}$$

is diagonalisable.

Theorem 5.25. Let G be a compact, semisimple lie group, and let $G_\mathbb{C}$ denote its complexification. Then almost every element (with respect to the measure induced by the invariant metric on G) is semisimple.

We now start the classification result. It will turn out that for solutions to the perturbed flat equations will be reducible when restricted to the support of the perturbation, which will be the solid torus $S^1 \times D^2$. Typically the structure group of the reducible connection will be $\mathbb{C}^*$. The following result is useful when studying such connections. Let μ be a non-zero two form on the interior of D^2, equal to 0 on the boundary and f be an arbitrary function.

Lemma 5.26. Let $\mathfrak{A}$ be a $G_{\mathbb{C}}$ connection on a bundle $P_{\mathbb{C}} \to S^1 \times D^2$. Consider the section U of the bundle

$$P_{\mathbb{C}} \times_{\mathrm{Ad}} G_{\mathbb{C}},$$

given by the holonomy of the connection around the S^1 direction, and suppose that for some $x \in S^1 \times D^2$, $\mathrm{Stab}(U) \cong (\mathbb{C}^*)^r$. Denote by $V \in \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}})$ the trace free part of U. Suppose also that

$$F_{\mathfrak{A}} = -f'(\mathrm{tr}(U_{\mathfrak{A}}))V_{\mathfrak{A}}\mu. \tag{5.5}$$

Then the connection reduces to a $(\mathbb{C}^*)^r$ connection, such that

$$m_i = \exp\left(f'\left(\mathrm{tr}(U)\right) l_i\right),$$

where the vector defined by the elements m_i is the holonomy around the ∂D^2 direction, and similarly l_i describes the holonomy around the S^1 direction.

Proof. We follow the proof of [BD95, Lemma 4]. Let v be a vector field on $S^1 \times D^2$ in the direction of D^2 (transverse to the S^1 direction) and vanishing on the boundary, and consider $\Phi_t : S^1 \times D^2 \to S^1 \times D^2$ to be the one parameter family of diffeomorphisms generated by v. We can lift these diffeomorphisms to the bundle $P_{\mathbb{C}}$ using the connection $\mathfrak{A}$. To do this we use the definition of the connection as a horizontal distribution in $P_{\mathbb{C}}$ and consider the one parameter family of diffeomorphisms generated by the lift of the vector field. We can then consider a family of connections

$$t \mapsto \mathfrak{A}_t = \Phi_t^*(\mathfrak{A})$$

Differentiating this family at $t = 0$ gives

$$\left.\frac{d}{dt}\Phi_t^*(\mathfrak{A})\right|_{t=0} = i_v F_{\mathfrak{A}}.$$

on $P_{\mathbb{C}}$. The contraction of the curvature with v gives a one form with no component in the S^1 direction, due to (5.5). As above, we have that the variation in the holonomy U is given by

$$\frac{dU_t}{dt} = \int_0^1 U(1;\tau)\frac{d\mathfrak{A}}{dt}U(\tau;0)d\tau$$

Hence, if

$$\frac{d\mathfrak{A}_t}{dt} = \frac{d\Phi_t^*(\mathfrak{A})}{dt} = 0,$$

then the holonomy is constant along the family. Since we defined Φ using the connection, this is saying that $\nabla_{\mathfrak{A}} U = 0$ in the directions of D^2. But it is also true that $\nabla_{\mathfrak{A}} U = 0$ in the S^1 direction by the definition of Holonomy. Hence U is covariantly constant and so defines a bundle reduction on the solid torus. Since we assume that the stabiliser of U is equal to $\mathbb{C}^*$ at a point, it must be equal to $\mathbb{C}^*$ on the whole of $S^1 \times D^2$.

To get the constraints on the holonomy, we have that the curvature and trace free holonomy of the connection will take values in the relevant subset of the Lie algebra. If $V = 0$ then the connection is flat and so the holonomy around D^2 would be trivial.

Since all $\mathbb{C}^r$ bundles over $S^1 \times D^2$ are trivial, we can trivialise the bundle and so we have that

$$F_{\mathfrak{A}} = -f'(\mathrm{tr}(U)) \begin{pmatrix} l_1 \\ \vdots \\ l_r \end{pmatrix} \mu,$$

where λ_j is the jth component of the holonomy in the S^1 direction. If we now consider the holonomy around ∂D^2, we have that

$$\frac{dU}{dt} = -\mathfrak{A}U,$$

and since $(\mathbb{C}^*)^r$ is Abelian, using Stokes' theorem on the trivialisation, we have that

$$\frac{d}{dt}(\log U) = -\int_{\partial D^2} \mathfrak{A} = -\int_{D^2} F_{\mathfrak{A}} = f'(\mathrm{tr}(U)) \begin{pmatrix} l_1 \\ \vdots \\ l_r \end{pmatrix},$$

and so

$$\begin{pmatrix} m_1 \\ \vdots \\ m_r \end{pmatrix} = \exp\left(f'\left(\sum_r l_r\right) \begin{pmatrix} l_1 \\ \vdots \\ l_r \end{pmatrix}\right). \tag{5.6}$$

Here we apply the exponential function to each element of the vector. $\square$

Theorem 5.27. There is a one to one correspondence between

1. Critical points of the perturbed Chern-Simons functional (as defined above), with $d_A^*\phi = 0$,
2. Solutions to $F_{\mathfrak{A}} = -f'(\mathrm{tr}(U_{\mathfrak{A}}))V_{\mathfrak{A}}\mu$, $d_A^*\phi = 0$.
3. Flat connections on the complement of $\mathrm{supp}(\mu)$ satisfying $d_A^*\phi = 0$, with holonomy on the torus boundary satisfying (5.6).
4. Flat connections on the complement of $\mathrm{supp}(\mu)$ with reductive holonomy representation and with holonomy on the torus boundary in $(\mathbb{C}^*)^r$ satisfying (5.6).

Proof. The equivalence between 1 and 2 follows Corollary 5.20. For 2 implies 3 we can use Lemma 5.26. We can always assume that we are in the generic case where the holonomy gives a reduction $(\mathbb{C}^*)^r$. If not then the reducible connection is not reductive and so we will not be able to solve $d_A^*\phi = 0$ on the manifold $S^1 \times D^2$. For 3 if and only if 4 we can use Proposition 4.12, where we split up the manifold into $S^1 \times D^2$ and its complement. We deduce that the representation is reductive on the solid torus and its complement by the existence of the solution $d_A^*\phi = 0$.

To prove 3 implies 2, we show how to fill in a connection with the given holonomy properties. To do this we can work on the trivialisation of the $(\mathbb{C}^*)^r$ bundle over $S^1 \times D^2$. The equation we want to solve is

$$d(A + i\phi) = -f'\left(\sum_r l_r\right)\begin{pmatrix} l_1 \\ \vdots \\ l_r \end{pmatrix}\mu, \quad d^*(A + i\phi) = 0.$$

This is just a vector of Poisson's equation for differential forms. By, for example, [Sch95, Theorem 3.2.5] this equation can be solved for any choice of boundary data

$$(A + i\phi)|_{\partial X} = \psi$$

that satisfies the integrability condition

$$\left\langle -f'\left(\sum_r l_j\right) l_i\mu, \eta \right\rangle = \int_{S^1\times\partial D^2} \psi_i \wedge *\eta,$$

where η is the harmonic representative of the homology class defined by μ that is trivial when restricted to the boundary. This condition is a holonomy constraint on the boundary connection. It is implied by the condition

$$m_i = \exp(f'(\mathrm{tr}\left(\sum_j l_j\right) l_i)$$

on the holonomy of the boundary connection.

Hence, if we're given a flat connection on the complement of the perturbation, satisfying the holonomy conditions then we can pick a trivialisation of the bundle over the boundary in which the connection is diagonal, extend this trivialisation to the support of the perturbation and then add in the solution to the above problem to get a full solution to the perturbed equations. Note that the gauge transformation required to trivialise the bundle on the boundary will generally be in $G_{\mathbb{C}}$, but in Proposition 4.12 we show that for any choice of boundary trivialisation, one can find a gauge transformation, trivial on the boundary, to solve $d_A^*\phi = 0$. □

5.5.1 Multiple Loops

To end our definition of perturbations to the complex Chern-Simons functional, we consider the case of multiple loops. The previous results for considering only one loop carry over very easily to this case.

Definition 5.28. Let $\gamma_1, \dots, \gamma_n$ be a collection of disjoint loops with non pairwise intersecting solid torus neighbourhoods $N(\gamma_i) \cong S^1 \times D^2$, and let $D_i \subseteq N(\gamma_i)$ be the disks corresponding to $t = 0$ on S^1. Pick holomorphic functions $f_1, \dots f_n$ and two forms, μ_i, with integral 1 on each D_i. Define the sections $U^i_{\mathfrak{A}}$ as above to be the holonomy of the connection $\mathfrak{A}$ around the S^1 factor of$N(\gamma_i)$ starting at a point $x \in N_{\gamma_i}$. Using this data, we define a perturbation to the Chern-Simons functional, η, to be

$$\mathfrak{C}_\eta = \mathfrak{C} + \eta = \mathfrak{C} + \sum_i \int_{D_i} f(\mathrm{tr}(U_i))\mu_i.$$

From now on, unless stated otherwise, the symbol η will denote such a perturbation.

5.6 Energy Estimates

A crucial quantity to control when studying connections, $\mathfrak{A}$, on complexified bundles is the L^2 norm of ϕ, where we write

$$\mathfrak{A} = A + i\phi.$$

This follows from the fact that from the Weitzenbock formula we have bounds of the form

$$\|\nabla_A \phi\|_{L^2} + \|[\phi, \phi]\|_{L^2}^2 \leq C\|\phi\|_{L^2}^2$$

and from this we can derive bounds on the curvature of the connection A. Given bounds on the curvature of A we can then find appropriate gauges using Uhlenbecks theorem and derive regularity results.

When studying solutions to the perturbed flat equations, we can derive estimates by considering the equations over two regions: the support of the perturbation and its complement. On the support of the perturbation we have that the curvature of the connection is bounded by the holonomy around the loops defining the perturbation. Away from the perturbation, the connection is flat. Moreover, we have that the metric h defining the G reduction minimises the harmonic map energy

$$E = \int_{Y \setminus N(k)} \|\phi\|_{L^2}^2.$$

over the metrics satisfying the given boundary conditions. Hence, we can get a bound on the norm of ϕ simply by calculating $\|\phi\|_{L^2}^2$ for any choice of G reduction. In this section, we study producing a bound on the minimum harmonic energy by calculating the energy for a given representative. This will allow us to then bound $\|\phi\|_{L^2}^2$ for the flat equation, by a function of the holonomy.

Lemma 5.29. Let $\mathfrak{A}$ be a flat $G_{\mathbb{C}}$ connection on a compact manifold, Y (possibly with boundary), and let $\rho : \pi_1(Y) \to G_{\mathbb{C}}$ be the holonomy representation. Pick a G reduction defined by a smooth metric of the bundle P. This defines an equivariant map

$$h : \tilde{Y} \to G_{\mathbb{C}}/G.$$

with energy bounded by a function of the representation ρ, the choice of G reduction and the topology of Y.

Proof. Pick a finite cover, $U_1, \ldots, U_n$ of Y with disjoint intersections under the action of π_1 when pulled back to the universal cover. For each element of the cover we can pick a representative $\tilde{U}_i$ in $\tilde{Y}$ such that $\pi : \tilde{Y} \to Y$ restricts to a homeomorphism on $\tilde{U}_i$. The union of all the $\tilde{U}_i$ is a bounded set in $\tilde{Y}$ and hence contained in a compact set. Moreover h differentiable on the U_i, with derivative depending on the representation ρ and initial metric h_0. We can bound the norm squared of the derivative h by a constant, K say, on each $\tilde{U}_i$. We then have that

$$\|\phi\|_{L^2}^2 \leq \sum_i \int_{U_i} |\nabla h|^2 \leq Kn.$$

□

Corollary 5.30. Let $\mathfrak{A}$ be a η–perturbed flat connection with finite holonomy representation, in L_1^2 with respect to some background connection. Then $\mathfrak{A}$ is gauge equivalent to a smooth connection, with bounds on derivatives depending only on $\|\phi\|_{L^2}$.

Proof. We can split the manifold into two sections, $Y \setminus \operatorname{supp}(\eta)$ and $\operatorname{supp}(\eta)$. On compact subsets of the support of η, we can bound the solution by elliptic estimates for the operator $d \oplus d^*$ and use a bootstrapping argument, since we know that the holonomy is bounded. For compact subsets of $Y \setminus \operatorname{supp}(\eta)$ we can use elliptic regularity of the Yang mills equations. To do this, note that the L^2 norm of ϕ bounds the L^2 norm of $\phi \wedge \phi$ and hence the curvature of F_A. We can then pick sufficiently small subsets to bound the L_1^2 norm of the connection one form and the L_1^2 norm of ϕ after a gauge transformation and then apply the regularity result for solutions to the complex Yang-Mills equation. We are left needing to consider an estimate of the L_1^2 norm of $\mathfrak{A}$ on the boundary torus. We can do this by picking a collar neighbourhood of the torus and considering interior estimates for the boundary value problem imposed by identifying the bundle over the neighbourhood. □

5.7 Perturbed Kapustin-Witten Equations

We end this chapter by defining the perturbed Kapustin-Witten equations. In the following, let $\eta : \mathcal{A}_{\mathbb{C}} \to \mathbb{C}$ be a holonomy perturbation functional, and denote by η' its gradient.

Definition 5.31. Let $\theta \in [0, \pi)$. Then the θ–Kapustin-Witten equations on the manifold $\mathbb{R} \times Y$ are defined to be the negative gradient flow equations of

$$\mathfrak{C} + \eta$$

For a path of connections $\mathfrak{A}$, with $\mathfrak{A}(\partial_t)$, this is the equation

$$\frac{d\mathfrak{A}}{dt} = - * e^{2i\theta} F_{\mathfrak{A}} + \eta'(\mathfrak{A}), d_A^* \phi = 0$$

Written in four-manifold notation, this is the equation

$$(e^{i\theta}(F_{\mathfrak{A}} + \hat{\eta}'))^+ = 0, d_A^* \phi = 0,$$

where $\hat{\eta}'(\mathfrak{A})$ is the pullback of $\eta'(\mathfrak{A})$ on each slice. That is, if γ is a loop defining the perturbation, and μ is the pull back of $\mu \in \Omega^2(N(\gamma))$ to $\mathbb{R} \times N(\gamma)$, then

$$\hat{\eta}' = -f'(U_{\mathfrak{A}_t}) V_{\mathfrak{A}_t} \mu.$$

Remark 5.32. The θ_v–Kapustin-Witten equations can be defined analogously for a function $\theta_v : \mathbb{R} \to \mathbb{R}$.

Chapter 6

Flat Connections on Heegaard Splittings

In this chapter we study perturbed flat $SL(2;\mathbb{C})$ connections on three-manifolds where the perturbation is defined using a Heegaard splitting. A Heegaard splitting gives a natural way to represent the fundamental group of a three-manifold. We use these representations of the fundamental group to prove a transversality result for our perturbed flat connections with these perturbations. Heegaard splittings are also used to define the Casson invariant.

6.1 Heegaard Splittings

We include these definitions here for completeness, but this is standard material in low dimensional topology.

Definition 6.1. A 0–handlebody is homeomorphic to S^3. An $n+1$–handlebody is homeomorphic to an n–handlebody connected sum a solid torus.

Definition 6.2. A Heegaard splitting of a closed three manifold M is a decomposition

$$M = V \cup_f W,$$

where V, W are n–handle bodies and $f : \partial V \to \partial W$ is a homeomorphism with which the boundaries are identified.

Example 6.3. The three sphere has admits a Heegaard splitting as

$$S^3 = D^2 \times S^1 \cup D^2 \times S^1,$$

where the meridian of each boundary torus gets identified with the longitude of the other.

Theorem 6.4. Any three manifold admits a Heegaard splitting.

Proof. See, for example, [Sch00, Section 2]. □

Given a Heegaard splitting of a three manifold, this can be trivially extended by direct summing two tori which we identify as S^3, so the resulting gluing gives the connected sum of the manifold with S^3.

Definition 6.5. We say that two Heegaard splittings are isotopic if there is a homotopy of homeomorphism sending the boundary in one of the splittings to the other. Two splittings are stably isotopic if they are isotopic after being trivially extended.

Theorem 6.6. Any two Heegaard splittings of the same manifold are stably isotopic.

Proof. See, for example, [Sch00, Theorem 7.1]. □

We now explain how the fundamental group of the three-manifold relates to the fundamental groups of the components V, W of a Heegaard splittings.

Lemma 6.7. The fundamental group of a g–handlebody is the free group on g letters.

Proof. This follows from the formula for the fundamental group under direct sums, or by noting that the handlebody deformation retracts onto the wedge product of g circles. □

The generators of the fundamental group of the g–handlebody may be visualised as loops going round each of the handles.

Lemma 6.8. The fundamental group of the orientable surface of genus g may be presented as

$$\pi_1(\Sigma_g) = \left\langle \lambda_1, \dots \lambda_g, \mu_1, \dots \mu_n \,\middle|\, \prod_{i=1}^{g} [\lambda_i, \mu_i] = 1 \right\rangle .$$

Proof. See, for example, [Hat02, Discussion after Proposition 1.26]. □

The generators of this group may be visualised as going around the handles of the boundary of the handle body.

6.2 Perturbed Representations on Handlebodies

In this section we study the representation theory side of the perturbed flat connection correspondence in Theorem 5.27 when we use a Heegaard splitting to pick our loops. We restrict to the group $SL(2;\mathbb{C})$. Strictly speaking what we describe here doesn't fit into the framework because our loops will contain a common base-points coming from the definition of the fundamental group, but by perturbing our loops near their intersections we can assume that our loops are disjoint. Changing the base point of our connections will conjugate the holonomy and in principle one could add account for these perturbations in the system of equations that. We don't do this however and assume that the classification results of flat connections from Chapter 5 carry over even in the case of intersecting loops.

Our starting point for this analysis is to pick a collection of curves γ_i and ξ_i which generate the fundamental groups of each of the handlebodies. We think of γ_i going around the ith handle, and the same for ξ_i. If we remove a neighbourhood of each of these loops then the fundamental group of each of the drilled out handlebodies is given by

$$\langle \lambda_1, \dots, \lambda_g, \mu_g, \dots, \mu_g | [\lambda_i, \mu_i] = 1 \rangle.$$

The fundamental group of the three manifold with the neighbourhoods of the loops removed can then be written as

$$\tau_g(Y) = \left\langle \lambda_1^1, \lambda_2^1, \dots, \lambda_g^2, \mu_1^1, \dots, \mu_g^2 | \mathcal{R} \right\rangle$$

where $\mathcal{R}$ is a collection of relations coming from the embeddings

$$V', W' \hookrightarrow Y,$$

along with the commutation relations, where V' and W' are the drilled out handlebodies.

Lemma 6.9. For the group $\tau_g(Y)$ associated to a Heegaard splitting, we can pick $\mathcal{R}$ so that it consists of $4g$ constraints. Let α_i be defined such that

$$\lambda_i^2 - \alpha_i(\lambda_1^1, \dots, \mu_g^1)) = 0,$$

$$\mu_i^2 - \alpha_{g+i}(\lambda_1^1, \dots, \mu_g^1) = 0$$

Then the map $\lambda_1^1, \dots, \mu_n^1 \to \lambda_1^2 \dots, \mu_n^2$ is invertible.

Proof. It is a standard result in the theory of Heegaard splittings that for the Heegaard decomposition V, W, f the fundamental group is presented by the generators λ_i^j and $2g$ relations. Hence we can pick our $4g$ relations by adding in the commutativity constraints coming from the torus boundaries. Fix a base point on the intersection boundary of V' and W' to define the fundamental groups. Note that the generators of the fundamental groups of V' and W' are all homotopic to loops on the intersection. This gives a bijective map $f : \pi_1(V') \to \pi_1(W')$, and this bijection is invertible. This allows us to write α_i explicitly and define the inverse map. □

We want to frame the problem of finding representations that satisfy our holonomy constraints in terms of the regular value theorem. We will prove that we can pick our parameters such that 0 is a regular value and hence the moduli space of solutions is 0 dimensional. As such we can consider the error of the relations considered in the previous proof being zero. We do this by considering a representation of the free group on $4g$ letters

$$\hat{\rho} : F_{4g} \to SL(2; \mathbb{C}). \tag{6.1}$$

This is nothing more than a choice of $4g$ matrices. Writing out the relations we have

$$\hat{\rho}(\mu_i^1) - \exp(f_i(\hat{\rho}(\lambda_i^1))) = 0, \tag{6.2}$$

$$\hat{\rho}(\mu_i^2) - \exp(h_i'(\hat{\rho}(\lambda_i^2))) = 0, \tag{6.3}$$

$$\hat{\rho}(\lambda_i^2) - \alpha_i(\hat{\rho}(\lambda_1^1), \ldots, \hat{\rho}(\mu_g^1))) = 0, \tag{6.4}$$

$$\hat{\rho}(\mu_i^2) - \alpha_{g+i}(\hat{\rho}(\lambda_1^1), \ldots, \hat{\rho}(\mu_g^1)) = 0 \tag{6.5}$$

If these relations are satisfied then ρ descends to a representation of $\tau_g(Y)$. The functions f_i, h_i' are the functions in the definition of the perturbation – to ease the notation we have decided not to define these as maps from $SL(2; \mathbb{C}) \to \mathbb{C}$ rather than applying a holomorphic function on $\mathbb{C}$ to the trace. We are slightly abusing notation here and the first two relations should be read as the condition that $\hat{\rho}(\mu_i^j)$ commutes with $\hat{\rho}(\lambda_i^1)$ and the eigenvalues satisfy the exponential relation.

Definition 6.10. Let the map

$$\mathcal{F} : SL(2; \mathbb{C})^{4g} \to SL(2; \mathbb{C})^{4g} \tag{6.6}$$

be the error coming from a choice of $\hat{\rho}(\lambda_i^j), \hat{\rho}(\mu_i^j)$.

Lemma 6.11. We can pick polynomials $f_i, h_i : \mathbb{C} \to \mathbb{C}$ such that 0 is a regular value of the map $\mathcal{F}$ in (6.6).

Proof. Consider the map from $\mathbb{C}$ to maps between holomorphic functions where we have $z \mapsto (f \mapsto f + z)$. The derivative of the map $\mathcal{F}$ can be seen as a section of $\mathcal{F}^*TSL(2;\mathbb{C})^{4g}$ valued one forms. Correspondingly, we can define a section of a suitable bundle over to $SL(2;\mathbb{C})^{4g} \times \mathbb{C}^{2g}$, where at the points with value z in the $\mathbb{C}^{2g}$ factor, the section is defined by defined by differentiating the map $\mathcal{F}_z$, which is just the map $\mathcal{F}$ defined using the functions $f_i + z_i$ and $g_i + z_{g+i}$. Now pick arbitrary holomorphic functions f_i. We claim that the derivative of the function is surjective at all points. To see this note that when for any vector in the tangent space of the first two conditions we can select z_i such that the derivative of the first two equations is equal to the given vector. For the second two conditions we have that the maps α_i give smooth bijections, and so any tangent vector in the second two equations can be achieved by varying the λ_i^2 and μ_i^2. Hence, in particular, 0 is a regular value of this map, and we can use the implicit function theorem to study □

As a result of this we get the following theorem.

Theorem 6.12. There exist perturbations of the complex Chern-Simons functional such that the moduli space of critical points if 0 dimensional.

Lemma 6.13. There can be infinitely many solutions to the perturbed flat $G_{\mathbb{C}}$ connection equation.

Proof. The simplest example is to consider S^3 as the union of two 1 handlebodies. Under the identification of the fundamental groups of the handlebodies, the relations α_i are given by

$$\lambda^2 = \alpha_1(\lambda^1, \mu^1) = \mu^1$$
$$\mu^2 = \alpha_2(\lambda^1, \mu^1) = \lambda^1.$$

Hence the condition on the holonomies, in terms of only λ_1, is

$$\lambda^1 - \exp(h'(\exp f'(\lambda^1))) = 0.$$

It follows that there are infinitely many solutions since the left hand side holomorphic function has an essential singularity at infinity, and so will take almost every value infinitely many times. □

The fact that there are infinitely many solutions is an artefact of the fact that the condition on the holonomies around each boundary torus is given in terms of the exponential function, which is transcendental. If polynomial conditions were given, then one would

Definition 6.14. We call a perturbation regular if all the perturbed flat connections are isolated.

Definition 6.15. We call a perturbation coercive if, whenever the holonomy of a connection $\mathfrak{A}_n$ tends to infinity along a sequence of regular critical points, the perturbed Chern-Simons functional also tends to infinity.

Proposition 6.16. We can choose our perturbation to be coercive.

Proof. If the holonomy is tending to infinity along a sequence of regular points, then we must have that there is an i such that one of the eigenvalues of $\hat{\rho}(\lambda_i)$ tends to infinity. We can make sure $f_i(\hat{\rho}(\lambda_i))$ tends to infinity along this sequence by picking the f_i to be polynomials. Now we show that the perturbed Chern-Simons functional is tending to infinity. It could be the case that both terms in the sum

$$\mathfrak{C}(\mathfrak{A}_n) + \eta(\mathfrak{A}_n)$$

are unbounded but the sum is still bounded. However, note that the norm of $\mathfrak{C}(\mathfrak{A})$ is bounded by $\|\mathfrak{A}_n\|_{L^2}\|\mathfrak{A}_n\|_{L^2_1} + \|\mathfrak{A}_n\|^3_{L^3}$, and hence by $\|\phi\|^2_{L^2}$. We can also bound $\|\phi\|_{L^2}$ by the logarithm of the holonomy, as in Lemma 5.29 and hence, as the holonomy tends to infinity, we will have that the Chern-Simons functional is growing as the log of the holonomy and hence the sum cannot be finite. ☐

Remark 6.17. Note that the above argument works just as well for reducible connections as it does for irreducible ones. Reducible connections are classified by maps

$$\rho : H^1(Y) \to \mathbb{C}^*,$$

and hence will exist in $\dim H^1(Y)$ dimensional families.

We now prove a result about the perturbed Chern-Simons functional of that we will need to prove a compactness result for solutions to the Kapustin-Witten equations with bounded energy.

Proposition 6.18. Suppose $\mathfrak{A}_n$ is a sequence of infinitely many distinct perturbed flat connections for a regular perturbation η. Then the values

$$z_n = \mathfrak{C}_\eta(\mathfrak{A}_n)$$

are such that

$$\liminf_{n\to\infty} |z_n| = \infty.$$

Proof. Since the critical points are isolated, it follows that the along the sequence, we must have that for some i the values

$$\hat{\rho}(\lambda_i^1) \to \infty.$$

otherwise we could pick a convergent subsequence of bounded holonomies that would converge in the bounded subset to a solution (since the map is continuous). $\square$

Chapter 7

Knot Theory

7.1 Relevant Knot Theory

In this section we review some concepts from knot theory that we will need in Chapter 8. Typically, for the descriptions of established knot invariants we will be thinking of knots in S^3. One of the advantages of gauge theoretic invariants is that they can be defined on a wider class of three manifolds. In this section we will define the invariants for as wide a class of manifolds that we can. Since the ideas from knot theory we will need are fairly rudimentary, we include all the definitions so that no prior knot theory knowledge is required to understand the thesis.

Definition 7.1. Let Y be a three manifold. A knot in Y is an embedded loop $S^1 \hookrightarrow Y$. We will always assume that our knots are smoothly embedded. A link in Y is a finite collection of knots, that is an embedding $\coprod S^1 \hookrightarrow Y$. A link that is the image of n disjoint copies of S^1 is an n component link.

Definition 7.2. Two knots are equivalent (isotopic) if there exists a homeomorphism

$$h : [0,1] \times Y \to Y \tag{7.1}$$

Such that $h_0 = \mathrm{Id}$ and $h_1(k_1) = k_2$.

We will only consider smooth knots. For smooth knots, two knots are isotopic if and only if there exists a diffeomorphism in the connected component of the identity sending one of the knots to the other. The Jones Polynomial Given a knot or link in a three manifold, a particularly useful way to study it is to study the three manifold where the knot has been removed.

Definition 7.3. Let $L : \coprod S^1 \hookrightarrow Y$ be a link (or knot) in a three manifold Y. The link (or knot) complement is the three manifold with boundary

$$Y \setminus N(L), \tag{7.2}$$

where $N(L)$ is a neighbourhood of the image of L, that is sufficiently small so that it can be seen as an embedding of solid tori. The link complement of an n component link has boundary given by n disjoint tori.

Remark 7.4. Sometimes the terms knot and link will be used interchangeably. In most cases, unless explicitly stated, we will use the term knot to mean either a knot or link and imagine links with one component, but the extension to general n component links will be straightforward.

Definition 7.5. The link (or knot) group is the fundamental group of the link (or knot) complement.

Many knot and link invariants are constructed by considering the topology of the knot complement.

7.1.1 Examples of Knots

In Chapter 8 we will study critical points of $\mathfrak{C}_\eta$ with the perturbations defined along certain knots. As such, it is useful to define some example knots, and give some explicit descriptions.

Definition 7.6 (The Unknot)**.** The unknot in S^3 is the isotopy class of knots equivalent to the standard embedding $S^1 \hookrightarrow \mathbb{R}^2 \hookrightarrow \mathbb{R}^3 \hookrightarrow S^3$. Such an element $k : S^1 \hookrightarrow S^3$ is characterised by the fact that it bounds an embedded disk. For a general three manifold we define *an unknot* to be any knot that bounds an embedded disk. In general, there could be many isotopy classes of such knots.

Definition 7.7 (Torus Knots)**.** Let $k : S^1 \to Y$ be a knot, and consider the solid torus neighbourhood $N(k)$ of k. On the boundary torus, we can consider the map, for $(p, q) \in \mathbb{Z} \times Z \setminus \{(0, 0)\}$

$$f_{p,q} = S^1 \to S^1 \times S^1, \quad e^{i\theta} \mapsto (e^{ip\theta}, e^{iq\theta}). \tag{7.3}$$

Here we pick the first S^1 factor of T^2 to correspond to the meridian of the knot neighbourhood. The (p, q) torus knot with respect to k, $(p, q)_k$, is the isotopy class of knots given by the image of $f_{p,q}$. The (p,q) torus knot in S^3 is the (p,q) torus knot with respect to the unknot.

Example 7.8. The (right handed) trefoil in S^3 is the $(2,3)$ torus knot.

Lemma 7.9. The knot group of the (p,q) torus knot in S^3 is

$$\langle x, y \mid x^p = y^q \rangle \tag{7.4}$$

Definition 7.10 (Hyperbolic Knots)**.** A knot k in S^3 is hyperbolic if its knot complement $S^3 \setminus N(k)$ admits a hyperbolic metric - that is a metric of constant negative sectional curvature. The (hyperbolic) volume of a hyperbolic knot is the volume of its knot complement with the hyperbolic metric.

Definition 7.11. The Hopf link is two unknots linked together.

Lemma 7.12. The link group of the Hopf link is $\mathbb{Z} \times \mathbb{Z}$

Proof. We can decompose S^3 as two solid tori (which we think of as $S^1 \times D^2$), identified along the torus boundary. The Hopf link can be seen as the union of the two curves $S^1 \times \{0\}$. Hence the link complement deformation retracts onto the boundary torus. □

Some knot invariants require extra data to define. In this case the knot invariants may not be true invariants but might also be invariants of the knot and the additional data. We now give some examples of some extra data that we will use to define knot invariants.

Definition 7.13. A framing of a link $l : L \hookrightarrow Y$ is a section of the bundle l^*TY that is nowhere colinear to the link direction, i.e. a non-zero section of the quotient i^*TY/TL.

Definition 7.14. The blackboard framing of a link diagram is the framing of a link where the framing vector is in the direction of the kernel of the derivative of the projection that gives the knot diagram.

Definition 7.15. The winding number (vector) of a framed knot (link) is the number of times the framing vector rotates around the knot (or each component of the link).

7.1.2 The A Polynomial

The main knot invariant we will need when describing perturbations is the A polynomial. For more details on the A and $\hat{A}$ polynomials, see [BC16].

Definition 7.16. A representation $\rho : G \to GL(n;\mathbb{C})$ is reductive if every G invariant subspace has a G invariant complement.

Reductive representations are a generalisation of irreducible representations and provide a notion of stability when considering representations of the fundamental group.

Lemma 7.17. Let M be a compact manifold with, possibly empty, boundary. The set of representations

$$\rho : \pi_1(M) \to SL(2;\mathbb{C}) \tag{7.5}$$

is an algebraic variety.

Proof. Since the fundamental group of M is finitely presented, the condition of finding a homomorphism into $SL(2;C)$ can be written as a system of polynomial equations on the entries of the value of ρ for each of the generators. This system can in turn be written as a system of polynomial equations in the matrix coefficients. Hence the set of representations forms an algebraic variety. □

Definition 7.18. Let M be a compact manifold, with or without boundary. The $SL(2;\mathbb{C})$ representation variety of M, $R(M)$, is the set of reductive representations of $\pi_1(M)$ into $SL(2;\mathbb{C})$ up to conjugation.

$$R(M) = \{f : \pi_1(M) \to SL(2;\mathbb{C}) | f \text{ is reductive}\} / SL(2;\mathbb{C}) \tag{7.6}$$

The character variety of M is the set variety defined by characters (traces) of representations.

Example 7.19. The character variety of T^2 is isomorphic to $(\mathbb{C}^*)^2/\mathbb{Z}_2$. This is because $\pi_1(T^2) = \mathbb{Z}^2$. Commuting elements of $SL(2;C)$ can be simultaneously upper triangularised. If this representation is reductive then we can view the representation as a pair of upper diagonal matrices:

$$M = \begin{pmatrix} \mu & 0 \\ 0 & \mu^{-1} \end{pmatrix}, \quad L = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix} \tag{7.7}$$

Here M represents the value of the representation on loops homotopic to the meridian of the torus (that is, transverse to the knot), and L represents the value on loops homotopic to the longitude (in the direction of the knot). The action of $\mathbb{Z}_2$ comes from simultaneously identifying M with M^{-1} and L with L^{-1}. This identification comes from conjugating the diagonal matrices by

$$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \in SL(2;\mathbb{C}) \tag{7.8}$$

which swaps the elements on the diagonal.

If we have a knot in a three manifold Y, then the embedding of the boundary of the knot complement induces a homomorphism

$$i_* : \pi_1\left(T^2\right) \to \pi_1\left(Y \setminus N(K)\right) \tag{7.9}$$

As a result we get a subset

$$\check{A} = \left\{g|_{\pi_1(T^2)}, g \in \chi\left(Y \setminus N(K)\right)\right\} \subseteq \chi\left(T^2\right). \tag{7.10}$$

Definition 7.20. We can consider the set $\check{A} \subseteq \chi(T^2)$ as the image

$$\tau : \left(\mathbb{C}^*\right)^2 \to \chi(T^2), \tag{7.11}$$

and define $\tilde{A} := \tau^{-1}(\check{A})$. The Zariski closure of $\tilde{A}$, $\bar{A} \subseteq \mathbb{C}^2$ is an affine variety, and hence the vanishing set of a polynomial. The A polynomial is this polynomial. That is, the A polynomial of a knot K is the polynomial $A(\lambda, \mu)$ such that

$$\bar{A} = A^{-1}(0). \tag{7.12}$$

Remark 7.21. Put another way, the A polynomial encodes the algebraic constraint of being able to extend a representation of the torus boundary fundamental group to the knot group.

Remark 7.22. The A ideal of a link is the generalisation of the A polynomial to links. An n component link will have n boundary tori, and so given a representation of the fundamental group of the link complement we can consider a point in $(\mathbb{C}^*)^{2n}$ corresponding to the eigenvalues of the boundary holonomy around each boundary component. As in the knot case, we can then take the Zariski closure of the set of these points. This can be defined by an ideal in the polynomial ring over $\mathbb{C}$ with $2n$ variables.

Example 7.23. In S^3, the A polynomial of the unknot is $1-L$. This can be seen by the fact that the longitude is contractible in the knot complement, and the fundamental group of the unknot is $\mathbb{Z}$. For any knot in S^3, the A polynomial contains a factor of $1-L$ coming from the subgroup of the knot group generated by the meridian. These representations are called the Abelian representations. The A polynomial of the trefoil is $(1-L)(L+M^6)$.

The $\hat{A}$ polynomial, defined in [BZ01], is a refined version of the A polynomial. Instead of defining $\hat{A}$ as the polynomial that defines a certain algebraic variety, one can consider each one dimensional component X_j of $\chi(Y \setminus N(K))$ that restricts to a one dimensional component in $\chi(T^2)$ and then define the $\hat{A}$ polynomial as the product of those factors. Up to repeated factors, A and $\hat{A}$ agree.

Definition 7.24. Let X_j denote the one-dimensional components of $\chi(Y \setminus N(K)$ which are one-dimensional when the associated representations are restricted to the torus boundary. Let A_j be the polynomial defining the restriction. We define the $\hat{A}$ polynomial to be the product

$$\hat{A} = \prod_j A_j. \tag{7.13}$$

Chapter 8

Example Critical Points

In this chapter we study the critical points of the complexified Chern-Simons functional and the perturbed Chern-Simons functional. As with Floer theory, these critical points serve will serve as asymptotic boundary conditions when studying the Kapustin-Witten equations on manifolds with tubular ends.

8.1 Facts about $SL(2;\mathbb{C})$

We will mainly consider examples with structure group $SL(2;\mathbb{C})$, and so before considering the geometric examples we recall some relevant facts about $SL(2;\mathbb{C})$.

We now recall some important facts about $SL(2;\mathbb{C})$ matrices.

Lemma 8.1 (Jordan Normal Form)**.** Let $M \in SL(2;\mathbb{C})$. Then either

1. We have $\mathrm{tr}(M) \neq \pm 2$. Then M is conjugate to the matrix

$$\begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix},$$

   where $\lambda + \lambda^{-1} = \mathrm{tr}(M)$, by an element of $SL(2;\mathbb{C})$. Or,

2. $\mathrm{tr}(M) = \pm 2$ in this case either M is equal to $\pm\,\mathrm{Id}$ or is conjugate, by an element of $SL(2;\mathbb{C})$ to

$$\pm \begin{pmatrix} 1 & \pm 1 \\ 0 & 1 \end{pmatrix}.$$

Proof. By the Jordan normal form theorem, for any matrix $M \in SL(2;\mathbb{C})$ there exists a $P \in GL(2;\mathbb{C})$ such that

$$PMP^{-1}$$

is in Jordan normal form. Furthermore, by multiplying P by $\det P^{-1}$ we can assume that $P \in SL(2;\mathbb{C})$. $\square$

In some of the examples we will consider paths of flat connections on knot complements that converge to reductive representations. The following lemma is instructive to understand the geometry of the representation variety along these paths.

Lemma 8.2. Let $\mathcal{P}(t) \subseteq SL(2;\mathbb{C})$, for $t \neq 0$, be the one parameter family of the sets of change of basis matrices that diagonalise a given matrix $M(t)$, where $M(t)$ is a smooth path of matrices. Suppose that at $t = 0$, the Jordan normal form of M is

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

and for all other values of t, M is diagonalisable. Define

$$n(t) = \inf\{\|P\| \mid P \in \mathcal{P}(t)\}.$$

Then $n(t) \to \infty$ as $t \to 0$.

Proof. For the matrix $M(t)$, all matrices that diagonalise $M(t)$ have the form

$$P = [v_1, v_2]$$

where v_1, v_2 are eigenvectors of $M(t)$. Hence we have

$$n(t)^2 = \inf_{v_1, v_2} \{\|v_1\|^2 + \|v_2\|^2\}$$

and the infimum is achieved when the norms of v_1 and v_2 are equal (note that P is constrained to have determinant 1 and so we just minimise over $(v_1', v_2') = (\lambda v_1, \lambda^{-1} v_2)$). When $t \to 0$ the vectors $v_1(t), v_2(t)$ become increasingly colinear, and hence for $\det P(t)$ to equal 1, we must have that the minimum norm of P tends to infinity. $\square$

Definition 8.3. A unipotent element in $SL(2;\mathbb{C})$ is an element with Jordan normal form

$$J = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

8.2 Representation Theory

In this section we recall a few results from representation theory that will be useful when studying our equations.

Lemma 8.4 (Schur's Lemma)**.** Let $\rho : G \to GL(n;\mathbb{C})$ be an irreducible representation of a group G. Then any linear map $M \in GL(n;\mathbb{C})$ that commutes with $\rho(g)$ for all $g \in G$ is a scalar multiple of the identity.

Proof. Consider a non-zero eigenspace, E_λ, of M. Then

$$\rho(g)E_\lambda \subseteq E_\lambda$$

since $M\rho(g)e = \rho(g)Me = \lambda\rho(g)e$ for all $e \in E_\lambda$. Hence $E_\lambda = \mathbb{C}^n$ as the representation is irreducible. □

Corollary 8.5. In an irreducible representation of a Lie group G, any central element will be mapped to a scalar multiple of the identity.

8.3 Flat Connections

In this section we consider the spaces of flat $SL(2;\mathbb{C})$ connections on various manifolds that satisfy

$$d_A^*\phi = 0$$

with respect to some choice of $SU(2)$ reduction.

8.3.1 Example: Flat $SL(2;\mathbb{C})$ connections on T^2 and $T^2\backslash\{*\}$

The first examples we will consider are connections on the torus and connections on the torus minus a point. This is a dimensionally reduced version of the perturbed flat connection on three manifolds.

Flat connections $SL(2;\mathbb{C})$ on T^2 correspond to maps $\pi(T^2) = \mathbb{Z}^2 \to SL(2;C)$, and hence up to conjugation, are given by pairs of commuting matrices. For these to be reductive representations, we require the these matrices to be diagonal. Conjugating by the matrix

$$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

swaps the entries on the diagonals. Elements in $SL(2;\mathbb{C})$ have one of two Jordan normal forms:

$$D_\lambda = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix} \text{ for } \lambda \in \mathbb{C}, \text{ or } \begin{pmatrix} \pm 1 & 1 \\ 0 & \pm 1 \end{pmatrix}.$$

Now, the centraliser of a generic ($\lambda \neq \pm 1$) diagonal matrix, $Z_{SL(2;\mathbb{C})}(D_\lambda)$, is just the diagonal matrices and, and hence isomorphic to $\mathbb{C}^*$. The commutator of the Jordan block is the set of matrices of the form

$$\begin{pmatrix} \pm 1 & a \\ 0 & \pm 1 \end{pmatrix}$$

which is isomorphic to the group $\mathbb{C} \times \{\pm 1\}$. Obviously the centraliser of $\pm I$ is just $SL(2;\mathbb{C})$. By Proposition 3.57, the moduli space of flat connections is given by the diagonal representations of $\mathbb{Z}^2$. Hence, the moduli space of flat $SL(2;\mathbb{C})$ connections may be described as

$$\left(D^2 \setminus \{0\}\right) \times \mathbb{C}^*.$$

Note that here we do not need to consider the representations with the images of generators given by a Jordan block since these representations will not be reductive and so are not in the moduli space that we are considering.

We may want to compactifiy this space by describing the limiting connections as the holonomies go to infinity (or, equivalently due to the embedding of $\mathbb{C}^* \to SL(2;C)$, to 0). This can be done in the sense of [Tau15] by considering normalising the connection one form, that is, considering

$$\phi \mapsto \frac{\phi}{\|\phi\|}.$$

For the torus with the flat metric, it is fairly straightforward to see that these one forms converge to the harmonic forms $d\theta_1$ and $d\theta_2$, where θ_1 and θ_2 are the coordinates on the respective S^1 factors of T^2, and we identify our generators of π_1 with the loops in these directions.

We can also consider connections on $T^2 \setminus N(*)$. In this case we need to impose a connection on the boundary. The condition $d^*\phi = 0$ on the boundary implies that the connection one-form is constant relative to the product connection.

The fundamental group of $T^2 \setminus N(*)$ is the free group on two generators. Hence, flat connections on T^2 correspond to any choice of matrices, A, B (up to conjugacy). To satisfy the boundary condition, we will have that

$$ABA^{-1}B^{-1} = P$$

where P is the choice of matrix defining the boundary conditions. To be able to solve the equation $d_A^*\phi = 0$, we require that A, B are not simultaneously upper triangular with respect to some choice of basis. This analysis gives the following proposition

Proposition 8.6. The moduli space of stable flat $SL(2;\mathbb{C})$ connections on the punctured torus is given by pairs of matrices A, B such that

$$ABA^{-1}B^{-1} = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}$$

for some λ. Solutions to this equation, up to conjugacy, are the family

$$A = \pm \begin{pmatrix} \sqrt{\lambda} & 0 \\ 0 & \sqrt{\lambda}^{-1} \end{pmatrix}, B = \begin{pmatrix} 0 & t \\ -t^{-1} & 0 \end{pmatrix}.$$

Proof. The content of the proposition is that the right hand side has to be a diagonal matrix and cannot have an upper triangular Jordan normal form. Some algebra shows that all the solutions to the equation

$$ABA^{-1}B^{-1} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

are of the form

$$A = \begin{pmatrix} a & \frac{1}{a(1-\lambda^2)} \\ 0 & a \end{pmatrix}, B = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}$$

up to conjugation and hence define a reductive representation. Hence these representations may be discarded, and we are left with the diagonal matrices, and working through the matrix algebra gives the result. □

8.3.2 Flat $SL(2;\mathbb{C})$ Connections on Knot Complements

Before describing solutions to the perturbed Flat connections, where the perturbation is defined over a knot, it is worth considering flat $SL(2;\mathbb{C})$ connections on the knot complement. The solutions to the full perturbed equations will then be given by filling in the knot complement.

By the results of Chapter 4, a reductive representation of the knot group

$$\pi_1(Y \setminus N(K)) \to SL(2;\mathbb{C})$$

may be written as $A + i\phi = 0$, $d_A^*\phi = 0$ for some choice of G reduction, provided the representation of $\pi_1(T^2)$ induced by $\partial N(K) \hookrightarrow Y$ is generic (in the sense that it is diagonalisable). This is implied by a generic choice of perturbation function f.

One way to get an insight into such representations for specific knots is to consider the A polynomial (see Definition 7.20). The A polynomial, $A_K(l,\mu)$ describes all the generic representations of the boundary's fundamental group that extend to representations of the knot group. Hence, the existence of a solution to the system of equations

$$A_K(l,\mu) = 0, \mu = \exp\left(f'\left(\frac{l+l^{-1}}{2}\right)\frac{l-l^{-1}}{2}\right)$$

implies the existence of a flat connection on the knot complement that satisfies the relevant holonomy conditions to be able to be extended to a solution of the perturbed flat equations.

8.4 Example Perturbed Flat $SL(2;\mathbb{C})$ Connections on Three Manifolds

8.4.1 T^3

Example 8.7. Consider $T^3 = S^1 \times T^2$, and consider one of the S^1 fibres. The knot complement of this fibre is $Y_U = S^1 \times \{T^2 \setminus \{*\}\}$. Hence the fundamental group is given by

$$\pi_1(Y_U) = \mathbb{Z} \times F_2 = \langle a, b, c \mid ab = ba, ac = ca\rangle .$$

With this presentation, the fundamental group of the boundary is

$$\pi_1(\partial Y_U) = \langle a, bcb^{-1}c^{-1}\rangle = \mathbb{Z}^2.$$

In the case where a is mapped to a generic element, the solutions to the flat equation on $T^3 \setminus S^1$ reduce to the flat connections on $T^2 \setminus \{*\}$.

8.4.2 Flat Connections on $S^1 \times S^2$

The fundamental group of $S^1 \times S^2$ is $\mathbb{Z}$. Hence each element of $SL(2;C)$ determines a flat connection: for each element of $SL(2;\mathbb{C})$ there exists a unique flat connection (up to $SL(2;\mathbb{C})$ gauge) such that the holonomy of the connection around the S^1 factor is the given element.

We now consider the perturbed flat connections on $S^1 \times S^2$ where we pick the supports of the perturbation to be a disjoint union of the S^1 fibers.

Lemma 8.8. Let $p_1, \ldots, p_k$ be $k > 1$ disjoint points in S^2. The fundamental group of the manifold

$$M_k = S^1 \times \left(S^2 \setminus \{p_1, \ldots, p_k\}\right)$$

is

$$\pi_1(M_k) = \mathbb{Z} \times F_{k-1}.$$

Proof. This follows quickly from the fact that the fundamental group of S^2 with k points removed is F_{k-1}, where F_n is the free group on n generators, and F_0 is the trivial group. $\square$

This allows us to classify the flat $SL(2;\mathbb{C})$ connections on M_k.

Lemma 8.9. There are two classes of flat $SL(2;\mathbb{C})$ connections on M_k. The centre of $\pi_1(M_k)$ is $\mathbb{Z}$. Call the generator of this z. Then for a representation $\rho : \pi_1(M_k) \to SL(2;\mathbb{C})$ either, up to conjugacy,

$$\rho(z) = \pm \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

$\rho(z)$ is a generic diagonal element with trace not 2, or

$$\rho(z) = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

In the first case, there are no constraints on the representation $\rho : F_{k-1} \to SL(2;\mathbb{C})$. In the second case, F_{k-1} must be mapped to a diagonal subgroup, and in the third case F_{k-1} is mapped to unipotent elements.

Proof. This is just a classification of the representations of $\mathbb{Z} \times F_{k-1}$ into $SL(2;\mathbb{C})$. Either the element z gets mapped to a central element ($\pm\,\mathrm{Id}$) or it does not. If it is central then there are no constraints on the representation of F_{k-1}. If it is not central then the image of F_{k-1} must be in the centraliser of the image of z. These are the latter two cases. $\square$

For the perturbed flat connection equation on this manifold, the perturbation introduces an additional constraint on the representations. We require that

$$m_j = \exp f_j(l),$$

where l is the eigenvalue of $\rho(z)$.

In the Abelian case where $\lambda \neq \pm 1$ we have that each $\rho(g_i)$ for each generator g_i is the matrix with the same eigenvectors as $\rho(z)$ but with eigenvalues given by m_j and m_j^{-1}. Thus this moduli space is in one to one correspondence with the choice of $\rho(z)$ up to conjugation, which is just $\mathbb{C}^* \setminus \{\pm 1\}$. For the non-Abelian representations, the moduli space of solutions is given by a choice of $k-1$ elements in $SL(2;\mathbb{C})$ with eigenvectors given by m_j. Up to conjugation we can fix one of these matrices to be diagonal, and so we have a $4(k-2)$ dimensional family of solutions.

8.5 Flat $SL(2;\mathbb{C})$ connections on S^3 with Knot Singularities

Following on from the previous sections, the general procedure for finding the moduli space of solutions to the perturbed, stable, flat connection equations is as follows.

1. Find the knot group and consider the stable (non-reductive) representations of the group.
2. Select the representations that satisfy the holonomy conditions, given by the choice of holomorphic function in the definition of the perturbations.
3. Each such representation admits a solution which can be constructed using the techniques of Chapter 4.

Example 8.10. Consider the unknot in S^3. Let $f : \mathbb{C} \to \mathbb{C}$ be a function. The A polynomial of the unknot is $1 - l$, and hence the moduli space of solutions is given by the solution $\mu = \exp f(1)$.

In the following example we see a family of representations of the knot group of the trefoil that converge to a non-reductive representation.

Example 8.11. Consider the trefoil in S^3. This has A polynomial given by

$$A(\lambda, \mu) = (\lambda - 1)(\lambda + \mu^6).$$

By Lemma 7.9, the knot group of the trefoil is

$$\langle x, y \mid x^2 = y^3 \rangle,$$

and the meridian and longitude of the boundary are represented as $m = xy^{-1}$ and $l = x^2(xy^{-1})^6$, [Coo+94, Proposition 2.7]. We can then define a representation as

$$\rho(x) = \begin{pmatrix} i & t \\ 0 & -i \end{pmatrix}, \quad \rho(y) = \begin{pmatrix} \omega & 0 \\ s & \omega^{-1} \end{pmatrix},$$

where $\omega = e^{\frac{2i\pi}{3}}$. If we conjugate this representation by $\mathrm{diag}(\lambda, \lambda^{-1})$ then $t \mapsto \lambda^2 t$, $s \mapsto \lambda^{-2} s$, and so the quantity $k = st$ is preserved under this operation. One can calculate that the eigenvalues of $\rho(m)$ are

$$\frac{-k - \sqrt{(k - \sqrt{3})^2 - 4} + \sqrt{3}}{2}$$

and its reciprocal. Due to the algebraic closure of $\mathbb{C}$, $\rho(m)$ can take any value and the eigenvalues of the longitude are then $\rho(l) = -\rho(m)^6$. Hence the solutions to the perturbed flat connection equation correspond to the solutions to the function

$$z = \exp(f'(z^6)).$$

These are finite dimensional and discrete, but there are infinitely many solutions by Picard's little theorem.

Chapter 9

Linear Analysis and Deformations

In this chapter we consider the linearisation of the Kapustin-Witten operator in the case of a compact manifolds and of manifolds of the form $\mathbb{R} \times Y$. In both cases we show that the index is 0. In the compact case we can use the index of the well studied operators $d_A^{\pm} + d_A^*$, where as in the tubular manifold case we use the

9.1 The Kapustin-Witten Operator on Compact Manifolds

Let X be a Riemannian 4-manifold, and let $P \to X$ be a principal G–bundle, for a compact group G. Consider the complexified bundle $P_{\mathbb{C}} \to X$ and its adjoint bundle $\mathfrak{g}_{P_{\mathbb{C}}}$. Then for a connection $B = \mathfrak{A}$, where $A = \mathrm{Re}(B)$, the θ–Kapustin-Witten equation is

$$\left(e^{i\theta} F_B\right)^+ = 0, \quad d_A^* \phi = 0.$$

and we consider this with respect to the Kapustin-Witten gauge fixing condition

$$\mathrm{Re}(d_{\mathfrak{A}}^* a + i\varphi) = 0.$$

To look at the linearisation of the equation around B, we can consider the connections $A + ta$ and the one-forms $\phi + t\varphi$.

Definition 9.1. The Kapustin-Witten operator at the connection $\mathfrak{A}$ is the map

$$\mathrm{KW} : \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}}) \to \Omega^2_+(\mathfrak{g}_{P_{\mathbb{C}}}) \oplus \Omega^0(\mathfrak{g}_{P_{\mathbb{C}}})$$

given by

$$\mathrm{KW}(a + i\varphi) = \left(\left(e^{i\theta} F_{\mathfrak{A}}\right)^+, d_{\mathfrak{A}}^*(a + i\varphi) + *[*a, \varphi]\right).$$

Lemma 9.2. The linearisation of the Kapustin-Witten operator is given by

$$L_\theta^{\mathfrak{A}} = e^{i\theta} d_{\mathfrak{A}}^+ \oplus d_{\mathfrak{A}}^*.$$

Note that this is not complex linear.

Proposition 9.3. The linearisation of the Kapustin-Witten operator is elliptic.

Proof. Up to multiplying the first factor of the codomain by $e^{-i\theta}$, this is essentially two copies of the familiar operator $d_A^+ \oplus d_A^*$, which is elliptic. Alternativley we can calculate the principal symbol to give a direct proof. On a trivialisation, the first order component of the operator is just

$$e^{i\theta} d^+ \oplus d^*.$$

If we take an orthonormal frame at a point $p \in X$ and $\square$

Theorem 9.4. The index of the Kapustin-Witten operator on a compact for manifold X is

$$-\chi(X)\dim(\mathfrak{g}).$$

Proof. The index operator is a locally constant map from the space of Fredholm operators to the integers. Hence, on the smooth family of operators

$$\theta \mapsto L_\theta^{\mathfrak{A}}$$

the index is constant. So the index of $L_\theta^{\mathfrak{A}}$ is the same as $L_0^{\mathfrak{A}}$, and so it is sufficient to just calculate the index of $L_0^{\mathfrak{A}}$. But this is just the sum of the indicies of $d^+ \oplus d^*$ and $d^- \oplus +d^*$. This gives the result. $\square$

9.2 Fredholm Theory on $\mathbb{R} \times Y$

In this section we study the linearisation of the Kapustin-Witten equations on tubes. This section is mostly a modification of the results in [Don02] to the case of the Kapustin-Witten equations. We first introduce the notion of adapted bundles and connections. Throughout this section we let Y be an arbitrary three manifold, denote by X the product $\mathbb{R} \times Y$ and use t to denote the $\mathbb{R}$ coordinate. Note that, since the Hessian of the perturbations defined in Chapter 5 is compact, for the results in this section we can restrict our attention to the unperturbed θ_v–Kapustin-Witten equations as Fredholm theory of an operator is unchanged if it is modified by a compact operator.

Definition 9.5. An adapted complexified bundle is a bundle $P_{\mathbb{C}} \to X$ with a choice of two flat $SL(2,\mathbb{C})$ connections $\mathfrak{A}_{\pm}$, which satisfy

$$d^*_{A_{\pm}}\phi_{\pm} = 0.$$

The connections $A_{\pm} + i\phi_{\pm}$ are called the limiting connections of the bundle. An adapted connection on $P_{\mathbb{C}}$ is a connection $\mathfrak{A}$ such that

$$\lim_{t\to\pm\infty} \mathfrak{A} = \mathfrak{A}_{\pm}.$$

Two adapted bundles are said to be equivalent if there exists a real bundle isomorphism between them that preserves adapted connections.

Lemma 9.6. Pick two limiting connections $\mathfrak{A}_{\pm}$. There is a one to one correspondence between the equivalence classes of adapted bundles with limiting connections $A_{\pm} + i\phi_{\pm}$ and the integers.

Proof. Given two such bundles we, $P_{\mathbb{C}}$ and $P'_{\mathbb{C}}$ we can apply a constant gauge transformation to $P_{\mathbb{C}}$ such that

$$\mathfrak{A}_{-} = A'_{-} + i\phi'_{-}.$$

Now pick two adapted connections, $\mathfrak{A}$, $A' + i\phi'$ and define

$$I' = \int_X \mathrm{tr}(F^2_{\mathfrak{A}}) - \mathrm{tr}(F^2_{A'+i\phi'}) = CS(\mathfrak{A}_{+}) - CS(A'_{+} + i\phi'_{+}) \in \mathbb{Z}.$$

This quantity is integer valued as it is the ambiguity in the Chern-Simons functional of the gauge equivalent limiting connections. If this quantity is 0 then the two connections $\mathfrak{A}_{+}$ and $A'_{+} + i\phi'_{+}$ are gauge equivalent by a gauge transformation $g \in \mathcal{G}_0$, where $\mathcal{G}_0$ is the connected component of the gauge group containing the identity. Hence by interpolating between this gauge transformation and the identity we may construct a bundle isomorphism that maps the flat limits. Moreover, if we have a third adapted bundle $P''_{\mathbb{C}}$, with $I' = I''$, then we have

$$CS(A'_{+} + i\phi'_{+}) = CS(A''_{+} + i\phi''_{+}),$$

and so $P'_{\mathbb{C}} \cong P''_{\mathbb{C}}$. □

9.2.1 Linearised Equations

For an adapted connection $\mathfrak{A}$ on $I \times Y$, we want to consider the linearisation of the θ_v–Kapustin-Witten operator.

$$D_{\mathfrak{A}} : \Omega^1 \oplus \Omega^1 \to \Omega^0 \oplus \Omega^+ \oplus \Omega^0 \oplus \Omega^-.$$

Our aim is to show that this operator is Fredholm. Since the connection is adapted we can split the real line into three sections, $(-\infty, -T), [-T, T]$ and (T, ∞) and consider the operator there. On the ends, for large T, the operator is constant and so it makes sense to first study connections that are constant.

To begin then we consider a connection that is constant in t. Pick a connection $\mathfrak{A}$ on a complexified bundle $P_{\mathbb{C}} \to Y$. Consider the operator

$$D = \frac{d}{dt} + L_{\mathfrak{A}}$$

where $L_{\mathfrak{A}} : \Omega^0(\mathfrak{g}_{P_{\mathbb{C}}}) \oplus \Omega^1(\mathfrak{g}_P) \oplus \Omega^1(\mathfrak{g}_P) \to \Omega^0(\mathfrak{g}_{P_{\mathbb{C}}}) \oplus \Omega^1(\mathfrak{g}_P) \oplus \Omega^1(\mathfrak{g}_P)$ given by the matrix

$$\begin{aligned} L_{\mathfrak{A}} &= \begin{pmatrix} 0 & e^{2i\theta_v} d^*_{\mathfrak{A}} \\ e^{-2i\theta_v} d_{\mathfrak{A}} & \bar{*} e^{i2\theta_v} d_{\mathfrak{A}} \end{pmatrix} \\ &= \begin{pmatrix} 0 & e^{2i\theta_v} d^*_{\mathfrak{A}} & i e^{2i\theta_v} d^*_{\mathfrak{A}} \\ \operatorname{Re}(e^{-2i\theta_v} d_{\mathfrak{A}}) & \cos 2\theta_v * d_A - \sin 2\theta_v * [\phi, \cdot] & -\sin 2\theta_v * d_A - \cos 2\theta_v * [\phi, \cdot] \\ -\operatorname{Im}(e^{-2i\theta_v} d_{\mathfrak{A}}) & -\sin 2\theta_v * d_A - \cos 2\theta_v * [\phi, \cdot] & \sin 2\theta_v * [\phi, \cdot] - \cos 2\theta_v * d_A \end{pmatrix}. \end{aligned}$$

Here in the first line we have made the identification $\Omega^1(\mathfrak{g}_P) \oplus \Omega^1(\mathfrak{g}_P) \cong \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}})$. Note that this matrix is not $\mathbb{C}$ linear due to the $*$ operator in the bottom right corner.

Remark 9.7. The bottom right corner of $L_{\mathfrak{A}}$ can be written in the slightly nicer form

$$\cos 2\theta_v \begin{pmatrix} *d_A & - * [\phi, \cdot] \\ - * [\phi, \cdot] & - * d_A \end{pmatrix} + \sin 2\theta_v \begin{pmatrix} *[\phi, \cdot] & - * d_A \\ - * d_A & *[\phi, \cdot] \end{pmatrix}.$$

Written in this way, it is clear that $L_{\mathfrak{A}}$ is self adjoint.

Remark 9.8. Note that when $\theta_v = 0$, $L_{\mathfrak{A}}$ is also the linearisation of the curvature operator from Lemma 3.63. This operator is set up to be very similar to the operator $\Omega^0 \oplus \Omega^1 \to \Omega^0 \oplus \Omega^1$,

$$\begin{pmatrix} 0 & d_A^* \\ d_A & *d_A \end{pmatrix},$$

studied in Floer theory. As such, many of the proofs in this section follow those given in [Don02, Chapter 3].

Throughout this section we assume that the function θ_v depends only on the t coordinate.

Lemma 9.9. For any connection $\mathfrak{A}$ on $P_{\mathbb{C}} \to Y$ the operator $L_{\mathfrak{A}}$ is self adjoint. Moreover, if $\mathfrak{A}$ is an acyclic flat connection then the operator $L_{\mathfrak{A}}^{\theta_v}$ is invertible.

Proof. From Remark 9.7 it is clear that $L_{\mathfrak{A}}^{\theta_v}$ is self adjoint. Hence to prove the moreover statement it is sufficient to prove that the kernel of $L_{\mathfrak{A}}^{\theta_v}$ is empty.

Suppose that $L_{\mathfrak{A}}(a, b) = 0$, for $a \in \Omega^0(\mathfrak{g}_{P_{\mathbb{C}}})$ and $b \in \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}})$. Then we have that

$$\bar{*}e^{2i\theta_v} d_{\mathfrak{A}} b = e^{-2i\theta} d_{\mathfrak{A}} a$$

Considering $\|d_{\mathfrak{A}} a\|_{L^2}^2$ we see that $d_{\mathfrak{A}} a = 0$ and so $a = 0$ since $\mathfrak{A}$ is acyclic. Hence $d_{A+i\phi} b = d_{A+i\phi}^* b = 0$ and so $b = 0$ as $A + i\phi$ is non-degenerate. □

Definition 9.10. We denote by $L_{\mathfrak{A}}^{\theta}$ the operator $L_{\mathfrak{A}}$ with $\theta_v = \theta$.

The spectral theories of $L_{\mathfrak{A}}^{\theta}$ for varying θ are very closely related.

Lemma 9.11. Suppose that $(a, b) \in \Omega^0(\mathfrak{g}_{P_{\mathbb{C}}}) \times \Omega^1(\mathfrak{g}_{P_{\mathbb{C}}})$ is an eigenvector of $L_{\mathfrak{A}}^0$ with eigenvalue λ. Then

$$(e^{i\theta} a, e^{-i\theta} b) \text{ and } (ie^{i\theta} a, -ie^{-i\theta} b)$$

are eigenvectors of $L_{\mathfrak{A}}^{\theta}$ with eigenvalues λ and $-\lambda$ respectively.

Proof. This is a direct calculation. We have that

$$L_{\mathfrak{A}}^{\theta} \begin{pmatrix} e^{i\theta} a \\ e^{-i\theta} b \end{pmatrix} = \begin{pmatrix} e^{i2\theta} d_{\mathfrak{A}} e^{-i\theta} b \\ e^{-i2\theta} d_{\mathfrak{A}} e^{i\theta} a + \bar{*} e^{i2\theta} d_{\mathfrak{A}} e^{-i\theta} b \end{pmatrix} = \lambda \begin{pmatrix} e^{i\theta} a \\ e^{-i\theta} b \end{pmatrix}$$

where $L_{\mathfrak{A}}^0(a, b) = \lambda(a, b)$. The second calculation is similar. □

Corollary 9.12. The operator $D_{\mathfrak{A}}$ has index 0

Proof. We prove that the operator $D_{\mathfrak{A}}$ is Fredholm in Theorem 9.18. We can define the index via spectral flow. Doing so, we see that every upward crossing of an eigenvalue λ is matched by a downwards crossing coming from $-\lambda$. □

In the following we use the following definition of the Sobolev norm L^2_1

$$\|\rho\|_{L^2_1} = \|D\rho\|_{L^2} + \|\rho\|_{L^2}$$

for smooth compactly supported sections sections, and then define the space L^2_1 to be the completion of these sections with respect to this norm.

Proposition 9.13. Define δ to be

$$\delta = \min_{\lambda \in \mathrm{Spec}(L)} |\lambda|.$$

Then the operator

$$D : L^2_1 \to L^2$$

has an inverse $Q : L^2 \to L^2_1$ such that for all $\rho \in L^2$,

$$\|Q\rho\|_{L^2_1} \leq \delta^{-1} \|\rho\|_{L^2}.$$

Proof. See [Don02, Proposition 3.4]. □

Lemma 9.14. Let ρ be a compactly supported smooth section of the bundle. Then there exists an f such that

$$Df = \rho,$$

and $\|f\|_{L^2} \leq \delta^{-1} \|\rho\|_{L^2}$

Proof. In the case of varying θ_v can we write $f = \sum f_\lambda e_\lambda$ and $\rho = \sum \rho_\lambda e_\lambda$, where

$$e_\lambda = (e^{i\theta_v} a, e^{-i\theta} b)$$

for an eigenvector (a, b) of eigenvalue λ, then the equation $Df = \rho$ becomes

$$\sum f'_\lambda + \lambda f_\lambda e_\lambda + \theta'_v f_\lambda e_{-\lambda} = \sum \rho_\lambda e_\lambda.$$

where $e_{-\lambda} = (ie^{i\theta_v} a, -ie^{-i\theta} b)$. Hence to solve for f_λ, for each λ we need to consider the system

$$\frac{d}{dt} \begin{pmatrix} f_\lambda \\ f_{-\lambda} \end{pmatrix} + \begin{pmatrix} \lambda & \theta'_v \\ \theta'_v & -\lambda \end{pmatrix} \begin{pmatrix} f_\lambda \\ f_{-\lambda} \end{pmatrix} = \begin{pmatrix} \rho_\lambda \\ \rho_{-\lambda} \end{pmatrix}. \tag{9.1}$$

We are assuming that $\theta_v \to \alpha_\pm$ as $t \to \pm\infty$ and that $\theta'_v = 0$ for $|t| > T$. Hence for large $|t|$ we have that the system of ODEs is given by

$$\frac{d}{dt}\begin{pmatrix} f_\lambda \\ f_{-\lambda} \end{pmatrix} + \begin{pmatrix} \lambda & 0 \\ 0 & -\lambda \end{pmatrix}\begin{pmatrix} f_\lambda \\ f_{-\lambda} \end{pmatrix} = \begin{pmatrix} \rho_\lambda \\ \rho_{-\lambda} \end{pmatrix}.$$

We can solve this system explicitly and get a unique bounded solution on $(-\infty, -T]$ and another unique bounded solution on $[T, \infty)$. Call these solutions u_λ and $u_{-\lambda}$ respectivley. Now by the theory of first order differential equations, we can find two unique solutions to (9.1) with $\rho_{\pm\lambda} = 0$ on the whole of $\mathbb{R}$ which are equal to $u_{\pm\lambda}$ when restricted to the relevant subsets of $\mathbb{R}$. We call these solutions $v_{\pm\lambda}$. So $v_\lambda = u_\lambda$ for large positive t and $v_{-\lambda} = u_{-\lambda}$ for large negagive T.

Now consider the matrix

$$M = \begin{pmatrix} u_\lambda & u_{-\lambda} \end{pmatrix}$$

with columns given by the $u_{\pm\lambda}$. This matrix is such that

$$\frac{d}{dt}M + \begin{pmatrix} \lambda & \theta'_v \\ \theta'_v & -\lambda \end{pmatrix} M = 0,$$

and hence if

$$v = M\left(\int_{-\infty}^{t} M^{-1}\begin{pmatrix} \rho_\lambda \\ \rho_{-\lambda} \end{pmatrix} + C\right),$$

then v solves (9.1). Now u_λ decays as $t \to -\infty$ and $u_{-\lambda}$ decays as $t \to \infty$. Using this, we can pick C such that v is bounded. Let

$$\begin{pmatrix} a \\ b \end{pmatrix} = \int_{-\infty}^{\infty} M^{-1}\begin{pmatrix} \rho_\lambda \\ \rho_{-\lambda} \end{pmatrix}.$$

Then if

$$C = \begin{pmatrix} -a \\ 0 \end{pmatrix},$$

$v = -au_\lambda$ for $t < -T$ and $v = bu_{-\lambda}$ for $t > T$. Hence v is bounded.

We now consider the L^2_1 norm of v. We have at each t

$$\left\|\frac{dv}{dt}\right\|^2 + \left(\lambda^2 + {\theta'_v}^2\right)\|v\|^2 = \|\rho_{\pm\lambda}\|^2 - \sqrt{\lambda^2 + {\theta'_v}^2}\,\frac{d}{dt}\langle v, v\rangle.$$

Hence we can bound the L^2 norm of v by the L^2 norm of $\rho_\pm$. Putting all this together for gives the result.

□

We now restrict to the case where θ_v is constant and prove some results about the operator D, namely that it is elliptic. It is straightforward to extend these results to θ_v where $\theta_v \to \theta_\pm$ as $t \to \infty$ since at the ends of $\mathbb{R}$ we can take θ to be constant and then in the middle compact region we can use standard elliptic theory and the patching arguments in the proof of Theorem 9.18. For constant θ these results can be proved using the eigenfunction expansions for the operator L, and proofs can be found in [Don02, Section 3.1]. In the following we define the norm

Lemma 9.15. Suppose that $Df = 0$ on $(-T, T) \times Y$. There is a constant C_δ, independent of T, such that on $(-T, T) \times Y$,

$$\int_{-T}^{T} \|f\|_{L^2} \leq C_\delta \left(\int_{-T}^{-T+1} \|f\|_{L^2} + \int_{T-1}^{T} \|f\|_{L^2} \right)$$

Corollary 9.16. Suppose that $Df = 0$ on $(0, \infty) \times Y$ and $f \in L^2$. Then

$$\int_0^\infty \|f\|_{L^2} \leq C_\delta \int_0^1 \|f\|_{L^2}.$$

Proof. The proof of Lemma 9.15 is unchanged if we relabel the interval to $(0, 2T)$. In doing so we get

$$\int_0^{2T} \|f\|_{L^2} \leq C_\delta \left(\int_0^1 \|f\|_{L^2} + \int_{2T-1}^{2T} \|f\|_{L^2} \right).$$

Since f is in L^2, the second summand tends to 0 as $T \to \infty$ to give the result. □

Corollary 9.17. Suppose that $\rho \in L^2(\mathbb{R} \times Y)$. Then there exists a solution

$$Df = \rho$$

on $\mathbb{R} \times Y$ and $\|f\|_{L^2} \leq \delta^{-1} \|\rho\|$.

Proof. Let ρ_i be a sequence of smooth, compactly supported sections converging to ρ in L^2 and suppose that the support of ρ_i is contained in ρ_j for $i < j$, and let f_i be the associated solutions extended to the whole interval by 0. Let f_{ij} be the solution to $Df = \rho_i - \rho_j$ again extended to the whole interval by 0. Now consider $g_{ij} = f_i - f_j - f_{ij}$. The L^2 norm of g_{ij} is uniformly bounded by a multiple of $\|\rho\|_{L^2}$. Moreover, the $g_{ij} \to 0$ on compact intervals by Lemma 9.15. It follows from the dominated convergence theorem that $g_{ij} \to 0$, and hence the sequence f_i is Cauchy. $\square$

Theorem 9.18. Let $\mathfrak{A}$ be an adapted connection on an adapted complexified principal bundle $P_{\mathbb{C}} \to Y$ such that on each end

$$\|\mathfrak{A} - (A_{\pm} + i\phi_{\pm})\|_{L^2_3(I\times Y)} < \infty.$$

Here $I_{-} = (-\infty, 0]$ and $I^{+} = [0, \infty)$. Then the operator

$$D : L^2_1 \to L^2$$

is Fredholm. That is, it has finite dimensional kernel and cokernel.

Proof. We follow [Don02, Proposition 3.6], adding in some details. As outlined above, we split the real numbers into three sections $(-\infty, -T), [-T, T]$, and (T, ∞). Define

$$\delta_{\pm} = \mathfrak{A} - (\mathfrak{A}_{\pm}).$$

Then by the Sobolev Embedding theorem, we have that for compact intervals I, there is a constant depending only on $|I|$ such that

$$\|\delta_{\pm}\|_{C^{0,\alpha}(I\times Y)} \leq C\|\delta_{\pm}\|_{L^2_3(I\times Y)}.$$

It follows that $\delta_{\pm} \in C^{0,\alpha}(I_{\pm} \times Y)$. Now suppose that we have a solution to

$$D_{\mathfrak{A}} f = 0.$$

To relate this to the limiting connections of $P_{\mathbb{C}}$ we can write

$$D_{\mathfrak{A}} f = D_{\mathfrak{A}_+} f + [\delta_+, f].$$

Hence f is a solution to

$$D_{\mathfrak{A}_+} f = -[\delta_+, f].$$

Let g be the solution to

$$D_{\mathfrak{A}_+} g = -[\delta_+, \chi f].$$

given by Corollary 9.17. Then $D_{\mathfrak{A}_+}(f - g) = 0$. Hence we have, using Corollary 9.16,

$$\int_T^\infty \|f\|_{L^2} \leq \int_T^\infty \|f-g\|_{L^2} + \|g\|_{L^2} \leq C_\delta \int_T^{T+1} \|f-g\|_{L^2} + \int_T^\infty \|[\delta_+, f]\|_{L^2}.$$

To prove that the kernel is finite dimensional we can prove sequential compactness of the unit sphere. Let f_i be a sequence of solutions to $D_{\mathfrak{A}} f_i = 0$ with $\|f_i\|_{L^2} = 1$. The space L^2_1 is weakly sequentially compact and so we may assume that $f_i \rightharpoonup f_\infty$ in L^2_1. Moreover, this is a weak solution to $D_{\mathfrak{A}} f_\infty = 0$. Now fix an interval $[-T, T]$. Then on the compact manifold $[-T-\varepsilon, T+\varepsilon] \times Y$, where $\varepsilon > 0$, we have the standard elliptic estimate

$$\begin{aligned} \|f_i - f_\infty\|_{L^2_1([-T,T]\times Y)} &\leq C_\varepsilon \left(\|D(f_i - f_\infty)\|_{L^2([-T-\varepsilon,T+\varepsilon]\times Y)} + \|f_i - f_\infty\|\right) \\ &\leq C_\varepsilon \|f_i - f_\infty\|_{L^2([-T-\varepsilon,T+\varepsilon]\times Y)}, \end{aligned}$$

where we can pick C_ε to be independent of T. The embedding $L^2_1 \to L^2$ is compact over compact manifolds, and hence we may extract a L^2 convergent subsequence, f_i^T which will be convergent in L^2_1 by the elliptic estimate. Now let $T_j = j$, and for each T_j pick a sequence f_i^j as above, starting with the sequence f_i^{j-1}. Taking a further subsequence we can assume that

$$\|f_i^j - f_\infty\|_{L^2([-j,j]\times Y)} \leq 2^{-(i+j)}.$$

Take the diagonal sequence $h_i = f_i^i$, and let g_i solve $D_{\mathfrak{A}_+} g_i = -[\delta_+, h_i - f_\infty]$. Then we have

$$\begin{aligned} \|h_i - f_\infty\|_{L^2((0,\infty)\times Y)} &\leq \int_0^i \|h_i - f_\infty\|_{L^2} + \int_i^\infty \|h_i - f_\infty\|_{L^2} \\ &\leq 2^{-2i} + C_\delta \int_{i-1}^i \|h_i - f_\infty - g\|_{L^2} + \int_i^\infty \|[\delta_+, f_i - f_\infty]\| \\ &\leq 2^{-2i} + C_\delta 2^{-2i} + \|\delta_+\|_{C^0([i,\infty)\times Y)} \left(\|f_i\|_{L^2} + \|f_\infty\|_{L^2}\right). \end{aligned}$$

This tends to 0 as $i \to \infty$ due to the decay of δ_+. We can do the same analysis with the negative half interval. This proves that the kernel is finite dimensional.

For the cokernel, we first prove that the image of $D_{\mathfrak{A}}$ is closed. To do this, note that we can invert the operator $D_{\mathfrak{A}}$ on $(T, \infty) \times Y$ for large T. To do this we use the inverse Q of $D_{\mathfrak{A}_+}$. Then to solve

$$D_{\mathfrak{A}} f = \rho$$

we need to solve

$$D_{\mathfrak{A}_+} f + [\delta_+, f] = \rho.$$

To do this we can form an iterative sequence, where we let $f_1 = Q\rho$ and then let

$$f_{i+1} = Q(\rho - [\delta_+, f_i]).$$

Then we have that

$$\|f_{i+1} - f_i\|_{L^2_1} = \|Q([\delta_+, f_i] - [\delta_+, f_{i-1}])\|_{L^2_1} = \|Q([\delta_+, f_i - f_{i-1}])\|_{L^2_1} \leq \delta^{-1}\|\delta_+\|_{C^0}\|f_i - f_{i-1}\|_{L^2}.$$

Hence, provided $\delta^{-1}\|\delta_+\|_{C^0} \leq 1$, the sequence f_i will be Cauchy in L^2_1 and hence converges to some f. This f solves

$$D_{\mathfrak{A}} f = \rho.$$

Moreover, we have the bound

$$\|f\|_{L^2_1} \leq C\|\rho\|_{L^2}.$$

Hence, if we let $\hat{Q}\rho = f$ then we have a bounded inverse to $D_{\mathfrak{A}}$ on $(T, \infty) \times Y$, where T is chosen to make $\|\delta_+\|$ small. We can do the same for the negative interval. Now we want to construct a right inverse on the remaining compact part $[-T, T] \times Y$. On small geodesic balls we can always solve

$$D_{\mathfrak{A}} f = \rho,$$

since if $s_r : B(0,1) \to B(0,r)$ is the map $x \to rx$ in $\mathbb{R}^4$ restricted to the unit ball, the pullback of a connection one form a by s_r is given by

$$s_r^* a = ra(ra).$$

Hence for smaller and smaller r, the pullback of the operator $D_{\mathfrak{A}}$ converges to $d^* + d^+$. This operator is invertible on $B(0,1)$ with the trivial Dirichlet boundary condition. Hence, we can cover the compact region with finitely many balls over which the operator is invertible. Call these U_i, and let the inverse operator be Q_i. Suppose that $B_i(0, \frac{1}{2}r_i)$ covers the compact region. Now let β_i be a partition of unity subordinate to the U_i and let

$$P = \sum \beta_i Q_i.$$

Note that on U_1

$$P(\rho) = Q_1(\rho) + \sum_{i=2}^{\infty} \beta_i (Q_i - Q_1)(\rho) =$$

Hence we have that

$$DP\rho = \rho + D\left(\sum \beta_i f_i\right).$$

where $f_i = (Q_i - Q_1)\rho$. Note that $Df_i = 0$. From elliptic regularity, we have that

$$\|Df_i\|_{L^2_k} \leq C_k \|f_i\|_{L^2}$$

and so the f_i are smooth. Hence, if we let

$$S(\rho) = DP\rho - \rho = D\left(\sum \beta_i f_i\right),$$

then we can bound all the derivatives of $S(\rho)$, given sufficient regularity of $\mathfrak{A}$. It follows from the Arzela-Ascoli theorem that this operator is compact. Now we show that the image of $D_{\mathfrak{A}}$ is closed. First note that the image of $D_{\mathfrak{A}}$ differs from the image of DP by a finite dimensional subspace. This follows from the compactness of S: if $\rho = Df$, then

$$Df = DP\rho - S(\rho)$$

and the image of S is finite dimensional. Hence, if $\mathrm{Im}DP$ is closed, then so too is the image of $D_{\mathfrak{A}}$. But we have that

$$DP = I + S$$

where S is compact. Suppose that $x = \lim (I+S)(f_i)$, for some $f_i \in \ker D^{\perp}$. If the sequence f_i is bounded then we have that along a subsequence $S(f_i) \to z$. Now we have that along this subsequence

$$f_i = f_i + S(f_i) - S(f_i) \to x - z,$$

and so $(I+S)(x-z) = \lim(I+S)(f_i) = x$. If the f_i are unbounded, then we can consider $\hat{f}_i = f_i / \|f_i\|_{L^2}$ after passing to a subsequence with $\|f_i\|_{L^2} \to \infty$. Then $(I+S)(\hat{f}_i) \to 0$. On the other hand, $S(\hat{f}_i) \to y$ along some subsequence. Since

$$0 = \lim \hat{f}_i + S(f_i)$$

we see that $\|S(y)\| = \|\hat{f}_i\| = 1$. But also, taking the limit we see that $(I+S)(y) = 0$. Hence $y \in \ker(I+S)$. But y is the limit of the f_i which all lie in the complement of $\ker(I+S)$, so this is a contradiction. Hence the range of D is closed.

Finally we prove that the dimension of the cokernel is finite. Since $\mathrm{Im} D$ is closed, we have that

$$L^2 = \mathrm{Im} D \oplus \mathrm{Im} D^{\perp},$$

and we have an isomorphism

$$\mathrm{Im} D^{\perp} \cong L^2/\mathrm{Im\ D}.$$

However, if we have that $g \in \mathrm{Im} D^{\perp}$ then

$$\langle D^* g, f \rangle = 0$$

for all $f \in L^2$ and so $g \in \ker D^*$. Also, if $g \in \ker D^*$ then $g \in \mathrm{Im}\ D^{\perp}$. Hence we have that D is Fredholm if and only if $\ker D^*$ is finite dimensional. But this follows from the fact that D is self adjoint and D has finite dimensional kernel. $\square$

Chapter 10

Solutions to the Kapustin-Witten equations on Tubes

In this chapter we consider compactness of perturbed flat $SL(2;\mathbb{C})$ connections on three manifolds and apply this result to consider compactness problems for the moduli space of solutions to the perturbed Kapustin-Witten equations.

We make a conjecture, Conjecture 10.7, on the behaviour of solutions to perturbed the Kapustin-Witten equations in order to rule out the situation where $\|\phi\|_{L^2} \to \infty$.

We start by considering the compactness result by Taubes from [Tau15] which is a generalisation of Uhlenbeck's theorem to the case of $PSL(2;\mathbb{C})$ and $SL(2,\mathbb{C})$ connections on three-manifolds.

From now on, in light of Remark 3.15, we make the assumption that $\phi(\partial_t) = 0$. One can also justify this assumption for solutions with constant θ (as opposed to the θ_v–Kapustin-Witten equations) via a maximum principle argument on the norm of $\phi(\partial_t)$, since we want to define our moduli spaces to be solutions that interpolate between flat connections. See, [He19, Corollary 4.7].

Definition 10.1. In this chapter we consider the gradient flow of the perturbed Chern-Simons functional, defined in Chapter 5. We suppose that we have chosen a regular coercive perturbation, η, (see Definitions 6.14 and 6.15), and we consider the gradient flow of

$$\mathfrak{C}_\eta = \mathfrak{C} + \eta.$$

We write

$$\eta'(\mathfrak{A}) = \nabla\eta(\mathfrak{A}).$$

10.1 Taubes' Compactness Theorems

In [Tau15], Taubes considers sequences of connections $\mathfrak{A}_n = A_n + i\phi_n$ on a $PSL(2;\mathbb{C})$ bundle $P_{\mathbb{C}} \to Y$ such that there exists a constant M and a bound

$$\|F_{\mathfrak{A}_n}\|_{L^2}^2 + \|d_{A_n}^* \phi_n\|_{L^2}^2 \leq M. \tag{10.1}$$

For any sequence of connections satisfying a bound (10.1), there is a dichotomy as to whether this sequence contains a convergent subsequence depending on the value of $\liminf_{n\to\infty} \|\phi_n\|_{L^2}$.

Theorem 10.2 ([Tau15, Theorem 1.1])**.** *Let $\mathfrak{A}_n = (A_n + i\phi_n)$ be a sequence of connections on a $SL(2;\mathbb{C})$ or $PSL(2;\mathbb{C})$ bundle $P \to Y$ over a three manifold, satisfying the bound (10.1). Then either $\|\phi_n\|_{L^2}$ contains a bounded subsequence or it does not. If $\|\phi_n\|_{L^2}$ contains a bounded sequence then there is an L_1^2 connection $A + i\phi$ satisfying (10.1) and a subsequence of $\mathfrak{A}_n$ that converges to a limit, $\mathfrak{A}$, weakly in L_1^2 up to L_2^2 gauge transformations. If, on the other hand,*

$$\liminf_{n\to\infty} \|\phi\|_{L^2} = \infty,$$

then there exists a subset Z of Hausdorff dimension at most 1, an $L^2_{1,\mathrm{loc}}$ connection A on $P|_{Y\setminus Z}$, a real line bundle $l \to Y \setminus Z$ equipped with a harmonic l valued one form, v, and a covariantly constant section of $l^* \otimes \mathfrak{g}_P$, σ, such that along a subsequence

1. *there exists gauge transformations such that $g_n^* A_n \rightharpoonup_{L_1^2} A$ on $P_{Y\setminus Z}$,*
2. *The sequence $\|\phi_n\|^{-1} g_n^* \phi_n \rightharpoonup_{L_1^2} v\sigma$*
3. *$|v| : Y \setminus Z \to \mathbb{R}$ extends to a Hölder continuous function in L_1^2 on all of Y with $Z = |v|^{-1}(0)$*

The proof of the first part of the theorem follows quickly from the Weitzenböck formula (3.3) which gives a bound on L^2 bound on F_A. Uhlenbeck's theorem then implies the existence of a limiting connection up to gauge transformations, and the connections converge weakly in L_1^2 after applying the gauge transformations. The Weitzenböck formula also gives a bound on

$$\|\nabla_{A_n} \phi_n\|_{L^2} \leq C.$$

The result then follows from weak sequential compactness of L_1^2 and the following lemma.

Remark 10.3. Note that if our perturbation function were bounded (for example, if one only perturbed using the holonomy of the connection A, or used bounded functions to define the perturbation) then the perturbations would not change the compactness result of Taubes. Since our perturbations are unbounded on the space of connections, the compactness theorem of Taubes does not directly apply.

Lemma 10.4. Suppose $\|\nabla_{A_n}\phi_n\| \leq C$ and $A_n \rightharpoonup A$ in L^2_1, then there exists a smooth connection B such that

$$\|\nabla_B \phi_n\|_{L^2_1} \leq C.$$

Proof. This follows from the equivalence of the L^2_1 Sobolev norms defined by L^2_1 connections on three manifolds. Suppose that $B = A + a$ for a one form $a \in L^2_{1,B}(\Omega^1(\mathfrak{g}_P))$. Then we have

$$\begin{aligned}\|\nabla_A \phi\|_{L^2} &\leq \|\nabla_B \phi\|_{L^2} + \|[a,\phi]\|_{L^2} \\ &\leq \|\nabla_B \phi\|_{L^2} + \|a\|_{L^2_{1,B}} \|\phi\|_{L^2_{1,B}},\end{aligned}$$

using the Sobolev embedding $L^2_1 \hookrightarrow L^4$. Hence we have

$$\|\phi\|_{L^2_{1,A}} \leq C\|\phi\|_{L^2_{1,B}}$$

and since the argument is symmetrical with A, B replaced the norms are equivalent. □

We now state another result of Taubes that we will use. This is an Uhlenbeck style compactness theorem for solutions to the θ–Kapustin-Witten equations on $I \times M$, where M is compact three manifold. In particular, M can have boundary. This follows from the fact that the following theorem is proven using local estimates on balls in the manifold.

Theorem 10.5. [Tau13, Theorem 1.2, first bullet point] Let $I \subseteq \mathbb{R}$ denote a closed, bounded interval of positive length and let Y be a compact, oriented Riemannian three-manifold. Let $P_{\mathbb{C}} \to I \times Y$ be a principal $SL(2;\mathbb{C})$ bundle, and suppose that, with respect to an $SU(2)$ reduction,

$$\mathfrak{A}_n = A_n + i\phi_n$$

is a sequence of solutions to the θ_n Kapustin-Witten equations of finite energy. Suppose in addition that $\|\phi\|_{L^2}$ is bounded. Then there exists a finite set $\Theta \subseteq$

$I \times Y$, gauge transformations, g_n, of $P|_{I\times Y\setminus\Theta}$ and a connection $\mathfrak{A}_0$ on $P_{\mathbb{C}}$ such that

$$g_n^*(\mathfrak{A}_n) \to \mathfrak{A}_0$$

on a subsequence in C^∞.

10.2 Perturbed Solutions with Bounded L^2 Energy

We now consider solutions to the Kapustin-Witten equations with bounded energy. The key ideas we can use here are that bounded energy implies that $\|\phi\|_{L^2(Y)}$ of the limiting connections is bounded. By analogy to Floer theory, [Don02, Section 4], we want to study solutions to the θ_v–Kapustin-Witten equations on $\mathbb{R}\times Y$ such that the total perturbed curvature

$$\|F_{\mathfrak{A}} + \eta'(\mathfrak{A})\|_{L^2} < \infty. \tag{10.2}$$

Since we have that

$$\left(e^{i\theta}\left(F_{\mathfrak{A}} + \eta'(\mathfrak{A})\right)\right)^+ = 0$$

the condition (10.2) is equivalent to

$$\left\|\left(e^{i\theta}\left(F_{\mathfrak{A}} + \eta'(\mathfrak{A})\right)\right)^-\right\|_{L^2} < \infty \tag{10.3}$$

A key identity for solutions to the gradient flow is that

$$\Delta(\mathfrak{C}+\eta)(\mathfrak{A}) = \int_{-\infty}^{\infty} \|F_{\mathfrak{A}} + \eta'(\mathfrak{A})\|^2_{L^2}. \tag{10.4}$$

Hence if the curvature is bounded then the change in the perturbed Chern-Simons functional is also bounded.

Lemma 10.6. Let η be a regular, coercive perturbation and suppose $\mathfrak{A} \in L^2_{1,\mathrm{loc}}$ is a finite energy solution to the η–perturbed Kapustin-Witten equations. If $\mathfrak{A}$ converges to a perturbed flat connection as $t \to \infty$, then the holonomy of the limiting perturbed flat connection is bounded.

Proof. Since the connection has finite energy and is in $L^2_{1,\mathrm{loc}}$, it follows that at time t, $\mathfrak{C}_\eta(\mathfrak{A}(t))$ is bounded, and so $\mathfrak{C}_\eta(\mathfrak{A}(t))$ is bounded for all t. Since η is a coercive perturbation and $\mathfrak{A}$ converges to a critical point, it follows that the critical point has bounded holonomy (otherwise $\mathfrak{C}_\eta$ would tend to infinity along the curve). $\square$

Conjecture 10.7. Let $\mathfrak{A}$ be bounded energy a solution to the perturbed Kapustin-Witten equations on $\mathbb{R}\times Y$. Then $t\mapsto \|\phi(t)\|_{L^2(Y)}$ is bounded.

We now explain why we make this conjecture. If we have that $\mathfrak{C}_\eta(t)$ is bounded at some point, then we must have that the holonomy of the connection is bounded on the boundary of suppη. From Theorem 10.2 we get that the connections $\mathfrak{A}(t)|_Y$ converge along a subsequence either to a flat connection if ϕ is L^2 bounded, or $\phi\to\infty$ and the rescaled connection converges to some limit. This limit solves is the solution to a Dirac equation on the three manifold. Now if we consider the rescaled solution on the whole of Y, then the

$$\varphi(t)=\frac{\phi(t)}{\|\phi(t)\|_{L^2}}$$

is converging to 0 along suppη, and the limiting object will satisfy a Dirac equation on Y. However, since suppη contains an open set, it follows that $\phi=0$ almost everywhere in Y. This is a contradiction on $\|\phi\|_{L^2}=1$.

Theorem 10.8. Assume Conjecture 10.7. Let $\mathfrak{A}$ be a finite energy solution to the perturbed Kapustin-Witten equations on $\mathbb{R}\times Y$. Then as $t\to\infty$, $\mathfrak{A}$ converges to a perturbed flat connection with finite holonomy.

Proof. For this proof, denote by $\mathfrak{A}_T$ the translate of the solution by $-T$ restricted to $(0,1)$. We are going to prove convergence of the connections $\mathfrak{A}_T$ on the bundle over $(0,1)\times Y$. Hence the starting point for this proof is the sequence of connections $\mathfrak{A}_T$ which are such that

$$\|F_{\mathfrak{A}}+\eta'\|_{L^2((0,1))}\to 0$$

and also satisfy the Kapustin-Witten equations. From Theorem 10.5, we can extract a subsequence of $\mathfrak{A}$ when restricted to $(0,1)\times(Y\setminus\eta)$. We either have that $\|\phi_T\|_{L^2}$ is bounded, or it is not. In the bounded case then we have that $\mathfrak{A}$ converges in C^∞ to a flat connection on $Y\setminus\eta$.

Let $\mathfrak{A}=A+i\phi$ and for each T, let $\mathfrak{A}_T=A_T+i\phi_T$. Since we know that $t\mapsto\|\phi\|_{L^2(Y)}$ is bounded from Conjecture 10.7, we can use Theorem 3.39 and Lemma 10.4 to get that

$$\|\phi_T\|_{L^2_1(0,1)}\le C$$

for some constant C. From this we can deduce that the one forms ϕ_T have a weak L^2_1 convergent subsequence as well as a bound on $\|[\phi,\phi]_T\|_{L^2}$. Since we have the finite curvature assumption we have a uniform bound on the curvature of the connections A_T. Hence, using Uhlenbeck's theorem, we can find a limiting connection $\mathfrak{A}_0 = A_0 + i\phi_0$ that is the weak L^2_1 limit of a subsequence of the connections $\mathfrak{A}_T$. This connection is such that

$$\|F_{\mathfrak{A}_0}\|_{L^2(S)} \leq \liminf_{T' \to \infty} \|F_{\mathfrak{A}}\|_{L^2(S)} = 0,$$

where S is any subset of $(0,1)\times Y$ with closure contained in $(0,1)\times Y \setminus \mathrm{supp}(\eta)$.

In particular, away from the support of η, $\mathfrak{A}_0$ defines a flat $SL(2;\mathbb{C})$ connection. We can thus use the interior estimates for the complex Yang-Mills equation, Theorem 3.44, to get that $\mathfrak{A}_0$ is C^∞ away from the support of η.

Now we can pick $\mathfrak{A}_0$ as our limiting connection and consider the differences

$$\mathfrak{A}_T - \mathfrak{A}_0 = (a + i\varphi)_T.$$

We have L^2_1 convergence of $\varphi \to \phi$ from Theorem 3.39. Moreover, from Taubes' compactness theorem, Theorem 10.5, we have that there exists gauge transformations such that

$$\|a_T\|_{L^2_1} \to 0.$$

Hence we can bootstrap the convergence to C^∞, on the set S, using the estimate for the Yang-Mills equations. Hence we have, away from $\mathrm{supp}(\eta)$, convergence to a C^∞ connection up to gauge transformations.

We now show that the convergence is C^∞ on the whole of the $(0,1)\times Y$. To do this, pick generators of $\pi_1(Y \setminus \mathrm{supp}(\eta))$. Since we have C^∞ convergence, we can consider the holonomy on these generators and this holonomy converges along the subsequence. Moreover, we have that the holonomy satisfies

$$\|F_{\mathfrak{A}_T} + \eta(\mathfrak{A}_T)\|_{L^2} \to 0.$$

and so talking the limit gives that $\mathfrak{A}_0$ is a perturbed flat connection. $\square$

Remark 10.9. A similar method is used in [Don02, Proposition 4.1] to prove the corresponding result for the anti-self-dual equations, although that case is simpler since one can just apply Uhlenbeck's theorem.

10.3 Exponential Decay

We now consider the decay of the perturbed Kapustin-Witten equations to the critical points. Since we know that the solutions converge to flat connections pointwise, most of the analysis here is similar to that given in [Don02, Chapter 4].

In this section we consider the properties of solutions to the θ_v–Kapustin-Witten equations on $(0,\infty)\times Y$ as $t\to\infty$. We assume that $\theta_v(t)$ is such that $\theta_v(t)=\theta_+$ for $t>T_+$, and $\theta_v(t)=\theta_-$ for $T<T_-$.

Proposition 10.10. Assume Conjecture 10.7. Let $\mathfrak{A}$ be a solution to the Kapustin-Witten equations on $(0,\infty)\times Y$ such that

$$\|F_{\mathfrak{A}}+\eta'(\mathfrak{A})\|_{L^2((0,\infty)\times Y)}\in L^p.$$

Then, up to gauge equivalence, $\mathfrak{A}_T|_{\frac{1}{2}\times Y}$ converges in C^∞ to a perturbed flat connection on $P_{\mathbb{C}}\to Y$.

Proof. Consider again the sequence of connections $\mathfrak{A}_T$ on $P_{\mathbb{C}}\to[0,1]\times Y$ which are the translations of $\mathfrak{A}$ by $-T$. So we have that $\mathfrak{A}_T$ is a family of connections with a bound

$$\|F_{\mathfrak{A}_T}\|_{L^2}\le\|F_{\mathfrak{A}_T}\|_{L^p}\to 0 \text{ as } T\to\infty.$$

Then Theorem 10.8 applies and we have that $\|\phi_T\|_{L^2}$ is bounded. But then $\mathfrak{A}$ converges, up to gauge transformations and taking a subsequence, to a perturbed flat connection on $\mathfrak{A}$ on compact subsets. In particular, on the slice $\left\{\frac{1}{2}\right\}\times Y$, the connections converge in C^∞, and we may extract a limit flat connection ρ on Y. ◻

Proposition 10.11. Assume Conjecture 10.7. Suppose we have a solution to the perturbed θ_v–Kapustin-Witten equations, $\mathfrak{A}$, on an adapted bundle $P_{\mathbb{C}}\to\mathbb{R}\times Y$ with limits $\mathfrak{A}_\pm$ and $F_{\mathfrak{A}}+\eta'(\mathfrak{A})\in L^p(\mathbb{R}\times Y)$. Then there exists a constant C such that

$$\|F_{\mathfrak{A}}+\eta'(\mathfrak{A})\|_{L^\infty(Y_t)}\le Ce^{-\delta t},$$

where δ is the norm of the least eigenvalue of $D_{\mathfrak{A}_+}$.

Proof. Since we know that $\|\phi\|_{L^2}$ is bounded, the perturbation is bounded by elliptic estimates. The result then follows from exponential decay for Floer theory. For completeness we give the proof. We follow the outline of [Don02, Theorem 4.2] and apply it to solutions of the perturbed Kapustin-Witten equations. First suppose that $F_{\mathfrak{A}} \in L^2$. We consider the function

$$J(T) = \int_T^{\infty} \|F_{\mathfrak{A}} + \eta'(\mathfrak{A})\|^2_{L^2(Y_t)},$$

and the idea is to prove that for sufficiently large T we have that

$$\frac{dJ}{dt} \leq -\delta J + C(T).f(J), \tag{10.5}$$

for some $\delta, \alpha > 0$, and function $C(T).f(J)$, such that $C \to 0$ as $T \to \infty$ and for J small we have $f(J) \leq J$. Since J is bounded, solutions to this equation will eventually solve the differential inequality

$$\frac{dJ}{dt} \leq -\delta' J$$

for any $\delta' < \delta$, and so the perturbed curvature will decay exponentially.

Note that from (10.4), for T sufficiently large, we have that

$$J(T) = \frac{1}{2}\mathrm{Re}\left(e^{i2\theta_+}(\mathfrak{C}(\mathfrak{A}(T)) - \mathfrak{C}(\mathfrak{A}_+)\right) = e^{2i\theta_+}\left(\mathfrak{C}(\mathfrak{A}(T)) - \mathfrak{C}(\mathfrak{A}_+)\right) \tag{10.6}$$

with the right hand equality following from the fact that the imaginary part of $e^{i2\theta_+}\mathfrak{C}$ is fixed under the gradient flow. We also have that

$$\frac{dJ}{dT} = -\|F_{\mathfrak{A}} + \eta'(\mathfrak{A})\|^2_{L^2(Y)}.$$

Since $\mathfrak{A} \to \mathfrak{A}_+$ in C^∞, we can pick T sufficiently large so that $\mathfrak{A} - \mathfrak{A}_+$ is as small as we like in L^2_2, and so at any given particular T, we can use Proposition 3.13 gauge transform so that we are in the complex Coulomb gauge. To get the required estimates, can write a Taylor expansion

$$|\mathfrak{C}_\eta(\mathfrak{A}) - \mathfrak{C}_\eta(\mathfrak{A}_+)| \leq \langle \mathfrak{a}, (L_{\mathfrak{A}_+} + H_\eta)\mathfrak{a}\rangle + R(\mathfrak{a}). \tag{10.7}$$

Note that we have the second derivative in the Taylor expansion (which is just the Hessian) satisfies

$$\langle \mathfrak{a}, (L_{\mathfrak{A}_+} + H_\eta)\mathfrak{a}\rangle \leq \delta^{-1}\|(L_{\mathfrak{A}_+} + H_\eta)\mathfrak{a}\|^2_{L^2}.$$

This inequality follows from the linear analysis

$$\begin{aligned}\langle \mathfrak{a}, (L_{\mathfrak{A}_+} + H_\eta)\mathfrak{a}\rangle &= \sum_{\lambda \in \mathrm{Spec}(L_{\mathfrak{A}_+})} \lambda \langle f_\lambda e_\lambda, f_\lambda e_\lambda\rangle \\ &\leq \sum_{\lambda} \lambda^{-1}\langle \|\lambda e_\lambda\|^2_{L^2} \leq \delta^{-1}\|(L_{\mathfrak{A}_+} + H_\eta)\mathfrak{a}\|^2_{L^2},\end{aligned}$$

see Chapter 9. Now,

$$\|(L_{\mathfrak{A}_+} + H_\eta)\mathfrak{a}\|_{L^2} \leq \|F_{\mathfrak{A}} + \eta'(\mathfrak{A}) - (F_{\mathfrak{A}_+} + \eta'(\mathfrak{A}_+))\|_{L^2} + O(\|a\|^2_{L^2_1}) \tag{10.8}$$

the remainder term satisfies a bound

$$R(\mathfrak{a}) \leq C\|F_{\mathfrak{A}} + \eta'(\mathfrak{A})\|^2_{L^2} \tag{10.9}$$

with $C \to 0$ as $t \to \infty$. This can be shown by estimating the difference in the holonomy integrals by elliptic estimates for $\mathfrak{a}$, and using the fact that when $\mathfrak{a}$ is L^2 small we can estimate the L^2_1 norm of $\mathfrak{a}$ with the curvature.

For $p > 2$ we can use weighted spaces to prove the estimate (and then it turns out that all the solutions are in L^2). See [Don02, Theorem 4.2]. $\square$

Remark 10.12. In the case that $\theta_v = \frac{\pi}{4}$, [He19, Theorem 7.1] gives exponential decay for solutions to the $\frac{\pi}{4}$–Kapustin-Witten equations that converge to non-degenerate flat connections without assuming that $\phi_0(\partial_t) = 0$.

Corollary 10.13. Assume Conjecture 10.7. Suppose that $\mathfrak{A}$ is a solution to the perturbed θ_v–Kapustin-Witten equations on $[0, \infty) \times Y$ with $\phi_0 = 0$. Then there exists a gauge transformation g and a perturbed flat connection $\mathfrak{A}_+$ such that

$$\|g^*(\mathfrak{A}) - \mathfrak{A}_+\|_{L^2_k(Y_t)} \leq C_k e^{-\frac{1}{2}\delta^{-1}t}.$$

Proof. This proof is much the same as the proof of [Don02, Proposition 4.3]. Write $\mathfrak{A} = \mathfrak{a} + \mathfrak{A}_+$. We may assume that a is in temporal gauge so that $a(\partial_t) = 0$. Then the Kapustin-Witten equations are

$$\frac{d\mathfrak{a}}{dt} = *e^{2i\theta}\left(F_{\mathfrak{A}} + \eta(\mathfrak{A})\right)$$

Hence we have that

$$\frac{d}{dt}\|\mathfrak{a}\|^2_{L^2_k(Y)} < \|F_{\mathfrak{A}} + \eta'(\mathfrak{A})\|_{L^2_k(Y_t)}\|\mathfrak{a}\|_{L^2_k(Y)} \leq C_k e^{-\frac{1}{2}\delta t}\|\mathfrak{a}\|_{L^2_k(Y)}.$$

From this it follows that $\mathfrak{a}$ converges at an exponential rate in L^2_k for all k to a limit $\mathfrak{a}_\infty$. This limit is gauge equivalent to $\mathfrak{A}_+$ and so for each t we can find a smooth gauge transformation g_t such that

$$g_t(\mathfrak{A}) \to \mathfrak{A}_+.$$

Taking a subsequence if necessary these g_t have a smooth limit g which is such that

$$g(\mathfrak{a} + \mathfrak{A}_+) = \mathfrak{A}_+.$$

If we then apply this gauge transformation for all t (i.e. apply the constant gauge transform g) then we have that the limit $\lim_{t\to 0}\left(g^*(\mathfrak{A}) - \mathfrak{A}\right) = 0$ and the exponential decay properties follow. $\square$

Part II

Witten's Conjectures and The Nahm Pole Boundary Condition

Chapter 11

Singular Solutions to the Kapustin-Witten equations with bounds on F_A

In this chapter we study the singular solutions to the Kapustin-Witten equations by studying the Kapustin-Witten equations as an ODE. We only consider structure group $SU(2)$ or $SO(3)$, so that the Lie algebra is isomorphic to $\mathfrak{su}(2)$.

Theorem 11.1. Let $U \subseteq Y$ be a subset with compact closure, and suppose that $A + i\phi$ is a solution to the Kapustin-Witten equations on $[0,1] \times Y$ such that $F_A \in L^\infty([0,1] \times U)$. Then ϕ is either bounded on $[0,1] \times U$ or admits a Nahm pole singularity at at least one end of the interval.

The proof of the theorem is broken into three steps. The first is to consider some linear algebra estimates for operators on the lie algebra $\mathfrak{su}(2)$. This is the reason that we restrict to specific choices of gauge groups. The second step is to consider the Kapustin-Witten equations as an ODE problem for each point $p \in Y$. We can do this due to the assumptions on F_A. The final step is to use the pointwise analysis to obtain global bounds on $d_A\varphi$, where $\varphi = \frac{\phi}{|\phi|}$ on U and deduce that it tends to 0.

11.1 Linear Algebra Estimates

We start by studying the linear algebra of operator

$$\phi \mapsto *(\phi \wedge \phi)$$

that occurs in the Kapustin-Witten equations. We will study this operator in the special case that $\mathfrak{g} = \mathfrak{su}(2)$ and ϕ is a three form on a three manifold, as the application we have in mind is solutions on $I \times Y^3$, where I is an interval. We pick a basis for $\mathfrak{su}(2)$, t_i, such that

$$[t_i, t_j] = \varepsilon_{ijk} t_k.$$

where ε is the Levi-Civita symbol. We pick (and fix) a metric on $\mathfrak{su}(2)$ such that these t_i are orthonormal. Then if we pick normal coordinates at some point $p \in Y$, we can write

$$\phi = a^i{}_j t_i dx^j$$

Definition 11.2. We call the matrix $(a^i{}_j)$ the matrix of ϕ at p.

Note that we can pick the coordinates such that $a^1{}_j = a^1{}_1 \delta^1_j$.

Lemma 11.3. The matrix of $*(\phi \wedge \phi)$ at the point p is given by the matrix $b^i{}_j$, where $b^i{}_j$ is the matrix of minors of $a^i{}_j$. That is, $b^i{}_j$ is equal to the determinant of the 2×2 matrix obtained from deleting the ith row and jth column from a, multiplied by $(-1)^{i+j}$.

Proof. This is a direct calculation. We will just prove that this is the case for the entry $b^1{}_1$. If we consider the rows of a as three vectors, v_1, v_2, v_3, then the first row of b is the cross product of v_2 and v_3. If we have

$$v_2 = (l, m, n), \quad v_3 = (p, q, r)$$

then $b^1{}_1 = mq - rn$, as required. □

The first application of this result is the following.

Lemma 11.4. The operator $\phi \mapsto *(\phi \wedge \phi)$ satisfies

$$| * (\phi \wedge \phi)| \leq \frac{1}{\sqrt{3}} |\phi|^2.$$

Proof. With out loss of generality, assume that the matrix of ϕ is given by

$$a = \begin{pmatrix} l & 0 & 0 \\ m & c & d \\ e & f & g \end{pmatrix},$$

then the matrix of $*(\phi \wedge \phi)$ is given by

$$b = \begin{pmatrix} cg - df & de - mg & mf - ec \\ 0 & lg & -lf \\ 0 & -ld & lc \end{pmatrix}. \tag{11.1}$$

Note that the bottom two rows of the matrix b have norms less than $|v_1||v_3|$ and $|v_1||v_2|$, respectively, where v_i is the ith row of the matrix of ϕ. This is a special case of the more general fact that for $v, w \in \mathfrak{su}(2)$ with our choice of inner product

$$|[v, w]| \leq |v||w|.$$

This follows, for example, from the formula for the norm of the cross product in $\mathbb{R}^3$. Hence we have, by considering the norms of each row in the matrix b, that

$$\begin{aligned} |*\phi \wedge \phi|^2 &\leq |v_1|^2|v_2|^2 + |v_2|^2|v_3|^2 + |v_3|^2|v_1|^2 \\ &\leq \frac{1}{3}\left(2\left(|v_1|^2|v_2|^2 + |v_2|^2|v_3|^2 + |v_3|^2|v_1|^2\right) + |v_1|^4 + |v_2|^4 + |v_3|^4\right) \\ &= \frac{1}{3}\left(|v_1|^2 + |v_2|^2 + |v_3|^2\right)^2 = \frac{1}{3}|\phi|^4. \end{aligned}$$

The second inequality follows from the estimate $cd < \frac{1}{2}(c^2 + d^2)$ for any real numbers c, d. The result follows. □

The next result proves that the previous inequality is sharp.

Proposition 11.5. Suppose $|\phi| = 1$ and that

$$*(\phi \wedge \phi) = \lambda\phi.$$

Then $\lambda = \pm\frac{1}{\sqrt{3}}$ or $\lambda = 0$. Moreover, solutions to this equation at the point p exist, and if $\lambda = \pm\frac{1}{\sqrt{3}}$ then the matrix of ϕ is a multiple of an orthonormal matrix.

Proof. For the moreover statement, a solution with $\lambda = 0$ is given by taking ϕ to have matrix $a^i_j = \delta^i_1\delta^1_j$. For $\lambda = 1/\sqrt{3}$ we can pick

$$\phi = \frac{1}{\sqrt{3}}(t_1dx^1 + t_2dx^2 + t_3dx^3)$$

then

$$*(\phi \wedge \phi) = \frac{1}{3}(t_1dx^1 + t_2dx^2 + t_3dx^3) = \frac{1}{\sqrt{3}}\phi.$$

To get $\lambda = -1/\sqrt{3}$ we can just take $-\phi$.

From the form of b in (11.1), if $*(\phi \wedge \phi) = \lambda\phi$ then we can write the matrix for ϕ as

$$a = \begin{pmatrix} l & 0 & 0 \\ 0 & m & c \\ 0 & d & e \end{pmatrix}.$$

We can then see that the matrix for $*(\phi \wedge \phi)$ takes the form

$$b = \begin{pmatrix} me - cd & 0 & 0 \\ 0 & le & -ld \\ 0 & -lc & lm \end{pmatrix}.$$

Hence, if $\lambda\phi = *(\phi \wedge \phi)$, comparing coefficients gives, for example, that

$$\lambda m = le, \lambda e = lm \Rightarrow \frac{\lambda}{l} = \frac{m}{e} = \frac{l}{\lambda}$$

and so

$$l = \pm\lambda, m = \pm e, c = \mp d.$$

Hence the lower right 2×2 matrix of ϕ has orthogonal columns and its norm squared $b^2 + d^2 + c^2 + d^2 = 1 - l^2$. This implies that

$$\lambda l = \pm l^2 = be - cd = b^2 + d^2 = \frac{1 - l^2}{2}.$$

We can solve this for l if and only if $l = \lambda$. This gives $l = \pm\frac{1}{\sqrt{3}}$, $\lambda = \pm\frac{1}{\sqrt{3}}$, $m = e$, $c = -d$. Finally we can consider the matrix

$$\sqrt{3}\begin{pmatrix} m & -d \\ d & m \end{pmatrix}.$$

The norm of this matrix is $3(1 - \frac{1}{3}) = 2$. It follows that the determinant is 1 and so this is a rotation matrix. Putting this all together, it follows that if $*(\phi \wedge \phi) = \lambda\phi$, there exists some θ such that the matrix of ϕ is given by

$$a = \pm\frac{1}{\sqrt{3}}\begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta \\ 0 & \sin\theta & \cos\theta \end{pmatrix}.$$

□

Remark 11.6. Note that previous result implies that if at any point $*(\phi \wedge \phi)$ and

ϕ are colinear, with $\phi \wedge \phi \neq 0$ then it follows that

$$*(\phi \wedge \phi) = \frac{|\phi|}{\sqrt{3}}\phi,$$

and the coefficients of ϕ with respect to an orthonormal frame on the manifold form a normal basis of $\mathfrak{su}(2)$. In other words, $\phi_p \in T_p^*M \otimes \mathfrak{su}(2)$ describes a linear isometry.

The final result of this section is a lower bound for $|*\phi \wedge \phi|$.

Proposition 11.7. Suppose $|\phi| = 1$. Then

$$\frac{1}{\sqrt{3}}|\langle \phi, *\phi \wedge \phi\rangle| \leq |*\phi \wedge \phi|^2. \tag{11.2}$$

Proof. We will prove that

$$\langle \phi, *\phi \wedge \phi\rangle^2 \leq 3|*\phi \wedge \phi|^4.$$

First, using Lemma 11.3, we have that the first row of the matrix of $*(\phi \wedge \phi)$ is the cross product of the second and third rows of the matrix of ϕ. There are also similar statements for the other rows given by cycling the indices. Hence

$$\langle \phi, *\phi \wedge \phi\rangle = 3\langle v_1, v_2 \wedge v_3\rangle,$$

where v_1, v_2, v_3 are the rows of the matrix of ϕ. We can write

$$\langle v_1, v_2 \wedge v_3\rangle^2 = |v_1|^2 |v_2^\perp \wedge v_3^\perp|^2,$$

where $v_i^\perp := v_i - |v_1|^{-2}\langle v_i^\perp, v_1\rangle v_1$ is the component of v_i orthogonal to v_1. The norm of the cross product in the previous expression is given by

$$|v_2^\perp \wedge v_3^\perp|^2 = |v_2^\perp|^2|v_3^\perp|^2 - \langle v_2^\perp, v_3^\perp\rangle^2 \leq |v_2^\perp|^2|v_3^\perp|^2.$$

Note that we also have

$$|v_1 \wedge v_2|^2 = |v_1|^2|v_2^\perp|^2,$$

and hence

$$|v_1|^2\langle v_1, v_2 \wedge v_3\rangle \leq |v_1 \wedge v_2|^2|v_1 \wedge v_3|^2.$$

Using that $|\phi| = 1$, and that $\langle a, b \wedge c\rangle = \langle b, c \wedge a\rangle$, we have

$$\langle v_1, v_2 \wedge v_3\rangle^2 = \sum_{i=1}^{3} |v_i|^2 \langle v_i, v_{i+1} \wedge v_{i+2}\rangle^2,$$

where we consider the subscript indices modulo 3. Putting all this together gives that

$$\frac{1}{9}\langle\phi, *\phi \wedge \phi\rangle^2 \leq |v_1 \wedge v_2|^2 |v_2 \wedge v_3|^2 + |v_2 \wedge v_3|^2 |v_3 \wedge v_1|^2 + |v_2 \wedge v_3|^2 |v_3 \wedge v_1|^2 =: \beta.$$

Using the estimate $2ab \leq a^2 + b^2$ for $a, b \in \mathbb{R}$ we have a bound

$$\beta \leq |v_1 \wedge v_2|^4 + |v_2 \wedge v_3|^4 + |v_3 \wedge v_1|^4.$$

Also

$$\begin{aligned} |*\phi \wedge \phi|^4 &= \left(|v_1 \wedge v_2|^2 + |v_2 \wedge v_3|^2 + |v_3 \wedge v_1|^2\right)^2 \\ &= |v_1 \wedge v_2|^4 + |v_2 \wedge v_3|^4 + |v_3 \wedge v_1|^4 + 2\beta \\ &\geq 3\beta. \end{aligned}$$

Hence

$$\frac{1}{9}\langle\phi, *\phi \wedge \phi\rangle^2 \leq \frac{1}{3}|*\phi \wedge \phi|^4,$$

as required. □

11.2 Pointwise Estimates

We start this section by considering pointwise estimates for solutions to the $\frac{\pi}{4}$–Kapustin-Witten equations $A + i\phi$ on an $SU(2)$ bundle $P \to I \times Y$, under the assumption that $\|F_A\|_{L^\infty}$ is finite and $\phi(\partial_t) = 0$. If we fix $x \in Y^3$, then we can consider the ODE for $A + i\phi$ evaluated at x. We view this as an evolution equation and will refer to the t coordinate as time. In particular we have

$$\frac{d\phi_x}{dt} = *\left(\phi_x \wedge \phi_x - F_{Ax}\right).$$

where the subscript x denotes evaluation at x, and so we consider ϕ_x as a function $I \to \mathbb{R}^{3^*} \otimes \mathfrak{g}$. In the following we omit the subscripts to ease the notation.

Remark 11.8. In the case where $|F_A|$ remains small under the gradient flow and

$|\phi|$ becomes large, the leading order behaviour of ϕ is governed by the equation

$$\frac{d\phi}{dt} = *(\phi \wedge \phi). \tag{11.3}$$

If we write $\phi = \sum_{i=1}^3 \phi_i dx^i$ for a choice of normal coordinates at x then we can rewrite (11.3) as the system

$$\frac{d\phi_i}{dt} = [\phi_j, \phi_k] \tag{11.4}$$

where (i, j, k) is an even permutation. The equations (11.4) are known as Nahm's equations (or the Nahm equations for consistent equation naming), and so in the situation where $|F_A|$ remains bounded the behaviour of ϕ is similar to that of solutions to the Nahm equations. Indeed, in the Generalised Seiberg-Witten approach to the Kapustin-Witten equations (see [Hay13, Section 4.4]), a solution to the Kapustin-Witten equations on $I \times Y^3$ can be viewed as an equation on Y^3 for a family of data that satisfies a perturbed version of the Nahm equations.

We now consider the Nahm equations with perturbation by a bounded function. That is, we study

$$\frac{d\phi}{dt} = *(\phi \wedge \phi) + f \tag{11.5}$$

for some bounded function $f : \mathbb{R} \to \mathbb{R}^{3^*} \otimes \mathfrak{g}$.

Definition 11.9. Throughout this section, for a solution ϕ of (11.5) on an interval $[a, b]$ we will denote by s the norm

$$s = |\phi(t)|$$

and denote by φ the normalised solution

$$\varphi = \frac{\phi}{s}.$$

We also define the quantity

$$\alpha = \langle *(\varphi \wedge \varphi), \varphi \rangle,$$

which we consider as a function on $[a, b]$.

In this section we will mainly study how the quantities s and α vary, and so it is useful to calculate their time derivatives which are implied by (11.5).

Lemma 11.10. Let ϕ be a solution to (11.5) on the interval $[a, b]$. Then

$$\begin{aligned}\frac{ds}{dt} &= s^2\langle *\varphi \wedge \varphi, \varphi\rangle + \langle f, \varphi\rangle \\ \frac{d\varphi}{dt} &= s\left(*\varphi \wedge \varphi - \langle *\varphi \wedge \varphi, \varphi\rangle\varphi\right) + \frac{1}{s}\left(f - \langle f, \varphi\rangle\varphi\right) \\ \frac{d\alpha}{dt} &= 3s\left(|\varphi \wedge \varphi|^2 - \alpha^2\right) + \frac{3}{s}\langle *\varphi \wedge \varphi - \alpha\varphi, f\rangle.\end{aligned} \tag{11.6}$$

Proof. For s we can write

$$\frac{ds^2}{dt} = 2s\frac{ds}{dt} = \langle \phi, *\phi \wedge \phi + f\rangle = s^3\langle \varphi, *\varphi \wedge \varphi\rangle + s\langle \varphi, f\rangle.$$

This proves the first equation in Equation (11.6). For the second equation we have

$$\frac{d\phi}{dt} = \frac{d}{dt}(s\varphi) = \frac{ds}{dt}\varphi + s\frac{d\varphi}{dt}$$

which implies that

$$\frac{d\varphi}{dt} = \frac{1}{s}\left(\frac{d\phi}{dt} - \frac{ds}{dt}\right).$$

Substituting in the values for the derivatives on the right hand side gives the result. Finally, for α, note that for $\mathfrak{g}$ valued one forms a, b, c, $\langle *(a \wedge b), c\rangle = \langle *(b \wedge c), a\rangle$ and $*(a \wedge b) = *(b \wedge a)$. Hence the quantity

$$\langle *a \wedge b, c\rangle$$

is invariant under permutations of a, b, c. It follows that

$$\begin{aligned}\frac{d\alpha}{dt} = \frac{d}{dt}\langle *\varphi \wedge \varphi, \varphi\rangle &= 3\left\langle *\varphi \wedge \varphi, \frac{d\varphi}{dt}\right\rangle \\ &= 3\left\langle *\varphi \wedge \varphi, s\left(*\varphi \wedge \varphi - \alpha\varphi\right) + \frac{1}{s}\left(f - \langle f, \varphi\rangle\varphi\right)\right\rangle \\ &= 3s\left(| * \varphi \wedge \varphi|^2 - \alpha^2\right) + \frac{3}{s}\left\langle *\varphi \wedge \varphi, f - \langle f, \varphi\rangle\varphi\right\rangle.\end{aligned}$$

For the $\frac{1}{s}$ coefficient, we can write

$$\langle *\varphi \wedge \varphi, f - \langle f, \varphi\rangle\varphi\rangle = \langle *\varphi \wedge \varphi, f\rangle - \alpha\langle \varphi, f\rangle = \langle *\varphi \wedge \varphi - \alpha\varphi, f\rangle.$$

and the result follows. ∎

We now prove that s becomes monotonic in t if s becomes large enough. This a key proposition to analyse singular solutions. Looking at (11.6), this is clear if we have a bound $\alpha > c > 0$. Hence, the main content of the following is proving that α is bounded below by a positive constant.

Proposition 11.11. Suppose ϕ is a smooth solution to (11.5) on $[a, b)$ for $-\infty \leq a < b < \infty$, such that

$$s = |\phi| \to \infty$$

as $t \to b$. Then there exists a $t_0 \in [a, b)$ such that for $t > t_0$, s is monotonic in t.

Proof. The function s is increasing provided

$$s^2\alpha > |f|,$$

where $|f|$ denotes the maximum value of f. By the mean value theorem the value of $s^2\alpha$ is unbounded. To see this, note that for any ε we have that there is a $\xi \in [a, b)$ such that

$$\frac{ds}{dt}(\xi) = s^2\alpha - \langle f, \varphi\rangle = \frac{s(b-\varepsilon) - s(a)}{b - \varepsilon - a}.$$

The right hand side is unbounded as $\varepsilon \to 0$, where as $\langle f, \varphi\rangle$ is bounded and so $s^2\alpha$ is unbounded. To prove that $\frac{ds}{dt}$ is always positive after a particular t we need to consider how α varies. Now, for α we have that

$$\frac{d\alpha}{dt} \geq 3\left(s\beta^2 - \frac{|f|\beta}{s}\right),$$

where $\beta^2 = |\varphi \wedge \varphi|^2 - \alpha^2 = |*(\varphi \wedge \varphi) - \alpha\varphi|^2$. From Proposition 11.7 we have that

$$\beta^2 \geq \frac{1}{\sqrt{3}}\alpha - \alpha^2$$

and so α is increasing if

$$s^2\beta > |f| \Leftrightarrow s^4\left(\frac{1}{\sqrt{3}}\alpha - \alpha^2\right) > |f|^2.$$

Figure 11.1 shows the curves

$$s^4 = \frac{|f|^2}{\alpha\left(\frac{1}{\sqrt{3}} - \alpha\right)}, s^2\alpha = |f| \text{ and } s^2\alpha = 2|f|,$$

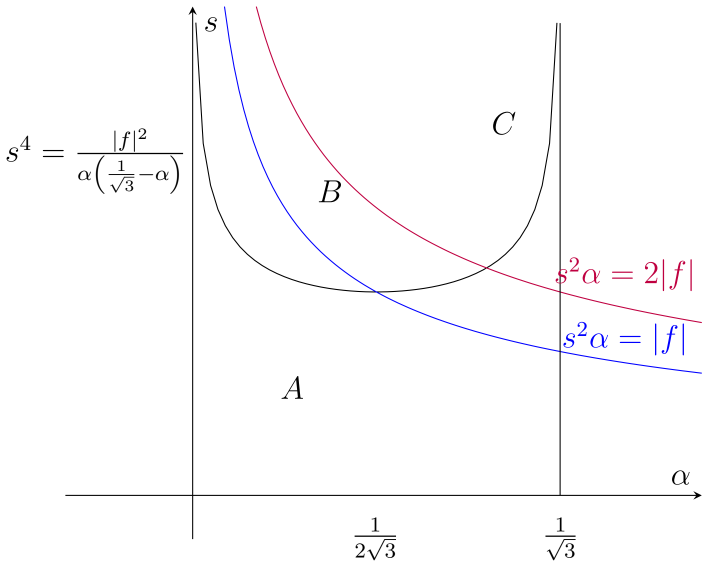

Figure 11.1: Plot of relevant curves in the (α, s) plane for the analysis in Proposition 11.11

.

and three labelled regions, A, B, C. The region A is where $s^4 < \frac{|f|^2}{\alpha\left(\frac{1}{\sqrt{3}}-\alpha\right)}$, B is where $s^2\alpha > |f|$ and C is where $s^2|\alpha| > 2|f|$. At each time t, we consider the point $p(t) = (\alpha(t), s(t))$ in the (α, s) plane. If p is in the complement of region A, α is increasing as a function of t, where as if p is in A, α may be decreasing or increasing. If p is in region B then s is increasing.

Since $s^2\alpha$ is unbounded as a function of t, there is a t_0 where the projection of the solution first enters region C (or it is in region C initially, in which case take $t_0 = a$). We either have that $p(t_0)$ is in A or it's complement. If $p(t_0)$ is in A then α may decrease to a value $\alpha_{\min}$ but s will be increasing. The value $\alpha_{\min}$ is the supremum of the α on the curve $s^2\alpha = 2|f|$ which is in the complement of region A. We can calculate that

$$\alpha_{\min} = \frac{1}{2\sqrt{3}}.$$

It follows that $p(t)$ is always in region B for $t > t_0$, and so s is always increasing after t_0. $\square$

We now study what happens to α as $|\phi| \to \infty$.

Proposition 11.12. Suppose ϕ is a smooth solution to (11.5) on $[a, b)$ with $a < b < \infty$ such that

$$\phi \to \infty$$

as $t \to b$. Then

$$\alpha \to \frac{1}{\sqrt{3}} \text{ as } t \to b.$$

Moreover, there exists a constant $E > 0$ such that

$$\left(\frac{1}{\sqrt{3}} - \alpha\right) < \frac{E}{s^2}$$

where E depends on $|f|$ and the value of α at the earliest time when $s^2\alpha \geq 3|f|$.

Proof. The proof of this is similar to the proof of the previous proposition. Let t_1 be the time where we first have

$$s^2\alpha \geq 3|f|,$$

that is

$$t_1 = \inf_{t\in[a,b)} \{t|s^2\alpha(t) > 3|f|\}.$$

For convenience let $\alpha_1 = \alpha(t_1)$. Then we have for $t \in [t_1, b)$ that

$$\alpha(t) > \eta\alpha_1,$$

where $\eta > \frac{4}{5}$ is some constant. This can be seen by looking at Figure 11.1: Either $(\alpha_1, s(t_1))$ is in the complement of region A or it is not. In the former case we have $\alpha(t) \geq \alpha_1$ for $t \in [t_1, b)$ since α is increasing in this region. In the latter case, α may decrease to a minimum of α_{min} which is where the line $s = s(t_1)$ intersects the curve bounding region A. The ratio

$$\frac{\alpha_{\text{min}}}{\alpha_1} \geq \eta$$

where η is the ratio of the α value intersecting the line $s^2 = 3\sqrt{3}|f|$ and the maximum possible value of α, $\frac{1}{\sqrt{3}}$. This constant is

$$\eta = \frac{1 + \frac{\sqrt{5}}{3}}{2} > \frac{87}{100} > \frac{4}{5}.$$

We now use this to bound $\frac{d\alpha}{dt}$ from below. We have, from Proposition 11.7 and

Lemma 11.10,

$$\begin{aligned}\frac{d\alpha}{dt} &= 3s\left(|\varphi\wedge\varphi|^2-\alpha^2\right)+\frac{3}{s}\langle *(\varphi\wedge\varphi)-\alpha\varphi, f\rangle\\ &\geq 3s\left(\frac{1}{\sqrt{3}}\alpha-\alpha^2\right)-\frac{3}{s}|f|.\end{aligned}$$

After t_1 we have that $\alpha > \frac{4}{5}\alpha_1$. We also have that s is increasing, see Proposition 11.11. Hence we can consider α as a function of s. We have that

$$s^2\alpha(t) > \frac{4}{5}s^2(t_1)\alpha_1 > \frac{12}{5}|f|$$

$$\begin{aligned}\frac{ds}{dt} &= s^2\alpha - \langle f, \varphi\rangle < s^2\alpha + |f|\\ &\leq \left(1+\frac{5}{12}\right)s^2\alpha.\end{aligned}$$

Hence

$$\frac{d\alpha}{ds} = \frac{d\alpha}{ds}\left(\frac{ds}{dt}\right)^{-1} > \frac{12}{17}\left(\frac{3}{s}\left(\frac{1}{\sqrt{3}}-\alpha\right)-\frac{|f|}{s^3\alpha_{\min}}\right).$$

This can be integrated to get that

$$\alpha > \frac{1}{\sqrt{3}} - \frac{D|f|}{s^2\alpha_{\min}} - \frac{C}{s^k}$$

where $k = \frac{36}{17} > 2$ and $C > 0$ is chosen so that the right hand side is equal to α_1 at $s = t_1$ and D is some constant that is independent of the initial data. The result follows. □

Corollary 11.13. We can improve the estimate in Proposition 11.12 to

$$\left(\frac{1}{\sqrt{3}}-\alpha\right) < \frac{D}{s^k}$$

for $k = \frac{36}{17}$.

Proof. This is a single iteration bootstrapping argument. Note that

$$\frac{3}{s}\left(*(\varphi\wedge\varphi)-\alpha\varphi, f\right) > -\frac{3}{s}|f|(|\varphi\wedge\varphi|^2-\alpha^2)^{\frac{1}{2}} > -\frac{C}{s^2},$$

using the result from the proposition. Hence, if we go through the proof of the proposition again, we can improve the estimate for the derivative of α to

$$\frac{d\alpha}{ds} > \frac{k}{s}\left(\frac{1}{\sqrt{3}}-\alpha\right)-\frac{C'}{s^4}.$$

When we integrate this we get

$$\alpha > \frac{1}{\sqrt{3}} - \frac{C''}{s^3} - \frac{D'}{s^k}$$

with the constant $D' > 0$ again depending on the value α_1. □

Remark 11.14. The proof of the previous proposition and its corollary give the precise control

$$\left(\frac{1}{\sqrt{3}} - \alpha\right) \leq \frac{C}{s^{2+\varepsilon}} \tag{11.7}$$

for any $\varepsilon < 1$. (To improve the exponent we can consider the first time when $s^2\alpha > n|f|$ for $n > 3$). This constant depends only on the value of α_1. Hence if we have a family of solutions to (11.5) with a uniform lower bound on α_1 then the bound (11.7) will be uniform. Also note that a uniform lower bound on α_1 is the same thing as a uniform upper bound on s^2 at t_1 by the definition of t_1. Hence if $a, b < \infty$ and s is initially bounded, then $s(t_1) < s(a) + 4|f|$. This analysis is used in Section 11.2.1 to prove uniform estimates for solutions to the Kapustin-Witten equations on a thee manifold times an interval.

Proposition 11.15. Let ϕ be a solution to (11.5) on $[a, b]$. Suppose $s \to \infty$ as $t \to b$. Then φ extends to a differentiable function on $[a, b]$ with derivative 0 at b.

Proof. We have

$$\frac{d\varphi}{dt} = s\left(*(\varphi \wedge \varphi) - \alpha\varphi\right) + \frac{1}{s}\left(f - \langle f, \varphi\rangle\varphi\right)$$

If $s \to \infty$ as $t \to \varepsilon$ then the second term on the right tends to 0. The first term can be bounded by

$$s\left(*(\varphi \wedge \varphi) - \alpha\varphi\right) < ks\left(\frac{1}{\sqrt{3}} - \alpha\right)^{\frac{1}{2}} < ks^{-\varepsilon}$$

for some $\varepsilon > 0$. The result follows. □

We now prove the main result of this section on the blow up rate of unbounded solutions. This is a key feature of the Nahm pole singularities described in [Wit11].

Theorem 11.16. Assume $s \to \infty$ in finite time. Relabelling t if necessary, suppose that $s \to \infty$ as $t \to 0$. Then we have

$$s(t) = -\frac{\sqrt{3}}{t} + \hat{s}(t) \tag{11.8}$$

where $t^{-1}\hat{s}(t)$ is bounded as $t \to 0$.

Proof. Since s is strictly increasing on the interval $[t_0, 0)$ we can either consider s as a function of t or t as a function of s. We can also consider f as a function of s instead. The condition (11.8), is equivalent to being able to write

$$t = -\frac{\sqrt{3}}{s - \tilde{s}(s)} \tag{11.9}$$

where $s\tilde{s}(s)$ is bounded as $s \to \infty$. (We define $\tilde{s}(s) = \hat{s}(t(s))$.) To see this, note that

$$\begin{aligned} s\tilde{s} &= \left(-\frac{\sqrt{3}}{t} + \hat{s}\right)\tilde{s} \\ &= \left(\left(-\frac{\sqrt{3}}{t}\right)\hat{s}\right)\left(1 + \frac{t}{\sqrt{3}}\hat{s} + \left(\frac{t}{\sqrt{3}}\hat{s}\right)^2 + \dots\right). \end{aligned}$$

To go the other way, that is if we have that (11.9) holds and $s\tilde{s}$ is bounded, then we can similarly write

$$\begin{aligned} -t^{-1}\tilde{s} &= \sqrt{3}^{-1}(s - \tilde{s})\,\tilde{s} \\ &= \sqrt{3}^{-1}(s\tilde{s})\,(1 + s\tilde{s} + \dots). \end{aligned}$$

Recall that s satisfies the differential equation

$$\frac{ds}{dt} = \alpha s^2 + \langle f, \varphi \rangle = \alpha s^2 + \hat{f}.$$

where we let $\hat{f}(s) = \langle f(t(s)), \varphi(t(s)) \rangle$. This separable differential equation has solution

$$t(s) = \int_{s_0}^{s} \frac{1}{\alpha u^2 + \hat{f}} du + C,$$

for sufficiently large s_0 as to avoid dividing by 0 in the integrand. The condition that $t \to 0$ as $s \to \infty$ implies that $C = -\int_{s_0}^{\infty} \frac{1}{\alpha u^2 + \hat{f}}$ and so

$$t(s) = -\int_{s}^{\infty} \frac{1}{\alpha u^2 + \hat{f}} du.$$

for $s > s_0$. To prove that there is a function $\tilde{s}$ satisfying (11.9), we can estimate the integral for t and show that

$$t = -\frac{\sqrt{3}}{s} + h(s)$$

for some function h. A little algebra gives that

$$-\frac{1}{s+\tilde{s}} = -\frac{1}{s} + h \Leftrightarrow \tilde{s} = \frac{s^2 h(s)}{1 - sh(s)}.$$

Hence we want to prove that $s^3 h(s) \to L$ for some L as $s \to \infty$ to get that $s^1 \tilde{s} \to L'$. First recall that $\hat{f}$ is bounded and $\alpha \to 1$ as $t \to \infty$. Hence for large u,

$$\frac{1}{\alpha u^2 + \hat{f}} \sim \frac{1}{u^2}.$$

It makes sense to rewrite the integral for t as

$$t(s) = -\int_s^\infty \frac{1}{u^2}\left(\sqrt{3} + \frac{1}{\alpha + \hat{f}/u^2} - \sqrt{3}\right) du = -\frac{\sqrt{3}}{s} - \int_s^\infty \frac{\sqrt{3}}{u^2}\left(\frac{1}{\sqrt{3}\left(\alpha + \hat{f}/u^2\right)} - 1\right) du,$$

and so the function $h(s)$ is given by the expression

$$h(s) = -\int_s^\infty \frac{\sqrt{3}}{u^2}\left(\frac{1}{\sqrt{3}(\alpha + \hat{f}/u^2)} - 1\right) du.$$

For large u, say $u > s_0$ we will have that there is some $0 < \varepsilon < \frac{1}{\sqrt{3}}$ such that

$$\alpha + \frac{\hat{f}}{u^2} > \varepsilon$$

for each $u > s_0$. Hence we get the bound

$$\begin{aligned}
|h(s)| &\leq \frac{\sqrt{3}}{\varepsilon}\int_s^\infty \frac{1}{u^2}\left(\frac{1}{\sqrt{3}} - \alpha(s) + \frac{\|\hat{f}\|_{L^\infty}}{u^2}\right) \\
&\leq \frac{\sqrt{3}}{\varepsilon}\left(\frac{\|\hat{f}\|_{L^\infty}}{s^3} + \int_s^\infty \frac{C}{u^4}\right) < \frac{D}{s^3}.
\end{aligned}$$

where the constant C is as in Proposition 11.12. It follows that $|s^3 h(s)|$ is bounded as $s \to \infty$ and so $s\tilde{s}(s)$ is bounded as $s \to \infty$. Hence considering $\tilde{s}$ as a function of t we get that

$$t = -\frac{\sqrt{3}}{s + \tilde{s}(t)}.$$

From this we have

$$s = -\frac{\sqrt{3}}{t} - \tilde{s}(t),$$

and

$$\frac{1}{t}\tilde{s}(t) = \frac{1}{\sqrt{3}}\left(s\tilde{s}(t) + \tilde{s}^2\right).$$

The right hand side is bounded as $t \to 0$.

□

Corollary 11.17. Let ϕ be as in the theorem. There exists a constant $C > 0$ such that

$$|t\phi - \sqrt{3}\varphi| < Cs^2$$

is bounded.

Proof. We have that, in the notation from the previous proof

$$\begin{aligned} t\phi - \sqrt{3}\varphi &= \left(t - \frac{\sqrt{3}}{s}\right) s\varphi \\ &= sh\varphi. \end{aligned}$$

The result follows from the fact that $s^3 h$ is bounded and $|\varphi| = 1$. □

11.2.1 Uniform Estimates

We now consider a solution A, ϕ to the Kapustin-Witten equations on the manifold $[a, b] \times Y^3$. We extend the definition of α in the previous section to a function on Y. In this section we prove that there is a uniform constant M such that

$$\left(\frac{1}{\sqrt{3}} - \alpha\right) \leq (\varepsilon - t)^{2+\delta} M. \tag{11.10}$$

for all $p \in Y$. The constant M depends on the initial data at $t = a$.

Lemma 11.18. Let A, ϕ be a solution to the Kapustin-Witten equations on $[0, \varepsilon) \times Y$. Then there exists a constant

$$\alpha_1 > 0$$

such that for all $p \in Y$, either the value

$$t_0(p) = \inf_{t \in [0,\varepsilon)} \{s^2\alpha(t_0) > 3\|f\|_{L^\infty}\},$$

exists and then we have

$$\alpha(t_0, p) > \alpha_1.$$

or t_0 does not exist and $|\phi(t,p)| < 4\varepsilon\|f\|_{L^\infty} + \|\phi\|_{L^\infty(Y_0)}$ for all $t \in [0.\varepsilon)$.

Proof. Suppose $s^2\alpha < 3\|f\|_{L^\infty}$. Then we have a bound

$$\frac{ds}{dt} = s^2\alpha + f \leq 4\|f\|_{L^\infty(X)}.$$

Hence, $s(t_0) < 4\varepsilon\|f\|_{L^\infty(X)} + \|\phi\|_{L^\infty(Y_0)} = k$, and

$$\alpha(t_0) = \frac{3\|f\|_{L^\infty}}{s(t_0)^2} > \frac{3\|f\|_{L^\infty}}{k^2} =: \alpha_1.$$

□

Lemma 11.19. Consider a solution to (11.5) on the interval $[a,b]$. Let

$$\alpha_m = \inf_{t \in [a,b]} \alpha.$$

Suppose that at $t = a$ we have $s^2\alpha > 3|f|$. Then

$$\alpha_m > \eta\alpha(a).$$

Proof. This is proved in Proposition 11.12. □

Definition 11.20. For a solution $A + i\phi$ to the Kapustin-Witten equations on $[0,\varepsilon) \times Y$ with $|\varphi(t)| \to \infty$ as $t \to \varepsilon$, we define the constant $\alpha_{\min}$ to be

$$\alpha_{\min} := \eta\alpha_1,$$

with α_1 defined as in Lemma 11.18 and η defined as in Lemma 11.19.

Proposition 11.21. Let ϕ be a solution to (11.5) on $[0,\varepsilon)$ Suppose that at some time t_0, $\alpha(t_0) > c > 0$ and $s^2(t_0)\alpha(t_0) > 3|f|$. Then for all $t \in [t_0, b)$,

$$s^k\left(\frac{1}{\sqrt{3}} - \alpha\right) \leq M(c, |f|).$$

for some $k > 2$. The constant in the bound is such that $M \to \infty$ as $c \to 0$ and $M \to 0$ as $c \to \infty$.

Proof. The first part of the statement is just Corollary 11.13. For the last part note that if we have

$$s^2\alpha = 3|f|$$

and $\alpha \to 0$ then $s^2 \to 0$. Hence for the constant C defined in the proof of Corollary 11.13 as

$$\alpha \geq \frac{1}{\sqrt{3}} - \frac{D}{s^3} - \frac{C}{s^k},$$

C becomes increasingly large as $\alpha(t_0) \to 0$ for equality to be satisfied at t_0. $\square$

Theorem 11.22. Let A, ϕ be a solution to the Kapustin-Witten equations on $[0, \varepsilon) \times Y$. Then there exists a constant M, depending on $\varepsilon, \|F_A\|_{L^\infty}, A_0, \phi_0$ such that

$$\sup_{p \in Y} |\phi(t,p)|^2 \left(\frac{1}{\sqrt{3}} - \alpha(t,p) \right) \leq M.$$

Proof. For each p we can split the interval in two

$$[0, \varepsilon) = [0, t_0) \cap [t_0, \varepsilon),$$

where

$$t_0(p) = \inf_{t \in [0,\varepsilon)} \{s^2\alpha(t_0) > 3\|f\|_{L^\infty}\}.$$

In the first interval

$$s < s_0 + 4\|f\|_{L^\infty} = k$$

and so

$$s^2 \left(\frac{1}{\sqrt{3}} - \alpha \right) \leq \frac{2}{\sqrt{3}} k^2.$$

For the second interval Proposition 11.21 gives the required bound, provided we can bound $\alpha(t_0)$ below. The constant bound depends on $\alpha(t_0)$. By Lemma 11.18 we have that this is bounded below and so putting these together proves the theorem. $\square$

Corollary 11.23. Let A, ϕ be as in the theorem. Suppose at $p \in Y$, $\phi(p) \to \infty$, then there is a constant M' depending only on $A_0, \phi_0, \|F_A\|_{L^\infty}$ such that

$$\frac{1}{(\varepsilon - t)^k} \left(\frac{1}{\sqrt{3}} - \alpha \right) \leq M'.$$

for some $k > 2$.

Proof. This follows from the theorem and the estimate

$$s \geq \frac{C}{\varepsilon - t}.$$

To prove estimate, note that

$$\frac{ds}{dt} \leq \frac{1}{\sqrt{3}} s^2 + \|f\|_{L^\infty}.$$

We can solve this differential inequality to get that

$$s = A \tan(B(t + c))$$

for some constants A, B, c in an interval $[\varepsilon - \delta, \varepsilon)$ and take c to be such that $s \to \infty$ as $t \to \varepsilon$ to get the required inequality. The constant δ is does not depend on $p \in Y$ and so we can extend the estimate using the boundedness of (A, ϕ) in the interior of $[0, \varepsilon) \times Y$. $\square$

Chapter 12

Nahm Pole Solutions with Bounded Real Curvature

We now apply the pointwise results from the previous section to study solutions to the Kapustin-Witten equations on $[0, \varepsilon) \times Y$ under the assumption that $\|F_A\|_{L^\infty}$ remains bounded and $\phi(\partial_t) = 0$. We first show that we can extend A to a connection $\bar{A}$ on $\bar{X} = [0, \varepsilon] \times Y$ with $A \in C^\alpha(\bar{X})$. To do this we show that $\lim_{t \to \varepsilon} A$ exists and is in L^p_1 for all $p \geq 1$ and then use the Sobolev embedding theorems on Y. We start by considering the limit $\|A\|_{L^p(Y)}$ using Uhlenbeck's theorem.

Theorem 12.1 ([Uhl82, Theorem 1.5])**.** Let A_i be a sequence of connections on a bundle $P \to Y$, such that there is some $M > 0$ with

$$\|F_{A_i}\|_{L^p} < M.$$

Then, along a subsequence, there exists a sequence of gauge transformations $\varphi_i \in L^p_2$ such that and a connection $A \in L^p_1$, with $\|F_A\|_{L^p} \leq M$, such that

$$\varphi_i^* A_i \rightharpoonup A$$

weakly in L^p_1.

Lemma 12.2. Let $A + i\phi$ be a solution to the Kapustin-Witten equations on $(0, \varepsilon) \times Y$, such that $\|F_A\|_{L^\infty}$ is bounded and $\|\phi\|_{L^\infty} \to \infty$ as $t \to \varepsilon$. Then $(\varepsilon - t)\phi \in L^\infty(Y)$ and

$$(\varepsilon - t) * (\phi \wedge \phi) - \phi \in L^\infty.$$

Proof. We have from that for each $p \in Y$, the function

$$s_p(t) = |\phi(t, p)|$$

is either bounded, and so $(\varepsilon - t)s$ is bounded, or it is unbounded. In the bounded case, from Theorem 11.16 we have that

$$s_p(t) = \frac{\sqrt{3}}{\varepsilon - t} + \hat{s}_p$$

and $|\hat{s}_p| < M(\varepsilon - t)$, with M independent of the choice of p. In the unbounded case we can assume that $s_p(t)$ satisfies a differential inequality of the form

$$\frac{ds_p}{dt} > s_p^2 \beta$$

for some small $\beta > 0$ and $t > t_0$. We can make this assumption since if $s^2\alpha(t_0) < 3\|F_A\|_{L^\infty}$, say, then $s < M$ with M independent of p, see Lemma 11.18. If $s\alpha^2(t_0) > 3\|F_A\|_{L^\infty}$ for some t_0 then we have that for $t > t_0$ $s^2\alpha > 2\|F_A\|_{L^\infty}$ and so we can take $\beta = \frac{\alpha_1}{2}$, with α_1 as in Lemma 11.18. Note that α_1 is a does not depend on p. Hence we have

$$\frac{1}{s_p} < \beta(C - t)$$

for some C such that

$$\frac{1}{s_p(\varepsilon)} = \beta(C - \varepsilon)$$

It then follows that

$$(\varepsilon - t)s_p = \frac{\varepsilon - t}{\beta(C - t)} \leq \frac{1}{\beta}.$$

Hence we have a uniform bound on s_p, and so $\varphi \in L^\infty$.

From Corollary 11.17, we have that at $p \in Y$

$$|(\varepsilon - t) - \sqrt{3}|\phi(p)|^{-1}| \leq C|\phi(p)|^{-2}$$

with C independent of p, and so if we write $s = |\phi|$ then

$$(\varepsilon - t) * (\phi \wedge \phi) - \phi = s\left(\sqrt{3} * \varphi \wedge \varphi - \varphi\right) + O(1).$$

Now we can calculate

$$\begin{aligned} \left|s\sqrt{3}(\varphi \wedge \varphi) - \varphi\right|^2 &= s^2\left(3|*(\varphi \wedge \varphi)|^2 + |\varphi|^2 - 2\sqrt{3}\alpha\right) \\ &\leq 2s^2\left(1 - \sqrt{3}\alpha\right) \to 0 \end{aligned}$$

as $t \to 0$. This proves the result. $\square$

Lemma 12.3. Let $A + i\phi$ be a solution to the Kapustin-Witten equations on $(0,\varepsilon) \times Y$, such that $\|F_A\|_{L^\infty}$ is bounded and $\|\phi\|_{L^\infty} \to \infty$ as $t \to \varepsilon$. Then

$$\int_0^\varepsilon \|d_{A_t}\phi\|_{L^2(Y)} < \infty.$$

Proof. Let $\varphi = (\varepsilon - t)\phi$. Then by the Lemma 12.2, $\varphi \in L^\infty$. Now consider the function

$$f(t) = \langle *F_{A_t}, \varphi(t)\rangle_{L^2(Y)}.$$

This function is bounded on $(0,\varepsilon)$ since F_A and φ are. The derivative of $f(t)$ is given by

$$\begin{aligned}\frac{df}{dt} &= \langle *F_A, -\phi + *(\varepsilon - t)(\phi \wedge \phi - F_A)\rangle - \langle d_A^* d_A\phi, \varphi\rangle \\ &= \langle *F_A, -\phi + *(\varepsilon - t)(\phi \wedge \phi - F_A)\rangle - \|d_A\phi\|_{L^2}.\end{aligned}$$

From Lemma 12.2 we have that the first inner product on the right hand side is bounded. It then follows that

$$\|d_{A_t}\phi\|_{L^2} \in L^1([0,\varepsilon)).$$

□

Proposition 12.4. Let $t \mapsto A_t$ be a smooth family of smooth connections on a bundle $P \to Y^3$ for $t \in (0,\varepsilon)$, with $\|F_{A_t}\| \leq M$ and Y compact. And suppose that $A_t + i\phi_t$ is a solution to the Kapustin-Witten equations on $(0,\varepsilon) \times Y$. Then up to gauge there exists a unique connection $A \in C^{0,\alpha}$, for $\alpha \in (0,1)$ such that for any sequence t_i with $t_i \to \varepsilon$, there exists a sequence of gauge transformations such that

$$\varphi_i^* A_{t_i} \to A_i$$

in $C^{0,\alpha}$.

Proof. We use that $L^\infty \hookrightarrow L^p$ for compact Y to get that F_{A_t} is uniformly bounded in L^p for all p. We also have the compact Sobolev embedding

$$L_1^p \hookrightarrow C^{0,\alpha} \text{ when } \alpha < 1 - \frac{3}{p}.$$

Hence, taking p large we can get a limiting connection $A \in C^{0,\alpha}$ using Theorem 12.1. The gauge transformations φ_{t_i} are in L_2^p, and so have a $C^{1,\alpha}$ convergent subsequence for large p. Take the limiting gauge transformation φ and apply it to the solution $A + i\phi$ by viewing it as a gauge transformation constant in t. The image of the solution $A + i\phi$ under φ^* is still, of course, a solution to the Kapustin-Witten equations in temporal gauge. Hence we have that

$$\|d_{\varphi^* A_t} \varphi^* \phi\|_{L^2} \in L^2([0, \varepsilon))$$

from Lemma 12.3. The Kapustin-Witten equations for $A + i\phi$ give that

$$\frac{da}{dt} = - * d_A \phi$$

and hence for $s_1, s_2 \in (0, \varepsilon)$,

$$\|a(s_2) - a(s_1)\|_{L^2} \leq \int_{s^1}^{s^2} \|d_A \phi\|_{L^2} \to 0 \text{ as } |s_2 - s_1| \to 0.$$

Hence, the sequence a_{t_i} is Cauchy in L^2 and so $\varphi^* A_{t_i}$ must converge to A in L^2. In particular,

$$\varphi^* A_t \to A \text{ as } t \to \varepsilon.$$

and so A is unique up to gauge transformations. □

Theorem 12.5. Let $A_t + i\phi_t$ be a solution to the Kapustin-Witten equations with $\|\phi\|_{L^\infty} \to \infty$ as $t \to \varepsilon$ and $\|F_A\|_{L^\infty}$ bounded. Denote by $A + i\phi$ the limiting connection $A_t + i\phi_t$, and let $\varphi = (\varepsilon - t)\phi$. Then $\varphi \in L^2_2(Y)$ and is a solution to

$$d_A \varphi = 0.$$

Proof. Since $\|d_A \phi\|_{L^2} \in L^1((0, \varepsilon))$ we have that $(\varepsilon - t)\|d_A \varphi\|_{L^2} \in L^1$. Hence we can find a sequence t_i such that

$$d_{A_{t_i}} \varphi(t_i) \to 0$$

in L^2. Otherwise $\|d_A \varphi\|_{L^2} > \delta$ for some δ which contradicts $\|d_A \phi\| \in L^2$. Now, if $\eta \in \Omega^2(Y)$ we have that

$$\begin{aligned} \langle \varphi, d_A^* \eta \rangle &= \lim_{i \to \infty} \langle \varphi(t_i), d_A \eta \rangle \\ &= \lim_{i \to \infty} \langle d_{A_i} \varphi(t_i), \eta \rangle = 0. \end{aligned}$$

We then have, from elliptic estimates from the operator $d_A + d_A^*$, that

$$\|\varphi(\varepsilon)\|_{L_1^2} \leq C\|\varphi(\varepsilon)\|_{L^2}.$$

If we pick a smooth connection B, such that $A = B + a$, then we use the elliptic estimates for $d_B + d_B^*$ to write

$$\begin{aligned}
\|\varphi\|_{L_2^2} &\leq C\left(\|d_B\varphi\|_{L_1^2} + \|d^*B\varphi\|_{L_1^2} + \|\varphi\|_{L^2}\right) \\
&\leq C\left(\|d_A\varphi\|_{L_1^2} + \|d_A^*\varphi\|_{L_1^2} + \|[a,\varphi]\|_{L_1^2} + \| * [a, *\varphi]\|_{L_1^2} + \|\varphi\|_{L^2}\right).
\end{aligned}$$

But we have that $d_A\varphi = 0$ and $d_A^*\varphi = 0$. Moreover, $a \in C^{0,\alpha} \cap L_1^p$ and $\varphi \in L^\infty$. Hence we have

$$\begin{aligned}
\|[a,\varphi]\|_{L_1^2} &\leq \|a\|_{L^\infty}\|\varphi\|_{L^2} + \|\nabla[a,\varphi]\|_{L^2} \\
&\leq \|a\|_{L^\infty}\|\varphi\|_{L^2} + \|\{\nabla a,\varphi\}\|_{L^2} + \|\{a,\nabla\varphi\}\|_{L^2} \\
&\leq \|a\|_{L^\infty}\|\varphi\|_{L^2} + \|\varphi\|_{L^\infty}\|a\|_{L_1^2} + \|a\|_{L^\infty}\|\varphi\|_{L_1^2} < \infty
\end{aligned}$$

Similarly, $\| * [a, *\varphi]\| \leq C$. Hence we have that $\varphi \in L_2^2$. □

Corollary 12.6. The section $\varphi = (\varepsilon - t)\phi$ is continuous on $[0,\varepsilon] \times Y$.

Proof. We have that φ is continuous on $[0,\varepsilon) \times Y$ by interior regularity of solutions to the Kapustin-Witten equations and from the theorem, $\lim_{t\to\varepsilon}\varphi$ is continuous since $L_2^2 \hookrightarrow C^{0,\alpha}$ for $\alpha < \frac{1}{2}$. Moreover from the analysis in Section 11.2, $t \to \varepsilon(t,p)$ is continuous for each t. Hence if $(t,p) \in [0,\varepsilon) \times Y$ and $(\varepsilon, p') \in Y$, we have

$$|\varphi(t,p) - \varphi(\varepsilon,p')| \leq |\varphi(t,p) - \varphi(\varepsilon,p)| + |\varphi(\varepsilon,p) - (\varepsilon,p')| \leq Cd\left((t,p),(\varepsilon,p')\right).$$

and this proves that φ is continuous on $[0,\varepsilon] \times Y$. □

Proposition 12.7. Suppose $A + i\phi$ is as in the theorem. Then $\varphi = (\varepsilon - t)\phi$ is nowhere zero.

Proof. We have that $d_A\varphi = 0$. Hence if $\varphi(p) \neq 0$ then

$$d|\varphi|^2 = 2\langle d_A\varphi, \varphi\rangle = 0,$$

and so $|\varphi| > 0$ on Y. Hence we just need to prove that there is a point where $\varphi \neq 0$. Suppose on the other hand that

$$\|\varphi\|_{L^\infty} \to 0$$

as $t \to \varepsilon$. If we pick a point $p \in Y$, then in the notation of Section 11.2 we have that

$$(\varepsilon - t)s \to 0.$$

Considering the ODE for s, (11.6), we have

$$\begin{aligned} \frac{d}{dt}\left((\varepsilon - t)s\right) &= (\varepsilon - t)\frac{ds}{st} - s \\ &= (\varepsilon - t)s^2\alpha + (\varepsilon - t)\langle f, \varphi\rangle - s \\ &= s\left((\varepsilon - t)s\alpha - 1\right) + (\varepsilon - t)\langle f, \varphi\rangle \end{aligned} \tag{12.1}$$

Now, since $\alpha < \frac{1}{\sqrt{3}}$ and $(\varepsilon - t)s \to 0$, we can assume that

$$((\varepsilon - t)s\alpha - 1) \leq -\frac{1}{2},$$

and so $x = (\varepsilon - t)s$ satisfies the differential inequality

$$\frac{dx}{dt} \leq -\frac{1}{2}\frac{x}{\varepsilon - t} + C$$

for some $C > 0$ where $(\varepsilon - t)\langle f, \varphi\rangle < C$. Integrating this gives

$$(\varepsilon - t)s \leq D\,(\varepsilon - t)^{\frac{1}{2}}\, e^{Ct}. \tag{12.2}$$

We can then put this estimate back into (12.1) to get that $x = (\varepsilon - t)s$ satisfies the differential inequality

$$\frac{dx}{dt} \leq x\left(K(\varepsilon - t)^{-\frac{1}{2}} - \frac{1}{\varepsilon - t}\right) + C.$$

where $K > De^{C\varepsilon\alpha}$. The differential inequality has solution

$$\log x \leq K(\varepsilon - t)^{\frac{1}{2}} + \log(\varepsilon - t) + Ct + E.$$

Hence

$$(\varepsilon - t)s \leq (\varepsilon - t)e^{K(\varepsilon - t)^{\frac{1}{2}} + Ct + E},$$

and so s is bounded. Moreover this bound can be taken to be uniform on Y. To see this, note that $\varphi = (\varepsilon - t)\phi$ is continuous, and so the constant D in (12.2), and also the constants K and E can also be taken to be uniform. $\square$

Lemma 12.8. Suppose $A + i\phi$ is a solution to the Kapustin-Witten equations on

$[0,\varepsilon)\times Y$ with $\|\phi\|_{L^\infty}=\infty$. Then

$$\lim_{t\to\varepsilon}\mathrm{Im}(CS(A+i\phi))=-\infty.$$

Proof. We have

$$\|F_{A+i\phi}\|_{L^2}^2=\mathrm{Im}\left(CS((A+i\phi)(0))-CS((A+i\phi)(\varepsilon))\right),$$

but on the other hand we have

$$\|F_{A+i\phi}\|_{L^2}^2=\|F_A-\phi\wedge\phi\|_{L^2}^2+\|d_A\phi\|_{L^2}^2\to\infty \text{ as } t\to\varepsilon.$$

and $CS(A+i\phi)(0)$ is finite. □

We sum up the results from this section in the following theorem.

Theorem 12.9. Let $A+i\phi$ be a solution to the Kapustin-Witten equations on $X=[0,\varepsilon)\times Y$, such that $\|F_A\|_{L^\infty(X)}$ is finite. Then either $\mathrm{Im}\left(CS(A+i\phi)\right)$ is bounded or it is not. In the unbounded case, ϕ has a Nahm pole singularity at ε. If $\mathrm{Im}(CS(A+i\phi))$ is bounded then so too is ϕ.

Proof. The section ϕ is either in L^∞ or it is not. In the case where ϕ is unbounded, $\mathrm{Im}CS$ is unbounded from the Lemma 12.8. If ϕ is bounded then

$$\frac{d}{dt}\|\phi\|_{L^2}^2=\langle\phi,*\phi\wedge\phi-F_A\rangle$$

is bounded. We can then use Lemma 3.41 to write

$$-d^*d|\phi|^2=*d*d|\phi|^2=|\nabla_A\phi|^2+\frac{1}{2}|[\phi,\phi]|^2-\langle\phi\circ\mathrm{Ricci},\phi\rangle$$

Applying Stokes' theorem we see that

$$\left(\int_{Y_\varepsilon}-\int_{Y_0}\right)\langle\phi,*(\phi\wedge\phi)-*F_A\rangle=\int_X|\nabla_A\phi|^2+\frac{1}{2}|[\phi,\phi]|^2-\langle\phi\circ\mathrm{Ricci},\phi\rangle.$$

Hence if ϕ is bounded then $d_A\phi$ is bounded and so

$$\begin{aligned}\mathrm{Im}CS(A+i\phi)(t)&=\mathrm{Im}CS(A+i\phi)(0)-\|F_{A+i\phi}\|_{L^2}^2\\&=\mathrm{Im}CS(A+i\phi)(0)-\|F_A-\phi\wedge\phi\|_{L^2}^2-\|d_A\phi\|_{L^2}^2<\infty.\end{aligned}$$

□

Remark 12.10. Since Re$CS(A+i\phi)$ is preserved under the gradient flow, we can rephrase the statement of the previous theorem to use the whole CS functional, instead of only its imaginary part.

Remark 12.11. The Haydys-Witten equations on $W = I \times X^4$, with the connection in temporal gauge, may be written in the form, for a connection $A = A_0 + a$ and self dual two form $\phi \in \Omega^2_+(\mathfrak{g}_P)$,

$$\begin{aligned}
\frac{da}{dt} &= -d_A^{+*}\phi \\
\frac{d\phi}{dt} &= \{\phi, \phi\} - F_A^+
\end{aligned} \tag{12.3}$$

As such, the analysis from Section 11.2 carries over to this case, and so we can consider solutions to (12.3) where $\|\phi\|_{L^\infty(X_t)} \to \infty$.

12.1 Example Solutions

In this section we derive some symmetric solutions to the Kapustin-Witten equations on $I \times SU(2)$. We look for solutions that are equivariant with respect to the group action. We first make some considerations about the choice of gauge before we introduce an ansatz for the form of the solutions. This ansatz reduces the Kapustin-Witten equations (up to a complex gauge transformation) to a plane autonomous system of ODEs. We only get one explicit solution, but the plane autonomous system is sufficient to prove the existence of an infinite family of solutions to the Kapustin-Witten equations satisfying the Nahm pole boundary condition at both ends.

12.1.1 The Ansatz

Definition 12.12. Let G be a Lie group with a lie algebra $\mathfrak{g}$. Then the Maurer-Cartan form $\omega : TG \to \mathfrak{g}$ given by

$$\omega(g)(v) = (L_{g^{-1}})_* v,$$

for $v \in T_g G$, defines an element $\omega \in \Omega^1(G, \mathfrak{g})$. We can interpret this as a connection on the trivial G-bundle $P \to G$, which is the Maurer-Cartan connection.

Lemma 12.13. The Maurer-Cartan connection is flat, and is gauge equivalent to the trivial connection.

Proof. Consider the gauge transformation of P, φ, given by

$$\varphi(g)p = pg.$$

Then $d\varphi\varphi^{-1} = \omega$ and hence ω is the image of the trivial connection under the gauge transformation φ. $\square$

Corollary 12.14. The Maurer-Cartan form satisfies the equation

$$d\omega + \frac{1}{2}[\omega, \omega] = 0.$$

We will assume that our solutions to the Kapustin-Witten equations are of the form $A = (x+iy)\omega$ for some functions $x, y : I \times SU(2) \to \mathbb{R}$ which only depend on the I coordinate, t. Here we identify solutions to the Kapustin-Witten equations with complex connections on the complexified bundle $P_{\mathbb{C}} \to I \times SU(2)$.

Lemma 12.15. The curvature of the connection $(x+iy)\omega$ is given by

$$F_A = (x+iy)d\omega + \frac{1}{2}(x+iy)^2[\omega, \omega] = \frac{1}{2}(x+iy-1)(x+iy)[\omega, \omega].$$

Proof. This is clear from the formula $F_A = dA + \frac{1}{2}[A, A]$ and Corollary 12.14. $\square$

Lemma 12.16. The connection $A = (x+iy)\omega$ satisfies the moment map equation

$$d^*_{x\omega} y\omega = 0.$$

Proof. Since everything is invariant on $SU(2)$ we may work at e. We have that

$$\omega = idx^1 + jdx^2 + kdx^3$$

where $dx^1 = i, dx^2 = j$ and $dx^3 = k$ in $T_eSU(2)^* \cong \mathfrak{su}(2)$. Then $*\omega = idx^2 \wedge dx^3 + jdx^3 \wedge dx^1 + kdx^1 \wedge dx^2$. Hence we have

$$d_{x\omega}(*\omega) = \left(\frac{d}{dt} idx^1(\exp ti) + \frac{d}{dt} jdx^2(\exp tj) + \frac{d}{dt} kdx^3(\exp tk)\right)\Bigg|_{t=0} dx^1 \wedge dx^2 \wedge dx^3 = 0$$

$\square$

Proposition 12.17. The connection $z = (x+iy)\omega$ satisfies the Kapustin-Witten equations on $I \times SU(2)$ if

$$\frac{dx}{dt} = y - 2xy, \frac{dy}{dt} = y^2 - x^2 + x, \tag{12.4}$$

where we view x, y as functions of the coordinate on I, t.

Proof. Note that

$$F_{A+i\phi} = (F_A - \phi \wedge \phi) + i d_A \phi$$

If $z = (x+iy)\omega$ then

$$F_z = \frac{1}{2}(x+iy-1)(x+iy)[\omega,\omega] = \frac{1}{2}(x^2-y^2-x)[\omega,\omega] + i(2xy-y)[\omega,\omega].$$

Hence, comparing real and imaginary parts, we have that

$$F_A - \phi \wedge \phi = \frac{1}{2}(x^2-y^2-x)[\omega,\omega]$$

and

$$d_A\phi = \frac{1}{2}(2xy-y)[\omega,\omega].$$

The result follows from the fact that $*[\omega,\omega] = 2\omega$. □

If we have a solution (x, y) to (12.4) then we can view this as a parametrised curve in the (x, y) plane. That is, one defines a curve $\gamma : I \to \mathbb{R}^2$ given by

$$\gamma(t) = (x(t), y(t)).$$

The curve γ is an integral curve of the vector field $X = (y - 2xy, y^2 - x^2 + x)$. Hence we can study these symmetric solutions to the Kapustin-Witten equations using the theory of plane autonomous systems.

12.1.2 The Autonomous System

We now study solutions $v : I \to \mathbb{R}^2$ of the differential equation

$$\frac{dv}{dt} = X(v), \tag{12.5}$$

where X is given by

$$X = (y - 2xy, y^2 - x^2 + x). \tag{12.6}$$

The results of this analysis are summarised in the following proposition. We call a solution v maximal if there does not exist another smooth solution $w : I' \to \mathbb{R}^2$ with $I \subseteq I'$ such that $v(t) = w(t)$ for $t \in I$.

Proposition 12.18. There exist maximal solutions $v : I \to \mathbb{R}^2$ to (12.5). If $v = (x, y)$ and $x(t)$ is uniformly bounded then v is gauge equivalent to one of two solutions. Either I is compact, in which case the connection $(x + iy)\omega$ is a solution to the Kapustin-Witten equations with Nahm pole boundary conditions at both ends of the interval and $|I| > \pi$, or I is a half interval and the solution admits a Nahm pole at the boundary. When I is non-compact the solution decays exponentially as $t \to \infty$ to the trivial connection, up to gauge equivalence.

Remark 12.19. There exist solutions (x, y) such that x is not uniformly bounded. To bound the real part of the connection one could apply a complex gauge transformation, but this would break the moment map equation in the Kapustin-Witten equations.

To study the plane autonomous system with respect to the vector field X we first note that the vector field is symmetric in the line $x = \frac{1}{2}$. This symmetry corresponds to applying the gauge transformation φ as described above. Hence, if we have a solution $\gamma = (x(t), y(t))$, then

$$\hat{\gamma}(t) = (1 - x(t), -y(t))$$

is another solution that is gauge equivalent to γ. Hence we can restrict our attention to the half plane $x < \frac{1}{2}$.

We first look for the points where X vanishes, the critical points. We see that the only critical points of this system are given by $(x, y) = (0, 0)$ and $(1, 0)$. Both of these solutions correspond to the unique flat connection on $SU(2)$: $(0, 0)$ is the trivial connection and $(1, 0)$ is the Maurer-Cartan connection described above. Because of the symmetry we can ignore the point at $(1, 0)$.

The linearisation of the vector field is given by

$$L_X = \begin{pmatrix} 2y & 2x - 1 \\ 2x - 1 & -2y \end{pmatrix}.$$

At $x = y = 0$, this has eigenvalues $1, -1$, with respective eigenvectors

$$\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ -1 \end{pmatrix}$$

and so a general solution near the trivial connection is given by

$$\begin{pmatrix} x \\ y \end{pmatrix} = Ae^{t} \begin{pmatrix} 1 \\ 1 \end{pmatrix} + Be^{-t} \begin{pmatrix} 1 \\ -1 \end{pmatrix}.$$

In the language of plane autonomous systems, this is a saddle point, and so typical solutions that get close to 0 move away, see Figure 12.1.

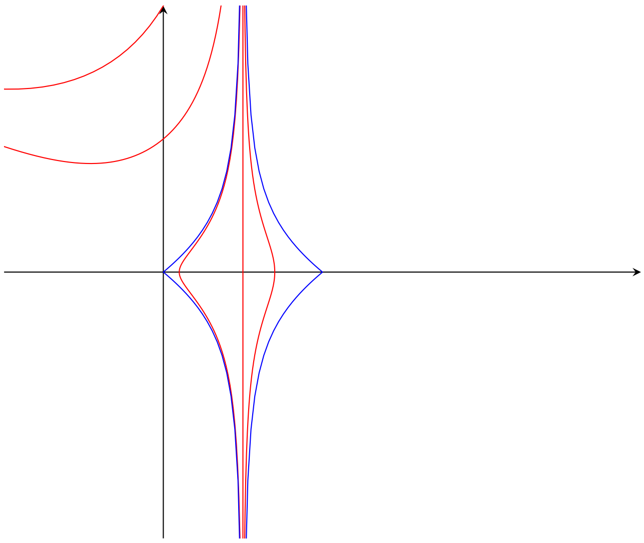

Figure 12.1: Plot of some integral curves of the vector field X, defined in (12.6). Red lines indicate that the solution is on a compact interval, blue lines indicate non compact intervals.

.

12.1.3 The line $x = \frac{1}{2}$

Looking at Figure 12.1, one can spot a solution where $x = \frac{1}{2}$ and y is some function of t. This comes from the fact that

$$\frac{dx}{dt} = 2y\left(x - \frac{1}{2}\right).$$

The connection given by $A = \frac{1}{2}\omega$ corresponds to the Levi-Civita connection on S^3 (with respect to the round metric). That is, if we identify $TSU(2)$ with our trivial $SU(2)$ vector bundle using the invariant vector fields, the connection one form for the Levi-Civita connection is given by $\frac{1}{2}\omega$. We can also see that all other solutions end up tending towards this line at at least one end of their interval, so studying the limiting behaviour of this solution helps understanding the general limiting behaviour.

To find y, we just need to solve the differential equation

$$\frac{dy}{dt} = -y^2 - \frac{1}{4}.$$

This can be integrated, resulting in

$$y = -\frac{1}{2} \tan \left(\frac{t + c}{2} \right).$$

We can set $c = 0$ as changing c is equivalent to relabelling the interval coordinate. Hence this solution is well defined (and gives a maximal solution) when $I = [-\pi, \pi]$. Moreover, we have the following.

Proposition 12.20. The solution to (12.5) with $x = \frac{1}{2}$, defined on $I = [-\pi, \pi]$, has a Nahm pole at both ends of the interval. That is, as $t \to \pm\pi$,

$$y \sim (t \mp \pi)^{-1}$$

Proof. The solution is symmetric, up to a change of sign under $t \to -t$ so we just consider $t \to \pi$. The power series expansions of $\sin \frac{t}{2}$ an $\cos \frac{t}{2}$ at $t = \pi$ are given by

$$\sin \frac{\pi + h}{2} = 1 + o(h^2), \quad \cos \frac{\pi + h}{2} = \frac{h}{2} + o(h^3).$$

The result follows. □

One can prove that the other solutions in Proposition 12.18 admit Nahm pole singularities directly. However, we can apply Theorem 12.9.

Bibliography

[AF12] Robert Adams and John Fournier. *Sobolev Spaces*. Oct. 31, 2012. ISBN: 978-0-12-044143-3. URL: https://shop.elsevier.com/books/sobolev-spaces/adams/978-0-12-044143-3 (visited on 07/08/2025).

[AM19] Mohammed Abouzaid and Ciprian Manolescu. *A Sheaf-Theoretic Model for SL(2,C) Floer Homology*. Jan. 25, 2019. DOI: 10.48550/arXiv.1708.00289. arXiv: 1708.00289 [math]. URL: http://arxiv.org/abs/1708.00289 (visited on 08/17/2024). Pre-published.

[AM90] Selman Akbulut and John D. McCarthy. *Casson's Invariant for Oriented Homology Three-Spheres: An Exposition. (MN-36)*. Princeton University Press, 1990. JSTOR: j.ctt7ztjbw. URL: https://www.jstor.org/stable/j.ctt7ztjbw (visited on 06/22/2025).

[BC16] Hans U. Boden and Cynthia L. Curtis. "The SL(2, C) Casson Invariant for Knots and the Â-polynomial". In: *Canadian Journal of Mathematics* 68.1 (Feb. 2016), pp. 3–23. ISSN: 0008-414X, 1496-4279. DOI: 10.4153/CJM-2015-025-5. URL: https://www.cambridge.org/core/journals/canadian-journal-of-mathematics/article/sl2-c-casson-invariant-for-knots-and-the-apolynomial/4E78D730A02CD329DAB5439C22E75B65 (visited on 07/03/2024).

[BD95] P. J. Braam and S. K. Donaldson. "Floer's Work on Instanton Homology, Knots and Surgery". In: *The Floer Memorial Volume*. Birkhäuser Basel, 1995, pp. 195–256. ISBN: 978-3-0348-9217-9. DOI: 10.1007/978-3-0348-9217-9_10. URL: https://link.springer.com/chapter/10.1007/978-3-0348-9217-9_10 (visited on 06/10/2025).

[BJ06] Armand Borel and Lizhin Ji. *Compactifications of Symmetric and Locally Symmetric Spaces*. Mathematics: Theory & Applications. Boston: Birkhäuser-Verlag, 2006. ISBN: 978-0-8176-3247-2. DOI: 10.1007/0-8176-4466-0. URL: http://link.springer.com/10.1007/0-8176-4466-0 (visited on 06/09/2025).

[BL81] Jean-Pierre Bourguignon and H. Blaine Lawson. “Stability and Isolation Phenomena for Yang-Mills Fields”. In: *Communications in Mathematical Physics* 79.2 (Mar. 1, 1981), pp. 189–230. ISSN: 1432-0916. DOI: 10.1007/BF01942061. URL: https://doi.org/10.1007/BF01942061 (visited on 06/26/2023).

[BZ01] Steven Boyer and Xingru Zhang. “A Proof of the Finite Filling Conjecture”. In: *Journal of Differential Geometry* 59.1 (Sept. 2001), pp. 87–176. ISSN: 0022-040X. DOI: 10.4310/jdg/1090349281. URL: https://projecteuclid.org/journals/journal-of-differential-geometry/volume-59/issue-1/A-Proof-of-the-Finite-Filling-Conjecture/10.4310/jdg/1090349281.full (visited on 06/24/2025).

[Coo+94] D. Cooper et al. “Plane Curves Associated to Character Varieties of 3-Manifolds”. In: *Inventiones mathematicae* 118.1 (1 Dec. 1, 1994), pp. 47–84. ISSN: 1432-1297. DOI: 10.1007/BF01231526. URL: https://link.springer.com/article/10.1007/BF01231526 (visited on 04/22/2025).

[Cor88] Kevin Corlette. “Flat G-Bundles with Canonical Metrics”. In: *Journal of Differential Geometry* 28.3 (Jan. 1988), pp. 361–382. ISSN: 0022-040X. DOI: 10.4310/jdg/1214442469. URL: https://projecteuclid.org/journals/journal-of-differential-geometry/volume-28/issue-3/Flat-G-bundles-with-canonical-metrics/10.4310/jdg/1214442469.full (visited on 03/01/2025).

[Cur01] Cynthia L. Curtis. “An Intersection Theory Count of the SL2(C)<math><mtext>SL</
Representations of the Fundamental Group of a 3-Manifold”. In: *Topology* 40.4 (July 1, 2001), pp. 773–787. ISSN: 0040-9383. DOI: 10.1016/S0040-9383(99)00083-X. URL: https://www.sciencedirect.com/science/article/pii/S004093839900083X (visited on 06/23/2025).

[DK97] S. K. Donaldson and P. B. Kronheimer. *The Geometry of Four-Manifolds*. Oxford Mathematical Monographs. Oxford, New York: Oxford University Press, Dec. 4, 1997. 456 pp. ISBN: 978-0-19-850269-2.

[Don02] S. K. Donaldson. *Floer Homology Groups in Yang-Mills Theory*. Cambridge Tracts in Mathematics. Cambridge: Cambridge University Press, 2002. ISBN: 978-0-521-80803-3. DOI: 10.1017/CBO9780511543098. URL: https://www.cambridge.org/core/books/floer-homology-groups-in-yangmills-theory/AC2BC9FE7B25EE8339E80C22335680F3 (visited on 01/28/2023).

[Don83] S. K. Donaldson. "An Application of Gauge Theory to Four-Dimensional Topology". In: *Journal of Differential Geometry* 18.2 (Jan. 1983), pp. 279–315. ISSN: 0022-040X. DOI: 10.4310/jdg/1214437665. URL: https://projecteuclid.org/journals/journal-of-differential-geometry/volume-18/issue-2/An-application-of-gauge-theory-to-four-dimensional-topology/10.4310/jdg/1214437665.full (visited on 01/28/2023).

[Don85] S. K. Donaldson. "Anti Self-Dual Yang-Mills Connections Over Complex Algebraic Surfaces and Stable Vector Bundles". In: *Proceedings of the London Mathematical Society* s3-50.1 (1985), pp. 1–26. ISSN: 1460-244X. DOI: 10.1112/plms/s3-50.1.1. URL: https://onlinelibrary.wiley.com/doi/abs/10.1112/plms/s3-50.1.1 (visited on 10/24/2023).

[DU77] Joseph Diestel and John Jerry Uhl. *Vector Measures*. American Mathematical Soc., June 1, 1977. 338 pp. ISBN: 978-0-8218-1515-1. Google Books: fEKZAwAAQBAJ.

[Ebe97] Patrick B. Eberlein. *Geometry of Nonpositively Curved Manifolds*. Chicago Lectures in Mathematics. Chicago, IL: University of Chicago Press, Apr. 1997. 456 pp. ISBN: 978-0-226-18198-1. URL: https://press.uchicago.edu/ucp/books/book/chicago/G/bo3633915.html (visited on 06/09/2025).

[ES64] James Eells and J. H. Sampson. "Harmonic Mappings of Riemannian Manifolds". In: *American Journal of Mathematics* 86.1 (1964), pp. 109–160. ISSN: 0002-9327. DOI: 10.2307/2373037. JSTOR: 2373037. URL: https://www.jstor.org/stable/2373037 (visited on 04/29/2025).

[Flo88] Andreas Floer. "An Instanton-Invariant for 3-Manifolds". In: *Communications in Mathematical Physics* 118.2 (Jan. 1988), pp. 215–240. ISSN: 0010-3616, 1432-0916. URL: https://projecteuclid.org/journals/communications-in-mathematical-physics/volume-118/issue-2/An-instanton-invariant-for-3-manifolds/cmp/1104161987.full (visited on 12/02/2023).

[Flo91] Andreas Floer. "Instanton Homology, Surgery and Knots". In: *Geometry of Low-Dimensional Manifolds: Volume 1: Gauge Theory and Algebraic Surfaces*. Ed. by C. B. Thomas and S. K. Donaldson. Vol. 1. London Mathematical Society Lecture Note Series. Cambridge: Cambridge University Press, 1991, pp. 97–114. ISBN: 978-0-521-39978-4. DOI: 10.1017/CBO9780511629334.008. URL: https://www.cambridge.org/core/books/geometry-of-lowdimensional-manifolds/instanton-homology-surgery-and-knots/6CF7E2BFB8B97E3BB5D7577FE274BF08 (visited on 07/04/2025).

[GS92] Michael Gromov and Richard Schoen. "Harmonic Maps into Singular Spaces and p-Adic Superrigidity for Lattices in Groups of Rank One". In: *Publications Mathématiques de l'IHÉS* 76 (1992), pp. 165–246. ISSN: 1618-1913. URL: https://www.numdam.org/item/PMIHES_1992__76__165_0/ (visited on 04/29/2025).

[GSZ23] Sudipta Ghosh, Steven Sivek, and Raphael Zentner. *Rational Homology 3-Spheres and SL(2,$\mathbb{C}$) Representations*. Version 1. Oct. 27, 2023. DOI: 10.48550/arXiv.2310.17965. arXiv: 2310.17965 [math]. URL: http://arxiv.org/abs/2310.17965 (visited on 07/04/2025). Pre-published.

[GU12] Michael Gagliardo and Karen Uhlenbeck. "Geometric Aspects of the Kapustin-Witten Equations". In: *Journal of Fixed Point Theory and Applications* 11.2 (June 2012), pp. 185–198. ISSN: 1661-7738, 1661-7746. DOI: 10.1007/s11784-012-0082-3. arXiv: 1401.7366 [math]. URL: http://arxiv.org/abs/1401.7366 (visited on 03/16/2023).

[Hat02] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002. 572 pp. ISBN: 978-0-521-79540-1. Google Books: DFGU6HxQIVUC.

[Hay13] Andriy Haydys. *Dirac Operators in Gauge Theory*. Version 2. Nov. 5, 2013. DOI: 10.48550/arXiv.1303.2971. arXiv: 1303.2971 [math]. URL: http://arxiv.org/abs/1303.2971 (visited on 01/28/2023). Pre-published.

[Hay15] Andriy Haydys. "Fukaya-Seidel Category and Gauge Theory". In: *Journal of Symplectic Geometry* 13.1 (2015), pp. 151–207. ISSN: 15275256, 15402347. DOI: 10.4310/JSG.2015.v13.n1.a5. arXiv: 1010.2353 [hep-th]. URL: http://arxiv.org/abs/1010.2353 (visited on 02/20/2023).

[Hay19] Andriy Haydys. *Introduction to Gauge Theory*. Oct. 23, 2019. DOI: 10.48550/arXiv.1910.10436. arXiv: 1910.10436 [math]. URL: http://arxiv.org/abs/1910.10436 (visited on 01/28/2023). Pre-published.

[He19] Siqi He. "A Gluing Theorem for the Kapustin-Witten Equations with a Nahm Pole". In: *Journal of Topology* 12.3 (Sept. 2019), pp. 855–915. ISSN: 1753-8416, 1753-8424. DOI: 10.1112/topo.12102. arXiv: 1707.06182 [math]. URL: http://arxiv.org/abs/1707.06182 (visited on 02/09/2023).

[JW84] Alf Jonsson and Hans Wallin. *Function Spaces on Subsets of Rn*. Harwood Academic Publishers, 1984. 240 pp. ISBN: 978-3-7186-0128-8. Google Books: dBKoAAAAIAAJ.

[KM10] P. B. Kronheimer and T. S. Mrowka. *Khovanov Homology Is an Unknot-Detector*. May 24, 2010. DOI: 10.48550/arXiv.1005.4346. arXiv: 1005.4346 [math]. URL: http://arxiv.org/abs/1005.4346 (visited on 01/28/2023). Pre-published.

[KM11] P. B. Kronheimer and T. S. Mrowka. "Knot Homology Groups from Instantons". In: *Journal of Topology* 4.4 (Dec. 2011), pp. 835–918. ISSN: 17538416. DOI: 10.1112/jtopol/jtr024. arXiv: 0806.1053 [math]. URL: http://arxiv.org/abs/0806.1053 (visited on 02/10/2023).

[KW07] Anton Kapustin and Edward Witten. *Electric-Magnetic Duality And The Geometric Langlands Program*. Jan. 16, 2007. DOI: 10.48550/arXiv.hep-th/0604151. arXiv: hep-th/0604151. URL: http://arxiv.org/abs/hep-th/0604151 (visited on 01/28/2023). Pre-published.

[Muñ25] Juan Muñoz-Echániz. *Counting $SL(2, \mathbb{C})$ Connections on Seifert-fibered Spaces*. Mar. 20, 2025. DOI: 10.48550/arXiv.2503.16370. arXiv: 2503.16370 [math]. URL: http://arxiv.org/abs/2503.16370 (visited on 03/21/2025). Pre-published.

[MW13] Rafe Mazzeo and Edward Witten. *The Nahm Pole Boundary Condition*. Dec. 3, 2013. DOI: 10.48550/arXiv.1311.3167. arXiv: 1311.3167 [hep-th]. URL: http://arxiv.org/abs/1311.3167 (visited on 01/28/2023). Pre-published.

[MW18] Rafe Mazzeo and Edward Witten. *The KW Equations and the Nahm Pole Boundary Condition with Knots*. Feb. 27, 2018. DOI: 10.48550/arXiv.1712.00835. arXiv: 1712.00835 [hep-th]. URL: http://arxiv.org/abs/1712.00835 (visited on 01/28/2023). Pre-published.

[Nic11] Liviu Nicolaescu. *An Invitation to Morse Theory*. Universitext. New York, NY: Springer, 2011. ISBN: 978-1-4614-1104-8 978-1-4614-1105-5. DOI: 10.1007/978-1-4614-1105-5. URL: https://link.springer.com/10.1007/978-1-4614-1105-5 (visited on 06/16/2025).

[Sch00] Martin Scharlemann. *Heegaard Splittings of Compact 3-Manifolds*. Version 1. July 24, 2000. DOI: 10.48550/arXiv.math/0007144. arXiv: math/0007144. URL: http://arxiv.org/abs/math/0007144 (visited on 07/07/2025). Pre-published.

[Sch95] Günter Schwarz. *Hodge Decomposition—A Method for Solving Boundary Value Problems*. Vol. 1607. Lecture Notes in Mathematics. Berlin, Heidelberg: Springer, 1995. ISBN: 978-3-540-60016-9 978-3-540-49403-4. DOI: 10.1007/BFb0095978. URL: http://link.springer.com/10.1007/BFb0095978 (visited on 06/10/2025).

[Sei08] Paul Seidel. *Fukaya Categories and Picard–Lefschetz Theory*. June 26, 2008. DOI: 10.4171/063. URL: https://ems.press/books/zlam/2 (visited on 01/28/2023).

[Tau13] Clifford Henry Taubes. *Compactness Theorems for SL(2;C) Generalizations of the 4-Dimensional Anti-Self Dual Equations*. arXiv.org. July 24, 2013. DOI: 10.48550/arXiv.1307.6447. URL: https://arxiv.org/abs/1307.6447v5 (visited on 02/08/2023).

[Tau15] Clifford Henry Taubes. *$PSL(2;\mathbb{C})$ Connections on 3-Manifolds with L2 Bounds on Curvature*. Aug. 4, 2015. DOI: 10.48550/arXiv.1205.0514. arXiv: 1205.0514 [math]. URL: http://arxiv.org/abs/1205.0514 (visited on 02/09/2023). Pre-published.

[Tau90] Clifford Henry Taubes. "Casson's Invariant and Gauge Theory". In: *Journal of Differential Geometry* 31.2 (Jan. 1990), pp. 547–599. ISSN: 0022-040X. DOI: 10.4310/jdg/1214444327. URL: https://projecteuclid.org/journals/journal-of-differential-geometry/volume-31/issue-2/Cassons-invariant-and-gauge-theory/10.4310/jdg/1214444327.full (visited on 01/28/2023).

[Uhl82] Karen K. Uhlenbeck. “Connections with L^p Bounds on Curvature”. In: *Communications in Mathematical Physics* 83.1 (Jan. 1982), pp. 31–42. ISSN: 0010-3616, 1432-0916. URL: https://projecteuclid.org/journals/communications-in-mathematical-physics/volume-83/issue-1/Connections-with-Lp-bounds-on-curvature/cmp/1103920743.full (visited on 05/21/2023).

[Wan22] Donghao Wang. *The Complex Gradient Flow Equation and Seidel's Spectral Sequence*. Sept. 6, 2022. DOI: 10.48550/arXiv.2209.02810. arXiv: 2209.02810 [math]. URL: http://arxiv.org/abs/2209.02810 (visited on 07/07/2025). Pre-published.

[Wit10] Edward Witten. *Analytic Continuation Of Chern-Simons Theory*. Aug. 28, 2010. DOI: 10.48550/arXiv.1001.2933. arXiv: 1001.2933 [hep-th]. URL: http://arxiv.org/abs/1001.2933 (visited on 01/28/2023). Pre-published.

[Wit11] Edward Witten. *Fivebranes and Knots*. Aug. 11, 2011. DOI: 10.48550/arXiv.1101.3216. arXiv: 1101.3216 [hep-th]. URL: http://arxiv.org/abs/1101.3216 (visited on 01/28/2023). Pre-published.

[Wit88] Edward Witten. “Topological Quantum Field Theory”. In: *Communications in Mathematical Physics* 117.3 (Sept. 1988), pp. 353–386. ISSN: 0010-3616, 1432-0916. DOI: 10.1007/BF01223371. URL: http://link.springer.com/10.1007/BF01223371 (visited on 01/28/2023).

[Wit89] Edward Witten. “Quantum Field Theory and the Jones Polynomial”. In: *Communications in Mathematical Physics* 121.3 (Jan. 1989), pp. 351–399. ISSN: 0010-3616, 1432-0916. URL: https://projecteuclid.org/journals/communications-in-mathematical-physics/volume-121/issue-3/Quantum-field-theory-and-the-Jones-polynomial/cmp/1104178138.full (visited on 07/31/2024).