Some boundedness properties of solutions to the Vafa–Witten equations on closed four-manifolds

Yuuji Tanaka

Abstract

We consider a set of gauge-theoretic equations on closed oriented four-manifolds, which was introduced by Vafa and Witten. The equations involve a triple consisting of a connection and extra fields associated to a principal bundle over a closed oriented four-manifold. They are similar to Hitchin's equations over compact Riemann surfaces, and as part of the resemblance, there is no L^2-bound on the curvature without an L^2-bound on the extra fields. In this article, however, we observe that under the particular circumstance where the curvature does not become concentrated and the limiting connection is not locally reducible, one obtains an L^2-bound on the extra fields.

1 Introduction

In [VW], Vafa and Witten introduced a set of gauge-theoretic equations on a four-manifold, motivated by the study of S-duality conjecture in a topologically twisted $N = 4$ supersymmetric Yang–Mills theory. The moduli space of solutions to the equations is expected to produce a possibly new invariant of some kind. Physicists believe that the partition functions which arise are the generating functions of the "Euler characteristic" of the ASD instanton moduli spaces, and the S-duality conjecture predicts that the partition functions satisfy an enchanting modularity property [VW]. The equations were also considered by Haydys [Ha] and Witten [Wi] from a different point of view. They suggest a programme of "categorification" of the Khovanov homology by using a five-dimensional gauge theory, where the Vafa–Witten equations are a dimensional reduction of their five-dimensional gauge-theoretic equation, and the solution spaces for the Vafa–Witten equations are expected to be used in constructing a Casson–Floer type theory.

The techniques developed by Taubes [T1], [T2], [T3] and [T4] have direct analogues in the Vafa–Witten case [T5], and we use them in this article.

The equations we consider involve a triple consisting of a connection and other extra fields associated to a principal bundle over a four-manifold. Let X be a closed, oriented, smooth Riemannian four-manifold with Riemannian metric g, and let $P \to X$ be a principal G-bundle over X with G being a compact Lie group. We denote by $\mathcal{A}_P$ the set of all connections of P, and by $\Omega^+(X, \mathfrak{g}_P)$ the set of self-dual two-forms valued in the adjoint bundle $\mathfrak{g}_P$ of P. We consider the following equations for a triple $(A, B, \Gamma) \in \mathcal{A}_P \times \Omega^+(X, \mathfrak{g}_P) \times \Omega^0(X, \mathfrak{g}_P)$,

$$d_A \Gamma + d_A^* B = 0,$$

$$F_A^+ + \frac{1}{8}[B.B] + \frac{1}{2}[B, \Gamma] = 0,$$

where $[B.B] \in \Omega^+(X, \mathfrak{g}_P)$ is defined in Section 2. We call these equations the *Vafa–Witten equations*. We impose the additional simplifying assumption $\Gamma = 0$, which is not a significant restriction. As explained in Section 2 later, if (A, B, Γ) is a solution, then so is $(A, B, 0)$. Moreover, the field Γ vanishes if the connection is irreducible. Thus we consider only solution (A, B) for which Γ is assumed to vanish.

Mares studied analytic aspects of these equations in his Ph.D thesis [M]. As in the case of Hitchin's equation [Hi], we do not have an L^2-bound on the curvature F_A of a connection A which satisfies the equations in general. In the case where the L^2-norm of B is bounded, the following identity leads to a bound on the L^2-norm of F_A for solutions.

$$\frac{1}{2}||d_A^* B||_{L^2}^2 + ||F_A^+ + \frac{1}{8}[B.B]||_{L^2}^2$$

$$= \int_X \left(\frac{1}{12} s |B|^2 - \frac{1}{2} W^+ \langle B \odot B \rangle \right)$$

$$+ \frac{1}{2}||F_A||_{L^2}^2 + \frac{1}{4}||\nabla_A B||_{L^2}^2 + \frac{1}{16}||[B.B]||_{L^2}^2 + \frac{1}{2} \int_X \langle F_A \wedge F_A \rangle,$$

where $\odot$ denotes some bilinear operation on $\Omega^+(X, \mathfrak{g}_P) \otimes \Omega^+(X, \mathfrak{g}_P)$ (see [M, §B.4] for more details), s is the scalar curvature of the metric, and W^+ is the self-dual part of the Weyl curvature of the metric. The last term on the right hand side of the identity is a topological term. When (A, B) is a solution, the left hand side vanishes. If L^2-norm of B is bounded, then the integral in the middle line is bounded. Hence, using positivity of the middle two terms on the last line, one gets an upper bound on $||F_A||_{L^2}^2$.

Mares obtained the following compactness theorem, under the assumption that the L^2 norm of B is bounded.

Theorem 1.1 ([M]). *Let $\{(A_i, B_i)\}$ be a sequence of solutions to the Vafa–Witten equations with $\int_X |B_i|^2 \leq C$. Then, by taking a subsequence of $\{(A_i, B_i)\}$ if necessary, $\{(A_i, B_i)\}$ converges in C^∞ on compact subsets to a smooth solution (A, B) of the Vafa–Witten equations outside a finite set $\{x_1, \ldots, x_k\} \subset X$ after some gauge transformations on $X \setminus \{x_1, \ldots, x_k\}$.*

In this article, we observe that there is a particular circumstance in which one obtains an L^2-bound on the extra field B. The argument implies a non-constructive L^2-bound on $B \in \Omega^+(X, \mathfrak{g}_P)$ from the assumption that a limiting connection A is not locally reducible. Here, we say that a connection A on a principal $SU(2)$ or $SO(3)$ bundle P is *locally reducible* if the vector bundle $\mathfrak{g}_P$ has a 1-dimensional subbundle that is A-covariantly constant (see Definition 2.1 in Section 2).

Remark 1.2. Note that a connection on a principal $SU(2)$ or $SO(3)$ bundle P being locally reducible is the same as being honestly reducible if X is simply-connected.

In order to properly state the main observation of this paper, we require a technical assumption. We denote by δ the injectivity radius of X, so that geodesic balls $B_r(x)$ of radius r and centered at x are well-defined whenever $r < \delta$.

For a sequence of connections $\{A_i\}$ on P, we define

$$S(\{A_i\}) := \bigcap_{r \in (0,\delta)} \{x \in X \mid \liminf_{i \to \infty} \int_{B_r(x)} |F_{A_i}|^2 \, d\mu_g \geq \varepsilon_\diamond^2\}, \qquad (1.1)$$

where $\varepsilon_\diamond > 0$ is a positive constant (it is, in particular, independent of $\{A_i\}$ and less than the constant of Uhlenbeck's gauge-fixing lemma). This number $\varepsilon_\diamond$ is defined as follows: Let $\kappa > 1$ denote the number that appears in Proposition 4.4. Let ε and c denote the numbers appearing in Appendix A.2's equation (A.7). Set $\varepsilon_\diamond = \min(\kappa^{-1}, \varepsilon c^{-1})$. This set $S(\{A_i\})$ describes the singular set of a sequence of connections $\{A_i\}$. In particular, if $x \notin S$ then there is a small ball around x for which some subsequence of connections has curvature below the Uhlenbeck gauge fixing constant. Our technical assumption is that

$$S(\{A_i\}) = \emptyset.$$

Consequently, Uhlenbeck's theorem gurantees that there is some ball around every $x \in X$, and a subsequence such that the connections can be put into Coulomb gauge on this ball, with a uniform L^2-bound on the connection.

The main observation of this paper can be stated as follows.

Theorem 1.3. *Let X be a closed, oriented, smooth Riemannian four-manifold with Riemannian metric g, and let $P \to X$ be a principal G-bundle over X with G being $SU(2)$ or $SO(3)$. Let $\{(A_i, B_i)\}_{i \in \mathbb{N}}$ be a sequence of solutions to the Vafa–Witten equations with $S(\{A_i\})$ being empty. Then, there exist a subsequence $\Xi \subset \mathbb{N}$ and a sequence of gauge transformations $\{g_i\}_{i \in \Xi}$ such that $\{g_i(A_i)\}_{i \in \Xi}$ converges weakly in the L_1^2-topology. If the limit is not locally reducible, then there exists a positive number C such that $\int_X |B_i|^2 d\mu_g \leq C$ for all $i \in \Xi$, and $\{(g_i(A_i), g_i(B_i))\}_{i \in \Xi}$ converges in the C^∞-topology to a pair that obeys the Vafa–Witten equations.*

A modification of recent work of Taubes [T2] and [T3] is used to prove this theorem.

As a particular case of Theorem 1.3, we have an L^2-bound on the extra fields in the "fibre" direction at a connection A_0 which is not locally reducible. Namely,

Corollary 1.4. *Let X be a closed, oriented, smooth Riemannian four-manifold with Riemannian metric g, and let $P \to X$ be a principal G-bundle over X with G being $SU(2)$ or $SO(3)$. Then, for any sequence of solutions $\{(A_0, B_i)\}_{i \in \mathbb{N}}$ of the Vafa–Witten equations with A_0 being not locally reducible, there exist a subsequence $\Xi \subset \mathbb{N}$ and a positive constant $C > 0$ such that $\int_X |B_i|^2 d\mu_g \leq C$ for all $i \in \Xi$.*

A proof of Theorem 1.3 is given in Section 3. In Section 2, we describe some properties of the equations such as the linearization, and also mention some remarks on reducible solutions to the equations.

In the rest of this section, let us mention an observation which led to a formulation of Theorem 1.3. To make the argument rather simple, we assume that X is simply-connected so that a connection being locally reducible is in fact honestly reducible. The point lies in relating the L^2-bound of a sequence of the extra fields $\{B_i\}$ to the irreducibility of the limiting connection.

Let $\{(A_i, B_i)\}_{i \in \mathbb{N}}$ be a sequence of solutions to the Vafa–Witten equations (2.4) and (2.5), and suppose that $||B_i||_{L^2} \to \infty$ as $i \to \infty$. Put $r_i := ||B_i||_{L^2}$ and $\beta_i := B_i/r_i$ for $i \in \mathbb{N}$. Then (A_i, β_i) satisfies $d_{A_i}^* \beta_i = 0$, $F_{A_i}^+ + \frac{1}{8} r_i^2 [\beta_i . \beta_i] = 0$ and $||\beta_i||_{L^2} = 1$.

The technicalities of convergence of such a sequence are very delicate, and occupy the bulk of this paper. But for now, we assume that there is some hypothesis which ensures the convergence of $\{(A_i, \beta_i)\}$ to $\{(A_0, \beta)\}$ which satisfies $d_{A_0}^* \beta = 0$, $[\beta . \beta] = 0$ and $||\beta||_{L^2} = 1$. After some further

assumptions, this would imply that the connection A_0 is reducible by the following argument.

Firstly, we recall a notion of rank of a section $B \in \Omega^+(X, \mathfrak{g}_P)$.

Definition 1.5. A $\mathfrak{g}_P$-valued self-dual two-form $B \in \Omega^+(X, \mathfrak{g}_P)$ is said to be of *rank r* if, when considered as a section of $\mathrm{Hom}(\Lambda^{+,*}, \mathfrak{g}_P)$, $B(x)$ has rank less than or equal to r at every point $x \in X$ with equality at some point.

If the structure group of the principal bundle P is either $SU(2)$ or $SO(3)$, then the possibilities for the rank of B are $0, 1, 2$ or 3. We next recall the following from [M, § 4.11].

Lemma 1.6 ([M]). *Let $P \to X$ be a principal G-bundle over a simply-connected Riemannian four-manifold X with structure group G being either $SU(2)$ or $SO(3)$. Let $B \in \Omega^+(X, \mathfrak{g}_P)$. If $[B.B] = 0$. then the rank of B is at most one.*

Then, we invoke Lemma 4.3.25 in [DK].

Lemma 1.7 ([DK] Lem. 4.3.25). *Let X be a simply-connected oriented Riemannian four-manifold, and let $P \to X$ be a principal G-bundle with structure group G being either $SU(2)$ or $SO(3)$. Let ϕ in $\Omega^+(X, \mathfrak{g}_P)$. Assume that $d_A \phi = 0$. Then on any simply-connected open subset in X, where ϕ has rank 1 and is nowhere vanishing, the connection A is reducible.*

As $\beta \in \Omega^+(X, \mathfrak{g}_P)$, it follows from $d_{A_0}^* \beta = 0$ that $d_{A_0} \beta = 0$. From Lemma 1.6, β is at most rank one. Thus, by Lemma 1.7, A_0 is reducible on an open subset in X. We now further assume that $r_i^2 [\beta_i.\beta_i] \to 0$ so that A_0 is anti-self-dual. We then invoke the following.

Lemma 1.8 ([DK], Lem. 4.3.21). *If A is an irreducible $SU(2)$ of $SO(3)$ anti-self-dual connection on a bundle E over a simply-connected four manifold X, then the restriction of A to any non-empty open set in X is also irreducible.*

From this, we deduce that A_0 is reducible on X under the assumptions.

2   The Vafa–Witten equations on four-manifolds

In this section, we describe the precise form of the equations and their linearization, and then mention some remarks on reducible solutions.

Self-dual two-forms.   For an oriented n-dimensional Riemannian manifold X with Riemannian metric g, there is an operator $* : \Lambda^p \to \Lambda^{n-p}$, called the *Hodge $*$-operator*, defined by $\alpha \wedge *\beta = (\alpha, \beta)d\mu_g$, where $(,)$ denotes the natural metric on the forms defined by the Riemannian metric, and $d\mu_g$ is the Riemannian volume element.

Setting $n = 4$, the Hodge star operator on the space of two forms on X induces the following splitting into ± 1 eigenspaces:

$$\Lambda^2(X) = \Lambda^+(X) \oplus \Lambda^-(X).$$

Accordingly, the space of $\mathfrak{g}_P$-valued two-forms $\Omega^2(X, \mathfrak{g}_P)$, where $P \to X$ is a principal G-bundle over X, splits as

$$\Omega^2(X, \mathfrak{g}_P) = \Omega^+(X, \mathfrak{g}_P) \oplus \Omega^-(X, \mathfrak{g}_P).$$

The splitting of the space of two-forms can be viewed as that corresponding to the isomorphism $\mathfrak{so}(4) \cong \mathfrak{so}(3) \oplus \mathfrak{so}(3)$ through the embedding $\Lambda^2 T_x^* X \subset \mathrm{End}(T_x^* X)$ as the subalgebra of skew-symmetric endomorphisms. Hence Λ^+ is equipped with a pointwise Lie-algebraic structure with the bracket denoted by $[\cdot, \cdot]_{\Lambda^+}$. Using this together with that of $\mathfrak{g}_P$, we define a bilinear map $[\cdot \cdot] : (\Lambda^+ \otimes \mathfrak{g}_P) \otimes (\Lambda^+ \otimes \mathfrak{g}_P) \to \Lambda^+ \otimes \mathfrak{g}_P$ by $\frac{1}{2}[\cdot, \cdot]_{\Lambda^+} \otimes [\cdot, \cdot]_{\mathfrak{g}_P}$.

Equations.   Let X be a closed, oriented, smooth four-manifold, and let $P \to X$ be a principal G-bundle over X with G being a compact Lie group.

We denote by $\mathcal{A}_P$ the set of all connections of P, by $\Omega^+(X, \mathfrak{g}_P)$ the set of self-dual two-forms valued in the adjoint bundle $\mathfrak{g}_P$ of P.

We consider the following equations for a triple $(A, B, \Gamma) \in \mathcal{A}_P \times \Omega^+(X, \mathfrak{g}_P) \times \Omega^0(X, \mathfrak{g}_P)$,

$$d_A\Gamma + d_A^* B = 0, \tag{2.1}$$

$$F_A^+ + \frac{1}{8}[B.B] + \frac{1}{2}[B, \Gamma] = 0. \tag{2.2}$$

We call (2.1) and (2.2) the *Vafa–Witten equations*.

The left action of gauge group $\mathcal{G}_P$ is given by $(u(A), \mathrm{ad}(u)B, \mathrm{ad}(u)\Gamma)$ for $u \in \mathcal{G}_P$ and $(A, B, \Gamma) \in \mathcal{A}_P \times \Omega^+(X, \mathfrak{g}_P) \times \Omega^0(X, \mathfrak{g}_P)$.

Linearization. Following [M, §3.2.1], let $(A, B, \Gamma) \in \mathcal{A}_P \times \Omega^+(X, \mathfrak{g}_P) \times \Omega^0(X, \mathfrak{g}_P)$ be a solution to the Vafa–Witten equations (2.1) and (2.2). Then, the infinitesimal deformation at (A, B, Γ) is given by the following.

$$0 \longrightarrow \Omega^0(X, \mathfrak{g}_P) \xrightarrow{d_{(A,B,\Gamma)}} \Omega^1(X, \mathfrak{g}_P) \oplus \Omega^+(X, \mathfrak{g}_P) \oplus \Omega^0(X, \mathfrak{g}_P)$$

$$\xrightarrow{d_{(A,B,\Gamma)}^+} \Omega^+(X, \mathfrak{g}_P) \oplus \Omega^1(X, \mathfrak{g}_P) \longrightarrow 0,$$

$$\tag{2.3}$$

where the maps $d_{(A,B,\Gamma)}$ and $d_{(A,B,\Gamma)}^+$ are defined for $\xi \in \Omega^0(X, \mathfrak{g}_P)$ and $(\alpha, \beta, \gamma) \in \Omega^1(X, \mathfrak{g}_P) \oplus \Omega^+(X, \mathfrak{g}_P) \oplus \Omega^0(X, \mathfrak{g}_P)$ by

$$d_{(A,B,\Gamma)}(\xi) := (-d_A\xi, [\xi, B], [\xi, \Gamma]),$$

$$d_{(A,B,\Gamma)}^+(\alpha, \beta, \gamma)$$
$$:= \left(d_A^+\alpha + \frac{1}{2}[B, \gamma] + \frac{1}{4}[B.\beta] - \frac{1}{2}[\Gamma, \beta], d_A\Gamma - [\Gamma, \alpha] - [B, \alpha] + d_A^*\beta \right).$$

The principal symbol of the operator $(d_{(A,B,\Gamma)}^+, d_{(A,B,\Gamma)}^*)$ is self-adjoint, so the index of the complex (2.3) is always zero.

Local reducibility.

Definition 2.1. A connection is *locally reducible* if there is an open cover of X such that on each of the open subsets, there is a non-zero, covariantly constant section of $\mathfrak{g}_P$.

In the case considered here, the span of these local covariantly constant sections patch together to define a covariantly constant line subbundle $\mathcal{I}$ in $\mathfrak{g}_P$ on the whole of X. This implies in particular that the holonomy of the connection lies in an $O(2)$ subgroup of G, which we are assuming $SO(3)$ or $SU(2)$. Appendix B to this paper reviews the relation between holonomy and constant sections of associated bundles. We also refer to the definition of irreducible connection explained in Appendix B.

The equations in the irreducible semisimple case. As described in [M, Th.2.1.1], if X is closed, the equations (2.1) and (2.2) are equivalent to the following.

$$d_A \Gamma = 0,\ d_A^* B = 0,$$

$$F_A^+ + \frac{1}{8}[B.B] = 0,\ [B, \Gamma] = 0.$$

If (A, B, Γ) is a solution, then it is clear from these equations that $(A, B, 0)$ is also a solution. Furthermore, $\Gamma \in \Omega^0(X, \mathfrak{g}_P)$ lies in the kernel of $d_A : \Omega^0(X, \mathfrak{g}_P) \to \Omega^1(X, \mathfrak{g}_P)$, which is the infinitesimal stabilizer of A. If G is semisimple, then the centre of G is discrete, and so the infinitesimal stabilizer of any irreducible A is zero. In this case, Γ vanishes. Since we are mainly interested in the irreducible semisimple case, we set $\Gamma = 0$, namely we consider only the following equations in the rest of this article.

$$d_A^* B = 0, \tag{2.4}$$

$$F_A^+ + \frac{1}{8}[B.B] = 0. \tag{2.5}$$

We end this section by noting the following. If A is a locally reducible connection on an $SU(2)$ or $SO(3)$ bundle over X, then the adjoint bundle $\mathfrak{g}_P$ decomposes as $\mathfrak{g}_P = \mathcal{I} \oplus C$, where $\mathcal{I}$ is a real line bundle fixed by A, and C is the complementary $O(2)$ subbundle, as explained in Appendix B. Hence, there are corresponding decompositions of $B = B_{\mathcal{I}} + B_C$ and $F_A^+ = (F_A^+)_{\mathcal{I}} + (F_A^+)_C$, where the subscript $\mathcal{I}$ denotes projection onto $\Omega^+(X, \mathcal{I})$ and the subscript C denotes projection onto $\Omega^+(X, C)$. The connection A is trivial on the $\mathcal{I}$ part, so $(F_A^+)_{\mathcal{I}} = 0$. According to [M, §4.1.1], this implies that B_C has rank at most one, and hence B has rank at most 2. In this case, $[B.B]$ and hence F_A^+ have rank at most one.

3   Proof of Main Theorem 1.3

In this section, we show that Theorem 1.3 is a corollary of the following theorem.

Theorem 3.1 ([T5]). *Let $\{(A_i, B_i)\}_{i \in \mathbb{N}}$ be a sequence of solutions to the Vafa–Witten equations with $S(\{A_i\})$ being empty. Put $r_i := ||B_i||_{L^2}$ for $i \in \mathbb{N}$, and assume that $\{r_i\}_{i \in \mathbb{N}}$ has no bounded subsequence. Then there exists a subsequence $\Xi \subset \mathbb{N}$ and a sequence of gauge transformation $\{g_i\}_{i \in \Xi}$ such that $\{g_i(A_i)\}_{i \in \Xi}$ converges in the weak L_1^2-topology on X to a limit that is anti-self-dual and locally reducible.*

The author learned from Clifford Taubes [T5] that one can prove Theorem 3.1 by adjusting some of his results in [T2] to the situation here and also by using Theorem 1.2 of [T3].

Proof of Theorem 1.3, assuming Theorem 3.1. Suppose the hypotheses of Theorem 1.3 are satisfied, namely that the limit connection is not locally reducible. We deduce firstly the L^2 bound on B_i, and secondly the C^∞ convergence. The first proceeds by contradiction, assuming that the bound claimed in Theorem 1.3 does not hold. Then the hypotheses of Theorem 3.1 are satisfied, namely $\{r_i\}$ has no bounded subsequence. The conclusion of Theorem 3.1 that the limit is locally reducible contradicts the hypothesis that the limit is not locally reducible in Theorem 1.3. Therefore the bound claimed in Theorem 1.3 must hold. For the second claim in Theorem 1.3 of C^∞ convergence, appeal to either Theorem 1.1 or a modification of the first bullet of Proposition 2.1 in [T2] while invoking that $S(\{A_i\})$ is empty. This implies the C^∞-convergence of $\{(g_i(A_i), g_i(B_i))\}_{i \in \Xi}$ with a limit satisfying the Vafa–Witten equations (2.4) and (2.5). $\qquad\square$

To complete the proof of Theorem 1.3, it remains only to prove Theorem 3.1.

4   Proof of Theorem 3.1 *

In this section, we give the proof of Theorem 3.1. The proof must show that a limiting connection exists (this connection is denoted by A_0). This is the content of Proposition 4.1, whose proof is given in Appendix A.

Proposition 4.1 (Weak L_1^2 compactness of non-concentrating connections modulo gauge [T5]). *Let $\{A_i\}_{i \in \mathbb{N}}$ denote a sequence of smooth connections*

---

*The details of the proof presented here were conveyed to the author by Clifford Taubes.

on the principal $SU(2)$ bundle P with $S(\{A_i\}) = \emptyset$. Then, there exists a subsequence $\Pi \subset \mathbb{N}$ and a sequence of automorphisms $\{g_i\}_{i \in \Pi}$ such that the sequence $\{g_i(A_i)\}_{i \in \Pi}$ converges weakly in the L_1^2 topology to a limit, which is a Sobolev class L_1^2 connection on P, denoted by A_0.

By way of a reminder, to say that $\{g_i(A_i\}_{i \in \Pi}$ converges weakly in the L_1^2 topology to A_0 means the following: Fix a smooth connection on P to be denoted by $\hat{A}$. Then write each $i \in \Pi$ version of $g_i(A_i)$ as $\hat{A} + a_i$ with a_i being a section of $T^*M \otimes \mathfrak{g}_P$ with $\mathfrak{g}_P$ denoting here (and subsequently) the Lie algebra of $SU(2)$. The weak L_1^2 convergence assertion requires that the elements in the sequence $\{a_i\}_{i \in \Pi}$ have uniformly bounded L_1^2 norm and that this sequence converges weakly in the Hilbert space of L_1^2 sections of $T^*M \otimes \mathfrak{g}_P$. Let a_0 denote the limit L_1^2 section. Then $A_0 = \hat{A} + a_0$.

With A_0 in hand, it remains to prove that A_0 has anti-self-dual curvature tensor and that it is locally reducible. Local reducibility means, as per Definition 2.1, that there is a rank one, real line subbundle $\mathcal{I} \subset \mathfrak{g}_P$, that is preserved by A_0-parallel transport. An equivalent formulation of local reducibility is as follows: There is a real line bundle $\mathcal{I}$ and a nonzero A_0-covariantly constant section (denoted by σ_0) of $\operatorname{Hom}(\mathcal{I}, \mathfrak{g}_P)$. The existence of $\mathcal{I}$ and σ_0 is asserted by Proposition 4.6. The intermediate Propositions 4.2–4.4 and Theorem 4.5 are needed for the construction of $\mathcal{I}$ and σ_0.

The task of proving local reducibility starts by considering the limit of the sequence $\{\beta_i := B_i/r_i\}_{i \in \Pi}$ when $\{r_i\}_{i \in \Pi}$ has no bounded subsequence. We show that away from the zero locus Z of the limit, and up to gauge, a subsequence of β_i converges to a tensor of rank one. The $\mathfrak{g}_P$ component of this tensor yields a section σ_Δ of $\operatorname{Hom}(\mathcal{I}, \mathfrak{g}_P)$ defined over $X \setminus Z$. Finally we show that σ_Δ extends over Z to give the section σ_0 which is defined over all of X and is A_0-covariantly constant, demonstrating that A_0 is locally reducible.

Firstly we have the following adjustment of Proposition 2.1 in [T2] to the Vafa–Witten case.

Proposition 4.2 ([T5])**.** *Assume that the sequence $\{r_n\}_{n \in \mathbb{N}}$ has no bounded subsequence. Then there exists a subsequence $\Lambda \subset \mathbb{N}$ such that the following items a)–e) listed below hold.*

a) *The sequences $\{\int_X (|d|\beta_i||^2 + |\beta_i|^2)\}_{i \in \Lambda}$ and $\{\sup_X |\beta_i|\}_{i \in \Lambda}$ are bounded by a number that depends only on the geometric data.*

b) *The sequence $\{|\beta_i|\}_{i \in \Lambda}$ converges weakly in the L_1^2-topology and strongly in all $p < \infty$ version of the L^p-topology. The limit function is de-*

noted suggestively by $|\hat{\beta}_\diamond|$, this being a priori an L^∞-function, despite the fact that $\hat{\beta}_\diamond$ remains undefined as a section of any bundle. The function $|\hat{\beta}_\diamond|$ is defined at each point in X by the rule whereby $|\hat{\beta}_\diamond|(p) = \limsup_{i \in \Lambda} |\beta_i|(p)$ for each $p \in X$.

c) The sequence $\{\langle \beta_i \otimes \beta_i \rangle\}_{i \in \Lambda}$ converges strongly in any $q < \infty$ version of L^q-topology on the space of sections of $\Lambda^+ \otimes \Lambda^+$. The limit section is denoted suggestively by $\langle \hat{\beta}_\diamond \otimes \hat{\beta}_\diamond \rangle$ and its trace is the function $|\hat{\beta}_\diamond|^2$.

d) Use f to denote any given C^0-function.

 i) The sequence $\{r_i^{-2} \int_X f |F_{A_i}|^2\}_{i \in \Lambda}$ converges.
 ii) The sequences $\{\int_X f |\nabla_{A_i} \beta_i|^2\}_{i \in \Lambda}$ and $\{2r_i^2 \int_X f |[\beta_i . \beta_i]|^2\}_{i \in \Lambda}$ converge. The limit of the first sequence is denoted by $Q_{\nabla,f}$ and that of the second by $Q_{\wedge,f}$. These are such that

$$\frac{1}{2} \int_X (d^* d\, f) |\hat{\beta}_\diamond|^2 + Q_{\nabla,f} + Q_{\wedge,f} + \int_X f\, \mathcal{R}\left(\langle \hat{\beta}_\diamond \otimes \hat{\beta}_\diamond \rangle \right) = 0,$$

 where $\mathcal{R}$ is an endomorphism associated to the Riemann curvature.

e) Fix $p \in X$ and let G_p denote the Green's function with pole at p for the operator $d^* d + 1$. The sequence $\{\int_X G_p(|\nabla_{A_i} \beta_i|^2 + 2r_i^2 |[\beta_i . \beta_i]|^2)\}_{i \in \Lambda}$ is bounded by a number that depends only on the geometric data. Let $Q_{\diamond,p}$ denote the lim-inf of this sequence. The function $|\hat{\beta}_\diamond|^2$ obeys the equation

$$\frac{1}{2} |\hat{\beta}_\diamond|^2 + Q_{\diamond,p} = -\int_X G_p \left(\frac{1}{2} |\hat{\beta}_\diamond|^2 - \mathcal{R}(\langle \hat{\beta}_\diamond \otimes \hat{\beta}_\diamond \rangle) \right).$$

The following is an analogue of Proposition 6.1 in [T2], but under the assumption that $S(\{A_i\})$ is empty, one does not need to follow the full argument of Section 5 in [T2].

Proposition 4.3 ([T5]). *The function $|\hat{\beta}_\diamond|$ of Proposition 4.2 is continuous on X and smooth at points where it is positive. Moreover, the sequence $\{|\beta_i|\}_{i \in \Lambda}$ converges to $|\hat{\beta}_\diamond|$ in the C^0-topology on X.*

Section 3 of [T2] introduces two positive numbers that measure the size of the curvature near any given point, these denoted by $r_{c\wedge}$ and r_{cF}. Much of the analysis in Sections 3–6 in [T2] is meant to deal with the a priori fact

that these numbers can be very small if the curvature becomes concentrated near the chosen point. However, in the present context, the assumption that $S(\{A_i\}) = \emptyset$ in Theorems 1.3 and 3.1 imply that both radii are bounded away from zero. Consequently, the assertions in Proposition 7.2 and Proposition 8.1 apply here. They are summarized as follows.

Proposition 4.4 ([T5]). *There exists $\kappa > 1$ with the following significance: Let $\{(A_i, \beta_i)\}_{i \in \mathbb{N}}$ be as described in Proposition 4.2, let $\Lambda \subset \mathbb{N}$ be the subsequence from that proposition, and let δ denote the injectivity radius of X. Suppose that there exists $\rho \in (0, \delta)$ such that*

$$\int_{B_p(\rho)} |F_A|^2 < \kappa^{-2}$$

for each $p \in X$ and $A \in \{A_i\}_{i \in \Lambda}$. Then, there exist a closed set $Z \subset X$ with empty interior, a real line bundle $\mathcal{I}$ over $X \backslash Z$, a section $\nu \in \Gamma(X \backslash Z; \mathcal{I} \otimes \Lambda^+)$, a connection A_Δ on $P|_{X \backslash Z}$ and an isometric bundle homomorphism $\sigma_\Delta : \mathcal{I} \to \mathfrak{g}_P$. Their properties are listed below.

- *Z is the zero locus of $|\hat{\beta}_\diamond|$.*

- *$|\hat{\beta}_\diamond|$ belongs to the Hölder space $C^{0,1/\kappa}(X)$.*

- *$|\nu|$ is the restriction of $|\hat{\beta}_\diamond|$ to $Z \subset X$.*

- *The section ν is harmonic in the sense that $d\nu = 0$.*

- *$|\nabla \nu|$ is an L^2-function on $X \setminus Z$ that extends as an L^2-function on X.*

- *The curvature tensor of A_Δ is anti-self-dual.*

- *The homomorphism σ_Δ is A_Δ-covariantly constant. There is, in addition, a subsequence $\Xi \subset \Lambda$ and a sequence $\{g_i\}_{i \in \Xi}$ of automorphisms from P such that*

  - *(i) $\{g_i(\beta_i)\}_{i \in \Xi}$ converges to $\sigma_\Delta \circ \nu$ in the L^2_1-topology on compact subsets in $X \setminus Z$ and in C^0-topology on X.*

  - *(ii) $\{g_i(A_i)\}_{i \in \Xi}$ converges on compact subsets of $X \setminus Z$ to A_Δ in the L^2_1-topology.*

The next theorem is a special case of Theorem 1.2 in [T3]. It implies among other things that Z has measure zero. (There are closed sets with empty interior and positive measure.) This upcoming theorem uses the

13

notion of a point of discontinuity for $\mathcal{I}$; this being a point in Z with the following property: if U is any neighbourhood of the point, then the restriction of $\mathcal{I}$ to $U \setminus Z$ is not isomorphic to the product bundle.

Theorem 4.5 ([T3]). *Let Z and $\mathcal{I}$ be as described in Proposition 4.4. The set Z has Hausdorff dimension at most 2, and moreover, the set of the points of discontinuity for $\mathcal{I}$ (defined in the preceding paragraph) are the points in the closure of an open subset of Z that has the structure of a 2-dimensional C^1-submanifold in X denoted by Σ.*

We are finally ready to use the above results in the following proposition, of which items 1 and 3 conclude the proof of Theorem 3.1.

Proposition 4.6 ([T5]). *Let $\mathcal{I}$, σ_Δ and A_Δ be as described in Proposition 4.4, so that σ_Δ and A_Δ are defined over $X \setminus Z$. Then*

1. *There exists a smooth anti-self-dual connection A_0 defined over all of X, and a continuous Sobolev-class L_2^2 gauge transformation g_0 defined over $X \setminus Z$, such that $g_0(A_0)$ is the restriction to $X \setminus Z$ of A_0. Defining $\sigma_0 := g_0(\sigma_\Delta)$ over $X \setminus Z$, then $\nabla_{A_0}\sigma_0 = 0$.*

2. *The bundle $\mathcal{I}$ over $X \setminus Z$ extends to a bundle defined over all of X, which we again denote by $\mathcal{I}$.*

3. *There exist extensions of both $\nu \in \Gamma(\mathcal{I} \otimes \Lambda^+)$ and $\sigma_0 : \mathcal{I} \to \mathfrak{g}_P$ to all of X. We again denote these by ν and σ_0. The extensions satisfy $d\nu = 0$ and $\nabla_{A_0}\sigma_0 = 0$.*

Proof. We first prove item 1.

Apply Proposition 4.1 to the subsequence $\{A_i\}_{i\in\Xi\subset\Lambda}$ of Proposition 4.4. This yields a subsequence $\Pi \subset \Xi$, a sequence of gauge transformations $\{g_i\}_{i\in\Pi}$, and a connection A_0 which is the weak L_1^2 limit of $\{g_i(A_i)\}_{i\in\Pi}$ over all of X.

Recall from Proposition 4.4 that A_Δ is the L_1^2 limit over compact subsets of $X \setminus Z$ of gauge transformations of $\{A_i\}_{i\in\Xi}$. In particular, both A_0 and A_Δ are weak L_1^2 limits over $X \setminus Z$ of gauge-equivalent connections. Since weak L_1^2 limits preserve L_2^2 gauge equivalence, it follows that there exists a Sobolev-class L_2^2 gauge transformation g_0 such that $g_0(A_\Delta) = A_0$.

Note that A_Δ is anti-self-dual and gauge-equivalent over the complement of Z to A_0. Thus, A_0 is an L_1^2 connection whose self-dual curvature is L^2, and vanishes on the complement of Z, which by Proposition 4.5 is a set of measure zero. Hence the curvature of A_0 is anti-self-dual, and so a standard

elliptic regularity argument (see e.g. [DK, §4.4] and Appendix A) can be used to prove that there is an L_2^2 and C^0 automorphism of P that tranforms A_0 into a smooth connection with anti-self-dual curvature. After possibly composing g_0 with such an automorphism, we may assume without loss of generality that A_0 is smooth and that g_0 is continuous.

That $\nabla_{A_0}\sigma_0 = 0$ follows from Proposition 4.4 since σ_Δ is A_Δ-covariantly constant. This establishes the item 1.

We next prove the item 2, that $\mathcal{I}$ extends over Z. Let $\Sigma \subset Z$ denote the C^1 submanifold that is described by Theorem 4.5. It is enough to prove that Σ is empty. For this purpose, assume to the contrary that $\Sigma \neq \emptyset$ and let $S \subset \Sigma$ be a component. This is a C^1 embedded, 2-dimensional surface. Fix a point $p \in S$. Since S is C^1, there is an embedded disk $D \subset X$ whose closure intersects S tranversally at a single point which is p. This is also its only intersection point with Z since S is an open subset of Z. Since p is a point of discontinuity for the bundle $\mathcal{I}$, the restriction of $\mathcal{I}$ to $D \setminus p$ is not isomorphic to the product line bundle. In particular, parallel transport by A_Δ of σ_Δ along any circle in $D \setminus p$ which wraps once around p gives $-\sigma_\Delta$. However, A_Δ is gauge-equivalent to a connection which is smooth over all of D. Thus parallel transport around sufficiently small circle in D will be arbitrarily close to $+\sigma_\Delta$, which is a contradiction.

Finally, we prove item 3 by showing that both ν and σ_0 extend to all of X as L_1^2 sections. Granted this extension, we may argue as in item 1 that both $d\nu$ and $\nabla_{A_0}\sigma_0$ are L^2 sections which vanish almost everywhere, and hence by elliptic regularity, ν and σ_0 are smooth and satisfy $d\nu = 0$ and $\nabla_{A_0}\sigma_0 = 0$ over all of X.

As a warning, note that not all L_1^2 functions extend over sets of measure zero. The simplest example is perhaps the Heaviside theta function, which is L_1^2 on $(-1,0) \cup (0,1)$, but has no L_1^2 extension to $(-1,1)$. Nevertheless, we will show that Z is negligibly small in the sense that it has zero capacity and thus it follows that L_1^2 extensions over Z always exist.

The section ν extends over X as desired if there exists a sequence $\{u_n\}_{n\in\{1,2,\dots\}}$ of smooth functions on X with pointwise norm bounded by 1 that vanish on Z and converge in the L_1^2 topology to the constant function 1. Indeed, given such a sequence, we may pass to a subsequence such that $\{u_n\}_{n\in\{1,2,\dots\}}$ converges to 1 almost everywhere. By the dominated convergence theorem in L^2, it follows that $\{u_n\nu\}_{n\in\{1,2,\dots\}}$ defines a Cauchy sequence of L_1^2 sections over the whole of X. The limit is equal to ν on $X \setminus Z$, and hence is our desired L_1^2 extension over X. The same argument also gives the desired L_1^2 extension of σ_0.

To find the desired sequence $\{u_n\}_{n\in\{1,2,\dots\}}$ of L_1^2 approximations to 1,

we first digress to give a definition. Fix a ball $B \subset X$ of small radius and suppose that V is a subset of B. The capacity of the pair (B, V) is denoted by $\mathrm{cap}(B, V)$ and it is defined as follows:

$$\mathrm{cap}(B, V) = \inf_{u \in C^\infty(B)} \left\{ \int_X |du|^2 : u = 1 \text{ on } X \setminus B \text{ and } u = 0 \text{ on } V \right\}.$$

For the present purposes, the significance of $\mathrm{cap}(B, V)$ is that if it is zero, then there is a sequence of functions on X with L^∞ norm 1 that vanish on V but converge to the constant function 1 in $L_1^2(X)$. Now, the definition is such that the capacity of a countable union of sets is no greater than the sum of the capacities of the sets. See, e.g. Theorem 1.3 in [R] by Reshetnyak or the more accessible reference of Theorem 3.2 in [C] by Costea. Meanwhile, Lemma 3 in [V] proves that a closed set in a 4-dimensional ball with finite 2-dimensional Hausdorff measure has zero capacity. Granted these last two observations, it is enough to prove that Z is contained in a countable union of sets with finite 2-dimensional Hausdorff measure. [†] To see that this is the case, note that because $\mathcal{I}$ extends over Z, it is locally isomorphic to the product bundle. As a consequence, the section ν of $\mathcal{I} \otimes \Lambda^+$ extends to the whole of X as a smooth, harmonic section. One can then use Taylor's theorem in the manner of [HHL] to see that Z is contained in a countable union of 2-dimensional, Lipshitz submanifolds. Since dimension 2, Lipshitz submanifolds have locally finite 2-dimensional Hausdorff measure, it follows that Σ is contained in a countable union of submanifolds with finite 2-dimensional Hausdorff measure. $\qquad\square$

A   Weak L_1^2 compactness of non-concentrating connections modulo gauge [‡]

In this section we prove Proposition 4.1.

We begin with a brief outline of the proof idea. Cover X by finitely many small balls $\{B_r(p_i)\}_i$ of fixed radius r such that Uhlenbeck's Theorem applies for a subsequence of $\{A_\alpha\}$. For any fixed α in this subsequence, let P_α denote the principal bundle associated to a cocycle of trivializations over the $\{B_r(p_i)\}_i$ which put A_α into Coulomb gauge. We assume that A_α is a smooth connection so that P_α is a smooth principal bundle. Denote by η_α

[†]We know from Theorem 4.5 only that the Hausdorff dimension of Z is at most 2, but this does not guarantee that the Hausdorff measure is locally finite. For instance, the Hausdorff dimension of the graph of $\sin(1/x)$ has Hausdorff dimension 1.

[‡]The proof presented here was also conveyed to the author by Clifford Taubes.

the canonical isomorphism of P_α with P. There is a subsequence such that in each ball, the pull-backs of the local connection forms converge weakly in L_1^2. There is a further subsequence such that the cocycles determining P_α converge weakly in L_2^2 to a C^0 cocycle which determines a bundle denoted by $P^\diamond$.

For α sufficiently large, so that the cocycle of $P^\diamond$ is sufficiently close to the cocycle of P_α, there is an isomorphism $\phi_\alpha : P^\diamond \to P_\alpha$. Next we construct a family of smoothing P^δ of $P^\diamond$, parametrized by small positive δ, and let P' denote a particular choice of P^δ. This comes with an isomorphism $\lambda : P' \to P^\diamond$. Finally, composing all these isomorphisms we get the composite isomorphism

$$\Phi_\alpha : P' \xrightarrow{\lambda} P^\diamond \xrightarrow{\phi_\alpha} P_\alpha \xrightarrow{\eta_\alpha} P.$$

These isomorphisms are constructed so that along the appropriate subsequence, $\Phi_\alpha^*(A_\alpha)$ is weakly L_1^2 convergent in the smooth bundle P'. Choosing some fixed index α_0, it follows that $\{\Phi_\alpha \Phi_{\alpha_0}^{-1}\}$ is a sequence of automorphisms such that $\{(\Phi_\alpha \Phi_{\alpha_0}^{-1})^*(A_\alpha)\}$ is weakly L_1^2 convergent on the original bundle P, as desired.

A.1 Background for analysis for the proof

To set the notation for what is to come, suppose that X is a Riemannian 4-manifold. Let $r_0 > 0$ be a small number chosen to be much less than the injectivity radius of X. (We assume that X is either compact, or if not, then the part of X of interest has a lower bound for its injectivity radius.) Gaussian coordinate charts can be defined on balls of radius r_0 in X. Supposing that $p \in X$ and $r \in (0, r_0)$, we use $B_r(p)$ to denote the open, radius r ball centered at p. If the point p is immaterial to the discussion, we will use just B_r to denote this ball. The boundary of the closure of B_r is denoted by ∂B_r. One more piece of notation: we use c_0 to denote a number that is greater than 1 and independent of any relevant choices made, such as a point p, radius r or a given connection or other geometric object. The precise value of c_0 can be assumed to increase between successive appearances. What follows next is a review of some of the basic analysis that is needed later. Supposing that p is a given point in X and $r \in (0, r_0)$, let f denote a given Sobolev class L_1^2 function on B_r. This is to say that f and its differential df are square integrable on B_r. The following are the basic function space inequalities that will be used:

- $\left(\int_{B_r} |f|^4 \right)^{\frac{1}{2}} \leq c_0 \int_{B_r} \left(|df|^2 + r^{-2} |f|^2 \right).$

- $\int_{B_r} \frac{1}{\text{dist}(\cdot,p)^2}|f|^2 \leq c_0 \int_{B_r} \left(|df|^2 + r^{-2}|f|^2\right).$

- If f has compact support in B_r or if the integral of f over B_r is zero, then the preceding two inequalities hold without the $r^{-2}|f|^2$ term in the integrand.

- $\int_{B_r} f^2 \leq c_0 \left(r^2 \int_{B_r} |df|^2 + r \int_{\partial B_r} f^2\right).$

$$(A.1)$$

The first, third and fourth inequalities in this list are basic Sobolev inequalities. The second inequality is called Hardy's inequality.

To set the stage for a second set of basic facts, suppose that n is a positive integer and let $B^n \subset \mathbb{R}^n$ denote the radius 1, open ball at the origin. Now let m denote a second positive integer and f denote a given smooth map from the closure of B^n to B^m mapping 0 to 0. Suppose that x is a Sobolev class L_2^2 map from X to B^n.

FACT 1: The composition of first x and then f defines a Sobolev class L_2^2 map from X to B^m whose L_2^2 norm is a priori bounded by c_f times the square of the L_2^2 norm of x with c_f being a number that depends only on f.

FACT 2: If x is a Sobolev class L_2^2 map from X to B^n and if $\{x_\alpha\}_{\alpha \in \{1,2,\dots\}}$ is a sequence of Sobolev class L_2^2 maps from X to B^n with an a priori L_2^2 norm bound that converges weakly in the L_2^2 topology to x, then the sequence $\{f \circ x_\alpha\}_{\alpha \in \{1,2,\dots\}}$ converges weakly in the L_2^2 topology to $f \circ x$.

The assertion made by FACT 1 can be proved by using the basic calculus to write

- $|(f \circ x)| \leq c_f |x|.$

- $\nabla(f \circ x) = (df)|_x \cdot \nabla x.$

- $\nabla(df \cdot dx) = (\nabla df)|_x \cdot (dx \otimes dx) + (df)|_x \cdot \nabla dx.$

$$(A.2)$$

The desired bound on the L^2 norms of $f \circ x$ and its first derivatives follow directly from the top two bullets in (A.2). The desired bound on the L^2 norm of the second derivatives of $f \circ x$ follow from the third bullet of (A.2) and the first bullet of (A.1). To prove FACT 2, first use FACT 1 to conclude that the sequence $\{f \circ x_\alpha\}_{\alpha \in \{1,2,\dots\}}$ is uniformly bounded in the L_2^2 topology for maps from X to B^m. It follows as a consequence that this sequence has a subsequence that converges weakly in the L_2^2 topology. It is sufficient to show that $f \circ x$ is necessarily the weak limit and to this end, it

is sufficient to show that $f \circ x$ is the strong L^2 limit. That can be proved by first using the fundamental theorem of calculus to write

$$(f \circ x) - (f \circ x_\alpha) = (x - x_\alpha) \int_0^1 \nabla f|_{x+t(x-x_\alpha)} dt \qquad (\text{A.3})$$

and then invoking the fact that $\{x_\alpha\}_{\alpha \in \{1,2,\dots\}}$ converges strongly in the L^2 topology to x.

Uhlenbeck's paper [U] plays a central role. To review some of what is done in this paper, fix a point $p \in X$ and a number $r \in (0, r_0)$ The product $SU(2)$ bundle over B_r is the bundle $B_r \times SU(2)$. The corresponding product connection is denoted by θ_0. Any connection on $B_r \times SU(2)$ can be written as $\theta_0 + \mathfrak{a}_A$ with $\mathfrak{a}_A$ being a section over B_r of the product bundle $B_r \times \mathfrak{su}(2)$. Here and below, we use $\mathfrak{su}(2)$ to denote the Lie algebra of $SU(2)$. The curvature 2-form of A is denoted by F_A. The group of automorphisms of the bundle $B_r \times SU(2)$ is the space of maps from B_r to $SU(2)$. If s is an automorphism of this product bundle and A is a connection, then s can be used to pull back A to get a new connection. Writing A as $\theta_0 + \mathfrak{a}_A$ and s^*A as $\theta_0 + \mathfrak{a}_{s^*A}$, the $\mathfrak{su}(2)$-valued 1-forms $\mathfrak{a}_A$ and $\mathfrak{a}_{s^*A}$ are related by the rule

$$\mathfrak{a}_{s^*A} = s^{-1}\mathfrak{a}_A s + s^{-1}ds \qquad (\text{A.4})$$

with s here viewed as a map from B_r to $SU(2)$.

The first basic result is Corollary 2.2 in [U]. A slightly weaker version is stated next as a lemma.

Lemma A.1. *There exists a universal number (independent of X and its metric) to be denoted by c which is greater than 1 and which has the following significance: Suppose that $p \in X$ and that $r \in (0, r_0)$. Let A denote a smooth connection on B_r that obeys*

$$\int_{B_r} |F_A|^2 \le c^{-2}$$

Then there exists an automorphism of $B_r \times SU(2)$ to be denoted by s such that s^*A can be written as $\theta_0 + \mathfrak{a}_{s^*A}$ obeying

- $d^*\mathfrak{a}_{s^*A} = 0.$

- $x \cdot \mathfrak{a}_{s^*A} = 0$ *on* $\partial B_r.$

- $\int_{B_r} (|\nabla \mathfrak{a}_{s^*A}|^2 + r^{-2}|\mathfrak{a}_{s^*A}|^2) \le c_0^2 \int_{B_r} |F_A|^2.$

Note that Uhlenbeck's Corollary 2.2 is stated using Sobolev class L_1^2 connections. We only need the version with the connection being smooth. (One can get Uhlenbeck's L_1^2 version in a straightforward way from the version above by taking limits.) By way of a parenthetical remark, the number c that appears in Lemma A.1 can be chosen so that $\mathfrak{a}_{s^*A}$ also obeys a reverse of sorts to the third bullet's inequality:

$$\int_{B_r} |F_A|^2 \le c_0^2 \int_{B_r} |\nabla \mathfrak{a}_{s^*A}|^2. \qquad (A.5)$$

That this can be done follows from the integration by parts equality at the bottom of page 35 in [U] and the inequalities in (A.1). The number c is chosen in what follows so that (A.5) holds. Let $p \in X$ and $r \in (0, r_0)$. Suppose now that $\{A_\alpha\}_{\alpha \in \{1,2,\dots\}}$ is a sequence of connections on $B_r \times SU(2)$, each with the L^2 norm on B_r of the curvature of each A_α being less than c^{-1}. Let $\{s_\alpha\}_{\alpha \in \{1,2,\dots\}}$ denote the corresponding sequence of automorphisms of $B_r \times SU(2)$ that is supplied by Lemma A.1. For each $\alpha \in \{1, 2, \dots\}$, write $s_\alpha^* A_\alpha$ as $\theta_0 + \mathfrak{a}_\alpha$ with $\mathfrak{a}_\alpha$ obeying the $A = A_\alpha$ version of the three bullets in Lemma A.1. In particular, it follows from the third bullet that the sequence $\{\mathfrak{a}_\alpha\}_{\alpha \in \{1,2,\dots\}}$ is uniformly bounded in the L_1^2 topology on B_r. Therefore, this sequence has a weakly convergent subsequence. Fix such a subsequence and let Θ denote the set of positive integers for the corresponding indexing set in $\{1, 2, \dots\}$. The limit of $\{\mathfrak{a}_\alpha\}_{\alpha \in \Theta}$ is a Sobolev L_1^2 section over B_r of $T^*B_r \times \mathfrak{su}(2)$. Denote this limit by $\mathfrak{a}$. A priori, this limit 1-form obeys

- $d^*\mathfrak{a} = 0$.

- $x \cdot \mathfrak{a} = 0$ on ∂B_r.

- $\int_{B_r} \left(|\nabla \mathfrak{a}|^2 + r^{-2}|\mathfrak{a}|^2 \right) \le c_0^2 \liminf_{\alpha \in \Theta} \int_{B_r} |F_{A_\alpha}|^2$.

$$(A.6)$$

Note with regards to the second bullet that the components of the restriction of $\mathfrak{a}$ to ∂B_r are functions in the fractional Sobolev space $L_{1/2}^2(\partial B_r)$.

A.2   Limits of sequences on X

Let $P \to X$ denote a principal $SU(2)$ bundle. We assume the structure group of P is just $SU(2)$ below, but a similar argument works for $SO(3)$ bundles as well. Suppose that $r \in (0, r_0)$, that N is a positive integer and that $\{p_i\}_{i \in \{1,2,\dots N\}}$ is a set of points in X such that the corresponding set $\{B_r(p_i)\}_{i \in \{1,2,\dots N\}}$ is a cover of X by balls of radius r.

Let c be the constant defined in the previous section, and suppose that $\{A_\alpha\}_{\alpha \in \{1,2,\dots\}}$ is a sequence of connections on P that obey the following condition: If $B_r \in \{B_r(p_i)\}_{i \in \{1,2,\dots\}}$, then

$$\int_{B_r} |F_{A_\alpha}|^2 \leq \varepsilon^2 c^{-2}. \tag{A.7}$$

The number ε is determined near the end of this appendix.

Fix an index $i \in \{1, \dots N\}$ and let B_r denote for the moment the ball $B_r(p_i)$. Fix an isomorphism $B_r \times SU(2)$ to $P|_{B_r}$. Denote this isomorphism by ι_i. Invoke Lemma A.1 for the sequence $\{\iota_i^* A_\alpha\}$ on B_r. Denote the corresponding sequence of product bundle automorphisms by $\{s_{i,\alpha}\}_{\alpha \in \{1,2,\dots\}}$ and for each index α, write $s_{i,\alpha}^* \iota_i^* A_\alpha$ as $\theta_0 + \mathfrak{a}_{i,\alpha}$. A subsequence in $\{1, 2, \dots\}$ to be denoted by Θ can be chosen with the following significance: For each index $i \in \{1, \dots, N\}$, the corresponding sequence $\{\mathfrak{a}_{i,\alpha}\}_{\alpha \in \Theta}$ converges weakly in the L_1^2 topology. Use $\mathfrak{a}_i$ to denote the limit of this sequence. The next lemma says that $\{\mathfrak{a}_i\}_{i \in \{1,\dots,N\}}$ define a Sobolev class L_1^2 connection on a principal $SU(2)$ bundle over X. By way of a reminder, a connection on a principal bundle $P' \to X$ is said to be of Sobolev class L_1^2 if it can be written as $A_0 + \hat{\mathfrak{a}}$ with A_0 being a smooth connection and with $\hat{\mathfrak{a}}$ being an L_1^2 section of the vector bundle $T^* X \otimes (P' \times_{SU(2)} \mathfrak{su}(2))$

Proposition A.2. *There is a smooth principal $SU(2)$ bundle $P' \to X$, a Sobolev class L_1^2 connection on P' to be denoted by A, and, for each $i \in \{1, \dots N\}$, a C^0 and Sobolev class L_2^2 isomorphism from $B_r(p_i) \times SU(2)$ that pulls back A as $\theta_0 + \mathfrak{a}_i$.*

Proof. The proof has six steps.

Step 1: Fix for the moment indices $i, k \in \{1, 2, \dots, N\}$ such that $B_r(p_i) \cap B_r(p_k) \neq \emptyset$. Use g_{ki} to denote the automorphism $\iota_k^{-1} \circ \iota_i$ of the product $SU(2)$ bundle over $B_r(p_i) \cap B_r(p_k)$. For index $\alpha \in \Theta$, let $s_{ki,\alpha} = s_{k,\alpha}^{-1} g_{ki} s_{i,\alpha}$. The definition is such that the two connections satisfy

$$\theta_0 + \mathfrak{a}_{i,\alpha} = s_{ki,\alpha}^* (\theta_0 + \mathfrak{a}_{k,\alpha}).$$

Thus,

$$
\begin{array}{ccc}
(B_r(p_i) \cap B_r(p_k)) \times SU(2) \xrightarrow{s_{i,\alpha}} (B_r(p_i) \cap B_r(p_k)) \times SU(2) \xrightarrow{\iota_i} P|_{B_r(p_i) \cap B_r(p_k)} \\
\downarrow{\scriptstyle s_{ki,\alpha}} \qquad\qquad \downarrow{\scriptstyle g_{ki}} \qquad\qquad \downarrow{\scriptstyle \mathrm{id}} \\
(B_r(p_i) \cap B_r(p_k)) \times SU(2) \xrightarrow{s_{k,\alpha}} (B_r(p_i) \cap B_r(p_k)) \times SU(2) \xrightarrow{\iota_k} P|_{B_r(p_i) \cap B_r(p_k)}
\end{array}
$$

commutes, and

$$s_{ki,\alpha}\mathfrak{a}_{i,\alpha} = \mathfrak{a}_{k,\alpha}s_{ki,\alpha} + ds_{ki,\alpha}. \tag{A.8}$$

This equation implies among other things that the sequence $\{s_{ki,\alpha}\}_{\alpha\in\Theta}$ is uniformly bounded in the L_2^2 topology on the space of maps from $B_r(p_i) \cap B_r(p_k)$ to $SU(2)$. It follows as a consequence that there is a subsequence of Θ such that the corresponding subsequence of $\{s_{ki,\alpha}\}_{\alpha\in\Theta}$ converges weakly in the L_2^2 topology to a Sobolev L_2^2 map from $B_r(p_i)\cap B_r(p_k)$ to $SU(2)$. We can assume without loss of generality that the same subsequence works for all pairs $(i,k) \in \{1,\ldots N\}$. We denote the latter sort of subsequence by Λ. With this understood, let s_{ki} denote the limit map.

Step 2: As remarked above, the sequence $\{s_{ki,\alpha}\}_{\alpha\in\Lambda}$ converges weakly in the L_2^2 topology to an L_2^2 map s_{ki} from $B_r(p_i) \cap B_r(p_k)$ to $SU(2)$. The next lemma lists three crucial properties of the collection $\{s_{ki}\}_{i,k\in\{1,\ldots,N\}}$.

Lemma A.3. *The collection $\{s_{ki}\}_{i,k\in\{1,\ldots,N\}}$ has the following properties.*

- *$s_{ki}\mathfrak{a}_k = \mathfrak{a}_i s_{ki} + ds_{ki}$ for each pair (i,k) such that $B_r(p_i) \cap B_r(p_k) \neq \emptyset$.*

- *$s_{ki}s_{ni}^{-1}s_{nk} = \mathbb{I}$ for each triple (i,k,n) such that $B_r(p_i) \cap B_r(p_k) \cap B_{r(p_n)} \neq \emptyset$.*

- *s_{ki} is continuous for each pair (i,k) such that $B_r(p_i) \cap B_r(p_k) \neq \emptyset$.*

$$\tag{A.9}$$

Proof. The first bullet asserts an equality between L_1^2 functions and the second asserts an equality between L_2^2 functions. They follow from the fact that the sequences $\{\mathfrak{a}_{i,\alpha}\}_{\alpha\in\Lambda}, \{\mathfrak{a}_{k,\alpha}\}_{\alpha\in\Lambda}$ converge weakly in L_1^2 and the sequence $\{s_{ki,\alpha}\}_{\alpha\in\Lambda}$ converges weakly in L_2^2. Note in this regard that the product of a bounded, L_2^2 function and an L_1^2 function is an L_1^2 function, and that the products of bounded L_2^2 functions are L_2^2 functions. However, if s denotes the relevant bounded L_2^2 function, then the respective norm bounds of these products depend on both the L_2^2 norm and L^∞ norm of s.

The third bullet in (A.9) is proved using the fact that $d^*\mathfrak{a}_i = 0$ and $d^*\mathfrak{a}_k = 0$. What with the first bullet in (A.9), these norm bounds imply that s_{ki} obeys the equation

$$d^\dagger ds_{ki} = \langle \mathfrak{a}_k, ds_{ki} \rangle - \langle ds_{ki}, \mathfrak{a}_i \rangle, \tag{A.10}$$

where $d^\dagger$ is short hand for $- * d*$ which is the formal L^2 adjoint of d; and where $\langle \alpha, \beta \rangle$ is shorthand for $*(\alpha \wedge *\beta)$.

To exploit this identity, fix an open set in $B_r(p_i) \cap B_r(p_k)$ with compact closure in $B_r(p_i) \cap B_r(p_k)$ to be denoted by U_{ki}. Then fix a smooth, non-negative function on $B_r(p_i) \cap B_r(p_k)$ to be denoted by χ_{ki} that is equal to 1 on a neighborhood of the set U_{ik} and 0 on a neighborhood of the boundary of $B_r(p_i) \cap B_r(p_k)$. Given a point $p \in B_r(p_i) \cap B_r(p_k)$, let G_p denote the Dirichlet Green's function on $B_r(p_i) \cap B_r(p_k)$ for the operator $d^\dagger d$ with pole at p. Formally, $d^\dagger d G_p = \delta_p$ with δ_p being the Dirac delta-function with unit mass at p. Note in any event that

$$|G_p|(\cdot) \leq c_0 \frac{1}{\operatorname{dist}(\cdot, p)^2} \text{ and } |\nabla G_p| \leq c_0 \frac{1}{\operatorname{dist}(\cdot, p)^3} \qquad \text{(A.11)}$$

It follows from (A.10) that if $p \in U_{ki}$, then

$$
\begin{aligned}
s_{ki}(p) = &\int_{B_r(p_i) \cap B_r(p_k)} G_p \chi_{ki}(\langle \mathfrak{a}_k, ds_{ki} \rangle - \langle ds_{ki}, \mathfrak{a}_i \rangle) \\
&+ \int_{B_r(p_i) \cap B_r(p_k)} G_p((d^\dagger d \chi_{ki}) s_{ki} - 2\langle d\chi_{ki}, ds_{ki} \rangle).
\end{aligned}
\qquad \text{(A.12)}
$$

Note that the absolute value of the left most integral on the right hand side of (A.12) can be bounded using the left most inequality in (A.11), Cauchy–Schwarz and the second bullet in (A.1) by

$$\left| \int_{B_r(p_i) \cap B_r(p_k)} G_p \chi_{ki}(\langle \mathfrak{a}_k, ds_{ki} \rangle - \langle ds_{ki}, \mathfrak{a}_i \rangle) \right| \leq c_0(||\mathfrak{a}_i||_* + ||\mathfrak{a}_k||_*)||ds_{ki}||_*,$$
$$\text{(A.13)}$$

where we introduce the notation whereby $||\alpha||_*$ is the sum of the L^2 norm of $\nabla\alpha$ and the L^4 norm of α on $B_r(p_i) \cap B_r(p_k)$. Meanwhile, the right most integral in (A.12) is no greater than

$$c_0 \operatorname{dist}(p, \partial U_{ki})^{-3}(||d\chi||_{L^1} + ||d^\dagger d\chi||_{L^1} \operatorname{dist}(p, \partial U_{ki})),$$

this being a consequence of integration by parts, both inequalities in (A.11) and the fact that s_{ki} has norm bounded by c_0. The formula in (A.13) is the key to proving that s_{ki} is continuous. To say more, fix for the moment $\rho \in (0, r)$ so that there is a ball of radius ρ contained entirely in U_{ki} with the distance from any point in this ball to ∂U_{ki} being greater than 100ρ. Let B_ρ denote such a ball. Suppose now that p and q are two points in B_ρ. Use (A.12) to write

$$s_{ki}(p) - s_{ki}(q) = \mathcal{K}_+ + \mathcal{K}_- \qquad \text{(A.14)}$$

where the $\mathcal{K}_+$ is the contribution to the left most integral in (A.12) from $B_{3\rho}$ and where $\mathcal{K}_-$ is the rest of this integral and the right most integral in (A.12). One finds using (A.12) and the second bullet of (A.1) that

- $|\mathcal{K}_+| \leq c_0(||\mathfrak{a}_i||_{*,3\rho} + ||\mathfrak{a}_k||_{*,3\rho})||ds_{ki}||_{*,3\rho}$.

- $|\mathcal{K}_-| \leq ||G_p - G_q||_{L^\infty_-}(||\chi||_{L^2}(||\mathfrak{a}_i||_* + ||\mathfrak{a}_k||_*)||ds_{ki}||_*$
$$+ ||d^\dagger d\chi||_{L^1}||s_{ki}||_{L^\infty} + ||d\chi||_{L^{4/3}}||ds_{ki}||_*).$$
$$\text{(A.15)}$$

Here, we define $||\alpha||_{*,3\rho}$ to denote the sum of the L^2 norm of $\nabla\alpha$ over $B_{3\rho}$ and the L^4 norm of α over $B_{3\rho}$ and L^∞_- to denote the L^∞ norm over the complement of $B_{3\rho}$.

With (A.14) and (A.15) understood, let $\delta > 0$ be arbitrary. Since $\mathfrak{a}_i, \mathfrak{a}_k$ and ds_{ki} are all L^2_1 functions, there exists some ρ_δ such that their L^2_1 norms and L^4 norms on any ball of radius $3\rho_\delta$ or less is smaller than δ. It follows as a consequence that if p, q are in such a ball, then $|\mathcal{K}_+| \leq c_0\delta^2$. Finally, once ρ is fixed, we observe that $||G_p - G_q||_{L^\infty_-} \leq c_0\rho^{-2}\,\mathrm{dist}(p,q)$, and so $|\mathcal{K}_-| \to 0$ as $\mathrm{dist}(p,q) \to 0$. Since $|\mathcal{K}_+|$ and $|\mathcal{K}_-|$ can simultaneously be made arbitrarily small, this proves that s_{ki} is continuous. □

Step 3: Since the set $\{s_{ki}\}_{i,k\in\{1,\dots,N\}}$ are continuous maps from intersections of the cover $\{B_r(p_i)\}_{i\in\{1,2,\dots\}}$ to $SU(2)$ (the third bullet of Lemma A.3) and they obey the cocycle condition (the second bullet of Lemma A.3), they a priori define a C^0 principal G-bundle over X. Denote this bundle by $P^\diamond$ The bundle $P^\diamond$ has more than just a C^0 principal bundle structure because its transition functions are also in the Sobolev space L^2_2. This implies in particular that one can talk about L^2_2 sections of associated bundles to $P^\diamond$ and L^2_1 connections. In particular, an L^2_1 connection is given by cocycle data consisting of a collection of L^2_1 Sobolev class 1-forms with values in $\mathfrak{su}(2)$, one for each set in $\{B_r(p_i)\}_{i\in\{1,2,\dots,N\}}$, with the collection subject to the cocycle condition that is depicted in the top bullet of Lemma A.3.

Step 4: Every C^0 principal bundle is isomorphic to some C^∞ principal bundle, and thus $P^\diamond$ is also. This being the case, the proof of Proposition A.2 is completed with an argument to the effect that there is an isomorphism from some smooth principal bundle to $P^\diamond$ that respects the L^2_2 structure; and in particular, that pulls back L^2_1 connections to L^2_1 connections. One way to construct the bundle and the desired isomorphism is to first view $P^\diamond$ as the pull-back of the tautological $SU(2)$ bundle over $\mathbb{HP}^{N-1}$ with the latter being the quaternionic projective space $(\times_N \mathbb{H} \setminus 0)/\mathbb{H}^*$, where $\mathbb{H}$ denotes the quaternions, N is a large integer and $\mathbb{H}^*$ denotes the non-zero quaternions. To define the map, choose a partition of unity for the cover $\{B_r(p_i)\}_{i\in\{1,\dots,N\}}$, this denoted by $\{\varphi_i\}$. The map in question is denoted by Φ. Given an index $i \in \{1,\dots,N\}$, and a point $p \in B_r(p_i)$, the point $\Phi(p)$ is the equivalence class in $\mathbb{HP}^{N-1}$ of a vector in $\mathbb{H}$ whose coordinates are as follows:

- The i'th coordinate of $\Phi(p)$ is $\varphi_i(p)1$.

- If $k \in \{1, \ldots, N\}$ and $k \neq i$ then the k'th coordinate of $\Phi(p)$ is

$a)$ $\varphi_k(p)s_{ki}(p)1$ if $B_r(p_i) \cap B_r(p_k) \neq \emptyset$.

$b)$ 0 if $B_r(p_i) \cap B_r(p_k) = \emptyset$.

$$(A.16)$$

This definition is consistent with regards to points in intersecting balls because of the cocycle condition from the second bullet of Lemma A.3. The map Φ is observedly an L_2^2 and C^0 map because this is so for the function $\{s_{ki}\}_{i,k \in \{1,2,\ldots,N\}}$. The tautological $SU(2)$ bundle over $\mathbb{HP}^{N-1}$ is denoted in what follows by $\mathcal{P}_{N-1}$. It follows by construction that the bundles $P^\diamond$ and $\Phi^*\mathcal{P}_{N-1}$ are given by the same cocycle data and so they are one and the same bundle.

Step 5: A standard way to smooth a C^0 map from one manifold to a second compact manifold is to embed the second manifold in some finite dimensional Euclidean space. One then smooths the embedding map as a map to the Euclidean space. If the smoothing is sufficiently close to the original map, then its image will lie in a tubular neighborhood of the embedding and so the tubular neighborhood projection will compose with the smoothing to give a perturbation of the original map that is smooth. Let X denote the domain manifold. In particular, if the smoothing operation is defined by convolution with a smoothing kernel with support very close to the diagonal in $X \times X$, then the resulting map to the Euclidean space will be very close to the original. Moreover, the smoothed map will also be very close to the original in a given Sobolev topology if the original map is in the given Sobolev topology. (This follows from FACT 1 in Appendix A.1.) Taking the image manifold to be $\mathbb{HP}^{N-1}$, one then obtains a smooth map from X to $\mathbb{HP}^{N-1}$ that is as close as desired to the original in the C^0 topology and also in the L_2^2 topology. The pull-back of the tautological bundle by such a smoothing is denoted here by P^δ with δ being an upper bound for the C^0 and L_2^2 distances between the map Φ and the smoothing. Note in particular that δ can be as small as desired. Note also that it follows from this construction that P^δ has a product structure on each ball $B_r(p_i)$ with smooth transition functions $\{g_{ki}^\delta\}_{k,i \in \{1,2,\ldots,N\}}$ that are δ close in the C^0 and L_2^2 topology to those of $P^\diamond$. (This again follows from FACT 1 in Appendix A.1.)

Step 6: Granted the preceding, one can invoke Proposition 3.2 in [U] and then copy almost verbatim the proof of Corollary 3.3 in [U] to prove the following lemma:

Lemma A.4. *Fix a refinement $\{U_i\}_{i\in\{1,\ldots,N\}}$ of the cover $\{B_r(p_i)\}$ such that each U_i has compact closure in the corresponding $B_r(p_i)$. If δ is sufficiently small, then there is a collection $\{\rho_i : U_i \to SU(2)\}_{i\in\{1,2,\ldots\}}$ of $C^0 \cap L_2^2$ maps such that $g_{ki}^\delta \rho_i = \rho_k s_{ki}$ for all pairs of indices i, k corresponding to intersecting balls. This collection $\{\rho_i\}$ therefore defines a C^0 and L_2^2 isomorphism from $P^\diamond$ to P^δ.*

Proof. As noted above, the proof amounts to invoking Proposition 3.2 in [U] and copying almost verbatim the arguments for Corollary 3.3 in [U]. The only changes are to take $p = 2$ and to use the a priori fact that the transition functions are close in the C^0 topology. In Uhlenbeck's proof, she uses the fact that they are close in the Sobolev topology to deduce C^0 closeness. Here, one has C^0 closeness a priori. The basic observation in this regard is that the exponential map on a Lie group is invertible on some fixed neighborhood of the identity element. Moreover, the inverse on this neighborhood has uniform bounds on its derivatives to any desired order. Thus, if one knows a priori that a map has image in this neighborhood, then derivatives of the map can be bounded knowing the derivative of its inverse under the exponential map, and vice-versa. (This is, again, an application of Fact 1 in Appendix A.1.)

We take P' to be P^δ for any fixed δ which is sufficiently small to satisfy the hypothesis of Lemma A.4. Let $\lambda : P' \to P^\delta$ denote the corresponding isomorphism of Lemma A.4. We take A to be the connection of P' which corresponds under λ to the connection on $P^\diamond$ determined by $\{\mathfrak{a}_i\}_{i\in\{1,\ldots,N\}}$. This concludes the proof of Proposition A.2. $\qquad\square$

$\square$

A.3   The convergence to the limit

Proposition A.2 asserts that the limit data $\{\mathfrak{a}_i\}_{i\in\{1,\ldots,N\}}$ comes from an L_1^2 connection on a smooth principal bundle over X. To restate matters, there is a smooth bundle $P' \to X$, an L_1^2 connection A on P', and, over each $B_r(p_i)$, an isomorphism $\tau_i : B_r(p_i) \times SU(2) \to P'$ such that $\tau_i^* A = \theta_0 + \mathfrak{a}_i$. We first check that P' is isomorphic to the bundle P where the original sequence $\{A_\alpha\}_{\alpha\in\Lambda}$ lives. To see that this is so, we use the fact that principal $SU(2)$ bundles are classified by their second Chern number which can be computed using curvature integrals. In particular, the second Chern number of P' is

$$-\frac{1}{4\pi^2} \int_X \mathrm{tr}(F_A \wedge F_A), \tag{A.17}$$

with the trace being the trace from the 2-dimensional representation of $SU(2)$. Likewise, the second Chern class of P is any $\alpha \in \Lambda$ version of

$$-\frac{1}{4\pi^2} \int_X \mathrm{tr}(F_{A_\alpha} \wedge F_{A_\alpha}). \tag{A.18}$$

To compute these numbers, we use the fact that $\mathrm{tr}(F_A \wedge F_A)$ on a set U_i from the refined cover can be written as $d(\mathrm{tr}(\mathfrak{a}_i \wedge d\mathfrak{a}_i + \frac{1}{3}\mathfrak{a}_i \wedge \mathfrak{a}_i \wedge \mathfrak{a}_i))$. A similar formula holds for the integrand in (A.18) on U_i, the only change being that $\mathfrak{a}_{i,\alpha}$, replaces $\mathfrak{a}_i$. Now let $\{\varphi_i\}_{i\in\{1,\ldots,N\}}$ be a partition of unity subordinate to $\{U_i\}_{i\in\{1,\ldots N\}}$. Then (A.17) is equal to

$$-\frac{1}{4\pi^2} \sum_{i\in\{1,\ldots,N\}} \int_X \varphi_i d\mathrm{tr}\left(\mathfrak{a}_i \wedge d\mathfrak{a}_i + \frac{1}{3}\mathfrak{a}_i \wedge \mathfrak{a}_i \wedge \mathfrak{a}_i\right), \tag{A.19}$$

which is, after an integration by parts,

$$\frac{1}{4\pi^2} \sum_{i\in\{1,\ldots,N\}} \int_X d\varphi_i \mathrm{tr}\left(\mathfrak{a}_i \wedge d\mathfrak{a}_i + \frac{1}{3}\mathfrak{a}_i \wedge \mathfrak{a}_i \wedge \mathfrak{a}_i\right). \tag{A.20}$$

A similar formula holds for (A.18) with the replacement of $\mathfrak{a}_i$ by $\mathfrak{a}_{i,\alpha}$. Since φ_i is smooth, and since $\{\mathfrak{a}_{i,\alpha}\}_{\alpha\in\Lambda}$ converges weakly in L_1^2 and strongly in L^2, it follows that for each i, the Λ-indexed sequences with α term given by

$$\int_X d\varphi_i \mathrm{tr}\left(\mathfrak{a}_{i,\alpha} \wedge d\mathfrak{a}_{i,\alpha} + \frac{1}{3}\mathfrak{a}_{i,\alpha} \wedge \mathfrak{a}_{i,\alpha} \wedge \mathfrak{a}_{i,\alpha}\right) \tag{A.21}$$

converges as α gets ever larger to the i'th term in (A.20).

Now that we know that P and P' are isomorphic, we need to construct a sequence $\{g_\alpha\}_{\alpha\in\Lambda}$ of smooth isomorphisms from P' to P with the property that $\{g_\alpha^* A_\alpha\}_{\alpha\in\Lambda}$ converges in the weak L_1^2 topology to A. The construction has seven steps.

Step 1: To set the stage for this, fix a pair $i,k \in \{1,\ldots,N\}$ whose corresponding balls intersect. Then fix an index α. The transition function $s_{ki,\alpha}$ obeys (A.8) and thus it obeys the analogue of (A.10):

$$d^\dagger ds_{ki,\alpha} = \langle \mathfrak{a}_{k,\alpha}, ds_{ki,\alpha}\rangle - \langle ds_{ki,\alpha}, \mathfrak{a}_{i,\alpha}\rangle \tag{A.22}$$

As a consequence, it is given on U_{ki} by the analogue of (A.12):

$$
\begin{aligned}
s_{ki,\alpha}(p) = &\int_{B_r(p_i)\cap B_r(p_k)} G_p \chi_{ki}(\langle \mathfrak{a}_{k,\alpha}, ds_{ki,\alpha}\rangle - \langle ds_{ki,\alpha}, \mathfrak{a}_{i,\alpha}\rangle) \\
&+ \int_{B_r(p_i)\cap B_r(p_k)} G_p\left(-d^\dagger d\chi_{ki}s_{ki,\alpha} + 2\langle d\chi_{ki}, ds_{ki,\alpha}\rangle\right).
\end{aligned} \tag{A.23}
$$

Now, if $p \in U_{ki}$, then the right most integral in (A.23) converges uniformly as α gets ever larger to the corresponding integral in (A.12) because G_p is smooth on the support of the integrand as long as $p \in U_{ki}$. The left most integral in (A.23) need not converge as α gets larger. In any event, it follows from (A.11) and from (A.1) using Lemma A.1 that the absolute value of this integral is no greater than $c_0(||\mathfrak{a}_{k,\alpha}||_* + ||\mathfrak{a}_{i,\alpha}||_*)||ds_{ki,\alpha}||_*$. From (A.8) and (A.10) it follows that

$$||ds_{ki,\alpha}||_* \leq c_0(||\mathfrak{a}_{k,\alpha}||_* + ||\mathfrak{a}_{i,\alpha}||_*)^2.$$

Using (A.7), it follows that the left most integral in (A.23) is bounded by $c_0 c^{-3} \varepsilon^3$.

Step 2: Since this same bound holds for the norm of the left most integral in (A.12), it follows that on the whole of U_{ki}, if α is sufficiently large, then

$$|s_{ki,\alpha} - s_{ki}| \leq c_0 c^{-3} \varepsilon^3. \tag{A.24}$$

This bound will determine the choice for ε. If one can arrange for ε to be small with a fixed size N to the number of elements in the cover $\{B_r(p_i)\}$, then one could invoke Uhlenbeck's Proposition 3.2 and the arguments of Corollary 3.3 (using Facts 1 and 2 in Appendix A.1) to obtain the desired sequence of isomorphisms. Unfortunately, our assumptions do not lead to an a priori bound on the cover. Even so, one this difficulty can be surmounted to obtain the desired conclusion. The next few steps explain how this is done.

Step 3: The balls $\{B_r(p_i)\}_{i \in \{1,\ldots,N\}}$ can be chosen in any event so that at most c_0 of them intersect at any given point, c_0 being independent of N. More to the point, one can use Lemma 16 in [D] to partition $\{1,\ldots,N\}$ into disjoint subsets $I_1 \cup I_2 \cup \cdots \cup I_\ell$ with $\ell \leq c_0$ such that if i and j are in the same subset, then $B_r(p_i)$ has distance greater than $8r$ from $B_r(p_j)$. Construct such a partition, and then for a $\in \{1, 2, \ldots, \ell\}$, let $\mathcal{U}_\mathrm{a}$ a denote the union of the balls $B_r(p_i)$ for $i \in I_\mathrm{a}$.

Step 4: The collection $\{\mathcal{U}_\mathrm{a}\}$ form an open cover of X with the number of elements being independent of ε and the sequence $\{A_\alpha\}$. The intersection between $\mathcal{U}_\mathrm{a}$ and $\mathcal{U}_\mathrm{b}$ consist of the union of the intersections of the balls with index $i \in I_\mathrm{a}$ with those with index $i \in I_\mathrm{b}$. Note that a given $B_r(p_i)$ for $i \in I_\mathrm{a}$ intersects at most one ball $B_r(p_j)$ for $j \in I_\mathrm{b}$ because the distance between the balls labeled by $i \in I_\mathrm{b}$ is greater than the diameter of the balls. Thus, $\mathcal{U}_\mathrm{a} \cap \mathcal{U}_\mathrm{b}$ is the disjoint union of the sets $\{B_r(p_i) \cap B_r(p_k) : i \in I_\mathrm{a}$ and $k \in I_\mathrm{b}\}$. There is one other point to note: If $\{\chi_i\}_{i \in \{1,\ldots,N\}}$ is a partition of unity

subordinate to the cover $\{B_r(p_i)\}_{i\in\{1,2,\ldots,N\}}$, then the collection

$$\left\{\beta_\mathrm{a} = \sum_{i\in I_\mathrm{a}} \chi_i\right\}_{\mathrm{a}\in\{1,\ldots,\ell\}} \tag{A.25}$$

is a partition of unity subordinate to the cover $\{\mathcal{U}_\mathrm{a}\}_{\mathrm{a}\in\{1,\ldots,\ell\}}$.

$\underline{\text{Step 5}}$: Supposing that $\alpha \in \Lambda$, define a principal bundle over X to be denoted by P_α with P_α being defined by the following cocycle data. It is given on any $\mathrm{a} \in \{1,\ldots,\ell\}$ version of $\mathcal{U}_\mathrm{a}$ by $\mathcal{U}_\mathrm{a} \times SU(2)$. If $i \in I_\mathrm{a}$ and $k \in I_\mathrm{b}$ are such that $B_r(p_i) \cap B_r(p_k) \neq \emptyset$, then $s_{ki,\alpha}$ is the transition on this component of $\mathcal{U}_\mathrm{a} \cap \mathcal{U}_\mathrm{b}$. The bundle $P^\diamond$ also has a product structure on each set $\mathcal{U}_\mathrm{a}$. If $i \in I_\mathrm{a}$ and $k \in I_\mathrm{b}$ are such that $B_r(p_i) \cap B_r(p_k) \neq \emptyset$, then s_{ki} is the transition on this component of $\mathcal{U}_\mathrm{a} \cap \mathcal{U}_\mathrm{b}$.

$\underline{\text{Step 6}}$: It follows from (A.24) in the case when ε is sufficiently small that Uhlenbeck's Proposition 3.2 and the arguments in her Corollary 3.3 (with Facts 1 and 2 in Appendix A.1) can be invoked using the cover $\{\mathcal{U}_\mathrm{a}\}_{\mathrm{a}\in\{1,\ldots,\ell\}}$ as input to find the following data when the index $\alpha \in \Lambda$ is large: First, a refinement $\{\mathcal{V}_\mathrm{a}\}_{\mathrm{a}\in\{1,\ldots,\ell\}}$; and second, a set of continuous, L_2^2 maps $\{\rho_{\mathrm{a},\alpha} : \mathcal{V}_\mathrm{a} \to SU(2)\}_{\mathrm{a}\in\{1,\ldots,\ell\}}$ that obey

- $s_{ki}\rho_{\mathrm{a},\alpha} = \rho_{\mathrm{b}\alpha}s_{ki,\alpha}$ on $B_r(p_i) \cap B_r(p_k)$ when $i \in I_\mathrm{a}$ and $k \in I_\mathrm{b}$ and $B_r(p_i) \cap B_r(p_k) \neq \emptyset$.

- The L_2^2 and C^0 norms of $\{\rho_{\mathrm{a},\alpha}\}_{\mathrm{a}\in\{1,\ldots,\ell\}}$ are less than ε times an α-independent number.

- The L_1^2 norms of $\{\rho_{\mathrm{a},\alpha}\}_{\mathrm{a}\in\{1,\ldots,\ell\}}$ have limit zero as $\alpha \to \infty$.

$$\tag{A.26}$$

Note in this regard that the third bullet follows from Uhlenbeck's construction and FACTS 1 and 2 in Appendix A.1 given that any (i,k) version of the sequence $\{s_{ki,\alpha}\}_{\alpha\in\Lambda}$ converges strongly to s_{ki} in the L_1^2 topology (because it converges weakly in the L_2^2 topology.) In fact, there is strong convergence to zero in any L_1^p topology for $p < 4$.

$\underline{\text{Step 7}}$: Supposing that the index $\alpha \in \Lambda$ is large, then the collection $\{\rho_{\mathrm{a},\alpha}\}_{\mathrm{a}\in\{1,\ldots,\ell\}}$ define a C^0 and Sobolev class L_2^2 isomorphism $\phi_\alpha : P^\diamond \to P_\alpha$ with an a priori upper bound on its L_2^2 norm. Meanwhile, the bundle P_α is, by definition, isomorphic to P via an isomorphism that identifies A_α on $B_r(p_i)$ for $i \in I_\mathrm{a}$ with $\theta_0 + \mathfrak{a}_{i,\alpha}$. Let $\eta_\alpha : P_\alpha \to P$ denote the latter isomorphism. Appendix A.2 defines a smooth principal bundle P' with a C^0 and Sobolev class L_2^2 isomorphism $\lambda : P' \to P^\diamond$ with the property that λ^{-1} identifies Proposition A.2's connection A on $B_r(p_i)$ with $\theta_0 + \mathfrak{a}_i$. Let

Φ_α denote the isomorphism from P' to P that is obtained by composing first λ, then ϕ_α and then η_α. This is a C^0 and L_2^2 isomorphism with the property that $\{\Phi_\alpha^* A_\alpha\}_{\alpha \in \Lambda}$ is a sequence of L_1^2 connections that converges weakly in the L_1^2 topology to the connection A. Since each Φ_α is C^0 and L_2^2, we can use much the same argument as used in Steps 5 and 6 of the proof of Proposition A.2 to smooth each large α version of Φ_α slightly so that the result is a smooth isomorphism from P' to P (to be denoted by Ψ_α) such that the corresponding sequence $\Psi_\alpha^* A_\alpha$ converges weakly in the L_1^2 topology to the connection A. $\qquad\square$

B   Reducibility and holonomy §

In this appendix, we discuss the relation between reducibility of connections on a principal G-bundle and holonomy in the cases $G = SU(2)$ or $SO(3)$.

As usual, we define the stabilizer Γ_A of A in the gauge group $\mathcal{G}_P$ by

$$\Gamma_A := \{u \in \mathcal{G}_P \mid u(A) = A\}.$$

Given a connection A on a principal G-bundle $P \to X$, there are various conflicting notions of reducibility and irreducibility. Before giving our particular definitions for the cases $G = SU(2)$ or $SO(3)$ in Definition B.1, we review the standard definitions. According to [DK, §4.2.2], a connection A is *reducible* if the holonomy of A is a proper subgroup of G. It is *irreducible* if the stabilizer of A under the group of gauge transformations coincides with the center of G.

According to these definitions, when X fails to be simply-connected, there are sometimes examples of connections which are simultaneously reducible and irreducible. Such an example is provided by any $SU(2)$ connection A_0 whose holonomy is a proper non-abelian subgroup of $SU(2)$. Since the holonomy of A_0 is assumed to be a proper subgroup of $SU(2)$, A_0 is certainly reducible. In order to show that it is also irreducible, we must show that the stabilizer of A_0 coincides with the center $\mathbb{Z}_2$ of $SU(2)$. The stabilizer of A_0 is the centralizer of its holonomy. Being non-abelian, the holonomy is not contained in any $U(1)$ subgroup. Thus the stabilizer of A_0 is a centralizer which is strictly smaller than the centralizer of $U(1)$. The centralizer of any $U(1)$ subgroup of $SU(2)$ is the same $U(1)$ subgroup. The only subgroups of $SU(2)$ which arise as centralizers are $SU(2), U(1)$, and $\mathbb{Z}_2$. Since the only centralizer subgroup of $SU(2)$ smaller than $U(1)$ is the

---

§This appendix is due to the referee. The author is grateful to the referee for allowing him to freely use it in this article.

center $\mathbb{Z}_2$, it follows that A_0 must be irreducible. Such a holonomy subgroup can not occur when X is simply-connected, since in this case the holonomy of any connection must be a closed and connected Lie subgroup. Thus when $G = SU(2)$ and X is simply-connected, there are no intermediate holonomy subgroups between $U(1)$ and $SU(2)$.

The most useful definition of reducibility for Vafa–Witten theory over non-simply-connected manifolds X is the following, which we use in this article.

Definition B.1. A connection A on a principal $SU(2)$ or $SO(3)$ bundle $P \to X$ is *reducible* if one of the following equivalent conditions is satisfied.

- The stabilizer of A under the group of gauge transformations has positive dimension.

- There exists a nonzero $\Gamma \in \Omega^0(X; \mathfrak{g}_P)$ such that $d_A\Gamma = 0$.

- The holonomy of A is contained in some $SO(2)$ subgroup.

A connection is said to be *irreducible* if it is not reducible.

Equivalence of the first two conditions in Definition B.1 follows from the fact that the Lie algebra of the stabilizer is the kernel of $d_A : \Omega^0(X; \mathfrak{g}_P) \to \Omega^1(X; \mathfrak{g}_P)$. Equivalence of the last two conditions in Definition B.1 follows from the decomposition $\mathfrak{g}_P = \underline{\mathbb{R}} \oplus C$, where $\underline{\mathbb{R}}$ is the span of Γ, and C is the complementary $SO(2)$ subbundle.

Remark B.2. When $G = SU(2)$, this definition of reducibility is equivalent to the holonomy of A being abelian. However, there are irreducible abelian connections when $G = SO(3)$. For example, if A preserves a splitting of $\mathfrak{g}_P$ as $\mathfrak{g}_P = \mathcal{I}_1 \oplus \mathcal{I}_2 \oplus (\mathcal{I}_1 \otimes \mathcal{I}_2)$ with each summand being a nontrivial real line bundle, then the holonomy is $\mathbb{Z}_2 \times \mathbb{Z}_2$, which can not arise as a subgroup of $SO(2)$.

With regards to the local reducibility introduced in Section 2 of this article (Definition 2.1), we have the following.

Proposition B.3. *A connection A on a principal $SU(2)$ or $SO(3)$ bundle $P \to X$ is locally reducible if and only if the holonomy of A is contained in some $O(2)$ subgroup.*

Proof. Every $O(2)$ subgroup of $SU(2)$ or $SO(3)$ is conjugate to the set of matrices of the form

$$\begin{pmatrix} * & 0 \\ 0 & * \end{pmatrix} \text{ and } \begin{pmatrix} 0 & * \\ * & 0 \end{pmatrix} \subset SU(2), \quad \begin{pmatrix} * & * & 0 \\ * & * & 0 \\ 0 & 0 & * \end{pmatrix} \subset SO(3),$$

which is the subgroup such that the adjoint action preserves the subspace

$$\sqrt{-1}\mathbb{R}\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \subset \mathfrak{su}(2), \quad \mathbb{R}\begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \subset \mathfrak{so}(3).$$

Whenever a 1-dimensional A-covariantly constant subbundle of $\mathfrak{g}_P$ exists, the holonomy must fix the subbundle, and hence must be contained in such a $O(2)$ subgroup. Conversely if holonomy is contained in an $O(2)$ subgroup, parallel transport of the fixed subspace gives the desired subbundle of $\mathfrak{g}_P$.

$\square$

Remark B.4. If A is locally reducible, then the restriction of A to any simply-connected subset of X is reducible. This is because the holonomy subgroup over a simply-connected region must be connected, and hence contained in an $SO(2)$ subgroup.

Remark B.5. If A is a locally reducible $SU(2)$ or $SO(3)$ connection, then the following conditions are equivalent.

- A is reducible.

- One of the rank one subbundles $\mathcal{I}$ of $\mathfrak{g}_P$ which is preserved by A admits a trivialization.

- The first Stiefel–Whitney class vanishes for some rank one subbundle $\mathcal{I}$ which is preserved by A.

Whenever $\pi_1(X)$ has no subgroup of index two, it follows that $H^1(X;\mathbb{Z}_2) = 0$ and thus every locally reducible connection is reducible.

References

[C]   S. Costea, *Sobolev capacity and Hausdorff measures in metric measure spaces*, Annales Academiae Scientiarum Fennicae Mathematica **34** (2009), 179–194.

[D]   S. K. Donaldson, *Symplectic submanifold in almost complex geometry*, Jour. Diff. Geom. **44** (1996), 666–705.

[DK]   S. K. Donaldson and P. B. Kronheimer, *The Geometry of Four-Manifolds*, Oxford University Press, 1990.

[Ha]   A. Haydys, *Fukaya–Seidel category and gauge theory*, J. Symplectic Geom. **13** (2015), 151–207.

[Hi]   N. J. Hitchin, *The self-duality equations on a Riemann surface*, Proc. London Math. Soc. **55** (1987), 59–126.

[HHL]   Q. Han, R. Hardt and F. Lin, *Geometric measure of singular sets of elliptic equations*, Comm. Pure Appl. Math. **51** (1998), 1425–1443.

[M]   B. Mares, *Some Analytic Aspects of Vafa–Witten Twisted $\mathcal{N} = 4$ Supersymmetric Yang–Mills theory*, Ph.D thesis, M.I.T., 2010.

[R]   Y. G. Reshetnyak, *The concept of capacity in the theory of functions with generalized derivatives*, Sibirskii Matematicheskii Zhurnal, **10** (1969), 1109–1138.

[T1]   C. H. Taubes, $PSL(2;\mathbb{C})$-*connections on 3-manifolds with L^2 bounds on curvature*, Camb. J. Math. **1** (2013), 239–397; and *Corrigendum to "$PSL(2;\mathbb{C})$ connections on 3-manifolds with L^2 bounds on curvature"*, Camb. J. Math. **3** (2015), 619–631.

[T2]   C. H. Taubes, *Compactness theorems for $SL(2;\mathbb{C})$ generalizations of the 4-dimensional anti-self-dual equations*, arXiv:1307.6447v4.

[T3]   C. H. Taubes, *The zero loci of $\mathbb{Z}/2$ harmonic spinors in dimension 2, 3 and 4*, arXiv:1407.6206.

[T4]   C. H. Taubes, *On the behavior of sequences of solutions to $U(1)$ Seiberg–Witten systems in dimension 4*, arXiv:1610.07163.

[T5]   C. H. Taubes, *private communications*.

[U]   K. K. Uhlenbeck, *Connections with L^p bounds on curvature*, Comm. Math. Phys. **83** (1982), 31–42.

[V]   J. Väisälä, *Capacity and Measure*, Michigan Math. Journal **22** (1975), 1–4.

[VW]   C. Vafa and E. Witten, *A strong coupling test of S-duality*, Nucl. Phys. B, **432**, (1994), 484–550.

[Wi]   E. Witten, *Fivebranes and Knots*, Quantum Topol. **3** (2012), 1–137.

Graduate School of Mathematics, Nagoya University, Furo-cho,
Chikusa-ku, Nagoya, 464-8602, Japan
yu2tanaka@gmail.com