

EXISTENCE OF mASD CONNECTIONS ON 4-MANIFOLDS WITH CYLINDRICAL ENDS

DAVID L. DUNCAN AND IAN HAMBLETON

ABSTRACT. Taubes' gluing theorems establish the existence of ASD connections on closed, oriented 4-manifolds. We extend these gluing results to the mASD connections of Morgan–Mrowka–Ruberman on oriented 4-manifolds with cylindrical ends. As a corollary, we obtain an ASD-existence result in the presence of degenerate asymptotic flat connections.

1. INTRODUCTION

The results of Taubes [19, 20] on “gluing” establish the existence of non-trivial anti-self dual (ASD) connections on closed, oriented 4-manifolds, provided one works with an $SU(2)$ -bundle with sufficiently high second Chern class. This was extended by Donaldson [3] to a general gluing theorem for connected sums; see also [2, 9]. These gluing results have direct extensions to cylindrical end 4-manifolds, provided one works with ASD connections having a *non-degenerate* flat connection as an asymptotic limit [4]. However, in the absence of such non-degeneracy assumptions, the space of ASD connections on a cylindrical end 4-manifold is generally not well-controlled (e.g., the ASD operator is not Fredholm) and this now-standard gluing formalism breaks down. Nevertheless, the question of existence for ASD connections in this degenerate cylindrical end setting remains well-posed. One of our main results, Theorem A below, establishes one such ASD-existence result in the degenerate setting.

To state this, suppose X is a connected, oriented 4-manifold with cylindrical ends. Thus, we can write $X = X_0 \cup \text{End } X$, where X_0 is a compact 4-manifold with boundary N , and $\text{End } X \cong [0, \infty) \times N$ is diffeomorphic to a cylinder. We refer to X_0 as the *compact part* and to $\text{End } X$ as the *cylindrical ends*. Unless otherwise stated, we allow the case where N has multiple components, or is empty. Fix a metric g on X that is asymptotically cylindrical in the sense described in Section 2A.

Theorem A. *Assume $b^+(X) \leq 1$. Assume further that the 3-manifold N is connected and satisfies one of the following:*

- (i) *N is a circle bundle over a surface with positive Euler class: $e(N) > 0$; or*
- (ii) *N has first Betti number at most one: $b_1(N) \leq 1$.*

Date: February 28, 2023.

Research partially supported by NSERC Discovery Grant A4000. The authors would like to thank the Fields Institute for Research in Mathematical Sciences in Toronto, and the Max Planck Institut für Mathematik in Bonn for their hospitality and support while parts of this work were done.

Then, for any integer $\ell \geq b^+(X) + 1$, the manifold X admits an irreducible ASD-connection A on a principal $SU(2)$ -bundle over X , and A satisfies

$$\int_X |F_A|^2 \, d\text{vol} = 8\pi^2 \ell.$$

We prove this in Section **6D**. As a concrete example, the hypotheses of Theorem A hold when X_0 is diffeomorphic to the total space of a positive Euler class disk bundle over a surface. To the authors' knowledge, Theorem A (and its extension, Theorem 6.9) is the first general ASD existence result for cylindrical end manifolds that allows for a degenerate flat limit down the end.

Our approach to Theorem A is to (locally) embed the space of ASD connections into the larger space of *modified ASD* (mASD) connections of Morgan, Mrowka, and Ruberman [17]. This larger space is obtained by modifying the ASD operator in such a way that one obtains a Fredholm operator whose zero set contains an open set in the space of finite-energy ASD connections; it may also contain some new solutions. It is shown in [17] that, by allowing the auxiliary choices in this construction to vary, every finite-energy ASD connection belongs to some mASD space of connections defined in this way. The other main results of the present paper, stated below, show that the gluing results of Taubes and Donaldson for connected sums have extensions to this mASD setting. We then arrive at Theorem A as a consequence of these mASD-gluing results; the topological hypotheses on N imply that the mASD connections thus obtained are in fact ASD.

Before stating these mASD-gluing results, we give several remarks to help provide further context for this mASD setting.

Remark 1.1. (a) Our primary motivation for developing these gluing results was to use the Morgan–Mrowka–Ruberman “moduli space” of mASD connections to study the action on X of a finite group π . Even in the ASD setting, generic perturbations are usually not π -equivariant, so the standard transversality arguments are not available, and one must appeal to some other approach to handle singularities in the moduli space. As a sequel to this paper, we planned to study the π -equivariant compactification of the “mASD moduli space” as was done in [11], [12], and [13] for the ASD moduli space.

Unfortunately, the mASD operator *fails* to be gauge equivariant in any reasonable sense (see Remark 2.12). This appears to be an oversight in the original text [17] (e.g., see [17, p. 125]), and at present we do not know how to define a suitable gauge quotient of the space of mASD connections that one might call the “mASD moduli space”. It is a fundamental and interesting open problem to construct an appropriate mASD-replacement for the ASD moduli space. See Section **2B.5** for more information about this issue.

(b) The foundational work of Mrowka [18], Morgan–Mrowka–Ruberman [17], and Taubes [21, 22] concerning instantons on cylindrical end 4-manifolds was done shortly before the Seiberg–Witten revolution in gauge theory. One of their striking results in this setting is that a finite-energy ASD connection has a well-defined limiting flat connection on the 3-manifold N “at infinity”.

At that time, a central problem was to understand the behaviour of ASD connections under neck-stretching within a closed 4-manifold, as well as the reverse operation in which ASD connections on non-compact 4-manifolds with matching data on their cylindrical ends could be glued together. Indeed, the authors of [17, p. 12] state: “The [mASD] moduli space seems to provide the correct geometric context for a general gluing theorem for ASD connections, although we do not treat this topic in this book” (see Remark 7.2). This point of view was a main ingredient in a paper of Fintushel and Stern [7] (and in unpublished work of Morgan and Mrowka [16]). An account of gluing along cylindrical ends from the perspective of Floer homology was later provided by Donaldson [4], simplified by assuming the presence of a perturbation to avoid degeneracies (see (c), below). We note, however, that the gluing results of the present paper take place on the compact part X_0 , and not on the ends.

(c) Researchers have worked around the technical issues involved in gluing in the degenerate setting by various methods. Of these methods, one of the most popular is to perturb the ASD equation on the ends in such a way that all perturbed-ASD connections are asymptotic to non-degenerate perturbed-flat connections [8], [4]. However, this approach has several drawbacks. For one, ASD connections are generally not solutions of perturbed-ASD equations of this type; this can obscure the geometric information one can infer from an abstract existence result for perturbed-ASD connections (e.g., to what extent do these connections depend on the perturbation?). Another drawback is that these perturbation schemes are not well-behaved in the presence of reducible flat connections (e.g., the trivial flat connection), and this limits the applicability of such approaches. For example, a full $SU(2)$ -instanton Floer theory for 3-manifolds N with $b_1(N) \geq 1$ is still lacking, and even the existing instanton Floer theory for integer homology spheres handles the trivial flat connection separately. In summary, a more in-depth understanding of ASD connections with degenerate limits is desired, and we view the results of this paper as being a step in that direction.

To state our gluing results for mASD connections, let G be a compact Lie group and fix a principal G -bundle $E \rightarrow X$. We assume that E is translation-invariant on the end; that is, we assume the diffeomorphism $\text{End } X \cong [0, \infty) \times N$ is covered by a bundle isomorphism $E|_{\text{End } X} \cong [0, \infty) \times Q$ for some principal G -bundle $Q \rightarrow N$. We also fix a flat connection Γ on Q .

Given a connection A on E that converges sufficiently fast down the end, one can define a quantity

$$\kappa(E, A|_{\text{End } X}) := -\frac{1}{8\pi^2} \int_X \langle F_A \wedge F_A \rangle \in \mathbb{R}$$

that we call the *relative characteristic number of the adapted bundle* $(E, A|_{\text{End } X})$; see Section 6A for more details. If A is ASD, then $\kappa(E, A|_{\text{End } X}) = (8\pi^2)^{-1} \int_X |F_A|^2$ equals the usual energy of the connection A . The upshot for us is that the quantity $\kappa(E, A|_{\text{End } X})$ is well-defined for a much larger class of connections than those with finite energy. Indeed, this relative characteristic number depends only on the topological type of the adapted bundle $(E, A|_{\text{End } X})$, and it is a lift of the Chern–Simons value of the connection on Q to which A is asymptotic. Note that if $\kappa(E, A|_{\text{End } X}) \neq 0$, then A is not flat. When X is

closed, then this relative characteristic number is actually an integer that depends only on E , and we will simply write it as $\kappa(E)$ (e.g., if $G = \mathrm{SU}(r)$, then $\kappa(E) = c_2(E)[X]$ is the second Chern number). We will primarily use $\kappa(E, A|_{\mathrm{End} X})$ to keep track of the topological data in our gluing operations, just as the second Chern class keeps track of the underlying bundle type when gluing in the standard $\mathrm{SU}(2)$ -setting for ASD connections on closed 4-manifolds.

By making several auxiliary choices, collectively called *thickening data*, one can define the *modified ASD (mASD) operator*, which is a non-linear Fredholm map s defined on a suitable space of connections on E (see Section 2 for definitions). In particular, we note that this space of connections is defined so that all elements are asymptotic to connections close to the flat connection Γ fixed above. By definition, the mASD connections are those in the zero set of s , and we say that an mASD connection A is *regular* if the linearization of s at A is surjective when restricted to a Coulomb slice.

For $k = 1, 2$, suppose X_k is an oriented, cylindrical end 4-manifold equipped with a principal G -bundle $E_k \rightarrow X_k$ and thickening data, as above. Let $X = X_1 \# X_2$ be a connected sum of these manifolds, taken at points in the compact parts of the X_k . Then the E_k can be used to form a connected sum bundle $E \rightarrow X$, and we equip this with the thickening data induced from those of the E_k ; see Section 3A. Our basic gluing result can be stated as follows.

Theorem B. *For $k = 1, 2$, suppose A_k is a regular mASD connection on E_k . Then for any $\epsilon > 0$, the bundle $E = E_1 \# E_2$ admits an mASD connection A with the property that the distance between $A|_{X_k \cap X}$ and $A_k|_{X_k \cap X}$ is less than ϵ for $k = 1, 2$. Moreover,*

$$(1.2) \quad \left| \kappa(E, A|_{\mathrm{End} X}) - \sum_{k=1}^2 \kappa(E_k, A_k|_{\mathrm{End} X}) \right| < \epsilon.$$

In the statement of Theorem B, the distance is relative to a $L_2^2(N) \times L_\delta^{p^*}(X)$ -metric on the space of connections; see (3.6) for a precise statement (the connection A of Theorem B is what is called $\mathcal{J}(A_1, A_2)$ in (3.6)). Theorem B is a special case of Theorem 3.3, which works in the broader setting where the A_k are not necessarily regular. In this broader setting, the connection A need not be mASD, but its failure to be mASD is expressed through an obstruction map. In Theorem 5.1, we extend Theorem B to a gluing result for *families* of regular mASD connections. These results are mASD-extensions of results familiar from the ASD setting; see [5, Section 7.2].

As an application of Theorems B and 5.1, we establish the following existence result, extending that of Taubes [19, 20] to the present cylindrical end mASD situation.

Theorem C. *Assume $G = \mathrm{SU}(2)$ and $b^+(X) \leq 1$, and fix an integer $\ell \geq b^+(X) + 1$. Then for every $\epsilon > 0$, there is a principal $\mathrm{SU}(2)$ -bundle $E \rightarrow X$ and an mASD connection A on E that is irreducible, and satisfies*

$$(1.3) \quad |\kappa(E, A|_{\mathrm{End} X}) - \ell| < \epsilon.$$

If $b^+(X) = 0$, then the connection A is regular.

The cases $b^+(X) = 0$ and $b^+(X) = 1$ are special cases of Theorem 6.2 and Theorem 6.3, respectively. Structurally, our proof strategies for these are very similar to the

analogous statements in the closed case [19, 20] by realizing X as a trivial connected sum $X \cong X \# S^4$. Under the assumption that $b^+(X) = 0$, it follows that the trivial flat connection on X is regular as an mASD connection (see Remark 2.19). It is well-known that the 4-sphere admits irreducible ASD connections of every positive second Chern class, and these are necessarily regular for topological reasons. Then Theorem C for $b^+(X) = 0$ follows from the general gluing result of Theorem B and adjacent results designed to handle gauge transformations (more below). We note also that Theorem 6.2 (the more general version of Theorem C) is proved for an arbitrary compact Lie group G , under mild hypotheses on ℓ and G .

Remark 1.4. If X is simply-connected and $b^+(X) = 0$, we expect that a modified gluing construction will produce an open subset of the space of mASD connections that are in Coulomb gauge relative to some fixed connection. The issues involved in carrying out this improvement are briefly indicated in Remark 5.3(b).

The strategy for our proof of Theorem C when $b^+(X) = 1$ is similar, albeit more involved since the trivial flat connection on X is no longer regular. Thus a careful analysis of the obstruction map of Theorem 3.3 is required. Just as in [20], we glue ASD connections on S^4 at several sites instead of one, and this is sufficient to show that the obstruction vanishes for some choice of gluing parameters. In this analysis, we use the assumption that $G = \text{SU}(2)$. As Taubes mentions [20, p. 518], it is likely that the restriction to $G = \text{SU}(2)$ can be removed, but that would call for a different approach.

We prove our general existence results only for $b^+ \leq 1$ because (i) these are the cases of interest for our applications, and (ii) extending the discussion to higher values of b^+ would add considerable length to the paper (this can already be seen in [20]). For similar reasons we also carry out our analysis with connections that are locally in Sobolev class L_1^p as opposed to, say, L_k^p for $k \geq 1$. We do not see any inherent obstruction to extending our results to higher Sobolev spaces and presumably such extensions would recover our L_1^p -results by elliptic regularity. We leave the details of such extensions to interested parties.

The appearance of $\epsilon > 0$ in the statements of Theorems B and C is new to this mASD setting. To explain it, we note that in the standard set-up of gluing ASD connections on a closed 4-manifold, the inequality (1.2) would be replaced by the equality $\kappa(E_1 \# E_2) = \kappa(E_1) + \kappa(E_2)$; likewise (1.3) would be replaced by $\kappa(E) = \ell$. The presence of an *inequality* for us reflects a need to freely vary the asymptotic values in order to obtain the mASD connection A . Indeed, as discussed further in Remark 2.19, the ability to freely vary these asymptotic values is at the heart of what makes the mASD set-up a viable candidate for the type of existence statement in Theorem C and thus Theorem A. For example, when $b^+(X) = 0$, the trivial flat connection is regular only because the mASD operator allows for this variation in the asymptotic values.

In Section 7, we have included a discussion of how Theorem C for $b^+(X) = 0$ provides a “partial compactification” of the space of mASD connections. We also discuss why this compactification is only partial, and what a more complete compactification would require.

As mentioned above, the lack of gauge-equivariance for the mASD operator means that we are *not* free to pass to the quotient modulo gauge. Indeed, to obtain a Fredholm problem for the gluing constructions, we work entirely within a fixed Coulomb slice. Since the natural Coulomb slice varies as the connections vary, this dependence becomes relevant when we glue over families of connections, which is necessary for Theorem C. This is a central obstacle with which we must contend in the present paper: In the usual ASD setting, one could apply suitable gauge transformations that put all nearby ASD connections into the same slice. However, in this mASD setting, the gauge-transformed mASD connections would no longer be mASD. To account for this, we establish a pair of gauge fixing results, Proposition 4.3 and Theorem 4.5, that show that, by making an additional perturbation, an mASD connection in one Coulomb slice can be perturbed to an mASD connection in a nearby Coulomb slice.

Apart from the failure of gauge equivariance in the mASD setting, the main difference between the mASD and ASD settings is that we now need to handle the additional nonlinearities that arise from the term modifying the ASD operator. The key observation we use for handling this term is that it factors through a finite-dimensional manifold.

Finally, we mention that if Γ is non-degenerate, then every mASD connection with asymptotic value near Γ is in fact ASD. E.g., this non-degeneracy hypothesis is satisfied when N is a rational homology 3-sphere and Γ is the trivial connection. As such, our results recover standard gluing results for ASD connections on cylindrical end manifolds with non-degenerate asymptotic limits; see Sections 2C.2 and 6D for more details. More interestingly, there are situations for which Γ is degenerate, but for which every mASD connection with asymptotic value near Γ is ASD. In such cases, our mASD gluing theorem produces an ASD connection. Theorem A is one result of this type.

CONTENTS

1. Introduction	1
2. Background on the thickened moduli space	7
2A. Auxiliary choices	7
2B. Gauge theory	10
2C. Special cases	17
3. Gluing two mASD connections	21
3A. Set-up for gluing	21
3B. Gluing two connections	23
3C. Extensions to $p < 2$	34
4. Gauge fixing and the mASD condition	36
4A. Gauge fixing	36
4B. Recovering the mASD condition within a slice	38
5. Gluing regular families	41
5A. Immersion lemmas	44
5B. Proof of Theorem 5.1	50
6. Existence results	51

6A.	Relative characteristic classes and adapted bundles	52
6B.	Existence when $b^+(X) = 0$	53
6C.	Existence when $b^+(X) = 1$	55
6D.	An ASD existence result and a proof of Theorem A	58
7.	Partial compactification—the Taubes boundary	60
7A.	The Taubes Boundary	60
7B.	Compactification Issues	62
	References	63

2. BACKGROUND ON THE THICKENED MODULI SPACE

In this section we give a rapid review of the relevant background material from [17]; we also expand on some of the results of [17], which will assist in our discussion of gluing below. With a few exceptions, we use much of the same notation and set-up established in [17]. To allow for a more streamlined discussion, we assume throughout that the 3-manifold end N is nonempty; however, see Section 2C for an extension to the case $N = \emptyset$.

We will write $\mathcal{A}(E)$ and $\mathcal{G}(E)$ for, respectively, the spaces of smooth connections and gauge transformations on $E \rightarrow X$. When the bundle is clear from context, we will simply write $\mathcal{A}(X)$ and $\mathcal{G}(X)$. Given a connection A , we denote by F_A its curvature, which is a 2-form on X with values in the adjoint bundle \mathfrak{g}_E . We will write $\Omega^\ell(X)$, and sometimes Ω^ℓ , for the space of smooth adjoint bundle-valued ℓ -forms on X that are rapidly decaying.

To touch base with constants associated with characteristic classes below, we work relative to an inner product on \mathfrak{g} obtained as follows. Fix a Lie group homomorphism

$$(2.1) \quad G \longrightarrow \mathrm{SU}(r)$$

that is also an immersion. Then the induced map $\mathfrak{g} \hookrightarrow \mathfrak{su}(r)$ is an embedding of Lie algebras. Let $\langle \cdot, \cdot \rangle : \mathfrak{g} \otimes \mathfrak{g} \rightarrow \mathbb{R}$ denote the inner product on \mathfrak{g} obtained by pulling back the inner product $A \otimes B \mapsto -\mathrm{tr}(AB)$ on $\mathfrak{su}(r)$. This inner product is Ad-invariant, and so determines a metric on the adjoint bundle \mathfrak{g}_E .

Notation such as $L_k^p(\Omega^\ell(X), g)$ will denote the L_k^p -Sobolev completion of $\Omega^\ell(X)$, relative to a metric g on X and the above-defined metric on \mathfrak{g}_E . When X or g are clear from context, or not relevant, we may drop them from the notation.

2A. Auxiliary choices.

2A.1. The center manifold. Recall from the introduction that we have fixed a bundle $Q \rightarrow N$ as well as a smooth flat connection Γ on Q . Fix a metric g_N on N . Let $U_\Gamma \subseteq L_2^2(\mathcal{A}(N))$ be a coordinate patch centered at Γ , in the sense of [17, Def. 2.3.1]; for our purposes, it suffices to know that U_Γ is a small open neighborhood of Γ in the Coulomb slice $\{\Gamma\} + \ker(d_\Gamma^*)$. As in [17, Lemma 2.5.1], there is a unique $\mathrm{Stab}(\Gamma)$ -equivariant map

$$\Theta : U_\Gamma \longrightarrow L_2^2(\Omega^0(N))$$

with $\Theta(a) \in (\ker \Delta_\Gamma)^\perp$ and

$$d_\Gamma^*(F_a - d_a \Theta(a)) = 0.$$

It follows from this last equation, and uniqueness, that if a has higher regularity then so too does $\Theta(a)$.

We will be interested in the densely-defined vector field

$$\nabla f_\Gamma: U_\Gamma \longrightarrow TU_\Gamma \quad a \longmapsto \nabla f_\Gamma(a) := - * F_a + d_a \Theta(a).$$

Note that the zeros of ∇f_Γ are precisely the flat connections in U_Γ . (As described in [17, Lemma 2.5.1(1)], this vector field is the (negative) gradient of the restriction to U_Γ of the Chern–Simons functional, where the gradient is taken relative to a certain inner product that takes into account the possibility of a non-trivial stabilizer of Γ .)

For $m \geq 2$, let $\mathcal{H} = \mathcal{H}_\Gamma \subseteq U_\Gamma$ be a $\text{Stab}(\Gamma)$ -invariant \mathcal{C}^m -center manifold for the vector field ∇f_Γ , as in [17, Cor. 5.1.4]. In particular, this means that

- \mathcal{H}_Γ is a finite-dimensional \mathcal{C}^m -manifold containing Γ ,
- the tangent space to \mathcal{H}_Γ at Γ is the Γ -harmonic space

$$H_\Gamma^1 := \ker(d_\Gamma \oplus d_\Gamma^*) \subseteq \Omega^1(N),$$

- ∇f_Γ is tangent to \mathcal{H}_Γ , and
- every zero of ∇f_Γ sufficiently close to Γ is contained in \mathcal{H}_Γ .

We denote by $\Xi = \Xi_\Gamma$ the restriction of ∇f_Γ to \mathcal{H}_Γ .

Fix a compactly supported cutoff function $\beta: \mathcal{H} \rightarrow [0, 1]$ that is identically 1 near Γ . The *trimmed vector field* is given by

$$\Xi^{tr}(h) := \beta(h)\Xi(h).$$

Set $\mathcal{H}_{in} = \beta^{-1}(1)$ and $\mathcal{H}_{out} = \beta^{-1}((0, 1])$.

Fix a real number $T \geq 1$. The trimmed vector field is complete and so, for each $h \in \mathcal{H}$, there is a unique solution $h_T: \mathbb{R} \rightarrow \mathcal{H}$ to the flow

$$\frac{d}{dt}h_T(t) = \Xi^{tr}(h_T(t)) \quad h_T(T) = h.$$

We then set

$$\alpha(h) = \alpha_T(h) := h_T(t) + \Theta(h_T(t))dt.$$

Depending on context, we may view α as a connection on the submanifold $\text{End } X \cong [0, \infty) \times N$, or on the cylinder $\mathbb{R} \times N$.

Lemma 2.2. *For all $h \in \mathcal{H}$, the connection $\alpha(h)$ is in $L_{2,loc}^2(\mathbb{R} \times N) \cap \mathcal{C}^0(\mathbb{R} \times N)$, and hence in $L_{1,loc}^p(\mathbb{R} \times N) \cap \mathcal{C}^0(\mathbb{R} \times N)$ for any $1 \leq p < 4$. Moreover, the map $\alpha: h \mapsto \alpha(h)$ is \mathcal{C}^m relative to the $L_2^2(N)$ -topology on the domain and the $\mathcal{C}^0(\mathbb{R} \times N)$ -topology on the codomain.*

Proof. The initial condition h is in $L_2^2(N) \subset \mathcal{C}^0(N)$, by assumption. It then follows from standard regularity arguments for flows that the path h_T is in $L_{2,loc}^2 \cap \mathcal{C}^0$ on $\mathbb{R} \times N$. Hence $\alpha(h)$ is in the same space as well, since the regularity of $\Theta(h_T)$ is controlled by that of h_T . That α is \mathcal{C}^m relative to these topologies follows from a similar argument

applied to its k th derivative for $1 \leq k \leq m$. The assertion about L_1^p follows from the embedding $L_{2,loc}^2 \hookrightarrow L_{1,loc}^p$, which holds provided $1 \leq p < 4$. \square

In \mathcal{H}_{in} , the flow defining $h_T(t)$ is gauge equivalent to the Chern–Simons gradient flow. It follows that if $h_T(t) \in \text{int}(\mathcal{H}_{in})$ lies in the interior for some t , then $\alpha(h)$ is ASD in a neighborhood of $\{t\} \times N$. This has the following useful linear analogue for the linearization $(D\alpha)_\Gamma$ of α at Γ .

Lemma 2.3. *If $\eta \in T_\Gamma \mathcal{H} = H_\Gamma^1$, then*

$$(D\alpha)_\Gamma \eta = \pi^* \eta$$

where $\pi : \mathbb{R} \times N \rightarrow N$ is the projection.

Proof. Since Γ is flat, we have $\Theta(\Gamma) = 0$ and so $\alpha(\Gamma) = \pi^* \Gamma$. This is a flat connection on $\mathbb{R} \times N$ and so it is ASD:

$$F_{\alpha(\Gamma)}^+ = F_{\pi^* \Gamma}^+ = 0.$$

Linearizing this in the direction of $\eta \in T_\Gamma \mathcal{H}$ implies

$$d_{\pi^* \Gamma}^+((D\alpha)_\Gamma \eta) = 0.$$

Note that, since η is Γ -harmonic, we also have

$$d_{\pi^* \Gamma}^+(\pi^* \eta) = 0.$$

In general, we can view the kernel of d_A^+ on the cylinder $\mathbb{R} \times N$ as defining a flow on the space $\Omega^1(N) \times \Omega^0(N)$. In the case of $(D\alpha)_\Gamma \eta$ and $\pi^* \eta$, these both take values in the graph

$$\text{Graph}((D\Theta)_\Gamma|_{T_\Gamma \mathcal{H}}) \subseteq T_\Gamma \mathcal{H} \times L_2^2(\Omega^0(N))$$

of the linearization of Θ . Observe two things: (i) this graph is finite-dimensional, and (ii) the 1-forms $(D\alpha)_\Gamma \eta$ and $\pi^* \eta$, viewed as paths in the graph, both equal η at time T . Then $(D\alpha)_\Gamma \eta = \pi^* \eta$ follows by the uniqueness for flows on finite-dimensional spaces. \square

2A.2. The choice of metric. We will use $t : \text{End } X \rightarrow [0, \infty)$ to denote the projection relative to the identification $\text{End } X \cong [0, \infty) \times N$. With the use of a cutoff function, we can view t as a smooth real-valued function defined on all of X , which we will denote by the same symbol.

Fix a smooth cylindrical end metric g_0 on X ; this means that the restriction

$$g_0|_{\text{End } X} = dt^2 + g_N$$

is a product metric, where g_N is the fixed metric on N . Let B be a \mathcal{C}^3 -neighborhood of g_0 in the space of $\mathcal{C}^{\max(m,3)}$ -metrics on X so that the conclusions of [17, Theorem 2.6.3] hold (the proof of Theorem 2.6.3 shows that such a set exists; the details of the theorem will not play an active role in the discussion that follows). Let μ_Γ^\pm be the smallest positive eigenvalue of $\mp * d_\Gamma : \Omega^1(N) \rightarrow \Omega^1(N)$. The sign convention here is to agree

with that of [17, Def. 2.1.1]. Then we will say that a metric g on X is *asymptotically cylindrical* if $g \in B$ and

$$\|g - g_0\|_{\mathcal{C}^1(\{t\} \times N)} \leq e^{-\max(\mu_\Gamma^-, \mu_\Gamma^+)t}$$

for all $t \geq 0$. (This is effectively Condition A3 of [17, p. 116].) Throughout, we will always assume our metrics are asymptotically cylindrical in this sense. Note that every cylindrical end metric is automatically asymptotically cylindrical.

2A.3. Thickening data. Fix data as in [17, Section 7.2]; we will refer to this as the *thickening data* and denote it by \mathcal{T}_Γ . In particular, this includes the choice of positive numbers ϵ_0 and δ , that we will describe momentarily. The details of the remaining data in \mathcal{T}_Γ will not play an active role in our discussion. For convenience, we also assume that \mathcal{T}_Γ includes the choice of the fixed $T \geq 1$ from above.

The key feature for us regarding ϵ_0 is that $|\mathcal{CS}(h_2) - \mathcal{CS}(h_1)| < \epsilon_0/2$ for all $h_1, h_2 \in \text{supp}(\beta)$, where \mathcal{CS} is the Chern–Simons function. For any $\epsilon_0 > 0$, this inequality can be arranged by shrinking the support of β , if necessary. The remaining requirements for ϵ_0 will not be directly relevant to us, but see [17, Definition 4.3.2] for more details. As for δ , we assume $\delta > \mu_\Gamma^-$ and that $\delta/2$ is not an eigenvalue of $-*d_\Gamma$. By shrinking the size of the coordinate patch U_Γ , if necessary, we may assume further that $\delta/2$ is not an eigenvalue of $-*d_a$ for any $a \in U_\Gamma$. At various times, we may place additional restrictions on δ .

2A.4. Weighted spaces. We define the space $L_{k,\delta}^p(X)$ to be the completion of the set of compactly supported smooth forms f on X , relative to the weighted Sobolev norm:

$$\|f\|_{L_{k,\delta}^p} := \|e^{\delta t/2} f\|_{L_k^p}$$

When $p = 2$, this recovers the family of norms used in [17]. The subspace of ℓ -forms will be denoted by $L_{k,\delta}^p(\Omega^\ell)$ or $L_{k,\delta}^p(\Omega^\ell(X))$. Following standard conventions, when $k = 0$, we will write L_δ^p for $L_{0,\delta}^p$. Note that the norm $\|f\|_{L_{k,\delta}^p}$ is equivalent to the norm:

$$\sum_{0 \leq j \leq k} \|e^{\delta t/2} \nabla^j f\|_{L^p}$$

In particular, we can use this equivalence to transfer Sobolev embedding results for L_k^p to the weighted setting; e.g., see the proof of Lemma 3.9.

2B. Gauge theory.

2B.1. The space of connections. For $1 \leq p < 4$, define $\mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ to consist of the connections A on E satisfying the following:

- A has regularity $L_{1,loc}^p$,
- there is some $h \in \mathcal{H}_{out}$ so that $A - \alpha(h) \in L_{1,\delta}^p(\Omega^1(\text{End } X))$,
- for each $t \geq T$, the connection $A|_{\{t\} \times N}$ is gauge equivalent to a connection in the coordinate patch U_Γ centered at Γ .

This space of connections is generally *not* an affine subspace of $L_{1,loc}^p(\mathcal{A}(E))$; this reflects the nonlinearities in the definition of the map $h \mapsto \alpha(h)$. We give $\mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ the structure of a \mathcal{C}^m -Banach manifold, as in [17, Section 7.2.2]. (Equivalently, this \mathcal{C}^m -Banach manifold structure is precisely the one for which the map ι , defined in (2.5) below, is a \mathcal{C}^m -diffeomorphism.) By [17, Prop. 7.2.3] (see also [17, p. 120]), given $A \in \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$, the element $h \in \mathcal{H}_{out}$ from the second bullet point is uniquely determined; this uses the assumption that $\delta > \mu_\Gamma^-$. As such, there is a well-defined map

$$p_T: \mathcal{A}^{1,p}(\mathcal{T}_\Gamma) \longrightarrow \mathcal{H}_{out}$$

that is \mathcal{C}^m -smooth.

Remark 2.4. (a) Note that our space $\mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ consists of connections with weaker regularity than the one in [17, Ch. 7], which is modeled on L_2^2 instead of L_1^p . This changes little as far as the exposition of [17] is concerned; the only significant exception to this is the gauge group, which we will discuss in the next section.

(b) The restriction to p less than 4 is, at the moment, coming from Lemma 2.2. A much deeper reason for restricting to p less than 4 will appear in our gluing analysis of Section 3 (in particular, (3.12)), where this condition ensures we have a right inverse to our linearized operator that is uniformly bounded; see also [5, p. 293].

Fix a smooth cutoff function β'' on X supported on $[T - 1/2, \infty) \times N$ and identically 1 on $[T, \infty) \times N$. Fix also a smooth reference connection A_{ref} on E ; we assume this belongs to the space $\mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$. Using these objects, we can form the map

$$\begin{aligned} i: \mathcal{H}_{out} &\longrightarrow \mathcal{A}^{1,p}(\mathcal{T}_\Gamma) \\ h &\longmapsto A_{ref} + \beta''(\alpha(h) - A_{ref}) \end{aligned}$$

where $\alpha(h) = h_T(t) + \Theta(h_T(t))dt$ is as above. This map i is \mathcal{C}^m -smooth. As in [17, Lemma 10.1.1], it is convenient to introduce the map

$$(2.5) \quad \begin{aligned} \iota: \mathcal{H}_{out} \times L_{1,\delta}^p(\Omega^1(X)) &\longrightarrow \mathcal{A}^{1,p}(\mathcal{T}_\Gamma) \\ (h, V) &\longmapsto \iota(h, V) := i(h) + V \end{aligned}$$

This map ι is a \mathcal{C}^m -diffeomorphism with inverse given by $A \mapsto (p_T(A), A - i(p_T(A)))$. It follows immediately from the definitions that

$$p_T(i(h)) = p_T(\iota(h, V)) = h$$

for all $h \in \mathcal{H}_{out}$ and $V \in L_{1,\delta}^p(\Omega^1)$. We view ι as providing something of a coordinate system on the space of connections.

The tangent space to $\mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ at A is the space of all 1-forms $W \in L_{1,loc}^p(\Omega^1(X))$ so that there is some $\eta \in T_{p(A)}\mathcal{H}$ with

$$W - (Di)_{p_T(A)}\eta \in L_{1,\delta}^p(X)$$

where $(Di)_h$ is the linearization at $h \in \mathcal{H}_{out}$ of the map $i: \mathcal{H}_{out} \rightarrow \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$. Linearizing ι at (h, V) , we obtain a Banach space isomorphism

$$(2.6) \quad \begin{aligned} (D\iota)_{(h,V)}: T_h \mathcal{H}_{out} \times L_{1,\delta}^p(\Omega^1(X)) &\longrightarrow T_A \mathcal{A}^{1,p}(\mathcal{T}_\Gamma) \\ (\eta, W) &\longmapsto (Di)_h \eta + W \end{aligned}$$

The 1-form $(Di)_h \eta$ vanishes on X_0 , so the operator norm of $(D\iota)_{(h,V)}$ is independent of the metric on X_0 .

2B.2. The gauge group. When $2 < p < 4$, we will write $\mathcal{G}_\delta^{2,p}(\Gamma)$ for the set of bundle automorphisms u of E with the property that $u^*A \in \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ for all $A \in \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$. The condition that p be less than 4 is a carry-over from Remark 2.4 (b). The condition that p be larger than 2 ensures that $\mathcal{G}_\delta^{2,p}(\Gamma)$ is a well-defined Banach Lie group that acts C^m -smoothly on $\mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$; these claims follow from the Sobolev multiplication maps $W_{loc}^{2,p} \times W_{loc}^{2,p} \rightarrow W_{loc}^{2,p}$ and $W_{loc}^{2,p} \times W_{loc}^{1,p} \rightarrow W_{loc}^{1,p}$ being well-defined in dimension 4 for $p > 2$. See, for example, [23, Lemma A.6]. We will only consider $\mathcal{G}_\delta^{2,p}(\Gamma)$ for p satisfying $2 < p < 4$.

The proof of [17, Lemma 7.2.7] carries over to this setting to imply that the group $\mathcal{G}_\delta^{2,p}(\Gamma)$ is equal to the space of $L_{2,loc}^p$ -gauge transformations with the property that there is some $\tau_u \in \text{Stab}(\Gamma)$, viewed as a t -invariant gauge transformation on $\text{End } X$, so that $u|_{\text{End } X} \circ \tau_u^{-1}$ is in $L_{2,\delta}^p(\text{End } X)$. The gauge transformation τ_u is uniquely determined by u , and we denote by $\mathcal{G}_\delta^{2,p} \subseteq \mathcal{G}_\delta^{2,p}(\Gamma)$ the (normal) subgroup of all gauge transformations u with $\tau_u = \text{I}$ equal to the identity. Thus, we have a short exact sequence of groups:

$$(2.7) \quad \{e\} \longrightarrow \mathcal{G}_\delta^{2,p} \longrightarrow \mathcal{G}_\delta^{2,p}(\Gamma) \longrightarrow \text{Stab}(\Gamma) \longrightarrow \{e\}$$

We will write $\text{Stab}(A)$ for the stabilizer of A under the action of $\mathcal{G}_\delta^{2,p}(\Gamma)$. The center $Z(G)$ of G embeds into $\mathcal{G}_\delta^{2,p}(\Gamma)$ as the set of constant maps $X \rightarrow Z(G)$, and we will identify $Z(G)$ with its image in the gauge group. Note that $Z(G)$ is also the center of $\mathcal{G}_\delta^{2,p}(\Gamma)$ and $Z(G) \subseteq \text{Stab}(A)$. We will say that A is *irreducible* if $Z(G)$ and $\text{Stab}(A)$ have the same dimension (equivalently, if they have isomorphic Lie algebras). Note that the term “irreducible” is only defined when $2 < p < 4$.

Lemma 2.8. *Fix $2 < p < 4$, and assume $A \in \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ is irreducible. Then there is a neighborhood $U \subseteq \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ of A so that A' is irreducible for all $A' \in U$.*

Proof. We begin with a few preliminaries. Set $\mathcal{A} := \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ and $\mathcal{G} := \mathcal{G}_\delta^{2,p}(\Gamma)$. Linearizing the gauge group action at $A \in \mathcal{A}$, we obtain a map

$$d_A: \text{Lie}(\mathcal{G}) \longrightarrow T_A \mathcal{A} \quad \phi \longmapsto d_A \phi.$$

Then a connection $A \in \mathcal{A}$ is irreducible if and only if the kernel of d_A equals the Lie algebra $\mathfrak{z} := \text{Lie}(Z(G))$ of the center of \mathcal{G} . It follows from the definition of the topologies on \mathcal{A} and \mathcal{G} , as well as from standard elliptic estimates for δ -decaying spaces, that the operator d_A is bounded. Moreover, this operator has a range that is closed and admits a complement in $T_A \mathcal{A}$.

Let $H_\Gamma^0 := \ker(d_\Gamma) \subseteq \Omega^0(N)$ be the Lie algebra of $\text{Stab}(\Gamma)$. The center \mathfrak{z} is naturally a subalgebra of H_Γ^0 , so we can write

$$H_\Gamma^0 = \mathfrak{z} \oplus \mathfrak{z}^\perp$$

where \mathfrak{z}^\perp is the $L^2(N)$ -orthogonal complement of \mathfrak{z} . Then for $\tau \in H_\Gamma^0$, we will write $\tau^\perp \in \mathfrak{z}^\perp$ for its projection.

Just as the map ι provides “coordinates” for \mathcal{A} , there is an analogous Banach space isomorphism

$$\iota_{\Omega^0} : H_\Gamma^0 \times L_{2,\delta}^p(\Omega^0) \longrightarrow \text{Lie}(\mathcal{G}) \quad (\tau, \xi) \longmapsto (\tau - \tau^\perp) + \beta'' \tau^\perp + \xi$$

where we are viewing $\tau - \tau^\perp \in \mathfrak{z}$ as a 0-form on X . This map ι_{Ω^0} takes $\mathfrak{z} \times \{0\}$ isomorphically to $\mathfrak{z} \subseteq \text{Lie}(\mathcal{G})$. Let $Y \subseteq \text{Lie}(\mathcal{G})$ be the image under ι_{Ω^0} of the complement $\mathfrak{z}^\perp \times L_{2,\delta}^p(\Omega^0)$ to $\mathfrak{z} \times \{0\}$. Then we have a direct sum decomposition

$$\text{Lie}(\mathcal{G}) = \mathfrak{z} \oplus Y.$$

The key point is that A is irreducible if and only if the restriction

$$d_A^\perp := d_A|_Y : Y \longrightarrow T_A \mathcal{A}$$

is injective. We will want to view this operator as a function of A , and for this it would be convenient if d_A^\perp were to have a codomain that is independent of A . Though this is not the case presently, we can arrange for A -independence of the codomain as follows: Let $(D\iota)_A : T_h \mathcal{H} \times L_{1,\delta}^p(\Omega^1) \rightarrow T_A \mathcal{A}$ be the linearization of the coordinate map ι . The L^2 -inner product for 1-forms on N provides a Riemannian metric on the finite-dimensional subspace $\mathcal{H} \subseteq L^2(\mathcal{A}(N))$ of connections on N . Use this Riemannian metric to define the parallel transport map $\text{PT}_h : T_h \mathcal{H} \rightarrow T_\Gamma \mathcal{H}$. Letting I denote the identity on $L_{1,\delta}^p(\Omega^1)$, we will be interested in the operator

$$D_A := (\text{PT}_h \times I) \circ (D\iota)_A^{-1} \circ d_A^\perp : Y \longrightarrow T_\Gamma \mathcal{H} \times L_{1,\delta}^p(\Omega^1).$$

This is a bounded linear map, and expansions of the form $d_A = d_{A'} + [(A - A'), \cdot]$ show it depends continuously on $A \in \mathcal{A}$ in the operator norm topology on the space $\mathcal{B}(Y, T_\Gamma \mathcal{H} \times L_{1,\delta}^p(\Omega^1))$ of bounded linear maps from Y to $T_\Gamma \mathcal{H} \times L_{1,\delta}^p(\Omega^1)$. Since Y has finite codimension, and d_A has closed range, the operator D_A has closed range as well.

It follows from the construction that D_A is injective if and only if A is irreducible. Assume that A is irreducible. Then the fact that $\text{im}(d_A)$ has a complement in $T_A \mathcal{A}$ implies that D_A admits a bounded left inverse, which we denote by L_A . In summary, the map

$$A \longrightarrow \mathcal{B}(Y, Y) \quad A' \longmapsto L_A D_{A'}$$

is a continuous map into the space of bounded linear operators on the Banach space Y . It is clearly invertible at $A' = A$. Since the set of invertible bounded linear maps on a Banach space is open, there is some neighborhood $U \subseteq \mathcal{A}$ of A so that $L_A D_{A'}$ is invertible for all $A' \in U$. Thus if $A' \in U$, then A' is irreducible. \square

Remark 2.9. Completing $L_{2,\delta}^p(Y)$ to $L_{1,\delta}^p(Y)$, the map D_A extends to a bounded linear operator of the form

$$D_A : L_{1,\delta}^p(Y) \longrightarrow T_\Gamma \mathcal{H} \times L_\delta^p(\Omega^1(X)).$$

Let $p^* = 4p/(4-p)$ be the Sobolev conjugate of $p \in (2, 4)$. Then one can show that the map $A \mapsto D_A \in \mathcal{B}(L_{1,\delta}^p(Y), T_\Gamma \mathcal{H} \times L_\delta^p(\Omega^1(X)))$ is continuous in $A = i(h) + V$ relative to the topology $(h, V) \in \mathcal{C}^0(N) \times L_\delta^{p^*}(X)$. The proof we gave for Lemma 2.8 carries over to show that $A' = \iota(h', V')$ is irreducible whenever $A = \iota(h, V)$ is irreducible, and $\|h - h'\|_{\mathcal{C}^0} + \|V - V'\|_{L_\delta^{p^*}}$ is sufficiently small.

2B.3. The mASD equation. Fix a cut off function β' on X that is identically 1 on the cylinder $[T + 1/2, \infty) \times N$ and supported on the slightly larger cylinder $[T, \infty) \times N$. Consider the map

$$s : \mathcal{A}^{1,p}(\mathcal{T}_\Gamma) \longrightarrow L_\delta^p(\Omega^+(X)) \quad A \longmapsto F_A^+ - \beta' F_{i(p_T(A))}^+.$$

We will call s the *modified ASD (mASD) operator*. The equation $s(A) = 0$ is the *modified ASD (mASD) equation*, and any A satisfying $s(A) = 0$ will be called *modified ASD (mASD)*. The map s is \mathcal{C}^m in the specified topologies; see [17, Lemma 7.1.1] and use the fact that the composition of \mathcal{C}^m functions is again \mathcal{C}^m .

The following will help us understand the linearization of s .

Lemma 2.10. *If $A = \iota(h, V)$ for $(h, V) \in \mathcal{H}_{out} \times L_{1,\delta}^p(\Omega^1(X))$, then*

$$(2.11) \quad s(A) = s(\iota(h, V)) = (1 - \beta') F_{i(h)}^+ + d_{i(h)}^+ V + \frac{1}{2} [V \wedge V]^+.$$

Proof. This follows from the identity $F_{i(h)+V}^+ = F_{i(h)}^+ + d_{i(h)}^+ V + \frac{1}{2} [V \wedge V]^+$ and the fact that $p_T(\iota(h, V)) = h$. \square

Remark 2.12. Unfortunately, when G is not abelian, the mASD operator s is not generally well-behaved under any suitable gauge group; e.g., it is not equivariant relative to the action of the gauge group of Section 2B. The issue is that the term $F_{i(p_T(A))}$ is gauge equivariant relative to the *trivial* G -action on \mathfrak{g} , while the term F_A is gauge equivariant relative to the *adjoint* G -action on \mathfrak{g} . Consequently, any non-trivial linear combination of these (e.g., as in the above formula for s) is not equivariant relative to either G -action. This issue is apparent even in the smooth compactly supported setting, and hence persists regardless of which Sobolev completion we choose. See Section 2B.5 for some comments about the effect of this issue on the results of [17].

2B.4. A Coulomb slice. To obtain a Fredholm operator, we will restrict the operator s to a Coulomb (gauge) slice

$$\mathcal{SL}(A') := \left\{ \iota(h, V) \mid h \in \mathcal{H} \text{ and } V \in \ker(d_{A'}^{*,\delta}) \subseteq L_{1,\delta}^p(\Omega^1) \right\}$$

for some fixed connection A' . Here $d_{A'}^{*,\delta} = e^{-t\delta} d_A^* e^{t\delta}$ is the adjoint relative to the L_δ^2 -inner product. For more details on this slice, see [17, Prop. 10.3.1]. Since $\delta/2$ is not in

the spectrum of $- * d_\Gamma$ on 1-forms, it follows from [17, Lem. 8.3.1] that the restriction $s|_{\mathcal{SL}(A')}$ of s to this slice is a Fredholm map.

We set

$$\widehat{\mathcal{M}} = \widehat{\mathcal{M}}(\mathcal{T}_\Gamma, A') := s^{-1}(0) \cap \mathcal{SL}(A')$$

which we refer to as the *space of mASD-connections*. For us, this will play the role that the ASD moduli space usually plays in the closed setting (though, as discussed in the introduction, this is less than satisfying for global considerations due to its dependence on A'). Elliptic regularity implies that any element of $\widehat{\mathcal{M}}$ has regularity \mathcal{C}^m .

Consider the restriction of the mASD operator s to this slice $\mathcal{SL}(A')$. Then the linearization at $A \in \mathcal{SL}(A')$ of this restriction is a bounded linear map

$$(Ds|_{\mathcal{SL}(A')})_A : T_A \mathcal{SL}(A') \longrightarrow L_\delta^p(\Omega^+).$$

We will say that an mASD connection A is *A' -regular* if this operator is surjective. When $A' = A$, we drop the A' and say that A is *regular* if it is A -regular. We will also be interested in connections that are not regular, and for these we will need to consider the cokernel

$$(2.13) \quad H_{A,\delta}^+ := \text{coker} (Ds|_{\mathcal{SL}(A)})_A.$$

Clearly $H_{A,\delta}^+ = 0$ if and only if A is regular.

We will write

$$\widehat{\mathcal{M}}_{reg} = \widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A') \subseteq \widehat{\mathcal{M}}(\mathcal{T}_\Gamma, A')$$

for the subset of A' -regular mASD connections. It follows from the implicit function theorem that $\widehat{\mathcal{M}}_{reg}$ is a \mathcal{C}^m -smooth manifold.

2B.5. Navigating Morgan–Mrowka–Ruberman [17]. Since the operator s is not gauge equivariant, we do not currently know how to define a suitable gauge quotient of $s^{-1}(0) \subset \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ in order to obtain the mASD moduli space $\mathcal{M}_{\ell,w}(\mathcal{T}_\Gamma, T_0, g)$ envisaged in [17, p. 125]. Though this appears to be a significant error in [17], the major results of [17] remain intact. For the reader's benefit, we will now review [17], highlighting those areas that need adjustment. In this section, we refer freely to the notation established in [17], and the phrase “thickened moduli space” will refer to any of the following spaces

$$\mathcal{M}_{\ell,w}(\mathcal{T}_\Gamma, T_0, g), \quad \mathcal{M}_{\ell,w}^0(\mathcal{T}_\Gamma, T_0, g), \quad \mathcal{M}_{\ell,w}(h, \mathcal{T}_\Gamma, T_0, g)$$

(note that sometimes the w or the g are dropped from the notation in [17]).

The thickened moduli space does not appear in any significant way in Chapters 1–6, 10–13, 15–16, and the results therein remain intact as stated. In Chapters 7–9, it is often the case that claims about the thickened moduli space remain intact if one interprets the term as meaning an object defined by a gauge slice as opposed to a quotient by the gauge group (e.g., replace the based version $\mathcal{M}_{\ell,w}^0(\mathcal{T}_\Gamma, T_0, g)$ by what we called $\widehat{\mathcal{M}}(\mathcal{T}_\Gamma, A')$ above). Such is the case with the patching results of § 7.4, the index calculations of Ch. 8 (more details below), and most of the generic metrics results of Ch. 9. However, it is not immediately clear to us how best to interpret the claims in Lemma

7.5.3 and §9.4 since they refer to the μ -map on the thickened moduli space, which is an inherently global object.

At first sight, Chapter 14 appears to have issues but, upon closer inspection, the results of this chapter remain intact as well, with the possible exception of the claims of [17, Remark 14.0.5] (which don't appear to be used elsewhere in [17]). Here are some more details regarding Chapter 14: The main goal of this chapter is to prove Theorem 14.0.1, which is a structural result that gives dimension formulas for various ASD moduli spaces. This theorem has two cases. In the situation of Case I, as the authors point out explicitly (p. 202), every mASD connection is actually ASD and so the “thickened moduli space” is indeed well-defined because it is really (an open set of) the ASD moduli space (this coincidence is one that we too exploit, and is discussed more in Section 6D). Thus, there is no problem with this case. The remaining Case II in Theorem 14.0.1 does not appeal to the thickened moduli space at all, and instead works with the space $\mathcal{B}_{\ell,\delta}^0(X)$ (despite being a gauge group quotient, the space $\mathcal{B}_{\ell,\delta}^0(X)$ is indeed well-defined as it does not make reference to the mASD operator s). Thus, the statement and proof of Theorem 14.0.1 need no adjustment and appear to be correct as written. However, we are not sure how to interpret the argument of Remark 14.0.5, which uses the thickened moduli space to deduce a Whitney stratification on the ASD-moduli space. In the setting of this remark, there are mASD connections that are not ASD, so the above-referenced coincidence does not apply.

We end this section with a more detailed discussion of Chapter 8 in [17], with the aim of salvaging the index calculation of Proposition 8.5.1. As in [17, Ch. 8], in addition to our usual hypotheses on δ , we also assume that $\delta < 2\mu_\Gamma^-$. We also restrict to the case where $G = \mathrm{SO}(3)$ or $G = \mathrm{SU}(2)$; the index for more general compact G can be computed using the strategy outlined in [5, Section 7.1] (e.g., when G is simple and simply-connected, use the data from [1, Table 8.1] to pin down the constants specific to G).

The first issue appears on p. 139, since the image of the map $D_\omega m$ is not generally contained in the kernel of $D_\omega s$ (this is the linear version of the fact that s is not gauge-equivariant). As a consequence, $E_\delta(\omega)$ is not a complex in the usual sense. Nevertheless, much of what is desired of $E_\delta(\omega)$ can be salvaged by “wrapping it up” and considering the operator

$$D := (d_\omega^{*\delta}, D_\omega s) : T_\omega \mathcal{A}_{\ell,\omega}(\mathcal{T}_\Gamma, T_0) \longrightarrow T_e \mathcal{G}_\delta(\Gamma) \times \Omega_{+,1,\delta}^2(X)$$

where $d_\omega^{*\delta} = e^{-t\delta} d_\omega^* e^{t\delta}$ is the L_δ^2 -adjoint of d_ω and $D_\omega s$ is the notation used in [17] for the linearization of s at a connection ω (what they call ω is what we call A). Then Proposition 8.5.1 can be interpreted as saying that this operator D is Fredholm with index given by the formula:

$$\mathrm{Ind}(D) = 8\ell - \frac{3}{2}(\chi(X) + \sigma(X)) + \frac{h_\Gamma^1 - h_\Gamma^0}{2} + \frac{\rho(\Gamma)}{2}.$$

There is a “based” version of this that replaces $\mathcal{G}_\delta(\Gamma)$ by the normal subgroup \mathcal{G}_δ . Wrapping in this case produces a map

$$D^0 := (d_\omega^{*\delta}, D_\omega s) : T_\omega \mathcal{A}_{\ell, \omega}(\mathcal{T}_\Gamma, T_0) \longrightarrow T_e \mathcal{G}_\delta \times \Omega_{+, 1, \delta}^2(X).$$

This operator is also Fredholm and, in light of the sequence (2.7), its index is given by

$$\text{Ind}(D^0) = \text{Ind}(D) + h_\Gamma^0 = 8\ell - \frac{3}{2}(\chi(X) + \sigma(X)) + \frac{h_\Gamma^1 + h_\Gamma^0}{2} + \frac{\rho(\Gamma)}{2}$$

since $h_\Gamma^0 = \dim(\text{Stab}(\Gamma))$. Restricting D^0 to a slice $\mathcal{S}\mathcal{L}(A')$ and projecting to the $\Omega_{+, 1, \delta}^2(X)$ -component, we obtain the operator that we called $(Ds|_{\mathcal{S}\mathcal{L}(A')})_A$ above. When $A = A'$ and this connection is flat down the end, it follows readily that the index of $(Ds|_{\mathcal{S}\mathcal{L}(A)})_A$ equals that of D^0 . Since the index remains unchanged under addition of compact operators, it follows that

$$(2.14) \quad \text{Ind}((Ds|_{\mathcal{S}\mathcal{L}(A')})_A) = 8\ell - \frac{3}{2}(\chi(X) + \sigma(X)) + \frac{h_\Gamma^1 + h_\Gamma^0}{2} + \frac{\rho(\Gamma)}{2}$$

for all $A \in \mathcal{S}\mathcal{L}(A')$. The right-hand side of (2.14) is the expected dimension of the mASD space $\widehat{\mathcal{M}}(\mathcal{T}_\Gamma, A')$, and it is the actual dimension of the \mathcal{C}^m -manifold $\widehat{\mathcal{M}}_{\text{reg}}(\mathcal{T}_\Gamma, A')$ of A' -regular mASD connections.

2C. Special cases.

2C.1. Flat connections. In this section, we study the linearized operator $(Ds)_A$ and its cokernel in the special case when A is flat and asymptotic to Γ . To simplify the discussion, we assume A is in temporal gauge on the end [4, p. 15] (though we continue to work in the general setting where the metric is asymptotically cylindrical). It follows that, for each $t \geq T$, the restriction $A|_{\{t\} \times N} = \Gamma$ is constantly equal to Γ on the end. Then $A \in \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ and $p(A) = \Gamma$. The associated flow $\alpha(\Gamma) = A$ recovers the flat connection A on End X . This implies A is mASD.

The operator s is defined in terms of the map ι , and we recall that the definition of ι required the choice of a reference connection A_{ref} on X . Likewise, the space of mASD connections is defined by restricting s to a slice $\mathcal{S}\mathcal{L}(A')$ for some choice of connection A' . It is convenient to take $A_{\text{ref}} := A$ and $A' := A$; the reader can check that any other choice of A_{ref} does not affect the outcome of the discussion that follows, though different choices of A' may. In particular, our choice of A_{ref} gives

$$A = i(\Gamma) = \iota(\Gamma, 0)$$

Let $b^+(X, A)$ be the dimension of a maximal positive definite subspace for the pairing map $q_A : \hat{H}^2(\bar{X}, \text{ad}(A)) \otimes \hat{H}^2(\bar{X}, \text{ad}(A)) \rightarrow \mathbb{R}$, as in [17, Section 8.7], where \bar{X} is the natural compactification of X obtained by adding a copy of N at infinity. For example, when $A = A_{\text{triv}}$ is the trivial connection on the trivial G -bundle, then $b^+(X, A_{\text{triv}}) = \dim(G)b^+(X)$ is a multiple of the usual self-dual Betti number of X . We will need the following result.

Proposition 2.15. *Assume $0 < \delta/2 < \mu_\Gamma^-$, where μ_Γ^- is as in Section 2A. Then the cokernel $H_{A,\delta}^+ = (Ds|_{\mathcal{SL}(A)})_A$ has dimension $b^+(X, A)$.*

This is proved in [17, Prop. 8.7.1(4)], however the discussion there does not deal with the linearized operator $(Ds)_A$ directly. In preparation for our gluing arguments below, we will summarize the argument given in [17, Prop. 8.7.1(4)], but from the present perspective. Our proof is sketched below, after we give some preliminary computations that will be useful in their own right.

The restriction $\iota|: \mathcal{H} \times \ker(d_A^{*,\delta}) \rightarrow \mathcal{SL}(A)$ is a diffeomorphism, essentially by definition. To understand the cokernel $H_{A,\delta}^+$ it suffices to understand the cokernel of the linearization of $s \circ \iota$. Towards this end, differentiating (2.11) at $A = \iota(\Gamma, 0)$ in the direction of $(\eta, V) \in T_\Gamma \mathcal{H}_{out} \times \ker(d_A^{*,\delta})$ gives

$$(2.16) \quad (Ds)_A \circ (D\iota)_{(\Gamma,0)}(\eta, V) = (1 - \beta')d_A^+(Di)_\Gamma \eta + d_A^+ V.$$

Next, it follows from Lemma 2.3 that $(Di)_\Gamma \eta = \beta'' \eta$. Since we also have $d_\Gamma \eta = 0$ for $\eta \in T_\Gamma \mathcal{H} = H_\Gamma^1$, it follows that

$$(2.17) \quad d_A^+(Di)_\Gamma \eta = (\partial_t \beta'')(dt \wedge \eta)^+.$$

This is zero everywhere except on $(T - 1/2, T) \times N$, where it vanishes if and only if $\eta = 0$. Combining this with (2.16), we therefore have the formula

$$(2.18) \quad (Ds)_A \circ (D\iota)_{(\Gamma,0)}(\eta, V) = (1 - \beta')(\partial_t \beta'')(dt \wedge \eta)^+ + d_A^+ V.$$

This shows that, relative to the coordinates afforded by ι , the leading order term of $(Ds|_{\mathcal{SL}(A)})_A$ is $d_A^+|_{\ker(d_A^{*,\delta})}$. The remaining term is compactly-supported and of order zero.

Proof of Proposition 2.15 (sketch). A maximal positive definite subspace for q_A can be realized as the space $H^+(X, \text{ad}(A))$ of self-dual 2-forms $W \in L^2(\Omega^+(X, \text{ad}(A)))$ satisfying $d_A W = 0$ and so that the restriction to any slice $\{t\} \times N$ has trivial Γ -harmonic part; note that this definition does not depend on δ , but see also Lemma 2.20. Setting

$$D := (Ds|_{\mathcal{SL}(A)})_A$$

we can similarly represent the cokernel $H_{A,\delta}^+$ of D as the L_δ^2 -orthogonal complement $(\text{im } D)^\perp$ to the image. Then Proposition 2.15 follows by showing that the map

$$j: (\text{im } D)^\perp \longrightarrow H^+(X, \text{ad}(A)) \quad W \longmapsto e^{\delta t} W$$

is well-defined and bijective. That the map j is well-defined follows from the formula in (2.18). Indeed, if $W \in (\text{im } D)^\perp$, then the identity (2.16) implies

$$0 = (W, D \circ (D\iota)_{(\Gamma,0)}(0, V))_\delta = (W, d_A^+ V)_\delta = (d_A^{*,\delta} W, V)_\delta$$

where $(\cdot, \cdot)_\delta$ is the L_δ^2 -inner product. This holds for all $V \in L_{1,\delta}^2(\Omega^1)$, so it follows that $d_A j(W) = -e^{t\delta} * d_A^{*,\delta} W = 0$. Similarly, we have

$$0 = (W, D \circ (D\iota)_{(\Gamma,0)}(\eta, 0))_\delta = (W, (1 - \beta')(\partial_t \beta'')(dt \wedge \eta)^+)_\delta.$$

Since $(1 - \beta')\partial_t\beta''$ is non-zero on the cylinder $(T - 1/2, T) \times N$, and since $\eta \in T_\Gamma\mathcal{H} = H_\Gamma^1$ is allowed to roam freely over the Γ -harmonic space, it follows that the harmonic part of $W|_{\{t\} \times N}$ must vanish for any $t \in (T - 1/2, T)$.

It is clear that the map j is injective; this already gives $\dim(H_{A,\delta}^+) \leq b^+(X, A)$. Surjectivity of j is equivalent to the reverse inequality holding, and this follows from a dimension count: By [17, Prop. 8.7.1(4)] (which uses the assumption $0 < \delta/2 < \mu_\Gamma^-$), the integer $b^+(X, A)$ is equal to the dimension of the cokernel of $(Ds)_A : T_A\mathcal{A}^{1,p}(\mathcal{T}_\Gamma) \rightarrow L_\delta^p(\Omega^+)$. The result follows because cokernels are non-decreasing under domain restriction:

$$\begin{aligned} b^+(X, A) &= \dim \operatorname{coker}((Ds)_A : T_A\mathcal{A}^{1,p}(\mathcal{T}_\Gamma) \rightarrow L_\delta^p(\Omega^+)) \\ &\leq \dim \operatorname{coker}((Ds|_{\mathcal{S}\mathcal{L}(A)})_A : T_A\mathcal{S}\mathcal{L}(A) \rightarrow L_\delta^p(\Omega^+)) \\ &= \dim H_{A,\delta}^+. \end{aligned}$$

□

Remark 2.19. Proposition 2.15 is the observation that makes the mASD-operator—and *not the ASD-operator*—a viable candidate for our existence results, including Theorem A which is purely in the ASD setting. Indeed, consider the case where $b^+(X) = 0$ and A is the trivial flat connection. Then $b^+(X, A) = b^+(X) = 0$ and so Proposition 2.15 implies A is regular as an mASD connection. However, when Γ is degenerate, the trivial flat connection A is *not* regular as an ASD connection: The proof just given shows that, for the operator (2.18), the image of $\eta \mapsto (1 - \beta')(\partial_t\beta'')(dt \wedge \eta)^+$ is *not* contained in the image of d_A^+ . Thus, the additional degree of freedom afforded by $\eta \in T_\Gamma\mathcal{H}$ in (2.18) is necessary (and sufficient) to obtain surjectivity of $(Ds|_{\mathcal{S}\mathcal{L}(A)})_A$ when A is the trivial flat connection, Γ is degenerate, and $b^+(X) = 0$.

We end with the following exponential decay estimate that we will use in Section 6C.

Lemma 2.20. *For each $W \in H^+(X, \operatorname{ad}(A))$, there is some C so that the restriction $W|_{\{t\} \times N}$ satisfies*

$$\|W|_{\{t\} \times N}\|_{C^0(N)} \leq Ce^{-\mu_\Gamma^+ t}$$

for all $t \geq 0$. In particular, $H^+(X, \operatorname{ad}(A)) \subseteq L_\delta^2(X)$ for any $\delta < 2\mu_\Gamma^+$.

Proof. It suffices to establish the estimate of the lemma under the assumption that the metric g is cylindrical. Since A is in temporal gauge on the end, its covariant derivative decomposes as $d_A = dt \wedge \partial_t + d_\Gamma$. Standard elliptic estimates for the operator $\Delta_\Gamma = d_\Gamma^*d_\Gamma + d_\Gamma d_\Gamma^*$ on N provide a uniform constant C so that

$$\|v\|_{C^0(N)} \leq C(\|v\|_{L^2(N)} + \|\Delta_\Gamma v\|_{L^2(N)})$$

for all $v \in \Omega^1(N)$.

Fix $W \in H^+(X, \operatorname{ad}(A))$. On End X , we can write $W = dt \wedge v + *_N v$ for some path $v = v(t) \in \Omega^1(N)$ of 1-forms. The condition $d_A W = 0$ implies $d_\Gamma *_N v = 0$ and

$*_N d_\Gamma v = \partial_t v$. In particular, the above elliptic estimate implies

$$\|v(t)\|_{C^0(N)} \leq C(\|v(t)\|_{L^2(N)} + \|\partial_t^2 v(t)\|_{L^2(N)}).$$

It suffices to show that $f(t) := \|v(t)\|_{L^2(N)}^2 + \|\partial_t^2 v(t)\|_{L^2(N)}^2$ decays exponentially in t at a rate of $2\mu_\Gamma^+$. To see this, differentiate twice to get

$$\begin{aligned} f''(t) &= 2\|\partial_t v(t)\|_{L^2(N)}^2 + 2\|\partial_t^3 v(t)\|_{L^2(N)}^2 + 2(\partial_t^2 v(t), v(t)) + 2(\partial_t^4 v(t), \partial_t^2 v(t)) \\ &= 4\|d_\Gamma v(t)\|_{L^2(N)}^2 + 4\|d_\Gamma \partial_t^2 v(t)\|_{L^2(N)}^2 \end{aligned}$$

where we used $\partial_t v = *_N d_\Gamma v$ and integration by parts. By definition of $H^+(X, \text{ad}(A))$, $v(t)$ is orthogonal to the Γ -harmonic space and $d_\Gamma *_N v(t) = 0$. It follows that $v(t)$ lies in the image of $*_N d_\Gamma$. Moreover, the 2-form W , and hence v , is in L^2 , by definition. This combines with the equation $*_N d_\Gamma v = \partial_t v$ to imply that $v(t)$ lies in the span of the negative eigenspaces of $*_N d_\Gamma$: express $v(t)$ as a t -dependent linear combination of an orthonormal basis of eigenvectors, use a separation of variables argument to show that each coefficient decays or grows down the end according to whether the eigenvalue is negative or positive, and then use the fact that v is in L^2 and so must decay down the end.

Likewise, $\partial_t^2 v(t)$ always lies in the negative eigenspace of $*_N d_\Gamma$. Recall that μ_Γ^+ is the smallest positive eigenvalue of $-*_N d_\Gamma$. Thus, it is the absolute value of the largest negative eigenvalue of $*_N d_\Gamma$, and so

$$f''(t) = 4\|d_\Gamma v(t)\|_{L^2(N)}^2 + 4\|d_\Gamma \partial_t^2 v(t)\|_{L^2(N)}^2 \geq 4(\mu_\Gamma^+)^2 f(t).$$

Since $f(t)$ is non-negative and converges to 0 as t approaches ∞ , it follows from this estimate that $f(t) \leq Ce^{-2\mu_\Gamma^+ t}$ (e.g., see [6, p. 623]). \square

2C.2. Non-degenerate Γ . Here we assume that Γ is non-degenerate in the sense that the harmonic space $H_\Gamma^1 = \{0\}$ is trivial. Then any center manifold necessarily consists of a single point and so there is a unique choice of cutoff function β . Then the mASD operator is the ASD operator. Moreover, non-degeneracy implies that any finite-energy ASD connection asymptotic to Γ decays exponentially on the end at a rate of $e^{-\delta t}$ for any $\delta/2 < \mu_\Gamma^-$, where μ_Γ^- is as in Section 2A (the proof of this assertion is similar to the proof of Lemma 2.20, but μ_Γ^- appears in the present setting since the curvature is *anti*-self dual). We also assumed that δ is greater than μ_Γ^- , but that was only to construct the map p_Γ , whose existence is trivial in the non-degenerate setting since the center manifold consists of a single point. Thus, in the non-degenerate case this lower bound restriction on δ can be dropped, though we still need to retain the assumption that $\delta/2$ is not an eigenvalue of $-*_N d_\Gamma$. In particular, it follows that whenever $-\mu_\Gamma^+ < \delta/2 < \mu_\Gamma^-$, the resulting mASD/ASD space is independent of the choice of this δ , and the dimension of the space is given by (2.14) (apart from this paragraph, we always assume $\delta > 0$).

In summary, when Γ is non-degenerate, there is an essentially canonical choice of thickening data $\overline{\mathcal{T}}_{\Gamma, \text{can}}$. Moreover, if A' is any connection defining a slice, then the

associated space of mASD connections

$$\widehat{\mathcal{M}}(\mathcal{T}_{\Gamma, \text{can}}, A') = \left\{ A \in \mathcal{A}(X) \mid F_A^+ = 0, \quad d_{A'}^*(A - A') = 0, \quad \lim_{t \rightarrow \infty} A|_{\{t\} \times N} = \Gamma \right\}$$

is the set of ASD connections in the A' -Coulomb slice that are asymptotic to Γ .

2C.3. Closed X . Here we assume that X is closed. View X as a cylindrical end 4-manifold with an empty end. Then one can check that it makes sense to choose the empty set $\mathcal{T}_{\emptyset} = \emptyset$ of thickening data, and that, e.g., the mASD space $\widehat{\mathcal{M}}(\mathcal{T}_{\emptyset}, A')$ is exactly the set of ASD connections on E in the A' -Coulomb slice.

3. GLUING TWO mASD CONNECTIONS

Here we state and prove our first gluing result, which discusses gluing together mASD connections over the compact parts of two cylindrical end 4-manifolds. When the connections are not regular, the resulting connection may not be mASD, and its failure to be mASD is captured by a suitable obstruction map. Our set-up is very similar to that of ASD gluing outlined in Donaldson–Kronheimer [5, Section 7.2], to which we refer the reader for more details at various points. When introducing new terms for the analysis, we have tried to keep our notation as consistent with that of [5] as possible. Our emphasis below will be on the new features that arise in the mASD setting. In the present section, the only serious new features arise from the fact that the mASD operator s has a nonlinear term not present in the usual ASD setting; these features manifest themselves in the proofs of the claims appearing in the proof of Theorem 3.3.

3A. Set-up for gluing. Let X_1 and X_2 be oriented cylindrical end 4-manifolds equipped with asymptotically cylindrical metrics as in Section 2A.2. We will write X_{k0} for compact part of X_k and we set $N_k := \partial X_{k0}$. We will need parameters $\lambda > 0$ and $L > 1$ so that $b := 4L\lambda^{1/2} \ll 1$. The constant L will be fixed later, but we will ultimately be interested in allowing λ to be arbitrarily small. For each k , fix a point

$$x_k \in B_b(x_k) \subset \text{int}(X_{k0})$$

in the interior of the compact part. To simplify the discussion, we assume that the metric on X_k is flat over $B_b(x_k)$; see [3, Section IV(vi)] for how to extend the discussion to handle more general metrics.

Following the approach in [5, Section 7.2.1], we glue along the annuli

$$\Omega_k := B_{L\lambda^{1/2}}(x_k) \setminus B_{L^{-1}\lambda^{1/2}}(x_k)$$

using an “inversion” map $f_\lambda : \Omega_1 \rightarrow \Omega_2$ to produce a connected sum

$$X = X(L, \lambda) := (X_1 \setminus B_{L^{-1}\lambda^{1/2}}(x_1)) \cup_{f_{L, \lambda}} (X_2 \setminus B_{L^{-1}\lambda^{1/2}}(x_2)).$$

Then X is an oriented cylindrical end 4-manifold with asymptotic 3-manifold $N = N_1 \sqcup N_2$. We will write X_0 for the compact part of X ; this is formed by analogously gluing the compact parts X_{k0} of the X_k . The metrics on the X_k can be glued to form a metric on X , and we assume this is done as outlined at the end of p. 293 in [5]. We denote this metric by $g_{L, \lambda}$. Since we are interested in the limiting behavior of this for

small λ , we will include the metric in the notation for our various norms and spaces of connections, forms, etc. whenever it is relevant.

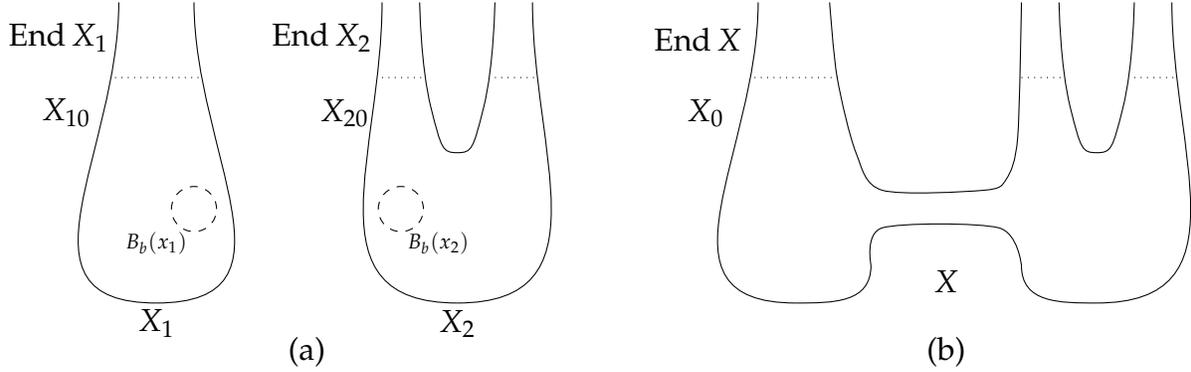


FIGURE 1. Illustrated above are the manifolds X_1, X_2 in (a), and their connected sum X in (b). The 3-manifolds N_1, N_2 , and N are unlabeled, but are illustrated as dotted lines in the figure above.

Fix principal G -bundles $E_k \rightarrow X_k$ and flat connections $\Gamma_k \in \mathcal{A}(N_k)$ for $k = 1, 2$. These induce a bundle over N as well as a flat connection Γ on N . Fix $\delta > 0$ as in Section 2A.3 associated to this flat connection Γ . It follows that, for $k = 1, 2$, the quantity $\delta/2$ is not in the spectrum of $-*d_{\Gamma_k}$ on 1-forms. Let $\mathcal{T}_{k, \Gamma_k}$ be thickening data for E_k with this δ .

Fix an isomorphism $\rho: (E_1)_{x_1} \rightarrow (E_2)_{x_2}$ of G -spaces, as well as flat connections $A_{b,k}$ for E_k over $B_b(x_k)$. Using these flat connections and radial parallel transport, we can extend ρ to a bundle isomorphism $E_1|_{\Omega_1} \cong E_2|_{\Omega_2}$ covering f_λ . It is with this bundle isomorphism that we glue the E_k over the Ω_k to obtain a bundle

$$E = E(\rho, L, \lambda) \longrightarrow X(L, \lambda).$$

Since the gluing takes place away from the cylindrical end, the thickening data $\mathcal{T}_{k, \Gamma_k}$ for the E_k induce thickening data \mathcal{T}_Γ for E .

Fix $1 \leq p < 4$ and suppose that, for each k , we have a smooth mASD connection

$$A_k \in \mathcal{A}^{1,p}(\mathcal{T}_{k, \Gamma_k})$$

on X_k . By performing the cutting off procedure described in Sections 7.2.1 and 4.4.5 of [5], we can form a connection A'_k on E_k that is equal to A_k outside of the ball $B_b(x_k)$ and equal to the flat connection $A_{b,k}$ inside of the ball $B_{b/2}(x_k)$. Then the A'_k patch together to determine a smooth connection $A' = A'(A_1, A_2)$ on E ; this depends on ρ, L, λ and the A_k . It follows that A' is equal to A_k in $X_k \setminus B_b(x_k) \subseteq X$ and that A' is approximately mASD:

$$(3.1) \quad \|s(A')\|_{L^p_\delta(X, g_{L, \lambda})} \leq C_{(3.1)} b^{4/p}$$

where $C_{(3.1)}$ is a constant independent of L, λ (see (7.2.36) in [5]). We will refer to A' as the *preglued connection*. We define the maps i and ι of Section 2B.1 by taking $A_{ref} := A'$.

Remark 3.2. Assume $2 < p < 4$ and set $p^* = 4p/(4-p)$. By [5, Eq. (7.2.37)], as $b \rightarrow 0$, the connections A'_k converge in $L_\delta^{p^*}$ to A_k . In particular, by Remark 2.9, if A_k is irreducible, then so too is A'_k provided b is sufficiently small. The stabilizer group of A' is contained in that of A'_k and so it follows that A' is irreducible when either of A_1 or A_2 is irreducible and b is sufficiently small.

On X_k , use the slice $\mathcal{SL}(A_k)$ defined by A_k and denote by $H_k^+ := H_{A_k, \delta}^+$ the cokernel of the linearized operator $D_k := (Ds|_{\mathcal{SL}(A_k)})_{A_k}$, as in (2.13). As described in [5, p. 290], we can choose lifts

$$\sigma_k: H_k^+ \longrightarrow L_\delta^p(\Omega^+(X_k))$$

so that the operator $D_k \oplus \sigma_k$ is surjective. Moreover, we can do this in such a way that, for every $v \in H_k^+$, the form $\sigma_k(v)$ is supported in the complement of the ball $B_{2b}(x_k)$. Set

$$H^+ := H_1^+ \oplus H_2^+$$

and consider the linear map

$$\sigma := \sigma_1 \oplus \sigma_2: H^+ \longrightarrow L_\delta^p(\Omega^+(X), g_{L,\lambda}).$$

Relative to the $L_\delta^p(X, g_{L,\lambda})$ -norm on H^+ , this map σ is bounded with a bound independent of L and λ .

3B. Gluing two connections. The main result of this section is as follows.

Theorem 3.3. *Assume $2 \leq p < 4$ and set $p^* = 4p/(4-p)$. Fix ρ, δ , thickening data, and mASD connections A_1, A_2 as in Section 3A. Then there are constants $C, L, \lambda_0 > 0$ so that the following holds for each $0 < \lambda < \lambda_0$.*

Let $A' = A'(A_1, A_2)$ be the preglued connection constructed from ρ, L, λ , and the A_k .

- (a) *There is a C^m -map $J_{A_1, A_2}: L_\delta^p(\Omega^+(X), g_{L,\lambda}) \rightarrow \mathcal{SL}(A') \subseteq \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ that satisfies $J_{A_1, A_2}(0) = A'$. The first m derivatives of $\xi \mapsto J_{A_1, A_2}(\xi)$ are bounded in operator norm by a bound that is independent of λ .*
- (b) *There is a linear map $\pi: L_\delta^p(\Omega^+(X), g_{L,\lambda}) \rightarrow H^+$ satisfying $\sigma \circ \pi \circ \sigma = \sigma$ and*

$$\|\pi \xi\|_{H^+} \leq C \|\xi\|_{L_\delta^p(X, g_{L,\lambda})} \quad \forall \xi \in \Omega^+(X).$$

- (c) *There is a unique 2-form $\xi(A_1, A_2) \in L_\delta^p(\Omega^+(X))$ so that*

$$(3.4) \quad \|\xi(A_1, A_2)\|_{L_\delta^p(X, g_{L,\lambda})} \leq C b^{4/p}$$

and so that the connection $\mathcal{J}(A_1, A_2) := J_{A_1, A_2}(\xi(A_1, A_2))$ satisfies

$$(3.5) \quad s(\mathcal{J}(A_1, A_2)) = -\sigma \pi \xi(A_1, A_2).$$

In particular, for $k = 1, 2$ the connection $\mathcal{J}(A_1, A_2)$ is close to A_k on $X_k \setminus B_{L\lambda^{1/2}}(x_k) \subseteq X$ in the sense that

$$(3.6) \quad \|\iota^{-1}(\mathcal{J}(A_1, A_2)) - \iota^{-1}(A_k)\|_{L_\delta^2(N_k) \times L_\delta^{p^*}(X_k \setminus B_{L\lambda^{1/2}}(x_k))} \leq C b^{4/p}.$$

If A_1 and A_2 are regular, then the connection $\mathcal{J}(A_1, A_2)$ is regular. In this case, $\mathcal{J}(A_1, A_2)$ is mASD and the maps $(A_1, A_2, \xi) \mapsto J_{A_1, A_2}(\xi)$ and $(A_1, A_2) \mapsto \xi(A_1, A_2)$ are both C^m -smooth, relative to the specified topologies.

If $p > 2$ and either A_1 or A_2 is irreducible, then so is $\mathcal{J}(A_1, A_2)$.

Before getting to the proof, we briefly discuss the maps appearing in this theorem, and their analogues in the standard ASD theory; precise definitions of these maps are given in the proof, below. First, the most interesting part of the theorem is part (c), with parts (a) and (b) serving to set up (c). The map π from (b) is a measure of the failure of A' to be regular. It serves the same role here and enjoys the same properties as the map of the same name [5, pp. 290–291] in the ASD setting. As for J in (a), this map is formed from a near-right inverse P of $(Ds|_{S\mathcal{L}(A')})_A \oplus \sigma$, pre-composed by an exponential map (see Claim 1 below for a precise statement). As an example, in the special case where Γ is non-degenerate (so mASD = ASD) the relevant space of connections is an affine space, and this exponential map is simply given by the affine action. In this case, the map J_{A_1, A_2} simplifies to $J_{A_1, A_2}(\xi) = A' + P\xi$, just as in the usual ASD setting [5, p. 289].

The object $\mathcal{J}(A_1, A_2)$ is the glued connection we are after. Equation (3.5) expresses the degree to which this connection is mASD. In particular, the obstruction map mentioned in the introduction can be taken to be the map $(A_1, A_2) \mapsto \pi \circ \xi(A_1, A_2)$. Finally we mention that the diffeomorphism ι^{-1} appears in (3.6) only to make explicit the sense in which $\mathcal{J}(A_1, A_2)$ approximates the A_k away from the gluing points.

The proof of Theorem 3.3 that we adopt relies on the following two lemmas.

Lemma 3.7. *Let $S: B \rightarrow B$ be a C^m -map on a Banach space B with $S(0) = 0$ and*

$$(3.8) \quad \|S(\xi_1) - S(\xi_2)\| \leq \kappa(\|\xi_1\| + \|\xi_2\|)\|\xi_1 - \xi_2\|,$$

for some $\kappa > 0$ and all $\xi_1, \xi_2 \in B_1(0) \subset B$ in the unit ball. Then for each $\eta \in B$ with $\|\eta\| < 1/(10\kappa)$, there is a unique $\xi \in B$ with $\|\xi\| \leq 1/(5\kappa)$ such that

$$\xi + S(\xi) = \eta.$$

Moreover, if $\eta = \eta(a)$ depends C^m -smoothly on a parameter a , then the solution $\xi = \xi(a)$ depends C^m -smoothly on this parameter as well.

Proof. The existence and uniqueness claims follow from the contraction mapping principle and is carried out in [5, Lemma 7.2.23]. The C^m -smooth dependence of ξ on a follows from, e.g., the discussion in [14, Section I.5]. \square

The remaining lemma will be used to establish the nonlinear estimate (3.8) in our mASD setting.

Lemma 3.9. *Assume $2 \leq p < 4$ and set $p^* = 4p/(4 - p)$. There is a constant $C_{(3.9)}$ so that if $L, \lambda > 0$ are any constants for which the connected sum X is defined, then*

$$\|fg\|_{L_\delta^p(X, g_{L, \lambda})} \leq C_{(3.9)} \|f\|_{L_\delta^{p^*}(X, g_{L, \lambda})} \|g\|_{L_\delta^{p^*}(X, g_{L, \lambda})}$$

for all real-valued functions $f, g \in L_\delta^{p^*}(X)$.

Proof. Writing $X = X_0 \cup \text{End } X$, it suffices to show that there is a uniform constant C so that

$$\begin{aligned} \|fg\|_{L^p(X_0, g_{L,\lambda})} &\leq C \|f\|_{L^{p^*}(X_0, g_{L,\lambda})} \|g\|_{L^{p^*}(X_0, g_{L,\lambda})} \\ \|fg\|_{L^p_\delta(\text{End } X, g_{L,\lambda})} &\leq C \|f\|_{L^{p^*}_\delta(\text{End } X, g_{L,\lambda})} \|g\|_{L^{p^*}_\delta(\text{End } X, g_{L,\lambda})}. \end{aligned}$$

We begin with the estimate over $\text{End } X$. Note that the metric $g_{L,\lambda}$ is independent of L, λ over this region, so we do not need to worry about showing that any such constant C is independent of L, λ . To obtain the estimate, use the definition of the δ -dependent norms, together with Hölder's inequality to get

$$\begin{aligned} \|fg\|_{L^p_\delta(\text{End } X)} &= \|e^{t\delta/2} fg\|_{L^p(\text{End } X)} \\ &= \|(e^{-t\delta/2} e^{t\delta/2} f)(e^{t\delta/2} g)\|_{L^p(\text{End } X)} \\ &\leq \|e^{-t\delta/2} (e^{t\delta/2} f)\|_{L^4(\text{End } X)} \|e^{t\delta/2} g\|_{L^{p^*}(\text{End } X)} \end{aligned}$$

Since $2 \leq p < 4$, we have $4 \leq p^* < \infty$. Hence there is some $4 < r \leq \infty$ with $4^{-1} = r^{-1} + (p^*)^{-1}$. Then we can use Hölder's inequality again to continue the above:

$$\|fg\|_{L^p_\delta(\text{End } X)} \leq \|e^{-t\delta/2}\|_{L^r(\text{End } X)} \|f\|_{L^{p^*}_\delta(\text{End } X)} \|g\|_{L^{p^*}_\delta(\text{End } X)}.$$

Then the requisite estimate holds with $C = \|e^{-t\delta/2}\|_{L^r(\text{End } X)}$, which is plainly finite.

As for the estimate over X_0 , the same type of argument gives

$$\|fg\|_{L^p(X_0, g_{L,\lambda})} \leq \text{vol}(X_0, g_{L,\lambda})^{1/r} \|f\|_{L^{p^*}(X_0, g_{L,\lambda})} \|g\|_{L^{p^*}(X_0, g_{L,\lambda})}.$$

As discussed on p. 293 of [5], the condition $p \geq 2$ implies that $\text{vol}(X_0, g_{L,\lambda})$ can be taken to be independent of L, λ , provided $L\lambda^{1/2}$ is uniformly bounded from above (which is necessarily the case whenever the connected sum is defined). \square

Proof of Theorem 3.3. Our intention is to apply Lemma 3.7. To do this, we need to recast solving $s(A) = 0$ for A into solving an equation for a self-mapping S of a Banach space. Ultimately, the Banach space will be the codomain of the mASD operator s , and S will essentially be the quadratic part of s .

We begin this process by passing to a local chart on $\mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ (recall from Section 2B.1 that this space of connections is generally *not* an affine space). For this, write

$$A' = \iota(h', V') = i(h') + V'$$

for $(h', V') \in \mathcal{H}_{out} \times L^p_{1,\delta}(\Omega^1(X))$. Let $\exp_{h'}: B_\epsilon(0) \subset T_{h'}\mathcal{H} \rightarrow \mathcal{H}$ be the exponential map associated to the $L^2_2(N)$ -metric on $\mathcal{H} := \mathcal{H}_\Gamma$; here $\epsilon > 0$ is small enough so that the exponential is well-defined. This is all taking place on the 3-manifold N , and so this exponential and this ϵ are manifestly independent of L, λ . Coupling this exponential on \mathcal{H} with the exponential on $\Omega^1(X)$ given by the affine action, we obtain a map

$$\begin{aligned} \exp_{(h', V')} : B_\epsilon(0) \times L^p_{1,\delta}(\Omega^1(X)) &\longrightarrow \mathcal{H} \times L^p_{1,\delta}(\Omega^1(X)) \\ (\eta, V) &\longmapsto (\exp_{h'}(\eta), V' + V). \end{aligned}$$

The chart for $\mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ that we will use is $\iota \circ \exp_{(h', V')}$. We note that $\iota \circ \exp_{(h', V')}$ identifies the “slice” $B_\epsilon(0) \times L^p_{1,\delta}(\ker(d^*_{A'}))$ with a neighborhood of A' in $\mathcal{S}\mathcal{L}(A')$.

Remark 3.10. Throughout the proof that follows, we will work with the $L_2^2(N)$ -norm on $T_{h'}\mathcal{H}$; we will often not keep track of this in the notation. Note that this choice of norm is effectively immaterial since \mathcal{H} is finite-dimensional and so any two norms are equivalent, provided they are well-defined.

Consider the map

$$\begin{aligned} \tilde{s}: B_\epsilon(0) \times L_{1,\delta}^p(\ker(d_{A'}^{*,\delta})) &\longrightarrow L_\delta^p(\Omega^+(X), g_{L,\lambda}) \\ (\eta, V) &\longmapsto s\left(\iota(\exp_{(h',V')}(\eta, V))\right) \end{aligned}$$

which is the map $s|_{\mathcal{S}\mathcal{L}(A')}$ relative to the chart just described. This satisfies $\tilde{s}(0,0) = s(A')$ and so (3.1) gives

$$(3.11) \quad \|\tilde{s}(0,0)\|_{L^p(X,g_{L,\lambda})} \leq C_{(3.1)} b^{4/p}.$$

Write $(D\tilde{s})_{(\eta,V)}$ for the linearization of \tilde{s} at (η, V) . As the following claim indicates, the definition of σ implies that the operator $(D\tilde{s})_{(0,0)} \oplus \sigma$ is surjective.

Claim 1: For $2 \leq p < 4$, there are constants $C_{(3.12)}, \lambda_0 > 0$, and $L > 1$, as well as linear maps

$$\begin{aligned} P: L_\delta^p(\Omega^+(X), g_{L,\lambda}) &\longrightarrow T_{h'}\mathcal{H} \times L_{1,\delta}^p(\ker(d_{A'}^{*,\delta}), g_{L,\lambda}) \\ \pi: L_\delta^p(\Omega^+(X), g_{L,\lambda}) &\longrightarrow H^+ \end{aligned}$$

so that $P \oplus \pi$ is a right inverse to $(D\tilde{s})_{(0,0)} \oplus \sigma$ that, for all $0 < \lambda < \lambda_0$, satisfies

$$(3.12) \quad \|(P \oplus \pi)\xi\|_{(L_2^2(N) \times L_\delta^{p*}(X, g_{L,\lambda})) \oplus L_\delta^p(X, g_{L,\lambda})} \leq C_{(3.12)} \|\xi\|_{L_\delta^p(X, g_{L,\lambda})} \quad \forall \xi \in \Omega^+(X).$$

This claim also has an extension to some $p < 2$; see Corollary 3.26 for details. We will prove Claim 1 shortly. At the moment, we will show how we use it to finish the proof of the theorem. To prove Theorem 3.3 (a), set

$$J(\xi) = J_{A_1, A_2}(\xi) := \iota(\exp_{(h',V')}(\xi)) \in \mathcal{S}\mathcal{L}(A').$$

Clearly $J(0) = A'$ and J is an immersion near 0. That the derivatives of J are bounded uniformly (λ -independent) follows from the fact that P is uniformly bounded (by the claim) and the fact that the maps ι and $\exp_{(h',V')}$ are uniformly bounded. The map in Theorem 3.3 (b) is the map π from Claim 1.

We now prove Theorem 3.3 (c). Define $S: L_\delta^p(\Omega^+(X), g_{L,\lambda}) \rightarrow L_\delta^p(\Omega^+(X), g_{L,\lambda})$ by

$$S(\xi) := \tilde{s}(P\xi) - (D\tilde{s})_{(0,0)}P\xi - \tilde{s}(0,0).$$

This is \mathcal{C}^m and is the nonlinear part of the map $\tilde{s} \circ P$. The following claim says that this map satisfies the requisite nonlinear estimates.

Claim 2: The map S satisfies the hypotheses of Lemma 3.7 with a constant κ that is independent of $0 < \lambda < \lambda_0$.

Once again, we defer the proof of this claim until after we have finished the argument for Theorem 3.3. It follows from Claim 2, Lemma 3.7, and the estimate (3.11) that, provided λ is sufficiently small, there is a unique $\xi = \xi(A_1, A_2) \in L_\delta^p(\Omega^+)$ satisfying

$$(3.13) \quad \xi + S(\xi) = -\tilde{s}(0,0) \quad \text{and} \quad \|\xi\|_{L_\delta^p(X)} \leq 1/(5\kappa).$$

Setting $\|\cdot\|_{L_\delta^p} := \|\cdot\|_{L_\delta^p(X, g_{L,\lambda})}$ and using (3.8), we get

$$\|\xi\|_{L_\delta^p} \leq \|\tilde{s}(0,0)\|_{L_\delta^p} + \|S(\xi)\|_{L_\delta^p} \leq C_{(3.1)} b^{4/p} + \frac{1}{5} \|\xi\|_{L_\delta^p}.$$

This implies the requisite estimate on ξ . Unraveling the definitions, we also have

$$s(\mathcal{J}(A_1, A_2)) = \tilde{s}(P\xi) = \tilde{s}(0,0) + (D\tilde{s})_{(0,0)} P\xi + S(\xi) = -\xi + (D\tilde{s})_{(0,0)} P\xi = -\sigma\pi\xi,$$

where $\mathcal{J}(A_1, A_2) := J(\xi)$, by definition. This finishes the proof of (c).

To prove (3.6), note that

$$\iota^{-1}(\mathcal{J}(A_1, A_2)) - \iota^{-1}(A') = \exp_{(h', V')} (P\xi) - \exp_{(h', V')} (0).$$

From our definition of the exponential appearing on the right, we see that the difference $\exp_{(h', V')} (P\xi) - \exp_{(h', V')} (0)$ equals $P\xi$ plus some higher order terms supported only on the center manifold component, and thus in a component with a norm independent of λ . In particular, since ξ is small, these higher order terms can be uniformly controlled to yield a uniform first order estimate of the form

$$\|\iota^{-1}(\mathcal{J}(A_1, A_2)) - \iota^{-1}(A')\|_{L_2^2(N) \times L_\delta^{p^*}(X, g_{L,\lambda})} \leq C_1 \|P\xi\|_{L_2^2(N) \times L_\delta^{p^*}(X, g_{L,\lambda})}$$

We can then combine this with (3.12) to get

$$(3.14) \quad \|\iota^{-1}(\mathcal{J}(A_1, A_2)) - \iota^{-1}(A')\|_{L_2^2(N) \times L_\delta^{p^*}(X, g_{L,\lambda})} \leq C_1 C_{(3.12)} \|\xi\|_{L_\delta^p(X, g_{L,\lambda})} \leq C_2 b^{4/p}$$

where the second inequality comes from the estimates of the previous paragraph. The estimate (3.6) follows from this and the fact that A_k agrees with $A' = A'(A_1, A_2)$ on $X_k \setminus B_{L\lambda^{1/2}}(x_k)$.

When the A_k are regular, the map π is the zero map so $\mathcal{J}(A_1, A_2)$ is automatically mASD by (3.5). In this case, the operator $(D\iota)_{(h', V')} \circ P$ is a right inverse to $(Ds|_{\mathcal{S}\mathcal{L}(A')})_{A'}$, essentially by definition; in particular, $(Ds|_{\mathcal{S}\mathcal{L}(A')})_{A'}$ is surjective. Thus $(Ds|_{\mathcal{S}\mathcal{L}(A)})_A$ is surjective whenever A is close to A' . In particular, by (3.14), this is the case with $A = \mathcal{J}(A_1, A_2)$ and so $\mathcal{J}(A_1, A_2)$ is regular. The \mathcal{C}^m -smooth dependence of J on the A_k follows from Remark 3.19 (a) below, and the \mathcal{C}^m -smoothness of ξ follows from the \mathcal{C}^m -smoothness assertion of Lemma 3.7.

Finally, assume A_1 or A_2 is irreducible (assume $p > 2$ so this term is defined). It follows from Remark 3.2 that A' is irreducible as well. The irreducibility of $\mathcal{J}(A_1, A_2)$ then follows from (3.14) and Lemma 2.8.

To finish the proof of Theorem 3.3, it therefore suffices to verify the claims; we begin with Claim 1. Let \mathcal{H}_k be the space \mathcal{H}_{out} for the connection Γ_k , and set

$$h_k = p_T(A_k) \in \mathcal{H}_k$$

where p_T is the map from Section **2B.1**. Similar to what we did above over X , for each k , we can form a map

$$\tilde{s}_k: T_{h_k} \mathcal{H}_k \times L_{1,\delta}^p(\ker(d_{A_k}^{*,\delta})) \longrightarrow L_\delta^p(\Omega^+(X_k))$$

by precomposing s with ι and the exponential \exp_{h_k} for \mathcal{H}_k based at h_k . Linearizing at $(0,0)$, and coupling with σ_k , we obtain a bounded linear map

$$\tilde{D}_k := (D\tilde{s}_k)_{(0,0)} \oplus \sigma_k: \left(T_{h_k} \mathcal{H}_k \times L_{1,\delta}^p(\ker(d_{A_k}^{*,\delta})) \right) \oplus H_k^+ \longrightarrow L_\delta^p(\Omega^+(X_k)).$$

This is surjective by the construction of σ_k . Standard elliptic theory for δ -decaying spaces [15] and the finite-dimensionality of \mathcal{H}_k imply that \tilde{D}_k restricts to a bounded map of the form

$$\tilde{D}_k: \left(T_{h_k} \mathcal{H}_k \times L_{\ell+1,\delta}^p(\ker(d_{A_k}^{*,\delta})) \right) \oplus H_k^+ \longrightarrow L_{\ell,\delta}^p(\Omega^+)$$

for each $\ell \geq 0$, and this restriction remains surjective. In particular, the ‘‘Laplacian’’

$$\tilde{D}_k \tilde{D}_k^{*,\delta}: L_{2,\delta}^p(\Omega^+) \longrightarrow L_\delta^p(\Omega^+)$$

is a Banach space isomorphism, where $\tilde{D}_k^{*,\delta}$ is the adjoint of \tilde{D}_k relative to the δ -decaying L^2 -inner products on the domain and codomain. It follows from these observations that the formula

$$P_k := \tilde{D}_k^{*,\delta} (\tilde{D}_k \tilde{D}_k^{*,\delta})^{-1}: L_\delta^p(\Omega^+) \longrightarrow \left(T_{h_k} \mathcal{H}_k \times L_{1,\delta}^p(\ker(d_{A_k}^{*,\delta})) \right) \oplus H_k^+$$

defines a bounded right inverse for \tilde{D}_k . Coupling this with the embedding $L_{1,\delta}^p \hookrightarrow L_\delta^{p*}$ it follows that there is a constant c_k with

$$(3.15) \quad \|P_k \tilde{\xi}\|_{(L_2^2(N_k) \times L_\delta^{p*}(X_k)) \oplus L_\delta^p(X_k)} \leq c_k \|\tilde{\xi}\|_{L_\delta^p(X_k)}, \quad \forall \tilde{\xi} \in \Omega^+(X_k).$$

The argument at this stage is almost identical to that given in [5, Section 7.2.3] (see also [5, Prop. 7.2.18]); however, we supply some of the details since we will refer to them below. Following [5, p. 288], the operators P_1, P_2 can be glued together to produce an operator

$$Q: L_\delta^p(\Omega^+(X), g_{L,\lambda}) \longrightarrow \left(T_{h'} \mathcal{H} \times L_{1,\delta}^p(\Omega^1(X), g_{L,\lambda}) \right) \oplus H^+$$

that satisfies

$$\|Q \tilde{\xi}\|_{(L_2^2(N) \times L_\delta^{p*}(X, g_{L,\lambda})) \oplus L_\delta^p(X, g_{L,\lambda})} \leq (c_1 + c_2) \|\tilde{\xi}\|_{L_\delta^p(X, g_{L,\lambda})} \quad \forall \tilde{\xi} \in \Omega^+(X).$$

Somewhat more explicitly, we have

$$Q = \beta_1 P_1 \gamma_1 + \beta_2 P_2 \gamma_2$$

where the $\{\beta_1, \beta_2\}$ and $\{\gamma_1, \gamma_2\}$ are partitions of unity on X . The only property we need about these partitions is that the derivatives $\nabla \beta_k$ are supported in the gluing region and satisfy $\|\nabla \beta_k\|_{L^4(X)} \rightarrow 0$ as $L \rightarrow \infty$ and $b \rightarrow 0$. For future reference we note

that since A' equals the A_k away from the gluing region and the P_k take values in the A_k -slice, there is a uniform constant C_3 so that

$$(3.16) \quad \|d_{A'}^{*,\delta} Q\tilde{\xi}\|_{L_\delta^p} \leq C_3(\|\nabla\beta_1\|_{L^4} + \|\nabla\beta_2\|_{L^4})\|Q\tilde{\xi}\|_{L_\delta^{p*}}.$$

It also follows that Q is an approximate right inverse to $(D\tilde{s})_{(0,0)} \oplus \sigma$ in the sense that

$$((D\tilde{s})_{(0,0)} \oplus \sigma) \circ Q = I + R$$

for some R satisfying

$$\|R(\tilde{\xi})\|_{L_\delta^p} \leq \epsilon(L, b, p)\|\tilde{\xi}\|_{L_\delta^p}$$

where $\epsilon(L, b, p) \rightarrow 0$ as $L \rightarrow \infty$ and $b \rightarrow 0$ (the assumption $p \geq 2$ is used here to establish this decay property for $\epsilon(L, b, p)$, see [5, pp. 293,294]).

At this stage, Donaldson and Kronheimer [5] take their right inverse to be $Q(I + R)^{-1}$. However, this is not sufficient for us since we want our right inverse to take values in the slice, and the operator $Q(I + R)^{-1}$ generally does not. We thus want to modify the construction above, and we do this by simply projecting to the slice. Specifically, let Π be the L_δ^2 -orthogonal projection of $\Omega^1(X)$ to the kernel of $d_{A'}^{*,\delta}$ (note that $d_{A'}^{*,\delta}$ is injective on the image of $I - \Pi$). This projection Π extends to the codomain of Q by acting on the $\Omega^1(X)$ -factor only. Thus, the map

$$\tilde{Q} := \Pi \circ Q$$

takes values in the slice. We want to use this \tilde{Q} in place of Q , but for this we need to port the estimates on Q over to \tilde{Q} . To achieve this goal, note that the difference $Q - \tilde{Q}$ takes values in the image of $I - \Pi$, on which the operator $d_{A'}^{*,\delta}$ is injective. From this we have the following Sobolev-type estimate

$$\|Q\tilde{\xi} - \tilde{Q}\tilde{\xi}\|_{(L_2^2(N) \times L_\delta^{p*}(X, g_{L,\lambda})) \oplus L_\delta^p(X, g_{L,\lambda})} \leq C_4 \|d_{A'}^{*,\delta}(Q\tilde{\xi} - \tilde{Q}\tilde{\xi})\|_{L_\delta^p}$$

for all $\tilde{\xi} \in L_\delta^p(\Omega^+(X))$. As in [5, Section 7.2.3], our range restriction on p implies that this constant C_4 can be taken to be uniform in the neck-scaling parameter b . By construction, \tilde{Q} takes values in the kernel of $d_{A'}^{*,\delta}$ and so we can combine the above with (3.16) to get

$$(3.17) \quad \|Q\tilde{\xi} - \tilde{Q}\tilde{\xi}\|_{(L_2^2(N) \times L_\delta^{p*}(X, g_{L,\lambda})) \oplus L_\delta^p(X, g_{L,\lambda})} \leq C_3 C_4 (\|\nabla\beta_1\|_{L^4} + \|\nabla\beta_2\|_{L^4}) \|Q\tilde{\xi}\|_{L_\delta^{p*}}$$

By taking L large and b sufficiently small, we can now transfer our estimate from Q to \tilde{Q} to conclude that \tilde{Q} is uniformly bounded and is an approximate right inverse to $(D\tilde{s})_{(0,0)} \oplus \sigma$ in the sense that

$$((D\tilde{s})_{(0,0)} \oplus \sigma) \circ \tilde{Q} = I + \tilde{R}$$

for some \tilde{R} satisfying

$$(3.18) \quad \|\tilde{R}(\tilde{\xi})\|_{L_\delta^p} \leq \tilde{\epsilon}(L, b, p)\|\tilde{\xi}\|_{L_\delta^p}$$

where $\tilde{\epsilon}(L, b, p) \rightarrow 0$ as $L \rightarrow \infty$ and $b \rightarrow 0$. The whole point, of course, is that

$$\tilde{Q}: L_\delta^p(\Omega^+(X), g_{L,\lambda}) \longrightarrow \left(T_{h'}\mathcal{H} \times L_{1,\delta}^p(\ker(d_{A'}^{*,\delta})) \right) \oplus H^+$$

takes values in the slice where Q may not have.

Choose $L > 1, \lambda_0 > 0$ so that $\tilde{\epsilon}(L, 4L\lambda_0^{1/2}, p) < 1/3$ and $C_3C_4(\|\nabla\beta_1\|_{L^4} + \|\nabla\beta_2\|_{L^4}) < 1$. Then $\tilde{Q}(I + \tilde{R})^{-1}$ is a right inverse to $\tilde{D} \oplus \sigma$ and has operator norm at most $3(c_1 + c_2 + 1)$. Then we can write this right inverse as $\tilde{Q}(I + \tilde{R})^{-1} = P \oplus \pi$, where the splitting corresponds to the direct sum decomposition of the codomain of \tilde{Q} . The estimate (3.12) immediately follows, so this finishes the proof of Claim 1.

Remark 3.19. (a) It is not hard to show from the construction outlined in [5] that, when the A_k are regular, then the right inverse P depends \mathcal{C}^m -smoothly on the A_k . The key observation here is that, though many choices have been made in this construction (e.g., cutoff functions), the only ones that depend on the A_k are the choices of σ_k , but these can be taken to be zero when the A_k are regular.

(b) By construction, the connection $\mathcal{J}(A_1, A_2)$ of Theorem 3.3 naturally belongs to two slices: The one centered at itself, and the one centered at A' .

Now we move on to prove Claim 2. Fix L, λ_0 as in Claim 1 and we assume $\lambda \in (0, \lambda_0)$. We clearly have $S(0) = 0$, so it suffices to show that S satisfies the quadratic estimate (3.8) for a uniform constant κ . For this, note that by Lemma 2.2 and Taylor's Theorem, we can write

$$i(\exp_{h'}(\eta)) = i(h') + (Di)_{h'}\eta + q_{h'}(\eta)$$

where $q_{h'}: T_{h'}\mathcal{H} \rightarrow L_{1,loc}^p(X) \cap \mathcal{C}^0(X)$ vanishes to first order. Since \mathcal{H}_{out} is finite-dimensional, we can quantify this relative to any metric with respect to which the terms are well-defined. In particular, there is a constant $C_{(3.20)}$ so that

$$(3.20) \quad \|q_{h'}(\eta_1) - q_{h'}(\eta_2)\|_{\mathcal{C}^0(X)} \leq C_{(3.20)}(\|\eta_1\|_{L_2^2(N)} + \|\eta_2\|_{L_2^2(N)})\|\eta_1 - \eta_2\|_{L_2^2(N)}$$

for all η_1, η_2 in the unit ball in $T_{h'}\mathcal{H}$. Note that $q_{h'}(\eta)$ need not decay to zero down the ends of X since $i(h')$ and $i(\exp_{h'}(\eta))$ generally do not converge to the same connection at infinity. However, on the compact part we have

$$(3.21) \quad q_{h'}(\eta)|_{X_0} = 0.$$

Indeed, on X_0 the connection $i(h)$ equals the reference connection for all $h \in \mathcal{H}_{out}$, and i vanishes to all but the zeroth order on X_0 .

To verify (3.8), fix $\xi_1, \xi_2 \in L_\delta^p(\Omega^+(X), g_\lambda)$ with $\|\xi_j\|_{L_\delta^p} \leq 1$ and set

$$(\eta_j, V_j) := P\xi_j \in T_{h'}\mathcal{H} \times L_{1,\delta}^p(\Omega^1(X)).$$

Then using the definition of S and the formula (2.11), we can write

$$(3.22) \quad \begin{aligned} S(\xi_j) &= \frac{1}{2} [V_j \wedge V_j]^+ + \frac{1-\beta'}{2} [((Di)_{h'}\eta_j + q_{h'}(\eta_j)) \wedge ((Di)_{h'}\eta_j + q_{h'}(\eta_j))]^+ \\ &\quad + (1-\beta')d_{i(h')}^+ q_{h'}(\eta_j) + [V' \wedge q_{h'}(\eta_j)]^+ \\ &\quad + [V_j \wedge q_{h'}(\eta_j)]^+ + [V_j \wedge (Di)_{h'}\eta_j]^+. \end{aligned}$$

(These are the higher order terms in the mASD operator s , expressed in terms of V_j and η_j .) It suffices to show that each term on the right satisfies an estimate of the form (3.8). Below we set $\|\cdot\|_{L_\delta^p} := \|\cdot\|_{L_\delta^p(X, g_\lambda)}$.

We begin with the first term on the right of (3.22). This shows up in the ASD setting as well (see [5, p. 289]), but our argument is a bit more involved due to the non-compactness of X . For this, we use Lemma 3.9 to get

$$\begin{aligned} \frac{1}{2} \|[V_1 \wedge V_1]^+ - [V_2 \wedge V_2]^+\|_{L_\delta^p} &= \frac{1}{2} \|[(V_1 + V_2) \wedge (V_1 - V_2)]^+\|_{L_\delta^p} \\ &\leq \mathfrak{c}_g \|[V_1 + V_2][V_1 - V_2]\|_{L_\delta^p} \\ &\leq \mathfrak{c}_g C_{(3.9)} \left(\|V_1\|_{L_\delta^{p^*}} + \|V_2\|_{L_\delta^{p^*}} \right) \|V_1 - V_2\|_{L_\delta^{p^*}} \end{aligned}$$

where \mathfrak{c}_g is defined by

$$(3.23) \quad \mathfrak{c}_g := \sup_{(v_1, v_2, v_3)} |\langle v_1, [v_2, v_3] \rangle|$$

with the supremum running over all $v_j \in \mathfrak{g}$ with $|v_j| = 1$. Since V_j is a component of $P\tilde{\zeta}_j$, we can then use the estimate of Claim 1 to continue the above and get

$$\frac{1}{2} \|[V_1 \wedge V_1]^+ - [V_2 \wedge V_2]^+\|_{L_\delta^p} \leq \mathfrak{c}_g C_{(3.9)} C_{(3.12)}^2 \left(\|\tilde{\zeta}_1\|_{L_\delta^p} + \|\tilde{\zeta}_2\|_{L_\delta^p} \right) \|\tilde{\zeta}_1 - \tilde{\zeta}_2\|_{L_\delta^p}$$

which is the desired estimate.

Now we move on to the second term in (3.22). Set $r(\eta) := (Di)_{h'}\eta + q_{h'}(\eta)$, so we want to bound the L_δ^p -norm of

$$\frac{1-\beta'}{2} \left([r(\eta_1) \wedge r(\eta_1)]^+ - [r(\eta_2) \wedge r(\eta_2)]^+ \right) = \frac{1-\beta'}{2} \left[(r(\eta_1) + r(\eta_2)) \wedge (r(\eta_1) - r(\eta_2)) \right]^+$$

in terms of the right-hand side of (3.8). Note that this is supported on the compact cylinder $\text{Cyl}_0 := [T - 1/2, T + 1/2] \times N$, and so its L_δ^p -norm is bounded by a constant times

$$\begin{aligned} &\left(\sum_{j=1}^2 \|r(\eta_j)\|_{L_\delta^{p^*}(\text{Cyl}_0)} \right) \|r(\eta_1) - r(\eta_2)\|_{L_\delta^{p^*}(\text{Cyl}_0)} \\ &\leq \|e^{t\delta/2}\|_{L^{p^*}(\text{Cyl}_0)}^2 \left(\sum_{j=1}^2 \|r(\eta_j)\|_{C^0(\text{Cyl}_0)} \right) \|r(\eta_1) - r(\eta_2)\|_{C^0(\text{Cyl}_0)} \end{aligned}$$

By Lemma 2.2, this is bounded by a constant times

$$\left(\|\eta_1\|_{L_2^2(N)} + \|\eta_2\|_{L_2^2(N)} \right) \|\eta_1 - \eta_2\|_{L_2^2(N)} \leq C_{(3.12)}^2 \left(\|\tilde{\zeta}_1\|_{L_\delta^p} + \|\tilde{\zeta}_2\|_{L_\delta^p} \right) \|\tilde{\zeta}_1 - \tilde{\zeta}_2\|_{L_\delta^p}$$

as desired.

The estimate for the third term $(1 - \beta')d_{i(h')}^+ q_{h'}(\eta_j)$ is similar and we leave it to the reader. Moving on to the fourth term in (3.22), recall from (3.21) that $q_{h'}(\eta_j)$ vanishes

on X_0 . This observation combines with (3.20) and then Claim 1 to give

$$\begin{aligned}
& \left\| [V' \wedge q_{h'}(\eta_1)]^+ - [V' \wedge q_{h'}(\eta_2)]^+ \right\|_{L_\delta^p} \\
&= \left\| [V' \wedge (q_{h'}(\eta_1) - q_{h'}(\eta_2))]^+ \right\|_{L_\delta^p(\text{End } X)} \\
&\leq \mathbf{c}_g \|V'\|_{L_\delta^p(\text{End } X)} \|q_{h'}(\eta_1) - q_{h'}(\eta_2)\|_{C^0(\text{End } X)} \\
&\leq \mathbf{c}_g C_{(3.20)} \|V'\|_{L_\delta^p(\text{End } X)} \left(\|\eta_1\|_{L_2^2(N)} + \|\eta_2\|_{L_2^2(N)} \right) \|\eta_1 - \eta_2\|_{L_2^2(N)} \\
&\leq \mathbf{c}_g C_{(3.20)} C_{(3.12)}^2 \|V'\|_{L_\delta^p(\text{End } X)} \left(\|\xi_1\|_{L_\delta^p} + \|\xi_2\|_{L_\delta^p} \right) \|\xi_1 - \xi_2\|_{L_\delta^p}
\end{aligned}$$

This is the desired estimate for this term because $\|V'\|_{L_\delta^p(\text{End } X)}$ is plainly independent of λ and the ξ_j .

The remaining two terms are the most difficult to bound. This is because (i) these terms involve both the infinite-dimensional terms V_i as well as the finite-dimensional terms $q_{h'}(\eta_j)$ and $(Di)_{h'}\eta_j$, and (ii) neither of these finite-dimensional terms generally decays to zero at infinity (nor do the differences $q_{h'}(\eta_1) - q_{h'}(\eta_2)$ and $(Di)_{h'}\eta_1 - (Di)_{h'}\eta_2$). The main estimate we need is the following, which we will see is equivalent to the fact that the operator $d_{A'}^+$ is Fredholm (on the appropriate spaces) with our choice of δ .

Claim 3: There is some $T_1 \gg T + 3/2$ and a constant $C_{(3.24)}$ so that

$$(3.24) \quad \|V\|_{L_\delta^p([T_1, \infty) \times N)} \leq C_{(3.24)} \|d_{A'}^+ V\|_{L_\delta^p([T_1 - 1, \infty) \times N)}$$

for all $V \in L_{1, \delta}^p(\Omega^1(X))$ with $d_{A'}^{*, \delta} V|_{[T_1 - 1, \infty) \times N} = 0$.

We prove Claim 3 after we finish our estimates for the last two terms in (3.22). The argument we give applies to both of these last two terms, so we focus on establishing the estimate for the second-to-last term:

$$\begin{aligned}
& [V_1 \wedge q_{h'}(\eta_1)]^+ - [V_2 \wedge q_{h'}(\eta_2)]^+ \\
&= \frac{1}{2} \left([(V_1 - V_2) \wedge (q_{h'}(\eta_1) + q_{h'}(\eta_2))]^+ + [(V_1 + V_2) \wedge (q_{h'}(\eta_1) - q_{h'}(\eta_2))]^+ \right).
\end{aligned}$$

It suffices to bound the L_δ^p -norm of each term on the right by the right-hand side of (3.8); we will carry this out for $[V_1 + V_2 \wedge q_{h'}(\eta_1) - q_{h'}(\eta_2)]^+$, the other term is similar. Since the $q_{h'}(\eta_j)$ are supported on $\text{End } X$, we do not need to worry whether our constants are λ -dependent. With T_1 as in Claim 3, write

$$\begin{aligned}
(3.25) \quad & \left\| [(V_1 + V_2) \wedge (q_{h'}(\eta_1) - q_{h'}(\eta_2))]^+ \right\|_{L_\delta^p} \\
& \leq \left(\left\| [(V_1 + V_2) \wedge (q_{h'}(\eta_1) - q_{h'}(\eta_2))]^+ \right\|_{L_\delta^p([0, T_1] \times N)} \right. \\
& \quad \left. + \left\| [(V_1 + V_2) \wedge (q_{h'}(\eta_1) - q_{h'}(\eta_2))]^+ \right\|_{L_\delta^p([T_1, \infty) \times N)} \right).
\end{aligned}$$

Set $\text{Cyl}_1 := [0, T_1] \times N$ and estimate the first term on the right as follows:

$$\begin{aligned}
& \left\| \left[(V_1 + V_2) \wedge (q_{h'}(\eta_1) - q_{h'}(\eta_2)) \right]^+ \right\|_{L_\delta^p(\text{Cyl}_1)} \\
& \leq \mathfrak{c}_g \left(\|V_1\|_{L_\delta^{p^*}(\text{Cyl}_1)} + \|V_2\|_{L_\delta^{p^*}(\text{Cyl}_1)} \right) \|q_{h'}(\eta_1) - q_{h'}(\eta_2)\|_{L_\delta^4(\text{Cyl}_1)} \\
& \leq \mathfrak{c}_g \|e^{t\delta/2}\|_{L^4(\text{Cyl}_1)} \left(\|V_1\|_{L_\delta^{p^*}} + \|V_2\|_{L_\delta^{p^*}} \right) \|q_{h'}(\eta_1) - q_{h'}(\eta_2)\|_{C^0(\text{Cyl}_1)} \\
& \leq \mathfrak{c}_g C_{(3.20)} C_{(3.12)}^3 \|e^{t\delta/2}\|_{L^4(\text{Cyl}_1)} \left(\|\xi_1\|_{L_\delta^p} + \|\xi_2\|_{L_\delta^p} \right)^2 \|\xi_1 - \xi_2\|_{L_\delta^p}
\end{aligned}$$

which is the desired estimate for this terms since we have assumed $\|\xi_j\|_{L_\delta^p} \leq 1$, and so

$$\left(\|\xi_1\|_{L_\delta^p} + \|\xi_2\|_{L_\delta^p} \right)^2 \leq 2 \left(\|\xi_1\|_{L_\delta^p} + \|\xi_2\|_{L_\delta^p} \right).$$

As for the remaining term on the right of (3.25), note that it is bounded by a constant times

$$\begin{aligned}
& \left(\|V_1\|_{L_\delta^p([T_1, \infty) \times N)} + \|V_2\|_{L_\delta^p([T_1, \infty) \times N)} \right) \|q_{h'}(\eta_1) - q_{h'}(\eta_2)\|_{C^0([T_1, \infty) \times N)} \\
& \leq C_{(3.20)} C_{(3.12)}^2 \left(\|V_1\|_{L_\delta^p([T_1, \infty) \times N)} + \|V_2\|_{L_\delta^p([T_1, \infty) \times N)} \right) \left(\|\xi_1\|_{L_\delta^p} + \|\xi_2\|_{L_\delta^p} \right) \|\xi_1 - \xi_2\|_{L_\delta^p}
\end{aligned}$$

We will therefore be done with the proof of Claim 2 if we can show that the terms $\|V_j\|_{L_\delta^p([T_1, \infty) \times N)}$ are bounded. For this, note that the linearization of \tilde{s} can be written as

$$\begin{aligned}
(D\tilde{s})_{(0,0)}(\eta, V) &= d_{A'}^+((Di)_{h'}\eta + V) - \beta' d_{i(h')}^+(Di)_{h'}\eta \\
&= d_{A'}^+V + (1 - \beta') d_{i(h')}^+(Di)_{h'}\eta + [V' \wedge (Di)_{h'}\eta]^+.
\end{aligned}$$

Note also that, by the definition of P , the 1-form V_j lies in the kernel of $d_{A'}^{*,\delta}$. We can therefore use Claim 3 and the above formula for $(D\tilde{s})_{(0,0)}$ to write

$$\begin{aligned}
\|V_j\|_{L_\delta^p([T_1, \infty) \times N)} &\leq C_{(3.24)} \|d_{A'}^+ V_j\|_{L_\delta^p([T_1 - 1, \infty) \times N)} \\
&\leq C_{(3.24)} \left(\|(D\tilde{s})_{(0,0)}(\eta_j, V_j)\|_{L_\delta^p} + \|[V' \wedge (Di)_{h'}\eta_j]^+\|_{L_\delta^p} \right)
\end{aligned}$$

where we also used the fact that $\beta' = 1$ on $[T_1 - 1, \infty) \times N$. Since $(\eta_j, V_j) = P\zeta_j$ and $P \oplus \pi$ is a right-inverse for $(D\tilde{s})_{(0,0)} \oplus \sigma$, we can continue this as

$$\begin{aligned}
&\leq C_{(3.24)} \left(\|\zeta_j\|_{L_\delta^p} + \|\sigma\pi\zeta_j\|_{L_\delta^p} + \mathfrak{c}_g \|V'\|_{L_\delta^p} \|(Di)_{h'}\eta_j\|_{C^0} \right) \\
&\leq C_{(3.24)} \left(1 + C_{\sigma\pi} + \mathfrak{c}_g C_{(Di)_{h'}} \|V'\|_{L_\delta^p} \|\eta_j\|_{L_2^2(N)} \right) \\
&\leq C_{(3.24)} \left(1 + C_{\sigma\pi} + \mathfrak{c}_g C_{(Di)_{h'}} C_{(3.12)} \|V'\|_{L_\delta^p} \right)
\end{aligned}$$

where $C_{\sigma\pi}$ and $C_{(Di)_{h'}}$ are the operator norms of $\sigma\pi$ and $(Di)_{h'}$, respectively (we are viewing the latter as a map $L_2^2(N) \rightarrow C^0(X)$; see Lemma 2.2 and the definition of i). This is the uniform bound we are after, and thus finishes the proof of Claim 2.

Finally, we prove Claim 3. Let $h'_T: [T, \infty) \rightarrow \mathcal{H}_{out}$ denote the flow of the trimmed vector field Ξ^{tr} with $h'_T(T) = h'$. Let

$$h'_\infty := \lim_{t \rightarrow \infty} h'_T(t) \in \mathcal{H}_{out}$$

be the limiting connection of this flow. This is a connection on N , but we will view it as a connection on $\text{End } X = [0, \infty) \times N$ that is constant in the t -direction. Let \mathcal{X} be the $L^p_{1,\delta}$ -completion of the space of 1-forms on X supported on $\text{End } X$, and let \mathcal{Y} be the L^p_δ -completion of the elements of $\Omega^+ \oplus \Omega^0$ supported on $\text{End } X$ (so the elements of \mathcal{X} and \mathcal{Y} vanish on the compact part). Then the map

$$d_{h'_\infty}^+ \oplus d_{h'_\infty}^{*,\delta}: \mathcal{X} \longrightarrow \mathcal{Y}$$

is bounded and elliptic. We have assumed that $\delta/2$ is not in the spectrum of $-*d_h$, so it follows that the above operator has trivial kernel (it also has trivial cokernel, though we do not need this). In particular, there is a constant C_5 so that

$$\|V\|_{L^p_\delta} \leq \|V\|_{L^p_{1,\delta}} \leq C_5 \left(\|d_{h'_\infty}^+ V\|_{L^p_\delta} + \|d_{h'_\infty}^{*,\delta} V\|_{L^p_\delta} \right)$$

for all $V \in \mathcal{X}$. The connection A' is \mathcal{C}^0 -asymptotic to h'_∞ . In particular, we can choose T_1 large enough so that

$$\|A' - h'_\infty\|_{\mathcal{C}^0([T_1-1, \infty) \times N)} < 1/(6c_g C_4).$$

Then if $V \in \mathcal{X}$ is supported on $[T_1 - 1, \infty) \times N$ and in the kernel of $d_{A'}^{*,\delta}$ we have

$$\begin{aligned} \|V\|_{L^p_\delta([T_1-1, \infty) \times N)} &\leq C_5 \left(\|d_{h'_\infty}^+ V\|_{L^p_\delta([T_1-1, \infty) \times N)} + \|d_{h'_\infty}^{*,\delta} V\|_{L^p_\delta([T_1-1, \infty) \times N)} \right) \\ &= C_5 \left(\|d_{A'}^+ V + [(h'_\infty - A') \wedge V]^+ \|_{L^p_\delta([T_1-1, \infty) \times N)} \right. \\ &\quad \left. + \|[(h'_\infty - A') \wedge *V]\|_{L^p_\delta([T_1-1, \infty) \times N)} \right) \\ &\leq C_5 \|d_{A'}^+ V\|_{L^p_\delta([T_1-1, \infty) \times N)} + \frac{1}{3} \|V\|_{L^p_\delta([T_1-1, \infty) \times N)}. \end{aligned}$$

Then Claim 3 follows (with $C_{(3.24)} := 3C_5/2$) from this estimate and a cutoff function. \square

3C. Extensions to $p < 2$. In our existence result of Section 6C, we will need extensions to $p < 2$ of the estimates (3.4) and (3.12); we state and prove the relevant extensions here. In fact, all we will need is an extension to $p = 4/3$ (so $p^* = 2$); we leave any more general extensions to the interested reader. Throughout this section, we fix data as in the statement of Theorem 3.3.

For the first result, let $L > 1$, $\lambda_0 > 0$, and $P \oplus \pi$ be as in the statement of Claim 1 appearing in the proof of Theorem 3.3.

Corollary 3.26. *There is a constant C so the following holds for all $0 < \lambda < \lambda_0$ and $\xi \in \Omega^+(X)$:*

$$\|(P \oplus \pi)\xi\|_{(L^2_2(N) \times L^2_\delta(X, g_{L,\lambda})) \oplus L^{4/3}_\delta(X, g_{L,\lambda})} \leq C \|\xi\|_{L^{4/3}_\delta(X, g_{L,\lambda})}.$$

Proof. We refer to the notation established in the proof of Claim 1. Momentarily suppressing Sobolev completions, let

$$(P \oplus \pi)^{*,\delta} : \left(T_{h'} \mathcal{H} \times \ker(d_{A'}^{*,\delta}) \right) \oplus H^+ \longrightarrow \Omega^+(X)$$

be the formal adjoint of $P \oplus \pi$, relative to the L_δ^2 -inner product on X . By the duality isometries $(L_\delta^2(X))^* \cong L_\delta^2(X)$, $(L_\delta^4(X))^* \cong L_\delta^{4/3}(X)$, and $(L_2^2(N))^* \cong L_{-2}^2(N)$, we will be done if we can establish a uniform bound of the form

$$\|(P \oplus \pi)^{*,\delta}(\eta, V, \mu)\|_{L_\delta^4(X)} \leq C_1 \left(\|\eta\|_{L_{-2}^2(N)} + \|V\|_{L_\delta^2(X)} + \|\mu\|_{L_\delta^4(X)} \right).$$

Since $T_{h'} \mathcal{H}$ is finite-dimensional, there is a bound of the form $\|\eta\|_{L_{-2}^2(N)} \leq C_2 \|\eta\|_{L_{-2}^2(N)}$ for all $\eta \in T_{h'} \mathcal{H}$. It therefore suffices to show

$$(3.27) \quad \|(P \oplus \pi)^{*,\delta}(\eta, V, \mu)\|_{L_\delta^4(X)} \leq C_3 \left(\|\eta\|_{L_2^2(N)} + \|V\|_{L_\delta^2(X)} + \|\mu\|_{L_\delta^4(X)} \right)$$

for a uniform constant C_3 . This is precisely the estimate of Claim 1, except with the adjoint operator $(P \oplus \pi)^{*,\delta}$ in place of $P \oplus \pi$. We will show that the proof of Claim 1 can be sufficiently modified to hold for this adjoint.

Towards this end, note that the adjoint of $P_k = \tilde{D}_k^{*,\delta} (\tilde{D}_k \tilde{D}_k^{*,\delta})^{-1}$ is given by $P_k^{*,\delta} = (\tilde{D}_k \tilde{D}_k^{*,\delta})^{-1} \tilde{D}_k$ and so satisfies

$$\|P_k^{*,\delta}(\eta_k, V_k, \mu_k)\|_{L_\delta^4(X_k)} \leq c_k \left(\|\eta_k\|_{L_2^2(N_k)} + \|V_k\|_{L_\delta^2(X_k)} + \|\mu_k\|_{L_\delta^4(X_k)} \right).$$

Just as before, these can be glued together to form an operator $Q^{*,\delta}$ that satisfies

$$(3.28) \quad \|Q^{*,\delta}(\eta, V, \mu)\|_{L_\delta^4(X)} \leq (c_1 + c_2) \left(\|\eta\|_{L_2^2(N)} + \|V\|_{L_\delta^2(X)} + \|\mu\|_{L_\delta^4(X)} \right).$$

Moreover, it is not hard to see that this gluing can be done so that $Q^{*,\delta}$ is exactly the formal L_δ^2 -adjoint of the operator Q appearing in the proof of Claim 1. It follows that the formal L_δ^2 -adjoint of \tilde{Q} is the restriction of $Q^{*,\delta}$ to the slice (this just restricts V to line in the kernel of $d_{A'}^{*,\delta}$).

Then the defining formula $P \circ \pi = \tilde{Q}(I + \tilde{R})^{-1}$ implies

$$(3.29) \quad (P \circ \pi)^* = (I + \tilde{R}^{*,\delta})^{-1} \tilde{Q}^{*,\delta}$$

where $\tilde{R}^{*,\delta}$ is the formal adjoint of \tilde{R} and so satisfies $\|\tilde{R}^{*,\delta} \xi\|_{L_\delta^2(X)} = \|\tilde{R} \xi\|_{L_\delta^2(X)}$. Then the estimate (3.27) follows from (3.29), (3.28), and (3.18). \square

For the second and last of the extensions we need, let $\xi(A_1, A_2) \in L_\delta^2(\Omega^+(X))$ be as in the conclusion of Theorem 3.3 (c).

Corollary 3.30. *There are $C, \lambda'_0 > 0$, so that if $0 < \lambda < \lambda'_0$ then*

$$\|\xi(A_1, A_2)\|_{L_\delta^{4/3}(X, g_{L,\lambda})} \leq C \lambda^{3/2}.$$

Proof. Setting $\tilde{\zeta} := \zeta(A_1, A_2)$, the identity in (3.13) gives

$$\|\tilde{\zeta}\|_{L_\delta^{4/3}(X)} \leq \|\tilde{s}(0, 0)\|_{L_\delta^{4/3}(X)} + \|S(\tilde{\zeta})\|_{L_\delta^{4/3}(X)}.$$

The estimate (3.1) holds with $p = 4/3$, so the same is true of (3.11); that is,

$$\|\tilde{s}(0, 0)\|_{L_\delta^{4/3}(X)} \leq C_1 b^3$$

for a uniform constant C_1 , where $b = 4L\lambda^{1/2}$. To estimate $S(\tilde{\zeta})$, note that the formula (3.22) implies that $S(\tilde{\zeta})$ is quadratically bounded in $P\tilde{\zeta}$. Then we can argue as we did in the proof of Claim 2, but use Hölder's inequality $\|fg\|_{L_\delta^{4/3}} \leq \|f\|_{L_\delta^4} \|g\|_{L^2}$, to get a uniform estimate of the form

$$\|S(\tilde{\zeta})\|_{L_\delta^{4/3}(X)} \leq C_2 \|P\tilde{\zeta}\|_{L_\delta^4(X)} \|P\tilde{\zeta}\|_{L_\delta^2(X)}.$$

By (3.12) and Corollary 3.26, this implies

$$\|S(\tilde{\zeta})\|_{L_\delta^{4/3}(X)} \leq C_3 \|\tilde{\zeta}\|_{L_\delta^2(X)} \|\tilde{\zeta}\|_{L_\delta^{4/3}(X)}.$$

It follows from (3.4) that we can assume $\|\tilde{\zeta}\|_{L_\delta^2(X)} < (2C_3)^{-1}$, provided $\lambda > 0$ is sufficiently small. In summary, this implies

$$\|\tilde{\zeta}\|_{L_\delta^{4/3}(X)} \leq C_1 b^3 + \frac{1}{2} \|\tilde{\zeta}\|_{L_\delta^{4/3}(X)}$$

from which the corollary follows with $C = 128L^3C_1$. \square

4. GAUGE FIXING AND THE mASD CONDITION

In the next section, we will find ourselves in the situation where we have an mASD connection A and a nearby connection A_{ref} . We will want to find a gauge transformation u so that u^*A is in the Coulomb slice of A_{ref} . The issue is that, due to the failure of the mASD equation to be gauge invariant, the connection u^*A will no longer be mASD. Nevertheless, we will show in this section that, when A is regular, the connection u^*A is close to a unique mASD connection that lies in the A_{ref} -Coulomb slice. This is made precise in Theorem 4.5, which extends the discussion to handle connections A that are not regular by means of an obstruction map. To accomplish this, we first prove a general gauge fixing result that is tailored to our setting; this is stated in Proposition 4.3.

4A. Gauge fixing. We begin by refining our choices of δ and the cut-off function β used to create \mathcal{H}_{out} . For the former, we assume $\delta^2/4$ is not in the spectrum of the Laplacian Δ on real-valued functions. It then follows from Sobolev embedding that, for each $1 < q < 4$, there is a constant \mathfrak{c}_q so that

$$(4.1) \quad \|f\|_{L_\delta^q(X)} + \|f\|_{L_\delta^{q^*}(X)} \leq \mathfrak{c}_q \|df\|_{L_\delta^q(X)}$$

for all compactly supported real-valued smooth functions f , where $q^* = 4q/(4 - q)$ is the Sobolev conjugate.

As for the cutoff function $\beta: \mathcal{H} \rightarrow [0, 1]$, we assume this is chosen so that it has small support in the sense that

$$(4.2) \quad \sup_{h, h_0 \in \text{supp}(\beta)} \|h - h_0\|_{C^0(N)} + \|\Theta(h) - \Theta(h_0)\|_{C^0(N)} < \frac{1}{2c_2c_g}$$

where c_g is the constant from (3.23) and c_2 is the constant from (4.1) with $q = 2$.

The main gauge fixing result we will need is as follows.

Proposition 4.3. *Fix $2 < p < 4$, set $p^* = 4p/(4 - p)$, and assume δ, β are as above. There are constants $C, \epsilon > 0$ so that if $A = \iota(h, V)$ and $A_{ref} = \iota(h_{ref}, V_{ref})$ are in $\mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ and satisfy*

$$\|V - V_{ref}\|_{L_\delta^{p^*}(X)} + \|d_{A_{ref}}^{*,\delta}(V - V_{ref})\|_{L_\delta^p(X)} < \epsilon$$

then there is a unique $\mu = \mu(A, A_{ref}) \in L_{2,\delta}^p(\Omega^0(X))$ so that

$$\exp(\mu)^* A \in \mathcal{SL}(A_{ref}), \quad \text{and} \quad \|d_{A_{ref}}^{*,\delta} d_A \mu\|_{L_\delta^p(X)} \leq C \|d_{A_{ref}}^{*,\delta}(V - V_{ref})\|_{L_\delta^p(X)}.$$

Moreover, this 0-form $\mu(A, A_{ref})$ depends C^m -smoothly on the pair (A, A_{ref}) .

Proof. We will show below that

$$\|\mu\|_{\mathcal{X}} := \|d_{A_{ref}}^{*,\delta} d_A \mu\|_{L_\delta^p(X)}$$

defines a norm on the space $\Omega^0(X)$ of smooth rapidly decaying adjoint bundle-valued 0-forms. Assuming this for now, we denote by \mathcal{X} the completion of $\Omega^0(X)$ relative to $\|\cdot\|_{\mathcal{X}}$. Let \mathcal{Y} be the completion of $\Omega^0(X)$, but relative to the norm $\|\cdot\|_{\mathcal{Y}} := \|\cdot\|_{L_\delta^p(X)}$. Since $p > 2$, the map

$$\begin{aligned} \mathcal{F}: \mathcal{A}^{1,p}(\mathcal{T}_\Gamma) \times \mathcal{A}^{1,p}(\mathcal{T}_\Gamma) \times \mathcal{X} &\longrightarrow \mathcal{Y} \\ (A, A_{ref}, \mu) &\longmapsto d_{A_{ref}}^{*,\delta}(u^* A - i(p(u^* A)) - V_{ref}) \end{aligned}$$

is C^m -smooth, where we have set $u = \exp(\mu) \in \mathcal{G}_\delta^{2,p}$. Note that, relative to the product structure given by ι via (2.5), the quantity $u^* A - i(p(u^* A))$ is the $L_{1,\delta}^p(\Omega^1(X))$ -component (i.e., non-center manifold-component) of $u^* A$, and so $u^* A \in \mathcal{SL}(A_{ref})$ for $u = \exp(\mu)$ if and only if $\mathcal{F}(A, A_{ref}, \mu) = 0$. It therefore suffices to solve $\mathcal{F}(A, A_{ref}, \mu) = 0$ for μ . For this, we have that $\mu = 0$ is an approximate solution since

$$\mathcal{F}(A, A_{ref}, 0) = d_{A_{ref}}^{*,\delta}(V - V_{ref})$$

which we have assumed is bounded by ϵ . The linearization in the third component of \mathcal{F} at $(A, A_{ref}, 0)$ is the operator

$$\mu \longmapsto d_{A_{ref}}^{*,\delta} d_A \mu.$$

This has operator norm 1 relative to the norms on \mathcal{X} and \mathcal{Y} . In particular, it is invertible and so the proposition follows from the inverse function theorem (e.g., precompose \mathcal{F} in the third component with the inverse of $d_{A_{ref}}^{*,\delta} d_A$ and then use Lemma 3.7).

All that remains is to show that $\|\cdot\|_{\mathcal{X}}$ defines a norm; it suffices to show that the operator

$$d_{A_{ref}}^{*,\delta} d_A: L_{2,\delta}^p \longrightarrow L_{\delta}^p$$

is injective. For this, suppose μ lies in its kernel and let $(\cdot, \cdot)_{\delta}$ be the δ -dependent L^2 -inner product. Note that $\mu \in L_{1,\delta}^2$ by Sobolev's embedding theorem $L_2^p \hookrightarrow L_1^2$ on cylindrical end 4-manifolds [4, Prop. 3.20]. This justifies the following computation:

$$0 = (d_{A_{ref}}^{*,\delta} d_A \mu, \mu)_{\delta} = (d_A \mu, d_{A_{ref}} \mu)_{\delta} = \|d_{A_{ref}} \mu\|_{L_{\delta}^2}^2 + \left([(A - A_{ref}), \mu], d_{A_{ref}} \mu \right)_{\delta}.$$

Hence

$$\|d_{A_{ref}} \mu\|_{L_{\delta}^2} \leq \left\| [(A - A_{ref}), \mu] \right\|_{L_{\delta}^2}.$$

The definition of ι gives $A - A_{ref} = \beta''(h - h_{ref} + (\Theta(h) - \Theta(h_{ref}))dt) + V - V_{ref}$. Then Hölder's inequality and (4.2) allow us to continue the above inequality to get

$$\begin{aligned} \|d_{A_{ref}} \mu\|_{L_{\delta}^2} &\leq c_{\mathfrak{g}} (\|V - V_{ref}\|_{L^4} \|\mu\|_{L_{\delta}^4} + \|h - h_{ref} + (\Theta(h) - \Theta(h_{ref}))dt\|_{C^0} \|\mu\|_{L_{\delta}^2}) \\ &\leq c_{\mathfrak{g}} \|e^{-\delta t/2}\|_{L^r} \|V - V_{ref}\|_{L_{\delta}^{p^*}} \|\mu\|_{L_{\delta}^4} + \frac{1}{2c_2} \|\mu\|_{L_{\delta}^2} \end{aligned}$$

where r is defined by $r^{-1} + (p^*)^{-1} = 4^{-1}$. Using (4.1) with $f = |\mu|$, and then Kato's inequality $|d|\mu|| \leq |d_{A_{ref}} \mu|$ (which holds for arbitrary metric connections), we can use the above to get

$$\begin{aligned} \|\mu\|_{L_{\delta}^2} + \|\mu\|_{L_{\delta}^4} &\leq c_2 \|d|\mu|\|_{L_{\delta}^2} \\ &\leq c_2 \|d_{A_{ref}} \mu\|_{L_{\delta}^2} \\ &\leq c_2 c_{\mathfrak{g}} \|e^{-\delta t/2}\|_{L^r} \|V - V_{ref}\|_{L_{\delta}^{p^*}} \|\mu\|_{L_{\delta}^4} + \frac{1}{2} \|\mu\|_{L_{\delta}^2}. \end{aligned}$$

When $\epsilon < 1/(2c_2 c_{\mathfrak{g}} \|e^{-\delta t/2}\|_{L^r})$, this implies that $\mu = 0$. \square

Remark 4.4. The operator $d_{A_{ref}}^{*,\delta} d_A: L_{2,\delta}^p \rightarrow L_{\delta}^p$ is Fredholm, and we have just seen that it has trivial kernel under the hypotheses of the proposition. It then follows from the embedding $L_{2,\delta}^p \subseteq L_{\delta}^p \cap C^0$ that there is a constant C so that

$$\|\mu\|_{L_{\delta}^p(X)} + \|\mu\|_{C^0(X)} \leq C \|d_{A_{ref}}^{*,\delta} d_A \mu\|_{L_{\delta}^p(X)}$$

for all $\mu \in L_{2,\delta}^p(\Omega^0(X))$. It follows from arguments similar to those just used that this constant can be chosen to be independent of A and A_{ref} provided these connections satisfy the hypotheses of Proposition 4.3.

4B. Recovering the mASD condition within a slice. Throughout this section, we assume $2 < p < 4$, and δ, β are chosen as in Section 4A.

As suggested in the introduction to this section, we will use Proposition 4.3 to put mASD connections into a fixed nearby slice, but this process will generally not preserve the mASD condition. The following theorem is our main readjustment tool that will recover the mASD condition, while simultaneously preserving the slice condition. To

state it, use the L_δ^2 -inner product to identify the cokernel $H_{A,\delta}^+ = \text{coker}(Ds|_{\mathcal{SL}(A)})_A$ with the subset of A -harmonic self-dual forms in $L_\delta^p(\Omega^+(X))$. We denote by

$$\sigma_A: H_{A,\delta}^+ \longrightarrow L_\delta^p(\Omega^+(X)) \quad \pi_A: L_\delta^p(\Omega^+(X)) \longrightarrow H_{A,\delta}^+$$

the inclusion and L_δ^2 -orthogonal projection, respectively. (These maps will play a role analogous to the one played by σ and π in Section 3.) It follows that $(Ds|_{\mathcal{SL}(A)})_A \oplus \sigma_A$ maps surjectively onto $L_\delta^p(\Omega^+(X))$.

Theorem 4.5. *Fix $A_{ref} = \iota(h_{ref}, V_{ref}) \in \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$. Then there are constants $C, \epsilon > 0$ so that the following holds for all mASD connections $A = \iota(h, V)$ satisfying*

$$(4.6) \quad \|h - h_{ref}\|_{L^2_2(N)} + \|V - V_{ref}\|_{L^{p^*}_\delta(X)} + \|d_{A_{ref}}^{*,\delta}(V - V_{ref})\|_{L^p_\delta(X)} < \epsilon.$$

- (a) *There is a C^m -map $K_A: L_\delta^p(\Omega^+(X)) \rightarrow \mathcal{SL}(A_{ref})$ that restricts to an embedding on a neighborhood U of 0.*
- (b) *There is a unique 2-form $\zeta(A) \in U \subseteq L_\delta^p(\Omega^+(X))$ so that*

$$\|\zeta(A)\|_{L^p_\delta(X)} \leq C \|d_{A_{ref}}^{*,\delta}(V - V_{ref})\|_{L^p_\delta(X)}$$

and so that the connection $\mathcal{K}(A) := K_A(\zeta(A))$ satisfies

$$s(\mathcal{K}(A)) = -\sigma_A \pi_A \zeta(A).$$

In particular, the connection $\mathcal{K}(A)$ is close to A in the sense that there is a constant C' so that

$$\|\iota^{-1}(\mathcal{K}(A)) - \iota^{-1}(A)\|_{L^2_2(N) \times L^{p^*}_\delta(X)} \leq C' \|d_{A_{ref}}^{*,\delta}(V - V_{ref})\|_{L^p_\delta(X)}.$$

If either A or A_{ref} is regular, then all three connections A , A_{ref} and $\mathcal{K}(A)$ are regular and $\mathcal{K}(A)$ is A_{ref} -regular. In this case, the connection $\mathcal{K}(A)$ is mASD and the maps $(A, \zeta) \mapsto K_A(\zeta)$ and $A \mapsto \zeta(A)$ are both C^m -smooth, relative to the specified topologies. If A is irreducible, then so is $\mathcal{K}(A)$. The constants ϵ, C, C' can be chosen to vary continuously in A_{ref} .

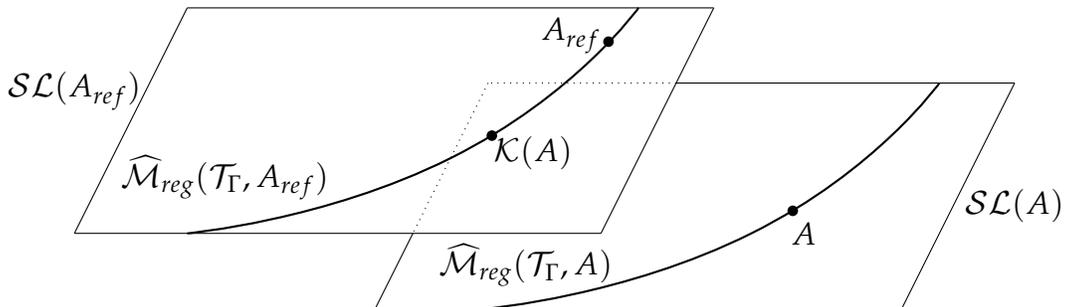


FIGURE 2. Pictured above is the special case of Theorem 4.5 where A is regular. The curved lines represent the spaces of regular mASD connections in the slices $\mathcal{SL}(A_{ref})$ and $\mathcal{SL}(A)$, respectively.

Proof. Take $\epsilon > 0$ to be no larger than the epsilon from the statement of Proposition 4.3. Then it follows from that proposition and Remark 4.4 that, given $A = \iota(h, V)$ with

$$s(A) = 0, \quad \text{and} \quad \|V - V_{ref}\|_{L_\delta^{p^*}(X)} + \|d_{A_{ref}}^{*,\delta}(V - V_{ref})\|_{L_\delta^p(X)} < \epsilon$$

there is a unique $\mu \in L_{2,\delta}^p(\Omega^0)$ so that $\exp(\mu)^* A \in \mathcal{SL}(A_{ref})$ and

$$\|\mu\|_{L_\delta^p} + \|\mu\|_{C^0} \leq C_1 \|d_{A_{ref}}^{*,\delta}(V - V_{ref})\|_{L_\delta^p}.$$

Set $u = \exp(\mu)$ and write $u^* A = \iota(h_A, V_A)$ for $h_A \in \mathcal{H}_{out}$ and $V_A \in L_{1,\delta}^p(\Omega^1)$. Let $\exp_{h_A} : B_\epsilon(0) \subseteq T_{h_A} \mathcal{H} \rightarrow \mathcal{H}$ be the exponential map for the center manifold based at h_A , and extend this to a map

$$\begin{aligned} \exp_{(h_A, V_A)} : B_\epsilon(0) \times L_{1,\delta}^p(\ker(d_{A_{ref}}^{*,\delta})) &\longrightarrow \mathcal{H} \times L_{1,\delta}^p(\ker(d_{A_{ref}}^{*,\delta})) \\ (\eta, V) &\longmapsto (\exp_{h_A}(\eta), V_A + V) \end{aligned}$$

which is a C^m -diffeomorphism in a neighborhood of $(0,0)$. Using this, define

$$\tilde{s} : T_{h_A} \mathcal{H} \times L_{1,\delta}^p(\ker(d_{A_{ref}}^{*,\delta})) \longrightarrow L_\delta^p(\Omega^+), \quad (\eta, V) \longmapsto s(\iota(\exp_{(h_A, V_A)}(\eta, V))).$$

By definition of σ_A , the operator $(Ds|_{\mathcal{SL}(A)})_A \oplus \sigma_A$ is surjective. The operators $(Ds)_A$ and $(Ds)_{u^*A}$ are approximately equal when u is C^0 -close to the identity (i.e., when $\|\mu\|_{C^0}$ is small). Thus, when ϵ is sufficiently small, the operator $(Ds|_{\mathcal{SL}(A_{ref})})_{u^*A} \oplus \sigma_A$ is also surjective, as is $(D\tilde{s})_{(0,0)} \oplus \sigma_A$. Then we can choose a right inverse to $(D\tilde{s})_{(0,0)} \oplus \sigma_A$ of the form $P \oplus \pi_A$, where π_A is the projection to $H_{A,\delta}^+$. For $\zeta \in L_\delta^p(\Omega^+(X))$, define

$$K_A(\zeta) := \iota(\exp_{(h_A, V_A)}(P\zeta)).$$

This proves (a) in the statement of the theorem, by taking $U \subseteq L_\delta^p(\Omega^+(X))$ to be small enough so that $P(U) \subseteq B_\epsilon(0) \times \ker(d_{A_{ref}}^{*,\delta})$.

To prove (b), we use the same implicit function theorem argument as in Theorem 3.3. Namely, set

$$S(\zeta) := \tilde{s}(P\zeta) - (D\tilde{s})_{(0,0)}P\zeta - \tilde{s}(0,0).$$

The argument of Claim 2 in the proof of Theorem 3.3 carries over to show that S satisfies the quadratic estimate of Lemma 3.7. We will show in a moment that there is a uniform constant C_2 so that

$$(4.7) \quad \|\tilde{s}(0,0)\|_{L_\delta^p} \leq C_2 \|d_{A_{ref}}^{*,\delta}(V - V_{ref})\|_{L_\delta^p}.$$

From this and Lemma 3.7 it follows that, by assuming ϵ is sufficiently small, there is a unique $\zeta(A)$ so that

$$\zeta(A) + S(\zeta(A)) = -\tilde{s}(0,0).$$

As we argued in the proof of Theorem 3.3, this $\zeta(A)$ satisfies the assertions of (b). The regularity and irreducibility assertions also follow as in Theorem 3.3.

It therefore suffices to verify (4.7). By definition of \tilde{s} , we have $\tilde{s}(0,0) = s(u^*A)$. Recall that the projection p_T to the center manifold is gauge invariant, so $p_T(u^*A) = p_T(A) = h$. This implies

$$\begin{aligned} \tilde{s}(0,0) = s(u^*A) &= F_{u^*A}^+ - \beta F_{i(h)}^+ \\ &= \text{Ad}(u^{-1})F_A^+ - \beta F_{i(h)}^+ \\ &= \text{Ad}(u^{-1})\beta F_{i(h)}^+ - \beta F_{i(h)}^+ \\ &= \beta(\text{Ad}(u^{-1}) - I)F_{i(h)}^+ \end{aligned}$$

where, in the penultimate equality, we used the assumption that $s(A) = F_A^+ - \beta F_{i(h)}^+ = 0$ vanishes. By shrinking ϵ further still, we may suppose $\|\mu\|_{C^0} \leq 1$. Then the Taylor expansion for the exponential $u = \exp(\mu)$ gives

$$\begin{aligned} \|\tilde{s}(0,0)\|_{L_\delta^p} &\leq C_3 \|F_{i(h)}^+\|_{C^0} \|\mu\|_{L_\delta^p} \\ &\leq C_1 C_3 \|F_{i(h)}^+\|_{C^0} \|d_{A_{ref}}^{*,\delta}(V - V_{ref})\|_{L_\delta^p}. \end{aligned}$$

The quantity $\|F_{i(h)}^+\|_{C^0}$ is bounded independent of h since \mathcal{H}_{out} is compact. \square

5. GLUING REGULAR FAMILIES

Throughout this section we work with the space $\mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ for fixed $2 < p < 4$. We assume that δ and the cutoff function β are chosen as in Section 4A. We also assume $\delta/2 < \mu_\Gamma^-$, so the index formula discussed in Section 2B.5 applies.

We freely refer to the notation of Section 3. For $k = 1, 2$, fix a precompact open set

$$G_k \subseteq \widehat{\mathcal{M}}_{reg}(\mathcal{T}_{k,\Gamma_k}, A_{ref,k})$$

of $A_{ref,k}$ -regular mASD connections on X_k relative to some reference connection $A_{ref,k}$. Since the G_k are precompact, we can fix $L, \lambda > 0$ so that conclusions of Theorem 3.3 hold for all $(A_1, A_2) \in G_1 \times G_2$. (In our applications of the material of this section, the values of L and λ will be fixed, so we do not keep track of them in the notation.) Then Theorem 3.3 produces a regular mASD connection $\mathcal{J}(A_1, A_2) \in \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$.

Ideally, we would want to view the mapping $(A_1, A_2) \mapsto \mathcal{J}(A_1, A_2)$ as a function from $G_1 \times G_2$ into a fixed mASD space. However, since the Coulomb slice to which $\mathcal{J}(A_1, A_2)$ belongs depends on (A_1, A_2) (cf. Remark 3.19 (b)), it is more natural to realize this mapping as a section of a bundle. Towards this end, set

$$\mathcal{E} := \left\{ (A_1, A_2, A) \mid A_k \in G_k, A \in \widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, \mathcal{J}(A_1, A_2)) \right\}.$$

Let $\Pi: \mathcal{E} \rightarrow G_1 \times G_2$ be the projection to the first two factors. Then the map

$$\Psi(A_1, A_2) := (A_1, A_2, \mathcal{J}(A_1, A_2))$$

is clearly a section of the map Π .

Theorem 5.1.

(a) For all sufficiently small $\lambda > 0$, there is a neighborhood $\mathcal{U} \subseteq \mathcal{E}$ of the image of Ψ so that the restriction $\Pi|_{\mathcal{U}}: \mathcal{U} \rightarrow G_1 \times G_2$ is a locally trivial fiber bundle. The fibers of $\Pi|_{\mathcal{U}}$ can be identified with open subsets of $\widehat{\mathcal{M}}_{reg}(\mathcal{T}_{\Gamma}, A_{ref})$ for some A_{ref} .

More specifically, every $(A_{10}, A_{20}) \in G_1 \times G_2$ is contained in an open neighborhood $\mathcal{V} \subseteq G_1 \times G_2$ so that the following holds. Let $A_{ref} = A'(A_{10}, A_{20})$ be the preglued connection, and consider the map \mathcal{K} from Theorem 4.5, defined relative to this reference connection A_{ref} . Then the map

$$(5.2) \quad \Pi^{-1}(\mathcal{V}) \cap \mathcal{U} \longrightarrow \mathcal{V} \times \widehat{\mathcal{M}}_{reg}(\mathcal{T}_{\Gamma}, A_{ref}) \quad (A_1, A_2, A) \longmapsto (A_1, A_2, \mathcal{K}(A))$$

is a well-defined C^m -diffeomorphism onto an open subset of the codomain, and this map produces a local trivialization of $\Pi|_{\mathcal{U}}$ over \mathcal{V} .

(b) The map

$$\Phi := \mathcal{K} \circ \mathcal{J} : \mathcal{V} \longmapsto \widehat{\mathcal{M}}_{reg}(\mathcal{T}_{\Gamma}, A_{ref})$$

is a C^m -embedding. If A_1 or A_2 is irreducible, then the connection $\Phi(A_1, A_2)$ is also irreducible.

Part (b) can be restated by saying that, relative to the local trivialization of (a), locally the section Ψ becomes a C^m -embedding onto an embedded C^m -submanifold of the fiber. This is an mASD version of the familiar result for ASD connections that gluing produces a parametrized family of connections in the ASD moduli space for a connected sum. See Figure 3 for an illustration of the fiber bundle in (a), and Figure 4 for an illustration of the specified trivialization, as well as the map Φ .

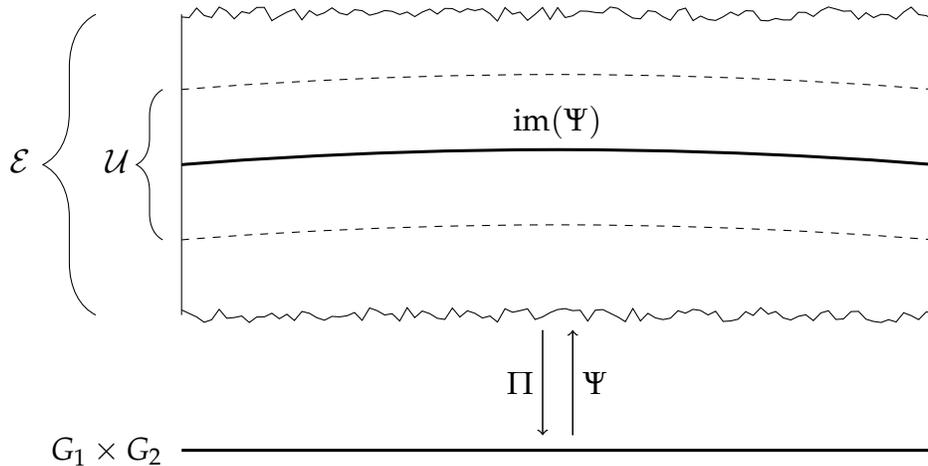


FIGURE 3. The above picture illustrates the fiber bundle $\Pi|_{\mathcal{U}}: \mathcal{U} \rightarrow G_1 \times G_2$ obtained by restricting the projection $\Pi: \mathcal{E} \rightarrow G_1 \times G_2$ to the open submanifold $\mathcal{U} \subseteq \mathcal{E}$. The fibers are C^m -diffeomorphic to open subsets of $\widehat{\mathcal{M}}_{reg}(\mathcal{T}_{\Gamma}, A_{ref})$.

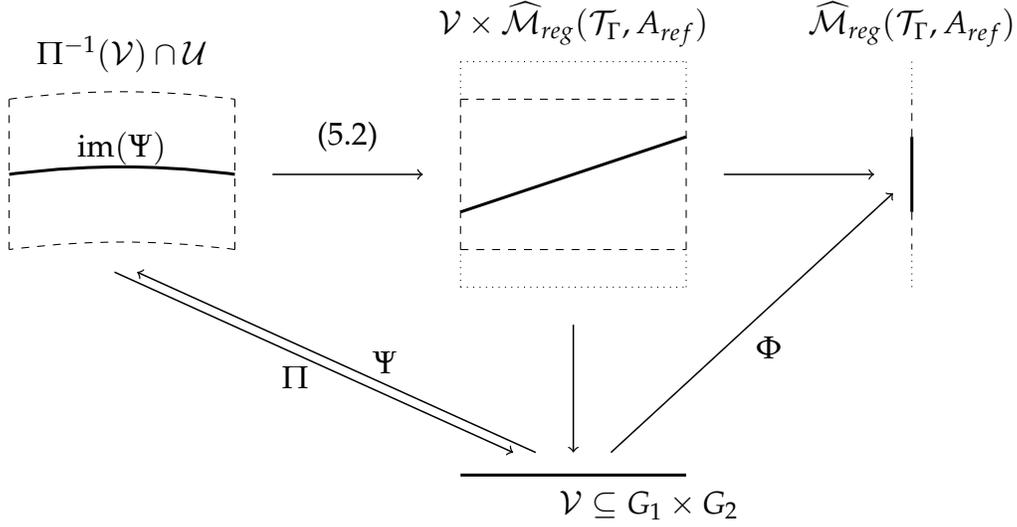


FIGURE 4. The horizontal arrow on the left is the \mathcal{C}^m -diffeomorphism (5.2) trivializing $\Pi|_{\mathcal{U}}$ over \mathcal{V} . This is a \mathcal{C}^m -diffeomorphism onto the region in $\mathcal{V} \times \widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A_{ref})$ represented by the dashed lines. It takes $\text{im}(\Psi)$, represented by the solid arc, \mathcal{C}^m -diffeomorphically onto the region in $\mathcal{V} \times \widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A_{ref})$ represented by the solid diagonal line. The horizontal arrow on the right is the projection of $\mathcal{V} \times \widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A_{ref})$ to the second factor. This takes the region represented by the dotted (resp., dashed) lines in the domain onto the region represented by the dotted (resp., dashed) line in the codomain. It restricts to a \mathcal{C}^m -diffeomorphism from the region represented by the solid diagonal line onto the region in the codomain represented by the solid line. The vertical arrow is the projection of $\mathcal{V} \times \widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A_{ref})$ to the first factor. This restricts to a \mathcal{C}^m -diffeomorphism from the region represented by the solid diagonal line onto \mathcal{V} . The map Φ is a \mathcal{C}^m -embedding and its image is represented by the solid line in the picture on the right (despite what the picture suggests, the dimension of the image of Φ may not equal the dimension of $\widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A_{ref})$; a higher dimensional illustration would be needed to accurately represent this phenomenon).

Remark 5.3. (a) We will also be interested in the case where G_2 consists of a single point (and so not necessarily an *open* set in the mASD space). In this case, Theorem 5.1 continues to hold verbatim; there is no significant change in the proofs to account for this extension.

(b) Recall the fiber isomorphism ρ from the beginning of Section 3. The usual ASD gluing results (e.g., those of [20, 19, 5]) allows ρ to vary as a “gluing parameter”. Presumably a similar construction could be carried out in the mASD setting, with

the expectation that a parametrized version of the map Φ from (a) would be a \mathcal{C}^m -diffeomorphism onto an open subset of $\widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A_{ref})$. Though we do not pursue the details of this parametrized gluing construction here, the isomorphism ρ will play an active role in our existence result of Section 6C.

(c) We have chosen to phrase Theorem 5.1 in terms of the reference connection $A_{ref} = A'(A_{10}, A_{20})$ given by the preglued connection. This is only in preparation for our applications below, and this specific choice is by no means necessary. Indeed, the proof will show the connection A_{ref} can be replaced by any connection that is sufficiently close to $\mathcal{J}(A_{10}, A_{20})$ in the sense that the coordinates of A_{ref} and $\mathcal{J}(A_{10}, A_{20})$ satisfy the estimate (5.12).

We begin by giving several technical lemmas in Section 5A, which are used to prove that Φ is an immersion. The proof of Theorem 5.1 is given in Section 5B.

5A. Immersion lemmas. Our ultimate goal is to show that the map Φ is an immersion. Recall this is made up of the maps \mathcal{J} and \mathcal{K} , and hence of the maps J, ξ, K, ζ of Theorems 3.3 and 4.5. Each of the four lemmas below establishes an estimate on the derivative of one of these latter four maps. To state the lemmas, we introduce the following seminorms on the tangent space $T_A \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$: Fix an open subset $U \subseteq X$ containing $\text{End } X$ and let $W \in T_A \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$. Using the isomorphism (2.6), we can identify W with a pair

$$(\eta, V) \in T_{p(A)} \mathcal{H}_\Gamma \times L_{1,\delta}^p(\Omega^1(X)).$$

Then set

$$\|W\|_{\mathcal{L}(U);A} := \|\eta\|_{L_2^2(N)} + \|V\|_{L_\delta^p(U)} + \|d_A^+ V\|_{L_\delta^p(U)} + \|d_A^{*,\delta} V\|_{L_\delta^p(U)}$$

where the derivatives defining the norm $\|\cdot\|_{L_2^2(N)}$ on $T_{p(A)} \mathcal{H}_\Gamma$ are defined using the connection Γ . Then $\|\cdot\|_{\mathcal{L}(U);A}$ is a continuous seminorm. These seminorms are gauge-invariant in the sense that

$$(5.4) \quad \|W\|_{\mathcal{L}(U);A} = \|\text{Ad}(u^{-1})W\|_{\mathcal{L}(U);u^*A}$$

for all gauge transformations $u \in \mathcal{G}_\delta^{2,p}(\Gamma)$; here, via a slight abuse of notation, we are writing $\text{Ad}(u^{-1})W$ for the linearization in the direction of W of the map $A \mapsto u^*A$. We use similar notation on the X_k . We note that if $U = X$, then $\|\cdot\|_{\mathcal{L}(X);A}$ is a norm that induces the topology on $T_A \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$. When the metric $g_{L,\lambda}$ is relevant, we will include it in the notation by writing $\|\cdot\|_{\mathcal{L}(X,g_{L,\lambda});A}$.

Our first lemma deals with the map $(A_1, A_2, \xi) \mapsto J_{A_1, A_2}(\xi)$. To first order, this map is the sum of the pregluing map $(A_1, A_2) \mapsto A'(A_1, A_2)$ together with a map that is bounded in ξ . We now quantify this to an extent that is sufficient for our purposes.

Lemma 5.5. *Fix connections $A_k \in G_k$ for $k = 1, 2$. There are constants $C, L, \lambda_0, \epsilon > 0$ so that the following holds for all $0 < \lambda < \lambda_0$ and all $\xi \in L_\delta^p(\Omega^+(X), g_{L,\lambda})$ with $\|\xi\|_{L_\delta^p(X, g_{L,\lambda})} < \epsilon$.*

Let $DJ_{(A_1, A_2, \xi)}(W_1, W_2, x)$ denote the linearization at (A_1, A_2, ξ) in the direction (W_1, W_2, x) of the map

$$(A_1, A_2, \xi) \longmapsto J_{A_1, A_2}(\xi)$$

from Theorem 3.3 (a). Then

$$\sum_{k=1}^2 \|W_k\|_{\mathcal{L}(X_k); A_k} \leq C \left(\|DJ_{(A_1, A_2, \xi)}(W_1, W_2, x)\|_{\mathcal{L}(X, g_{L, \lambda}), A'} + \|x\|_{L_\delta^p(X, g_{L, \lambda})} \right)$$

for all $W_k \in T_{A_k} G_k$ and all $x \in L_\delta^p(\Omega^+(X), g_{L, \lambda})$, where $A' := A'(A_1, A_2)$ is the preglued connection. The constants $C, L, \lambda_0, \epsilon$ can be chosen to depend continuously on the $A_k \in G_k$.

The next lemma shows that the map $(A_1, A_2) \mapsto \xi(A_1, A_2)$ depends minimally on the connections A_1, A_2 . In the next section, this will combine with the previous lemma to show that the map $\mathcal{J}(A_1, A_2) = J_{A_1, A_2}(\xi(A_1, A_2))$ is approximately the pregluing map $(A_1, A_2) \mapsto A'(A_1, A_2)$ for λ small; at this point it will follow that \mathcal{J} is an immersion.

Lemma 5.6. Fix $A_k \in G_k$ for $k = 1, 2$. Then there are constants $C, L, \lambda_0 > 0$ so that the following holds for all $0 < \lambda < \lambda_0$. Let $D\xi_{(A_1, A_2)}(W_1, W_2)$ denote the linearization at (A_1, A_2) in the direction (W_1, W_2) of the map

$$(A_1, A_2) \longmapsto \xi(A_1, A_2)$$

from Theorem 3.3 (c). Then this satisfies

$$\|D\xi_{(A_1, A_2)}(W_1, W_2)\|_{L_\delta^p(X, g_{L, \lambda})} \leq b^{4/p} C \sum_{k=1}^2 \|W_k\|_{\mathcal{L}(X_k); A_k}$$

for all $W_k \in T_{A_k} G_k$, where $b = 4L\lambda^{1/2}$. The constants C, L, λ_0 can be chosen to depend continuously on the $A_k \in G_k$.

These next two lemmas are analogues of the previous two, but for the map $\mathcal{K}(A) = K_A(\zeta(A))$ in place of $\mathcal{J}(A_1, A_2) = J_{A_1, A_2}(\xi(A_1, A_2))$.

Lemma 5.7. Fix a regular connection $A_{ref} \in \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$. Then there are constants $C, \epsilon' > 0$ so that the following holds for all connections $A \in \widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A_{ref})$ satisfying (4.6) with respect to any $0 < \epsilon < \epsilon'$. Let $DK_{(A, \zeta)}(W, z)$ denote the linearization at (A, ζ) in the direction (W, z) of the map

$$(A, \zeta) \longmapsto K_A(\zeta)$$

from Theorem 4.5 (a). Then this satisfies

$$\|W\|_{\mathcal{L}(X), A} \leq C \left(\|DK_{(A, \zeta)}(W, z)\|_{\mathcal{L}(X), A'} + \|z\|_{L_\delta^p(X)} \right)$$

for all $W \in T_A \widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A_{ref})$ and all $z \in L_\delta^p(\Omega^+(X))$. The constants C, ϵ can be chosen to depend continuously on A and A_{ref} .

Lemma 5.8. *Fix a regular connection $A_{ref} \in \mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$. Then there are constants $C, \epsilon' > 0$ so that the following holds for all connections $A \in \widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A_{ref})$ satisfying (4.6) with respect to any $0 < \epsilon < \epsilon'$. Let $D\zeta_A W$ denote the linearization at A in the direction W of the map*

$$A \longmapsto \zeta(A)$$

from Theorem 4.5 (b). Then this satisfies

$$\|D\zeta_A W\|_{L_\delta^p(X)} \leq C \|d_{A_{ref}}^{*,\delta}(V - V_{ref})\|_{L_\delta^p(X)} \|W\|_{\mathcal{L}(X),A}$$

for all $W \in T_A \widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A_{ref})$. The constants C, ϵ can be chosen to depend continuously on A_{ref} .

Now we give the proofs of Lemmas 5.5, 5.7, 5.6, and 5.8, in that order.

Proof of Lemma 5.5. The tangent space $T_{A_k} G_k$ is cut out by linear elliptic equations. In particular, unique continuation holds for the elements of this tangent space, and so the assignment

$$W_k \longmapsto \|W_k\|_{\mathcal{L}(X_k \setminus B_{L\lambda^{1/2}}(x_k)); A_k}$$

defines a norm on $T_{A_k} G_k$. Since $T_{A_k} G_k$ is a finite-dimensional vector space, any two norms are equivalent and so there is a constant C_1 so that

$$\|W_k\|_{\mathcal{L}(X_k); A_k} \leq C_1 \|W_k\|_{\mathcal{L}(X_k \setminus B_{L\lambda^{1/2}}(x_k)); A_k}$$

for all $W_k \in T_{A_k} G_k$. A simple contradiction argument shows that this constant can be taken to be independent of L, λ , provided $L \geq 1$ and λ is sufficiently small.

Now fix tangent vectors $W_k \in T_{A_k} G_k$ for $k = 1, 2$. Since A_k is regular, we can find a \mathcal{C}^m -smooth path $A_k(\tau)$ of A_k -regular mASD connections with $A_k(0) = A_k$ and $\frac{d}{d\tau}|_{\tau=0} A_k(\tau) = W_k$. Let $W' = \frac{d}{d\tau}|_{\tau=0} A'(A_1(\tau), A_2(\tau))$. Note that the construction of the preglued connection $A'(A_1, A_2)$ implies there is a uniform constant C_2 so that

$$\sum_k \|W_k\|_{\mathcal{L}(X_k \setminus B_{L\lambda^{1/2}}(x_k)); A_k} \leq C_2 \|W'\|_{\mathcal{L}(X); A'(A_1, A_2)}$$

provided $\lambda > 0$ is sufficiently small. Thus we have

$$(5.9) \quad \sum_k \|W_k\|_{\mathcal{L}(X_k); A_k} \leq C_1 C_2 \|W'\|_{\mathcal{L}(X); A'}$$

The next claim ties this in with the linearization of the map J at (A_1, A_2, ξ) when $\xi = 0$.

Claim 1: $DJ_{(A_1, A_2, 0)}(W_1, W_2, x) = W' + (D\iota)_{\iota^{-1}(A')} P x \quad \forall x \in L_\delta^p(\Omega^+(X))$.

Here P is the right-inverse from the proof Theorem 3.3. This depends on A_1, A_2 , so to emphasize this, we will temporarily write $P_{A_1, A_2} := P$. Consider the map

$$(5.10) \quad (A, \xi) \longmapsto \iota \circ \exp_{\iota^{-1}(A)} (P_{A_1, A_2} \xi)$$

where A ranges over all connections near $A' = A'(A_1, A_2)$ and ξ ranges over all self-dual 2-forms near 0. The linearization of (5.10) at $(A', 0)$ is the operator

$$(W, x) \longmapsto W + (D\iota)_{\iota^{-1}(A')} (P_{A_1, A_2} x).$$

Recall from the proof of Theorem 3.3 that

$$J_{A_1, A_2}(\xi) = \iota \left(\exp_{\iota^{-1}(A'(A_1, A_2))} (P_{A_1, A_2} \xi) \right).$$

That is, $(A_1, A_2, \xi) \mapsto J_{A_1, A_2}(\xi)$ is the map (5.10) precomposed with $A'(A_1, A_2)$ in the A -component. Then Claim 1 follows from the chain rule and the fact that we are differentiating at $\xi = 0$, which kills off all terms involving the A_k -derivatives of P_{A_1, A_2} .

In summary, we have

$$\begin{aligned} \sum_k \|W_k\|_{\mathcal{L}(X_k); A_k} &\leq C_1 C_2 \|W'\|_{\mathcal{L}(X); A'(A_1, A_2)} \\ &\leq C_1 C_2 \left(\|W' + (D\iota)_{\iota^{-1}(A')} Px\|_{\mathcal{L}(X); A'} + \|(D\iota)_{\iota^{-1}(A')} Px\|_{\mathcal{L}(X); A'} \right) \\ &= C_1 C_2 \left(\|DJ_{(A_1, A_2, 0)}(W_1, W_2, x)\|_{\mathcal{L}(X); A'} + \|(D\iota)_{\iota^{-1}(A')} Px\|_{\mathcal{L}(X); A'} \right). \end{aligned}$$

We will discuss each term on the right individually.

The first term on the right is almost satisfactory, except we linearized at $(A_1, A_2, 0)$ instead of (A_1, A_2, ξ) . To account for this, note that it follows from our regularity assumptions and Theorem 3.3 that J is C^m -smooth. In particular, Taylor's theorem gives

$$\begin{aligned} \|(DJ_{(A_1, A_2, \xi)} - DJ_{(A_1, A_2, 0)})(W_1, W_2, x)\|_{\mathcal{L}(X); A'} \\ \leq C_3 \|\xi\|_{L_\delta^p(X)} \left(\|x\|_{L_\delta^p(X)} + \sum_k \|W_k\|_{\mathcal{L}(X_k); A_k} \right) \end{aligned}$$

for some constant C_3 that depends continuously on the A_k and λ .

Claim 2: The constant C_3 can be taken to be independent of λ , provided λ is sufficiently small.

To see this, recall that the proof of Taylor's theorem shows that C_3 can be taken to be a constant multiple of the supremum of the operator norm of the second derivative of J at $(A_1, A_2, 0)$. By the chain rule, it therefore suffices to uniformly estimate the first two derivatives of ι , $\exp_{\iota^{-1}(A'(A_1, A_2))}$ and $P = P_{A_1, A_2}$. Obtaining such estimates for ι and the exponential map follow readily because the gluing region is in the complement of the cylindrical end (e.g., ι is affine-linear over this gluing region). That the derivatives of P are uniformly bounded is addressed in Remark 3.19 (a), above.

With this claim in hand, we have

$$\begin{aligned} \sum_k \|W_k\|_{\mathcal{L}(X_k); A_k} &\leq C_1 C_2 \left(\|DJ_{(A_1, A_2, \xi)}(W_1, W_2, x)\|_{\mathcal{L}(X); A'} + \|(D\iota)_{\iota^{-1}(A')} Px\|_{\mathcal{L}(X); A'} \right) \\ &\quad + C_1 C_2 C_3 \|\xi\|_{L_\delta^p(X)} \left(\|x\|_{L_\delta^p(X)} + \sum_k \|W_k\|_{\mathcal{L}(X_k); A_k} \right) \end{aligned}$$

When $\|\xi\|_{L_\delta^p(X)} < \epsilon := 1/2C_1 C_2 C_3$ this implies that $\sum_k \|W_k\|_{\mathcal{L}(X_k); A_k}$ is bounded by

$$2C_1 C_2 \left(\|DJ_{(A_1, A_2, \xi)}(W_1, W_2, x)\|_{\mathcal{L}(X); A'} + \|(D\iota)_{\iota^{-1}(A')} Px\|_{\mathcal{L}(X); A'} \right) + \|x\|_{L_\delta^p(X)}$$

The lemma now follows from the next claim.

Claim 3: There are constants $C_4, L, \lambda_0 > 0$ so that

$$\|(Dt)_{t^{-1}(A')}Px\|_{\mathcal{L}(X, \mathcal{G}_{L, \lambda}); A'} \leq C_4 \|x\|_{L^p_\delta(X, \mathcal{G}_{L, \lambda})}$$

for all x and all $0 < \lambda < \lambda_0$. These constants can be chosen to depend continuously on the $A_k \in \mathcal{A}^{1,p}(\mathcal{T}_{k, \Gamma_k})$.

We briefly sketch the proof, leaving the details to the reader. Use the fact that P is uniformly bounded to control the zeroth order terms appearing in the definition of $\|\cdot\|_{\mathcal{L}(X); A'}$. The term involving $d_{A'}^{*, \delta}$ vanishes because P takes values in the slice. To control the $d_{A'}^+$ term use the fact that P is a right inverse to an operator that is essentially $d_{A'}^+$ plus lower order terms. \square

Proof of Lemma 5.7. Fix A and W as in the lemma. Let $A(\tau)$ be a path in $\widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A_{ref})$ that is C^m -smooth and satisfies $A(0) = A$ and $\frac{d}{d\tau}|_{\tau=0}A(\tau) = W$. Let $\mu_\tau = \mu(A(\tau), A_{ref})$ be the 0-form from Proposition 4.3 associated to $A(\tau)$ and A_{ref} . Set

$$u_\tau := \exp(\mu_\tau), \quad A' := u_0^*A, \quad W' := \frac{d}{d\tau}\Big|_{\tau=0} u_\tau^*A(\tau).$$

By the product rule, we have

$$W' = \text{Ad}(u_0^{-1})W + d_{A'}\left(\frac{d}{d\tau}\Big|_{\tau=0}\mu_\tau\right).$$

Now the gauge invariance (5.4) and the definition of our norms give

$$\begin{aligned} \|W\|_{\mathcal{L}(X); A} &= \|\text{Ad}(u_0^{-1})W\|_{\mathcal{L}(X); A'} \\ &\leq \|W'\|_{\mathcal{L}(X); A'} + \left\|d_{A'}\left(\frac{d}{d\tau}\Big|_{\tau=0}\mu_\tau\right)\right\|_{\mathcal{L}(X); A'} \\ &= \|W'\|_{\mathcal{L}(X); A'} + \left\|d_{A'}\left(\frac{d}{d\tau}\Big|_{\tau=0}\mu_\tau\right)\right\|_{L^p_\delta(X)} \\ &\quad + \left\|F_{A'}^+\left(\frac{d}{d\tau}\Big|_{\tau=0}\mu_\tau\right)\right\|_{L^p_\delta(X)} + \left\|d_{A'}^{*, \delta}d_{A'}\left(\frac{d}{d\tau}\Big|_{\tau=0}\mu_\tau\right)\right\|_{L^p_\delta(X)}. \end{aligned}$$

Focusing on the second term on the right, we note that the operator $d_{A'}^{*, \delta}$ is injective on $\text{im}(d_{A'})$ so there is a bound of the form

$$\left\|d_{A'}\left(\frac{d}{d\tau}\Big|_{\tau=0}\mu_\tau\right)\right\|_{L^p_\delta(X)} \leq C_1 \left\|d_{A'}^{*, \delta}d_{A'}\left(\frac{d}{d\tau}\Big|_{\tau=0}\mu_\tau\right)\right\|_{L^p_\delta(X)}.$$

As for the third term on the right, the fact that A is mASD implies that F_A^+ is uniformly bounded in \mathcal{C}^0 ; the same is therefore true of $F_{A'}^+ = \text{Ad}(u_0^{-1})F_A^+$. Combining this with the fact that the operator $d_{A'}^{*, \delta}d_{A'}$ is injective on 0-forms, we obtain

$$\begin{aligned} \left\|F_{A'}^+\left(\frac{d}{d\tau}\Big|_{\tau=0}\mu_\tau\right)\right\|_{L^p_\delta(X)} &\leq C_2 \left\|\frac{d}{d\tau}\Big|_{\tau=0}\mu_\tau\right\|_{L^p_\delta(X)} \\ &\leq C_3 \left\|d_{A'}^{*, \delta}d_{A'}\left(\frac{d}{d\tau}\Big|_{\tau=0}\mu_\tau\right)\right\|_{L^p_\delta(X)}. \end{aligned}$$

In summary, we have

$$\|W\|_{\mathcal{L}(X);A} \leq \|W'\|_{\mathcal{L}(X);A'} + (1 + C_1 + C_3) \left\| d_{A'}^{*,\delta} d_{A'} \left(\frac{d}{d\tau} \Big|_{\tau=0} \mu_\tau \right) \right\|_{L_\delta^p(X)}$$

Our hypotheses imply that A' and A_{ref} differ by a term that is controlled by the \mathcal{C}^0 -norm of μ . This implies we have an estimate of the form

$$\left\| d_{A'}^{*,\delta} d_{A'} \left(\frac{d}{d\tau} \Big|_{\tau=0} \mu_\tau \right) \right\|_{L_\delta^p} \leq C_4 \left\| d_{A_{ref}}^{*,\delta} d_{A'} \left(\frac{d}{d\tau} \Big|_{\tau=0} \mu_\tau \right) \right\|_{L_\delta^p}.$$

To estimate this further, differentiate the defining identity $0 = d_{A_{ref}}^{*,\delta} (u_\tau^* A_\tau - A_{ref})$ at $\tau = 0$ to get

$$d_{A_{ref}}^{*,\delta} d_{A'} \left(\frac{d}{d\tau} \Big|_{\tau=0} \mu_\tau \right) = -d_{A_{ref}}^{*,\delta} W.$$

Thus

$$\begin{aligned} \left\| d_{A_{ref}}^{*,\delta} d_{A'} \left(\frac{d}{d\tau} \Big|_{\tau=0} \mu_\tau \right) \right\|_{L_\delta^p} &= \left\| d_{A_{ref}}^{*,\delta} W \right\|_{L_\delta^p} \\ &\leq \left\| d_A^{*,\delta} W \right\|_{L_\delta^p} + C_5 \|W\|_{L_\delta^p} \\ &= \left\| d_{A'}^{*,\delta} W' \right\|_{L_\delta^p} + C_5 \|W'\|_{L_\delta^p} \\ &\leq \max(1, C_5) \|W'\|_{\mathcal{L}(X);A'}. \end{aligned}$$

Hence

$$\|W\|_{\mathcal{L}(X);A} \leq C_6 \|W'\|_{\mathcal{L}(X);A'}$$

where $C_6 = 1 + (1 + C_1 + C_3)C_4 \max(1, C_5)$. To finish the proof of the lemma, argue exactly as we did in the proof of Lemma 5.5, starting after the estimate (5.9). \square

Proof of Lemma 5.6. Let $A_k(\tau)$ be a \mathcal{C}^m -smooth path in G_k satisfying $A_k(0) = A_k$ and $\frac{d}{d\tau} \Big|_{\tau=0} A_k(\tau) = W_k$. Set $\xi_\tau := \xi(A_1(\tau), A_2(\tau))$. Note that the τ -derivative

$$\frac{d}{d\tau} \Big|_{\tau=0} \xi_\tau = D\xi_{(A_1, A_2)}(W_1, W_2)$$

is the term that we are looking to bound.

The regularity hypotheses and Theorem 3.3 (c) imply that ξ_τ satisfies $0 = s(J_\tau(\xi_\tau))$ for all τ , where $J_\tau := J_{A_1(\tau), A_2(\tau)}$ is the map from Theorem 3.3 (a). Continuing to use a subscript τ for any term defined in terms of the $A_k(\tau)$ (and hence dependent on τ), we recall the definition of \tilde{s} from the proof of Theorem 3.3; in particular, this satisfies $s(J_\tau(\cdot)) = \tilde{s}(P_\tau(\cdot))$. The Taylor expansion of \tilde{s} therefore gives

$$0 = s(J_\tau(\xi_\tau)) = \tilde{s}(P_\tau(\xi_\tau)) = \tilde{s}_\tau(0, 0) + \xi_\tau + S_\tau(\xi_\tau).$$

Differentiate the right-hand side at $\tau = 0$ and rearrange to get

$$(5.11) \quad \frac{d}{d\tau} \Big|_{\tau=0} \xi_\tau = -\frac{d}{d\tau} \Big|_{\tau=0} \tilde{s}_\tau(0, 0) - \frac{d}{d\tau} \Big|_{\tau=0} S_\tau(\xi_0) - (DS_0)_{\xi_0} \left(\frac{d}{d\tau} \Big|_{\tau=0} \xi_\tau \right)$$

where $(DS_0)_{\xi_0}$ is the linearization at ξ_0 of S_0 . We will return to this after we estimate each term on the right individually.

For the first term on the right of (5.11), note that $\tilde{s}_\tau(0,0) = s(A'_\tau)$ depends on τ only through the preglued connection $A'_\tau := A'(A_1(\tau), A_2(\tau))$. Moreover, the proof of (3.1) shows that $s(A'_\tau)$ is equal to a product of a cutoff function supported in the gluing region, times the connection form for A'_τ in this region. In particular, differentiating this in τ , the same argument used for (3.1) allows us to conclude a uniform bound of the form

$$\begin{aligned} \left\| \frac{d}{d\tau} \Big|_{\tau=0} \tilde{s}_\tau(0,0) \right\|_{L^p_\delta(X, g_{L,\lambda})} &\leq C_1 b^{4/p} \left\| \frac{d}{d\tau} \Big|_{\tau=0} A'_\tau \right\|_{\mathcal{L}(X, g_{L,\lambda}), A'_0} \\ &\leq C_2 b^{4/p} \sum_k \left\| W_k \right\|_{\mathcal{L}(X); A_k} \end{aligned}$$

where the second inequality follows by differentiating the defining formula for the preglued connection $A'(A_1, A_2)$. This is the desired bound on the first term.

The second term on the right of (5.11) is similar, albeit a little more involved. The point here is that the quadratic estimates on S_τ give a uniform bound of the form

$$\left\| \frac{d}{d\tau} \Big|_{\tau=0} S_\tau(\tilde{\zeta}_0) \right\|_{L^p_\delta(X, g_{L,\lambda})} \leq C_3 \|\tilde{\zeta}_0\|_{L^p_\delta(X, g_{L,\lambda})} \sum_k \left\| W_k \right\|_{T_{A_k} \mathcal{A}}.$$

Theorem 3.3 (c) gives $\|\tilde{\zeta}_0\|_{L^p_\delta(X, g_{L,\lambda})} \leq C_4 b^{4/p}$, so the desired estimate for this term follows.

Turn now to the last term on the right of (5.11). By the estimate (3.8), the linearization $(DS_0)_{\tilde{\zeta}_0}$ satisfies

$$\|(DS_0)_{\tilde{\zeta}_0} \tilde{\zeta}'\|_{L^p_\delta(X, g_{L,\lambda})} \leq 2\kappa \|\tilde{\zeta}_0\|_{L^p_\delta(X, g_{L,\lambda})} \|\tilde{\zeta}'\|_{L^p_\delta(X, g_{L,\lambda})}$$

for all $\tilde{\zeta}'$. Since $\|\tilde{\zeta}_0\|_{L^p_\delta(X, g_{L,\lambda})} \leq C_4 b^{4/p}$, we may assume that $\|\tilde{\zeta}_0\|_{L^p_\delta(X, g_{L,\lambda})} < 1/4\kappa$, which gives

$$\|(DS_0)_{\tilde{\zeta}_0} \tilde{\zeta}'\|_{L^p_\delta(X, g_{L,\lambda})} \leq \frac{1}{2} \|\tilde{\zeta}'\|_{L^p_\delta(X, g_{L,\lambda})}.$$

To see that the above estimates imply the lemma, take the norm of each side of (5.11) and use the estimates just established to obtain

$$\left\| \frac{d}{d\tau} \Big|_{\tau=0} \tilde{\zeta}_\tau \right\|_{L^p_\delta(X, g_{L,\lambda})} \leq (C_2 + C_3 C_4) b^{4/p} \sum_k \left\| W_k \right\|_{T_{A_k} \mathcal{A}} + \frac{1}{2} \left\| \frac{d}{d\tau} \Big|_{\tau=0} \tilde{\zeta}_\tau \right\|_{L^p_\delta(X, g_{L,\lambda})}.$$

The corollary follows by subtracting the last term from both sides, and using the identity $D\tilde{\zeta}_{(A_1, A_2)}(W_1, W_2) = \frac{d}{d\tau} \Big|_{\tau=0} \tilde{\zeta}_\tau$. \square

Proof of Lemma 5.8. This follows from the same type of argument given for Lemma 5.6. \square

5B. Proof of Theorem 5.1. Let $\epsilon > 0$ be small enough so that Theorem 4.5 holds with this value of ϵ . Define \mathcal{U} to be the set of triples (A_1, A_2, A) with $A_k \in G_k$ and $A \in \widehat{\mathcal{M}}(\mathcal{T}_\Gamma, \mathcal{J}(A_1, A_2))$, and so that

$$(5.12) \quad \|h - h_0\|_{L^2_\delta(N)} + \|V - V_0\|_{L^p_\delta(X)^*} + \|d_{A_0}^{*,\delta}(V - V_0)\|_{L^p_\delta(X)} < \epsilon/3$$

where $A = \iota(h, V)$ and $A_0 = \iota(h_0, V_0) := \mathcal{J}(A_1, A_2)$. Since all elements of the G_k are regular, it follows from Theorem 3.3 that $\mathcal{J}(A_1, A_2)$ is regular, so any connection A satisfying (5.12) is automatically A_0 -regular. Thus $\mathcal{U} \subseteq \mathcal{E}$.

To show that $\Pi|_{\mathcal{U}}: \mathcal{U} \rightarrow G_1 \times G_2$ is locally trivial, fix $(A_{10}, A_{20}) \in G_1 \times G_2$, and set $A_{ref} := A'(A_{10}, A_{20})$. By (3.14) and the fact that $\mathcal{J}(A_{10}, A_{20})$ takes values in the A_{ref} -slice, by choosing λ sufficiently small, it follows that the coordinates of A_{ref} and $\mathcal{J}(A_{10}, A_{20})$ satisfy the estimate (5.12); in fact, this estimate is uniform in λ , in the sense that it holds for all sufficiently small λ . Fix any such λ ; we will refine this choice in the next paragraph. Take $\mathcal{V} \subseteq G_1 \times G_2$ to be a neighborhood of (A_{10}, A_{20}) that is small enough so that if $(A_1, A_2) \in \mathcal{V}$, then the components of A_{ref} and $A'(A_1, A_2)$ satisfy (5.12). Though we do not use this observation presently, we note that the set \mathcal{V} can also be chosen to be uniform in λ , provided $\lambda > 0$ is sufficiently small; this is due to the scaling properties the L^{p^*} -norm of 1-forms [5, p.293]. Use two applications of the triangle inequality to conclude that, for any triple $(A_1, A_2, A) \in \Pi^{-1}(\mathcal{V}) \cap \mathcal{U}$, the pair A, A_{ref} satisfies the hypotheses of Theorem 4.5. Thus the map (5.2) is well-defined and indeed provides a local trivialization of $\Pi|_{\mathcal{U}}$.

To finish the proof, it suffices to show that the map $\Phi = \mathcal{K} \circ \mathcal{J}: \mathcal{V} \rightarrow \widehat{\mathcal{M}}_{reg}(\mathcal{T}_\Gamma, A_{ref})$ is a \mathcal{C}^m -immersion; it can then be made into an embedding by further shrinking \mathcal{V} , if necessary. Since $\Phi = \mathcal{K} \circ \mathcal{J}$, it suffices to show that \mathcal{J} and \mathcal{K} are immersions for all $\lambda > 0$ sufficiently small. For \mathcal{J} , note that the linearization at (A_1, A_2) is the map

$$(W_1, W_2) \longmapsto DJ_{(A_1, A_2, \xi(A_1, A_2))}(W_1, W_2, D\xi_{(A_1, A_2)}(W_1, W_2)).$$

Suppose this vanishes at some (W_1, W_2) . Then by Lemmas 5.5 and 5.6, we would have

$$\sum_k \|W_k\|_{\mathcal{L}(X_k); A_k} \leq Cb^{4/p} \sum_k \|W_k\|_{\mathcal{L}(X_k); A_k}$$

By taking $\lambda > 0$ sufficiently small, we may assume $Cb^{4/p} < 1$ and so $W_k = 0$. Thus \mathcal{J} is an immersion. A similar argument, but using Lemmas 5.7 and 5.8, shows that \mathcal{K} is an immersion for small λ .

The irreducibility claims follow from the analogous claims appearing in Theorems 3.3 and 4.5. \square

6. EXISTENCE RESULTS

Let X be an oriented cylindrical end 4-manifold with $b^+(X) = 0$ or 1. In this section, we will show how to use the above framework to prove the existence of families of mASD connections on X ; the cases $b^+(X) = 0$ and $b^+(X) = 1$ are treated in Sections 6B and 6C, respectively. The ASD existence result Theorem A is proved in Section 6D.

Part of our existence results state that the connections we construct are topologically non-trivial in a certain sense. In the case of closed 4-manifolds, this non-triviality is captured by the non-vanishing of a characteristic class of the bundle supporting the connections. In the present cylindrical end setting, we will use a certain relative characteristic class to measure this non-triviality. The details of this are carried out in Section 6A.

6A. Relative characteristic classes and adapted bundles. This section reviews topological quantities associated to 4-manifolds with cylindrical ends. We begin with a review of characteristic classes in the closed (compact with no boundary) setting.

Suppose Z is a closed, oriented 4-manifold and $P \rightarrow Z$ is a principal G -bundle. We define

$$\kappa(P) := -\frac{1}{8\pi^2} \int_Z \langle F_A \wedge F_A \rangle$$

where $A \in \mathcal{A}(P)$ is any connection and $\langle F_A \wedge F_A \rangle$ is obtained by combining the wedge and the inner product on \mathfrak{g} defined via the immersion (2.1). Then $\kappa(P)$ is independent of the choice of A by the Bianchi identity. Topologically, $\kappa(P) = c_2(P \times_G \mathbb{C}^r)[Z]$ is the second Chern number of the \mathbb{C}^r -bundle associated to P via the map (2.1) and the standard action of $SU(r)$ on \mathbb{C}^r . In particular, $\kappa(P) \in \mathbb{Z}$ is an integer representing an obstruction to P being trivializable.

Now consider the bundle $Q \rightarrow N$ over the 3-manifold N , and fix a gauge transformation $u \in \mathcal{G}(Q)$. We can form the mapping torus $Q_u = [0, 1] \times Q / (0, u(q)) \sim (1, q)$ which is a principal G -bundle over $S^1 \times N$. Then we define the *degree* of u to be the integer

$$\deg(u) := \kappa(Q_u).$$

This depends only on the homotopy type of u and so descends to a group homomorphism

$$\deg: \pi_0(\mathcal{G}(Q)) \longrightarrow \mathbb{Z}$$

from the group of components of $\mathcal{G}(Q)$. The degree is an obstruction to extending u to a gauge transformation on X_0 (or equivalently X). We denote by $\mathcal{G}_0(Q)$ the subgroup of degree-zero gauge transformations. When G is simply-connected, the degree $\deg: \pi_0(\mathcal{G}(Q)) \rightarrow \mathbb{Z}$ is injective, and so $\mathcal{G}_0(Q)$ is exactly the identity component of $\mathcal{G}(Q)$.

Since the cylinder $\text{End } X$ deformation retracts to the 3-manifold N , we have a natural isomorphism $\pi_0(\mathcal{G}(\text{End } X)) \cong \pi_0(\mathcal{G}(N))$ and so the degree provides a homomorphism $\deg: \pi_0(\mathcal{G}(\text{End } X)) \rightarrow \mathbb{Z}$. We denote by $\mathcal{G}_0(\text{End } X)$ the degree-zero elements of $\mathcal{G}(\text{End } X)$.

We will be working with principal G -bundles on the cylindrical end 4-manifold X . Bundle isomorphism is too coarse of an equivalence relation to be useful in the cylindrical-end setting (e.g., when G is simply-connected, all principal G -bundles are trivializable since $H^4(X) = 0$). A more useful relation for our purposes deals with *adapted bundles*, which are pairs (E, A_{End}) , where $E \rightarrow X$ is a principal G -bundle, and A_{End} is a connection on the cylindrical end $\text{End } X$. Then we say that (E, A_{End}) is *equivalent* to (E', A'_{End}) if there is a bundle isomorphism from E to E' that carries A_{End} to A'_{End} . See Donaldson's book [4, Section 3.2] for more details; note that Donaldson only treats flat connections A_{End} , but our applications require that we extend the discussion.

By the above discussion, it follows that $\mathcal{G}_0(\text{End } X)$ consists of gauge transformations on $\text{End } X$ that have extensions to $E \rightarrow X$. Thus, any adapted bundle (E, A_{End}) depends on A_{End} only through its $\mathcal{G}_0(\text{End } X)$ -equivalence class. The next example illustrates an

interplay between the degree and the equivalence classes of adapted bundles; it will be relevant to our gluing discussion below.

Example 6.1. Fix an adapted bundle (E, A_{End}) and a point $x \in X$. Suppose $E_\ell \rightarrow S^4$ is a principal G -bundle with $\kappa(E_\ell) = \ell \in \mathbb{Z}$. Taking the connected sum of X and S^4 at x , we recover the same manifold back $X \cong X \# S^4$, up to diffeomorphism. At the bundle level, we can carry out a similar connected sum procedure to obtain a bundle $E' = E \# E_\ell$ over X . Provided $x \in X_0$ is not on the end, the connection A_{End} can be viewed as a connection on E' . Then the adapted bundles (E, A_{End}) and (E', A_{End}) are equivalent if and only if $\ell = 0$. More generally, there is a gauge transformation u on $\text{End } X$ with $\deg(u) = \ell$, and so that the adapted bundle $(E, u^* A_{\text{End}})$ is equivalent to (E', A_{End}) .

Assume that the connection A_{End} converges on the end in the sense that

$$\lim_{t \rightarrow \infty} A_{\text{End}}|_{\{t\} \times N} = \Gamma$$

for some connection Γ on N , where the limit is in $L^2_1(N)$, say. Let A be any connection on E that restricts on $\text{End } X$ to A_{End} . Then the quantity

$$\kappa(E, A_{\text{End}}) := \lim_{T \rightarrow \infty} -\frac{1}{8\pi^2} \int_{X_0 \cup [0, T] \times N} \langle F_A \wedge F_A \rangle$$

is well-defined and independent of the choice of A . We will call $\kappa(E, A_{\text{End}})$ the *relative characteristic number* of the adapted bundle (E, A_{End}) . It depends on A_{End} only through the value of Γ and the topological type of E . Indeed, if $E' = E \# E_\ell$ is as in Example 6.1, then

$$\kappa(E', A_{\text{End}}) = \kappa(E, A_{\text{End}}) + \ell.$$

Moreover, if A_{End} is asymptotic to Γ , then working modulo \mathbb{Z} , we recover

$$8\pi^2 \kappa(E, A_{\text{End}}) = \mathcal{CS}(\Gamma) \pmod{\mathbb{Z}}$$

the Chern–Simons value of Γ as defined in [17, Section 2.1] (here one should interpret the trace in [17] as the one induced from (2.1)).

6B. Existence when $b^+(X) = 0$. Let $E_{\text{triv}} \rightarrow X$ be the trivial bundle, A_{triv} the trivial connection on E_{triv} , and Γ_{triv} the trivial connection on the end. Fix thickening data $\mathcal{T}_{\Gamma_{\text{triv}}}$. Here we assume δ and β are chosen as in the beginning of Section 5. We recall from Section 2B that the thickening data also includes the choice of $\epsilon_0 > 0$ so that any two points in the center manifold have Chern–Simons values differing by $\epsilon_0/2$. For each $0 < \epsilon < \epsilon_0$, we will write $\mathcal{T}(\epsilon)$ for the same set of thickening data as $\mathcal{T}_{\Gamma_{\text{triv}}}$, but with ϵ in place of ϵ_0 .

Let $E_\ell \rightarrow S^4$ be a principal G -bundle with $\kappa(E_\ell) = \ell \in \mathbb{Z}$, where κ is the characteristic number of Section 6A. We will write

$$\mathcal{M}_\ell(S^4, G) := \{A \in \mathcal{A}(E_\ell) \mid F_A^+ = 0\} / \mathcal{G}(E_\ell)$$

for the moduli space of ASD connections on E_ℓ ; here we are working relative to the standard metric on S^4 . Let $\mathcal{M}_\ell^*(S^4, G) \subseteq \mathcal{M}_\ell(S^4, G)$ denote the subset of irreducible

ASD connections. The existence of irreducible ASD connections on S^4 was studied extensively in [1, Section 8]. For example, when $G = \mathrm{SU}(r)$ and the embedding (2.1) is the identity, then the space $\mathcal{M}_\ell^*(S^4, \mathrm{SU}(r))$ is nonempty if and only if $\ell \geq r/2$. The most famous situation is when $G = \mathrm{SU}(2)$ and $\ell = 1$, in which case $\mathcal{M}_1(S^4, \mathrm{SU}(2)) = \mathcal{M}_1^*(S^4, \mathrm{SU}(2))$ and this is diffeomorphic to the open unit ball in \mathbb{R}^5 . The dimension of $\mathcal{M}_\ell^*(S^4, G)$ for general simple, simply-connected G is given in [1, Table 8.1].

The following is the first of our main existence results for mASD connections; Theorem C with $b^+(X) = 0$ from the introduction is an immediate consequence.

Theorem 6.2 (Existence of mASD-connections when $b^+ = 0$). *Assume $b^+(X) = 0$ and A_ℓ is an irreducible ASD connection on $E_\ell \rightarrow S^4$ for some $\ell \in \mathbb{Z}$. Then for every $0 < \epsilon < \epsilon_0$, there is*

- (a) a neighborhood $\mathcal{V} \subseteq \mathcal{M}_\ell(S^4, G)$ of $[A_\ell] \in \mathcal{M}_\ell(S^4, G)$;
- (b) a trivializable principal G -bundle $E \rightarrow X$ that is canonically trivial on the end;
- (c) a connection A' on E that is flat in the complement of a compact set, asymptotic to Γ on the end, and satisfies $\kappa(E, A'|_{\mathrm{End} X}) = \ell$; and
- (d) a C^m -embedding

$$\Phi : \mathcal{V} \longrightarrow \widehat{\mathcal{M}}_{\mathrm{reg}}(\mathcal{T}(\epsilon), A').$$

The image of Φ consists of irreducible connections. In particular, there exist an irreducible, regular mASD connection A on E with $|\kappa(E, A|_{\mathrm{End} X}) - \ell| < \epsilon/2$.

For $\ell \neq 0$ and $0 < \epsilon < 2|\ell|$, the condition $|\kappa(E, A|_{\mathrm{End} X}) - \ell| < \epsilon/2$ implies that A is not flat. (The analogous statement in the case where X is closed and $G = \mathrm{SU}(r)$ is that A is supported on a bundle P with $c_2(P)[X] = \ell$.)

Proof of Theorem 6.2. View S^4 as a cylindrical end 4-manifold with no ends, and let \mathcal{T}_\emptyset be the empty set of thickening data as in Section 2C.3. Form the connected sum of S^4 and X at any point in S^4 and any point in X lying in the interior of the compact part. Note that all ASD connections on S^4 are regular (e.g., use the Weitzenböck formula [5, (7.1.23)]). In particular, since A_ℓ is irreducible and regular, there is a neighborhood $\mathcal{V} \subseteq \mathcal{M}_\ell^*(S^4, G)$ of $[A_\ell]$ that is diffeomorphic to a precompact open set $G_1 \subseteq \widehat{\mathcal{M}}_{\mathrm{reg}}(\mathcal{T}_\emptyset, A_\ell)$ containing A_ℓ ; that is, \mathcal{V} consists of gauge equivalence classes of regular, irreducible ASD connections on S^4 , and G_1 consists of their lifts to the Coulomb slice through A_ℓ . Using this diffeomorphism, we identify \mathcal{V} and G_1 .

By Proposition 2.15, the assumption $b^+(X) = 0$ implies that A_{triv} is regular; see also Remark 2.19. Then the singleton set $G_2 := \{A_{\mathrm{triv}}\}$ plainly consists of regular connections. By Remark 5.3 (a), we can apply Theorem 5.1 with this G_2 . Define $E := E_{\mathrm{triv}} \# E_\ell$ to be the connected sum bundle as in Section 6A, equipped with thickening data $\mathcal{T}(\epsilon)$. Take A' to be the preglued connection, which is plainly asymptotic to Γ_{triv} . In particular, $\kappa(E; A'|_{\mathrm{End} X}) = \ell$ due to the discussion of Section 6A. By possibly shrinking \mathcal{V} , if necessary, we define Φ to be the C^m -embedding of the same name from Theorem 5.1 (b) (here we are using the identifications $\mathcal{V} \cong G_1 \cong G_1 \times G_2$). If A is any connection in the image of Φ , then $|\kappa(E, A|_{\mathrm{End} X}) - \ell| < \epsilon/2$ follows from the definition of ϵ as

a parameter in the set of thickening data and the fact that $\kappa(E; A|_{\text{End } X})$ recovers the Chern–Simons value of the asymptotic limit of A . \square

6C. Existence when $b^+(X) = 1$. Here we consider the case where $b^+(X) = 1$ and $G = \text{SU}(2)$. We assume (2.1) is the identity (so $r = 2$); then the characteristic number κ from Section 6A is the second Chern number. Fix thickening data $\mathcal{T}_{\Gamma_{\text{triv}}}$ and assume δ and β are chosen as in the beginning of Section 5. Define $\mathcal{T}(\epsilon)$ as in Section 6B. The following is the second of our main existence results for mASD connections. Theorem C with $b^+(X) = 1$ is an immediate consequence.

Theorem 6.3 (Existence of mASD-connections when $b^+ = 1$). *Assume $b^+(X) = 1$ and fix an integer $\ell \geq 2$. Then for every $0 < \epsilon < \epsilon_0$, there is*

- (a) *a trivializable principal $\text{SU}(2)$ -bundle $E \rightarrow X$ that is canonically trivial on the end;*
- (b) *a connection A' on E that is flat in the complement of a compact set, asymptotic to Γ on the end, and satisfies $\kappa(E; A'|_{\text{End } X}) = \ell$; and*
- (c) *an irreducible mASD connection $A \in \widehat{\mathcal{M}}(\mathcal{T}(\epsilon), A')$ on E satisfying*

$$|\kappa(E, A|_{\text{End } X}) - \ell| < \epsilon/2.$$

Proof. Our proof follows that of [20, Section 7] and [3, pp. 327–334]. The assumption that $b^+ = 1$ implies that the cokernel of the operator $d^+ : L^p_{1,\delta}(\Omega^1(X, \mathbb{R})) \rightarrow L^p_{\delta}(\Omega^+(X, \mathbb{R}))$ is one-dimensional. As in Section 2C.1, this cokernel can be realized as the space $H^+(X, \mathbb{R})$ of closed self-dual 2-forms in $L^2(\Omega^+(X, \mathbb{R}))$ that restrict on each slice $\{t\} \times N$ to be orthogonal to the space of harmonic forms on N . Fix a non-zero element $\omega_0 \in H^+(X, \mathbb{R})$; this is unique up to scaling. By unique continuation for solutions of elliptic equations, it follows that the set of points in X where ω_0 does not vanish is open and dense. In particular, we can find two distinct points $x_1, x_2 \in \text{int}(X_0)$ with $\omega_0(x_1) \neq 0$ and $\omega_0(x_2) \neq 0$. When $\ell > 2$, choose additional points $x_3, \dots, x_{\ell-2} \in \text{int}(X_0)$; these can be arbitrarily chosen, provided the x_i are all distinct. The gluing Theorem 3.3 has a straightforward extension to handle gluing for multiple connected sums that we briefly describe now.

Fix scaling parameters $\lambda_1, \dots, \lambda_{\ell} > 0$, and set

$$\lambda := \max(\lambda_1, \dots, \lambda_{\ell}).$$

Here we will consider ℓ copies of S^4 ; denote these copies by S^4_1, \dots, S^4_{ℓ} , and fix points $x'_i \in S^4_i$. Then as we did in Section 3A, glue $x_i \in X_0$ to $x'_i \in S^4_i$ over balls with radii controlled by λ_i .

At the bundle level, let $E_{\text{triv}} \rightarrow X$ be the trivial $\text{SU}(2)$ -bundle, and let $E_1 \rightarrow S^4$ be the $\text{SU}(2)$ -bundle with $\kappa(E_1) = c_2(E_1) [S^4] = 1$. More concretely, we can take E_1 to be the frame bundle of $\Lambda^+ T^* S^4$ (then $\Lambda^+ T^* S^4$ is the adjoint bundle of E_1). In Section 3, gluing the bundles depended on the choice of fiber isomorphism ρ identifying the fibers of the principal bundles at the gluing points. In the present setting with ℓ gluing points, this corresponds to the choice of a fiber isomorphism

$$\rho_i \in \text{Gl}_i = \text{Hom}_{\text{SU}(2)}((E_{\text{triv}})_{x_i}, (E_1)_{x'_i})$$

for each $1 \leq i \leq \ell$.

Let A_{triv} be the trivial connection on E_{triv} . Let A' be the preglued connection on E obtained from A_{triv} and the standard “one-instanton” A_{st} on each of the bundles $E_1 \rightarrow S_i^4$ for $1 \leq i \leq \ell$. Then $\kappa(E, A'|_{\text{End } X}) = \ell$; note that A' depends on the λ_i and ρ_i . The proof of Theorem 3.3 extends to produce C, L, λ_0, J, π , and $\zeta \in L_\delta^p(\Omega^+(X))$, satisfying the conditions of Theorem 3.3 (a)–(c) and Corollary 3.26 whenever $0 < \lambda < \lambda_0$; though we suppress this in the notation, these quantities depend on the connections A_{triv} and A_{st} , as well as the isomorphisms ρ_i . In particular, the connection $A := J(\zeta)$ is irreducible and satisfies

$$s(A) = -\sigma\pi\zeta \quad \text{and} \quad |\kappa(E, A|_{\text{End } X}) - \ell| < \epsilon/2.$$

It suffices to show that the λ_i and ρ_i can be chosen so that $\sigma\pi\zeta = 0$, since this implies that A is mASD. For this, let X' be the complement in X of the $L\lambda_0^{1/2}$ -balls around the x_i ; we assume λ_0 is small enough so these balls do not intersect and are contained in X_0 . Note that the bundles E and E_{triv} are canonically identified over X' , and so over X' we can compare 2-forms on E_{triv} with 2-forms on E . The self-dual 2-form $\sigma\pi\zeta$ vanishes if and only if the integral

$$(6.4) \quad \int_{X'} \langle \omega \wedge \sigma\pi\zeta \rangle = 0$$

vanishes for all $\omega \in H^+(X, \text{ad}(A_{\text{triv}})) = H^+(X, \mathbb{R}) \otimes \mathfrak{g}$.

Claim:

$$(6.5) \quad \int_{X'} \langle \omega \wedge \sigma\pi\zeta \rangle = q_\omega^\ell(\{(\lambda_i, \rho_i)\}_i) + O(\lambda^3)$$

where

$$q_\omega^\ell(\{(\lambda_i, \rho_i)\}_i) := \sum_{i=1}^{\ell} \lambda_i^2 \text{tr}(\rho_i \omega(x_i)).$$

Here $\text{tr}(\rho_i \omega(x_i)) \in \mathbb{R}$ is the pairing of ρ_i and $\omega(x_i)$ as described in [3, Equation (5.3)]. We will prove this claim below, but first we will show how it is used to finish the proof of the theorem. From the discussion leading up to the claim, we are interested in the simultaneous system of equations

$$(6.6) \quad q_\omega^\ell(\{(\lambda_i, \rho_i)\}_i) = 0 \quad \forall \omega \in H^+(X, \text{ad}(A_{\text{triv}})).$$

When $\ell = 2$, the argument of [3, Section V(ii)] carries over verbatim to show that the solutions set of the system (6.6) is non-empty and cut out transversely, whenever $\max(\lambda_1, \lambda_2)$ is sufficiently small. This uses the assumption $\omega_0(x_1), \omega_0(x_2) \neq 0$. Note that Donaldson’s argument uses $b^+(X) = 1$. (Alternatively, the reader could follow the original argument of Taubes [20, Prop. 7.1], but our notation is more inline with that of [3].) When $\ell > 2$, it was pointed out by Taubes [20, Prop. 6.2] that by taking $\max(\lambda_3, \dots, \lambda_\ell)$ sufficiently small relative to $\max(\lambda_1, \lambda_2)$, any transverse zero of q_ω^2 implies the existence of a transverse zero of q_ω^ℓ . In summary, for each $\ell \geq 2$, there are

$\lambda'_0 > 0$ and $\mu \in (0, 1)$ so that the system (6.6) has a nonempty, transverse solution set, for all $\lambda_1, \dots, \lambda_\ell > 0$ with

$$\max(\lambda_1, \lambda_2) < \lambda'_0 \quad \text{and} \quad \max(\lambda_3, \dots, \lambda_\ell) < \mu \max(\lambda_1, \lambda_2).$$

For any such $\lambda_1, \dots, \lambda_\ell$, since q_ω^ℓ is $O(\lambda^2)$, it then follows from the transversality of $q_\omega^\ell = 0$ and the identity (6.5) that the solution sets to (6.4) and (6.6) are diffeomorphic, provided λ is sufficiently small. In particular, there is a simultaneous zero $\{(\lambda_i, \rho_i)\}_{i=1}^\ell$ of the solution set to (6.4). For this collection of gluing data, the glued connection A is therefore mASD, as desired.

It therefore suffices to verify the above Claim. We will first unpack the notation. Note that the preglued connection A' restricts on X' to equal the trivial flat connection. Let $A'(\lambda_0)$ be the preglued connection defined using λ_0 at every gluing site, and the same ρ_i as was used to define A' (so the only difference between A' and $A'(\lambda_0)$ is that the former uses λ_i at the gluing site x_i , while the latter uses λ_0 at all gluing sites). Define the map i (and hence ι) using $A'(\lambda_0)$ as a reference connection. Write $\Gamma = \Gamma_{\text{triv}}$ for the trivial connection on N . Then we can write $A' = \iota(\Gamma_{\text{triv}}, V') = i(\Gamma_{\text{triv}}) + V' = A'(\lambda_0) + V'$ for some 1-form V' . It follows that $V'|_{X'} = 0$, and we note also that $A'(\lambda_0)|_{X'} = A_{\text{triv}}$.

Next, recall the map $P : L_\delta^p(\Omega^+) \rightarrow T_\Gamma \mathcal{H} \times L_{1,\delta}^p(\Omega^1)$ from Claim 1 in the proof of Theorem 3.3, and write $P\xi = (\eta, V)$. The definition of the map $J = J_{A_{\text{triv}}, A_{\text{st}}}$ gives

$$J(\xi) = i(\exp_\Gamma(\eta)) + V' + V$$

where $\exp_\Gamma : T_\Gamma \mathcal{H} \rightarrow \mathcal{H}$ is the exponential. The observations of the previous paragraph combine with the formula (2.11) to give that the restriction of $s(A)$ takes the following form:

$$s(A)|_{X'} = (1 - \beta')F_{i(\exp_\Gamma(\eta))}^+ + d_{i(\exp_\Gamma(\eta))}^+ V + \frac{1}{2} [V \wedge V]^+.$$

Returning to the integral (6.4), we can use the defining property of ξ and the above identity for $s(A)$ to get

$$\begin{aligned} \int_{X'} \langle \omega \wedge \sigma \pi \xi \rangle &= - \int_{X'} \langle \omega \wedge s(A) \rangle \\ (6.7) \quad &= - \int_{X'} (1 - \beta') \langle \omega \wedge F_{i(\exp_\Gamma(\eta))}^+ \rangle - \int_{X'} \langle \omega \wedge d_{i(\exp_\Gamma(\eta))}^+ V \rangle \\ &\quad - \frac{1}{2} \int_{X'} \langle \omega \wedge [V \wedge V]^+ \rangle. \end{aligned}$$

Focus on the last term on the right. Recall from Lemma 2.20 that ω decays in \mathcal{C}^0 like $e^{-\mu_\Gamma^+ t}$. In particular, ω is bounded and so

$$\int_{X'} \left| \langle \omega \wedge [V \wedge V]^+ \rangle \right| \leq C_1 \|V\|_{L^2(X)}^2 \leq C_1 \|V\|_{L_\delta^2(X)}^2 \leq C_1 \|P\xi\|_{L_\delta^2(X)}^2$$

for some constant C_1 . By Corollaries 3.26 and 3.30, this term decays like λ^3 :

$$-\frac{1}{2} \int_{X'} \langle \omega \wedge [V \wedge V]^+ \rangle = O(\lambda^3).$$

We can control the nonlinear parts of the other two terms in (6.7) similarly. Indeed, use $i(\exp_\Gamma(\eta)) = A_{\text{triv}} + (Di)_\Gamma \eta + O(\eta^2)$ and the expansion formulas for the curvature and covariant derivative, to get

$$\int_{X'} \langle \omega \wedge \sigma \pi \zeta \rangle = - \int_{X'} (1 - \beta') \langle \omega \wedge d^+(Di)_\Gamma \eta \rangle - \int_{X'} \langle \omega \wedge d^+ V \rangle + O(\lambda^3)$$

where $d = d_{A_{\text{triv}}}$. Focus on the first term on the right (there is no analogue of this term in the standard ASD framework). It follows from the definitions of β' and i that $(1 - \beta')(d^+(Di)_\Gamma \eta)$ is supported on $[T - 1/2, T + 1/2] \times N$. Using the formula (2.17), we have

$$- \int_{X'} (1 - \beta') \langle \omega \wedge d^+(Di)_\Gamma \eta \rangle = - \int_{X'} (1 - \beta') (\partial_t \beta'') \langle \omega \wedge (dt \wedge \eta)^+ \rangle = 0$$

where the last equality uses the facts that (i) $\eta \in H_{\Gamma_{\text{triv}}}^1$ is in the harmonic space on N , and (ii) elements of $H^+(X, \text{ad}(A_{\text{triv}}))$ restrict on each slice to be orthogonal to $H_{\Gamma_{\text{triv}}}^1$. In summary, this gives

$$\begin{aligned} \int_{X'} \langle \omega \wedge \sigma \pi \zeta \rangle &= - \int_{X'} \langle \omega \wedge d^+ V \rangle + O(\lambda^3) \\ &= \int_{\partial X'} \langle \omega \wedge V \rangle + O(\lambda^3). \end{aligned}$$

What remains is to estimate the integral $\int_{\partial X'} \langle \omega \wedge V \rangle$. This is an integral taking place at the boundary of the disks centered at the gluing sites x_1, \dots, x_ℓ . In particular, this integral is identical to the analogous term that arises when gluing in the standard ASD setting (e.g., see the top of [3, p. 328]). Then the argument of [3, pp. 328—331] carries over verbatim to give

$$\int_{\partial X'} \langle \omega \wedge V \rangle = q_\omega^\ell(\{(\lambda_i, \rho_i)\}) + O(\lambda^3).$$

This proves (6.5). □

6D. An ASD existence result and a proof of Theorem A. Recall from Section 2A the definition of the vector field Ξ_Γ on the center manifold. We will be interested in the case where the flat connection Γ satisfies the following hypothesis:

Hypothesis H. *There is a neighborhood $U \subseteq \mathcal{H}_\Gamma$ of Γ so that every $a \in U$ flows under Ξ_Γ to a flat connection in U .*

Example 6.8.

(a) Recall from Section 2A that U_Γ is a neighborhood of Γ in the Coulomb slice through Γ . Suppose the set of flat connections in U_Γ is smooth in a neighborhood $U' \subseteq U_\Gamma$ of Γ and has the same dimension as \mathcal{H}_Γ . Then $U := U' \cap \mathcal{H}_\Gamma$ satisfies Hypothesis H.

(b) The assumption of (a) trivially holds when Γ is non-degenerate, since \mathcal{H}_Γ consists of a single point. More generally, the assumption of (a) also holds when the Chern–Simons function is Morse–Bott in a neighborhood of Γ (though there is no assumption in part (a) about the nondegeneracy of the Hessian in the normal directions).

(c) Suppose $N = T^3$, and let Γ be a flat connection on the trivial $SU(2)$ -bundle. If Γ is not gauge equivalent to the trivial connection, then Γ satisfies the assumption in (a), and hence Hypothesis H; see [10, Lemma 14.2(i)]. However, the trivial connection on T^3 does *not* satisfy Hypothesis H.

The main usefulness of Hypothesis H for us is through the following theorem.

Theorem 6.9. *Consider the situation of Theorem 3.3, and assume A_1 and A_2 are regular. In addition, assume that Γ_1 and Γ_2 each satisfy Hypothesis H. Let $\lambda_0 > 0$ be the constant from Theorem 3.3. Then there is some $0 < \lambda'_0 \leq \lambda_0$ so that for all $\lambda \in (0, \lambda'_0)$ the mASD connection $\mathcal{J}(A_1, A_2)$ guaranteed by Theorem 3.3 (and hence by Theorems B and C) is in fact ASD.*

Proof. Fix $0 < \lambda < \lambda_0$ and let $A_\lambda := \mathcal{J}(A_1, A_2)$ be the mASD connection from Theorem 3.3 associated to this value of λ . Recall that $\Gamma = \Gamma_1 \sqcup \Gamma_2$, and so Γ satisfies Hypothesis H since the Γ_k do. It follows from (3.6) that $p_T(A_\lambda) \in \mathcal{H}_{in}$ converges to Γ as λ approaches 0. In particular, by taking λ sufficiently small, Hypothesis H implies that the Ξ_Γ -flow line beginning at $p_T(A_\lambda)$ lies in \mathcal{H}_{in} for all positive time. This implies $i(p_T(A_\lambda))$ is ASD (see the paragraph just before the statement of Lemma 2.3), and so

$$F_{A_\lambda}^+ = F_{A_\lambda}^+ - \beta' F_{i(p_T(A_\lambda))}^+ = s(A_\lambda) = 0.$$

□

Now we can prove our application from the introduction.

Proof of Theorem A. Take Γ to be the trivial flat connection on the trivial $SU(2)$ -bundle. We will show that the two conditions on N stated in Theorem A each imply that Γ satisfies Hypothesis H; it will then be immediate that the mASD connection guaranteed by Theorem C is in fact ASD, as desired.

First assume N is a circle bundle over a surface with positive Euler class. Then [17, Corollary 13.2.2] implies that Γ satisfies the condition of Example 6.8 (a), and thus Hypothesis H.

Now assume that $b_1(N) \leq 1$. If $b_1(N) = 0$, then $H_\Gamma^1 = H^1(N) \otimes \mathfrak{g} = 0$ and so Γ is nondegenerate. Thus, Γ again satisfies Hypothesis H, but this time by Example 6.8 (b).

Finally, suppose $b_1(N) = 1$. We will show here that Γ satisfies the condition of Example 6.8 (a). Since $b_1(N) = 1$, there is a loop $\gamma : S^1 \rightarrow N$ and a harmonic 1-form $\eta \in \Omega^1(N, \mathbb{R})$ so that $\int_\gamma \eta = 1$. For each $\xi \in \mathfrak{g}$, let

$$a_\xi := \Gamma + \xi \otimes \eta.$$

We claim that a_ξ lies in the center manifold \mathcal{H}_Γ for all sufficiently small ξ . To see this, first note that a_ξ is flat, since

$$F_{a_\xi} = F_\Gamma + \xi \otimes (d\eta) + \frac{1}{2}[\xi, \xi] \otimes \eta \wedge \eta = 0.$$

This connection also lies in the Coulomb gauge slice for Γ , since

$$d_{\Gamma}^*(a_{\bar{\zeta}} - \Gamma) = \bar{\zeta} \otimes d^*\eta = 0.$$

Recall the map Θ and the vector field ∇f_{Γ} from Section 2A.1. Since $F_{a_{\bar{\zeta}}} = 0$, we have

$$\Theta(a_{\bar{\zeta}}) = 0.$$

Thus $\nabla_{a_{\bar{\zeta}}} f_{\Gamma} = 0$. One of the defining features of \mathcal{H}_{Γ} is that it contains all zeros of ∇f_{Γ} that are sufficiently close to Γ , so this proves the claim.

It thus follows that there is some $\epsilon > 0$ so that the map

$$B_{\epsilon}(0) \subseteq \mathfrak{g} \longrightarrow \mathcal{H}_{\Gamma} \quad \bar{\zeta} \longmapsto a_{\bar{\zeta}}$$

is well-defined. It is clearly an immersion, so a dimension count implies that it must be a local diffeomorphism; this uses the fact that $b_1(N) = 1$. This establishes the condition of Example 6.8 (a). \square

7. PARTIAL COMPACTIFICATION—THE TAUBES BOUNDARY

Here we give a more global formulation of the result of Theorem 6.2 in the case where $G = \mathrm{SU}(2)$ and $\ell = 1$. Fix a closed set X'_0 contained in the interior of the compact part X_0 . Let A_{st} be the standard one-instanton on the $\mathrm{SU}(2)$ -bundle $E_1 \rightarrow S^4$ with $c_2(E_1) [S^4] = 1$. For $x \in X'_0$, let X_x be the connected sum of X and S^4 obtained by gluing $x \in X$ to the north pole in S^4 . Similarly, glue the trivial $\mathrm{SU}(2)$ -bundle on X to $E_1 \rightarrow S^4$ and let $E_x \rightarrow X_x$ be the resulting bundle. Let $A'_x = A'(A_{\mathrm{triv}}, A_{\mathrm{st}})$ be the preglued connection on E_x , where A_{triv} is the trivial connection on X . Note that in the present situation, all auxiliary gluing data can be chosen to be independent of x . For example, the fiber isomorphism ρ of Section 3A can be taken to be independent of x since we are starting with the trivial bundle on X .

7A. The Taubes Boundary. Fix $\epsilon > 0$, and let $\mathcal{T}(\epsilon)$ be thickening data with this choice of ϵ , as in Section 6B. By Theorem 3.3, there are $\epsilon_0, \lambda_0 > 0$ so the following holds: For all $0 < \epsilon < \epsilon_0$ and $0 < \lambda \leq \lambda_0$, there is an irreducible, regular mASD connection

$$A(x, \lambda) := \mathcal{J}(A_{\mathrm{triv}}, A_{\mathrm{st}}) \in \mathcal{A}^{1,p}(\mathcal{T}(\epsilon))$$

with the property that $A(x, \lambda) - A_{\mathrm{triv}}|_{X \setminus \mathrm{nbhd}(x)}$ goes to zero in λ in the sense of (3.6). This ϵ_0 depends only on the trivial flat connection on the 3-manifold N ; hence ϵ_0 is independent of x . Since X'_0 is compact, we can assume this λ_0 is independent of x as well.

We want to allow x to vary, and for this, we form the space

$$\mathcal{E}' := \left\{ (x, \lambda, A) \in X'_0 \times (0, \lambda_0] \times \mathcal{A}^{1,p}(\mathcal{T}(\epsilon)) \mid A \in \widehat{\mathcal{M}}_{\mathrm{reg}}(\mathcal{T}(\epsilon), A(x, \lambda)) \right\}.$$

Let $\Pi' : \mathcal{E}' \rightarrow X'_0 \times (0, \lambda_0]$ be the projection to the first two factors. Then the assignment $\Psi'(x, \lambda) := (x, \lambda, A(x, \lambda))$ defines a section of Π' . Just as in Theorem 5.1, there is an open neighborhood $\mathcal{U}' \subseteq \mathcal{E}'$ of the image of Ψ' so that the restriction $\Pi'|_{\mathcal{U}'}$ is a locally-trivial \mathcal{C}^m -fiber bundle over $X'_0 \times (0, \lambda_0]$. By construction, the fiber over (x, λ) is an open subset of $\widehat{\mathcal{M}}_{\mathrm{reg}}(\mathcal{T}(\epsilon), A(x, \lambda))$ containing $A(x, \lambda)$.

Remark 7.1. Here we describe a sense in which Theorem 6.2 can be viewed as a local version of this fiber bundle construction. Fix a small neighborhood $U_x \subseteq X_0$ around x . The gluing procedure of Section 3A identifies this with a small neighborhood of the north pole in S^4 . The standard description [9, Ch. 6] of the ASD moduli space $\mathcal{M}_1(S^4, \text{SU}(2))$ gives an embedding $S^4 \times (0, \lambda_0] \rightarrow \mathcal{M}_1(S^4, \text{SU}(2))$ with the S^4 -component specifying the center of mass and $(0, \lambda_0]$ parametrizing the scale of the curvature; here the energy-density of the curvature is concentrating, as λ approaches 0, to a Dirac delta measure supported at the center of mass. Combining these, we have a diffeomorphism

$$f : U_x \times (0, \lambda_0] \longrightarrow \mathcal{V} \subseteq \mathcal{M}_1(S^4, \text{SU}(2))$$

onto an open set \mathcal{V} . It follows from this construction that there is a local trivialization of the fiber bundle $\Pi'|_{\mathcal{U}'}$ relative to which Ψ' takes the form $(y, \lambda) \mapsto (y, \lambda, \Phi(f(y, \lambda)))$ where Φ is the map of Theorem 6.2. In fact, by possibly shrinking U_x further, this local trivialization can be taken to be over the full cylinder $U_x \times (0, \lambda_0]$; this due to the fact that the constructions in the proof of Theorem 5.1 can be taken to be uniform in λ . This construction is exploiting a coupling between the parameter λ and the “scale” parameter for the concentration of instantons on $E_1 \rightarrow S^4$; see [5, p. 323] for a related discussion.

Now we consider the behavior of this section Ψ' near $\lambda = 0$. For this, suppose $(x_n, \lambda_n) \in X'_0 \times (0, \lambda_0]$ is a sequence with $\lambda_n \rightarrow 0$; we will call $\Psi'(x_n, \lambda_n)$ a *bubbling sequence in X'_0* . By passing to a subsequence, we may assume the x_n converge to some $x_\infty \in X'_0$. It follows from a straight-forward Uhlenbeck-type compactness argument and (3.6) that, after passing to a subsequence, the associated connections $A(x_n, \lambda_n)$ converge *weakly* to the ideal connection (A_{triv}, x_0) in the sense that the energy densities $|F_{A(x_n, \lambda_n)}|^2$ converge in measure to the delta measure supported at x_0 , and

$$\lim_{n \rightarrow \infty} \left\| \iota^{-1}(A(x_n, \lambda_n)) - \iota^{-1}(A_{\text{triv}}) \right\|_{L^2_2(N) \times L^p_{1,\delta}(X \setminus B_r(x_0))} = 0$$

for all $r > 0$; see [5, Section 4.4.1] for the analogous ASD case.

Following the lead of [5, Section 4.4.1], the discussion of the previous paragraph can be framed geometrically as follows. Consider the set

$$I(\mathcal{U}') := \mathcal{U}' \cup (X'_0 \times \{A_{\text{triv}}\})$$

which we view as coupling the connections in $\mathcal{U}' \subseteq \mathcal{E}'$ into the same space as the above-mentioned ideal connections. We can extend $\Pi'|_{\mathcal{U}'}$ to a map $I(\Pi') : I(\mathcal{U}') \rightarrow X'_0 \times [0, \lambda_0]$ by declaring it to send (x, A_{triv}) to $(x, 0)$. Give $I(\mathcal{U}')$ any topology (more below) for which the map $I(\Pi')$ is continuous and so that the notion of weak convergence from the previous paragraph implies convergence in $I(\mathcal{U}')$; we assume also that this topology is first countable. Then the observations of the previous paragraph imply the section Ψ' extends continuously over $X'_0 \times \{0\}$ to a section $\bar{\Psi}'$ of $I(\Pi')$. It is due to this that we may view $I(\mathcal{U}')$ as a “partial compactification” for mASD connections: The bubbling sequences in X'_0 converge in $I(\mathcal{U}')$.

7B. Compactification Issues. We end this section with several comments about the construction of the partial compactification $I(\mathcal{U}')$, as well as some of its limitations. This partial compactification is constructed only so that bubbling sequences in X'_0 converge—our assumptions on the topology on $I(\mathcal{U}')$ do not necessarily imply subsequential convergence of other types of sequences. The simple reason for this is that we do not yet know how such sequences behave, and what additional limiting objects we would need to include in $I(\mathcal{U}')$ to ensure their subsequential convergence. What we are presently lacking is a sufficiently strong version of Uhlenbeck’s compactness theorem for mASD connections. In the end, such a theorem would need to (at least) address the following:

- (a) *Bubble formation on the end:* To what extent is the mASD condition preserved under Uhlenbeck limits where the curvature concentrates at a point in $\text{End } X$? More fundamentally, is the connections space $\mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ suitably closed under such limits? This is related to (c) below.

From the gluing perspective, we avoided these questions altogether by only gluing at points in the compact part where mASD connections are ASD; that is, $I(\mathcal{U}')$ only corresponds to the points in the “Taubes boundary” that corresponds to bubbles in $X'_0 \subseteq X \setminus \text{End } X$ (and relative to a *fixed* gluing parameter ρ ; see Remark 5.3 (b)). A more thorough investigation would require not only an understanding of the mASD condition under Uhlenbeck limits, but also an understanding of how to glue at points on $\text{End } X$.

- (b) *Energy escaping down the end:* One example of this is bubbling on the end, as discussed in (a). Another example is where a non-trivial amount of energy escapes down the end. This is familiar in the ASD setting, where compactification can be achieved by including spaces of translationally-invariant ASD connections on $\mathbb{R} \times N$ (spaces of “Floer trajectories”); see [8, 4]. In the mASD setting, one would likely need to include spaces of mASD connections on $\mathbb{R} \times N$ to account for energy escaping. The details of this appear to be subtle, since the energy values of such connections are not governed by topological quantities, as is the case in the ASD setting. (In the discussion above, where we considered sequences in the image of Ψ' , we were able to exclude non-trivial energy on the end by appealing to (3.6).)
- (c) *Failure of the slice-wise gauge fixing condition:* In the definition of the space $\mathcal{A}^{1,p}(\mathcal{T}_\Gamma)$ from Section 2B.1, we restricted attention to connections that restrict on each time slice $\{t\} \times N$, for $t \geq T$, to be gauge equivalent to a connection in the gauge slice U_Γ . This is an open condition in the space of all $L^p_{1,loc}$ connections, and we do not see a reason why this condition should be retained through limits of mASD connections.

It is clear from these observations that $I(\mathcal{U}')$ is by no means the end of the story when it comes to compactification. It is due to this that we have avoided defining a specific topology on $I(\mathcal{U}')$ above, choosing instead to axiomatize a minimal set of desirable properties.

Remark 7.2. A challenging problem is to establish a general “gluing theorem on the ends” for mASD-connections (X_1, A_1) and (X_2, A_2) with “matching boundary conditions” on two given 4-manifolds with cylindrical ends. Such a theorem was obtained in [16], for ASD-connections assuming that their flat limits in the common 3-manifold N end where irreducible smooth points in the representation variety of N . In this case, the gluing was unobstructed. More generally, in order to glue two mASD-connections, the matching conditions should at least include (i) an identification $\text{End } X_1 \cong \text{End } X_2 \cong N \times [0, \infty)$, (ii) the same flat reference connection Γ on N and center manifold \mathcal{H}_Γ , and (iii) the same flow lines $\hat{\omega}_h$ in \mathcal{H}_Γ , where $h = p_T(A_1) = p_T(A_2)$. We would expect the glued-up connections to provide an embedded submanifold of connections on the closed 4-manifold $X = X_1 \cup X_2$ obtained by identifying along their cylindrical ends. In this general mASD case, the glued-up connection would presumably satisfy some version of the mASD equation on X that equals the ASD equation on the complement of the neck. How do the ASD connections on X compare to these “mASD connections” on X ? For example, do these “mASD connections” on X form some sort of local thickening of the ASD moduli space, as is the case for cylindrical end manifolds? Since X is compact, it seems likely that the ASD operator differs from the “mASD operator” by a compact operator. Can this operator be scaled in some way to show that ASD and mASD spaces on X are, in some sense, cobordant?

REFERENCES

- [1] M. Atiyah, N. Hitchin, and I. Singer, *Self-duality in four-dimensional Riemannian geometry*, Proc. R. Soc. London, **362** (1978), 425–461.
- [2] S. K. Donaldson, *An application of gauge theory to four dimensional topology*, J. Differential Geom. **18** (1983) 279–315.
- [3] ———, *Connections, cohomology and the intersection forms of four manifolds*. J. Differential Geom. **24** (1986) 275–341.
- [4] ———, *Floer homology groups in Yang–Mills theory*, Cambridge Tracts in Mathematics, vol. 147, Cambridge University Press, Cambridge, 2002, With the assistance of M. Furuta and D. Kotschick.
- [5] S. K. Donaldson and P. B. Kronheimer, *The geometry of four-manifolds*, Oxford Mathematical Monographs, The Clarendon Press Oxford University Press, New York, 1990, Oxford Science Publications.
- [6] S. Dostoglou and D. Salamon, *Self-dual instantons and holomorphic curves*, Ann. of Math., **139** (1994), 581–640.
- [7] R. Fintushel and R. J. Stern, *Surgery in cusp neighborhoods and the geography of irreducible 4-manifolds*, Invent. Math **117** (1994), 455–523.
- [8] A. Floer, *An instanton-invariant for 3-manifolds*, Comm. Math. Phys. **118** (1988), 215–240.
- [9] D. S. Freed and K. K. Uhlenbeck, *Instantons and four-manifolds*, second ed., Mathematical Sciences Research Institute Publications, vol. 1, Springer-Verlag, New York, 1991.
- [10] R. Gompf and T. Mrowka, *Irreducible four manifolds need not be complex*, Annals of Math., **138** (1993), 61–111.
- [11] I. Hambleton and R. Lee, *Perturbation of equivariant moduli spaces*, Math. Ann. **293** (1992), 17–37.
- [12] ———, *Smooth group actions on definite 4-manifolds and moduli spaces*, Duke Math. J. **78** (1995), 715–732.
- [13] I. Hambleton and M. Tanase, *Permutations, isotropy and smooth cyclic group actions on definite 4-manifolds*, Geom. Topol. **8** (2004), 475–509.

- [14] S. Lang. *Differential and Riemannian Manifolds*, Graduate Texts in Mathematics, vol. 160, Springer-Verlag, New York, NY, 1995.
- [15] R. Lockhart and R. McOwen, *Elliptic differential operators on non-compact manifolds*, Annali di Scuola Norm. Sup. de Pisa (1985), 409–448.
- [16] J. W. Morgan and T. Mrowka, *On the gluing theorem for instantons on manifolds containing long cylinders*, unpublished preprint (1994?).
- [17] J. W. Morgan, T. Mrowka, and D. Ruberman, *The L^2 -moduli space and a vanishing theorem for Donaldson polynomial invariants*, Monographs in Geometry and Topology, II, International Press, Cambridge, MA, 1994.
- [18] T. S. Mrowka, *A local Mayer-Vietoris principle for Yang-Mills moduli spaces*, Ph.D Thesis (1988), University of California, Berkeley.
- [19] C. H. Taubes, *Self-dual Yang-Mills connections over non-self-dual 4-manifolds*, J. Differential Geom. **17** (1982), 139–170.
- [20] ———, *Self-dual connections on 4-manifolds with indefinite intersection matrix*, J. Differential Geom. **19** (1984), 517–560.
- [21] ———, *Gauge theory on asymptotically periodic 4-manifolds*, J. Differential Geom. **25** (1987), 363–430.
- [22] ———, *L^2 moduli spaces on 4-manifolds with cylindrical ends*, International Press, Cambridge, MA, 1993.
- [23] K. Wehrheim, *Uhlenbeck Compactness*, European Mathematical Society Publishing House, Zürich, 2004.

DEPARTMENT OF MATHEMATICS & STATISTICS, JAMES MADISON UNIVERSITY, HARRISONBURG, VIRGINIA 22807

Email address: duncandl@jmu.edu

DEPARTMENT OF MATHEMATICS & STATISTICS, MCMASTER UNIVERSITY, HAMILTON, ONTARIO L8S 4K1, CANADA

Email address: hambleton@mcmaster.ca