

EQUIVARIANT INSTANTON HOMOLOGY

S. MICHAEL MILLER

ABSTRACT. We define four versions of equivariant instanton Floer homology (I^+ , I^- , I^∞ and \tilde{I}) for a class of 3-manifolds and $SO(3)$ -bundles over them including all rational homology spheres. These versions are analogous to the four flavors of monopole and Heegaard Floer homology theories. This construction is functorial for a large class of 4-manifold cobordisms, and agrees with Donaldson's definition of equivariant instanton homology for integer homology spheres. Furthermore, one of our invariants is isomorphic to Floer's instanton homology for admissible bundles, and we calculate I^∞ in all cases it is defined, away from characteristic 2.

The appendix, possibly of independent interest, defines an algebraic construction of three equivariant homology theories for dg-modules over a dg-algebra, the equivariant homology $H^+(A, M)$, the coBorel homology $H^-(A, M)$, and the Tate homology $H^\infty(A, M)$. The constructions of the appendix are used to define our invariants.

CONTENTS

Introduction	3
Survey of the homology theories	8
Organization	11
1. Configuration spaces and their reducibles	13
1.1. Framed connections and the framed configuration space	14
1.2. The equivalent $U(2)$ model	15
1.3. Reducibles on 3-manifolds	16
1.4. Configurations on the cylinder	20
1.5. Configurations on a cobordism	23
2. Analysis of configuration spaces	26
2.1. Tangent spaces	26
2.2. Hilbert manifold completions	27
2.3. The 4-dimensional case	29
3. Critical points and perturbations	33
3.1. Holonomy perturbations	33
3.2. Linear analysis	37
3.3. Reducible critical points	38
3.4. Transversality for critical points	40
3.5. Signature data of reducible critical points	42
4. Moduli spaces of instantons	46
4.1. Moduli spaces for cylinders and cobordisms	46
4.2. Linear analysis and index theory	52
4.3. Uhlenbeck compactness for framed instantons	56
4.4. Reducible instantons on the cylinder and cobordisms	65
4.5. Index calculations	72
4.6. Transversality for the cylinder and cobordisms	82
4.7. Gluing	87
4.8. Families of metrics and perturbations	102
5. Orientations	111
5.1. Orientability and determinant line bundles	111
5.2. A gluing operation for determinant lines	115
5.3. The gluing diffeomorphism is orientation-preserving	118
5.4. Canonical orientations for determinant lines	124
6. Floer homology	130
6.1. Geometric homology	130
6.2. Equivariant instanton homology	139
6.3. The index filtration	154
6.4. Four flavors of instanton homology	157
7. Examples, calculations, and comparisons	162
7.1. Equivariant instanton homology for admissible bundles	162
7.2. Comparison with Donaldson's theory	163
7.3. Instanton Tate homology	180
7.4. Examples of the I^\bullet and the index spectral sequence	183
7.5. Orientation reversal and equivariant cohomology	189
Appendix	192
A. Equivariant homology of dg-modules	192
A.1. Bar constructions	192

A.2. Invariance	195
A.3. The dualizing complex and Tate homology	200
A.4. Periodicity in Tate homology	204
A.5. Simplifying the twisted Borel homology	207
A.6. Spectral sequences	209
A.7. Group algebras	210
A.8. Periodic homological algebra	216
References	219

INTRODUCTION

In [Flo88], Andreas Floer introduced the instanton homology groups $I(Y)$, $\mathbb{Z}/8$ -graded abelian groups associated to integer homology 3-spheres. These form a sort of TQFT in which oriented cobordisms $W : Y_0 \rightarrow Y_1$ induce homomorphisms on the corresponding instanton homology groups. Since then, similar TQFT-style invariants have found themselves a powerful tool in 3- and 4-dimensional topology, especially the related monopole Floer homology of [KM07] and the Heegaard Floer homology of [OS04].

In ideal circumstances, the instanton homology groups are defined by a chain complex generated by *irreducible flat $SU(2)$ connections* up to isomorphism (equivalently, representations $\pi_1(Y) \rightarrow SU(2)$ whose image is non-abelian, modulo conjugacy by elements of $SU(2)$); the component of the differential between two flat connections α_-, α_+ is given by an algebraic count of solutions on the cylinder $\mathbb{R} \times Y$ to the *ASD equation*

$$F_{\mathbf{A}}^+ = 0,$$

where \mathbf{A} is a connection on the trivial $SU(2)$ -bundle over the cylinder which is asymptotically equal to the α_{\pm} . We can think of this as the “Morse chain complex” of the Chern-Simons functional on the space of irreducible connections modulo gauge equivalence, $\mathcal{B}_Y^* = \mathcal{A}_Y^*/\mathcal{G}$, whose critical points are the flat connections and gradient flow equation is (formally) the ASD equation. While these equations depend on a choice of metric on the 3-manifold Y , the homology groups are an invariant of the Y itself.

Floer’s theory is constrained to homology 3-spheres because of the presence of *reducible connections*. While the instanton chain complex above is still a chain complex for rational homology spheres, the proof that its homology is independent of the choice of metric fails in the presence of *$U(1)$ -reducible connections* (corresponding to representations $\pi_1(Y) \rightarrow SU(2)$ with image lying inside a circle subgroup). One would need to take these reducible connections into account in the definition of the chain complex, but this cannot be done naively: while the instanton chain complex is a Morse complex for the space of irreducible connections, which is an infinite-dimensional manifold, the gauge group \mathcal{G} does not act freely on the entire space of connections, and so the configuration space of *all* connections modulo gauge, \mathcal{B}_Y , is not a manifold.

Austin and Braam in [AB96] resolve this difficulty for a class of 3-manifolds (including all rational homology spheres) by defining an invariant called the *equivariant instanton homology* of Y , a $\mathbb{Z}/8$ -graded \mathbb{R} -vector space $I_*^G(Y)$ with an action

of $\mathbb{R}[U] = H^*(BSO(3); \mathbb{R})$, to be a form of $SO(3)$ -equivariant Morse theory on an infinite-dimensional $SO(3)$ -manifold $\tilde{\mathcal{B}}_Y$ with

$$\tilde{\mathcal{B}}_Y/SO(3) = \mathcal{B}_Y.$$

The manifold $\tilde{\mathcal{B}}_Y$ might be called the configuration space of *framed connections* on (the trivial $SO(3)$ -bundle over) Y . Their invariant is defined using the equivariant de Rham complex as a model for the equivariant (co)homology of a smooth G -manifold, and thus inherently uses real coefficients.

Floer also defined, in [Flo91], instanton homology groups for $SO(3)$ -bundles E over 3-manifolds Y satisfying the admissibility criterion that $w_2(E) \in H^2(Y; \mathbb{Z}/2)$ lifts to a non-torsion class in $H^2(Y; \mathbb{Z})$; in particular, $b_1(Y) > 0$. In this case, there are no reducible connections, and the homology of the Floer complex is a well-defined invariant of the pair (Y, E) . This case is important for his work on surgery triangles in instanton homology.

Using this, Kronheimer and Mrowka introduce *framed instanton homology groups* $I^\#(Y, E)$ for an arbitrary $SO(3)$ -bundle over a 3-manifold in [KM11b] by studying the instanton homology of $(Y \# T^3, E \# Q)$ for a certain admissible bundle Q over T^3 . This is meant to be a version of the (non-equivariant) Morse homology of the space of framed connections $\tilde{\mathcal{B}}_E$.

In this paper, we jointly generalize Floer's work on admissible bundles and Austin-Braam's work for rational homology spheres; to speak of both in the same breath, we say that an $SO(3)$ -bundle E over a 3-manifold Y is *weakly admissible* if either $w_2(E)$ has no lifts to a torsion class in $H^2(Y; \mathbb{Z})$, or if $b_1(Y) = 0$.

We take an alternate approach to Kronheimer and Mrowka's to the framed instanton homology groups: instead of taking a connected sum with T^3 , we work on the space of framed connections $\tilde{\mathcal{B}}_E$ itself. We do this with a sort of Morse-Bott complex for a smooth G -manifold equipped with an equivariant Morse function; our definition is partly inspired by the Morse-Bott complex introduced for monopole Floer homology in [Lin18]. This uses Lipyanskiy's notion of the *geometric chain complex* $C_*^{\text{gm}}(X; R)$ of a smooth manifold X , introduced in [Lip14], whose homology gives the usual singular homology of X .

While there are technical obstructions to carrying this out for all 3-manifolds, this has the advantage of providing more structure: for (Y, E) a weakly admissible bundle and R a commutative ground ring, we can define a $\mathbb{Z}/8$ -graded chain complex of R -modules, $\tilde{CI}(Y, E, \pi; R)$, which carries the action of the differential graded algebra $C_*(SO(3); R)$. (This is what we find to be the cleanest notion of a chain complex with an action of the Lie group $SO(3)$.) This chain complex depends on further data π , including a metric on the 3-manifold itself and a perturbation of the functional defining the Morse complex, but this turns out to be *mostly* inessential: associated to a perturbation π on a pair (Y, E) is an element of a finite set $\sigma(Y, E)$ of *signature data*; this will be defined in section 3.5. For concreteness, we remark that if Y is a rational homology sphere whose universal abelian cover \tilde{Y} has $H^1(\tilde{Y}; \mathbb{C}) = 0$, then $\sigma(Y, E)$ consists of a single element for all E . The set $\sigma(Y, E)$ corresponds precisely to the "natural classes of perturbations" stated in the main theorem of [AB96].

In fact, the TQFT structure of the usual instanton Floer homology groups can be lifted to the level of the homology groups of \tilde{CI} . This is the main theorem of this paper.

Theorem 1. *There is a category $\text{Cob}_{3,b}^{U(2),w}$ of based 3-manifolds (Y, E, σ, b) equipped with weakly admissible $U(2)$ -bundles and signature data, whose morphisms are certain ‘weakly admissible’ oriented cobordisms $(W, \tilde{\mathbf{E}})$ equipped with a path between the basepoints on the ends. There is also a category of relatively $\mathbb{Z}/8$ -graded R -modules with an action of $H_*(SO(3); R)$, and a functor*

$$\tilde{I} : \text{Cob}_{3,b}^{U(2),w} \rightarrow \text{Mod}_{H_*(SO(3);R)}^{r,\mathbb{Z}/8}.$$

The relatively graded group $\tilde{I}(Y, E, \sigma; R)$ is called the framed instanton homology of the triple (Y, E, σ) . When Y is a rational homology sphere equipped with the trivial bundle, the relative grading may be lifted to an absolute $\mathbb{Z}/8$ -grading.

Note the sudden shift to $U(2)$ bundles in this statement. This is to pin down signs in the definition of the differential and cobordism maps, and nothing else: the majority of this text is written in the context of $SO(3)$ -bundles, and we only make the passage to $U(2)$ -bundles in Section 5.4, in our discussion of orientations on the moduli spaces. In particular, if we had chosen to work over a coefficient ring where $1 = -1$, we may omit discussion of $U(2)$ -bundles entirely. Furthermore, if we only care about the underlying isomorphism class of group $\tilde{I}(Y, E, \pi)$, this only depends on the underlying $SO(3)$ -bundle: see Corollary 6.19.

In the most important cases (including trivial bundles), the weak admissibility condition includes a condition called ρ -monotonicity, which is defined in terms of the Atiyah-Patodi-Singer ρ invariant for flat connections and the signature data σ . Roughly, their sum should always increase across the cobordism. This condition has to do with achieving transversality normal to the reducible locus; when it fails, it is not clear how to try to define the cobordism maps.

For a more precise statement, see Section 3.5 for Definition 3.3 of signature data; the ρ -invariant and ρ -monotonicity condition are introduced at the end of Section 4.5, and Definition 4.14 gives the definition of weakly (and fully) admissible bundles. Finally, Section 6.2 contains Definition 6.10 of the weakly admissible cobordism category (and below it, two relatives) as well as Theorem 6.17 defining the framed instanton functor.

It is our expectation that the notion of signature data above is inconsequential:

Conjecture 1. *If Y is a rational homology sphere, $\tilde{CI}_*(Y, E, \sigma; R)$ is independent of the choice of signature data σ up to graded $C_*(SO(3); R)$ -equivariant quasi-isomorphism.*

From here, if R is a principal ideal domain, we may construct a complex $CI_*^+(Y, E, \sigma; R)$ using the *bar construction* (topologically, the Borel construction) on \tilde{CI} . Its homology groups are denoted $I_*^+(Y, E, \sigma; R)$ and form a relatively graded module under the action of $H^*(BSO(3); R)$. This construction is standard in algebraic topology, and reviewed in the appendix.

One appealing feature of the monopole and Heegaard Floer theories (which are in some sense S^1 -equivariant homology theories) is the existence of two other variants of the Floer homology groups that fit into an exact triangle. (Largely, we choose our notation to fit with that of Heegaard Floer theory.)

In the appendix, we describe algebraic constructions $C^\bullet(A, M)$, where $\bullet = +, -, \infty$, where A is a dg-algebra and M a dg-module. $C^-(A, M)$ is a *cobar construction*, dual to the above bar construction, and $C^\infty(A, M)$ is the *Tate complex*,

constructed as a comparison between a variant C^+ and C^- . Thus up to a sort of twist (often only a grading shift), we have an exact sequence

$$0 \rightarrow C^+(A, M) \rightarrow C^-(A, M) \rightarrow C^\infty(A, M) \rightarrow 0,$$

leading to an exact triangle of homology groups. The Tate homology groups $H^\infty(A, M)$ satisfy a short list of axioms, including that $H^\infty(A, A) = 0$.

Applying a version of this for $\mathbb{Z}/8$ -graded complexes to the framed instanton complex \widetilde{CI} as a module over $C_*(SO(3); R)$, and taking homology, we arrive at the equivariant instanton homology groups

$$I^+(Y, E, \sigma; R), \quad I^-(Y, E, \sigma; R), \quad I^\infty(Y, E, \sigma; R).$$

As desired, these are all modules over $H^{-*}(BSO(3); R)$, and fit into an exact triangle.

Theorem 2. *There are functors from the cobordism category of pointed 3-manifolds equipped with weakly admissible $U(2)$ bundles and signature data to the category of relatively $\mathbb{Z}/8$ -graded $H^{-*}(BSO(3); R)$ -modules,*

$$I_*^\bullet : \text{Cob}_{3,b}^{U(2),w} \rightarrow \text{Mod}_{H^{-*}(BSO(3);R)}^{r,\mathbb{Z}/8}$$

$\bullet = +, -, \infty$. *When Y is a rational homology sphere, $I^\bullet(Y, E, \sigma; R)$'s relative grading lifts to an absolute $\mathbb{Z}/8$ -grading. Furthermore, there is a long exact sequence*

$$\dots \rightarrow I^+ \xrightarrow{[3]} I^- \rightarrow I^\infty \xrightarrow{[-4]} I^+ \rightarrow \dots$$

where $[n]$ denotes that the map increases grading by n .¹

The main tool we use to establish invariance properties and perform calculations and compare with existing theories is a collection of spectral sequences that calculate these equivariant homology groups.

Theorem 3. *Let (Y, E, σ) be a closed oriented 3-manifold equipped with weakly admissible $U(2)$ -bundle and signature datum. Supposing π is a regular perturbation for (Y, E, σ) , we denote by \mathfrak{C}_π the set of critical $SO(3)$ -orbits.*

Then there are $\mathbb{Z}/8 \times \mathbb{Z}$ -graded spectral sequences

$$\bigoplus_{\alpha \in \mathfrak{C}_\pi} H_{SO(3)}^\bullet(\alpha; R) \rightarrow I^\bullet(Y, E, \sigma; R).$$

These spectral sequences converge strongly for \widetilde{I}, I^+ , and I^- , meaning they can be used to calculate the homology groups, and converge weakly for I^∞ , meaning that we can detect quasi-isomorphisms as isomorphisms on some finite page E^r .

This is proved in two parts, Theorems 6.20 and 6.23, first for nonequivariant homology and then for equivariant homology.

Using the fact that Tate homology of a free A -module vanishes, the groups $I^\infty(Y, E)$ are especially computable. For an admissible bundle, they vanish, and we can compare I^+ to Floer's invariant $I(Y, E; R)$ for admissible bundles.

Theorem 4. *If E is an admissible bundle over a 3-manifold Y , then*

$$I^\infty(Y, E; R) = 0, \quad \text{and} \quad I^+(Y, E; R) \cong I(Y, E; R).$$

¹Note that the grading shift is only meaningful when Y is a rational homology sphere, for which the grading is absolute.

If we also have $\frac{1}{2} \in R$, the action of $H^*(BSO(3); R) = R[U]$ on $I^+(Y, E; R)$ is taken by the isomorphism to the standard Floer-theroetic U -map.

The isomorphisms are given in Theorem 7.1 and Corollary 7.9.

Let $\frac{1}{2} \in R$. In [Don02], Donaldson introduced three chain complexes for Y an integer homology sphere equipped with the trivial $SU(2)$ -bundle: the framed complex $\widetilde{CF}(Y; R)$, the equivariant homology complex $\overline{CF}(Y; R)$, and the equivariant cohomology complex $\underline{CF}(Y; R)$. The first complex has an action of the exterior algebra $\Lambda(u)$, with $|u| = 3$, and the second two complexes have an action of $R[U]$, where $|U| = -4$. In Section.2, we prove the following.

Theorem 5. *For (Y, E) a 3-manifold equipped with a weakly admissible bundle, there is a finite dimensional $\Lambda(u)$ -module $DCI(Y, E; R)$ and finite type $R[U]$ -modules $\overline{DCI}^\pm(Y, E; R)$, so that there are equivariant quasi-isomorphisms*

$$\begin{aligned} DCI(Y, E; R) &\simeq \widetilde{CI}(Y, E; R) \\ \overline{DCI}^+(Y, E; R) &\simeq CI^+(Y, E; R) \\ \overline{DCI}^-(Y, E; R) &\simeq CI^-(Y, E; R) \end{aligned}$$

Furthermore, when Y is an integer homology sphere, we have equalities

$$\begin{aligned} DCI(Y, E; R) &= \widetilde{CF}(Y, E; R) \\ \overline{DCI}^+ &= \overline{CF}(Y, E; R) \\ \overline{DCI}^- &= \underline{CF}(Y, E; R), \end{aligned}$$

up to a rescaling of basis.

Here finite type means that it is a direct sum of finitely many simple pieces: for \overline{DCI}^+ , they are R with trivial U -action and $R[[U^*]]$, where U contracts against U^* ; for \overline{DCI}^- , they are R with trivial U -action and $R[U]$ with canonical U -action.

It is by passing through this isomorphism that we show that the isomorphisms Theorem 4 above preserve the U -action; this is Corollary 7.7.

While more complicated than the Tate calculations for admissible bundles, we are able to exploit the isomorphisms above to calculate instanton Tate homology for an arbitrary rational homology 3-sphere in Section.3.

Theorem 6. *Let (Y, E, σ) be a rational homology 3-sphere equipped with $U(2)$ -bundle and signature datum, and suppose R is a PID in which 2 is invertible.*

We write $R[H^2(Y)]$ to mean the group algebra of the finite group $H^2(Y; \mathbb{Z}) \cong H_1(Y; \mathbb{Z})$. If $c = c_1 E \in H^2(Y)$, there is an action of $\mathbb{Z}/2$ on $H^2(Y)$ given by the involution $x \mapsto c - x$. Further define an action of $\mathbb{Z}/2$ on the ring $R[U^{1/2}, U^{-1/2}]$, acting on the basis by $U^{n/2} \mapsto (-1)^n U^{n/2}$. Here $|U^{1/2}| = -2$.

Then there is a canonical isomorphism of $\mathbb{Z}/8$ -graded modules over $R[U]$,

$$I^\infty(Y, E, \sigma; R) \cong R[U^{1/2}, U^{-1/2}] \otimes_{R[\mathbb{Z}/2]} R[H^2(Y)].$$

Remark 1. In [AB96], an equivariant instanton Floer homology is associated to any 3-manifold with $b_1(Y) = 0$ or $H_1(Y; \mathbb{Z}) \cong \mathbb{Z}^a \oplus (\mathbb{Z}/2)^b$ equipped with the trivial $SO(3)$ bundle, though their analysis is unchanged in the case of $b_1(Y) = 0$ and nontrivial $SO(3)$ bundle.

In the case that (Y, E) supports an $SO(2)$ -reducible flat connection — that is, if

$$\langle w_2(E), H_2(Y; \mathbb{Z}) \bmod 2 \rangle, \text{ and } 2\text{Tors}(H_1(Y; \mathbb{Z})) \neq 0,$$

there is a topological obstruction to equivariant transversality.

When this is the case, the unperturbed moduli space of $SO(2)$ -reducible flat connections is a disjoint union of tori $T^{b_1(Y)}$. After applying an $SO(3)$ -invariant perturbation, the space of $SO(2)$ -invariant flowlines between various reducible critical points lying in the same T^{b_1} are smooth manifolds of dimension at most $b_1(Y) - 1$; they have a map to a space of Fredholm operators, given by sending a reducible ASD connection splitting as $\mathbf{A} \cong \theta \oplus A$ to the ASD operator corresponding to A ; in this specific case, this operator has index 0. This map must be made disjoint from the loci of operators with cokernel of dimension at least 1; this is a subspace of codimension 2. Thus when $b_1(Y) \geq 3$, we may have some non-trivial intersection with this locus which cannot be removed by a small perturbation. A similar argument shows that there is a further obstruction to achieving invariance than those we have already identified.

Austin and Braam's invariant in the case that $b_1(Y) > 0$ is defined by careful (non-generic) choices of equivariant perturbation. We expect that the instanton complex defined in this paper can be extended to that level of generality, as well as the proof of its invariance up to signature data. However, the analysis is somewhat more delicate, and the choice of perturbation is necessarily non-generic. The cases with $b_1(Y) > 0$ also add some difficulty in defining a cobordism category that instanton homology is functorial on, as many cobordisms for which the cobordism maps are naturally and essentially uniquely defined do not necessarily compose to cobordisms in the same class.

Survey of the homology theories. As a number of different chain complexes are introduced in this paper and exist in the literature, we survey here the definitions and relationships.

Here (Y, E) is a closed oriented 3-manifold with weakly admissible $U(2)$ -bundle, π is a regular perturbation and R is a PID. This means, in particular, that there are finitely many critical orbits of the perturbed Chern-Simons functional on $\tilde{\mathcal{B}}_E^e$. The finite set whose elements are connected components of the critical set is written \mathfrak{C}_π ; an element $\alpha \in \mathfrak{C}_\pi$ is an $SO(3)$ -space, either a point, S^2 , or $SO(3)$ itself. If E is equipped with a trivialization, each $\alpha \in \mathfrak{C}_\pi$ has an associated grading $i(\alpha) \in \mathbb{Z}/8$; otherwise, we instead have relative gradings $i(\alpha, \beta) \in \mathbb{Z}/8$. For uniformity of notation we often write the relative grading as if it were given by an absolute grading.

The first chain complex introduced is $\widetilde{CI}(Y, E, \pi; R)$. As an R -module, this is given by

$$\bigoplus_{\alpha \in \mathfrak{C}_\pi} C_*^{\text{gm}}(\alpha; R)[i(\alpha)].$$

The individual terms $C_*^{\text{gm}}(\alpha; R)$ are the *geometric chain complexes* of the orbits α described in Section 6.1. A generic basis element of $C_*^{\text{gm}}(\alpha)$ is given by a “strong δ -chain”, but it is helpful to imagine that a generator of $C_*^{\text{gm}}(\alpha; R)$ is given by a smooth map $\sigma : P \rightarrow \alpha$, where P is a compact smooth manifold with corners, and σ is considered up to diffeomorphism of the domain. (This is essentially a special case of the more general notion of δ -chain.)

The boundary operator on $\widetilde{CI}(Y, E, \pi; R)$ is given as the sum of the geometric boundary operator (sending $\sigma : P \rightarrow \alpha$ to $\partial P \rightarrow \alpha$) and a fiber product map with the moduli spaces whose properties are discussed in the first half of this paper, taking a map $\sigma : P \rightarrow \alpha$ to a map

$$(*) \quad P \times_{\alpha} \overline{\mathcal{M}}(\alpha, \beta) \rightarrow \beta.$$

This chain complex carries the action of a $C_*^{\text{gm}}(SO(3); R)$ -module.

After this, one applies the results of the appendix to construct chain complexes

$$CI^+(Y, E, \pi; R), \quad CI^-(Y, E, \pi; R), \quad CI^{\infty}(Y, E, \pi; R).$$

These are constructed algebraically, and we will not give detailed descriptions here of a generic element of these chain complexes. Writing \bullet for one of $\{+, -, \infty\}$, suffice it to say that if G is a connected Lie group and H a connected subgroup, then for $\alpha = G/H$, there are chain complexes $\hat{C}_G^{\bullet}(\alpha; R)$ (the completed group homology complexes) so that $H_G^+(\alpha)$ computes group homology of H , while $H_G^-(\alpha)$ computes group cohomology of H in negative degrees and $H_G^{\infty}(H)$ computes Tate homology of H . These all having some degree shift and an action of $H^{-*}(BG; R)$ induced by restriction to cohomology of BH . Further,

$$CI^{\bullet}(Y, E, \pi; R) = \bigoplus_{\alpha \in \mathfrak{C}_{\pi}} \hat{C}_{SO(3)}^{\bullet}(\alpha; R).$$

Next when $\frac{1}{2} \in R$ we define a finite-dimensional chain complex $DCI(Y, E, \pi; R)$, given as an R -module by

$$DCI(Y, E, \pi; R) = \bigoplus_{\alpha \in \mathfrak{C}_{\pi}} H_*(\alpha; R)[i(\alpha)],$$

with an action of $H_*(SO(3); R)$; this is now an exterior algebra on a single degree-3 generator, the fundamental class of $SO(3)$. One may think of this as contributing, for each irreducible orbit α , a copy of R in degree $i(\alpha)$ and $i(\alpha) + 3$ with the action of $H_*(SO(3); R)$ taking the first to the second, for each $SO(2)$ -reducible a copy of R in degrees $i(\alpha)$ and $i(\alpha) + 2$, and for each full reducible a copy of R in degree $i(\alpha)$. The differential counts 0-dimensional moduli spaces between irreducibles, as well as the degrees of maps of moduli spaces between different orbits. One of these maps is often called the U -map, which we write U_{FI} to distinguish from later algebraic terms, also written U .

There is a $C_*(SO(3); R)$ -equivariant quasi-isomorphism

$$DCI(Y, E, \pi; R) \simeq \widetilde{CI}(Y, E, \pi; R).$$

When Y is an integer homology sphere and E is the trivial bundle, $DCI(Y, \pi; R)$ is isomorphic to the complex Donaldson writes as $\widetilde{HF}(Y, \pi; R)$ in [Don02, Section 7.3.3]; though Donaldson only writes his for rational coefficients, there is no difficulty extending the definition to the broader case $\frac{1}{2} \in R$.

One may immediately apply the constructions of the appendix to $DCI(Y, E, \pi; R)$. Applying the plus-homology construction we arrive at $DCI^+(Y, E, \pi; R)$, and applying the minus-homology construction we arrive at $DCI^-(Y, E, \pi; R)$. If U is a degree -4 element and U^* is a degree 4 element, the underlying R -modules are given by

$$DCI^+(Y, E, \pi; R) = DCI(Y, E, \pi; R)[[U^*]]$$

and

$$DCI^-(Y, E, \pi; R) = DCI(Y, E, \pi; R)[U].$$

One thinks of each irreducible here as contributing what looks like a copy of $ESO(3)$.

Applying a trick of Seidel and Smith from [SS10], we pass from these to complexes \overline{DCI}^+ and \overline{DCI}^- . The first has underlying R -module given by

$$\bigoplus_{\alpha \in \mathfrak{C}_\pi} H_*^{SO(3)}(\alpha; R)[i(\alpha)]$$

and the second by

$$\bigoplus_{\alpha \in \mathfrak{C}_\pi} H_{SO(3)}^{\dim \alpha - *}(\alpha; R)[i(\alpha)].$$

That is, each irreducible contributes a copy of R (in degree $i(\alpha)$ or $i(\alpha) + 3$), while each $SO(2)$ -reducible contributes a tower $R[[U^{*/2}]]$ or $R[U^{1/2}][2]$ respectively, where $|U^{*/2}| = 2$ and $|U^{1/2}| = -2$. Lastly, each full reducible contributes a copy of $R[[U^*]]$ or $R[U]$, respectively. The differentials involve large powers of U_{F1} .

We have $C^{-*}(BSO(3); R)$ -equivariant quasi-isomorphisms

$$CI^+(Y, E, \pi; R) \simeq DCI^+(Y, E, \pi; R) \simeq \overline{DCI}^+(Y, E, \pi; R)$$

and

$$CI^-(Y, E, \pi; R) \simeq DCI^-(Y, E, \pi; R) \simeq \overline{DCI}^-(Y, E, \pi; R).$$

For an integer homology sphere Y equipped with the trivial bundle E , the complex $\overline{DCI}^+(Y, E, \pi; R)$ is isomorphic as a U -module to the complex Donaldson writes as $\overline{CF}(Y; R)$, and similarly $\overline{DCI}^-(Y, E, \pi; R) \cong \underline{CF}(Y; R)$.

Donaldson also defines complexes $\overline{CF}(Y; R)$ and $\underline{CF}(Y; R)$. These are best understood as the quotient and fixed points of the chain complex $DCI(Y, \pi; R) \cong \overline{CF}(Y, \pi; R)$ under the $H_*(SO(3); R)$ -action. We do not use these, but do notice them as appearing in certain spectral sequences for integer homology spheres. In that context, we write the resulting homology groups as \overline{I} and \underline{I} .

There are also Froyshov's reduced instanton Floer homology groups, which he writes \widetilde{HF} and Donaldson writes \widehat{HF} . We identify these in Section 7.4 as being the image of $I^+(Y; R)$ inside $I^-(Y; R)$ (or more precisely, we see this at the level of an E^∞ page of a spectral sequence calculating those), but otherwise do not use them. If we were to give these groups a notation, we could call them \widehat{I} .

There is also Floer's original version of instanton homology, $I(Y)$, which is defined only for integer homology spheres and does not use the reducibles. This, too, may be seen in terms of the spectral sequence for I^+ (ignoring the reducible piece).

Finally, Kronheimer and Mrowka define instanton homology groups $I^\#(Y, E; R)$ for all pairs of 3-manifolds and $SO(3)$ -bundles, by taking the connected sum with a pair (T^3, E) where $w_2(E)$ is Poincaré dual to $T^2 \times \{*\}$, also called 'framed instanton homology'. These are not the same as the groups $\widetilde{I}(Y, E; R)$. A calculation of [Sca15] using Fukaya's connected sum theorem shows that when Y is an integer homology sphere, the group $I^\#(Y; R)$ may be calculated using a chain complex very much like that defining $DCI(Y, \pi; R)$, but they differ in one component of the matrix defining the differential: the term U_{F1} in ∂_{DCI} is instead given by $U_{F1} - 8$ in the complex defining $I^\#(Y; R)$. So one finds instead that $I^\#(Y, E; R)$ is a sort of deformation of $\widetilde{I}(Y, E; R)$. We do not discuss this relationship further here.

In summary, there are nine versions of instanton homology for integer homology spheres,

$$\tilde{I}, I^+, I^-, I^\infty, \bar{I}, \underline{I}, \hat{I}, I^\#, I;$$

all but the last of these are also defined for pairs of a rational homology sphere and $U(2)$ -bundle. In this paper we will work most extensively with the first four of these. We will briefly mention the next three ‘reduced’ homology theories \bar{I} , \underline{I} , and \hat{I} when discussing spectral sequence calculations for integer homology spheres. We will not mention $I^\#$ and I any further.

In addition to the complexes defining each of these homology theories, we will also use five important complexes for calculation, the Donaldson models

$$DCI, DCI^+, \overline{DCI}^+, DCI^-, \overline{DCI}^-.$$

The first of these has homology naturally isomorphic to \tilde{I} , the second two of these have homology naturally isomorphic to I^+ , and the last two of these have homology naturally isomorphic to I^- .

Organization. Section 1 introduces the framed configuration spaces on which we attempt to do Morse theory, and discusses the action of $SO(3)$ on them. Section 2 explains how to complete these topological spaces using Sobolev spaces and obtain the structure of a Banach manifold.

The technical heart of the paper is in Sections 3 and 4, where we use the standard holonomy perturbations in instanton Floer theory to show that we can achieve equivariant transversality: a generic perturbation gives rise to a finite set of critical $SO(3)$ -orbits, and the space of trajectories between them forms a smooth manifold. As long as the dimension is sufficiently small, this has a compactification to a topological manifold with corners, and we use these compactifications to define the complex \tilde{CI} . This includes a calculation of the reducible perturbed instantons in Section 4.4 and their indices in Section 4.5, which is then used to give a condition which guarantees we can achieve transversality normal to the reducible locus. This condition uses the notion of signature data introduced in Section 3.5.

Sections 1-5 provide us with the technical machine (the moduli spaces and their properties) we need. Sections 6 and 7 are devoted to defining and calculating equivariant instanton homology given this machine, and may be read independently of the first part of the paper.

In Section 6, we define the invariants CI^\bullet . First we review Lipyanskiy’s geometric homology, which is a crucial technical tool in our definition of \tilde{CI} : it allows us to define the chain complex without triangulating the moduli spaces appearing in (*), and only requires the use of moduli spaces of small dimension; in particular, small enough that the Uhlenbeck bubbling phenomenon does not arise. The chain complex $\tilde{CI}(Y, E, \pi; R)$, which depends on a choice of metric and perturbation, is defined in Section 6.2. This construction comes with cobordism maps, and these provide us with the usual invariance properties, giving Theorem 1.

In Section 6.3 we explain the notion of index filtration for \tilde{CI} , which gives rise to a $\mathbb{Z}/8 \times \mathbb{Z}$ -graded spectral sequence. This is not a filtration on \tilde{CI} in the standard sense, but rather what we call a *periodic filtration*. In Section 6.4 we use the idea of periodic filtrations to define the equivariant instanton homology complexes CI^\bullet , following the construction of Section A.8.

In Section 7, we use the index spectral sequences of sections 6.3 and 6.4 to carry out some calculations and comparisons.

In Section 7.1 we warm up with a calculation of $I^\bullet(Y, E)$ for admissible bundles, giving Theorem 4 above.

After this, in Section 7.2 we compare \widetilde{CI} to Donaldson's complex DCI (written in [Don02] as \widetilde{CF}), and in particular show that their homologies are isomorphic, justifying the notation $H(\widetilde{CI}) = \widetilde{I}$. We are then able to show $I^+(Y; R) \cong \overline{HF}(Y; R)$ for rings containing $1/2$, as well as $I^-(Y; R) \cong \underline{HF}(Y; R)$. In the same section, we extend the definition of Donaldson's chain complexes to rational homology spheres.

Section 7.3 uses the Donaldson model, and a localization/periodicity theorem for Tate homology, to give the explicit formula for I^∞ mentioned as Theorem 6 above. Section 7.4 contains a calculation of I^\bullet for the 3-sphere, lens spaces, and the Poincaré homology sphere, as well as a description of the spectral sequences for integer homology spheres and a brief comment on the connection to Frøyshov's reduced groups. We conclude in Section 7.5 with a discussion of orientation-reversal and duality in $CI^\bullet(Y)$.

The appendix describes the algebraic constructions $C^\bullet(A, M)$ and their invariance properties, as well as providing calculational tools. In particular, if A is the group algebra $C_*(G; R)$, where G is a compact Lie group, then $C^\bullet(G, G/H)$ is calculated for all \bullet and any closed subgroup $H \subset G$. Section A.8 concludes with an extension of this machinery to the $\mathbb{Z}/8$ -graded case, when our complexes come equipped with well-behaved periodic filtrations.

Acknowledgements. I would like to thank my advisor, Ciprian Manolescu, for his constant advice and support. I also thank Ali Daemi, Kim Frøyshov, Tom Mrowka, Chris Scaduto, and Matt Stoffregen for useful discussions about gauge theory, and Sucharit Sarkar for conversations on equivariant homology. Furthermore, I would also like to thank Kevin Carlson, Tyler Lawson, and Mike Hill, for helpful conversations during the preparation of the appendix, and especially Aaron Royer for many patient discussions. The author was partially supported by NSF grant number DMS-1708320.

1. CONFIGURATION SPACES AND THEIR REDUCIBLES

Let Y be an oriented closed 3-manifold equipped with a Riemannian metric, and $E \rightarrow Y$ an $SO(3)$ -bundle over Y ; it is worth mentioning that $SO(3)$ -bundles over a 3-complex are determined up to isomorphism by their second Stiefel-Whitney class $w_2(E) \in H^2(Y; \mathbb{Z}/2)$. Associated to E by the adjoint representation $SO(3) \curvearrowright \mathfrak{so}(3)$ is the adjoint bundle $\mathfrak{g}_E \subset \text{End}(E)$, the subbundle given by skew-adjoint endomorphisms. The space \mathcal{A}_E of orthogonal connections on E is affine over $\Omega^1(\mathfrak{g}_E)$. There is a gauge group, \mathcal{G}_E , the set of smooth bundle automorphisms of E that cover the identity; equivalently, this is the set of smooth sections of the non-principal $SO(3)$ bundle $\text{Aut}(E)$, the associated bundle to E under the conjugation action $SO(3) \curvearrowright SO(3)$. Using the isomorphism $\text{Inn}(SU(2)) \cong SO(3)$, we may form the bundle

$$\widetilde{\text{Aut}}(E) = \text{Aut}(E) \times_{SO(3)} SU(2).$$

We say that $\sigma \in \mathcal{G}_E$ is an *even gauge transformation* if σ lifts to a section of $\widetilde{\text{Aut}}(E)$, and denote the group of even gauge transformations by \mathcal{G}_E^e . Obstruction theory applied to sections of $\text{Aut}(E)$ provides a short exact sequence of groups

$$1 \rightarrow \mathcal{G}_E^e \rightarrow \mathcal{G}_E \rightarrow H^1(Y; \mathbb{Z}/2) \rightarrow 0.$$

There is a map, the *Chern-Simons functional*,

$$\text{cs} : \mathcal{A}_E \rightarrow \mathbb{R},$$

defined as follows. Pick a compact oriented 4-manifold X , equipped with an $SO(3)$ bundle E_X , such that $\partial(X, E_X) = (Y, E)$. Given a connection A on E , extend it to a connection A_X on E_X . Then define

$$\text{cs}(A) = \int_X \text{Tr}(F_{A_X}^2).$$

This gives a function that depends only on the choice (X, E_X) , not on the extension A_X . If $(X', E_{X'}, A_{X'})$ is another extension, we may define $M = X \cup_Y \overline{X'}$ with the obvious choice of $SO(3)$ -bundle and connection over it and invoke the Chern-Weil formula

$$-2\pi^2 p_1(E') = \int_M \text{Tr}(F_{A'}^2).$$

Thus cs is defined in general up to an $8\pi^2\mathbb{Z}$ ambiguity (using that $4 \mid p_1(E')$). If $(X', E_{X'}) = (X, E_X)$, then $p_1(E_X \cup E_{X'}) = 0$, giving us a well-defined functional cs on \mathcal{A}_E conditional on that choice of bounding 4-manifold (X, E_X) .

The gauge group \mathcal{G}_E acts on \mathcal{A}_E by $\sigma(A) = A - (\nabla_A \sigma)\sigma^{-1}$, where we take the covariant derivative of σ by considering it as a section of $\text{End}(E) = E \otimes E^*$, where E is considered now as an oriented vector bundle equipped with metric. We denote the quotient by this group action

$$\mathcal{B}_E = \mathcal{A}_E / \mathcal{G}_E,$$

the configuration space of connections on E , and also write the *even configuration space* as the quotient

$$\mathcal{B}_E^e := \mathcal{A}_E / \mathcal{G}_E^e.$$

Immediately from the definition we see that the stabilizer at A by the action of \mathcal{G}_E or \mathcal{G}_E^e is precisely the subset of (even) A -parallel gauge transformations; thus elements of the stabilizer are determined by their value at a single point, and evaluating at a point $b \in Y$ gives an isomorphism to a subgroup of $SO(3) = \text{Aut}(E_b)$.

The Chern-Simons functional $\text{cs} : \mathcal{A}_E \rightarrow \mathbb{R}$ does not descend naively to \mathcal{B}_E — its value may change after applying an element of gauge group \mathcal{G}_E . However, if u is a gauge transformation,

$$\text{cs}(u(A)) = \text{cs}(A) + 8\pi^2 k$$

for some integer k , so cs descends to a continuous map $\mathcal{B}_E \rightarrow \mathbb{R}/8\pi^2\mathbb{Z}$. This statement is little more than saying that there is always a 4-manifold with $SO(3)$ -bundle (X, E) over which the gauge transformation u extends, for then one sees that

$$\text{cs}(u(A)) = \int_X \text{tr}(F_{u(A)}^2) = \int_X \text{tr}(F_A^2);$$

that the integrals are equal is Stokes' theorem. It is this circle-valued functional, on a slightly modified space, that we hope to do Morse theory with.

1.1. Framed connections and the framed configuration space. The configuration space \mathcal{B}_E is in no sense a manifold, because \mathcal{G}_E does not act freely on the space of connections. To free up the action of the gauge group on the space of connections so that the quotient by gauge is a manifold (at least heuristically, at this point), we pick a basepoint $b \in Y$ and consider the space of *framed connections*

$$\tilde{\mathcal{A}}_E = \mathcal{A}_E \times E_b.$$

We call a point $p \in E_b$ a framing because it determines an isomorphism $SO(3) \cong E_b$ sending the identity to the point p ; here we are thinking of E as a principal bundle, not a vector bundle. A gauge transformation σ evaluates to $\sigma(b) \in \text{Aut}(E_b)$, which acts on E_b on the left by applying the automorphism. (By definition, to say $f \in \text{Aut}(E_b)$ means $f(pg) = f(p)g$.) This gives us an action of \mathcal{G}_E on the space of framed connections $\tilde{\mathcal{A}}_E$. There is further a natural *right* $SO(3)$ action on $\tilde{\mathcal{A}}_E$ by acting on E_b by translation. If $\sigma \in \mathcal{G}_E^e$, $(A, p) \in \tilde{\mathcal{A}}_E$, and $g \in SO(3)$, we have that

$$\sigma \cdot ((A, p) \cdot g) = (\sigma \cdot (A, p)) \cdot g,$$

because the gauge group acts by automorphisms of the right G -set E_b .

Because any gauge transformation in the stabilizer of A is A -parallel, it is trivial if its value at any point is the identity. As a consequence, \mathcal{G}_E^e acts *freely* on $\tilde{\mathcal{A}}_E$. Its quotient under the gauge group action, denoted $\tilde{\mathcal{B}}_E^e$, retains the right $SO(3)$ -action. We call this the *framed configuration space* of connections on E . The stabilizer of $[A, p] \in \tilde{\mathcal{B}}_E^e$ in $SO(3)$ is the isomorphic image of the stabilizer of $A \in \mathcal{A}_E$ under the action of the even gauge group following evaluation $\mathcal{G}_E^e \rightarrow \text{Aut}(E_b)$ and the natural isomorphism $\text{Aut}(E_b) \cong SO(3)$. There is a map

$$\tilde{\mathcal{B}}_E^e \rightarrow \mathcal{B}_E^e,$$

given by quotienting by the leftover $SO(3)$ action or equivalently by forgetting the framing. Said another way, orbits of the $SO(3)$ -action on $\tilde{\mathcal{B}}_E$ or $\tilde{\mathcal{B}}_E^e$ correspond to gauge equivalence classes of connections, and a point on an orbit $\alpha \subset \tilde{\mathcal{B}}_E$ corresponding to $[A]$ is an equivalence class of framings, with $\alpha \cong E_b/\text{Stab}([A])$.

We will soon define Hilbert manifold completions of these spaces, in the context of which we will find that $\tilde{\mathcal{B}}_E^e$ is a smooth Hilbert manifold with smooth $SO(3)$ -action. These remarks also apply to the full gauge group, with quotient $\tilde{\mathcal{B}}_E = \tilde{\mathcal{A}}_E/\mathcal{G}_E$.

Remark 1.1. One could also define $\tilde{\mathcal{B}}_E^e = \mathcal{A}_E / \tilde{\mathcal{G}}_E^{e,b}$, quotienting by the group of *based* gauge transformations (gauge transformations with $\sigma(b) = \text{Id}$). This works just as well here, and the right $SO(3)$ action corresponds to the inverse of the left $SO(3)$ action given by $\mathcal{G}_E^e / \mathcal{G}_E^{e,b}$. However, we find later discussions of the 4-manifold configuration spaces and the restriction maps to their ends clearer in the language of framed connections, so this is our preferred model.

1.2. The equivalent $U(2)$ model. We would like to understand the reducible subspaces of $\tilde{\mathcal{B}}_E^e$ under the $SO(3)$ action. The $SO(2)$ fixed points are more easily understood if we introduce an auxiliary construction. Pick a principal $U(2)$ -bundle \tilde{E} and an isomorphism

$$\tilde{E} \times_{U(2)} SO(3) \cong E.$$

This is possible because $U(2)$ -bundles on 3-manifolds are classified by their first Chern class (and $SO(3)$ bundles by their second Steifel-Whitney class), so we only need to know that we can pick a lift of w_2 to an integral cohomology class; that this is possible follows from the Bockstein long exact sequence and the fact that $H^3(Y; \mathbb{Z}) \cong \mathbb{Z}$ has no 2-torsion. Thinking of \tilde{E} instead as a complex vector bundle and explicitly identifying the quotient homomorphism $U(2) \rightarrow SO(3)$, the construction $\tilde{E} \times_{U(2)} SO(3) \cong E$ produces the oriented 3-plane bundle $\mathfrak{su}(\tilde{E}) \subset \text{End}(\tilde{E})$ of skew-Hermitian endomorphisms of \tilde{E} as the associated $SO(3)$ -bundle.

Fix a connection A_0 on the determinant complex line bundle $\det(\tilde{E}) := \lambda$ (if $c_1(\tilde{E})$ is finite order, we can choose this to be the flat connection, unique up to gauge transformation). We consider the space of connections with fixed determinant connection

$$\mathcal{A}_{\tilde{E}}^{\text{det}} = \{A \in \mathcal{A}_{\tilde{E}} \mid \text{tr}(A) = A_0 \in \Omega^1(Y; i\mathbb{R})\};$$

this is also affine over $\Omega^1(\mathfrak{g}_E)$. To every connection on \tilde{E} there is an associated connection on E , giving us a map $\mathcal{A}_{\tilde{E}} \rightarrow \mathcal{A}_E$ which is a bijection when restricted to $\mathcal{A}_{\tilde{E}}^{\text{det}}$. To enhance this to a bijection of spaces of framed connections, let

$$\tilde{\mathcal{A}}_E^{\text{det}} = \mathcal{A}_{\tilde{E}}^{\text{det}} \times (\tilde{E}_b \times_{U(2)} SO(3));$$

this carries a natural right action by $SO(3)$. Furthermore, the natural gauge group acting on $\mathcal{A}_{\tilde{E}}^{\text{det}}$ is the set of gauge transformations whose (pointwise) determinant is 1, denoted

$$\mathcal{G}_{\tilde{E}}^{\text{det}} = \Gamma(\text{Aut}(\tilde{E}) \mid \det \gamma = 1).$$

We can further identify this latter bundle of groups as isomorphic to $\widetilde{\text{Aut}}(E)$ and thus there is a surjective homomorphism

$$\mathcal{G}_{\tilde{E}}^{\text{det}} \rightarrow \mathcal{G}_E^e.$$

It is not a bijection: recall that the latter group is defined as a subset of $\mathcal{G}_E = \Gamma(\text{Aut}(E))$. Its kernel is the 2-element set of gauge transformations whose pointwise values are $\pm 1 \in SU(2)$. This subgroup acts trivially on $\tilde{\mathcal{A}}_E^{\text{det}}$.

Now it is easy to verify that the bijection $\tilde{\mathcal{A}}_E^{\text{det}} \rightarrow \tilde{\mathcal{A}}_E$ is equivariant under the actions of the gauge groups, and thus after quotienting we have an equivariant diffeomorphism

$$(1) \quad \tilde{\mathcal{B}}_E^{\text{det}} \xrightarrow{\cong} \tilde{\mathcal{B}}_E^e.$$

The group acting on the former space is $SU(2)$, but ± 1 act trivially, and so passing to the quotient $PSU(2) \cong SO(3)$ we identify these two configuration spaces as $SO(3)$ -spaces. Furthermore, the natural definition of Chern-Simons functional on this first space is sent to our Chern-Simons functional on the second space, and there is no difference in the resulting gauge theory. We only consider this construction auxiliary because it depends on the unnecessary input data of \tilde{E} and A_0 . When working on the $U(2)$ -bundle, we usually prefer to speak of the $SU(2)$ action, as $SU(2)$ naturally sits inside the $U(2)$ gauge group, even though ± 1 act trivially on the framed configuration space.

In the simple case that $\tilde{E} = \eta_1 \oplus \eta_2$, the associated $SO(3)$ -bundle is $i\mathbb{R} \oplus (\eta_1 \otimes \eta_2^{-1})$, and a connection on the former respecting the splitting is taken under the bijection $\mathcal{A}_E^{\det} \rightarrow \mathcal{A}_E$ to a connection which respects the latter splitting.

1.3. Reducibles on 3-manifolds. We can explicitly describe the stabilizers that arise under the action of the even gauge group on \mathcal{A}_E , and hence the orbit types in $\tilde{\mathcal{B}}_E^e$. In understanding these, it's also convenient to carry out the same analysis for the full gauge group \mathcal{G}_E and its quotient $\tilde{\mathcal{B}}_E$. Recall that

$$\mathcal{G}_E/\mathcal{G}_E^e \cong H^1(Y; \mathbb{Z}/2),$$

identifying the latter with the group of obstructions to lifting a section of $\text{Aut}(E)$ to a section of $\widetilde{\text{Aut}}(E)$. This leaves an $H^1(Y; \mathbb{Z}/2)$ action on $\tilde{\mathcal{B}}_E^e$, whose quotient is $\tilde{\mathcal{B}}_E$, the quotient by all gauge transformations. This is a useful gadget to keep track of, especially in light of the following lemma (a version of [DK90, Lemma 4.28]), calculating the stabilizers of connections under both \mathcal{G}_E^e and \mathcal{G}_E .

Lemma 1.1. *Let A be an $SO(3)$ -connection on E . Then the stabilizer Γ_A under the action of the full gauge group \mathcal{G}_E is*

$$C(H_A) \subset SO(3),$$

the centralizer of the holonomy group of A at some choice of basepoint b . Let $\pi : SU(2) \rightarrow SO(3)$ be the projection. $\Gamma_A \cap \mathcal{G}_E^e$ is

$$C_{SU(2)}(\pi^{-1}H_A) \subset SO(3),$$

the set of elements of $SO(3)$ that fix $\pi^{-1}H_A$ under conjugation, considering $SO(3) = \text{Inn}(SU(2))$.

Proof. If $SO(3)$ acts smoothly on some manifold M , consider the M -bundle $E \times_{SO(3)} M$. If A is an $SO(3)$ -connection on E , there is a natural connection induced on the associated M -bundle. Then $m \in M_x$ extends to a parallel section of this M -bundle if and only if it is fixed under the action of H_A on M . Taking M to be $SO(3)$ and $SU(2)$ equipped with the conjugation action gives the desired result. ■

A group-theoretic calculation shows that the only subgroups of $SU(2)$ that arise as centralizers are the center $\mathbb{Z}/2$, the circle subgroups $U(1)$, and the full group $SU(2)$. So the only possible stabilizers of the action of \mathcal{G}_E^e on \mathcal{A}_E are

$$\{I\}, SO(2), SO(3).$$

A similar calculation shows that the subgroups of $SO(3)$ that arise as centralizers are additionally $O(2)$, the diagonal subgroup V_4 , and $\mathbb{Z}/2$. (Calculate centralizers of elements first, then calculate the possible intersections of these.) Comparing stabilizers, we see that the action of $H^1(Y; \mathbb{Z}/2)$ on \mathcal{B}_E^e giving rise to the quotient

by the full gauge group \mathcal{B}_E is free except at reducibles with full stabilizer $O(2)$, V_4 , or $\mathbb{Z}/2$, where the action of $H^1(Y; \mathbb{Z}/2)$ has stabilizer isomorphic to $\mathbb{Z}/2$, V_4 , and $\mathbb{Z}/2$, respectively.

When the stabilizer of a connection A on E is $SO(3) \subset \mathcal{G}_E^e$, the bundle must be trivial and the connection gauge equivalent to the trivial connection; however, it needn't be equivalent by an *even* gauge transformation. In fact, because $H^1(Y; \mathbb{Z}/2)$ acts freely on points in \mathcal{B}_E^e with full stabilizer and there is only one such in the full quotient, there are $H^1(Y; \mathbb{Z}/2)$ different elements of $\tilde{\mathcal{B}}_E^e$ with stabilizer $SO(3)$.

Inside the $U(2)$ -model, [SS17, Section 3] identifies the action of $H^1(Y; \mathbb{Z}/2)$ as sending a connection on \tilde{E} to the corresponding connection on $\tilde{E} \otimes \xi_{\mathbb{C}}$, where $\xi_{\mathbb{C}}$ is the complexification of a real line bundle (equivalently, a complex line bundle equipped with a unitary connection with holonomy in ± 1 .)

We will ultimately be interested in reducible *orbits* of the $SO(3)$ -action. To describe the reducible orbits in a G -space X with orbit type isomorphic to G/H , it suffices to describe the set of H -fixed points X^H , and the action of the *Weyl group*

$$W(H) = N_G(H)/H$$

on X^H : every orbit Gx of a point x with stabilizer conjugate to H intersects X^H nontrivially; and in fact, if $x \in X^H$, then

$$Gx \cap X^H = W(H)x.$$

Thus, for instance, a G -invariant function on the subspace of points whose stabilizer contains a conjugate of H is determined uniquely by a $W(H)$ -invariant function on X^H . The description in terms of fixed subspaces and Weyl groups tends to be easier to state and prove, so we largely prefer that language as long as possible. The only case we actually use is that of $SO(2) \subset SO(3)$, which has Weyl group isomorphic to $\mathbb{Z}/2$.

Proposition 1.2. *Suppose E is an $SO(3)$ -bundle over a closed 3-manifold Y . Denote \tilde{E} a $U(2)$ bundle with $c_1(\tilde{E}) \bmod 2 = w_2(E)$ and write $\lambda = \det \tilde{E}$, as in Section 1.2.*

- (1) *The set of points of $\tilde{\mathcal{B}}_E^e$ fixed by $SO(2) \subset SO(3)$ is identified with*

$$\bigsqcup_{\eta} \mathcal{B}_{\eta} := \bigsqcup_{\eta} \mathcal{A}_{\eta} / \mathcal{G}_{\eta}$$

where η varies over isomorphism classes of complex line bundles, \mathcal{A}_{η} is the configuration space of unitary connections on η , affine over $\Omega^1(Y; i\mathbb{R})$, and the gauge group \mathcal{G}_{η} is the space of sections of $\text{Aut}(\eta)$, which is the same as the space of maps $\text{Map}(Y, S^1)$. The $SO(3)$ orbit of the $SO(2)$ -fixed point space is the set of all reducible connections. The action of the Weyl group sends a class of connection $[A]$ on η to the class of $[A_0 - A]$ on $\lambda \otimes \eta^{-1}$.

- (2) *If E is trivial, $\tilde{\mathcal{B}}_E^e$ has $SO(3)$ fixed point set in bijection with $H^1(Y; \mathbb{Z}/2)$; if E is nontrivial, the even configuration space has no $SO(3)$ fixed points.*

Proof. Following equation (1), it suffices to find the fixed subspaces in $\tilde{\mathcal{B}}_E^{\text{det}}$. Consider the subset of framed connections (\tilde{A}, p) in $\tilde{\mathcal{A}}_E^{\text{det}}$ whose connection term \tilde{A} has stabilizer consisting of gauge transformations with $\sigma(b)$ in the diagonal subgroup

$$S(U(1) \times U(1)) \subset SU(2),$$

using the framing p to specify the isomorphism

$$\text{Aut}(\tilde{E}_b) \cong U(2).$$

All other circle stabilizers in $\mathcal{G}_{\tilde{E}}^{\det}$ are conjugate to this subgroup, and this fixed subspace of the space of framed connections projects to the subset of $SO(2)$ -fixed connections in $\tilde{\mathcal{B}}_{\tilde{E}}^{\det}$. Our assumption implies that \tilde{A} has holonomy contained in the diagonal subgroup

$$U(1) \times U(1) \subset U(2).$$

This leaves invariant a splitting

$$\tilde{E}_b \cong \mathbb{C}^2 \cong \mathbb{C} \oplus \mathbb{C},$$

the first isomorphism given by the framing p , and hence gives a parallel splitting $\tilde{E} \cong \eta \oplus \eta'$. Sending (\tilde{A}, p) to the corresponding connection A' on η gives us the map from this subset to \mathcal{A}_η , and because any gauge transformation of η extends to a determinant-1 gauge transformation of \tilde{E} (act by the inverse in the η' coordinate), this descends to a well-defined map

$$\left(\tilde{\mathcal{B}}_{\tilde{E}}^{\det}\right)^{U(1)} \rightarrow \bigsqcup_{\eta} \mathcal{B}_\eta$$

modulo gauge. Conversely, any connection on η induces a connection on \tilde{E} of the specified form, and any gauge transformation of η is induced by a unique determinant-1 gauge transformation of \tilde{E} . Any framing on η then induces a framing on \tilde{E} , but there is a determinant-1 gauge transformation changing any framing on η to another, so after modding out by gauge the choice of framing on η didn't matter. This gives the stated bijection for those connections whose induced splitting gives a line bundle of topological type η in the first coordinate; but note that on a 3-manifold, a $U(2)$ -bundle is determined by its first Chern class, and so we have an isomorphism

$$\tilde{E} \cong \eta \oplus (\lambda \otimes \eta^{-1}).$$

Thus every complex line bundle η arises in such a splitting.

The action of the Weyl group is to swap the coordinates of the framing; this then swaps the components $[A' \oplus (A_0 - A')]$ of the connection, as stated in the lemma. In particular, it only fixes A' if there is an isomorphism

$$\eta \cong \lambda \otimes \eta^{-1}$$

and this isomorphism sends A' to $A_0 - A'$. This is only possible if λ is twice an integral class, and so $w_2(E) = 0$; but if one such choice is made, all others are affine over $H^1(Y; \mathbb{Z}/2)$, tensoring the whole bundle (and hence each component) with $\xi_{\mathbb{C}}$, the complexification of a real line bundle, or equivalently a complex line bundle with holonomy in ± 1 . \blacksquare

Corollary 1.3. *Let E and \tilde{E} be as in the previous proposition; write $c = c_1(\lambda) = c_1(\tilde{E})$. Then the reducible subspace of $\tilde{\mathcal{B}}_{\tilde{E}}^e$, consisting of framed connections with nontrivial stabilizer, is a disjoint union over connected components labeled by pairs $\{z_1, z_2\} \subset H^2(Y; \mathbb{Z})$ with $z_1 + z_2 = c$, where the z_i are cohomology classes corresponding to complex line bundles η_i .*

If $z_1 \neq z_2$, this connected component is a fiber bundle over $\mathcal{B}_{\eta_1} \cong \mathcal{B}_{\eta_2}$ with fiber S^2 , where $SO(3)$ acts trivially on the base and via the standard action on the fiber. If $z_1 = z_2$, this connected component contains a unique fully reducible connection.

Proof. Consider the subspace of framed reducibles corresponding to the splitting $\tilde{E} \cong \eta_1 \oplus \eta_2$; we write this space as $\tilde{\mathcal{B}}_{\eta_1, \eta_2}^{\text{red}}$. The quotient of this space by the $SO(3)$ -action is the same as the quotient of its $SO(2)$ -fixed subspace by the action of $\mathbb{Z}/2$. In the case that $z_1 \neq z_2$, the $SO(2)$ -fixed subspace is $\mathcal{B}_{\eta_1} \sqcup \mathcal{B}_{\eta_2}$, and the action of the Weyl group identifies these. Thus the desired fiber bundle is the quotient map

$$\tilde{\mathcal{B}}_{\eta_1, \eta_2}^{\text{red}} \rightarrow \mathcal{B}_{\eta_1} \cong \mathcal{B}_{\eta_2}.$$

When $z_1 = z_2 =: z$, corresponding to a complex line bundle η , the $SO(2)$ -fixed subspace is \mathcal{B}_η . Its quotient by the action of the Weyl group is connected, so $\tilde{\mathcal{B}}_{\eta_1, \eta_2}^{\text{red}}$ is connected. In this case, E is trivial, so we may choose \tilde{E} trivial for convenience of discussion; then the action of the Weyl group may be described as complex conjugation, using the isomorphism $\eta \cong \xi_{\mathbb{C}}$ for a unique real line bundle ξ ; the induced connection is the fully reducible connection. \blacksquare

Remark 1.2. The full group of gauge transformations preserves the reducible set, so $\mathcal{G}_E/\mathcal{G}_E^e \cong H^1(Y; \mathbb{Z}/2)$ acts on $\text{Red}(Y, E)$. If $\beta : H^1(Y; \mathbb{Z}/2) \rightarrow H^2(Y; \mathbb{Z})$ is the Bockstein homomorphism, then in the notation of the above corollary, the action on $\text{Red}(Y, E)$ is given by sending

$$x \cdot \{z_1, z_2\} \mapsto \{z_1 + \beta x, z_2 + \beta x\}.$$

In particular, this action is free and transitive on $\text{Red}_*(Y, E)$.

If x stabilizes $\{z_1, z_2\}$, then $z_2 = z_1 + \beta x$, and we see that $c = 2z_1 + \beta x$, and we see that we may rechoose c to be βx ; then the above corresponds to the reduction

$$\mathbf{E} \cong \xi_{\mathbb{C}} \oplus \mathbb{R} \cong \xi \oplus \xi \oplus \mathbb{R},$$

where ξ is the real line bundle corresponding to x .

The discussion after Lemma 1.1 shows that the $H^1(Y; \mathbb{Z}/2)$ action has stabilizer equal to $\mathbb{Z}/2$ at an $O(2)$ connection, and otherwise acts freely on the reducible set. So in fact we see that $H^1(Y; \mathbb{Z}/2)$ acts freely on the set of reducible components that do not contain an $O(2)$ -connection, and has stabilizer $\mathbb{Z}/2$ on those that do.

We summarize the content of this section as notation:

Definition 1.1. *Let Y be a closed oriented 3-manifold and E an $SO(3)$ -bundle over Y . We write $\text{Red}(Y, E)$ for the set of connected components of the reducible subspace of $\tilde{\mathcal{B}}_E^e$. This may be written as*

$$\text{Red}(Y, E) = \text{Red}_*(Y, E) \sqcup \text{Red}_{SO(2)}(Y, E),$$

where the first term refers to those components containing a fully reducible orbit, and the second refers to those components entirely consisting of framed connections whose stabilizers are conjugate to $SO(2)$.

If we fix a choice of $c \in H^2(Y; \mathbb{Z})$ that reduces mod 2 to $w_2(E)$, we are furnished with a bijection between $\text{Red}(Y, E)$ and the set of unordered pairs $\{z_1, z_2\} \subset H^2(Y; \mathbb{Z})$ with $z_1 + z_2 = c$. The set $\text{Red}_*(Y, E)$ is sent to the 1-element sets², and the set $\text{Red}_{SO(2)}(Y, E)$ is sent to the 2-element sets $\{z_1, z_2\}$ with $z_1 \neq z_2$.

We write elements of $\text{Red}(Y, E)$ as $\{z_1, z_2\}$, making the choice of c implicit.

In the case that $b_1(Y) = 0$, we determine the cardinality of these sets as a simple exercise in algebraic topology.

²that is, those pairs with $z_1 = z_2$; equivalently the set of z with $2z = c$

Proposition 1.4. *Let Y be a closed oriented 3-manifold with $b_1(Y) = 0$, and E an $SO(3)$ -bundle over it. We may determine the number of reducible components as follows.*

- If E is trivial, then

$$\begin{aligned} |\mathrm{Red}_*(Y, E)| &= |H^1(Y; \mathbb{Z}/2)| \\ |\mathrm{Red}_{SO(2)}(Y, E)| &= (|H^2(Y; \mathbb{Z})| - |H^1(Y; \mathbb{Z}/2)|) / 2. \end{aligned}$$

- If E is nontrivial, $\mathrm{Red}_*(Y, E)$ is empty, and

$$|\mathrm{Red}_{SO(2)}(Y, E)| = |H^2(Y; \mathbb{Z})| / 2.$$

Proof. If E is trivial, we choose the integral lift λ of $w_2(E) = 0$ to be 0. Then $\mathrm{Red}_*(Y, E)$ is in bijection with 2-torsion elements of $H^2(Y; \mathbb{Z})$. The Bockstein exact sequence

$$0 = H^1(Y; \mathbb{Z}) \xrightarrow{\text{mod } 2} H^1(Y; \mathbb{Z}/2) \xrightarrow{\beta} H^2(Y; \mathbb{Z}) \xrightarrow{\times 2} H^2(Y; \mathbb{Z})$$

shows that the set of 2-torsion elements is in bijection with $H^1(Y; \mathbb{Z}/2)$, the bijection sending $x \in H^1(Y; \mathbb{Z}/2)$ to $\beta x \in H^2(Y; \mathbb{Z})$. The set $\mathrm{Red}_{SO(2)}(Y, E)$ is in bijection with the set of pairs

$$\{z, -z\} \in H^2(Y; \mathbb{Z})$$

for z not 2-torsion. This proves both equalities.

If E is nontrivial, and c is an integral lift of $w_2 E$, then a set $\{z_1, z_2\}$ with $z_1 + z_2 = c$ cannot be a singleton: if $c = 2z_1$, then c reduces to 0 mod 2, and thus $w_2 E = 0$, so $\mathrm{Red}_*(Y, E)$ is empty. Thus every pair (z_1, z_2) with $z_1 + z_2 = c$ consists of distinct elements, and the $\mathbb{Z}/2$ action swapping z_1 and z_2 has no fixed points. Because a pair (z_1, z_2) is determined completely by $z_1 \in H^2(Y; \mathbb{Z})$, the desired equality is proven.

As a sanity check, observe that $|H^2(Y; \mathbb{Z})|$ is divisible by $|H^1(Y; \mathbb{Z}/2)|$, which is a power of 2, and we assumed $H^1(Y; \mathbb{Z}/2)$ to be nontrivial in assuming that $w_2(E)$ is nontrivial, so $|H^2(Y; \mathbb{Z})|/2$ is indeed an integer. \blacksquare

1.4. Configurations on the cylinder. There are two more kinds of configuration spaces we should consider. First, we should have a configuration space of connections on the bundle π^*E over $\mathbb{R} \times Y$. We write $\mathcal{A}_{\pi^*E}^{(4)}$ as the space of connections which, when restricted sufficiently far on each end, are pullbacks of connections on Y . (We will denote any gauge group or space of connections on a 4-manifold with a (4) unless there is no risk of confusion.) When we later give a Hilbert manifold modification of this construction, it will instead be replaced by an appropriate space of connections which have exponential decay to certain constant connections on the end. The space of framed connections on π^*E is

$$\tilde{\mathcal{A}}_{\pi^*E}^{(4)} := \mathcal{A}_{\pi^*E}^{(4)} \times (\pi^*E)_{(0,b)},$$

with the framing chosen at $(0, b) \in \mathbb{R} \times Y$. The gauge group acting on this is $\mathcal{G}_{\pi^*E}^{(4),e}$ — even gauge transformations which are constant on the ends. As usual, the framed configuration space is the space

$$\tilde{\mathcal{B}}_{\pi^*E}^{(4),e} := \tilde{\mathcal{A}}_{\pi^*E}^{(4)} / \mathcal{G}_{\pi^*E}^{(4),e},$$

with the right $SO(3)$ action on $\pi^*E_{(0,b)}$. The ‘restriction to the ends’ map is slightly more complicated than what one might do naively: there is a map

$$\text{ev}_- : \tilde{\mathcal{A}}_{\pi^*E}^{(4)} \rightarrow \tilde{\mathcal{A}}_E$$

given by sending (\mathbf{A}, p) to $(\mathbf{A}_{-\infty}, \gamma_{-\infty}(\mathbf{A}, p))$, where in the latter term we send p to its parallel transport under \mathbf{A} along γ traversed backwards from $(0, b)$. When the constant framing of the path (assigning the same $p \in \pi^*E_{(t,b)}$ for all t) is \mathbf{A} -parallel — if \mathbf{A} is in temporal gauge, for instance — this map is particularly simple: $\gamma_{-\infty}(\mathbf{A}, p) = p$. This map is equivariant on the left under the restriction-to- $(-\infty)$ map

$$\mathcal{G}_{\pi^*E}^{(4),e} \rightarrow \mathcal{G}_E^e,$$

and on the right under the $SO(3)$ -action, because parallel transport is an isomorphism

$$\pi^*E_{(0,b)} \rightarrow \pi^*E_{(-\infty,b)}$$

of right $SO(3)$ -sets. There is also a corresponding map ev_+ .

Thus the evaluation maps descend to a map of right $SO(3)$ -spaces

$$[\text{ev}_-] \times [\text{ev}_+] : \tilde{\mathcal{B}}_{\pi^*E}^{(4),e} \rightarrow \tilde{\mathcal{B}}_E^e \times \tilde{\mathcal{B}}_E^e.$$

We are most frequently interested in the spaces of connections on $\mathbb{R} \times Y$ with specified limits $[\text{ev}_\pm] \in \alpha_\pm$ inside previously specified $SO(3)$ -orbits. Furthermore, we should pick a relative homotopy class $z \in \pi_1(\mathcal{B}_E^e, [\alpha_-], [\alpha_+])$. Because \mathcal{A}_E is contractible and \mathcal{G}_E^e acts with connected stabilizers, these are in bijection with

$$\pi_0 \mathcal{G}_E^e \cong H^3(Y; \pi_3 \widetilde{\text{Aut}}(E_b)) \cong \mathbb{Z}$$

by an obstruction theory argument. Given limiting framed connections $[A_\pm, p_\pm] \in \alpha_\pm$, we choose lifts (denoted A_- and A_+) so that a path from A_- to A_+ in \mathcal{A}_E projects to the relative homotopy class z in $\tilde{\mathcal{B}}_E^e$. We then define

$$\tilde{\mathcal{A}}_{\pi^*E,z}(A_-, A_+) = \{(\mathbf{A}, p) \in \tilde{\mathcal{A}}_{\pi^*E}^{(4)} \mid \mathbf{A}_{\pm\infty} = A_{\pm\infty}\}.$$

(Note that the endpoint conditions only include the connection and nothing about the framing.) The natural gauge group here is the group of even gauge transformations that are harmonic on the ends:

$$\mathcal{G}_{\pi^*E}^{e,h}(A_-, A_+) = \{\sigma \in \mathcal{G}_{\pi^*E}^{(4),e} \mid d_{\pi^*A_\pm} \sigma = 0 \text{ on the corresponding end}\}.$$

The following is straightforward, and justifies this simplification.

Lemma 1.5. *The natural map*

$$\tilde{\mathcal{A}}_{\pi^*E,z}(A_-, A_+) / \mathcal{G}_{\pi^*E}^{e,h}(A_-, A_+) \rightarrow [\text{ev}]^{-1}(\alpha_- \times \alpha_+)$$

*is an isomorphism of $SO(3)$ -spaces, where the final space denotes the corresponding subset of $\tilde{\mathcal{B}}_{\pi^*E}^{(4),e}$.*

While the latter space is clearly the object of interest, the former is easier to describe, so we prefer to use it as our working definition of the configuration space of framed connections between two orbits. We record this as a definition:

Definition 1.2. If α_{\pm} are $SO(3)$ -orbits in $\tilde{\mathcal{B}}_E^e$, the naive configuration space³ of smooth framed connections from α_- to α_+ is

$$\tilde{\mathcal{B}}_{\pi^*E,z}^e(\alpha_-, \alpha_+) := \tilde{\mathcal{A}}_{\pi^*E,z}(A_-, A_+)/\mathcal{G}_{\pi^*E}^{e,h}(A_-, A_+),$$

where A_- and A_+ are representatives for the connections in the α_{\pm} , and a path from A_- to A_+ projects to the homotopy class z . This carries a continuous $SO(3)$ action, acting by translation on $\pi^*E_{(0,b)}$ on the right. Identifying

$$\alpha_{\pm} = [A_{\pm}, p] \cong E_b/\Gamma_{[A_{\pm}]}$$

with p varying over all possible framings at the basepoint, the endpoint maps

$$[\text{ev}_{\pm}] : \tilde{\mathcal{B}}_{\pi^*E,z}^e(\alpha_-, \alpha_+) \rightarrow \alpha_{\pm}$$

are induced by $\tilde{\mathcal{A}}_{\pi^*E}(A_-, A_+) \rightarrow E_b$, sending (\mathbf{A}, p) to $\gamma_{\pm\infty}(\mathbf{A}, p)$, respectively. There is a left translation action of \mathbb{R} , induced by the family of diffeomorphisms $\tau_t : \mathbb{R} \rightarrow \mathbb{R}$ with $\tau_t(s) = s - t$, given by sending

$$\tau_t^*(\mathbf{A}, p) = (\tau_t^*\mathbf{A}, \gamma_t(\tau_t^*\mathbf{A}, \tau_t^*p)).$$

In the framing coordinate, we are parallel transporting τ_t^*p forward by t along γ , so that it is sent back to the basepoint.

We have a corresponding version of Proposition 1.2 for the configuration spaces on cylinders. We do not repeat the setup of the determinant-1 $U(2)$ -model, but refer to it freely in the proof below. Lemma 1.1 still gives that the only possible stabilizers of elements in the 4-dimensional configuration space are the trivial group, the full group $SO(3)$, and circle subgroups conjugate to $SO(2)$.

Proposition 1.6. Let E be an $SO(3)$ -bundle over a 3-manifold Y , and α_{\pm} are $SO(3)$ -orbits in $\tilde{\mathcal{B}}_E^e$. We have the following descriptions of the reducible subspaces in the configuration space $\tilde{\mathcal{B}}_{\pi^*E,z}^e(\alpha_-, \alpha_+)$ under the $SO(3)$ action:

- (1) If the α_{\pm} are not both reducible orbits belonging to the same component $\text{Red}(Y, E)$, or after concatenating a path from α_+ to α_- in the (simply connected) reducible component they lie in, the homotopy class z is non-trivial, then $SO(3)$ acts freely on the configuration space.
- (2) If $\alpha_{\pm} \in \text{Red}_{SO(2)}(Y, E)$ are $SO(2)$ -reducible orbits lying in the same connected component of reducibles which contains no fully reducible point, labelled by $\{z_1, z_2\} \subset H^2(Y)$, and the homotopy class z is trivial, then the reducible subspace is a fiber bundle with base $\mathcal{B}_{\eta_1}(\alpha_-, \alpha_+)$ and fiber S^2 , and in particular consists of one connected component. There are no fully reducible points.
- (3) If $\alpha_- = \alpha_+$ are the same unique fully reducible orbit in a fixed component $\{z_1, z_1\} \in \text{Red}_*(Y, E)$, and z is the trivial homotopy class, there is a unique fully reducible point in the configuration space on the cylinder, lying in the unique connected component of reducible orbits.

Proof. It's easier to find the reducibles inside the larger space $\tilde{\mathcal{B}}_{\pi^*E}^{(4),e}$ and afterwards take the intersection with $[\text{ev}]^{-1}(\alpha_- \times \alpha_+)$. As in Proposition 1.2, working with connections on $\pi^*\tilde{E}$ with fixed trace, a connection representing an $SO(2)$ -fixed point has holonomy inside $U(1) \times U(1)$ and hence induces a global splitting $\pi^*\tilde{E} \cong \eta_1 \oplus \eta_2$.

³We use the term naive to contrast this space of smooth connections which are constant near ∞ to the later *Hilbert* spaces of connections which decay exponentially at ∞ .

As before, this gives a natural correspondence between the $SO(2)$ -fixed points and $\bigsqcup_{\eta} \mathcal{B}_{\eta_1}^{(4)} \sqcup \mathcal{B}_{\eta_2}^{(4)}$ when $\eta_1 \neq \eta_2$ and with the single space $\mathcal{B}_{\eta}^{(4)}$ when $\eta_1 = \eta_2$. In the first case the involution given by the Weyl group action swaps the two spaces and in the second case it acts with a single fixed point.

Write $\tilde{\mathcal{B}}_{\eta}^{(4),\text{red}}$ for the subspace of $\tilde{\mathcal{B}}_{\pi^*E}^{(4),e}$ consisting of reducible connections for which the corresponding line bundle is of topological type η ; in particular, we do not specify the stabilizer. If η is an $SO(2)$ -reducible component, the above shows that the map

$$\tilde{\mathcal{B}}_{\eta}^{(4),\text{red}} \rightarrow \tilde{\mathcal{B}}_{\eta}^{(4),\text{red}}/SO(3)$$

is a constant rank submersion, with fiber S^2 . The base may be identified with $\mathcal{B}_{\eta}^{(4)}$: it is identified as the $SO(2)$ -fixed point space modulo the Weyl group action.

To calculate the intersection with $[\text{ev}]^{-1}(\alpha_- \times \alpha_+)$, observe that the line bundles η_i restrict to isomorphic line bundles on both ends (because the inclusion $Y \hookrightarrow \mathbb{R} \times Y$ is a homotopy equivalence), and so the pair of line bundles associated to α_- and α_+ must be the same. Furthermore, after being put in temporal gauge, the relative homotopy class z traced out by one of these induced connections lies inside the simply connected space $\tilde{\mathcal{B}}_{\eta}^e(\alpha_-, \alpha_+)$. This gives us the restrictions in (2)-(3). ■

Remark 1.3. If the $SO(3)$ -bundle E is nontrivial, there are no full reducibles, and so only cases (1) and (2) arise.

When the limiting connections are flat, we can distinguish the relative homotopy classes z by the values of the integral calculating the “relative Pontryagin number” of the connection \mathbf{A} by

$$p_1(\mathbf{A}) := \frac{1}{8\pi^2} \int_{\mathbb{R} \times Y} F_{\mathbf{A}}^2;$$

this takes on a well-defined value modulo \mathbb{Z} , and the different discrete values it can take in \mathbb{R} parameterize the set of connected components $\pi_0 \tilde{\mathcal{B}}_{\pi^*E}^e(\alpha_-, \alpha_+)$. When the limits are not flat, one could pick a base connection A_0 on the cylinder interpolating between them and compute $\int (F_A^2 - F_{A_0}^2)$; alternatively, later the limits will be critical points of a perturbed Chern-Simons functional and we will define the action of a connection in terms of this perturbed functional.

1.5. Configurations on a cobordism. Suppose we are given a complete Riemannian manifold W with two cylindrical ends (i.e., specified isometries to $(-\infty, 0] \times Y_1$ and $[0, \infty) \times Y_2$) and an $SO(3)$ bundle \mathbf{E} with specified isomorphisms over the ends to bundles $E_i \rightarrow Y_i$. We think of this as a cobordism $Y_1 \rightarrow Y_2$. Here $E_i \rightarrow Y_i$ is an $SO(3)$ bundle, framed over a basepoint $b_i \in Y_i$. Furthermore, we pick a smooth embedded path $\gamma : \mathbb{R} \rightarrow W$ with $\gamma(t) = (b_1, t)$ for $t < 0$ sufficiently small and $\gamma(t) = (b_2, t)$ for t sufficiently large, and a trivialization of γ^*E restricting to the given trivializations on the ends. This path γ will serve the role $\mathbb{R} \times \{b\}$ did for the cylinder. Given orbits $\alpha_i \subset \tilde{\mathcal{B}}_{E_i}^e$, we construct the configuration space $\tilde{\mathcal{B}}_{\mathbf{E},z}^e(\alpha_1, \alpha_2)$ much as we did before. The basepoint is instead $\gamma(0)$, and the framing portion of the endpoint maps given by parallel transport to $\pm\infty$ along γ . There is still a decomposition of the configuration space into disconnected components, which we still label by z , but this no longer has a description in terms of relative homotopy classes. Our above discussion on relative Pontryagin numbers, however, still does, as will the definition of an action using the perturbed Chern-Simons functional.

As in the previous sections, we analyze the reducible points in terms of the connected components they lie in.

Definition 1.3. *Let W be a 4-manifold with two cylindrical ends, as above, equipped with an $SO(3)$ -bundle \mathbf{E} . We write $\text{Red}(W, \mathbf{E})$ to denote the set of connected components of the reducible subspace of $\tilde{\mathcal{B}}_{\mathbf{E}}^e$. This may be written as the disjoint union*

$$\text{Red}_*(W, \mathbf{E}) \sqcup \text{Red}_{SO(2)}(W, \mathbf{E}),$$

where the former denotes components that include some fully reducible orbit. If α_- and α_+ label orbits in the configuration spaces $\tilde{\mathcal{B}}_{E_i}^e$ of the ends, we write $\text{Red}(W, \mathbf{E})(\alpha_-, \alpha_+)$ for the set of connected components of the reducible subspace $\tilde{\mathcal{B}}_{\mathbf{E}}^e(\alpha_-, \alpha_+)$.

We enumerate the reducible components in the following.

Proposition 1.7. *Suppose W is a Riemannian 4-manifold equipped with $SO(3)$ bundle \mathbf{E} , with one incoming cylindrical end (Y_1, E_1) and one outgoing cylindrical end (Y_2, E_2) . The reducible subspaces of the $SO(3)$ -action on the configuration space $\tilde{\mathcal{B}}_{\mathbf{E}}^e(\alpha_-, \alpha_+)$ is as follows.*

- (1) *If either of the α_{\pm} are irreducible, or $\beta w_2 \mathbf{E} \in H^3(W; \mathbb{Z})$ is nonzero, then there are no reducible points in the configuration space.*
- (2) *If $\beta w_2 \mathbf{E} = 0$, then fix an integral lift c of $w_2 \mathbf{E}$, and use the restriction of c to the ends to determine the bijection between $\text{Red}(Y_i, E_i)$ and 2-element sets of cohomology classes in $H^2(Y_i; \mathbb{Z})$ as in Definition 1.1. Then the connected components $\text{Red}_{SO(2)}(W, \mathbf{E})$ are in bijection with pairs $\{z_1, z_2\} \subset H^2(W; \mathbb{Z})$ with $z_1 + z_2 = c$ and $z_1 \neq z_2$. The set $\text{Red}_{SO(2)}(W, \mathbf{E})(\alpha_-, \alpha_+)$ is the subset of $\text{Red}_{SO(2)}(W, \mathbf{E})$ consisting of pairs that restrict to the $\{x_1, x_2\}$ corresponding to each α_{\pm} .⁴*
- (3) *If \mathbf{E} is nontrivial there are no fully reducible points. Otherwise, fix a trivialization; this produces a bijection of $\text{Red}_*(Y_i, E_i)$ with 2-torsion cohomology classes in $H^2(Y_i; \mathbb{Z})$. Then the components $\text{Red}_*(W, \mathbf{E})$ are in bijection with 2-torsion cohomology classes in $H^2(W; \mathbb{Z})$. If α_{\pm} are both fully reducible orbits, then $\text{Red}_*(W, \mathbf{E})(\alpha_-, \alpha_+)$ is the subset of those 2-torsion classes in $H^2(W; \mathbb{Z})$ that restrict to the 2-torsion classes in $H^2(Y_i; \mathbb{Z})$ labelled by the α_{\pm} . In each component with a fully reducible point, there is a unique such point.⁵*

Proof. If $\beta w_2 \mathbf{E} \neq 0$, then \mathbf{E} cannot be written as the direct sum of a trivial line bundle and an oriented 2-plane bundle: $\beta w_2(\mathbf{E}) = e(\mathbf{E})$, the Euler class, which is the obstruction to finding a nonvanishing section; but an $SO(2)$ -reducible point induces such a splitting (there is a parallel section). Note that this is also the obstruction to finding an integral lift of $w_2(\mathbf{E})$. So suppose now that $c \in H^2(W; \mathbb{Z})$ is an integral lift of $w_2 \mathbf{E}$.

Almost identically to Proposition 1.6, the $SO(2)$ -fixed point space of $\tilde{\mathcal{B}}_{\mathbf{E}}^e$ is identified with the disjoint union over configuration spaces of connections on line

⁴Note in particular that taking the intersection of the reducible subspace with $[\text{ev}]^{-1}(\alpha_- \times \alpha_+)$ — that is, specifying the limits — does not increase the number of connected components. Furthermore, either of α_{\pm} may be $SO(2)$ -reducible or fully reducible.

⁵It may be that W has more full reducibles than the Y_i ; this is true if, for instance, each Y_i is an integer homology sphere but $H_1(W; \mathbb{Z})$ has 2-torsion.

bundles over W : $\sqcup_{\eta} \mathcal{B}_{\eta}^W$. If $\eta_1 \oplus \eta_2$ has first Chern class c , then the Weyl group acts on \mathcal{B}_{η_i} by sending a connection $[A]$ on one to the connection $[A - A_0]$ on the other, where A_0 is a fixed connection on the complex line bundle λ . Quotienting by the Weyl group we are only left with components labelled by pairs $\{z_1, z_2\}$ with $z_1 + z_2 = c$, and if $z_1 = z_2$, the Weyl group fixes a unique point in \mathcal{B}_{η}^W , the corresponding fully reducible point. ■

2. ANALYSIS OF CONFIGURATION SPACES

2.1. Tangent spaces. Before we introduce the Hilbert manifold versions of these constructions, we investigate the tangent spaces and differentials involved in the naive constructions (where all connections and gauge transformations are smooth).

The space \mathcal{A}_E of all smooth connections on E is an affine space over $\Omega^1(\mathfrak{g}_E)$; the natural Riemannian metric on $T_A\mathcal{A}_E = \Omega^1(\mathfrak{g}_E)$ is the (incomplete!) L^2 inner product on forms, using a bi-invariant inner product on the Lie algebra bundle \mathfrak{g}_E to define the Hodge star. With this, we can define ∇cs as a vector field on \mathcal{A}_E ; we may identify $(\nabla\text{cs})(A) = *F_A$. The action of the gauge group \mathcal{G}_E^e at $A \in \mathcal{A}_E$ has differential

$$d_A : \Omega^0(\mathfrak{g}_E) \cong T_e\mathcal{G}^e(E) \rightarrow T_A\mathcal{A}_E \cong \Omega^1(\mathfrak{g}_E).$$

The kernel of this map is the space of A -harmonic sections of \mathfrak{g}_E . The differential of the action of \mathcal{G}_E^e on the fiber E_b (whose tangent spaces are all canonically identified with \mathfrak{g}_b , the fiber of \mathfrak{g}_E at b) is given by evaluation at the basepoint $\text{ev}_b : \Omega^0(\mathfrak{g}_E) \rightarrow \mathfrak{g}_b$. We write

$$\tilde{d}_A : \Omega^0(\mathfrak{g}_E) \rightarrow \Omega^1(\mathfrak{g}_E) \oplus \mathfrak{g}_b$$

for the differential of the action of \mathcal{G}_E^e on $\tilde{\mathcal{A}}_E$ at (A, p) . Unfortunately, there is no natural identification of the cokernel of \tilde{d}_A — which we hope to eventually identify with $T_{[A,p]}\tilde{\mathcal{B}}_E^e$ — with the L^2 orthogonal complement of its image, since the image is not L^2 closed. Its closure is $\text{Im}(d_A) \oplus \mathfrak{g}_b$, the essential difficulty here being that point-evaluation is not L^2 -continuous. Rather, this orthogonal complement is precisely the same as $\ker(d_A^*)$. However, we can identify

$$(\text{Im}(d_A) \oplus \mathfrak{g}_b) / \text{Im}(\tilde{d}_A) \cong \mathfrak{g}_A^\perp,$$

the orthogonal complement inside \mathfrak{g}_b of the subspace \mathfrak{g}_A of elements that extend to A -parallel sections. Equivalently, \mathfrak{g}_A is the tangent space to the image of Γ_A in $\text{Aut}(E_b)$.

Picking a choice of (necessarily non-orthogonal) complement of $\text{Im}(\tilde{d}_A)$ inside $\text{Im}(d_A) \oplus \mathfrak{g}_b$ to identify with \mathfrak{g}_A^\perp , we may still decompose

$$T_{(A,p)}\tilde{\mathcal{A}}_E = \ker(d_A^*) \oplus \text{Im}(\tilde{d}_A) \oplus \mathfrak{g}_A^\perp.$$

We will henceforth write this as

$$\mathcal{T}_{A,p} = \mathcal{K}_A \oplus \mathcal{G}_A \oplus \mathfrak{g}_A^\perp,$$

giving us the isomorphism

$$T_{[A,p]}\tilde{\mathcal{B}}_E^e \cong \mathcal{K}_A \oplus \mathfrak{g}_A^\perp.$$

This can not be a decomposition into a direct sum of locally trivial bundles (the last factor has dimension varying between 0, 2, and 3), but it is when restricted to the submanifold of $\tilde{\mathcal{B}}_E^e$ consisting of framed connections with stabilizer of *fixed* conjugacy class (irreducible, conjugate to $SO(2)$, or fully reducible), where \mathfrak{g}_A^\perp actually defines a locally trivial vector bundle.

Because cs is, up to a constant, fixed under the action of \mathcal{G}_E^e and not just $\mathcal{G}_E^{e,b}$, its L^2 gradient on \mathcal{A}_E must take values in \mathcal{K}_A . It is clear from

$$\nabla\text{cs}(A) = *F_A$$

that the critical points of cs are precisely the flat connections.

If A_- and A_+ are connections on the ends of π^*E over $\mathbb{R} \times Y$ (and α_\pm the corresponding orbits in the configuration space), we constructed the 4-dimensional configuration space in Definition 1.2 as a quotient of $\tilde{\mathcal{A}}_{\pi^*E,z}(A_-, A_+)$. If we pick a base connection A_0 in this space, we get a description

$$\tilde{\mathcal{A}}_{\pi^*E,z}(A_-, A_+) = \{A_0 + \Omega_c^1(\pi^* \mathfrak{g}_E)\} \times (\pi^*E)_{0,b},$$

and in particular of the space of unframed connections as an affine space. The Lie algebra of the gauge group $\mathcal{G}_{\pi^*E}^{e,h}(A_-, A_+)$ is $\Omega_h^0(\mathfrak{g}_E)$, the space of sections of \mathfrak{g}_E that are A_\pm -harmonic on the corresponding ends. The differential of the gauge group action at some (\mathbf{A}, p) is given by $\sigma \mapsto (d_{\mathbf{A}}\sigma, \sigma(b))$. Picking a complement \mathfrak{g}_A^\perp of $\text{Im}(d_{\mathbf{A}})$, we again have our decomposition

$$T_{(\mathbf{A},p)}\tilde{\mathcal{A}}_E = \ker(d_{\mathbf{A}}^*) \oplus \text{Im}(d_{\mathbf{A}}) \oplus \mathfrak{g}_A^\perp.$$

We also have the endpoint maps to $(A_\pm, E_b) \cong SO(3)$. Before modding out by gauge, the positive endpoint map sends $\text{ev}_+(\mathbf{A}, p) = \gamma_\infty(\mathbf{A}, p)$, the parallel transport of p to ∞ using \mathbf{A} . The differential of this map at (\mathbf{A}, p) is a linear map

$$\Omega_c^1(\mathfrak{g}_E) \oplus \mathfrak{g}_{0,b} \rightarrow \mathfrak{g}_b,$$

equal to the integral

$$\int_0^\infty \gamma_{\mathbf{A}}^* \sigma(t, b) dt.$$

Here we use parallel transport backwards along γ using \mathbf{A} so that $\sigma(t, b)$ is taken to the single vector space $\mathfrak{g}_{0,b}$, where it makes sense to integrate. Observe that if \mathbf{A} is in temporal gauge, this is just $\int_0^\infty \sigma(t, b) dt$.

Nothing changes in the above description when passing to a general Riemannian 4-manifold W with cylindrical ends.

2.2. Hilbert manifold completions. To have a useful transversality theory, we must replace the naive “infinite-dimensional manifolds” \mathcal{A} , \mathcal{G} , and so on with certain completions which are manifolds modeled on Hilbert spaces. If E is an $SO(3)$ -bundle over a 3-manifold, define $\mathcal{A}_{E,k}$ to be the set of L_k^2 connections on E ; more precisely, if Ω_k^1 denotes the Hilbert space of L_k^2 1-forms, then

$$\mathcal{A}_{E,k} = A_0 + \Omega_k^1(\mathfrak{g}_E)$$

for any choice of smooth base connection A_0 . The resulting set does not depend on the choice of smooth connection. Similarly, we may define

$$\mathcal{G}_{E,k+1} = \{\sigma \in L_{k+1}^2(\text{End}(E)) \mid \sigma(x) \text{ is an isomorphism for all } x\},$$

given the constraint that $k+1 \geq 2$ so that gauge transformations are automatically continuous and we may take point evaluations as in the definition. Then $\mathcal{G}_{E,k+1}^e$ is the subset of sections which lift continuously to $\widetilde{\text{Aut}}(E)$. As in [DK90, Section 4.2], the space $\mathcal{G}_{E,k+1}^e$ becomes a Banach Lie group, and we have a smooth action of $\mathcal{G}_{E,k+1}^e$ on $\tilde{\mathcal{A}}_{E,k} = \mathcal{A}_{E,k} \times E_b$; the same argument as the case of smooth gauge transformations shows that this action is free. The following is [Bou08, Proposition 3.1.5.10]. To facilitate comparison to the version stated there, note that Bourbaki’s definition of immersion includes the assumption that the differential has closed image with a closed complement.

Proposition 2.1. *Suppose G is a Banach Lie group and X a Banach manifold on which G acts. Suppose that G acts freely and properly on X , and such that if $\rho(x) : T_x G \rightarrow T_x X$ is the differential of the action at the point x , the image of $\rho(x)$ is closed and has some closed complement for all x . Then the quotient topology on X/G is Hausdorff, and there is a unique smooth structure on X/G such that $\pi : X \rightarrow X/G$ is a submersion. Furthermore, $X \rightarrow X/G$ is a principal G -bundle.*

We should see that all these assumptions apply to the action of $\mathcal{G}_{E,k+1}^e$ on $\tilde{\mathcal{A}}_{E,k}$, giving us the quotient manifold $\tilde{\mathcal{B}}_{E,k}^e$. We record this as a lemma.

Lemma 2.2. *$\tilde{\mathcal{B}}_{E,k}^e$ is a Hausdorff Hilbert manifold if $k \geq 1$.*

Proof. Properness of this action would follow from properness of the action on the unframed space of connections, because the map $\tilde{\mathcal{A}}_{E,k} \rightarrow \mathcal{A}_{E,k}$ is proper (because E_b is compact). That the action of $\mathcal{G}_{E,k+1}$ on $\mathcal{A}_{E,k}$ is proper is proved in [DK90, Section 2.3.7], by a bootstrapping argument.

Because $\tilde{\mathcal{A}}_{E,k}$ is a Hilbert manifold, we only need to verify that the image of the differential is closed; write this explicitly as the sum

$$(d_A, \text{ev}_b) : \Omega_{k+1}^0(\mathfrak{g}_E) \rightarrow \Omega_k^1(\mathfrak{g}_E) \oplus \mathfrak{g}_b.$$

First note that $\text{Im}(d_A)$ is closed: there is an L^2 -orthogonal decomposition

$$\Omega_k^1(\mathfrak{g}_E) = \ker(d_A^*) \oplus \text{Im}(d_A);$$

that this is true in Ω_k^1 and not just Ω_0^1 follows from elliptic regularity (if $d_A \sigma = \eta$ where $\eta \in L_k^2$, then $\sigma \in L_{k+1}^2$).

Then suppose $(\eta_n, g_n) \rightarrow (\eta, g)$ is a convergent sequence, for which there is a sequence of $\sigma_n \in \Omega_{k+1}^0$ with $d_A \sigma_n = \eta_n$ and $\sigma_n(b) = g_n$. Because $\text{Im}(d_A)$ is closed, there is some σ with $\eta = d_A \sigma$. Thus there is some sequence $\psi_n \in \ker(d_A)$ so that $\sigma_n + \psi_n \rightarrow \sigma$. As long as $k+1 \geq 2$, point-evaluation is continuous, and so $g_n + \psi_n(b) \rightarrow g$; because $g_n \rightarrow g$, we see that $\psi_n(b) \rightarrow 0$. Thus $\sigma(b) = g$ and $d_A \sigma = \eta$ as desired. \blacksquare

This provides somewhat less inspiring charts than the case of irreducible connections and unbased gauge group, where one can explicitly present slices for the action of the gauge group as those connections with

$$d_{A_0}^*(A - A_0) = 0.$$

One can at least use the decomposition

$$\mathcal{T}_{A,k} = \mathcal{K}_{A,k} \oplus \mathcal{G}_{A,k} \oplus \mathfrak{g}_A^\perp$$

to provide slices for the based gauge group as exponentials of $\mathcal{K}_{A,k} \oplus \mathfrak{g}_A^\perp$. In addition to the natural tangent bundle $\mathcal{T}\mathcal{A}_{E,k} = \Omega_k^1(\mathfrak{g}_E)$, we may also define completed tangent bundles $\mathcal{T}_j \tilde{\mathcal{A}}_{E,k} = \Omega_j^1(\mathfrak{g}_E)$ for any $j \leq k$.

After passing to this completion, the analogue of Proposition 1.2 remains true. Whenever we construct a gauge transformation via parallel transport, it has one greater order of differentiability than the connection we used to define parallel transport.

2.3. The 4-dimensional case. In contrast to the case of homology 3-spheres, where to define Floer's instanton homology we only needed to study instantons on the cylinder with irreducible limits on the ends, the same flavor of Sobolev completion will not suffice to define the configuration space of 4-manifolds with general limits. Instead, we will need to use *weighted* Sobolev spaces. Recall π^*E is an $SO(3)$ -bundle on $\mathbb{R} \times Y$, given as the pullback of some bundle E on Y under the projection map. Let f_δ be a smooth positive function on \mathbb{R} such that $f_\delta(t) = e^{\delta|t|}$ for $|t| \geq 1$.

Suppose we have fixed limiting connections A_- and A_+ on E . Here we are fixing *connections* and not gauge equivalence classes; we write the corresponding orbits in $\tilde{\mathcal{B}}_E^e$ as α_\pm . Pick a smooth base connection A_0 on π^*E over $\mathbb{R} \times Y$ that agrees with the pullback connections of A_- and A_+ near $\pm\infty$ respectively; it traces out a relative homotopy class $z \in \pi_1(\tilde{\mathcal{B}}_E^e, A_-, A_+)$. Then one may define the space of $L_{k,\delta}^2$ sections of π^*E as those sections for which

$$\int_{\mathbb{R} \times Y} f^2 (|\sigma|^2 + |\nabla_{A_0}\sigma|^2 + \dots + |\nabla_{A_0}^k\sigma|^2) < \infty;$$

this is the $L_{k,\delta}^2$ norm with respect to A_0 . Then we say

$$\mathcal{A}_{\pi^*E,z,k,\delta}(A_-, A_+) := \{A_0 + \Omega_{k,\delta}^1(\mathfrak{g}_E)\}.$$

This only depends on A_0 through the choice of z , but it clearly depends on the limits A_\pm ; if two pairs (A_\pm^i) of limits are gauge equivalent by gauge transformations in the same homotopy class, then there is a gauge transformation taking

$$\mathcal{A}_{\pi^*E,z,k,\delta}(A_-^1, A_+^1) \cong \mathcal{A}_{\pi^*E,z,k,\delta}(A_-^2, A_+^2).$$

Similarly, this set does not depend on f , only δ .

We need a Banach Lie group $\mathcal{G}_{\pi^*E,k+1,\delta}^{e,h}(A_-, A_+)$. Set

$$\Gamma_\pm = \Gamma(\alpha_\pm) \subset \mathcal{G}_{E,k+1}^e;$$

this is precisely the set of α_\pm -harmonic gauge transformations. As before, we make sense of the $L_{k+1,\delta}^2$ norm by considering gauge transformations as certain sections of $\pi^*(\text{End}(E))$. If we only cared about constant limits (as is the case when the α_\pm are irreducible), we could take the Lie group

$$\mathcal{G}_{\pi^*E,k+1,\delta} = \{\sigma \in \Gamma(\text{Aut}(\pi^*E)) \subset \Gamma(\text{End}(\pi^*E)) \mid (\sigma - 1) \in L_{k+1,\delta}^2(\text{End}(\pi^*E))\}.$$

That is, roughly, gauge transformations in this group are sections of $\pi^*(E \otimes E^*)$ that are pointwise automorphisms and decay exponentially quickly to the identity. This is again a Banach Lie group that acts smoothly on $\mathcal{A}_{\pi^*E,z,k,\delta}(A_-, A_+)$; this is essentially the content of [Don02, Section 4.3]. We write

$$\begin{aligned} \mathcal{G}_{\pi^*E,k+1,\delta}^h(\alpha_\pm) = \\ \{\sigma \in \Gamma(\text{Aut}(E)) \mid \exists \psi \text{ with parallel ends and } \sigma - \psi \in L_{k+1,\delta}^2(\text{End}(E))\}. \end{aligned}$$

This is again a Banach Lie group, but now we have the surjective endpoint evaluation map

$$\mathcal{G}_{\pi^*E,k+1,\delta}^h(A_-, A_+) \rightarrow \Gamma_- \times \Gamma_+.$$

There is also the natural subgroup $\mathcal{G}_{\pi^*E,k+1,\delta}^{e,h}(A_-, A_+)$, those transformations that lift to sections of $\widetilde{\text{Aut}}(\pi^*E)$. Following Definition 1.2, we use this to define the configuration space as follows.

Definition 2.1. *The configuration space of framed connections on π^*E between α_- and α_+ is defined to be*

$$\tilde{\mathcal{B}}_{\pi^*E,z,k,\delta}^e(A_-, A_+) = \tilde{\mathcal{A}}_{\pi^*E,z,k,\delta}(A_-, A_+) / \mathcal{G}_{\pi^*E,k+1,\delta}^{e,h}(A_-, A_+).$$

This carries an action of $SO(3)$ by acting on the fiber above the basepoint on the right. The endpoint maps given by parallel transport along $\mathbb{R} \times \{b\}$ are right $SO(3)$ maps to $[\alpha_{\pm}] \cong E_b / \Gamma_{\pm}$; these maps are smooth and furthermore submersions. Equivariant with respect to these is a (left) \mathbb{R} action by pullback of the connection and parallel transport of the framing.

Again, we should check that this group action is well-behaved.

Lemma 2.3. *$\tilde{\mathcal{B}}_{\pi^*E,z,k,\delta}^e(\alpha_-, \alpha_+)$ carries the natural structure of a smooth Hilbert manifold when $k \geq 2$.*

Proof. We need to check that the action on $\tilde{\mathcal{A}}_{\pi^*E,z,k,\delta}(A_-, A_+)$ is proper, and that the tangent spaces to orbits are closed.

Let (σ_n, A_n) be a sequence of $L_{k+1,\delta}^2$ gauge transformations and $L_{k,\delta}^2$ connections so that $\sigma_n^* A_n \rightarrow B$ in the $L_{k,\delta}^2$ topology, and so that $A_n \rightarrow A$ in the $L_{k,\delta}^2$ topology. Write $B_n = \sigma_n^* A_n$. We want to show that σ_n then has a convergent subsequence in the $L_{k+1,\delta}^2$ topology.

On any chart with compact closure, choose an arbitrary trivialization of π^*E , so that we represent A_n and B_n as matrices of 1-forms and σ_n as a map to $SO(3)$. Then as in [DK90, Section 2.3.7] we have the formula

$$d\sigma_n = \sigma_n A_n - B_n \sigma_n;$$

from this it follows that an L_j^2 bound on σ_n implies an L_{j+1}^2 bound as long as $j \leq k$, and an L^2 bound follows from compactness of $SO(3)$ and of the chart itself. Therefore $\|\sigma_n\|_{L_{k+1}^2}$ is uniformly bounded; by compactness, we may choose a convergent subsequence $\sigma_n \rightarrow \sigma$ in L_k^2 . But then we may identify $\sigma^* A = B$, and therefore we have

$$d\sigma - d\sigma_n = \sigma A - \sigma_n A_n - B\sigma + B_n \sigma_n,$$

which may be rewritten as

$$(\sigma - \sigma_n)(A) + \sigma_n(A - A_n) - B(\sigma - \sigma_n) - (B - B_n)(\sigma_n).$$

Then because all of $\sigma - \sigma_n$, $A - A_n$, and $B - B_n$ go to 0 in L_k^2 , we see that $d\sigma - d\sigma_n \rightarrow 0$ in L_k^2 , and therefore $\sigma_n \rightarrow \sigma$ in L_{k+1}^2 .

This shows that $\sigma_n \rightarrow \sigma$ in $L_{k+1,\text{loc}}^2$. We should check convergence on the ends. We may write A_n and B_n as the flat connection α plus an $L_{k,\delta}^2$ 1-form valued in \mathfrak{g}_E . Choose a trivialization of E on a chart U of Y , and extend that to a trivialization of π^*E over $\mathbb{R} \times U$. Consider this on one end at a time; for convenience we say $(-\infty, 0] \times U$. We write the flat connection α as $d + \omega_\alpha$ in this chart, where ω_α is a particular matrix of 1-forms, and then the connections in this trivialization take the form $d + \omega_\alpha + A_n$ for an $L_{k,\delta}^2$ 1-form A_n .

The above formula now gives us

$$d\sigma_n = \sigma_n(\alpha + A_n) - (\alpha + B_n)\sigma_n = \sigma_n A_n - B_n \sigma_n + (\sigma_n \alpha - \alpha \sigma_n);$$

more simply stated, this is

$$d_\alpha \sigma_n = \sigma_n A_n - B_n \sigma_n.$$

For the limit $\sigma \in L_{k+1, \text{loc}}^2$, we then have $(d_\alpha \sigma) = \sigma A - B\sigma$. Because the $L_{0, \delta}^2$ norm of σA agrees with that of A , we see that $(d_\alpha \sigma)$ has an $L_{0, \delta}^2$ bound; because σ is the limit of gauge transformations which are asymptotic to 1 on the ends, σ is asymptotic to 1 on the ends, and so σ is $L_{0, \delta}^2$.

We also have

$$d_\alpha \sigma - d_\alpha \sigma_n = \sigma(A - A_n) + (\sigma - \sigma_n)A_n - (B - B_n)(\sigma) - B_n(\sigma - \sigma_n).$$

Therefore

$$\|d_\alpha \sigma - d_\alpha \sigma_n\|_{L_{0, \delta}^2} \leq \|A - A_n\| + 2\|A_n\| + \|B - B_n\| + 2\|B_n\|,$$

and in particular $\|\sigma_n\|_{L_{1, \delta}^2}$ is uniformly bounded. By the compactness theorem for weighted Sobolev spaces, $\sigma_n \rightarrow \sigma$ in $L_{1, \delta'}^2$ for $\delta' < \delta$. Inducting with the previous formula, we may see that $\|\sigma_n\|_{L_{k+1, \delta}^2}$ is uniformly bounded. Therefore $\sigma_n \rightarrow \sigma$ in $L_{k+1, \delta'}^2$ for $\delta' < \delta$. The above formula again shows using Sobolev multiplication that $\|\sigma - \sigma_n\|_{L_{k+1, \delta}^2} \rightarrow 0$ on this chart, and by running this argument a few times $\sigma_n \rightarrow \sigma$ in $L_{k+1, \delta}^2$ on all of $\mathbb{R} \times Y$, as desired.

The extension to the entire gauge group (with gauge transformations that are exponentially decaying to α -harmonic gauge transformations) is formal, as the quotient by the subgroup used above is compact, equal to the space of harmonic gauge transformations on the two ends.

To see that tangent spaces are closed, it suffices to show that $\Omega_{k, \delta}^1(\mathfrak{g}_E)$ admits a closed splitting $\text{Im}(d_{\mathbf{A}}) \oplus \ker(d_{\mathbf{A}}^*)$; we want to show that every element may be written uniquely as a sum of that form.

That these are closed subspaces follows from the closed range theorem for densely defined operators (thinking of this as a densely defined operator in L^2): the range of $d_{\mathbf{A}}$ is closed if and only if its image is the L^2 orthogonal complement of $\ker(d_{\mathbf{A}}^*) \cap L_{-k, -\delta}^2$.

If $\psi \in \ker(d_{\mathbf{A}}^*) \cap L_{-k, -\delta}^2$, then by separation of variables on the end, using that the signature operator on Y has no nonzero eigenvalues of magnitude less than δ , we see that $\psi \in L_{-k, \delta, \text{ext}}^2$: it is the sum of a section which is constant and α - or β -harmonic on the ends and an $L_{-k, \delta}^2$ section; similarly elliptic regularity and the assumption that \mathbf{A} is $L_{k, \delta}^2$ implies that $\psi \in L_{k, \delta, \text{ext}}^2$. Now we may apply the usual integration by parts trick to show that $\langle d_{\mathbf{A}} \sigma, \psi \rangle = \langle \sigma, d_{\mathbf{A}}^* \psi \rangle = 0$ for all σ .

What is left is to see that the equation $\Delta_{\mathbf{A}} \sigma = -d_{\mathbf{A}}^* \psi$ has a solution for any $\psi \in \Omega_{k, \delta}^1$; to see this, choose ψ so that $d_{\mathbf{A}}^* \psi$ is in the orthogonal complement of $\text{Im} \Delta_{\mathbf{A}}$. But as before we may see that if

$$0 = \langle \Delta_{\mathbf{A}} \sigma, d_{\mathbf{A}}^* \psi \rangle = \langle d_{\mathbf{A}} \sigma, d_{\mathbf{A}} d_{\mathbf{A}}^* \psi \rangle,$$

then in particular $d_{\mathbf{A}} d_{\mathbf{A}}^* \psi = 0$, and as such $d_{\mathbf{A}}^* \psi$ is parallel. Because ψ is asymptotic to 0 on the ends, we see that $d_{\mathbf{A}}^* \psi = 0$, and hence indeed

$$\Delta_{\mathbf{A}} \sigma = -d_{\mathbf{A}}^* \psi$$

is solvable for all ψ . Uniqueness follows because if $\Delta_{\mathbf{A}} \sigma = 0$, then because σ is asymptotic to 0 on the ends, by the usual integration by parts trick $d_{\mathbf{A}} \sigma = 0$; again because σ is asymptotically zero, we see that σ is globally zero. \blacksquare

The discussion of tangent bundles and reducible configurations (Proposition 1.6) from the naive case carry over without change.

Finally, suppose we are given a complete Riemannian manifold W with two cylindrical ends (i.e., specified isometries to $(-\infty, 0] \times Y_1$ and $[0, \infty) \times \overline{Y_2}$), an $SO(3)$ bundle \mathbf{E} with specified isomorphisms over the ends to bundles $E_i \rightarrow Y_i$, and specified α_i on the \overline{E}_i . Further suppose W is equipped with an embedded path γ in W , cylindrical and agreeing with specified basepoints b_i on the ends. Then the construction of configuration spaces $\tilde{\mathcal{B}}_{\mathbf{E}, k, \delta}^e(\alpha_1, \alpha_2)$ with equivariant endpoint maps to $(E_i)_{b_i}/\Gamma_{\alpha_i}$ carries through with essentially no change from above. When defining cobordism maps in instanton homology, the data of the embedded path will only ultimately matter up to homotopy of such paths, which reduces to the data of a relative homotopy class $(D^1, \{0\}, \{1\}) \rightarrow (\overline{W}, Y_1, Y_2)$.

3. CRITICAL POINTS AND PERTURBATIONS

3.1. Holonomy perturbations. In order to carry out the construction of a Morse-like homology theory on our configuration space $\tilde{\mathcal{B}}_{E,k}^e$, the Chern-Simons functional will usually need to be perturbed so that the critical sets are isolated orbits of $SO(3)$, and the moduli spaces of trajectories are smooth manifolds of the appropriate dimension. Here we recall the standard perturbations used in this situation, as well as their basic properties, from [KM11b, Section 3.2]; they work relative to a knot K , but we may simply take $K = \emptyset$. These perturbations originate in Floer's definition of instanton homology and have been used consistently in the development of the subject.

Suppose we are given a collection of immersions

$$q_i : S^1 \times D^2 \looparrowright Y, \quad i = 1, \dots, N,$$

such that the q_i all agree in a small neighborhood of $\{1\} \times D^2$. Further pick a conjugation-invariant smooth function $h : G^N \rightarrow \mathbb{R}$ and a nonnegative 2-form μ on D^2 that integrates to 1 and vanishes near the boundary. Taking the holonomy of a connection A along the family of curves $q_i(e^{2\pi it}, z)$ parameterized by i and z gives a map

$$\text{Hol}_q(A) : D^2 \rightarrow SO(3)^N;$$

we are interested in the functions (called $SO(3)$ -cylinder functions)

$$f_{q,h,\mu}(A) = \int_{D^2} h(\text{Hol}_q(A))\mu.$$

These functions $f_{q,h,\mu}$ have smooth extensions to the Hilbert manifolds $\mathcal{A}_{E,k}$, where (as on \mathcal{A}_E) they are invariant under the action of the gauge group. By ignoring the framing, we can extend this trivially to a map $f_{q,h,\mu} : \tilde{\mathcal{A}}_{E,k} \rightarrow \mathbb{R}$ invariant under both the left action of the *full* gauge group and the right action of $SO(3)$. Thus they descend to $SO(3)$ -invariant smooth functions $f_{q,h,\mu} : \tilde{\mathcal{B}}_{E,k}^e \rightarrow \mathbb{R}$. Note that they are furthermore invariant under the remaining $H^1(Y; \mathbb{Z}/2)$ action.

In the context of the $U(2)$ -model, we choose an integral cohomology class $c \in H^2(Y; \mathbb{Z})$ reducing mod 2 to $w_2 \mathbf{E}$ and hence a lift of \mathbf{E} to the $U(2)$ -bundle $\lambda \oplus \mathbb{C}$, where $c_1 \lambda = c$. Fixing a base connection A_0 , the configuration space is the set of connections for which $\text{tr}(A) = A_0$. Associated to any connection with fixed trace is the holonomy map

$$\text{Hol}_A : \Omega Y \rightarrow U(2),$$

where our loops are based at b ; we also have, associated to A_0 , a map

$$\text{Hol}_{A_0} : \Omega Y \rightarrow U(1).$$

Because $\text{tr}(A) = A_0$, we have

$$\det \text{Hol}_A = \text{Hol}_{A_0};$$

thus we may define a relative holonomy map

$$\text{Hol}'_A := \text{Hol}_A \text{Hol}_{A_0}^{-1} : \Omega Y \rightarrow SU(2).$$

This map is equivariant under the action of the gauge group of determinant-1 automorphisms (acting on $SU(2)$ by translation by $\sigma(b)$). If one further acts on

connections by the operation of tensoring with $\chi_{\mathbb{C}}$, where χ is a real line bundle, this describes the action of

$$\mathcal{G}_E/\mathcal{G}_E^e = H^1(Y; \mathbb{Z}/2)$$

on this space. Abusing notation to write $\chi : \pi_1 W \rightarrow \pm 1$ for the holonomy of the unique flat connection on this real line bundle, the corresponding map $\chi_{\mathbb{C}} : \pi_1 W \rightarrow U(1)$ has determinant 1; so

$$\text{Hol}'_{A \otimes \chi}(\gamma) = \chi(\gamma) \text{Hol}'_A.$$

That is, Hol' is not invariant under the $H^1(Y; \mathbb{Z}/2)$ action. Now as above we may define cylinder functions using this relative holonomy map; we call these $SU(2)$ -cylinder functions. Note that $SO(3)$ -cylinder functions are the special case of $H^1(Y; \mathbb{Z}/2)$ -invariant $SU(2)$ -cylinder functions.

We're interested in the gradient flow equation of the Chern-Simons functional on $\tilde{\mathcal{B}}$; the relevant perturbations (called ‘‘holonomy perturbations’’) will be the formal gradients of the cylinder functions described above, so we will need to know that the formal gradients are well-behaved. The most convenient way to do so is to calculate them explicitly. We write

$$\nabla h : SO(3)^N \rightarrow \mathfrak{so}(3)^N$$

for the gradient of h using the Lie group trivialization of the tangent bundle and the standard inner product on $\mathfrak{so}(3)$; write

$$\nabla^j h : SO(3)^N \rightarrow \mathfrak{so}(3)$$

for the j th component. Kronheimer and Mrowka give a formula for this in [KM11b, Equation (74)] along the base $[-\varepsilon, \varepsilon] \times D^2 \hookrightarrow Y$ of the embeddings q_i as

$$\nabla_f := (\nabla f_{q,h,\mu})(A) = * \left(\sum_{i=1}^N (q_i)_* ((\nabla^i h)(\text{Hol}_{\mathbf{q}}(A))\mu) \right),$$

where $(q_i)_*$ is the pushforward of differential forms on each tangent space. This is defined on the rest of $\text{Im}(\mathbf{q})$ via parallel transport, and then extended by 0 to the rest of Y . The following properties of the formal gradient ∇_f are enumerated as in [KM11b, Proposition 3.5].

Theorem 3.1. *Let $f : \mathcal{A}_{E,k} \rightarrow \mathbb{R}$ be a fixed $SU(2)$ -cylinder function as above. Then the formal gradient ∇_f defines a smooth section of the tangent bundle $\mathcal{T}_k \mathcal{A}_{E,k}$, and its first derivative $D\nabla_f$, considered as a section of the bundle $\text{Hom}(\mathcal{T}_k, \mathcal{T}_k)$, extends to $\text{Hom}(\mathcal{T}_j, \mathcal{T}_j)$ for all $j \leq k$.*

Further, fixing a reference connection A_0 , we have the following pointwise norm bounds for constants C, C_k, C' :

$$\begin{aligned} \|\nabla_f(A)\|_{L^\infty} &\leq C \\ \|\nabla_f(A)\|_{L_k^2} &\leq C_k(1 + \|A - A_0\|_{L_k^2})^k \\ \forall \infty \geq p \geq 2, \|\nabla_f(A) - \nabla_f(A')\|_{L^p} &\leq C' \|A - A'\|_{L^p}. \end{aligned}$$

More generally, for any n we may find continuous increasing functions k_n (depending on f) such that

$$\|(D^n \nabla_f)(A)\|_{L_k^2} \leq k_n(\|A - A_0\|_{L_k^2}),$$

where on the left side we're taking an operator norm. Furthermore, f (and thus ∇_f) are invariant under the action of the full gauge group.

Suppose now we have a collection of $SU(2)$ -cylinder functions f_i ; we may associate an increasing sequence of constants (K_i) so that when $\sum_i K_i |a_i|$ converges, $\sum a_i \nabla_{f_i}$ converges in C^k on compact sets. One may take K_i to be, for instance,

$$K_{i-1} + \sum_{n=1}^i k_n(f_i)(i).$$

Then we may define the K -weighted ℓ^1 Banach space as the space of sequences (a_i) with $\|a\|_{\mathcal{P}} = \sum K_i |a_i| < \infty$. There is a map $\mathcal{P} \rightarrow C^\infty(\tilde{\mathcal{B}}_{E,k}^e, \mathbb{R})$ given by sending

$$\pi = \sum_i a_i \mapsto \sum_i a_i f_i =: f_\pi$$

and a map to smooth sections of the tangent bundle by sending π to its formal gradient, which we still call ∇_π . We call \mathcal{P} a Banach space of perturbations and the induced functions f_π (which are now L^1 sums of $SU(2)$ cylinder functions, but not necessarily cylinder functions themselves) *holonomy perturbations*.

The results of the previous theorem extend to smoothness of the \mathcal{P} -parameterized section of the tangent bundle, and parameterized inequalities for $\nabla_\pi(A)$ where now π may vary within \mathcal{P} (with an extra factor of $\|\pi\|_{\mathcal{P}}$ in the right-hand sides); this is [KM11b, Proposition 3.7].

The following is the relevant geometric fact about $SO(3)$ and $SU(2)$ -cylinder functions — and therefore also holonomy perturbations. It is well-known, though usually stated nonequivariantly and for irreducible connections. Note that in the first statement we have quotiented by the entire (as opposed to even) based gauge group.

Lemma 3.2. *Given any compact $SO(3)$ -invariant submanifold $S \subset \tilde{\mathcal{B}}_{E,k}$, the restriction of $SO(3)$ -cylinder functions to S are dense in the space of smooth invariant functions $C^\infty(S)^{SO(3)}$. This is furthermore true even if we demand the cylinder functions vanish in a small invariant open neighborhood \mathcal{O} of a finite set of orbits.*

Similarly, given a compact $SO(3)$ -invariant submanifold $S \subset \tilde{\mathcal{B}}_E^{\det} \cong \tilde{\mathcal{B}}_E^e$, the restriction of $SU(2)$ -cylinder functions to S are dense in the space of smooth invariant functions on S .

Proof. If $y \in Y$ is the basepoint, consider the maps $\mathcal{A}_{E,k} \rightarrow SO(3)$ given by $\text{Hol}(\gamma_i)$ where γ_i are some family of immersed loops, based at y , all with the same germ there. ('Immersed' is implicit for all loops in the rest of this argument.) These maps are equivariant with respect to the action of the gauge group on $SO(3)$ (conjugating by the value at the basepoint), and hence descend to $SO(3)$ -equivariant maps $\tilde{\mathcal{B}}_{E,k} \rightarrow SO(3)$. The claim is that these maps separate points and tangent vectors in $\tilde{\mathcal{B}}_{E,k}$; then, by compactness of S there is some finite collection γ_i , $1 \leq i \leq N$, such that $\text{Hol}_A(\gamma_i) : S \rightarrow SO(3)^N$ is an embedding.

If two connections have the same holonomy along every immersed loop γ based at y with specified germ, they have the same holonomy along every piecewise smooth immersed loop: we may write any piecewise smooth loop γ based at y as being the composition of an arbitrarily small loop and a loop with the specified germ, and so $\text{Hol}_A \gamma = \text{Hol}_{A'} \gamma$ for any piecewise smooth based loop γ . For any point $p \in Y$, pick

a path α from y to p , and define

$$\sigma(p) = \text{Hol}_A(\alpha)\text{Hol}_{A'}^{-1}(\alpha);$$

that this is well-defined follows from the assumption that the (based) holonomies always agree along closed loops. Taking derivatives along any smooth curve based at y we see that $A - d_A\sigma = A'$, and clearly $\sigma(y) = 1$. So based holonomy separates connections modulo based gauge. Equivalently, holonomy separates framed connections modulo gauge.

At the level of tangent spaces, we run the same argument with any ω such that $d\text{Hol}_A(\gamma)(\omega) = 0$ for all γ to see that $\omega = d_A\sigma$ for some $\sigma \in \Omega^0(\mathfrak{g}_E)$. Pick a path α from y to an arbitrary point p ; then we have a natural isomorphism

$$\mathfrak{g}_p \cong T\text{Isom}(E_y, E_p)$$

coming from the framing on E_y , and so we may define

$$\sigma(p) = d\text{Hol}_A(\alpha)(\omega);$$

as before this is independent of the choice of α and it is easy to verify from this definition that $d_A\sigma = \omega$. Thus based holonomy separates tangent vectors in $\tilde{\mathcal{B}}_{E,k}$.

Now choose some finite set of embedded based curves γ_i with the same germ so that $\text{Hol}_A(\gamma)$ embeds S into $SO(3)^N$. Given $h : S \rightarrow \mathbb{R}$ an $SO(3)$ -invariant function, extend it to $\tilde{h} : SO(3)^N \rightarrow \mathbb{R}$ arbitrarily using the equivariant tubular neighborhood theorem. We may then approximate h by holonomy perturbations by picking a sequence of embeddings q_i^j of solid tori of small radius around the curves so that the radii go to zero as $j \rightarrow \infty$; then the cylinder functions

$$f_{q^j, \tilde{h}, \mu} : S \rightarrow \mathbb{R}$$

approach h in $C^\infty(S)$ for any choice of 2-form μ , as desired.

The final claim about $SO(3)$ -cylinder functions follows by applying the previous paragraphs to embed the submanifold S and the disjoint finite set of orbits C into $SO(3)^N$; if \mathcal{O} is, for instance, the inverse image of a small neighborhood of the image of C in $SO(3)^N$ (whose closure is disjoint from the image of S). Then choose the extended function \tilde{h} to be zero on this small neighborhood.

The corresponding fact for $SU(2)$ -cylinder functions is a straightforward modifications of the same proof. \blacksquare

Following the lemma, we choose our Banach space of perturbations \mathcal{P}_E to be generated by a choice of a countable set of holonomy perturbations $\{\pi_i\}$, generated by (h, γ, μ) . Here

- h varies over a countable dense set of $SO(3)$ -invariant functions on $SU(2)^N$ which vanish in a neighborhood of $(\pm 1)^N$,
- γ varies over a countable dense subset of the space of N immersions of $S^1 \times D^2$ that share a germ around $\{1\} \times D^2$,
- μ is specified, and
- N varies over all positive integers.

We say that a Banach space of perturbations \mathcal{P}_E arising in this way is *sufficiently large*. The first condition will be useful in technical arguments later, so that the Hessian at fully reducible connections is the same no matter the perturbation.

The particular countable family of holonomy perturbations is inessential (other than these denseness conditions). While the particular choice of countable dense

set to choose is noncanonical, the union of any two gives another; any theorem on independence of perturbation in \mathcal{P}_E extends to show that the particular choice of \mathcal{P}_E is irrelevant.

In practice we will need to make further perturbations, but we want to do so without changing the existing functional at a finite set C of ‘acceptable’ critical orbits. We say that holonomy perturbations vanishing on a small open neighborhood \mathcal{O} of this set C *adapted* to \mathcal{O} ; the lemma shows these are in large supply. If \mathcal{P}_E is a Banach space of perturbations and π_0 is some fixed perturbation, we denote

$$\mathcal{P}_{E,\mathcal{O}} := \{\pi \in \mathcal{P}_E \mid f_\pi|_{\mathcal{O}} = f_{\pi_0}|_{\mathcal{O}}\}.$$

It is well-known that the space of flat connections is compact. The set of critical points of the unperturbed Chern-Simons functional framed flat connections, modulo even gauge; it’s equivalently described as a sort of framed projective representation variety of E . We recall the following compactness principle for critical points from [KM11b, Lemma 3.8]; their argument is unchanged for the framed moduli space, modulo even gauge.

Lemma 3.3. *Let \mathcal{P} be a Banach space of perturbations. The map*

$$F : \mathcal{P} \times \tilde{\mathcal{B}}_{E,k}^e \rightarrow \mathcal{P} \times \mathcal{T}_{k-1} \tilde{\mathcal{B}}_{E,k}^e,$$

$(\pi, [A]) \mapsto (\pi, \nabla(\text{cs} + f_\pi)([A]))$, *is proper. In particular, if $\mathfrak{C}_\bullet \subset \mathcal{P} \times \tilde{\mathcal{B}}_{E,k}^e$ denotes the subset whose fiber over $\pi \in \mathcal{P}$ is the set of critical points of $\text{cs} + f_\pi$ — that is, $\mathfrak{C}_\bullet = F^{-1}(\mathcal{P} \times \{0\})$ — then the projection $\mathfrak{C}_\bullet \rightarrow \mathcal{P}$ is proper.*

3.2. Linear analysis. It follows immediately from [Don02, Equation 2.18], modified to fit our normalization of cs , that

$$\text{cs}(A + \omega) = \text{cs}(A) + \int_Y \text{tr}(d_A \omega \wedge \omega + \frac{2}{3} \omega \wedge \omega \wedge \omega)$$

that the Hessian of the Chern-Simons functional on \mathcal{A}_E is precisely

$$*d_A : \Omega_k^1(\mathfrak{g}_E) \rightarrow \Omega_{k-1}^1(\mathfrak{g}_E).$$

(When working with Sobolev completions, the Hessian is a smooth map between tangent bundles of different regularity.) Because the tangent space at a framed connection $[A, p] \in \tilde{\mathcal{B}}_{E,k}^e$ can be decomposed as $\mathcal{K}_{A,k} \oplus \mathfrak{g}_A^\perp$ (recall that $\mathcal{K}_{A,k}$ is defined to be the Coulomb slice $\ker(d_A^*) \subset \Omega_k^1(\mathfrak{g}_E)$), this identifies the Hessian of the Chern-Simons functional *at a flat connection* as

$$\text{Hess}_A = *d_A \oplus 0 : \mathcal{K}_{A,k} \oplus \mathfrak{g}_A^\perp \rightarrow \mathcal{K}_{A,k-1} \oplus \mathfrak{g}_A^\perp.$$

We call the summand $*d_A : \mathcal{K}_{A,k} \rightarrow \mathcal{K}_{A,k-1}$ the normal Hessian, Hess_A^ν , which is the Hessian restricted to the normal space to the orbit. The Hessian is a summand of a larger elliptic operator, the extended Hessian

$$\widehat{\text{Hess}}_A : \Omega_k^0(\mathfrak{g}_E) \oplus \Omega_k^1(\mathfrak{g}_E) \rightarrow \Omega_{k-1}^0(\mathfrak{g}_E) \oplus \Omega_{k-1}^1(\mathfrak{g}_E)$$

by the matrix

$$\begin{pmatrix} 0 & -d_A^* \\ -d_A & *d_A \end{pmatrix}.$$

It is clear that $\widehat{\text{Hess}}_A^2 = \Delta_A$ on $\Omega^0 \oplus \Omega^1$ and thus $\widehat{\text{Hess}}_A$ is a self-adjoint elliptic operator. At a flat connection, where $(*d_A)^2 = 0$, the extended Hessian has the

further decomposition under

$$\Omega^1(\mathfrak{g}_E) = \mathcal{K}_{A,k} \oplus \mathcal{G}_{A,k}$$

as the orthogonal direct sum of Hess'_A on $\mathcal{K}_{A,k}$ and the signature operator S_A on $\Omega^0 \oplus \mathcal{G}_{A,k}$. The kernel of the Hessian is identified with the kernel of the Laplacian on 0- and 1-forms (and is thus invertible if the Laplacian is).

Adding a perturbing term, the Hessian of the perturbed Chern-Simons functional $\text{Hess}_{A,\pi}$ differs only by the gradient of the vector field ∇_π at $[A]$, written $D_A \nabla_\pi$. We define the perturbed extended Hessian the same way. The perturbation of $\widehat{\text{Hess}}_A$ is a compact self-adjoint operator, and so the resulting perturbed extended Hessian operator is still self-adjoint first order elliptic. At critical points of $\text{cs} + f_\pi$, using the decomposition $\Omega^0 = \mathcal{H}_A^0 \oplus \text{Im}(d_A^*)$, we see that

$$\widehat{\text{Hess}}_{A,\pi} = 0 \oplus \text{Hess}'_{A,\pi} \oplus S_A$$

on

$$\mathcal{H}_A^0 \oplus (\text{Im}(d_A^*) \oplus \mathcal{G}_{A,k}) \oplus \mathcal{K}_{A,k}.$$

Thus at critical points A of the perturbed Chern-Simons functional, $\text{Hess}'_{A,\pi}$ is a summand of a self-adjoint elliptic operator, and so it has discrete eigenvalues and the space it acts on enjoys a direct sum decomposition over these (finite-dimensional) eigenspaces.

Definition 3.1. *A critical orbit $[A, p] \in \tilde{\mathcal{B}}_{E,k}^e$ of the perturbed function $\text{cs} + f_\pi$ is called nondegenerate if $\text{Hess}'_{[A],\pi}$ is invertible, or, equivalently, the kernel of $\widehat{\text{Hess}}_{A,\pi}$ on $\Omega_k^0(\mathfrak{g}_E) \oplus \Omega_k^1(\mathfrak{g}_E)$ consists only of harmonic 0-forms.*

This terminology agrees with [Don02]; if the connection A is furthermore irreducible, and so its extended Hessian operator is invertible, Donaldson calls A *acyclic*.

3.3. Reducible critical points. Because of the linear nature of the $SO(2)$ -fixed subspace of $\tilde{\mathcal{B}}_E^e$,

$$\left(\tilde{\mathcal{B}}_E^e\right)^{SO(2)} = \bigsqcup_{\eta} \mathcal{B}_\eta,$$

it's especially easy to determine the critical points of the Chern-Simons functional. When $b_1(Y) = 0$, there is a unique flat connection in each component (corresponding to the calculation $\text{Hom}(\pi_1(Y), \mathbb{Z}) = 0$), at which the Hessian inside the $SO(2)$ -fixed subspace is simply the restriction of $*d : \Omega_k^1(Y; i\mathbb{R}) \rightarrow \Omega_{k-1}^1(Y; i\mathbb{R})$ to

$$\text{Hess}_\theta = *d : \ker(d^*)_k \rightarrow \ker(d^*)_{k-1}.$$

This depends only on the underlying metric on Y , not the connection A . Because $b_1(Y) = 0$, Hodge theory guarantees that this is an isomorphism, and so these critical points are nondegenerate in the $SO(2)$ -fixed locus.

Proposition 3.4. *Suppose Y is a Riemannian 3-manifold with $b_1(Y) = 0$, equipped with an $SO(3)$ -bundle E . There is a positive constant ε so that for any perturbation $\pi \in \mathcal{P}$ with $\|\pi\| \leq \varepsilon$, there is precisely one critical point of $\text{cs} + f_\pi$ in each component $\mathcal{B}_{\eta,k}$ of the $SO(2)$ -fixed point space, and the Hessian of each inside $\mathcal{B}_{\eta,k}$ is nondegenerate.*

Proof. The compactness principle Lemma 3.3 also applies to $\mathcal{B}_{\eta,k}$. The set of non-degenerate critical points in \mathfrak{C}_\bullet (which we call \mathfrak{C}_\bullet^*) is open, and hence the subset of degenerate points is closed; because proper maps are closed, the set of perturbations π for which some perturbed critical point $[A]$ is cut out *non-transversely* is closed. Conversely, regular perturbations (written \mathcal{P}^*) form an open set. Thus, because every reducible is cut out transversely in the fixed locus for $\pi = 0$, the same is true for $\|\pi\| \leq \varepsilon$ for some small ε . Finally, if $p : \mathfrak{C}_\bullet \rightarrow \mathcal{P}$ is the projection, the map

$$p : p^{-1}(\mathcal{P}^*) \rightarrow \mathcal{P}^*$$

is still a proper local diffeomorphism and hence the inverse function theorem guarantees that (for sufficiently small $\|\pi\|$) there is precisely one critical point of $cs + f_\pi$ in each component $\mathcal{B}_{\eta,k}$ of the $SO(2)$ -fixed point set. \blacksquare

There is a class of $SO(3)$ -bundles over 3-manifolds with $b_1(Y) > 0$ for which we can avoid reducibles entirely, called admissible bundles (sometimes ‘non-trivial admissible’ bundles). These were introduced in Floer’s work on Dehn surgery in [Flo95], and were extensively used in Kronheimer and Mrowka’s study of instanton knot homology, beginning with [KM11b]. In defining instanton Floer homology we will restrict to these two somewhat orthogonal cases: $b_1 = 0$ or E admissible.

Definition 3.2. *An $SO(3)$ -bundle over a 3-manifold Y is admissible if every lift of $w_2(E)$ to a class in $H^2(Y; \mathbb{Z})$ is non-torsion. We say that an $SO(3)$ -bundle E over Y is weakly admissible if either E is admissible or $b_1(Y) = 0$.*

Ultimately, the instanton chain complex is described in terms of critical *orbits*, not fixed points, so we give a description of these. The class of weakly admissible bundles allows for a succinct description of the class of reducible critical orbits.

Proposition 3.5. *Let Y be a Riemannian 3-manifold equipped with weakly admissible $SO(3)$ -bundle E , equipped with a perturbation π small enough that Proposition 3.4 applies. Then we have the following description of the reducible critical $SO(3)$ -orbits of $cs + f_\pi$ in $\tilde{\mathcal{B}}_{E,k}^e$.*

- If $b_1(Y) = 0$, then there is a unique critical orbit lying in each reducible component $\{z_1, z_2\} \in \text{Red}(Y, E)$ of the configuration space. In the components $\text{Red}_*(Y, E)$ that contain a fully reducible orbit, the full reducible is the critical orbit.
- If E is admissible, there are no reducible critical orbits.

Proof. First suppose $b_1(Y) = 0$. By Proposition 3.4, there is a unique critical point in each connected component of the $SO(2)$ -fixed subspace. The components of the $SO(2)$ -fixed subspaces were enumerated in Corollary 1.3.

In general, a critical orbit of cs in the component $\{z_1, z_2\} \in \text{Red}(Y, E)$ is a gauge equivalence class of flat connection respecting the isomorphism

$$E \cong i\mathbb{R} \oplus \eta_1 \otimes \eta_2^{-1}.$$

In particular, $\eta_1 \otimes \eta_2^{-1}$ would support a flat connection, and thus have torsion first Chern class. Because $c = z_1 + z_2$, this implies that $c - 2z_2$ is a torsion cohomology class; as this also represents $w_2(E)$, this implies that E is not admissible. Because E supports no reducible critical points for the unperturbed Chern-Simons functional, the same is true of all sufficiently small perturbations. \blacksquare

3.4. Transversality for critical points. Before discussing the 4-dimensional case, where we will assume that the limiting connections are nondegenerate, we should verify that this situation is achievable! This leads us to the first transversality theorem of this paper; its proof is a model for the rest of the transversality results we will need. It essentially follows the corresponding proof in [AB96].

Theorem 3.6. *Let Y be a Riemannian 3-manifold, equipped with an $SO(3)$ -bundle E . Suppose either that $b_1(Y) = 0$ or that E is admissible. Given a sufficiently large Banach space of perturbations \mathcal{P} , there is a dense open set of $\pi \in \mathcal{P}$ for which there are finitely many critical orbits of the perturbed Chern-Simons functional $\tilde{\mathcal{B}}_{E,k}^e$, on each of which the normal Hessian is invertible.*

The proof proceeds by achieving transversality in each locus inductively. In the course of achieving transversality over the reducible locus, we need a simple functional analysis lemma.

Lemma 3.7. *Let H be a Hilbert space, equipped with a densely defined self-adjoint Fredholm operator $T : H \rightarrow H$; write $H_0 = \ker(T)$ and $H_1 = \text{im}(T)$ for the associated orthogonal splitting of H . Let $K : H \rightarrow H$ be a bounded self-adjoint operator with $K_{00} : H_0 \rightarrow H_0$ injective. Then for sufficiently small positive ε , the map $T + \varepsilon K$ is injective; in particular, an isomorphism.*

Proof. If $D = H_0 \oplus D_1$ is the domain of T , then because $T_{11} : D_1 \rightarrow H_1$ is an isomorphism, we have a bound $\|T_{11}x_1\| \geq C_T\|x_1\|$ for some $C_T > 0$ (where the norms are defined on the appropriate domains). Similarly we have $\|K_{00}x_0\| \geq C_K\|x_0\|$ for some $C_K > 0$. Now suppose $(T + \varepsilon K)(x_0 + x_1) = 0$ for some $x = (x_0, x_1) \in H$. This means precisely that

$$\begin{aligned} K_{00}x_0 + K_{10}x_1 &= 0, \\ T_{11}x_1 + \varepsilon(K_{01}x_0 + K_{11}x_1) &= 0. \end{aligned}$$

These two equations give us the bounds

$$\begin{aligned} C_K\|x_0\| &\leq \|K_{00}x_0\| = \|K_{10}x_1\| \leq \|K\|\|x_1\|, \\ C_T\|x_1\| &\leq \|T_{11}x_1\| = \varepsilon\|K_{01}x_0 + K_{11}x_1\| \leq \varepsilon\|K\|(\|x_0\| + \|x_1\|). \end{aligned}$$

Combining these two inequalities we then have

$$C_T\|x_1\| \leq \varepsilon\|K\|(1 + \|K\|/C_K)\|x_1\|.$$

As soon as

$$\varepsilon < \frac{C_T}{\|K\|(1 + \|K\|/C_K)},$$

this implies that $x_1 = 0$, and hence that $x_0 = 0$. ■

Proof of Theorem 3.6. By the discussion in the proof of Proposition 3.4, there is an open subset U of \mathcal{P} so that all reducibles are cut out transversely in their own locus. For $\pi \in U$, there are thus a finite number of reducible π -critical orbits. Recall that cs and $SU(2)$ -holonomy perturbations are invariant under the action of $SO(3)$, but not of $H^1(Y; \mathbb{Z}/2)$.

The normal Hessian to an $SO(2)$ -reducible fixed point might have nontrivial kernel, a finite-dimensional subspace of the tangent space.

Enumerate an element of each $SO(2)$ -reducible critical orbit that is not fully reducible as $[A_i]$, and let D_i denote a small disc inside the kernel of $\text{Hess}_{A_i}^\nu = Z_i$.

The manifold D_i is an $SO(2)$ -manifold; we may identify the orbit through D_i as a vector bundle over S^2 with generic fiber D_i ; call this neighborhood V_i . An $SO(3)$ -invariant function on V_i is the same thing as an $SO(2)$ -invariant function on D_i . At an $SO(2)$ -reducible corresponding to a reduction $E \cong \mathbb{R} \oplus \zeta$, the kernel of this operator is $H^1(Y; \zeta)$. The $SO(2)$ -action is by scalar multiplication on this complex vector space. The restriction of some $SO(2)$ -invariant nondegenerate quadratic function q on Z_i gives a smooth function on D_i with Hessian at the origin equal to the identity. Note that this implies that there are an even number of positive and negative eigenvalues.

We will use holonomy perturbations approximating this smooth function on D_i to correct the fact that the Hessian has kernel. More precisely, if γ_j is a sequence of curves for which $\text{Hol}_A(\gamma)$ gives an equivariant embedding of the neighborhoods V_i into $SU(2)^N$, we choose a function \tilde{h} on $SU(2)^N$ which agrees in a neighborhood of each A_i with $v \mapsto q(v)$ (identifying a neighborhood of each A_i with the tangent space at that point). Now pick a sequence of holonomy perturbations π_n arising from approximations of \tilde{h} and γ_j . For sufficiently large n , the restriction of $D_{A_i} \nabla \pi_n$ to $\ker(\text{Hess}_{A_i})$ is very close to q .

If we decompose $\mathcal{K}_{A_i, k}$ as the direct sum of Z_i and its orthogonal complement, then one may write $\text{Hess}_{A_i, \pi}$ in block-matrix form with respect to this splitting; the component $Z_i \rightarrow Z_i$ is the Hessian restricted to the submanifold D_i at A_i . Therefore, we can apply Lemma 3.7 to $T = \text{Hess}_{A_i, \pi}$ and $K = D_{A_i} \nabla \pi_n$ for large n to see that $\pi + \varepsilon \pi_n$ is a regular perturbation at reducible connections for sufficiently large n and sufficiently small ε . A similar argument works for the fully reducible connections, now using a nondegenerate $SO(3)$ -invariant quadratic form.

Similarly to Proposition 3.4, nondegeneracy for reducible critical points is an open condition⁶ in perturbations \mathcal{P} , and we've just seen that it is also a dense condition.

Now we have reduced ourselves to furthermore showing that perturbations for which the irreducible critical points are nondegenerate are dense in $\mathcal{P}_{\leq \varepsilon}$. This is an application of the Sard-Smale theorem, using that $\widehat{\nabla}_\pi(\mathbf{A})$ separates points and tangent vectors as π varies over \mathcal{P} . (For the irreducible case, see [Flo88] and [Don02, Section 5.5.1].) The open-ness is slightly more delicate, and uses that nondegenerate critical points are isolated; therefore, for $\pi \in U$, reducible critical points are isolated. This implies that the projection $p^{-1}(U)^* \rightarrow U$ is proper, where the asterisk denotes that we restrict to irreducible connections. This is enough to conclude, as in Proposition 3.4. \blacksquare

Remark 3.1. There is not much additional difficulty in proving this theorem for $SO(3)$ -holonomy perturbations, which are invariant under the $H^1(Y; \mathbb{Z}/2)$ action. However, one needs to do more bookkeeping at the other kinds of reducibles; those with stabilizer $\mathbb{Z}/2$ are the most difficult, being labelled by their Euler classes in twisted cohomology.

Remark 3.2. In the above proof, we only used the assumption that E is weakly admissible — that is, either $b_1(Y) = 0$ or E is admissible — to ensure that the reducible flat connections are cut out transversely in the reducible locus. This is

⁶Note that 3.4 shows that nondegeneracy *in the reducible locus* is an open condition, and here we want nondegeneracy in $\tilde{\mathcal{B}}_E^e$; however, the proof requires only trivial modifications.

not a serious issue: one may easily first choose a perturbation for which the reducible π -flat connections *are* cut out transversely in the reducible locus. However, transversality for the moduli spaces of flowlines is much more restrictive, and forces the weak admissibility assumption on us — if $b_1(Y) > 0$ and E admits reducible connections, then there will be trajectories which have higher expected dimension in the reducible locus than the irreducible locus. That in mind, we see no harm in making the assumption a little early.

Note that the critical set of $cs + f_\pi$ is canonically identified with that in any Hilbert manifold completion of lower regularity by an elliptic regularity argument; we are not being remiss in leaving the regularity k out from our notation \mathfrak{C}_π for the critical sets.

3.5. Signature data of reducible critical points. Suppose Y is a rational homology sphere. Fix a lift c of $w_2(E)$ to an integral cohomology class; then the connected components $\text{Red}(Y, E)$ of the configurations of reducible connections on E are in bijection with pairs $\{z_1, z_2\} \subset H^2(Y; \mathbb{Z})$ with $z_1 + z_2 = c$. (See Definition 1.1 for the definition of this set and a discussion of its enumeration.)

The component labeled by $\{z_1, z_2\}$ corresponds to equivalence classes of framed connections (with fixed determinant connection) on $\tilde{E} = \eta_1 \oplus \eta_2$ respecting the direct sum decomposition, where $c_1(\eta_i) = z_i$ and $\det \tilde{E} = \lambda$, with $c_1(\lambda) = c$. In

$$E \cong \mathbb{R} \oplus (\eta_1 \otimes \eta_2^{-1})$$

these again correspond to connections respecting the direct sum decomposition. In particular, if A represents the unique class of flat connection on E that preserves the splitting, we may identify $\widehat{\text{Hess}}_A(\mathfrak{g}_E)$ with the direct sum

$$\widehat{\text{Hess}}_\theta(\mathbb{R}) \oplus \widehat{\text{Hess}}_\rho(\eta_1 \otimes \eta_2^{-1}),$$

where ρ is the flat connection given by restricting A to $\eta_1 \otimes \eta_2^{-1}$. (and thus, the unique equivalence class of flat connection on this oriented 2-plane bundle). In particular, because the space of harmonic forms for a flat connection is isomorphic to the cohomology groups for the local system of coefficients defined by its holonomy representation, we have

$$\ker \widehat{\text{Hess}}_A^\nu \cong H^1(Y; \rho).$$

Our next goal is to obtain a more accessible calculation of this space.

Lemma 3.8. *Let Y be a rational homology sphere and fix an abelian representation $\rho : \pi_1 Y \rightarrow U(1)$. If \tilde{Y} is the universal abelian cover of Y — the closed 3-manifold arising as the covering space of Y corresponding to $[\pi_1 Y, \pi_1 Y]$ — then we may identify $H^1(Y; \rho)$ with the subspace of $H^1(\tilde{Y}; \mathbb{C})$ on which $gx = \rho(g)x$ for all $g \in H_1(Y)$ (the ‘ ρ -eigenspace’). In particular, if $H^1(\tilde{Y}; \mathbb{C}) = 0$, then $H^1(Y; \rho) = 0$ for all local coefficient systems ρ .*

Proof. Let $C_*(Y; \mathbb{C})$ be the cellular chain complex of some finite CW-decomposition of Y , and $C_*(\tilde{Y}; \mathbb{C})$ the cellular chain complex with $H_1(Y)$ -action. The cohomology groups $H^*(Y; \rho)$ are (essentially by definition) isomorphic to the homology groups of the cochain complex

$$\text{Hom}_{H_1 Y} \left(C_*(\tilde{Y}; \mathbb{C}), \mathbb{C} \right),$$

where $H_1(Y)$ acts on \mathbb{C} via ρ . By definition, this cochain complex is isomorphic to the ρ -eigenspace of $C^*(\tilde{Y}; \mathbb{C})$. Because the chain complex $C^*(\tilde{Y}; \mathbb{C})$ splits as a direct

sum of its eigenspaces, labelled by homomorphisms $H_1(Y) \rightarrow U(1)$, we see that the cohomology $H^*(\tilde{Y}; \mathbb{C})$ splits as a direct sum of the direct sum of its eigenspaces $H_\rho^*(\tilde{Y}; \mathbb{C})$, each of which is isomorphic to the homology of the eigenspace $C_\rho^*(\tilde{Y}; \mathbb{C})$, which is then isomorphic to $H^*(Y; \rho)$. \blacksquare

Proposition 3.9. *Let Y be a rational homology sphere equipped with $SO(3)$ -bundle E . For $\varepsilon > 0$ sufficiently small, the set of perturbations π with nondegenerate reducible critical points and $\|\pi\| < \varepsilon$ is an open set, which decomposes as the disjoint union of nonempty open subspaces, each labelled by a function*

$$N_\pi : \text{Red}_{SO(2)}(Y, E) \rightarrow \mathbb{Z}_{\geq 0}$$

with $N_\pi(\{z_1, z_2\}) \leq \dim_{\mathbb{C}} H^1(Y; \eta_1 \otimes \eta_2^{-1})$.

Proof. That the set of such perturbations π is open is a mild extension of Proposition 3.4.

Given any ε -small perturbation π , let $\beta : [0, 1] \rightarrow [0, 1]$ be a function equal to 1 near 0 and equal to 0 near 1, and consider the path $\pi(t) = \beta(t)\pi$ from π to 0; we have $\|\pi(t)\| < \varepsilon$ for all t . If

$$r = \{z_1, z_2\} \in \text{Red}_{SO(2)}(Y, E)$$

denotes a connected component of the reducible subset of the configuration space (here $z_1 + z_2$ is a fixed integral lift of $w_2 E$ and $z_1 \neq z_2$), by Proposition 3.4, there is a unique $\pi(t)$ -critical orbit in the component labeled by r , and the proof shows it varies smoothly in t . Choose a particular path $A(t)$ of framed connections so that $[A(t)]$ lies in the $\pi(t)$ -critical orbit in the reducible component labelled by r . We obtain an associated continuous path of self-adjoint Fredholm operators

$$\widehat{\text{Hess}}_{A(t), \pi(t)} : \Omega_k^0 \oplus \Omega_k^1 \rightarrow \Omega_{k-1}^0 \oplus \Omega_{k-1}^1.$$

There is an associated spectral flow, $\text{sf} \left(\widehat{\text{Hess}}_{A(t), \pi(t)} \right) \in \mathbb{Z}$, defined as the intersection number of the eigenvalues with the line $\lambda = -c$ for sufficiently small positive c ; it is essentially the number of eigenvalues that go from negative to nonnegative, counted with sign. (This definition allows for the operators at $t = 0, 1$ to have kernel, as opposed to intersecting with the line $\lambda = 0$.) In fact, $\widehat{\text{Hess}}_{A(t), \pi(t)}$ splits as a direct sum corresponding to the reductions $\mathbb{R} \oplus (\eta_1 \otimes \eta_2^{-1})$ (here $c_1(\eta_i) = z_i$); the operator in the first component is constant, and in the second component is complex linear, so this spectral flow is actually an even integer. The spectral flow of this path depends only on the ending perturbation π , so we define

$$N_\pi(r) = \frac{1}{2} \text{sf} \left(\widehat{\text{Hess}}_{A(t), \pi(t)} \right).$$

The kernel of $\widehat{\text{Hess}}_{A(t), \pi(t)}$ splits as the direct sum of harmonic 0-forms and a space of 1-forms. Note that $A(t)$ is $SO(2)$ -reducible for all t (and never fully reducible), which implies that the dimension of the space of harmonic 0-forms is constant in t , and so does not contribute to the spectral flow.

The spectral flow depends on π continuously (that is, sufficiently close perturbations give the same integer) when $\widehat{\text{Hess}}_{A(0), \pi}$ has no kernel other than the harmonic 0-forms. To say that $\widehat{\text{Hess}}_{A(0), \pi}$ has no kernel other than harmonic 0-forms is precisely to say that π is regular at the critical point $A(0)$. Therefore, the spectral

flow to 0 is a continuous function on the open set of interest, the space of regular perturbations.

As long as ε is chosen small enough, the only eigenvalues that are close enough to 0 to cross above or below it are the zero eigenvalues in $\ker \widehat{\text{Hess}}_{A(1)}$, the unperturbed flat connection in this reducible component. The eigenvalues corresponding to harmonic 0-forms stay constant, so we only need to think about the other zero eigenvalues.

As we are counting the number of eigenvalues that go from negative to nonnegative as we pass from $\widehat{\text{Hess}}_{A(0),\pi}$ to the unperturbed extended Hessian, the spectral flow must be nonnegative (said another way, the only relevant eigenvalues are the zero eigenvalues at $t = 1$, and those can only contribute to spectral flow by going down; reversing the direction, the only way to contribute to the spectral flow is to have an eigenvalue go upwards to 0 as we approach $t \rightarrow 1$). Because there are

$$\dim_{\mathbb{R}} H^1(Y; \eta_1 \otimes \eta_2^{-1})$$

of these eigenvalues, this spectral flow is bounded above by that number.

Thus we have a continuous map from the space of reducible-regular perturbations to the discrete space

$$\text{Map}(\text{Red}_{SO(2)}(Y, E), \mathbb{Z}_{\geq 0})$$

satisfying the stated bound; this decomposes the open set of reducible-regular perturbations as the disjoint union of open subsets labelled by these functions

$$N_{\pi} : \text{Red}(Y, E) \rightarrow \mathbb{Z}_{\geq 0}$$

with upper bounds as in the statement of the proposition.

Given any function N satisfying these bounds, we may construct a small perturbation π so that N arises as the spectral-flow function N_{π} . Choose an S^1 -invariant quadratic function with $2N_{\pi}(r)$ positive eigenvalues on the space $\ker \text{Hess}'_A$, where A is the unique critical point in the component labelled by r ; note that on any given complex line, this function takes the form cr^2 where r is the radius, for some constant c .

Pick a holonomy perturbation that well approximates that quadratic function, thinking of $\ker \text{Hess}'_A$ as an S^1 -submanifold of $\tilde{\mathcal{B}}_{E,k}^e$ by exponentiating. Because holonomy perturbations are invariant under the full gauge group, the same quadratic function is introduced at the reducible labelled $x \cdot r$. Doing this as we vary over $\text{Red}_{SO(2)}(Y, E)$ gives us a perturbation which restricts to a nondegenerate quadratic function, with the appropriate number of eigenvalues at each reducible critical point. Then if we choose c sufficiently small, the perturbation $c \cdot \pi$ will be nondegenerate at each reducible and have $c \cdot \|\pi\| < \epsilon$.

Note that there is no analogous discussion in the components $\text{Red}_{*}(Y, E)$ containing a fully reducible point, because this point is the unique critical point in that component and is nondegenerate for all small π : the spectral flow is trivial. ■

Remark 3.3. It is likely the case that the open set labelled by $N_{\pi} : \text{Red}_{SO(2)}(Y, E) \rightarrow \mathbb{Z}_{\geq 0}$ is *connected*, but we will not find this necessary, so do not prove it here. In a sense similar to Remark 6.3, the ‘space of regular perturbations with fixed signature data’ is contractible. This will not be true in the 4-dimensional setting.

Remark 3.4. If one works with $SO(3)$ -holonomy perturbations, which are invariant under $H^1(Y; \mathbb{Z}/2)$, one needs to do this bookkeeping for additional types of reducibles, each with different demands on the signature data.

The essential count turns out not to be the number of negative eigenvalues, but rather the *signature* of the corresponding real vector space. As long as we remember the dimension of the underlying vector space, this is equivalent information.

Definition 3.3. *Let (Y, E) be an $SO(3)$ -bundle over a rational homology sphere, and fix a complex line bundle λ with $w_2(E) = c_1(\lambda) \pmod{2}$, which induces a bijection of $\text{Red}_{SO(2)}(Y, E)$ with pairs $\{z_1, z_2\} \subset H^2(Y; \mathbb{Z})$ with $z_1 + z_2 = c_1(\lambda)$; the choice of λ is essentially irrelevant. We write η_1 and η_2 for the corresponding complex line bundles.*

We say that a signature datum on (Y, E) is a choice of function

$$\sigma : \text{Red}_{SO(2)}(Y, E) \rightarrow 2\mathbb{Z}$$

with

$$|\sigma_\pi(\{z_1, z_2\})| \leq \dim_{\mathbb{R}} H^1(Y; \eta_1 \otimes \eta_2^{-1})$$

and

$$\sigma_\pi(\{z_1, z_2\}) \equiv \dim_{\mathbb{R}} H^1(Y; \eta_1 \otimes \eta_2^{-1}) \pmod{4}.$$

Given a small regular perturbation π , the associated signature datum σ_π is

$$\sigma_\pi(\{z_1, z_2\}) = \dim_{\mathbb{R}} H^1(Y; \eta_1 \otimes \eta_2^{-1}) - 4N_\pi(\{z_1, z_2\}).$$

The set of signature data is denoted $\sigma(Y, E)$.

The name ‘signature datum’ refers to the fact that σ essentially chooses the symmetric bilinear form on each complex vector space $\ker \text{Hess}'_A$ that we perturb in the direction of (which are determined by their signature). The constraints are precisely those that the signature of an S^1 -invariant bilinear form on a complex vector space must satisfy.

Remark 3.5. If the universal abelian cover \tilde{Y} of Y (corresponding to the subgroup $[\pi_1 Y, \pi_1 Y] \subset \pi_1 Y$) is a rational homology sphere, then $|\sigma(Y, E)| = 1$ for all E . If the cover \tilde{Y}' corresponding to the subgroup $[\pi_1 Y, \pi_1 Y]^2$ is a rational homology sphere, then $|\sigma(Y, \text{triv})| = 1$.

4. MODULI SPACES OF INSTANTONS

4.1. Moduli spaces for cylinders and cobordisms. In Definition 2.1, we defined the configuration space of framed connections on a bundle π^*E over a cylinder with specified limits as a quotient

$$\tilde{\mathcal{B}}_{k,\delta}^e(\alpha_-, \alpha_+) := \tilde{\mathcal{A}}_{k,\delta}(A_-, A_+)/\mathcal{G}_{k+1,\delta}^{e,h}(A_-, A_+).$$

The trivial vector bundle over this configuration space with fiber $\Omega_{k-1,\delta}^{2,+}(\text{End}(\pi^*E))$ carries a compatible linear action of $\mathcal{G}_{\pi^*E,k+1,\delta}^{e,h}(A_-, A_+)$, given by the action of the gauge group on $\pi^*\mathfrak{g}_E$; it therefore descends to a vector bundle $\mathcal{S}_{k-1,\delta}$ over the quotient (\mathcal{S} for ‘self-dual’) with fibers isomorphic to $\Omega_{k-1,\delta}^{2,+}$. The smoothly varying section $\mathbf{A} \mapsto F_{\mathbf{A}}^+$, which is equivariant under the gauge group action on the left and invariant under the $SO(3)$ action on the right, descends to an $SO(3)$ -invariant smooth section of $\mathcal{S}_{k-1,\delta}$.

Definition 4.1. *The (unperturbed) (even) moduli space of framed instantons on the bundle $\pi^*E \rightarrow \mathbb{R} \times Y$ with limits α_-, α_+ is the $SO(3)$ -invariant subspace*

$$\tilde{\mathcal{M}}_E(\alpha_-, \alpha_+) \subset \tilde{\mathcal{B}}_{k,\delta}^e(\alpha_-, \alpha_+),$$

consisting of gauge equivalence classes of framed connections $(\mathbf{A}, p) \in \tilde{\mathcal{A}}_{k,\delta}(A_-, A_+)$ with $F_{\mathbf{A}}^+ = 0$, where A_{\pm} are connections whose gauge equivalence class is α_{\pm} . Equivalently, it is the zero set of the section $F_{\mathbf{A}}^+$ of $\mathcal{S}_{k-1,\delta}$. There are equivariant endpoint maps $\text{ev}_{\pm} : \tilde{\mathcal{M}}_E(\alpha_-, \alpha_+) \rightarrow \alpha_{\pm}$. We denote the quotient as

$$\mathcal{M}_E(\alpha_-, \alpha_+) := \tilde{\mathcal{M}}_E(\alpha_-, \alpha_+)/SO(3).$$

In our notation for the moduli space, we drop the symbols k, δ , and e . The first two are unnecessary, because as long as the perturbation π is of regularity L_j^2 , the moduli spaces for all $1 < k \leq j$ are canonically identified by an elliptic regularity result. The requirement $k > 1$ is necessary to define the moduli spaces, as we require gauge transformations in L_{k+1}^2 to have continuous point-evaluation maps, which requires $k+1 > 2$. Similarly $\delta > 0$ is irrelevant so long as it is taken less than the absolute value of any nonzero eigenvalue of any of the eigenvalues of the limiting operator $\widehat{\text{Hess}}_{\alpha}$. The even moduli space is not the same as its quotient by the full gauge group, but we will never work with non-even moduli spaces, so we drop the superscript anyway.

Note that if the A_{\pm} are not flat connections, there are no unperturbed instantons with those limits and L^2 curvature.

If π is a regular perturbation, so that all critical points of $cs + f_{\pi}$ are nondegenerate, we instead work with perturbed moduli spaces of instantons; we will need to do this to guarantee these moduli spaces are smooth manifolds.

Write \mathcal{P}_E for the Banach space of perturbations of Section 3.1. Given $\pi \in \mathcal{P}_E$ corresponding to $f_{\pi} : \tilde{\mathcal{B}}_E^e \rightarrow \mathbb{R}$, the natural perturbation on the connections \mathbf{A} over $\mathbb{R} \times Y$ is

$$(\widehat{\nabla}_{\pi})(\mathbf{A}) = (dt \wedge (\nabla f_{\pi})(i_t^* \mathbf{A}))^+,$$

where i_t^* is pullback to the slice $\{t\} \times Y$.

The operator $\widehat{\nabla}_{\pi}$ gives a well-defined, smoothly varying section of the trivial bundle

$$\tilde{\mathcal{A}}_{k,\delta}(A_-, A_+) \times \Omega_{k-1,\delta}^{2,+}(\text{End}(\pi^*E)) \rightarrow \tilde{\mathcal{A}}_{k,\delta}(A_-, A_+),$$

equivariant under the action of $\mathcal{G}_{k+1,\delta}^{e,h}(A_-, A_+)$, and the rest of Theorem 3.1 carries over with the obvious modifications (including a version with dependence on π).

In particular, because the above perturbations are invariant under the action of the full 4-dimensional gauge group $\mathcal{G}_{k+1,\delta}^h(A_-, A_+)$, as well as the $SO(3)$ -action, $F_{\mathbf{A}}^+ + \widehat{\nabla}_{\pi}(\mathbf{A})$ defines a smooth section

$$(2) \quad \mathcal{P} \times \widetilde{\mathcal{B}}_{k,\delta}^e(\alpha_-, \alpha_+) \rightarrow \mathcal{P} \times \mathcal{S}_{k-1,\delta};$$

the fact that $F_{\mathbf{A}}^+ + \widehat{\nabla}_{\pi}(\mathbf{A}) \in \Omega_{k-1,\delta}^{2,+}(\text{End}(\pi^*E))$ follows from the assumption that \mathbf{A}_t decays exponentially on the ends to connections with $*F_{A_{\pm}} = -\widehat{\nabla}_{\pi}(A_{\pm})$.

As in Definition 4.1, we say that the *perturbed moduli space* of framed instantons for the perturbation π , written $\widetilde{\mathcal{M}}_{E,\pi}(\alpha_-, \alpha_+)$, is the set of gauge equivalence classes of framed connections with $F_{\mathbf{A}}^+ = -\widehat{\nabla}_{\pi}(\mathbf{A})$. As in the unperturbed case, \mathbf{A} can only have L^2 curvature if $*F_{A_{\pm}} = -\nabla_{\pi}(A_{\pm})$, and so the limiting connections are critical points for $\text{cs} + f_{\pi}$. We call these π -flat connections.

Using the gradient equation, we can define two useful notions of energy for an instanton on the cylinder.

Proposition 4.1. *For a perturbed connection \mathbf{A} limiting to π -flat connections on the ends, we call the expression*

$$\int \|F_{\mathbf{A}} + *_3 \nabla_{\pi}(\mathbf{A}(t))\|^2$$

the analytic energy of \mathbf{A} . If \mathbf{A} is a π -perturbed instanton, the analytic energy of \mathbf{A} depends only on the endpoints α_-, α_+ of \mathbf{A} , and the homotopy class that \mathbf{A} traces out, written $z \in \pi_1(\widetilde{\mathcal{B}}_E, \alpha_-, \alpha_+)$. If A_+ and A_- are the limiting connections of \mathbf{A} , it is equal to $2(\text{cs}_{\pi}(\alpha_-) - 2\text{cs}_{\pi}(\alpha_+))$.

Correspondingly, given a pair of π -flat connections α_-, α_+ and a homotopy class z between them, we call

$$\mathcal{E}_z^{\pi}(\alpha_-, \alpha_+) = 2(\text{cs}_{\pi}(\alpha_-) - \text{cs}_{\pi}(\alpha_+))$$

the topological energy of (α_-, α_+, z) .

Whenever there is an instanton going from α_- to α_+ in the homotopy class z , we have $\mathcal{E}_z^{\pi}(\alpha_-, \alpha_+) \geq 0$ with equality if and only if \mathbf{A} is constant. If z_i is a homotopy class from α_i to α_{i+1} , then

$$\mathcal{E}_{z_1 * z_2}^{\pi}(\alpha_1, \alpha_3) = \mathcal{E}_{z_1}^{\pi}(\alpha_1, \alpha_2) + \mathcal{E}_{z_2}^{\pi}(\alpha_2, \alpha_3).$$

If $1 \in \pi_1(\widetilde{\mathcal{B}}_E)$ is the positive generator, then $\mathcal{E}_{z+1}^{\pi}(\alpha_-, \alpha_+) = 64\pi^2 + \mathcal{E}_z^{\pi}(\alpha_-, \alpha_+)$.

Proof. The given expression is gauge invariant, so we may assume \mathbf{A} is in temporal gauge and write $\mathbf{A} = d/dt + \mathbf{A}(t)$. The curvature is $F_{\mathbf{A}} = F_{\mathbf{A}(t)} + dt \wedge A'(t)$, and so we may write this integral as

$$\int \|A'(t)\|^2 + \|F_{\mathbf{A}(t)} + \nabla_{\pi}(\mathbf{A}(t))\|^2 dt.$$

Because

$$F_{\mathbf{A}} + dt \wedge \nabla_{\pi}(\mathbf{A}(t)) = F_{\mathbf{A}} + (*_3 \nabla_{\pi}(\mathbf{A}(t)))$$

is anti-self-dual, and $*(dt \wedge \omega) = *_3 \omega$ for a 1-form ω pulled back from Y , we see that

$$A'(t) = - *_3 F_{\mathbf{A}(t)} - \nabla_{\pi}(\mathbf{A}(t)) = -\nabla(\text{cs} + f_{\pi})(\mathbf{A}(t)).$$

Using that $\mathbf{A}(t)$ is a gradient flowline, we may rewrite the integral as

$$\begin{aligned}
& \int \|A'(t) + \nabla(\text{cs} + f_\pi)(\mathbf{A}(t))\|^2 - 2\langle A'(t), \nabla(\text{cs} + f_\pi)(\mathbf{A}(t)) \rangle dt \\
&= -2 \int \langle A'(t), \nabla(\text{cs} + f_\pi)(\mathbf{A}(t)) \rangle dt \\
&= -2 \int d(\text{cs} + f_\pi)(A'(t)) \\
&= -2((\text{cs} + f_\pi)(A_+) - (\text{cs} + f_\pi)(A_-)) \in \mathbb{R},
\end{aligned}$$

the desired result. Now it is clear that this only depends on the connected component of the connection \mathbf{A} . (Here we are using the path to pin down the real lift of this difference of Chern-Simons values.)

That $\mathcal{E}_0^\pi(\alpha, \alpha) = 0$ is clear from the existence of the constant trajectory. Additivity is clear by picking a connections \mathbf{A}_i for each that are constant sufficiently far down the cylindrical ends, and gluing those together; the difference in $\text{cs} + f_\pi$ is additive.

Lastly, we need to determine $\mathcal{E}_1^\pi(\alpha, \alpha)$; by definition, this is the same as determining

$$2(\text{cs} + f_\pi)(g(A)) - 2(\text{cs} + f_\pi)(A)$$

for a gauge transformation g generating $\pi_0 \mathcal{G}_E^e$. Because f_π is invariant under the full gauge group, we are only asking to determine the difference in the Chern-Simons functional. Pick a trajectory \mathbf{A} going from A to $g(A)$, constant near the ends. By definition, this is given by computing $2 \int_{I \times Y} \text{Tr}(F_{\mathbf{A}}^2)$. Because A differs from $g(A)$ by a gauge transformation, we may glue these together to get a connection on an $SO(3)$ -bundle \mathbf{E} over $S^1 \times Y$, the bundle \mathbf{E} being the mapping torus of the automorphism g ; noting that $F_{\mathbf{A}}^2$ is zero near the ends, this implies we may compute the curvature integral just as well over this closed 4-manifold, where it is equal to $-2 \cdot 8\pi^2 p_1(\mathbf{E})$. We choose the positive generator of $\pi_1(\tilde{\mathcal{B}}_E^e)$ to be the one that makes this integral positive; it remains to compute it. This is computed more generally in [KM11b, Equation (24)], and in this case the minimal $p_1(\mathbf{E})$ is 4, as desired. Because of the slight difference that we use specifically even gauge transformations, we provide a short proof, given in the lemma that follows. \blacksquare

Lemma 4.2. *Let E be an $SO(3)$ -bundle over a 3-manifold Y . Given an even gauge transformation $g \in \Gamma(\text{Aut}(E))$, the minimal first Pontryagin class of the mapping torus of g on $S^1 \times Y$ is 4.*

Proof. The bundle \mathbf{E} has

$$w_2 \mathbf{E} \in H^2(S^1 \times Y; \mathbb{Z}/2) \cong H^2(Y; \mathbb{Z}/2) \oplus H^1(Y; \mathbb{Z}/2).$$

The first factor of $w_2 \mathbf{E}$ is precisely the second Stiefel-Whitney class of its restriction to Y , which is $w_2 E$. The second term of $w_2 \mathbf{E}$ is precisely the obstruction class $o(g) \in H^1(Y; \mathbb{Z}/2)$ for lifting a section of $\text{Aut}(E)$ to a section of $\widetilde{\text{Aut}}(E)$.

To see this, pick a loop γ in Y and consider the corresponding torus in $S^1 \times Y$. The map $\gamma \mapsto w_2(\mathbf{E})|_{S^1 \times \gamma}$ defines a homomorphism $H_1(Y; \mathbb{Z}) \rightarrow \mathbb{Z}/2$. If $o(g) = 0$, then the restriction of g to γ lifts to $\widetilde{\text{Aut}}(E)$, and so is homotopic to the identity (because $\pi_1 \widetilde{\text{Aut}}(E_x) = 0$), so the corresponding bundle on the torus $S^1 \times \gamma$ is trivial.

Conversely, every bundle over $S^1 \times Y$ may be constructed as the mapping torus of a bundle over Y , and both $w_2(E)$ and $o(g)$ of the bundle and mapping are

determined by $w_2\mathbf{E}$. Given fixed $w_2(\mathbf{E}) = w_2(E) \oplus o(g)$, the Dold-Whitney theorem identifies the possible values of $p_1(\mathbf{E})$ as those elements $p \in H^4(S^1 \times Y; \mathbb{Z})$ with

$$p \equiv (w_2\mathbf{E})^2 \pmod{4},$$

where the mod 4 indicates we use the *Pontryagin square* on even-dimensional cohomology $H^2(X; \mathbb{Z}/2) \rightarrow H^4(X; \mathbb{Z}/4)$. Expanding this, we obtain

$$p_1(\mathbf{E}) = [S^1] \cdot 2o(g) \cdot w_2(E) \in H^4(S^1 \times Y; \mathbb{Z}/4).$$

In particular, as long as $o = 0$, the only condition is $p_1(\mathbf{E}) \equiv 0 \pmod{4}$. ■

Corollary 4.3. *Given a 3-manifold Y equipped with $SO(3)$ -bundle E and a regular perturbation π , if α is a π -critical orbit, the constant connection is the unique π -perturbed instanton \mathbf{A} in the component labelled by the trivial homotopy class $0 \in \pi_1(\hat{\mathcal{B}}_E^e, \alpha)$.*

Proof. The constant connection at a π -flat connection is always a solution to the perturbed ASD equations. Because $\mathcal{E}_0^\pi(\alpha, \alpha) = 0$, and the only connections with analytic energy equal to zero are constant, so is any solution of the perturbed ASD equations in this component of the space of connections. ■

For a Riemannian manifold W , with two cylindrical ends, orientably isometric to $(-\infty, 0] \times Y_1$ and $[0, \infty) \times Y_2$, written as $W : Y_1 \rightarrow Y_2$, we can no longer consider constant perturbations (pulled back from Y). Further, we need to be able to interpolate between perturbations, and have enough available to prove transversality results. As before, the precise formulation of the perturbations here is similar to that in [KM11b, Section 3.2].

Fix once and for all a smooth function $\beta_0 : [0, \infty) \rightarrow [0, 1]$ with $\beta_0(x) = 0$ for $x \leq 1$ and $\beta_0(x) = 1$ for $x \geq 2$.

Given a perturbation $\pi \in \mathcal{P}_{E_2}$, we may define the associated 4-dimensional perturbation $\hat{\nabla}_\pi$ on $[0, \infty) \times Y_2$ as

$$(\hat{\nabla}_\pi)(\mathbf{A}) = \beta_0(t)(dt \wedge \nabla_\pi(\mathbf{A}(t)))^+,$$

and identically zero elsewhere, and similarly for a perturbation in \mathcal{P}_{E_1} .

An *end perturbation* on W is labelled by is

$$(\pi_-, \pi_+) \in \mathcal{P}_{E_1} \times \mathcal{P}_{E_2} =: \mathcal{P}_{\text{end}}^{(4)}.$$

The terms π_\pm denote the constant perturbations on the ends, and the π_i are perturbing terms on a compact piece of the end, as above. (While $\mathcal{P}_{\text{end}}^{(4)}$ depends on W and on the choice of isometry between the ends and $[0, \infty) \times Y$, this dependence is implicit and should be clear from context.)

These end perturbations are the same as the perturbations used in [KM11b]; they introduce further perturbations supported in $(1, 2) \times Y$, but as those are insufficient for later purposes, we use an altogether different source for our holonomy perturbations on the interior.

We will need one further type of perturbation, which is mainly useful to achieve transversality at flat connections and reducible connections. These were introduced in the setting of closed 4-manifolds in [Kro05], and are similar to the notion of holonomy perturbations along thickened loops in [Don87] and [Frø04].

Definition 4.2. Let W be a compact Riemannian manifold with boundary, equipped with an $SO(3)$ -bundle \mathbf{E} , and W° its interior. A collection of thickened loops in W consists of a choice of closed ball $B \subset W^\circ$, a finite collection q_1, \dots, q_n of smooth submersions $q_i : S^1 \times B \rightarrow W^\circ$ so that $q_i(1, b) = b$ and $q_i(-, b)$ is an immersion for all $1 \leq i \leq n$ and $b \in B$; we will write \vec{q} for the entire collection.

If $\gamma : S^1 \rightarrow X^\circ$ is a curve, we write $\text{Hol}_\gamma(\mathbf{A}) \in \widetilde{\text{Aut}}(\mathbf{E}_{\gamma(1)})$ for the parallel transport around the loop, thought of as based at the image of $1 \in S^1$; the lift to Aut is made canonical by the choice that $\text{Hol}_\gamma(\mathbf{A}) = 1$ if \mathbf{A} is fully reducible above γ . This is preserved by even based gauge transformations. Similarly we write $\text{Hol}_{\vec{q}}(\mathbf{A}) \in \widetilde{\text{Aut}}(\mathbf{E}_{\vec{q}(1)})^n$ for an n -tuple of curves with a common basepoint.

An interior holonomy perturbation is induced by the data of a collection of thickened loops q_1, \dots, q_n , as well as a choice of smooth, (even) gauge-equivariant bundle map

$$r : \widetilde{\text{Aut}}(\mathbf{E})^N|_B \rightarrow \mathfrak{g}_{\mathbf{E}}|_B$$

and a smooth self-dual 2-form ω on W with support in the interior of B .

The interior holonomy perturbation corresponding to (\vec{q}, r, ω) is the map

$$\begin{aligned} \widehat{\nabla}_\pi : \mathcal{A}_{\mathbf{E}} &\rightarrow \Omega^{2,+}(W; \mathfrak{g}_{\mathbf{E}}) \\ \widehat{\nabla}_\pi(\mathbf{A})(p) &= r(\text{Hol}_{\vec{q}_i(p,-)}(\mathbf{A})) \otimes \omega(p). \end{aligned}$$

Note that $\omega(p) = 0$ when $p \notin B$, making this formula well-defined.

We will apply these perturbations on the submanifold

$$W' = W \setminus ([1, \infty) \times Y).$$

Being gauge-equivariant, these perturbations descend to sections of the bundle $\mathcal{V}^+ \rightarrow \widehat{\mathcal{B}}$. We will shortly discuss extensions to the configuration spaces modeled on Sobolev spaces.

Soon, we will choose a countable dense set of the possible choices of data above.

These perturbations satisfy the following properties, analogous to those of [KM11b, Proposition 3.7] or [Lin18, Definition 3.2.7].

Proposition 4.4. Suppose π_\pm are 3-manifold perturbations on (Y_\pm) , with fixed critical orbits α_\pm . Suppose π is a 4-dimensional perturbation on the Riemannian manifold W with cylindrical ends defined above, including both interior and end perturbations. The map $\widehat{\nabla}_\pi$ extends to smooth maps

$$\widehat{\nabla}_\pi : \mathcal{A}_{\mathbf{E},k,\delta}(\alpha_-, \alpha_+) \rightarrow \Omega_{k,\delta}^{2,+}(W; \mathfrak{g}_{\mathbf{E}})$$

for any $k \geq 2$ and $\delta > 0$ sufficiently small. They satisfy the following properties.

- (1) Fixing a base connection \mathbf{A}_0 , for each $n \geq 0$ and $k \geq 2$, there is an increasing continuous function c_n so that

$$\|D^n \widehat{\nabla}_\pi(\mathbf{A})\|_{L_{k,\delta}^2} \leq c_n \left(\|\mathbf{A} - \mathbf{A}_0\|_{L_{k,\delta}^2} \right);$$

when $n > 0$ this should be interpreted as the operator norm.

- (2) $\|\widehat{\nabla}_\pi(\mathbf{A})\|_{L^\infty} \leq K$, independent of \mathbf{A} .
- (3) There is a constant C so that for all $1 \leq p \leq \infty$, and all L_k^2 connections \mathbf{A}, \mathbf{A}' , we have

$$\|\widehat{\nabla}_\pi(\mathbf{A}) - \widehat{\nabla}_\pi(\mathbf{A}')\|_{L^p} \leq C \|\mathbf{A} - \mathbf{A}'\|_{L^p}.$$

- (4) Suppose σ is a gauge transformation of \mathbf{E} , defined away from some finite set of points; and suppose \mathbf{A}_i are globally defined connections with $\sigma^*\mathbf{A}_1 = \mathbf{A}_2$ where defined. Then

$$\widehat{\nabla}_\pi((\mathbf{A}_1)) = \widehat{\nabla}_\pi(\mathbf{A}_2).$$

- (5) Suppose we have a finite set $\mathbf{x} \subset W$, a sequence of gauge transformations σ_n defined on $W \setminus \mathbf{x}$ which are $L^2_{k+1, \delta}$ on the ends and locally L^2_k on the complement of \mathbf{x} , and a sequence of connections \mathbf{A}_n on \mathbf{E} , so that over any compact set K for any $p \geq 2$, the connection $\sigma_n(\mathbf{A}_n)$ is bounded in L^p_1 over K . Then $(\sigma_n)_*\widehat{\nabla}_\pi(\mathbf{A}_n)$ has a subsequence which is Cauchy in $L^p(W)$. If there is an L^p connection \mathbf{A} on \mathbf{E} so that $\sigma_n(\mathbf{A}_n) \rightarrow \mathbf{A}$ in L^p on compact sets away from \mathbf{x} , then the limit of this Cauchy subsequence is $\widehat{\nabla}_\pi(\mathbf{A})$.

The proof must be divided into two parts, depending on the contribution from the end perturbations and the interior perturbations. In the former case, everything except the last point is given in [KM11b, Proposition 3.7]. These calculations for interior perturbations are given in [Kro05, Section 3]; the last point is argued in [Kro05, Lemma 3.4], and that argument works with only slight modifications for end perturbations as well.

As with the 3-dimensional perturbations and the end perturbations, then, we see that for any countable set of perturbations π_i there is an increasing sequence of constants $C_i > 0$ so that for any sequence $d_i \in \mathbb{R}$ with

$$\sum |d_i| C_i < \infty,$$

the limit $\sum d_i \widehat{\nabla}_{\pi_i}$ has the same properties as above; one defines the Banach space of perturbations to be the L^1 space weighted by C_i (that is, the space of such sequences d_i , with norm $\|(d_i)\| = \sum |d_i| C_i$). All inequalities in the above may then be made dependent on π , adding a factor of $\|\pi\|$ on the right-hand side.

Definition 4.3. *In this section, and in what follows, we used the following notation for Banach spaces of perturbations.*

- On a 3-manifold Y equipped with $SO(3)$ -bundle E , the space \mathcal{P}_E given by weighted L^1 sums of a countable dense set of data defining cylinder functions (with derivatives vanishing to all orders at the fully reducible connections),
- On a 4-manifold (W, \mathbf{E}) , with cylindrical ends (Y_1, E_1) and (Y_2, E_2) , the space

$$\mathcal{P}_{\text{end}}^{(4)} = \mathcal{P}_{E_1} \oplus \mathcal{P}_{E_2},$$

where the spaces \mathcal{P}_{E_i} contribute a perturbation of the form

$$\beta_0(t) (dt \wedge \nabla_\pi(\mathbf{A}(t)))^+$$

on the end, β_0 being a fixed cutoff function supported in $(1, \infty)$ (with W left implicit in the notation $\mathcal{P}^{(4)}$),

- On the same 4-manifold (W, \mathbf{E}) , the space $\mathcal{P}_{\text{int}}^{(4)}$ is given by L^1 sums of a countable dense set of interior holonomy perturbations, and we write

$$\mathcal{P}^{(4)} = \mathcal{P}_{\text{end}}^{(4)} \oplus \mathcal{P}_{\text{int}}^{(4)}.$$

- The affine subspace $\mathcal{P}_\epsilon^{(4)} \subset \mathcal{P}^{(4)}$ of perturbations with $(\pi_-, \pi_+) \in \mathcal{P}_{E_1} \oplus \mathcal{P}_{E_2}$ fixed (but left implicit in notation); this subspace is affine over \mathcal{P}_{int} .

Fix $\epsilon > 0$ so that for any reducible flat connection α , the only eigenvalues of $\widehat{\text{Hess}}_{\alpha,0}^\nu$ with absolute value at most ϵ are the elements of the kernel; one may fix a connected neighborhood $0 \in U_\epsilon \subset \mathcal{P}_E$ so that this remains true for all $\pi \in U_\epsilon$ and any reducible π -flat connection α .

This implies that $\dim \text{Eig}_{\pm\epsilon}(\widehat{\text{Hess}}_{\alpha,\pi}^\nu) = 0$ for all $\pi \in U_\epsilon$ and α a π -flat connection. In particular, this means that

$$\dim \text{Eig}_{-\epsilon \leq \lambda \leq \epsilon}(\widehat{\text{Hess}}_{\alpha,\pi}^\nu) = \dim \ker \widehat{\text{Hess}}_{\alpha,0}^\nu$$

for all such (π, α) .

If we let δ be a constant $0 < \delta \ll \epsilon$, then we denote by $\mathcal{P}_{E,\delta} \subset U_\epsilon$ the subset of those $\pi \in U_\epsilon$ so that $\widehat{\text{Hess}}_{\alpha,\pi}^\nu$ has no eigenvalues of absolute value at most δ , for any π -flat connection α (not necessarily reducible). This subset does not include 0, but by Theorem 3.6, $\bigcup_{\delta \rightarrow 0} \mathcal{P}_{E,\delta}$ is dense in U_ϵ . There are corresponding open subspaces $\mathcal{P}_\delta^{(4)}$ and $\mathcal{P}_{\text{end},\delta}^{(4)}$ for which the perturbations on the 3-manifold at ∞ lie in $\mathcal{P}_{E_i,\delta}$. These spaces will be important for the later weighted Sobolev theory.

So long as δ is sufficiently small, there is a union of connected components of $\mathcal{P}_{E,\delta}$ so that the enumeration of reducible flat connections of Proposition 3.4 holds. This follows because modifying the count of reducible flat connections would require one of them to become critical, and so the count is the same at every point in the component; further, there is some $\pi \in \mathcal{P}_{E,\delta}$ arbitrarily close to zero, so long as δ is chosen sufficiently small relative to the distance to zero.

Additionally, we may decompose $\mathcal{P}_{E,\delta}$ into connected open sets depending on the spectral flow of $\widehat{\text{Hess}}_{\alpha_t,\pi_t}^\nu$ for a generic path $0 \rightarrow \pi \in \mathcal{P}_{E,\delta}$ at the reducible π_t -flat connections α , precisely as in Proposition 3.9; this will be relevant later when calculating certain operator indices. We still have a notion of analytic energy of a configuration on W , but it is no longer a topological invariant, and will mostly be useful in keeping track of what happens during compactification.

Definition 4.4. Let W be a Riemannian manifold with cylindrical ends, equipped with an $SO(3)$ -bundle and a perturbation $\pi = (\pi_-, \pi_1, \pi_2, \pi_+)$ as above. Let α_\pm be critical orbits with respect to π_\pm . The analytic energy of a connection $\mathbf{A} \in \widetilde{\mathcal{B}}_{E,k,\delta}^e(\alpha_-, \alpha_+)$ is defined to be

$$\mathcal{E}_{\text{an}}^\pi(\mathbf{A}) = \int_W \|F_{\mathbf{A}}^+ + \widehat{\nabla}_\pi(\mathbf{A})\|^2.$$

This gives a real-valued lift of $\text{cs}_{\pi_+}(\alpha_+) - \text{cs}_{\pi_-}(\alpha_-)$ for π -ASD connections which only depends on their homotopy class z , and agrees with

$$\int \|F_{\mathbf{A}} + *_3 \nabla_\pi(\mathbf{A}(t))\|^2$$

for connections on the cylinder with constant perturbation.

4.2. Linear analysis and index theory. Throughout this section, a perturbation π on W is fixed so that the limiting perturbations π_i on (Y_i, E_i) have finitely many critical orbits in \mathfrak{C}_{π_i} , all nondegenerate and such that each $\widehat{\text{Hess}}_{\alpha,\pi}^\nu$ has all eigenvalues of absolute value larger than a fixed constant δ (independent of α).

We fix two such critical orbits, $\alpha_{\pm} = [A_{\pm}, E_b]$. Here recall that the elements of α_{\pm} vary over gauge equivalence classes of pairs (A_{\pm}, p) , where $p \in E_b$ is a point in the fiber of E above the basepoint b , thought of as a framing of E above that point.

To analyze the local structure of the moduli spaces around a framed instanton (\mathbf{A}, p) on $\mathbb{R} \times Y$ framed above $(0, b)$, we should restrict to the *framed Coulomb slice* $(\mathbf{A} + \ker(d_{\mathbf{A}}^*)) \times (\pi^*E)_b \subset \tilde{\mathcal{A}}_{k,\delta}(A_-, A_+)$: this is just the usual Coulomb slice with an additional $SO(3)$ coordinate for the framing. Every framed connection sufficiently close to (\mathbf{A}, p) is gauge equivalent to one in the framed Coulomb slice, and the representation is unique up to the action of $\Gamma_{\mathbf{A}} \hookrightarrow \text{Aut}(E_b) \cong SO(3)$ (the last isomorphism depending on the choice of framing at the basepoint). This acts trivially on the connection coordinate, and by translation in the framing coordinate.

We also find it convenient to pass to a smaller $\Gamma_{\mathbf{A}}$ -invariant open set in the framing coordinate, so that our subset is identified with

$$(3) \quad \ker(d_{\mathbf{A}}^*) \times N(\Gamma_{\mathbf{A}}) \cong \mathcal{K}_{\mathbf{A},k}^{(4)} \times \Gamma_{\mathbf{A}} \times \mathfrak{g}_{\mathbf{A}}^{\perp};$$

we embed the normal bundle in $SO(3)$ by exponentiating. We refer to this as the *extended Coulomb slice*. The action of $\Gamma_{\mathbf{A}}$ is again identified with left translation in the framing coordinate. Here $\mathfrak{g}_{\mathbf{A}}^{\perp} \subset \mathfrak{g}_b$ is the orthogonal complement to $\mathfrak{g}_{\mathbf{A}}$, the subset of \mathfrak{g}_b which extends to \mathbf{A} -parallel sections; equivalently, this is the tangent space to $\Gamma_{\mathbf{A}}$.

Inside the extended Coulomb slice around (\mathbf{A}, p) , (say, $\tilde{\mathcal{K}}_{k,\delta}^{(4)}$), there is a $\Gamma_{\mathbf{A}}$ -equivariant trivialization of the bundle $\Omega_{k-1,\delta}^{2,+}(\text{End}(E))$, descending to a local trivialization of $\mathcal{S}_{k-1,\delta}$ over $\tilde{\mathcal{K}}_{k,\delta}/\Gamma_{\mathbf{A}} \subset \tilde{\mathcal{B}}_{k,\delta}^e(\alpha_-, \alpha_+)$. Using the second expression in (3), an element $q \in N(\Gamma_{\mathbf{A}})$ can be expressed uniquely as $u \cdot (e^{\xi}p)$ for $\xi \in \mathfrak{g}_{\mathbf{A}}^{\perp}$ and $u \in \Gamma_{\mathbf{A}}$; the equivariant trivialization sends

$$(\mathbf{A} + a, ue^{\xi}p, \omega) \mapsto ((\mathbf{A} + a, ue^{\xi}p), u\omega).$$

In this equivariant trivialization, the section defined by (2) is given as

$$(\pi, \mathbf{A} + a, ue^{\xi}p) \mapsto u(d_{\mathbf{A}}a + a \wedge a)^+ + u(\widehat{\nabla}_{\pi}(\mathbf{A} + a) - \widehat{\nabla}_{\pi}(\mathbf{A})).$$

We can thus identify the derivative of the perturbed map (for fixed π and at (\mathbf{A}, p) a π -perturbed framed instanton) $\tilde{\mathcal{K}}_{k,\delta}^{(4)} \rightarrow \Omega_{k-1,\delta}^2$ as the *perturbed ASD operator*,

$$d_{\mathbf{A}}^+ a + (D_{\mathbf{A}} \nabla_{\pi})(a) \oplus 0 : \mathcal{K}_{k,\delta}^{(4)} \oplus \mathfrak{g}_{\mathbf{A}}^{\perp} \rightarrow \Omega_{k-1,\delta}^{2,+};$$

we write this as $D_{\mathbf{A},\pi}$, and the corresponding perturbed *normal ASD operator* as

$$D_{\mathbf{A},\pi}^{\nu} : \mathcal{K}_{k,\delta}^{(4)} \rightarrow \Omega_{k-1,\delta}^{2,+}.$$

Just as in Section 3.2, we gain control over the normal ASD operator with domain $\mathcal{K}_{k,\delta}^{(4)}$ by expressing it as a summand of a larger elliptic operator and including a gauge fixing condition. Define $Q_{\mathbf{A},\pi}$, the perturbed *extended ASD operator*, by

$$Q_{\mathbf{A},\pi} : \Omega_{k,\delta}^1(\pi^* \mathfrak{g}_E) \rightarrow \Omega_{k-1,\delta}^0(\pi^* \mathfrak{g}_E) \oplus \Omega_{k-1,\delta}^{2,+}(\pi^* \mathfrak{g}_E),$$

given by

$$Q_{\mathbf{A},\pi}(a, \xi) = (d_{\mathbf{A}}^* a, d_{\mathbf{A}}^+ a + D_{\mathbf{A}} \widehat{\nabla}_{\pi}(a)).$$

This is the operator usually used in the instanton theory for integer homology spheres, as applied to weighted Sobolev spaces.

Now recall that we have an L^2 -orthogonal splitting

$$\Omega_{k,\delta}^1(\mathfrak{g}_{\mathbf{E}}) = \text{Im}(d_{\mathbf{A}}) \oplus \ker(d_{\mathbf{A}}^*);$$

that every element may be written uniquely as a sum in this way is essentially the statement that the map

$$\Delta_{\mathbf{A}} : \Omega_{k+1,\delta}^0(\mathfrak{g}_{\mathbf{E}})/\mathfrak{g}_{\mathbf{A}} \rightarrow \Omega_{k-1,\delta}^0(\mathfrak{g}_{\mathbf{E}})/\mathfrak{g}_{\mathbf{A}}$$

is an isomorphism, assuming δ is not an eigenvalue of the Laplacian of the limiting connections, which is true as long as δ is chosen sufficiently small.

Now in this splitting $\Omega^1 = \text{Im}(d_{\mathbf{A}}) \oplus \ker(d_{\mathbf{A}}^*)$, and rewriting the first term as $\Omega_{k+1,\delta}^0/\mathfrak{g}_{\mathbf{A}}$ under the isomorphism $d_{\mathbf{A}}^{-1}$, we may write

$$Q_{\mathbf{A},\pi} = \begin{pmatrix} \Delta_{\mathbf{A}} & 0 \\ 0 & D_{\mathbf{A},\pi}^{\nu} \end{pmatrix}.$$

The fact that $D_{\mathbf{A},\pi}^{\nu}$ takes values in $\ker(d_{\mathbf{A}}^*)$ is the linearization of the statement that $F_{\mathbf{A}}^+ + \widehat{\nabla}_{\pi}(\mathbf{A})$ is a gauge invariant quantity. In particular, the normal ASD operator is Fredholm, and because the index of the top left operator is $-\dim \mathfrak{g}_{\mathbf{A}}$, the index of the normal ASD operator $D_{\mathbf{A},\pi}^{\nu}$ is the index of $Q_{\mathbf{A},\pi}$ plus the dimension of $\mathfrak{g}_{\mathbf{A}}$.

We recall the basic Fredholm property from [Don02, Sections 3.2 and 3.3.1].

Proposition 4.5. *Suppose W is a Riemannian 4-manifold with cylindrical ends, isometric to $(-\infty, 0] \times Y_1$ and $[0, \infty) \times \bar{Y}_2$, equipped with an $SO(3)$ -bundle \mathbf{E} restricting to the pullbacks of fixed $SO(3)$ -bundles E_i on the ends. In this situation, we say that (W, \mathbf{E}) is a cobordism from (Y_1, E_1) to (Y_2, E_2) . Suppose π is a fixed perturbation on \mathbf{E} , restricting to regular perturbations π_i on the ends, and A_i are fixed nondegenerate critical points of $\text{cs}_{Y_i} + f_{\pi_i}$.*

Let \mathbf{A} be a choice of connection in $\mathcal{A}_{\mathbf{E},k,\delta}(A_1, A_2)$.

If $\delta > 0$ is less than the absolute value of any eigenvalue of $\widehat{\text{Hess}}_{A_i,\pi}^{\nu}$, then $Q_{\mathbf{A},\pi}$ is Fredholm, and has index independent of such δ . If the A_i are irreducible, we may even take $\delta = 0$.

There are two essential points. The first is that when a connection $\mathbf{A} = \mathbf{A}(t)$ on the cylinder is in temporal gauge, the operator $Q_{\mathbf{A},\pi}$ can be written in the form

$$Q_{\mathbf{A},\pi} = \frac{d}{dt} + \widehat{\text{Hess}}_{\mathbf{A}(t),\pi}.$$

(See [Don02, Section 2.5].) Secondly, if $e^{\sigma(t)} = f_{\delta}(t)$ is the function used to define the weighted Sobolev spaces, multiplication by f_{δ} is an isometry $L_{k,\delta}^2 \rightarrow L_k^2$, and a first order linear differential operator D of the form $\frac{d}{dt} + L_t$ is taken under this isometry to $\frac{d}{dt} + L_t - \sigma'(t)$. So to study our operator $Q_{\mathbf{A},\pi}^{\nu}$ on weighted Sobolev spaces, we should equivalently study

$$\frac{d}{dt} + \widehat{\text{Hess}}_{\mathbf{A}(t),\pi} - \sigma'(t)$$

where $\sigma(t) = -\delta t$ for $t \ll 0$ and $\sigma(t) = \delta t$ for $t \gg 0$. (This is well-explained in [Don02, Section 3.3.1] and the beginning of [Lin18, Section 3.3].) Once we've done this, the above result for the ASD operator is a consequence of general theory; it is proven for L_t a family of almost self-adjoint first order differential operators as [KM07, Proposition 14.2.1].

An important consequence of this description is that for $\mathbf{A} = \mathbf{A}(t)$ in temporal gauge, we can describe the index of $Q_{\mathbf{A},\pi}$ as the *spectral flow* of the family of operators $\widehat{\text{Hess}}_{\mathbf{A}(t),\pi} - \sigma'(t)$ between $\widehat{\text{Hess}}_{A_1,\pi} + \delta I$ and $\widehat{\text{Hess}}_{A_2,\pi} - \delta I$, the (algebraic) intersection number of the paths the eigenvalues take with $0 \in \mathbb{R}$.

Because each $\mathcal{A}_{\mathbf{E},k,\delta,z}(A_1, A_2)$ is connected (even contractible) and index is a homotopy invariant of self-adjoint Fredholm operators, the index of $Q_{\mathbf{A},\pi}$ only depends on the homotopy class z , not the actual choice of connection \mathbf{A} . For the same reason, this also agrees with the index of $Q_{\mathbf{A},\pi}'$ for any other perturbation which restricts to the same perturbation sufficiently far on the ends.

It is perhaps worth observing that the index of $D_{\mathbf{A},\pi}'$, defined to be the derivative of the section operator $\widetilde{\mathcal{B}}_{\mathbf{E},k,\delta}^e \rightarrow \mathcal{S}_{k-1,\delta}$ normal to an orbit, *does* have its index jump as we pass from irreducibles to reducibles. This makes sense, thinking of D^ν as measuring the expected codimension of the orbit through \mathbf{A} in the entire moduli space: the codimension is larger at smaller-dimensional orbits.

Following the definition after [KM11b, Lemma 3.13], we use this to give the following definition.

Definition 4.5. *In the situation of Proposition 4.5, let z denote a connected component of $\widetilde{\mathcal{A}}_{\mathbf{E},k,\delta}(\alpha, \beta)$. We write the unframed grading $\overline{\text{gr}}_z^W(\alpha, \beta) = \text{ind}(Q_{\mathbf{A},\pi}'$ for any choice of \mathbf{A} in the component z , and the relative grading between the orbits α , β , with respect to the path z , is $\text{gr}_z(\alpha, \beta) = \overline{\text{gr}}_z(\alpha, \beta) + 3 - \dim \alpha$.*

When W is the cylinder with the constant perturbation, we drop the superscript W .

The relative grading here will be quite natural in the definition of the framed instanton Floer complex, whose differential is defined in terms of fiber products with moduli spaces; $\text{gr}_z^W(\alpha, \beta)$ is the expected dimension of a fiber of the map

$$\text{ev}_- : \widetilde{\mathcal{M}}_{\mathbf{E},\pi,z}(\alpha, \beta) \rightarrow \alpha.$$

It is worth noting here that $3 - \dim \alpha = \dim \mathfrak{g}_\alpha = \dim \Gamma_\alpha$.

The more immediate point of this definition is that the relative grading is *additive*. In the particular case that W is isometric to the cylinder $\mathbb{R} \times Y$, z corresponds to a relative homotopy class between $[\alpha]$ and $[\beta]$ in $\mathcal{B}_{E,k}^e$. If w is a path from $[\beta]$ to $[\gamma]$, then

$$\text{gr}_z(\alpha, \beta) + \text{gr}_w(\beta, \gamma) = \text{gr}_{z \cdot w}(\alpha, \gamma)$$

($z \cdot w$ the concatenated path). This follows by computing these as a spectral flow: if $\mathbf{A}_z(t) = \mathbf{A}_z \in \widetilde{\mathcal{A}}_{k,\delta,z}(\alpha, \beta)$ is in temporal gauge and $\mathbf{A}_z(t) = \alpha$ (resp. β) for $t \ll 0$ (resp $t \gg 0$), the index of $Q_{\mathbf{A}_z,\pi}$ is the spectral flow of the path

$$\widehat{\text{Hess}}_{A_z(t),\pi} - \sigma'(t),$$

where $\sigma'(t) = -\delta$ for $t \ll 0$ and $\sigma'(t) = \delta$ for $t \gg 0$. Making a similar choice of \mathbf{A}_w , if we glue together \mathbf{A}_z and \mathbf{A}_w sufficiently far out on the ends, we can find a connection $\mathbf{A}_{z \cdot w}$ in the component corresponding to $z \cdot w$ as the concatenation of the paths $\mathbf{A}_z(t)$ and $\mathbf{A}_w(t)$. However, we cannot concatenate the corresponding paths of self-adjoint operators yet; for large t , the first ends at $\widehat{\text{Hess}}_{\beta,\pi} - \delta$ and the second begins at $\widehat{\text{Hess}}_{\beta,\pi} + \delta$. To actually concatenate them, we must traverse the path

$$\widehat{\text{Hess}}_{\beta,\pi} - (1 - 2t)\delta;$$

doing so changes $(3 - \dim \beta) = \dim \ker \left(\widehat{\text{Hess}}_{\beta, \pi} \right)$ negative eigenvalues to positive, and so

$$\overline{\text{gr}}_z(\alpha, \beta) + (3 - \dim \beta) + \overline{\text{gr}}_w(\beta, \gamma) = \overline{\text{gr}}_{z \cdot w}(\alpha, \gamma).$$

Additivity of this grading for general cobordisms is also true. We record this as a proposition.

Proposition 4.6. *Suppose we have cobordisms (W_1, \mathbf{E}_1) from (Y_1, E_1) to (Y_2, E_2) and (W_2, \mathbf{E}_2) from (Y_2, E_2) to (Y_3, E_3) , equipped with paths $\gamma_1 : \mathbb{R} \rightarrow W_1$ and $\gamma_2 : \mathbb{R} \rightarrow W_2$ between the basepoints of the Y_i . Suppose the W_i are equipped with perturbations $\hat{\pi}_i$, which restrict to regular perturbations π_j on each Y_j , and furthermore suppose each (Y_j, E_j, π_j) is equipped with some π_j -flat connection A_j .*

*We can define the composed cobordism $(W_{12}^T, \mathbf{E}_{12})$ from (Y_1, E_1) to (Y_3, E_3) , identifying $(T, \infty) \times \overline{Y}_2$ with $(-\infty, -T) \times Y_2$ on the cylindrical ends; there is a corresponding perturbation $\hat{\pi}$ interpolating between π_1 and π_3 . We denote by z_i a component of $\tilde{\mathcal{A}}_{\mathbf{E}_i, k, \delta}(A_i, A_{i+1})$; there is a component $z_1 * z_2$ corresponding to gluing representative connections of these components along the ends. In this situation, we have*

$$\text{gr}_{z_1}^{W_1}(\alpha_1, \alpha_2) + \text{gr}_{z_2}^{W_2}(\alpha_2, \alpha_3) = \text{gr}_{z_1 * z_2}^{W_{12}^T}(\alpha_1, \alpha_3).$$

This follows from the additivity theorem of the index when the limiting operators over the ends have no kernel ([Don02, Proposition 3.9]) and the relation of operators on weighted Sobolev spaces to unweighted spaces, given by conjugating by the weighting function: the operator $Q_{\mathbf{A}_1, \pi_1}$ on its positive end, after conjugating by the weighting function e^σ , takes the form $\frac{d}{dt} + \widehat{\text{Hess}}_{A_2, \pi_2} - \delta$ and on the negative end of W_2 takes the form $\frac{d}{dt} + \widehat{\text{Hess}}_{A_2, \pi_2} + \delta$. To glue these we first need to interpolate between $-\delta$ and $+\delta$, moving the

$$(3 - \dim \alpha_2) = \dim \ker \left(\widehat{\text{Hess}}_{A_2, \pi_2} \right)$$

negative eigenvalues across 0; this is observed as [Don02, Proposition 3.10], identifying the index on weighted spaces with what he denotes $\text{ind}^+(P)$.

4.3. Uhlenbeck compactness for framed instantons. The following definition is precisely [KM07, Definition 16.1.1]. Note that there are no framings involved yet.

Definition 4.6. *Let Y be a Riemannian 3-manifold, equipped with $SO(3)$ -bundle E and regular perturbation π ; then π has finitely many critical points in $\mathcal{B}_{E, k}^e$, which we write a generic point of as α . We say that a trajectory from α_- to α_+ is an equivalence class of nonconstant π -perturbed instanton $\mathbf{A} \in \mathcal{M}_E(\alpha_-, \alpha_+) \subset \mathcal{B}_{\pi * E, k, \delta}^e(\alpha_-, \alpha_+)$ on $\mathbb{R} \times Y$ under the translation action.*

The homotopy class of a trajectory is the element of $\pi_1(\mathcal{B}_{E, k}^e, \alpha_-, \alpha_+)$ it traces out; these are in noncanonical bijection with \mathbb{Z} . The topological energy of a trajectory was defined above as

$$2((\text{cs} + f_\pi)(\text{ev}_+ \mathbf{A}) - (\text{cs} + f_\pi)(\text{ev}_- \mathbf{A}));$$

even though the Chern-Simons functional of an individual connection is only defined in $\mathbb{R}/8\pi^2\mathbb{Z}$, this difference of boundary components of a connection on a cylinder is defined in \mathbb{R} .

A broken trajectory from α_- to α_+ consists of a finite sequence of π -perturbed instantons \mathbf{A}_i on π^*E over $\mathbb{R} \times Y$ (say $1 \leq i \leq n$), with

$$\text{ev}_- \mathbf{A}_1 = \alpha_-, \quad \text{ev}_+ \mathbf{A}_n = \alpha_+, \quad \text{and } \text{ev}_+ \mathbf{A}_i = \text{ev}_- \mathbf{A}_{i+1} \text{ for } 0 < i < n.$$

The homotopy class of a broken trajectory is the composite of the homotopy classes given by the individual trajectories; the energy of a broken trajectory is defined as

$$\mathcal{E}^\pi(\mathbf{A}_1, \dots, \mathbf{A}_n) = \sum_{i=1}^n \mathcal{E}^\pi(\mathbf{A}_i).$$

We topologize this exactly as in [KM07, Page 276]. As per the author's taste, we present this in terms of sequences, following [Don02, Page 116]: give the space of broken trajectories the final topology so that the following sequences of unbroken instantons, and their natural generalizations to sequences of broken instantons, converge to their stated limits: \mathbf{A}_i , where $i \in \mathbb{N}$, converges to the broken trajectory $(\mathbf{B}_1, \dots, \mathbf{B}_n)$, if there is a sequence (T_i^1, \dots, T_i^n) of real numbers with $T_i^j \leq T_i^{j+1}$ for $0 < j < n$, and the successive differences $T_i^{j+1} - T_i^j \rightarrow \infty$ as $i \rightarrow \infty$, so that the pullbacks $\tau_{T_i^j}^* \mathbf{A}$ converge as $i \rightarrow \infty$ to \mathbf{B}_j in the $L_{k,\delta}^2$ topology. Note that the energy of a trajectory is continuous with respect to this 'chain convergence'.

This form of noncompactness, trajectories breaking into a composite of lower-index trajectories, is familiar in Morse theory. There is another kind of noncompactness familiar in the instanton theory: Uhlenbeck bubbling.

Definition 4.7. *Let (Y, E, π) be as above. An ideal instanton is a solution \mathbf{A} to the π -ASD equations on $(\mathbb{R} \times Y, \pi^*E)$, along with a finite (possibly empty) collection of points $x_i \in \mathbb{R} \times Y$ and integer weights $k_i \geq 1$ at each point. An ideal trajectory is an \mathbb{R} -equivalence class of nontrivial ideal instantons, where nontrivial means either that \mathbf{A} is nonconstant or that $\{x_i\}$ is nonempty; equivalently, a nontrivial ideal instanton is one so that $\tau_t^*(\mathbf{A}, x, k)$ is not gauge equivalent to (\mathbf{A}, x, k) for any t . Similarly, a broken ideal trajectory is a sequence of ideal trajectories $(\mathbf{A}_i)_{i=1}^n$ with $\text{ev}_+ \mathbf{A}_i = \text{ev}_- \mathbf{A}_{i+1}$.*

The energy of an ideal instanton (\mathbf{A}, x_i, k_i) , where $1 \leq i \leq n$, is

$$\mathcal{E}^\pi(\mathbf{A}) + 16\pi^2 \sum_i k_i.$$

If the homotopy class of the instanton \mathbf{A} is z , the homotopy class of the ideal instanton (\mathbf{A}, x_i, k_i) is $z + \sum k_i$. The energy and homotopy class of a broken ideal instanton are defined to be additive under concatenation.

We say that a sequence of π -perturbed instantons \mathbf{A}_n converges to an ideal instanton (\mathbf{A}, x_i, k_i) if there is a sequence of gauge transformations σ_n defined on $(\mathbb{R} \times Y) \setminus \{x_i\}$ such that $\sigma_n^* \mathbf{A}_n \rightarrow \mathbf{A}$ in the L_1^p topology on compact subsets of $(\mathbb{R} \times Y) \setminus \{x_i\}$ for all $\infty > p \geq 2$, and such that the density measures converge:

$$2|F_{\mathbf{A}_n}|^2 \rightarrow 2|F_{\mathbf{A}}|^2 + \sum_{i=1}^n 64\pi^2 k_i \delta_{x_i}.$$

There is then a natural extension of this to a definition of convergence to (and of) broken ideal trajectories. The space of broken ideal trajectories from α to β in the homotopy class z is written $\mathcal{M}_{E,z,\pi}(\alpha, \beta)$. Note that this does *not* include constant trajectories, or broken 'trajectories' for which one of the components is constant.

As observed in [Kro05], the non-local nature of holonomy perturbations means that we cannot expect better convergence than L_1^p under bubble-limits.

The Uhlenbeck compactness theorem for the ASD equations on the cylinder is the following.

Proposition 4.7. *The subspace of broken ideal trajectories in $\check{\mathcal{M}}_{E,z,\pi}(\alpha,\beta)$ with a fixed energy bound $\mathcal{E}^\pi(\mathbf{A}) \leq C$ is compact.*

We do not repeat the proof, which can be seen in [Don02, Section 5.1]; the corresponding fact for compact cylinders is [KM11b, Proposition 3.20].

A somewhat stronger statement, Proposition 4.8 below, is true; we will only use it briefly, but find it to be somewhat interesting.

As in Definition 4.3, for any real numbers $\epsilon \gg \delta > 0$, we let $\mathcal{P}_{E,\delta} \subset \mathcal{P}_E$ be the open subset of perturbations on (Y, E) so that, for each π -critical point α , all eigenvalues of $\widehat{\text{Hess}}_{\alpha,\pi}^\nu$ have absolute value larger than δ , and so that for every pair (π, α) of perturbation and π -flat *reducible* connection, the operator $\widehat{\text{Hess}}_{\alpha,\pi}^\nu$ has no eigenvalues of absolute value ϵ . (In particular, we assume π is a regular perturbation.)

Because the projection $\mathfrak{C}_\pi \rightarrow \mathcal{P}_{E,\delta}$ of the parameterized critical set (in $\check{\mathcal{B}}_E^\epsilon$) to the space of perturbations is a proper submersion, it is in particular a locally trivial fiber bundle. Therefore, for some small open set U around any $\pi_0 \in \mathcal{P}_{E,\delta}$, we have a canonical bijection $\mathfrak{C}_{\pi_0} \cong \mathfrak{C}_\pi$ for any $\pi \in U$; in fact, if we fix $\alpha \in \mathfrak{C}_{\pi_0}$, we may choose a smooth map $s_\alpha : U \rightarrow \mathcal{A}_{E,k}$, so that $s(\pi)$ is a π -flat connection which is identified under the above bijection with α . Choose once and for all, for each homotopy class z , a smooth map

$$r_z : \mathcal{A}_{E,k} \times \mathcal{A}_{E,k} \rightarrow \mathcal{A}_{\pi^*E,z,k,\delta}^{(4)},$$

sending (A_-, A_+) to a connection which is constant at A_- for $t \leq -1$ and constant at A_+ for $t \geq 1$ and in the homotopy class z .

We may use the s_α to define the parameterized configuration space $\mathcal{P}_\delta \mathcal{A}_{\pi^*E,z,k,\delta}^{(4)}$, whose elements $(\pi, \alpha_\pm, \mathbf{A})$ consist of a perturbation $\pi \in \mathcal{P}_{E,\delta}$, a choice of two π_\pm -flat connections α_\pm , and a connection $\mathbf{A} \in \mathcal{A}_{\pi^*E,z,k,\delta}^{(4)}(\alpha_-, \alpha_+)$.

This set is given the structure of a smooth Banach manifold by patching together charts of the form

$$U \times \mathcal{A}_{\pi^*E,z,k,\delta}(\alpha_-, \alpha_+) \cong r_z(s_{\alpha_-} \pi, s_{\alpha_+} \pi) + \Omega_{k,\delta}^1(\mathfrak{g}_E).$$

Write

$$\mathcal{P}_\delta \tilde{\mathcal{A}}_{\pi^*E,z,k,\delta}^{(4)} = \mathcal{P}_\delta \mathcal{A}_{\pi^*E,z,k,\delta}^{(4)} \times \pi^*E_{(0,b)}$$

for the parameterized space of framed connections, which inherits a smooth structure and a smooth right action of $SO(3)$. It carries a smooth projection map

$$\mathcal{P}_\delta \tilde{\mathcal{A}}_{\pi^*E,z,k,\delta}^{(4)} \rightarrow \mathcal{P}_{E,\delta}.$$

It carries the action of a bundle of Banach Lie groups over $\mathcal{P}_{E,\delta}$; the fiberwise quotient gives a topological space $\mathcal{P}_\delta \tilde{\mathcal{B}}_{\pi^*E,z,k,\delta}^e$. We may define the parameterized moduli space

$$\mathcal{P}_\delta \tilde{\mathcal{M}} \subset \mathcal{P}_\delta \mathcal{B}_{\pi^*E,z,k,\delta}^e$$

as the equivalence classes of triples (π, \mathbf{A}, p) , where $\pi \in \mathcal{P}_{E,\delta}$ is a perturbation, \mathbf{A} is a π -perturbed instanton, and p is a framing.

Because the perturbed ASD equations do not depend on the framing, this set inherits the right $SO(3)$ -action. The quotient of $\mathcal{P}_\delta \widetilde{\mathcal{M}}$ by this $SO(3)$ action is the parameterized moduli space of instantons, which we denote $\mathcal{P}_\delta \mathcal{M}$. If we take the quotient by the \mathbb{R} action, throw out the constant trajectories, and incorporate bubble-limits into the topology, we may extend this to a space $\mathcal{P}_\delta \check{\mathcal{M}}$, the parameterized moduli space of ideal broken trajectories. (We will soon discuss the version of this appropriate to the framed setting, which is slightly more subtle.)

Proposition 4.8. *Let $\mathcal{P}_\delta \check{\mathcal{M}}^{\leq C}$ be the subspace of $\mathcal{P}_\delta \check{\mathcal{M}}$ consisting of those pairs (π, \mathbf{A}) so that $\mathcal{E}^\pi(\mathbf{A}) \leq C$. Then the projection map $\mathcal{P}_\delta \check{\mathcal{M}}^{\leq C} \rightarrow \mathcal{P}_\delta$ is proper.*

Proof. Suppose we have a sequence (π_n, \mathbf{A}_n) of perturbations and unbroken instantons so that $\pi_n \rightarrow \pi$. We want to show that there is a subsequence of \mathbf{A}_n which converges to a broken ideal π -trajectory. (The general case where \mathbf{A}_n is itself a broken ideal trajectory provides no further difficulty.) First, because the possible ev_\pm take values in a finite set, choose a subset of \mathbf{A}_n so that, for n large, $\text{ev}_\pm \mathbf{A}_n$ correspond to $\alpha_\pm \in \mathfrak{C}_\pi \cong \mathfrak{C}_{\pi_n}$. We may now apply [Don02, Lemma 4.3], which establishes that any π -instanton with sufficiently small energy and $\text{ev}_- \mathbf{A} = \alpha$ is gauge equivalent on $(-\infty, 0) \times Y$ to $\alpha + a$ for some a with a uniform bound on $|a(t)|e^{-\delta t}$, as well as the derivatives $|\nabla^{(\ell)}(a)(t)|e^{\delta t}$ for $\ell \leq k$. The constants in these uniform bounds are bounded for the convergent sequence $\pi_n \rightarrow \pi$, and δ is fixed. In particular, for our solutions \mathbf{A}_n , gauge equivalent to $\alpha_n + a_n$, this is enough for the Arzela-Ascoli and dominated convergence argument in [Don02, Lemma 5.1] to imply that $a_n \rightarrow a$ for some function a satisfying the same bounds; in fact we must have

$$F_{\alpha+a}^+ + \widehat{\nabla}_\pi(\alpha + a) = 0$$

as this is the pointwise limit of the corresponding equations

$$F_{\mathbf{A}_n}^+ + \widehat{\nabla}_\pi(\mathbf{A}_n) = 0.$$

Now that we have control over the ends, everything else is standard: a uniform bound on $\mathcal{E}^{\pi_n}(\mathbf{A}_n)$ implies a uniform bound on the L^2 norm of $F_{\mathbf{A}}$ on compact sets, so one has a limit on compact sets after accounting for bubbling, and then a limit on the whole line to a broken ideal trajectory; this uses that there is a uniform positive lower bound on the minimal energy of a nontrivial π_n -instanton. To see this, recall that the energy of an instanton may be written as $2 \int |\mathbf{A}'(t)|^2$ as in the proof of Proposition 4.1. If \mathbf{A}_n is a sequence of nontrivial π_n -instantons with $\mathcal{E}^{\pi_n}(\mathbf{A}_n) \rightarrow 0$, this implies that the distance between the endpoints α_\pm^n goes to zero. But $\alpha_\pm^n \rightarrow \alpha_\pm$, so $\alpha_- = \alpha_+$; but then because α_\pm^n is sent to α_\pm under the bijection $\mathfrak{C}_{\pi_n} \cong \mathfrak{C}_\pi$, we see that $\alpha_-^n = \alpha_+^n$ for large n . However, this would imply that $\mathcal{E}^{\pi_n}(\mathbf{A}_n)$ is a multiple of $64\pi^2$ for large n , and also may be made arbitrarily close to zero by taking n large; therefore $\mathcal{E}^{\pi_n}(\mathbf{A}_n) = 0$ for large n . However, this implies $\mathbf{A}'_n(t) = 0$, but we assumed \mathbf{A}_n was nontrivial.

The reason we restrict to perturbations in $\mathcal{P}_{E,\delta}$ is for the definition of the parameterized moduli space as $L_{k,\delta}^2$ connections; we always want to take δ less than the eigenvalues of the extended Hessian, and otherwise would need to choose δ depending on the perturbation. \blacksquare

We conclude the discussion of compactness for cylinders by defining the object of interest to us: the compactification of the *framed* moduli space of trajectories.

Definition 4.8. *Let Y be a Riemannian 3-manifold, equipped with $SO(3)$ -bundle E and regular perturbation π . We say that a framed ideal trajectory from α_- to α_+ is an equivalence class of π -perturbed framed instanton $(\mathbf{A}, p) \in \tilde{\mathcal{B}}_{\pi^*E, k, \delta}^e(\alpha_-, \alpha_+)$ equipped with a (possibly empty) points of points $x \in \mathbb{R} \times Y$ and positive integer weights k_x , where none of the x are the basepoint $(0, b)$. We demand the ideal instanton is nontrivial, in the sense that either the set of points x_i is nonempty or the trajectory \mathbf{A} is nonconstant.*

A deframed ideal trajectory from α_- to α_+ is an ideal trajectory from α_- to α_+ so that $(0, b)$ is a weight-point x with $k_{(0, b)} > 0$.⁷

A framed broken trajectory from α_- to α_+ is a finite sequence whose elements are either framed ideal trajectories $(\mathbf{A}_i, p_i, x_i, k_i)$ or deframed ideal trajectories (\mathbf{A}, x_i, k_i) with

$$\text{ev}_-(\mathbf{A}_1, p_1) \in \alpha_-, \quad \text{ev}_+(\mathbf{A}_n, p_n) \in \alpha_+, \quad \text{and}$$

$$\text{ev}_+(\mathbf{A}_i, p_i) = \text{ev}_-(\mathbf{A}_{i+1}, p_{i+1}) \text{ for } 0 < i < n, \text{ when } \mathbf{A}_i, \mathbf{A}_{i+1} \text{ are both framed.}$$

The set of framed broken trajectories in the homotopy class z from α_- to α_+ is written $\overline{\mathcal{M}}_{E, z, \pi}(\alpha_-, \alpha_+)$. We write a generic element as (\mathbf{A}, p) , even though \mathbf{A} may be broken and not every piece must be framed.

This set is topologized as follows. A sequence (\mathbf{A}_n, p_n) of framed instantons converges to a framed ideal instanton (\mathbf{A}, p, x_i, k) if there is a sequence of gauge transformations σ_n , defined on $(\mathbb{R} \times Y) \setminus \{x_i\}$, and in particular defined on the basepoint $(0, b)$, so that $\sigma_n^* \mathbf{A}_n \rightarrow \mathbf{A}$ converge in the sense of ideal instantons above, and $\sigma_n^* p_n \rightarrow p$.

However, if the underlying trajectories \mathbf{A}_n converge to an ideal trajectory with nontrivial weight at $(0, b)$, so that the σ_n are undefined at $(0, b)$, then it doesn't make sense to compare $\sigma_n^* p_n$ and p . In this situation we lose the framing via bubbling at the basepoint. In this case, the sequence (\mathbf{A}_n, p_n) of framed instantons converges to a deframed ideal trajectory. Incorporating this into the topology on broken trajectories is straightforward.

Proposition 4.9. *The space $\overline{\mathcal{M}}_{E, z, \pi}(\alpha_-, \alpha_+)$ is compact.*

Proof. We begin by recalling that there is a surjective map

$$\overline{\mathcal{M}}_{E, z, \pi}(\alpha_-, \alpha_+) \rightarrow \check{\mathcal{M}}_{E, z, \pi}(\alpha_-, \alpha_+).$$

This map is proper: if a sequence (\mathbf{A}_n, p_n) of framed broken trajectories has underlying sequence of broken ideal trajectories converge to \mathbf{A} , the framings are either incomparable (and that limit component of \mathbf{A}_n is a deframed ideal trajectory), or (an appropriate sequence of translations of) $\sigma_n^* p_n$ are all defined, and live in the compact space $E_b \cong SO(3)$; so some subsequence converges, as desired. Doing this for the finitely many components of the limit \mathbf{A} constructs an element of $\overline{\mathcal{M}}$ that a subsequence of (\mathbf{A}_n, p_n) converges to. Because we have constructed a proper map to a compact space, the total space $\overline{\mathcal{M}}$ is compact. \blacksquare

Similarly, the projection from the parameterized moduli space $\mathcal{P}_\delta \overline{\mathcal{M}}^{\leq C}$ to $\mathcal{P}_{E, \delta}$ is proper.

⁷Here deframed is meant to indicate that the framing has been removed via the placement of a δ -mass at the basepoint; not all the weight-sets are allowed, and in particular a non-ideal instanton without a framing is not a 'deframed ideal trajectory'.

There are straightforward extensions of these to moduli spaces on cobordisms, which we now state. In what follows, we write $\mathcal{P}_\delta^{(4)}$ for the subspace

$$\mathcal{P}_{E_1, \delta} \oplus \mathcal{P}_{\text{int}}^{(4)} \oplus \mathcal{P}_{E_2, \delta}.$$

(As usual, the manifold W is left implicit in the notation.)

Definition 4.9. *Let W be an oriented Riemannian 4-manifold, with cylindrical ends oriented isometric to $(-\infty, 0] \times Y_1$ and $[0, \infty) \times Y_2$, equipped with an $SO(3)$ -bundle \mathbf{E} and specified isomorphisms to the pullback of bundles on Y_i over the ends. Suppose W is equipped with an embedding $\gamma : \mathbb{R} \hookrightarrow W$ which agrees for $|t|$ sufficiently large with (t, b_1) or (t, b_2) , depending on the sign of t , and write $b = \gamma(0)$ as the basepoint of W . Suppose W is equipped with a perturbation π , restricting to fixed regular perturbations π_i on the ends. Connected components of the space of connections on W limiting to α_\pm , critical points of π_\pm , are in bijection with \mathbb{Z} , and labeled by z .*

Then a broken ideal W -trajectory from α_- to α_+ in the homotopy class z is a triple of a broken ideal π_- -trajectory on Y_1 from α_- to some β (possibly constant, if $\alpha_- = \beta$), an ideal π -instanton on W from β to some γ , and a broken ideal π_+ -trajectory on Y_2 from γ to α_+ (possibly constant, if $\gamma = \alpha_+$). The homotopy class is the composite of the corresponding homotopy classes of non-ideal instantons, then summing the weights. We denote the set of broken ideal W -trajectories from α_- to α_+ in the homotopy class z as $\mathcal{M}_{\mathbf{E}, z}^{W, \pi}(\alpha_-, \alpha_+)$.

A framed broken ideal W -trajectory is the same, but with in addition a framing on each component where there is no weight at the basepoint of Y_i or W , and so that the neighboring evaluations agree:

$$\text{Hol}_{\mathbf{A}_i, \gamma_i}^{0 \rightarrow \infty}(p_i) = \text{Hol}_{\mathbf{A}_{i+1}, \gamma_{i+1}}^{0 \rightarrow -\infty}(p_{i+1}).$$

We denote the set of framed broken ideal W -trajectories from α_- to α_+ in the homotopy class z as $\mathcal{M}_{\mathbf{E}, z, \pi}(\alpha_-, \alpha_+)$. We similarly denote the parameterized space by

$$\mathcal{P}_{L, \delta} \overline{\mathcal{M}}_{\mathbf{E}, z}.$$

Note in particular that we include the ‘deframed trajectories’ here as in the cylindrical case, and at these, we have lost the framing at the basepoint. Here (W, \mathbf{E}) is collapsed in the notation to simply the bundle \mathbf{E} .

A sequence of instantons \mathbf{A}_n on W from α_- to α_+ converges to a broken trajectory $(\mathbf{A}_-, \mathbf{A}, \mathbf{A}_+)$ if the following hold. First, for each end, we demand that there is a sequence of real numbers T_n^j with

$$\begin{aligned} \lim_{j \rightarrow \infty} T_n^{j+1} - T_n^j &= \pm \infty \\ \lim_{n \rightarrow \infty} T_n^0 &= \pm \infty; \end{aligned}$$

the limit should be negative on the negative end and positive on the positive end. These real numbers should be such that the translates $\tau_{T_n^j}^* \mathbf{A}_n$ converge to $(\mathbf{A}_-)_j$ in

$$\mathcal{B}_{\pi^* E, k, \delta}((-\infty, N] \times Y_1)$$

for all N , and similarly for the positive end. Note that the assumption that $T_n^0 \rightarrow \pm \infty$ implies that these translates are eventually defined on all of $(-\infty, N] \times Y_1$ for any fixed N , and similarly $[-N, \infty) \times Y_2$.

Second, we demand that $\mathbf{A}_n \rightarrow \mathbf{A}$ in $\mathcal{B}_{\mathbf{E},k,\delta}$.

One similarly accounts for the bubbling phenomenon to define limits to broken *ideal* trajectories: there should be a sequence of gauge transformations defined away from the bubble points so that convergence of $\sigma_n(\mathbf{A}_n)$ is L^p_1 on compact sets away from the bubble-points for all $\infty > p \geq 2$.

We have the following analogue of Proposition 4.8 for W -trajectories, which will be useful to us. The proof is a verbatim combination of what was already said in the case of the ends (that is, in the cylindrical case), as well as a discussion of what happens on compact subsets of W .

Proposition 4.10. *Let $\mathcal{P}_\delta \check{\mathcal{M}}_{\mathbf{E}}^{\leq C}$ be the subspace of $\mathcal{P}_\delta \check{\mathcal{M}}_{\mathbf{E}}$ consisting of those pairs (π, \mathbf{A}) so that $\mathcal{E}^\pi(\mathbf{A}) \leq C$. Then the projection map $\mathcal{P}_\delta \check{\mathcal{M}}_{\mathbf{E}}^{\leq C} \rightarrow \mathcal{P}_\delta^{(4)}$ is proper.*

Proof. Fix (\mathbf{A}_n, π_n) with $\pi_n \rightarrow \pi$. We have already dealt with the ends in the discussion of the cylinder, and so we will focus on W itself. We need to show that there is some sequence of gauge transformations, defined on the complement of some finite set, so that there is a connection \mathbf{A} and a subsequence (still written \mathbf{A}_n) so that $\sigma_n^* \mathbf{A}_n \rightarrow \mathbf{A}$. What we need is to show that a uniform energy bound gives, over any given compact set, a uniform bound on the L^2 norm of $F_{\mathbf{A}_n}$. Because

$$F_{\mathbf{A}_n}^+ = \widehat{\nabla}_{\pi_n}(\mathbf{A}_n)$$

is bounded in L^∞ independent of n , Uhlenbeck's compactness theorem will then guarantee that there is a sequence of gauge transformations σ_n defined on the complement of a finite set \mathbf{x} , and an L^p_1 connection \mathbf{A} in Coulomb gauge with respect to some reference smooth connection, so that $\sigma_n(\mathbf{A}_n) \rightarrow \mathbf{A}$ in L^p_1 on compact sets of $W \setminus \mathbf{x}$. Proposition 4.4 (5) guarantees that

$$\widehat{\nabla}_{\pi_n}(\sigma_n)_*(\mathbf{A}_n) \rightarrow_{L^p} \widehat{\nabla}_\pi(\mathbf{A}),$$

and so

$$0 = (\sigma_n)_* \left(F_{\mathbf{A}_n}^+ + \widehat{\nabla}_{\pi_n}(\mathbf{A}_n) \right) \rightarrow_{L^p} F_{\mathbf{A}}^+ + \widehat{\nabla}_\pi(\mathbf{A});$$

therefore, \mathbf{A} is indeed a π -instanton. Because \mathbf{A} is in Coulomb gauge with respect to a reference smooth connection, the usual bootstrapping arguments imply \mathbf{A} is smooth.

We will further need to ensure that if \mathbf{A}_n converges to an ideal perturbed instanton, the weight of the δ -mass at the ideal point x is a multiple of $64\pi^2$.

First, we will resolve the question of curvature. Consider the restriction of \mathbf{A} to the complement of the ends $[1, \infty) \times Y$, a compact manifold W' with boundary Y . The perturbations vanish on the boundary of W' (and on the interior are the previously discussed interior holonomy perturbations). On the ends the perturbations π_n are of the form (eg, for $[0, \infty) \times Y_1$)

$$\beta_0(t)(\pi_-)_n,$$

where $\beta_0 : [0, \infty) \rightarrow [0, 1]$ is zero for $t \leq 1$ and 1 for $t \geq 2$.

Remember that for a π -instanton \mathbf{A} (these chosen arbitrarily), its analytic energy $\mathcal{E}_\pi^{\text{an}}(\mathbf{A})$ was defined to be equal to

$$\mathcal{E}_\pi^{\text{top}}(\mathbf{A}) = 2 \left((cs + f_{\pi_-})(\alpha_-) - (cs + f_{\pi_+})(\alpha_+) \right).$$

Cutting this off at the boundary of the ends (we say \mathbf{A} restricted to $\{t\} \times Y_1$ is $\mathbf{A}_-(t)$, and restricted to $\{t\} \times Y_2$ is $\mathbf{A}_+(t)$), where the perturbations are all zero, this decomposes into three pieces:

$$\begin{aligned}
& 2((\text{cs} + f_{\pi_-})(\alpha_-) - (\text{cs})(\mathbf{A}_-(0))) \\
& + 2((\text{cs})(\mathbf{A}_-(0)) - (\text{cs})(\mathbf{A}_+(0))) \\
& + 2((\text{cs})(\mathbf{A}_+(0)) - (\text{cs} + f_{\pi_+})(\alpha_+))
\end{aligned}$$

First we study how the value of $\text{cs} + f_{\pi(t)}$ changes along the ends: consider the function $(\text{cs} + f_{\pi(t)})(\mathbf{A}_-(t))$. Here, $\pi(t) = \beta_0(t)\pi_-$. Then

$$\int d((\text{cs} + f_{\pi(t)})(\mathbf{A}_-(t))) = \text{cs}(\mathbf{A}_-(0)) - (\text{cs} + f_{\pi_-})(\alpha_-).$$

We may expand the integral as

$$\int (\text{cs} + f_{\pi(t)})'(\mathbf{A}_-(t))dt + \int d(\text{cs} + f_{\pi(t)}) \left(\frac{d}{dt} \mathbf{A}_-(t) \right).$$

Note that the second integral is the same as

$$\int \left\langle *F_{\mathbf{A}_-(t)} + \nabla_{\pi(t)}(\mathbf{A}_-(t)), \frac{d}{dt} \mathbf{A}_-(t) \right\rangle;$$

by assumption that $\mathbf{A}_-(t)$ satisfies the time-dependent gradient flow equations, this integral is $\int \left\| \frac{d}{dt} \mathbf{A}_-(t) \right\|^2 dt$. Because $f_{\pi(t)} = \beta_0(t)f_{\pi_0}$, the first integral is the same as

$$\int \beta_0'(t) f_{\pi_-}(\mathbf{A}_-(t)).$$

This is uniformly bounded above and below because the support of β_0' is compact, and the function f_{π_-} is bounded. Thus up to an amount bounded by a continuous function of π , the first part of the Chern-Simons difference above is $2 \int \left\| \frac{d}{dt} \mathbf{A}_-(t) \right\|^2$; a similar discussion implies the same of the the third part of the Chern-Simons difference and $2 \int \left\| \frac{d}{dt} \mathbf{A}_+(t) \right\|^2$.

Because $\pi_n \rightarrow \pi$, we obtain a bound on the energy of any π_n -instanton \mathbf{A}_n on the ends, uniform in n .

Altogether we see that an energy bound on \mathbf{A} gives a bound on the L^2 norm of the $F_{\mathbf{A}_n}$ on W' , as well as a bound on the L^2 norms of $\mathbf{A}_{\pm}(t)$. Because $\mathbf{A}_-(t)$ satisfies the equations

$$\frac{d}{dt} \mathbf{A}_-(t) = - *F_{\mathbf{A}_-(t)} - \nabla_{\pi(t)}(\mathbf{A}_-(t)),$$

and $\|\nabla_{\pi(t)} \mathbf{A}\|_{L^\infty}$ is uniformly bounded, we get a uniform L^2 bound on curvature on any compact piece of the negative end of W ; similarly with the positive end. Because this bound holds for all \mathbf{A}_n and is uniform in n , as in Proposition 4.8 we may conclude the existence of the appropriate gauge transformations.

What remains is to check is that the energy lost at the points is as expected.

Now suppose \mathbf{A}_n is a sequence of π_n -perturbed instantons on W and suppose there is a sequence of gauge transformations σ_n^* , defined on the complement of a finite set of points, so that $\sigma_n(\mathbf{A}_n) \rightarrow \mathbf{A}$ in L^p_1 on compact sets in the complement of those points, for $p \geq 2$. Call one of these points x ; we will show the energy lost at x is a non-negative multiple of $64\pi^2$; call this energy c . Abuse notation and rewrite $\mathbf{A}_n = \sigma_n(\mathbf{A}_n)$ for convenience, write $S(r)$ for the geodesic sphere of radius r around x , and write $\mathbf{A}_n(r)$ for its restriction to $S(r)$.

Because each $\mathbf{A}_n(r) \rightarrow \mathbf{A}(r)$ in the $L^2_{1/2}$ topology, we see that

$$\lim_n \text{cs}(\mathbf{A}_n(r)) \equiv \text{cs}(\mathbf{A}(r)) \pmod{64\pi^2\mathbb{Z}},$$

We also have, by definition,

$$\text{cs}(\mathbf{A}_n(r)) = \int_{B(r)} \text{Tr}(F_{\mathbf{A}_n}^2);$$

because \mathbf{A}_n is a π -instanton, we have

$$\text{cs}(\mathbf{A}_n(r)) = \int_{B(r)} \|F_{\mathbf{A}_n}\|^2 - 2 \int_{B(r)} \|\widehat{\nabla}_{\pi_n}(\mathbf{A}_n)\|.$$

Because $\|\widehat{\nabla}_{\pi}(\mathbf{A})\|_{L^\infty} \leq C\|\pi\|$, the second term goes to zero as $r \rightarrow 0$, and the first term limits to c , the energy lost at x .

Putting this together, we find

$$c = \lim_{r \rightarrow 0} \lim_{n \rightarrow \infty} \text{cs}(\mathbf{A}_n(r)) \equiv \lim_{r \rightarrow 0} \text{cs}(\mathbf{A}(r)) = 0 \pmod{64\pi^2\mathbb{Z}},$$

as desired. ■

Corollary 4.11. *The space $\overline{\mathcal{M}}_{\mathbf{E},\pi}^{\leq C}(\alpha_-, \alpha_+)$ of broken ideal trajectories on (W, \mathbf{E}) with uniform energy bound is compact.*

The following theorem summarizes the content of this section.

Theorem 4.12. *Let (Y, E) be a 3-manifold equipped with an $SO(3)$ -bundle and regular perturbation π . There is a natural compactification of the space of unparameterized trajectories*

$$\widetilde{\mathcal{M}}_{E,z,\pi}^0(\alpha, \beta) \subset \overline{\mathcal{M}}_{E,z,\pi}(\alpha, \beta)$$

as a compact $SO(3)$ -space equipped with equivariant endpoint maps.⁸ The added strata $\overline{\mathcal{M}} \setminus \widetilde{\mathcal{M}}$ are given as a union of fiber products of moduli spaces of lower dimension (in codimension equal to the number of intermediary orbits between α and β) and strata corresponding to Uhlenbeck bubbling.

The same is true for a cobordism (W, \mathbf{E}) equipped with a perturbation π which is regular on the bounding manifolds (Y_i, E_i) : the space of framed instantons $\widetilde{\mathcal{M}}_{\mathbf{E},z,\pi}(\alpha, \beta)$ has a natural compactification to the compact $SO(3)$ -space $\overline{\mathcal{M}}_{\mathbf{E},z,\pi}(\alpha, \beta)$.

We conclude with some discussions on energy bounds. On a cylinder $\mathbb{R} \times Y$, equipped with a regular perturbation π , there are finitely many critical orbits α . For a fixed pair α, β , the space of connections from α to β has

$$\pi_1 \widetilde{\mathcal{B}}_E^e(\alpha, \beta) \cong \mathbb{Z}$$

components. We may define two functions of $z \in \pi_1 \widetilde{\mathcal{B}}_E^e(\alpha, \beta)$. The first is the topological energy $\mathcal{E}^\pi(z)$, defined by taking a connection A in the component labelled by z , and evaluating

$$2((\text{cs} + f_\pi)(\beta) - (\text{cs} + f_\pi)(\alpha)) \in \mathbb{R};$$

we use the path A and the homotopy lifting property of the covering $\mathbb{R} \rightarrow \mathbb{R}/8\pi^2$ to pin down a real lift of this difference. The second is $\text{gr}_z(\alpha, \beta) \in \mathbb{Z}$. These are both

⁸Here recall that $\widetilde{\mathcal{M}}^0$ is $\widetilde{\mathcal{M}}/\mathbb{R}$, less the constant trajectories.

affine functions; writing $z \in \mathbb{Z}$, we have $\mathcal{E}^\pi(z) = 64\pi^2 z + c$ and $\text{gr}_z(\alpha, \beta) = 8z + c'$. In particular, we see that

$$\mathcal{E}^\pi(z) - 8\pi^2 \text{gr}_z(\alpha, \beta)$$

is constant. Therefore an energy bound is equivalent to an upper bound on $\text{gr}_z(\alpha, \beta)$. Suppose that the cobordism

$$(W, \mathbf{E}) : (Y_1, E_1) \rightarrow (Y_2, E_2)$$

is equipped with a perturbation π so that its restriction π_i to the ends is regular, and all π -perturbed instanton moduli spaces up to some energy bound C are cut out transversely. Unfortunately, we do not have a notion of perturbed topological energy for a cobordism, so we must be slightly more precise in our phrasing.

The set $\pi_0(\tilde{\mathcal{B}}_{\mathbf{E}}^e(\alpha, \beta))$ of components of the space of trajectories from α to β is affine over both $\pi_1(\tilde{\mathcal{B}}_{E_1}^e, \alpha)$ and $\pi_1(\tilde{\mathcal{B}}_{E_2}^e, \beta)$, and hence we may write the set of components as affine over \mathbb{Z} . Then $\text{gr}_z(\alpha, \beta)$ is affine with distortion 8. Now, on this same set, we have a *partially defined* function to \mathbb{R} , given by the analytic energy

$$\mathcal{E}_{\text{an}}^\pi(\mathbf{A}) = \int_W \|F_{\mathbf{A}} + \hat{\nabla}_\pi(\mathbf{A})\|^2$$

for a π -instanton \mathbf{A} in the homotopy class z . Of course, this is only defined when such π -instantons exist, but it is independent of the choice when they do. Furthermore, when defined, this function is affine with distortion $64\pi^2$. This is enough to say that an energy bound $\mathcal{E}_{\text{an}}^\pi(\mathbf{A}) \leq C$ is equivalent to a grading bound $\text{gr}_z(\alpha, \beta) \leq C'$ for π -instantons \mathbf{A} .

In what follows we will often only care about the case $\text{gr}_z(\alpha, \beta) \leq 10 - \dim \alpha$; there is a corresponding energy bound for π -instantons, for π fixed (or for π chosen from some compact set).

Because components of the moduli space with $\text{gr}_z(\alpha, \beta) < 3 - \dim \alpha$ are empty when cut out transversely, and bubbling decreases $\text{gr}_z(\alpha, \beta)$ by multiples of 8, we have the following.

Corollary 4.13. *When π is a regular perturbation and $\text{gr}_z(\alpha, \beta) \leq 10 - \dim \alpha$, the compactification $\overline{\mathcal{M}}_{\mathbf{E}, z, \pi}(\alpha, \beta)$ has no strata corresponding to Uhlenbeck bubbling.*

4.4. Reducible instantons on the cylinder and cobordisms. The goal of this section is to characterize the reducible π -ASD connections on a cobordism W . We begin with the especially simple case of the cylinder.

Proposition 4.14. *Let E be an $SO(3)$ -bundle over a rational homology sphere Y equipped with a Riemannian metric. Let \mathcal{P}_E denote a Banach space of perturbations on (Y, E) . By Proposition 3.5, if $\pi \in \mathcal{P}_E$ has $\|\pi\|$ sufficiently small, then there is a unique critical orbit α in every reducible component $\text{Red}(Y, E)$, which is the unique fully reducible orbit if the reducible component has a fully reducible point.*

Consider moduli spaces of reducible instantons on the cylinder $\mathbb{R} \times Y$, equipped with the constant perturbation π on the cylinder. By Proposition 1.6, the only moduli spaces which can possibly be nonempty are $\widetilde{\mathcal{M}}_{E, 0, \pi}(\alpha, \alpha)$, the spaces of trajectories from α to α in the trivial homotopy class. There is exactly one instanton in each of these: the constant trajectory at α . Therefore, the moduli space $\mathcal{M}_{E, 0, \pi}(\alpha, \alpha) = \emptyset$, by definition.

Proof. In other words, the constant trajectory (which is always a solution) remains the only solution. Because each component in $\text{Red}(Y, E)$ lies inside $\tilde{\mathcal{B}}_{E,0,k,\delta}^e(\alpha, \alpha)$ (the zero denoting the trivial homotopy class), and an instanton in this component has energy equal to zero, the trajectory must be constant, as in Corollary 4.3. ■

In fact, even in the case of cobordisms (or cylinders with nonconstant perturbations), the enumeration will be in terms of topological information.

Proposition 4.15. *Let W be an oriented Riemannian 4-manifold equipped with an $SO(3)$ -bundle \mathbf{E} and with one incoming cylindrical end (Y_1, E_1) and one outgoing cylindrical end (Y_2, E_2) . Suppose $b^+(W) = 0$.*

Consider the space of perturbations on W so that the perturbations π_{\pm} on the ends lie in $\mathcal{P}_{E,\delta}^{\text{red}}$, the subspace of perturbations so that the only eigenvalues of $\widehat{\text{Hess}}_{\alpha, \pi_{\pm}}^{\nu}$ acting on $\Omega^1(Y; i\mathbb{R})$ have absolute value at least δ ; abusing notation⁹, we still call this space $\mathcal{P}_{\delta}^{(4)}$.

If $\beta(w_2(\mathbf{E})) \neq 0$ or one of (Y_i, E_i) is admissible, then (W, \mathbf{E}) admits no π -perturbed reducible instantons in $\tilde{\mathcal{M}}_{\mathbf{E},k,\delta}$ whatsoever.

If \mathbf{E} is trivial and $b_1(W) = 0$, there is a neighborhood $\mathcal{P}_{\text{end}}^{(4)} \subset U \subset \mathcal{P}^{(4)}$ so that, for all $\pi \in U$, all fully reducible connections are cut out transversely. We write $U_{\delta} = U \cap \mathcal{P}_{\delta}^{(4)}$; this set contains 0 so long as δ is sufficiently small.

Fix $(\pi_-, \pi_+) \in \mathcal{P}_{\text{end}}^{(4)}$; write $U_{\delta,c} = U_{\delta} \cap \mathcal{P}_c^{(4)}$ for the portion contained inside the affine space through (π_-, π_+) over $\mathcal{P}_{\text{int}}^{(4)}$.

For all π in a residual subset of $U_{\delta,c}$ and for every component in $\text{Red}(W, \mathbf{E})$, the reducible π -instantons in this component comprise a finite set of orbits, which are cut out transversely inside the reducible locus.

If $\pi \in \mathcal{P}^{(4)}$ is sufficiently small (with respect to some constant $\epsilon_{W,C}$), then for every component in $\text{Red}(W, \mathbf{E})$ with $\mathcal{E}^{\pi}(\mathbf{A}) \leq C$, there is further a unique orbit of π -perturbed reducible instantons in that component; for those components $\text{Red}_{}(W, \mathbf{E})$ containing a fully reducible orbit, that orbit is the unique reducible instanton in that component.*

For bundles with $\beta w_2 \neq 0$ or with admissible ends, Proposition 1.7 shows that there are not even reducible *connections* with L^2 curvature. Thus there is only something interesting to say for cobordisms with rational homology sphere ends.

There is also a version when $b^+(W) > 0$. Our goal is to avoid reducibles in this case, as they cannot be cut out transversely.

Proposition 4.16. *Let W be an oriented Riemannian 4-manifold equipped with an $SO(3)$ -bundle \mathbf{E} , with one incoming cylindrical end (Y_1, E_1) and one outgoing cylindrical end (Y_2, E_2) . Suppose $b^+(W) > 0$.*

If $\beta(w_2(\mathbf{E})) \neq 0$ or one of (Y_i, E_i) is admissible, then (W, \mathbf{E}) admits no π -perturbed reducible instantons in $\tilde{\mathcal{M}}_{\mathbf{E},\pi}$ whatsoever.

If $b_1(W) < b^+(W)$, and \mathbf{E} is nontrivial but neither E_i is admissible, for any fixed $\pi_{\pm} \in \mathcal{P}_{\text{end},\delta}^{(4)}$, there exists a residual set of the affine slice $\pi \in \mathcal{P}_c^{(4)}$ for which there are no reducible π -ASD connections in any component z .

⁹This space is larger than the usual $\mathcal{P}_{\delta}^{(4)}$ defined in Definition 4.3, as the requirement is easier to satisfy; as the results will apply to the usual subset $\mathcal{P}_{\delta}^{(4)}$, and this larger set will not be used outside this section, we do not feel this abuse of notation is harmful.

If \mathbf{E} is trivial, each fully reducible component always has the full reducible as a solution, no matter the perturbation; if $b^+(W) > b_1(W)$, these fully reducible solutions are never cut out transversely in the $SO(2)$ -fixed locus, for any $\pi \in \mathcal{P}_\delta^{(4)}$.

We prove these simultaneously, much like Proposition 3.5: we define a Banach manifold of reducible solutions to the ASD equations, equipped with a map to U_δ (or $\mathcal{P}_\delta^{(4)}$, in the case that there are no fully reducible connections), which is proper; when $b_1(W) = b^+(W) = 0$, the projection is a local diffeomorphism above points near 0. Recall from Proposition 1.6 that the $SO(2)$ -fixed subspace is a disjoint union over copies of $\mathcal{B}_\eta(\alpha_-, \alpha_+)$, as η varies over certain complex line bundles. Note that this configuration space is unframed, despite being a subset of the framed configuration space of $SO(3)$ -connections on E . Setting this up takes some small amount of work because the Hilbert manifold the equation is defined on, $\mathcal{B}_\eta(\alpha_-, \alpha_+)$, depends on the limiting orbits, and hence depends on π (which determines the reducible critical orbits). While we discussed such a configuration space in the previous section, we did not show that this (unframed) quotient was a Banach manifold: that is usually not true, but is in the special case of $SO(2)$ -bundles.

Lemma 4.17. *Let W be an oriented Riemannian 4-manifold equipped with an $SO(3)$ -bundle \mathbf{E} and with one incoming cylindrical end (Y_1, E_1) and one outgoing cylindrical end (Y_2, E_2) , both rational homology spheres. Fix a lift of \mathbf{E} to a $U(2)$ -bundle $\tilde{\mathbf{E}}$; if none exist, then W admits no reducible connections whatsoever by Proposition 1.7. Write λ for the complex line bundle $\det(\tilde{\mathbf{E}})$ and fix a connection \mathbf{A}_0 on λ .*

There is a Banach manifold $\mathcal{P}_\delta \mathcal{B}_{\eta, k, \delta}$, whose objects are pairs (π, \mathbf{A}) defined as follows. First, $\pi \in \mathcal{P}_\delta^{(4)}$ is a perturbation as defined above; second, \mathbf{A} is an $L_{k, \delta}^2$ connection on $\tilde{\mathbf{E}}$ which respects a fixed splitting

$$\tilde{\mathbf{E}} \cong \eta \oplus (\lambda \otimes \eta^{-1})$$

and has π_\pm -flat limits on the ends, considered up to equivalence by the group of $L_{k+1, \delta, \text{ext}}^2$ maps $W \rightarrow S^1$, meaning that these maps exponentially decay to constant maps on the ends.

This Banach manifold comes equipped with a smooth submersion $\mathcal{P}_\delta \mathcal{B}_{\eta, k, \delta} \rightarrow \mathcal{P}_\delta^{(4)}$. There is a smooth vector bundle $\mathcal{S}_{k-1, \delta} \rightarrow \mathcal{P}_\delta \mathcal{B}_{\eta, k, \delta}$ with fiber isomorphic to $\Omega_{k-1, \delta}^{2, +}(i\mathbb{R})$, and a smooth Fredholm section $s(\pi, \mathbf{A}) = F_{\mathbf{A}}^+ + \widehat{\nabla} \pi(\mathbf{A})$.

Proof. Denote by η_i the restrictions of η to the ends. Write $\mathcal{P}_\delta \mathcal{C}_\eta$ for the set whose elements are triples $(\pi, \alpha_1, \alpha_2)$, where $\pi \in \mathcal{P}_\delta^{(4)}$ is a perturbation restricting to π_i on the ends, and $\alpha_i \in \mathcal{B}_{\eta_i}$ are equivalence classes of π_i -flat connections on η_i . With minimal change, Lemma 3.3 asserts that the map $\mathcal{P}_\delta \mathcal{C}_\eta \rightarrow \mathcal{P}_\delta^{(4)}$, sending the parameterized critical set of gauge equivalence classes of connections on η to the space of perturbations, is proper. That $\mathcal{P}_\delta \mathcal{C}_\eta$ is a manifold and that the projection above is a covering map follow quickly from the assumption that every $\pi \in \mathcal{P}_\delta^{(4)}$ is such that $\widehat{\text{Hess}}_{A_i, \pi_i}^\nu$ never has eigenvalues of absolute value at most δ . Thus, $\mathcal{P}_\delta \mathcal{C}_\eta$ has charts given by local sections ψ of the projection to $\mathcal{P}_\delta^{(4)}$. Fix such a chart, given by $V \subset \mathcal{P}_\delta^{(4)}$.

We may lift the above section $\psi : V \rightarrow \mathcal{P}_\delta \mathfrak{C}_\eta$ to a smooth map

$$p : V \rightarrow \mathcal{A}_{\eta_1, k} \times \mathcal{A}_{\eta_2, k},$$

so that $p(\pi_1, \cdot, \cdot, \pi_2)$ is a pair of connections in the same gauge equivalence as $\psi(\pi_1, \cdot, \cdot, \pi_2)$. That is, this is a smooth map picking out critical points for the end-perturbations π_i . This entire discussion depends only on the connections over the ends.

For this chart V , there are two possible situations: it could be that there are no reducible connections over the cobordism restricting to the configuration $\psi(\pi_1, 0, 0, \pi_2)$ on the ends, in which case $\mathcal{P}_\delta \mathcal{B}_{\eta, k, \delta}$ is empty above this chart and is tautologically a manifold.

More interestingly, if there is some reducible connection extending that configuration, then may choose a smooth map $e : \mathcal{A}_{\eta_1, k} \times \mathcal{A}_{\eta_2, k} \rightarrow \mathcal{A}_{\eta, k, \delta}^{(4)}$ so that $e(A_1, A_2)$ is a connection \mathbf{A} on η so that \mathbf{A} is constant and equal to A_i on the ends. Composing these, $e \circ p : \mathcal{P}_\delta^{(4)} \rightarrow \mathcal{A}_{\eta, k, \delta}^{(4)}$ is a smooth map choosing a connection $(e \circ p)(\pi)$ which is constant and equal to the π_\pm -critical connections $p(\pi)$ on the ends.

We may use this to define

$$V\mathcal{A}_{\eta, k, \delta}^{(4)} = \{(\pi, \mathbf{A}) \mid \pi \in V, \mathbf{A} - (e \circ p)(\pi) \in \Omega_{k, \delta}^1(W; i\mathbb{R})\}.$$

In particular, this Banach manifold is diffeomorphic to $V \times \Omega_{k, \delta}^1(W; i\mathbb{R})$. Clearly it comes equipped with a smooth projection to V . There is the trivial bundle $\mathcal{S}_{k-1, \delta}^+ = \Omega_{k-1, \delta}^{2,+}(W; i\mathbb{R})$ over $V \times \mathcal{A}_{\eta, k, \delta}^{(4)}$. Given a connection \mathbf{A} on η , the induced connection on $\eta \oplus (\lambda \otimes \eta^{-1})$ has curvature equal to $2F_{\mathbf{A}} - F_{\mathbf{A}_0}$, the non-central part of the curvature of $\eta \oplus (\lambda \otimes \eta^{-1})$. Then the section $s : V\mathcal{A}_{\eta, k, \delta}^{(4)} \rightarrow \mathcal{S}_{k-1, \delta}$ is given by

$$(\pi, \mathbf{A}) \mapsto 2F_{\mathbf{A}}^+ - F_{\mathbf{A}_0}^+ + \widehat{\nabla}_\pi(\mathbf{A}).$$

This space a smooth action by the Banach Lie group $\mathcal{G}_{\eta, k+1, \delta, \text{ext}}$, preserving the projection map to $\mathcal{P}_\delta^{(4)}$ and the section s . While the action is not free, the stabilizer is the same at every point: it is the group of constant maps $W \rightarrow S^1$. The action of $\mathcal{G}_{\eta, k+1, \delta}$ factors through its quotient by the subgroup of constant maps, and the action of *this* quotient group is free. The fact that ‘ \mathbf{A} -harmonic gauge transformations’ are the same for all \mathbf{A} is a convenient and unique aspect of the case of $SO(2)$ -bundles.

Now it is easy to verify that this action is proper, such that each orbit has closed complemented tangent space, giving us a quotient manifold $V\mathcal{B}_{\eta, k, \delta}$ with a projection map and a section of a vector bundle $\mathcal{S}_{k-1, \delta}$ (still the trivial bundle). That the section on $V\mathcal{B}_{\eta, k, \delta}$ is Fredholm follows from the same fact fiberwise, on a fixed Banach manifold $\mathcal{B}_{\eta, k, \delta}(\alpha_1, \alpha_2)$. Patching these together for different charts V of $\mathcal{P}_\delta \mathfrak{C}_\eta$ gives us a smooth Banach manifold structure on $\mathcal{P}_\delta \mathcal{B}_{\eta, k, \delta}$. \blacksquare

Note that this Banach manifold is precisely the same as a component of the $SO(2)$ -fixed point space of $\mathcal{P}_\delta \widetilde{\mathcal{B}}_{\mathbf{E}, k, \delta}$, expressed in terms of the $U(2)$ model.

In the case of reducible trajectories, the compactness properties are especially strong. Note that there are no broken or ideal trajectories in the following statement.

Lemma 4.18. *The map $\mathcal{P}_\delta \mathcal{M}_{k, \delta} \rightarrow \mathcal{P}_\delta^{(4)}$ is proper.*

Proof. This is an application of Proposition 4.10. Start by choosing a sequence (π_n, \mathbf{A}_n) of perturbed reducible instantons so that $\pi_n \rightarrow \pi$.

First, observe that because every fiber of $\mathcal{P}_\delta \mathcal{B}_{\eta,k,\delta} \rightarrow \mathcal{P}_\delta^{(4)}$ is *connected*, the analytic energy $\mathcal{E}^\pi(\mathbf{A})$, for a reducible π -instanton \mathbf{A} , depends only on π , and depends continuously on π at that. In particular, we have $\mathcal{E}^\pi(\mathbf{A}_n) \leq C$ for some constant C (without already assuming an energy bound on the \mathbf{A}_n).

In particular, by the compactness theorem there is a subsequence of \mathbf{A}_n and a broken ideal trajectory \mathbf{A} such that a subsequence of \mathbf{A}_n converges to \mathbf{A} as a broken bubble-limit. Because all of the \mathbf{A}_n are reducible, so is \mathbf{A} . Bubbling cannot change the topological type of η (a line bundle over a manifold is determined by its restriction to the complement of any finite set, so long as the manifold is dimension at least 2), and so for large n we have that \mathbf{A}_n lies in the same component as \mathbf{A} . Thus, we know that $\mathcal{E}^{\pi_n}(\mathbf{A}_n) \rightarrow \mathcal{E}^\pi(\mathbf{A})$, implying no energy is lost at bubble-points.

So the only noncompactness can arise from breaking of trajectories. We already know by Proposition 4.14 that on the cylinder $\mathbb{R} \times Y$, where Y is equipped with a weakly admissible bundle and regular perturbation, reducible trajectories are constant. So after passing to a subsequence and applying suitable gauge transformations, the configuration \mathbf{A}_n converges in $L^2_{k,\delta}$ an honest reducible connection on η with the same limits at $\pm\infty$, as desired, with no energy lost at the ends. ■

The following unique continuation lemma is quoted essentially verbatim from [KM07, Lemma 7.1.3]; we will use it frequently.

Lemma 4.19. *Let H be a real Hilbert space and $I = [t_1, t_2]$ a closed interval, equipped with a family $L(t)$ of unbounded operators $D \rightarrow H$ with common dense domain $D \subset H$, so that $L(t) = L_+(t) + L_-(t)$ for L_+ symmetric on D and $L_-(t)$ skew-adjoint and bounded on the whole of H . Further suppose that the time-derivative $L'(t)$ is a well-defined operator $D \rightarrow H$, defined pointwise, which has a bound*

$$\|L'(t)x\| \leq C_1(\|L(t)x\| + \|x\|)$$

uniform in $t \in I$ and $x \in D$.

Let $f : I \rightarrow H$ be a continuous map and $z : I \rightarrow D$ be a solution of the equation $z' + Lz = f$, where we have the bound on the inhomogeneous term $\|f(t)\| \leq C_2\|z(t)\|$, uniform in t . Then if $z(t) = 0$ for some $t \in I$, then z is identically zero.

We will first justify the open set U_δ used in the statement of Proposition 4.15.

Lemma 4.20. *Suppose $b_1(W) = 0$. If $b^+(W) = 0$ or \mathbf{E} is nontrivial, then all fully reducible connections are cut out transversely for any end perturbation; therefore, there is a connected open set*

$$\mathcal{P}_{\text{end}}^{(4)} \subset U \subset \mathcal{P}^{(4)}$$

for which every fully reducible connection is cut out transversely. When $\delta > 0$ is sufficiently small, the intersection $U_\delta = U \cap \mathcal{P}_\delta^{(4)}$ contains 0.

Proof. If \mathbf{E} is nontrivial, there are no fully reducible connections; if $b^+(W) = 0$, then for any fully reducible connection \mathbf{A} , we have

$$\text{coker}(Q_{\mathbf{A},\pi}^\vee) = \mathcal{H}_{\mathbf{A}}^+ = \mathcal{H}^+ = 0$$

by the Hodge theorem. Therefore every fully reducible connection is cut out transversely above 0; furthermore, because we assume the 3-dimensional holonomy perturbations f_π vanish in a neighborhood of fully reducible connections, we necessarily

have $D_A \nabla_\pi = 0$ for any fully reducible connection A . In particular, because the end perturbations take the form

$$\widehat{\nabla}_\pi \mathbf{A} = \beta_0(t)(dt \wedge \nabla_\pi \mathbf{A}(t)),$$

we see that $D_{\mathbf{A}} \widehat{\nabla}_\pi = 0$ for any full reducible \mathbf{A} on W and any end perturbation π ; so $Q_{\mathbf{A},\pi}$ is surjective iff $Q_{\mathbf{A},0}$ is.

Because the projection from the parameterized configuration space of fully reducible connections to $\mathcal{P}_\delta^{(4)}$ is a proper local diffeomorphism, we see that every fully reducible connection is cut out transversely in a neighborhood of $\mathcal{P}_{\text{end}}^{(4)}$, and we call this neighborhood U .

To say that $0 \in \mathcal{P}_\delta^{(4)}$ just means that for the reducible flat connections α , all eigenvalues of $\widehat{\text{Hess}}_\alpha^\nu$, acting on $\Omega^1(W; i\mathbb{R})$, have absolute value larger than δ . Because the kernel of this operator is $\mathcal{H}^1(W) = 0$, and the eigenvalues form a discrete closed set, we see that for the finitely many reducible flat connections α , all eigenvalues have absolute value larger than some $\delta > 0$, as desired. \blacksquare

When $b_1(W) = 0$ and $b^+(W) > 0$ but \mathbf{E} is trivial, the set U_δ is empty: full reducibles are always π -perturbed ASD connections, regardless of π , but the index of $D_{\mathbf{A},\pi}$ is always negative, so they cannot be cut out transversely.

Lemma 4.21. *Let (W, \mathbf{E}) be a Riemannian manifold with two cylindrical ends, both of which are modelled on rational homology spheres, with $b_1(W) = 0$. Suppose either that $b^+(W) = 0$ or that \mathbf{E} is nontrivial, and fix a component of the $SO(2)$ -fixed point space, labeled by a line bundle η . We write $U_\delta \subset \mathcal{P}_\delta^{(4)}$ for the open set in which all fully reducible connections are cut out transversely; in the latter case, when \mathbf{E} is non-trivial, this containment is an equality.*

Then the zero set

$$U_\delta \mathcal{M}_\eta \subset U_\delta \mathcal{B}_{\eta,k,\delta}$$

is a smooth Banach submanifold.

Similarly, for any fixed $(\pi_-, \pi_+) \in \mathcal{P}_{\text{end},\delta}^{(4)}$, the zero set $U_{\delta,c} \mathcal{M}_\eta$ is a Banach submanifold of $U_\delta \mathcal{B}_{\eta,k,\delta}$.

Proof. By assumption, we need only consider the $SO(2)$ -reducible ASD connections which are not fully reducible.

We want to show that

$$Q'_{\mathbf{A},\pi} = (Q_{\mathbf{A},\pi}, \widehat{\nabla}_{\pi'}(\mathbf{A})) : \Omega_{k,\delta}^1(W; i\mathbb{R}) \oplus \mathcal{P}^{(4)} \rightarrow \Omega_{k-1,\delta}^{2,+}(W; i\mathbb{R})$$

is surjective for all reducible π -instantons \mathbf{A} in the component corresponding to η (of class $L_{k,\delta}^2$).

To see this, consider $\psi \in \text{coker}(Q'_{\mathbf{A},\pi})$; necessarily, $\psi \in \text{coker}(Q_{\mathbf{A},\pi})$, and in particular $Q_{\mathbf{A},\pi}^* \psi = 0$. We may write this as

$$Q_{\mathbf{A},0}^* \psi = -(D_{\mathbf{A}} \widehat{\nabla}_\pi)^* \psi;$$

because $D_{\mathbf{A}} \widehat{\nabla}_\pi$ extends to L^2 over all of W (but to all L_j^2 with $-k \leq j \leq k$ on the ends), the usual elliptic bootstrapping shows that $\psi \in L_1^2(W)$, and furthermore that $\psi \in L_{k,\delta}^2$ on the ends.

Recall that the interior holonomy perturbations are supported inside a compact submanifold $W' \subset W$, whose complement is the portion $[1, \infty) \times Y$ of the cylindrical

ends. Now, the argument in [Kro05, Lemma 3.7] shows that so long as π is not fully reducible, the image of $\widehat{\nabla}_{\pi'}(\mathbf{A})$ is dense in continuous sections of $\Lambda^{2,+}(i\mathbb{R})$ over W' which vanish on the boundary; in particular, because $\psi \in \text{coker}(Q')$, it should be L^2 orthogonal to all of these, and hence vanish on the whole of W' .

If π is fully reducible, these perturbations are identically zero. The hypothesis of $SO(2)$ -reducibility is used to ensure that, at any point in W' , one may find a collection of loops so that $\text{Hol}_{\mathbf{A}} \notin (\pm 1)^N$; at these points of $U(1)^N \subset SU(2)^N$, any equivariant map $r : SU(2)^N \rightarrow \mathfrak{su}(2)$ must necessarily send $r((\pm 1)^N) = 0$; points in $U(1)^N \setminus (\pm 1)^N$ are sent into $\mathfrak{u}(1) \subset \mathfrak{su}(2)$, but are not otherwise constrained.

Now ψ satisfies the perturbed ASD equations on the end, written $[0, \infty) \times Y$, with $\psi \in L_k^2$ and $\psi|_{\{0\} \times Y} = 0$. So long as $k \geq 2$, the map ψ defines an element of

$$C^1([0, \infty), \Omega_0^1(Y)) \cap C^0([0, \infty), \Omega_1^1(Y));$$

here the subscripts k, δ indicate we are taking 1-forms of Sobolev class L_k^2 . Further, it satisfies the equations

$$d^*\psi(t) = 0\psi'(t) - *d\psi(t) = D_{\mathbf{A}(t)}\nabla_{\pi(t)}\psi(t).$$

This ODE satisfies the conditions of Lemma 4.19 so long as $\pi(t)$ and $\mathbf{A}(t)$ are C^1 paths; the fact that $\mathbf{A}(t)$ is a C^1 path in $\Omega_0^1(Y)$ follows because \mathbf{A} is L_2^2 , and $\pi(t) = \beta_0(t)\pi_\infty$ is smooth.

Therefore $\psi(0) = 0$ implies that $\psi = 0$ on this entire end. Repeating for both ends, we find that any $\psi \in \text{coker}(Q'_{\mathbf{A}, \pi})$ is necessarily zero, as desired. \blacksquare

Proof of Propositions 4.15 and 4.16. Until the final paragraph of this proof, we assume that $b_1(W) = 0$, and that \mathbf{E} is nontrivial or $b^+(W) = 0$.

Fix $(\pi_-, \pi_+) \in \mathcal{P}_{\text{end}, \delta}^{(4)}$.

The set $U_{\delta, c}\mathcal{M}_{\eta, k, \delta}$ is a smooth submanifold of $U_{\delta, c}\mathcal{B}_{\eta, k, \delta}$. For fixed π , the set of all π -instantons \mathbf{A} are cut out transversely for all in $\mathcal{B}_{\eta, k, \delta}$ if and only if the projection

$$p_\eta : U_{\delta, c}\mathcal{M}_\eta \rightarrow U_{\delta, c}$$

has π as a regular value. By the Sard-Smale theorem, such π are in large supply; they form a dense set. Furthermore, the projection p_η is proper by Lemma 4.18. Therefore, the regular values form an open set of $U_{\delta, c}$.

In the case that $b^+(W) > 0$, this means that $\mathcal{M}_{\eta, \pi} = \emptyset$ for π a regular value, as the index of the operator $D_{\mathbf{A}, \pi}$ is negative. In the case that $b^+(W) = 0$, the moduli space $\mathcal{M}_{\eta, \pi}$ for regular π is a finite set, and 0 is a regular value.

What remains is to see what the solutions are when $\pi = 0$; then by properness, we immediately understand the solutions for all sufficiently small π . But for a cobordism (W, \mathbf{E}) with $b^+(W) = 0$ but $b_1(W) = n$, the space of reducible unperturbed ASD connections in each reducible component is the torus T^n ; this is [Dae15, Lemma 1.6], and again follows from Hodge theory, as the ASD equation is affine in the $SO(2)$ -reducible case. See also the related [Frø02, Lemma 2]. In the case we most care about, $b_1(W) = b^+(W) = 0$, this means that there is a unique π -perturbed reducible ASD connection in each reducible component, so long as π is sufficiently small.

Intersecting over all line bundles η , these open dense sets become residual sets. For the statement of Proposition 4.15, we want the enumeration above to hold for all line bundles up to a certain energy level; the assumption that $b^+(W) = 0$ implies that there are only finitely many line bundles whose topological energy is bounded

above by a given constant C , and so intersecting over the finitely many relevant open sets, we find an open set containing 0 so that the enumeration of reducibles up to energy C is correct for all π in this open set.

If \mathbf{E} is trivial, the trivial connection \mathbf{A} is always an ASD connection, no matter the perturbation. For $\pi \in \mathcal{P}_\delta^{(4)}$, we have $\text{ind}(D_{\mathbf{A},\pi}) = \text{ind}(D_{\mathbf{A},0})$. When $b^+(W) > b_1(W)$, this index is negative, and so \mathbf{A} cannot be cut out transversely. ■

4.5. Index calculations. With a strong grasp on the regularity properties of reducible instantons (internal to the reducible locus), we move on to understanding the index of the ASD operator *normal* to the reducible locus.

Let (W, \mathbf{E}) be a cobordism from (Y_1, E_1) to (Y_2, E_2) , thought of as a manifold with cylindrical ends, and let π be a perturbation so that the ends π_\pm are regular perturbations on Y_i . Suppose δ is larger than the least nonzero eigenvalue of the extended Hessian operators of the π_\pm -critical points. Choose π_\pm sufficiently small that Proposition 3.4 holds; this is guaranteed by choosing $\pi \in \mathcal{P}_\delta^{(4)}$.

If \mathbf{A} is a π -perturbed instanton, we would like to compute the index of

$$Q_{\mathbf{A},\pi} : \Omega_{k,\delta}^1(W; \mathfrak{g}_E) \rightarrow \Omega_{k-1,\delta}^{2,+} \oplus \Omega_{k-1,\delta}^0.$$

In fact, by invariance properties of the index, this only depends on the component \mathbf{A} sits inside and on the value of π on the ends. Correspondingly, we assume π has no component corresponding to an interior perturbation: we suppose that π is supported on the ends.

The most useful tool for computing the index of a differential operator on a compact manifold with boundary is the Atiyah-Patodi-Singer index theorem. Let X be a compact manifold with boundary, with metric of product type near the boundary. Let D be an elliptic differential operator acting on sections of bundles V_1, V_2 over W' (thought of as a compact manifold with boundary) with specified isomorphisms near the boundary components $[0, \infty) \times Y$ from the given operator to $\frac{d}{dt} + A$, where A is a self-adjoint elliptic operator whose *nonzero* eigenvalues c have $|c| > \delta$. Then one may consider D as an operator

$$D_{APS} : L_1^2(X, V_1; P) \rightarrow L^2(X, V_2),$$

where here P means that we demand that *the restriction of $\sigma \in L_1^2(W, V_1; P)$ to the boundary lies in the subspace spanned by the negative eigenvalues*: that is, if P is the projection operator corresponding to the eigenspaces with $\lambda \geq 0$, we demand that $P\sigma = 0$. Then even if A , the operator at the boundary, has kernel, D_{APS} is a Fredholm operator with a well-defined index. If

$$\hat{X} := X \cup_{\partial X} [0, \infty) \times \partial X$$

with the appropriate product metric on the ends, we may just as well consider D as an operator on \hat{X} . We write $I(D)$ to mean the index of D , computed as a map between weighted Sobolev spaces of weight $\delta > 0$, and $I_{APS}(D)$ to mean the index computed as D_{APS} .

Lemma 4.22. *As long as δ is less than any nonzero eigenvalue of the boundary operators A_i , we have $I_{APS}(D) = I(D)$.*

Proof. This is essentially [APS75a, Proposition 3.11]. When thought of as a map on weighted Sobolev spaces, $\ker(D)$ clearly consists of L^2 solutions of Df on \hat{X} . Conversely, every L^2 solution has as much smoothness as the operator itself, and

in the asymptotic expansion $\sum a_\lambda e^{-t\lambda} \phi_\lambda$ on the ends, a_λ must be zero when $\lambda \leq 0$ if the solution is to be L^2 . So a solution decays exponentially with weight at least $\delta > 0$, where δ is less than any positive eigenvalue of A .

The kernel of D^* is computed in weighted Sobolev spaces as a subspace of $L^2_{k,-\delta}$. To decay exponentially slower than $e^{-\delta t}$, we must demand that $a_\lambda = 0$ for $\lambda \geq \delta$. Because positive eigenvalues are at least as large as δ , we see that a solution is, on the ends, a sum of an exponentially decaying solution and a solution which is constant in time. This is precisely the kernel of D^* on extended L^2 sections, as in the computation of $I(D_{APS})$ given in [APS75a, Proposition 3.11]. Because the two kernels and cokernels agree, we have $I_{APS}(D) = I(D)$. ■

So it suffices to compute $I_{APS}(Q_{\mathbf{A},\pi})$, and for this, we have the Atiyah-Singer index theorem. We will need some simple lemmas first before we can calculate this index.

Lemma 4.23. *Suppose W is a compact manifold with boundary of product type, equipped with an elliptic operator D that is of product type near the boundary. Suppose further that there is an oriented closed submanifold $Y \subset W$ with a neighborhood of product type so that the operator may be written as $d/dt + A$ on this neighborhood. Write D_1 for the operator on the compact manifold whose positive boundary contains Y , and D_2 for the operator on the compact manifold whose negative boundary contains Y ; say A is the operator at Y . Then*

$$I_{APS}(D) = I_{APS}(D_1) + I_{APS}(D_2) + \dim \ker(A).$$

Proof. This follows immediately from the index theorem itself, [APS75a, Theorem 3.10]. The only term which is not additive is $-h/2$, where h is the kernel of the boundary operator. Because we are contributing two extra copies of $-\dim \ker(A)/2$ on the right, we counterbalance that by adding $\dim \ker(A)$. ■

Finally, we will need to know how the index of spectral flow is computed.

Lemma 4.24. *Let $D = d/dt + A_t$ be an operator on $[0, 1] \times Y$, where A_t is a time-dependent self-adjoint elliptic operator, possibly with kernel. We define the spectral flow $\text{sf}(A_t)$ to be the intersection number of the graph of the spectra with the line $\lambda(A_t) = -\delta$, which A_0 and A_1 do not intersect. This is the aggregate number of eigenvalues that go from $\leq -\delta$ to ≥ 0 , counted with sign.*

Then

$$I_{APS}(D) = \text{sf}(A_t) - \dim \ker(A_1).$$

Proof. This may be proved using separation of variables. A similar formula is stated below [APS96, Theorem 7.4], only giving the spectral flow term; their argument uses the periodic boundary conditions on $[0, 1] \times Y$, which corresponds to the projection to nonnegative eigenvalues at $t = 0$ and positive eigenvalues at $t = 1$, whereas the APS boundary conditions stated above use the spectral projection to *nonnegative* eigenvalues at $t = 1$. The second operator has smaller domain, of codimension $\dim \ker(A_1)$, having included the demand that an element f of the domain projects nontrivially to $\ker(A_1)$. Because index is additive under composition and the inclusion of this subspace has index $-\dim \ker(A_1)$, the theorem follows. ■

Now W may be decomposed as the union of a compact manifold X and the two cylindrical ends. This decomposition will provide the desired calculation. Before stating the result, we recall the definition of one of the terms that will appear.

Definition 4.10. Let α be an orthogonal (resp. unitary) flat connection on a real (complex) vector bundle E of dimension n over a closed manifold Y . The Atiyah-Patodi-Singer ρ -invariant of α is defined to be

$$\eta_\alpha(0) - n \cdot \eta_\theta(0).$$

Here η_α is the Atiyah-Patodi-Singer η invariant associated to the α -twisted signature operator

$$\begin{aligned} B : \Omega^{\text{ev}}(W, E) &\rightarrow \Omega^{\text{ev}}(W, E) \\ B\omega &= (-1)^{\deg \omega} (*d_\alpha - d_\alpha*)\omega, \end{aligned}$$

as is studied in [APS75b]. This constant is denoted $\rho(\alpha)$.

If \mathbf{E} has a reduction labelled by $\{\zeta_1, \zeta_2\}$, where ζ_i are complex line bundles and

$$c_1(\zeta_1) + c_1(\zeta_2) = \lambda,$$

we define $\rho(\{\zeta_1, \zeta_2\})$ to be the ρ -invariant of the induced $SO(3)$ -flat connection on \mathbf{E} , given as $\mathbb{R} \oplus (\zeta_1 \otimes \zeta_2^{-1})$.

The following equivalent computations of the ρ -invariant given here are listed in the discussion in [HK11, Section 2.9]. We state them without proof.

Lemma 4.25. Suppose $E \rightarrow Y$ is an $SO(3)$ -bundle, equipped with a reduction $E \cong \mathbb{R} \oplus (\zeta_1 \otimes \zeta_2^{-1})$. Write $\kappa = \zeta_1 \otimes \zeta_2^{-1}$, thought of as a flat complex line bundle. If $z_i = c_1(\zeta_i)$, then $\rho(\{z_1, z_2\})$ coincides with $\rho(\kappa) + \rho(\kappa^{-1})$. Further, if

$$(W, \mathbf{E}) : (Y_1, E_1) \rightarrow (Y_2, E_2)$$

is a cobordism equipped with a compatible reduction

$$\mathbf{E} \cong \mathbb{R} \oplus (\zeta_1^W \otimes (\zeta_2^W)^{-1}).$$

Again, write $\kappa = \zeta_1^W \otimes (\zeta_2^W)^{-1}$, write κ_i for the restrictions to each component, and write $z_i^W = c_1(\zeta_i^W)$. Then

$$\rho(\kappa_2) - \rho(\kappa_1) = \text{Sign}(W) - \text{Sign}_\kappa(W),$$

where $\text{Sign}_\kappa(W)$ is the index of the twisted signature operator on W .

Now we calculate the Atiyah-Patodi-Singer index of the operator $Q_{\mathbf{A}, \pi}$, where \mathbf{A} is a reducible trajectory on the 4-manifold W . By a homotopy, put \mathbf{A} in a form so that it is constant sufficiently far on the ends, and constant at the boundary $\{0\} \times Y$ of the ends, constant at the unique unperturbed flat connection on Y in the corresponding component of reducibles. For the statement of the following theorem, recall Definition 3.3 of signature data on a pair (Y, E) , and in particular signature data associated to a perturbation.

Proposition 4.26. Suppose $(W, \mathbf{E}) : (Y_1, E_1, \pi_1) \rightarrow (Y_2, E_2, \pi_2)$ is a cobordism (with cylindrical ends) between rational homology spheres, equipped with a perturbation $\pi \in \mathcal{P}_\delta^{(4)}$ which is regular on the ends. Then by Definition 4.3, the perturbations on the ends are sufficiently small that Proposition 3.4 applies.

Let \mathbf{A} be a reducible connection on \mathbf{E} which, sufficiently far on the ends, is constant and equal to a π -critical point; suppose \mathbf{A} is in the component labeled by $r \in \text{Red}(W, \mathbf{E})$; let $r_i \in \text{Red}(Y_i, E_i)$ be the restrictions to the ends. Suppose r corresponds to the pair of complex line bundles $\{\eta_1, \eta_2\}$, corresponding to cohomology classes $\{z_1, z_2\}$. Because the ends are rational homology spheres, we may write

$z_i \in H_c^2(W; \mathbb{Q})$; the compactly supported cohomology ring has a cup-product with values in \mathbb{Q} .

Let $S(r) = 1$ if r is a component of $SO(2)$ -reducibles and $S(r) = 3$ if r contains a full reducible.

Then

$$I_{APS}(Q_{\mathbf{A}, \pi}) = -2(z_1 - z_2)^2 + 3(b^1 - b^+) + \frac{\rho(r_2) - \rho(r_1)}{2} + \frac{\sigma_{\pi_2}(r_2) - \sigma_{\pi_1}(r_1)}{2} - \frac{S(r_1) + S(r_2)}{2}.$$

Proof. Let W_N be the compact submanifold of W given by including the first $[0, N]$ of each end; for N sufficiently large, the operator $Q_{\mathbf{A}}^\nu$ is of product type near the boundary.

By splitting W_N into three pieces, $[-N, 0] \times Y_1 \cup X \cup [0, N] \times Y_2$, we may decompose the operator $Q_{\mathbf{A}, \pi}$ into its pieces on these three corresponding ends. Write Q_\pm for the pieces of $Q_{\mathbf{A}, \pi}$ on the corresponding ends, Q for the piece on W' , and A_i the connections \mathbf{A} restricts to on $\{0\} \times Y_i$.

First we calculate $I_{APS}(Q_-)$. Write $\mathbf{A}(t)$ for the restriction of \mathbf{A} to the negative end, where $t \in [-N, 0]$. The path $\mathbf{A}(t)$ is homotopic to a path $\mathbf{A}_f(t)$ so that $\mathbf{A}_f(t)$ is in the unique gauge equivalence class of $\pi(t)$ -flat connection in its reducible component: This is the path used in the proof of Proposition 3.9, and our assumption that π_i on the ends lie in the sets $\mathcal{P}_{E_i, \delta}$ is to ensure this doesn't go awry. (See the remark immediately after Definition 4.3.) In particular, because $Q_{\mathbf{A}, \pi}$ is the operator $\frac{d}{dt} + \widehat{\text{Hess}}_{\mathbf{A}(t), \pi(t)}$, this spectral flow is by definition equal to the function $2N_{\pi_1}(r_1)$ defined in that proposition. Thus by Lemma 4.24,

$$I_{APS}(Q_-) = 2N_{\pi_1}(r_1) - \dim \ker \widehat{\text{Hess}}_{A_1}.$$

A similar discussion gives

$$I_{APS}(Q_+) = -S(r_2) - 2N_{\pi_2}(r_2).$$

What remains is to apply the index theorem to D on X . If we write Q_θ to mean the corresponding ASD operator for the trivial connection, this can be read off from [MMR94, Proposition 8.4.1] as giving

$$I_{APS}(D) - 3I_{APS}(D_\theta) = -2p_1(\mathbf{A}) + \frac{\rho(r_2) - \rho(r_1)}{2} + 3 - \frac{h_1 + h_2}{2}.$$

Here $h_i = \dim \ker \widehat{\text{Hess}}_{A_i}$. A detailed computation is provided in [HK11, Proposition 2.6], but note that our sign conventions on boundary orientations and the definition of p_1 are the negative of theirs. Hodge theory provides the equality

$$I_{APS}(Q_\theta) = -(1 - b^1 + b^+),$$

so we obtain

$$I_{APS}(Q) = -2p_1(\mathbf{A}) + 3(b^1 - b^+) + \frac{\rho(r_2) - \rho(r_1)}{2} - \frac{h_1 + h_2}{2}.$$

Summing over these and including boundary kernel terms as in Lemma 4.23, we obtain

$$-2p_1 \mathbf{E} + 3(b^1 - b^+) + \frac{\rho(r_2) - \rho(r_1)}{2} + 2N_{\pi_1}(r_1) - \frac{h_1}{2} + \frac{h_2}{2} - 2N_{\pi_2}(r_2) - S(r_2).$$

Now \mathbf{A} induces a reduction $\mathbf{E} \cong \mathbb{R} \oplus (\eta_1 \otimes \eta_2^{-1})$. Pontryagin classes are preserved under stabilization, so we want to compute $p_1(\eta_1 \otimes \eta_2^{-1})$. For a complex line bundle

ζ , we have $p_1\zeta = c_1(\zeta)^2$, and considering the classes z_i in $H_c^2(W; \mathbb{Q})$ corresponding to η_i , we obtain $p_1\mathbf{E} = (z_1 - z_2)^2$.

We focus now on the last few terms. If we write

$$D(r_i) = \dim_{\mathbb{R}} H^1(Y_i; \eta_1 \otimes \eta_2^{-1}),$$

then we have $h_1 = D(r_1) + S(r_1)$ and $h_2 = D(r_2) + S(r_2)$. Because the dimension of a vector space equipped with a nondegenerate symmetric bilinear form is the number of positive eigenvalues plus the number of negative eigenvalues, we see that $D(r) - 4N_{\pi}(r)$ is the number of positive eigenvalues less the number of negative eigenvalues, and hence

$$2N_{\pi_1}(r_1) - \frac{h_1}{2} + \frac{h_2}{2} - 2N_{\pi_2}(r_2) - S(r_2) = \frac{\sigma_{\pi_2}(r_2) - \sigma_{\pi_1}(r_1)}{2} - \frac{S(r_1) + S(r_2)}{2}.$$

■

The above calculation did not at all depend on the fact that the connection \mathbf{A} was reducible. In general, the same formula holds, where if the restriction of \mathbf{A} to one of the ends is irreducible, we write $S(r_1) = 0$ for that end; and the first term should be read $-2p_1(\mathbf{E})$, defined as a curvature integral for a connection \mathbf{A} on the compact manifold W which restricts to the relevant flat connections on a neighborhood of the boundary.

Note that if we choose a different connection \mathbf{A}' that restricts to the same flat connections on the boundary, the only thing that can possibly change in the index formula is $-2p_1(\mathbf{E})$. Fix a base connection \mathbf{A}_0 , equal to the desired flat connections near the boundary. Consider the double of W , with \mathbf{A}_0 on one half and, on the other half, an arbitrary connection \mathbf{A} restricting to the desired flat connections on ∂W . Clearly $-2p_1$ of this new connection on a closed manifold is $-2p_1(\mathbf{A}_0) - 2p_1(\mathbf{A})$. Note that $-2p_1\mathbf{E}$ is constant mod 8 on a closed manifold: it reduces to

$$2(w_2\mathbf{E})^2 \in H^4(X; 2\mathbb{Z}/8) = 2\mathbb{Z}/8,$$

where here we take the Pontryagin square to write $w_2^2 \in \mathbb{Z}/4$. Then we immediately have the following corollary.

Corollary 4.27. *The relative grading $\text{gr}_z(\alpha, \beta) \in \mathbb{Z}$ is independent, modulo 8, of the choice of z . Therefore, we may unambiguously write $\text{gr}(\alpha, \beta) \in \mathbb{Z}/8$.*

The above index calculation in mind for reducible connections, we combine the ρ and signature terms into a single function.

Definition 4.11. *Let (Y, E) be a rational homology sphere equipped with an $SO(3)$ -bundle and small regular perturbation π . We define the perturbed ρ -invariant to be $\rho_{\pi}(r) = \rho(r) + \sigma_{\pi}(r)$, where σ_{π} is the signature datum associated to π as in Definition 3.3.*

Now if we write $D_{\mathbf{A}, \pi}^{\nu}$ for the normal ASD operator, the linearization of the section defining the moduli spaces in $\tilde{\mathcal{B}}_{\mathbf{E}, z, k, \delta}^e(\alpha_1, \alpha_2)$ restricted to the normal space to an orbit, then we see by the discussion at the beginning of Section 4.2 that

$$I(D_{\mathbf{A}, \pi}^{\nu}) = I_{APS}(Q_{\mathbf{A}}^{\nu}) + S(r),$$

where $S(r) = \dim \mathfrak{g}_{\mathbf{A}}$ is the dimension of the space of \mathbf{A} -parallel gauge transformations. Further, at an $SO(2)$ -reducible connection \mathbf{A} , the normal ASD operator

splits as a sum of a ‘reducible part’ and an ‘irreducible part’, corresponding to the splitting of $\mathfrak{g}_E \cong E \cong \mathbb{R} \oplus \lambda$; write

$$D_{\mathbf{A},\pi}^\nu = D_{\mathbf{A},\pi}^{\text{red}} \oplus D_{\mathbf{A},\pi}^{\text{irred}};$$

one has $I(D^{\text{red}}) = b^1 - b^+$. This is the component that $S(r) = 1$ contributes to, as the operator Q^{irred} is the same as the operator Q_θ written in the above proof, and hence has index $-1 + b^1 - b^+$.

Correspondingly, we see that at an $SO(2)$ -reducible \mathbf{A} we have

$$I(D_{\mathbf{A},\pi}^{\text{irred}}) = -2(z_1 - z_2)^2 + 2(b^1 - b^+) + \frac{\rho_{\pi_2}(r_2) - \rho_{\pi_1}(r_1)}{2} + 1 - \frac{S(r_1) + S(r_2)}{2}.$$

Definition 4.12. Let (W, \mathbf{E}) be a cobordism $(Y_1, E_1) \rightarrow (Y_2, E_2)$ equipped with some small perturbation π , regular at the ends. We say that an $SO(2)$ -reducible r on (W, \mathbf{E}) is good if $I(D_{\mathbf{A},\pi}^{\text{irred}}) \geq 0$, and bad otherwise.

Remark 4.1. In fact, $D_{\mathbf{A},\pi}^{\text{irred}}$ is a complex linear operator, so its index is even.

There is a natural condition on a cobordism-with-perturbation (W, \mathbf{E}, π) that reduces the class of bad reducibles to a simple, sometimes avoidable, set.

Definition 4.13. Let $(W, \mathbf{E}, \pi) : (Y_1, E_1, \pi_1) \rightarrow (Y_2, E_2, \pi_2)$ be a cobordism. For an $SO(2)$ -reducible component r on (W, \mathbf{E}) , we write its restriction to the two ends as r_i . We say that (W, \mathbf{E}, π) is ρ -monotonic if for every $SO(2)$ -reducible component r , we have

$$\rho_{\pi_1}(r_1) \leq \rho_{\pi_2}(r_2).$$

For a ρ -monotonic cobordism, one of the most mysterious terms in the index formula is nonnegative, so we may focus on the rest.

Lemma 4.28. For a ρ -monotonic cobordism (W, \mathbf{E}, π) with $b_1(W) = b^+(W) = 0$ and rational homology sphere ends, the only bad reducibles are $\{z_1, z_2\}$ where $z_1 - z_2$ is a torsion class on $H^2(W; \mathbb{Z})$ that restricts trivially to the ends. In particular, bad reducibles can only exist if $H_1 Y_1 \oplus H_1 Y_2 \rightarrow H_1 W$ fails to be surjective.

Proof. The assumption of ρ -monotonicity means

$$\frac{\rho_{\pi_2}(r_2) - \rho_{\pi_1}(r_1)}{2} \geq 0,$$

so

$$I(D_{\mathbf{A},\pi}^{\text{irred}}) \geq -2p_1 \mathbf{A} + 1 - \frac{S(r_1) + S(r_2)}{2}.$$

If both r_i are $SO(2)$ -reducible but not fully reducible,

$$1 - \frac{S(r_1) + S(r_2)}{2} = 0;$$

if precisely one of the r_i is fully reducible then that same term is -1 , and if both of the r_i are fully reducible then the final term is -2 . Because $b^+ = 0$, we have $-2p_1 \mathbf{A} = -2(z_1 - z_2)^2 \geq 0$. Therefore for a ρ -monotonic cobordism, $I(D_{\mathbf{A},\pi}^{\text{irred}}) \geq -2$, with equality if and only if both r_i are fully reducible and $-2p_1 \mathbf{A} = 0$.

Because $H_c^2(W; \mathbb{Q}) \cong H^2(W; \mathbb{Q})$, and the intersection form is nondegenerate negative definite on H_c^2 , we see that

$$-2p_1 \mathbf{A} = -2(z_1 - z_2)^2 = 0$$

iff $z_1 - z_2 = 0 \in H^2(W; \mathbb{Q})$; this is the same as saying that $z_1 - z_2$ is a torsion class. The assumption that the restriction to the ends is fully reducible is precisely the same as saying that $z_1 - z_2$ restricts trivially to the ends. Applying the universal coefficient theorem, we obtain a class in $\text{Ext}(H_1W, \mathbb{Z})$ which restricts trivially to $\text{Ext}(H_1Y_1, \mathbb{Z}) \oplus \text{Ext}(H_1Y_2, \mathbb{Z})$, and so

$$\text{Ext}(H_1W) \rightarrow \text{Ext}(H_1Y_1) \oplus \text{Ext}(H_1Y_2)$$

is not injective. The natural isomorphism

$$\text{Ext}(A, \mathbb{Z}) \cong \text{Hom}(A, S^1)$$

for finite abelian groups A implies that if the map on Ext is not injective, then the map $H_1Y_1 \oplus H_1Y_2 \rightarrow H_1W$ fails to be surjective.

Because $I(D_{\mathbf{A}, \pi}^{\text{irred}})$ is an even integer, if it is larger than -2 , it is nonnegative, so r is a good reducible. \blacksquare

Definition 4.14. *Let (W, \mathbf{E}) be a cobordism between 3-manifolds with $SO(3)$ -bundles and signature data (Y_1, E_1, σ_1) and (Y_2, E_2, σ_2) . We say that (W, \mathbf{E}) is weakly admissible if one of the following holds.*

- *The negative end (Y_1, E_1) is admissible, meaning that $w_2(E_1)$ only lifts to non-torsion classes in $H^2(Y_1; \mathbb{Z})$.*
- *$\beta w_2(\mathbf{E}) \neq 0 \in H^3(W; \mathbb{Z})$, where β is the integral Bockstein homomorphism.*
- *$b_1(W) = b^+(W) = 0$, for every $\beta \in \text{Red}(W, \mathbf{E})$ restricting to β_i on the ends,*

$$\rho(\beta_2) - \rho(\beta_1) + \sigma_{\pi_2}(\beta_2) - \sigma_{\pi_1}(\beta_1) \geq 0,$$

and $H_1(Y_1) \oplus H_1(Y_2) \rightarrow H_1(W)$ is surjective. That is, (W, \mathbf{E}) is ρ -monotonic and supports no bad reducibles.

- *\mathbf{E} is non-trivial, $b_1(W) = 0$ and $b^+(W) > 0$.*

Later we will need a similar notion for $U(2)$ -bundles. If $\tilde{\mathbf{E}}$ is a $U(2)$ -bundle on the cobordism, we say that it is weakly admissible if its reduction to an $SO(3)$ -bundle \mathbf{E} is weakly admissible.

The weakly admissible $U(2)$ -bundles correspond to the first, third, and fourth cases above; the second case precisely means that \mathbf{E} admits no lift to a $U(2)$ -bundle, which implies that it admits no reducible connections.

We will soon see that every item on this list admits a regular perturbation. It should be noted that in fact we can achieve regular perturbations when (Y_2, E_2) is admissible but the negative boundary component is not. We do not include these in our definition of weakly admissible bundles as they do not glue together well in general. The composite of a cobordism from a rational homology sphere to an admissible bundle, and then back to a rational homology sphere, need not be weakly admissible.

For the definition above, we have the following.

Lemma 4.29. *The composite of two weakly admissible cobordisms remains weakly admissible.*

Proof. First, it is clear that the composite of the first type of weakly admissible cobordism with any other remains weakly admissible. That the composite of any cobordism and one with $\beta w_2 \mathbf{E} \neq 0$ still has $\beta w_2 \mathbf{E} \neq 0$ follows immediately from the naturality of the Bockstein and the fact that w_2 is natural under restriction. Suppose we have weakly admissible cobordisms of the third or fourth type that are

not of the first or second. Then their boundaries are rational homology spheres, and the composite also has $b_1(W) = 0$; the term $b^+(W)$ is additive under gluing along rational homology spheres, so the composite of any of the third or fourth type with the fourth type is again of the fourth type.

So what remains to check is that the composite of any cobordism of the third type with another of that type remains of that same type. It is clear that the ρ -monotonicity condition is additive; the interesting thing is to ask about is the homological condition. At first glance, it is mysterious why composites of cobordisms satisfying this condition should still satisfy this condition; it is made more clear by remembering the point of that condition.

Let \mathbf{A} be a reducible on the cobordism W with $b_1(W) = 0$, corresponding to cohomology classes $\{z_1, z_2\}$; write r_i for the restriction of \mathbf{A} to the corresponding boundary component. The ρ -monotonicity condition and homological condition, combined, are equivalent to the following: the index

$$I(D_{\mathbf{A},\pi}^{\text{irred}}) = -2(z_1 - z_2)^2 - 2b^+(W) + \frac{\rho_{\pi_2}(r_2) - \rho_{\pi_1}(r_1)}{2} + 1 - \frac{S(r_1) + S(r_2)}{2}$$

is nonnegative for all (\mathbf{A}, π) . (This is the index computed above Definition 4.12.)

Now let W_1 and W_2 be weakly admissible cobordisms between rational homology spheres equipped with reducibles \mathbf{A}_i ; write \mathbf{A}_{12} for the reducible on the composite cobordism W , and r_1, r_2, r_3 for the restrictions of s to the successive 3-manifolds that serve as boundary components of W_1 and W_2 . It is easy to see that

$$I(D_{\mathbf{A}_{12},\pi}^{\text{irred}}) = I(D_{\mathbf{A}_1,\pi}^{\text{irred}}) + I(D_{\mathbf{A}_2,\pi}^{\text{irred}}) + S(r_2) - 1.$$

Now the index $I(D_{\mathbf{A}_i,\pi}^{\text{irred}})$ is non-negative for all reducibles; so for the composite, we find $I(D_{\mathbf{A}_{12},\pi}^{\text{irred}}) \geq S(r_2) - 1$. But r_2 is a reducible, and so $S(r_2)$ is either 1 or 3; thus we obtain $I(D_{\mathbf{A}_{12},\pi}^{\text{irred}}) \geq 0$, as desired. \blacksquare

Consider the situation that W_1 is a cobordism from a homology sphere to a rational homology sphere, and W_2 is a cobordism from a rational homology sphere back to a homology sphere, so that $H_1 W_i$ are both isomorphic to the middle boundary component. Then clearly the composite cobordism fails the homological condition, even though it tautologically satisfies ρ -monotonicity (all reducibles are full reducibles, where the ρ -invariant and signature data are both zero). So even though they satisfy the homological condition, one of W_1 and W_2 must fail to be ρ -monotonic.

For every weakly admissible cobordism, all reducible configurations that exist for generic perturbation π have non-negative normal index. (In the admissible case, we may choose the perturbation π so that there are no π -ASD connections on the cobordism.)

In the first three cases, this remains true when considering paths of perturbations. In the final case, reducibles only appear generically in families of perturbations of dimension $b^+(W)$, but do not necessarily have non-negative index; when they arise, we may not achieve transversality. Strengthening the ρ -monotonicity requirement does not solve this problem.

We conclude this section with some remarks on gradings.

Lemma 4.30. *Let $(W, \mathbf{E}) : (Y_1, E_1, \pi_1) \rightarrow (Y_2, E_2, \pi_2)$ be a cobordism between 3-manifolds equipped with weakly admissible bundles and regular perturbations. If α_i*

and β_i are choices of critical orbits, then

$$\mathrm{gr}^W(\alpha_1, \alpha_2) - \mathrm{gr}^W(\beta_1, \beta_2) = \mathrm{gr}(\alpha_1, \beta_1) - \mathrm{gr}(\alpha_2, \beta_2).$$

This follows immediately from the additivity property of the grading and the fact that $\mathrm{gr}(\beta, \beta) = 0$.

Corollary 4.31. *Let $(W, \mathbf{E}) : (Y_1, \pi_1) \rightarrow (Y_2, \pi_2)$ be a cobordism between rational homology spheres equipped with the trivial bundle and small regular perturbations. Write θ_i for the corresponding trivial connections. We have*

$$w_2(\mathbf{E}) \in H^2(W, \partial W; \mathbb{Z}/2),$$

and so we may use the Pontryagin square to write

$$w_2(\mathbf{E})^2 \in H^4(W, \partial W; \mathbb{Z}/4) = \mathbb{Z}/4,$$

and $2w_2(\mathbf{E})^2 \in 2\mathbb{Z}/8$.

Then

$$\mathrm{gr}^W(\alpha_1, \alpha_2) = -2w_2(\mathbf{E})^2 + 3(b^1 - b^+) + \mathrm{gr}(\theta_2, \alpha_2) - \mathrm{gr}(\theta_1, \alpha_1).$$

Suppose $(W, \mathbf{E}, \mathbf{A})$ is a cobordism from (Y, E, π, α) to itself, where Y is a rational homology sphere equipped with a regular perturbation π and π -flat connection α . Then

$$\mathrm{gr}^W(\alpha, \alpha) = -2w_2(\mathbf{E})^2 + 3(b^1(W) - b^+(W)).$$

Proof. Recall Definition 4.5 that $\overline{\mathrm{gr}}(\alpha, \beta)$ is given by $\mathrm{ind}(Q_{\mathbf{A}, \pi}^\nu)$ for any connection \mathbf{A} connecting α and β with perturbation π limiting to the fixed perturbations on the ends, and

$$\mathrm{gr}(\alpha, \beta) = \overline{\mathrm{gr}}(\alpha, \beta) + 3 - \dim \alpha.$$

Proposition 4.26 then provides us with

$$\overline{\mathrm{gr}}^W(\theta_1, \theta_2) = -2w_2(\mathbf{E})^2 - 3(1 - b^1 + b^+),$$

using the fact that $p_1 \equiv w_2^2 \pmod{4}$ and $S(\theta) = 3$, as well as the vanishing of the invariants $\rho(\theta) = \sigma_\pi(\theta) = 0$. Therefore

$$\mathrm{gr}^W(\theta_1, \theta_2) = -2w_2(\mathbf{E})^2 + 3(b^1 - b^+).$$

Then the conclusion is simply a special case of the previous lemma.

The final claim follows from the same index calculation we used to calculate $\mathrm{gr}^W(\theta_1, \theta_2)$; the key points are that the $\sigma_\pi(\alpha)$ terms cancel, and that $S(\alpha)$ is $3 - \dim \alpha$. \blacksquare

We may say something about the gradings of reducibles. We begin with the fully reducible connections.

Proposition 4.32. *Let Y be a rational homology 3-sphere equipped with a trivial bundle and small regular perturbation π . If Θ and Θ' are fully reducible connections, then $\mathrm{gr}(\Theta, \Theta') \in 4\mathbb{Z}/8$.*

Proof. It suffices to prove this when $\pi = 0$; the spectral flow description above, as well as the assumption that π is sufficiently small that no eigenvalues cross the weight δ , implies the grading is the same for arbitrary perturbation.

Now recall that there is an action of $\mathcal{G}_E/\mathcal{G}_E^e = H^1(Y; \mathbb{Z}/2)$ on both the 3-dimensional configuration space $\tilde{\mathcal{B}}_{E,k}^e$. The unperturbed Chern-Simons functional is invariant under the full gauge group, and so the critical set is preserved by this

action. Therefore, we also have an action on the 4-dimensional configuration space $\sqcup_{\alpha, \beta} \tilde{\mathcal{B}}_{\pi^*E, k, \delta}^e(\alpha, \beta)$; we take the disjoint union here because the action of x takes $\tilde{\mathcal{B}}_{\pi^*E, k, \delta}^e(\alpha, \beta)$ to $\tilde{\mathcal{B}}_{\pi^*E, k, \delta}^e(x \cdot \alpha, x \cdot \beta)$.

The identification in Proposition 1.4 implies that $H^1(Y; \mathbb{Z}/2)$ acts freely and transitively on the set of fully reducible critical points (or what is the same, the set of fully reducible points). So there is a unique $x \in H^1(Y; \mathbb{Z}/2)$ with $x \cdot \Theta = \Theta'$. Now pick a connection $\mathbf{A} \in \tilde{\mathcal{B}}_{\pi^*E, k, \delta}^e(\Theta, \Theta')$. Because the action of $H^1(Y; \mathbb{Z}/2)$ preserves the ASD equations, the index of $Q_{\mathbf{A}}$ agrees with the index of $Q_{x \cdot \mathbf{A}}$.

Therefore, $\text{gr}(\Theta, \Theta') = \text{gr}(\Theta', \Theta)$; here we use that $2x = 0$, so

$$x \cdot \Theta' = (x + x) \cdot \Theta = \Theta.$$

But we know

$$2\text{gr}(\Theta, \Theta') = \text{gr}(\Theta, \Theta') + \text{gr}(\Theta', \Theta) = \text{gr}(\Theta, \Theta) = 0.$$

Because these take values in $\mathbb{Z}/8$, we see that $\text{gr}(\Theta, \Theta')$ is a multiple of 4. \blacksquare

We may make a similar observation about general reducibles (both $SO(2)$ - and fully reducible connections), with a completely different proof.

Proposition 4.33. *Let Y be a rational homology 3-sphere equipped with a weakly admissible bundle E and a small regular perturbation π . If α and β are reducible critical orbits, then $\text{gr}(\alpha, \beta)$ is even.*

Proof. First, we show that there exists a cobordism $(W, \mathbf{E}) : (Y, E) \rightarrow (Y, E)$ so that there is a reducible connection on (W, \mathbf{E}) restricting to α and β on the corresponding ends. To see this, recall from Proposition 1.6 that components of reducibles are classified by pairs of cohomology classes $\{z_1, z_2\} \subset H^2(W; \mathbb{Z})$ so that $z_1 + z_2$ is a fixed integral lift of $w_2 \mathbf{E}$ (and, in particular, if there are reducible components such an integral lift exists). We see, therefore, that it suffices to show that every oriented closed 3-manifold equipped with a pair of cohomology classes $(z_1, z_2) \subset H^2(Y; \mathbb{Z})$ is null-bordant through an oriented 4-manifold equipped with a pair of cohomology classes (z_1^W, z_2^W) that restrict to the z_i on the boundary. Because pairs of cohomology classes are classified by maps to $\mathbb{C}\mathbb{P}^\infty \times \mathbb{C}\mathbb{P}^\infty$, we are asking that $\Omega_3^{\text{or}}(\mathbb{C}\mathbb{P}^\infty \times \mathbb{C}\mathbb{P}^\infty) = 0$. This follows immediately from the existence of the Atiyah-Hirzebruch spectral sequence, the fact that $\mathbb{C}\mathbb{P}^\infty \times \mathbb{C}\mathbb{P}^\infty$ has cohomology only in even degrees, and $\Omega_i^{\text{or}} = 0$ for $1 \leq i \leq 3$, as this implies the terms $E_{k, 3-k}^2 = 0$ for all k , and hence the same is true of E^∞ .

Now if η_i^W are the associated complex line bundles to (z_1^W, z_2^W) , the $SO(3)$ -bundle is

$$\mathbf{E} \cong \mathbb{R} \oplus \eta_1^W \otimes (\eta_2^W)^{-1}$$

and the reducible component is labelled by $\{z_1^W, z_2^W\}$; choosing such a bounding manifold for both (Y, α) and (\bar{Y}, β) , we may simply take the connected sum to obtain the desired cobordism $(Y, E, \alpha) \rightarrow (Y, E, \beta)$.

Now recall from Lemma 4.30 that

$$\text{gr}^W(\beta, \alpha) - \text{gr}^W(\beta, \beta) = \text{gr}(\alpha, \beta).$$

To show that $\text{gr}(\alpha, \beta)$ is even, our goal is to show that $\text{gr}^W(\beta, \alpha) \equiv \text{gr}^W(\beta, \beta) \pmod{2}$. Pick a reducible connection \mathbf{A} , asymptotic to β at $-\infty$ and α at $+\infty$; the previous discussion amounts to saying we may do so.

First, we remark that the \mathbf{A} enjoys a splitting $\theta \oplus A$, for a connection A on a complex line bundle and θ the trivial connection on the trivial real line bundle. We may thus write $I(Q_{\mathbf{A}}) = I(Q_{\theta}) + I(Q_A)$. Because the index of a complex linear operator is even, we see that $I(Q_{\mathbf{A}}) \equiv I(Q_{\theta}) \pmod{2}$ and

$$I(Q_{\theta}) = b^1(W) - 1 - b^+(W).$$

Therefore, because β is reducible and hence $3 - \dim \beta$ is odd, we find that

$$\mathrm{gr}^W(\beta, \alpha) \equiv b^1(W) - b^+(W) \pmod{2}.$$

As for $\mathrm{gr}^W(\beta, \beta)$, the final part of Corollary 4.31 says that

$$\mathrm{gr}^W(\beta, \beta) = -2w_2(\mathbf{E})^2 + 3(b^1(W) - b^+(W)).$$

Reducing modulo 2, we find that $\mathrm{gr}^W(\alpha, \alpha) \equiv b^1(W) - b^+(W) \pmod{2}$, and so

$$\mathrm{gr}^W(\alpha, \alpha) \equiv \mathrm{gr}^W(\beta, \alpha) \pmod{2},$$

as desired. ■

4.6. Transversality for the cylinder and cobordisms.

Theorem 4.34. *Let E be a weakly admissible bundle over a 3-manifold Y . Suppose a perturbation $\pi_0 \in \mathcal{P}_{E,\delta}$ has been chosen. Then by Definition 4.3 of this space of perturbations, the enumeration of reducibles in Proposition 3.4 holds, and the critical set \mathfrak{C}_{π_0} is a finite set of nondegenerate $SO(3)$ -orbits. If \mathcal{O} is a small $SO(3)$ -invariant neighborhood of these orbits, and $\mathcal{P}_{E,\mathcal{O},\delta}$ the space of perturbations $\pi \in \mathcal{P}_{E,\delta}$ with $f_{\pi}|_{\mathcal{O}} = f_{\pi_0}|_{\mathcal{O}}$, then for a residual set of $\pi \in \mathcal{P}_{E,\mathcal{O},\delta}$, the \mathbb{R} -reduced moduli spaces of framed instantons $\widetilde{\mathcal{M}}_{E,z,\pi}^0(\alpha, \beta)$ between any two critical orbits are cut out nondegenerately. Hence these are smooth $SO(3)$ -manifolds of dimension $\mathrm{gr}_z(\alpha, \beta) + \dim \alpha - 1$ unless $\alpha = \beta$ and z is trivial, in which case $\widetilde{\mathcal{M}}_{E,0,\pi}^0(\alpha, \alpha) = \alpha$.*

Proof. First, we remark on the reducibles: by Proposition 4.14, the only reducible solutions are constant. We want to verify that the solutions are nondegenerate at constant trajectories; this amounts to saying that for the operator

$$\frac{d}{dt} + \widehat{\mathrm{Hess}}_{\mathbf{A},\pi} : \Omega^0 \oplus \Omega^1 \rightarrow \Omega^0 \oplus \Omega^1$$

has no cokernel other than the constant trajectories at $\ker(\Delta_{\alpha}) \subset \Omega^0(\mathfrak{g}_E)$. This is true by applying a standard separation of variables argument to the adjoint operator.

Now the only points to worry about are irreducibles, for which this theorem is standard: see [Don02, Section 5.5.1] or [KM11b]. (This is precisely where we use the perturbations in $\mathcal{P}_{E,\mathcal{O}}$ which agree with our original perturbations in a neighborhood of the π -flat connections.) ■

Before we continue the proof of transversality for cobordisms, we will need the following lemma.

The key assumption we started with was that I is nonnegative, as the following lemma makes clear. The non-equivariant index 0 case is written in [Sal99, Appendix A.3].

Lemma 4.35. *Fix a compact Lie group G and two separable G -Hilbert spaces X and Y . The Banach manifold $\mathcal{F}_G(X, Y)$ of G -equivariant Fredholm maps $X \rightarrow Y$ decomposes into disjoint open sets $\mathcal{F}_G^I(X, Y)$, where $I \in RO(G)$ is an element of*

the Grothendieck group on finite-dimensional G -representations, labelling the index of an operator.

Each $\mathcal{F}_G^I(X, Y)$ is stratified by locally closed subsets $\mathcal{F}_G^{V, W}(X, Y)$, where V and W are finite-dimensional G -representations with $[V] - [W] = I$. These are defined to consist of those operators $T \in \mathcal{F}_G(X, Y)$ with $\ker(T) \cong V$ and $\operatorname{coker}(T) \cong W$.

Each $\mathcal{F}_G^{V, W}(X, Y)$ is a smooth submanifold of $\mathcal{F}_G(X, Y)$; the normal space at an operator $T \in \mathcal{F}_G^{V, W}(X, Y)$ is isomorphic to $\operatorname{Hom}_G(\ker T, \operatorname{coker} T)$.

Proof. Let $T \in \mathcal{F}_G^{V, W}(X, Y)$, and split $X = X_0 \oplus X_1$ and $Y = Y_0 \oplus Y_1$, where $X_0 = \ker(T)$, $X_1 = X_0^\perp$, and $Y_1 = \operatorname{Im}(T)$, while $Y_0 = Y_1^\perp$; then by assumption T_{11} is an isomorphism. A neighborhood of T in $\mathcal{F}_G^{V, W}(X, Y)$ consists of equivariant maps $T' = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$, written in block-matrix form, where

$$A : X_0 \rightarrow X_0, \quad B : X_1 \rightarrow X_0, \quad C : X_0 \rightarrow X_1, \quad D : X_1 \rightarrow X_1,$$

where D is an isomorphism. If $T' \in \mathcal{F}_G^{V, W}(X, Y)$ and is sufficiently close to T , then the projection map $\ker(T') \rightarrow \ker(T) = X_0$ is an isomorphism (and otherwise it cannot possibly be). So suppose this projection map is an isomorphism. Write the equivariant inverse $v : X_0 \rightarrow \ker(T')$ and write it in components as $(x, v_1(x))$. The defining property of v is that

$$(Ax + Bv_1(x), Cx + Dv_1(x)) = 0.$$

Because D is an isomorphism, we may write $v_1(x) = -D^{-1}Cx$, and then we see that $Ax = BD^{-1}Cx$. Conversely, it is easy to see that if $A = BD^{-1}C$, then the projection map $\ker(T') \rightarrow X_0$ is an isomorphism. So $A = BD^{-1}C$ is a defining equation for $\mathcal{F}_G^{V, W}(X, Y)$ near zero. The map $\mathcal{F}_G^{V, W}(X, Y) \rightarrow \operatorname{Hom}_G(X_0, Y_0)$ given by

$$T' \mapsto A - BD^{-1}C$$

has derivative $T' \mapsto A$ at zero, which is a surjective linear map, and hence we see that $\mathcal{F}_G^{V, W}(X, Y)$ is a smooth manifold near T . \blacksquare

For the most difficult part of the argument — transversality normal to the reducible locus — we will need the following lemma.

Lemma 4.36. *Let π be a regular perturbation on (Y, E) , and consider the π -perturbed ASD operator on $[0, \infty) \times Y$ with Coulomb condition. Writing a 1-form on $[0, \infty) \times Y$ as $\psi(t) + dt \wedge \phi(t)$ for time-varying 0- and 1-forms ϕ and ψ , and a self-dual 2-form as $(dt \wedge * \omega(t))^+$ for a time-varying 1-form, this operator is*

$$Q_{\mathbf{A}, \pi}(\phi_t, \psi_t) = \left(\phi'_t - d_{\mathbf{A}(t)}^* \psi_t, \psi'_t - d_{\mathbf{A}(t)} \phi_t + \beta_0(t) \cdot (D_{\mathbf{A}(t)}) (\operatorname{cs} + f_\pi) \psi_t \right);$$

its adjoint is the same, with the signs on the time-derivatives negated. Henceforth, we denote

$$L(t) = \beta_0(t) D_{\mathbf{A}(t)} (\operatorname{cs} + f_\pi).$$

Then the following are true.

- (1) *If ψ is a time-dependent 1-form with $Q_{\mathbf{A}, \pi}^* \psi = 0$, and furthermore $\psi(0) = 0$, then in fact $\psi = 0$ on the whole of $[0, \infty) \times Y$.*

- (2) If $Q_{\mathbf{A},\pi}(\phi, \psi) = 0$, and furthermore $\psi(0) = -d_{\mathbf{A}(0)}\sigma$ for some 0-form σ , then one may extend σ to a time-dependent 0-form σ_t with

$$(\phi_t, \psi_t) = (\sigma'_t, -d_{\mathbf{A}(t)}\sigma_t);$$

that is, if an element of the kernel is exact on the boundary, it is globally exact.

Proof. The first statement follows immediately from Lemma 4.19; the second is more subtle.

First, set $\sigma_t = \Delta_{\mathbf{A}(t)}^{-1}d_{\mathbf{A}(t)}^*\psi_t$. Adding $d_{\mathbf{A}}\sigma$ to the given element of the kernel, we obtain a new one with $d_{\mathbf{A}(t)}^*\psi_t = 0$ for all t . We aim to show that this (ϕ, ψ) is itself the differential of some 0-form on $[0, \infty) \times Y$. This destroys the Coulomb condition, but preserves the linearized ASD equation.

First, we write the ASD equation explicitly: it now takes the form

$$\psi'_t + L(t)\psi_t = d_{A(t)}\phi_t.$$

Fix a C^1 isomorphism

$$e_t : \ker(d_{\mathbf{A}(0)}^*) \rightarrow \ker(d_{\mathbf{A}(t)}^*).$$

We use this to write $\psi_t = e_t(f_t)$, where f_t is a time-dependent element of $\ker(d_{\mathbf{A}(0)}^*)$. Then we may write $\psi'_t = e_t(f'_t) + e'_t(f_t)$, where $e'_t : \ker(d_{\mathbf{A}(0)}^*) \rightarrow \Omega^1$ is a continuous linear map, varying continuously in time.

Projecting to $\ker(d_{\mathbf{A}(t)}^*)$ (and calling the projection operator Π_t), we obtain the following equation on $\ker(d_{\mathbf{A}(0)}^*)$:

$$f'_t + (e_t^{-1}\Pi_t L(t)e_t)(f_t) = -(e_t)^{-1}(\Pi_t e'_t)(f_t).$$

The operator $e_t^{-1}(\Pi_t L(t)e_t)$ differs from a symmetric and densely defined operator by a bounded operator, and the operator $e_t^{-1}(\Pi_t e'_t)$ is bounded, so this satisfies the conditions of Lemma 4.19; therefore, $f_t = 0$ for all time, and so $\psi = 0$. The ASD equation becomes

$$d_{\mathbf{A}(t)}\phi_t = 0.$$

We need to show that $(\phi_t, 0)$ is exact. But if $\sigma_t = \int_0^t \phi_s ds$, we have

$$d_{\mathbf{A}}\sigma_t = (\phi_t, \int_0^t d_{\mathbf{A}(s)}\sigma_s ds) = (\phi_t, 0).$$

So any element of the kernel which is exact on the boundary is exact on the whole of $[0, \infty) \times Y$. ■

Theorem 4.37. *Let $(W, \mathbf{E}, \pi) : (Y_1, E_1, \pi_1) \rightarrow (Y_2, E_2, \pi_2)$ be a weakly admissible cobordism, where $\pi \in \mathcal{P}_\delta^{(4)}$; this means, in particular, that the π_i are regular perturbations, and the enumeration of reducibles Proposition 3.4 applies. Write, as usual, $\mathcal{P}_c^{(4)}$ for the space of perturbations which agree with the π_i on the ends.*

For any fixed constant $C > 0$, there is an open neighborhood $(\pi_-, \pi_+) \in U_C \subset \mathcal{P}_c^{(4)}$ with the following significance.

There is a residual set of $\pi \in U_C$ so that for any π in this residual set, every moduli space $\widetilde{\mathcal{M}}_{\mathbf{E}, z, \pi'}(\alpha, \beta)$ with energy at most C is cut out transversely, and hence is a smooth $SO(3)$ -manifold of dimension $\text{gr}_z(\alpha, \beta) + \dim \alpha$ if it contains

irreducible connections. In particular, regular perturbations¹⁰ exist arbitrarily close to $(\pi_1, 0, \pi_2)$.

Proof. As before, we argue inductively on reducibility type. We have already seen in Lemma 4.20 that the set of perturbations for which the fully reducible connections are cut out transversely in a neighborhood $(\pi_-, \pi_+) \subset U$ of any end perturbation (π_-, π_+) on a weakly admissible cobordism. Furthermore, we saw in the proof of Proposition 4.15 and Proposition 4.16 that for a residual set of $\pi \in U$, we achieve transversality internally to the reducible locus. The same argument applies just as well to irreducible connections. The only remaining difficulty is achieving transversality *normal* to the reducible locus.

Recall that above the neighborhood U , we have a Hilbert manifold of parameterized reducible instantons $UM_{k,\delta}^{\text{red}}$, equipped with a projection to U ; this projection is proper so long as you restrict to π -instantons with energy bounded by a fixed constant C . A fixed component of this space corresponds to instantons which respect a splitting of topological type $\mathbf{E} \cong \mathbb{R} \oplus \eta$.

Over this Hilbert manifold there is a Banach bundle, whose fiber above (\mathbf{A}, π) is the space of $SO(2)$ -equivariant Fredholm maps

$$(\ker(d_{\mathbf{A}}^*)_{k,\delta}) \rightarrow \Omega_{k-1,\delta}^{2,+}(W; \eta);$$

a more invariant way of specifying the domain is $N_{\mathbf{A}}\tilde{\mathcal{B}}_{\mathbf{A},k,\delta}^{\text{red}}$, the normal bundle to the reducible locus.

Write $U\mathcal{F}$ for this Banach bundle with projection

$$U\mathcal{F} \rightarrow UM_{k,\delta}^{\text{red}}.$$

There is a stratified sub-bundle of this bundle, which is a locally closed union of manifolds: the subspace $U\mathcal{F}^{\geq k}$ of those Fredholm operators with cokernel isomorphic to \mathbb{C}^j , for some $j \geq k$, with $SO(2)$ acting with weight one. By Lemma 4.35, the normal space to $U\mathcal{F}^j$ at $(\mathbf{A}, \pi, D_{\mathbf{A},\pi}^{\text{irred}})$ is

$$\text{Hom}_{SO(2)}(\ker(D_{\mathbf{A},\pi}^{\text{irred}}), \text{coker}(D_{\mathbf{A},\pi}^{\text{irred}})).$$

Here D^{irred} is the ASD operator restricted to the bundle η ; it has no gauge fixing condition, or rather has the gauge fixing condition built into the domain of the map.

Our next goal is to show that the section $D_{\mathbf{A},\pi}^{\text{irred}} : UM_{k,\delta}^{\text{red}} \rightarrow U\mathcal{F}$ is transverse to $U\mathcal{F}^{\geq 1}$ along $\mathcal{M}_{(\pi_-, \pi_+), k, \delta}^{\text{red}}$ — that is, we are only showing that this map D^{irred} is transverse to $U\mathcal{F}^{\geq 1}$ along the fiber above 0. This is true for all fully reducible connections by the definition of U , so we may as well assume \mathbf{A} is $SO(2)$ -reducible.

Now suppose $(\mathbf{A}, (\pi_-, \pi_+))$ is a reducible perturbed instanton whose perturbation has no interior part. Recall that if W has cylindrical ends $[0, \infty) \times Y$, we define the interior holonomy perturbations on the interior of the submanifold

$$W' = W \setminus ((1, \infty) \times Y).$$

Our first claim is that for any point p in the interior of W' , there is a collection of loops $\gamma_1, \dots, \gamma_N$ based at p so that the equivariant map

¹⁰Here we use regular in the sense that moduli spaces up to some energy bound are cut out regularly; we will only need those moduli spaces with $\text{gr}_z \leq 10$, which corresponds to an energy bound.

$$d_{\mathbf{A}}\text{Hol}^{\vec{\gamma}} : \ker(D_{\mathbf{A},\pi}^{\text{irred}}) \rightarrow (\mathfrak{g}_{E_p})^N$$

is an embedding.

This follows because ω must be exact over the interior of W' by Lemma 3.2, and therefore by the second point of Lemma 4.36 ω is exact along $(0, \infty) \times Y$ as well. The issue is that if σ_1, σ_2 are the 0-forms with $d_{\mathbf{A}}\sigma_i = \omega$ over the interior and ends, respectively, it's not clear that $\sigma_1 = \sigma_2$ over $(0, 1) \times Y$; however, their difference is a closed 0-form on $(0, 1) \times Y$.

If \mathbf{A} is reducible over $(0, 1) \times Y$, then a unique continuation argument shows that it is reducible of the same class over all of $(0, \infty) \times Y$; in particular, the closed 0-form $\sigma_1 - \sigma_2$ extends to a closed 0-form on the whole of $(0, \infty) \times Y$. Adding this to σ_2 , we find that indeed there is a globally defined 0-form $\text{sin}\Omega_{k+1, \delta, \text{ext}}^0(W; \mathfrak{g}_{\mathbf{E}})$ so that $\omega = d_{\mathbf{A}}\sigma$, and thus ω is zero as a tangent vector in $\tilde{\mathcal{B}}_{\mathbf{E}, k, \delta}$.

In fact, because this derivative map is $SO(2)$ -equivariant, its image must lie in some complement to the $SO(2)$ -fixed subspace $(i\mathbb{R})^N$.

The second claim is that one may choose a finite set of points p_1, \dots, p_m in the interior of W' so that the evaluation map

$$\text{coker}(D_{\mathbf{A},\pi}^{\text{irred}}) \rightarrow \bigoplus_{i=1}^m \Lambda^{2,+}(\mathfrak{g}_{\mathbf{E}_{p_i}})$$

is injective; again as above we may restrict to $\Lambda^{2,+}(\mathfrak{g}/i\mathbb{R})$.

This is where we use, in an essential way, that π does not have a component which is an interior perturbation — to run the elliptic regularity argument that identifies $\psi \in L_{k, \delta}^2$ requires that the perturbations $D_{\mathbf{A}}\widehat{\nabla}_{\pi}$ extend to Sobolev spaces $L_{j, \text{loc}}^2$ for all $-k \leq j \leq k$, whereas we only have that extension through $0 \leq j$ when the perturbation includes an interior part. When there is no interior part, however, we have enough regularity that we may point-evaluate as above.

So suppose that we have enough regularity to point-evaluate, and that the above map was *not* injective, for any choice of points p_i in the interior of W' . This would imply that there was some $\psi \in \text{coker}(D_{\mathbf{A},\pi}^{\text{irred}})$ which was identically zero on the whole of W' . But then, by the first point of Lemma 4.36, ψ is zero on all of W . Thus some collection of points p_i suffices.

Now \mathbf{A} , being $SO(2)$ -reducible but not fully reducible, we may assume that the curves above (at each point p_i) are chosen so that

$$\text{Hol}^{\vec{\gamma}_j}(\mathbf{A}) \in U(1)^{N_j} \setminus (\pm 1)^{N_j} :$$

here we identify $\Gamma_{\mathbf{A}} \subset \widetilde{\text{Aut}}(\mathbf{E}_b)$ with $U(1)$. Call this element H_j for convenience.

This means that when we choose the equivariant map

$$r_j : \widetilde{\text{Aut}}(\mathbf{E}_{p_j})^N \rightarrow \Lambda^{2,+}(\mathfrak{g}_{\mathbf{E}_{p_j}})$$

so that the derivative $d_{H_j}r_j$ is an arbitrary $SO(2)$ -equivariant map

$$\mathfrak{g}_{\mathbf{E}_{p_j}}^{N_j} \rightarrow \Lambda^{2,+}(\mathfrak{g}_{\mathbf{E}_{p_j}}),$$

and in particular can restrict to an arbitrary chosen $SO(2)$ -equivariant map

$$\left(\mathfrak{g}_{\mathbf{E}_{p_j}}/i\mathbb{R}\right)^{N_j} \rightarrow \Lambda^{2,+}(\mathfrak{g}_{\mathbf{E}_{p_j}}/i\mathbb{R}).$$

Thus, finally, we may choose these maps r_j so that, upon restricting to the image of $\ker(D_{\mathbf{A},\pi}^{\text{irred}})$ and projecting to the image of $\text{coker}(D_{\mathbf{A},\pi}^{\text{irred}})$, the maps r_j give an arbitrary chosen $SO(2)$ -equivariant map

$$\ker(D_{\mathbf{A},\pi}^{\text{irred}}) \rightarrow \text{coker}(D_{\mathbf{A},\pi}^{\text{irred}}).$$

As usual, now, we may approximate these collection of loops and maps r_j by elements of the fixed countable dense set of interior holonomy perturbations π_i , choosing the self-dual 2-forms supported in a neighborhood of the points to very well approximate δ -masses at those points.

So we may approximate an arbitrary $SO(2)$ -equivariant map from the kernel to the cokernel as a map induced by $D_{\mathbf{A}}^{\text{irred}} \widehat{\nabla}_{\pi_i}$.

Thus, the map

$$D^{\text{irred}} : U\mathcal{M}_{k,\delta}^{\text{red}} \rightarrow U\mathcal{F}$$

is indeed transverse to $U\mathcal{F}^{\geq 1}$ above $(\pi_-, \pi_+) \in U$. Now recall from Lemma 4.18 that, so long as you restrict to those instantons of energy at most C , the map

$$U\mathcal{M}_{k,\delta}^{\text{red}, \leq C} \rightarrow U$$

is proper, and in particular closed, so the tube lemma applies: there is a neighborhood $(\pi_-, \pi_+) \in U_C \subset U$ so that the restriction of D^{irred} to $U_C\mathcal{M}^{\text{red}}$ is transverse to $U_C\mathcal{F}$ for those π -instantons with energy at most C . It is safe to assume that C is not the energy of any reducible π -instanton above U_C , so that this is still a smooth manifold; we drop the C from the notation from \mathcal{M}^{red} for convenience.

Write

$$(D^{\text{irred}})^{-1}(U_C\mathcal{F}^{\geq 1}) = U_C\mathcal{M}_{k,\delta}^{\text{red}, \geq 1};$$

this is a union of a countable collection of locally closed submanifolds of positive codimension. By Sard's theorem, the regular values of the projection

$$U_C\mathcal{M}_{k,\delta}^{\text{red}, \geq 1} \rightarrow U_C$$

form a residual set; for π to be regular simply means that $\text{coker}(D_{\mathbf{A},\pi}^{\text{irred}}) = 0$ for all reducible π -instantons \mathbf{A} .

This is what we set out to prove. \blacksquare

It is worth recalling in the above that we may choose a constant C so that for all $\pi \in U_C$, any moduli space of instantons with $\text{gr}_z(\alpha, \beta) \leq 10$ consists of instantons with energy at most C . In particular, there is no need to take C arbitrarily large, and no risk of the open sets U_C shrinking to zero as one attempts to do so.

4.7. Gluing. We follow the approach to gluing given in [KM07, Chapter 19], and in particular we need to briefly discuss weighted Sobolev spaces and the perturbed ASD equations on compact cylinders.

We write $Z^T = [-T, T] \times Y$ and $Z^\infty = [0, \infty) \times Y \sqcup (-\infty, 0] \times Y$; we view Z^∞ as a limit of the Z^T , stretching until what used to be 0 becomes the point at ∞ . We uniformly have $\partial Z^T = Y \sqcup \bar{Y}$. The function spaces of interest to us are the Sobolev spaces of sections of vector bundles over Z^T . The L_k^2 Sobolev space is the completion of the space of compactly supported smooth sections on Z^T ; notice that there are no boundary conditions on these sections. As a remark before continuing, the restriction map to the boundary takes value in the $L_{k-1/2}^2$ Sobolev space.

We define the *weighted* Sobolev spaces on Z^T as in [Lin18, Page 72]: let $\sigma_\delta : \mathbb{R} \rightarrow [-\delta, \delta]$ be an odd smooth function such that $\sigma_\delta(t) = -\delta$ for $t \geq 1$ and $\sigma_\delta(t) = \delta$ for $t \leq -1$. For each finite $T \geq 2$ we let $g_{T,\delta}$ be the positive smooth even function on

$[-T, T]$ which is equal to 1 on the boundary and has $\sigma_\delta = \log(g_{T,\delta})'$. For $t \in [1, T]$, we have $g_{T,\delta} = e^{-\delta(t-T)}$, an exponentially decreasing function with final value 1. Then for sections on the finite cylinder Z^T we set

$$\|f\|_{L_{k,\delta}^2} = \|g_{T,\delta} \cdot f\|_{L_k^2}.$$

For $T = \infty$ then as before we will use $g_{\infty,\delta} = e^{|t|\delta}$. For finite T , this is equivalent to the usual L_k^2 norm, with implicit constants growing exponentially in T .

Now as before we may introduce the moduli space of instantons on the finite cylinder Z^T in the same component as the constant solution γ_A in the usual way: we consider the configuration space $\gamma_A + \Omega_{k,\delta}^1(Z^T; \mathfrak{g}_E)$ and quotient by the space $\mathcal{G}_{E,k+1,\delta}^{e,h}$ of gauge transformations. In the case $T = \infty$, the h (for harmonic) denotes that these gauge transformations should be asymptotic on the noncompact end to elements of Γ_A , possibly different on each component. These moduli spaces with no boundary conditions are infinite-dimensional and depend on the Sobolev index k , as the restriction of an L_k^2 ASD connection to the boundary may only be of Sobolev class $L_{k-1/2}^2$, in any $L_{k+1/2}^2$ gauge equivalence class.

When defining *framed* moduli spaces on the finite cylinder Z^T , we set

$$\widetilde{\mathcal{M}}_{A,k} := \mathcal{A}_{A,k}(Z^T) \times E_{(-T,b)} / \mathcal{G}_{\pi^*E,k+1,\delta}^e.$$

Here A denotes that we are based at the constant trajectory γ_A . Because these moduli spaces depend on the Sobolev index k , we include it in our notation.

In the infinite case Z^∞ , there should be *two* framings, one for each component, which have the same value in the orbit through A in $\widetilde{\mathcal{B}}_E^e$ when we take the holonomy to $\pm\infty$. To mirror the case of finite cylinders, we write this as

$$\widetilde{\mathcal{A}}_{A,k,\delta}(Z^\infty) := \mathcal{A}_{A,k,\delta}(Z^\infty) \times E_{(0^+,b)} \times E_{(0^-,b)};$$

the notation $(0^+, b)$ signifies that this lies in the component $[0, \infty) \times Y$.

Write $\widetilde{\mathcal{M}}'_{A,k,\delta}(Z^\infty)$ for the quotient. Now we may take \mathbf{A} -holonomy to (∞, b) or $(-\infty, b)$, respectively, and project the framing factors to E_b/Γ_α ; and we write $\widetilde{\mathcal{M}}_{A,k,\delta}(Z^\infty)$ for the subset on which the two framings project to the same element of E_b/Γ_α .

For all T , there is a restriction map

$$R^T : \widetilde{\mathcal{M}}_{A,k,\delta}(Z^T) \rightarrow \widetilde{\mathcal{B}}_{E,k-1/2}^e(Y \sqcup \bar{Y}).$$

When $T = \infty$, an element of $\widetilde{\mathcal{M}}_{A,k,\delta}(Z^\infty)$ is an equivalence class of connection with a framing at each basepoint on the two boundary components. Therefore, R^∞ may be given by literally restricting the connection and framing to the boundary (up to gauge equivalence).

For finite T this is a little more complicated. On the left boundary component Y , this is the usual restriction map (the framed basepoint is on Y), but on \bar{Y} , it is given by restriction in the connection coordinate and parallel transport from $-T$ to $+T$ along $\mathbb{R} \times \{b\}$ in the framing coordinate.

Now we turn to slices for the gauge group action so that we may compute neighborhoods of $\widetilde{\mathcal{M}}_{A,k,\delta}$ as an equation on a linear space.

As before, we may consider the Coulomb slice $\ker(d_{\gamma_A}^*) \subset \Omega_{k,\delta}^1(Z^T; \mathfrak{g}_E)$ where γ_A denotes the constant trajectory at A ; in particular

$$d_{\gamma_A}^*(dt \wedge \sigma(t) + \omega(t)) = -\sigma'(t) + d_A^* \omega(t).$$

The Coulomb slice is no longer a slice for the gauge group action: if we try to solve the equation $d_{\gamma_A}^*(d_{\gamma_A}\sigma + \omega) = 0$, to write an arbitrary element of Ω^1 as a sum of an element of $\text{Im}(d_{\gamma_A})$ and an element of $\ker(d_{\gamma_A}^*)$, we find that

$$d_{\gamma_A}^* d_{\gamma_A} \sigma = -d_{\gamma_A}^* \omega$$

has a unique solution $d_{\gamma_A}\sigma$ for each ω and *fixed boundary values* $d_A\sigma|_{Y \sqcup \bar{Y}}$.

We obtain the *Coulomb-Neumann slice* around γ_A , written $\mathcal{CN}_{A,k,\delta}$: the subset of $\omega \in \Omega_{k,\delta}^1(Z^T; \mathfrak{g}_E)$ on which $d_{\gamma_A}^*\omega = 0$ and, writing $\omega = \phi + dt \wedge \psi$, where ψ is a time-dependent 0-form, we demand that

$$\psi|_{\partial Z^T} = 0.$$

Every connection on Z^T sufficiently close to γ_A is gauge equivalent to one on the Coulomb-Neumann slice, and the only remaining ambiguity is that $\mathcal{CN}_{A,k,\delta}$ carries the action of the stabilizer Γ_{γ_A} in the gauge group, and $\gamma_A + a$ is gauge equivalent to $\gamma_A + u(A)$. In particular, observe that $\mathcal{CN}_{A,k,\delta} \times_{\Gamma_{\gamma_A}} SO(3)$ gives an $SO(3)$ -invariant neighborhood of γ_A in $\tilde{\mathcal{B}}(Z^T)$.

Before moving on, we recall Lin's abstract Morse-Bott gluing theorem. Let E be a vector bundle over a closed oriented 3-manifold Y , and let L be a self-adjoint elliptic operator acting on $L_k^2(Y; E)$ with kernel H_0 . In what follows, we consider the operator $D := d/dt + L$ acting on finite cylinders $Z^T = [-T, T] \times Y$ and infinite cylinders $Z^\infty = [0, \infty) \times Y \sqcup (-\infty, 0] \times Y$.

On the infinite cylinder, we write

$$L_{k,\delta,\text{ext}}^2(Z^\infty; \pi^*E) := H_0 + L_{k,\delta}^2(Z^\infty; \pi^*E),$$

where H_0 indicates sections s_{h_0} in $\ker(D)$ on Z^∞ which are constant in time at some element h_0 in $H_0 = \ker(L)$. In fact, we write

$$\mathcal{E}_\delta^T = L_{k,\delta}^2(Z^T; \pi^*E), \quad \mathcal{E}_\delta^\infty = L_{k,\delta,\text{ext}}^2(Z^\infty, \pi^*E)$$

and

$$\mathcal{F}_\delta^T = L_{k-1,\delta}^2(Z^T; \pi^*E), \quad \mathcal{F}_\delta^\infty = L_{k-1,\delta}^2(Z^\infty, \pi^*E).$$

We have a projection map $\Pi_0^\infty : \mathcal{E}_\delta^\infty \rightarrow H_0$ given by taking the asymptotic value at $\pm\infty$ and projection maps $\Pi_0^T : \mathcal{E}_\delta^T \rightarrow H_0$ given by projection onto the subspace of constant sections in the $L_{k,\delta}^2$ norm (equivalently, in the $L_{0,\delta}^2$ norm).

Suppose we are given a bounded linear operator

$$\Pi : L_{k-1/2}^2(Y \sqcup \bar{Y}; E) \rightarrow H$$

for some Hilbert space H ; by restriction to the boundary this induces maps

$$\Pi : L_{k,\delta}^2(Z^T; \pi^*E) \rightarrow H$$

$$\Pi : L_{k,\delta,\text{ext}}^2(Z^\infty; \pi^*E) \rightarrow H.$$

We now assume, crucially, that the map

$$(D, \Pi_0^\infty, \Pi) : \mathcal{E}_\delta^\infty \rightarrow \mathcal{F}_\delta^\infty \oplus H_0 \oplus H$$

is an isomorphism, which implies the same for δ' sufficiently close to δ . Now suppose we are given a *non-linear* map

$$\alpha : C^\infty(Z^T; \pi^*E) \rightarrow L_{\text{loc}}^2(Z^T; \pi^*E),$$

obtained by restriction to slices from a map $\alpha_0 : C^\infty(Y; \pi^*E) \rightarrow L_{\text{loc}}^2(Y; \pi^*E)$. We assume that α extends to a smooth map

$$L_k^2([-1, 1] \times Y; \pi^*E) \rightarrow L_{k-1}^2([-1, 1] \times Y; \pi^*E)$$

with $\alpha(h_0)$ for every $h_0 \in H_0$ and we assume that α is purely nonlinear, in the sense that its derivative at $0 \in L_k^2([-1, 1] \times Y; \pi^*E)$ is zero. This implies that α defines smooth maps $\mathcal{E}_\delta^T \rightarrow \mathcal{F}_\delta^T$ with the same property.

Now write

$$F^T = D + \alpha : \mathcal{E}_\delta^T \rightarrow \mathcal{F}_\delta^T$$

and

$$M(T) = (F^T)^{-1}(0) \subset \mathcal{E}_\delta^T.$$

The following is [Lin18, Proposition 3.5.15].

Proposition 4.38. *For $T \in [T_0, \infty]$, the sets $M(T)$ are Hilbert submanifolds of $\mathcal{E}_\delta^\infty$ in a neighborhood of 0. There exist $\eta > 0$ and smooth maps*

$$u(T, -) : B_\eta(H_0 \oplus H) \rightarrow M(T)$$

which are diffeomorphisms onto their image and have

$$(\Pi_0, \Pi)u(T, (h_0, h)) = (h_0, h) \quad \text{and} \quad u(T, (h_0, 0)) = s_{h_0}.$$

For $T \in [T_0, \infty]$, we have a map

$$\mu_T : B_\eta(H_0 \oplus H) \rightarrow L_{k-1/2}^2(Y \sqcup \bar{Y}; E)$$

obtained as a composition of $u(T, -)$ with restriction to the boundary. For each T , μ_T is a smooth embedding; μ_T is smooth as a function on $[T_0, \infty) \times B_\eta(H_0 \oplus H)$, and $\mu_T \rightarrow \mu_\infty$ as $T \rightarrow \infty$ in the C_{loc}^∞ topology on maps $B_\eta(H_0 \oplus H) \rightarrow L_{k-1/2}^2(Y \sqcup \bar{Y}; E)$. Finally, there is an $\eta' > 0$ independent of T so that the images of the maps $u(T, -)$ contains all solutions $u \in M(T)$ with $\|u\|_{L_{k,\delta}^2} \leq \eta'$.

We will now state our main gluing lemma. In what follows, we will use Lin's abstract gluing theorem to prove our lemma in three cases, depending on the reducibility type of the orbit α , along with some mild changes (the flavor of the changes depending on how reducible α is).

Lemma 4.39. *Let (Y, E, π) be a closed oriented 3-manifold equipped with a weakly admissible $SO(3)$ -bundle and regular perturbation π . We write α for a π -critical orbit; if A is a connection in the gauge equivalence class of α , we write the corresponding Coulomb slice*

$$\mathcal{K}_\alpha = \ker(d_A^*) \cap \Omega_{k-1/2}^1(Y; \mathfrak{g}_E),$$

and $B(\mathcal{K}_\alpha)$ for its unit ball. The Hilbert space \mathcal{K}_α carries the action of the stabilizer of A in the gauge group; thinking of \mathcal{K}_α as the normal space to a point in the orbit α in $\tilde{\mathcal{B}}_{k-1/2}^e(Y)$, this is the same as the action of Γ_α , the stabilizer of a point of α in $SO(3)$.

We may thus extend \mathcal{K}_α to a vector bundle

$$\tilde{\mathcal{K}}_\alpha := SO(3) \times_H \mathcal{K}_\alpha$$

over the orbit α , equipped with an $SO(3)$ action (acting on the factor of $SO(3)$ on the left). We write the associated unit disc bundle as $B(\tilde{\mathcal{K}}_\alpha)$.

There is a T_0 so that for all $T \in [T_0, \infty]$, we may find smooth, $SO(3)$ -equivariant maps

$$\tilde{u}(T, -) : B(\tilde{\mathcal{K}}_\alpha) \rightarrow \tilde{\mathcal{M}}_{\alpha, k, \delta}(Z^T)$$

which are diffeomorphisms onto neighborhoods of the constant solution γ_α , and such that the map

$$\tilde{\mu}_T : B(\tilde{\mathcal{K}}_\alpha) \rightarrow \tilde{\mathcal{B}}_{E, k-1/2}^e(Y \sqcup \bar{Y}),$$

given by composing $u(T, -)$ with the restriction map

$$R^T : \tilde{\mathcal{M}}_{\alpha, k, \delta}(Z^T) \rightarrow \tilde{\mathcal{B}}_{E, k-1/2}^e(Y \sqcup \bar{Y})$$

described above using holonomy in the framing coordinate for finite T , has the following properties.

First, $\tilde{\mu}_T$, being the composition of equivariant maps, is $SO(3)$ -equivariant; $\tilde{\mu}_T$ is a smooth embedding of $B(\tilde{\mathcal{K}}_\alpha)$ for all $T \in [T_0, \infty]$, and $\tilde{\mu}_T$ is smooth as a function of $[T, \infty) \times B(\tilde{\mathcal{K}}_\alpha)$. Though not smooth as a function on $[T_0, \infty]$, we at least have $\tilde{\mu}_T \rightarrow \tilde{\mu}_\infty$ in the C_{loc}^∞ topology on $B(\tilde{\mathcal{K}}_\alpha)$ as $T \rightarrow \infty$.

Finally, there is an $\eta > 0$, independent of T , so that the images of $\tilde{u}(T, -)$ contain all solutions $[\gamma] \in \tilde{\mathcal{M}}_{\alpha, k, \delta}(Z^T)$ such that

$$\|\gamma - \gamma_\alpha\|_{L_{k, \delta}^2(Z^T)} \leq \eta.$$

The vector space \mathcal{K}_α has an eigenspace decomposition $\mathcal{K}_\alpha^- \oplus \mathcal{K}_\alpha^+ = \mathcal{K}_\alpha$; correspondingly we have an $SO(3)$ -equivariant fiber product decomposition

$$B(\tilde{\mathcal{K}}_\alpha) \cong B(\tilde{\mathcal{K}}_\alpha^-) \times_\alpha B(\tilde{\mathcal{K}}_\alpha^+).$$

We also have the decomposition

$$\tilde{\mathcal{M}}_{\alpha, k, \delta}(Z^\infty) = \tilde{\mathcal{M}}_{\alpha, k, \delta}(\mathbb{R}^{\geq 0} \times Y) \times_\alpha \tilde{\mathcal{M}}_{\alpha, k, \delta}(\mathbb{R}^{\leq 0} \times Y).$$

The map $\tilde{u}(\infty, -)$ respects this decomposition: it is a fiber bundle map over α , and on each fiber it respects the product structure.

The easiest case is when α is irreducible, as then we may solve the gluing problem in a Morse (not Morse-Bott) setting.

Proof of Lemma 4.39 in the irreducible case. We will be studying the ASD operator with a gauge fixing condition; the linear operator L on Y is the perturbed operator $\widehat{\text{Hess}}_{A, \pi}$. In the irreducible case, this operator has no kernel, and Lin's gluing theorem reduces to the abstract gluing theorem [KM07, Theorem 18.3.5].

First we will apply the abstract gluing lemma to see the corresponding statement about *unframed* moduli spaces: we will see that there are maps

$$u(T, -) : B(\mathcal{K}_\alpha) \rightarrow \mathcal{M}_{\alpha, k, \delta}(Z^T)$$

parameterizing solutions of the perturbed ASD equations in a neighborhood of γ_A satisfying the same conditions (without the equivariance); the corresponding restriction maps R^T are given by the actual restriction of connections to the boundary (as opposed to before, where it involved a global quantity, the holonomy).

The argument in [KM07, Section 18.4] applying the abstract gluing theorem to the Seiberg-Witten equations readily applies all the same to the ASD equations.

We sketch the argument, as we will use essentially the same argument in the reducible case. Fix a connection A on Y in the gauge equivalence class of α . We identify a neighborhood of α in $\mathcal{B}_{E, k-1/2}^e$ as $\ker(d_A^*)$ via the Coulomb slice.

On the cylinder Z^T , we identify the spaces

$$\Omega_{k,\delta}^1(Z^T; \mathfrak{g}_E)$$

and

$$(\Omega^0 \oplus \Omega^{2,+})_{k,\delta}(Z^T; \mathfrak{g}_E)$$

with

$$L_{k,\delta}^2(Z^T; \pi^*(\mathbb{R} \oplus T^*Y) \otimes \mathfrak{g}_E);$$

the first by writing every 1-form as $\psi + dt \wedge \sigma$ for ψ a time-dependent 1-form on Y , and the second by writing every self-dual 2-form as $(dt \wedge \psi)^+$.

Now note from this isomorphism that when we have $a \in \Omega_{k,\delta}^1(Z^T; \mathfrak{g}_E)$, we may keep track of more information at the boundary than just $a|_{Y \sqcup \bar{Y}}$; this restriction kills any term of the form $dt \wedge \sigma$. So we write

$$r : \Omega_{k,\delta}^1(Z^T; \mathfrak{g}_E) \rightarrow \Omega_{k-1/2}^1(Y \sqcup \bar{Y}; \mathfrak{g}_E) \oplus \Omega_{k-1/2}^0(Y \sqcup \bar{Y}; \mathfrak{g}_E),$$

to record both the restriction of a and its normal value $\sigma(\pm T, b)$, or $\sigma(0^\pm, b)$ in the case of the infinite cylinder.

On $\Omega_{k-1/2}^1(Y; \mathfrak{g}_E)$, we may project to $\ker(d_A^*)$, which has an eigenspace decomposition $\mathcal{K}_\alpha^- \oplus \mathcal{K}_\alpha^+$ for the action of $D_{A,\pi}$. (The positive and negative eigenspaces swap upon orientation-reversal.) Write

$$\Pi : \Omega_{k-1/2}^1(Y \sqcup \bar{Y}; \mathfrak{g}_E) \oplus \Omega_{k-1/2}^0(Y \sqcup \bar{Y}; \mathfrak{g}_E) \rightarrow \mathcal{K}_\alpha^- \oplus \Omega_{k-1/2}^0(Y; \mathfrak{g}_E) \oplus \mathcal{K}_\alpha^+ \oplus \Omega_{k-1/2}^0(\bar{Y}; \mathfrak{g}_E);$$

this is the spectral projection on the $\ker(d_A^*)$ term (projecting to negative eigenvalues on Y and positive eigenvalues on \bar{Y}), and records the data of $\sigma|_{\partial Z^T}$.

Now, on the linear space $\Omega_{k,\delta}^1(Z^T; \mathfrak{g}_E)$, consider the equations

$$\begin{aligned} d_{\gamma_A}^+ a + (a \wedge a)^+ + (dt \wedge \nabla_\pi(\gamma_A + a))^+ &= 0 \\ d_{\gamma_A}^* a &= 0 \\ \Pi(a|_{\partial Z^T}) &= c \end{aligned}$$

Here

$$c \in \mathcal{K}_\alpha^- \oplus \Omega_{k-1/2}^0(Y; \mathfrak{g}_E) \oplus \mathcal{K}_\alpha^+ \oplus \Omega_{k-1/2}^0(\bar{Y}; \mathfrak{g}_E);$$

we henceforth write this space as $H = H^- \oplus H^+$ to simplify notation. The above equations are simply the ASD equations with boundary conditions corresponding to the spectral projection of the $\Omega^1(Y)$ component and restriction of the $\Omega^0(Y)$ component.

The linearization of these equations is

$$\begin{aligned} Q_{\gamma_A, \pi} a &= 0 \\ \Pi(a|_{\partial Z^T}) &= 0, \end{aligned}$$

so the most important thing to check is that *these linearized equations determine an isomorphism*

$$\Omega_k^1(Z^\infty; \mathfrak{g}_E) \rightarrow (\Omega^{2,+} \oplus \Omega^0)_{k-1}(Z^\infty; \mathfrak{g}_E) \oplus H;$$

equivalently it suffices to show that

$$\Omega_k^1([0, \infty) \times Y; \mathfrak{g}_E) \rightarrow (\Omega^{2,+} \oplus \Omega^0)_{k-1}([0, \infty) \times Y; \mathfrak{g}_E) \oplus H^-$$

is an isomorphism. To do this, we write $Q_{\gamma_A, \pi} = \frac{d}{dt} + \widehat{\text{Hess}}_{A, \pi}$, and split

$$\Omega^0(Y) \oplus \Omega^1(Y) = \Omega^0 \oplus \text{Im}(d_A) \oplus \ker(d_A^*);$$

rewriting $\text{Im}(d_A) \cong \Omega_{k+1/2}^0(Y; \mathfrak{g}_E)$, we have

$$\widehat{\text{Hess}}_{A, \pi} = \begin{pmatrix} 0 & -\Delta_A \\ -1 & 0 \end{pmatrix},$$

which is invertible because A is irreducible. On the other component, our operator takes the form $D_{A, \pi}$, which is invertible because π is a regular perturbation.

First let us see that $(Q_{\gamma_A, \pi}, \Pi)$ is surjective. We may write an arbitrary element of the domain as

$$(\sigma_s, \psi_s, h, c) \in (\Omega^{2,+} \oplus \Omega^0)_{k-1}([0, \infty) \times Y; \mathfrak{g}_E) \oplus H^-,$$

where σ_s is a time-dependent 0-form and ψ_s is a time-dependent 1-form, thought of as the self-dual 2-form via $(dt \wedge \psi_s)^+$; $h \in \mathcal{K}_\alpha^-$ and $c \in \Omega_{k-1/2}^0(Y; \mathfrak{g}_E)$.

We may decompose $\Omega^0(Y)$ into the eigenspaces of Δ_A and $\ker(d_A^*)$ into the eigenspaces of $D_{A, \pi}$. For $\phi \in \Omega_k^1([0, \infty) \times Y; \mathfrak{g}_E)$ we write $\phi_s = \omega_s + d_A \eta_s$, where $\omega_s \in \ker(d_A^*)_{k-1/2}$. We will explicitly write out $(Q_{\gamma_A, \pi}, \Pi)$ in terms of these.

Using the eigenspace decompositions, we may write

$$\begin{aligned} \sigma_s &= \sum b_\lambda(s) \sigma_\lambda \\ \eta_s &= \sum c_\lambda(s) \sigma_\lambda \\ \psi_s &= \sum d_\lambda(s) \omega_\lambda, \end{aligned}$$

where the coefficients range over a basis of eigenfunctions.

Now we find that these must satisfy $\sigma'(s) = \Delta_A \eta(s)$ and $\eta'(s) = \sigma(s)$. These immediately give us

$$\begin{aligned} b_\lambda &= c'_\lambda, \\ b'_\lambda(s) &= \lambda c_\lambda(s). \end{aligned}$$

Combining these we solve to find

$$c_\lambda(t) = c_\lambda(0) \cdot e^{\pm\sqrt{\lambda}t}.$$

However, under the assumption that $\phi \in L_k^2$ and hence $\eta \in L_{k+1}^2$, the sign must be negative. By specifying c above, we specify the values of $c_\lambda(0)$ for all λ , and hence both b and c for all time. Because $\sum c_\lambda(0)$ defines an $L_{k-1/2}^2$ function on Y and the sum $\sum e^{-\sqrt{\lambda}t} c_\lambda(0)$ satisfies an elliptic equation, regularity implies that this defines an L_k^2 function on $[0, \infty) \times Y$ (and similarly for b).

Now for ψ_s we get the formula $\omega'_\lambda(s) = \lambda \omega_\lambda(s)$. This can only contribute to the solution if $\lambda < 0$, as otherwise ω is not L^2 . We specified the values of $\omega_\lambda(0)$ for $\lambda < 0$ in the choice of $h \in H^-$, and thus have specified ω for all time; as above it has the desired regularity. We have constructed a solution to the equations of the desired regularity and with desired boundary values, so (Q, Π) is surjective.

To see that (Q, Π) is injective, we remark that the above argument always produced *unique* solutions.

This in hand, the abstract gluing theorem gives us a map

$$u(T, -) : B_\eta \left(\mathcal{K}_\alpha^- \oplus \mathcal{K}_\alpha^+ \oplus \Omega_{k-1/2}^0(Y \sqcup \bar{Y}) \right) \rightarrow \Omega_{k, \delta}^1(Z^T; \mathfrak{g}_E)$$

parameterizing solutions in a neighborhood of 0. To use this to get a parameterization of γ_A in $\mathcal{M}_{A,k,\delta}(Z^T)$, we observe that a neighborhood of γ_A in the latter is given by the Coulomb-Neumann slice $\mathcal{CN}_{A,k,\delta}$, and this parameterization is compatible with the restriction maps; so to obtain the desired parameterizations, we simply need to restrict the domain of u to $B_\eta(\mathcal{K}_\alpha^- \oplus \mathcal{K}_\alpha^+)$, imposing the Neumann gauge condition that the $dt \wedge \sigma$ component is zero on the boundary.

Recall that projection $\widetilde{\mathcal{M}}(Z^T) \rightarrow \mathcal{M}(Z^T)$ forms a principal $SO(3)$ -bundle over the irreducible instantons. Our goal is to choose a section of this over $B(\mathcal{K}_\alpha)$, sitting inside $\widetilde{\mathcal{M}}(Z^T)$ via $u(T, -)$, which will automatically give us an equivariant map $SO(3) \times B(\mathcal{K}_\alpha) \rightarrow \widetilde{\mathcal{M}}(Z^T)$, as desired. But we should be careful in how we choose this lift so that the maps μ_T have the desired properties.

We identified a neighborhood of γ_A in $\mathcal{M}_{A,k,\delta}(Z^T)$ with a subset of $\mathcal{CN}_{A,k,\delta}(Z^T)$. We may thus define the section above $\mathcal{CN}_{A,k,\delta}(Z^T)$; choose the section to be

$$\begin{aligned} \mathcal{CN}_{A,k,\delta}(Z^T) &\rightarrow SO(3) \times \mathcal{CN}_{A,k,\delta}(Z^T) \\ \mathbf{A} &\mapsto (p, \mathbf{A}) \end{aligned}$$

for p a fixed framing.

Then writing $B(\widetilde{\mathcal{K}}_\alpha) = SO(3) \times \widetilde{\mathcal{K}}_\alpha$, our parameterization

$$\tilde{u}(T, -) : SO(3) \times B(\mathcal{K}_\alpha) \rightarrow \widetilde{\mathcal{M}}(Z^T)$$

is given as $(\tilde{u})(g, h) = \tilde{s}(u(h)) \cdot g$ in the ‘framed Coulomb-Neumann slice’ $SO(3) \times \mathcal{CN}_{A,k,\delta}$. That this is a smooth embedding for all T follows from the corresponding fact for unframed moduli spaces, and because μ_T is assumed smooth in finite T , so is $\tilde{\mu}_T$; similarly $\tilde{\mu}_T \rightarrow \tilde{\mu}_\infty$ in the C_{loc}^∞ topology.

Finally, because the unframed map $u(\infty, -)$ respects the product structure and we defined \tilde{u} to be a fiber bundle map, the last statement is true. \blacksquare

We quickly discuss linear models for the configuration space $\widetilde{\mathcal{B}}_{A,k,\delta}(Z^T)$ when A is fully reducible. When T is finite, choosing a base framing p , a neighborhood of 0 in the Coulomb slice

$$\mathcal{CN}_{A,k,\delta}(Z^T) \subset \mathcal{A}_{A,k,\delta}(Z^T) \times \{p\} \subset \widetilde{\mathcal{A}}_{A,k,\delta}(Z^T)$$

projects to a chart in $\widetilde{\mathcal{B}}_{\pi^*E,k,\delta}(Z^T)$ around γ_A , so we may effectively study neighborhoods of γ_A in the moduli space on finite cylinders by studying a subset of the equations in Coulomb-Neumann gauge. We will consider the Coulomb-Neumann slice as a subset of $\Omega_{k,\delta}^1(Z^T)$ and work there. In this case, the space $H_0 = \ker(L)$ above is the Lie algebra \mathfrak{g}_b ; when considered as a subset of $\Omega^1(Z^T)$, these consist of the 1-forms $dt \wedge \sigma$, where σ is a fixed A -parallel section of \mathfrak{g}_b .

The discussion of the Coulomb-Neumann slice remains correct on the infinite cylinder: everything in a neighborhood of γ_A in $\widetilde{\mathcal{B}}_{A,k,\delta}(Z^\infty)$ is gauge equivalent to an asymptotically decaying connection in Coulomb-Neumann gauge. Recall here that our definition of $\widetilde{\mathcal{B}}_{A,k,\delta}(Z^\infty)$ involves *two framings*: one on each boundary component. Applying A -parallel gauge transformations on each component, we may change those framings arbitrarily while obtaining another connection in Coulomb-Neumann gauge.

In fact, while we will consider the space $\mathcal{E}_\delta^\infty = \Omega_{k,\delta,\text{ext}}^1(Z^\infty; \mathfrak{g}_E)$ of 1-forms which asymptotically decay to $dt \wedge \sigma$, in the case of the infinite cylinder the extra kernel arising as $dt \wedge \sigma$ is somehow illusory: applying a gauge transformation we may

change any connection in this extended space to a connection in Coulomb-Neumann gauge.

Similarly, for any $a \in \mathcal{CN}_{A,k,\delta}$, we have by definition of Coulomb-Neumann gauge the integration by parts formula

$$0 = \langle d_{\gamma_A}^* a, b \rangle_{L^2} = \langle a, d_{\gamma_A} b \rangle_{L^2}.$$

Above, the projection Π_0 was projection onto the constant sections $dt \wedge \sigma$ in the $L_{0,\delta}^2$ inner product; then this is defined to be the same as $\langle a, g_{T,\delta}^2 dt \wedge \sigma \rangle_{L^2}$, where $g_{T,\delta}$ is the weight function in the definition of Sobolev spaces on finite cylinders.

But we may choose an antiderivative $\frac{d}{dt} G_{T,\delta} = g_{T,\delta}^2$, and then the right-hand side of this inner product is $d_{\gamma_A}(G_{T,\delta}\sigma)$, because σ is A -parallel. Therefore, for any solution in Coulomb-Neumann gauge, $\Pi_0^T a = 0$. With this, we begin the proof.

Proof of Lemma 4.39 in the fully reducible case. The equations we look to study with Lin's abstract gluing theorem are nearly the same as last time: now they are

$$\begin{aligned} d_{\gamma_A}^+ a + (a \wedge a)^+ + (dt \wedge \nabla_\pi(\gamma_A + a))^+ &= 0 \\ d_{\gamma_A}^* a &= 0 \\ \Pi(a|_{\partial Z^T}) &= c \\ \Pi_0^T a &= h_0. \end{aligned}$$

That our operator (D, Π_0, Π) satisfies the invertibility assumption follows by the same separation of variables argument as in the irreducible case.

Lin's theorem, therefore, gives us a parameterization

$$u(T, h_0, h) : B_\eta(H_0 \oplus H) \rightarrow \Omega_{k,\delta}^1(Z^T)$$

of solutions to the above equations on Z^T with $(h_0, h) \in B_\eta(H_0 \oplus H)$; recall that here H_0 consists of 1-forms of the form $dt \wedge \sigma$ for $\sigma \in \ker(\Delta_A)$ and

$$H = \mathcal{K}_A^- \oplus \Omega^0(Y; \mathfrak{g}_E) \oplus \mathcal{K}_A^+ \oplus \Omega^0(\bar{Y}; \mathfrak{g}_E).$$

For his application, Lin restricts to the subdomain $H_0 \oplus \mathcal{K}_A^- \oplus \mathcal{K}_A^+$; based on the discussion above, we see that in our case we should actually restrict to the subdomain $\mathcal{K}_A^- \oplus \mathcal{K}_A^+$ and ignore the terms coming from H_0 . We abuse notation and write $u(T, h) = u(T, 0, h)$. As long as T is finite, all solutions thus obtained are in Coulomb-Neumann gauge, and this defines a parameterization of a small neighborhood of γ_A in $\widetilde{\mathcal{M}}_{A,k}(Z^T)$.

Before going on, we should explain why this map u is equivariant; this follows as a consequence of uniqueness of solutions. That is, if g is a gauge transformation preserving γ_A , then $u(gh_0, gh)$ is the unique solution a to the equations

$$\begin{aligned} d_{\gamma_A}^+ a + (a \wedge a)^+ + (dt \wedge \nabla_\pi(\gamma_A + a))^+ &= 0 \\ d_{\gamma_A}^* a &= 0 \\ \Pi(a|_{\partial Z^T}) &= gc \\ \Pi_0^T a &= 0 \end{aligned}$$

Here $c = (h, 0, 0)$ and Π is the projection operator to

$$\mathcal{K}^- \oplus \mathcal{K}^+ \oplus \Omega^0(Y \sqcup \bar{Y}; \mathfrak{g}_E) \oplus H_0.$$

But this equation is also satisfied by $gu(T, h)$, because the projections are equivariant and the ASD equations are invariant under gauge transformations.

So $u(T, h)$ defines a map

$$B_\eta(\mathcal{K}_\alpha) \rightarrow \mathcal{CN}_{A,k,\delta}(Z^T) \times \{p\} \subset \tilde{\mathcal{A}}_{A,k,\delta}(Z^T)$$

for T finite, where p is a fixed framing; projecting to $\tilde{\mathcal{B}}_{A,k,\delta}(Z^T)$, this gives a neighborhood of γ_A in $\tilde{\mathcal{M}}_{A,k,\delta}(Z^T)$.

So we have our parameterization of a neighborhood of γ_A for all $T \in [T_0, \infty]$; we should check the desired properties of the restriction maps. The abstract gluing theorem guarantees certain properties of the maps $\mu_T^\pm : B_\eta(H) \rightarrow \ker(d_A^*)_{k-1/2}$ obtained as the restriction of $\tilde{u}(T, h)$ to the corresponding boundary component. However, our gluing lemma asks for properties of a restriction map *whose definition includes the framing*, written $\tilde{\mu}_T$. In the case of restriction to the leftmost boundary component, there is no change, as the framing is already fixed on the left boundary component; there is no need to perform any holonomy. But the right endpoint map is more complicated.

Let's be precise. Let (\mathbf{A}, p) be a framed connection on Z^T , where the framing is chosen at $(-T, b)$. Then the definition of the right endpoint map is

$$(\mathbf{A}, p) \mapsto \left(\mathbf{A}|_{\bar{Y}}, \text{Hol}_{\mathbf{A}}^{(-T,b) \rightarrow (T,b)} p \right);$$

we then pass to the quotient by the gauge group action (as this is an equivariant map).

Now, we identified a neighborhood of A in $\mathcal{B}_{E,k-1/2}(\bar{Y})$ with the projection of the Coulomb slice $\ker(d_A^*) \times \{p\} \subset \tilde{\mathcal{A}}_{E,k-1/2}(\bar{Y})$, where p is the same fixed framing as above. Therefore, identifying a neighborhood of γ_A in $\tilde{\mathcal{M}}_{A,k,\delta}(Z^T)$ with the corresponding subset of the Coulomb-Neumann slice (and a neighborhood of A with the corresponding Coulomb slices), the restriction $\tilde{\mu}_T^-$ to the left boundary component is still given by the same

$$\mu_T^- : \mathcal{CN}_{A,k,\delta}(Z^T) \rightarrow \ker(d_A^*)_{k-1/2}(Y)$$

as in the abstract gluing theorem, because our framing is chosen fixed at the left boundary component to begin with. However, the restriction to the right boundary component is now given by

$$\tilde{\mu}_T^+(a) = \sigma^T(a) \cdot \mu^T(a) \in \ker(d_A^*)_{k-1/2}(\bar{Y}),$$

where we have identified $\sigma^T(a) \in SO(3)$ with an A -parallel gauge transformation.

Precisely, the map σ^T is given by

$$\sigma^T(a) = \left(\text{Hol}_{\gamma_{A+a}}^{(-T,b) \rightarrow (T,b)} \right)^{-1},$$

using the same framing p at $\pm T$ to identify this isomorphism $E_{(-T,b)} \rightarrow E_{(T,b)}$ with an element of $SO(3)$. We have $\sigma(0) = 1$: the connection γ_A is in temporal gauge and so there is no holonomy across $\mathbb{R} \times \{b\}$.

What we need to do first is see that $\sigma^T(u(T, h))$ converges in the C_{loc}^∞ topology. We will then use this to define the map $\tilde{u}(\infty, -)$, as an appropriate modification of $u(\infty, 0, -)$.

First we recall how holonomy is defined. First, given $a \in \gamma_A + \Omega_{k,\delta}^1(Z^T)$, there is a natural restriction map to

$$d/dt + \Omega_{k-3/2,\delta}^1([-T, T], \mathfrak{g}_b) = d/dt + dt \wedge \Omega_{k-3/2,\delta}^0([-T, T], \mathfrak{g}_b)$$

to the line $\mathbb{R} \times \{b\}$. The holonomy along the path only depends on the values of a on this path, and more precisely, if we decompose $a = dt \wedge \eta(t) + \psi(t)$, it only depends on $\eta(b, t)$. Precisely, let $\gamma(t) : [-T, T] \rightarrow SO(3)$ be the unique solution to the differential equation

$$\gamma'(t)\gamma^{-1}(t) = \eta(b, t) \quad \text{and} \quad \gamma(-T) = 1,$$

guaranteed as long as $k-3/2 \geq 1$, so that η is C^1 . Solving this differential equation gives a smooth map

$$\text{Lift} : L_{k-3/2,\delta}^2([-T, T], \mathfrak{g}_b) \rightarrow L_{k-5/2,\delta}^2([-T, T], SO(3)),$$

and then evaluation at T is clearly smooth.

If we write $u_+(h)$ as the component of $u(\infty, h)$ on the cylinder $[0, \infty) \times Y$, and similarly for $u_-(h)$, we write

$$U(T, h) = \tau_T^* u_+(h) + \tau_{-T}^* u_-(h).$$

We see from [Lin18, Lemma 3.5.19] that $U(T, h) - \tilde{u}(T, h)$ goes to 0 in the C_{loc}^∞ topology. So it suffices to check that $\sigma^T(U(T, h))$ has a limit in C_{loc}^∞ .

Now we may be very explicit. By translation, consider instead the interval $[0, 2T]$; we may see that $\sigma^T(U(T, h))^{-1}$ is the solution to the differential equation $\gamma'(t)\gamma^{-1}(t) = u_+(h)(t, b) + u_-(h)(t - 2T, b)$ with $\gamma(0) = 1$, evaluated at $2T$. We may instead consider this as the product $g_2 g_1$ of two elements of $SO(3)$: first, g_1 is the time- T value of that same equation; second, g_2 is the time- T value of the solution to the differential equation

$$\gamma'(t)\gamma^{-1}(t) = u_-(h)(-t, b) + u_+(h)(2T - t, b) \quad \text{and} \quad \gamma(0) = 1.$$

That $g = g_2 g_1$ is just the statement that the holonomy of two paths, traversed in succession, is the product of their holonomies and that the holonomy of a path traversed in reverse is the inverse of the holonomy.

We examine g_1 ; the analysis of the other term is very similar. Notice that the stated differential equation only needs the values of the given functions on $[0, T]$, so we introduce a cutoff function β which has $\beta(t) = 1$ for $t \leq 0$ and $\beta(t) = 0$ for $t \geq 1/2$, and $\beta_T(t) = \beta(t - T)$; write

$$q_T(h) = u_+(h)(t, b) + \beta_T(t) (u_-(h)(t - 2T, b)).$$

It is clear that $U(T, h)$ and $q_T(h)$ give the same time- T holonomy map, as they agree on $[0, T]$. Our first claim, to be proved below, is that the time-dependent map $B_\eta(\mathcal{K}_\alpha) \rightarrow L_{k-3/2,\delta}^2([0, \infty), \mathfrak{g}_b)$ given by sending $h \mapsto \beta_T(t)u_-(h)(t - 2T, b)$ converges to 0 in C_{loc}^∞ ; given that this is true, our goal becomes finding the C_{loc}^∞ limit as $T \rightarrow \infty$ of

$$\text{ev}_T \text{Lift}(u_+(h)(t, b)).$$

Now there is nothing left to do here, as u_+ is time-independent! We see that the output converges in the C_{loc}^∞ topology to

$$\text{Hol}_{\gamma_A + u_+(h)}^{(0^+, b) \rightarrow (\infty, b)}.$$

Similarly we may identify the limit of g_2 as the inverse of the corresponding $\text{Hol}_{u_-(h)}^{(0^-,b)\rightarrow(-\infty,b)}$.

We conclude, therefore, that $\sigma^T(u(T, h))$ converges in C_{loc}^∞ to $\sigma^\infty(h)$, which is defined to be the inverse of

$$\left(\text{Hol}_{u_-(h)}^{(0^-,b)\rightarrow(-\infty,b)}\right) \left(\text{Hol}_{\gamma_A+u_+(h)}^{(0^+,b)\rightarrow(\infty,b)}\right).$$

This makes some intuitive sense; the appearance of σ^T corresponds to the twist in framing as we move the left endpoint framing to the right endpoint. So on the infinite cylinder, we should be forced to take holonomy from left endpoint to right endpoint.

Based on the above, we define the map

$$\tilde{u}(\infty, -) : B_\eta(\mathcal{K}_\alpha) \rightarrow \mathcal{CN}_{A,k,\delta}(Z^\infty) \times SO(3) \times SO(3) \subset \tilde{\mathcal{A}}_{A,k,\delta}(Z^\infty)$$

to be

$$\tilde{u}(\infty, h) = (u(\infty, h), p, \sigma^\infty(h) \cdot p);$$

we project from this framed Coulomb slice to the configuration space $\tilde{\mathcal{M}}(Z^\infty)$. Then the above discussion shows that indeed, $\tilde{\mu}_T \rightarrow \tilde{\mu}_\infty$ as $T \rightarrow \infty$.

Now to prove the claim that $\beta_T(t)(u_+(h)(t-2T, b))$ converges to 0 in C_{loc}^∞ . The point is that the application of the inverse function theorem defining u_+ can be factored through a Sobolev space with smaller exponent, so that one may find pointwise bounds on all of the derivatives of u_+ :

$$\|\mathcal{D}_m u_+(h)(t-2T, b)\| \leq C_m(h) e^{-\delta'(t-2T)},$$

where the C_m are continuous functions of h depending on m ; the same is true when adding in the bump function (whose derivatives all have compact support).

Because the support of $\beta_T u_+(t-2T)$ is contained in $[0, T+1]$, we obtain bounds by $C_m(h) e^{-\delta'(1-T)}$; taking $T \rightarrow \infty$ we see that this goes to 0 in C_{loc}^∞ . \blacksquare

What remains is the case of $SO(2)$ -reducible critical orbits. We use a mix of the above techniques; we work in the Coulomb-Neumann slice as a model for the normal space to γ_A , and then extend this to an $SO(3)$ -invariant parameterization.

Proof of Lemma 4.39 in the $SO(2)$ -reducible case. We study the same equations as above, where γ_A is an $SO(2)$ -reducible constant trajectory. We start by choosing $p \in S^2 = E_{(0,b)}/\Gamma_A$, a choice of framing *modulo* the action of Γ_A ; restrict to the subspace $\tilde{\mathcal{M}}'(Z^T)$ given by those pairs (\mathbf{A}, q) with framings that project to p in the quotient. (This is the space that has a neighborhood modelled by the solutions to the equations in the Coulomb-Neumann slice.) We find just as above that Lin's abstract gluing theorem again provides us with an S^1 -equivariant parameterization $B_\eta(\mathcal{K}_\alpha) \rightarrow \tilde{\mathcal{M}}(Z^T)$. The issue is in verifying that the restriction maps to the right end, which depend on holonomy, converge as $T \rightarrow \infty$ in C_{loc}^∞ , as before.

We may choose an arbitrary lift of p to an actual framing \tilde{p} . Note that because γ_A is in temporal gauge, the holonomy from $-T$ to T is the identity. Now choosing the domain of the parameterization u small enough, we may demand that $\text{Hol}_{\gamma_A+u(T,0,h)}^{(-T,b)\rightarrow(T,b)} \tilde{p}$ lies in an S^1 -invariant neighborhood of $\Gamma_A \cdot \tilde{p}$. Exponentiating the normal bundle, we may choose an S^1 -equivariant diffeomorphism $U \cong S^1 \times D^2$. This gives us an S^1 -equivariant projection map $P : U \rightarrow S^1$. We write

$$\overline{\text{Hol}}_{\gamma_A+u(T,0,h)}^{(-T,b)\rightarrow(T,b)} = P \text{Hol}_{\gamma_A+u(T,0,h)}^{(-T,b)\rightarrow(T,b)} \tilde{p};$$

this is precisely the amount of holonomy in the ‘ S^1 -direction’, and the analogue of $\sigma^T(u)^{-1}$ for the $SO(2)$ -reducible case (that is, the inverse of this term is precisely the rotation appearing in the right-most restriction map.)

We saw above that $\text{Hol}_{u(T,0,h)}^{(-T,b) \rightarrow (T,b)}$ does converge in C_{loc}^∞ as $T \rightarrow \infty$, and to

$$\left(\text{Hol}_{u(\infty,0,h)}^{(0^-,b) \rightarrow (-\infty,b)} \right)^{-1} \text{Hol}_{u(\infty,0,h)}^{(0^+,b) \rightarrow (\infty,b)}.$$

Because the projection $P : U \rightarrow S^1$ is smooth, the same is true of $\overline{\text{Hol}}_{u(T,0,h)}^{(-T,b) \rightarrow (T,b)}$; we use this to define the map $u(\infty, 0, h)$ on $B_\eta(\mathcal{K}_\alpha) \rightarrow \widetilde{\mathcal{M}}'_{A,k,\delta}(Z^\infty)$ as before, and see that it is $SO(2)$ -equivariant.

All that is left is to extend this to an $SO(3)$ -equivariant parameterization of the whole of $\widetilde{\mathcal{M}}_{A,k,\delta}(Z^T)$. We write the neighborhood of (γ_A, \tilde{p}) as

$$(p \cdot S^1) \times_{S^1} \mathcal{CN}_{\gamma_A,k,\delta}(Z^T).$$

That is, we allow ourselves to vary over all framings that lie above the fixed $p \in S^2$, and then quotient by the natural action of Γ_A . Then there is a canonical S^1 -equivariant map

$$S^1 \times_{S^1} \mathcal{CN}_{\gamma_A,k,\delta}(Z^T) \rightarrow SO(3) \times_{S^1} \mathcal{CN}_{\gamma_A,k,\delta}(Z^T).$$

This is our desired section. We thus extend our parameterization to an $SO(3)$ -equivariant parameterization $SO(3) \times_{S^1} B_\eta(\mathcal{K}_\alpha) \rightarrow \widetilde{\mathcal{M}}(Z^T)$ of a neighborhood of the orbit through γ_A , as desired. \blacksquare

With this, the proof of the gluing theorem follows essentially the same lines as in [KM07, Chapters 19 and 24.7]. We outline the procedure and show what mild modifications are necessary. For the rest of this section, if (W, \mathbf{E}, π) is a manifold with cylindrical ends and perturbation π , we demand that π is regular.

First, the moduli spaces $\widetilde{\mathcal{M}}_{E,z,\pi,k}(Z^T)$ of π -perturbed instantons on the cylinder, finite or infinite, are cut out transversely inside $\widetilde{\mathcal{B}}_{E,z,\pi,k}^e(Z^T)$, and therefore form a smooth Hilbert manifold no matter the perturbation π . The proof of this is no different than [KM07, Theorem 17.3.1]; the point is that we have a unique continuation result for zeroes of the adjoint operator $Q_{\mathbf{A},\pi}^*$, as the equations are of gradient-flow type.

Fix a compact manifold W with boundary $Y_1 \sqcup \overline{Y_2}$, with cylindrical metric near the boundary, and $SO(3)$ -bundle \mathbf{E} which restricts to bundles E_i over the respective ends Y_i . Suppose we have fixed small perturbations π_i on (Y_i, E_i) so that there are finitely many nondegenerate critical orbits of $cs_{Y_i} + f_{\pi_i}$. Further, we assume (W, \mathbf{E}) is weakly admissible, so it admits *some* perturbation achieving transversality on all moduli spaces of bounded dimension on the infinite cobordism \hat{W} .

On the cylindrical ends of the infinite manifold \hat{W} , we fix the constant perturbation $(dt \wedge \nabla_{\pi_\pm}(\mathbf{A}))^+$. This is dampened by a cutoff function on the end

$$[0, \infty) \times (Y_1 \sqcup \overline{Y_2}),$$

equal to 0 for $x \leq 1$ and equal to 1 for $x \geq 2$. There is a further interior holonomy perturbation supported in the complement of $[1, \infty) \times (Y_1 \sqcup \overline{Y_2})$.

For $L \geq 2$, if W^L is the complement of $(L, \infty) \times (Y_1 \sqcup \overline{Y_2})$, then the moduli space of instantons $\widetilde{\mathcal{M}}_{\mathbf{E},\pi,k}(W^L)$ on the compact manifold W^L is a smooth Hilbert manifold in a neighborhood of $\widetilde{\mathcal{M}}_{\mathbf{E},k,\delta}(\hat{W})$. This is a matter of showing that if \mathbf{A} is

a π -instanton on \widehat{W} , then any nonzero element of the kernel of $Q_{\mathbf{A},\pi}^*$ on W^L must restrict nontrivially to the boundary.

For suppose to the contrary that we had $\psi \in \Omega_{-k}^{2,+}(W^L; \mathfrak{g}_{\mathbf{E}})$ with $\psi|_{\partial W^L} = 0$. Elliptic regularity allows one to obtain $\psi \in L_k^2$ on $W^L \setminus W$, where the perturbation arises from the gradient of cylinder functions on slices. For the same reason, we can argue that $\psi|_{\partial W^L} = 0$ means by unique continuation that $\psi|_{\partial W^L \setminus W} = 0$. In particular, ψ extends by 0 to an element

$$\widehat{\psi} \in \Omega_{-k,\delta}^{2,+}(\widehat{W}; \mathfrak{g}_{\mathbf{E}}), \quad \text{with } Q_{\mathbf{A},\pi}^* \widehat{\psi} = 0.$$

The assumption that π is a regular perturbation then means that $\widehat{\psi} = 0$ globally, and so $\psi = 0$, as desired.

We write

$$\widetilde{\mathcal{M}}_{\mathbf{E},k,\delta} \subset \widetilde{V}_{\mathbf{E},\pi,k}^L \subset \widetilde{\mathcal{M}}_{\mathbf{E},\pi,k}(W^L),$$

the middle term denoting an open set of π -instantons \mathbf{A} on W^L for which $Q_{\mathbf{A},\pi} : \Omega_k^1(W^L; \mathfrak{g}_{\mathbf{E}}) \rightarrow \Omega_{k-1}^2(W^L; \mathfrak{g}_{\mathbf{E}})$ is surjective.

From here, we may describe moduli spaces on the manifold W with infinite cylindrical ends as a fiber product. Let I_j^+ and I_ℓ^- be a finite sequence of intervals ($0 \leq j \leq n, 0 \leq \ell \leq m$), where the first negative interval is $I_0^- = (-\infty, 0]$, the first positive interval is $I_0^+ = [0, \infty)$, and all other intervals are finite. Write $\widetilde{\mathcal{B}}^\pm = \widetilde{\mathcal{B}}_{E,k-1/2}^e(Y^\pm)$. There are evaluation maps

$$\text{ev}_\pm : \widetilde{V}_{\mathbf{E},\pi}^L \rightarrow \widetilde{\mathcal{B}}^- \times \widetilde{\mathcal{B}}^+,$$

and similarly for $\widetilde{\mathcal{M}}_{\mathbf{E},\pi,k}(Z^{I_j^\pm}) \rightarrow \widetilde{\mathcal{B}}^\pm \times \widetilde{\mathcal{B}}^\pm$. For the infinite ends, we only have one evaluation map. We assemble all of this into a map

$$(R_-, R_+) : \widetilde{V}_{\mathbf{E},\pi}^L \times_{j=0}^n \widetilde{\mathcal{M}}_{E,\pi,k}(Z^{I_j^+}) \times_{\ell=0}^m \widetilde{\mathcal{M}}_{E,\pi,k}(Z^{I_\ell^-}) \rightarrow \left(\times_{j=1}^n \widetilde{\mathcal{B}}^+ \times_{\ell=1}^m \widetilde{\mathcal{B}}^- \right)^2.$$

Lemma 4.40. *Let W be the cobordism with cylindrical ends attached. There is a natural map from the fiber product $\text{Fib}(R_-, R_+)$ to $\widetilde{\mathcal{M}}_{\mathbf{E},\pi,k}(\widehat{W})$, which is a homeomorphism. Because any instanton \mathbf{A} in the image is assumed to be cut out transversely, the corresponding element of $\widetilde{V}_{\mathbf{E},\pi}^L \times \prod \widetilde{\mathcal{M}}(Z^I)$, the map $R_- \times R_+$ is transverse to the diagonal.*

We do not repeat the proof, which follows essentially as in [KM07, Theorem 19.1.4].

We now want a canonical way to cut up the real line (or multiple copies of the real line, in the case of broken trajectories) so that every instanton is written as a fiber product, as above, where most intervals are of a fixed large length $2L$ (where most of the energy is supported), and otherwise are of large variable length $T \geq T_0$. The mechanism for this is given in [KM07, Section 19.2]. For $A \in \widetilde{\mathcal{B}}_{E_i,k-1/2}^e(Y_i)$, we write $c_i(A) = \|\nabla_A(cs + f_\pi)\|_{L^2}$. Fix $\epsilon > 0$ so that any nontrivial instanton on $\mathbb{R} \times Y_i$ has $c_i(\mathbf{A}(t)) \leq \epsilon$ for some $t \in \mathbb{R}$. Let β be a cutoff function equal to 1 for $x \geq \epsilon$ and 0 for $x \leq \epsilon/2$.

Definition 4.15. *Let $I = [-L, L]$. An instanton on $I \times Y$ is said to be centered if*

- (1) $c(\mathbf{A}(t)) \leq \epsilon/2$ for $t \in [-L, -L+1] \cup [L-1, L]$,

(2) $c(\mathbf{A}(t)) \geq \epsilon$ for some $t \in [-L, L]$, and

(3) The center of mass

$$\int t\beta(c(\mathbf{A}(t)))dt / \int \beta(c(\mathbf{A}(t)))dt$$

is zero.

The space of centered instantons forms a smooth Hilbert manifold. Then we have the following analogue of [KM07, Proposition 24.7.3], with the same proof.

Lemma 4.41. *For any compact subset*

$$K \subset \widetilde{\mathcal{M}}_{E_1, k, \delta}(\alpha_0, \alpha_1) \times_{\alpha_1} \cdots \times \widetilde{\mathcal{M}}_{E_1, k, \delta}(\alpha_{n-1}, \alpha_n) \\ \times_{\alpha_n} \widetilde{\mathcal{M}}_{\mathbf{E}, k, \delta}(\alpha_n, \beta_0) \times_{\beta_0} \cdots \times_{\beta_{m-1}} \widetilde{\mathcal{M}}_{E_2, k, \delta}(\beta_{m-1}, \beta_m),$$

there is an L_0 so that for all $\infty > L \geq L_0$, we have a neighborhood $K \subset V(K, L) \subset \overline{\mathcal{M}}_{\mathbf{E}, k, \delta}(\alpha_0, \beta_m)$ for which

(1) For any $\mathbf{A} \in V(K, L)$, the restriction of \mathbf{A} to ∂W^L has $c(\mathbf{A}|_{\partial W^L}) \leq \epsilon/2$,

(2) Any $\mathbf{A} \in V(K, L)$ admits a unique collection of cylinders of length $2L$ in the complement of W^L so that $\mathbf{A}|_I$ is a centered instanton, while

(3) The complement of both W^L and these intervals consists of cylinders of length at least T_0 , where T_0 is larger than the least T for which Lemma 4.39 applies to each of Y_1 and Y_2 .

As a corollary, we have a map $d : V(K, L) \rightarrow (0, \infty]^{n+m}$ measuring the distance between the centers of successive centered intervals; for the ends closest to W , we consider ∂W to be the location of that corresponding center. It is easy to see that the values of d do not depend on K or L ; in particular, taking the union over all $V(K, L)$, we obtain a map $V \rightarrow (0, \infty]^{n+m}$ from a neighborhood of the stratum in $\overline{\mathcal{M}}$, the neighborhood V consisting of all instantons that may be partitioned as above.

At the same time, we have a map

$$\text{ev}_T : (0, \infty]^{n+m} \times \widetilde{\mathcal{M}}_{E_1, k, \delta}^{(-\infty, 0] \times Y_1}(\alpha_0) \times_{i=1}^n B_\eta(\tilde{\mathcal{K}}_{\alpha_i}) \times^n \widetilde{\mathcal{M}}_{\text{cen}}(Z_1^L) \times \tilde{V}_{\mathbf{E}, \pi, k}^L \\ \times_{j=0}^{m-1} B_\eta(\tilde{\mathcal{K}}_{\beta_j}) \times^m \widetilde{\mathcal{M}}_{\text{cen}}(Z_2^L) \times \widetilde{\mathcal{M}}_{E_2, k, \delta}^{[0, \infty) \times Y_2}(\beta_m) \rightarrow \left(\times^n \tilde{\mathcal{B}}_{E_1, k-1/2} \times^m \tilde{\mathcal{B}}_{E_2, k-1/2} \right)^2,$$

the map in every case given by restriction; perhaps most notably here, for the $B_\eta(\mathcal{K}_A)$ factors we are using the maps $\tilde{u}(T, -)$ of Lemma 4.39 to parameterize a neighborhood of γ_A in $\widetilde{\mathcal{M}}_{\gamma_{\alpha_i}, k, \delta}(Z^T)$, where $T \in [T_0, \infty]$, and then restricting to the boundary; Lemma 4.39 included the fact that these restriction maps converge in C_{loc}^∞ as $T \rightarrow \infty$.

We call the domain $(0, \infty]^{n+m} \times \mathcal{M}$ and the codomain \mathcal{N} . This map is smooth on each stratum of $(0, \infty]^{n+m}$, and for a convergent sequence $T_n \in (0, \infty]^{n+m}$, the maps ev_{T_n} converge to the appropriate limit in the C_{loc}^∞ topology. We may identify $\text{ev}_T^{-1}(\Delta)$ in the above fiber product as the neighborhood $V(L)$ in $\overline{\mathcal{M}}_{\mathbf{E}, k, \delta}(\alpha_0, \beta_m)$. Now we may apply [KM07, Lemma 19.3.3].

Theorem 4.42. *Let W be a compact 4-manifold equipped with two boundary components, cylindrical metric near the boundary, an $SO(3)$ -bundle \mathbf{E} restricting to the pullback of fixed bundles over the ends, and a curve γ which is cylindrical at the basepoints on the ends. Further let W be equipped with a perturbation π so that $\overline{\mathcal{M}}_{\mathbf{E},\pi,z}(\alpha,\beta)$ is cut out transversely whenever the expected dimension is at most 10.*

Then suppose we are given any open stratum $\sigma \subset \overline{\mathcal{M}}_{\mathbf{E},\pi,z}(\alpha,\beta)$, where $\text{gr}_z(\alpha,\beta) \leq 10 - \dim \alpha$. In that case, we may find a neighborhood $V(\sigma)$ and an $SO(3)$ -invariant map $V(\sigma) \rightarrow (0,\infty]^{n+m}$ which is stratum-preserving, a submersion on each stratum, and a topologically trivial fiber bundle for the inverse image of a neighborhood of ∞ . In particular, $\overline{\mathcal{M}}_{\mathbf{E},\pi,z}(\alpha,\beta)$ is a topological manifold with corners and a smooth structure on each stratum.

Proof. Above we constructed a map $\text{ev}_T : (0,\infty]^{n+m} \times \mathcal{N} \rightarrow \mathcal{R}$, where \mathcal{N} and \mathcal{R} are Hilbert manifolds; here $T \in (0,\infty]^{n+m}$. For fixed T , the map ev_T is smooth, and $T \mapsto \text{ev}_T$ is convergent in the C_{loc}^∞ topology because the same is true of the restriction maps μ_T on $B(\tilde{\mathcal{K}}_{\alpha_i})$ and $B(\tilde{\mathcal{K}}_{\beta_j})$. Thus the aforementioned lemma tells us that $\text{ev}^{-1}(\Delta) \rightarrow (0,\infty]^{n+m}$ is a topological submersion, which means precisely that $V(\sigma) \rightarrow (0,\infty]^{n+m}$ is a trivial fiber bundle near ∞ . That the projection is smooth and a submersion on each stratum follows because μ_T is actually smooth in T on each stratum. \blacksquare

We summarize the results of this section as follows.

Proposition 4.43. *For a regular perturbation π , and $\text{gr}_z(\alpha,\beta) \leq 10 - \dim \alpha$ as in Corollary 4.13, $\overline{\mathcal{M}}_{E,z,\pi}(\alpha,\beta)$ can be given the natural structure of a compact topological $SO(3)$ -manifold with corners with a smooth structure on each open stratum. We have the following decomposition of the boundary:*

$$\partial \overline{\mathcal{M}}_{E,z,\pi}(\alpha,\beta) = \bigsqcup_{\gamma; z_1 * z_2 = z} \overline{\mathcal{M}}_{E,z_1,\pi}(\alpha,\gamma) \times_\gamma \overline{\mathcal{M}}_{E,z_2,\pi}(\gamma,\beta).$$

The same is true on a cobordism W as long as $\text{gr}_z^W(\alpha,\beta) \leq 10 - \dim \alpha$. In that case, we have the decomposition

$$\begin{aligned} \partial \overline{\mathcal{M}}_{\mathbf{E},z,\pi}^W(\alpha,\beta) &= \bigsqcup_{\gamma \in \mathfrak{C}_{\pi_1}; z_1 * z_2 = z} \overline{\mathcal{M}}_{E_1,z_1,\pi_1}(\alpha,\gamma) \times_\gamma \overline{\mathcal{M}}_{\mathbf{E},z_2,\pi}^W(\gamma,\beta) \\ &\quad \bigsqcup_{\zeta \in \mathfrak{C}_{\pi_2}; z_1 * z_2 = z} \overline{\mathcal{M}}_{\mathbf{E},z_1,\pi}^W(\alpha,\zeta) \times_\zeta \overline{\mathcal{M}}_{E_2,z_2,\pi_2}(\zeta,\beta). \end{aligned}$$

4.8. Families of metrics and perturbations. Let S be a compact smooth manifold with corners. Let (W, \mathbf{E}) be a cylindrical end cobordism $(Y_1, E_1) \rightarrow (Y_2, E_2)$, and suppose each (Y_i, E_i) is equipped with a regular metric and perturbation (g_i, π_i) . A family of metrics and perturbations parameterized by S , written π_S , is the data of:

- a smooth metric on the bundle p_W^*TW on $S \times W$, restricting to the product metric $dt^2 + g_i$ on $\mathbb{R} \times Y_i$ on each end, and
- a smooth map $S \rightarrow \mathcal{P}_c^{(4)}$, the latter being the space of perturbations on W which agree with π_i on the corresponding ends. We usually restrict to a small neighborhood U_C of the perturbation (π_-, π_+) with no interior part, as in Theorem 4.37.

We may then define the parameterized moduli spaces $\widetilde{\mathcal{M}}_{\mathbf{E},z,S}(\alpha, \beta)$ of pairs (s, \mathbf{A}) , where \mathbf{A} is a π_s -perturbed $L^2_{k,\delta}$ instanton going between critical orbits α and β in a fixed component of trajectories z . Fiberwise compactifying by ideal instantons and broken trajectories we obtain $\overline{\mathcal{M}}_{\mathbf{E},z,S}(\alpha, \beta)$; the result is compact by a version of Proposition 4.10 which allows for variations of metric on the interior (which requires no change in the argument).

There is a bundle over $S \times \widetilde{\mathcal{B}}_{\mathbf{E},z,k,\delta}^e(\alpha, \beta)$, written \mathcal{S}_S , whose fiber above (s, \mathbf{A}) is $\Omega_{k-1,\delta}^{2,+s}(W; \mathfrak{g}_E)$; the notation $+, s$ indicates that we are taking the self-dual 2-forms with respect to the metric g_s .

Definition 4.16. *At each perturbed instanton $(s, \mathbf{A}) \in \widetilde{\mathcal{M}}_{\mathbf{E},z,S}(\alpha, \beta)$, taking the derivative of the $SO(3)$ -equivariant map $\sigma : S \times \widetilde{\mathcal{B}}_{\mathbf{E},z,k,\delta}^e(\alpha, \beta) \rightarrow \mathcal{S}$ defined by the instanton equation induces a map on the normal space to each orbit,*

$$(d\sigma)_{\mathbf{A},s} : T_s S \times \ker(d_{\mathbf{A}}^{*,s})_{k,\delta} \rightarrow \Omega_{k-1,\delta}^{2,+s}(W; \mathfrak{g}_E).$$

We say the parameterized moduli spaces $\widetilde{\mathcal{M}}_{\mathbf{E},z,S}(\alpha, \beta)$ are cut out regularly if $(d\sigma)_{\mathbf{A},s}$ is surjective for all perturbed instantons (s, \mathbf{A}) in the homotopy class z . In this case the parameterized moduli space $\widetilde{\mathcal{M}}_{\mathbf{E},z,S}(\alpha, \beta)$ is a smooth $SO(3)$ -manifold of dimension $\text{gr}_z(\alpha, \beta) + \dim S$.

We say the compactified parameterized moduli spaces $\overline{\mathcal{M}}_{\mathbf{E},z,S}(\alpha, \beta)$ are cut out regularly if every moduli space $\widetilde{\mathcal{M}}_{\mathbf{E},w,S}(\gamma, \gamma')$ appearing in the compactification is cut out regularly.

If $\overline{\mathcal{M}}_{\mathbf{E},z,S}(\alpha, \beta)$ is cut out regularly for all (α, β, z) with

$$\text{gr}_z(\alpha, \beta) + \dim S \leq 10 - \dim \alpha,$$

we say that π_S is regular.

Recall Definition 4.14 of weakly admissible cobordism. In the first two cases, no reducibles arise for any perturbation, and still do not arise in families. In the third case, there is a unique reducible of each type for every perturbation, and in particular for the moduli spaces parameterized by a family of perturbations there is a $\dim S$ -dimensional space of reducibles of each type.

In the final case, we have $b^+(W) > 0$ and \mathbf{E} is nontrivial. In this case, for a *generic* perturbation, there are no reducibles; the expected dimension of the space of reducibles is $-b^+(W)$. This means that they *do* arise in generic families when $\dim S \geq b^+(W)$. Something else needs to be asserted to guarantee that in families, reducibles are still cut out transversely. Instead of adding more homological conditions (akin to ρ -monotonicity, but even more strict), we will simply endeavor to avoid the reducibles.

Proposition 4.44. *Suppose (W, \mathbf{E}) is equipped with a metric, a weakly admissible $SO(3)$ -bundle and regular perturbations π_{\pm} on the ends. There is an open set $(\pi_-, 0, \pi_+) \in U_C \subset \mathcal{P}_c^{(4)}$ of the space of perturbations which equal π_{\pm} on the ends. The set U_C is the small neighborhood around $(\pi_-, 0, \pi_+)$ given in Theorem 4.37, in which all fully reducible connections are cut out transversely and the parameterized space of reducibles $U_C \mathcal{M}^{\text{red}}$ is transverse to the space of Fredholm operators with non-trivial cokernel at those reducibles with energy $\leq C$. We may demand that U_C is contractible by passing to a small ball around $(\pi_-, 0, \pi_+)$.*

Suppose $\pi_0, \pi_1 \in U_C$ are regular perturbations.

If either $b^+(W) \neq 1$ or that \mathbf{E} admits no $SO(2)$ -reducible connections of class $L_{k,\delta}^2$ with π_{\pm} -flat limits, then there is a regular family π_t of metrics and perturbations $\pi_t : [0, 1] \rightarrow U_C$, all with the same values on the ends.

If $b^+(W) = 1$ and \mathbf{E} supports $SO(2)$ -reducibles, suppose $\pi_t : [0, 1] \rightarrow U_C$ is a path of metrics and perturbations such that no π_t supports a reducible instanton. Then we may modify the perturbations π_t by an arbitrarily small amount on the interior of $[0, 1]$ so that π_t forms a regular family of metrics and perturbations.

If K is any compact family of metrics on W , constant on the ends, then replacing the open set U_C above with a smaller open set $U_{C,K}$, everything above still applies while allowing us to vary the metric in the family K .

Proof. This follows from the same strategy as in Theorem 4.37. In the first case (that $b^+(W) \neq 1$ or \mathbf{E} supports no reducibles), pick an arbitrary path $\pi_t : [0, 1] \rightarrow U_C$; in the second case we are supplied such a path. By the assumption that $\pi_t \in U_C$, all fully reducible connections are cut out transversely for all t . We have a parameterized moduli space of reducibles, and the map

$$U_C \mathcal{M}^{\text{red}} \rightarrow U_C$$

has index $-b^+(W)$; we have already assumed that π_0, π_1 are regular values, so by a small perturbation we may assume that π_t is transverse to this projection. That is to say that for the family π_t , the reducibles are cut out transversely inside the reducible locus. When $b^+(W) > 1$, this means there are no reducible π -ASD connections; when $b^+(W) = 1$, we already assumed there were no reducible π -ASD connections; when $b^+(W) = 0$, there is a finite set of π_t -ASD connections for each t .

Now as in Theorem 4.37 we have a countable union of submanifolds of even codimension

$$U_C \mathcal{M}^{\text{red}, \geq 1} \subset U_C \mathcal{M}^{\text{red}},$$

restricting to those reducible connections of energy $\leq C$. This is the space of reducible connections supporting a splitting $\mathbf{E} \cong \mathbb{R} \oplus \eta$ so that the normal ASD operator

$$D_{\mathbf{A}, \pi}^{\text{irred}} : \Omega_{k,\delta}^1(W; \eta) \rightarrow \Omega_{k-1,\delta}^{2,+}(W; \eta)$$

has non-trivial cokernel. The codimension of each of these is even because the normal space is isomorphic to $\text{Hom}_{\mathbb{C}}(\mathbb{C}^{i+j}, \mathbb{C}^i)$ for $i \geq 0, j > 0$, which is a non-trivial complex vector space; the fact that the index i is positive for each reducible \mathbf{A} is part of the assumption that W is weakly admissible.

By assumption, π_0 and π_1 are regular values of the projection

$$U_C \mathcal{M}^{\text{red}, \geq 1} \rightarrow U_C,$$

so by a small perturbation we may assume that π_t is transverse to this projection as well. Because the codimension of each manifold in $U_C \mathcal{M}^{\text{red}, \geq 1}$ is at least 2, this means that for any reducible π_t -instanton \mathbf{A} , the normal ASD operator is surjective.

That is to say, for the path π_t , any reducible π_t -instanton \mathbf{A} of energy $\leq C$ is cut out transversely, and not just internally to its own locus.

Now we conclude by playing the same game with the irreducible connections: we have a parameterized moduli space $U_C \mathcal{M}^{\text{irred}}$, and we demand that the path π_t is transverse to the projection

$$U_C \mathcal{M}^{\text{irred}} \rightarrow U_C,$$

which we can make true by a small perturbation on the interior of $[0, 1]$.

The set U_C depends on the chosen metric on W , but as in the statement of the proposition, we may vary over a compact set of metrics by passing to a smaller open set $U_{C,K}$. \blacksquare

Remark 4.2. Except in the case where \mathbf{E} admits no reducible connections whatsoever, there are obstructions to extending the above theorem to arbitrary families of metrics and perturbations, even when $b^+ = 0$. Generically, one expects the set of reducibles for which the operator $Q_{\mathbf{A},\pi}^{\text{irred}}$ has nontrivial cokernel to form a codimension $I + 2$ family, where I is the index of this operator. The definition of regular family of metrics and perturbations does nothing to help this: the map we want to be surjective,

$$\widehat{\nabla}_{s'(\pi)} \mathbf{A} : T_s S \rightarrow \text{coker}(Q^{\text{irred}}),$$

is an *equivariant* map. The former space has trivial circle action, and the latter space has circle action of weight 1: the only equivariant map is zero. So if our family of metrics and perturbations supports a reducible which is not cut out transversely in the usual sense, it won't be cut out transversely in the family sense, either.

When \mathbf{E} admits no reducible connections, there is no difficulty applying the usual ideas: there is a smooth map $\pi : \mathcal{P}\mathcal{M} \rightarrow \mathcal{P}$, and we choose a map $S \rightarrow \mathcal{P}$ transverse to π extending a given map $\partial S \rightarrow \mathcal{P}$ transverse to π .

Remark 4.3. Suppose we are given metrics g_0, g_1 on W , which give open sets $U_{C,i}$ depending on the metric; and suppose we are given regular $\pi_i \in U_{C,i}$. Choosing a path g_t between these metrics, we see that we may choose a generic path between the two perturbations, so long as they are chosen in a smaller set $U_{C,g_t} \subset U_{C,0} \cap U_{C,1}$.

To actually construct a generic path between the perturbations, then, first choose a perturbation $\pi'_i \in U_{C,g_t}$, and choose a generic path from π_i to π'_i inside of $U_{C,i}$. Then choose a generic path from π'_0 to π'_1 inside of U_{C,g_t} .

Assuming S is a regular family of metrics and perturbations, we have smooth moduli spaces $\widetilde{\mathcal{M}}_{\mathbf{E},z,S}(\alpha, \beta)$ with smooth projection maps $\widetilde{\mathcal{M}}_{\mathbf{E},z,S} \rightarrow S$. The space $\widetilde{\mathcal{M}}_{\mathbf{E},z,S}$ is not compact but the fiberwise Uhlenbeck compactification $\overline{\mathcal{M}}_{\mathbf{E},z,S}$ is. As long as

$$\text{gr}_z^W(\alpha, \beta) + \dim \alpha + \dim S \leq 10,$$

this being at least the dimension of $\widetilde{\mathcal{M}}_S/SO(3)$, there is no Uhlenbeck bubbling arising in this compactification. Note that here it is possible for $\widetilde{\mathcal{M}}_{\mathbf{E},z,S}$ to be nonempty even if $\overline{\text{gr}}_z(\alpha, \beta)$ is negative. In this case, the set of $s \in S$ so that π_s is nonempty is a compact submanifold with corners of codimension $-\overline{\text{gr}}_z^W(\alpha, \beta)$.

We then have the following analogue of Theorem 4.42.

Proposition 4.45. *If π_S is a regular family of metrics and perturbations indexed by S , and $\text{gr}_z^W(\alpha, \beta) + \dim S \leq 10 - \dim \alpha$, then $\overline{\mathcal{M}}_{\mathbf{E},z,S}$ has the natural structure of a compact topological $SO(3)$ -manifold with corners and a smooth structure on each stratum. Its boundary has the decomposition*

$$\begin{aligned} \partial \overline{\mathcal{M}}_{\mathbf{E},z,S}(\alpha, \beta) = & \bigcup_{\substack{\gamma \in \mathfrak{C}_{Y_1} \\ z_1 * z_2 = z}} \overline{\mathcal{M}}_{E_1, z_1, \pi}(\alpha, \gamma) \times_{\gamma} \overline{\mathcal{M}}_{\mathbf{E}, z_2, S}(\gamma, \beta) \\ & \bigcup_{\substack{\zeta \in \mathfrak{C}_{Y_2} \\ z_1 * z_2 = z}} \overline{\mathcal{M}}_{\mathbf{E}, z_1, S}(\alpha, \zeta) \times_{\zeta} \overline{\mathcal{M}}_{E_2, z_2, \pi}(\zeta, \beta) \cup \overline{\mathcal{M}}_{\mathbf{E}, z, \partial S}(\alpha, \beta). \end{aligned}$$

In particular, suppose $S = [0, 1]$, and write W as the composite of cobordisms with cylindrical ends W_1 and W_2 , obtained by splitting together the positive end of W_1 with the negative end of W_2 . Write the broken perturbation $\pi(1)$ as π_i on each of the W_i , where π_- is the perturbation on the negative end of W_1 and π_+ is the perturbation on the positive end of W_2 ; write the common perturbation in the center as π_b . So long as $\text{gr}_z^W(\alpha, \beta) \leq 9 - \dim \alpha$, we have a decomposition

$$\begin{aligned} \partial \overline{\mathcal{M}}_{\mathbf{E}, S, z}(\alpha, \beta) = & \bigcup_{\substack{\gamma \in \mathfrak{C}_{\pi_-} \\ z_1 * z_2 = z}} \overline{\mathcal{M}}_{E_-, z_1, \pi_-}(\alpha, \gamma) \times_{\gamma} \overline{\mathcal{M}}_{\mathbf{E}, z_2, S}(\gamma, \beta) \\ & \bigcup_{\substack{\eta \in \mathfrak{C}_{\pi_b} \\ z_1 * z_2 = z}} \overline{\mathcal{M}}_{\mathbf{E}_1, z_1, \pi_1}(\alpha, \eta) \times_{\eta} \overline{\mathcal{M}}_{\mathbf{E}_2, z_2, \pi_2}(\eta, \beta) \\ & \bigcup_{\zeta \in \mathfrak{C}_{\pi_+}; z_1 * z_2 = z} \overline{\mathcal{M}}_{\mathbf{E}, z_1, S}(\alpha, \zeta) \times_{\zeta} \overline{\mathcal{M}}_{E_+, z_2, \pi_+}(\zeta, \beta) \\ & \cup \overline{\mathcal{M}}_{\mathbf{E}, z, \pi(0)}(\alpha, \beta). \end{aligned}$$

Proof. Fix an open stratum $\sigma \subset \overline{\mathcal{M}}_{\mathbf{E}, z, S}(\alpha, \beta)$ corresponding to n breakings along Y_1 and m breakings along Y_2 . The result will follow essentially as in Theorem 4.42. There we described a gluing map

$$\text{ev}_T : (0, \infty]^{n+m} \times \mathcal{N} \rightarrow \mathcal{R}$$

between Hilbert manifolds. The manifold \mathcal{N} consisted of products of spaces parameterizing n long (length L) broken pieces of framed instantons on $[-L, L] \times Y_1$, and similarly m on (Y_2, E_2) , a space parameterizing instantons on cylinders $[-T, T] \times Y_i$ for all $T_0 \leq T \leq \infty$ sufficiently large, and a space parameterizing instantons on the compact manifold W^L (which is W with the cylindrical ends $[0, \infty) \times Y$ truncated to $[0, L] \times Y$). The codomain \mathcal{R} consisted of restrictions of these to $2n$ copies of the configuration space of framed connections on Y_1 and $2m$ copies of the configuration space of framed connections on Y_2 . This is smooth on each stratum of $(0, \infty]^{n+m}$, and the individual maps ev_T converge in C_{loc}^{∞} as

$$T_k \rightarrow T \in (0, \infty]^{n+m}.$$

The moduli space of instantons on W , compactified by broken trajectories, is identified with $\text{ev}^{-1}(\Delta)$, the subset of instantons on each piece so that the restrictions to corresponding boundary pieces agree.

When we allow the metric and perturbation to vary in some family S , this variation occurs in a compact part of the manifold with cylindrical ends \hat{W} . As long as L above is taken large enough, this variation only takes place in W^L , and hence only affects that term of the product; we may replace the term

$$\widetilde{\mathcal{M}}_{\mathbf{E}, \pi, k, \delta}(\hat{W}) \subset \widetilde{V}_{\mathbf{E}, \pi}^L \subset \widetilde{\mathcal{M}}_{\mathbf{E}, \pi}(W^L)$$

instead with a parameterized version $\widetilde{V}_{\mathbf{E}, S}^L$ of instantons on W^L for which

$$\begin{aligned} \Omega^1(W^L) \oplus T_s S & \rightarrow \Omega^{2,+}(W^L) \\ (\omega, \pi') & \mapsto D_{\mathbf{A}, \pi} + \widehat{\nabla}_{\pi'}(\mathbf{A}) \end{aligned}$$

is surjective.

The modified evaluation map above is written

$$(0, \infty]^{n+m} \times \mathcal{N}_S \rightarrow \mathcal{R},$$

and is again smooth on each stratum of $(0, \infty]^{n+m}$, with convergent sequences of $T_k \in (0, \infty]^{n+m}$ giving sequences of maps ev_{T_k} which converge in C_{loc}^∞ . Again we have that $\overline{\mathcal{M}}_{\mathbf{E}, z, S} = \text{ev}^{-1}(\Delta)$.

Therefore we may apply [KM07, Lemma 19.3.3], which asserts that

$$\text{ev}^{-1}(\Delta) \rightarrow (0, \infty]^{n+m}$$

is a topological submersion, meaning that it is a topologically trivial fiber bundle near ∞ . In particular, this gives a chart for a neighborhood of $\sigma \subset \overline{\mathcal{M}}_{S, z}$, an open embedding $(0, \infty]^{n+m} \times \sigma \rightarrow \overline{\mathcal{M}}_{S, z}$ that is the identity on σ , is stratum-preserving, and is a diffeomorphism on each open stratum. \blacksquare

We will need a version of these results for more general families of metrics and perturbations termed ‘families of broken metrics’ (see [Dae15] and [KM11a], where this notion was used to study spectral sequences from Khovanov homology).

Suppose are given a 4-manifold W with cylindrical ends $(-\infty, 0] \times Y_1$ and $[0, \infty) \times Y_2$, and a sequence of Riemannian 3-manifolds M_1, \dots, M_ℓ , equipped with an isometric embedding

$$\sqcup \varphi_i : (-1 - \epsilon, 1 + \epsilon) \times M_i \rightarrow W.$$

Each interval $[-1, 1] \times M_i$ will be equipped with a perturbation $\widehat{\nabla}_{\pi_i}$ which restricts to some neighborhood of $(-c_i, c_i) \times M_i$ as $(dt \wedge \nabla_{\pi_i})^+$ for some regular perturbation ∇_{π_i} on the pair (M_i, \mathbf{E}_i) (that is, so that the critical orbits are cut out nondegenerately and all moduli spaces up to some sufficiently large dimension are also cut out nondegenerately). We say that the perturbation is adapted to the cut.

To define the restriction map for the framed instanton moduli spaces, we need to assume that the path $\gamma : \mathbb{R} \rightarrow W$ is adapted to the cuts, in the sense that there is a sequence t_1, \dots, t_ℓ of real numbers with $t_i + 2 < t_{i+1}$ so that for the interval $t \in [t_i - 1, t_i + 1]$, we have $\gamma(t_i + t) = \varphi_i(t - t_i, x_i)$ for some fixed $x_i \in M_i$. In particular, γ passes through the M_i in the order listed.

The canonical family of broken metrics on W is parameterized by $[0, \infty]^\ell$; for a fixed element (T_1, \dots, T_ℓ) , one replaces the isometric copies of $[-1, 1] \times M_i$ with isometric copies of

$$\begin{aligned} &[-1 - T_i, 1 + T_i] \times M_i \text{ if } T_i < \infty, \text{ or} \\ &([-1, \infty) \sqcup (-\infty, 1]) \times M_i \text{ if } T_i = \infty, \end{aligned}$$

and replaces the image of γ appropriately. The perturbation $\widehat{\nabla}_{\pi_i}$ associated to (T_1, \dots, T_ℓ) is defined by assuming that its value on the interval $(-c_i - T_i, c_i + T_i)$ is equal to $(dt \wedge \nabla_{\pi_i})^+$ as above; this remains true in the infinite case. We say that the metric parameterized by (∞, \dots, ∞) is an ℓ -times broken metric on W , with cuts along M_1, \dots, M_ℓ .

A family of broken metrics and perturbations is given by a compact smooth manifold with corners S equipped with a family of metrics and perturbations on W , possibly containing broken metrics and perturbations. If $s \in S$ corresponds to one of these broken metrics, cut along M_1, \dots, M_ℓ , then we demand there is a chart $\varphi : (T, \infty]^\ell \times U \rightarrow S$, where U is some topological manifold with corners and φ is a stratum-preserving open embedding, so that $\varphi(T_1, \dots, T_\ell, x)$ is the element of

the canonical family of broken metrics and perturbations associated to (T_1, \dots, T_ℓ) described above. The variance in metric and perturbation parameterized by $x \in U$ occurs in the complement of the canonical intervals $I \times M_i$. The union over the images of $\varphi(U)$ in S , as the φ vary over charts for neighborhoods of a specific sequence of cuts M_1, \dots, M_ℓ , is called the *cut stratum* for M_1, \dots, M_ℓ .

If S parameterizes a family of broken metrics and perturbations, we define the parameterized moduli space $\widetilde{\mathcal{M}}_S$ as follows. For a fixed broken metric and perturbation π , cut along M_1, \dots, M_ℓ , we are left with a sequence W_0, \dots, W_ℓ of manifolds with two cylindrical ends; other than the outer two, we have W_i a manifold with incoming cylindrical end $(-\infty, 0] \times M_i$ and outgoing cylindrical end $[0, \infty) \times M_{i+1}$. Each of these has a corresponding collection of moduli spaces

$$\bigsqcup_{\substack{\alpha_i, \alpha_{i+1} \in \mathfrak{C}_{\pi_i} \\ z_i \in \pi_1(\widetilde{\mathfrak{B}}_E^e(\alpha_i, \alpha_{i+1}))}} \widetilde{\mathcal{M}}_{\mathbf{E}, z_i, \pi}^{W_i}(\alpha_i, \alpha_{i+1});$$

we say that an element of $\widetilde{\mathcal{M}}_{\mathbf{E}, z, \pi}^W$ is a sequence of framed instantons

$$\mathbf{A} \in \widetilde{\mathcal{M}}_{\mathbf{E}, z_i, \pi}^{W_i}(\alpha_i, \alpha_{i+1})$$

so that the concatenation $z_1 * \dots * z_\ell$ is z . That is, writing the space given by the disjoint union of π_i -critical orbits on M_i as $\widetilde{R}_{\pi_i}(M_i)$, we have

$$\widetilde{\mathcal{M}}_{\mathbf{E}, \pi}^W := \widetilde{\mathcal{M}}_{\mathbf{E}, \pi}^{W_0} \times_{\widetilde{R}_{\pi_1}(M_1)} \dots \times_{\widetilde{R}_{\pi_\ell}(M_\ell)} \widetilde{\mathcal{M}}_{\mathbf{E}, \pi}^{W_\ell},$$

which then may be decomposed as a disjoint union according to the paths (z_1, \dots, z_ℓ) .

The moduli space of the family is then defined to be

$$\widetilde{\mathcal{M}}_{\mathbf{E}, z, S}(\alpha_0, \alpha_n) = \cup_{s \in S} \widetilde{\mathcal{M}}_{\mathbf{E}, z, \pi_s}(\alpha_0, \alpha_n),$$

topologized so that (s_n, \mathbf{A}_n) converges to (s, \mathbf{A}) if $s_n \rightarrow s$ and $\mathbf{A}_n \rightarrow \mathbf{A}$ in L_k^2 on compact sets.

We say that the individual moduli space $\widetilde{\mathcal{M}}_{\mathbf{E}, z, \pi}(\alpha_0, \alpha_n)$ is *regular* if each component moduli space $\widetilde{\mathcal{M}}_{\mathbf{E}, z_i, \pi}(\alpha_i, \alpha_{i+1})$ is cut out regularly in dimensions at most 10.

We say that a family of broken metrics and perturbations parameterized by S is *regular* if, for each cut stratum σ of S parameterizing ℓ -broken metrics with cuts along M_1, \dots, M_ℓ , but varies metric and perturbation elsewhere on W , we have that each of the component moduli spaces $\widetilde{\mathcal{M}}_{\mathbf{E}, z_i, \sigma}(\alpha_i, \alpha_{i+1})$ is regular. For $\ell = 0$, this just says that the open submanifold of S parameterizing unbroken metrics should be regular in the sense of Definition 4.16.

One may then define the fiberwise Uhlenbeck compactification $\overline{\mathcal{M}}_S$ as before, by setting $\overline{\mathcal{M}}_s$ to be the fiber products

$$\overline{\mathcal{M}}_{\mathbf{E}_0, \pi_0} \times_{\widetilde{R}_{\pi_1}(M_1)} \dots \times_{\widetilde{R}_{\pi_\ell}(M_\ell)} \overline{\mathcal{M}}_{\mathbf{E}_\ell, \pi_\ell}$$

of Uhlenbeck compactifications (by ideal instantons and broken trajectories); $\overline{\mathcal{M}}_{\mathbf{E}, \pi, S}$ is the union of these over $s \in S$, topologized as above. One easily verifies that $\overline{\mathcal{M}}_S$ is compact by the same argument as the case of families of unbroken metrics; now if $s_n \rightarrow s$ is a sequence that limits to a metric and perturbation with more breakings, energy slides off the newly-infinite ends. If $\overline{\mathcal{M}}_S$ is regular and of dimension at most 10, then this compactification only includes broken trajectories, not ideal instantons.

On the composite cobordism W , we will want to consider the composite perturbation, but this involves a term along the neck which doesn't come from what we call 'interior perturbations'. We will need to dampen this perturbation out if we hope to compare the moduli spaces on the composite, with generic perturbation, to those of the individual manifolds W_i .

We resolve this by considering a 2-parameter family of perturbations. Suppose S is a family of broken metrics and perturbations; let $(\pi, t) \in S \times [0, 1]$ denote the perturbation which is the same as π , but instead of being $\beta_0(s)\widehat{\nabla}_{\pi_i}(\mathbf{A})$ on the neck, we instead use $t\beta_0(s)\widehat{\nabla}_{\pi_i}(\mathbf{A})$; for $t = 1$, this is what we already had, and for $t = 0$, this is identically zero.

To see why this is necessary or desirable, recall where the open set U_C of Theorem 4.37 comes from: it is a set of perturbations π so that the map

$$D^{\text{irred}} : U_C \mathcal{M}^{\text{red}} \rightarrow U_C \mathcal{F}$$

sending a reducible π -instanton to its normal ASD operator, is transverse to those operators with non-trivial cokernel. This open set is guaranteed because we see that D^{irred} is transverse to $U_C \mathcal{F}$ along the perturbations $(\pi_-, 0, \pi_+)$ with no part coming from 'interior holonomy perturbations'.

The point is the following. Suppose W is a manifold with two cylindrical ends, arising from gluing the W_i together. Suppose further that $\pi_-, \pi_1, \dots, \pi_+$ is a sequence of regular perturbations on the 'cutting' 3-manifolds, and we consider the perturbation $(\pi_-, 0, t\pi_1, 0, \dots, t\pi_n, 0, \pi_+)$ on the composite, where we dampen the perturbation on the necks as above. Then the same argument that shows D^{irred} is transverse along $(\pi_-, 0, \pi_+)$ shows that D^{irred} is transverse just as well along $(\pi_-, 0, \dots, 0, \pi_+)$ - this is guaranteed so long as there is no 'interior holonomy perturbation' component.

Given this, suppose we have a family of broken metrics and perturbations S ; we will demand that the perturbation on the necks takes a 'dampened' form such as the above. Then if $\mathcal{P}_{c,i}^{(4)}$ are the spaces of interior perturbations on W_i , there is an open set $U_{C,S,i} \subset \mathcal{P}_{c,i}^{(4)}$ so that D^{irred} is transverse to the operators with non-trivial cokernel along any

$$(\pi_-, \pi_0^{\text{int}}, \pi_1, \dots, \pi_n^{\text{int}}, \pi_+)$$

with

$$\pi_i^{\text{int}} \in U_{C,S,i}$$

for all i .

The reason that this is an open condition is the following. Suppose \mathbf{A}_n is a sequence of instantons on W^L , with gluing factors L finite but going to infinity, with \mathbf{A}_n converging to a broken instanton \mathbf{A} . If all of the components of \mathbf{A} are regular, then so is \mathbf{A}_n , for large n .

Proposition 4.46. *Let S parameterize a family of broken metrics, and $\pi : S \rightarrow \prod_i U_{C,S,i}$ parameterize a family of perturbations on (W, \mathbf{E}) so that ∂S is a regular family.*

Then the family of broken metrics on S only consists of unbroken metrics and perturbations on the interior of S , and one may modify S by an arbitrarily small amount on the interior to make it into a regular family. In particular, given a broken metric and perturbation $\pi_0 \in \prod_i U_{C,S,i}$, there is a regular path π_t of metrics and perturbations, with π_t unbroken for $t > 0$, and so that the perturbation π_1 a

regular perturbation on the composite cobordism W with zero perturbation on the ‘neck’.

Now suppose that S parameterizes a regular family of broken metrics and perturbations on (W, \mathbf{E}) . Let $\sigma \subset S$ be a cut stratum of codimension ℓ , so parameterizing metrics with a sequence of cuts along M_1, \dots, M_ℓ . For any stratum $\tilde{\sigma} \subset \overline{\mathcal{M}}_{\mathbf{E}, \sigma, z}(\alpha, \beta)$ of codimension $n + k_1 + \dots + k_\ell + m$, parameterizing broken s -perturbed instantons for $s \in \sigma$ with n breakings along the incoming end Y_1 , with k_i breakings along the internal end M_i , and with m breakings along the outgoing end Y_2 . Then there is a neighborhood

$$V(\sigma) \subset \overline{\mathcal{M}}_{\mathbf{E}, z, S}(\alpha, \beta)$$

and an $SO(3)$ -invariant map $V(\sigma) \rightarrow (0, \infty]^{n+k_1+\dots+k_\ell+m}$ which is stratum-preserving, a submersion on each stratum, and a topologically trivial fiber bundle for the inverse image of a neighborhood of ∞ . In particular, $\overline{\mathcal{M}}_{\mathbf{E}, S, z}(\alpha, \beta)$ is a compact topological manifold with corners and a smooth structure on each stratum, with the same decomposition as in Proposition 4.45.

Proof. The transversality claim follows without real change as in Proposition 4.44. By the discussion before this proposition, if we have a broken perturbation π_0 , we may choose a path π_t so that π_1 is of the form (π_-, π, π_+) , where the compactly supported perturbation π is an interior holonomy perturbation — it has nothing corresponding to the holonomy perturbation on the neck.

The rest follows the same lines as the proof of the gluing theorem: first, we may use Lemma 4.39 to build charts $\mu_T^i : B_{\eta_i}(\tilde{H}_i) \rightarrow \widetilde{\mathcal{M}}_{\alpha_i, k, \delta}(Z_i^T)$ near the constant solution γ_{α_i} , where $Z_i^T = [-T, T] \times M_i$ for finite T and

$$([0, \infty) \cup (-\infty, 0]) \times M_0$$

for $T = \infty$. For uniformity of notation, set $M_0 = Y_1$ and $M_{\ell+1} = Y_2$. Then as in the proof of Lemma 4.41, for sufficiently large fixed L may assemble a map with domain the product

$$\begin{aligned} & (0, \infty]^{n+k_1+\dots+k_\ell+m} \times \widetilde{\mathcal{M}}_{E_1, k, \delta}^{(-\infty \times Y_1)}(\alpha_0) \times^n B_{\eta_0}(\tilde{H}_0) \times^n \widetilde{\mathcal{M}}_{\text{cen}}(Z_0^L) \\ & \quad \times \widetilde{V}_{\sigma, \mathbf{E}, k}^{W'_0} \times^{k_1} B_{\eta_1}(\tilde{H}_1) \times^{k_1} \widetilde{\mathcal{M}}_{\text{cen}}(Z_1^L) \\ & \quad \dots \\ & \quad \times {}_\sigma \widetilde{V}_{\sigma, \mathbf{E}, k}^{W'_\ell} \times^{k_\ell} B_{\eta_\ell}(\tilde{H}_\ell) \times^{k_\ell} \widetilde{\mathcal{M}}_{\text{cen}}(Z_\ell^L) \\ & \quad \times^m B_{\eta_{\ell+1}}(\tilde{H}_{\ell+1}) \times^m \widetilde{\mathcal{M}}_{\text{cen}}(Z_{\ell+1}^L) \times \widetilde{\mathcal{M}}_{E_2, k, \delta}^{[0, \infty) \times Y_2}(\beta). \end{aligned}$$

Recall that $\widetilde{V}_{\mathbf{E}, k}^W \subset \widetilde{\mathcal{M}}_{\mathbf{E}, k}^{W'}$ is an open neighborhood of those instantons which extend to the whole of W .

Here the \times_σ indicates that the moduli spaces

$$\widetilde{V}^{W'_0} \times \widetilde{V}^{W'_1} \times \dots \times \widetilde{V}^{W'_\ell}$$

should then be considered as those moduli spaces altogether parameterized by σ , but not that we parameterize each by S separately and take the product. Secondly, here we mean by W_i the compact manifolds given by truncating the cylindrical ends of W'_i , written $\text{End}(W_i) = [0, \infty) \times \partial W'_i$, at some large finite $[0, L]$.

The codomain of the map is restriction to the boundary of each piece (this uses parallel transport along the path γ):

$$\left(\times^n \tilde{\mathcal{B}}_{Y_1, k-1/2}^e \times^{k_1} \tilde{\mathcal{B}}_{M_1, k-1/2}^e \times \cdots \times^{k_\ell} \tilde{\mathcal{B}}_{M_\ell, k-1/2}^e \times^m \tilde{\mathcal{B}}_{Y_2, k-1/2}^e \right)^2.$$

Call the domain $(0, \infty]^{n+k_1+\cdots+k_\ell+m} \times \mathcal{N}$ and the codomain \mathcal{R} . Mapping to an element of the diagonal in the codomain means that the corresponding instantons on each piece may be glued to a (possibly broken) instanton on the sequence of manifolds with cylindrical ends W_0, \dots, W_ℓ . The assumption that ∂S is regular means, in particular, that σ is a regular family of metrics and perturbations, and so the map $\text{ev}_\infty : \mathcal{N} \rightarrow \mathcal{R}$ is transverse to the diagonal. By convergence in C_{loc}^∞ , the same is true for other nearby ev_T , and in particular the families (T, σ) are regular families of metrics and perturbations for T sufficiently large. This is the first part of the proposition.

Now [KM07, Lemma 19.3.3] gives the second part: the fact that this map

$$(0, \infty]^{n+k_1+\cdots+k_\ell+m} \times \mathcal{N} \rightarrow \mathcal{R}$$

is transverse to the diagonal at ∞ means that the projection

$$\text{ev}^{-1}(\Delta) \rightarrow (0, \infty]^{n+k_1+\cdots+k_\ell+m}$$

is a topological submersion at ∞ , which is smooth on each stratum, and hence we may obtain a corresponding chart for a neighborhood of $\text{ev}_\infty^{-1}(\Delta) \subset \text{ev}^{-1}(\Delta)$, which is smooth on each stratum. This space $\text{ev}_\infty^{-1}(\Delta)$ contains a neighborhood $V(\sigma)$ of the stratum of \mathcal{M}_S described in the statement, and so we may conclude. \blacksquare

5. ORIENTATIONS

5.1. Orientability and determinant line bundles. We begin by briefly recalling some generalities on determinant lines. Let H_1 and H_2 be Hilbert spaces, and let $\mathcal{F}(H_1, H_2)$ be the space of Fredholm maps. We would like to claim that there is a natural line bundle over $\mathcal{F}(H_1, H_2)$, called the *determinant line bundle*, whose fiber over T is isomorphic to

$$\Lambda^*(\ker(T)) \otimes \Lambda^*(\text{coker}(T))^*;$$

we intuit this as being the determinant line of the virtual vector space $\ker(T) - \text{coker}(T)$.

If one attempts to do this naively (with literally those fibers), the result does not have continuous transition functions due to jumps in the kernel. Instead, given any Fredholm operator T_0 , choose a finite-dimensional subspace $J \subset H_2$ so that

$$T_0 \oplus \text{Id}_J : H_1 \oplus J \rightarrow H_2$$

is surjective, and instead define the determinant line bundle to be

$$\Lambda^* \ker(T \oplus \text{Id}_J) \otimes \Lambda^*(J^*)$$

for T near T_0 . Because $T + \text{Id}_J$ is surjective for T close enough to T_0 , this gives a well-defined line bundle over this chart; one should argue that the line bundles are identical for different choices of J , and that doing this on an open cover defines a legitimate line bundle on $\mathcal{F}(H_1, H_2)$.

The idea of the determinant line bundle originates in [Qui85] in the context of Riemann surfaces; the operators $\bar{\partial}_J$ define a family of Fredholm operators over the moduli space of all Riemann surfaces. Our application is similar, as with most applications of this construction. Also see [KM07, Section 20.2], in which the

authors describe the construction of the determinant line bundle, and explain how to construct a canonical associative isomorphism

$$q : \det(\mathcal{F}(H_1, H_2)) \otimes \det(\mathcal{F}(K_1, K_2)) \rightarrow \det(\mathcal{F}(H_1 \oplus K_1, H_2 \oplus K_2)),$$

its definition involves signs that depend on the choice of J in a local description as above. All in all, we have smooth real line bundles over each $\mathcal{F}(H_1, H_2)$ which have canonical direct-sum isomorphisms.

Now let (W, \mathbf{E}) be a Riemannian 4-manifold equipped with weakly admissible $SO(3)$ -bundle and regular perturbation π , with two cylindrical ends: one incoming (Y_1, E_1, π_1) and one outgoing (Y_2, E_2, π_2) . If α is a π_1 -flat connection on Y_1 , and similarly β is π_2 -flat, then the space of connections $\mathcal{A}_{\mathbf{E}, k, \delta}(\alpha, \beta)$ admits a smooth map

$$\mathcal{A}_{\mathbf{E}, k, \delta}(\alpha, \beta) \rightarrow \mathcal{F}\left(\Omega_{k, \delta}^1, \Omega_{k-1, \delta}^{2,+} \oplus \Omega_{k-1, \delta}^0\right),$$

sending a connection \mathbf{A} to the perturbed ASD operator

$$Q_{\mathbf{A}, \pi} = (d_{\mathbf{A}}^+ + D_{\mathbf{A}} \widehat{\nabla}_{\pi}, d_{\mathbf{A}}^*);$$

pulling back the determinant line bundle on the space of Fredholm operators, we find that we have a line bundle over $\mathcal{A}_{\mathbf{E}, k, \delta}^{(4)}$.

Furthermore, because $Q_{\mathbf{A}}$ is invariant under the gauge group, we see that Q descends to an $SO(3)$ -invariant map $\widetilde{\mathcal{B}}_{\mathbf{E}, k, \delta}^e \rightarrow \mathcal{F}$, and hence that there is an $SO(3)$ -equivariant line bundle $\det(Q)$ over this space.

Lemma 5.1. *The line bundle $\det(Q)$ over $\widetilde{\mathcal{B}}_{\mathbf{E}, k, \delta}^e$ is trivializable.*

Proof. For convenience of notation we drop the super/subscripts on \mathcal{B} ; everything here is modulo even gauge.

We begin with Donaldson's stabilization trick, adapted for $SO(3)$ -bundles (as opposed to $SU(2)$ -bundles). There is a natural map $\mathcal{A}_{\mathbf{E}, k, \delta} \rightarrow \mathcal{A}_{\mathbf{E} \oplus \mathbb{C}, k, \delta}$ given by taking the direct sum with the trivial connection; this is equivariant for the even gauge group (where for an $SO(n)$ -bundle, 'even' means that the gauge transformation lifts to a section of $\text{Aut}(E) \times_{SO(n)} \text{Spin}(n)$, the action by conjugation). Therefore, this map descends to a map on the quotients $\widetilde{\mathcal{B}}_{\mathbf{E}}^e \rightarrow \widetilde{\mathcal{B}}_{\mathbf{E} \oplus \mathbb{C}}^e$. Write $\det(Q')$ for the determinant line bundle for connections over $\mathbf{E} \oplus \mathbb{C}$; for a connection of the form $\mathbf{A} \oplus \theta$, the operator Q' splits as a direct sum of the ASD deformation complex for the three bundles \mathfrak{g}_E , $\mathbb{C}^* \otimes E$, and \mathbb{R} , where on the first term we use the connection induced by \mathbf{A} , on the second term the connection $\theta \otimes \mathbf{A}$, and on the final term the trivial connection. Note that the second term is a *complex linear* operator, and hence has canonically trivial determinant.

Thus this stabilization map pulls back $\det(Q')$ to $\det(Q) \oplus \det(Q_{\theta})$, where the final term is a trivial line bundle. In particular, if $\det(Q')$ is trivializable, then $\det(Q)$ is as well. Write $\mathcal{B}_{(\ell)}$ for the configuration space of connections on $\mathbf{E} \oplus \mathbb{C}^{\ell}$, modulo even gauge, and write $\widetilde{\mathcal{B}}_{(\ell)}$ for the smooth $SO(2\ell + 3)$ -manifold of framed configurations on $\mathbf{E} \oplus \mathbb{C}^{\ell}$, whose quotient is $\widetilde{\mathcal{B}}_{(\ell)}$.

Next, because the reducible subspace is of infinite codimension, transversality implies that the inclusion of the irreducible subspace $\widetilde{\mathcal{B}}_{(\ell)}^* \hookrightarrow \widetilde{\mathcal{B}}_{(\ell)}$ is a weak homotopy equivalence. In particular the restriction map

$$[\widetilde{\mathcal{B}}_{(\ell)}, \mathbb{R}\mathbb{P}^{\infty}] \rightarrow [\widetilde{\mathcal{B}}_{(\ell)}^*, \mathbb{R}\mathbb{P}^{\infty}]$$

is a bijection, so any line bundle which is trivializable over $\tilde{\mathcal{B}}_{(\ell)}^*$ is necessarily globally trivializable. Now our bundle is crucially an $SO(3)$ -equivariant line bundle over $\tilde{\mathcal{B}}_{(\ell)}^*$, so it is the pullback of some line bundle over $\mathcal{B}_{(\ell)}^*$, the space of irreducible connections modulo gauge (with no framing). Our goal now is to show that this line bundle is trivial, which will start from a calculation of $\pi_1 \mathcal{B}_{(\ell)}^*$. By definition, this space is the quotient of \mathcal{A}^* by the even gauge group $\mathcal{G}_{\mathbf{E} \oplus \mathbf{C}^\ell}^{e,h}$; because \mathcal{A}^* is again the complement of a union of infinite codimension submanifolds, we see that it too is contractible, and so $\pi_1 \mathcal{B}_{(\ell)}^* = \pi_0 \mathcal{G}_{\mathbf{E} \oplus \mathbf{C}^\ell}^{e,h}$. Next, we must calculate this group of components. This is where stabilization is first useful.

Fix a principal G -bundle P over a CW complex Y . It is not difficult to show that the pointed mapping space $\text{Map}_*^P(Y, BG)$ gives a model for $B\mathcal{G}(P)$, where the superscript indicates we are only interested in the component which classifies the bundle P . Now, because BSO is an H -space (with product structure given by direct sum of vector bundles), in fact every component of $\text{Map}_*(Y, BSO)$ is homotopy equivalent to any another. For n sufficiently large,

$$\pi_i \text{Map}_*^P(Y, BSO) \cong \text{Map}_*^P(Y, BSO(n))$$

(this is true for $i = 1$ already as soon as $n = 7$). To compute the fundamental group of the quotient by all gauge transformations, it suffices to compute the fundamental group of the identity component of $\text{Map}_*(Y, BSO)$. As before, because we are only interested in the quotient by even gauge transformations, we actually need only compute $\pi_1 \text{Map}_*(Y, B\text{Spin})$. Because there is a fibration

$$F \rightarrow B\text{Spin} \rightarrow K(\mathbb{Z}, 4)$$

with F 6-connected, we see that

$$\pi_1 \text{Map}_*(Y, B\text{Spin}) \cong \pi_1 \text{Map}_*(Y, K(\mathbb{Z}, 4)).$$

Now, it is a theorem of Thom [Tho57] that we have a splitting

$$\text{Map}_*(Y, K(G, n)) = \prod_{i=1}^n K(H^{n-i}(Y; G), i),$$

and in particular, $\pi_1 \text{Map}_*(Y, K(\mathbb{Z}, 4)) \cong H^3(Y; \mathbb{Z})$.

In our case, the gauge group $\mathcal{G}^{e,h}$ consists of even gauge transformations which are asymptotically parallel; this fits into an exact sequence $\mathcal{G}^e \rightarrow \mathcal{G}^{e,h} \rightarrow \Gamma_\alpha \times \Gamma_\beta$, where the groups Γ are the stabilizers of the corresponding connections and \mathcal{G}^e is the gauge group of asymptotically trivial gauge transformations. In particular, we see that $\pi_0 \mathcal{G}^e \rightarrow \pi_0 \mathcal{G}^{e,h}$ is a surjection. The homotopy type of the group of gauge transformations which are asymptotically trivial is the same as the relative gauge group over $(W, \partial W)$; the discussion of the above paragraph goes through without difficulty to this relative setting, replacing cohomology with relative cohomology. We conclude that

$$\pi_0 \mathcal{G}^e = H^3(W, \partial W; \mathbb{Z}) \cong H_1(W; \mathbb{Z})$$

by Poincaré duality, and in particular $H_1(W; \mathbb{Z}) \rightarrow \pi_1 \tilde{\mathcal{B}}_{\mathbf{E} \oplus \mathbf{C}^\ell}$ is surjective.

Now that we know what the loops in this space are, the goal is to show that $w_1(\det(Q'))$ pairs trivially against them. Now we are in standard territory, and the proof follows exactly as in [Don87, Lemma 3.23] (modifying the Poincaré duality argument to account for the boundary): if one chooses a base connection \mathbf{A} , a concentrated charge-one instanton J on S^4 , and a loop $[\gamma] \in H_1(W)$ in the interior

of W , then one performs a family connected-sum construction along the loop to construct a loop of connections. One may identify both that the corresponding loop in $\pi_1 \tilde{\mathcal{B}}_{\mathbf{E} \oplus \mathbb{C}^e}$ is the image of $[\gamma]$, and that the determinant line bundle over this loop of connections is trivial (if one so desires, this can be carried out using the index-gluing constructions along the boundary of a 4-manifold, which will be described later in this section).

Therefore, the pairing of $w_1(\det(Q'))$ against $H_1(W; \mathbb{Z})$ is identically zero, and so $\det(Q')$ is trivializable. Being the pullback of the sum of this bundle and a trivial line bundle, the same is true of $\det(Q)$, as desired. \blacksquare

Now choose a framed π -instanton $[\mathbf{A}, p] \in \tilde{\mathcal{M}}_{\mathbf{E}, z, \pi}$, where π is a regular perturbation. The tangent space $T_{[\mathbf{A}, p]} \tilde{\mathcal{M}}_{\mathbf{E}, z, \pi}$ is identified with the first cohomology of the deformation complex

$$\Omega_{k+1, \delta}^0(\mathfrak{g}_E) \xrightarrow{d_{\mathbf{A}, \text{ev}_b}} \Omega_{k, \delta}^1(\mathfrak{g}_E) \oplus \mathfrak{g} \xrightarrow{D_{\mathbf{A}, \pi}} \Omega_{k-1, \delta}^{2,+}(\mathfrak{g}_E),$$

where \mathfrak{g} means the fiber of \mathfrak{g}_E above b , and $\text{ev}_b \sigma = \sigma(b)$. We denote this complex $C_{\mathbf{A}}$. Here we require k to be sufficiently large that every L^2_{k+1} function is continuous, so that evaluation is a continuous operation. Write $\mathfrak{g}_{\mathbf{A}}$ to denote the subspace of \mathfrak{g} that extends to \mathbf{A} -parallel gauge transformations (so that $\mathfrak{g}_{\mathbf{A}}$ is the Lie algebra of $\Gamma_{\mathbf{A}}$), and write $e : \mathfrak{g}_{\mathbf{A}} \rightarrow \Omega_{k+1, \delta, \text{ext}}^0$ for the operator that takes a point-value and sends it to the corresponding parallel section.

Take a map $p_{\mathbf{A}} : \mathfrak{g}_{\mathbf{A}} \rightarrow \Omega^0(\mathfrak{g}_E)$ which has $\langle p_{\mathbf{A}}(h), e(h) \rangle > 0$ for all $h \neq 0 \in \mathfrak{g}_{\mathbf{A}}$. We write the operator $\tilde{Q}_{\mathbf{A}, \pi} : \Omega_{k, \delta}^1 \oplus \mathfrak{g} \rightarrow \Omega_{k-1, \delta}^{2,+} \oplus \Omega_{k-1, \delta}^0$ for $\tilde{Q} = Q + p_{\mathbf{A}}$ (suppressing the $p_{\mathbf{A}}$ from notation; as we shall see shortly, this is not so depraved).

Lemma 5.2. *The composite map $\ker(\tilde{Q}) \hookrightarrow \ker(D_{\mathbf{A}, \pi}) \rightarrow H^1(C_{\mathbf{A}})$ is an isomorphism if $p_{\mathbf{A}}$ is sufficiently small.*

Proof. By assumption that π is a regular perturbation, the complex $C_{\mathbf{A}}$ has cohomology concentrated in degree one, and $\text{coker}(Q_{\mathbf{A}, \pi}) = \mathfrak{g}_{\mathbf{A}}$. By assumption that $p_{\mathbf{A}}(h)$ pairs nontrivially with $e(h)$ for any $h \in \mathfrak{g}_{\mathbf{A}}$, we see that $\tilde{Q}_{\mathbf{A}, \pi}$ is surjective. Because the index of the operator \tilde{Q} agrees with the index of the complex C , we see that $\ker(\tilde{Q})$ and $H^1(C_{\mathbf{A}})$ have the same dimension; so it suffices to show surjectivity.

So suppose $D_{\mathbf{A}, \pi} \omega = 0$; our goal is to show there is some $\sigma \in \Omega_{\text{ext}}^0$ so that

$$\Delta_{\mathbf{A}} \sigma + p_{\mathbf{A}}(\sigma(b)) = -d_{\mathbf{A}}^* \omega.$$

The map $\Delta_{\mathbf{A}} : \Omega_{k+1, \text{ext}, \delta}^0 \rightarrow \Omega_{k-1, \delta}^0$ is a self-adjoint operator with kernel $\mathfrak{g}_{\mathbf{A}}$.

In particular, because $p_{\mathbf{A}} : \ker(\Delta_{\mathbf{A}}) \rightarrow \text{coker}(\Delta_{\mathbf{A}})$ is assumed to be an isomorphism, the map $\Delta_{\mathbf{A}} + tp_{\mathbf{A}}$ is surjective for sufficiently small t . So (replacing $p_{\mathbf{A}}$ with $tp_{\mathbf{A}}$ if necessary, t small) we see that $\Delta_{\mathbf{A}} + p_{\mathbf{A}} : \Omega_{k+1, \delta, \text{ext}}^0 \rightarrow \Omega_{k-1, \delta}^0$ is an isomorphism. Therefore we may indeed solve the desired equation, so that we see the map $\ker(\tilde{Q}) \rightarrow H^1(C_{\mathbf{A}})$ is surjective (and hence an isomorphism), as desired. \blacksquare

In particular, we have a canonical isomorphism *in a neighborhood of* $[\mathbf{A}, p]$ between $\det(T\tilde{\mathcal{M}})$ and $\det(\tilde{Q})$. The sign of this isomorphism is independent of the auxiliary choice of $p_{\mathbf{A}}$. What we look to investigate is, then, the latter operator.

Recall now that the index of Fredholm operators is invariant under homotopy through Fredholm operators; because the map $Q_{\mathbf{A}, \pi} + tp_{\mathbf{A}}$, for $t \in [0, 1]$, gives

a homotopy through Fredholm operators (as \mathfrak{g} is finite-dimensional), we see that $I(\tilde{Q}_{\mathbf{A},\pi}) \cong I(Q_{\mathbf{A},\pi}) \oplus \mathfrak{g}$, and in particular

$$\det(\tilde{Q}_{\mathbf{A},\pi}) \cong \det(Q_{\mathbf{A},\pi}) \otimes \det(\mathfrak{g}).$$

Therefore, a choice of orientation of the Lie group $SO(3)$ and an orientation of Q canonically gives us an orientation of \tilde{Q} . This is completely independent of the choice of r , and we may as well remove r from the discussion of the index theory, and instead study $Q_{\mathbf{A},\pi}$ itself.

In particular, a choice of orientation of the line bundle $\det(Q_{\mathbf{A},\pi})$ for \mathbf{A} in some component of $\tilde{\mathcal{B}}_{\mathbf{E}}^e$, and a fixed orientation of $\mathfrak{so}(3)$, therefore gives a global orientation of the framed moduli space $\tilde{\mathcal{M}}$ sitting inside that component.

If $\mathbf{A} \in \tilde{\mathcal{B}}_{\mathbf{E},z,k,\delta}^e(\alpha, \beta)$, where here $z \in \pi_0 \tilde{\mathcal{B}}_{\mathbf{E},k,\delta}^e$ labels a path-component, we write $\Lambda_z^W(\alpha, \beta)$ for the two-element set of orientations of $\det(Q_{\mathbf{A}}) \otimes \det(\mathfrak{g})$. As above, a choice of element of $\Lambda_z^W(\alpha, \beta)$ induces an orientation on the moduli space $\tilde{\mathcal{M}}_{\mathbf{E},z,\pi}(\alpha, \beta)$.

5.2. A gluing operation for determinant lines. Let (W_i, \mathbf{E}_i) be two cobordisms (considered as manifolds with two cylindrical ends), so that the positive end of (W_1, \mathbf{E}_1) is isometric to the negative end of (W_2, \mathbf{E}_2) . Here we include the possibility that either or both of the W_i are cylinders; write $(W_{12}^T, \mathbf{E}_{12})$ for the composite cobordism, glued by cutting off the positive end of W_1 (say, $[0, \infty) \times Y_2$) at $\{T\} \times Y_2$ and similarly for the negative end of W_2 . If it is not necessary for clarity, we drop the T from notation. If z_1 and z_2 are components of $\tilde{\mathcal{B}}_{\mathbf{E}_1,k,\delta}^e(\alpha, \beta)$ and $\tilde{\mathcal{B}}_{\mathbf{E}_2,k,\delta}^e(\beta, \gamma)$, respectively, then denote the apparent ‘composite component’ as z_{12} .

What we would like is to take an element of each of $\Lambda_{z_1}^{W_1}(\alpha, \beta)$ and $\Lambda_{z_2}^{W_2}(\beta, \gamma)$, and output an element of $\Lambda_{z_{12}}^{W_{12}}(\alpha, \gamma)$ (possibly depending on some extra input coming from β alone).

To do this, we will use the mechanism of gluing indices along boundary components. We begin by setting up some notation.

Let X_1 and X_2 be 4-manifolds equipped with $SO(3)$ -bundles \mathbf{E}_i , possibly with cylindrical ends, so that ∂X_i is nonempty and the metric is of product type near the boundary. We decompose each ∂X_i into a union $\partial_- X_i \cup \partial_+ X_i$ of connected components (either term possibly empty) which we call ‘positive’ or ‘negative’, and orient them so that a neighborhood of $\partial_+ X_i$ is isometric to $(-t, 0] \times \partial_+ X_i$ and a neighborhood of $\partial_- X_i$ is isometric to $[0, t) \times \partial_- X_i$. Suppose furthermore that $(\partial_+ X_1, \mathbf{E}_1)$ is oriented isometric to $(\partial_- X_2, \mathbf{E}_2)$. Write

$$Y_1 = \partial_- X_1, \quad Y_2 = \partial_+ X_1 = \partial_- X_2, \quad Y_3 = \partial_+ X_2,$$

and write $Y = Y_1 \cup Y_2 \cup Y_3$.

Then we may form the composite $X_{12} = X_1 \cup_{Y_2} X_2$, equipped with $SO(3)$ -bundle \mathbf{E}_{12} .

Suppose each (X_i, \mathbf{E}_i) is equipped with a choice of 4-manifold perturbation, regular on the ends, which is also constant in time near the boundary. Suppose each is equipped with a connection \mathbf{A}_i of regularity $L_{k,\delta}^2$, decaying towards π -flat connections on the ends, and suppose that \mathbf{A}_{12} is a connection on X_{12} of the same regularity which agrees with each \mathbf{A}_i when restricted to each X_i . Write A for the restriction of \mathbf{A}_{12} to Y ; it is of regularity $L_{k-1/2}^2$.

We are interested in studying the perturbed ASD operator on these compact manifolds with boundary. We should introduce a restriction on the boundary values to make this a Fredholm operator, as in Section 4.7. However, the choice there (the Coulomb-Neumann gauge) is not appropriate for us here; that boundary value problem is not continuously varying as we change the reducibility type of A . Instead, we use the standard Atiyah-Patodi-Singer spectral boundary value problem. On the composite X_{12} , the ASD operator takes the form $d/dt + L_{A,\pi}$ on $(-t, t) \times Y$ for the signature operator

$$L_{A,\pi} : \Omega_{k-1/2}^0(Y) \oplus \Omega_{k-1/2}^1(Y) \rightarrow \Omega_{k-3/2}^0(Y) \oplus \Omega_{k-3/2}^1(Y),$$

written in matrix form as

$$L_{A,\pi} = \begin{pmatrix} 0 & -d_A^* \\ -d_A & D_{A,\pi} \end{pmatrix},$$

where $D_{A,\pi} = *d_A + D_A \nabla_\pi$.

Henceforth we write $\Omega^0(Y) \oplus \Omega^1(Y)$ as H , only writing the Sobolev indices if necessary. When there are multiple Y_i , we write

$$H_i := \Omega^0(Y_i) \oplus \Omega^1(Y_i).$$

Suppose $\lambda \in \mathbb{R}$ is not an eigenvalue of $L_{A,\pi}$ for any A arising as the restriction of \mathbf{A} to a boundary component. Then we may split H as $H^{>\lambda} \oplus H^{<\lambda}$ as the closure of the linear span of eigenvectors with eigenvalue greater than, or less than, λ . Then we define the λ -weighted ASD operator on X_1 to be

$$Q_{\mathbf{A},\pi}^\lambda = (D_{\mathbf{A},\pi}, d_{\mathbf{A}}^*, \Pi_1^{<\lambda} r, \Pi_2^{>\lambda} r) : \Omega_{k,\delta}^1(X_1) \rightarrow \Omega^{2,+}(X_1) \oplus \Omega^0(X_1) \oplus H_1^{<\lambda} \oplus H_2^{>\lambda},$$

where r is restriction to the appropriate boundary components and the operators Π are the $L_{k-1/2}^2$ orthogonal projections onto the corresponding eigenspaces. The operator written on X_2 is the same, with the boundary value projections changed to $\Pi_2^{<\lambda}$ and $\Pi_3^{>\lambda}$, and similarly with X_{12} .

With this notation in hand, the following is essentially the content of the beginning of [KM07, Section 20.3], especially pages 383-384.

Lemma 5.3. *Let K be any compact family of connections \mathbf{A}_{12} and perturbations as above, where λ is never an eigenvalue for any $L_{A,\pi}$ for any A given by restriction of $\mathbf{A} \in K$ to the boundary. Then the operators $Q_{\mathbf{A}_i,\pi}^\lambda$ form a family of Fredholm operators, and there is a canonical isomorphism of line bundles over K*

$$\det(Q_{\mathbf{A}_1,\pi}^\lambda) \otimes \det(Q_{\mathbf{A}_2,\pi}^\lambda) \rightarrow \det(Q_{\mathbf{A}_{12},\pi}^\lambda),$$

given by a homotopy between the operators $Q_{\mathbf{A}_1,\pi}^\lambda \oplus Q_{\mathbf{A}_2,\pi}^\lambda$ and an operator whose kernel and cokernel are canonically identified with that of $Q_{\mathbf{A}_{12},\pi}^\lambda$.

What we really care about are connections $\mathbf{A} \in \tilde{\mathcal{B}}_{\mathbf{E},z,k,\delta}(\alpha, \beta)$ on a manifold W with two cylindrical ends and no boundary components. To make use of the above lemma, we should break these into pieces. For T large, write \mathbf{N} for the restriction of \mathbf{A} to $(-\infty, -T] \times Y_1$, \mathbf{P} for the restriction to $[T, \infty) \times Y_2$, and \mathbf{C} for the restriction to the compact piece lying between these.

We fix a small constant $\epsilon > 0$. What we demand of our choice of T , writing $\mathbf{N} = p^* \alpha + a$ for $p^* \alpha$ the pullback connection $d/dt + d_\alpha$ and $a \in \Omega_{k,\delta}^1((-\infty, -T] \times Y_1)$, is the following. First, that the signature operator $L_{\alpha+ta(-T)}$ never has an eigenvalue of absolute value ϵ for any $t \in [0, 1]$. Second, that the only eigenvalue of $L_{\alpha,\pi}$ in

$[-\epsilon, \epsilon]$ is 0, which corresponds to $\ker(L_{\alpha, \pi}) = \mathfrak{g}_\alpha$, the space of α -parallel sections of \mathfrak{g}_E ; this is the entire kernel by the assumption that α is a nondegenerate critical orbit. We demand similarly for \mathbf{P} . This is a smallness condition on \mathbf{N} and \mathbf{P} , or rather the size of the differences between those and constant connections. In particular, the operators $Q_{\mathbf{N}, \pi}^{-\epsilon}$ and $Q_{\mathbf{P}, \pi}^{-\epsilon}$ are both Fredholm operators, and there are canonical isomorphisms

$$\det(Q_{\mathbf{N}, \pi}^{-\epsilon}) \cong \det(Q_{p^* \alpha, \pi}^{-\epsilon}), \quad \det(Q_{\mathbf{P}, \pi}^{-\epsilon}) \cong \det(Q_{\beta, \pi}^{-\epsilon}).$$

This remains true for families of such connections, so long as T is chosen to satisfy the given properties for all connections in the family; this is always possible if the family is compact.

Then by Lemma 5.3, we have an isomorphism

$$\det(Q_{p^* \alpha, \pi}^{-\epsilon}) \otimes \det(Q_{\mathbf{C}, \pi}^{-\epsilon}) \otimes \det(Q_{p^* \beta, \pi}^{-\epsilon}) \cong \det(Q_{\mathbf{A}, \pi}).$$

To reduce this to studying \mathbf{C} alone, we should understand explicitly the operators corresponding to each end.

Lemma 5.4. *Write $Z^- = (-\infty, 0] \times Y_1$ and $Z^+ = [0, \infty) \times Y_2$.*

The operator

$$Q_{p^* \alpha, \pi}^{-\epsilon} : \Omega_{k, \delta}^1(Z^-) \rightarrow \Omega_{k-1, \delta}^{2,+}(Z^-) \oplus \Omega_{k-1, \delta}^0(Z^-) \oplus H_1^{<-\epsilon}$$

is an isomorphism, and so there is a canonical trivialization $\det(Q_{p^ \alpha, \pi}^{-\epsilon}) \cong \mathbb{R}$. The operator*

$$Q_{p^* \beta, \pi}^{-\epsilon} : \Omega_{k, \delta}^1(Z^+) \rightarrow \Omega_{k-1, \delta}^{2,+}(Z^+) \oplus \Omega_{k-1, \delta}^0(Z^+) \oplus H_1^{>-\epsilon}$$

is injective with cokernel canonically identified with \mathfrak{g}_β , the space of β -parallel sections of \mathfrak{g}_E , and so we have a canonical isomorphism $\det(Q_{p^ \beta, \pi}^{-\epsilon}) \cong \det(\mathfrak{g}_\beta)^*$, giving a trivialization as soon as we orient the vector space \mathfrak{g}_β .*

Proof. Because the operator

$$(D_{p^* \alpha, \pi}, d_{p^* \alpha}^*) = (d_{p^* \alpha}^+ + D_\alpha \widehat{\nabla}_\pi, d_{p^* \alpha}^*)$$

takes the form $d/dt + L_{\alpha, \pi}$, and identically with β , we may solve the equation $(D_{p^* \alpha, \pi}, d_{p^* \alpha}^*)\omega = 0$ by separation of variables, decomposing $\Omega^0(Y_1) \oplus \Omega^0(Y_2)$ into the eigenspaces of $L_{\alpha, \pi}$. A solution with boundary value $\phi \in L_{k-1/2}^2$ exists if and only if ϕ is in the closure of the span of eigenvectors with negative eigenvalues. Because $H_1^{<-\epsilon} = H_1^{<0}$ by assumption on ϵ , we see that the operator $Q_{p^* \alpha, \pi}^{-\epsilon}$ is invertible, as desired. However, for the other cylindrical end, we have $H_2^{>-\epsilon} = H_2^{\geq 0}$, we see that we have cokernel equal to $H_2^0 = \ker(L_{\alpha, \pi}) = \mathfrak{g}_\alpha$, by the assumption that α is nondegenerate. \blacksquare

We use these facts about the index to construct a comparison map between certain determinant line bundles.

Definition 5.1. *Let \mathbf{A}_1 and \mathbf{A}_2 be connection on cobordisms W_1 and W_2 as above, going between α and β or β and γ , respectively. If \mathbf{A}_{12} is a connection on the composite W_{12}^T which is uniformly close to \mathbf{A}_1 on $W_1^{\leq T}$ and \mathbf{A}_2 on $W_2^{\geq -T}$, then there is a canonical isomorphism*

$$\rho_{\text{an}} : \det(Q_{\mathbf{A}_1, \pi}) \otimes \det(\mathfrak{g}_\beta) \otimes \det(Q_{\mathbf{A}_2, \pi}) \cong \det(\widetilde{Q}_{\mathbf{A}_{12}, \pi}),$$

given by decomposing

$$\det(Q_{\mathbf{A}_1, \pi}) \cong \det(Q_{\mathbf{C}_1, \pi}^{-\epsilon}) \otimes \det(\mathfrak{g}_\beta)^*$$

and

$$\det(Q_{\mathbf{A}_2, \pi}^{-\epsilon}) \cong \det(Q_{\mathbf{C}_2, \pi}) \otimes \det(\mathfrak{g}_\gamma)^*,$$

as above, and applying the canonical isomorphism $\det(\mathfrak{g}_\beta)^* \det(\mathfrak{g}_\beta) \cong \mathbb{R}$, as well as the isomorphism

$$\det(Q_{\mathbf{C}_1, \pi}^{-\epsilon}) \otimes \det(Q_{\mathbf{C}_2, \pi}^{-\epsilon}) \cong \det(Q_{\mathbf{C}_{12}, \pi}^{-\epsilon})$$

given by Lemma 5.3, by first applying a small homotopy taking \mathbf{C}_i to the restriction of \mathbf{C}_2 .

We call ρ_{an} the analytic gluing map. There is a corresponding version of this map for the operators \tilde{Q} which have a factor of \mathfrak{g} to account for the $SO(3)$ -action; this is written

$$\tilde{\rho}_{\text{an}} : \det(\tilde{Q}_{\mathbf{A}_1, \pi}) \otimes \det(\mathfrak{g}_\beta^\perp)^* \otimes \det(\tilde{Q}_{\mathbf{A}_2, \pi}) \cong \det(\widetilde{Q_{\mathbf{A}_{12}, \pi}}),$$

given by decomposing

$$\det(\tilde{Q}_{\mathbf{A}_i, \pi}) \cong \det(Q_{\mathbf{A}_i, \pi}) \otimes \det(\mathfrak{g})$$

and using the canonical isomorphism $\det(\mathfrak{g}) \cong \det(\mathfrak{g}_\beta) \otimes \det(\mathfrak{g}_\beta^\perp)$, where here we implicitly demand that $\mathfrak{g} = \mathfrak{g}_\beta \oplus \mathfrak{g}_\beta^\perp$ is an oriented splitting if we choose an orientation on any two of these.

In particular, this gives an isomorphism

$$\Lambda_{z_1}^{W_1}(\alpha, \beta) \times_{\mathbb{Z}/2} \Lambda(\mathfrak{g}_\beta) \times_{\mathbb{Z}/2} \Lambda_{z_2}^{W_2}(\beta, \gamma) \cong \Lambda_{z_{12}}^{W_{12}}(\alpha, \gamma),$$

where $\Lambda(\mathfrak{g}_\beta)$ is the set of orientations of \mathfrak{g}_β .

Given any compact family of \mathbf{A}_i as above, the analytic gluing map ρ_{an} gives an isomorphism of line bundles over the parameter space, and identically with $\tilde{\rho}_{\text{an}}$.

Note that this is usually relevant to us in the case where one of W_i is a cylinder. If neither are cylinders, the orientations that actually appear in practice are of the parameterized ASD operator, with parameter given by some family of broken metrics.

5.3. The gluing diffeomorphism is orientation-preserving. In the previous section, we defined a gluing isomorphism ρ_{an} between certain determinant lines. We also have a similar isomorphism coming from a different source: the inverse function theorem. Recall from Section 4.7 that, for certain connected $SO(3)$ -invariant open subsets $U_{\alpha\beta} \subset \tilde{\mathcal{M}}_{\mathbf{E}_1, z_1}(\alpha, \beta)$ and $U_{\beta\gamma} \subset \tilde{\mathcal{M}}_{\mathbf{E}_2, z_2}(\beta, \gamma)$, Proposition 4.43 gives a diffeomorphism

$$\text{gl} : U_{\alpha\beta} \times_\beta U_{\beta\gamma} \rightarrow U_{\alpha\gamma}.$$

In the compactification by broken trajectories, the closure of $\text{gl}(U_{\alpha\beta} \times_\beta U_{\beta\gamma})$ includes the open face

$$\tilde{\mathcal{M}}_{\alpha\beta} \times_\beta \tilde{\mathcal{M}}_{\beta\gamma} \subset \partial \overline{\mathcal{M}}_{\alpha\gamma}.$$

We remind the reader here that when considering $\tilde{\mathcal{M}}_{\mathbf{E}, k, \delta}$ in the case where W is a cylinder, this is the moduli space of *parameterized* trajectories: we do not reduce by the \mathbb{R} -action.

In particular, suppose we are given a pair of framed instantons $([\mathbf{A}_1, p_1], [\mathbf{A}_2, p_2])$ which project to the same framing at β . Taking the derivative of gl , we obtain an isomorphism

$$\tilde{\rho}_{\text{gm}} : \det(\tilde{Q}_{\mathbf{A}_1, \pi}) \otimes \det(\mathfrak{g}_\beta)^* \otimes \det(\tilde{Q}_{\mathbf{A}_2, \pi}) \cong \det(\tilde{Q}_{\mathbf{A}_{12}, \pi});$$

note that the $SO(3)$ -action on a point gives an isomorphism $T_p\beta \cong \mathfrak{g}_\beta^\perp$, which explains this factor. Because the gluing map is equivariant, this descends to an isomorphism

$$\rho_{\text{gm}} : \det(Q_{\mathbf{A}_1, \pi}) \otimes \det(\mathfrak{g}_\beta^\perp) \otimes \det(Q_{\mathbf{A}_2, \pi}) \cong \det(Q_{\mathbf{A}_{12}, \pi}).$$

We call ρ_{gm} the *geometric gluing map*. If the open sets $U_{\alpha\beta}$ and $U_{\beta\gamma}$ are chosen so that the restriction of $\mathbf{A} \in U_{\alpha\beta}$ to each end is close enough to the constant trajectory, and similarly for the restriction of $\mathbf{A} \in U_{\beta\gamma}$ to each end, that we may apply Lemma 5.3. Then we obtain *two* isomorphisms

$$\rho : \det(Q_{\mathbf{A}_1, \pi}) \otimes \det(\mathfrak{g}_\beta) \otimes \det(Q_{\mathbf{A}_2, \pi}) \cong \det(Q_{\mathbf{A}_{12}, \pi}).$$

The next part of this section is dedicated to showing that these two isomorphisms are positive scalar multiples of one another, and hence both induce the same orientation on $\det(Q_{\mathbf{A}_{12}, \pi})$.

The map ρ_{gm} (or rather, $\tilde{\rho}_{\text{gm}}$) is not difficult to understand at the level of kernels and cokernels; it is less clear how to understand $\tilde{\rho}_{\text{an}}$ at this level, so this is our next immediate goal.

Lemma 5.5. *Let V, W , and W_b be Hilbert spaces; suppose we have a continuous family of Fredholm maps*

$$A_t = (A, r_t) : V \rightarrow W \oplus W_b,$$

with $A : V \rightarrow W$ surjective. Suppose that A_0 is surjective, and that the map

$$r_t - r_0 : \ker(A) \rightarrow W_b$$

has sufficiently small norm, uniformly in t . Let $J \subset W_b$ be a finite-dimensional subspace so that $\tilde{A}_t : V \oplus J \rightarrow W \oplus W_b$, given by $\begin{pmatrix} A & 0 \\ r_t & 1 \end{pmatrix}$, is surjective for all t .

Then the composite

$$\ker(\tilde{A}_0) \hookrightarrow V \oplus J \twoheadrightarrow \ker(\tilde{A}_t)$$

is an isomorphism, where the first map is inclusion and the second map is orthogonal projection. Similarly, the map $\ker(A_0) \rightarrow \ker(A_1)$ is injective, and the above gives us an isomorphism $\ker(A_0) \oplus \text{coker}(A_1) \cong \ker(A_1)$. This gives an isomorphism between orientation sets $\Lambda(A_0) \cong \Lambda(A_1)$, the same as that induced by the homotopy through Fredholm operators A_t .

Proof. First, we argue in slightly more generality. Suppose we are given two Fredholm maps $T_i : X \rightarrow Y$ of Hilbert spaces, and that $\|T_2|_{\ker(T_1)}\| < \epsilon$. Write $\pi_2 : X \rightarrow \ker(T_2)$ for the orthogonal projection, and $C_2 = \ker(\pi_2)$, the orthogonal complement to $\ker(T_2)$. Then the map $T_2 : C_2 \rightarrow Y$ is an isomorphism onto its (closed) image, and in particular enjoys an inequality of the form $\|T_2 v\|_Y \geq c\|v\|_{C_2}$ for some $c > 0$. For any $v \in \ker(T_1)$, we have

$$\|(1 - \pi_2)v\| \leq \frac{1}{c}\|T_2(1 - \pi_2)v\| = \frac{1}{c}\|T_2 v\| \leq \frac{\epsilon}{c}\|v\|;$$

the equality follows because $T_2\pi_2v = 0$, by definition of π_2 . As soon as $\epsilon < c$, we see that π_2v must be nonzero. In particular, so long as $\|T_2|_{\ker(T_1)}\|$ is sufficiently small, the composite $\ker(T_1) \hookrightarrow X \twoheadrightarrow \ker(T_2)$ is injective.

Given an element (v, j) of $\ker(\tilde{A}_0)$, we have

$$\tilde{A}_t(v, j) = (Av, r_tv + j) = (0, (r_t - r_0)v);$$

because $v \in \ker(A)$, we see that

$$\|\tilde{A}_t|_{\ker(\tilde{A}_0)}\| \leq \|(r_t - r_0)|_{\ker(A)}\|.$$

So if $\|(r_t - r_0)|_{\ker(A)}\|$ is sufficiently small, we see that $\ker(\tilde{A}_0) \rightarrow \ker(\tilde{A}_t)$ as defined above is injective, for all t ; because the operators \tilde{A}_t are all homotopic through Fredholm operators and assumed to be surjective, these kernels all have the same dimension, and hence this injection is an isomorphism.

By the assumption that A_0 is surjective, we may identify $\ker(\tilde{A}_0)$ as $\ker(A_0) \oplus J$. If we write $\text{coker}(\tilde{A}_t) \oplus J'_t \cong J$, then choosing a section of the projection $\ker(\tilde{A}_t) \rightarrow J'_t$, we may similarly identify $\ker(\tilde{A}_t) \cong \ker(A_t) \oplus J'_t$. The isomorphism

$$\ker(A_0) \oplus J \cong \ker(A_t) \oplus J'_t$$

is isotopic to the direct sum of an isomorphism $\ker(A_0) \oplus \text{coker}(A_t) \rightarrow \ker(A_t)$ with the identity map on J'_t . Here, the map $\text{coker}(A_t) \rightarrow \ker(A_t)$ is given (up to a small homotopy) by first choosing a section s of the surjective map

$$\ker(A) \xrightarrow{-r_0} W_b \rightarrow \text{coker}(A_t),$$

which supplies us with a map

$$\begin{aligned} s' : \text{coker}(A_t) &\rightarrow \ker(A_0) \oplus J \\ s'(j) &= (s(j), r_0s(j)); \end{aligned}$$

because the space of sections is contractible this map only depends on the choice of section up to an isotopy. Then as usual we apply orthogonal projection to $s(j)$ to obtain a map $\ker(A_0) \oplus J \rightarrow \ker(A_t)$.

This, then, is identified as an isomorphism of virtual vector spaces

$$(\ker(A_0), 0) \rightarrow (\ker(A_t), \text{coker}(A_t)),$$

the map on kernels just being the composite $\ker(A_0) \rightarrow V \rightarrow \ker(A_t)$ as above.

Given that we have a family of such isomorphisms, starting with the identity, identifies this map with the isomorphism $\Lambda(A_0) \rightarrow \Lambda(A_1)$ induced by the homotopy through Fredholm operators. \blacksquare

Let W_1 and W_2 be cobordisms, with the positive end of W_1 and the negative end of W_2 modeled on (Y, E) . Suppose we have regular framed instantons $(\mathbf{A}_1, p_1), (\mathbf{A}_2, p_2)$ on the respective cobordisms (possibly cylinders) which limit to the critical orbit β ; the gluing map provides a framed instanton (\mathbf{A}_{12}^T, p) on the glued-up cobordism

$$W_{12}^T = W_1^{\leq T} \cup_Y W_2^{\geq -T},$$

obtained by truncating the two cobordisms at some large $\{\pm T\} \times Y$ and pasting them together. The defining property of (\mathbf{A}_{12}^T, p) , arising from the inverse function theorem, is as follows. First, for large enough T the restriction to $W_1^{\leq T/2}$ and $W_2^{\geq -T/2}$ is uniformly close to (\mathbf{A}_1, p_1) and (\mathbf{A}_2, p_2) , respectively (here we use

parallel transport to move the framing between various basepoints). Second, the restriction to the region in between these two pieces, isometric to $[-T/2, T/2] \times Y$, is uniformly small. We write $[-T/2, T/2] \times Y =: Z^{T/2}$. If $\pi : E_b \rightarrow E_b/\Gamma_\beta$ is the projection and q the framing at some basepoint b on $Z^{T/2}$, then πq is uniformly close to the common limiting value q of p_1 and p_2 .

In this context, we consider

$$V = \Omega_{k,\delta}^1(W_1^{\leq T/2}) \oplus \Omega_{k,\delta}^1(Z^{T/2}) \oplus \Omega_{k,\delta}^1(W_2^{\geq -T/2}),$$

while

$$W = (\Omega^{2,+} \oplus \Omega^0)_{k-1,\delta}(W_1^{\leq T/2}) \oplus (\Omega^{2,+} \oplus \Omega^0)_{k-1,\delta}(Z^{T/2}) \oplus (\Omega^{2,+} \oplus \Omega^0)_{k-1,\delta}(W_2^{\geq -T/2}),$$

and

$$W_b = (\Omega^0 \oplus \Omega^1)_{k-1/2}(\{-T/2, T/2\} \times Y).$$

The operator $A : V \rightarrow W$ is given by (d^*, d^+) , while there are two restriction operators $V \rightarrow W_b$ of interest to us: one is $r_0 = (\text{res}, -\text{res})$, taking the restriction on each positive boundary component and negative the restriction on each negative boundary component. The other is $r_1 = (\Pi^{>-\epsilon}, \Pi^{<-\epsilon})$, the spectral projection onto the $> -\epsilon$ eigenspaces on positive ends, and the $< \epsilon$ eigenspaces on negative ends. The operator A_1 has the property that it splits as a direct sum of operators corresponding to the pieces $W_1^{\leq T/2}, Z^{T/2}, W_2^{\geq -T/2}$, while the operator A_0 has kernel and cokernel naturally isomorphic to $Q_{\mathbf{A}_{12},\pi}$ itself. Here the appropriate choice of $L_{k,\delta}^2$ norm on the direct sum in fact has a weight on the $Z^{T/2}$ component, as in Section 4.7, essentially weighted by a symmetric function equal to $e^{\delta(t+T/2)}$ on $[-T/2, -t)$ for some small $t > 0$.

We may interpolate through these projections via

$$r_t = (\Pi^{>-\epsilon} + t\Pi^{<-\epsilon}, -t\Pi^{>-\epsilon} - \Pi^{<-\epsilon});$$

as discussed in [KM07, Section 20.3], this gives a homotopy $A_t = (A, r_t)$ through Fredholm operators. Furthermore, if T is chosen large enough, then $\|(r_t - r_0)|_{\ker(A)}\|$ is arbitrarily small, uniformly in t . Observe here that

$$r_t - r_0 = t(\Pi^{<-\epsilon}, -\Pi^{>-\epsilon}),$$

so we may focus attention on the case $t = 1$. For convenience, we only pay attention to the factor $\Pi^{<-\epsilon}$ as we approach $\{T/2\} \times Y$ from the left.

Write $\omega_{T/4} = \omega|_{\{T/4\} \times Y}$ and $\omega_{T/2} = \omega|_{\{T/2\} \times Y}$. For T very large, we have that $\mathbf{A}_{12}|_{W_1^{[T/4, T/2]}}$ is uniformly close to the constant trajectory at the π -flat connection β , as follows because \mathbf{A}_1 decays exponentially. For some constant C , independent of T , we have $\|\omega_{T/4}\|_{k-1/2} \leq C\|\omega\|_{k,\delta}$ for all $\omega \in L_{k,\delta}^2$; this is just the claim that the trace along a hypersurface is a continuous map (and the fact that a neighborhood of $\{T/4\} \times Y$ is isometric to $(-t, t) \times Y$, independent of T). Writing the ASD operator for the constant trajectory at β as an ODE in the eigenvalues of

$$L_{\beta,\pi} : (\Omega^0 \oplus \Omega^1)_{k-1/2}(Y) \rightarrow (\Omega^0 \oplus \Omega^1)_{k-3/2}(Y),$$

and writing $\omega_{T/2} = \omega_{T/2}^+ \oplus \omega_{T/2}^-$ for the spectral decomposition, we see that

$$\|\omega_{T/2}^-\| \leq \frac{e^{-\delta T/4}}{2} \|\omega_{T/4}^-\| \leq \frac{e^{-\delta T/4}}{C} \|\omega\|$$

for any $\omega \in \ker(A)$, where δ is the absolute value of the least eigenvalue of $L_{\beta,\pi}$, and the factor of 2 is simply a fudge factor to account for the fact that \mathbf{A}_1 is not *literally* the constant trajectory at β in the relevant portion of the cobordism. Here we use that $-\epsilon < 0$ to conclude that this component of $r_t - r_0$ is uniformly small.

For the negative boundary components, a similar argument applies, but now one must exploit the exponential weights in the definition of the $L_{k,\delta}^2$ norm on the compact cylinder and the fact that ϵ is chosen less than δ to get the desired bound (now instead there is a factor of $e^{(\epsilon-\delta)T/4}$).

This is almost sufficient to apply the above lemma, except for the condition that A_0 is surjective; rather, we have a canonical identification $\text{coker}(A_0) \cong \mathfrak{g}_{\mathbf{A}_{12}}$, the tangent space to the stabilizer $\Gamma_{\mathbf{A}_{12}}$. This is easy enough to dispatch; recall the definitions of the extended operators $\tilde{Q}_{\mathbf{A},\pi}$ from earlier in this section, depending on a certain choice of map $p : \mathfrak{g} \rightarrow \Omega^0(W)$. These were chosen precisely so that $\tilde{Q}_{\mathbf{A},\pi}$ is surjective for a regular instanton \mathbf{A} . In this context, take $\tilde{V} = V \oplus \mathfrak{g} \oplus \mathfrak{g}$ (thinking of each \mathfrak{g} as being a choice of framing on one of the components $W_1^{\leq T/2}$ or $W_2^{\geq -T/2}$, based at say $\gamma(-T)$ and $\gamma(T)$ for the base curve γ), while $\tilde{W} = W \oplus \mathfrak{g}$ and $\tilde{W}_b = W_b$. The map $\mathfrak{g} \oplus \mathfrak{g} \rightarrow \tilde{W}$ is the expected maps p_i on the $\Omega^0(W_i)$ factors, while the map $\mathfrak{g} \oplus \mathfrak{g} \rightarrow \mathfrak{g}$ is the identity in the first factor, and $-\text{Hol}_{\mathbf{A}_{12}}^{\gamma(T) \rightarrow \gamma(-T)}$ in the second factor; if the points $\gamma(\pm T)$ are chosen to lie on the boundary of the two pieces, then the composite of this map with the projection $\mathfrak{g} \rightarrow \mathfrak{g}_\beta^\perp$ is very close to the identity, because \mathbf{A}_{12} is sufficiently close to the constant trajectory at β .

Then the map $\tilde{A} : \tilde{V} \rightarrow \tilde{W}$ is surjective, as is $\tilde{A}_0 : \tilde{V} \rightarrow \tilde{W} \oplus \tilde{W}_b$, and the maps p_i contribute minimally to the boundary-evaluation, so these satisfy the assumptions of the lemma.

Corollary 5.6. *If T is sufficiently large, the index-theoretic isomorphism $\tilde{\Lambda}(\mathbf{A}_{12}) \cong \tilde{\Lambda}(\mathbf{A}_1^{\leq T/2})\Lambda(\mathfrak{g})^*\tilde{\Lambda}(\mathbf{A}_2^{\geq -T/2})$ is given, at the level of kernels and cokernels, by an injection*

$$\ker(\tilde{Q}_{\mathbf{A}_{12}}) \rightarrow \ker(\tilde{Q}_{\mathbf{A}_1}^{\leq T/2}) \oplus_{\mathfrak{g}} \ker(\tilde{Q}_{\mathbf{A}_2}^{\geq -T/2}),$$

obtained first by restriction to $\Omega^1(W_1^{\leq T/2} \sqcup W_2^{\geq -T/2})$ and second by projection to the kernel. (Here $\oplus_{\mathfrak{g}}$ means we take the kernel of the natural projection of the direct sum to \mathfrak{g} .)

Similarly, for sufficiently large cutoffs T , the index-theoretic isomorphism

$$\tilde{\Lambda}(\mathbf{A}_1)\Lambda(\mathfrak{g}_\beta^\perp)^*\tilde{\Lambda}(\mathbf{A}_2) \rightarrow \tilde{\Lambda}(\mathbf{A}_1^{\leq T/2})\Lambda(\mathfrak{g})^*\tilde{\Lambda}(\mathbf{A}_2^{\geq -T/2})$$

may be described at the level of kernels and cokernels as a map

$$\ker(\tilde{Q}_{\mathbf{A}_1}) \oplus_{\mathfrak{g}_\beta^\perp} \ker(\tilde{Q}_{\mathbf{A}_2}) \rightarrow \ker(\tilde{Q}_{\mathbf{A}_1}^{\leq T/2}) \oplus \ker(\tilde{Q}_{\mathbf{A}_2}^{\geq -T/2}),$$

which may be described explicitly again as restriction to the corresponding manifolds with boundary, and then projection to the latter kernel.

Proof. The only point not outlined above is a discussion of the middle piece, isometric to $[-T/2, T/2] \times Y$. Given that \mathbf{A}_{12} is uniformly close to the constant trajectory at β for large enough T , and that for large enough T the ASD map on the cylinder (with boundary conditions) is an isomorphism for the constant trajectory at β , we see that its determinant line is canonically trivial, and does not enter into the discussion. \blacksquare

Now consider the diagram

$$\begin{array}{ccc} \tilde{\Lambda}(\mathbf{A}_1)\Lambda(\mathfrak{g}_\beta^\perp)*\tilde{\Lambda}(\mathbf{A}_2) & \longrightarrow & \tilde{\Lambda}(\mathbf{A}_1^{\leq T/2})\Lambda(\mathfrak{g})*\tilde{\Lambda}(\mathbf{A}_2^{\geq -T/2}) \\ \downarrow & \nearrow & \\ \tilde{\Lambda}(\mathbf{A}_{12}) & & \end{array}$$

The vertical map may be taken to be either $\Lambda(\tilde{\rho}_{\text{gm}})$ or $\Lambda(\tilde{\rho}_{\text{an}})$. The horizontal map and upper-right map are as discussed above: they maybe understood either as the inverses of index-theoretic gluing isomorphisms, or via projection maps between various kernels.

If the vertical map is taken to be $\tilde{\rho}_{\text{an}}$, then this diagram commutes (in fact, ρ_{an} is *defined* so that this diagram commutes). If we take the vertical map to be $\tilde{\rho}_{\text{gm}}$, this diagram still commutes: we are claiming that the composite map

$$(\mathbf{A}_1, p_1, T, \mathbf{A}_2, p_2) \rightarrow (\mathbf{A}_{12}^T, p) \rightarrow (\mathbf{A}_{12}^{\leq T/2}, p'_1, \mathbf{A}_{12}^{\geq -T/2}, p'_2)$$

is uniformly close to restriction, and hence this map commutes (up to a small homotopy) at the level of kernels. Because the diagram commutes with either choice, $\Lambda(\tilde{\rho}_{\text{gm}}) = \Lambda(\tilde{\rho}_{\text{an}})$: the two gluing maps are the same.

If we have fixed an orientation of \mathfrak{g}_β , then via the recipe given by the analytic gluing map, an orientation of $\Lambda_z^{W_1}(\alpha, \beta)$ and an orientation of $\Lambda_w^{W_2}(\beta, \gamma)$ induces an orientation of $\Lambda_{zw}^{W_1 W_2}(\alpha, \gamma)$; it also naturally induces an orientation of the fiber product

$$\tilde{\mathcal{M}}_{\mathbf{E}_1, z, k, \delta}(\alpha, \beta) \times_\beta \tilde{\mathcal{M}}_{\mathbf{E}_2, w, k, \delta}(\beta, \gamma),$$

and what we learned above is that the geometric gluing map, which gives rise to a diffeomorphism between an open subset of this fiber product and $\tilde{\mathcal{M}}_{\mathbf{E}_{12}, zw, k, \delta}(\alpha, \gamma)$ is *orientation-preserving*, having oriented the latter via the recipe given by the analytic gluing map.

Suppose that W_1 is a cylinder. Then if $\overline{\mathcal{M}}'_{\mathbf{E}_1, z}(\alpha, \beta)$ is the ‘compactification’ of $\tilde{\mathcal{M}}_{\mathbf{E}_1, z}(\alpha, \beta)$ by broken trajectories,¹¹ quotienting by \mathbb{R} gives the usual $\overline{\mathcal{M}}_{\mathbf{E}_1, z}(\alpha, \beta)$ of unparameterized broken trajectories. We always orient this so that the diffeomorphism

$$\mathbb{R} \times \overline{\mathcal{M}}_{\mathbf{E}_1, z, k, \delta}(\alpha, \beta) \cong \overline{\mathcal{M}}'_{\mathbf{E}_1, z, k, \delta}(\alpha, \beta)$$

is orientation-preserving.

The notation in the following is slightly different than in Proposition 4.43, to allow for a more uniform discussion.

Proposition 5.7. *If (Y, E) is a 3-manifold equipped with regular perturbation, then if $\overline{\mathcal{M}}_{E, k, \delta}(\alpha, \beta)$ is the compactified moduli space of unparameterized flowlines, then Proposition 4.43 gives a decomposition of the boundary*

$$\partial \overline{\mathcal{M}}_{E, z, \pi}(\alpha, \beta) = \bigcup_{\substack{\gamma \in \mathfrak{C}_\pi \\ z_1 * z_2 = z}} \overline{\mathcal{M}}_{E, z_1, \pi}(\alpha, \gamma) \times_\gamma \overline{\mathcal{M}}_{E, z_2, \pi}(\gamma, \beta).$$

¹¹To make sense of this, the broken trajectories are not parameterized on each component, but rather the stratum corresponding to a k -broken trajectory is quotiented by the action of $\mathbb{R}^{k-1} \subset \mathbb{R}^k$, sitting inside as the subset with zero sum.

Suppose we orient $\Lambda_z(\alpha, \beta)$ using the analytic gluing map and fixed orientations of the three of $\Lambda_{z_1}(\alpha, \gamma)$, $\Lambda_{z_2}(\gamma, \beta)$, and \mathfrak{g}_γ . If d is the dimension of $\widetilde{\mathcal{M}}_{E, z_1, \pi}(\alpha, \gamma)$, then the boundary orientation on

$$\overline{\mathcal{M}}_{E, z_1, \pi}(\alpha, \gamma) \times_\gamma \overline{\mathcal{M}}_{E, z_2, \pi}(\gamma, \beta)$$

differs from the fiber product orientation by a factor of $(-1)^d$.

Now let (W, \mathbf{E}) be a cobordism equipped with regular perturbation. By Proposition 4.43, each compactified moduli space $\overline{\mathcal{M}}_{\mathbf{E}, z, \pi}$ has boundary given as the union

$$\begin{aligned} \partial \overline{\mathcal{M}}_{\mathbf{E}, z, \pi}(\alpha, \beta) = & \bigcup_{\substack{\gamma \in \mathcal{C}_{\pi_1} \\ z_1 * z_2 = z}} \overline{\mathcal{M}}_{E_1, z_1, \pi_1}(\alpha, \gamma) \times_\gamma \overline{\mathcal{M}}_{\mathbf{E}, z_2, \pi}(\gamma, \beta) \\ & \bigcup_{\substack{\zeta \in \mathcal{C}_{\pi_2} \\ z_1 * z_2 = z}} \overline{\mathcal{M}}_{\mathbf{E}, z_1, \pi}(\alpha, \zeta) \times_\zeta \overline{\mathcal{M}}_{E_2, z_2, \pi_2}(\zeta, \beta). \end{aligned}$$

This decomposes the boundary into two types of components: whether breaking occurs at the negative end or at the positive end.

Then if we orient $\Lambda_z(\alpha, \beta)$ using the analytic gluing map and fixed orientations of the three of $\Lambda_{z_1}(\alpha, \gamma)$, $\Lambda_{z_2}(\gamma, \beta)$, and \mathfrak{g}_γ , the orientation of

$$\overline{\mathcal{M}}_{E_1, z_1, \pi_1}(\alpha, \gamma) \times_\gamma \overline{\mathcal{M}}_{\mathbf{E}, z_2, \pi}(\gamma, \beta)$$

as a boundary stratum agrees with the orientation induced by the fiber product. If the dimension of $\overline{\mathcal{M}}_{\mathbf{E}, z_1, \pi}$ is d , then the orientation of

$$\overline{\mathcal{M}}_{\mathbf{E}, z_1, \pi}(\alpha, \zeta) \times_\zeta \overline{\mathcal{M}}_{E_2, z_2, \pi_2}(\zeta, \beta)$$

as a boundary stratum is $(-1)^d$ that of the orientation given as a fiber product.

Proof. The argument is no different from [KM07, Proposition 20.5.2] and [KM07, Proposition 25.1.1], respectively. We write $\widetilde{\mathcal{M}}_{E, z, \pi}^0(\alpha, \beta)$ for the quotient of the space of parameterized flowlines by the \mathbb{R} action, not compactified. Then in the first case, the sign arises from the orientation-preserving local homeomorphism

$$\widetilde{\mathcal{M}}_{E, z_1, \pi}^0(\alpha, \gamma) \times [0, \infty) \times_\gamma \times \widetilde{\mathcal{M}}_{E, z_2, \pi}^0(\gamma, \beta) \rightarrow \overline{\mathcal{M}}_{E, z, \pi}(\alpha, \beta);$$

to compare to the boundary orientation, we should commute the $[0, \infty)$ factor across $\widetilde{\mathcal{M}}_{E, z_1, \pi}^0(\alpha, \gamma)$, which has dimension $d - 1$, as above; the final sign $(-1)^d$ comes because we restrict to the negative boundary component $\{0\}$.

The argument in the case of a cobordism is identical, and the difference in signs arises because in the case of breakings along the negative end, we do not need to commute anything. \blacksquare

5.4. Canonical orientations for determinant lines. What remains is to find a recipe so that, given some choices for each α and β , and a choice depending only on the underlying cobordism (in the case that it is not just a cylinder), we are given natural orientations of each $\Lambda_z^W(\alpha, \beta)$. These must compose appropriately under gluing. If we change one of the choices for α then the sign of the orientation on each $\Lambda_z(\alpha, \beta)$ should change, and similarly with swapping the choice for β . If one changes the choice for W , then every orientation should change uniformly. This assignment of orientations is what we call in the title of this section ‘canonical orientations’.

To do this, we will ultimately require that the cobordism (or possibly some larger cobordism) admits a reducible connection. Already this places the demand that $\beta w_2(E) = 0$, but this is not enough to pin down orientations.

Supposing that \mathbf{A} is an $SO(2)$ -reducible, we have a splitting $Q_{\mathbf{A},\pi} \cong Q_{\eta,\pi} \oplus Q_{\theta,\pi}$. If we chose a complex structure on η , then the term $Q_{\eta,\pi}$ is complex linear with respect to this complex structure, which would give an isomorphism

$$\det(Q_{\mathbf{A},\pi}) \cong \det(Q_{\theta,\pi}),$$

thus reducing discussion to the case of the trivial connection (possibly perturbed), but a different choice of complex structure will give a different isomorphism. We need to ensure that a choice of such a complex structure is part of our data.

For the first time in this text we need to choose (for each 3-manifold Y and each cobordism W) not an $SO(3)$ -bundle over the manifold, but rather a $U(2)$ -bundle. An $SO(3)$ -bundle has a lift to a $U(2)$ -bundle if and only if $\beta w_2(E) = 0 \in H^3(-; \mathbb{Z})$. If $\tilde{\mathbf{E}}$ is a rank 2 complex vector bundle with fixed connection on its determinant line, then a reducible connection \mathbf{A} induces a splitting $\tilde{\mathbf{E}} \cong \eta \oplus \zeta$, where $\eta \otimes \zeta \cong \det(\tilde{\mathbf{E}})$ and η, ζ are complex line bundles. Thus for the trace-0 adjoint bundle we obtain an isomorphism $\mathfrak{g}_{\tilde{\mathbf{E}}} \cong \mathbb{R} \oplus (\eta \otimes \zeta^{-1})$, respecting the splitting induced by \mathbf{A} ; thus we have fixed a complex structure on the second component of the above splitting.

We begin this with some preliminaries.

Lemma 5.8. *Let (W, \mathbf{E}) be a complete Riemannian manifold equipped with a $U(2)$ -bundle with no boundary components and some number of cylindrical ends, modelled on either $(-\infty, 0] \times Y_i$ or $[0, \infty) \times Y_i$; in the former case we say Y_i is a negative end, in the latter a positive end. Suppose W is equipped with a perturbation which is regular on each end, and a choice of $\alpha_i \in \mathfrak{C}_\pi(Y_i)$ for each 3-manifold Y_i the ends are modelled on. Write $\tilde{\mathcal{B}}_{\mathbf{E},k,\delta}^e(\alpha)$ for the configuration space of $L_{k,\delta}^2$ connections, asymptotic to the α_i on the corresponding ends.*

Then $\pi_0 \mathcal{B}_{\mathbf{E},k,\delta}^e(\alpha) \cong \mathbb{Z}$, this isomorphism affine over $\pi_1 \mathcal{B}_{\mathbf{E}_i}^e \cong \mathbb{Z}$ for any end (Y_i, \mathbf{E}_i) .

If all of the α_i are trivial, then there is a unique component $z \in \pi_0 \mathcal{B}_{\mathbf{E},k,\delta}^e(\alpha)$ which supports a reducible connection.

Proof. The statement about components in the space of connections is little more than the classification of $U(2)$ -bundles over a compact 4-manifold with boundary with fixed isomorphism class on the boundary in terms of their first Chern class $c_1 \in H^2(W; \mathbb{Z})$ and Pontryagin class $p_1 \in \Lambda \subset H^4(W, \partial W; \mathbb{R})$. Here Λ is a subset affine over $8\pi^2\mathbb{Z}$, the latter defined by a curvature integral with respect to some connection with fixed boundary components. Gluing in the nontrivial positive (determined by the same curvature integral) generator of $\pi_1 \mathcal{B}_{\mathbf{E}_i}^e$ increases this by $8\pi^2$, as expected. We denote this operation as $z \mapsto z + 1$.

The fact that one and only one component supports a reducible connection follows from the enumeration of reducible components in Proposition 1.7 (or rather, a version allowing more ends). \blacksquare

Lemma 5.9. *In the situation above, there is a canonical isomorphism $\Lambda_z^W(\alpha) \cong \Lambda_{z+1}^W(\alpha)$, compatible with the gluing maps.*

Suppose all of the α_i are trivial, and z is the corresponding component supporting a reducible connection, and the perturbation is zero on the ends of W . Then if we

write

$$Q_\theta^W : \Omega_{k,\delta}^1(W; \mathbb{R}) \rightarrow \Omega_{k,\delta}^{2,+}(W, \mathbb{R}) \oplus \Omega_{k,\delta}^0(W; \mathbb{R}),$$

for (d^+, d^*) , we have an isomorphism $\Lambda_z^W(\alpha) \cong \Lambda^W(\theta)$, where the latter is the set of orientations of Q_θ^W . This is true for any choice of Sobolev indices on the various ends, so long as they never coincide with eigenvalues of $\widehat{\text{Hess}}_{\alpha_i}$ or $\widehat{\text{Hess}}_\theta$.

Proof. The first isomorphism follows by gluing charge-1 instantons on S^4 into the cobordism. By the assumption that the perturbation is compactly supported in W , we may modify $Q_{\mathbf{A},\pi}^W$ by a homotopy through Fredholm operators to $Q_{\mathbf{A}}^W$.

We discussed the second claim above for $SO(2)$ -reducible connections on $SO(3)$ -bundles. When we have a lift of the $SO(3)$ -bundle to a $U(2)$ -bundle, if the reducible corresponds to some splitting $\mathbf{E} \cong \mathbb{R} \oplus \xi$, then we obtain a canonical complex orientation on the complement of the \mathbb{R} -factor, giving the desired claim. For $SO(3)$ -reducible connections, instead we have $Q_{\mathbf{A}}^W = Q_\theta^W \otimes \mathfrak{g}$, and so an orientation of one canonically induces an orientation of the other (supposing an orientation of \mathfrak{g} is given). \blacksquare

Thus we drop z from the notation $\Lambda^W(\alpha, \beta)$.

If one of α or β is irreducible, we cannot appeal to reducible connections on the cobordism as above. So we need to 'cap off' the two ends, as appropriate. Being precise about this is a little intricate, as described in the following definition.

Definition 5.2. *Let $(W, \tilde{\mathbf{E}})$ be a cobordism $(Y_1, E_1) \rightarrow (Y_2, E_2)$ equipped with $U(2)$ -bundle, a choice of perturbation on the ends (possibly zero, allowing for it to fail to be regular), and a choice of critical orbits α, β on the negative and positive end.*

The symbol $\Lambda^W(\alpha, \beta)$ always means the usual two-element set of orientations of $\det(Q_{\mathbf{A},\pi})$, where \mathbf{A} is a connection asymptotic to α and β on the appropriate ends, and $Q_{\mathbf{A},\pi}$ has domain $\Omega_{k,\delta}^1$ and codomain $\Omega_{k-1,\delta}^{2,+} \oplus \Omega_{k-1,\delta}^0$. (If one of these terms α, β is trivial, we sometimes do not write the corresponding θ .) When unadorned, the symbol $\Lambda^W(\theta)$ denotes the same for the trivial connection on the trivial bundle.

If $(W, \tilde{\mathbf{E}})$ is a cobordism from some (Y_-, triv) to (Y_1, E_1) , for which the perturbation is zero on the incoming end, we say W is an incoming cap; if (Y_2, E_2) is the incoming end and (Y_+, triv) is the outgoing end, with zero perturbation on the outgoing end, we call $(W, \tilde{\mathbf{E}})$ an outgoing cap, and otherwise call W an intermediate cobordism.

If W is an incoming, intermediate, or outgoing cobordism respectively, we write $\Lambda_-^W(\theta)$, $\Lambda_i^W(\theta)$, and $\Lambda_+^W(\theta)$ for the two-element sets of orientations of the following three operators:

$$\begin{aligned} Q_\theta^- &: \Omega_{\delta,-\delta}^1 \rightarrow (\Omega_{\delta,-\delta}^{2,+} \oplus \Omega_{\delta,-\delta}^0) \\ Q_\theta^i &: \Omega_{-\delta}^1 \rightarrow (\Omega_{-\delta}^{2,+} \oplus \Omega_{-\delta}^0)_{-\delta,\delta} \\ Q_\theta^+ &: \Omega_{-\delta,\delta}^1 \rightarrow (\Omega_{-\delta,\delta}^{2,+} \oplus \Omega_{-\delta,\delta}^0) \end{aligned}$$

Here the subscripts indicate the Sobolev weights on the ends; if only one subscript appears, it is the Sobolev weight on both ends, and if two appear, they are the Sobolev weights on the negative and positive ends, in that order.

Whenever we have two-element sets Λ and Λ' , we write $\Lambda\Lambda'$ to denote $\Lambda \times_{\mathbb{Z}/2} \Lambda'$, the $\mathbb{Z}/2$ action the canonical free involution on both sets.

The point of the choices of Sobolev indices is that these operators (for the unperturbed trivial connection) enjoy an additivity property, immediate from spectral flow arguments.

Lemma 5.10. *Let W_0, W, W_1 denote incoming, intermediate, and outgoing cobordisms equipped with $U(2)$ -bundles, as above. If \hat{W} denotes the composite of these, then we have canonical isomorphisms*

$$\Lambda^{W_0}(\alpha)\Lambda(\mathfrak{g}_\alpha)\Lambda^W(\alpha, \beta)\Lambda(\mathfrak{g}_\beta)\Lambda^{W_1}(\beta) \cong \Lambda^{\hat{W}}(\theta)$$

and

$$\Lambda_-^{W_0}(\theta)\Lambda_i^W(\theta)\Lambda_+^{W_1}(\theta) \cong \Lambda^{\hat{W}}(\theta).$$

Given a 3-manifold with $U(2)$ -bundle (Y, \tilde{E}) , given any two incoming caps $(W_0, \tilde{\mathbf{E}})$ and $(W'_0, \tilde{\mathbf{E}}')$, the sets $\Lambda(\mathfrak{g}_\alpha)\Lambda^W(\alpha)\Lambda_-^W(\theta)$ and $\Lambda(\mathfrak{g}_\alpha)\Lambda^{W'}(\alpha)\Lambda_-^{W'}(\theta)$ are canonically isomorphic. The appropriate modification is true for the outgoing caps, as well.

Proof. These follow from the usual gluing of operators on manifolds with cylindrical ends (for which the Sobolev weights on the ends match up appropriately, or there is a corresponding term to account for the spectral flow). The composite of the terms in the first displayed equation are naturally isomorphic, rather, to the orientation set of a nontrivial connection over \hat{W} , but as above we may reduce this to the case of a reducible connection and then to the trivial connection, as above.

Once we have this, using twice the orientation-reversal $\overline{W_0}$ as a positive cap, using these canonical isomorphisms we further get canonical isomorphisms

$$\Lambda(\mathfrak{g}_\alpha)\Lambda^{W_0}(\alpha)\Lambda_-^{W_0}(\theta) \cong \Lambda^{\overline{W_0}}(\alpha)\Lambda_+^{\overline{W_0}}(\theta) \cong \Lambda(\mathfrak{g}_\alpha)\Lambda^{W'_0}(\alpha)\Lambda_-^{W'_0}(\theta),$$

as desired. ■

The collection of isomorphisms above between the different possible 2-element sets $\Lambda(\mathfrak{g}_\alpha)\Lambda^W(\alpha)\Lambda_-^W(\theta)$ determine an equivalence relation on the set

$$\bigsqcup_{\substack{(W, \tilde{\mathbf{E}}) \\ \text{W incoming cap}}} \Lambda(\mathfrak{g}_\alpha)\Lambda^{W_0}(\alpha)\Lambda_-^{W_0}(\theta),$$

whose quotient is a 2-element set we write $\Lambda_-(\alpha)$. Similarly the 2-element set $\Lambda_+(\alpha)$ is a quotient of various copies of $\Lambda_+^{W_1}(\theta)\Lambda^{W_1}(\alpha)\Lambda(\mathfrak{g}_\alpha)$ as W_1 varies over positive caps.

We thus have from the first part of the above lemma a natural isomorphism

$$\Lambda_-(\alpha)\Lambda_i^W(\theta)\Lambda_+(\beta) \cong \Lambda^W(\alpha, \beta).$$

This was essentially our goal, though we are not quite done.

Lemma 5.11. *There is a canonical isomorphism $\Lambda_+(\beta) \cong \Lambda(\mathfrak{g}_\beta)\Lambda_-(\beta)$.*

Proof. Writing this out explicitly, we have chosen incoming and outgoing caps W_0, W_1 , we are trying to construct a trivialization of

$$\Lambda_+^{W_1}(\theta)\Lambda^{W_1}(\beta)\Lambda(\mathfrak{g}_\beta)\Lambda^{W_0}(\beta)\Lambda_-^{W_0}(\theta),$$

here exploiting the isomorphism between orientations on a vector space and their dual to reverse the order of the first few terms. By gluing the middle three terms, we obtain an isomorphism to $\Lambda_+^{W_1}(\theta)\Lambda^{\hat{W}}(\theta)\Lambda_-^{W_0}(\theta)$; splitting up the middle term, this is isomorphic to $\Lambda_+^{W_1}(\theta)\Lambda_-^{W_0}(\theta)\Lambda_+^{W_1}(\theta)\Lambda_-^{W_0}(\theta)$. As this takes the form Λ^2 for some 2-element set Λ , it is canonically trivial. ■

Therefore, if we simply write $\Lambda(\alpha) := \Lambda_+(\alpha)$, we have above found a canonical isomorphism $\Lambda(\mathfrak{g}_\alpha)\Lambda(\alpha)\Lambda_i^W(\theta)\Lambda(\beta) \cong \Lambda^W(\alpha, \beta)$; then we obtain the gluing isomorphism

$$\Lambda^W(\alpha, \beta)\Lambda(\mathfrak{g}_\beta)\Lambda^{W'}(\beta, \gamma) \cong \Lambda^{W \circ W'}(\alpha, \gamma)$$

simply by paing off adjacent like terms.

The last thing to be clear about is precisely what the middle orientation set $\Lambda_i^W(\theta)$ is.

Definition 5.3. *A homology orientation of W , a cobordism with cylindrical ends and incoming end $(-\infty, 0] \times Y_1$, is an orientation of the real vector space $H^1(W) \oplus H^{2,+}(W) \oplus H^1(Y_1)$.*

Because W and Y are connected, note that we have a canonical isomorphism $H^0(W) \oplus H^0(Y_1) \cong \mathbb{R}^2$, and in particular carries a canonical orientation. Because we have

$$\begin{aligned} \ker(Q_\theta^i) &\cong H^1(W) \\ \text{coker}(Q_\theta^i) &\cong H^{2,+}(W) \oplus H^1(Y_1) \oplus H^0(W) \oplus H^0(Y_1), \end{aligned}$$

a homology orientation canonically induces an orientation of $\Lambda_i^W(\theta)$.

Remark 5.1. Homology orientations have a natural associative composition law, given by the index-gluing described above. However, the composition law may be described explicitly, without passing to a discussion of Fredholm operators. An explicit formula for this law was found in [Sca15], where an explicit understanding was crucial to discuss the differentials in a spectral sequence to Khovanov homology. We do not need this here, and so will not discuss Scaduto's results in any more detail.

We assemble the content of this section into a proposition.

Proposition 5.12. *For any 3-manifold Y equipped with $U(2)$ -bundle \tilde{E} and regular perturbation π , there are canonical 2-element sets $\Lambda(\alpha)$ for each critical orbit $\alpha \in \mathfrak{C}_\pi$, and we have the canonical isomorphism $\Lambda(\alpha, \beta) \cong \Lambda(\mathfrak{g}_\alpha)\Lambda(\alpha)\Lambda(\beta)$. Therefore, the moduli space $\tilde{\mathcal{M}}_z(\alpha, \beta)$ may be given an orientation if we choose an element of $\Lambda(\alpha)$ and $\Lambda(\beta)$, and an orientation of the orbit α . Choosing the other element of either set will negate the orientation of $\tilde{\mathcal{M}}(\alpha, \beta)$.*

If one fixes a choice of element of each of $\Lambda(\alpha), \Lambda(\beta), \Lambda(\gamma)$, and an orientation of α , then the orientation these induce on $\overline{\mathcal{M}}_{E,\pi}(\alpha, \gamma) \times_\gamma \overline{\mathcal{M}}_{E,\pi}(\gamma, \beta)$ via the fiber product differs from the orientation induced as the fiber of a component of $\partial \overline{\mathcal{M}}_{E,\pi}^{\text{fib}}(\alpha, \beta)$ by a sign of $(-1)^d$, where d is the dimension of $\tilde{\mathcal{M}}_{E,\pi}(\alpha, \gamma)$.

Let $(W, \tilde{\mathbf{E}})$ be a cobordism from (Y_1, E_1) to (Y_2, E_2) , equipped with a regular perturbation π that restricts to π_i on the ends. Then a choice of element of each $\Lambda(\alpha)$ and $\Lambda(\beta)$, as well as a homology orientation of W and an orientation of α , give rise to an orientation of $\overline{\mathcal{M}}_{\mathbf{E},\pi}(\alpha, \beta)$; swapping any one of these elements will negate this orientation.

The boundary components of $\overline{\mathcal{M}}_{\mathbf{E},\pi}(\alpha, \beta)$ arise in two pieces: those of the form

$$\overline{\mathcal{M}}_{E_1, z_1, \pi_1}(\alpha, \gamma) \times_\gamma \overline{\mathcal{M}}_{\mathbf{E}, z_2, \pi}(\gamma, \beta),$$

where $\gamma \in \mathfrak{C}_{\pi_1}$, and those of the form

$$\overline{\mathcal{M}}_{\mathbf{E}, z_1, \pi}(\alpha, \zeta) \times_\zeta \overline{\mathcal{M}}_{E_2, z_2, \pi_2}(\zeta, \beta),$$

where $\zeta \in \mathfrak{C}_{\pi_2}$.

Given a choice of element of each $\Lambda(\alpha), \Lambda(\beta), \Lambda(\gamma)$, and $\Lambda(\zeta)$ as above, as well as a choice of homology orientation on W and orientation of α , the orientations on $\partial\overline{\mathcal{M}}$ and on the fiber products agree in the first case; they differ by a sign of $(-1)^{d_W}$ in the second case, where d_W is the dimension of $\widetilde{\mathcal{M}}_{\mathbf{E}, z_1, \pi}(\alpha, \zeta)$.

Similarly, if S is a 1-parameter family of metrics on W abutting to a broken metric, the decomposition of the boundary of $\overline{\mathcal{M}}_{\mathbf{E}, z, S}(\alpha, \beta)$ of Proposition 4.45 gives the boundary components orientations which differ by the orientation induced by analytic gluing by a factor of $1, (-1)^{d_1}, (-1)^{d_{12}}$, and -1 , respectively (the last term coming from the boundary orientation of $\{0\} \subset [0, 1]$), where d_1 is the dimension of $\widetilde{\mathcal{M}}_{\mathbf{E}, z_1}(\alpha, \gamma)$, while d_{12} is the dimension of $\widetilde{\mathcal{M}}_{\mathbf{E}, z_1, S}(\alpha, \eta)$.

6. FLOER HOMOLOGY

6.1. Geometric homology. It will be useful in what follows to have a chain complex computing the singular homology of a smooth manifold X , whose generators are smooth maps from oriented manifolds with corners (as opposed to continuous maps from simplices). Fix a principal ideal domain R .

The definitions below are modeled on [Lin18, Chapter 4.1], which in turn follow the simpler [Lip14]. For technical reasons, we must introduce the notion of strong δ -chain. The reader will not be led astray in what follows by pretending every occurrence of “strong δ -chain” means “compact smooth manifold with corners”; if we could achieve that level of smoothness on the instanton moduli spaces, this sequence of definitions would not be necessary to set up our homology theory.

Definition 6.1. *A d -dimensional strong δ -chain is a compact topological space P with a stratification*

$$P^d \supset P^{d-1} \supset \dots \supset P^0 \supset \emptyset$$

by closed subsets, so that $P^e \setminus P^{e-1}$ is decomposed as a finite disjoint union of smooth manifolds of dimension e , written $\sqcup_{i=1}^{m_e} M_i^e$; the top stratum $P \setminus P^{d-1}$ has only one open face $M_1^d = P \setminus P^{d-1}$ in its decomposition. (Note that M_i^e need not be connected!) We denote the closure of any one of these manifolds an e -dimensional face, and write it as Δ . We write Δ° for the interior of a face (one of the manifolds M_i^e in the disjoint union $P^e \setminus P^{e-1}$) and call it an open face.

We demand that whenever a codimension e face Δ_0 is contained in a codimension $(e-2)$ face Δ_2 , there are exactly two codimension $(e-1)$ faces Δ' with $\Delta_0 \subset \Delta' \subset \Delta_2$.

Whenever $\Delta_1 \subset \Delta_2$, we assign a set $n(\Delta_1, \Delta_2)$ (which we will write as n , or n_{12} when the faces are not implicit) with cardinality $\dim \Delta_1 - \dim \Delta_2$, an open neighborhood $\Delta_1^\circ \subset W(\Delta_1, \Delta_2) \subset \Delta_2$, and a map $r : W \rightarrow [0, \varepsilon]^n$ so that $\Delta_1^\circ = r^{-1}(0)$.

We also assign a space $W \hookrightarrow EW(\Delta_1, \Delta_2)$, where EW is a topological manifold with corners and a smooth structure on each stratum of dimension $d + m(\Delta_1, \Delta_2)$, for which the map from W is smooth on each stratum, and equipped with a map $\tilde{r} : EW(\Delta_1, \Delta_2) \rightarrow [0, \varepsilon]^n$ extending $r : W \rightarrow [0, \varepsilon]^n$. We demand the following.

- There is a vector bundle $V(\Delta_1, \Delta_2) \rightarrow EW$ of rank $m(\Delta_1, \Delta_2)$, and a section σ of V , smooth and transverse to the zero section on each stratum, so that $\sigma^{-1}(0) = W$.
- Write EW^k for the inverse image of $R_k \subset [0, \varepsilon]^n$, the set of points for which exactly k coordinates are zero; we demand that the map $EW^k \rightarrow R_k$ is a smooth submersion.
- \tilde{r} is a fiber bundle projection. The restriction of σ to $\tilde{r}^{-1}(0)$ is transverse to the zero section, whose fiber above zero is Δ_1° . In particular, EW is diffeomorphic to $[0, \varepsilon]^{n_{12}} \times \tilde{r}_{12}^{-1}(0)$, and $\tilde{r}_{12}^{-1}(0)$ is diffeomorphic to a neighborhood of the zero section in the restriction of V to Δ_1° .

These are compatible in the following sense. Associated to a sequence of inclusions $\Delta_1 \subset \Delta_2 \subset \Delta_3$ we have an inclusion of sets $n(\Delta_1, \Delta_2) \hookrightarrow n(\Delta_1, \Delta_3)$, an embedding $EW(\Delta_1, \Delta_2) \hookrightarrow EW(\Delta_1, \Delta_3)$, as well as an embedding of vector bundles $V(\Delta_1, \Delta_2) \hookrightarrow V(\Delta_1, \Delta_3)$ covering this; the section σ_{12} is the restriction of σ_{13} . The map $n_{12} \rightarrow n_{13}$ of sets induces a stratum-preserving embedding $[0, \varepsilon]^{n_{12}} \hookrightarrow [0, \varepsilon]^{n_{13}}$ so that, with respect to this inclusion, \tilde{r}_{12} is the restriction of \tilde{r}_{13} . Finally, we

demand that with respect to these embeddings,

$$EW(\Delta_1, \Delta_2) = \tilde{r}_{13}^{-1}([0, \varepsilon]^{n_{12}}).$$

This complicated definition is in fact more or less forced on us by a few simple requirements. First, our chains should include compact topological manifolds with corners with a smooth structure on each stratum (satisfying the combinatorial condition). Second, they should be closed under transverse intersections (in particular, the inverse image of a regular value should be a chain). This is the property we need to ensure that excision holds in our coming homology theory based on maps from strong δ -chains.

Already we run into trouble: because the smooth structure on the strata do not interact, all we see is that the transverse intersection of two topological manifolds with smooth structures on each stratum is a space *stratified* by smooth manifolds. For instance, it is not hard to construct a θ -shaped graph as the zero set of a continuous function f on $[0, 1] \times \mathbb{R}$ which is smooth on each stratum and for which 0 is a regular value on each stratum.

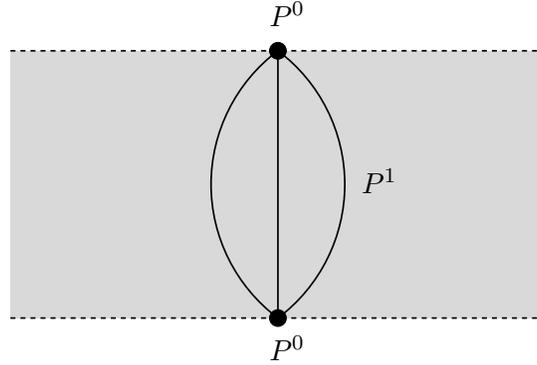


FIGURE 1. The δ -chain $P \subset [0, 1] \times \mathbb{R}$ in this figure is the zero set of a continuous function $f : [0, 1] \times \mathbb{R} \rightarrow \mathbb{R}$ which is smooth and transverse to zero on each stratum. $P^0 = f^{-1}(0) \cap \{0, 1\} \times \mathbb{R}$ is two points, while $P^1 \setminus P^0$ is three open arcs.

This is what forces us to think of spaces equipped with “local thickenings” which actually *are* manifolds. In the case of the θ -shaped graph, the three open arcs of the graph form the unique top-dimensional open face, whose local thickening is a small open neighborhood of these in the strip; the vector bundle is the trivial line bundle, and σ is the function f . The stratum P^0 consists of two points, for which the neighborhood $W(\Delta_0, \Delta_1)$ is a small neighborhood of one of these points (which looks like a chicken foot). The local thickenings are small 2-dimensional neighborhoods $[0, \varepsilon] \times (-t, t)$ and the map \tilde{r} is projection to $[0, \varepsilon]$. Again, the vector bundle is the trivial line bundle and $\sigma = f$.

The third requirement is that faces of a chain should also be chains, and thus we are forced to define local thickenings not just for the whole space, but also for inclusions $\Delta_1 \subset \Delta_2$. Finally, it is crucial that the combinatorial boundary operator satisfies $\partial^2 = 0$; this is why we demand the combinatorial condition on faces. At the level of manifolds with corners, it says that a bigon is a manifold with corners, but a teardrop (unigon) is not.

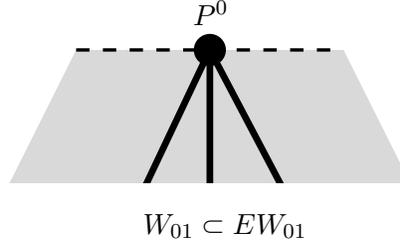


FIGURE 2. The stratum $P^0 \subset P^1$ has a neighborhood $P^0 \subset W(P^0, P^1) \subset P^1$ (the ‘chicken feet’), along with a further neighborhood $W(P^0, P^1) \subset EW(P^0, P^1)$ so that $EW(P^0, P^1)$ is a smooth surface with corners and $Ef_{01} : EW \rightarrow \mathbb{R}$ is a smooth map with $W_{01} = Ef_{01}^{-1}$. Here, we may take $EW(P^0, P^1)$ to a neighborhood of W_{01} inside of $[0, 1] \times \mathbb{R}$, and Ef_{01} to simply be $f|_{EW}$. In this figure we depict the top half of EW_{01} (which should include a neighborhood of both the top and bottom point).

Essentially, the definition of strong δ -chain is one possible approach to formalize a space with local charts modelling a neighborhood of each point as the zero set of a smooth function (smooth section of a vector bundle) on a compact topological manifold with smooth structures on each stratum, and compatibility relations between the charts.

The notion of strong δ -chain here is a slight modification of the notion of δ -chain given in [KM07, Definition 24.7.1] and [Lin18, Definition 4.1.1].

There are two differences. First, for a strong δ -chain the set $W \subset EW$ is cut out as the zero set of a transverse (on each stratum) section of a vector bundle over EW ; for a δ -chain, instead the map is simply to some vector space \mathbb{R}^k . The definition of δ -chain is attempting to formalize the kind of spaces that appear as the monopole moduli spaces, while the definition of strong δ -chain is meant to capture all transverse intersections of compact topological manifolds with corners and smooth structures on each stratum. These transverse intersections may be written as the inverse image of a diagonal $\Delta \subset Y \times Y$; the vector bundle in question is $N\Delta$, and if Y has interesting topology this bundle may be nontrivial, so we need to include zero sets of sections of vector bundles. Locally near any point, these are the same notion, but they may not be the same in a neighborhood of an entire stratum.

Second, the map $EW \rightarrow \mathbb{R}^k$ in the definition of δ -chain is required to satisfy certain positivity properties, depending on the values of $\tilde{r}(EW) \subset [0, \varepsilon]^n$; we do not make any such demand for strong δ -chains. This requirement originates from the boundary-obstructedness phenomenon in monopole moduli spaces, which we do not encounter.

Definition 6.2. *Suppose we are given the following data.*

- For each pair of faces $\Delta_1 \subset \Delta_2$, open subsets $U_{12} \subset W_{12}$ and $U'_{12} \subset W'_{12}$, as well as open subsets EU_{12}, EU'_{12} of the extensions,
- homeomorphisms $\varphi_{12} : U_{12} \rightarrow U'_{12}$, $\tilde{\varphi}_{12} : EU_{12} \rightarrow EU'_{12}$, which are diffeomorphisms over each stratum,

- an isomorphism of vector bundles $V_{12} \cong V'_{12}$ over these open subsets EU and EU' , covering the diffeomorphism $\tilde{\varphi}_{12}$,

all of which respect all the associated structure and are compatible for triples $\Delta_0 \subset \Delta_1 \subset \Delta_2$. Then we say the two strong δ -chains are germ equivalent.

Suppose now that we are given two strong δ -chains with the same P, W , and maps r , but different thickenings $EW, E'W$. Suppose we are given the following.

- A vector bundle F_{12} over EW_{12} for each inclusion of faces,
- an embedding $F_{12} \rightarrow F_{13}$ covering the embedding $EW_{12} \rightarrow EW_{13}$ corresponding to any triple of faces,
- a diffeomorphism $\varphi_{12} : E'W_{12} \cong F_{12}$ (the total space of the vector bundle),
- an isomorphism $\psi_{12} : \pi^*(V_{12} \oplus F_{12}) \cong V'_{12}$ covering φ_{12} , taking $\sigma \oplus 0$ to $\sigma' \oplus 0$,

all of which are compatible for triples of faces, then we say that the strong δ -chains are stably equivalent.

This definition fits with the above intuition for strong δ -chains as spaces equipped with a special kind of chart modeling them as a zero set: if we pass to smaller open subsets containing the zero set, or stabilize, this should still present a perfectly good, equivalent chart at a point.

Next, we will show the important property that strong δ -chains are, in a sense, closed under fiber products. Stable equivalence was partly introduced as a means to this end. To start, we should define what a map from a strong δ -chain is.

Definition 6.3. *Suppose X is a smooth manifold, and P is a strong δ -chain. We say a map $f : P \rightarrow X$ is a continuous map $f : P \rightarrow X$ on the underlying topological space which is smooth on each stratum, and for each inclusion of faces $\Delta_1 \subset \Delta_2$, an extension $Ef_{12} : EW_{12} \rightarrow X$ which is smooth on each stratum. Given a sequence of faces $\Delta_0 \subset \Delta_1 \subset \Delta_2$, recall that $EW_{02} \hookrightarrow EW_{12}$; we demand that Ef_{02} is the restriction of Ef_{12} .*

If $f_1 : P_1 \rightarrow X$ and $f_2 : P_2 \rightarrow X$ are two maps from strong δ -chains to X , we say that they are transverse if they are transverse on each stratum, and for any pair of faces in P_1 and any pair of faces in P_2 , the extensions Ef_1 and Ef_2 are transverse on each stratum in a neighborhood of $W_1 \times W_2$ in $EW_1 \times EW_2$.

We say that a map f is submersive if Ef_{12} is a submersion on each stratum.

Note that the definition of submersive says nothing about the underlying map f . For a map to be submersive EW has to be quite large, but we may do this at the cost of a stabilization; this is why we introduced stable equivalence above.

Lemma 6.1. *Let $f : P \rightarrow X$ be a map from a strong δ -chain to a smooth manifold X without boundary. Then f extends to a submersive map from the stabilization of P by f^*TX (that is, the same underlying space P and open sets W but $E'W_{12}$ is diffeomorphic to a neighborhood of the zero section in Ef_{12}^*TX).*

Proof. Equip X with a complete Riemannian metric. Then $E'f_{12}(p, v)$ is $\exp(Ef_{12}(p), v)$, the exponential of the tangent vector $v \in T_{Ef(p)}$ using the metric on X . This map is submersive so long as we restrict $E'f_{12}$ to a neighborhood of EW_{12} , because the derivative of the exponential map at $0 \in T_{Ef(p)}$ is the identity. ■

Proposition 6.2. *Suppose $f_i : P_i \rightarrow X$ are transverse maps from germ equivalence classes of strong δ -chains P_1, P_2 to a complete Riemannian manifold X , and suppose*

that f_2 is submersive. Then $P_1 \times_X P_2$ has the natural structure of a germ equivalence class of strong δ -chain.

Proof. The compact space $P_1 \times_X P_2$ is still stratified by smooth manifolds (the fiber products of the original open faces). For a pair of inclusions $\Delta_1^i \subset \Delta_2^i$ of faces of P_i , the open set W corresponding to

$$\Delta_1^1 \times_X \Delta_1^2 \subset \Delta_2^1 \times_X \Delta_2^2$$

is the fiber product of the corresponding open sets. EW is the fiber product $EW_1 \times_X EW_2$; this is still a topological manifold with corners and a smooth structure on each stratum by the assumption that Ef_2 is a submersion on each stratum. All of the additional structures (the vector bundle V , the section σ , the map r , and so on) is given as the fiber product of the structures on the EW_i ; for instance the vector bundle V is $V_1 \times_X V_2$, whose fibers above $EW_1 \times_X EW_2$ are isomorphic to the direct sum of fibers $V_1 \oplus V_2$. ■

To use coefficients other than $\mathbb{Z}/2$, we must introduce a notion of orientation.

Definition 6.4. *Let P be a strong δ -chain.*

Suppose we have an orientation on the top stratum of P , and for a codimension 1 face $\Delta_1 \subset P$, suppose we have oriented the open manifold Δ_1° . We may find an isomorphism $\det(T EW_{12}) \cong \det V_{12}$ above W_{12}° using the fiber bundle isomorphism $EW \cong [0, \varepsilon) \times \tilde{r}_{12}^{-1}(0)$, using that $\tilde{r}_{12}^{-1}(0)$ is isomorphic to a neighborhood of the zero section of the restriction of V_{12} to Δ_1° , and that we have oriented Δ_1° . Then

$$W_{12}^\circ := r_{12}^{-1}(0, \varepsilon)^n = \tilde{r}_{12}^{-1}(0, \varepsilon)^n \cap \sigma^{-1}(0)$$

may be oriented as the zero set of the section $EW_{12}^\circ \rightarrow V_{12}$ using the above isomorphism of determinant bundles. This gives an orientation on an open subset of $P^d \setminus P^{d-1}$. If the same as the orientation already given, we say that Δ_1 has the boundary orientation; we express that this is possible for some orientation on Δ_1 by saying that Δ_1 is consistently orientable with respect to P .

For any sequence of codimension 1 faces $\Delta_k \subset \dots \subset \Delta_1 \subset P$, we may thus inductively define the notion of any sequence being consistently orientable with respect to P .

We say that P is an oriented δ -chain if P is equipped with an orientation on its top stratum $P^d \setminus P^{d-1}$, satisfying the following conditions:

- *Any sequence of codimension 1 faces of P is consistently orientable with respect to P ,*
- *Given a codimension 2 inclusion $\Delta_0 \subset \Delta_2$, there are two intermediate faces Δ_1 and Δ'_1 ; we demand that if we fix an orientation on Δ_2 (induced from P by some sequence of codimension 1 faces as above), then the boundary orientation that arises on Δ_0 from the sequence $\Delta_0 \subset \Delta_1 \subset \Delta_2$ is negative that from the sequence $\Delta_0 \subset \Delta'_1 \subset \Delta_2$.*

The definition of boundary orientation above is relevant for the following calculation for strong δ -chains. This is proved for δ -chains as [KM07, Theorem 21.3.2]; instead of subdividing our strong δ -chains into pieces that may be given the structure of δ -chains in the sense of Kronheimer-Mrowka, instead we spell out their proof in simpler language as an exercise in understanding the definition of strong δ -chains. The proof follows similar lines as a discussion in [SS10, Section 2.3].

Lemma 6.3. *If P is a 1-dimensional oriented strong δ -chain, then P^0 consists of a finite number of oriented points whose count, signed by the boundary orientation, is equal to zero.*

Proof. If $x \in P^0$, there is an open subset $W_x \subset P$ with a map to $[0, \varepsilon)$, as well as a larger space $EW_x \cong [0, \varepsilon) \times V_x$ so that $W_x \hookrightarrow EW_x$, with x mapping to $(0, 0)$, and W_x° mapping into $(0, \varepsilon) \times V_x$; the neighborhood $W_x \subset P$ is the zero set of a map $\sigma : [0, \varepsilon) \times V_x \rightarrow V_x$ which is smooth and transverse to zero on each stratum. Write $\sigma_t(x) = \sigma(t, x)$. That σ is transverse to zero on the boundary implies that $\sigma^{-1}(0) \cap (\{0\} \times V_x)$ is discrete (and in particular discrete inside W). In particular, P^0 is a discrete set in the compact space P , so it is finite. For convenience, after passing to a possibly smaller neighborhood and reparameterizing, we assume $\sigma_0 = \text{Id}$ in our coordinates.

We also see that σ does not vanish on a sufficiently small sphere $\{0\} \times S_\delta(V_x)$ in the boundary of the local thickening near x ; the inclusion map from this sphere to $V_x \setminus \{0\}$ is degree 1 (regardless of how we orient V_x , as long as we orient these compatibly). Extending this to a disc D_x in $[0, \varepsilon) \times V_x$ which only intersects $\{0\} \times V_x$ in $S_\delta(V_x)$ and is transverse to $\sigma^{-1}(0)$, we see that $D_x \cap \sigma^{-1}(0)$ is a finite set of points whose oriented sum is 1, orienting them as $\det d(\sigma_t)(y)$ at an intersection point (t, y) .

D_x bounds a ball one dimension larger; deleting the part of W_x contained in the interior of this ball for each x , what is left of P^1 is a compact oriented 1-manifold L with boundary $\bigsqcup_{x \in P^0} (D_x \cap \sigma^{-1}(0))$. The boundary orientation convention is that a point (t, y) in this intersection is oriented positively as the boundary of L if the sign of $\det(d\sigma_t)(y)$ agrees with the sign of x . In particular, the signed sum of points in ∂L is

$$\sum_{x \in P^0} \text{sgn}(x) \# (D_x \cap \sigma^{-1}(0)) = \sum_{x \in P^0} \text{sgn}(x),$$

where the signed count is as in the previous paragraph. The signed count of points in the boundary of an oriented compact 1-manifold is zero, and thus the desired count is zero. \blacksquare

Now that these objects have been introduced, we can introduce the *geometric chain complex* computing singular homology. First we will write down the generating set.

Definition 6.5. *Let X be a smooth manifold. Fix once and for all a Hilbert space H equipped with a Hilbert bundle V_H .*

A basic chain of degree d in X is a stable germ equivalence class of maps $\sigma : P \rightarrow X$, where $P \subset H$ is a compact connected subspace equipped with the structure of an oriented strong δ -chain of dimension d , whose thickenings EW_{12} are compatibly embedded subspaces of H , and whose vector bundles V_{12} are subbundles of V_H .

Two basic chains $\sigma_i : P_i \rightarrow X$ are isomorphic if there is an orientation-preserving homeomorphism $f : P_1 \rightarrow P_2$ which is a diffeomorphism on each stratum, with $\sigma_2 f = \sigma_1$, and an extension for each $\Delta_1 \subset \Delta_2$ to $Ef_{12} : EW_{12}^1 \rightarrow EW_{12}^2$ which are diffeomorphisms on each stratum compatible with restriction. A basic chain $\sigma : P \rightarrow X$ is achiral if it is isomorphic to itself with the opposite orientation, and chiral if it is not achiral. A basic chain $\sigma : P \rightarrow X$ has small image if there is map $g : Q \rightarrow X$ from a δ -chain Q of strictly smaller dimension.

Write $\tilde{C}_*(X; R)$ for the graded vector space freely generated in degree d by the basic chains of degree d . One step closer to the desired chain complex, there is a quotient $\overline{C}_*(X; R)$ obtained from the relations

$$-[\sigma : P \rightarrow X] = [\sigma : \overline{P} \rightarrow X],$$

and if $2 \neq 0 \in R$ we further set any achiral basic chains σ equal to zero.

This is still a free graded R -module, freely generated by isomorphism classes of chiral basic chains, with exactly one choice of orientation for each nontrivial basic chain appearing as a generator. (The achiral chains are precisely the basis elements of $\tilde{C}_*(X; R)$ which are sent to zero in $\overline{C}_*(X; R)$; if $2 = 0 \in R$ then no basis elements are set to zero.)

There is a geometric boundary operator $\tilde{\partial} : \tilde{C}_*(X; R) \rightarrow \tilde{C}_{*-1}(X; R)$, sending each (isomorphism class of) basic chain $\sigma : P \rightarrow X$ to the sum of the faces of P^{d-1} equipped with their boundary orientation. Because this is compatible with orientation-reversal, it descends to an operator $\overline{\partial} : \overline{C}_*(X; R) \rightarrow \overline{C}_{*-1}(X; R)$.

Lemma 6.4. *We have $\overline{\partial}^2 = 0$.*

Proof. This immediately follows from the combinatorial conditions on the definition of δ -chain and orientation thereof.

Given any face $\Delta_2 \subset P$ of codimension 2, there are two intermediate faces of codimension 1. The boundary of ∂P then contains two copies of the face Δ_2 : once from each of the intermediate faces. The assumption on orientations were that these two copies were oriented negatively to one another, so they sum to zero. Therefore, $\overline{\partial}^2 \sigma = 0 \in \overline{C}_*(X; R)$. ■

As is familiar already from cubical definitions of singular homology, the homology of $\overline{C}_*(\text{pt}; R)$ is not obviously concentrated in degree zero unless we impose some further degeneracy requirements. (These degeneracy requirements will later turn out to be essential in our definition of the instanton Floer complex.) This is furnished by Lipyanskiy's notion of small image. However, the basic chains of small image do not span a subcomplex of $\overline{C}_*(X; R)$: for instance, a basic chain of dimension larger than X is automatically of small image, but if the chain is of dimension $\dim X + 1$, there is no such guarantee that its boundary has small image. This inspires us to make the following definition.

Definition 6.6. *A basic chain $\sigma : P \rightarrow X$ is degenerate if both σ has small image and $\partial \sigma$ is a disjoint union of basic chains of small image and an achiral chain. The span of the degenerate basic chains forms a subcomplex $D_*(X; R) \subset \overline{C}_*(X; R)$; this is precisely the submodule spanned by basic chains with small image and whose boundary in $\overline{C}_*(X; R)$ is a sum of elements of small image. We define the geometric chain complex of X (with coefficients in R) as*

$$C_*^{\text{gm}}(X; R) := \overline{C}_*(X; R)/D_*(X; R).$$

Again, $C_*^{\text{gm}}(X; R)$ is degreewise R -free. It is functorial under smooth maps, and is supported in degrees $[0, \dim X + 1]$; the map $\partial : C_{\dim X + 1}^{\text{gm}}(X; R) \rightarrow C_{\dim X}^{\text{gm}}(X; R)$ is an isomorphism onto its image. It is also well-behaved with respect to transverse intersections.

Definition 6.7. *Suppose $f_i : P_i \rightarrow X$ is a countable family \mathcal{F} of maps from δ -chains to X . The subcomplex of $C_*^{\text{gm}}(X; R)$ spanned by nondegenerate chains transverse to all of the f_i is written $C_*^{\text{gm}, \mathcal{F}}(X; R)$.*

The following lemma can be proved as an inductive application of transversality theorems to each stratum, using that $P^e \setminus P^{e-1}$ is a manifold and P^e is compact. A proof in the only slightly different setting of δ -chains (not strong) is given in [Lin18, Lemma 4.1.14].

Lemma 6.5. *The inclusion of $C_*^{\text{gm}, \mathcal{F}}(X) \hookrightarrow C_*^{\text{gm}}(X)$ is a quasi-isomorphism.*

Most importantly for us, these chain complexes have chain-level fiber product maps, as long as we're potentially willing to pass to a quasi-isomorphic subcomplex. First we need an additional definition.

Definition 6.8. *Let $f : Z \rightarrow X$ be a map from a strong δ -chain to a smooth manifold X which is a submersion on each stratum. We say that a fiber orientation of f is an orientation of the δ -chain $f^{-1}(x)$ for any $x \in X$.*

Lemma 6.6. *Suppose X is a connected smooth manifold. Let Z be a strong δ -chain, and suppose we have maps $e_- : Z \rightarrow X, e_+ : Z \rightarrow Y$, for which e_- is a submersion on each stratum of Z . Further, suppose e_- is equipped with a fiber orientation. Then there is an induced map $C_*^{\text{gm}}(X; R) \rightarrow C_*^{\text{gm}}(Y; R)$ of degree $(\dim Z - \dim X)$ given by sending a basic chain $\sigma : P \rightarrow X$ to $e_+ : P \times_{e_-} Z \rightarrow Y$, where*

$$P \times_{e_-} Z = \{(p, z) \in P \times Z \mid \sigma(p) = e_-(z)\}$$

and $e_+(p, z) = e_+(z)$. We denote this chain $\sigma \times_{e_-} Z$. Here we orient the fiber product so that

$$T_p P \oplus T_z^{\text{fib}} Z \cong T_{p,z}(P \times_X Z)$$

is an oriented splitting.

If $\dim Z > \dim X + \dim Y + 1$, this chain map is identically zero. These maps satisfy the formula

$$\partial(\sigma \times_{e_-} Z) = \partial\sigma \times_{e_-} Z + (-1)^{\dim P + \dim X} \sigma \times_{e_-} \partial Z.$$

The second statement follows because $C_*^{\text{gm}}(Y; R)$ vanishes in degrees larger than $\dim Y + 1$, and the final statement because taking the boundary in the final case involves commuting the desired outward-pointing normal across the first two factors (of P and the base X).

Now we should justify the claim that this is just a fancy way to write down singular homology with desirable chain-level properties. To make sense of the Eilenberg-Steenrod axioms for *smooth manifolds*, we work with the notion of *admissible pairs* (X, A) , where X is a smooth manifold without boundary, and A is a closed (in the sense of point-set topology) submanifold of X of codimension zero. These were introduced in [Sch93] to prove the equivalence between Morse homology and singular homology, and used in [Lin18, Theorem 4.1.13] to prove that a very similar homology theory to ours (using δ -chains, instead of strong δ -chains) agrees with singular homology; our proof follows similar lines as his.

Theorem 6.7. *The functor $H_*^{\text{gm}}(X; R)$ from smooth manifolds and maps to graded R -modules satisfies the Eilenberg-Steenrod axioms for a homology theory:*

- (1) *The induced map $H_*^{\text{gm}}(X; R) \rightarrow H_*^{\text{gm}}(Y; R)$ is a homotopy invariant,*
- (2) *There is a natural relative long exact sequence relating*

$$H_*^{\text{gm}}(X, Y) := H(C_*^{\text{gm}}(X)/C_*^{\text{gm}}(Y))$$

to $H_*^{\text{gm}}(X)$ and $H_*^{\text{gm}}(Y)$,

- (3) If pt denotes the one-point space, then $H_*^{\text{gm}}(\text{pt}; R)$ is a copy of R concentrated in degree zero, and
- (4) If A is codimension zero, and $Y = \partial A$ is a codimension 1 submanifold of X , bounding A on one side and B on the other, then the inclusion induces an isomorphism $H_*^{\text{gm}}(B, Y) \cong H_*^{\text{gm}}(X, A)$.

As a result, there is an isomorphism $H_*^{\text{sing}}(X; R) \cong H_*^{\text{gm}}(X; R)$ for all X , natural under smooth maps.

Proof. We prove each property separately, and then explain why this restricted class of Eilenberg-Steenrod axioms is enough.

Homotopy invariance follows the expected strategy: write down a chain homotopy of the induced map, given by sending a basic chain P to $P \times I$, with map given by composing the map $\sigma : P \rightarrow X$ with the homotopy f_t .

The relative long exact sequence is a matter of homological algebra (it is induced by a short exact sequence of chain complexes).

To check the first nontrivial axiom, observe that $C_*^{\text{gm}}(\text{pt}; R)$ is concentrated in degrees 0 and 1. We saw in Lemma 6.3 that the boundary map $C_1^{\text{gm}}(\text{pt}; R) \rightarrow C_0^{\text{gm}}(\text{pt}; R) = R$ is identically zero; but a chain of small image with zero boundary has boundary of small image, and is in particular degenerate, so $C_1^{\text{gm}}(\text{pt}; R) = 0$. Therefore

$$H_*^{\text{gm}}(\text{pt}) = C_*^{\text{gm}}(\text{pt}) = R.$$

Excision is the hardest property to verify. If we write $C_*^{\text{gm}}(A \cup B)$ for the image of $C_*^{\text{gm}}(A) \oplus C_*^{\text{gm}}(B)$ in $C_*^{\text{gm}}(X)$, then there is a diagram of exact sequences

$$\begin{array}{ccccccccc} 0 & \longrightarrow & C_*^{\text{gm}}(A) & \longrightarrow & C_*^{\text{gm}}(A \cup B) & \longrightarrow & C_*^{\text{gm}}(B, Y) & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & C_*^{\text{gm}}(A) & \longrightarrow & C_*^{\text{gm}}(X) & \longrightarrow & C_*^{\text{gm}}(X, A) & \longrightarrow & 0. \end{array}$$

By the five lemma and the induced map on homology long exact sequences, if we can show the middle vertical arrow is a quasi-isomorphism, so too will be the last arrow, as desired. Now, if we write $C_*^{\text{gm}, Y}(X)$ for the quasi-isomorphic subcomplex of strong δ -chains transverse to the submanifold Y , there is a map

$$\rho : C_*^{\text{gm}, Y}(X) \rightarrow C_*^{\text{gm}}(A \cup B)$$

given by sending each basic chain $\sigma : P \rightarrow X$ to the sum of $P \cap A$ and $P \cap B$; these are again δ -chains because of the transversality hypothesis, and their boundary is a subdivision of the original boundary into $(\partial P) \cap A$ and $(\partial P) \cap B$, respectively, and so this is a chain map. If P is a cycle, then $\rho(P)$ is homologous to P itself, the bounding chain given by $P \times I$, with the usual δ -structure on one end, and the ‘broken’ δ -structure $P = (P \cap A) \cup (P \cap B)$ on the other. In particular, we see that the inclusion $C_*^{\text{gm}}(A \cup B) \rightarrow C_*^{\text{gm}}(X)$ is surjective on homology. Injectivity is similar: given a chain whose boundary is transverse to P , we can represent it instead by a chain transverse to P with homologous boundary still transverse to P ; cutting it into two pieces, we get a chain in $C_*^{\text{gm}}(A \cup B)$ whose boundary is homologous to the original.

Now we should explain why the homotopy category of admissible pairs of smooth manifolds (with the homotopy type of a finite CW complex) is equivalent to the homotopy category of pairs of finite CW complexes. Putting a relative CW structure on each admissible pair via Morse theory, we have a natural inclusion from the

category of admissible pairs to the category of finite CW complexes; by smooth and (relative) smooth approximation, this is a fully faithful functor, and it suffices to show that it's surjective on objects. Given a pair (X, A) , find an embedding of this in some large Euclidean space; then a small open neighborhood U_X of the image of X (called a regular neighborhood) is homotopy equivalent to X itself and so that ∂U_X is a manifold, and we may choose an open neighborhood U_A of A satisfying the same property, with $\bar{U}_A \subset U_X$. Then the pair (U_X, \bar{U}_A) is the desired admissible pair. ■

If R has characteristic 2, we can simplify definitions by removing all reference to orientations; this is even desirable, as it allows us to refer to the identity map of a non-oriented compact manifold as a chain.

6.2. Equivariant instanton homology. We are now ready to define the framed instanton chain complex $\widetilde{CI}(Y, E, \pi; R)$, using the instanton moduli spaces. The material of sections 1-4 amounts to the following package of theorems about these instanton moduli spaces. Recall from Definition 3.2 that for a $U(2)$ -bundle E over a 3-manifold Y to be weakly admissible, either $b_1(Y) = 0$ or $c_1(E)$ is not twice some element of $H^2(Y; \mathbb{Z})$.

In the following, the regular perturbation π is chosen from an open subset $\mathcal{P}_{E, \delta}$ of the Banach space \mathcal{P}_E of perturbations, defined in Definition 4.3.

Theorem 6.8. *Let Y be a 3-manifold equipped with a weakly admissible $U(2)$ -bundle E , a basepoint b , and a choice of metric and regular perturbation π (which always exists). Then:*

- (1) *The collection of critical orbits of the perturbed Chern-Simons function $cs + \pi$ on the configuration space $\widetilde{\mathcal{B}}_E$ is a finite set of $SO(3)$ -orbits. We write this set of orbits as \mathfrak{C}_π .*
- (2) *Write $c_1(E) = \lambda \in H^2(Y; \mathbb{Z})$. Define $\text{Pair}(H^2)$ to be the set of unordered pairs of integral cohomology classes on Y ; that is, the set $H^2(Y; \mathbb{Z}) \times H^2(Y; \mathbb{Z})$ modulo the relation $(z, w) \sim (w, z)$. The set of reducible critical orbits may be identified with the set*

$$\mathcal{R}_\lambda(Y, E) \subset \text{Pair}(H^2)$$

given by those (z, w) with $z + w = \lambda$. The fully reducible critical orbits are those for which $z = w$, and there are hence $H^1(Y; \mathbb{Z}/2)$ of them if Y is a rational homology sphere with $c_1(E) = 0$, and none otherwise.

- (3) *There is a number $\text{gr}_z(\alpha, \beta) \in \mathbb{Z}$, assigned to each pair of critical orbits α, β and homotopy class of path z between them in $\widetilde{\mathcal{B}}_E$. If w is a path from β to γ and $z * w$ is the concatenation of paths, we have*

$$\text{gr}_{z*w}(\alpha, \gamma) = \text{gr}_z(\alpha, \beta) + \text{gr}_w(\beta, \gamma).$$

For different paths z, w , we have $\text{gr}_z(\alpha, \beta) - \text{gr}_w(\alpha, \beta) \in 8\mathbb{Z}$, so that $\text{gr}(\alpha, \beta) \in \mathbb{Z}/8$ is well-defined.

- (4) *If α and β are reducible, then $\text{gr}(\alpha, \beta)$ is even. If α and β are fully reducible, then $\text{gr}(\alpha, \beta)$ is divisible by 4.*

- (5) Associated to each pair (α, β) of critical orbits and homotopy class z is a smooth $SO(3)$ -manifold (possibly empty) $\widetilde{\mathcal{M}}_{E,z,\pi}^0(\alpha, \beta)$ of dimension $\text{gr}_z(\alpha, \beta) + \dim \alpha - 1$. It comes equipped with equivariant smooth maps

$$\alpha \xleftarrow{e_-} \widetilde{\mathcal{M}}_{E,z,\pi}^0(\alpha, \beta) \xrightarrow{e_+} \beta.$$

- (6) For each critical orbit α , there is an associated 2-element set $\Lambda(\alpha)$. A choice of element of each of $\Lambda(\alpha)$ and $\Lambda(\beta)$ induces an orientation on the fiber of

$$e_- : \widetilde{\mathcal{M}}_{E,z,\pi}^0(\alpha, \beta)$$

for all z ; in the language of Definition 6.8, this is a fiber orientation of e_- . Negating either one of these choices negates the resulting fiber orientation.

- (7) If $\text{gr}_z(\alpha, \beta) \leq 10 - \dim \alpha$, there is a natural compactification

$$\widetilde{\mathcal{M}}_{E,z,\pi}^0(\alpha, \beta) \subset \overline{\mathcal{M}}_{E,z,\pi}(\alpha, \beta)$$

into a compact topological $SO(3)$ -manifold with corners and a smooth structure on each stratum. The endpoint maps extend to equivariant maps from $\overline{\mathcal{M}}^0$ which are smooth on each stratum; we will use the same notation e_{\pm} for these extended maps. In the trivial case $\alpha = \beta$ with homotopy class $z = 0$, we consider $\overline{\mathcal{M}}_{E,0,\pi}(\alpha, \alpha)$ to be empty.

- (8) The action of $SO(3)$ on $\overline{\mathcal{M}}_{E,z,\pi}(\alpha, \beta)$ is free.

- (9) Given any choice of element of each of $\Lambda(\alpha)$, $\Lambda(\beta)$, and $\Lambda(\gamma)$, so that all of the relevant moduli spaces are fiber oriented, there is a decomposition respecting fiber orientations

$$(-1)^{\dim \alpha} \partial \overline{\mathcal{M}}_{E,z,\pi}(\alpha, \beta) \cong \bigsqcup_{\substack{\gamma, w_1, w_2 \\ w_1 * w_2 = z}} (-1)^{\text{gr}_{w_1}(\alpha, \gamma)} \overline{\mathcal{M}}_{E,w_1,\pi}(\alpha, \gamma) \times_{\gamma} \overline{\mathcal{M}}_{E,w_2,\pi}(\gamma, \beta).$$

Observe that the relation $\text{gr}_z(\alpha, \beta) = \dim \overline{\mathcal{M}}_{E,z,\pi}(\alpha, \beta) - \dim \alpha + 1$ is compatible with the dimensions of the manifolds involved in the gluing formula.

Proof. For a fixed metric g , the existence of an admissible perturbation π is the combination of Theorem 3.6 and Theorem 4.34; these guarantee that the critical orbits are isolated, and that the moduli spaces of trajectories between them are smooth manifolds of the appropriate dimension, respectively. That there are only finitely many critical orbits follows from Lemma 3.3 (that the derivative of our perturbed Chern-Simons functional is a proper map). The enumeration of reducible critical orbits is a combination of Corollary 1.3 and Proposition 3.5; as mentioned after Definition 4.3, one of the reasons we demand $\pi \in \mathcal{P}_{E,\delta}$ is so that this enumeration remains true. The integer $\text{gr}_z(\alpha, \beta) = \overline{\text{gr}}_z(\alpha, \beta) - \dim \alpha$ is defined in Definition 4.5, where we also explain why $\overline{\text{gr}}_z$ counts the dimension of the fiber of $\widetilde{\mathcal{M}} \cong \mathbb{R} \times \widetilde{\mathcal{M}}^0$ above a point in α . That the grading is well-defined modulo 8 is Corollary 4.27. The calculation of relative gradings between full reducibles is Proposition 4.32 and arbitrary reducibles is Proposition 4.33. The existence of a compactification only in terms of fiber products of lower-dimensional moduli spaces is Corollary 4.13; that the resulting object is a topological manifold with corners with a smooth structure on each stratum is Proposition 4.43. That the $SO(3)$ -action is free means that

there are no reducible solutions; this is guaranteed by Proposition 4.14. The moduli spaces carry compatible fiber orientations by Proposition 5.12, writing the sign in terms of the grading function gr_z . \blacksquare

We will also need a similar package in the case of cobordisms W . We include this here as well. Recall the definition of weakly admissible bundle over a cobordism from Definition 4.14, as well as Definition 5.3 of homology orientations of a cobordism W .

In the following, we choose the 4-manifold perturbation π from a contractible open set $U \subset \mathcal{P}_c^{(4)}$ of perturbations which restrict to fixed elements of $\mathcal{P}_{E_i, \delta}$ on the ends; this open set was defined in the process of proving Theorem 4.37 (when comparing, take the constant C large enough that all instantons \mathbf{A} in a trajectory z with $\text{gr}_z(\alpha, \beta) \leq 10$ have energy less than C). It contains the perturbation $(\pi_1, 0, \pi_2)$ supported on the ends, with no interior part.

If $\pi \in U$, we say that π is W -small.

Theorem 6.9. *Suppose (W, \mathbf{E}) is an oriented 4-manifold with two cylindrical ends, with incoming end modeled on (Y_1, E_1) and outgoing end modeled on (Y_2, E_2) ; suppose the E_i are weakly admissible, and that we have chosen regular perturbations $\pi_i \in \mathcal{P}_{E_i, \delta}$. Furthermore suppose that \mathbf{E} is a weakly admissible $U(2)$ -bundle over W . Then we have the following.*

- (1) *For every pair of critical orbits $\alpha \subset \mathfrak{C}_{\pi_1}(Y_1)$ and $\beta \subset \mathfrak{C}_{\pi_2}(Y_2)$, there is a set of configurations of framed connections on W from α to β , denoted $\tilde{\mathcal{B}}_{\mathbf{E}}(\alpha, \beta)$. The set of components is written $\pi_0 \tilde{\mathcal{B}}_{\mathbf{E}}(\alpha, \beta)$, and carries a free and transitive action of both $\pi_1 \tilde{\mathcal{B}}_{E_1}(\alpha)$ and similarly with E_2 ; in particular, it carries an affine identification*

$$\pi_0 \tilde{\mathcal{B}}_{\mathbf{E}}(\alpha, \beta) \cong \mathbb{Z}.$$

For each z in this set, there is an integer $\text{gr}_z^W(\alpha, \beta)$, satisfying the same additivity formula as in the previous theorem, which is independent of z after reducing modulo 8.

- (2) *Let π be a W -small regular perturbation on W , restricting to regular perturbations π_i on the cylindrical ends corresponding to Y_i . For each pair of critical orbits α, β , with z as above, we associate a smooth $SO(3)$ -manifold $\tilde{\mathcal{M}}_{\mathbf{E}, z, \pi}(\alpha, \beta)$, which is of dimension $\text{gr}_z^W(\alpha, \beta) + \dim \alpha$ so long as $\tilde{\mathcal{M}}_z(\alpha, \beta)$ contains an irreducible connection. The term $\text{gr}_z^W(\alpha, \beta)$ is independent of z modulo 8. When W is equipped with a smoothly embedded path $\gamma : \mathbb{R} \rightarrow W$ which agrees with $(-t, b_1) \in (-\infty, 0] \times Y_1$ and $(t, b_2) \in [0, \infty) \times Y_2$ for large enough $|t|$, then $\tilde{\mathcal{M}}$ is imbued with smooth equivariant endpoint maps to α and β . A residual set of $\pi \in U$ are regular.*
- (3) *If $\text{gr}_z^W(\alpha, \beta) \leq 10 - \dim \alpha$, then there is a natural compactification of this $SO(3)$ -manifold to a compact topological $SO(3)$ -manifold with corners and a smooth structure on each stratum, $\overline{\mathcal{M}}_{\mathbf{E}, z, \pi}^W(\alpha, \beta)$. The endpoint maps extend to $\overline{\mathcal{M}}^W$.*

- (4) If W is given a homology orientation as in Definition 5.3, and we choose generators of the 2-element sets $\Lambda(\alpha)$ and $\Lambda(\beta)$ from the previous proposition, these induce a fiber orientation on $e_- : \overline{\mathcal{M}}_{\mathbf{E},z,\pi}(\alpha, \beta) \rightarrow \alpha$. Negating any one of these choices negates the corresponding orientation.
- (5) Suppose $b_1(W) = b^+(W) = 0$. So long as all three of the perturbations π_1, π, π_2 are taken sufficiently small, the set of reducible orbits in $\overline{\mathcal{M}}_{\mathbf{E},\pi}(\alpha_1, \alpha_2)$ may be identified as the following set, written $\text{Red}(W, \mathbf{E})$. Write $c_1(\mathbf{E}) = \lambda \in H^2(W; \mathbb{Z})$, and write $\lambda_i = \lambda|_{Y_i}$. There is an induced map $\text{Pair}(H^2W) \rightarrow \text{Pair}(H^2Y_i)$ given by restriction; the orbit α_i corresponds to some

$$r_i \in \mathcal{R}_{\lambda_i}(Y_i, E_i) \subset \text{Pair}(H^2Y_i).$$

Then $\text{Red}(W, \mathbf{E})$ is the subset of $\text{Pair}(H^2W)$ consisting of pairs $\{z, w\}$ restricting to r_i on the corresponding ends, and with $z + w = \lambda$. The set of full reducibles is taken to the subset with $z = w$. If $b^+(W) \neq 0$ and π is a regular perturbation, then no reducibles arise in the moduli space $\overline{\mathcal{M}}$.

- (6) Assuming we have chosen a homology orientation of W and an element of all relevant orientation sets, there is a fiber oriented decomposition

$$(-1)^{\dim \alpha} \partial \overline{\mathcal{M}}_{\mathbf{E},z,\pi}^W(\alpha, \beta) \cong \bigcup_{\substack{\gamma \subset \mathfrak{C}_{\pi_2, w_1, w_2} \\ w_1 * w_2 = z}} (-1)^{\text{gr}_{w_1}^W(\alpha, \gamma)} \overline{\mathcal{M}}_{\mathbf{E}, w_1, \pi}^W(\alpha, \gamma) \times_{\gamma} \overline{\mathcal{M}}_{E_2, w_2, \pi_2}(\gamma, \beta) \\ \bigcup_{\substack{\gamma \subset \mathfrak{C}_{\pi_1, w_1, w_2} \\ w_1 * w_2 = z}} \overline{\mathcal{M}}_{E_2, w_1, \pi_1}(\alpha, \gamma) \times_{\gamma} \overline{\mathcal{M}}_{\mathbf{E}, w_2, \pi}^W(\gamma, \beta).$$

Proof. That such a regular perturbation exists, and that they are generic in U , is Theorem 4.37. The grading function was defined in Definition 4.5 and seen to be well-defined mod 8 in Corollary 4.27, the compactification with no bubbling is given in Corollary 4.13, and the manifold-with corners structure (and decomposition of the boundary) given by Proposition 4.43. The orientations are provided by Proposition 5.12. The identification of reducible orbits was given by Proposition 1.7 and Proposition 4.15.

The path γ is necessary to choose a basepoint in W (as $\gamma(1/2)$, say) and then to parallel transport the chosen framing to the basepoints $\gamma(0), \gamma(1)$ on the boundary components.

The fiber orientations and oriented decomposition were provided by Proposition 5.12. ■

Lastly, to prove invariance of these induced maps, we will want a version of this package for *families* of perturbations π . Recall both the definition of *broken metric* and *regular family of metrics and perturbations* from Section 4.8. We only use this notation for the case of families indexed by the interval $[0, 1]$.

Also recall that Definition 4.14 of weakly admissible cobordisms partitions these into three classes: first, the admissible cobordisms for which one end is an admissible bundle (so supports no reducible flat connections). Second, there are the weakly admissible cobordisms with rational homology sphere ends, $b_1(W) = 0$, and $b^+(W) > 0$ and \mathbf{E} non-trivial. Third, there are those with $b_1(W) = b^+(W) = 0$

and which satisfy the ‘ ρ -monotonicity inequality’

$$\rho_{\pi_1}(r_1) \leq \rho_{\pi_2}(r_2)$$

and support no bad reducibles, meaning that $H_1 Y_1 \oplus H_1 Y_2 \rightarrow H_1 W$ is surjective. (The second condition Definition 4.14 is never satisfied for $U(2)$ -bundles.)

Theorem 6.10. *Suppose $(W, \mathbf{E}) : (Y_1, E_1, \pi_-) \rightarrow (Y_2, E_2, \pi_+)$ is a weakly admissible cobordism between 3-manifolds equipped with regular perturbations $\pi_{\pm} \in \mathcal{P}_{E_i, \delta}$.*

Suppose $I = [0, 1]$ parameterizes a family of metrics on W (with fixed cylindrical ends); the metrics g_0 and g_1 may be cut along a separating 3-dimensional submanifold of W , as in Section 4.8, itself equipped with a regular perturbation. Then there is a contractible open set U_g with

$$(\pi_-, 0, \pi_+) \subset U_g \subset U \subset \mathcal{P}_c^{(4)}$$

such that the following are true. Suppose (W, \mathbf{E}) is equipped with a pair of regular perturbations $\pi(0), \pi(1) \in U_g \subset \mathcal{P}_c^{(4)}$ which equal the π_{\pm} on the ends.

- (1) *If $b^+(W) \neq 1$ or \mathbf{E} supports no reducibles, there is a path $\pi : [0, 1] \rightarrow U_g$ so that $\pi(t)$ and g_t form a regular family of metrics and perturbations.*
- (2) *If $b^+(W) = 1$, then for π_0, π_1 as above and any path $\pi : [0, 1] \rightarrow U_g$ that supports no reducible instantons, we may perturb $\pi(t)$ on the interior to form a regular family of metrics and perturbations.*
- (3) *If $\pi(t)$ is a regular family of perturbations on $[0, 1] = I$, there are smooth $SO(3)$ -manifolds $\widetilde{\mathcal{M}}_{\mathbf{E}, z, \pi(t)}(\alpha, \beta)$ of dimension $\text{gr}_z(\alpha, \beta) + \dim \alpha + 1$.*
- (4) *A homology orientation of W and a choice of element of each of $\Lambda(\alpha)$ and $\Lambda(\beta)$ induces a fiber orientation on*

$$e_- : \widetilde{\mathcal{M}}_{\mathbf{E}, z, \pi(t)}(\alpha, \beta) \rightarrow \alpha,$$

as before. These orientations negate under orientation-reversal of any one of these choices.

- (5) *As long as $\text{gr}_z^W(\alpha, \beta) \leq 9 - \dim \alpha$, the moduli space $\widetilde{\mathcal{M}}_{\mathbf{E}, z, \pi(t)}(\alpha, \beta)$ has a natural compactification $\overline{\mathcal{M}}_{\mathbf{E}, z, \pi(t)}(\alpha, \beta)$ satisfying the same properties as before.*
- (6) *If an element is chosen from each relevant orientation set as above, then there is a fiber oriented decomposition*

$$\begin{aligned} \partial \overline{\mathcal{M}}_{\mathbf{E}, z, \pi(t)}(\alpha, \beta) = & \bigcup_{\gamma \in \mathcal{C}_{\pi_-; z_1 * z_2 = z}} \overline{\mathcal{M}}_{E_1, z_1, \pi_-}(\alpha, \gamma) \times_{\gamma} \overline{\mathcal{M}}_{\mathbf{E}, z_2, \pi(t)}(\gamma, \beta) \\ & \bigcup_{\zeta \in \mathcal{C}_{\pi_+; z_1 * z_2 = z}} (-1)^{\text{gr}_{z_1}^W(\alpha, \zeta)} \overline{\mathcal{M}}_{\mathbf{E}, z_1, \pi(t)}(\alpha, \zeta) \times_{\zeta} \overline{\mathcal{M}}_{z_2}(\zeta, \gamma) \\ & \cup \overline{\mathcal{M}}_{\mathbf{E}, z, \pi(1)}^W(\alpha, \beta) \cup -\overline{\mathcal{M}}_{\mathbf{E}, z, \pi(0)}^W(\alpha, \beta). \end{aligned}$$

If the perturbation π is broken along some submanifold Y with limiting perturbation π_b along Y , and restricting to perturbations π_i on the piece

(W_i, \mathbf{E}_i) , then we have the fiber oriented decomposition

$$\partial \overline{\mathcal{M}}_{\mathbf{E}, z, \pi}^W(\alpha, \beta) = \bigcup_{\eta \in \mathfrak{C}_{\pi_b}; z_1 * z_2 = z} (-1)^{\text{gr}_{z_1}^{W_1}(\alpha, \eta)} \overline{\mathcal{M}}_{\mathbf{E}_1, z_1, \pi_1}^{W_1}(\alpha, \eta) \times_{\eta} \overline{\mathcal{M}}_{\mathbf{E}_2, z_2, \pi_2}(\eta, \beta).$$

Proof. This is precisely the content of Section 4.8; in particular, the transversality and gluing results in the absence of broken metrics are given by Proposition 4.44 and Proposition 4.45, respectively; the extension to the case of broken metrics, including the third item above, is given by Proposition 4.46. \blacksquare

We will use the fiber product maps of Lemma 6.6 associated to the $SO(3)$ -equivariant endpoint maps

$$e_- : \overline{\mathcal{M}}_{E, z, \pi}(\alpha, \beta) \rightarrow \alpha, \quad e_+ : \overline{\mathcal{M}}_{E, z, \pi}(\alpha, \beta) \rightarrow \beta$$

to define the framed instanton differential, as opposed to the usual counting of points in 0-dimensional moduli spaces. (Note that the degree of the fiber product map is precisely $\text{gr}_z(\alpha, \beta) - 1$ for moduli spaces on the cylinder, and precisely $\text{gr}_z^W(\alpha, \beta)$ on a cobordism.) Lemma 6.6 guarantees that when

$$\dim \overline{\mathcal{M}}(\alpha, \beta) \geq \dim \alpha + \dim \beta + 2$$

the map $\sigma \mapsto \sigma \times_{e_-} \overline{\mathcal{M}}(\alpha, \beta)$ is identically zero on the chain level. Because all orbits have dimension at most 3, and

$$\dim \overline{\mathcal{M}}(\alpha, \beta) = \text{gr}_z(\alpha, \beta) + \dim \alpha + 1,$$

we see that these fiber product maps vanish when $\text{gr}_z(\alpha, \beta) \geq 6$. As a result, we need not panic about the Uhlenbeck bubbling arising in large-dimensional moduli spaces. This same observation arose in the original definition of equivariant instanton homology given in [AB96]. In their case, the differential was defined via pullback and integration-over-the-fiber of differential forms, and whenever the differential form is in degree lower than the dimension of the fiber, the integral is zero.

Definition 6.9. A relative $\mathbb{Z}/8$ grading on a set S is defined by a function

$$r : S \times S \rightarrow \mathbb{Z}/8,$$

with $r(x, y) + r(y, z) = r(x, z)$. Such an additive function is equivalent to a function $i : S \rightarrow \mathbb{Z}/8$ considered up to translation; that is, $i \sim i'$ if $i(x) = i'(x) + c$, for some fixed c . Given a function i , one obtains an additive function from its differences:

$$r_i(x, y) = i(y) - i(x).$$

This gives the same function for any equivalent i, i' . Conversely, evaluation defines functions $i_s(x) = r(x, s)$, with $i_s(x) = i_t(x) + i(t, s)$, so that $i_s \sim i_t$.

A relatively graded chain complex splits as a direct sum $C = \bigoplus_{i \in \mathbb{Z}/8} C_i$, where the differential is the direct sum over $d_i : C_i \rightarrow C_{i-1}$.

The relevant relative grading to us is $\text{gr} : \mathfrak{C}_{\pi} \times \mathfrak{C}_{\pi} \rightarrow \mathbb{Z}/8$. Picking a critical orbit ρ arbitrarily, we write $i(\alpha) = \text{gr}(\rho, \alpha)$. If \mathbf{E} is trivial, we set $i(\alpha) = i(\theta, \alpha)$, where θ is the trivial connection. This gives us an absolute grading when \mathbf{E} is trivial (though changing the trivialization may change i by a multiple of 4).

Just as there is a relative grading on critical orbits, the framed instanton chain complex \widetilde{CI} has a relative grading. As a relatively graded R -module, \widetilde{CI} is defined by

$$\widetilde{CI}_*(Y, E, \pi; R) := \bigoplus_{\alpha \subset \mathfrak{C}_\pi} C_*^{\text{gm}}(\alpha; R)[i(\alpha)] \otimes_{R[\mathbb{Z}/2]} R[\Lambda(\alpha)].$$

Here $\Lambda(\alpha)$ is the 2-element orientation set discussed in Theorem 6.8 (6). The group $\mathbb{Z}/2$ acts on $\Lambda(\alpha)$ by swapping the two elements, and on $C_*^{\text{gm}}(\alpha; R)$ by negation.

If $\sigma_1 : P \rightarrow \alpha$ and $\sigma_2 : Q \rightarrow \beta$ are basic chains, the relative grading $|\sigma_1| - |\sigma_2| \in \mathbb{Z}/8$ is given as

$$|\sigma_1| - |\sigma_2| = (\dim P + i(\alpha)) - (\dim Q + i(\beta)) \pmod{8}.$$

Because $i(\alpha) - i(\beta) = \text{gr}(\rho, \alpha) - \text{gr}(\rho, \beta)$, this simplifies using the additivity formula of Theorem 6.8 (2) to

$$|\sigma_1| - |\sigma_2| = \dim P - \dim Q + \text{gr}(\alpha, \beta).$$

This further implies $(|\sigma_1| - |\sigma_2|) + (|\sigma_2| - |\sigma_3|) = |\sigma_1| - |\sigma_3|$, as expected.

The differential is given (termwise on basic chains $\sigma : P \rightarrow \alpha$) by

$$\partial_{CI}\sigma = \partial\sigma + \sum_{\substack{\beta \subset \mathfrak{C}_\pi \\ z, \text{gr}_z(\alpha, \beta) \leq 5}} (-1)^{\dim \sigma} \sigma \times_{e_-} \overline{\mathcal{M}}_{E, z, \pi}(\alpha, \beta).$$

Here $\partial\sigma$ denotes the differential inside $C_*^{\text{gm}}(\alpha)$ and the chain $\sigma \times_{e_-} \mathcal{M}$ is defined as in Lemma 6.6; we orient the moduli spaces using choices of elements of $\Lambda(\alpha)$ and $\Lambda(\beta)$. Because swapping elements of these orientation sets negates the orientation on the fiber product, the map $C_*^{\text{gm}} \otimes R[\Lambda(\alpha)] \rightarrow C_*^{\text{gm}} \otimes R[\Lambda(\beta)]$ descends to the quotient under the two $\mathbb{Z}/2$ actions. We drop the orientation sets from notation as much as is reasonably possible.

The index demand on the sum ensures that all of the moduli spaces $\overline{\mathcal{M}}$ appearing in the sum are compact oriented topological manifolds with corners and a smooth structure on each stratum (as a consequence of Theorem 6.8 (7)). The reason we can do this without concern is that the fiber product with any larger-dimensional moduli spaces is *identically zero* — we can consider our sum as being a formal truncation of what “should be” the instanton differential ∂_{CI} , where we throw out moduli spaces of dimension too large to contribute. This happens precisely when the degree of the fiber product map is larger than $\dim \beta + 1$; because $SO(3)$ -orbits have dimension at most 3, this is true when $\text{gr}_z(\alpha, \beta) - 1 > 4$.

First observe that the differential decreases the relative grading by one: if $\sigma : P \rightarrow \alpha$ is a basic chain, taking the fiber product gives a basic chain $\sigma' : P \times_{e_-} \overline{\mathcal{M}}_z \rightarrow \beta$. The relative grading between these is

$$|\sigma| - |\sigma'| = \dim P - \dim(P \times_{e_-} \overline{\mathcal{M}}_z) + \text{gr}_z(\alpha, \beta).$$

Lemma 6.6 tells us that the dimension of the fiber product is $\dim P + \dim \overline{\mathcal{M}}_z - \dim \alpha$. Combining these with $\dim \overline{\mathcal{M}}_z = \text{gr}_z(\alpha, \beta) + \dim \alpha - 1$, we see that the relative grading

$$|\sigma| - |\sigma'| = \dim P - (\dim P + \text{gr}_z(\alpha, \beta) - 1) + \text{gr}_z(\alpha, \beta) = 1.$$

Lemma 6.11. $\widetilde{CI}_*(Y, E, \pi; R)$ is a chain complex. That is, $\partial_{CI}^2 = 0$. Furthermore, the right action of $C_*^{\text{gm}}(SO(3); R)$ on \widetilde{CI} , acting on each $\bigoplus C_*^{\text{gm}}(\alpha; R)$ on the right (induced by the right action of $SO(3)$ on α) gives \widetilde{CI} the structure of a dg-module.

Proof. It is clear that, for a basic chain $\sigma : P \rightarrow \alpha$,

$$\begin{aligned} \partial_{CI}^2 \sigma &= \partial^2 \sigma + \sum_{\substack{\beta \in \mathfrak{C}_\pi \\ z, \text{gr}_z(\alpha, \beta) \leq 5}} (-1)^{\dim \sigma - 1} (\partial \sigma) \times_{e_-} \overline{\mathcal{M}}_{E, z, \pi}(\alpha, \beta) \\ &+ \sum_{\substack{\beta \in \mathfrak{C}_\pi \\ z, \text{gr}_z(\alpha, \beta) \leq 5}} (-1)^{\dim \sigma} (\sigma \times_{e_-} \overline{\mathcal{M}}_{E, z, \pi}(\alpha, \beta)) \\ &+ \sum_{\substack{\gamma, \beta \in \mathfrak{C}_\pi \\ z, \text{gr}_z(\alpha, \gamma) \leq 5 \\ w, \text{gr}_w(\gamma, \beta) \leq 5}} (-1)^{2 \dim \sigma + \text{gr}(\alpha, \gamma) - 1} \sigma \times_{e_-} (\overline{\mathcal{M}}_{E, z, \pi}(\alpha, \gamma) \times_\gamma \overline{\mathcal{M}}_{E, z, \pi}(\gamma, \beta)). \end{aligned}$$

First, $\partial^2 \sigma = 0$ because $C_*^{\text{gm}}(\alpha; R)$ is a chain complex. Using the decomposition

$$\partial(\sigma \times_{e_-} \overline{\mathcal{M}}) = \partial \sigma \times_{e_-} \overline{\mathcal{M}} + (-1)^{\dim \sigma + \dim \alpha} \sigma \times_{e_-} \overline{\mathcal{M}}$$

of Lemma 6.6, and cancelling the $(\partial \sigma) \times_{e_-} \overline{\mathcal{M}}$ terms, this reduces to

$$\begin{aligned} \partial_{CI}^2 \sigma &= \sum_{\substack{\beta \in \mathfrak{C}_\pi \\ z, \text{gr}_z(\alpha, \beta) \leq 5}} (-1)^{\dim \alpha} \sigma \times_{e_-} (\partial \overline{\mathcal{M}}_{E, z, \pi}(\alpha, \beta)) \\ &+ \sum_{\substack{\gamma, \beta \in \mathfrak{C}_\pi \\ z, \text{gr}_z(\alpha, \gamma) \leq 5 \\ w, \text{gr}_w(\gamma, \beta) \leq 5}} (-1)^{\text{gr}(\alpha, \gamma) - 1} \sigma \times_{e_-} (\overline{\mathcal{M}}_{E, z, \pi}(\alpha, \gamma) \times_\gamma \overline{\mathcal{M}}_{E, z, \pi}(\gamma, \beta)). \end{aligned}$$

The terms in the second sum can only be nonzero when $\text{gr}_z(\alpha, \gamma) + \text{gr}_w(\gamma, \beta) \leq 5$, which is to say that $\text{gr}_{z*w}(\alpha, \beta) \leq 5$. After eliminating terms in the sum which vanish for dimension reasons, we're left with

$$\begin{aligned} (-1)^{\dim \alpha} \partial_{CI}^2 \sigma &= \sum_{\substack{\beta \in \mathfrak{C}_\pi \\ z, \text{gr}_z(\alpha, \beta) \leq 5}} \sigma \times_{e_-} (\partial \overline{\mathcal{M}}_{E, z, \pi}(\alpha, \beta)) \\ &+ \sum_{\substack{\gamma, \beta \in \mathfrak{C}_\pi \\ z, w, \text{gr}_{z*w}(\alpha, \beta) \leq 5}} (-1)^{\dim \alpha + \text{gr}(\alpha, \gamma) - 1} \sigma \times_{e_-} (\overline{\mathcal{M}}_{E, z, \pi}(\alpha, \gamma) \times_\gamma \overline{\mathcal{M}}_{E, z, \pi}(\gamma, \beta)). \end{aligned}$$

This is zero by the decomposition of the boundary given in Theorem 6.8 (9).

That the action of $C_*^{\text{gm}}(SO(3); R)$ makes \widetilde{CI} into a $C_*^{\text{gm}}(SO(3); R)$ -module is clear from the fact that each summand $C_*^{\text{gm}}(\alpha; R)$ is, and that the fiber product map of Lemma 6.6 is a $C_*^{\text{gm}}(SO(3); R)$ -module homomorphism. \blacksquare

Suppose we have a weakly admissible bundle \mathbf{E} over a cobordism W from (Y_1, E_1) to (Y_2, E_2) . To fix the orientations of these moduli spaces, we need to choose a homology orientation on W ; we will suppress this from notation and just refer to “a cobordism W ”.

Lemma 6.12. *Suppose*

$$(W, \mathbf{E}, \pi) : (Y_1, E_1, \pi_1) \rightarrow (Y_2, E_2, \pi_2)$$

is a cobordism where \mathbf{E} is a weakly admissible bundle and π is a choice of metric and regular perturbation restricting to the π_i on the ends. Furthermore suppose W

is equipped with an embedded path γ between the basepoints b_1 and b_2 . Then there is an induced $C_*^{\text{gm}}(SO(3); R)$ -equivariant chain map

$$\tilde{F}_{W, \mathbf{E}, \pi, \gamma} : \tilde{CI}(Y_1, E_1, \pi_1; R) \rightarrow \tilde{CI}(Y_2, E_2, \pi_2; R).$$

Proof. The map is defined analogously to the differential itself; we only need to define its value on a basic chain $\sigma : P \rightarrow \alpha$, where $\alpha \subset \mathfrak{C}_{\pi_1}$ is a critical orbit on Y_1 . Here the value is

$$\tilde{F}_{W, \mathbf{E}, \pi, \gamma}(\sigma) = \sum_{\substack{\beta \in \mathfrak{C}_{\pi_2} \\ z, \text{gr}_z^W(\alpha, \beta) \leq 4}} (-1)^{\dim \sigma} \sigma \times_{e_-} \overline{\mathcal{M}}_{\mathbf{E}, z, \pi},$$

where the map to β is defined using the positive endpoint map $e_+ : \overline{\mathcal{M}}_{\mathbf{E}, z, \pi} \rightarrow \beta$. Defining this endpoint map is the essential place the path $\gamma : \mathbb{R} \rightarrow W$ is used. Here the sum is up to $\text{gr}_z(\alpha, \beta) = 4$ instead of 5, because for cobordisms $\dim \overline{\mathcal{M}} = \text{gr} + \dim \alpha$, whereas there is an extra factor of +1 on the cylinder, account for the translation action.

That this is an equivariant chain map follows as in the previous lemma from the decomposition of the moduli space $\overline{\mathcal{M}}_{\mathbf{E}, z, \pi}$ of Theorem 6.9 (6). \blacksquare

These maps are essentially independent of the perturbation chosen on the cobordism, except in the case that W is weakly admissible and $b^+(W) = 1$. (In the statement of the following, recall the definition of weakly admissible from Definition 4.14.)

Lemma 6.13. *Given two W -small regular perturbations π_0, π_1 on a weakly admissible cobordism (W, \mathbf{E}) , there is a regular family of W -small perturbations $\pi(t)$ parameterized by $[0, 1]$ interpolating between them, restricting to the same fixed perturbations on the ends for all t , or possibly some chain of them, so long as either $b^+(W) \neq 1$ or \mathbf{E} supports no reducible connections. This induces a $C_*^{\text{gm}}(SO(3); R)$ -equivariant chain homotopy between*

$$\tilde{F}_{W, \mathbf{E}, \pi_0, \gamma} \simeq \tilde{F}_{W, \mathbf{E}, \pi_1, \gamma}.$$

*If (W, \mathbf{E}) is weakly admissible and $b^+(W) = 1$, then consider the set of metrics and regular perturbations on W , equipped with the equivalence relation $(g_0, \pi_0) \sim (g_1, \pi_1)$ if there is a path $*g_t \pi_t$ between them so that no (g_t, π_t) supports a reducible π_t -ASD connection. If $\pi_0 \sim \pi_1$, then we have a chain homotopy as above.*

Proof. When W supports no reducible connections, or has $b^+(W) \neq 1$, then the existence of such a regular family of perturbations is Theorem 6.10 (1), and when $b^+(W) = 1$ it is point (2). Some care should be taken when varying the metric; first we should modify the perturbations π_i to be small relative to the path of metrics g_t , and then we should construct a regular path between them.

That the moduli spaces of relevant dimension have compactifications to manifolds with corners and a smooth structure on each stratum follows from Theorem 6.10 (5). The chain map is given on a basic chain $\sigma : P \rightarrow \alpha$ by sending

$$\sigma \mapsto \sum_{\substack{\beta \in \mathfrak{C}_{Y_2} \\ z, \text{gr}_z^W(\alpha, \beta) + 1 \leq 4}} (-1)^{\dim \sigma} \sigma \times_{e_-} \overline{\mathcal{M}}_{\mathbf{E}, z, \pi(t)}.$$

That this is a chain homotopy from $\tilde{F}_{W, \mathbf{E}, \pi_0, \gamma}$ to $\tilde{F}_{W, \mathbf{E}, \pi_1, \gamma}$ follows from the decomposition of the moduli space given by Theorem 6.10 (6).

The main interesting point to make here is that these moduli spaces may be nonempty even though $\text{gr}_z(\alpha, \beta)$ is negative. This corresponds to the existence of $\pi(t)$ -instantons for $t \in (0, 1)$ that are regularly cut out in the *family* — but because of the index they cannot possibly be cut out regularly (considered as an instanton for the single perturbation $\pi(t)$). ■

We need to relate maps arising from the geometric composition of cobordisms and the composition of chain maps.

Lemma 6.14. *Suppose we have two weakly admissible cobordisms $(W_i, \mathbf{E}_i, \pi_i)$ from Y_i to Y_{i+1} with regular perturbations $\pi_i \in \mathcal{P}_{W_i, L_i, \delta}^{(4)}$. Further suppose each W_i is equipped with a path γ_i , so that the positive end of W_1 agrees with the negative end of W_2 , and the paths concatenate to form a smooth path γ in the composite cobordism. Denote the composite (W, \mathbf{E}) . Then there is a regular family of metrics and perturbations $\pi(t)$ so that $\pi(0)$ is a regular perturbation and metric on W , while $\pi(1)$ is the broken metric (and perturbation) corresponding to the obvious way to glue the two cobordisms. This induces a $C_*^{\text{gm}}(SO(3); R)$ -equivariant chain homotopy*

$$\tilde{F}_{W_2, \mathbf{E}_2, \pi_2, \gamma_2} \circ \tilde{F}_{W_1, \mathbf{E}_1, \pi_1, \gamma_1} \simeq \tilde{F}_{W, \mathbf{E}, \pi(0), \gamma}.$$

Proof. Recall that the moduli spaces for the broken metric are by definition given by

$$\overline{\mathcal{M}}_{\mathbf{E}, z, \pi(1)}(\alpha, \gamma) = \bigcup_{\substack{\beta \in \mathfrak{C}_{Y_2} \\ z_1 * z_2 = z}} \overline{\mathcal{M}}_{\mathbf{E}_1, z_1, \pi_1} \times_{\beta} \overline{\mathcal{M}}_{\mathbf{E}_2, z_2, \pi_2}.$$

Thus the chain map given by a broken metric is precisely the composite of its component cobordisms.

Lemma 4.29 guarantees that the composite of weakly admissible cobordisms remains weakly admissible. The conditions in Theorem 6.10 (2)-(3) were that $\pi(t)$ admits $SO(2)$ -reducible instantons if and only if $b^+(W) = 0$, and all fully reducible instantons are cut out transversely. These are assumed to be true for the two pieces of the broken metric. If $\pi(t)$ is the path from a broken metric to a non-broken metric, the same will hold for sufficiently small t . Some care should be taken to ensure that the perturbation $\pi(0)$ is W -small and has no part on the neck corresponding to the perturbation that used to lie ‘at infinity’, as in the discussion preceding Proposition 4.46. Then by a small perturbation, we may ensure that $\pi(t)$ is a regular family of perturbations.

Given that, we may define the induced map as

$$\sigma \mapsto \sum_{\beta} (-1)^{\dim \sigma} \sigma \times_{e_-} \overline{\mathcal{M}}_{\mathbf{E}, z, \pi(t)}(\alpha, \beta).$$

There are no additional difficulties in verifying this is a chain homotopy between the map induced by the perturbation $\pi(0)$ and the map induced by the broken metric/perturbation $\pi(1)$. ■

The chain maps \tilde{F} are invariant under diffeomorphisms of the cobordisms.

Lemma 6.15. *Suppose W_1 and W_2 are cobordisms equipped with weakly admissible bundles \mathbf{E}_i , regular perturbations π_i , and paths γ_i between the basepoints of the ends. Furthermore suppose that the ends of the two cobordisms $(W_i, \mathbf{E}_i, \pi_i, \gamma_i)$ agree (so that we think of them as cobordisms between the same manifolds). Suppose*

$\varphi : W_1 \rightarrow W_2$ is a diffeomorphism, equal to the identity on the ends and with $\varphi(\gamma_1) = \gamma_2$, and $\Psi : \varphi^*\mathbf{E}_2 \cong \mathbf{E}_1$ an isomorphism of $U(2)$ -bundles so that Ψ takes $\varphi^*\pi_2$ to π_1 . Furthermore suppose φ preserves the homology orientations on the W_i . Then the chain maps $\tilde{F}_{W_i, \mathbf{E}_i, \pi_i, \gamma_i}$ are identical.

Proof. This data induces a diffeomorphism

$$\Psi' : \overline{\mathcal{M}}_{\mathbf{E}_1, \pi_1}(\alpha, \beta) \rightarrow \overline{\mathcal{M}}_{\mathbf{E}_2, \pi_2}(\alpha, \beta).$$

That this diffeomorphism preserves the endpoint maps follows because $\varphi(\gamma_1) = \gamma_2$; it preserves orientation because φ preserves the homology orientation of the cobordisms. Given any basic chain $\sigma : P \rightarrow \alpha$, the map

$$P \times_{e_-} \overline{\mathcal{M}}_{\mathbf{E}_1} \rightarrow P \times_{e_-} \overline{\mathcal{M}}_{\mathbf{E}_2}$$

induced on the fiber product is a diffeomorphism preserving $e_+ : P \times_{e_-} \overline{\mathcal{M}}_{\mathbf{E}_i} \rightarrow \beta$. By definition, this means these two basic chains are isomorphic, hence equal in $C_*^{\text{gm}}(\beta; R)$. \blacksquare

Remark 6.1. In dimension 4, any two embedded paths which agree near the ends and are homotopic relative to their boundary are in fact isotopic relative to their boundary. The isotopy extension theorem then provides a diffeomorphism of W fixing the boundary and taking $\varphi(\gamma_2) = \gamma_1$. Hence from the previous lemma $\tilde{F}_{W, \mathbf{E}, \pi, \gamma_1} = \tilde{F}_{W, \mathbf{E}, \varphi^*\pi, \gamma_2}$. So homotopic paths induce the same map in homology.

Combining all of these, we see that we have a functor from a sort of cobordism category to a homotopy category of chain complexes. We make this precise in the following definition. Recall Definition 3.3 of the finite set $\sigma(Y, E)$ of *signature data* on (Y, E) .

There are two categories of cobordisms relevant to us. The first includes the data of the perturbations; the second removes it as much as possible.

Definition 6.10. Let $\text{Cob}_{3,b}^{U(2),w,\pi}$ denote the following category.¹² The objects are closed oriented Riemannian 3-manifolds Y , equipped with a choice of basepoint $b \in Y$, a weakly admissible $U(2)$ -bundle $E \rightarrow Y$, and a regular perturbation $\pi \in \mathcal{P}_{E,\delta}$ for some $\delta > 0$.

The morphisms in $\text{Cob}_{3,b}^{U(2),w,\pi}$ will be formal concatenations of morphisms of the following type (called ‘unbroken cobordisms’).

An unbroken cobordism $(Y_1, E_1, \pi, b_1) \rightarrow (Y_2, E_2, \pi_2, b_2)$ is given by the data of $(W, \tilde{\mathbf{E}}, \pi, \gamma, \varphi, \psi)$. Here W is an oriented, homology oriented, Riemannian 4-manifold with two cylindrical ends, with chosen oriented isometries of these ends to $(-\infty, 0] \times Y_1$ and $[0, \infty) \times Y_2$, respectively, while \mathbf{E} is a weakly admissible $U(2)$ -bundle with specified isomorphisms to the pullback of E_i on the cylindrical ends. The data π is a W -small regular perturbation restricting to the perturbations π_i on the ends, and $\gamma : \mathbb{R} \rightarrow W$ is an embedded path which agrees with (t, b_i) on the ends, following the specified isometries above.

A morphism in $\text{Cob}_{3,b}^{U(2),w,\pi}$ is simply a finite sequence of morphisms W_i (for $1 \leq i \leq n$) so that the target of W_i is the source of W_{i+1} when $1 \leq i < n$. These morphisms are thought of as Riemannian manifolds with ‘broken metric and perturbation’.

¹²The notation b is meant to stand for based, and w for weakly admissible; the π signifies that the perturbation data are included in an essential way.

Let $\text{Cob}_{3,b}^{U(2),w}$ denote the following category. Its objects are closed oriented 3-manifolds Y equipped with a basepoint $b \in Y$, a weakly admissible $U(2)$ -bundle $E \rightarrow Y$, and a signature datum $\sigma \in \sigma(Y, E)$. A morphism $(Y_1, \tilde{E}_1, \sigma_1, b_1) \rightarrow (Y_2, \tilde{E}_2, \sigma_2, b_2)$ is given by the data of $(W, \tilde{\mathbf{E}}, [g, \pi], \gamma, \varphi, \psi)$. Here W is a compact oriented, homology oriented, 4-manifold; $\tilde{\mathbf{E}}$ is a weakly admissible $U(2)$ -bundle; $[g, \pi]$ is an equivalence class as in Lemma 6.13 of metric g and (W, g) -small regular perturbation π^{13} restricting to perturbations π_i on the ends with associated signature data $\sigma_{\pi_i} = \sigma_i$; and $\gamma : [0, 1] \rightarrow W$ is an embedded path which is cylindrical in a collar of the boundary; $\varphi : \partial W \rightarrow Y_1 \sqcup \bar{Y}_2$ is a diffeomorphism sending $\gamma(i)$ to b_{i+1} ; ψ is an isomorphism $E_1 \sqcup E_2 \rightarrow \varphi^* \tilde{\mathbf{E}}$.

Two such cobordisms $(W_i, \tilde{\mathbf{E}}_i, [\pi_i], \gamma_i, \varphi_i)$ are to be considered the same morphism in $\text{Cob}_{3,b}^{U(2),w}$ if there is a diffeomorphism $\phi : W_1 \rightarrow W_2$ with $\varphi_2 \phi = \varphi_1$, a bundle isomorphism $\Psi : \tilde{\mathbf{E}}_1 \rightarrow \phi^* \tilde{\mathbf{E}}_2$ taking $\Psi^*[\pi_2] = [\pi_1]$ and a homotopy relative to the endpoints between γ_1 and $\phi^{-1}(\gamma_2)$.

The target category is the following.

Definition 6.11. Let $\text{Kom}_{C_*SO(3);R}^{r,\mathbb{Z}/8}$ denote the category with objects dg -modules over $C_*^{\text{gm}}(SO(3); R)$ equipped with a relative $\mathbb{Z}/8$ grading, and whose morphisms are relatively graded $C_*^{\text{gm}}(SO(3); R)$ -equivariant chain maps. There is an equivalence relation on the morphisms in this category — $C_*^{\text{gm}}(SO(3); R)$ -equivariant chain homotopy — and there is a category

$$\text{Ho} \left(\text{Kom}_{C_*SO(3);R}^{r,\mathbb{Z}/8} \right),$$

the homotopy category of right $C_*^{\text{gm}}(SO(3); R)$ -modules, whose morphisms are equivalence classes of equivariant chain maps. If we denote the category of relatively $\mathbb{Z}/8$ -graded R -modules with a graded action of $H_*(SO(3); R)$ as $\text{Mod}_{H_*SO(3);R}^{r,\mathbb{Z}/8}$, then taking homology gives a functor

$$\text{Ho} \left(\text{Kom}_{C_*SO(3);R}^{r,\mathbb{Z}/8} \right) \rightarrow \text{Mod}_{H_*SO(3);R}^{r,\mathbb{Z}/8}.$$

The following is essentially immediate from definitions, and we record it as a lemma.

Lemma 6.16. Sending (Y, E, b, π) to the framed instanton complex $\widetilde{CI}(Y, E, \pi; R)$ defines a functor

$$\text{Cob}_{3,b}^{U(2),w,\pi} \rightarrow \text{Kom}_{C_*SO(3);R}^{r,\mathbb{Z}/8}.$$

Proof. Because morphisms $\text{Cob}_{3,b}^{U(2),w,\pi}$ are given by formal concatenations of a given generating set of morphisms, this merely asserts that we have defined induced maps for that generating set. This was done in Lemma 6.12. \blacksquare

Finally, we see that framed instanton homology is functorial on the category of cobordisms *without* perturbation data.

Theorem 6.17. There is a functor

$$\tilde{I} : \text{Cob}_{3,b}^{U(2),w} \rightarrow \text{Mod}_{H_*SO(3);R}^{r,\mathbb{Z}/8}$$

¹³By definition, unless $b^+(W) = 1$, all perturbations are equivalent, and so this additional data is vacuous!

so that for any regular perturbation π on (Y, E) , the group $\tilde{I}(Y, E, \sigma; R)$ is canonically isomorphic to the homology groups $\tilde{I}(Y, E, \pi; R)$ of the chain complexes defined above. Following this canonical isomorphism, the induced map of a cobordism $(W, \mathbf{E}, [\pi])$ is given by the induced map on homology of the map of chain complexes defined above.

Proof. The previous lemmas showed that \tilde{CI} is a complex of the appropriate type, and showed that cobordisms W in this category — when equipped with a perturbation π — induce chain maps, as long as π is sufficiently small. An isomorphism between $(W, \tilde{\mathbf{E}}, \gamma, \varphi, \psi)$ (here, taking the path γ_1 diffeomorphically to γ_2) takes the moduli spaces (and their endpoint maps) of the first to those of the second, and hence the chain maps are the same on the nose. In particular, given regular perturbations π_1, π_2 on (Y, E) with signature data σ , any path π_i between them gives a sufficiently small perturbation on $\mathbb{R} \times Y$ (possibly with time-varying metric), and hence gives an induced map $\tilde{F}_{\mathbb{R} \times Y, \pi^* E, \pi_i, b}$; by Lemma 6.13, any two such chain maps are homotopic.

We will use this to pin down the group $I(Y, E, \sigma)$ as something independent of the choice of metric and perturbation.

Write $\rho_{12} : \tilde{I}(Y, E, \pi_1) \rightarrow \tilde{I}(Y, E, \pi_2)$ for the unique isomorphism between the groups associated to the perturbations π_i (abbreviated for the moment $\tilde{I}(\pi_i)$) defined by choosing a perturbation along $\mathbb{R} \times Y$ as above. Set

$$\tilde{I}(Y, E, \sigma; R) := \left(\bigoplus_{\pi; \sigma_\pi = \sigma} \tilde{I}(Y, E, \pi; R) \right) / \left((x \in \tilde{I}(\pi_1) \sim \rho_{12} x \in \tilde{I}(\pi_2)) \right).$$

The map $\tilde{I}(Y, E, \pi; R) \rightarrow \tilde{I}(Y, E; R)$ induced by inclusion into the direct sum is an isomorphism for any regular perturbation π .

A cobordism with cylindrical ends (W, \mathbf{E}, π) restricting to π_i on the appropriate ends induces a map on each summand $\tilde{F}_{W, \tilde{\mathbf{E}}, \gamma} : \tilde{I}(Y_1, E_1, \pi_1) \rightarrow \tilde{I}(Y_2, E_2, \pi_2)$. The invariance result of Lemma 6.13 implies that this map is the same, independent of the choice of π extending the π_i on the ends; furthermore, Lemma 6.14 implies that these compose as expected with the canonical isomorphisms ρ_{12} above. Therefore, this defines an induced map $\tilde{F}_{W, \mathbf{E}, \gamma} : \tilde{I}(Y_1, E_1, \sigma_1) \rightarrow \tilde{I}(Y_2, E_2, \sigma_2)$ by attaching cylindrical ends on each boundary component and choosing some sufficiently small regular perturbation.

We saw that the chain maps were independent of the perturbation π and homotopy class γ up to chain homotopy in Lemma 6.13 (unless $b^+(W) = 1$, in which case there is an equivalence relation on perturbations), and that they compose as a functor should up to chain homotopy in Lemma 6.14. \blacksquare

The choice of $U(2)$ bundle is necessary to pin down signs, both in the differential on the chain complex and in the cobordism maps. But at least at the level of Floer homology groups, the $U(2)$ -bundle is mostly inessential, as we will see in the following.

Lemma 6.18. *Let J be a nonempty finite set and A be an abelian group. Write $a\delta_{ij}$ for the function $J^2 \rightarrow A$ which equals a if $i = j$ and is zero otherwise; then denote $a\delta_i : J^2 \rightarrow A$ by $(a\delta_i)(x, y) = (a\delta_{xi}) - (a\delta_{yi})$.*

Suppose $s : J^2 \rightarrow A$ be a set map satisfying

$$s(x, y) + s(y, z) = s(x, z).$$

Then there is a choice $a_j \in A$ for each $j \in J$ so that

$$s = \sum_j a_j \delta_j;$$

the choice of a_j is unique up to a uniform change $a'_j = a_j + b$, where b is independent of j .

Proof. Choose a basepoint $p \in J$, and set $a_j = s(p, j)$; then $s' = s - \sum_j a_j \delta_j$ still satisfies $s'(x, y) + s'(y, z) = s'(x, z)$, and now further satisfies $s'(p, j) = 0$ for all j . Therefore for any $i, j \in J$, we have

$$s'(p, i) + s'(i, j) = s'(p, j),$$

and so s' is identically zero, as desired.

The essential uniqueness of a_j follows from knowing the kernel of the map $(a_j) \mapsto \sum_j a_j \delta_j$. If $\sum_j a_j \delta_j(i, k) = 0$ for all i, k , then we obtain from the definition of δ_j that $a_i - a_k = 0$ for all i, k , as desired. \blacksquare

Corollary 6.19. *Let (Y, E, π) be a 3-manifold equipped with a weakly admissible $SO(3)$ -bundle and regular perturbation π . Suppose we are given any two $U(2)$ -bundles \tilde{E}_i with an isomorphism $\tilde{E}_i \times_{U(2)} SO(3) \cong E$; we may define the four flavors of instanton Floer homology groups $I^\bullet(Y, \tilde{E}_i, \pi)$.*

Then there is an isomorphism $f : I^\bullet(Y, \tilde{E}_1, \pi) \rightarrow I^\bullet(Y, \tilde{E}_2, \pi)$ of relatively graded R -modules with an action of either $H_(SO(3); R)$ or $H^{-*}(BSO(3); R)$, as appropriate to the flavor \bullet . While this isomorphism is not canonically defined, it is defined up to a sign; that is, $\{f, -f\}$ is a canonically defined pair of isomorphisms between these two a priori different Floer homology groups.*

Proof. Note that the underlying graded R -module (or graded $C_*(SO(3); R)$ -module, as appropriate) is independent of the choice of $U(2)$ -bundle; the only contribution the choice of $U(2)$ -bundle makes is an orientation of the moduli spaces $\overline{\mathcal{M}}(\alpha, \beta)$, or rather the fiber above a point in α , assuming we choose an element of two orientation sets Λ_α and Λ_β . If we change the $U(2)$ -bundle, the orientation sets change. Choosing an isomorphism $\Lambda_\alpha \rightarrow \Lambda'_\alpha$ for each α gives an identification $\overline{\mathcal{M}}_{\alpha\beta} \cong \overline{\mathcal{M}}'_{\alpha\beta}$, which is either oriented or unoriented, depending on a sign $s(\alpha, \beta) \in \pm 1$ (note that this sign is the same for any connected component of the moduli space); because the orientations are compatible with gluing, we find that $s(\alpha, \beta)s(\beta, \gamma) = s(\alpha, \gamma)$.

We may abstract somewhat. The situation we have (for any of the instanton homology complexes) is that there is a decomposition

$$CI^\bullet(Y, \tilde{E}_i, \pi) \cong \bigoplus_{j \in J} C_j \times_{\pm 1} \Lambda_j^i$$

as graded R -modules; here $J = \mathfrak{C}_\pi$, the set of π -flat orbits, and Λ_j^i is a 2-element orientation set for each j , which depends on the choice of bundle \tilde{E}_i . The differential may be written in matrix form as (d_{jk}^i) ; pinning down an actual map $C_j \rightarrow C_k$ requires choosing an element of Λ_j^i and Λ_k^i .

Choose an arbitrary isomorphism $r_j : \Lambda_j^1 \cong \Lambda_j^2$ for each j ; this gives a map

$$f_r : CI^\bullet(Y, \tilde{E}_1, \pi) \rightarrow CI^\bullet(Y, \tilde{E}_2, \pi).$$

This map satisfies

$$f_r d_{jk}^1 = s(j, k) f_r d_{jk}^2,$$

with as above $s(j, k)s(k, \ell) = s(j, \ell)$. Following Lemma 6.18, there are exactly two functions $t : J \rightarrow \{\pm 1\}$ with $t(j)t(k) = s(j, k)$. If we change the chosen identifications $r_j : \Lambda_j^1 \rightarrow \Lambda_j^2$ to

$$(tr)_j := t(j)r_j : \Lambda_j^1 \rightarrow \Lambda_j^2,$$

correspondingly we find that the new map f_{tr} is in fact a chain map, and hence an isomorphism of chain complexes (with action of whatever dg-algebra is appropriate). Because there are only two choices of r so that f_r is a chain map (and the two choices are negatives of one another), we have a canonical pair of isomorphisms $\{f, -f\}$ between $I^\bullet(Y, \tilde{E}_1, \pi)$ and $I^\bullet(Y, \tilde{E}_2, \pi)$, where \tilde{E}_i are two $U(2)$ -lifts of a given $SO(3)$ -bundle. \blacksquare

Remark 6.2. There is a natural partial ordering on $\sigma(Y, E)$: an element σ defines a function $f_\sigma : \text{Red}(Y, E) \rightarrow \mathbb{Z}_{\geq 0}$, and we say that $\sigma \leq \sigma'$ if $f_\sigma(\alpha) \leq f_{\sigma'}(\alpha)$ for all α . The problem with attempting to show that the instanton homology groups are invariant of the signature datum is that the cylinder is only a weakly admissible cobordism $(Y, E, \sigma) \rightarrow (Y, E, \sigma')$ if $\sigma \leq \sigma'$. Any attempt to resolve this needs somehow to cope with reducible connections on the cylinder which cannot be made to be cut out transversely; perhaps the obstruction bundle technique, as described briefly at the end of [Don02] and used to great effect in [Tau84], is one such tool.

Remark 6.3. Suppose we work instead in the category $\text{Cob}_{3,b}^{SO(3),a}$ of admissible cobordisms. Then in fact, a much stronger invariance property of the maps $\tilde{F}_{W,\mathbf{E}}$ is true: the homotopies between induced maps are homotopic, and so on.

A precise formulation of this statement uses the language of *quasicategories* (sometimes called $(\infty, 1)$ -categories); a good introduction to the language is [HLS16]. There is a quasicategory $\pi\text{Cob}_{3,b}^{U(2),a}$ of 3-manifolds equipped with perturbations, whose morphism spaces are simplicial sets whose n -simplices are Δ^n -indexed admissible families of (possibly broken) perturbations on cobordisms. The framed instanton chain complex lifts to a functor from this quasicategory to the quasicategory of chain complexes, and furthermore the forgetful functor $\pi\text{Cob}_{3,b}^{U(2),a} \rightarrow \text{Cob}_{3,b}^{U(2),a}$ to the homotopy category is an equivalence of quasicategories. This is essentially the statement that the space of admissible perturbations on a given cobordism is contractible, and a version of this would be provided by a natural extension of Theorem 6.10 for simplices; the only obstructions to such an extension lies at reducible connections, which the admissibility assumption is used to avoid. Using this, it's possible to find a chain-level version of \tilde{CI} that is functorial under cobordisms on the nose, as opposed to up to homotopy, and well-defined up to essentially unique natural equivalence. In fact, this is done using the quasicategorical Kan extension; the definition of $I(Y, E, \sigma)$ above is a very special case of a Kan extension. We don't see any need for this structure, and so leave the details to the interested reader.

Any extension of such a result to all weakly admissible $(W, \tilde{\mathbf{E}})$ would require a substantially different notion of “regular family”, as there are obstructions to achieving transversality in the standard sense at reducible connections for families.

In this framework, the space of perturbations on a fixed 3-manifold form a quasi-category (actually, a Kan complex): 0-simplices are pairs (Y, \tilde{E}) of a weakly admissible 3-manifold and regular perturbation π , and an n -simplex starting at (Y, \tilde{E}, π_1) and ending at (Y, \tilde{E}, π_2) is a regular family of perturbations on $\mathbb{R} \times Y$ parameterized by the $(n - 1)$ -dimensional simplex, so that for all perturbations in

this family, the restriction of π to the the corresponding end is the fixed perturbation π_i . Then 6.10 (1) guarantees that the space of perturbations *with fixed signature perturbation* is contractible, and that this quasi-category is equivalent to the poset whose elements are signature data on Y , and there is a morphism $\sigma \rightarrow \sigma'$ iff $\sigma \leq \sigma'$.

6.3. The index filtration. Recall that the relatively $\mathbb{Z}/8$ -graded complex

$$\widetilde{CI}(Y, E, \pi; R)$$

is defined as a graded $C_*^{\text{gm}}(SO(3); R)$ -module to be

$$\bigoplus_{\alpha \in \mathfrak{C}_\pi} C_*^{\text{gm}}(\alpha; R)[i(\alpha)],$$

where $i(\alpha) = \text{gr}(\rho, \alpha) \in \mathbb{Z}/8$; this grading is well-defined up to a translation (arising from choosing a different base orbit ρ), and thus defines a relative grading on \widetilde{CI} . Here we drop the orientation sets $\Lambda(\alpha)$ from notation; if so desired, there is no harm in making a choice of orientation for each critical orbit.

Now if z is a homotopy class of path from ρ to α , write

$$i(\alpha, z) = \text{gr}_z(\rho, \alpha) \in \mathbb{Z}/8.$$

There is a \mathbb{Z} -graded complex $\widetilde{CI}_{\text{unr}}(Y, E, \pi; R)$ given as

$$\bigoplus_{\substack{\alpha \in \mathfrak{C}_\pi \\ z \in \pi_1(\widetilde{\mathcal{B}}_E, \rho, \alpha)}} C_*^{\text{gm}}(\alpha; R)[i(\alpha, z)].$$

If (α_i, z_i) are two such labelled critical orbits, there is a unique homotopy class of path w from α_1 to α_2 so that $z_2 \simeq w * z_1$. Then using the additivity of gr , we see that the degree difference between $\sigma_i : P_i \rightarrow \alpha_i$ is given as

$$\begin{aligned} |\sigma_1| - |\sigma_2| &= (\dim P_1 - i(\alpha_1, z_1)) - (\dim P_2 - i(\alpha_2, z_2)) \\ &= \dim P_1 - \dim P_2 + (\text{gr}_{z_2}(\rho, \alpha_2) - \text{gr}_{z_1}(\rho, \alpha_1)) \\ &= \dim P_1 - \dim P_2 + \text{gr}_w(\alpha_1, \alpha_2). \end{aligned}$$

We thus see that the grading induced by choosing ρ as a basepoint induces the expected relative grading.

We write a basic chain $\sigma : P \rightarrow \alpha$ corresponding to a critical orbit α labelled by z as (σ, z) . The differential of $\sigma : P \rightarrow (\alpha, z)$ is given by

$$\partial_{CI}(\sigma, z) = (\partial\sigma, z) + \sum_{\substack{\beta \in \mathfrak{C}_\pi \\ w \in \pi_1(\widetilde{\mathcal{B}}_E, \alpha, \beta) \\ \text{gr}_w(\alpha, \beta) \leq 5}} (\sigma \times_{e_-} \overline{\mathcal{M}}_{E, w, \pi}(\alpha, \beta), z * w).$$

Each critical orbit α has a unique homotopy class of path $1 \in \pi_1(\widetilde{\mathcal{B}}_E, \alpha)$ with $\text{gr}_1(\alpha, \alpha) = 8$; the element 1 is a cyclic generator of the fundamental group, which is isomorphic to \mathbb{Z} . The complex $\widetilde{CI}_{\text{unr}}(Y, E, \pi; R)$ has a periodicity isomorphism, sending the component $C_*^{\text{gm}}(\alpha; R)$ labelled by z identically to the component labelled by $z+1$. That this commutes with the differential above is only the statement $w+1 = 1+w$ for any path w .

The $8\mathbb{Z}$ -periodic, \mathbb{Z} -graded complex $\widetilde{CI}_{\text{unr}}(Y, E, \pi; R)$ is called the *unrolled complex* of $\widetilde{CI}(Y, E, \pi; R)$ in Section A.8; the quotient by the above periodicity isomorphism just forgets about the labelling by z .

$\widetilde{C}I_{\text{unr}}(Y, E, \pi; R)$ carries an honest filtration by index. We write

$$F_s \widetilde{C}I_{\text{unr}}(Y, E, \pi; R) = \bigoplus_{\substack{\alpha \in \mathfrak{C}_\pi \\ z \in \pi_1(\widetilde{\mathcal{B}}_E, \rho, \alpha) \\ i(\alpha, z) \leq s}} C_*^{\text{gm}}(\alpha; R)[i(\alpha, z)].$$

Because a nonempty moduli space $\overline{\mathcal{M}}_{E,w,\pi}(\alpha, \beta)$ can only exist if

$$\text{gr}_z(\rho, \alpha) > \text{gr}_{z*w}(\rho, \beta),$$

the differential decreases index and hence preserves the filtration. Furthermore, because the filtration is defined by taking a direct sum of $C_*^{\text{gm}}(\alpha; R)$ for (α, z) satisfying the index bound, this is a filtration by $C_*^{\text{gm}}(SO(3); R)$ -modules.

Because $\text{gr}_{z+1}(\rho, \alpha) = \text{gr}_z(\rho, \alpha) + 8$, this is a periodic filtration in the sense of Definition A.9. The associated graded module is

$$\text{gr}_p \widetilde{C}I_{\text{unr}}(Y, E, \pi; R) \cong \bigoplus_{\substack{\alpha \in \mathfrak{C}_\pi \\ i(\alpha) = p \bmod 8}} C_*^{\text{gm}}(\alpha; R)[p].$$

In particular, because $C_*^{\text{gm}}(X; R)$ is supported in degrees at most $\dim X + 1$, and all of our components α are $SO(3)$ -orbits, each associated graded piece is bounded, supported in degrees $[0, 4]$.

Now note that the instanton differential $\partial_{CI} : \widetilde{C}I_{\text{unr}} \rightarrow \widetilde{C}I_{\text{unr}}$ decomposes as $\partial_{CI} = \partial_0 + \partial_1 + \dots + \partial_5$, where $\partial_0(\sigma, z) = (\partial\sigma, z)$ and for $k > 0$,

$$\partial_k(\sigma, z) = \sum_{\substack{\beta \in \mathfrak{C}_\pi \\ w \in \pi_1(\widetilde{\mathcal{B}}_E, \alpha, \beta) \\ \text{gr}_w(\alpha, \beta) = k}} (\sigma \times_{e_-} \overline{\mathcal{M}}_{E,w,\pi}(\alpha, \beta), z * w).$$

This is the component of the differential that decreases filtration by k , and is given by those fiber product maps with moduli spaces which increase dimension by $k - 1$. The decomposition of ∂_{CI} into the ∂_k ends at ∂_5 because fiber products with moduli spaces with $\text{gr}_w(\alpha, \beta) > 5$ are identically zero.

Furthermore, observe that $\dim \overline{\mathcal{M}}_z(\alpha, \beta) = \dim \alpha + \text{gr}_z(\alpha, \beta) - 1$, while for any nonempty moduli space we have $\dim \overline{\mathcal{M}}_z(\alpha, \beta) \geq 3$ because by Theorem 6.8 (8), the $SO(3)$ -action is free. Therefore, if $\overline{\mathcal{M}}_z(\alpha, \beta)$ is nonempty and α is an $SO(2)$ -reducible, necessarily $\text{gr}_z(\alpha, \beta) > 1$; if α is fully reducible, then necessarily $\text{gr}_z(\alpha, \beta) > 3$. In particular, ∂_1 is identically zero on reducible orbits; in addition both ∂_2 and ∂_3 vanish on fully reducible orbits.

Therefore, we have the following, which a special case of Theorem A.29, item (3).

Theorem 6.20. *Let (Y, E) be a weakly admissible bundle and π a regular perturbation. There is a $(\mathbb{Z}/8, \mathbb{Z})$ -bigraded spectral sequence of $H_*(SO(3); R)$ -modules so that the $\mathbb{Z}/8$ -grading is relative, whose E^1 page is*

$$E_{p,q}^1 = \bigoplus_{\substack{\alpha \in \mathfrak{C}_\pi \\ \text{gr}(\alpha) = p}} H_q(\alpha; R)$$

this spectral sequence converges to $E_{p,q}^\infty \cong \text{gr}_p \widetilde{I}_{p+q}(Y, E, \pi; R)$, the p th component of the associated graded vector space of $\widetilde{I}(Y, E, \pi; R)_{p+q}$.

On the unrolled (\mathbb{Z}, \mathbb{Z}) -graded spectral sequence, we have

$$E_{p,q}^\infty \cong \text{gr}_p \widetilde{I}_{p+q}^{\text{unr}}.$$

For a class $[x] \in E^r$ which has a representative $x \in \widetilde{CI}$ with $d_i x = 0$ for $i < r$, the spectral sequence differential $d_r[x]$ may be identified with $[\partial_r x]$. In particular, the differential d_1 on E^1 is only nonzero on the irreducible components, where it is given by counting points in 0-dimensional moduli spaces $\overline{\mathcal{M}}_{E,z,\pi}(\alpha, \beta)/SO(3)$ (note that β need not be irreducible).

The index spectral sequence for \widetilde{I} degenerates on the E^5 page for dimension reasons.

Proof. The only conditions we need to check to apply Theorem A.29 is that the algebra $A = C_*^{\text{gm}}(SO(3); R)$ and each associated graded module $\text{gr}_p \widetilde{CI}_{\text{unr}}$ (for each fixed p) are bounded in degree. These follow by definition: $C_*^{\text{gm}}(X; R)$ is supported in degrees $[0, \dim X + 1]$, and $\text{gr}_p \widetilde{CI}_{\text{unr}}$ is a direct sum of $SO(3)$ -orbits with $i(\alpha) = p$, so is supported in degrees $[p, p + 4]$.

That the differentials on elements $[x]$ with lifts that have $d_i x = 0$ for $i < r$ are induced by ∂_r is an elementary diagram chase. In fact, that this is true holds more generally: Wall's notion of a *multicomplex* (see Definition A.4) is a complex so that the differential splits nicely into components d_r that decrease the filtration by r , and C is equipped with the induced differential

$$d = d_0 + d_1 + d_2 + \cdots$$

Observe that our choice of ∂_k is precisely the component of the differential that decreases the filtration level of \widetilde{CI} by k .

The invariants RE^∞ and W of Proposition A.3 are both zero for this spectral sequence, because it degenerates at the 5th page.

By the second part of Proposition A.3, and Definition A.3 of strong convergence, we see that $E^\infty(\widetilde{CI}) = \text{gr } \widetilde{I}$. \blacksquare

As an immediate corollary, we may define and calculate the Euler characteristic $\chi(\widetilde{I}(Y, E, \sigma))$, where σ is a signature datum.

Corollary 6.21. *Let (Y, E, σ) be a 3-manifold equipped with a weakly admissible bundle and signature datum σ . Then $\widetilde{I}(Y, E, \sigma)$ is finitely generated, and*

$$\chi(\widetilde{I}(Y, E, \sigma)) = |H^2(Y; \mathbb{Z})| = |H_1(Y; \mathbb{Z})|.$$

If $b_1(Y) > 0$, we interpret the right-hand-side to be zero.

This follows because

$$\chi(\widetilde{I}(Y, E, \sigma)) = \chi(\text{gr } \widetilde{I}(Y, E, \sigma));$$

because $\chi(SO(3)) = 0$, the only contribution to this Euler characteristic is from the reducibles. Then the result follows from the enumeration of reducibles given by Proposition 1.4, which have even index by the calculation of Proposition 4.33 (and so all contribute positively to the Euler characteristic sum). The isomorphism $H^2(Y; \mathbb{Z}) = H_1(Y; \mathbb{Z})$ is Poincaré duality.

Remark 6.4. In general it's rare for a multicomplex that a class $[x] \in E^r$ has a representative with all $d_i x = 0$ for $i < r$, but is relatively common in the \widetilde{I} spectral sequence. The discussion simplifies substantially if $\frac{1}{2} \in R$, so we make this assumption; then we know that $H_*(SO(3); R) = R \oplus R[3]$ and

$$H_*(SO(3)/SO(2); R) = H_*(S^2; R) = R \oplus R[2].$$

Choosing an arbitrary class $\rho \in \mathfrak{C}_\pi$, we write the absolute grading

$$i(\alpha) = \text{gr}(\rho, \alpha).$$

We define

$$\begin{aligned} C_p^{\text{irr}} &= \bigoplus_{\substack{\alpha \in \mathfrak{C}_\pi \\ \Gamma_\alpha = 1 \\ i(\alpha) = p}} R \\ C_p^{SO(2)} &= \bigoplus_{\substack{\alpha \in \mathfrak{C}_\pi \\ \Gamma_\alpha \cong SO(2) \\ i(\alpha) = p}} R \\ C_p^\theta &= \bigoplus_{\substack{\alpha \in \mathfrak{C}_\pi \\ \Gamma_\alpha = SO(3) \\ i(\alpha) = p}} R \end{aligned}$$

With this notation, the E^1 page is given as

$$E_{p,q}^1 = \begin{cases} C_p^{\text{irr}} \oplus C_p^{SO(2)} \oplus C_p^\theta & q = 0 \\ C_p^{SO(2)} & q = 2 \\ C_p^{\text{irr}} & q = 3 \end{cases}$$

The E^1 differential on C_p^{irr} counts $SO(3)$ -orbits in 3-dimensional moduli spaces. For $\alpha \in E_{p,0}^1$ the first page differential is given by counting points in unframed moduli spaces with $\text{gr}(\alpha, \beta) = 1$;

$$d_1 \alpha = \sum_{\beta, \text{gr}(\alpha, \beta) = 1} (\#\overline{\mathcal{M}}(\alpha, \beta)/SO(3)) \beta,$$

and $d_1 \alpha$ is zero if α is reducible.

Because the E^1 differential vanishes on the $SO(2)$ -irreducibles, we see that the differential d_2 on $[\alpha] \in E_{p,2}^1$, for $\alpha \in C_p^{SO(2)} \subset E_{p,2}^1$, is given by a similar formula as above: counting, for each β with $\text{gr}(\alpha, \beta) = 2$, the (signed) number of $SO(3)$ -orbits of trajectories between them: $n(\alpha, \beta) = \#\overline{\mathcal{M}}(\alpha, \beta)/SO(3)$ for β irreducible; so

$$d_2[\alpha] = \sum_{\beta, \text{gr}(\alpha, \beta) = 2} n(\alpha, \beta)[\beta]_{p-2,3}.$$

d_2 vanishes everywhere else for dimension reasons. Finally, d_4 on $[\theta] \in E_{0,0}^4$ (when Y is a rational homology sphere and so has trivial connections, and $0 \in V$ makes sense) counts $SO(3)$ -orbits $\overline{\mathcal{M}}(\theta, \beta)/SO(3)$ between an irreducible and an irreducible with $\text{gr}(\theta, \beta) = 4$. (This is the map Donaldson calls D_2 in [Don02, Section 7.1].)

Once we develop the equivalence to the Donaldson model DCI in Section 7.2, the terms in this spectral sequence become much easier to calculate, and visible at the chain level.

6.4. Four flavors of instanton homology. We will soon apply Theorem A.29 to define three additional flavors of equivariant instanton homology, $I^+(Y)$, $I^-(Y)$, and $I^\infty(Y)$. To explain what these different flavors represent, it is useful to analogize to the situation of finite G -CW complexes. To correctly state certain stabilization phenomena, we must ensure everything in sight has a basepoint. Let X_+ denote

the disjoint union of a G -space X with a disjoint fixed basepoint. The smash product of pointed spaces is

$$X \wedge Y = (X \times Y)/(X \times *_Y, *_X \times Y).$$

A pointed G -CW complex is a G -space assembled from cells of the form

$$D^i \wedge (G/H)_+ = (D^i \times G/H)/(* \times G/H),$$

where the basepoint is in the boundary sphere and attaching maps preserve the basepoint. Given a pointed G -CW complex X , its most well-known homological invariants are its nonequivariant homology $H_*(X)$ and its Borel equivariant homology

$$H_*^G(X) = H_*((X \wedge EG_+)/G),$$

both taken relative to a basepoint (equivalently, taking reduced homology). These are both invariants of X up to equivariant homotopy equivalence. In fact, write ΣX denotes the reduced suspension

$$S^1 \wedge X = ([0, 1] \times X)/([0, 1] \times *, \{0\} \times X, \{1\} \times X)$$

of a pointed G -space. Because $H_{*+1}(\Sigma X) \cong H_*(X)$ and $\Sigma X \wedge_G EG_+ = \Sigma(X \wedge_G EG_+)$, the equivariant homology groups H_*^G are invariants of X up to *stable* equivariant homotopy equivalence.

In the stable homotopy category of finite G -CW complexes, where the suspension operator Σ has an inverse and there is a sphere S^n for every $n \in \mathbb{Z}$, there is a contravariantly functorial Spanier-Whitehead duality operator D_G sending

$$D_G(S^n \wedge (G/H)_+) \cong S^{-n} \wedge (G/H)_+.$$

Explicitly, if $V \cong \mathbb{R}^N$ is a G -representation, and $X \hookrightarrow S^{V+1}$ is an equivariant embedding into the one-point compactification of $V \oplus \mathbb{R}$, then $D_G X$ is the desuspension $\Sigma^{-V}(S^{V+1} \setminus X)$.

In the nonequivariant case, the remarkable Alexander duality theorem identifies $H^{-*}(DX) \cong H_*(X)$, by

$$H^{-*}(DX) = H^{-*}(\Sigma^{-n}(S^{n+1} \setminus X)) = H^{n-*}(S^{n+1} \setminus X) \cong H_*(X).$$

Though the double dual involved in writing $H^{-*}(DX)$ outputs a homology theory (a composition of two contravariant functors is a covariant functor), Alexander duality identifies it as the homology theory we started with, so this operation doesn't produce anything new. Note that we used negatively graded cohomology, H^{-*} , so that the gradings in this formula would work out.

However, in the equivariant case, something altogether new happens: applying the above formula to orbits,

$$H_G^{-*}(D_G((G/H)_+)) = H_G^{-*}((G/H)_+) = H^{-*}(BH),$$

which is not $H_*^G((G/H)_+) = H_*(BH)$, nor some degree shift of it: $H^{-*}(BH)$ is concentrated in nonpositive degrees, and $H_*(BH)$ is concentrated in nonnegative degrees. Thus $H_G^{-*}(D_G X)$ is an altogether different homology theory, which deserves to be called *coBorel homology*. [GM95] and [Man16] write this $cH_*^G(X)$. Following [Jon87], and because $cH_*^G(X)$ is usually supported in the negative direction, we prefer to denote it $H_G^-(X)$. Correspondingly, we write $H_G^+(X) = H_*^G X$, as this is usually supported in the positive direction.

Using the cap product of cohomology and homology, and pulling back cohomology classes from BG , we find that $H_*^+(X)$ is a module over $H^{-*}(BG)$ (note the negative grading, as cohomology classes contract against homology, decreasing degree). Using the cup product instead, $H_*^-(X)$ is also a module over $H^{-*}(BG)$.

If G is finite or connected, is a homomorphism relating these two homology theories, the *norm map*:

$$N_G : H_G^+(X)[\dim G] \rightarrow H_G^-(X).$$

Note the degree shift by $\dim G$; we think of the norm map as a sort of averaging operator. For a general compact Lie group there is a twist involving the character $\pi_0 G \rightarrow \pm 1$ given by the determinant of the adjoint representation (see Theorem A.22).

There is a final equivariant homology theory, *Tate homology*, written $H_*^\infty(X)$ and fitting into a long exact sequence

$$\dots \rightarrow H_{*-\dim G}^+(X) \xrightarrow{N_G} H_*^-(X) \rightarrow H_*^\infty(X) \rightarrow H_{*-\dim G-1}^+(X) \rightarrow \dots$$

Remarkably, the norm map is an $H^{-*}(BG)$ -module homomorphism, and Tate homology also has the structure of an $H^{-*}(BG)$ -module for which all maps in the above exact sequence are module homomorphisms.

The other essential property of Tate homology is that $H_G^\infty(G) = 0$, so that Tate homology of a finite G -CW complex can be calculated (essentially) from its subcomplex of points with nontrivial stabilizer. This is extremely useful for calculation, as we will see later in the context of I^∞ .

The following table summarizes the analogies between equivariant homology of G -spaces and the various Floer homologies of 3-manifolds. The monopole Floer homology is defined to be a sort of S^1 -equivariant Floer homology, but the S^1 symmetry is less visible in Heegaard Floer theory. In any case, the monopole Floer and Heegaard Floer homology groups are modules over $H^{-*}(BS^1; \mathbb{Z}) = \mathbb{Z}[U]$, where $|U| = -2$.

Heegaard Floer	Monopole Floer	Instanton Floer	Equivariant homology
$HF_*^+(Y)$	$\widehat{HM}_*(Y)$	$I_*^+(Y)$	$H_*^+(X) = H_*(EG_+ \wedge_G X)$
$HF_*^-(Y)$	$\widehat{HM}_*(Y)$	$I_*^-(Y)$	$H_*^-(X) = H_G^{-*}(D_G X)$
$HF_*^\infty(Y)$	$\widehat{HM}_*(Y)$	$I_*^\infty(Y)$	$H_*^\infty(X)$
$\widehat{HF}_*(Y)$	$\widehat{HM}_*(Y)$	$\tilde{I}_*(Y)$	$H_*(X)$

In our setting, we have a functorial dg-module with dg-algebra action

$$\widetilde{CI}(Y, \tilde{E}, \pi; R) \curvearrowright C_*^{\text{gm}}(SO(3); R),$$

acting from the right because $\tilde{\mathcal{B}}_E$ carries a right $SO(3)$ -action; these are right orbits. This is a reasonable notion of ‘chain complex with action of $SO(3)$ ’, and we would like to obtain invariants that behave like equivariant homology of a G -space X . The majority of the appendix develops chain complexes $C^\bullet(A; M)$ whose homology gives us equivariant homology theories for \mathbb{Z} -graded dg-modules over a dg-algebra.

The definitions for $\mathbb{Z}/8$ -graded complexes are more delicate, and carried out in Section A.8. Here we take the point of view that the periodic filtration on the module M should give rise to a periodic filtration on $C_{SO(3)}^\bullet(M)$, and to do so we

require that $\mathrm{gr}_p \tilde{M}$ is bounded. Our complexes $\tilde{C}I$ are equipped with a periodic filtration, and each $\mathrm{gr}_p \tilde{C}I_{\mathrm{unr}}(Y, \tilde{E}, \pi; R)$ is bounded. In particular, Theorem A.29 is readymade for us to apply to $\tilde{C}I$.

Theorem 6.22. *Suppose R is a PID. There is a dg-algebra $\hat{C}^- \left(C_*^{\mathrm{gm}}(SO(3); R) \right)$ whose homology algebra is graded isomorphic to $H^{-*}(BSO(3); R)$. We denote this dg-algebra by C^- .*

There are functors

$$CI^\bullet : \mathrm{Cob}_{3,b}^{U(2),w,\pi} \rightarrow \mathrm{Ho} \left(C^- \text{-Mod}_R^{r,\mathbb{Z}/8} \right),$$

where $\bullet \in \{+, -, \infty\}$, given by applying the constructions

$$C^\bullet \left(C_*^{\mathrm{gm}}(SO(3); R), \tilde{C}I_{\mathrm{unr}}(Y, \tilde{E}, \pi; R) \right)$$

of Theorem A.29 for periodically filtered modules with finite support on each associated graded piece. After taking homology, we then have functors

$$I^\bullet : \mathrm{Cob}_{3,b}^{U(2),w} \rightarrow \mathrm{Mod}_{H^{-*}(BSO(3);R)}^{r,\mathbb{Z}/8}.$$

These fit into an exact triangle

$$\cdots \rightarrow I^+(Y, E; R) \xrightarrow{[3]} I^-(Y, E; R) \rightarrow I^\infty(Y, E; R) \xrightarrow{[-4]} I^+(Y, E; R) \rightarrow \cdots$$

where all arrows are $H^{-*}(BSO(3); R)$ -module maps.

Given any other bounded, non-negatively graded dg-algebra A and right A -module M , equipped with a periodic filtration each of whose associated graded groups $\mathrm{gr}_p M$ is bounded, there are A -homology groups $H_A^\bullet(M)$. If there is an algebra map which is a quasi-isomorphism

$$f : C_*^{\mathrm{gm}}(SO(3); R) \rightarrow A,$$

and an f -equivariant filtered module map $g : \tilde{C}I(Y, \tilde{E}, \pi; R) \rightarrow M$ whose induced map on associated graded complexes is a quasi-isomorphism, then there is an induced canonical isomorphism $I^\bullet \cong H_A^\bullet(M)$, equivariant under the actions of $H_A^-(R)$ under the induced isomorphism $H^{-}(BSO(3); R) \cong H_A^-(R)$.*

We also have spectral sequence for the equivariant homology groups. As with the non-equivariant homology, we state this in slightly more detail than in Theorem A.29, point (4).

Theorem 6.23. *For each of $\bullet \in \{+, -, \infty\}$, if (Y, E) is a weakly admissible bundle equipped with regular perturbation π , there is a $(\mathbb{Z}/8, \mathbb{Z})$ -bigraded spectral sequence of $H^{-*}(BSO(3); R)$ -modules from*

$$E_{p,q}^1 = \bigoplus_{\substack{\alpha \in \mathcal{C}_\pi \\ \mathrm{gr}(\alpha) = p}} H_{SO(3)}^\bullet(\alpha; R)$$

whose target is the associated equivariant instanton homology $I^\bullet(Y, E, \pi; R)$.

For I^- , we have an isomorphism of the unrolled (\mathbb{Z}, \mathbb{Z}) -bigraded spectral sequence $E_{p,q}^\infty \cong \mathrm{gr}_p I_{p+q}^-, \mathrm{unr}$. For any of the theories, a filtered chain map $CI^\bullet(Y, E, \pi) \rightarrow CI^\bullet(Y', E', \pi')$ which induces an isomorphism on any finite page E^r induces an isomorphism on I^\bullet .

Proof. The existence of this spectral sequence, and the fact that we may detect quasi-isomorphisms by isomorphisms on a finite page E^r , is precisely Theorem A.29, point (4). As for the specific calculation of the E^1 page, our associated graded complex

$$\mathrm{gr}_p \widetilde{CI}_{\mathrm{unr}}(Y, E, \pi; R) = \bigoplus_{\substack{\alpha \in \mathcal{C}_\pi \\ \mathrm{gr}(\alpha) = p}} C_*^{\mathrm{gm}}(\alpha; R)$$

is a bounded $C_*^{\mathrm{gm}}(SO(3); R)$ -module. By Proposition A.28, there is a chain of quasi-isomorphisms connecting $C_*^{\mathrm{gm}}(SO(3); R)$ to the algebra $C_*(SO(3); R)$, and similarly a chain of equivariant quasi-isomorphisms between $C_*^{\mathrm{gm}}(\alpha; R)$ to $C_*(\alpha; R)$. In particular, we may identify

$$H^\bullet(C_*^{\mathrm{gm}}(SO(3)), C_*^{\mathrm{gm}}(\alpha; R)) \cong H_{SO(3)}^\bullet(\alpha; R)$$

using Theorem A.26. ■

As before, these spectral sequences will become significantly more computable after the introduction of the Donaldson models \overline{DCI}^\pm for CI^- and CI^+ in Section 7.2.

Remark 6.5. It is important to note that while the homology theory for \mathbb{Z} -graded complexes $H_A^+(M)$ described at first in the appendix sends A -equivariant quasi-isomorphisms $M \rightarrow M'$ to isomorphisms on H_A^+ , this is *not true* for arbitrary filtered A -equivariant quasi-isomorphisms between $\mathbb{Z}/8$ -graded complexes. An acyclic $\mathbb{Z}/8$ -graded complex M whose associated graded complex is not acyclic need not have $H_A^+(M) = 0$.

What remains true, and will be used extensively, is that if M, M' are $\mathbb{Z}/8$ -graded A -modules equipped with a periodic filtration, and $f : M \rightarrow M'$ is a filtered dg-module map inducing an isomorphism on the associated graded homology groups (the E^1 page of the spectral sequence), then f induces an isomorphism on $H_A^+(M)$.

The reason this happens is that these objects are defined via a completion with respect to the given filtration, which ensures that the associated spectral sequences converge. However, quasi-isomorphisms between the uncompleted complexes which are not quasi-isomorphisms on the associated graded do not necessarily induce quasi-isomorphisms between the completions.

7. EXAMPLES, CALCULATIONS, AND COMPARISONS

For the entirety of Section 7, we choose elements in each orientation set $\Lambda(\alpha)$ arbitrarily, and suppress these orientation sets from notation.

7.1. Equivariant instanton homology for admissible bundles. When E is a nontrivial admissible $U(2)$ -bundle over an oriented 3-manifold Y , all critical orbits are free. The complex $\widetilde{CI}(Y, E; R)$ resembles the equivariant Morse complex of a finite-dimensional free $SO(3)$ -manifold M , so our heuristic is that the equivariant homology groups $I^\bullet(Y, E; R)$ should behave like the equivariant homology $H_G^\bullet(M)$:

$$\begin{aligned} H_G^-(M; R)[\dim M] &\cong H_G^+(M; R) \cong H(M/G; R) \\ H_G^\infty(M; R) &= 0. \end{aligned}$$

Theorem 7.1. *Suppose E is a nontrivial admissible bundle over a 3-manifold Y . Then $I^\infty(Y, E) = 0$, and if $CI(Y, E)$ is Floer's instanton chain complex for admissible bundles, there is a natural quasi-isomorphism $CI^+(Y, E; R) \rightarrow CI(Y, E; R)$. In particular, the equivariant Floer homology $I^+(Y, E)$ of a nontrivial admissible bundle is Floer's original instanton homology group $I(Y, E)$. Furthermore, there is also a functorial quasi-isomorphism $CI(Y, E; R)[3] \rightarrow CI^-(Y, E; R)$.*

Proof. We first prove the statement for CI^∞ . The E^1 page of the index spectral sequence is

$$\bigoplus_{\substack{\alpha \in \mathfrak{C}, n \in \mathbb{Z} \\ n = i(\alpha) \pmod{8}}} H_*^\infty(\alpha; R)[i(\alpha)].$$

Now all orbits are free orbits $\alpha \cong SO(3)$, and so we can calculate this as

$$\bigoplus H_*^\infty(SO(3); R)[i(\alpha)].$$

That $H_{SO(3)}^\infty(SO(3))$ vanishes is one of the defining features of Tate homology: see Proposition A.12 (4). So the E^1 page is identically zero, and by the I^∞ part of Theorem 6.23 we see that the map $0 \rightarrow CI^\infty(Y, E; R)$ is a quasi-isomorphism, as desired.

Floer's instanton complex $CI(Y, E)$ is defined as the free R -module generated by critical orbits $\alpha \in \mathfrak{C}$, the relative grading defined still as $\text{gr}_z(\alpha, \beta)$; the differential counts points in zero-dimensional moduli spaces of *unframed* instantons between α and β . As a graded \mathcal{C}^- -module,

$$\begin{aligned} CI^+ &= \bigoplus_{\alpha \in \mathfrak{C}} C^+(\alpha; R)[i(\alpha)], \\ CI &= \bigoplus_{\alpha \in \mathfrak{C}} R[i(\alpha)]. \end{aligned}$$

The differential is a sum of the induced differential on each factor $C^+(\alpha; R)$ and the maps $C^+(\alpha; R) \rightarrow C^+(\beta; R)$ induced by the $C_*^{\text{gm}}(SO(3))$ -equivariant chain maps

$$\sigma \mapsto \sigma \times_{e_-} \overline{\mathcal{M}}(\alpha, \beta).$$

Recall from Lemma A.21 that $H_{SO(3)}^+(SO(3); R) = R$, concentrated in degree zero. There is a chain map $\varphi : CI^+(Y, E; R) \rightarrow CI(Y, E; R)$ given on the summand for each critical orbit by the augmentation map $C_{SO(3)}^+(\alpha; R) \rightarrow R$ induced by the point-counting map $C_*(\alpha; R) \rightarrow R$; this map kills everything in degree larger than $i(\alpha)$ and sends $C_0^+(\alpha; R) = C_0(\alpha; R) \rightarrow H_{i(\alpha)}^+ = R$ by the augmentation.

To see that φ is a chain map, note that if $\sigma \in CI_k^+(Y, E; R)$ is not sent to zero, it can be written as a sum of points in

$$\bigoplus_{\substack{\alpha \in \mathfrak{C}_\pi \\ i(\alpha) = k}} C_0(\alpha; R).$$

So it suffices to check that $\varphi(\tilde{d}\sigma) = d\varphi(\sigma)$ for $\sigma = p$, where p is a point in α . The only components of $d\sigma \in CI^+$ not automatically sent to zero are those in some $C_0(\beta; R) \subset C^+(\beta; R)$, which arise from taking fiber products with moduli spaces $\overline{\mathcal{M}}(\alpha, \beta)$ of dimension 3. Because the $SO(3)$ action on $\overline{\mathcal{M}}$ is free, the count of points of $p \times_\alpha \overline{\mathcal{M}}(\alpha, \beta)$ is the same as that of $\overline{\mathcal{M}}(\alpha, \beta)/SO(3)$, Floer's unframed moduli space of instantons. In particular, we find that φ is a chain map, as desired.

Now, on the unrolled complexes, this map is a filtered chain map for the tautological periodic filtration of $CI(Y, E; R)$ by degree, and Theorem 6.23 says that the above map induces an isomorphism

$$E^1(CI^+(Y, E; R)) \rightarrow E^1(CI(Y, E; R)) = CI(Y, E; R).$$

Therefore, by the comparison theorem, the induced map

$$I^+(Y, E; R) \rightarrow I(Y, E; R)$$

is an isomorphism. That these isomorphisms fit into a commutative square with the corresponding cobordism maps is proved the same way.

The statement about I^- follows because of the exact triangle relating I^∞ , I^+ , and I^- and the fact that I^∞ vanishes; one may alternately define a chain map $CI[3] \rightarrow CI^-$ explicitly, almost precisely as above except including the fundamental class instead of a point class. \blacksquare

Suppose $\frac{1}{2} \in R$; then $H^{-*}(BSO(3); R) = R[U]$, where $|U| = -4$. Floer's $I(Y, E; R)$ also carries an action by this ring (when $\frac{1}{2} \in R$), and we should check that the above map preserves the U -action. Because the U -action carries things down vertically in the spectral sequence, but $E^2(CI^+)$ is here concentrated on a single horizontal line, this is not a theorem well-suited to a spectral sequence proof. Instead, we must get our hands dirty at the chain level. This will be delayed until the following section, where we construct a simpler chain-level model of \tilde{CI} .

7.2. Comparison with Donaldson's theory. In this section, we describe a finite-dimensional complex computing framed instanton homology \tilde{I} , and then explain how to use it to calculate the equivariant instanton groups. While it seems likely that this is possible over any PID R , the situation is drastically simpler when $\frac{1}{2} \in R$. So for the rest of this section, R is a PID in which 2 is invertible; the ring R will be dropped from the notation whenever possible.

If 2 is invertible in R , then $H_*(SO(3); R)$ is isomorphic to a single copy of R in degrees 0 and 3 and zero otherwise. We denote this R -algebra $\Lambda(u)$, the exterior algebra on a generator u in degree 3. There is a dg-algebra homomorphism

$$i : \Lambda(u) \rightarrow C_*^{\text{gm}}(SO(3); R)$$

given by picking out the identity in degree 0 and the fundamental class in degree 3; this map induces the identity on homology. By Proposition A.13, because i is a quasi-isomorphism, we have

$$H^\bullet(C_*^{\text{gm}}(SO(3); R), \tilde{CI}) \cong H^\bullet(\Lambda(u), \tilde{CI}).$$

For the purposes of computing the equivariant instanton homology groups I^+ , I^- , and I^∞ , it therefore suffices to consider \widetilde{CI} as a module over $\Lambda(u)$.

Our goal is to write down a differential on the $\mathbb{Z}/8$ -graded R -module

$$DCI(Y, E, \pi; R) = \bigoplus_\alpha H_*(\alpha; R)[i(\alpha)],$$

which has a periodic filtration by index of the orbit α . This differential should *decrease* the filtration, and there should be a map $\widetilde{CI}(Y, E, \pi) \rightarrow DCI$ which is the identity on the E^1 page, or something like it. To produce this, we use what is called the *homological perturbation lemma*. Recall its statement:

Lemma 7.2 (Homological perturbation lemma). *Suppose (C, d) and (C', d') are chain complexes, equipped with an inclusion $i : C \hookrightarrow C'$ and a projection $p : C' \rightarrow C$ so that $pi = 1$, both quasi-isomorphisms, as well as a degree 1 map $h : C' \rightarrow C$ serving as a homotopy witnessing this. That is,*

$$ip = 1 + d'h + hd'.$$

This data is called deformation retract data, and is depicted

$$(C, d) \begin{array}{c} \xrightarrow{i} \\ \xleftarrow{p} \end{array} (C', d') \circlearrowleft h.$$

Suppose that C' is equipped with a deformation: an additional map $\delta : C' \rightarrow C'$ so that $(d' + \delta)^2 = 0$. Suppose that $h^n = 0$ for sufficiently large n . Write $A = \sum_{n=0}^{\infty} (\delta h)^n \delta$.

Then $(C, d + pAi)$ is a chain complex, equipped with deformation retract data

$$(C, d + pAi) \begin{array}{c} \xrightarrow{i+hAi} \\ \xleftarrow{p+pAh} \end{array} (C', d' + \delta) \circlearrowleft h + hAh.$$

If C and C' are dg-modules over a dga, and all of i, p, h , and δ are dg-module homomorphisms, then the same is true of the perturbed deformation retract data. In particular $(C, d + pAi)$ is homotopy equivalent to $(C', d' + \delta)$ as a dg-module.

If C and C' are filtered, and i, p, h , and δ all preserve the filtration (in the sense that $f(F_k) \subset F'_k$ for all k), then the same is true for the perturbed deformation retract data. In particular, $(C, d + pAi)$ is a filtered complex, which is filtered homotopy equivalent to $(C', d' + \delta)$.

A very explicit reference is [Cra04].

To apply the homological perturbation lemma, first recall that we may write the differential on \widetilde{CI} as $d + d_M$. The first term is the usual boundary operator on geometric chains, and the second term is the contribution from fiber products with moduli spaces. We want to apply the lemma to $(DCI, 0)$ and (\widetilde{CI}, d) , with $\delta = d_M$.

Choose a basepoint $b_\alpha \in \alpha$ for each critical orbit α ; if α is an $SO(2)$ -reducible, we demand that the stabilizer of b_α is the standard $SO(2) \subset SO(3)$. Later, when considering the $SU(2)$ -model, this will translate into the demand that the stabilizer is the standard $U(1) \subset SU(2)$.

Our inclusion map

$$i : DCI \rightarrow \widetilde{CI}$$

is given on $H_0(\alpha; R) \rightarrow C_*^{\text{gm}}(\alpha; R)$ by sending the generator to the basepoint b_α , and on $H_{\dim \alpha}(\alpha; R) \rightarrow C_*^{\text{gm}}(\alpha; R)$ by sending a generator to the corresponding fundamental class of α . (Note that the choice of fundamental class depends on

an orientation of α ; we are making this choice by choosing a generator of top homology.)

We will proceed to the deformation retract data in two steps.

First, we will pass from \widetilde{CI} to a slightly different but quasi-isomorphic chain complex, for which we can explicitly define p and h . Then we will modify the differential on \widetilde{CI} itself to be chain-homotopy equivalent to the complex we started with, but in such a way that we have

$$p(d_M h)^n d_M i = 0 \text{ for } n > 0.$$

This implies that the perturbed differential pAi on DCI consists only of the term $pd_M i$, which has a concrete description.

To begin, note that the map $H_*(SU(2); R) \rightarrow H_*(SO(3); R)$ is an isomorphism. It will be important in what follows to first replace \widetilde{CI} with a quasi-isomorphic complex called the $SU(2)$ model. Throughout this section, in an attempt to minimize notational overload, we write $\overline{\mathcal{M}}_{\alpha\beta}$ for the compactified moduli space of framed instantons $\overline{\mathcal{M}}_{E,z,\pi}(\alpha, \beta)$ for the unique homotopy class z with

$$-2 \leq \text{gr}_z(\alpha, \beta) \leq 5,$$

as we ignore all larger moduli spaces. This space can only be nonempty if that grading is positive. Its quotient $\overline{\mathcal{M}}_{\alpha\beta}/SO(3)$ is denoted simply $\mathcal{M}_{\alpha\beta}$.

The main difficulty is that while we may define the holonomy from one point to another as an element of $SU(2)$ (using the $U(2)$ -model for the configuration space and the fact that we are only quotienting by even gauge transformations), it is not true that we can lift the endpoint maps $\overline{\mathcal{M}}_{\alpha\beta} \rightarrow \alpha \times \beta$ to maps to $SU(2)$ when α or β is irreducible. For instance, consider the case that $\mathcal{M}_{\alpha\beta}$ is a single point and α, β are irreducible; then $\overline{\mathcal{M}}_{\alpha\beta} \rightarrow \alpha \times \beta$ may be identified with the diagonal which is nontrivial on π_1 , and therefore cannot be lifted to codomain $SU(2) \times SU(2)$.

What we would like to do is the following. First, replace every irreducible $SO(3)$ -orbit α with its universal cover $\tilde{\alpha}$; this carries a canonical $SU(2)$ action lifting the $SO(3)$ action on α . Second, replace every moduli space $\overline{\mathcal{M}}_{\alpha\beta}$ with the $\pi_1(\alpha) \times \pi_1(\beta)$ cover corresponding to the homomorphism

$$e_* : \pi_1(\overline{\mathcal{M}}_{\alpha\beta}) \rightarrow \pi_1(\alpha \times \beta);$$

call this cover $\overline{\mathcal{M}}_{\alpha\beta}^{\text{cov}}$. This cover carries a lift to a map

$$\overline{\mathcal{M}}_{\alpha\beta}^{\text{cov}} \rightarrow \tilde{\alpha} \times \tilde{\beta},$$

unique up to isomorphism. However, when β is irreducible, the fiber product

$$\overline{\mathcal{M}}_{\alpha\beta}^{\text{cov}} \times_{\tilde{\beta}} \overline{\mathcal{M}}_{\beta\gamma}^{\text{cov}}$$

is *not* isomorphic to a boundary stratum of $\overline{\mathcal{M}}_{\alpha\gamma}^{\text{cov}}$. For convenience of notation, call this stratum $\partial_\beta \overline{\mathcal{M}}_{\alpha\gamma}^{\text{cov}}$. Instead, there is a canonical covering map

$$\overline{\mathcal{M}}_{\alpha\beta}^{\text{cov}} \times_{\tilde{\beta}} \overline{\mathcal{M}}_{\beta\gamma}^{\text{cov}} \rightarrow \partial_\beta \overline{\mathcal{M}}_{\alpha\gamma}^{\text{cov}}$$

with fiber of cardinality $|\pi_1(\beta)|$.

Suppose we had some rule allowing us to identify the fundamental class (in some sense) of a compact manifold with half of the fundamental class of its double cover. Formally, set

$$\overline{\mathcal{M}}'_{\alpha\beta} := \frac{1}{|\pi_1(\alpha)|} \overline{\mathcal{M}}_{\alpha\beta}^{\text{cov}}.$$

Then, given this ‘rule’, we would be able to write

$$\overline{\mathcal{M}}'_{\alpha\beta} \times_{\tilde{\beta}} \overline{\mathcal{M}}'_{\beta\gamma} = \partial_{\beta} \overline{\mathcal{M}}'_{\alpha\gamma},$$

giving us the relation necessary to define the instanton chain complex. We therefore aim to make this idea precise.

Recall the definition of the chain complex $C_*^{\text{gm}}(\alpha; R)$ from Section 6.1. Whenever $2 \in R$ is invertible, we will construct a further quotient of this suitable to our purposes, written $C_*^{\text{gm},2}(\alpha; R)$. Whenever we have a basic chain given as a composite

$$\sigma' : P' \xrightarrow{p} P \xrightarrow{f} \alpha,$$

where p is a covering map with fiber of order 2^n , we quotient by the additional relation

$$[\sigma'] = 2^n[\sigma].$$

Because the composite of two finite covers is again a finite cover, no new additional relations are imposed by chains of covers $P'' \rightarrow P' \rightarrow P$ that we have not already quotiented by. The resulting chain complex is again natural under smooth maps.

A few verifications are in order, all of which are straightforward enough that we state them without proof. First, $C_*^{\text{gm},2}(\alpha; R)$ satisfies the Eilenberg-Steenrod axioms for admissible pairs, with the proof essentially unchanged from Theorem 6.7. Second, the fiber product map

$$\times_X : C_*^{\text{gm}}(\alpha; R) \rightarrow C_*^{\text{gm}}(\beta; R)$$

induced by a pair of maps (with left map a submersion)

$$\alpha \leftarrow X \rightarrow \beta$$

descends to $C_*^{\text{gm},2}$: this follows because the fiber product with a covering space is a cover of the fiber product. Finally, the complex $C_*^{\text{gm},2}(\alpha; R)$ is still a free R -module: the set of chains naturally forms a partial ordered set with $\sigma' \geq \sigma$ if σ' covers σ with fiber of cardinality 2^n ; considering this as a graph by forgetting the direction of edges, a choice of basic chain in each path-component of this graph gives a free generating set. Here it was crucial that 2 is invertible in R .

Now we define the $SU(2)$ model exactly as suggested above. As a relatively graded $C_*^{\text{gm}}(SU(2); R)$ -module, it is given by

$$\widetilde{CI}^{SU(2)}(Y, E, \pi; R) = \bigoplus_{\alpha \in \mathfrak{C}_{\pi}} C_*^{\text{gm},2}(\tilde{\alpha}; R),$$

and the differential is the sum of the usual boundary operator and fiber products with the $\overline{\mathcal{M}}'_{\alpha\beta}$ defined above. Note that here the formal definition given above makes sense, with a factor of $\frac{1}{2}$ when $\alpha \cong SO(3)$; this means we take the usual fiber product with $\overline{\mathcal{M}}_{\alpha\beta}^{\text{cov}}$ and then multiply by $\frac{1}{2}$ if $\alpha \cong SO(3)$. Because the expected decomposition of $\partial \overline{\mathcal{M}}'_{\alpha\beta}$ holds, this defines a chain complex (in fact, a dg-module).

Lemma 7.3. *The projection $r : C_*^{\text{gm}}(SU(2); R) \rightarrow C_*^{\text{gm}}(SO(3); R)$ induced by the standard covering homomorphism $r : SU(2) \rightarrow SO(3)$ is a dg-algebra homomorphism which is a quasi-isomorphism. There is a zig-zag of dg-module quasi-isomorphisms between \widetilde{CI} and $\widetilde{CI}^{SU(2)}$.*

Proof. That r is a dg-algebra homomorphism follows immediately from the fact that the corresponding map $r : SU(2) \rightarrow SO(3)$ is a homomorphism of groups. To construct the zig-zag, first consider the $C_*^{\text{gm}}(SO(3); R)$ -module given by

$$\widetilde{CI}' = \bigoplus C_*^{\text{gm},2}(\alpha; R);$$

the natural projection $\widetilde{CI} \rightarrow \widetilde{CI}'$ is a module homomorphism and a quasi-isomorphism, because it is an isomorphism on the E^1 page of the corresponding index spectral sequence.

There is a map

$$f : \widetilde{CI}^{SU(2)} \rightarrow \widetilde{CI}',$$

defined as follows. If we denote $\pi_\alpha : \tilde{\alpha} \rightarrow \alpha$ for the covering map on each orbit, and $\sigma : P \rightarrow \tilde{\alpha}$ is a chain, then we set

$$f(\sigma) = \frac{1}{|\pi_1(\alpha)|} \pi_\alpha \sigma.$$

The map f is clearly equivariant with respect to r . Next we should see that π gives a chain map. Observe that if $\sigma : P \rightarrow \tilde{\alpha}$ is a chain, then

$$\sigma \times_{\tilde{\alpha}} \overline{\mathcal{M}}_{\alpha\beta}^{\text{cov}} = |\pi_1 \alpha| (\pi_\alpha \sigma) \times_\alpha \overline{\mathcal{M}}_{\alpha\beta}^{\text{cov}}$$

as chains equipped with maps to $\tilde{\beta}$; here the factor $|\pi_1 \alpha|$ indicates that the former is the disjoint union of that many copies of the latter. Next, we have the canonical covering map

$$(\pi_\alpha \sigma) \times_\alpha \overline{\mathcal{M}}_{\alpha\beta}^{\text{cov}} \rightarrow (\pi_\alpha \sigma) \times_\alpha \overline{\mathcal{M}}_{\alpha\beta},$$

with fibers of cardinality $|\pi_1 \alpha| \cdot |\pi_1 \beta|$. Combining these with the definition of

$$\overline{\mathcal{M}}'_{\alpha\beta} = \frac{1}{|\pi_1 \alpha|} \overline{\mathcal{M}}_{\alpha\beta}^{\text{cov}},$$

we have the relation

$$\sigma \times_{\tilde{\alpha}} \overline{\mathcal{M}}'_{\alpha\beta} = \frac{|\pi_1 \beta|}{|\pi_1 \alpha|} (\pi_\alpha \sigma) \times_\alpha \overline{\mathcal{M}}_{\alpha\beta}$$

as elements of $C_*^{\text{gm},2}(\beta; R)$.

This is precisely the statement that f is a chain map. ■

In a flagrant abuse of notation, from now on we will write $\tilde{\alpha}$ simply as α and the new moduli space $\overline{\mathcal{M}}'_{\alpha\beta}$ simply as $\overline{\mathcal{M}}_{\alpha\beta}$, just as before, unless otherwise stated. As we will primarily work in the $SU(2)$ -model for the rest of this section, we hope that this will cause no confusion, and will make the distinction clear when both moduli spaces are in play.

To define the map p , choose basepoints $q_\alpha \in \alpha$, with the property that each q_α corresponding to an $U(1)$ -reducible α has a different stabilizer, which is not the standard $U(1)$.

Recall from Definition 6.7 and Lemma 6.5 that, given a countable family \mathcal{F} of maps from δ -chains to a space X , the geometric chain complex has a quasi-isomorphic subcomplex $C_*^{\text{gm},\mathcal{F}}(X; R)$, spanned by those chains which are transverse on every stratum to the countable family of maps. If X is an $SO(3)$ -orbit, this is a $\Lambda(u)$ -invariant subcomplex, because the evaluation map $SO(3) \rightarrow X$ is a submersion. This means that the product of any chain with the fundamental class is transverse to any element of \mathcal{F} .

To define the map

$$p : C_*^{\text{gm}, \mathcal{F}}(\alpha; R) \rightarrow H_*(\alpha; R),$$

take \mathcal{F} to include the basepoint q_α . In degree zero, any 0-dimensional δ -chain has an underlying oriented point, and the projection is given by counting these with sign. In top degree, we send a δ -chain $\sigma : P \rightarrow \alpha$ to the signed count $\#\sigma^{-1}(q_\alpha)$. Note that to define the sign of this signed count, we need to choose an underlying orientation of α , which determines the generator of $H_{\dim \alpha}(\alpha; R)$ we should send 1 to. The homomorphism we obtain is independent of this choice of orientation of α .

Unfortunately, the fiber products with moduli spaces do not send chains transverse to q_α to chains transverse to q_β . So this submodule of \widetilde{CI} is *not* a subcomplex. To fix this, we enlarge the collection \mathcal{F} that our chains must be transverse to.

Denote $\overline{\mathcal{M}}_{\alpha\beta}^q$ for the fiber above q_β of the endpoint map

$$\overline{\mathcal{M}}_{\alpha\beta} \rightarrow \beta.$$

More generally, given a sequence of orbits $\gamma_1, \dots, \gamma_n$, denote $\overline{\mathcal{M}}_{\alpha\gamma\beta}^q$ denote the fiber above q_β of the map

$$\overline{\mathcal{M}}_{\alpha\gamma_1} \times_{\gamma_1} \cdots \times_{\gamma_n} \overline{\mathcal{M}}_{\gamma_n\beta} \rightarrow \beta.$$

For the orbit α , we let the collection \mathcal{F} consist of q_α and all of the endpoint maps of moduli spaces

$$\overline{\mathcal{M}}_{\alpha\gamma\beta}^q \rightarrow \alpha.$$

Because this collection is closed under fiber products with moduli spaces on the left, the condition of being transverse to \mathcal{F} is preserved by d_M . So for this collection \mathcal{F} , we have a $\Lambda(u)$ -subcomplex

$$\widetilde{CI}^{\mathcal{F}}(Y, E, \pi; R) \subset \widetilde{CI}(Y, E, \pi; R)$$

for which the projection map p above is defined. Running the index spectral sequence, and using that the inclusion of $C_*^{\text{gm}, \mathcal{F}}(\alpha; R) \rightarrow C_*^{\text{gm}}(\alpha; R)$ is a quasi-isomorphism for each orbit, we see that this subcomplex $\widetilde{CI}^{\mathcal{F}}$ is equivariantly quasi-isomorphic to the whole thing. Furthermore, it is a free R -submodule; its generators are precisely those generators of \widetilde{CI} for which all strata are transverse to the given family of δ -chains.

To ensure that the image of i is in this subcomplex, we need to demand that b_α is a regular value of $\overline{\mathcal{M}}_{\alpha\gamma\beta}^q \rightarrow \alpha$; one may symmetrically demand that q_β is the regular value of the endpoint map from the fiber above b_α . It is not difficult to ensure that the q_α satisfy this, by an application of Sard's theorem.

We will define the map h in the course of Lemma 7.5, but for intuition it is convenient to imagine it as follows. In degrees up to $\dim \alpha - 1$, one imagines that h is given by the 'cone to b_α ' map, using stereographic projection based at q_α . In degree $\dim \alpha$, it is chosen arbitrarily to satisfy the requirements that

$$ip - 1 = d'h + hd'$$

and $uh = hu$. The latter condition is automatically satisfied for degree reasons except possibly on 0-chains on an irreducible orbit; in that case, it becomes the demand $h([SU(2)]) = 0$. In degree $\dim \alpha + 1$, the map h is again zero for degree reasons.

The issue is that h , as defined here, does not preserve transversality to \mathcal{F} (only to q_α). We will resolve this when the time comes; for now, we discuss what would happen if h did indeed preserve transversality to \mathcal{F} .

This is enough to apply the homological perturbation lemma and get *some* result, a differential on $DCI(Y, E, \pi; R)$ which is $\Lambda(u)$ -equivariantly homotopy equivalent to $\widetilde{CI}^{\mathcal{F}}(Y, E, \pi; R)$. Unfortunately, the differential is still more complicated than we would like: it takes the form

$$p \sum_{n=0}^{\infty} (d_M h)^n d_M i,$$

where d_M is the part of the differential coming from fiber products with moduli spaces. It is not impossible that many components of this may be nonzero: one takes a fiber product of the basepoint b_α with some moduli space, pushes it forward to β , cones towards b_β , takes another fiber product, and then counts intersections with b_γ . One could easily imagine that, starting and ending at irreducibles α, γ , one takes the fiber product with a moduli space that increases the dimension by 2, cones towards b_β , and then takes a fiber product with a 0-dimensional moduli space and gets a result that has nonzero degree above b_γ .

To resolve this, we should reduce the influence of lower-dimensional moduli spaces, an approach suggested by [Don02, Section 7.3.2]. We will be modifying the endpoint maps, as in the following two lemmas. The proofs are somewhat technical, and it is not necessary to understand the arguments to read the rest of the section; we invite the reader to skip them on a first pass.

Lemma 7.4. *Suppose $e_{\alpha\beta} : \overline{\mathcal{M}}_{\alpha\beta} \rightarrow \alpha \times \beta$ are the equivariant endpoint maps associated to moduli spaces of instantons on a weakly admissible bundle (Y, E) with respect a fixed perturbation π .*

Suppose we choose a collection of $e_{\alpha\beta}^t : \overline{\mathcal{M}}_{\alpha\beta} \rightarrow \alpha \times \beta$, homotopies through equivariant maps that are smooth on each stratum, one homotopy for each pair (α, β) , so that $e_{\alpha\beta}^0 = e_{\alpha\beta}$ above. Furthermore demand that these are compatible in the sense that

$$e_{\alpha\beta}^t \times_\beta e_{\beta\gamma}^t : \overline{\mathcal{M}}_{\alpha\beta} \times_\beta \overline{\mathcal{M}}_{\beta\gamma} \rightarrow \alpha \times \gamma$$

agrees with the restriction of $e_{\alpha\gamma}^t$ to that stratum of $\overline{\mathcal{M}}_{\alpha\gamma}$.

We may define a chain complex $\widetilde{CI}_t(Y, E, \pi)$, identical as a $C_{\#}^{\text{gm}}(SO(3))$ -module to \widetilde{CI} , but whose differential uses the fiber product with the same moduli spaces but endpoint maps $e_{\alpha\beta}^t$. Then $\widetilde{CI}(Y, E, \pi)$ is $\Lambda(u)$ -equivariantly quasi-isomorphic to $\widetilde{CI}_1(Y, E, \pi)$.

This lifts to an equivalence between the $SU(2)$ -models. In fact, suppose one instead chooses $SU(2)$ -equivariant homotopies $e_{\alpha\beta}^t : \overline{\mathcal{M}}'_{\alpha\beta} \rightarrow \tilde{\alpha} \times \tilde{\beta}$, which are compatible in the sense that the restriction of $e_{\alpha\gamma}^t$ to the stratum corresponding to factorizations through β is a double cover of $e_{\alpha\beta} \times_{\tilde{\beta}} e_{\beta\gamma}$. Then one gets an $SU(2)$ -equivariant homotopy equivalence between the corresponding $SU(2)$ -models.

Proof. We identify the cube Δ^n with the subset of $\Delta^n \subset [0, 1]^n$ of nondecreasing sequences, with the induced orientation. We write Δ_i^n for the subset of (t_1, \dots, t_n) with $t_i = t_{i+1}$, interpreting $t_0 = 0$ and $t_{n+1} = 1$, equipping it with the induced boundary orientation. There is a canonical identification $\Delta^{n-1} \cong \Delta_i^n$, and this identification is oriented up to a sign of $(-1)^{i+1}$.

Let $\alpha, \gamma_1, \dots, \gamma_n, \beta$ be a sequence of critical orbits. Write

$$\begin{aligned} \overline{\mathcal{M}}_{\alpha\vec{\gamma}\beta}^{\Delta^{n+1}} &\subset \Delta^n \times \mathcal{M}_{\alpha\gamma_1} \cdots \mathcal{M}_{\gamma_n\beta} \\ \overline{\mathcal{M}}_{\alpha\vec{\gamma}\beta}^{\Delta^{n+1}} &= \{(t_1, \dots, t_{n+1}), \mathbf{A}_1, \dots, \mathbf{A}_{n+1} \mid e_+^{t_i} \mathbf{A}_i = e_-^{t_{i+1}} \mathbf{A}_{i+1}\}, \end{aligned}$$

whose fiber is oriented as

$$\Delta^{n+1} \times \text{Fib}(\mathcal{M}_{\alpha\gamma_1}) \times \cdots \times \text{Fib}(\mathcal{M}_{\gamma_n\beta}),$$

where $\text{Fib}(\mathcal{M}_{\alpha\beta})$ indicates a generic fiber above a point in α . Given any oriented submanifold $S \subset \Delta^{n+1}$, restriction to S gives another fiber-oriented manifold.

A straightforward but tedious computation shows that, as fiber-oriented manifolds, we have

$$\begin{aligned} (-1)^{\dim \alpha} \partial \overline{\mathcal{M}}_{\alpha\vec{\gamma}\beta}^{\Delta^{n+1}} &\cong \bigcup_{i=0}^{n+1} \mathcal{M}_{\alpha\vec{\gamma}\beta}^{\Delta_i^{n+1}} \\ &\bigcup_{j=0}^n \bigcup_{\eta} (-1)^{\text{gr}(\alpha, \eta) + n + 1} \overline{\mathcal{M}}_{\alpha, \vec{\gamma}_{<j}, \eta, \vec{\gamma}_{>j}, \eta}^{\Delta_{j+1}^{n+2}}. \end{aligned}$$

Write $\widetilde{\mathcal{C}}I_t$ for the chain complex given by using the moduli spaces with endpoint maps e_{\pm}^t ; the underlying R -modules are all the same, it is only the differential that changes. Consider the maps $H_n : \widetilde{\mathcal{C}}I_0 \rightarrow \widetilde{\mathcal{C}}I_1$, given for each $n > 0$ by

$$H_n(\sigma) = \sum_{\gamma_1, \dots, \gamma_n, \beta} (-1)^{\sum_{i=1}^n (\text{gr}(\alpha, \gamma_i) + i - 1)} \sigma \times_{\alpha} \widetilde{\mathcal{M}}_{\alpha\vec{\gamma}\beta}^{\Delta^{n+1}}.$$

For convenience, we write the sign here as $s(n, \alpha, \vec{\gamma}, \beta)$; it satisfies the recurrence

$$s(n-1, \alpha, \vec{\gamma}_{<j}, \vec{\gamma}_{>j}, \beta) \cdot (-1)^{\text{gr}(\alpha, \gamma_j) + n - 1} = s(n, \alpha, \vec{\gamma}, \beta).$$

Our desired chain map will be $\text{Id} + \sum_{n>0} H_n$. The map H_n is identically zero for degree reasons when $n > 3$, so this sum is finite. Because H_n increases the filtration depth, we see that the induced map on associated graded complexes is the identity; so as long as this is a chain map, the usual spectral sequence arguments will imply it is a quasi-isomorphism.

First we explicitly calculate $\partial_1 H_n - H_n \partial_0$. This is given by

$$\begin{aligned} (-1)^{\dim \sigma} (\partial_1 H_n - H_n \partial_0)(\sigma) &= \sum_{\vec{\gamma}, \beta} (-1)^{\dim \alpha} s(n, \alpha, \vec{\gamma}, \beta) \sigma \times_{\alpha} \partial \mathcal{M}_{\alpha, \vec{\gamma}, \beta}^{\Delta^{n+1}} \\ &\quad - \sum_{\eta, \vec{\gamma}, \beta} s(n, \eta, \vec{\gamma}, \beta) (-1)^{(\text{gr}(\alpha, \eta) + 1)(n+1)} \mathcal{M}_{\alpha, \eta, \vec{\gamma}, \beta}^{\Delta_0^{n+2}} \\ &\quad + \sum_{\vec{\gamma}, \eta, \beta} s(n, \alpha, \vec{\gamma}, \eta) (-1)^{\text{gr}(\alpha, \eta) + n + 1} \mathcal{M}_{\alpha, \vec{\gamma}, \eta, \beta}^{\Delta_{n+2}^{n+2}}. \end{aligned}$$

The claim, then, is that the sum $\partial_1 - \partial_0 + \sum_{n>0} \partial_1 H_n - H_n \partial_0$ is a telescoping sum. Expanding out the decomposition of $\partial \mathcal{M}_{\alpha, \vec{\gamma}, \beta}^{\Delta^{n+1}}$ above, we see that this telescoping amounts to the claims that

$$\begin{aligned}
s(n, \alpha, \vec{\gamma}, \beta) &= s(n-1, \alpha, \vec{\gamma}_{<i}, \vec{\gamma}_{>i}, \beta) \cdot (-1)^{\text{gr}(\alpha, \gamma_i) + n - 1} \quad 0 < i < n \\
s(n, \alpha, \vec{\gamma}, \beta) &= s(n-1, \gamma_1, \vec{\gamma}_{>1}, \beta) \cdot (-1)^{n \text{gr}(\alpha, \gamma_1) + n - 1} \\
s(n, \alpha, \vec{\gamma}, \beta) &= s(n-1, \alpha, \vec{\gamma}_{<n}, \gamma_n) \cdot (-1)^{\text{gr}(\alpha, \gamma_n) + n - 1}.
\end{aligned}$$

The first is just the recurrence stated above, and the last is essentially the same. The middle term in addition uses that $\text{gr}(\alpha, \gamma_1) + \text{gr}(\gamma_1, \eta) = \text{gr}(\alpha, \eta)$.

Thus $\text{Id} + \sum_{n>0} H_n : \widetilde{CI}_0 \rightarrow \widetilde{CI}_1$ is an $SO(3)$ -equivariant quasi-isomorphism. The extension to $\widetilde{CI}^{SU(2)}$ is obvious: replace H with fiber products with the lifted moduli spaces $\overline{\mathcal{M}}'$. \blacksquare

To be clear, there are no transversality conditions imposed on the above chain complexes. One should choose the basepoints (b_α, q_α) , and hence the map h , only *after* having specified the endpoint maps from the moduli spaces. Now what we need is the following.

Lemma 7.5. *We may find a coherent collection of $SU(2)$ -equivariant homotopies $e_\pm^t : \overline{\mathcal{M}}_{\alpha\beta} \rightarrow \alpha \times \beta$, and basepoints b_α, q_α , so that if d_M is the map corresponding to fiber products with moduli spaces using the time-1 maps e_\pm^1 , and p, h are the maps in the deformation retract data specified by the b_α and q_α , the map $pd_M h d_M$ is identically zero.*

Proof. If we write

$$X_{\alpha\beta} = (\alpha \times \beta) / SU(2),$$

there is an induced map $\mathcal{M}_{\alpha\beta} \rightarrow X_{\alpha\beta}$; we will soon demand that this map is constant as much as we can force it to be. If one of the orbits (say, α) is irreducible, then $X_{\alpha\beta} \cong \beta$; if one of the orbits is trivial, $X_{\alpha\beta}$ is a point; if both α and β are $U(1)$ -reducibles, $X_{\alpha\beta} \cong [-2, 2]$.

We induct on two different integers: the dimension i of $\mathcal{M}_{\alpha\beta}$, and the maximal codimension j of a stratum of $\mathcal{M}_{\alpha\beta}$; of course $i \geq j$.

We will start by explaining the argument in the case of integer homology spheres; the $U(1)$ -reducibles present substantial new issues that must be addressed, but the idea is largely the same.

Let $\overline{\mathcal{M}}_{\alpha\beta}$ be a moduli space without any boundary strata, and suppose both α and β are irreducible, with $\dim \mathcal{M}_{\alpha\beta} < 3$. Using the trivialization $\alpha \cong SU(2)$ given by the basepoint b_α , the left endpoint map to α gives a trivialization

$$\overline{\mathcal{M}}_{\alpha\beta} \cong \mathcal{M}_{\alpha\beta} \times SU(2),$$

as well as a canonical isomorphism

$$\alpha \times \beta \cong X_{\alpha\beta} \times SU(2).$$

Therefore, an equivariant homotopy of the endpoint map is the same as choosing a homotopy of the map

$$\mathcal{M}_{\alpha\beta} \rightarrow X_{\alpha\beta} \cong S^3.$$

Because S^3 is 2-connected, any such map is null-homotopic; choose a homotopy taking this map to the constant map at $[b_\alpha, b_\beta]$. Do nothing to the higher-dimensional moduli spaces, and do nothing to the moduli spaces for which one of α or β is fully reducible.

We now consider moduli spaces having at most a codimension 1 boundary stratum. Again, we first consider the case where both α and β are irreducible. If one of these strata corresponds to a factorization through the full reducible, then $\dim \mathcal{M}_{\alpha\beta} > 3$, and we will not worry about it yet.

Suppose $\dim \mathcal{M}_{\alpha\beta} < 3$. We have already chosen a specified null-homotopy $\partial \mathcal{M}_{\alpha\beta} \rightarrow X_{\alpha\beta}$. Again, because S^3 is 2-connected and the final stage of the homotopy on the boundary is constant at $[b_\alpha, b_\beta]$, we may extend this to a null-homotopy of the entire moduli space to $[b_\alpha, b_\beta]$.

In every other case (that is, either $\dim \mathcal{M}_{\alpha\beta} \geq 3$ or one of α or β is fully reducible), we take the existing homotopy of

$$\partial \overline{\mathcal{M}}_{\alpha\beta} \rightarrow \alpha \times \beta$$

and use the equivariant homotopy extension property to extend it arbitrarily to the whole of $\overline{\mathcal{M}}_{\alpha\beta}$.

Do the same for those moduli spaces of dimension 2 with codimension 2 strata (between irreducibles) and extend arbitrarily for higher-dimensional moduli spaces. Now inductively extend the existing homotopies on all moduli spaces with strata of the same maximal codimension j , for $j \geq 3$.

Having done this, we choose basepoints q_α in each orbit so that the map $\overline{\mathcal{M}}_{\alpha\tilde{\gamma}\beta} \rightarrow \alpha \times \beta$ is transverse to the orbit through (b_α, q_β) for all $\alpha, \tilde{\gamma}, \beta$, as in the setup before Lemma 7.4. Choose the homotopy h arbitrarily so that $h([b_\alpha]) = 0$ and $h([\alpha]) = 0$; its existence is guaranteed as i and p are quasi-isomorphisms between free R -modules whose composite is the identity on homology. We claim that $hd_M i$ is either zero, or of dimension larger than the orbit it lies in. This follows by a case-by-case analysis. The fiber product

$$b_\alpha \times_\alpha \overline{\mathcal{M}}_{\alpha\beta} \rightarrow \beta$$

is in one of the following five cases:

- The orbits α and β are both irreducible and $\dim \mathcal{M}_{\alpha\beta} = 0$; in this case, the fiber product is simply some number of copies of $[b_\beta]$, and applying h returns zero.
- The orbits α and β are both irreducible, while $0 < \dim \mathcal{M}_{\alpha\beta} < 3$, the fiber product is a constant map to b_α and hence a degenerate chain (its boundary is either also small image, or its boundary is a sum of points whose signed count must be zero), and hence zero at the chain level.
- The orbits α and β are both irreducible, while $\dim \mathcal{M}_{\alpha\beta} \geq 3$; taking the fiber product and then applying h leaves a chain of dimension larger than β .
- The orbit α is fully reducible; in this case $b_\alpha \times_\alpha \overline{\mathcal{M}}_{\alpha\beta} = \overline{\mathcal{M}}_{\alpha\beta}$, and hence has dimension at least as large as β . Applying h gives a chain of dimension larger than β .
- The orbit β is fully reducible; applying h to $d_M i$ gives something of dimension at least 1, which is then automatically degenerate.

Because the projection map p kills any chain of dimension larger than the orbit it lives on, we see that $p(d_M h)^n d_M i = 0$ for $n > 0$, as desired. What remains of the differential simply counts the preimages of $\mathcal{M}_{\alpha\beta} \rightarrow X_{\alpha\beta}$ above $[b_\alpha, q_\beta]$.

We now discuss the changes necessary in the presence of $U(1)$ -reducibles. The argument — crush maps down to constant maps, unless the dimensions are too large to ensure this is possible — largely applies, until the case of moduli spaces of

dimension 2 between irreducibles α, β , with a codimension 1 stratum corresponding to factorization through a $U(1)$ -reducible γ . Then, applying the existing homotopies on the boundary, we find that the stratum corresponding to that factorization is (a disjoint union of signed copies of) $U(1)$, and the map from these $U(1)$ s to β is the embedding of $U(1)$ onto the $U(1)$ -orbit through b_β . This is forced on us by equivariance. We will call these moduli spaces IUI type to indicate that they are between two irreducibles, with a stratum corresponding to a factorization through a $U(1)$ -reducible.

It is worth remarking that in the case of 1-dimensional moduli spaces between two $U(1)$ -reducibles, we need slightly more argument to see that we may homotope the endpoint maps down to the basepoint. To do this, we use that the $SU(2)$ action on $\overline{\mathcal{M}}$ is free, and hence there is an isomorphism $\overline{\mathcal{M}} \cong \mathcal{M} \times SU(2)$. Because \mathcal{M} is 1-dimensional, it is a collection of circles and arcs, and the $SU(2)$ -equivariant null-homotopy arises from the fact that $\alpha \times \beta \cong S^2 \times S^2$ is simply connected.

To get around the more serious issue with 2-dimensional moduli spaces, we use h to take something that was not previously of small image, and make it small image.

To be precise, for each irreducible orbit α , choose a disc bounding $b_\alpha \cdot U(1)$. Call this disc D_α . In the above situation, what we *can* do is homotope this moduli space $\mathcal{M}_{\alpha\beta}$, relative to the boundary, so that its image lies inside the disc D_α . (Here again we use that S^2 is 2-connected.) Once we have done this for all such moduli spaces, we continue the usual pattern of extending arbitrarily to everything whose dimension is 3 or larger. We will continue to use $\overline{\mathcal{M}}_{\alpha\beta}$ to denote the moduli spaces equipped with this new family of maps.

We will set

$$h(b_\alpha \cdot U(1)) = D_\alpha.$$

For convenience, following the isomorphism $\alpha \cong SU(2) \cong S^3$ determined by the basepoint b_α , we demand that this is the same disc for all irreducible orbits α .

For each moduli space $\overline{\mathcal{M}}_{\alpha\beta}$ of IUI type, we choose a chain $\sigma_{\alpha\beta}$ with

$$\partial\sigma_{\alpha\beta} = (b_\alpha \times_\alpha \overline{\mathcal{M}}_{\alpha\beta}) \cup nD_\alpha,$$

where n is the appropriate integer; that is, we choose a chain which cobounds the map from the moduli space and the appropriate number of discs (so that their sum is null-homologous). Furthermore, we demand that $\sigma_{\alpha\beta}$ factor through the disc D_α . That we may do all this follows because $H_2(D^2) = 0$. We will set

$$h(b_\alpha \times_\alpha \overline{\mathcal{M}}_{\alpha\beta}) = \sigma_{\alpha\beta}.$$

Lastly, we will let $h(D_\alpha) = 0$.

Next, we choose the basepoints q_α so that the following conditions hold. First, if ${}^b\overline{\mathcal{M}}_{\alpha\bar{\gamma}\beta}$ is the fiber product

$$b_\alpha \times_\alpha \overline{\mathcal{M}}_{\alpha\gamma_1} \times_{\gamma_1} \cdots \times_{\gamma_n} \overline{\mathcal{M}}_{\gamma_n\beta},$$

thought of as a chain in β , we demand that q_β is a regular value of ${}^b\overline{\mathcal{M}}_{\alpha\bar{\gamma}\beta}$.

Similarly, let ${}^D\overline{\mathcal{M}}_{\alpha\bar{\gamma}\beta}$ denote the fiber product with D_α , where α is an irreducible orbit, and let ${}^\sigma\overline{\mathcal{M}}_{\alpha\bar{\gamma}\beta}$ the fiber product with $\sigma_{\xi\alpha}$ as ξ varies. We further demand that q_β is a regular value of these chains in β .

Now set \mathcal{F} to consist of the chains in α given by

$$\alpha \leftarrow ({}^b\overline{\mathcal{M}}_{\alpha\bar{\gamma}\beta} \times_\beta q_\beta).$$

Then $\widetilde{CI}^{SU(2),\mathcal{F}}$ is a quasi-isomorphic, $\Lambda(u)$ -invariant subcomplex of $\widetilde{CI}^{SU(2)}$. The first condition above guarantees that i has image inside this subcomplex (that is, b_α is transverse to all of the chains in \mathcal{F}); similarly, D_α and σ_{ξ_α} are all transverse to \mathcal{F} , and hence our choice of $h(b_\alpha \cdot U(1))$ above is valid in this subcomplex, as is our choice of $h({}^b\overline{\mathcal{M}}_{\alpha\beta})$.

Additionally, we demand the conditions $h([b_\alpha]) = 0$ and $h([\alpha]) = 0$. Beyond these we extend h arbitrarily as a chain homotopy on the entire complex between ip and 1.

What remains is a case-by-case analysis that $p(d_M h)^n d_M i = 0$ for $n > 0$. Most cases are similar (if not identical) to those discussed above in the integer homology sphere case: automatically zero because the moduli space itself is a degenerate chain, or automatically zero because the dimension counts are wrong. Here we enumerate the cases that are fundamentally new.

- Consider the component of $pd_M hd_M i$ arising by first taking the fiber product with a moduli space $\overline{\mathcal{M}}_{\alpha\beta}$ of type IOI as above, applying h , and then taking the fiber product with a 0-dimensional moduli space between irreducibles β, γ . The output of $hd_M i$ here is the chain $\sigma_{\alpha\beta} \rightarrow \beta$, which is of small image; applying the appropriate component of d_M to this, we again obtain a chain of small image. Projecting, we get zero: the intersection number of any chain of small image with a regular value is always zero.
- Consider the component of $pd_M hd_M i$ arising by first taking the fiber product with a 0-dimensional moduli space $\alpha \rightarrow \beta$, where α is $U(1)$ -reducible and β is irreducible. The output of $hd_M i$ is some number of copies of the disc D_β . Suppose the moduli space corresponding to the second factor of d_M is a 1-dimensional moduli space $\beta \rightarrow \gamma$, both orbits irreducible. Then the corresponding fiber product is degenerate: it is of small image (its image is the image of the disc D_β), so the corresponding projection is zero.
- Consider the component $pd_M hd_M hd_M i$ arising nearly as above: the first d_M contributes a 0-dimensional moduli space $\alpha \rightarrow \beta$, the next factor contributes a 0-dimensional moduli space $\beta \rightarrow \gamma$, and the last contributes a 0-dimensional moduli space $\gamma \rightarrow \delta$, where all three of β, γ, δ are irreducible, while α is $U(1)$ -reducible. Then as above, $hd_M i$ outputs some number of copies of the disc D_β ; taking another fiber product with a 0-dimensional moduli space outputs D_γ . Applying h to this gives zero.
- Consider the component $pd_M hd_M i$ arising from a 0-dimensional moduli space $\alpha \rightarrow \beta$, and a 0-dimensional moduli space $\beta \rightarrow \gamma$, where α, γ are both $U(1)$ -reducible and β is irreducible. Because $hd_M i$ is some number of copies of the disc D_β , and the projection $SU(2) \rightarrow S^2$ sends the disc D_β to a class which generates the second homology of S^2 , we can *not* rule these out in the same way as above. However, this would imply that $\text{gr}(\alpha, \gamma) = 3$, which contradicts Proposition 6.8 (4), stating that the grading of any two reducibles is the same parity.

Thus, the basepoints q , the endpoint maps of the moduli spaces, and the homotopy h have been chosen as desired. ■

Take the complex $\widetilde{CI}_1(Y, E, \pi; R)$ equipped with the differential $d + d_M$ induced by the above fiber products with moduli spaces; we described the construction of

deformation retract data

$$(DCI, 0) \xrightleftharpoons[p]{i} (\widetilde{CI}_1, d + d_M) \circlearrowright h$$

above. Applying the homological perturbation lemma, and using the fact that

$$pd_M(hd_M)^n i = 0$$

for $n > 0$, we arrive at a chain complex $(DCI, pd_M i)$ with $\Lambda(u)$ -action, and an equivariant homotopy equivalence to $(\widetilde{CI}_1, d + d_M)$. We may finally describe the differential $pd_M i$.

Denote by C_*^{irr} the free R -module on the irreducible orbits $\alpha \subset \mathfrak{E}_\pi$ (graded the same), $C_*^{U(1)}$ the free R -module on $U(1)$ -reducible orbits, and C_*^θ the free R -module on the full reducibles. Set

$$DCI(Y, E, \pi) = C_*^{\text{irr}} \oplus C_*^{\text{irr}}[3] \oplus C_*^{U(1)} \oplus C_*^{U(1)}[2] \oplus C_*^\theta.$$

As graded R -modules, $DCI(Y, E, \pi) = \bigoplus H_{*-i(\alpha)}(\alpha)$.

The differential $pd_M i$ that DCI inherits is given by the matrix

$$\widetilde{\partial}_{CI} := \begin{pmatrix} \partial_1 & 0 & 0 & 0 & 0 \\ U_{\text{F1}} & -\partial_1 & V_4 & V_2 & D_2 \\ V_1 & 0 & 0 & 0 & 0 \\ V_3 & 0 & 0 & 0 & 0 \\ D_1 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

(We will momentarily define the terms in this matrix.) Here we write "dim σ " to mean 0 for the component C_*^{irr} and 3 for the component $C_*^{\text{irr}}[3]$.

The action of u on \widetilde{CI} is defined by the matrix whose only nonzero term is $\text{id} : C_*^{\text{irr}} \rightarrow C_*^{\text{irr}}[3]$. It is easy to see that $u\widetilde{\partial}_{CI} = \widetilde{\partial}_{CI}u = \partial_1$, considered as a map $C_*^{\text{irr}} \rightarrow C_{*-1}^{\text{irr}}[3]$.

As before, we denote

$$\mathcal{M}_{\alpha\beta} = \overline{\mathcal{M}}_{\alpha\beta}/SU(2);$$

these are the spaces of unframed flowlines between the underlying connections $[\alpha], [\beta]$, having forgotten the framings. Also recall we defined $X_{\alpha\beta} = (\alpha \times \beta)/SU(2)$ above. The operators in this matrix all arise of the form $T(\alpha) = \sum_\beta n(\alpha, \beta)\beta$, where the sum is taken over β in the right degree, and $n(\alpha, \beta)$ arises either by counting points (with orientation) in $\mathcal{M}_{\alpha\beta}$, or degree of the map $\mathcal{M}_{\alpha\beta} \rightarrow X_{\alpha\beta}$, as measured by the number of points lying above $[b_\alpha, q_\beta]$. We've chosen b_α, q_β to ensure that this is a regular value (and in the case that both α and β are $U(1)$ -reducibles, to have different stabilizer).

- ∂_1 counts points in $\mathcal{M}_{\alpha\beta}$, when α, β are both irreducible and $\text{gr}(\alpha, \beta) = 1$.
- D_1 counts points in $\mathcal{M}_{\alpha\beta}$ when α is irreducible, β is fully reducible, and $\text{gr}(\alpha, \beta) = 1$.
- D_2 counts points in $\mathcal{M}_{\alpha\beta}$ when α is fully reducible, β is irreducible, and $\text{gr}(\alpha, \beta) = 4$.¹⁴

¹⁴Recall that this means that $\overline{\mathcal{M}}$ is of dimension $\dim \alpha + \text{gr}(\alpha, \beta) - 1$; because α is fully reducible, this means that $\overline{\mathcal{M}}$ is 3-dimensional. Because the endpoint map to β is equivariant, and β is irreducible, $\overline{\mathcal{M}}$ is a finite set of orbits with no stabilizer, and \mathcal{M} is 0-dimensional.

- U_{F1} counts the number of points in the fiber above $[b_\alpha, q_\beta]$ of the map $\mathcal{M}_{\alpha\beta} \rightarrow X_{\alpha\beta}$, when α and β are both irreducible and $\text{gr}(\alpha, \beta) = 3$.
- V_1 counts points in $\mathcal{M}_{\alpha\beta}$ when α is irreducible, β is $U(1)$ -reducible, and $\text{gr}(\alpha, \beta) = 1$.
- V_2 counts points in $\mathcal{M}_{\alpha\beta}$ when α is $U(1)$ -reducible, β is irreducible, and $\text{gr}(\alpha, \beta) = 2$.
- V_3 counts points above $[b_\alpha, q_\beta]$ when α is irreducible, β is $U(1)$ -reducible, and $\text{gr}(\alpha, \beta) = 3$.
- V_4 counts points above $[b_\alpha, q_\beta]$ when α is $U(1)$ -reducible, β is irreducible, and $\text{gr}(\alpha, \beta) = 4$.

The derivation of this matrix is relatively self-explanatory; the sign arises from the factor $(-1)^{\dim \sigma}$ in d_M . The one small subtlety is why there is no matrix element corresponding to counting points above $[b_\alpha, q_\beta]$ when both α, β are $U(1)$ -reducible and $\text{gr}(\alpha, \beta) = 3$. By Proposition 6.8 (4), this is impossible: the relative grading of reducible connections is always even!

We summarize the result of the homological perturbation lemma discussion as follows.

Corollary 7.6. *The R -module with differential $(DCI, pd_M i)$ described above is a dg-module over $\Lambda(u)$, and is $\Lambda(u)$ -equivariantly homotopy equivalent to \widetilde{CI} .*

Remark 7.1. What we call DCI was described in [Don02, Section 7.3.3] for integer homology spheres, with the notation \widetilde{CF} . Donaldson’s complex includes the data of the filtration¹⁵ by degree of an element in $H_*\alpha$, as well as the action of $u \in \Lambda(u)$. In Section 7.3.2, Donaldson introduces a method for reducing the contribution of small-dimensional moduli spaces, so that this complex makes sense; his homotopy is an example of the many possible homotopies we could have used above. Thus, in particular, \widetilde{CI} is equivariantly homotopy equivalent to Donaldson’s \widetilde{CF} .

In what follows we use the complex DCI to produce small models which we write as \overline{DCI}^\pm for CI^+ and CI^- . In the case of integer homology spheres, these are essentially the same as what Donaldson writes \overline{CF} and \underline{CF} .

As described in the appendix, the *reduced bar construction* of an augmented dg-algebra A and a right dg-module M is given, as a graded R -module, as

$$B(M, A, R) := \bigoplus_{n=0}^{\infty} M \otimes \overline{A}[1]^n,$$

where $\overline{A} = \ker \varepsilon$ and ε is the augmentation. For us, $A = \Lambda(u)$, and \overline{A} is just a copy of R concentrated in degree 3. Therefore, we may write

$$B(M, A, R) = \bigoplus_{n=0}^{\infty} M[4n].$$

We write this as $M \otimes_R R[U^*]$, where $|U^*| = 4$.¹⁶ The differential d^+ is seen to be

$$d^+(m \otimes (U^*)^n) = um \otimes (U^*)^{n-1} + (-1)^n m \otimes (U^*)^n,$$

¹⁵Thinking of DCI as a $\mathbb{Z}/8 \times \mathbb{Z}$ -graded multicomplex with $d_0 = 0$, this is the filtration corresponding to the \mathbb{Z} factor, while ours is the periodic filtration corresponding to the $\mathbb{Z}/8$ factor.

¹⁶We choose the notation U^* here because U is reserved for the degree -4 operation coming from $H^{-4}(BSO(3); R)$.

where $(U^*)^{-1}$ should be interpreted as zero. Write U_{alg}^+ for the operator with

$$U_{\text{alg}}^+(m \otimes (U^*)^n) = (-1)^n um \otimes (U^*)^{n-1};$$

again, this means that $U_{\text{alg}}^+(1) = 0$.

Now, when we apply this to a $\mathbb{Z}/8$ -graded module M (with finite basis of finite-dimensional $\Lambda(u)$ -modules degreewise isomorphic to R) as in Section A.8, we need to *complete*; in particular,

$$CI^+ = \prod_{n \geq 0} \widetilde{CI}[4n] \cong \widetilde{CI} \widehat{\otimes}_R R[U^*] = \widetilde{CI} \otimes_R R[U^*].$$

The action of $R[U] \subset C_{\Lambda(u)}^-$ is such that $U \cdot (U^*)^k = (U^*)^{k-1}$; we must restrict to this subalgebra as the action of the entire algebra does not extend to the completion. We denote $DCI^+ = C_{\Lambda(u)}^+(DCI)$, equipped with its action of $R[U^*]$.

By convention, if A is an operator from one summand of DCI to another, then we use A to also refer to the corresponding map between summands of DCI^+ with

$$A(x \otimes (U^*)^n) = (-1)^n Ax \otimes (U^*)^n.$$

Then the differential of

$$DCI^+ \cong \left(C_*^{\text{irr}} \oplus C_*^{\text{irr}}[3] \oplus C_*^{U(1)} \oplus C_*^{U(1)}[2] \oplus C_*^\theta \right) \widehat{\otimes}_R R[U^*]$$

is given by the following matrix.

$$\partial_{DCI}^+ := \begin{pmatrix} \partial_1 & 0 & 0 & 0 & 0 \\ U_{\text{alg}}^+ + U_{\text{Fl}} & -\partial_1 & V_4 & V_2 & D_2 \\ V_1 & 0 & 0 & 0 & 0 \\ V_3 & 0 & 0 & 0 & 0 \\ D_1 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

Now we apply the following analogue to [SS10, Lemma 5] to our very similar situation. First, observe that the map $U_{\text{Fl}} + U_{\text{alg}}$ defines an isomorphism

$$U^* C_*^{\text{irr}}[U^*] \rightarrow C_*^{\text{irr}}[3][U^*];$$

here we have used the completeness in U^* . The inverse is explicitly given by

$$\sum_{n \geq 0} c_n (U^*)^n \mapsto \sum_{n \geq 0} \left(\sum_{j=0}^n (-1)^{j+1} U_{\text{Fl}}^{n-j} c_j \right) (U^*)^{n+1}.$$

We write P^+ for this inverse.

Consider the projection map $\epsilon_1 : DCI^+ \rightarrow C_*^{\text{irr}}[U^*][3]$; we let

$$\epsilon = (\epsilon_1, \epsilon_1 \partial^+) : DCI^+ \rightarrow C_*^{\text{irr}}[U^*][3] \oplus C_*^{\text{irr}}[U^*][2].$$

The codomain is given the differential $\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$, and ϵ is a chain map; in fact, it is

surjective, as (x, y) is in the image of $(P^+(y - \partial_1 x), x) \in DCI^+$. As the codomain is acyclic, the inclusion $\ker(\epsilon) \rightarrow DCI^+$ is a quasi-isomorphism.

Write

$$\overline{DCI}^+ = C_*^{\text{irr}} \oplus \left(C_*^{U(1)} \oplus C_*^{U(1)}[2] \oplus C_*^\theta \right) \otimes_R R[U^*]$$

as a graded R -module. There is also a projection $\pi : DCI^+ \rightarrow \overline{DCI}^+$, and the composite $\pi i : \ker \epsilon \rightarrow \overline{DCI}^+$ is an isomorphism of R -modules. We wish to describe the differential (and U -action) on \overline{DCI}^+ , pulling back via this isomorphism. The differential is written $(\pi i)\partial^+(\pi i)^{-1}$, and similarly the U -map (whose degree is -4) is $(\pi i)U(\pi i)^{-1}$.

Explicitly, $(\pi i)^{-1}x$ is the unique element $y \in DCI^+$ so that neither y nor ∂^+y have components in $C_*^{\text{irr}}[[U^*]][3]$ with $\pi(y) = x$. If

$$x = (x_0, x_1, x_2, x_3) \in C_*^{\text{irr}} \oplus \left(C_*^{U(1)} \oplus C_*^{U(1)}[2] \oplus C_*^\theta \right) \otimes_R R[[U^*]],$$

then

$$(i\pi)^{-1}(x) = (x_0 + P^+(U_{\text{F1}}x_0 + V_4x_1 + V_2x_2 + D_2x_3), 0, x_1, x_2, x_3) \in \left(C_*^{\text{irr}} \oplus C_*^{\text{irr}}[3] \oplus C_*^{U(1)} \oplus C_*^{U(1)}[2] \oplus C_*^\theta \right) R[[U^*]].$$

Applying ∂^+ , projecting onto \overline{DCI}^+ , and using that $\partial_1 P^+x$ is in $\ker(\epsilon)$, we find that the desired matrix is

$$\bar{\partial}_{DCI}^+ := \begin{pmatrix} \partial_1 & 0 & 0 & 0 \\ V_1 + V_1 P^+ U_{\text{F1}} & V_1 P^+ V_4 & V_1 P^+ V_2 & V_1 P^+ D_2 \\ V_3 + V_3 P^+ U_{\text{F1}} & V_3 P^+ V_4 & V_3 P^+ V_2 & V_3 P^+ D_2 \\ D_1 + D_1 P^+ U_{\text{F1}} & D_1 P^+ V_4 & D_1 P^+ V_2 & D_1 P^+ D_2 \end{pmatrix}$$

A remarkable number of these terms are zero for degree reasons. Recall from Proposition 4.33 that the relative grading of any pair of reducible orbits is even. Now, D_i and V_i have degree of the same parity as i , and all terms in the infinite sum P^+x have the same grading as x modulo 4; therefore the entries not in the left column are all of odd degree, and hence identically zero! Therefore, we have

$$\bar{\partial}_{DCI}^+ = \begin{pmatrix} \partial_1 & 0 & 0 & 0 \\ V_1 + V_1 P^+ U_{\text{F1}} & 0 & 0 & 0 \\ V_3 + V_3 P^+ U_{\text{F1}} & 0 & 0 & 0 \\ D_1 + D_1 P^+ U_{\text{F1}} & 0 & 0 & 0 \end{pmatrix}$$

Similarly, the action of $U \in H^{-4}(BSO(3); R)$ was previously contraction against U^* ; using $U \frown$ to denote contraction by U , the action \bar{U} on \overline{DCI}^+ is given by

$$\begin{pmatrix} -U_{\text{F1}} & -V_4 & -V_2 & -D_2 \\ 0 & U \frown & 0 & 0 \\ 0 & 0 & U \frown & 0 \\ 0 & 0 & 0 & U \frown \end{pmatrix}.$$

In the case that Y is an integer homology sphere and $\mathbb{Q} \subset R$, up to a change of basis (just a scaling of each coordinate) the chain complex $\overline{DCI}^+(Y; R)$ is identical as a U -module to Donaldson's $\overline{CF}(Y; R)$.

Now we may write

$$DCI^- = \left(C_*^{\text{irr}} \oplus C_*^{\text{irr}}[3] \oplus C_*^{U(1)} \oplus C_*^{U(1)}[2] \oplus C_*^\theta \right) [U],$$

with differential

$$\partial_{DCI}^- := \begin{pmatrix} \partial_1 & 0 & 0 & 0 & 0 \\ U_{\text{alg}}^- + U_{\text{Fl}} & -\partial_1 & V_4 & V_2 & D_2 \\ V_1 & 0 & 0 & 0 & 0 \\ V_3 & 0 & 0 & 0 & 0 \\ D_1 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

Unlike in the case of DCI^+ , the operators ∂_1, V_i, D_i do not pick up an extra sign when acting on $C[U]$. For $x \in C_*^{\text{irr}}$, we define the operator

$$U_{\text{alg}}^-(x \otimes U^k) = (-1)^{|x|+k} x \otimes U^{k+1} \in C_*^{\text{irr}}[3][U];$$

this is the only sign in the above differential hidden in the notation.

Now

$$U_{\text{Fl}} + U_{\text{alg}}^- : C_*^{\text{irr}}[U] \rightarrow (C_*^{\text{irr}}[3])[U]/(1 \cdot C_*^{\text{irr}}[3])$$

is an isomorphism; it is crucial here that we are working with a polynomial ring and not a power series ring. The inverse, P^- , is given on a basis element by

$$P^-(x \otimes U^{n+1}) = \sum_{i=0}^n (-1)^{(i+1)(|x|+n)+i(i+1)/2} U_{\text{Fl}}^i x \otimes U^{n-i}.$$

Now we follow [SS10] more closely, and instead of taking the kernel of a map to an acyclic complex, we quotient by an acyclic subcomplex. The subcomplex Z is spanned by $C_*^{\text{irr}}[U]$ and its image under ∂_{DCI}^- , which is an injective map. The quotient DCI^-/Z remains a U -module. If we set

$$\overline{DCI}^- = C_*^{\text{irr}}[3] \oplus \left(C_*^{U(1)} \oplus C_*^{U(1)}[2] \oplus C_*^\theta \right) \otimes_R R[U]$$

we have a natural inclusion map $i : \overline{DCI}^- \hookrightarrow DCI^-$ so that the composite

$$\pi i : \overline{DCI}^- \rightarrow DCI^-/Z$$

is an isomorphism. As above, we may compute the induced differential as

$$\overline{\partial}_{DCI}^- = \begin{pmatrix} -\partial_1 & V_4 - U_{\text{Fl}}(P^-V_4)_{U=0} & V_2 - U_{\text{Fl}}(P^-V_2)_{U=0} & D_2 - U_{\text{Fl}}(P^-D_2)_{U=0} \\ 0 & -V_1P^-V_4 & -V_1P^-V_2 & -V_1P^-D_2 \\ 0 & -V_3P^-V_4 & -V_3P^-V_2 & -V_3P^-D_2 \\ 0 & -D_1P^-V_4 & -D_1P^-V_2 & -D_1P^-D_2 \end{pmatrix}.$$

Here the term V_4 in the topmost row, and similarly with V_2 and D_2 in that same row, denotes the composite

$$C_*^{U(1)}[U] \xrightarrow{U=0} C_*^{U(1)} \xrightarrow{V_4} C_*^{\text{irr}}[3];$$

in the bottom right 3×3 block matrix, there is no projection component to the terms V_4, V_2 , or D_2 . When we write for instance $(P^-V_4)_{U=0}$, we mean the composite

$$C_*^{U(1)}[U] \xrightarrow{P^-V_4} C_*^{\text{irr}}[3][U] \xrightarrow{U=0} C_*^{\text{irr}}[3]:$$

that is, apply P^-V_4 then project to the constant term.

Again, most of these terms are zero for degree reasons, and we may in fact write

$$\bar{\partial}_{DCI}^- = \begin{pmatrix} -\partial_1 & V_4 - U_{\text{Fl}}(P^-V_4)_{U=0} & V_2 - U_{\text{Fl}}(P^-V_2)_{U=0} & D_2 - U_{\text{Fl}}(P^-D_2)_{U=0} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

Finally, the action of U is given on \overline{DCI}^- by

$$\begin{pmatrix} (-1)^{|x|+1}U_{\text{Fl}} & 0 & 0 & 0 \\ (-1)^{|x|+1}V_1 & U \smile & 0 & 0 \\ (-1)^{|x|+1}V_3 & 0 & U \smile & 0 \\ (-1)^{|x|+1}D_1 & 0 & 0 & U \smile \end{pmatrix}.$$

Again, when Y is an integer homology sphere and $\mathbb{Q} \subset R$, the chain complex $\overline{DCI}^+(Y; R)$ is identical as a U -module to Donaldson's $\underline{CF}(Y; R)$ up to scaling of basis.

When (Y, E) is equipped with an admissible bundle, there are no reducible connections; in this case, the equivalence of $CI^+(Y, E; R)$ with $\overline{DCI}^+(Y, E; R)$ as a $R[U]$ -module (and the same for the minus flavor) immediately gives is the following.

Corollary 7.7. *When $\frac{1}{2} \in R$, the equivalence $I^\pm(Y, E) \cong I(Y, E)$ for an admissible bundle (Y, E) takes the U -map to the U -map up to sign.*

7.3. Instanton Tate homology. We have the following theorems for Tate homology. Fix a ground ring R with $\frac{1}{2} \in R$; then we have a canonical isomorphism $H_{SO(3)}^-(R) \cong R[U]$ and a quasi-isomorphism $C_*(SO(3); R) \simeq \Lambda(u) := \Lambda$, where $\Lambda(u)$ denotes the exterior algebra on a degree 3 generator u .

Proposition 7.8. *The action of U^* is an isomorphism on $I^\infty(Y, E; R)$. In fact, the map $I^-(Y, E; R) \rightarrow I^\infty(Y, E; R)$ may be identified with the localization*

$$I^-(Y, E; R) \rightarrow I^-(Y, E; R) \otimes_{R[U]} R[U, U^{-1}];$$

even more explicitly, we may identify the map $CI^-(Y, E; R) \rightarrow CI^\infty(Y, E; R)$ on the chain level using the reduced Donaldson model as

$$\overline{DCI}^-(Y, E; R) \rightarrow \overline{DCI}^-(Y, E; R) \otimes_{R[U]} R[U, U^{-1}].$$

Proof. We follow much the same lines as in Proposition A.18 in the \mathbb{Z} -graded case, and work with the Donaldson model \overline{DCI}^- . First, to see that the action of U on I^∞ is invertible, we look at the spectral sequence corresponding to the index filtration of CI^∞ . The E^1 page is identified with a direct sum of copies of

$$H_\Lambda^\infty(\Lambda), H_\Lambda^\infty(R \oplus R[2]), \text{ and } H_{SO(3)}^\infty(R);$$

the first corresponds to irreducible orbits, the second corresponds to $SO(2)$ -reducibles and is just a sum of two copies of $H_\Lambda^\infty(R)$, and the last corresponds to full reducibles. In all cases, the action of U is an isomorphism on each factor: that the first group is zero is one of the axioms of Tate homology, and that the second two have invertible U action is precisely the calculation of Lemma A.17. Because the index filtration

is complete, a map which is an isomorphism on the E^1 page is an isomorphism on homology, and so the action of U is invertible in $I^\infty(Y, E; R)$.

Given a map $f : M \rightarrow M$ of degree k of a dg-module, there is a definition given before Proposition A.18 of a chain complex $M[f^{-1}]$; it is the mapping cone of $M[t] \xrightarrow{1-tf} M[t]$, where $|t| = -k$. This is defined so that

$$H(M[f^{-1}]) = H(M)[f^{-1}],$$

the latter notation meaning the usual module-theoretic sense of inverting an element or map.

If one applies this to define $CI^\infty(Y, E; R)[U^{-1}]$, one finds that the index filtration is *no longer complete*, even though the index filtration is complete on CI^∞ . One may then pass to the completion of $CI^\infty(Y, E; R)[U^{-1}]$; this is the filtered complex that the spectral sequence actually computes the homology of (and the E^1 pages of the spectral sequence for a filtered complex and its completion are the same). Precisely, if $F_s M$ is a filtered complex, the completion is

$$\widehat{M} = \lim_{p \rightarrow \infty} M/F_{-p}M.$$

We write this completion as $CI^\infty(Y, E; R)[\widehat{U^{-1}}]$. There is a natural map

$$CI^\infty(Y, E; R) \rightarrow CI^\infty(Y, E; R)[\widehat{U^{-1}}]$$

which we know to be an isomorphism on the E^1 page, because the action of U^* is an isomorphism on the E^1 page of $CI^\infty(Y, E; R)$.

Because completion is natural, we have maps

$$CI^-(Y, E; R)[\widehat{U^{-1}}] \rightarrow CI^\infty(Y, E; R)[\widehat{U^{-1}}];$$

again, the computations of Lemma A.17 imply that this is an isomorphism on the E^1 page, and because these filtrations are complete, the same is true at the level of homology.

To conclude we need to find a quasi-isomorphism

$$CI^-(Y, E; R)[\widehat{U^{-1}}] \simeq \overline{DCI}^-(Y, E; R) \otimes_{R[U]} R[U, U^{-1}].$$

This follows again via naturality of completion: there is a canonical map

$$\overline{DCI}^-(Y, E; R)[U^{-1}] \rightarrow \overline{DCI}^-(Y, E; R) \otimes_{R[U]} R[U, U^{-1}]$$

of filtered complexes for which the filtration is *not* complete (remember that the first term is defined to be the mapping cone of a certain map on $DCI^-[t]$). The same argument as before shows this map is an isomorphism on the E^1 page, so the corresponding map on completions

$$\overline{DCI}^-(Y, E; R)[\widehat{U^{-1}}] \rightarrow \overline{DCI}^-(Y, E; R) \widehat{\otimes}_{R[U]} R[U, U^{-1}]$$

is a quasi-isomorphism. We conclude by observing that we may explicitly identify the completion as

$$\overline{DCI}^-(Y, E; R) \widehat{\otimes}_{R[U]} R[U, U^{-1}] = \overline{DCI}^-(Y, E; R) \otimes_{R[U]} R[U, U^{-1}],$$

using the fact that \overline{DCI}^- is a finite direct sum of copies of R and $R[U]$. ■

Corollary 7.9. *Let (Y, E) be a rational homology 3-sphere with $U(2)$ -bundle and signature data, and write $c = c_1(E) \in H^2(Y; \mathbb{Z})$. Then there is a natural $\mathbb{Z}/2$ -graded isomorphism*

$$I^\infty(Y, E) \cong R[U^{1/2}, U^{-1/2}] \otimes_{R[\mathbb{Z}/2]} R[H^2Y].$$

Here the first term is Laurent series in $U^{-1/2}$, and the action of $\mathbb{Z}/2$ is $(-1) \cdot U^{k/2} = (-1)^k U^{k/2}$, while the action on $R[H^2Y]$ is given by $(-1) \cdot e^{[x]} = e^{[c-x]}$.

Proof. This is immediate from the computation of the differential on \overline{DCI}^- in the previous section: there is no component of the differential going from the reducible part of \overline{DCI}^- back to itself. At the same time, the reducible part is the only thing that survives to $\overline{DCI}^- \otimes_{R[U]} R[U, U^{-1}]$. Therefore the differential on this is identically zero, and the graded group

$$\overline{DCI}^- \otimes_{R[U]} R[U, U^{-1}]$$

itself calculates the Tate homology.

What we see, then, is that $I^\infty(Y, E; R)$ is a direct sum of copies of $R[U, U^{-1}]$, one for each full reducible and two for each $SO(2)$ -reducible (one of the copies shifted up by 2). Now, the content of Proposition 6.8 (2) is that pairs $\{z_1, z_2\}$ with $z_1 + z_2 = c$ and $z_1 \neq z_2$ are in bijection with the $SO(2)$ -reducible critical orbits, and that the pairs with $z_1 = z_2$ are in bijection with the full reducibles.

We choose the perturbation to be small enough so that this enumeration of reducible flat connections holds.

Contributing two copies of $R[U, U^{-1}]$, one shifted up by degree 2, is equivalent to contributing a single copy of

$$R[U^{1/2}, U^{-1/2}] = R[U, U^{-1}] \oplus U^{-1/2} R[U, U^{-1}],$$

recalling here that $|U^{-1/2}| = 2$.

The description in the corollary arises from the observation that we may give this a more succinct description, only in terms of the cohomology group $H^2(Y)$ and the class $c_1 E$. The point is that

$$R[U^{1/2}, U^{-1/2}] \otimes_{R[\mathbb{Z}/2]} R[H^2Y]$$

may be described explicitly: for a pair $\{z_1, z_2\}$ with $z_1 + z_2 = c$ and $z_1 \neq z_2$, the involution swaps the two corresponding copies of $R[U^{1/2}, U^{-1/2}]$, and therefore in the quotient we have identified these two towers; so we have a copy of $R[U^{1/2}, U^{-1/2}]$ for each such pair. If, on the other hand, $z_1 = z_2$, the action identifies

$$U^{k/2} \otimes e^{[z_1]} \sim (-1)^k U^{k/2} \otimes e^{[c-z_1]} = (-1)^k U^{k/2} \otimes e^{[z_1]}.$$

Because 2 is invertible in R , we see that this kills precisely the terms of the form $U^{k/2}$; so such a pair contributes a copy of $R[U, U^{-1}]$.

This is precisely the description given above as

$$\overline{DCI}^-(Y, E, \pi; R) \otimes_{R[U]} R[U, U^{-1}].$$

■

7.4. Examples of the I^\bullet and the index spectral sequence.

Example 7.2. Let $Y = S^3$ equipped with the trivial bundle. Then $\widetilde{CI}(Y; R) = R$ concentrated in degree 0. Therefore, $CI^+(S^3; R)$ is given as the completion $\widehat{C}_*(BSO(3); R)$. Recall that this means that for $i \in \mathbb{Z}/8$, we have

$$\widehat{C}_i(BSO(3); R) = \prod_{\substack{\mathbb{Z} \ni j \geq 0 \\ j \equiv i \pmod{8}}} C_j(BSO(3); R).$$

Therefore, we have

$$I_*^+(S^3; R) \cong \widehat{H}_*(BSO(3); R).$$

Similarly,

$$I_*^-(S^3; R) \cong H_{\text{fin}}^{-*}(BSO(3); R),$$

meaning that we take cohomology classes which correspond to functionals of finite support; that is, we do not complete in the direction of negative degrees:

$$H_{\text{fin}}^{-i}(BSO(3); R) = \bigoplus_{j \equiv i \pmod{8}}^{\mathbb{Z} \ni j \geq 0} H^{-j}(BSO(3); R).$$

Finally, as graded R -modules we have

$$I^\infty(S^3; R) = \widehat{H}_*(BSO(3); R)[3] \oplus H_{\text{fin}}^{-*}(BSO(3); R).$$

The calculations of these groups, and their module structures, is given in Example A.2 for $\frac{1}{2} \in R$ and $R = \mathbb{Z}/2$.

We are presented with a dichotomy. When $\frac{1}{2} \in R$, we see that

$$\begin{aligned} I^+(S^3; R) &\cong R[[U^{-1}]], \\ I^-(S^3; R) &\cong R[U] \\ I^\infty(S^3; R) &\cong R[U, U^{-1}], \end{aligned}$$

all as $R[U]$ -modules; this periodicity calculation in Tate homology is the content of Lemma A.17, and the other two are straightforward calculations using the definitions of bar and cobar constructions.

However, when 2 is not invertible, we do not have such a periodicity. The most dramatic case is when $R = \mathbb{Z}/2$; in that case,

$$I^+(S^3; R) = (\mathbb{Z}/2)[[w_2^*, w_3^*]],$$

and

$$I^-(S^3; R) = (\mathbb{Z}/2)[w_2, w_3].$$

Then $I^\infty(S^3; R)$ looks like a ‘bi-infinite cone’: the Tate homology $H_{SO(3)}^\infty(\mathbb{Z}/2)$ has rank growing roughly linearly in the degree $|k|$. In fact, the action of $H_{\text{fin}}^{-*}(BSO(3); \mathbb{Z}/2)$ on

$$\widehat{H}_*(BSO(3); \mathbb{Z}/2) \subset H_{SO(3)}^\infty(\mathbb{Z}/2)$$

is nilpotent on each element, as opposed to having an element that induces a periodicity isomorphism $I^\infty(S^3) \rightarrow I^\infty(S^3)$!

This suggests that in the case of $R = \mathbb{Z}/2$, it would be more appropriate to localize $\widetilde{CI}(BSO(3); \mathbb{Z}/2)$ to some specific exterior algebra $\Lambda(u_1)$ or $\Lambda(u_2)$ in $C_*(SO(3); \mathbb{F}_2)$. Mostly we will content ourselves with the case $\frac{1}{2} \in R$.

Example 7.3. If $p > q$ are coprime integers (not necessarily prime), let $L(p, q)$ be the lens space, which by our convention is the quotient of S^3 by the \mathbb{Z}/p action generated by $[1] \cdot (z, w) = (e^{2\pi/p}z, e^{2q\pi/p}w)$.

First, we work with the trivial $U(2)$ -bundle; the set of critical orbits now correspond to

$$\text{Hom}(\pi_1, SU(2))/\sim = \text{Hom}(\mathbb{Z}/p, SU(2))/\sim,$$

where the equivalence relation is conjugacy in $SU(2)$. Because \mathbb{Z}/p is abelian, and simultaneously commuting matrices are simultaneously diagonalizable, this is the same as $\text{Hom}(\mathbb{Z}/p, S^1)/\text{conj}$; here conj is complex conjugation. Identify $\text{Hom}(\mathbb{Z}/p, S^1)$ with the p th roots of unity, and hence with \mathbb{Z}/p again after fixing the generator $e^{2\pi i/p}$; thus reducibles correspond to $(\mathbb{Z}/p)/\pm 1 = [0, p/2]$, and full reducibles — those fixed by the conjugation action, which is ± 1 on \mathbb{Z}/p — correspond to $[0]$ for any p and $[p/2]$ when p is even. The $SO(2)$ -reducibles correspond to $0 < i < p/2$. Finally, 0 corresponds to the trivial connection.

Austin calculates in [Aus90] the expected dimension of the different components of the moduli space of unframed instantons (before quotienting by the translation action) between two flat connections on $L(p, q)$, and in particular the expected dimension of $\mathcal{M}(L(p, q), 0, i)$; the expected dimension of our $\overline{\mathcal{M}}(L(p, q), \theta, \alpha_i)$ is $\dim SO(3) - \dim \mathbb{R} = 2$ dimensions larger; the index $\text{gr}(\alpha_i) = \text{gr}(\theta, \alpha_i)$ differs from this dimension by subtracting $\dim \alpha_i$ and adding 1. Thus we should add either 3 or 1 to Austin’s result; we add 3 in the case that i is either $[0]$ or $[p/2]$, and we add 1 otherwise. Set $\varepsilon(i) = 1$ if $0 < i < p/2$ and $\varepsilon(0) = \varepsilon(p/2) = 0$, and write $0 < q' < p$ for the unique integer with $qq' = ap + 1$.

Then the grading function $\delta(p, q, i) = \text{gr}(L(p, q), i) \in \mathbb{Z}/8$ is given by

$$\delta(p, q, i) = \frac{8i^2q'}{p} - \varepsilon(i) + \frac{2}{p} \sum_{j=1}^{p-1} \cot\left(\frac{j\pi}{p}\right) \cot\left(\frac{jq\pi}{p}\right) \sin^2\left(\frac{2ij\pi}{p}\right) \pmod{8}.$$

Sasahira [Sas13, Corollary 5.3] has given formulas for the $\delta(p, q, i)$ in terms of counting solutions to congruences.

A few observations about this complicated-looking function are in order:

- We have $\delta(p, q, 0) = 0$, as we should.
- If p is even, $\delta(p, q, p/2) = 4q \cdot \frac{p}{2}$; because q' must be odd, $\delta(4k, q, 2k) = 0$ and $\delta(4k + 2, q, 2k + 1) = 4$.
- Because the sum only depends on q' ’s value modulo p , and \cot is odd, when $0 < i < p/2$ we have $\delta(p, q, i) = -\delta(p, p - q, i) - 2$. The factor of -2 comes from $-2\varepsilon(i)$. It is trivially true by the above calculations that $\delta(p, q, i) = -\delta(p, p - q, i)$ when $i = 0$ or $p/2$.
- As Austin observes, $\delta(p, q, i)$ is even for all i .
- If $qq' = ap + 1$, then $\delta(p, q, qi) = \delta(p, q', i)$. Here we are considering d as a function from the integers, but notice that $\delta(p, q, i) = -\delta(p, q, i)$ and $\delta(p, q, i + p) = \delta(p, q, i)$. Then the claimed equality follows because the summation in $\delta(p, q', i)$ is just the summation in $\delta(p, q, qi)$, but with index jq' instead of j ; as $\times q' : ((\mathbb{Z}/p)\setminus 0) \rightarrow ((\mathbb{Z}/p)\setminus 0)$ is a bijection, and j ’s value modulo p is all that’s relevant, the sum is therefore the same.

It seems plausible that $\delta(p, q, ki) = \delta(p, q', i)$ for all i iff $k = 1$ and $q' = q$ or $k = q$ and $q' = q^{-1} \pmod{p}$; calculation shows that this is true at least for

$p \leq 15$. This would be consistent with the classification of lens spaces up to oriented homeomorphism. However, we do not know a proof.

As for the differential, we saw in Lemma 7.4 that if we coherently homotope the endpoint maps of the compactified moduli spaces, the resulting chain complex is quasi-isomorphic to the original. The discussion following that ensures that between reducible orbits α and β , the endpoint maps may be chosen so that the corresponding fiber product maps $\times_{\alpha} \overline{\mathcal{M}}_{\alpha\beta}$ are identically zero on the chain level. While Lemma 7.5 proved this for the $SU(2)$ -model (which may only be applied for coefficient rings with $\frac{1}{2} \in R$), this passage to a double cover is only necessary for the argument when considering the irreducible orbits; here, we have none.

Thus, as a dg $C_*^{\text{gm}}(SO(3))$ -module, and with $L(p, q)$ equipped with the trivial $SO(3)$ -bundle, we see that we have a quasi-isomorphism

$$\widetilde{CI}(L(p, q); R) \simeq \begin{cases} R \oplus_{i=1}^{\lfloor p/2 \rfloor} C_*(S^2; R)[\delta(p, q, i)] & p \text{ odd} \\ R \oplus_{i=1}^{p/2-1} C_*(S^2; R)[\delta(p, q, i)] \oplus R[\delta(p, q, p/2)] & \text{else} \end{cases}$$

Thus to write down I^{\bullet} (where $\bullet = +, -, \infty$) we use Theorem A.26, which says that the equivariant homology groups of an orbit are given by the homology groups of the stabilizer (with a degree shift) as long as the stabilizer is connected. So to take I^{\bullet} , we replace every appearance of $C_*(S^2; R)$ with $H_{SO(2)}^{\bullet}(R)$, with a degree shift of 2 if $\bullet \in \{-, \infty\}$, and every appearance of $R[d]$ with $H_{SO(3)}^{\bullet}(R)[d]$. When $\frac{1}{2} \in R$, we have isomorphisms as $R[U] = H_{SO(3)}^-(R)$ -modules

$$\begin{aligned} H_{SO(2)}^+(R) &\cong R[U^{1/2}, U^{-1/2}]/U^{1/2}R[U^{1/2}] \\ H_{SO(2)}^-(R) &\cong R[U^{1/2}] \\ H_{SO(2)}^{\infty}(R) &\cong R[U^{1/2}, U^{-1/2}] \end{aligned}$$

If p is even, $L(p, q)$ carries a unique nontrivial $SO(3)$ -bundle; choose a lift of this to a $U(2)$ -bundle. We can identify the classes now with $\{a + \xi\}/\pm 1$, where $\xi \in \mathbb{Z}/p$ is a choice of odd number (the choice of $U(2)$ lift). This, then, may be identified with the odd numbers in $\mathbb{Z}/2p$ modulo ± 1 ; no points are fixed, so all are $SO(2)$ -reducibles. So we label the reducibles by $i = 1, 3, \dots, p-1$; there are $p/2$ of them. Their relative grading is also given in [Aus90], now as

$$\begin{aligned} \delta(p, q, i, i') &= \frac{2(i^2 - (i')^2)q'}{p} - 2 + \\ &\frac{2}{p} \sum_{i=1}^{p-1} \cot\left(\frac{j\pi}{p}\right) \cot\left(\frac{jq\pi}{p}\right) \left(\sin^2\left(\frac{i'j\pi}{p}\right) - \sin^2\left(\frac{ij\pi}{p}\right) \right) \pmod{8}. \end{aligned}$$

Again, this is always even.

The result is here independent of the choice of $U(2)$ -lift of the underlying $SO(3)$ -bundle because there are no differentials, and thus we may ignore the orientation of the moduli spaces; in fact, we expect this in general (up to noncanonical isomorphism, the noncanonicity coming from sign choices).

Example 7.4. Suppose Y is an integer homology sphere and E is the trivial bundle. We may use the results of Section 7.2 to determine the index spectral sequences explicitly.

We write $CI(Y)$ for the chain complex $(C_*^{\text{irr}}, \partial_1)$, Floer's original chain complex for integer homology spheres. Donaldson introduced in [Don02, Section 7.1] the complex $\overline{CI}(Y)$, given as $C_*^{\text{irr}} \oplus R$ with differential

$$\bar{\partial} = \begin{pmatrix} \partial_1 & 0 \\ D_1 & 0 \end{pmatrix}.$$

Corollary 7.6 shows that there is an equivariant filtered homotopy equivalence $\widetilde{CI}(Y) \simeq DCI(Y)$, and so to investigate the index spectral sequence computing \bar{I} , we may do the same for the index spectral sequence on the finite-dimensional complex $DCI(Y)$.

We consider $DCI(Y)$; this is, as an R -module, $CI(Y) \oplus CI(Y)[3] \oplus R$, with differential

$$\begin{pmatrix} \partial_1 & 0 & 0 \\ U_{\text{Fl}} & -\partial_1 & D_2 \\ D_1 & 0 & 0 \end{pmatrix}.$$

The differential splits as $d_1 + d_4$ into pieces which decrease the filtration the corresponding amount;

$$d_1 = \begin{pmatrix} \partial_1 & 0 & 0 \\ 0 & -\partial_1 & 0 \\ D_1 & 0 & 0 \end{pmatrix} \text{ and } d_4 = \begin{pmatrix} 0 & 0 & 0 \\ U_{\text{Fl}} & 0 & D_2 \\ 0 & 0 & 0 \end{pmatrix}.$$

Then the E^2 page of the index spectral sequence for $DCI(Y)$ is

$$(DCI(Y), d_1) = \overline{CI}(Y) \oplus CI(Y)[3] \oplus R,$$

and so the E^3 page is $\bar{I}(Y) \oplus I(Y)[3]$. The matrix d_4 defines a chain map

$$(U, D_2) : \overline{CI}(Y) \rightarrow CI(Y)[3];$$

writing the induced map $f : \bar{I}(Y) \rightarrow I(Y)$, we see that the E^5 page is

$$(\bar{I}(Y) \oplus I(Y)[3], f).$$

Therefore, the E^6 page (and all successive pages, for degree reasons) is $\ker(f) \oplus \text{coker}(f)$.

If R is a field, there is nothing left to say; there are no extension problems to resolve, and $\bar{I}(Y) \cong \ker(f) \oplus \text{coker}(f)$.

Now consider the CI^+ spectral sequence by passing through \overline{DCI}^+ . The complex \overline{DCI}^+ is, as a graded R -module, given by $CI(Y) \oplus R[U^*]$, where $|U^*| = 4$, and the differential is given by

$$\begin{pmatrix} \partial_1 & 0 \\ D_1 + D_1 P^+ U_{\text{Fl}} & 0 \end{pmatrix},$$

where recall that by definition when $x \in CI(Y)$, we have

$$(D_1 + D_1 P^+ U_{\text{Fl}})x = \sum_{i \geq 0} (-1)^i D_1 U_{\text{Fl}}^{i+1} x (U^*)^{i+1}.$$

We decompose this differential as $d_1 + d_5 + d_9 + \dots$, where $d_1 = \begin{pmatrix} \partial_1 & 0 \\ D_1 & 0 \end{pmatrix}$ and

$$d_{4i+1}x = \begin{pmatrix} 0 & 0 \\ (-1)^{i+1}D_1U_{\mathbb{F}_1}^i x \otimes (U^*)^i & 0 \end{pmatrix}.$$

Now this is a multicomplex in the sense of Wall, and we may identify the E^2 page as $\bar{I} \oplus U^*R[[U^*]]$.

The map d_5 defines a homomorphism $\bar{I} \rightarrow U^*R$; write \bar{I}_2 for its kernel. Inductively, there is a homomorphism

$$d_{4k+1} : \bar{I}_k \rightarrow (U^*)^k R,$$

and we write its kernel as \bar{I}_{k+1} . We may identify the $q = 0$ line in E^{4k-2} page with \bar{I}_k (identifying the successive differentials with zig-zags of the differentials d_{4k+1} , and using that any compositions of the differentials are zero). In particular, we identify the $q = 0$ line in the E^∞ page with

$$\bar{I}_\infty = \bigcap \ker(D_1U_{\mathbb{F}_1}^k) \subset \bar{I}.$$

A similar discussion presents itself for the CI^- spectral sequence, passing through $\overline{DCI^-}$. As a graded R -module,

$$\overline{DCI^-}(Y) = CI(Y)[3] \oplus R[U];$$

thinking of an element of $R[U]$ as a polynomial f , we write f_k for the coefficient of U^k ; this has differential

$$\bar{\partial}^- = \begin{pmatrix} -\partial_1 & \sum_{k \geq 0} (-1)^{k(k-1)/2} U_{\mathbb{F}_1}^k D_2 f_k \\ 0 & 0 \end{pmatrix}.$$

The top-right term means that

$$\bar{\partial}^-(U^k) = (-1)^{k(k-1)/2} U_{\mathbb{F}_1}^k D_2(1).$$

As before, we may split this into pieces $d_1 + d_5 + d_9 + \dots$. The E^1 page is given as $CI[3] \oplus R$ with differential

$$\begin{pmatrix} -\partial_1 & D_2 \\ 0 & 0 \end{pmatrix},$$

which (shifted by 3 and with a sign change) Donaldson calls $\underline{CI}(Y)$. The homology, written \underline{I} , is the $q = 0$ page of this spectral sequence. The differential d_5 (which is, up to sign, $U_{\mathbb{F}_1}D_2$) determines a homomorphism $R \rightarrow \underline{I}$; its cokernel might be written $\underline{I}^{(2)}$, and may be identified with the E^6 page. Inductively, the differential d_{4k+1} defines a homomorphism $R \rightarrow \underline{I}^{(k)}$, and its cokernel is written $\underline{I}^{(k+1)}$, and may be identified with the E^{4k+2} page.

If Λ is the exterior algebra on a generator of degree 3, the norm map is an isomorphism

$$H_\Lambda^+(\Lambda)[3] \rightarrow H_\Lambda^-(\Lambda)$$

(as a corollary of item (4) of Theorem A.12), and is zero for $H_\Lambda^+(R)[3] \rightarrow H_\Lambda^-(R)$ for degree reasons: the former is supported in degrees no less than 3, and the latter is supported in nonpositive degrees.

Therefore, the induced map of the norm map on the E^1 page of the I^+ and I^- spectral sequences is the identity on C_*^{irr} , and zero on $R[[U^*]]$. We thus identify the image $\bar{I}_{(\infty)} \rightarrow \underline{I}^{(\infty)}$ as the quotient group $\bar{I}_{(\infty)}/\langle U_{\mathbb{F}_1}^i D_2, 1 \rangle$, the second term in the quotient generating the reducible piece of $\bar{I}_{(\infty)}$. This¹⁷ is Frøyshov's reduced Floer group \hat{I} ; the first hints of this were introduced in the d^α homomorphisms of [Frø95] and used to introduce the h -invariant in [Frø02].

Example 7.5. For the sake of compactness of notation, we write R_n in this example to mean $R[n]$ (a copy of R in degree n), and similarly for bigradings $R_{p,q}$. Let us see the above in practice in the explicit case of $\Sigma(2, 3, 5)$. In this case, we have a reducible generator of \widetilde{CI} in degree 0, as well as an irreducible generator in each pair of degrees $(1, 4)$ and $(5, 0)$, with the action of u sending the lower degree generator to the generator 3 degrees larger. So as a u -module our complex is given as

$$R_0 \oplus (R_1 \oplus R_4) \oplus (R_5 \oplus R_0) = R_0 \oplus (R_1 \oplus R_5) \otimes \Lambda(u).$$

For degree reasons, $\partial_1 = 0$, as is $D_2 : \widetilde{CI}_0 \rightarrow \widetilde{CI}_7$. Through either the beautiful and explicit calculations of moduli spaces in [Aus95] or the general theory of [Frø02]¹⁸, we may identify that the map $D_1 : \widetilde{CI}_1 \rightarrow \widetilde{CI}_0$ is the identity map $R \rightarrow R$, while $U_{\mathbb{F}_1} : \widetilde{CI}_1 \rightarrow \widetilde{CI}_0$ and $U_{\mathbb{F}_1} : \widetilde{CI}_5 \rightarrow \widetilde{CI}_4$ are multiplication by one of ± 8 .

The E^1 page of the above spectral sequence is $R_{0,0} \oplus R_{1,0} \oplus R_{1,3} \oplus R_{5,0} \oplus R_{5,3}$, with multiplication-by- u map taking us up 3 vertical degrees. The differential d_1 comprises only the map D_1 , and therefore is the identity $R_{1,0} \rightarrow R_{0,0}$.

The E^2 page is therefore given by $R_{1,3} \oplus R_{5,0} \oplus R_{5,3}$. There are no differentials until the E^5 page, where the next differential is $U_{\mathbb{F}_1} : R_{5,0} \rightarrow R_{1,3}$, which we know to be multiplication-by-8. Since we are working over a ground ring with $\frac{1}{2} \in R$, we see that

$$E^\infty(\widetilde{CI}(\Sigma(2, 3, 5))) = R_{5,3};$$

because there can be no extension problems,

$$\widetilde{I}(\Sigma(2, 3, 5)) = R_0,$$

with trivial multiplication-by- u map.

For the CI^+ spectral sequence, instead we start with E^1 page equal to

$$R_{0,0}[[U^*]] \oplus R_{1,0} \oplus R_{5,0},$$

where $|U^*| = (0, 4)$. Then our first differential d_1 is given as the identity map $R_{1,0} \rightarrow R_{0,0}$; the E^5 page is then identified with $R_{0,4}[[U^*]] \oplus R_{5,0}$. We may identify the differential d_5 with the multiplication-by-8 map $R_{5,0} \rightarrow R_{0,4}$. The E^6 page, then, is simply $R_{0,8}[[U^*]]$, and there are no further differentials. As this is concentrated in a single vertical line, there are no extension problems, and we see

$$I^+(\Sigma(2, 3, 5); R) \cong (U^*)^2 R[[U^*]],$$

with U -action given by contraction against U^* . It is clear from this that the above 'reduced group' is

$$\hat{I}(\Sigma(2, 3, 5)) = 0.$$

¹⁷Despite the choice of notation, these groups are not related to the Tate homology groups.

¹⁸Note that Frøyshov works in cohomology, while we work in homology.

In the CI^- spectral sequence, on the other hand, all differentials past the E^2 page are identically zero, as they factor through the map D_2 . In fact,

$$I^-(\Sigma(2, 3, 5); R) \cong \overline{DCI}^-(\Sigma(2, 3, 5)) = R_0[U] \oplus (R_4 \oplus R_0)$$

as a chain complex; by the formula for the U -map from Section 7.2, which includes a term corresponding to D_1 , we see that the U -action is given by

$$U \cdot (U^k, x, y) = (U^{k+1} + D_1x, 0, 0).$$

Therefore we may write $I^-(\Sigma(2, 3, 5); R) \cong R_4[U] \oplus R_0$ as a U -module, where $|U| = -4$, identifying $1 \in R_0[U]$ with D_1x , where $x \in R_4$ is a generator.

7.5. Orientation reversal and equivariant cohomology. For a dg- A -module M equipped with a periodic filtration whose associated graded complex is bounded, we defined equivariant \bullet -homology chain complexes $C_G^\bullet(M)$, with an action of the ring $C_G^-(R) = \text{Hom}_{\text{fin}}(C_G^+(R), R)$, the negative chains of finite support.

We define the corresponding \bullet -cohomology, written $C_G^\bullet(M)$, to be

$$\text{Hom}(C_G^\bullet(M), R),$$

the literal dual of C_G^\bullet . We then define the instanton \bullet -cohomology, I_\bullet^* , to be $H_\bullet^{SO(3)}(\widetilde{CI})$, where precisely by $SO(3)$ we mean the dg-algebra $C_*^{\text{gm}}(SO(3); R)$. We have the following duality theorem.

Theorem 7.10. *Let $\frac{1}{2} \in R$. There is a canonical isomorphism of $\Lambda(u)$ -modules $\tilde{I}(Y, E) \cong \tilde{I}^*(\bar{Y}, E)$, and furthermore canonical isomorphisms of $H^{-*}(BSO(3); R) = R[U]$ -modules*

$$\begin{aligned} I_*^+(Y, E) &\cong I_-^*(\bar{Y}, E) \\ I_*^-(Y, E) &\cong I_+^*(\bar{Y}, E) \\ I_*^\infty(Y, E) &\cong I_\infty^*(\bar{Y}, E), \end{aligned}$$

Proof. We will show that $\text{DCI}(Y)^\vee$ is isomorphic, as a $\Lambda(u)$ -module, to $\text{DCI}(\bar{Y})$. This implies the rest of the results, as one may check that at the level of \mathcal{C}_A^- -modules, we have $(C_A^+(D))^\vee = C_A^-(D^\vee)$ as well as $(C_A^-(D))^\vee = C_A^+(D^\vee)$. These equalities preserve the norm map, so we get the same result for Tate homology.

Before doing this calculation, recall that if D is a chain complex, its dual D^\vee has the differential $\delta f = (-1)^{|f|+1}fd$. This is isomorphic as an A -module to the chain complex with $\delta f = fd$, with isomorphism given by $f \mapsto (-1)^{|f|(|f|+1)/2}f$. (Observe that this makes sense for any chain complex graded over $\mathbb{Z}/4N$ for any integer $N \geq 0$.) In what follows, we will use the second differential on the dual.

Recall the definition of the chain complex

$$\text{DCI}(Y) = C_*^{\text{irr}}(Y) \oplus C_*^{\text{irr}}(Y)[3] \oplus C_*^{U(1)}(Y) \oplus C_*^{U(1)}(Y)[2] \oplus C_*^\theta(Y)$$

stated before Corollary 7.6. Passing to the dual (and slightly rewriting), we obtain

$$\text{DCI}(Y)^\vee = C_{-*}^{\text{irr}}(Y) \oplus C_{-*}^{\text{irr}}(Y)[3] \oplus C_{-*}^{U(1)}(Y) \oplus C_{-*}^{U(1)}(Y)[2] \oplus C_{-*}^\theta(Y),$$

with differential

$$\widetilde{\partial}_{CI}^\vee := \begin{pmatrix} \partial_1 & 0 & 0 & 0 & 0 \\ U_{F1} & -\partial_1 & V_3 & V_1 & D_1 \\ V_2 & 0 & 0 & 0 & 0 \\ V_4 & 0 & 0 & 0 & 0 \\ D_2 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

Now if π is a regular perturbation on Y , then the same perturbation is regular on \overline{Y} . Using the time-reversal symmetry $\mathbb{R} \times \overline{Y} \cong \mathbb{R} \times Y$, if \mathbf{A} is a connection on the latter going from critical orbits α to β , then it is sent to a connection $\overline{\mathbf{A}}$ going from β to α on $\mathbb{R} \times \overline{Y}$; in particular, the indices of the corresponding deformation operators are equal: $\text{ind}(Q_{\mathbf{A},\pi}^\nu) = \text{ind}(Q_{\overline{\mathbf{A}},\pi}^\nu)$. Recalling the definition of grading from Definition 4.5, we see that

$$\text{gr}_Y(\alpha, \beta) + (\dim \alpha - \dim \beta) = \text{gr}_{\overline{Y}}(\beta, \alpha).$$

In particular, if β is the trivial connection, we see that $i_{\overline{Y}}(\alpha) = -i_Y(\alpha) - \dim \alpha$. Thus the graded basis for $DCI(Y)^\vee$ described above is precisely a graded basis for $DCI(\overline{Y})$.

Furthermore, this time-reversal symmetry gives isomorphisms

$$\overline{\mathcal{M}}_{Y,z}(\alpha, \beta) \cong \overline{\mathcal{M}}_{\overline{Y},-z}(\beta, \alpha).$$

Following this isomorphism, we may use the same coherent family of homotopies to define $DCI(\overline{Y})$ as for $DCI(Y)$. Once we do this, we have equalities of the operators

$$\begin{aligned} D_1(Y)^\vee &= D_2(\overline{Y}), \\ V_1(Y)^\vee &= V_2(\overline{Y}), \\ V_3(Y)^\vee &= V_4(\overline{Y}), \\ U_{F1}(Y)^\vee &= U_{F1}(\overline{Y}). \end{aligned}$$

Therefore the isomorphism of graded vector spaces $DCI(Y)^\vee \cong DCI(\overline{Y})$ described above, which also preserves the $\Lambda(u)$ -action, is in fact an isomorphism of chain complexes. \blacksquare

Remark 7.6. While this theorem is surely true for arbitrary principal ideal domains R , there are technical obstructions in proving it. One natural idea is to try to construct an equivariant pairing $\widetilde{CI}_*(Y) \otimes \widetilde{CI}_*(\overline{Y}) \rightarrow R$, or at least a pairing on a quasi-isomorphism $C_*^{\text{gm}}(SO(3); R)$ -submodule of that tensor product (one defined by the demand that the chains intersect transversely, as in [McC06]). To do this, one needs to be able to put a module structure on the tensor product. If A is a dg-algebra and we wish to endow the category of A -modules with a tensor product structure, we require a comultiplication $\Delta : A \rightarrow A \otimes A$. While this is easy in the model of $C_*(SO(3); R)$ with simplicial chains (given by the Alexander-Whitney map), it seems unlikely this is possible for $C_*^{\text{gm}}(SO(3))$; we have no way to cut a δ -chain into canonical pieces. Perhaps a modification of this model exists that admits the structure of a bialgebra (or even a Hopf algebra), but this is not clear to the author.

A more straightforward approach is to extend the Donaldson model to work with arbitrary principal ideal domains R , where instead of $C_*(SO(3); R) \simeq \Lambda(u)$,

one would instead use $C_*(SO(3); R) \simeq \mathcal{A}$, where \mathcal{A} is a dg-algebra with a generator a_i in each degree $0 \leq i \leq 3$, with the relations

$$d(a_2) = 2a_1, \quad a_1^2 = 0, \quad \text{and} \quad a_1 a_2 = a_3.$$

We expect that the above proof then generalizes easily, but we leave the details to an interested reader.

Appendix

A. EQUIVARIANT HOMOLOGY OF DG-MODULES

Let G be a finite group acting on a topological space X . We may define its *equivariant (co)homology* as the (co)homology of the *Borel space* $(X \times EG)/G$. An algebraic approach to this construction, which works more generally for any $\mathbb{Z}[G]$ -module, replaces $C_*(X)$ with a resolution in the category of chain complexes over $\mathbb{Z}[G]$, takes the quotient by $\mathbb{Z}[G]$, and then computes the homology (or cohomology) of its totalization.

We will want a generalization of this process that will work for compact Lie groups G in some sense “acting on a chain complex” C , which gives equivariant (co)homology in the case of G acting on $C_*(X)$. We model this by considering dg-modules over the dg algebra $C_*(G)$. After that, we will describe the dual homology theory, called *coBorel homology*, and a homology theory called *Tate homology* that compares the two. Our approach to the Borel/bar and coBorel/cobar constructions are strongly inspired by [GM74] and [BMR14], while the approach to Tate homology is essentially that of [Kle02].

We conclude with a discussion of extensions to the case of complexes graded over $\mathbb{Z}/2N$, equipped with an appropriate object called a *periodic filtration*, which will be used in the main body of the text to define the equivariant instanton invariants.

A.1. Bar constructions. We begin by fixing conventions. From here onwards, R is a principal ideal domain (PID) which will serve as the ground ring of all of our chain complexes.

A $\mathbb{Z}/2N$ -graded chain complex of R -modules, for $N \geq 0$, is a collection of R -modules C_a , indexed by $i \in \mathbb{Z}/2N$, equipped with an R -linear map $d_i : C_i \rightarrow C_{i-1}$ so that the composite $d_{i-1}d_i : C_i \rightarrow C_{i-2}$ is zero. The tensor product of $\mathbb{Z}/2N$ -graded chain complexes C and D is

$$(C \otimes D)_n = \bigoplus_{i+j=n} C_i \otimes D_j,$$

with differential acting on $a \otimes b \in C_i \otimes D_j$ as

$$d(a \otimes b) = (da) \otimes b + (-1)^i a \otimes (db)$$

and swap isomorphism

$$\tau(a \otimes b) = (-1)^{|a||b|} (b \otimes a).$$

Here it is necessary that we use $\mathbb{Z}/2N$ as opposed to some odd order finite ring, so that there is a reduction map $\mathbb{Z}/2N \rightarrow \mathbb{Z}/2$.

We will introduce the basic constructions for modules over a differential graded algebra that we use to define the various flavors of equivariant homology.

Let A be a homologically graded unital dg-algebra (over the ground ring R); so A is a \mathbb{Z} -graded chain complex of R -modules whose product satisfies the graded Leibniz rule $d(ab) = d(a)b + (-1)^{|a|}ad(b)$. Furthermore, assume A has an augmentation $\epsilon : A \rightarrow R$ with $\epsilon(1) = 1$; then we may identify

$$\ker(\epsilon) = A/\langle 1 \rangle =: \bar{A},$$

and give R the natural structure of an A (bi)-module. In a differential graded left A -module N , the same Leibniz rule must hold and the unit should act by the

identity; for right A -modules M , the sign in the Leibniz rule uses the grading of N instead of the grading of A .

Definition A.1. *The (two-sided) bar construction $B(M, A, N)$ is the totalization of the resolution*

$$M \otimes N \leftarrow M \otimes \bar{A} \otimes N \leftarrow M \otimes \bar{A}^{\otimes 2} \otimes N \leftarrow \dots$$

So as a graded module,

$$B(M, A, N) = \bigoplus_{n=0}^{\infty} M \otimes \bar{A}[1]^{\otimes n} \otimes N.$$

where $\bar{A}_k = \bar{A}[1]_{k+1}$. Writing a generic tensor product as $m[a_1 | \dots | a_k]n$ where possibly $k = 0$, its differential is given as

$$\begin{aligned} & (-1)^k \left(dm[a_1, | \dots | a_k]n + \sum_{i=1}^k (-1)^{|m|+\varepsilon_{i-1}} m[a_1 | \dots | da_i | \dots | a_k]n + (-1)^{|m|+\varepsilon_k} m[a_1 | \dots | a_k] dn \right) \\ & + \left(ma_1[a_2 | \dots | a_k]n + \sum_{i=1}^{k-1} (-1)^i m[a_1 | \dots | a_i a_{i+1} | \dots | a_k]n + (-1)^k m[a_1 | \dots | a_{k-1}] a_k n \right) \end{aligned}$$

Here $\varepsilon_i = |a_1| + \dots + |a_i| + i$.

These sign conventions are those of [GM74]. Gugenheim and May define $\mathrm{Tor}^A(M, N)$ for arbitrary pairs of dg-modules over a dga; their Corollary A.9 shows that

$$\mathrm{Tor}_*^A(M, N) = H_* B(M, A, N)$$

as long as M , A , and N satisfy appropriate flatness hypotheses. We will state, and then assume, the relevant hypotheses shortly; they are there to ensure that $B(M, A, N)$ takes quasi-isomorphisms to quasi-isomorphisms.

Most of the time we are interested in the special case $B(M, A, R)$. If one thinks of $A = C_*(G)$ as the canonical dg-algebra of interest, acting on the module $C_*(X)$ for some G -space X , then $B(M, A, R)$ models the Borel construction, also known as the ‘homotopy quotient’,

$$M_{hG} = M \times_G EG.$$

Note that whenever $a_k n$ appears in the differential above, this term is zero, as $a_k \in \bar{A} = \ker(\epsilon)$, acting on $n \in R$ via the augmentation ϵ .

The bar construction is functorial (as an A -module) under maps

$$f : A \rightarrow A', \quad g : M \rightarrow M', \quad h : N \rightarrow N'$$

with $g(ma) = g(m)f(a)$ and similarly for h . It is also functorial under homotopies of such maps. Most commonly one either fixes the modules or A when using this functoriality; for a map of triples

$$(M, A, N) \rightarrow (M', A', N'),$$

the corresponding map of bar constructions factors as

$$B(M, A, N) \rightarrow B(M', A, N') \rightarrow B(M', A', N').$$

There is a canonical map, natural in (M, A, N) ,

$$(4) \quad B(M, A, N) \rightarrow B(M, A, R) \otimes B(R, A, N)$$

given as

$$m[a_1 | \cdots | a_k]n \mapsto \sum_{i=0}^k (-1)^{(k-i)(|m|+\varepsilon_i)} m[a_1 | \cdots | a_i] \otimes [a_{i+1} | \cdots | a_k]n.$$

This endows $BA := B(R, A, R)$ with the structure of a coalgebra, and $B(M, A, R)$ the structure of a right BA -comodule.

Definition A.2. *The cobar construction of (N, A, M) , where A is a dg-algebra and both N and M are right A -modules, is the chain complex*

$$cB(N, A, M) = \text{Hom}_A(B(N, A, A), M).$$

This, too, is a special case Gugenheim and May's constructions; now we have $H_*cB(N, A, M) = \text{Ext}_A^{-*}(N, M)$ under suitable projectivity hypotheses. We will primarily be interested in the case $N = R$.

For concreteness, we note that $cB(N, A, M)$ is isomorphic as a graded R -module to

$$\prod_{i=0}^{\infty} \text{Hom}_R(N \otimes \overline{A}[1]^{\otimes i}, M) = \text{Hom}_R(N \otimes BA, M),$$

graded so that $|a| + |\eta| = |\eta(a)|$. This is the negative of the usual grading (which would make $\text{Hom}_R(C, D)$ into a *cochain* complex); for M bounded and A nonnegatively graded, the *chain* complex we define here is bounded above but unbounded in the negative direction! Passing through the above isomorphism, the differential of an element $\eta : B(N, A, R) \rightarrow M$ of degree d is given as

$$\begin{aligned} (d\eta)(n[a_1 | \cdots | a_p]) &= d_M(\eta(n[a_1 | \cdots | a_p])) - (-1)^d \eta(d_{B(N, A, R)}n[a_1 | \cdots | a_p]) \\ &\quad - (-1)^{d+p} \eta(n[a_1 | \cdots | a_{p-1}])a_p. \end{aligned}$$

The cobar construction is functorial for triples (M, A, N) ; it is covariant in M , but contravariant in A and N . The equivariance condition on morphisms is that we demand $f : A' \rightarrow A$ and $g : M \rightarrow M'$ satisfy $g(mf(a')) = g(m)a'$; the condition on A and N is the usual, that $h(n'a') = h(n')f(a')$.

Dual to equation (4),

$$cB(R, A, R) = \text{Hom}_R(BA, R)$$

is naturally an algebra. The chain complex $cB(R, A, M)$ is naturally a left module over this algebra, following the diagram

$$\begin{aligned} &\text{Hom}_R(B(R, A, R), R) \otimes \text{Hom}_A(B(R, A, A), M) \\ &\rightarrow \text{Hom}_A(B(R, A, R) \otimes B(R, A, A), M) \rightarrow \text{Hom}_A(B(R, A, A), M) = cB(R, A, M). \end{aligned}$$

Here the final map is dual to the BA -comodule structure on $B(R, A, A)$; note that the map $B(R, A, A) \rightarrow B(R, A, R) \otimes B(R, A, A)$ is a map of right A -modules, so it makes sense to apply Hom_A to this map.

Finally, observe that $B(M, A, R)$ is a *left* module over $cB(R, A, R)$, by applying the comodule structure and then the pairing. In the following map, we suggestively

rewrite $cB(R, A, R) = BA^\vee$:

$$\begin{aligned} BA^\vee \otimes B(M, A, R) &\xrightarrow{1 \otimes \Delta} BA^\vee \otimes B(M, A, R) \otimes BA \\ &\xrightarrow{\tau \otimes 1} B(M, A, R) \otimes BA^\vee \otimes BA \xrightarrow{1 \otimes \text{eval}} B(M, A, R). \end{aligned}$$

The map Δ is the comodule structure and τ is the swap map $x \otimes y \mapsto (-1)^{|x||y|} y \otimes x$. Here, and elsewhere in this text, we write M^\vee for an R -module M to be its dual as an R -module, $\text{Hom}_R(M, R)$; if M was a right A -module, then M^\vee carries the structure of a left A -module by $(a\eta)(m) = \eta(ma)$.

A.2. Invariance. Suppose we have a map of pairs $(M, A, N) \rightarrow (M', A', N')$ as above inducing a map $B(M, A, N) \rightarrow B(M', A', N')$. When is the induced map an isomorphism on homology? Because the bar construction is functorial under homotopies, this is true if the map of triples is a homotopy equivalence (of triples). More generally, we have the following theorem. These flatness restrictions are harmless for most purposes in topology and for use in this paper, but indicate that the approach taken here is too naive for general modules over a dga. This result, and much of [GM74], is put into the powerful general framework of model categories in [BMR14], but we will not need this language here.

Theorem A.1. *Suppose we have a map of triples*

$$(g, f, h) : (M, A, N) \rightarrow (M', A', N')$$

so that f , g and h all induce isomorphisms on homology (from here on we will say “are quasi-isomorphisms”). Suppose further that N, N', \bar{A} and \bar{A}' are flat (as graded R -modules). Then the induced map $B(M, A, N) \rightarrow B(M', A', N')$ is a quasi-isomorphism.

Note in particular that we may take $N = R$ in the above theorem. A similar result is true of $cB(N, A, M)$ (making slightly stronger assumptions).

Theorem A.2. *Suppose M, N are A -modules and N, N' are A' -modules. Let*

$$f : A' \rightarrow A, \quad g : M \rightarrow M' \quad h : N' \rightarrow N$$

be quasi-isomorphisms that are equivariant in the sense that

$$g(mf(a')) = g(m)a' \quad \text{and} \quad h(n'a') = h(n')f(a').$$

If all of A, A', N and N' are R -free, then $cB(N, A, M) \rightarrow cB(N', A', M')$ is a quasi-isomorphism.

These theorems are proved by appealing to natural spectral sequences associated to the bar and cobar constructions. Now, and later in this appendix, it's useful to use Boardman's language of *strongly convergent* and *conditionally convergent* spectral sequences from [Boa99]. We have slightly different assumptions and indexing than Boardman, chosen to fit with our preference for homological gradings.

Definition A.3. *Suppose C is a chain complex equipped with an increasing (possibly unbounded) filtration*

$$\dots \subset F_p C \subset \dots \subset C.$$

We will call this a ‘filtered complex’. If $\bigcup_p F_p C = C$, then we say that $F_\bullet C$ is an exhaustive filtration. We say that the filtration is Hausdorff if $\bigcap_p F_p C = 0$ and complete if the map $C \rightarrow \lim_p C/F_{-p} C$ is an isomorphism. We say that

the associated spectral sequence E^r converges conditionally to H_*C if the filtration $H_*(F_pC)$ of $H_*(C)$ is complete and Hausdorff. We say that it converges strongly to H_*C if furthermore we have isomorphisms $E_{p,q}^\infty \cong \text{gr}_p H_{p+q}C$, where $\text{gr}_p M$ for a filtered module M , we write the associated graded module

$$\text{gr}_p M = F_p M / F_{p-1} M.$$

We will only ever care about filtrations that are exhaustive and complete Hausdorff (but it is still important to check these conditions). Completeness in particular is often subtle and delicate.

Conditional convergence does not promise us that we can calculate H_*C from the E_∞ page, but it is both quite powerful and quite flexible. Its relevance comes from the following proposition ([Boa99, Theorems 8.2-8.3]):

Proposition A.3. *Suppose we have exhaustive filtered complexes C and C' such that associated spectral sequences converge conditionally to H_*C and H_*C' . If $f : C \rightarrow C'$ is a map of filtered complexes inducing an isomorphism on some finite page $E^r(F_*C) \rightarrow E^r(F_*C')$, then f is a quasi-isomorphism. There are invariants RE^∞ and $W(E)$ of conditionally convergent spectral sequences so that if $RE^\infty = W(E) = 0$, then the spectral sequence converges strongly.*

If E is a half-plane spectral sequence, then $W(E) = 0$. If there are only finitely many differentials leaving a given $E_{s,t}$, then $RE^\infty = 0$. If E degenerates on some finite page, then both $RE^\infty = W(E) = 0$.

So even though we may not be able to use the E^∞ page of a conditionally convergent spectral sequence to calculate homology, we can still use it to detect quasi-isomorphisms. To apply this, we will make frequent use of [Boa99, Theorems 9.2-9.3], which show that conditionally convergent spectral sequences are relatively common:

Proposition A.4. *If C is a filtered complex whose filtration is exhaustive and complete Hausdorff, then the associated spectral sequence, whose E^1 page given by $H(F_pC/F_{p+1}C)$, is conditionally convergent to $H_*(C)$. If the filtration of C is not exhaustive, complete, or Hausdorff, the associated spectral sequence instead converges conditionally to $\hat{C} = \lim_p \text{colim}_s F_s C / F_{-p} C$.*

If the filtration is bounded below, it is automatically both complete and Hausdorff; if it is bounded above, it is automatically exhaustive. Another situation that frequently arises is that the filtration is the totalization of a *multicomplex*; in this case, the completion is easy to describe. We state the definition here (first appearing in [Wal61]).

Definition A.4. *A multicomplex is a bigraded R -module $M_{s,t}$ with differentials $d_r : M_{s,t} \rightarrow M_{s-r,t+r-1}$ for $r \geq 0$ so that*

$$\sum_{i+j=n} d_i d_j = 0.$$

The associated (completed) filtered complex is the subcomplex

$$\hat{C}_n = \prod_{s \rightarrow -\infty} M_{s,n-s} \subset \prod_{s \in \mathbb{Z}} M_{n,n-s}$$

consisting of those sequences (x_s) with $x_s = 0$ for sufficiently large s . (That is to say, the product is only taken in the negative direction.) The differential is given as $\sum_{r \geq 0} d_r$.

Proposition A.5. *The completed complex \hat{C} of a multicomplex M is complete Hausdorff, and we can identify the E^1 page of the associated (conditionally converging) spectral sequence with $H(M, d_0)$ equipped with the differential $H(d_1)$.*

These tools in hand, we can prove the two invariance results we need.

Proof of Theorem A.1. Filter $B(M, A, N)$ by

$$F_p B(M, A, N) = M \otimes \left(\bigoplus_{i=0}^p \bar{A}^{\otimes i} \otimes N \right).$$

Then we calculate the E^1 page of the associated spectral sequence as the (bigraded) complex

$$H(M \otimes N) \leftarrow H(M \otimes \bar{A} \otimes N) \leftarrow \dots$$

Because the map $B(M, A, N) \rightarrow B(M', A', N')$ preserves the filtration, it induces a map of spectral sequences; if the maps

$$H(M \otimes \bar{A}^{\otimes i} \otimes N) \rightarrow H(M' \otimes \bar{A}'^{\otimes i} \otimes N')$$

are isomorphisms, then we would have proved that $B(g, f, h)$ induces an isomorphism on the E^1 page and hence all pages of the spectral sequence.

We prove this fact inductively on the number of tensor factors.

Lemma A.6. *If $F : X \rightarrow X'$ is a quasi-isomorphism of degreewise R -flat complexes, and $G : Y \rightarrow Y'$ is a quasi-isomorphism of arbitrary complexes, then $F \otimes G : X \otimes Y \rightarrow X' \otimes Y'$ is a quasi-isomorphism.*

This follows immediately from comparing the short exact sequences (which exist because X, X' are degreewise flat)

$$0 \rightarrow \bigoplus_{p+q=n} H_p(X) \otimes H_q(Y) \rightarrow H_n(X \otimes Y) \rightarrow \bigoplus_{p+q=n-1} \text{Tor}(H_p X, H_q Y) \rightarrow 0,$$

because the outside terms only depend on HX and HY . This is where we need R to be a PID; otherwise we would need to apply the Kunnetth spectral sequence, and would need boundedness assumptions. One may use the spectral sequence argument to extend this result to Dedekind domains, as in [KM95, Lemma 2.2].

The filtration of $B(M, A, N)$ is trivially complete and Hausdorff, as $F_{-1} = 0$. It is exhaustive because the infinite direct sum is the union of its finite direct summands.

Because N, N', \bar{A} , and \bar{A}' are degreewise R -flat, we have $B(g, f, h)$ is an isomorphism on the E^2 pages. Because the filtration is bounded below, the spectral sequences converge strongly by Proposition A.4 and thus $B(g, f, h)$ is a quasi-isomorphism by Proposition A.3. \blacksquare

Proof of Theorem A.2. Now the filtration is

$$F_{-p} cB(N, A, M) = \text{Hom}_A(B(N, A, A)/F_p B(N, A, A), M);$$

that is, it consists of those functionals that vanish on $F_p B(N, A, A)$. Let us abbreviate $B(N, A, A) =: B$. This filtration is now bounded above, and thus is automatically exhaustive. Note that the intersection

$$\lim_{\leftarrow p} \text{Hom}_A(B/F_p, M) = \text{Hom}_A(\text{colim}_p B/F_p, M).$$

The colimit is an exact functor, so

$$\text{colim}_p B/F_p = B/\text{colim}_p F_p = B/B = 0,$$

because the filtration of B was exhaustive, and thus the filtration is Hausdorff. Now, the short exact sequence $0 \rightarrow F_p \rightarrow B \rightarrow B/F_p \rightarrow 0$ is *split* as A -modules (not dg), identifying B/F_p with the summand $\bigoplus_{i>p} N \otimes \overline{A}^{\otimes i} \otimes A$ (but ignoring the differential). This implies that

$$\mathrm{Hom}_A(B, M)/\mathrm{Hom}_A(B/F_p, M) \cong \mathrm{Hom}_A(F_p, M)$$

as R -modules (Hom takes split exact sequences to split exact sequences). Applying

$$\begin{aligned} \lim \mathrm{Hom}_A(B, M)/\mathrm{Hom}_A(B/F_p, M) &\cong \lim \mathrm{Hom}_A(F_p, M) \\ &= \mathrm{Hom}_A(\mathrm{colim} F_p, M) = \mathrm{Hom}_A(B, M), \end{aligned}$$

we see that the filtration is complete (the last equality because the filtration on B was exhaustive).

Using the isomorphism of $cB(N, A, M)$ as a graded R -module to $\mathrm{Hom}(N \otimes BA, M)$, we identify the E^1 page as the totalization of the double complex

$$H(\mathrm{Hom}_R(N, M)) \rightarrow H(\mathrm{Hom}_R(N \otimes \overline{A}, M)) \rightarrow \cdots$$

We need to see that $\mathrm{Hom}(N \otimes \overline{A}^{\otimes i}, M) \rightarrow \mathrm{Hom}(N' \otimes \overline{A}'^{\otimes i}, M')$ is a quasi-isomorphism for all i ; then the theorem will be proved. We write this as a lemma: the only property we will use is that $N \otimes \overline{A}^{\otimes i}$ is free for any i (a tensor product of free modules is free). \blacksquare

Lemma A.7. *If X and X' are R -free, and $F : X' \rightarrow X$ is a quasi-isomorphism, then for any quasi-isomorphism $G : Y \rightarrow Y'$, the induced map $\mathrm{Hom}_R(X, Y) \rightarrow \mathrm{Hom}_R(X', Y')$ is a quasi-isomorphism.*

Proof. When both X and $d(X) \subset X$ are complexes of projective modules, we have a natural Kunnet short exact sequence

$$0 \rightarrow \prod_{p+q=n-1} \mathrm{Ext}^1(H_p X, H_{-q} Y) \rightarrow H_{-n} \mathrm{Hom}(X, Y) \rightarrow \prod_{p+q=n} \mathrm{Hom}(H_p X, H_{-q} Y) \rightarrow 0;$$

from this the theorem is clear, as long as every submodule $d(X)$ of a free module of arbitrary rank is free; this is true for PIDs. For the statement of the Kunnet theorem and that submodules of free modules are projective (and hence free, as projective modules over a PID are free), see [Wei95, Exercises 3.6.1-3.6.2]. \blacksquare

We will often be interested in thinking of one-variable versions of the bar and cobar constructions as providing homology theories for A -modules; we will introduce new notation for the sake of compactness. Our notation is chosen to fit with both [Jon87] and [OS04].

Definition A.5. *Suppose that A is an augmented dg-algebra with each A_n free over R , and M a right A -module. We define the positive and negative A -chains on M to be*

$$\begin{aligned} C_A^+(M) &:= B(M, A, R) \\ C_A^-(M) &:= cB(R, A, M). \end{aligned}$$

The homology of $C_A^+(M)$ is denoted $H_A^+(M)$, the (Borel, or positive) A -homology of M , while the homology of $C_A^-(M)$ is denoted $H_A^-(M)$, the coBorel (or negative) A -homology of M .

Note that both $C_A^+(M)$ and $C_A^-(M)$ are *covariant* in M ! When A is clear and we are not interested in varying it, we omit it from the notation.

In the $\mathbb{Z}/2N$ graded setting will need versions of these invariance results for the *completed* bar constructions and the *finitely supported* cobar constructions. By the nature of these completions, the spectral sequence arguments no longer work (the filtrations will fail to either be exhaustive or complete, respectively), and so we are left with significantly weaker results. They will, however, be enough for us.

Definition A.6. *Let A be a dg-algebra, and let M be a right A -module. If N is a left A -module, the completed bar construction $\hat{B}(M, A, N)$ is given as*

$$\prod_{i=0}^{\infty} M \otimes \bar{A}^{\otimes i} \otimes N,$$

with differential defined by the same formula as before, which extends to the completion because every term sends an element of $B_i = M \otimes \bar{A}^{\otimes i} \otimes N$ to either B_i or B_{i-1} . If N is a right A -module, the finitely supported cobar construction $\widehat{cB}(N, A, M)$ is given as $\text{Hom}_A^{\text{fin}}(B(N, A, A), M)$, the set of those A -equivariant maps which vanish on

$$\bigoplus_{i \geq k} M \otimes \bar{A}^{\otimes i} \otimes N$$

for some sufficiently large k .

Essentially, what we need is the following.

Lemma A.8. *Suppose $(g, f, h) : (M, A, N) \rightarrow (M', A', N')$ satisfy the conditions of Theorem A.1, and additionally, suppose that the algebras A, A' are supported in non-negative degrees, and all of M, M', N, N' are bounded below. Then*

$$\left(\hat{B}(M, A, N) \right)_k = B(M, A, N)_k;$$

in particular, the induced map $\hat{B}(g, f, h)$ is a quasi-isomorphism.

This is self-evident: if M and N are supported in degrees above m and n , respectively, then $M \otimes \bar{A}^{\otimes i} \otimes N$ is supported in degrees above $m + i + n$. So elements of $\hat{B}(M, A, N)$ in any fixed degree k in fact lie in $B(M, A, N)$.

We have the dual assertion for the cobar construction.

Lemma A.9. *Suppose $(g, f, h) : (M', A', N) \rightarrow (M, A, N')$ satisfy the conditions of Theorem A.1, and additionally, suppose that the algebras A, A' are supported in non-negative degrees, the modules N, N' are bounded below, and M, M' are bounded above. Then*

$$\left(\widehat{cB}(N, A, M) \right)_k = cB(N, A, M)_k;$$

in particular, the induced map $\widehat{cB}(h, f, g)$ is a quasi-isomorphism.

Here, because N is bounded below (by degree, say, n) and A is concentrated in non-negative degrees, we see that $\bigoplus_{i \geq k} N \otimes \bar{A}[1]^{\otimes i} \otimes A$ is supported in degrees at least $n + k$. So an element of $cB(N, A, M)$ of degree d is a map $B(N, A, A) \rightarrow M$ which sends an element of

$$\bigoplus_{i \geq k} N \otimes \bar{A}[1]^{\otimes i} \otimes A$$

to some M_p , where $n + k - d \leq p$. Because M is bounded above, as k increases without bound, we see that the target of such a map is zero. So any element of $cB(N, A, M)$ is finitely supported in the above sense.

A.3. The dualizing complex and Tate homology. Tate homology, constructed in this section for dg-modules over a dg-algebra and written as $H_A^\infty(M)$, has appeared in the literature in many forms. Our approach here is essentially a chain-level interpretation of [Kle02] (which was written in the context of Tate homology of G -spectra). Tate homology may be viewed, in some sense, as the homology theory that arises when you kill off free objects. The classic reference to Tate homology of spectra is [GM95]; we warn that if X is a G -space, our $H_{C_*(G)}^\infty(C_*(X); R)$ is more analogous to what they would call $t_{HR \wedge G}(HR)(X)$, not $t_G(HR)(X)$. (We are ‘chainifying’ the group from the start as well.) The idea of Tate homology is beautifully developed in an abstract homotopical setting in [Gre01], which surely includes as a special case the content of this section. A recent approach to Tate homology of G -spectra via localization of $(\infty, 1)$ -categories appears in [NS17]; this level of abstraction has the advantage of giving multiplicativity and uniqueness results that are not as easily available otherwise (in some cases, not available at all).

We take our current approach as it seems to minimize input energy, at the cost perhaps of some conceptual clarity and multiplicativity results (we do not construct a product on $C_A^\infty(R)$, for instance).

Because A is a left A -module, the complex $cB(R, A, A) = C_A^-(A)$ inherits the structure of a left A -module, by $(a\eta)([b_1 | \cdots | b_n]b) = a \cdot \eta([b_1 | \cdots | b_n]b)$.

Definition A.7. *The dualizing complex D_A of a degreewise R -free dg-algebra A is the left A -module $C_A^-(A)$.*

Definition A.8. *Let A be a degreewise R -free dg-algebra and M a left A -module. The map*

$$N_M : B(M, A, D_A) \rightarrow C_A^-(M)$$

given by

$$N_M(m[a_1 | \cdots | a_k]\psi) = \begin{cases} 0 & k > 0 \\ m \cdot \psi & k = 0 \end{cases}$$

is called the norm map of M . The product $m \cdot \psi$ makes ψ into an M -valued functional by using the fact that ψ is A -valued and M is a right A -module. The mapping cone of N_M is denoted $C_A^\infty(M)$, the Tate complex of M . Its homology is the Tate homology $H_A^\infty(M)$ of M .

The following is important enough to record as a lemma.

Lemma A.10. *D_A is degreewise R -flat.*

Proof. Recall that

$$D_A = cB(R, A, A) = \prod_{i=0}^{\infty} (\overline{A}^\vee)^{\otimes i} \otimes A;$$

A is degreewise R -free, and in particular torsion-free. Because R is a PID, R -flatness is equivalent to (R) -torsion-freeness. Duals of torsion-free modules are torsion-free, sums of torsion-free modules are torsion free, and products of torsion-free modules are torsion-free, so $cB(R, A, A)$ is degreewise torsion-free, and hence degreewise flat. ■

To justify the name ‘‘Tate homology’’, we show that this satisfies part of the corresponding versions of Klein’s axioms defining Tate cohomology [Kle02], skipping the complete verification that H_A^∞ and $H(B(-, A, D_A))$ are homology theories: the remaining axioms state that these preserve homotopy pullbacks and filtered homotopy colimits. These are not hard to verify, but we will not use them.

Theorem A.11. *The functor $C_A^\infty(M)$ from left A -modules to chain complexes satisfies Klein’s four axioms specifying the Tate homology of M :*

- (1) $C_A^\infty(M)$ preserves weak equivalences;
- (2) $H_A^\infty(X \otimes A) = 0$, where X is a chain complex and the right A -module structure is given by acting on A ;
- (3) There is a map $C_A^-(M) \rightarrow C_A^\infty(M)$, natural in M , whose homotopy fiber preserves weak equivalences.

Proof. It follows from Theorem A.1 that $B(M, A, D_A)$ preserves weak equivalences in M , because by Lemma A.10 above, D_A is flat. The Tate complex is the mapping cone (homotopy cofiber) of $N_M : B(M, A, D_A) \rightarrow C_*^-(M)$, and both the domain and the codomain of the norm map preserve weak equivalences, so $C_*^\infty(M)$ does as well. This is Axiom 1.

For Axiom 2, note that

$$B(X \otimes A, A, D_A) \cong X \otimes B(A, A, D_A).$$

Then the norm map is identified with the projection $X \otimes B(A, A, D_A) \rightarrow X \otimes D_A$; but $B(A, A, N) \rightarrow N$ is a homotopy equivalence for any left A -module N , so the mapping cone of the norm map is contractible.

The first part of Axiom 3 is obvious (it is the inclusion into a mapping cone), and the second part almost so: the homotopy fiber = mapping cocone is naturally equivalent to $B(M, A, D_A)$. ■

We henceforth write $H_A^{+, \text{tw}}(M)$ for $H(B(M, A, D_A))$, and call it the *twisted Borel homology*. We will investigate its relationship to $H_A^+(M)$ later.

We won’t prove Klein’s uniqueness theorem that these axioms do uniquely characterize Tate cohomology, but rather use it as motivation that we have the correct definition. (It seems likely that some variation of his argument works in this context.)

In addition to the above, it is important to observe that there is a natural left action of $C_A^-(R)$ on D_A , and hence on $B(M, A, D_A)$; the norm map N_M is easily seen to be $C_A^-(R)$ -equivariant. Therefore, the mapping cone inherits the structure of a left $C_A^-(R)$ -module, and the natural map $C_A^-(M) \rightarrow C_A^\infty(M)$ is a module homomorphism.

The following theorem summarizes everything we’ve assembled about the three A -homology functors $H^{+, \text{tw}}, H^-, H^\infty$.

Theorem A.12. *Let A be a dg-algebra over a commutative PID R which is $(R-)$ flat in each degree A_n . There are functors*

$$H_A^{+, \text{tw}}(M), \quad H_A^-(M), \quad H_A^\infty(M),$$

from dg A -modules to graded R -modules, satisfying the following properties.

- (1) *The functors send short exact sequences of A -modules to exact triangles and preserve weak equivalences.*

- (2) $H_A^-(R)$ is a ring, and each of these homology theories carry a natural left module structure over $H_A^-(R)$.
- (3) There is an exact triangle of $H_A^-(R)$ -modules

$$H_A^{+,tw}(M) \rightarrow H_A^-(M) \rightarrow H_A^\infty(M) \xrightarrow{[-1]} H_A^{+,tw}(M) \rightarrow \dots$$

- (4) $H_A^\infty(X \otimes A) = 0$ when X is a chain complex and $X \otimes A$ is given the canonical right action.

We state an invariance theorem for equivariant homology with respect to quasi-isomorphic dgas.

Proposition A.13. *Suppose $f : A \rightarrow A'$ is a quasi-isomorphism of algebras, each degreewise free over R . This induces a functor $F : \text{Mod}_{A'} \rightarrow \text{Mod}_A$ via restriction of scalars, and there are natural isomorphisms $H_A^\bullet(FM) \xrightarrow{\cong} H_{A'}^\bullet(M)$, for the homology theories $\bullet \in \{+, tw, -, \infty\}$. These natural isomorphisms induce an $H_A^-(R)$ -equivariant natural isomorphism of exact triangles.*

Proof. The fact that $cB(R, A, M) \leftarrow cB(R, A', M)$ is a quasi-isomorphism is Theorem A.2; for that reason, the maps

$$D_A = cB(R, A, A) \rightarrow cB(R, A, A') \leftarrow cB(R, A', A') = D_{A'}$$

are quasi-isomorphisms. This is the crucial place we need A and A' to be degreewise free.

To make the diagrams smaller, let us write

$$\begin{aligned} B_A &= B(D_A, A, FM), \\ cB_A &= C_A^-(FM) = \text{Hom}_A(B(A, A, R), FM), \\ B' &= B(C_{A'}^-(A'), A', M). \end{aligned}$$

Then the induced maps on bar constructions $B_A \rightarrow B' \leftarrow B_{A'}$ are quasi-isomorphisms by Theorem A.1, using that $D_A, C_{A'}^-(A')$, and $D_{A'}$ are all flat R -modules, which is Lemma A.10.

Comparing the Tate homology groups is more complicated. The following diagram commutes, which is an easy check left to the reader.

$$\begin{array}{ccccc} B_A & \longrightarrow & B' & \longleftarrow & B_{A'} \\ \downarrow N_A & & \swarrow N' & & \downarrow N_{A'} \\ cB_A & \longleftarrow & & \longrightarrow & cB_B \\ \downarrow & & \searrow & & \downarrow \\ C(N_A) & \dashrightarrow & C(N') & \longleftarrow \dashrightarrow & C(N_{A'}) \end{array}$$

Here $N' : \text{Hom}_A(B(A, A, R), A') \otimes_{A'} B(A', A', M) \rightarrow \text{Hom}_A(B(A, A, R), M')$ is obtained from the projection map $B(A', A', M) \rightarrow M$. The maps in the bottom row are the natural maps induced on a cone, and we denote the norm maps N_A and $N_{A'}$ instead of the usual N_M to reflect the different choice of algebra M is a module over.

Because all of the maps in the top two rows are quasi-isomorphisms, the maps on the bottom row are also quasi-isomorphisms by an application of the five lemma.

In particular, inverting the bottom-right quasi-isomorphism on homology, we have an isomorphism $H_A^\infty(FM) \rightarrow H_{A'}^\infty(M)$.

All of the maps in this diagram are module homomorphisms with respect to either $C_A^-(R)$ or $C_{A'}^-(R)$, as appropriate; this implies that the isomorphisms on homology are $H_A^-(R)$ -equivariant. \blacksquare

In what follows we briefly describe the completed version of Tate homology and its comparatively weak invariance properties. The finitely supported dualizing complex, \widehat{D}_A , is precisely

$$\widehat{cB}(R, A, A) = \text{Hom}_A^{\text{fin}}(B(R, A, A), A).$$

As before, there is a natural map $\widehat{B}(M, A, \widehat{D}_A) \rightarrow \widehat{cB}(R, A, M)$; we write $\widehat{C}_A^\infty(M)$ for its mapping cone, and (for now) call this the completed Tate homology.

To begin with, we can compare \widehat{D}_A and $\widehat{D}_{A'}$, using Lemma A.9.

Lemma A.14. *Suppose we are given a quasi-isomorphism $f : A \rightarrow A'$ of nonnegatively graded, R -free, and bounded above dg-algebras A, A' . Then*

$$\widehat{D}_A \rightarrow \widehat{cB}(R, A, A') \leftarrow \widehat{D}_{A'}$$

is a zig-zag of quasi-isomorphisms.

To compare the completed versions of the twisted Borel homology, we need a more subtle argument, as follows. We allow to vary the dg-algebra involved in the dualizing complex to make use of the zig-zag above.

Lemma A.15. *Suppose A and A' are non-negatively graded, R -free, and bounded above. Let $f : A \rightarrow A'$ be a dg-algebra quasi-isomorphism, and $f : M \rightarrow M'$ be a quasi-isomorphism of bounded below right A' -modules (thought of also as A -modules by restriction of scalars), then the maps*

$$\widehat{B}(M, A, \widehat{D}_A) \rightarrow \widehat{B}(M, A, \widehat{cB}(R, A, A')) \leftarrow \widehat{B}(M, A, \widehat{D}_{A'}) \rightarrow \widehat{B}(M', A', \widehat{D}_{A'})$$

are all quasi-isomorphisms.

Proof. All of these will fall prey to the same argument, so we investigate the leftmost one in particular.

The right A -module $\widehat{B}(R, A, A)$ (which agrees with $B(R, A, A)$ by Lemma A.8) is the increasing union of the bounded subcomplexes

$$B^{\leq k}(R, A, A) = \bigoplus_{i=0}^k \overline{A}[1]^{\otimes i} \otimes A;$$

this stabilizes in each degree for some finite k (but this $k = k(i)$ depends on the degree, and may be arbitrarily large). Here we used that A is bounded to see that these subcomplexes are bounded. Therefore there is also a sequence of bounded quotient complexes

$$\widehat{cB}^{\leq k}(R, A, A) = \text{Hom}_A^{\text{fin}}(B^{\leq k}(R, A, A), A),$$

and similarly with $\widehat{cB}^{\leq k}(A, A')$. We have

$$\lim_k \widehat{cB}^{\leq k}(R, A, A) = \widehat{cB}(R, A, A),$$

and this limit stabilizes degreewise at some finite k (again k depending on the degree). Further, observe that the map $\widehat{cB}^{\leq k}(R, A, A) \rightarrow \widehat{cB}^{\leq k}(R, A, A')$ is a quasi-isomorphism by the argument of Theorem A.2; now the filtration used there is finite.

Now $\widehat{B}(M, A, \widehat{D}_A^{\leq k})$ and $\widehat{B}(M, A, \widehat{cB}^{\leq k}(R, A, A'))$ both satisfy the conditions of Lemma A.8, so the map

$$\widehat{B}(M, A, \widehat{D}_A^{\leq k}) \rightarrow \widehat{B}(M, A, \widehat{cB}^{\leq k}(R, A, A'))$$

is a quasi-isomorphism for each k . Next, observe that

$$\widehat{B}(M, A, \widehat{D}_A) \cong \lim_k \widehat{B}(M, A, \widehat{D}_A^{\leq k}) :$$

there is a natural chain map from the leftmost complex to the right, which is an isomorphism at the level of R -modules because limits commute (the second relevant limit is $\widehat{B} = \lim B^{\leq k}$, which makes sense at the level of R -modules). Because this limit is over a tower of quotient maps, the tower in particular satisfies the Mittag-Leffler condition, and we may apply [Wei95, Theorem 3.5.8], which says that there is a short exact sequence

$$0 \rightarrow \lim^1 H\widehat{B}(M, A, \widehat{D}_A^{\leq k}) [1] \rightarrow H\widehat{B}(M, A, \widehat{D}_A) \rightarrow \lim H\widehat{B}(M, A, \widehat{D}_A^{\leq k}) \rightarrow 0.$$

Such a sequence is natural for maps of towers of chain complexes that satisfy the Mittag-Leffler condition.

The given assumptions imply that $\widehat{D}_A^{\leq k}$ is bounded; given the standing assumption that M is bounded below, we may apply Lemma A.8 to obtain a quasi-isomorphism between $\widehat{B}(M, A, \widehat{D}_A^{\leq k})$ and $\widehat{B}(M, A, \widehat{cB}^{\leq k}(R, A, A'))$; then applying the five lemma to this exact sequence, we see that the induced map on the non-truncated versions is a quasi-isomorphism.

A similar discussion applies to all of the other chain complexes in the statement of the lemma. \blacksquare

We may thus conclude using the same argument as Proposition A.13 that the completed Tate complexes are quasi-isomorphic (with quasi-isomorphism respecting the $\widehat{C}_A^-(R)$ action), under the conditions in the previous lemma.

Corollary A.16. *Suppose A and A' are bounded, non-negatively graded, R -free dg-algebras, for which $f : M \rightarrow M'$ is a quasi-isomorphism of bounded right A' -modules. Then there is a canonical isomorphism $\widehat{H}_A^\infty(M) \cong \widehat{H}_{A'}^\infty(M')$, equivariant under the action of $\widehat{H}_A^-(R) \cong \widehat{H}_{A'}^-(R)$.*

A.4. Periodicity in Tate homology. Tate homology, in some cases, is *periodic*: there is some class in H_A^- so that its action on $H_A^\infty(M)$ is an isomorphism when M is a finite A -module (in an appropriate sense). A beautiful reference for this phenomenon is [GM95, Section III.16] in the setting of genuine G -spectra, where G is a compact Lie group that acts freely on some sphere. Without this assumption, periodicity phenomena often fail; see for instance [BC92].

Here we endeavor only to prove an analogue of it in a simple case which will suffice for our purposes, and in a computational manner. We write $A = \Lambda := \Lambda(u_n)$, where $|u_n| = n$, and either n is odd or the ground ring R has $2 = 0$. We have $H_\Lambda^- = R[[U]]$, where $|U| = -n - 1$, and $H_\Lambda^+ = R[U^*]$ with $|U| = n + 1$, where

the action of U is contraction against U^* ; these are immediately clear from the definition of bar and cobar construction here, which have no nonzero differentials.

We see from the tautological exact triangle that, as a graded R -module, we have $H_\Lambda^\infty = R[[U, U^{-1}]]$ (where here we have suggestively rewritten $U^* = U^{-1}$); we know from the fact that the exact triangle are maps of H_Λ^- -modules that the action of U on these increases the power of U by 1, *except possibly for the action on U^{-1}* . The content of the following crucial lemma is that $U \cdot U^{-1} = 1$.

Lemma A.17. *Let Λ be the exterior algebra over R on an element in degree n , where either n is odd or R has characteristic 2. We have an H_Λ^- -module isomorphism $H_\Lambda^\infty \cong R[[U, U^{-1}]]$, with $U \cdot U^i = U^{i+1}$ for all $i \in \mathbb{Z}$.*

Proof. We write this out very explicitly; all of what follows is a transcription of definitions. First, we may write $D_\Lambda = R[[\eta]] \otimes \Lambda$, with differential

$$d(\eta^k \otimes x) = (-1)^{nk+1} \eta^{k+1} \otimes ux.$$

To avoid confusion later, we have used the notation η where previously we wrote U : it is the degree $(-n-1)$ functional on $B\Lambda$ which sends U^* to $1 \in R$. As graded R -modules we have

$$B(R, \Lambda, D_\Lambda) = R[U^*] \otimes R[[\eta]] \otimes \Lambda(u);$$

under this isomorphism, the differential of $B(D_\Lambda, \Lambda, R)$ is taken to

$$d((U^*)^k \otimes \eta^j \otimes 1) = (-1)^{n(j+k)+1} (U^*)^k \otimes \eta^{j+1} \otimes u + (-1)^k (U^*)^{k-1} \eta^j \otimes u,$$

and is zero on $(U^*)^k \otimes \eta^j \otimes u$. The action of U on this complex is

$$U^i \cdot ((U^*)^k \otimes \eta^j \otimes p) = (-1)^{(n+1)(j+k+1)i} (U^*)^k \otimes \eta^{j+i} \otimes p$$

for any $p \in \Lambda(u)$. Finally, the norm map $N_\Lambda : B(R, \Lambda, D_\Lambda) \rightarrow C_\Lambda^-(R)$ is given by $\eta^k \mapsto U^k$ for any $k \geq 0$ and is otherwise zero.

By definition, the Tate complex

$$C_\Lambda^\infty(R) = B(R, \Lambda, D_\Lambda)[1] \oplus C_\Lambda^-(R)$$

is the mapping cone of the norm map; this means that its differential is

$$d_\infty = \begin{pmatrix} -d & 0 \\ N_\Lambda & d \end{pmatrix}.$$

The action of U^* is the same as before on each component. Now note that

$$H_* B(R, \Lambda, D_\Lambda) \cong H_\Lambda^+[n]$$

as an H_Λ^- -module, and a chain in $B(R, \Lambda, D_\Lambda)$ representing the degree n element of $H_\Lambda^+[n]$ is given by $(U^*)^0 \otimes \eta^0 \otimes u$. Our goal, then, is to show that

$$U \cdot ((U^*)^0 \otimes \eta^0 \otimes u) = (-1)^{n+1} (U^*)^0 \otimes \eta \otimes u$$

is homologous in $C_\Lambda^\infty(R)$ to $U^0 \in C_\Lambda^-(R)$. But from the formula for the differential, we have

$$d_\infty((U^*)^0 \otimes \eta^0 \otimes 1) = (U^*)^0 \otimes \eta \otimes u + U^0;$$

this is precisely what we wanted, so long as n is odd or R has characteristic 2. ■

We can use this to prove the following localization theorem.

Proposition A.18. *Let M be a dg-module over Λ , degreewise free over R ; suppose M has a finite filtration $0 = F_{-1}M \subset F_0M \subset \cdots \subset F_nM = M$ so that each piece $F_kM/F_{k-1}M$ of the associated graded dg-module is quasi-isomorphic to a finite direct sum of copies of R and Λ . Then the action of $U \in H_\Lambda^-$ on $H_\Lambda^\infty(M)$ is an isomorphism. Therefore the natural map $H_\Lambda^-(M) \rightarrow H_\Lambda^\infty(M)$ factors through $H_\Lambda^-(M)[U^{-1}]$, and the map*

$$H_\Lambda^-(M)[U^{-1}] \rightarrow H_\Lambda^\infty(M)$$

is an isomorphism, natural for dg-module homomorphisms of such M .

To prove this, we first define the notion of *inverting an endomorphism* of a dg-module with respect to a dg-module homomorphism $f : M \rightarrow M$ of degree k , written $M[f^{-1}]$. This is **NOT** the strict notion of inversion of an element familiar in module theory,

$$\operatorname{colim} \left(M \xrightarrow{f} M \xrightarrow{f} M \xrightarrow{f} \cdots \right),$$

because this rarely plays well with taking homology. Instead it is a *homotopy colimit*. The simplest way to phrase this (taken from [HR96, Definition 24.5]) is that $M[f^{-1}]$ is the mapping cone of the map $1 - tf : M[t] \rightarrow M[t]$, where t is a polynomial generator in degree $-k$. Then immediate from the exact triangle on homology, and the fact that $(1 - tf)_* : H(M)[t] \rightarrow H(M)[t]$ is injective, we see that $H(M[t^{-1}]) = H(M)[t^{-1}]$, where on the right side we are inverting an element of a module in the usual sense. (Here we are using that the module-theoretic notion may be defined perfectly well as $M[t]/(1 - tf)$.)

Proof of Proposition A.18. First we show that the action of U^* on $H_\Lambda^\infty(M)$ is an isomorphism. This is equivalent to showing that the natural map $C_\Lambda^\infty(M) \rightarrow C_\Lambda^\infty(M)[U^{-1}]$ is a quasi-isomorphism. To see this, we use that the filtration F_kM induces a filtration $F_kC_\Lambda^\infty(M) = C_\Lambda^\infty(F_kM)$, and similarly a filtration on $C_\Lambda^\infty(M)[U^{-1}]$; these are complete because the filtration is finite. Now the E^1 page of the corresponding spectral sequence is given as a direct sum of copies of $H_\Lambda^\infty(R)$ and $H_\Lambda^\infty(\Lambda)$ (or, respectively, the results of inverting U), and the action of U on the E^1 page is the direct sum of the corresponding actions. But $H_\Lambda^\infty(\Lambda) = 0$, and the fact that U^* is an isomorphism on $H_\Lambda^\infty(R)$ was the content of Lemma A.17. So the natural (filtered) map $C_\Lambda^\infty(M) \rightarrow C_\Lambda^\infty(M)[U^{-1}]$ is an isomorphism on the E^1 page, and therefore a quasi-isomorphism. This proves the first part of the theorem.

The rest follows similar lines: there is a natural map

$$C_\Lambda^-(M)[U^{-1}] \rightarrow C_\Lambda^\infty(M)[U^{-1}],$$

and our goal is to show that this map is a quasi-isomorphism; because the map $C_\Lambda^-(M) \rightarrow C_\Lambda^-(M)[U^{-1}]$ is identified on homology with the result of inverting U , the desired result follows. But using the same filtration as the above, the E^1 page of the first spectral sequence is a direct sum of copies of $H_\Lambda^-(R)[U^{-1}]$ and $H_\Lambda^-(\Lambda)[U^{-1}]$, and the E^1 page of the second spectral sequence is a direct sum of corresponding copies of $H_\Lambda^\infty(R)$ and $H_\Lambda^\infty(\Lambda)$. Then the only thing to observe is that what we already know: $H_\Lambda^-(R)[U^{-1}] \rightarrow H_\Lambda^\infty(R)$ is an isomorphism, and that $H_\Lambda^-(\Lambda)$ is a copy of R concentrated in degree zero, so that $H_\Lambda^-(\Lambda)[U^{-1}] = 0$. Therefore the above map is an isomorphism on the E^1 page of the corresponding spectral sequence, and therefore a quasi-isomorphism. \blacksquare

A.5. Simplifying the twisted Borel homology. In this section, we give conditions under which there is a natural isomorphism $H_A^{+, \text{tw}}(M) \cong H_A^+(M)[n]$ for some degree shift n . We would, furthermore, like this isomorphism to preserve the action of $H_A^-(R)$.

The simplest assumption we can make that guarantees this is that A satisfies *Poincaré duality of degree n* , meaning that there is a quasi-isomorphism of right A -modules $A \simeq A^\vee[n]$. Here (as usual) $A^\vee := \text{Hom}_R(A, R)$ with the grading making this into a chain complex, with the right A -action induced by the left A -action on A . This is enough to get the desired R -module isomorphism; some slight additional conditions guarantee the $H_A^-(R)$ -module isomorphism as well.

Remark A.1. This is a very special case of what is called a *Gorenstein condition* in the literature, which usually amounts to something like the condition $\text{Hom}_A(R, A) \cong A[n]$, interpreted in a derived sense. Our condition implies this because of the chain of equivalences

$$\begin{aligned} \text{Hom}_A(R, A) &\simeq \text{Hom}_A(R, \text{Hom}_R(A, R))[n] = \text{Hom}_R(A \otimes_R R, \text{Hom}_A(A, R))[n] \\ &= \text{Hom}_R(A, R)[n] \simeq A. \end{aligned}$$

The second equality uses the tensor-hom adjunction for (bi)modules over two rings, and all appearances of “Hom” are derived. We will not use the Gorenstein condition except through the above Poincaré duality condition; this remark is purely motivational.

A particularly nice reference, which applies to the dg-algebra case, is [DGI06].

Theorem A.19. *If A is R -flat and satisfies Poincaré duality of degree n , then there is a quasi-isomorphism of left A -modules $D_A \simeq R[n]$. Thus, there is a natural isomorphism $H_A^{+, \text{tw}}(M) \cong H_A^+(M)[n]$. If furthermore A is non-negatively graded and $H_0 A \cong R[G]$ for some finite group G , this isomorphism preserves the action of $H_A^-(R)$.*

Proof. We have $D_A = \text{Hom}_A(B(R, A, A), A) \simeq \text{Hom}_A(B(R, A, A), A^\vee)[n]$ by assumption. That $D_A \simeq R[n]$ then follows from the fact that the projection

$$\pi : B(R, A, A) \rightarrow R$$

is an A -equivariant quasi-isomorphism of R -flat modules, and the elementary isomorphism of left A -modules

$$\text{Hom}_R(B(R, A, A), R) \cong \text{Hom}_A(B(R, A, A), A^\vee)$$

given on an element $\eta \in \text{Hom}_R(B(R, A, A), R)$ by $\eta \mapsto \eta'$ with $\eta'(b)(a) = \eta(ba)$.

This equivalence $D_A \simeq \text{Hom}_R(B(R, A, A), R)$ is equivariant under the actions of $C_A^-(R)$ essentially by definition: the product structure uses the left BA -comodule structure of $B(R, A, A)$, and all of the maps above only involved the rightmost factor. We write $B(R, A, A) = EA$ for convenience; this equivalence says $D_A \simeq (EA)^\vee$, this equivalence equivariant under both the left A -action and the left $C_A^-(R) = (BA)^\vee$ -action.

Now we have a map

$$B(M, A, (EA)^\vee) \rightarrow B(M, A, R) \otimes_R B(R, A, (EA)^\vee) = B(M, A, R) \otimes_R EA \otimes_A (EA)^\vee.$$

The pairing $(EA) \otimes (EA)^\vee \rightarrow R$ is A -equivariant, so it factors through $EA \otimes_A (EA)^\vee$, and thus we have a composite $B(M, A, (EA)^\vee) \rightarrow B(M, A, R)$. The claim is that this map is equivariant under the $C_A^-(R) = (BA)^\vee$ -action.

We write this computation down at the chain level, following the sign conventions in [Law]. Write ψ for an element of $(EA)^\vee$ and β for an element of $(BA)^\vee$; then a generic element of $B(M, A, (EA)^\vee) = B(M, A, A) \otimes_A (EA)^\vee$ is written as $m[a_1 | \cdots | a_n] \psi$. Following the given map and then applying β , we have

$$\begin{aligned} \beta \otimes m[a_1 | \cdots | a_n] \psi &\mapsto \beta \otimes \sum_{i=0}^n (-1)^{s_1} m[a_1 | \cdots | a_i] \otimes [a_{i+1} | \cdots | a_n] \psi \\ &\mapsto \beta \otimes \sum_{i=0}^n (-1)^{s_2} m[a_1 | \cdots | a_i] \psi ([a_{i+1} | \cdots | a_n]) \\ &\mapsto \beta \otimes \sum_{0 \leq j \leq i \leq n} (-1)^{s_3} m[a_1 | \cdots | a_j] \otimes [a_{j+1} | \cdots | a_i] \psi ([a_{i+1} | \cdots | a_n]) \\ &\mapsto \sum_{0 \leq j \leq i \leq n} (-1)^{s_4} m[a_1 | \cdots | a_j] \otimes \beta([a_{j+1} | \cdots | a_i]) \psi ([a_{i+1} | \cdots | a_n]). \end{aligned}$$

Here the exponents are

$$\begin{aligned} s_1 &= (n-i)(|m| + \epsilon_i) \\ s_2 &= (n-i)(|m| + \epsilon_i) + |\psi|(\epsilon_n - \epsilon_i) \\ s_3 &= (i-j)(|m| + \epsilon_j) + (n-i)(|m| + \epsilon_i) + |\psi|(\epsilon_n - \epsilon_i) \\ s_4 &= (|\beta| + i-j)(|m| + \epsilon_j) + (n-i)(|m| + \epsilon_i) + |\psi|(\epsilon_n - \epsilon_i). \end{aligned}$$

Here recall that $\epsilon_i := |a_1| + \cdots + |a_i| + i$ is the degree of $[a_1 | \cdots | a_j]$. We have used the rule that the swap map is $a \otimes b \mapsto (-1)^{|a||b|} b \otimes a$.

On the other hand, if we apply β first to ψ and then follow the given map, we have

$$\begin{aligned} \beta \otimes m[a_1 | \cdots | a_n] \psi &\mapsto (-1)^{t_1} m[a_1 | \cdots | a_n] (\beta \cdot \psi) \\ &\mapsto \sum_{0 \leq j \leq n} (-1)^{t_2} m[a_1 | \cdots | a_j] \otimes [a_{j+1} | \cdots | a_n] (\beta \cdot \psi) \\ &\mapsto \sum_{0 \leq j \leq n} (-1)^{t_3} m[a_1 | \cdots | a_j] (\beta \cdot \psi) ([a_{j+1} | \cdots | a_n]) \\ &= \sum_{0 \leq j \leq i \leq n} (-1)^{t_4} m[a_1 | \cdots | a_j] (\beta \otimes \psi) ([a_{j+1} | \cdots | a_i] \otimes [a_{i+1} | \cdots | a_n]) \\ &= \sum_{0 \leq j \leq i \leq n} (-1)^{t_5} m[a_1 | \cdots | a_j] \beta([a_{j+1} | \cdots | a_i]) \psi([a_{i+1} | \cdots | a_n]), \end{aligned}$$

where the exponents are

$$\begin{aligned} t_1 &= |\beta|(m + \epsilon_n) \\ t_2 &= (n-j)(|m| + \epsilon_j) + |\beta|(|m| + \epsilon_n) \\ t_3 &= (n-j)(|m| + \epsilon_j) + |\beta|(|m| + \epsilon_n) + (|\beta| + |\psi|)(\epsilon_n - \epsilon_j) \\ t_4 &= (n-j)(|m| + \epsilon_j) + |\beta|(|m| + \epsilon_n) + (|\beta| + |\psi|)(\epsilon_n - \epsilon_j) + (n-i)(\epsilon_i - \epsilon_j) \\ t_5 &= (n-j)(|m| + \epsilon_j) + |\beta|(|m| + \epsilon_n) + (|\beta| + |\psi|)(\epsilon_n - \epsilon_j) + (|\psi| + n-i)(\epsilon_i - \epsilon_j). \end{aligned}$$

The first and second expressions clearly agree, at least up to the signs on each factor. To check that $(-1)^{t_4} = (-1)^{s_5}$, observe the congruences mod 2

$$\begin{aligned} & (n-j)(|m| + \epsilon_j) + |\beta|(|m| + \epsilon_n) + (|\beta| + |\psi|)(\epsilon_n - \epsilon_j) + (|\psi| + n - i)(\epsilon_i - \epsilon_j) \\ & \equiv (n-j + |\beta|)|m| + |\psi|\epsilon_n + (|\psi| + n - i)\epsilon_i + (|\beta| + i - j)\epsilon_j \\ & \equiv (|\beta| + i - j)(|m| + \epsilon_j) + (n - i)(|m| + \epsilon_i) + |\psi|(\epsilon_n - \epsilon_i). \end{aligned}$$

These are easily seen by breaking the top and bottom formulas (which are the relevant exponents of -1) into the components labelled by $|m|$, ϵ_n , ϵ_i , and ϵ_j in the center formula.

The natural isomorphism $H_A^{+, \text{tw}}(M) \cong H_A^+(M)[n]$ then follows from the invariance theorem for bar constructions, Theorem A.1, and we have seen this is an isomorphism of $H_A^-(R)$ -modules. \blacksquare

The argument doesn't change at all after completion, making the usual assumptions that A is non-negatively graded and bounded above, and it follows that if such an algebra satisfies Poincaré duality of degree n , we have $\hat{H}_A^{+, \text{tw}}(M) \cong \hat{H}_A^+(M)[n]$ as $\hat{H}_A^-(R)$ -modules, so long as M is bounded below.

A.6. Spectral sequences. Making some further mild assumptions on the algebras, we have useful spectral sequences for calculating the various flavors of A -homology.

Proposition A.20. *If A is a non-negatively graded dg-algebra ($A_n = 0$ for $n < 0$), there is a projection of dg-algebras $\pi : A \rightarrow H_0A$, through which we can have A act on H_qM . Then for any $\bullet \in \{+, -, \infty\}$, there is a conditionally convergent spectral sequence of $H_A^-(R)$ -modules*

$$H_p^\bullet(A, H_qM) \rightarrow H_{p+q}^\bullet(A, M).$$

If $H_0A = R$, observe that the action of A on $H(M)$ is trivial; if $H(M)$ is flat over R , we may identify $E_{p,q}^2 \cong H_p^\bullet(A) \otimes H_q(M)$.

The spectral sequence for H^+ is always strongly convergent, the spectral sequence for H^- and H^∞ are strongly convergent if $H_(M)$ is bounded below.*

Proof. Recalling that $B(M, A, R) \cong M \otimes BA$ as graded modules, we filter

$$F_p B(M, A, R)_n = \bigoplus_{i \leq p} (M)_{n-i} \otimes BA_i;$$

that is, $F_p B(M, A, R)$ consists of elements with total BA -degree at most p . The differential on $B(M, A, R)$ can be written (ignoring signs) as three terms:

$$\begin{aligned} d(m[a_1 | \cdots | a_n]) &= ((dm)[a_1 | \cdots | a_n]) \\ &+ (ma_1[a_2 | \cdots | a_n]) \\ &+ (md([a_1 | \cdots | a_n])). \end{aligned}$$

The first and last term clearly preserve the filtration, and the second does because $A_n = 0$ for $n < -1$. The first differential of the associated spectral sequence is the differential on the associated graded complex, which is $d_0 = d_M \otimes 1_{BA}$. We can thus identify the E^1 page with $BA \otimes HM$ (remember BA is degreewise R -flat). The piece that decreases filtration by exactly 1 is d_1 , given by the differential in BA and multiplication by elements of A_0 . Therefore we identify that the E^1 page is given as $B(HM, A, R)$, where a positive degree element of A acts trivially (these

components of the differential decrease the filtration by at least 2) and the action of A_0 factors through $A_0 \rightarrow H_0A$. This filtration is a pleasant example of the filtration on the totalization of a *multicomplex*; see [Boa99, Theorem 11.3]. Because this is a complete exhaustive Hausdorff filtration, it gives rise to a conditionally convergent spectral sequence by Proposition A.4.

In Boardman's language this is a spectral sequence with exiting differentials because $(BA)_i = 0$ for negative i , and strong convergence is given by A.3.

Observe that the action of $C_A^-(R)$, which contracts against BA , is filtered; so this is a spectral sequence of $H_A^-(R)$ -modules.

The proof for the H^- and H^∞ spectral sequences follow mostly as before. As graded modules, we may write

$$cB(R, A, M)_n = \prod_{i \rightarrow -\infty} cB(R, A, R)_i \otimes M_{n-i},$$

where the symbol $\prod_{i \rightarrow -\infty}$ is as in Definition A.4, and filter as usual

$$F_p cB(R, A, M)_n = \prod_{i \leq p} cB(R, A, R)_i \otimes M_{n-i}.$$

(We again need A to be nonnegatively graded for this to be a filtration, and identify the E^2 page as before.) Because this filtration is complete, exhaustive, and Hausdorff, we thus get a conditionally convergent spectral sequence of $H_A^-(R)$ -modules by Proposition A.4, because the action of $C_A^-(R)$ is again filtered. As for $C_A^\infty(M)$, as a graded R -module we have

$$C_A^\infty(M)_n = \prod_{i \rightarrow -\infty} C_A^\infty(R)_i \otimes M_{n-i}.$$

The same filtration applies and is complete, giving a conditionally convergent spectral sequence of $H_A^-(R)$ -modules with the expected E^2 page.

If $H_*(M)$ is not bounded below, then in principle there could be infinitely many nonzero differentials leaving a given point in the spectral sequence for $H_A^-(M)$ or $H_A^\infty(M)$; but if $H_*(M)$ is bounded below, then each $E_{s,t}^p$ stabilizes at some finite p , and thus $RE^\infty = 0$. Because $H_*(M)$ is bounded below, this is a half-plane spectral sequence, and so $W(E) = 0$ as well, so these spectral sequences are strongly convergent. \blacksquare

A.7. Group algebras. The most important application of the A -homology functors is for $A = C_*(G; R)$, where G is a topological group (but most importantly a compact Lie group). The product structure is given as the composite

$$C_*(G) \otimes C_*(G) \xrightarrow{EZ} C_*(G \times G) \xrightarrow{\times} C_*(G),$$

where EZ is the Eilenberg-Zilber map. The unit is $[e] \in C_0(G)$ and the augmentation $C_0(G) \rightarrow R$ is the natural augmentation (add up points). These makes $C_*(G; R)$ into an associative augmented algebra. For the rest of this section, we frequently abuse notation and write G for $C_*(G; R)$ unless there is danger of confusion.

We begin with the following well-known fact to justify our definition of group homology; a more detailed proof may be found in [GM74, Theorem 3.9]. Here EG denotes a left G -space that G acts freely on, and for a right G -space X we denote $X_{hG} = (X \times EG)/G$ for the Borel construction on X .

Lemma A.21. *Let G be a compact Lie group and $M = C_*(X; R)$ for X a right G -space. Then $C_G^+(M) \simeq C_*((X \times EG)/G; R)$. We also have that $C_G^-(M^\vee) = C_G^-(C^{-*}(X))$ is the dual of $C_G^+(M)$, and in particular*

$$H_G^-(C^{-*}(X; R)) \simeq H_G^{-*}(X_{hG}; R).$$

Product structures are compatible with this: the action of $H_G^-(R)$ on $H_G^+(M)$ is that of $H^{-}(BG; R)$ on $H_*(X_{hG}; R)$.*

The first fact follows because $(EG \times X)/G$ can be given as the bar construction $B(G, X)$ in topological spaces, and the natural map

$$B(C_*(X), C_*(G), R) \rightarrow C_*(B(X, G, \text{pt}); R)$$

given levelwise by Eilenberg-Zilber maps is a homotopy-equivalence (with homotopy inverse given by a map built out of Alexander-Whitney maps). The second statement, that $C_A^-(M^\vee)$ is the dual of $C_A^+(M)$, is true for any algebra A . There is a slight subtlety in making sense of this: while M is a right A -module, instead M^\vee is a left A -module. So when we write $C_A^-(M^\vee)$, we mean $\text{Hom}_A(B(R, A, A), M^\vee)$.

Our first goal is to calculate what we need to make Theorem A.12 practically useful. The following is our first goal, which we prove in pieces.

Theorem A.22. *Let G be a compact Lie group G of dimension n , and suppose that the homomorphism $\det \text{Ad} : \pi_0 G \rightarrow \mathbb{Z}/2$ is trivial, where $\text{Ad} : G \rightarrow GL(\mathfrak{g})$ is the adjoint homomorphism. Then $C_*(G; R)$ is a Poincaré duality algebra of degree n , and in particular by Theorem A.19 we have $H_G^{+, \text{tw}}(M) \simeq H_G^+(M)[n]$.*

Lemma A.23. *If G is a compact Lie group, then $D_G \simeq \tilde{C}_*(S^{Ad}) =: \tilde{S}^{Ad}$ as left G -modules, where S^{Ad} is the one-point compactification of the adjoint representation on \mathfrak{g} and the \tilde{C} indicates we take reduced chains.*

Proof. As in [Kle01, Theorem 10.1], Atiyah duality gives a quasi-isomorphism of $C_*(G)$ -bimodules

$$C_*(G) \rightarrow \text{Hom}_R(C_*(G), \tilde{C}_*(S^{Ad})),$$

where the right module structure comes from the left action on $C_*(G)$ and the left module structure from the left G -module structure on \tilde{S}^{Ad} ; we thus have quasi-isomorphisms of left G -modules

$$D_G \rightarrow cB(R, G, \text{Hom}(G, \tilde{S}^{Ad})) \leftarrow cB(R, G, \text{Hom}(G, R)) \otimes \tilde{S}^{Ad}.$$

As in the proof of Theorem A.19, we conclude using the contractibility of $B(R, A, A)$ and the quasi-isomorphism of left A -modules

$$\text{Hom}(B(R, A, A), R) \simeq \text{Hom}_A(B(R, A, A), \text{Hom}(A, R)).$$

■

The above lemma is useful because $H_*(\tilde{S}^{Ad})$ is a copy of R concentrated in degree $d = \dim G$, and the action of $H_0 G = R[\pi_0 G]$ is given by $\det \text{Ad}$. Indeed, writing this G -module as λ , we have

Lemma A.24. *If G is a compact Lie group, then \tilde{S}^{Ad} is related to λ by a zig-zag of quasi-isomorphisms of left G -modules. Hence $C_*(G; R) \simeq C_*(G; R)^\vee \otimes \lambda[d]$.*

Proof. Let B^τ be the G -subcomplex of \tilde{S}^{Ad} given by

$$B_k^\tau = \begin{cases} 0 & k < d \\ \text{Im}(d_{k+1}) \subset C_k(S^{Ad}) & k = d \\ C_k(S^{Ad}) & k > d \end{cases}$$

Clearly B^τ is acyclic, so $\tilde{S}^{Ad} \rightarrow \tilde{S}^{Ad}/B^\tau$ is a quasi-isomorphism of left G -modules. Then there is a natural inclusion $\lambda \hookrightarrow \tilde{S}^{Ad}/B^\tau$ picking out the generator of homology (because $\lambda = Z_d(\tilde{S}^{Ad})/B_d(\tilde{S}^{Ad})$). This is clearly a quasi-isomorphism of left G -modules.

This, combined with the Atiyah duality equivalence of the previous lemma, then gives us an equivalence

$$C_*(G) \simeq C_*(G)^\vee \otimes \tilde{S}^{Ad} \simeq C_*(G)^\vee \otimes \lambda[d].$$

■

With this, we can assemble the last of basic groups H_G^+ , H_G^- , and H_G^∞ .

Corollary A.25. *Let G be a compact Lie group of dimension d , and suppose that $\det \text{Ad}$ is trivial. Then $H_*B(D_G, G, R) = H_{*-d}(BG; R)$. Thus, by Theorem A.12 (4),*

$$H_{G,k}^\infty(pt; R) = \begin{cases} H_{k-d-1}(BG; R) & k > d \\ H^{-k}(BG; R) & k \leq 0 \end{cases}.$$

If $\det \text{Ad}$ was not trivial, instead the result would be $H_{k-d-1}(BG; \lambda)$ in degrees $k > d$.

In particular, the long exact sequence of Theorem A.12 (4) reduces to

$$\cdots \rightarrow H_{*-d}^+(M) \rightarrow H_*^-(M) \rightarrow H_*^\infty(M) \rightarrow \cdots$$

In general, the first term would be replaced by $H_*^+(M \otimes \lambda)$ (the degree shift reflected in the fact that λ is concentrated in degree d).

We can generalize this calculation to that of the equivariant homologies of orbits G/H :

Theorem A.26. *If G is a compact Lie group and H is a closed subgroup, let G/H be the orbit of right cosets of H (which is thus a right G -space). Then there is an H -module λ_{NH} , given by a copy of R in degree $\dim G - \dim H$ with action given by a map $\pi_0 H \rightarrow \{\pm 1\}$ so that we have the following isomorphisms:*

- (1) $H_G^+(G/H) \cong H_H^+(R) \cong H_*(BH; R)$
- (2) $H_G^-(G/H) \cong H_H^-(\lambda_{NH}) \cong H^{-*}(BH; \lambda_{NH})$
- (3) $H_G^\infty(G/H) \cong H_H^\infty(\lambda_{NH})$.

Proof. The map $C_*(G) \rightarrow C_*(G/H)$ is $C_*(H)$ -invariant, inducing a quasi-isomorphism of right G -modules $B(R, H, G) \rightarrow C_*(G/H)$ (that this is a quasi-isomorphism follows from [GM74, Theorem 3.9]). There is further a canonical isomorphism of chain complexes

$$B(B(R, H, G), G, R) \cong B(R, H, B(G, G, R)).$$

(1) then follows because there is an equivalence $B(G, G, R) \rightarrow R$ of left G -modules (induced by the augmentation $\varepsilon : C_*(G) \rightarrow R$). The same argument in general identifies $H_G^+(G \otimes_H M)$ with $H_H^+(M)$.

(2) follows from a combination of this idea and the argument of Lemma A.23: exactly analogous to the argument of [Kle01, Theorem 10.1], there is an Atiyah duality quasi-isomorphism $C_*(G/H) \rightarrow \text{Hom}_{C_*H}(C_*G, \tilde{C}_*S^{NH})$ of right $C_*(G)$ -modules. Then we may show

$$\begin{aligned} \text{Map}_G(B(R, G, G), G/H) &\simeq \text{Map}_G\left(B(R, G, G), \text{Map}_H(G, \tilde{S}^{NH})\right) \\ &\cong \text{Map}_H(B(R, G, G), \tilde{S}^{NH}) \simeq \text{Map}_H(B(H, H, R), \tilde{S}^{NH}) = cB(R, H, \tilde{S}^{NH}). \end{aligned}$$

The second equivalence is given just as in Lemma A.23, via $\eta(b)(g) \mapsto \eta(bg)$. The second-to-last equivalence follows from [GM74, Page 11] and verifying that, in that language, $B(G, G, R)$ is a proper split Kunneth resolution of R as an H -module, which follows similar lines as verifying that $C_*(G)$ is split as an H -module, implicit in the proof of (1). The rest of the statement follows as in Lemma A.24, reducing \tilde{S}^{NH} to a 1-dimensional representation λ_{NH} .

We can combine the previous two parts and the long exact sequence of Theorem A.12 (4) to verify the isomorphism on Tate homology. This follows from the following homotopy commutative diagram, where every vertical arrow is a quasi-isomorphism and horizontal arrow is an appropriate modification of the norm map.

$$\begin{array}{ccc} B(G/H, G, D_G) & \longrightarrow & cB(R, G, G/H) \\ \downarrow & & \parallel \\ B(R, H, D_G) & \longrightarrow & cB(R, G, G/H) \\ \downarrow & & \downarrow \\ B(R, H, D_H \otimes \tilde{S}^{NH}) & \longrightarrow & cB(R, H, \tilde{S}^{NH}) \end{array}$$

The bottom-left vertical map is induced by the map of pairs (recall that cB is contravariant in the algebra and covariant in the module) $(G, G) \rightarrow (H, H \otimes \tilde{S}^{NH})$, the first map inclusion $H \hookrightarrow G$ and the second map a logarithm map

$$C_*(G) \rightarrow C_*(\text{Th}(NH)) \xrightarrow{AW} C_*(H) \otimes \tilde{S}^{NH}.$$

This is the explicit form Klein gives of the Atiyah duality isomorphism. As it is a quasi-isomorphism, the map on bar constructions is a quasi-isomorphism.

Finally, we replace the bottom left term with $B(D_H \otimes \lambda_{NH}, H, *)$ and commute λ_{NH} across with the following lemma to identify the bottom row with the norm map for the H -module λ_{NH} .

The isomorphisms to homology and (twisted) cohomology of BH follow from Lemma A.21, adapted to twisted coefficient systems. \blacksquare

Lemma A.27. *Let ρ be a G -bimodule with action (on both sides) induced by some group homomorphism $f : \pi_0 G \rightarrow R^\times$. Then $B(\rho, G, M) \cong B(G, \rho \otimes M)$.*

Proof. If $\pi : C_*(G; R) \rightarrow H_0 G \rightarrow R$ is the composition of the projection and the map $[g] \mapsto f(g)$, there is an isomorphism $B(R, G, \rho \otimes M) \rightarrow B(\rho, G, M)$ given by

$$\pi(a_1 \cdots a_k)[a_1 \mid \cdots \mid a_k]m \mapsto [a_1 \mid \cdots \mid a_k]m;$$

the inverse uses the inverse in $\pi_0 G$ (giving rise to the antipode in $H_0 G = R[\pi_0 G]$). This is part of a general isomorphism for moving factors across two-sided bar construction over a Hopf algebra, but the simple form here follows because the G -action factors through the action of a central algebra concentrated in degree 0. ■

Observe that everywhere thusfar we have twisted by some orientation character, if we had worked with R a ring of characteristic 2, the action of $\pi_0 G$ on R would in fact be trivial.

We conclude this section with some related results which we use in the main text.

Example A.2. We may apply this to calculate the three cases relevant to us in this text: $G = SO(3)$ and H one of the three subgroups $\{e\}, SO(2)$, and $SO(3)$. In every case H is connected, and so the representation λ_{NH} of $\pi_0 H$ is trivial and is precisely $R[3 - \dim H]$. Thus $H_{SO(3)}^\bullet(SO(3)/H; R) = H_H^\bullet(R)$, with a dimension shift if appropriate.

When $H = \{e\}$, the Tate homology is trivial, and $H_{\{e\}}^+(R) = H_{\{e\}}^-(R) = R$ concentrated in degree zero. (This is just the axiomatic property of Tate homology.)

When $H = SO(2)$, we are left with $H_{SO(2)}^+(R) = H_*(\mathbb{C}\mathbb{P}^\infty)$ and $H_{SO(2)}^-(R) = H^{-*}(\mathbb{C}\mathbb{P}^\infty)$ by Lemma A.21, and the Tate homology is a splicing of these. By Proposition A.18, as modules over $H_{SO(2)}^-(R)$, we may write the Borel homology, coBorel homology, and Tate homology respectively as

$$R[V], R[[V^{-1}]], R[V, V^{-1}],$$

where $|V| = 2$. Then, because $H_{SO(3)}^-(R) \rightarrow H_{SO(2)}^-(R)$ sends

$$p_1 \mapsto V^2 \in H^{-*}(BSO(2); R),$$

we learn the module structure of these over $\langle p_1 \rangle \in H^{-*}(BSO(3); R)$. When $\frac{1}{2} \in R$, we have

$$H^{-*}(BSO(3); R) = H^{-*}(BSU(2); R)R[[p_1]],$$

so this determines the module structure over $H_{SO(3)}^-(R)$.

On the other hand, we have that

$$H^{-*}(BSO(3); \mathbb{Z}/2) = (\mathbb{Z}/2)[[w_2, w_3]]$$

is a power series ring in two generators (where here one takes $H^{-*} = \prod H^{-k}$ as opposed to a direct sum; the latter would give a polynomial ring). This is standard, and proved in [Hat, Proposition 3.12]. A calculation with integral coefficients is given in [BJ82], and in particular shows that p_1 restricts to e^2 on oriented 2-plane bundles, where e is the Euler class.

Proposition A.28. *There is a chain of algebra quasi-isomorphisms*

$$C_*^{\text{gm}}(SO(3); R) \simeq C_*(SO(3); R).$$

Proof. We neglect to mention the coefficients R throughout, as they play no major role.

First we should replace singular chains with something more easily comparable with the degeneracy relations involved in C_*^{gm} . We define the chain complex of smooth singular chains $C_*^{\text{sm}}(M)$ to be the set of smooth maps from $\Delta^n \rightarrow M$; then it has a quotient $C_*^{\text{smd}}(M; R)$ after we carry out the previous identifications under orientation-preserving diffeomorphisms and quotienting by the same degeneracies.

There are, in this situation, two forgetful maps $C_*^{\text{smd}}(M) \leftarrow C_*^{\text{sm}}(M) \rightarrow C_*(M)$. The easiest way to see that all of these maps are quasi-isomorphisms is to prove that all of these theories are in fact homology theories on smooth manifolds M , and that the induced map on $M = \text{pt}$ is an isomorphism on homology. In fact the rightmost chain complexes are identical for $M = \text{pt}$. On the other hand, $C_*^{\text{smd}}(\text{pt}) = R$, a copy of the ground ring concentrated in degree zero. All higher-dimensional chains are degenerate. Because $H_*(\text{pt}) = R_{(0)}$, and the given maps do the obvious things to points (which are cycles generating $H_*(\text{pt})$), they are all homology isomorphisms. A proof that these are homology theories follows similar lines as Theorem 6.7.

It is also possible to show these quasi-isomorphisms extremely explicitly, using e.g. smooth approximation to find cycles in C_*^{sm} homologous to any given cycle in C_* .

So we now need to relate the chain complex of smooth simplices modulo degeneracy, $C_*^{\text{smd}}(M)$, to $C_*^{\text{gm}}(M)$.

The theory interpolating between these is the chain complex of *triangulated geometric chains* on a smooth manifold M . This chain complex $C_*^{\text{tgm}}(M)$ is functorial under smooth maps, and its homology groups define a homology theory (as before).

Precisely, a *triangulated basic chain* on M is a compact smooth oriented manifold with corners P equipped with a smooth triangulation (that is, a homeomorphism $f : |X| \rightarrow P$, where $|X|$ is the realization of a simplicial complex, so that f is a diffeomorphism from each closed simplex onto its image), and a smooth map $\sigma : P \rightarrow M$. Two triangulated basic chains are isomorphic if there is an orientation-preserving diffeomorphism $\varphi : P \rightarrow P'$ and an isomorphism of simplicial complexes $\psi : X \rightarrow X'$ so that $\varphi'f = f'\psi$ and $\sigma'f = f'\sigma$. The triangulated geometric chain complex $C_*^{\text{tgm}}(M)$ is defined following the same procedure as for $C_*^{\text{gm}}(M)$: identify orientation-reversals with their negative, and quotient by the subcomplex of degenerate chains (triangulated basic chains for which the images of $\sigma|_{\Delta^k}$ and $\partial\sigma|_{\Delta^k}$ are both contained in the image of some smooth manifold of smaller dimension than k , resp $k - 1$, for each component simplex Δ^k).

Now observe that there are natural chain maps

$$C_*^{\text{gm}}(M) \leftarrow C_*^{\text{tgm}}(M) \rightarrow C_*^{\text{smd}}(M).$$

The left map is given by forgetting the triangulation and thinking of a smooth manifold with corners as a very special kind of δ -chain, and the right map is given by sending a triangulated n -chain to the sum of its component simplices: $\sigma : P \rightarrow M$ to

$$\sum_{\Delta^n \subset X} \sigma|_{\Delta^n}.$$

As before, because these are homology theories, to show that these are quasi-isomorphisms in general it suffices to check that these maps are isomorphisms on $H_*(\text{pt})$. Each of these chain complexes (because of the nondegeneracy requirements) are simply a copy of R concentrated in degree zero, and the induced map between them is the identity.

Now, if G is a Lie group, $C_*^{\text{smd}}(G)$ is a dg-algebra using the Eilenberg-Zilber product (multiply the chains $\sigma \times \eta$ using the group structure and triangulate the result according to a standard triangulation of $\Delta^k \times \Delta^n$). So is $C_*^{\text{tgm}}(M)$: now if $\sigma : P \rightarrow M$ and $\eta : Q \rightarrow M$ are triangulated basic chains, their product is $\sigma \times \eta : P \times Q \rightarrow M$ equipped with the product triangulation (again, triangulate each component $\Delta^k \times \Delta^n$ in a standard way). With these algebra structures, the

maps

$$C_*^{\text{smd}}(M) \leftarrow C_*^{\text{tgm}}(M) \rightarrow C_*^{\text{gm}}(M)$$

are in fact dg-algebra homomorphisms. It is even easier to see that the map $C_*^{\text{smd}}(M) \leftarrow C_*^{\text{sm}}(M) \rightarrow C_*(M)$ are dg-algebra homomorphisms, giving us the desired zig-zag of dg-algebra equivalences. ■

A.8. Periodic homological algebra. In this section, we set up a version of the machinery before, built to work for complexes graded over $\mathbb{Z}/2N$ which are finite in a suitable sense. In the main text, we will frequently want to use a filtration resembling the index filtration on the Morse-Bott complex of a finite-dimensional compact smooth manifold, but because the index is only defined in $\mathbb{Z}/8$, this doesn't make sense! So if we want to use spectral sequences to check that maps are quasi-isomorphisms, we must do something at least slightly more subtle.

We resolve this by passing instead to the *unrolled complex* \tilde{C}_* of C , defined as $\tilde{C}_k = C_{k \bmod 2N}$, with the obvious differential. A $2N\mathbb{Z}$ -periodic complex is a complex \tilde{C} equipped with a periodicity isomorphism $\varphi : \tilde{C} \rightarrow \tilde{C}[2N]$. From here, we recover the $\mathbb{Z}/2N$ -graded complex by picking an interval $[i, i + 2N - 1] \subset \mathbb{Z}$, and for $k \in [i, i + 2N - 1]$, define

$$C_{k \bmod 2N} = \tilde{C}_k;$$

the differential is defined in the obvious way except on $C_{i \bmod 2N}$, where it is defined as $dx = \varphi(\tilde{d}x)$; this makes sense as $\varphi(\tilde{d}x)$ is in degree $i - 1 + 2N = i + 2N - 1$. The definitions immediately imply that the $\mathbb{Z}/2N$ -graded complex C did not depend on the choice of representative interval, and that $H_k(\tilde{C}) = H_{k \bmod 2N}(C)$.

This implies that if we're trying to show that a map $C \rightarrow C'$ of $\mathbb{Z}/2N$ -graded complexes is a quasi-isomorphism, this is equivalent to showing that the same is true for the map $\tilde{f} : \tilde{C} \rightarrow \tilde{C}'$ of $2N\mathbb{Z}$ -periodic complexes. We may more or less pass freely between these notions. To avoid notational irritation, we ignore the periodicity isomorphism φ — up to isomorphism of $2N\mathbb{Z}$ -periodic complexes, we can take $\tilde{C}_k = \tilde{C}_{k+2N}$ on the nose, and $\varphi = \text{Id}[2N]$.

Instead of attempting to filter a $\mathbb{Z}/2N$ -graded complex, we find that the appropriate notion seems to be the following.

Definition A.9. Let \tilde{C} be a $2N\mathbb{Z}$ -periodic complex. A periodic filtration on \tilde{C} is a filtration $\cdots \subset F_s \tilde{C} \subset F_{s+1} \tilde{C} \subset \cdots$ with

$$\bigcup_s (F_s \tilde{C}) = \tilde{C}, \quad \bigcap_s F_s \tilde{C} = 0,$$

and $F_{s+2N} \tilde{C} = (F_s \tilde{C})[2N]$; that is,

$$F_{s+2N} \tilde{C}_{t+2N} = F_s \tilde{C}_t.$$

So while the filtration on each of the individual abelian groups $F_s C_k$ may stabilize (so $F_N C_k = C_k$ for large N) — as indeed is frequently the case for us — the filtration itself is infinite in both directions (i.e., neither is the complex $F_{-N} \tilde{C}$ equal to zero for any $N \geq 0$, nor is $F_N \tilde{C} = \tilde{C}$ for any N).

We now pass to the spectral sequence $E_{s,t}^r$ of the filtered complex \tilde{C} . The associated graded complex is

$$E_{s,t}^0 \tilde{C} = F_s \tilde{C}_{s+t} / F_{s-1} \tilde{C}_{s+t},$$

and the fact that the filtration is periodic implies that $E_{s,t}^0 = E_{s+2N,t}^0$. This identification is the map induced by the periodicity isomorphism $\tilde{C} \rightarrow \tilde{C}[2N]$, which is a filtered chain map; because it induces an isomorphism on the E^0 page, the same is true for all pages E^r of the spectral sequence: there is a periodicity isomorphism $E_{s,t}^r \rightarrow E_{s+2N,t}^r$ preserving the differentials.

If so desired, we may thus view this filtration as inducing a $(\mathbb{Z}/2N, \mathbb{Z})$ -bigraded spectral sequence $E_{[s],t}^r$, with $[s] \in \mathbb{Z}/2N$. This may be pictured as a cylinder, where differentials wrap around.

If we have a map $f : \tilde{C} \rightarrow \tilde{C}'$ of unrolled complexes, compatible with a periodic filtration of each, it induces a map of the spectral sequences $E_{s,t}^r$. We would like to know when we can check that the map f is a quasi-isomorphism from the corresponding fact about the E^2 page of this spectral sequence. Because the ‘unrolled’ spectral sequence $E_{s,t}^r$ (considered with (\mathbb{Z}, \mathbb{Z}) -bigrading, $8\mathbb{Z}$ -periodic in the first grading) is a whole-plane spectral sequence, it is difficult to prove and often false that the E^∞ page actually *calculates* the homology groups. However, as long as the periodic filtration is *complete*, we at least still know from Theorem A.4 that we may detect quasi-isomorphisms from isomorphisms on the E^2 page; further, if the spectral sequence degenerates on some finite page, we may indeed calculate the associated graded homology groups from the E^∞ page.

Now let A be a bounded, non-negatively graded dg-algebra satisfying Poincaré duality of degree n , as in Section A.5. Let M be a right dg- A -module, graded over $\mathbb{Z}/2N$, equipped with a complete periodic filtration; suppose that the associated graded A -module $F_p \tilde{M} / F_{p-1} \tilde{M}$ is bounded.

We would like to apply the constructions of the previous sections to construct ‘equivariant homology’ complexes $C_A^\bullet(M)$, with corresponding periodic filtrations. The clear thing to try is to pass to the unfolded complex \tilde{M} and consider the corresponding complexes $C_A^\bullet(\tilde{M})$. However, because \tilde{M} is unbounded in both directions, these filtrations are not complete; BA has elements in arbitrarily large degrees, which by pairing with elements of \tilde{M} in arbitrarily low filtration may contribute a sequence of nonzero elements of $C_A^\bullet(\tilde{M})$ of arbitrarily low filtration but fixed degree. The spectral sequence will not converge to its homology, but to its completion’s.

For a right A -module M with periodic filtration as above, we thus write

$$\begin{aligned} C_A^+(M) &:= \widehat{B}(\tilde{M}, A, R). \\ C_A^-(M) &:= \widehat{cB}(A, \tilde{M}) = \text{Hom}_A^{\text{fin}}(B(R, A, A), \tilde{M}) \\ C_A^\infty(M) &:= \text{Cone}\left(\widehat{B}(\tilde{M}, A, \hat{D}_A) \rightarrow \widehat{cB}(A, \tilde{M})\right). \end{aligned}$$

The filtrations $F_s C_A^+(M) = \widehat{C}_A^\bullet(F_s \tilde{M})$ are now complete, exhaustive, and Hausdorff; these homology groups are what the spectral sequences induced by the filtration on $C_A^\bullet(\tilde{M})$ actually converge to.

We may assemble this into the following homology theories for periodically graded A -modules.

Theorem A.29. *Let A be a non-negatively graded, bounded above, and R -free dg-algebra satisfying Poincaré duality of degree n . If M is a $\mathbb{Z}/2N$ -graded right A -module, equipped with a periodic filtration whose associated graded complex $\text{gr}_p \tilde{M}$*

is bounded for each p , there are $\mathbb{Z}/2N$ -graded complexes $C_A^\bullet(M)$, for $\bullet \in \{+, -, \infty\}$, satisfying the following properties.

- (1) Each $C_A^\bullet(M)$ is equipped with a left action of $\hat{C}_A^-(R)$;
- (2) The $C_A^\bullet(M)$ are functorial under A -module maps $f : M \rightarrow M'$ of $\mathbb{Z}/2N$ graded complexes, which are filtered in the sense that $\tilde{f}(F_s \tilde{M}) \subset F_{s+d} \tilde{M}'$ for all s and a fixed constant d ;

- (3) There is a strongly convergent spectral sequence of $H(A)$ -modules,

$$\mathrm{gr}_p(H_q(M)) \rightarrow H_{p+q}(M),$$

which degenerates at some finite page;

- (4) There is a conditionally convergent spectral sequence of $H_A^-(R)$ -modules

$$H_A^\bullet(\mathrm{gr}(\tilde{M})) \rightarrow H_A^\bullet(M),$$

which is strongly convergent for $H_A^-(M)$.

- (5) A filtered A -module map $M \rightarrow M'$ which is a quasi-isomorphism on the associated graded complexes induces a quasi-isomorphism $C_A^\bullet(M) \rightarrow C_A^\bullet(M')$;
- (6) If the associated graded is free over A , in the sense that $\mathrm{gr}_p \tilde{M} \cong X_p \otimes A$ for each p and some R -free complex X_p , then $H_A^\infty(M) = 0$;
- (7) There is a long exact sequence of $H_A^-(R)$ -modules

$$\dots \rightarrow H_A^+(M) \xrightarrow{[n]} H_A^-(M) \rightarrow H_A^\infty(M) \xrightarrow{[-n-1]} H_A^+(M) \rightarrow \dots$$

natural under filtered A -module maps.

Proof. The actions of $\hat{C}_A^-(R)$ on $C_A^\bullet(M)$ were detailed in the sections the completed complexes were introduced. An A -module map $f : M \rightarrow M'$ induces an A -module map $\tilde{f} : \tilde{M} \rightarrow \tilde{M}'$, and therefore a map $C_A^\bullet(\tilde{M}) \rightarrow C_A^\bullet(\tilde{M}')$, so it induces a map on their completions — which are what we define to be the unwrapped, \mathbb{Z} -graded complexes of $C_A^\bullet(M)$ and $C_A^\bullet(M')$, when M, M' are $\mathbb{Z}/2N$ -graded.

By assumption, each $\mathrm{gr}_p(M_k) \cong \mathrm{gr}_{p+2N}(M_{k+2N})$ is bounded in k , which has the same bound for $k \equiv k' \pmod{2N}$. Because there are only finitely many equivalence classes modulo $2N$, there is a uniform bound in k on $\mathrm{gr}_p(M_{p+k})$, independent of p . In particular, the filtration is finite in each degree, and the associated spectral sequence is contained in some finite strip; if the strip has width ℓ , then the differentials on any $E_{s,t}^r$ must leave the strip if $r > \ell + 1$, and so the spectral sequence degenerates at $E^{\ell+2}$. This gives (3).

The existence of the spectral sequence was outlined above. If we also know that $\hat{H}_A^\bullet(\mathrm{gr}_p \tilde{M})$ takes quasi-isomorphisms of A -modules to isomorphisms of $\hat{H}_A^-(R)$ -modules, we may conclude item (5), as we have detected an isomorphism on the E^1 page. That this is true follows because A is R -free, non-negatively graded, and bounded above, while each $\mathrm{gr}_p \tilde{M}$ is bounded: for \hat{H}^+ this is Lemma A.8, for \hat{H}^- this is Lemma A.9, and for \hat{H}^∞ this is Corollary A.16.

The vanishing theorem A.12 (4) implies that if $\mathrm{gr} \tilde{M} \cong X \otimes A$, we have

$$H_A^\infty(\mathrm{gr} \tilde{M}) = 0,$$

and so the spectral sequence $H_A^\infty(\mathrm{gr} \tilde{M}) \rightarrow H_A^\infty(M)$ collapses at the E^1 page, giving $H_A^\infty(M) = 0$.

The long exact sequence is clear from the existence of a short exact sequence

$$0 \rightarrow \widehat{B}(\widetilde{M}, A, \widehat{D}_A) \rightarrow \text{Cyl}(\widehat{N}_{\widetilde{M}}) \rightarrow C_A^\infty(M) \rightarrow 0,$$

where $\text{Cyl}(\widehat{N}_{\widetilde{M}})$ is the mapping cylinder of the norm map, and so is naturally equivalent as $C_A^-(R)$ -modules to $C_A^-(M)$. In particular, by the completed analogue of Theorem A.19 (which is proved in an essentially identical manner), we have that the homology of the first term is naturally equivalent to $H_A^+(M)[n]$ as an $H_A^-(R)$ -module. The discussion there guarantees this for M bounded below, but if $M' \subset \widetilde{M}$ is some bounded below A -submodule which equals M in all sufficiently large degrees, one may identify $\widehat{H}_A^{+, \text{tw}}(M) = \lim_{n \rightarrow -\infty} \widehat{H}_A^{+, \text{tw}}(M'[-2N])$. ■

REFERENCES

- [AB96] David M. Austin and Peter J. Braam. Equivariant Floer theory and gluing Donaldson polynomials. *Topology*, 35(1):167–200, 1996.
- [APS75a] M.F. Atiyah, V.K. Patodi, and I.M. Singer. Spectral asymmetry and Riemannian geometry. I. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 77, pages 43–69. Cambridge University Press, 1975.
- [APS75b] M.F. Atiyah, V.K. Patodi, and I.M. Singer. Spectral asymmetry and Riemannian geometry. II. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 78, pages 405–432. Cambridge University Press, 1975.
- [APS96] M.F. Atiyah, V.K. Patodi, and I.M. Singer. Spectral asymmetry and Riemannian geometry.: III. In *Collected Papers Of V.K. Patodi*, pages 189–217. World Scientific, 1996.
- [Aus90] David M. Austin. $SO(3)$ -instantons on $L(p, q) \times \mathbf{R}$. *Journal of Differential Geometry*, 32(2):383–413, 1990.
- [Aus95] David M. Austin. Equivariant Floer groups for binary polyhedral spaces. *Mathematische Annalen*, 302(1):295–322, 1995.
- [BC92] David J Benson and Jon F Carlson. Products in negative cohomology. *Journal of pure and applied algebra*, 82(2):107–129, 1992.
- [BJ82] Edgar H Brown Jr. The cohomology of BSO_n and BO_n with integer coefficients. *Proceedings of the American Mathematical Society*, pages 283–288, 1982.
- [BMR14] Tobias Barthel, J.P. May, and Emily Riehl. Six model structures for DG-modules over DGAs: model category theory in homological action. *New York J. Math*, 20:1077–1159, 2014.
- [Boa99] J. Michael Boardman. Conditionally convergent spectral sequences. In *Homotopy Invariant Algebraic Structures*, volume 239 of *Contemporary Mathematics*, pages 49–84. American Mathematical Society, Baltimore, MD, 1999.
- [Bou08] N. Bourbaki. *Lie groups and Lie algebras: Chapters 1-3*. Bourbaki, Nicolas: Elements of mathematics. Springer Berlin Heidelberg, 2008.
- [Cra04] Marius Crainic. On the perturbation lemma, and deformations. *arXiv preprint math/0403266*, 2004.
- [Dae15] Aliakbar Daemi. Abelian gauge theory, knots and odd Khovanov homology. *ArXiv e-prints*, August 2015.
- [DGI06] W.G. Dwyer, J.P.C. Greenlees, and S. Iyengar. Duality in algebra and topology. *Advances in Mathematics*, 200(2):357–402, 2006.
- [DK90] S.K. Donaldson and P.B. Kronheimer. *The Geometry of four-manifolds*. Oxford mathematical monographs. Clarendon Press, 1990.
- [Don87] S. K. Donaldson. The orientation of Yang-Mills moduli spaces and 4-manifold topology. *Journal of Differential Geometry*, 26(3):397–428, 1987.
- [Don02] S.K. Donaldson. *Floer homology groups in Yang-Mills theory*, volume 147. Cambridge University Press, 2002.
- [Flo88] Andreas Floer. An instanton-invariant for 3-manifolds. *Comm. Math. Phys.*, 118(2):215–240, 1988.
- [Flo91] Andreas Floer. *Instanton homology, surgery and knots*, volume 1 of *London Mathematical Society Lecture Note Series*, page 97–114. Cambridge University Press, 1991.

- [Flo95] Andreas Floer. Instanton homology and Dehn surgery. In *The Floer memorial volume*, pages 77–97. Springer, 1995.
- [Frø95] Kim A. Frøyshov. *On Floer Homology and Four-Manifolds with Boundary*. PhD thesis, Balliol College, University of Oxford, 1995.
- [Frø02] Kim A. Frøyshov. Equivariant aspects of Yang–Mills Floer theory. *Topology*, 41(3):525 – 552, 2002.
- [Frø04] Kim A. Frøyshov. An inequality for the h-invariant in instanton Floer theory. *Topology*, 43(2):407–432, 2004.
- [GM74] V.K.A.M. Gugenheim and J.P. May. *On the theory and applications of differential torsion products*. Number no. 142 in Memoirs of the American Mathematical Society. American Mathematical Society, 1974.
- [GM95] J.P.C. Greenlees and J.P. May. *Generalized Tate Cohomology*. Number no. 543 in American Mathematical Society: Memoirs of the American Mathematical Society. American Mathematical Society, 1995.
- [Gre01] J. P. C. Greenlees. *Tate cohomology in axiomatic stable homotopy theory*, pages 149–176. Birkhäuser Basel, Basel, 2001.
- [Hat] Allen Hatcher. Vector Bundles and K-Theory. <https://pi.math.cornell.edu/hatcher/VBKT/VBpage.html>. Accessed: 2018-11-21.
- [HK11] Matthew Hedden and Paul Kirk. Chern-Simons invariants, $SO(3)$ instantons, and $\mathbb{Z}/2$ -homology cobordism. In J.E. Andersen, H. Boden, A. Hahn, and B. Himpel, editors, *Chern-Simons Gauge Theory: 20 Years After*, pages 83–114. American Mathematical Society, 2011.
- [HLS16] Kristen Hendricks, Robert Lipshitz, and Sucharit Sarkar. A simplicial construction of G-equivariant Floer homology. *arXiv preprint arXiv:1609.09132*, 2016.
- [HR96] Bruce Hughes and Andrew Ranicki. *Ends of complexes*. Number 123. Cambridge university press, 1996.
- [Jon87] John D.S. Jones. Cyclic homology and equivariant homology. *Inventiones mathematicae*, 87(2):403–423, 1987.
- [Kle01] John R. Klein. The dualizing spectrum of a topological group. *Mathematische Annalen*, 319(3):421–456, Mar 2001.
- [Kle02] John R. Klein. Axioms for generalized Farrell–Tate cohomology. *Journal of Pure and Applied Algebra*, 172(2):225 – 238, 2002.
- [KM95] Igor Kriz and J Peter May. *Operads, algebras, modules and motives*. Number 233 in Astérisque. Société mathématique de France, 1995.
- [KM07] Peter Kronheimer and Tomasz Mrowka. *Monopoles and Three-Manifolds*. New Mathematical Monographs. Cambridge University Press, 2007.
- [KM11a] P. B. Kronheimer and T. S Mrowka. Khovanov homology is an unknot-detector. *Publications mathématiques de l’IHÉS*, 113(1):97–208, 2011.
- [KM11b] P. B. Kronheimer and T. S. Mrowka. Knot homology groups from instantons. *Journal of Topology*, 4(4):835–918, 2011.
- [Kro05] P. B. Kronheimer. Four-manifold invariants from higher-rank bundles. *Journal of Differential Geometry*, 70(1):59–112, 2005.
- [Law] Tyler Lawson. In which I try to get the signs right for once. <http://www-users.math.umn.edu/tlawson/papers/signs.pdf>. Accessed: 2018-11-09.
- [Lin18] Francesco Lin. A Morse-Bott approach to monopole Floer homology and the Triangulation conjecture. *Memoirs of the American Mathematical Society*, 255(1221):162pp, 2018.
- [Lip14] Max Lipyanskiy. Geometric homology. *ArXiv e-prints*, September 2014.
- [Man16] Ciprian Manolescu. Pin(2)-equivariant Seiberg-Witten Floer homology and the triangulation conjecture. *Journal of the American Mathematical Society*, 29(1):147–176, 2016.
- [McC06] J.E. McClure. On the chain-level intersection pairing for PL manifolds. *Geometry & Topology*, 10(3):1391–1424, 2006.
- [MMR94] 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. International Press, 1994.
- [NS17] Thomas Nikolaus and Peter Scholze. On topological cyclic homology. *arXiv preprint arXiv:1707.01799*, 2017.
- [OS04] Peter Ozsváth and Zoltán Szabó. Holomorphic disks and topological invariants for closed three-manifolds. *Annals of Mathematics*, 159(3):1027–1158, 2004.

- [Qui85] Daniel Quillen. Determinants of Cauchy-Riemann operators over a Riemann surface. *Functional Analysis and Its Applications*, 19(1):31–34, 1985.
- [Sal99] Dietmar Salamon. *Spin geometry and Seiberg-Witten invariants*. 1999. Online; accessed Sept 16, 2018 at <https://people.math.ethz.ch/~salamon/PREPRINTS/witsei.pdf>.
- [Sas13] Hirofumi Sasahira. Instanton Floer homology for lens spaces. *Mathematische Zeitschrift*, 273(1-2):237–281, 2013.
- [Sca15] Christopher W Scaduto. Instantons and odd khovanov homology. *Journal of Topology*, 8(3):744–810, 2015.
- [Sch93] Matthias Schwarz. *Morse homology*. Progress in Mathematics. Birkhäuser, 1993.
- [SS10] Paul Seidel and Ivan Smith. Localization for involutions in Floer cohomology. *Geometric and Functional Analysis*, 20(6):1464–1501, 2010.
- [SS17] Christopher Scaduto and Matthew Stoffregen. Klein-four connections and the Casson invariant for nontrivial admissible $U(2)$ bundles. *Algebr. Geom. Topol.*, 17(5):2841–2861, 2017.
- [Tau84] Clifford Henry Taubes. Self-dual connections on 4-manifolds with indefinite intersection matrix. *Journal of Differential Geometry*, 19(2):517–560, 1984.
- [Tho57] René Thom. L’homologie des espaces fonctionnels. In *Colloque de topologie algébrique, Louvain*, pages 29–39, 1957.
- [Wal61] C.T.C. Wall. Resolutions for extensions of groups. *Mathematical Proceedings of the Cambridge Philosophical Society*, 57(2):251–255, 1961.
- [Wei95] C.A. Weibel. *An introduction to homological algebra*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1995.