

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Seiberg-Witten and Gromov invariants for self-dual harmonic 2-forms

Permalink

<https://escholarship.org/uc/item/1730608x>

Author

Gerig, Chris

Publication Date

2018

Peer reviewed|Thesis/dissertation

Seiberg-Witten and Gromov invariants for self-dual harmonic 2-forms

by

Christopher A. Gerig

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Michael Hutchings, Chair

Professor Denis Auroux

Professor Petr Hořava

Spring 2018

Seiberg-Witten and Gromov invariants for self-dual harmonic 2-forms

Copyright 2018
by
Christopher A. Gerig

Abstract

Seiberg-Witten and Gromov invariants for self-dual harmonic 2-forms

by

Christopher A. Gerig

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Michael Hutchings, Chair

For a closed oriented smooth 4-manifold X with $b_+^2(X) > 0$, the Seiberg-Witten invariants are well-defined. Taubes’ “SW=Gr” theorem asserts that if X carries a symplectic form then these invariants are equal to well-defined counts of pseudoholomorphic curves, Taubes’ Gromov invariants. In the absence of a symplectic form there are still nontrivial closed self-dual 2-forms which vanish along a disjoint union of circles and are symplectic elsewhere. This thesis describes well-defined counts of pseudoholomorphic curves in the complement of the zero set of such near-symplectic 2-forms, and it is shown that they recover the Seiberg-Witten invariants over $\mathbb{Z}/2\mathbb{Z}$. This is an extension of Taubes’ “SW=Gr” theorem to non-symplectic 4-manifolds.

The main results are the following. Given a suitable near-symplectic form ω and tubular neighborhood \mathcal{N} of its zero set, there are well-defined counts of pseudoholomorphic curves in a completion of the symplectic cobordism $(X - \mathcal{N}, \omega)$ which are asymptotic to certain Reeb orbits on the ends. They can be packaged together to form “near-symplectic” Gromov invariants as a map $\text{Spin}^c(X) \rightarrow \Lambda^* H^1(X; \mathbb{Z})$. They are furthermore equal to the Seiberg-Witten invariants with $\mathbb{Z}/2\mathbb{Z}$ coefficients, where ω determines the “chamber” for defining the latter invariants when $b_+^2(X) = 1$.

In the final chapter, as a non sequitur, a new proof of the Fredholm index formula for punctured pseudoholomorphic curves is sketched. This generalizes Taubes’ proof of the Riemann-Roch theorem for compact Riemann surfaces.

Contents

Contents	i
1 Taming the pseudoholomorphic beasts in $\mathbb{R} \times (S^1 \times S^2)$	1
1.1 Introduction	1
1.2 Review of pseudoholomorphic curve theory	8
1.3 Towards a near-symplectic Gromov invariant	14
1.4 Appendix	38
2 Seiberg-Witten and Gromov invariants	40
2.1 Introduction	40
2.2 S^1 -valued Morse theory	44
2.3 Review of gauge theory	46
2.4 Review of Taubes' isomorphisms	56
2.5 The blow-up formula	61
2.6 Equating ECH and HM cobordism counts	64
2.7 Relation of SW counts	75
2.8 Appendix – some explanation to Taubes' analysis	82
3 “Riemann-Roch” for punctured curves	83
3.1 Setup	83
3.2 L -flat approximation	85
3.3 Spectral flow for the perturbed asymptotic operator	85
3.4 Localization argument	88
3.5 Appendix	88
Bibliography	90

Acknowledgements

I am beyond grateful for knowing and studying under my amazing advisor, Michael Hutchings. Likewise, I am also indebted to Clifford Taubes, who among other things told me to meet Michael upon entering UC Berkeley. Let me elaborate. At the end of an undergraduate physics program at Cornell I was heading into a physics PhD program but hoping to switch into pure math, and I subsequently found out that Cliff had the same career trajectory. After I met him and got the aforementioned recommendation, I ditched the “open house” of UC Berkeley’s physics PhD program to go to Michael’s office. This was one year before entering the math program, which I imagine only worked out because Michael agreed to be my advisor from the beginning.

Michael’s and Cliff’s guidance and ideas were valuable, crucial, and appreciated. I thank Michael for the bluntness when it comes to advising me on mathematics, for his “hands off” style of guidance, for not allowing vague questions or hand-wavy answers in his office, for not minding during the times that I was lazy or had prioritized sports, and for the unlimited funding to travel to conferences. I also thank Cliff for the mathematical chats via phone calls and for quoting Feynman in order to give me advice.

I thank Dan Cristofaro-Gardiner and Chris Wendl, both for collaborations and for keeping me “in check” to become a better mathematician on multiple levels. I thank Denis Auroux, Rob Kirby, and Katrin Wehrheim for helpful conversations, advice, and support, both mathematically and personally. I thank Simon Donaldson not only for a recommendation letter, but for teaching me how to punt (on one of those canoes/boats) over in Cambridge, UK. I thank Tomasz Mrowka for helping me to understand certain aspects of monopole Floer homology... and for the inevitable future conversations about them.

I’m also glad to have had interactions with Keon Choi and Vinicius Ramos when I first entered the math program, and with Mihai Munteanu towards the end (all students of Michael). I’ve further enjoyed learning from Robert Bryant, Yasha Eliashberg, Dusa McDuff, Tim Nguyen, John Pardon, Tim Perutz, Wolfgang Schmaltz, Steven Sivek, Qiaochu Yuan, and many others who hopefully don’t mind that I’m not listing them.

While working on this thesis I was partially supported by NSF grants DMS-1406312, DMS-1344991, and DMS-0943745.

Chapter 1

Taming the pseudoholomorphic beasts in $\mathbb{R} \times (S^1 \times S^2)$

1.1 Introduction

Motivation

The Seiberg-Witten invariants, introduced by Witten [80], are defined for any closed oriented 4-manifold X with $b_+^2(X) \geq 1$. These invariants $SW_X(\mathfrak{s})$ are constructed by counting solutions to the Seiberg-Witten equations associated with a given spin-c structure $\mathfrak{s} \in \text{Spin}^c(X)$. When $b_+^2(X) > 1$ the invariants only depend on the smooth structure of X and a choice of “homology orientation” of X . When $b_+^2(X) = 1$ the invariants also depend on a choice of “chamber.”

The Gromov invariants, introduced by Taubes [54] who generalized the work of Gromov [13] and Ruan [47], are defined for any closed symplectic 4-manifold (X, ω) . These invariants $Gr_{X, \omega}(A)$ are constructed by counting pseudoholomorphic curves in X which represent a given homology class $A \in H_2(X; \mathbb{Z})$. The invariants only depend on the smooth structure of X .

Now, the existence of a symplectic structure implies that $b_+^2(X) > 0$, and when $b_+^2(X) = 1$ a given symplectic form determines a canonical “chamber” with which to define the Seiberg-Witten invariants of X . The symplectic form also determines a canonical “homology orientation” of X and it induces a canonical identification of $H_2(X; \mathbb{Z})$ with $\text{Spin}^c(X)$. Under these identifications it was shown by Taubes [63] that the Gromov invariants and the Seiberg-Witten invariants are equivalent. But what occurs when X does not have a symplectic structure? This was originally asked by Taubes during 1995 in a paper [71] from which we quote the ending:

“It is observed that there are manifolds with nonzero Seiberg-Witten invariants which do not admit symplectic forms. With this understood, one is led to ask whether there is any sort of “Gromov invariant” interpretation for the Seiberg-Witten invariants in the nonsymplectic world.”

In the absence of a symplectic form but keeping $b_+^2(X) > 0$, given a generic Riemannian metric on X there are still nontrivial closed self-dual 2-forms which vanish transversally along a disjoint union of circles in X and are symplectic elsewhere. These were studied intensively in unpublished work of Karl Luttinger during the early 1990's. Such a 2-form ω is called a *near-symplectic form*, and $X - \omega^{-1}(0)$ is a noncompact symplectic 4-manifold which has an almost complex structure J determined by ω and the Riemannian metric. It was noticed by Taubes that the estimates on symplectic 4-manifolds which were used to find J -holomorphic curves from given Seiberg-Witten solutions could also be used to find J -holomorphic curves when $\omega^{-1}(0) \neq \emptyset$.

Theorem 1.1.1 (Taubes [64]). *If X has a nonzero Seiberg-Witten invariant then there exists at least one J -holomorphic curve in $X - \omega^{-1}(0)$, homologically bounding $\omega^{-1}(0)$ in the sense that it has intersection number 1 with every linking 2-sphere of $\omega^{-1}(0)$.*

This provides evidence that there might be “Gromov invariants” for non-symplectic 4-manifolds, and that they might recover the Seiberg-Witten invariants. An elaboration can be found in [70, 69]. But there are difficulties in constructing well-defined counts of such curves because J becomes singular along $\omega^{-1}(0)$. That is, a suitable Fredholm theory for the moduli spaces of curves in $X - \omega^{-1}(0)$ is hard to establish (see [73, 60]). We can rephrase the problem (see [65]), for which it is standard to complete $X - \omega^{-1}(0)$ by attaching symplectization ends and then the J -holomorphic curves satisfy certain asymptotic conditions. The boundary Y of a tubular neighborhood of $\omega^{-1}(0)$ is a contact manifold, with contact form induced by ω , and the J -holomorphic curves are asymptotic to periodic orbits of the Reeb flow induced by the contact form on Y . Models of these curves in the symplectization of the contact manifold with $Y = S^1 \times S^2$ are studied in [51, 61, 62, 38], and the aforementioned difficulties are now caused by the contact dynamics: the existence of certain Reeb orbits permits the existence of non-transverse multiply covered curves in the relevant moduli spaces (see Remark 1.3.10).

As will be explained momentarily, we overcome the transversality difficulties by modifying the chosen neighborhood of $\omega^{-1}(0)$ and hence the contact dynamics on each component $S^1 \times S^2$ of Y . We then pick out the appropriate J -holomorphic curves to count using embedded contact homology (ECH), a Floer theory constructed by Hutchings [28]. ECH was originally motivated by the desire to find an analog of Taubes' equivalence $SW_X = Gr_{X,\omega}$ in dimension three, granted that a version of Seiberg-Witten Floer homology [33] existed on the gauge theory side. These two homologies are now known to be isomorphic by Taubes [55], using the same techniques that were used to prove $SW_X = Gr_{X,\omega}$. All of this machinery plays a crucial role in this thesis.

Near-symplectic geometry

Throughout this thesis, (X, g) denotes a closed connected oriented smooth Riemannian 4-manifold with $b_+^2(X) \geq 1$, where $b_+^2(X)$ denotes the dimension of any maximal positive-definite subspace of $H^2(X; \mathbb{R})$ under the intersection form on X . In particular, if X is

simply connected then the only spaces excluded by this assumption on $b_+^2(X)$ are 4-spheres and their blow-ups $\overline{\mathbb{CP}}^2 \# \cdots \# \overline{\mathbb{CP}}^2$. Let $\omega \in \Omega^2(X; \mathbb{R})$ be a nontrivial closed self-dual (hence harmonic) 2-form. These always exist by Hodge theory, and in fact, the set of such 2-forms determines a subspace $H_+^2(X; \mathbb{R}) \subset H^2(X; \mathbb{R})$ for which $b_+^2(X) = \dim H_+^2(X; \mathbb{R})$. A nice property of ω is that the complement of its zero set $Z := \omega^{-1}(0)$ is symplectic,

$$\omega \wedge \omega = \omega \wedge * \omega = |\omega|^2 d\text{vol}_g$$

but we cannot always expect $Z = \emptyset$, i.e. for X to be symplectic. For starters, if a symplectic form existed then an almost complex structure could be built from it and the metric, forcing $1 - b^1(X) + b_+^2(X)$ to have even parity by characteristic class theory. What we do know in general is that Z cannot contain any open subset of X [8, Corollary 4.3.23].

For generic metrics, Honda [16, Theorem 1.1] and LeBrun [36, Proposition 1] have shown that there exist nontrivial closed self-dual 2-forms that vanish transversally as sections of the self-dual 3-plane subbundle $\Lambda_+^2 T^* X$ of $\Lambda^2 T^* X \rightarrow X$. Such 2-forms are examples of

Definition 1.1.2. A closed 2-form $\omega : X \rightarrow \Lambda^2 T^* X$ is *near-symplectic* if for all points $x \in X$ either $\omega^2(x) > 0$, or $\omega(x) = 0$ and the rank of the gradient $\nabla \omega_x : T_x X \rightarrow \Lambda^2 T_x^* X$ is three.

It follows from [1, Proposition 1] that we can always find a metric for which a given near-symplectic form is self-dual, and we assume throughout this thesis that such metrics have been chosen. It follows from the definition that Z consists of a finite disjoint union of smooth embedded circles [46, Lemma 1.2], and Z is null-homologous while the individual zero-circles need not be [45, Proposition 1.1.26]. Moreover, the zero-circles are not all the same but come in two “types” depending on the behavior of ω near them:

Let N_Z denote the normal bundle of $Z \subset X$, identified with the orthogonal complement TZ^\perp such that $TX = TZ \oplus N_Z$. The gradient $\nabla \omega$ defines a vector bundle isomorphism $N_Z \rightarrow \Lambda_+^2 T^* X|_Z$, and we orient N_Z such that $\nabla \omega$ is orientation-reversing. The orientation of TX orients $\Lambda_+^2 T^* X$, so TZ is subsequently oriented. As described in [73, 46, 15], ω determines a particular subbundle decomposition

$$N_Z = L_Z \oplus L_Z^\perp$$

where $L_Z \rightarrow Z$ is a rank one line bundle. Explicitly, L_Z is the negative-definite subspace with respect to the induced quadratic form on N_Z ,

$$v \mapsto \langle \nabla_v \omega(\partial_0, \cdot), v \rangle$$

where $\partial_0 \in \Gamma(TZ)$ is the unit-length oriented vector field. A zero-circle of Z is called *untwisted* if L_Z restricted to that zero-circle is orientable, and *twisted* otherwise.¹

¹The literature uses the terminology “even and odd” as well as “orientable and non-orientable” to describe the zero-circles, but the author has found this to be very confusing.

By work of Luttinger, any given pair (ω, g) of closed self-dual 2-form and Riemannian metric can be modified so that Z has any positive number of components (see [46, 52]), but as noted by Gompf (see [46, Theorem 1.8]), the number of untwisted zero-circles must have the same parity as that of $1 - b^1(X) + b_+^2(X)$. We can use these modifications to get rid of all twisted zero-circles (see [46, Remark 1.9]), and explicit constructions of near-symplectic forms having only untwisted zero-circles are given in [9, 50] in terms of handlebody decompositions.

Main results

Let ω be a near-symplectic form on (X, g) whose zero set Z has $N \geq 0$ components, all of which are untwisted zero-circles. As we just mentioned, this can always be arranged.

Remark 1.1.3. The assumption that Z has only untwisted zero-circles can be weakened in this thesis to some extent. As will be clarified in the appendix, we may assume Z to also have twisted zero-circles which are *non-contractible* in X .

Let \mathcal{N} denote a union of arbitrarily small tubular neighborhoods of the components of $Z \subset X$. Using Moser-type results of Honda [15] (see also [65, §2e]), \mathcal{N} can be chosen so that the complement

$$(X_0, \omega) := (X - \mathcal{N}, \omega|_{X - \mathcal{N}})$$

is a symplectic manifold with contact-type boundary, where each boundary component is a copy of $(S^1 \times S^2, e^{-1} \lambda_{\text{Taubes}})$. Here, λ_{Taubes} is an overtwisted contact form which is described in Section 1.3, and studied intensively by Taubes in order to characterize the pseudoholomorphic “beasts” living in the symplectization $\mathbb{R} \times (S^1 \times S^2)$.

Remark 1.1.4. The orientation of $S^1 \times S^2$ as a contact 3-manifold disagrees with the orientation of $S^1 \times S^2$ as a boundary component of X_0 , so each boundary component of X_0 is *concave* (if the orientations agreed then it would be a *convex* boundary component). This is consistent with a well-known result of Eliashberg: an overtwisted contact 3-manifold cannot be the convex boundary of a symplectic 4-manifold, i.e. there are no symplectic fillings of this contact 3-manifold. But note that (\mathcal{N}, ω) is a near-symplectic filling.

There is a canonical spin-c structure \mathfrak{s}_ω on X_0 whose positive spinor bundle is $\mathbb{S}_+ = \underline{\mathbb{C}} \oplus K^{-1}$, where $\underline{\mathbb{C}} \rightarrow X_0$ denotes the trivial complex line bundle and K is the canonical bundle of (X_0, J) for any chosen ω -compatible almost complex structure J . Any other spin-c structure on X_0 differs from \mathfrak{s}_ω by tensoring with a complex line bundle on X_0 . It follows from Taubes’ work on near-symplectic geometry that ω also induces a canonical identification of $\text{Spin}^c(X)$ with the set

$$\text{Rel}_\omega(X) := \left\{ A \in H_2(X_0, \partial X_0; \mathbb{Z}) \mid \partial A = \mathbb{1} \in H_1(\partial X_0; \mathbb{Z}) \right\}$$

where $\mathbb{1}$ is the oriented generator on each component (the orientation conventions are specified in Section 1.3). This correspondence is given by restricting $\mathfrak{s} \in \text{Spin}^c(X)$ to X_0 , which

differs from \mathfrak{s}_ω by a unique complex line bundle $L_{\mathfrak{s}} \rightarrow X_0$, and taking the Poincaré-Lefschetz dual of $c_1(L_{\mathfrak{s}})$. Denote the resulting class by $A_{\mathfrak{s}} \in \text{Rel}_\omega(X)$.

Of relevance to the Seiberg-Witten invariants of X will be J -holomorphic curves in a completion of (X_0, ω) which represent classes in $\text{Rel}_\omega(X)$. The completion (\overline{X}, ω) of (X_0, ω) is obtained by attaching cylindrical ends to the components of ∂X_0 and extending ω over them (see Section 1.2). In order to obtain well-defined counts we need all curves to be transverse. Unfortunately, there are special Reeb orbits in $(S^1 \times S^2, \lambda_{\text{Taubes}})$ for which there exist non-transverse multiply covered curves in \overline{X} asymptotic to multiple covers of those orbits.² What can we do?

A given class $A \in \text{Rel}_\omega(X)$ determines an upper bound $\rho(A)$ of the “lengths” of the possible Reeb orbits which the relevant J -holomorphic curves are asymptotic to, so there are only finitely many orbits which permit the existence of non-transverse curves. The idea now is to go back and choose a different tubular neighborhood \mathcal{N} , determining a different contact form on $S^1 \times S^2$ having the same contact structure ξ_{Taubes} as λ_{Taubes} . In particular, we search for a particular contact form whose Reeb orbits of “lengths” less than $\rho(A)$ do not permit the existence of non-transverse curves representing A . This is the content of the following lemma, proved in Section 1.3.

Lemma 1.1.5. *For $A \in \text{Rel}_\omega(X)$, there exists a tubular neighborhood \mathcal{N} of Z in X such that the symplectic cobordism $(X - \mathcal{N}, \omega)$ has contact-type boundary, where each boundary component is a copy of $(S^1 \times S^2, \lambda_A)$. Here, λ_A denotes a nondegenerate overtwisted contact form with contact structure ξ_{Taubes} such that its Reeb orbits of symplectic action less than $\rho(A)$ are $\rho(A)$ -flat and are either positive hyperbolic or $\rho(A)$ -positive elliptic.*

The “ $\rho(A)$ -positive” condition on the elliptic orbits is a key ingredient to preventing the existence of (negative ECH index) non-transverse curves that represent A , and this was observed by Hutchings [26] for more general symplectic cobordisms. The “ $\rho(A)$ -flatness” condition on λ_A is needed in order to identify our counts of curves with the Seiberg-Witten invariants, as Taubes [55] did for the isomorphism between ECH and Seiberg-Witten Floer cohomology.

The relevant counts of curves will a priori depend on a suitably generic choice of J and they will be packaged together as elements in $\text{ECH}_*(-\partial X_0, \xi_{\text{Taubes}})$. Roughly speaking, the ECH chain complex over \mathbb{Z} is generated by “orbit sets” which are finite sets of pairs (γ, m) of Reeb orbits γ in $(-\partial X_0, \xi_{\text{Taubes}})$ and multiplicities $m \in \mathbb{N}$. To define invariants of X , we will first count J -holomorphic curves in \overline{X} to obtain a cycle in the ECH chain complex: the asymptotics of each curve define a generator Θ of the ECH chain complex, and the integer coefficient attached to each such Θ is an integrally weighted count of the curves which are asymptotic to Θ . Specifically,

Theorem 1.1.6. *With (X, ω, A) specified above, fix a nonnegative integer I and an ordered set of equivalence classes $[\bar{\eta}] := \{[\eta_1], \dots, [\eta_p]\} \subset H_1(X; \mathbb{Z})/\text{Torsion}$, where $0 \leq p \leq I$ such*

²The problem is worse: these non-transverse curves also have negative ECH index (see Remark 1.3.10).

that $I - p$ is even. For suitably generic J on \overline{X} , there is a well-defined element

$$Gr_{X,\omega,J}^I(A, [\bar{\eta}]) \in ECH_*(-\partial X_0, \xi_{\text{Taubes}})$$

concentrated in a single grading $g(A, I)$. This element is given by integrally weighted counts of (disjoint and possibly multiply covered) J -holomorphic curves in (\overline{X}, ω) which satisfy the following: They are asymptotic to Reeb orbits with respect to λ_A whose total homology class in each $S^1 \times S^2$ component is the oriented generator; they represent the relative class A ; they pass through a priori chosen base points in X_0 ; they pass through an a priori chosen ordered set of disjoint oriented loops in X_0 which represent $[\bar{\eta}]$; and they have ECH index I .

Here, the “ECH index” may be viewed as a formal dimension of the relevant moduli space of J -holomorphic currents, while the point/loop constraints carve out a zero-dimensional subset of the moduli space. These curves are explicitly specified in Proposition 1.3.17 and Proposition 1.3.19.

Now, a given $\mathfrak{s} \in \text{Spin}^c(X)$ determines not only the class $A_{\mathfrak{s}} \in \text{Rel}_{\omega}(X)$ but also an integer

$$d(\mathfrak{s}) := \frac{1}{4} (c_1(\mathfrak{s})^2 - 2\chi(X) - 3\sigma(X))$$

where $c_1(\mathfrak{s})$ denotes the first Chern class of the spin-c structure’s positive spinor bundle, $\chi(X)$ denotes the Euler characteristic of X , and $\sigma(X)$ denotes the signature of X . This integer $d(\mathfrak{s})$ is the formal dimension of the moduli space of Seiberg-Witten solutions on X with respect to \mathfrak{s} .

Using ideas from Seiberg-Witten theory,³ we will show in Chapter 2 that $g(A_{\mathfrak{s}}, d(\mathfrak{s}))$ is the lowest grading for which $ECH_*(-\partial X_0, \xi_{\text{Taubes}})$ is nonzero. After choosing an ordering of the components of Z , there is then an identification of the group $ECH_{g(A_{\mathfrak{s}}, d(\mathfrak{s}))}(\partial X_0, \xi_{\text{Taubes}})$ with \mathbb{Z} (see Proposition 1.3.2, Remark 1.3.22, and Proposition 1.3.30).

Definition 1.1.7. Fix an ordering of the zero-circles of ω . The *near-symplectic Gromov invariants*

$$Gr_{X,\omega,J} : \text{Spin}^c(X) \rightarrow \Lambda^* H^1(X; \mathbb{Z})$$

are defined as follows. The component of $Gr_{X,\omega,J}(\mathfrak{s})$ in $\Lambda^p H^1(X; \mathbb{Z})$, for $p \leq d(\mathfrak{s})$ such that $d(\mathfrak{s}) - p$ is even, is

$$[\eta_1] \wedge \cdots \wedge [\eta_p] \mapsto Gr_{X,\omega,J}^{d(\mathfrak{s})}(A_{\mathfrak{s}}, [\bar{\eta}]) \in \mathbb{Z}$$

and it is defined to be zero for all other integers p .

These invariants currently depend on the choice of ω and J . With $\mathbb{Z}/2\mathbb{Z}$ coefficients, we will show in Chapter 2 that they recover the Seiberg-Witten invariants of X , hence are smooth invariants.⁴

³We may also be able to establish this fact using J -holomorphic curves alone.

⁴We may also be able to establish invariance, with \mathbb{Z} coefficients, by varying (ω, J) and analyzing the resulting moduli spaces of curves.

Theorem 1.1.8 (Chapter 2). *Let (X, ω) be as specified above, and $\mathfrak{s} \in \text{Spin}^c(X)$. Then*

$$Gr_{X, \omega}(\mathfrak{s}) = SW_X(\mathfrak{s}) \in \Lambda^* H^1(X; \mathbb{Z}) \otimes \mathbb{Z}/2\mathbb{Z}$$

where ω determines the chamber for defining the Seiberg-Witten invariants when $b_+^2(X) = 1$.

Remark 1.1.9. We expect that Theorem 1.1.8 also holds with \mathbb{Z} coefficients, similarly to how Taubes' isomorphisms (between embedded contact homology and a version of Seiberg-Witten Floer homology) with $\mathbb{Z}/2\mathbb{Z}$ coefficients can be lifted to \mathbb{Z} coefficients (see [57, §3.b]). Note that the definition of SW_X depends on a choice of homology orientation of X , while the definition of $Gr_{X, \omega}$ depends on a choice of ordering of the zero-circles of ω . We expect that ω , with a fixed ordering of its zero-circles, canonically determines a homology orientation of X (see Chapter 2 for an elaboration).

Definition 1.1.7 is an extension of Taubes' Gromov invariants [54], which were constructed in the setting $Z = \emptyset$ and refined by McDuff [40] in the presence of symplectic embedded spheres of square -1 . To see how Taubes' invariants fit into the framework of Theorem 1.1.6, we first set $\partial X_0 = \emptyset$ so that $(\bar{X}, \omega) = (X, \omega)$ is a closed symplectic 4-manifold. We then make the identification $ECH_0(\emptyset, 0) \cong \mathbb{Z}$ whose generator is the empty set of orbits. For a given $\mathfrak{s} \in \text{Spin}^c(X)$, the corresponding absolute class $A_{\mathfrak{s}} \in H_2(X; \mathbb{Z})$ determines the ECH index

$$I(A_{\mathfrak{s}}) := c_1(TX) \cdot A_{\mathfrak{s}} + A_{\mathfrak{s}} \cdot A_{\mathfrak{s}} \in 2\mathbb{Z}$$

which equals $d(\mathfrak{s})$, and Taubes' invariants may be interpreted as

$$Gr_{X, \omega, J}^{I(A_{\mathfrak{s}})}(A_{\mathfrak{s}}, [\bar{\eta}]) \in ECH_0(\emptyset, 0)$$

since all J -holomorphic curves have no punctures.

Remark 1.1.10. If X is a symplectic manifold and ω_0 is a near-symplectic form, we may consider a deformation of ω_0 to a symplectic form ω_1 through near-symplectic forms (except at a finite number of times along the deformation). The near-symplectic Gromov invariants of (X, ω_0) must equal Taubes' Gromov invariants of (X, ω_1) in light of Theorem 1.1.8 and [71, Theorem 4.1]. The relation between the respective pseudoholomorphic curves is not pursued in this thesis, but we leave the reader with a question: When a zero-circle shrinks and dies, does a holomorphic plane decrease its symplectic area and vanish, or does it close up into a sphere, or else?

Remark 1.1.11. With respect to Theorem 1.1.1, we expect that the near-symplectic Gromov invariants are related to counts of J -holomorphic curves in $X - Z$ by shrinking \mathcal{N} to its core Z . These curves would homologically bound Z in the sense that they have algebraic intersection number 1 with every linking 2-sphere of Z . A curve may also “pinch off” along Z in the sense that a portion of the curve forms a multi-sheeted cone in a small neighborhood of a point on Z . These properties correspond to the following facts: a relevant

J -holomorphic curve in $\overline{X - \mathcal{N}}$ is asymptotic to an orbit set on each $S^1 \times S^2$ whose total homology class is the oriented generator of $H_1(S^1 \times S^2; \mathbb{Z})$, and some orbits in that orbit set may be contractible. This is in agreement with Taubes' study [73] of the structure of such curves in the vicinity of Z .

What follows is a brief outline of the remainder of this chapter. Section 1.2 consists of a review of ECH, with a clarification in Section 1.2 of “ L -flat approximations” which were defined and used in Taubes’ isomorphisms [55]. Section 1.3 consists of the relevant facts of the contact manifold $(S^1 \times S^2, \lambda_{\text{Taubes}})$. Section 1.3 specifies how λ_{Taubes} is to be modified, through a particular change in the tubular neighborhood \mathcal{N} of Z . Section 1.3 clarifies the a priori difficulties with constructing well-defined counts of curves in $X_0 = X - \mathcal{N}$; in particular, we have to separate the cases where the relative classes $A \in \text{Rel}_\omega(X)$ may be represented by multiply covered exceptional spheres. Section 1.3 specifies the relevant J -holomorphic curves when there are no multiply covered exceptional spheres, while Section 1.3 permits the existence of multiply covered exceptional spheres. Section 1.3 specifies the integer weights that are assigned to the J -holomorphic curves for the definition of $Gr_{X, \omega, J}^I(A, [\bar{\eta}])$. Section 1.3 establishes the independence of choices of base points and loops that are used to define the relevant moduli spaces of J -holomorphic curves. Section 1.3 clarifies what $Gr_{X, \omega, J}^I(A, [\bar{\eta}])$ looks like as a class in ECH. Lastly, the appendix clarifies what changes are to be made in this thesis when ω is chosen to have twisted zero-circles.

1.2 Review of pseudoholomorphic curve theory

The point of this section is to introduce most of the terminology and notations that appear in the later sections. Further information and more complete details are found in [28].

Orbits with λ

Let (Y, λ) be a closed contact 3-manifold, oriented by $\lambda \wedge d\lambda > 0$. Let $\xi = \text{Ker } \lambda$ be the contact structure, oriented by $d\lambda$. Equivalently, ξ is oriented by the *Reeb vector field* R determined by $d\lambda(R, \cdot) = 0$ and $\lambda(R) = 1$. A Reeb orbit is a map $\gamma : \mathbb{R}/T\mathbb{Z} \rightarrow Y$ for some $T > 0$ with $\gamma'(t) = R(\gamma(t))$, modulo reparametrization, which is necessarily an m -fold cover of an embedded Reeb orbit for some $m \geq 1$. A given Reeb orbit is *nondegenerate* if the linearization of the Reeb flow around it does not have 1 as an eigenvalue, in which case the eigenvalues are either on the unit circle (such γ are *elliptic*) or on the real axis (such γ are *hyperbolic*). Assume from now on that λ is *nondegenerate*, i.e. all Reeb orbits are nondegenerate, which is a generic property of contact forms.

An *orbit set* is a finite set of pairs $\Theta = \{(\Theta_i, m_i)\}$ where the Θ_i are distinct embedded Reeb orbits and the m_i are positive integers. An orbit set is *admissible* if $m_i = 1$ whenever

Θ_i is hyperbolic. Its symplectic action (or “length”) is defined by

$$\mathcal{A}(\Theta) := \sum_i m_i \int_{\Theta_i} \lambda \geq 0$$

and its homology class is defined by

$$[\Theta] := \sum_i m_i [\Theta_i] \in H_1(Y; \mathbb{Z})$$

For a given $\Gamma \in H_1(Y; \mathbb{Z})$, the ECH chain complex $ECC_*(Y, \lambda, J, \Gamma)$ is freely generated over \mathbb{Z} by equivalence classes of pairs (Θ, \mathfrak{o}) , where Θ is an admissible orbit set satisfying $[\Theta] = \Gamma$ and \mathfrak{o} is a choice of ordering of the positive hyperbolic orbits and a $\mathbb{Z}/2\mathbb{Z}$ choice for each such orbit, such that $(\Theta, \mathfrak{o}) = -(\Theta, \mathfrak{o}')$ if \mathfrak{o} and \mathfrak{o}' differ by an odd permutation. We will suppress the notation of the orientation choices \mathfrak{o} . The differential ∂_{ECH} will be defined momentarily.

Curves with J

Let (Y_{\pm}, λ_{\pm}) be two contact 3-manifolds as above, possibly disconnected or empty. A *strong symplectic cobordism* from (Y_+, λ_+) to (Y_-, λ_-) is a compact symplectic manifold (X, ω) with oriented boundary

$$\partial X = Y_+ \sqcup -Y_-$$

such that $\omega|_{Y_{\pm}} = d\lambda_{\pm}$. We can always find neighborhoods N_{\pm} of Y_{\pm} in X diffeomorphic to $(-\varepsilon, 0] \times Y_+$ and $[0, \varepsilon) \times Y_-$, such that $\omega|_{N_{\pm}} = d(e^{\pm s} \lambda_{\pm})$ where s denotes the coordinate on $(-\varepsilon, 0]$. We then glue symplectization ends to X to obtain the *completion*

$$\overline{X} := ((-\infty, 0] \times Y_-) \cup_{Y_-} X \cup_{Y_+} ([0, \infty) \times Y_+)$$

of X , a noncompact symplectic 4-manifold whose symplectic form is also denoted by ω . We will also use the notation \overline{X} to denote the symplectization $\mathbb{R} \times Y$ of (Y, λ) , with $\omega = d(e^s \lambda)$.

An almost complex structure J on a symplectization $(\mathbb{R} \times Y, d(e^s \lambda))$ is *symplectization-admissible* if it is \mathbb{R} -invariant; $J(\partial_s) = R$; and $J(\xi) \subseteq \xi$ such that $d\lambda(v, Jv) \geq 0$ for $v \in \xi$. An almost complex structure J on the completion \overline{X} is *cobordism-admissible* if it is ω -compatible on X and agrees with symplectization-admissible almost complex structures on the ends $[0, \infty) \times Y_+$ and $(-\infty, 0] \times Y_-$.

Remark 1.2.1. With respect to the Riemannian metric defined by the symplectic form and the admissible almost complex structure, the symplectic form is a self-dual harmonic 2-form.

Given a cobordism-admissible J on \overline{X} and orbit sets $\Theta^+ = \{(\Theta_i^+, m_i^+)\}$ in Y_+ and $\Theta^- = \{(\Theta_j^-, m_j^-)\}$ in Y_- , a *J-holomorphic curve \mathcal{C} in \overline{X} from Θ^+ to Θ^-* is defined as follows. It is a J -holomorphic map $\mathcal{C} \rightarrow \overline{X}$ whose domain is a possibly disconnected punctured compact Riemann surface, defined up to composition with biholomorphisms of the domain,

with positive ends of \mathcal{C} asymptotic to covers of Θ_i^+ with total multiplicity m_i^+ , and with negative ends of \mathcal{C} asymptotic to covers of Θ_j^- with total multiplicity m_j^- (see [28, §3.1]). The moduli space of such curves is denoted by $\mathcal{M}(\Theta^+, \Theta^-)$, but where two such curves are considered equivalent if they represent the same current⁵ in \overline{X} , and in the case of a symplectization $\overline{X} = \mathbb{R} \times Y$ the equivalence includes translation of the \mathbb{R} -coordinate. An element $\mathcal{C} \in \mathcal{M}(\Theta^+, \Theta^-)$ can thus be viewed as a finite set of pairs $\{(C_k, d_k)\}$ or formal sum $\sum d_k C_k$, where the C_k are distinct irreducible somewhere injective J -holomorphic curves and the d_k are positive integers.

Let $H_2(X, \Theta^+, \Theta^-)$ be the set of relative 2-chains Σ in X such that

$$\partial\Sigma = \sum_i m_i^+ \Theta_i^+ - \sum_j m_j^- \Theta_j^-$$

modulo boundaries of 3-chains. It is an affine space over $H_2(X; \mathbb{Z})$, and every curve \mathcal{C} defines a relative class $[\mathcal{C}] \in H_2(X, \Theta^+, \Theta^-)$.

A *broken J -holomorphic curve (of height n) from Θ^+ to Θ^-* is a finite sequence of holomorphic curves $\{\mathcal{C}_k\}_{1 \leq k \leq n}$ and orbit sets $\{\Theta_k^\pm\}_{1 \leq k \leq n+1}$, such that there exists an integer $1 \leq k_0 \leq n$ such that the following holds:

- $\{\Theta_k^+\}_{k \geq k_0}$ belong to (Y_+, λ_+) and $\{\Theta_k^-\}_{k \leq k_0}$ belong to (Y_-, λ_-) ,
- $\Theta_1^- = \Theta^-$ and $\Theta_{n+1}^+ = \Theta^+$ and $\Theta_k^- = \Theta_{k-1}^+$ for $k > 1$,
- $\mathcal{C}_k \in \mathcal{M}(\Theta_k^+, \Theta_k^-)$ with respect to $J|_{\mathbb{R} \times Y_+}$ for $k > k_0$ (*symplectization levels*),
- $\mathcal{C}_k \in \mathcal{M}(\Theta_k^+, \Theta_k^-)$ with respect to $J|_{\mathbb{R} \times Y_-}$ for $k < k_0$ (*symplectization levels*),
- $\mathcal{C}_{k_0} \in \mathcal{M}(\Theta_{k_0}^+, \Theta_{k_0}^-)$ with respect to J (*cobordism level*),
- \mathcal{C}_k is not a union of unbranched covers of \mathbb{R} -invariant cylinders for $k \neq k_0$.

The moduli space of such broken curves is denoted by $\overline{\mathcal{M}(\Theta^+, \Theta^-)}$. There is an analogous definition of a *broken J -holomorphic current*, with the further requirement that each current \mathcal{C}_k for $k \neq k_0$ is not a union of \mathbb{R} -invariant cylinders with multiplicities.

There are relevant versions of Gromov compactness for the aforementioned moduli spaces. Any sequence of J -holomorphic curves $\{\mathcal{C}^\nu\}_{\nu \geq 1} \subset \mathcal{M}(\Theta^+, \Theta^-)$ with fixed genus and uniform energy bound has a subsequence which converges in the sense of *SFT compactness* [3] to a broken J -holomorphic curve. Any sequence of J -holomorphic currents $\{\mathcal{C}^\nu\}_{\nu \geq 1} \subset \mathcal{M}(\Theta^+, \Theta^-)$ with uniform energy bound has a subsequence which converges in an appropriate sense to a broken J -holomorphic current $(\mathcal{C}_1, \dots, \mathcal{C}_n) \in \overline{\mathcal{M}(\Theta^+, \Theta^-)}$, such that

$$\sum_{k=1}^n [\mathcal{C}_k] = [\mathcal{C}^\nu] \in H_2(X, \Theta^+, \Theta^-)$$

for ν sufficiently large (see [73, Proposition 3.3] and [28, Lemma 5.11]).

⁵For example, if \mathcal{C} is a d -fold cover of an embedded curve C , then the associated current is the \mathbb{R} -valued functional on $\Omega^2(\overline{X})$ given by $\sigma \mapsto d \int_C \sigma$. In particular, all branching data of the cover has been lost.

ECH

Denote by τ a homotopy class of symplectic trivializations of the restrictions of $\xi_{\pm} = \text{Ker } \lambda_{\pm}$ to the embedded orbits appearing in the orbit sets Θ^{\pm} . The *ECH index* of a current, or more generally of a class in $H_2(X, \Theta^+, \Theta^-)$, is given by

$$I(\mathcal{C}) = c_{\tau}(\mathcal{C}) + Q_{\tau}(\mathcal{C}) + CZ_{\tau}^I(\mathcal{C})$$

where $c_{\tau}(\mathcal{C})$ denotes the relative first Chern class of $\det T\bar{X}$ over \mathcal{C} with respect to τ (see [28, §3.2]), $Q_{\tau}(\mathcal{C})$ denotes a relative self-intersection pairing with respect to τ (see [28, §3.3]), and

$$CZ_{\tau}^I(\mathcal{C}) = \sum_i \sum_{k=1}^{m_i^+} CZ_{\tau}(\Theta_i^{+k}) - \sum_j \sum_{k=1}^{m_j^-} CZ_{\tau}(\Theta_j^{-k})$$

is a sum of Conley-Zehnder indices (of covers of orbits in Θ^{\pm}) with respect to τ . The definition of the Conley-Zehnder index will not be reviewed here (see [28, §3.2]), but it is noted that we can adjust τ over a given embedded orbit γ so that

$$CZ_{\tau}(\gamma^m) = 0$$

when γ is positive hyperbolic,

$$CZ_{\tau}(\gamma^m) = m$$

when γ is negative hyperbolic, and

$$CZ_{\tau}(\gamma^m) = 2\lfloor m\theta \rfloor + 1$$

when γ is elliptic. Here, the linearization of the Reeb flow around an elliptic orbit is conjugate to a rotation by angle $2\pi\theta$ with respect to τ , and $\theta \in \mathbb{R} - \mathbb{Q}$ is the *rotation number*. The equivalence class of θ in \mathbb{R}/\mathbb{Z} is the *rotation class* of the elliptic orbit, which does not depend on τ .

The *Fredholm index* of a curve \mathcal{C} , having i^{th} positive end asymptotic to α_i with multiplicity m_i and j^{th} negative end asymptotic to β_j with multiplicity n_j , is given by

$$\text{ind}(\mathcal{C}) = -\chi(\mathcal{C}) + 2c_{\tau}(\mathcal{C}) + CZ_{\tau}^{\text{ind}}(\mathcal{C})$$

where

$$CZ_{\tau}^{\text{ind}}(\mathcal{C}) = \sum_i CZ_{\tau}(\alpha_i^{m_i}) - \sum_j CZ_{\tau}(\beta_j^{n_j})$$

If \mathcal{C} has no multiply covered components then we have the *index inequality*

$$\text{ind}(\mathcal{C}) \leq I(\mathcal{C}) - 2\delta(\mathcal{C})$$

where $\delta(\mathcal{C})$ denotes an algebraic count of singularities of \mathcal{C} with positive integer weights (see [28, §3.4]). If J is furthermore generic then $\mathcal{M}(\Theta^+, \Theta^-)$ is an $\text{ind}(\mathcal{C})$ -dimensional manifold near \mathcal{C} .

Denote by $\mathcal{M}_I(\Theta^+, \Theta^-)$ the subset of elements in $\mathcal{M}(\Theta^+, \Theta^-)$ that have ECH index I . In a symplectization $\overline{X} = \mathbb{R} \times Y$ there is a characterization of currents with low ECH index. That is, if J is generic and \mathcal{C} is a J -holomorphic current in the symplectization $\overline{X} = \mathbb{R} \times Y$ then

- $I(\mathcal{C}) \geq 0$, with equality if and only if \mathcal{C} is a union of \mathbb{R} -invariant cylinders,

- If $I(\mathcal{C}) = 1$ then $\mathcal{C} = \mathcal{C}_0 \sqcup \mathcal{C}_1$, where \mathcal{C}_1 is an embedded $\text{ind} = I = 1$ curve and $I(\mathcal{C}_0) = 0$.

See [28, Proposition 3.7] for a proof. Given admissible orbit sets Θ^\pm of (Y, λ) , the coefficient $\langle \partial_{\text{ECH}} \Theta^+, \Theta^- \rangle$ is the signed⁶ count of elements in $\mathcal{M}_1(\Theta^+, \Theta^-)$. If J is generic then ∂_{ECH} is well-defined and $\partial_{\text{ECH}}^2 = 0$. The resulting homology is independent of the choice of J , depends only on ξ and Γ , and is denoted by $ECH_*(Y, \xi, \Gamma)$.

If Y is connected then Taubes constructed a canonical isomorphism of relatively graded modules

$$ECH_*(Y, \xi, \Gamma) \cong \widehat{HM}^{-*}(Y, \mathfrak{s}_\xi + \Gamma) \quad (1.1)$$

where $\widehat{HM}^{-*}(\cdot)$ is a version of Seiberg-Witten Floer homology defined by Kronheimer and Mrowka [33] and \mathfrak{s}_ξ is a certain spin-c structure determined by ξ . Moreover, both homologies admit absolute gradings by homotopy classes of oriented 2-plane fields on Y and Taubes' isomorphism preserves these gradings (see [7]).

Gradings and U-maps

Assume Y to be connected in this section. For each Γ there is a canonical absolute $\mathbb{Z}/2\mathbb{Z}$ grading on $ECH_*(Y, \xi, \Gamma)$ by the parity of the number of positive hyperbolic Reeb orbits in an admissible orbit set Θ . The total sum

$$ECH_*(Y, \xi) := \bigoplus_{\Gamma \in H_1(Y; \mathbb{Z})} ECH_*(Y, \xi, \Gamma)$$

has an absolute grading by homotopy classes of oriented 2-plane fields on Y (see [30, §3]), the set of which is denoted by $J(Y)$. This grading of an admissible orbit set Θ is denoted by $|\Theta|$.

As described in [30, §3], [33, §28], and [12, §4], there is a well-defined map $J(Y) \rightarrow \text{Spin}^c(Y)$ with the following properties. If $H^2(Y; \mathbb{Z})$ has no 2-torsion then the Euler class of the given 2-plane field uniquely determines the corresponding spin-c structure. There is a transitive \mathbb{Z} -action on $J(Y)$ whose orbits correspond to the spin-c structures: If $[\xi] \in J(Y)$ then $[\xi] + n$ is the homotopy class of a 2-plane field which agrees with ξ outside a small

⁶The moduli spaces are coherently oriented in the sense of [22, §9].

ball $B^3 \subset Y$ and disagrees with ξ on B^3 by a map $(B^3, \partial B^3) \rightarrow (SO(3), \{1\})$ of degree $2n$.⁷ A given orbit $J(Y, \mathfrak{s})$ is freely acted on by \mathbb{Z} if and only if the corresponding Euler class is torsion.

With that said, $ECH_*(Y, \xi, \Gamma)$ has a relative $\mathbb{Z}/d\mathbb{Z}$ grading, where d denotes the divisibility of $c_1(\xi) + 2\text{PD}(\Gamma)$ in $H^2(Y; \mathbb{Z})/\text{Torsion}$. It is refined by the absolute grading and satisfies

$$|\Theta^+| - |\Theta^-| \equiv I(\Sigma) \pmod{d}$$

for any $\Sigma \in H_2(Y, \Theta^+, \Theta^-)$, thanks to the index ambiguity formula [28, Equation 3.6].

Similarly to the degree -1 ECH differential, there is a degree -2 chain map

$$U_y : ECH_*(Y, \lambda, \Gamma) \rightarrow ECH_{*-2}(Y, \lambda, \Gamma)$$

that counts ECH index 2 currents passing through an a priori chosen base point $(0, y) \in \mathbb{R} \times Y$, where y does not lie on any Reeb orbit (see [28, §3.8]). Such a current is of the form $\mathcal{C}_0 \sqcup C_2$, where $I(\mathcal{C}_0) = 0$ and C_2 is an embedded $\text{ind}(C_2) = I(C_2) = 2$ curve passing through $(0, y)$. On the level of homology this U -map does not depend on the choice of base point.

L-flat approximations

The symplectic action induces a filtration on the ECH chain complex. For a positive real number L , the L -filtered ECH is the homology of the subcomplex $ECC_*^L(Y, \lambda, J, \Gamma)$ spanned by admissible orbit sets of action less than L . The ordinary ECH is recovered by taking the direct limit over L , via maps induced by inclusions of the filtered chain complexes.

Let $u : C \rightarrow \overline{X}$ be an immersed connected J -holomorphic curve and denote by N_C its normal bundle. The linearization of the J -holomorphic equation for C defines its *deformation operator*, a 1st order elliptic differential operator

$$D_C : L_1^2(N_C) \rightarrow L^2(T^{0,1}C \otimes N_C), \quad \eta \mapsto \bar{\partial}\eta + \nu_C\eta + \mu_C\bar{\eta}$$

where the appropriate sections $\nu_C \in \Gamma(T^{0,1}C)$ and $\mu_C \in \Gamma(T^{0,1}C \otimes N_C^2)$ are determined by the covariant derivatives of J in directions normal to C , $\bar{\partial}$ is the d-bar operator arising from the Hermitian structure on N_C , and $\bar{\eta}$ denotes the conjugate of η in N_C^{-1} . Let N_C and $T^{0,1}C$ be suitably trivialized on an end of C asymptotic to a Reeb orbit γ . Then D_C is asymptotic (in the sense of [76, §2]) to the *asymptotic operator* associated with γ ,

$$L_\gamma : L_1^2(\gamma^*\xi) \rightarrow L^2(\gamma^*\xi), \quad \eta \mapsto \frac{i}{2}\partial_t\eta + \nu\eta + \mu\bar{\eta}$$

and the pair (ν_C, μ_C) is asymptotic to the pair (ν, μ) over γ .

For a fixed $L > 0$ it will be convenient to modify λ and J on small tubular neighborhoods of all Reeb orbits of action less than L , in order to relate J -holomorphic curves to

⁷This convention is opposite to that used in [33].

Seiberg-Witten theory most easily. The desired modifications of (λ, J) are called *L-flat approximations*, and were introduced by Taubes in [55, Appendix] and [59, Proposition 2.5]. They induce isomorphisms on the *L*-filtered ECH chain complex, but for the point of this thesis (see Lemma 1.3.9 and Section 1.3) we really only need to know that:

- The Reeb orbits of action less than L (and their action) are not altered,
- The C^1 -norm of the difference between the contact forms can be made as small as desired,
- The pair (ν, μ) associated with an elliptic orbit with rotation number θ in a given trivialization is modified to $(\frac{1}{2}\theta, 0)$, so that its asymptotic operator is complex linear.

Remark 1.2.2. The proof of Theorem 1.1.8 in Chapter 2 will make crucial use of Taubes' isomorphism between ECH and Seiberg-Witten Floer cohomology, and that makes use of *L*-flat approximations. The key fact here is that Taubes' isomorphism actually exists on the *L*-filtered chain complex level, for which *L*-flat orbit sets are in bijection with Seiberg-Witten solutions of “energy” less than $2\pi L$.

1.3 Towards a near-symplectic Gromov invariant

This section spells out the proof of Theorem 1.1.6. We introduce the contact form on $S^1 \times S^2$ that was studied in the past by Taubes, and then we find a different contact form on $S^1 \times S^2$ whose contact dynamics is “tame” in a certain sense. This new contact form is better suited for establishing well-defined counts of J -holomorphic curves in $X - Z$. We also find a tubular neighborhood \mathcal{N} of Z in X for which the induced contact dynamics on its boundary $\partial\mathcal{N}$ is given by our new contact form. Then we construct the relevant moduli spaces of J -holomorphic curves in $X - \mathcal{N}$ which are to be counted, including their integer weights, and from there we define the relevant class in ECH. We retain the same assumptions that are made in Section 1.1.

Taubes' contact form

The main focus of this thesis is $S^1 \times S^2$ equipped with *Taubes' contact form*

$$\lambda_{\text{Taubes}} = -(1 - 3 \cos^2 \theta)dt - \sqrt{6} \cos \theta \sin^2 \theta d\varphi \quad (1.2)$$

for coordinates $(t, \theta, \varphi) \in S^1 \times S^2$ such that $0 \leq t \leq 2\pi$ and $0 \leq \theta \leq \pi$ and $0 \leq \varphi \leq 2\pi$. In order for $\lambda_{\text{Taubes}} \wedge d\lambda_{\text{Taubes}}$ to be positive, $S^1 \times S^2$ is oriented by the 3-form

$$-\sin \theta dt d\theta d\varphi$$

The S^1 -factor will be oriented by the 1-form $-dt$, and the S^2 -factor will be oriented by the 2-form $\sin \theta d\theta d\varphi$. This contact manifold was originally studied in [65, 15, 45], and we now describe some details that will be of use later on.

The Reeb field associated with λ_{Taubes} is

$$\frac{-1}{1 + 3 \cos^4 \theta} R_{\text{Taubes}}, \quad R_{\text{Taubes}} = (1 - 3 \cos^2 \theta) \partial_t + \sqrt{6} \cos \theta \partial_\varphi$$

The Reeb orbits live in the constant $\theta = \theta_0$ slices of $S^1 \times S^2$ satisfying

$$\theta_0 \in \{0, \pi\} \quad \text{or} \quad \frac{\sqrt{6} \cos \theta_0}{1 - 3 \cos^2 \theta_0} \in \mathbb{Q} \cup \{\pm\infty\} \quad (1.3)$$

The two nondegenerate orbits at $\theta_0 \in \{0, \pi\}$ are elliptic. They are denoted by e_0, e_π and called *the exceptional orbits*. The remaining orbits are degenerate, there being an S^1 -family of orbits for each such θ_0 . In other words,

$$T(\theta_0) := \{\text{constant } \theta = \theta_0 \text{ slice}\} \subset S^1 \times S^2$$

is a torus foliated by orbits. Taubes' contact form is therefore not nondegenerate, but it is “Morse-Bott” in the sense of [4].

Remark 1.3.1. The contact structure $\xi_{\text{Taubes}} = \text{Ker } \lambda_{\text{Taubes}}$ is overtwisted, as pointed out in [65, §2.f] and [15, Proposition 9]. A well-known result of Hofer states that an overtwisted contact 3-manifold must have at least one contractible orbit, and indeed we see that $T(\arccos(\frac{1}{\sqrt{3}}))$ and $T(\arccos(-\frac{1}{\sqrt{3}}))$ consist of contractible orbits. The remaining orbits are all homologically nontrivial.

With the contact structure ξ_{Taubes} oriented by $d\lambda_{\text{Taubes}}$, we compute

$$c_1(\xi_{\text{Taubes}}) = -2 \in \mathbb{Z} \cong H^2(S^1 \times S^2; \mathbb{Z})$$

by using the section $\sin \theta \partial_\theta \in \Gamma(\xi)$ and noting that the orientation on ξ_{Taubes} disagrees with the orientation on the S^2 -factor at $\theta = 0$ and $\theta = \pi$. The spin-c structure \mathfrak{s}_ξ determined by ξ_{Taubes} satisfies

$$c_1(\mathfrak{s}_\xi + 1) = c_1(\mathfrak{s}_\xi) + 2 = c_1(\xi_{\text{Taubes}}) + 2 = 0$$

and so Taubes' isomorphism (1.1) reads

$$ECH_j(S^1 \times S^2, \xi_{\text{Taubes}}, 1) \cong \widehat{HM}^j(S^1 \times S^2, \mathfrak{s}_\xi + 1)$$

where $j \in J(S^1 \times S^2, \mathfrak{s}_\xi + 1) \cong \mathbb{Z}$ as \mathbb{Z} -sets. There is a unique class $j = [\xi_*]$ represented by an oriented 2-plane field ξ_* on $S^1 \times S^2$ which has vanishing Euler class and is invariant under rotations of the S^1 -factor.

Proposition 1.3.2. *If $\Gamma \in H_1(S^1 \times S^2; \mathbb{Z})$ is not the oriented generator 1 then*

$$ECH_*(S^1 \times S^2, \xi_{\text{Taubes}}, \Gamma) = 0$$

In the remaining case $\Gamma = 1$, $ECH_j(S^1 \times S^2, \xi_{\text{Taubes}}, 1)$ is zero in gradings below $j = [\xi_*]$, and for each $n \geq 0$

$$ECH_{[\xi_*]+n}(S^1 \times S^2, \xi_{\text{Taubes}}, 1) \cong \mathbb{Z}$$

The homotopy class $[\xi_*]$ has odd parity under the canonical absolute $\mathbb{Z}/2\mathbb{Z}$ grading on ECH. After perturbing λ_{Taubes} to a nondegenerate contact form, $[\xi_*]$ is generated by a single positive hyperbolic orbit wrapping positively once around the S^1 -factor in $T(\frac{\pi}{2}) = S^1 \times \{\text{equator}\}$.

Proof. The first two statements are proved in [33, §IX.36] for the relevant version of Seiberg-Witten Floer homology, so they follow for ECH via Taubes' isomorphism (1.1). The latter two statements follow from the relevant statements in [20, §12.2.1]. This reference uses a “twisted” version of ECH that remembers some information about the relative homology classes of the J -holomorphic curves in $S^1 \times S^2$, but the untwisted version may be obtained via a spectral sequence in [20, §8.1]. \square

In the upcoming section, Taubes' contact form will be modified in various ways. We now preemptively analyze the contact form $e^f \lambda_{\text{Taubes}}$ for a given smooth function $f : S^1 \times S^2 \rightarrow \mathbb{R}$ depending only on the θ coordinate, whose contact structure is nonetheless ξ_{Taubes} . Such a contact form can be written as

$$\lambda = a_1(\theta)dt + a_2(\theta)d\varphi \tag{1.4}$$

for some smooth pair

$$a = (a_1, a_2) : [0, \pi] \rightarrow \mathbb{R}^2 - \{(0, 0)\}$$

Let $a \times a' := a_1 a'_2 - a_2 a'_1$, where the tick-mark signifies the derivative with respect to θ . The condition for λ to be a positive contact form is then

$$\frac{a \times a'(\theta)}{\sin \theta} < 0$$

for all $\theta \in [0, \pi]$. For $\theta \in (0, \pi)$ the Reeb field of λ is

$$R = \frac{1}{a \times a'(\theta)} \left(a'_2(\theta) \frac{\partial}{\partial t} - a'_1(\theta) \frac{\partial}{\partial \varphi} \right)$$

and the condition (1.3), for which $T(\theta_0) \subset S^1 \times S^2$ is a torus foliated by orbits, is now given by

$$\frac{a'_1(\theta_0)}{a'_2(\theta_0)} \in \mathbb{Q} \cup \{\pm\infty\} \tag{1.5}$$

Every embedded orbit in $T(\theta_0)$ represents the same class in $H_1(T(\theta_0); \mathbb{Z})$ and they all have the same action $\mathcal{A}(\theta_0) > 0$. There are also two exceptional nondegenerate elliptic orbits at $\theta_0 \in \{0, \pi\}$.

Lemma 1.3.3. *The exceptional elliptic orbit at $\theta = \theta_0$, for $\theta_0 \in \{0, \pi\}$, has rotation class*

$$\left(\operatorname{sign} \lim_{\theta \rightarrow \theta_0} \frac{-a'_2(\theta)}{\sin \theta \cos \theta} \right) \frac{a'_1(\theta_0)}{a'_2(\theta_0)} \mod 1$$

In particular, the rotation class for either exceptional orbit of Taubes' contact form is $\sqrt{\frac{3}{2}} \mod 1$.

Proof. The Reeb field along each exceptional orbit is

$$\frac{1}{a_1(\theta_0)} \partial_t$$

Note that $a_1(\theta_0) > 0$, because Taubes' contact form satisfies it and e^f is a positive rescaling. Therefore, in a neighborhood of $\theta = \theta_0$ the Reeb flow is in the positive t -direction and wraps $-a'_1(\theta_0)/a'_2(\theta_0)$ times around the φ -coordinate circle after traversing once around the t -coordinate circle. The rotation class of each elliptic orbit is therefore

$$\epsilon \cdot \frac{-a'_1(\theta_0)}{a'_2(\theta_0)} \mod 1 \quad (1.6)$$

where $\epsilon = \pm 1$ depending on whether the φ -coordinate circle is positively or negatively oriented with respect to the orientation of the contact planes $T_{\theta_0} S^2$ given by $d\lambda$. To determine ϵ , we use the Cartesian coordinates

$$\begin{cases} (x, y) = (\sin \theta \cos \varphi, \sin \theta \sin \varphi) \\ dx dy = \sin \theta \cos \theta d\theta d\varphi \end{cases}$$

near each pole of S^2 . In these coordinates,

$$d\lambda|_{T_{\theta_0} S^2} = \left(\lim_{\theta \rightarrow \theta_0} \frac{a'_2(\theta)}{\sin \theta \cos \theta} \right) dx dy$$

and ϵ is precisely the sign of this parenthetical expression. \square

Changing Taubes' contact form

In order to obtain well-defined counts of J -holomorphic curves which represent a given class $A \in \operatorname{Rel}_\omega(X)$, we will need to ensure a bound on their energy as well as a bound on the symplectic action of their orbit sets. As explained in [27], these bounds are given by a particular quantity $\rho(A)$, defined as follows:

Let $u : \Sigma \rightarrow X_0$ be any given smooth map which represents A , where Σ is a compact oriented smooth surface with boundary and $u(\partial \Sigma) \subset \partial X_0$. Then

$$\rho(A) := \int_{\Sigma} \omega + \int_{\partial \Sigma} e^{-1} \lambda_{\text{Taubes}} \quad (1.7)$$

This quantity is additive with respect to composition of symplectic cobordisms, and it vanishes on exact symplectic cobordisms (recall that a symplectic cobordism is *exact* if the contact form on the boundary extends as a global primitive 1-form of the symplectic form). Subsequently, $\rho(A)$ does not change if the symplectic cobordism (X_0, ω) is modified by composing it with an exact symplectic cobordism.

Remark 1.3.4. We can view $\rho : \text{Rel}_\omega(X) \rightarrow \mathbb{R}$ as either a measure of the failure of exactness of ω , or as an energy. In particular, if (X, ω) were a closed symplectic manifold, i.e. $\partial X_0 = \emptyset$ and $A \in H_2(X; \mathbb{Z})$, then $\rho(A) = A \cdot [\omega]$.

Now, three modifications will be made to Taubes' contact form. First, we will want all Reeb orbits to be nondegenerate in order to define ECH. Second, we will want all Reeb orbits of action less than $\rho(A)$ to be $\rho(A)$ -flat in order to relate the J -holomorphic curves to Seiberg-Witten theory. Third, we will want the elliptic orbits of action less than $\rho(A)$, especially the exceptional orbits, to be “ $\rho(A)$ -positive” in order to guarantee transversality of the relevant moduli spaces of J -holomorphic curves (see Remark 1.3.10 below). As defined in [26], the quantifier “ $\rho(A)$ -positive” means the following:

Definition 1.3.5. Fix $L > 0$. Let γ be a nondegenerate embedded elliptic orbit with rotation class $\theta \in \mathbb{R}/\mathbb{Z}$ and symplectic action $\mathcal{A}(\gamma) < L$. Then γ is *L-positive* if $\theta \in (0, \mathcal{A}(\gamma)/L) \bmod 1$.

A key property of any L -positive elliptic orbit γ is that if L is much greater than $\mathcal{A}(\gamma)$, then $CZ_\tau(\gamma^m) = 1$ for $m < L/\mathcal{A}(\gamma)$ and a particular choice of trivialization τ of $\gamma^* \xi_{\text{Taubes}}$.

The next lemma below shows how to modify the Morse-Bott orbits, in the sense of [4] and adapted from [26, Lemma 5.4]. For a given positive contact form written as (1.4), the lemma requires the following technical condition

$$a' \times a''(\theta_0) < 0 \tag{1.8}$$

for all $\theta_0 \in (0, \pi)$ that satisfy (1.5). Note that Taubes' contact form satisfies the technical condition.

Lemma 1.3.6. Suppose the positive contact form $\lambda = a_1(\theta)dt + a_2(\theta)d\varphi$ satisfies the technical condition (1.8). Then for every $L > 0$ and sufficiently small $\delta > 0$, there exists a perturbation $e^{f_{\delta,L}} \lambda$ of λ satisfying the following properties:

- $f_{\delta,L} \in C^\infty(S^1 \times S^2)$ satisfies $\|f_{\delta,L}\|_{C^0} < \delta$,
- $e^{f_{\delta,L}} \lambda$ agrees with λ near the exceptional orbits at $\theta_0 \in \{0, \pi\}$,
- Each torus $T(\theta_0)$ with $\mathcal{A}(\theta_0) < L$ is replaced by a positive hyperbolic orbit and an L -positive elliptic orbit, both of action less than L and within δ of $\mathcal{A}(\theta_0)$,
- $e^{f_{\delta,L}} \lambda$ has no other embedded orbits of action less than L .

Proof. The function $f_{\delta,L}$ is given by Bourgeois' perturbation [4] of λ , which breaks up each $T(\theta_0)$ into two embedded nondegenerate orbits of action slightly less than $\mathcal{A}(\theta_0)$ in lieu of orbits of action greater than L . Namely, there is a positive hyperbolic orbit and an elliptic orbit e_{θ_0} , both representing the same class in $H_1(T(\theta_0); \mathbb{Z})$. For sufficiently small perturbations there cannot exist other orbits of action less than L , otherwise we would find a sequence $\{(\gamma_k, \delta_k)\}_{k \in \mathbb{N}}$ of such orbits of uniformly bounded action L and perturbations $\delta_k \rightarrow 0$ for which a subsequence converges to one of the original degenerate orbits (by the Arzelà-Ascoli theorem), yielding a contradiction.

It remains to compute the rotation class of the elliptic orbit created from each Morse-Bott family. Let $a^\perp := (a_2, -a_1)$. The basis $\langle \partial_\theta, a^\perp \rangle$ defines a trivialization τ of the contact structure ξ over $S^1 \times (S^2 - \{\text{poles}\})$, since

$$d\lambda(\partial_\theta, a^\perp) = -a \times a'(\theta) > 0$$

for $0 < \theta < \pi$. We then compute the Lie derivatives

$$\mathcal{L}_{\partial_\theta} R = -\frac{a' \times a''}{(a \times a')^2} a^\perp, \quad \mathcal{L}_{a^\perp} R = 0$$

to see that the linearized Reeb flow along $T(\theta_0)$ induces the linearized return map

$$1 + \begin{pmatrix} 0 & 0 \\ r(\theta_0) \mathcal{A}(\theta_0) & 0 \end{pmatrix}$$

on ξ in the chosen basis, where

$$r := -\frac{a' \times a''}{(a \times a')^2}$$

The linearized return map along e_{θ_0} is a perturbation of the original linearized return map along $T(\theta_0)$, so the rotation number of e_{θ_0} has the same sign as $r(\theta_0)$, i.e. it has the opposite sign of $a' \times a''(\theta_0)$. This rotation number can be made arbitrarily small by choosing δ sufficiently small. In other words, it follows from the technical condition (1.8) that each e_{θ_0} is L -positive. \square

Now we show how to modify the exceptional orbits while preserving the technical condition (1.8).

Lemma 1.3.7. *Given $c \geq 0$ and $0 < \varepsilon \leq \sqrt{\frac{3}{2}}$, there exists a smooth nonpositive function $f_{\varepsilon,c}$ on $S^1 \times S^2$ which only depends on the θ coordinate, such that $e^{f_{\varepsilon,c}} \lambda_{\text{Taubes}}$ satisfies the technical condition (1.8) and whose exceptional orbits e_0 and e_π have rotation classes both equal to $\varepsilon \bmod 1$.*

Proof. By Lemma 1.3.3, the rotation classes under consideration are

$$\left(\text{sign} \lim_{\theta \rightarrow \theta_0} (3 \cos \theta - \frac{1}{\cos \theta}) \right) \frac{\frac{\partial f_{\varepsilon,c}}{\partial \theta} \frac{3 \cos^2 \theta - 1}{\sin \theta} - 6 \cos \theta}{\sqrt{6}(1 - 3 \cos^2 \theta) - \frac{\partial f_{\varepsilon,c}}{\partial \theta} \sqrt{6} \cos \theta \sin \theta} \bmod 1$$

for $\theta_0 \in \{0, \pi\}$. After setting

$$\frac{\partial f_{\varepsilon,c}}{\partial \theta} := \tilde{f}_\varepsilon(\theta) \sin \theta$$

for a smooth function $\tilde{f}_\varepsilon : [0, \pi] \rightarrow \mathbb{R}$, it follows that the rotation classes of e_0 and e_π are respectively given by

$$\begin{aligned} -\frac{1}{\sqrt{6}} (\tilde{f}_\varepsilon(\theta) - 3 \cos \theta) \Big|_{\theta=0} &\mod 1 \\ \frac{1}{\sqrt{6}} (\tilde{f}_\varepsilon(\theta) - 3 \cos \theta) \Big|_{\theta=\pi} &\mod 1 \end{aligned}$$

These are both equal to ε upon setting $\tilde{f}_\varepsilon(\theta) = (3 - \sqrt{6}\varepsilon) \cos \theta$. We then pick the antiderivative of \tilde{f}_ε using c so that the desired function on $S^1 \times S^2$ is

$$f_{\varepsilon,c}(t, \theta, \varphi) = -\frac{3 - \sqrt{6}\varepsilon}{2} \cos^2 \theta - c \quad (1.9)$$

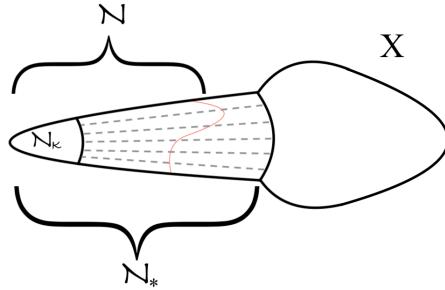
A brute force calculation shows that $e^{f_{\varepsilon,c}} \lambda_{\text{Taubes}}$ satisfies the technical condition (1.8). \square

Remark 1.3.8. As a sanity check, if $\varepsilon = \sqrt{\frac{3}{2}}$ and $c = 0$ then we recover Taubes' contact form.

We now move forward and show how exactly Taubes' contact form is to be modified to prove Lemma 1.1.5 in the introduction. Such modifications give us control over the orbits of low symplectic action, at the expense of producing new orbits of high symplectic action with unknown properties. This is sufficient for the purposes of this thesis, because for a given class $A \in \text{Rel}_\omega(X)$ only the orbit sets of symplectic action less than $\rho(A)$ are relevant to the tentative Gromov invariant.

Lemma 1.3.9. *Suppose (X, ω) is a near-symplectic 4-manifold such that all components of Z are untwisted zero-circles, and fix $A \in \text{Rel}_\omega(X)$. Then there is a choice of neighborhood \mathcal{N} of Z in X such that $(X - \mathcal{N}, \omega)$ is a symplectic manifold with contact-type boundary whose boundary components are copies of $(S^1 \times S^2, \lambda_A)$. Here, λ_A is a nondegenerate contact form with contact structure ξ_{Taubes} but whose orbits of symplectic action less than $\rho(A)$ are all $\rho(A)$ -flat and are either positive hyperbolic or $\rho(A)$ -positive elliptic.*

Proof. By Lemma 1.3.7, for any L sufficiently greater than $\rho(A)$ there is a smooth function $f_{\varepsilon,c}$ on $S^1 \times S^2$ for any $c \geq 0$ and sufficiently small choice of ε such that the exceptional orbits of $e^{f_{\varepsilon,c}} \lambda_{\text{Taubes}}$ are L -positive. By Lemma 1.3.6, we can perturb this new Morse-Bott contact form (à la Bourgeois) so that all orbits of action less than L are nondegenerate and L -positive when elliptic. As explained in Section 1.2, we can perturb the resulting contact form (à la Taubes) so that all orbits of action less than L are furthermore L -flat. This new contact form is still degenerate, but it can be perturbed to become nondegenerate without disturbing the existing orbits of action less than L and without introducing other orbits of

Figure 1.1: Hypersurface defined by F_A in red

action less than L . Explicitly, let $U \subset S^1 \times S^2$ be the union of L -flat neighborhoods of the orbits of action less than L , which do not contain any other orbits, and use [5, Lemma 2.2] to make a generic perturbation of the contact form on the complement of U .

The final result is a contact form λ_A which has the same contact structure as λ_{Taubes} , so for c sufficiently large there exists a smooth negative function $F_A \in C^\infty(S^1 \times S^2)$ such that the maximum value of F_A over $S^1 \times S^2$ is less than -1 and

$$\lambda_A = e^{F_A} \lambda_{\text{Taubes}}$$

It remains to show that the neighborhood N can be chosen so that the contact form on any boundary component of $X - N$ is λ_A , preserving compatibility with ω . Assume without loss of generality that Z is a single circle. We first choose the neighborhood N_* of Z as specified in Section 1.1. Moreover, this choice of N_* can be constrained further (using [65, Lemma 2.3]) so that for any fixed $\kappa > 1$, there is a neighborhood $N_\kappa \subset N_*$ of Z such that $(N_* - N_\kappa, \omega)$ is symplectomorphic to $([-\kappa, -1] \times S^1 \times S^2, d(e^s \lambda_{\text{Taubes}}))$, where s denotes the coordinate on $[-\kappa, -1]$. We choose $\kappa \geq -\inf F_A > 1$, because $(N_* - N_\kappa, \omega)$ then contains the contact hypersurface

$$\{(s, x) \in \mathbb{R} \times (S^1 \times S^2) \mid s = F_A(x)\}$$

with contact form λ_A . Therefore, our desired neighborhood is

$$N := N_\kappa \cup \{(s, x) \in [-\kappa, -1] \times (S^1 \times S^2) \mid s \leq F_A(x)\}$$

which is the subset of N_* “below the contact hypersurface” (a schematic is given by Figure 1.1). Note that the quantity $\rho(A)$ has not changed, since $(X - N, \omega)$ differs from $(X - N_*, \omega)$ by composition with an exact symplectic cobordism.

□

ECH cobordism maps

Fix $A \in \text{Rel}_\omega(X)$. Thanks to Lemma 1.3.9, we choose N so that (X_0, ω) is a strong symplectic cobordism from the empty set $(\emptyset, 0)$ to a disjoint union of N copies of the contact

3-manifold $(S^1 \times S^2, \lambda_A)$. Let \overline{X} denote its completion, and fix a cobordism-admissible almost complex structure J on (\overline{X}, ω) . As shown in [27], there are induced *ECH cobordism maps* of the form

$$\Phi_A : ECH_0(\emptyset, 0, 0) \rightarrow ECH_*(\bigsqcup_{k=1}^N S^1 \times S^2, \xi_{\text{Taubes}}, 1) \quad (1.10)$$

defined by suitable counts of Seiberg-Witten instantons on \overline{X} . Since $ECH_0(\emptyset, 0, 0) \cong \mathbb{Z}$ is generated by the empty set of orbits, and a choice of ordering of the components of ∂X_0 defines an identification of $ECH_*(\bigsqcup_{k=1}^N S^1 \times S^2, \xi_{\text{Taubes}}, 1)$ with $\bigotimes_{k=1}^N ECH_*(S^1 \times S^2, \xi_{\text{Taubes}}, 1)$, the map (1.10) should really be viewed as an element

$$\Phi_A \in \bigotimes_{k=1}^N ECH_*(S^1 \times S^2, \xi_{\text{Taubes}}, 1)$$

We now present a definition of Φ_A via counts of J -holomorphic curves in \overline{X} . In Chapter 2 we will show that these two definitions coincide.

Remark 1.3.10. As explained in [28, 27], the main problem with constructing ECH cobordism maps via J -holomorphic curves is that negative ECH index curves can arise and it is currently unknown how to count them appropriately. In this thesis, negative ECH index curves (with punctures) do not arise in the counts for Φ_A thanks to our choice of contact form for the boundary components of X_0 . If we used Taubes' contact form instead, the rotation classes of the exceptional orbits would allow the Conley-Zehnder index of their multiple covers to get large enough in magnitude to force the ECH index to become negative:

Example 1.3.11. Let $C \hookrightarrow \overline{X}$ be an embedded J -holomorphic plane asymptotic to the exceptional orbit e_0 with multiplicity 1, such that $0 = \text{ind}(C) = I(C) = 0$. Fix the trivialization τ of $e_0^* \xi_{\text{Taubes}}$ so that the rotation number of e_0 is $\sqrt{\frac{3}{2}} - 1$. Then $c_\tau(C) = 1$, $Q_\tau(C) = 0$, $CZ_\tau(e_0) = 1$, and

$$I(d \cdot C) = d \cdot c_\tau(C) + d^2 Q_\tau(C) + CZ_\tau^I(d \cdot C) = -2 \sum_{k=1}^d \lfloor k(\sqrt{\frac{3}{2}} - 1) \rfloor$$

which is negative for $d \geq 5$.

An *exceptional sphere* in X is an embedded smooth sphere of self-intersection -1 , and X is *minimal* if there are no exceptional spheres. As with Taubes' Gromov invariants for closed symplectic 4-manifolds, the trouble with multiply covered exceptional spheres is that they have negative ECH index: a holomorphic exceptional sphere with multiplicity $d > 1$ has index $-d(d-1)$. We will deal with this complication momentarily.

Let $\mathcal{E}_\omega \subset H_2(X_0, \partial X_0; \mathbb{Z})$ denote the set of classes represented by *symplectic* exceptional spheres in X_0 . Note that there is a nondegenerate bilinear pairing

$$H_2(X_0; \mathbb{Z}) \otimes H_2(X_0, \partial X_0; \mathbb{Z}) \xrightarrow{\cdot} \mathbb{Z}$$

induced by Poincaré-Lefschetz duality and the relative cap-product, and $e \cdot e = -1$ for any $e \in \mathcal{E}_\omega$.

Lemma 1.3.12. *For generic cobordism-admissible J on (\bar{X}, ω) , every symplectic exceptional sphere is isotopic through symplectically embedded spheres to the image of a (unique) J -holomorphic sphere.*

This technical lemma was originally proved for closed symplectic 4-manifolds in [41, Lemma 3.1] (see also [77, Theorem 5.1]), and we will prove it at the end of Section 1.3. By intersection positivity of J -holomorphic curves, a J -holomorphic exceptional sphere is the unique J -holomorphic curve in its homology class. We will denote the unique J -holomorphic representative of $e \in \mathcal{E}_\omega$ by E_e .

Definition 1.3.13. Let $\Theta = \{(\Theta_i, m_i)\}$ be an orbit set for $(-\partial X_0, \lambda_A)$. An element $\mathcal{C} \in \mathcal{M}(\emptyset, \Theta)$ is called *good* if the only exceptional sphere components have multiplicity one.

If there is a good element representing the class A then

$$e \cdot A \geq -1 \quad \forall e \in \mathcal{E}_\omega$$

Indeed, if C is a connected J -holomorphic curve then $e \cdot [C] \geq 0$ (by positivity of intersections) unless $C = E_e$, for which $e \cdot [C] = -1$. Thus there are no good elements when

$$e \cdot A \leq -2$$

for at least one $e \in \mathcal{E}_\omega$. In light of this, we will first assume in Section 1.3 that $e \cdot A \geq -1$ for all $e \in \mathcal{E}_\omega$, and then we will relax the assumption in Section 1.3.

When there are no multiply covered exceptional spheres

In this section we assume that $e \cdot A \geq -1$ for all $e \in \mathcal{E}_\omega$. To construct the counts of J -holomorphic curves for Theorem 1.1.6, we first make the following choices:

- an integer $I \geq 0$,
- an integer $p \in \{0, \dots, I\}$ such that $I - p$ is even,
- an ordered set of p disjoint oriented loops $\bar{\eta} := \{\eta_1, \dots, \eta_p\} \subset X_0$,
- a set of $\frac{1}{2}(I - p)$ disjoint points $\bar{z} := \{z_1, \dots, z_{(I-p)/2}\} \subset X_0 - \bar{\eta}$.

Denote by $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ the subset of elements in $\mathcal{M}_I(\emptyset, \Theta)$ which represent the class A and intersect all points \bar{z} and all loops $\bar{\eta}$. Define the chain

$$\Phi_{J, \bar{z}}^I(A, \bar{\eta}) := \sum_{\Theta} \sum_{\mathcal{C} \in \mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})} q(\mathcal{C}) \cdot \Theta \in \bigotimes_{k=1}^N ECC_*(S^1 \times S^2, \lambda_A, 1) \quad (1.11)$$

where Θ indexes over the admissible orbit sets, and $q(\mathcal{C}) \in \mathbb{Z}$ are weights that will be specified later in Section 1.3. This is the chain which is to be used for Theorem 1.1.6.

Since the existence of curves in $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ implies

$$\mathcal{A}(\Theta) \leq \rho(A)$$

there are only finitely many orbit sets Θ which can arise in $\Phi_{J, \bar{z}}^I(A, \bar{\eta})$. Therefore, the fact that this chain is well-defined (for the given choices of data $J, \bar{z}, \bar{\eta}$) follows from the following proposition.⁸

Proposition 1.3.14. *In the above setup, for generic J , the moduli space $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ is a finite set for each admissible orbit set Θ .*

Before we prove this proposition we must figure out what the curves in $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ look like. Given two distinct connected somewhere injective curves C_1 and C_2 , denote by $C_1 \cdot C_2$ the algebraic count of intersections. In our analysis we will make use of a “self-intersection” number that depends on $\rho(A)$ and stems from [26, Definition 4.7]. Given a connected somewhere injective curve C , define

$$C \cdot C := \frac{1}{2} [2g(C) - 2 + \text{ind}(C) + h(C) + 2e_A(C) + 4\delta(C)]$$

where g denotes its genus, e_A denotes the total multiplicity of all $\rho(A)$ -positive elliptic orbits, and h denotes the number of ends at hyperbolic orbits. Note that when C has no punctures, $C \cdot C$ is the algebraic count of self-intersections.

Next, it is crucial to know that for any multiply covered curve arising in our analysis, its Fredholm index is nonnegative. This is granted by the following lemma.

Lemma 1.3.15. *Suppose J is generic and $f : \tilde{C} \rightarrow C$ is a d -fold branched cover (with b branch points) of a somewhere-injective J -holomorphic curve $C \rightarrow \bar{X}$, where the ends of C are asymptotic to either positive hyperbolic orbits or d -positive elliptic orbits. Then $\text{ind}(\tilde{C}) \geq b$. Furthermore, if \tilde{C} is connected and C is a Fredholm index 0 plane whose end is asymptotic to a d -positive elliptic orbit, then $\text{ind}(\tilde{C}) > 0$ unless $f = \mathbb{1}$.*

Proof. The trivialization τ is assumed to be chosen so that the Conley-Zehnder indices of the orbits are as specified in Section 1.2. Let e (and \tilde{e}) denote the number of ends of C (and \tilde{C}) at elliptic orbits, and let b denote the number of branch points of f . Then $\text{ind}(C) = -\chi(C) + 2c_\tau(C) - e$ which is nonnegative by genericity of J , and

$$\text{ind}(\tilde{C}) = b - d \cdot \chi(C) + 2d \cdot c_\tau(C) - \tilde{e} = d \cdot \text{ind}(C) + b + (de - \tilde{e})$$

which is nonnegative because $\tilde{e} \leq de$.

⁸Proposition 1.3.14 is the analog of [68, Proposition 7.1] (also [54, Proposition 4.3]) for closed symplectic manifolds and [28, Lemma 5.10] for symplectizations.

If C is the plane satisfying the hypotheses in the statement of the lemma (in particular, $e = 1$), and g denotes the genus of \tilde{C} , then the Riemann-Hurwitz formula implies

$$b = d + 2g - 2 + \tilde{e}$$

and hence

$$\text{ind}(\tilde{C}) = 2(d + g - 1) \geq 0$$

which is zero if and only $f = 1$. \square

Finally, there will be certain pseudoholomorphic curves in \overline{X} that require special attention. We give them a name:

Definition 1.3.16. A J -holomorphic curve in \overline{X} is *special* if it has Fredholm/ECH index zero, and is either an embedded torus or an embedded plane whose negative end is asymptotic to an embedded elliptic orbit with multiplicity one.

Proposition 1.3.17. *For generic J , every current in $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ is good and takes the following form. Its underlying components are embedded, pairwise disjoint, and pairwise do not both have negative ends at covers of the same elliptic orbit. An embedded component which intersects l of the points in \bar{z} and l' of the loops in $\bar{\eta}$ has (ECH and Fredholm) index $2l + l'$. A component can be multiply covered only when it is special.*

Proof. Decompose a given current as

$$\mathcal{C} = \{(C_k, d_k)\} \cup \{(E_\sigma, m_\sigma)\}$$

where σ indexes the components which are (covers of) exceptional spheres. Then [26, Proposition 4.8] implies

$$\begin{aligned} I = I(\mathcal{C}) \geq \sum_k d_k I(C_k) + \sum_k d_k(d_k - 1) C_k \cdot C_k - \sum_\sigma m_\sigma(m_\sigma - 1) + \sum_{k \neq k'} d_k d_{k'} C_k \cdot C_{k'} \\ + 2 \sum_{\sigma, k} m_\sigma d_k E_\sigma \cdot C_k \end{aligned} \quad (1.12)$$

Since C_k is not an exceptional sphere and Θ contains only positive hyperbolic and $\rho(A)$ -positive elliptic orbits,

$$C_k \cdot C_k \geq 2\delta(C_k) \quad (1.13)$$

Indeed, C_k cannot be a plane whose end is at a positive hyperbolic orbit because $\text{ind}(C_k)$ would then be odd hence nonzero.

Because $A = \sum_k d_k[C_k] + \sum_\sigma m_\sigma[E_\sigma]$, the inequality (1.12) can be rewritten as

$$I \geq \sum_k d_k I(C_k) + 2 \sum_k d_k(d_k - 1) \delta(C_k) + \sum_\sigma m_\sigma(m_\sigma + 1) + \sum_{k \neq k'} d_k d_{k'} C_k \cdot C_{k'} + 2 \sum_\sigma m_\sigma [E_\sigma] \cdot A$$

By the index inequality $\text{ind}(C_k) \leq I(C_k) - 2\delta(C_k)$ and the assumption $e \cdot A \geq -1$ for all $e \in \mathcal{E}_\omega$, we can simplify the inequality further as

$$I \geq \sum_k d_k \text{ind}(C_k) + 2 \sum_k d_k^2 \delta(C_k) + \sum_\sigma m_\sigma(m_\sigma - 1) + \sum_{k \neq k'} d_k d_{k'} C_k \cdot C_{k'} \quad (1.14)$$

If C_k intersects $l \geq 0$ of the base points and $l' \geq 0$ of the loops then $\text{ind}(C_k) \geq 2l + l'$. The inequality (1.14) is thus an equality,

$$I = \sum_k d_k \text{ind}(C_k) + 2 \sum_k d_k^2 \delta(C_k) + \sum_\sigma m_\sigma(m_\sigma - 1) + \sum_{k \neq k'} d_k d_{k'} C_k \cdot C_{k'} \quad (1.15)$$

In particular, there are no negative index curves nor nodal curves.⁹ The remaining properties of the curves stated in the proposition can be read off from (1.15), and [26, Proposition 4.8] implies that C_i and C_j for $i \neq j$ do not both have negative ends at covers of the same $\rho(A)$ -positive elliptic orbit. To clarify the components which are multiply covered when $I(C_k) = 0$, note that the equality (1.12) implies

$$0 = C_k \cdot C_k$$

if $d_k > 1$. A zero self-intersection can only hold if C_k is a torus or an embedded plane with its end at an elliptic orbit with multiplicity one (or an embedded cylinder with ends at distinct hyperbolic orbits, but such cylinders cannot be multiply covered thanks to admissibility of Θ). \square

For the upcoming proof of Proposition 1.3.14, it will be crucial to know that a generic cobordism-admissible J satisfies certain Fredholm regularity properties. We state these properties as a definition.

Definition 1.3.18. For a given cobordism-admissible J , the moduli space $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ is said to be *cut out transversally* if the following is true for every current \mathcal{C} in the moduli space: For a component $(C, 1) \in \mathcal{C}$ its deformation operator D_C is surjective, and for a component $(C, d) \in \mathcal{C}$ with $d > 1$ the pull-back of D_C to any unbranched cover $f : \tilde{C} \rightarrow C$ with $\deg(f) \leq d$ is injective.

It turns out that regularity automatically holds for special planes (see Lemma 1.3.23). Likewise, Taubes was able to establish regularity for the special tori in his work on the Gromov invariants: The proof of [68, Proposition 7.1] (see also the proof of [54, Lemma 5.4]) explains how to perturb J on special tori to achieve regularity, with one caveat. Taubes requires J to be perturbed on the entire image of each torus, but the images of the tori appearing in $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ can intersect the symplectization ends of \bar{X} , for which J has additional constraints that do not appear in Taubes' work. Thankfully, this requirement of Taubes is weakened in [79]. In particular, the perturbations can be made local to the curve in the cobordism region X_0 .

⁹The existence of nodal curves is a priori possible because the symplectic form is not exact.

Proof of Proposition 1.3.14. We follow the proof of [28, Lemma 5.10] which argues that the ECH differential ∂_{ECH} is well-defined. Suppose otherwise that there are infinitely many such currents in $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$. By Gromov compactness for currents, any sequence of currents in $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ has a subsequence converging to a possibly broken ECH index I current from \emptyset to Θ , with the cobordism level intersecting the points \bar{z} and the loops $\bar{\eta}$. The symplectization levels of the broken holomorphic current do not have any closed components, thanks to exactness of the symplectic form (a sequence of tori sliding off the ends of \bar{X} would have been a source of noncompactness). The cobordism level \mathcal{C} is an element of $\mathcal{M}(\emptyset, \Theta'; A', \bar{z}, \bar{\eta})$ where Θ' is a potentially inadmissible orbit set that satisfies $\mathcal{A}(\Theta') \leq \rho(A)$, and $I \geq I(\mathcal{C})$. But the proof of Proposition 1.3.17 shows that in fact $I = I(\mathcal{C})$ and that \mathcal{C} has no nodes. With nodes now excluded, no exceptional spheres arise in the limit if they did not originally exist.

To show that the broken current is in fact unbroken, we will make use of SFT compactness (for curves) in lieu of Gromov compactness (for currents) as follows. We claim that there are only finitely many possibilities for the multiply covered components in our given sequence of currents. Assuming this claim for the moment, we can restrict to a subsequence $\{\mathcal{C}_\nu\} \subset \mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ so that the multiply covered components are the same. The remaining components of the subsequence are embedded and asymptotic to a fixed orbit subset $\Theta'' \subseteq \Theta$, and by the above convergence as currents we can restrict to a further subsequence so that they represent the same relative homology class in $H_2(\bar{X}, \emptyset, \Theta'')$. It then follows from [30, Corollary 6.10, Proposition 6.14] that there is a uniform bound on the genus of the embedded components, so we can restrict to a further subsequence so that the topological type of the embedded components is fixed. Now we invoke SFT compactness to obtain a convergent subsequence that converges as curves to a broken J -holomorphic curve.

By additivity of the ECH index, all symplectization levels have ECH index 0. Note that a sequence of closed components cannot break, because the only $I = 0$ curves in a symplectization are covers of \mathbb{R} -invariant cylinders. Using additivity of the Fredholm index, in lieu of nonnegativity of the Fredholm index by Lemma 1.3.15, the proof of [21, Lemma 7.19] concludes that there is in fact only one level (that is, there cannot exist symplectization levels of the broken curve consisting solely of nontrivially branched covers of \mathbb{R} -invariant cylinders). Thus the limiting curve is unbroken.

To complete the argument that $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ is a finite set, it remains to prove: 1) the aforementioned claim, and 2) that the limiting curve is isolated.

As for (1), if there were infinitely many possible multiply covered components (of a current representing the class A) then there would be infinitely many special tori or special planes (the underlying components). A formal repeat of the previous paragraphs shows that an infinite subsequence of such embedded curves must converge to an unbroken Fredholm index 0 curve. So it suffices to prove that this limiting curve is isolated, which is the goal of (2).

As for (2), such isolation is only violated if a sequence of embedded curves $\{C_k\}$ converges to a multiply covered curve $f : \tilde{C}_\infty \rightarrow C_\infty$, where C_∞ is a special torus or special plane. And this only occurs if $\text{ind}(C_k) = \text{ind}(\tilde{C}_\infty) = 0$ for all k , otherwise there would be point

constraints on C_k and hence point constraints on C_∞ , contradicting the fact that $\text{ind}(C_\infty) = 0$. Therefore, \tilde{C}_∞ is an unbranched cover of C_∞ by Lemma 1.3.15. Now, such a convergent sequence $\{C_k\}$ would produce a nonzero element in the kernel of the pull-back D_f of the deformation operator D_{C_∞} , but this contradicts the fact that $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ is cut out transversally for generic J . \square

We end this section with a proof of Lemma 1.3.12. We postponed the proof until now because it makes convenient use of the proof of Proposition 1.3.14.

Proof of Lemma 1.3.12. We follow the proof of [77, Theorem 5.1]. Given a symplectic exceptional sphere in the class $e \in \mathcal{E}_\omega$, we can always construct a (probably non-generic) cobordism-admissible almost complex structure J_0 on (\bar{X}, ω) which is integrable in a neighborhood of the sphere. This J_0 -holomorphic sphere has Fredholm/ECH index zero and is automatically transverse (see Lemma 1.3.23). Now we deform J_0 to $J_1 = J$ using a generic smooth 1-parameter family $\{J_s\}_{s \in [0,1]}$ of cobordism-admissible almost complex structures and deduce what happens to the J_0 -holomorphic sphere. That is, we analyze the parametric moduli space $\mathcal{M}_0(e, \{J_s\})$ consisting of pairs (C, s) such that $s \in [0, 1]$ and C is the unique J_s -holomorphic sphere representing e if it exists.

Due to the existence of the J_0 -holomorphic exceptional sphere we know that

$$c_1(e) = \chi(S^2) + e \cdot e = 1$$

and therefore all J_s -holomorphic spheres representing e are embedded, have Fredholm/ECH index 0, and are automatically transverse (see Lemma 1.3.23). It follows that the natural projection map $\mathcal{M}_0(e, \{J_s\}) \rightarrow [0, 1]$ is a submersion. It suffices to show that $\mathcal{M}_0(e, \{J_s\})$ is compact, for then there exists a (unique) J_s -holomorphic sphere representing e for every $s \in [0, 1]$. Compactness fails only if a sequence of J_s -holomorphic exceptional spheres (varying s) becomes nodal or breaks. But we already know that there are no nodal spheres representing e , and the proof of Proposition 1.3.14 can be copied verbatim to show that there are no broken curves representing e . \square

Allowing multiply covered exceptional spheres

In this section we do *not* assume that $e \cdot A \geq -1$ for all $e \in \mathcal{E}_\omega$, so we cannot rule out the existence of multiply covered exceptional spheres. When (X, ω) is a closed symplectic manifold, McDuff [40] showed how Taubes' construction of the Gromov invariants could be modified to count multiple covers of exceptional spheres. What we show now is that McDuff's modification also works in our setting of (X_0, ω) . The idea is to add extra base points to the set \bar{z} from the previous section, which "pin down" any positive index component of a J -holomorphic curve that exists to offset the negative index sphere components.

For each $e \in \mathcal{E}_\omega$, define its *algebraic multiplicity* to be the quantity $m_e(A) := \max(-e \cdot A, 0)$. Then denote

$$m(A) := \sum_{e \in \mathcal{E}_\omega} (m_e(A)^2 - m_e(A))$$

We formally repeat Section 1.3 using the moduli space $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}', \bar{\eta})$, except \bar{z}' denotes

- a set of $\frac{1}{2}(I - p + m(A))$ disjoint points $\bar{z}' := \{z_1, \dots, z_{(I-p+m(A))/2}\} \subset X_0 - \bar{\eta}$.

If $e \cdot A < -1$ for some $e \in \mathcal{E}_\omega$ then this moduli space does not contain any good elements, i.e. every element has a holomorphic exceptional sphere component with multiplicity at least 2 (see Section 1.3 and Lemma 1.3.12). Each such sphere does not pass through any base point in \bar{z}' , and the goal of this section is to demonstrate that we can effectively “isolate and ignore” such spheres. To start, we can mimic the proofs of Proposition 1.3.14 and Proposition 1.3.17 to show:

Proposition 1.3.19. *In the above setup, for generic J , the moduli space $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}', \bar{\eta})$ is a finite set for each admissible orbit set Θ . For each current in such a moduli space, its underlying components are embedded, pairwise disjoint, and pairwise do not both have negative ends at covers of the same elliptic orbit. A component which intersects l of the points in \bar{z}' and l' of the loops in $\bar{\eta}$ has (ECH and Fredholm) index $2l + l'$. A component can be multiply covered only when it is special or an exceptional sphere. For each $e \in \mathcal{E}_\omega$ such that $m_e(A) > 0$, every current must have $(E_e, m_e(A))$ as a component.*

Proof. Decompose a given current in $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}', \bar{\eta})$ as

$$\mathcal{C} = \{(C_k, d_k)\} \cup \{(E_\sigma, m_\sigma)\} \cup \{(F_{\sigma'}, m_{\sigma'})\}$$

where $E_\sigma \cdot A < -1$ and $F_{\sigma'} \cdot A \geq -1$ for $E_\sigma, F_{\sigma'} \in \mathcal{E}_\omega$. The proofs of Proposition 1.3.14 and Proposition 1.3.17 are only affected by dropping the constraint on A that $e \cdot A \geq -1$ for all $e \in \mathcal{E}_\omega$, for which (1.12) now implies

$$I = I(\mathcal{C}) \geq \sum_k d_k \text{ind}(C_k) + \sum_\sigma m_\sigma(m_\sigma + 1) + 2 \sum_\sigma m_\sigma [E_\sigma] \cdot A \quad (1.16)$$

It suffices to show that the right-hand-side of (1.16) is bounded below by I . Due to the point and loop constraints, the first term on the right-hand-side of (1.16) is bounded below by

$$2 \cdot \frac{1}{2}(I - p + m(A)) + p = I + m(A)$$

Note that, by definition, the third term on the right-hand-side of (1.16) is

$$-2 \sum_\sigma m_\sigma \cdot m_{[E_\sigma]}(A)$$

Therefore, it suffices to show that

$$m(A) + \sum_\sigma m_\sigma(m_\sigma + 1) - 2 \sum_\sigma m_\sigma \cdot m_{[E_\sigma]}(A) \geq 0$$

By definition of $m(A)$, we are reduced to showing that

$$m_{[E_\sigma]}(A)^2 - m_{[E_\sigma]}(A) + m_\sigma^2 + m_\sigma - 2m_\sigma \cdot m_{[E_\sigma]}(A) \geq 0 \quad (1.17)$$

for each index σ . Now, the left-hand-side of (1.17) is precisely

$$(m_\sigma - m_{[E_\sigma]}(A))^2 + (m_\sigma - m_{[E_\sigma]}(A))$$

so it suffices to show that

$$m_\sigma \geq m_{[E_\sigma]}(A)$$

for each index σ . To show this, we look at the decomposition of $A = [\mathcal{C}]$ into the components of \mathcal{C} and compute $[E_\sigma] \cdot A$ to get

$$-m_{[E_\sigma]}(A) = -m_\sigma + [E_\sigma] \cdot ([\mathcal{C}] - m_\sigma [E_\sigma]) \geq -m_\sigma$$

where the latter inequality follows from positivity of intersections of J -holomorphic curves. The desired result follows. \square

We now rephrase the result of Proposition 1.3.19 to remove the multiply covered exceptional spheres; this is the analog of [40, Lemma 3.3] for closed symplectic 4-manifolds. Consider the relative class

$$A' := A + \sum_{e \in \mathcal{E}_\omega \mid e \cdot A < -1} (e \cdot A) e$$

which satisfies $e \cdot A' \geq -1$ for all $e \in \mathcal{E}_\omega$, and the nonnegative integer

$$I' := I + m(A)$$

which has the same parity as I . Since each current in $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}', \bar{\eta})$ consists of a good element representing $A' \in \text{Rel}_\omega(X)$ and the same collection of disjoint multiply covered exceptional spheres representing $-\sum_{e \in \mathcal{E}_\omega \mid e \cdot A < -1} (e \cdot A) e$, we have the following corollary.

Corollary 1.3.20. *In the scenario of Proposition 1.3.19, there is a one-to one correspondence*

$$\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}', \bar{\eta}) \longleftrightarrow \mathcal{M}_{I'}(\emptyset, \Theta; A', \bar{z}', \bar{\eta}) \quad (1.18)$$

which simply ignores the exceptional sphere components having multiplicity greater than one.

Therefore, in the definition of the chain (1.11) we may use the finite sets $\mathcal{M}_{I'}(\emptyset, \Theta; A', \bar{z}', \bar{\eta})$.

Remark 1.3.21. It may be possible to establish a blow-up formula by counting J -holomorphic curves on \overline{X} and any of its blow-downs, along the lines of [37, Theorem 4.2] for closed symplectic manifolds, but the details have not been sorted out. Instead, the blow-up formula follows as a corollary from the blow-up formula for the Seiberg-Witten invariants (see Chapter 2).

Orientations and weights

This section defines the integer weight $q(\mathcal{C})$ attached to a current $\mathcal{C} = \{(C_k, d_k)\}$ which belongs to a given moduli space $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$, for a generic choice of J . This moduli space is assumed to be cut out transversally (see Definition 1.3.18) and $e \cdot A \geq -1$ for all $e \in \mathcal{E}_\omega$, but see Remark 1.3.25 for the case of multiply covered exceptional spheres. The total weight is of the form

$$q(\mathcal{C}) = \varepsilon(\mathcal{C}) \prod_k r(C_k, d_k) \in \mathbb{Z}$$

where $r(C, d)$ is an integer weight attached to each component (C, d) , and $\varepsilon(\mathcal{C})$ is a global sign which depends on the ordering of both the set $\bar{\eta}$ and the set \mathcal{P} of positive hyperbolic orbits in Θ . We specify these below.

Remark 1.3.22. Remember that Θ is an orbit set for the disjoint union of N copies of $S^1 \times S^2$. In order to identify ECH of this disjoint union with the tensor product of N copies of ECH of $S^1 \times S^2$, we need to choose an ordering of the N copies of $S^1 \times S^2$ (which is an ordering of the zero-circles of ω). Once this ordering is made we can decompose \mathcal{P} as $\mathcal{P}_1 \sqcup \dots \sqcup \mathcal{P}_N$, where \mathcal{P}_k is the set of positive hyperbolic orbits in the orbit set Θ_k for the k^{th} copy of $S^1 \times S^2$, and $(\Theta, \mathfrak{o}) = \pm(\Theta_1, \mathfrak{o}_1) \otimes \dots \otimes (\Theta_N, \mathfrak{o}_N)$. Here, the sign \pm is determined by whether the ordering of \mathcal{P} defined by \mathfrak{o} (dis)agrees with the ordering of $\mathcal{P}_1 \sqcup \dots \sqcup \mathcal{P}_N$ defined by $\mathfrak{o}_1 \otimes \dots \otimes \mathfrak{o}_N$.

Given orientations of the admissible orbit sets (\mathfrak{o} of Θ), each moduli space $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ is “coherently oriented” with the conventions in [22, §9.5] (based off of [2]). In particular, $\text{Ker } D_C$ is oriented for every component C of \mathcal{C} .

We first define $r(C, 1)$, following [54, §2] and [24, §2.5]. Recall that the set $\bar{\eta}$ is ordered, and labelled accordingly as $\{\eta_1, \dots, \eta_p\}$. The curve C intersects l of the loops for some $l \in \{0, \dots, p\}$, say $\{\gamma_1, \dots, \gamma_l\}$ written in the order that they appear in $\bar{\eta}$. Denote the intersection points by $\{w_1, \dots, w_l\}$. The curve C then must intersect $\frac{1}{2}(I(C) - l)$ points in \bar{z} , which without loss of generality are the points $\{z_1, \dots, z_{(I(C)-l)/2}\}$. From this data we build the $\mathbb{R}^{I(C)}$ -vector space

$$V_C := \bigoplus_{i=1}^{\frac{1}{2}(I(C)-l)} N_{z_i} \oplus \bigoplus_{i=1}^l \left(N_{w_i} / \pi(T_{w_i} \gamma_i) \right)$$

where N_x denotes the fiber of the normal bundle N_C of C over the point $x \in C$, and π denotes the canonical projection of $T\bar{X}$ onto N_C . Then V_C is oriented because the normal bundle of C is oriented, the loops γ_i are oriented, and the points w_i are ordered; the ordering of the points z_i does not matter because each N_{z_i} is 2-dimensional. For generic J , the restriction map

$$\text{Ker}(D_C) \rightarrow V_C$$

is an isomorphism, and $r(C, 1) = \pm 1$ depending on whether this restriction map is orientation-preserving or orientation-reversing.

Before defining $r(C, d)$ when $d > 1$, it will be useful to write out $r(C, 1)$ in the cases where $I(C) = 0$ and C is either an exceptional sphere, a special torus, or a special plane. Then $r(C, 1)$ is the modulo 2 count of spectral flow of a generic path of Fredholm first-order operators from the deformation operator D_C to a complex linear operator. The following “automatic transversality” lemma will help us compute this spectral flow.

Lemma 1.3.23. *Let C be a connected immersed J -holomorphic curve in \overline{X} with ends at nondegenerate Reeb orbits $\{\gamma_j\}$, let g denote the genus of C , and let h_+ denote the number of ends of C at positive hyperbolic orbits (including even multiples of negative hyperbolic orbits). Given any Cauchy-Riemann type operator*

$$D : L_1^2(N_C) \rightarrow L^2(T^{0,1}C \otimes N_C)$$

that is asymptotic to the fixed asymptotic operators L_{γ_j} of C , if

$$2g - 2 + h_+ < \text{ind}(D)$$

then D is surjective. In particular, if the inequality holds for the deformation operator D_C then C is transverse without any genericity assumption on J .

See [76, Proposition 2.2] for a proof of this lemma and the technical definition of “a Cauchy-Riemann type operator that is asymptotic to an asymptotic operator.” In the case where C is an exceptional sphere, Lemma 1.3.23 implies $r(C, 1) = 1$. In the case where C is a special torus, Lemma 1.3.23 unfortunately does not say anything. In the case where C is a special plane, its asymptotic operator over the elliptic orbit is L -flat and hence complex linear. Thus its deformation operator has the form

$$D_C = \bar{\partial} + \nu_C + \mu_C$$

with complex anti-linear term μ_C asymptotic to zero along the end of C . We then apply Lemma 1.3.23 to the path of Cauchy-Riemann type operators¹⁰

$$r \in [0, 1] \mapsto \bar{\partial} + \nu_C + (1 - r) \cdot \mu_C \tag{1.19}$$

The asymptotic operators of (1.19) are all the same, so the Fredholm index remains constant. Thus there is no spectral flow when deforming D_C to a complex linear operator and hence $r(C, 1) = 1$.

We now consider $r(C, d)$ when $d > 1$. Here, C is necessarily either a special torus or a special plane. In the case where C is a torus, the explicit description of $r(C, d)$ is found in [54, Definition 3.2] and will not be repeated here. Suffice to say, the weight assigned

¹⁰We stated “automatic transversality” for operators that aren’t necessarily deformation operators of J -holomorphic curves because it may not be possible for the path of operators (1.19) to come from a path of admissible J .

to a multiply covered torus is determined by the spectral flows of four operators on C and guarantees that Taubes' Gromov invariants are well-defined. Along a path of almost complex structures, multiple covers of tori may “pop off” to nearby honest curves, or two tori with opposing sign may collide and annihilate, and this must be accounted for in order to obtain an actual invariant.

In the case where C is a plane, we define

$$r(C, d) := r(C, 1) = 1$$

because it turns out that C is “automatically super-rigid.” To elaborate, given any connected (branched) multiple cover $f : \tilde{C} \rightarrow C$ and deformation operator $D_C = \bar{\partial} + \nu_C + \mu_C$ there is an induced pull-back deformation operator

$$D_f = \bar{\partial} + f^* \nu_C + f^* \mu_C$$

from $\Gamma(f^* N_C)$ to $\Gamma(T^{0,1} \tilde{C} \otimes f^* N_C)$. This Fredholm operator satisfies $\text{ind}(D_f) = \text{ind}(\tilde{C}) - 2b$, where b is the number of branch points of f and $\text{ind}(\tilde{C})$ is the Fredholm index of the J -holomorphic curve $\tilde{C} \rightarrow C \rightarrow \overline{X}$ (see [25, 79]). Granted that $\text{ind}(C) = 0$, we say that C is *d-nondegenerate* if $\text{Ker}(D_f) = 0$ for all unbranched covers of C of degree no greater than d , and that C is *super-rigid* if $\text{Ker}(D_f) = 0$ for all (branched) covers of C . For a special plane C , similarly to the setup when $d = 1$, we can deform D_f to a complex linear operator along a path of Cauchy-Riemann type operators such that the Fredholm index stays constant. The following lemma is a version of “automatic transversality” for such operators over \tilde{C} (see [25] or [76, Proposition 2.2] for a proof).

Lemma 1.3.24. *Let C be a connected immersed J -holomorphic curve in \overline{X} with ends at nondegenerate Reeb orbits, let $f : \tilde{C} \rightarrow C$ be a branched cover with b branch points, \tilde{g} denote the genus of \tilde{C} , and let \tilde{h}_+ denote the number of ends of \tilde{C} at positive hyperbolic orbits (including even multiples of negative hyperbolic orbits). Given any Cauchy-Riemann type operator*

$$D : L_1^2(f^* N_C) \rightarrow L^2(T^{0,1} \tilde{C} \otimes f^* N_C)$$

that is asymptotic to fixed asymptotic operators of \tilde{C} and satisfies $\text{ind}(D) = \text{ind}(\tilde{C}) - 2b$, if

$$2\tilde{g} - 2 + \tilde{h}_+ + \text{ind}(\tilde{C}) - 2b < 0$$

then D is injective.

For a special plane C and a given cover f ,

$$\text{ind}(\tilde{C}) = 2 - 2\tilde{g} + 2b - e - \sum_{i=1}^e CZ_\tau(\gamma_i)$$

where e is the number of punctures of \tilde{C} such that its i^{th} end is asymptotic to the elliptic orbit γ_i , and the trivialization τ is chosen so that the Conley-Zehnder indices are as specified in Section 1.2. Thus from Lemma 1.3.24 we see that not only is C super-rigid but there is no spectral flow when deforming D_f to a complex linear operator for any f .

Remark 1.3.25. Section 1.3 concerns the case that $e \cdot A < -1$ for some $e \in \mathcal{E}_\omega$. Define

$$r(E_e, m_e(A)) := r(E_e, 1) = 1$$

where E_e denotes the unique holomorphic sphere representative of e such that $m_e(A) \geq 1$, which makes the identification of the moduli spaces in the correspondence (1.18) orientation-preserving.

It remains to define $\varepsilon(\mathcal{C})$. First, note that an even (respectively, odd) Fredholm index component of \mathcal{C} has an even (respectively, odd) number of ends asymptotic to positive hyperbolic orbits, and it intersects an even (respectively, odd) number of loops. Second, note that an ordering of the components of \mathcal{C} determines a partition of \mathcal{P} whose ordering differs from the fixed ordering of \mathcal{P} by a permutation $\sigma_{\mathcal{P}}$, and it also determines a partition of $\bar{\eta}$ whose ordering differs from the fixed ordering of $\bar{\eta}$ by a permutation $\sigma_{\bar{\eta}}$. Third, note that any permutation σ has a sign $\varepsilon(\sigma)$ given by its parity. Then

$$\varepsilon(\mathcal{C}) := \varepsilon(\sigma_{\mathcal{P}})\varepsilon(\sigma_{\bar{\eta}})$$

which does not depend on the ordering of the components of \mathcal{C} .

Remark 1.3.26. Our definition of $\varepsilon(\mathcal{C})$ is consistent with that in Taubes' construction of the Gromov invariants [54, §2]. There are no Reeb orbits to deal with for closed symplectic 4-manifolds, so $\varepsilon(\mathcal{C})$ reduces to the sign associated with the ordering of the set of loops $\bar{\eta}$. That sign is well-defined, i.e. independent of the ordering of the components of \mathcal{C} , because every component of \mathcal{C} is closed and thus has even Fredholm index.

Equations for chain maps and chain homotopies

The goal of this section is to show that $\Phi_{J, \bar{z}}^I(A, \bar{\eta})$ defines an element in ECH, and to clarify its dependence on \bar{z} and $\bar{\eta}$. In light of Corollary 1.3.20, assume $e \cdot A \geq -1$ for all $e \in \mathcal{E}_\omega$.

Proposition 1.3.27. *For generic J , the chain $\Phi_{J, \bar{z}}^I(A, \bar{\eta})$ is a cycle,*

$$\partial_{\text{ECH}} \circ \Phi_{J, \bar{z}}^I(A, \bar{\eta}) = 0 \tag{1.20}$$

The proof relies on a gluing theorem, for which we need some information about the asymptotics of the curves that are to be glued. The multiplicities of the positive (respectively, negative) ends of a connected curve asymptotic to a given orbit γ defines a positive (respectively, negative) partition of the total multiplicity m of γ . A key fact is that a nontrivial component of a holomorphic current that contributes to the ECH differential or the U -map satisfies the *partition conditions* [28, §3.9], which means the partitions $p_\gamma^\pm(m)$ are uniquely determined by the orbit γ . For example, if γ is elliptic with rotation class $\theta \in (0, \frac{1}{m})$ mod 1 then $p_\gamma^+(m) = (1, \dots, 1)$ and $p_\gamma^-(m) = (m)$.

Proof of Proposition 1.3.27. To prove the chain map equation (1.20), we fix an admissible orbit set Θ and analyze the ends of $\mathcal{M}_{I+1}(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$. A broken curve arising as a limit of curves in $\mathcal{M}_{I+1}(\emptyset, \Theta; A, \bar{z}, \bar{\eta})$ cannot contain a multiply covered *punctured* component in the cobordism level, except for a cover of a special plane or an embedded cylinder¹¹ (asymptotic to hyperbolic orbits). This follows from the arguments in the proof of Proposition 1.3.14, for which the cobordism level \mathcal{C} now satisfies

$$I + 1 \geq I(\mathcal{C}) \geq I$$

If there were no such multiply covered punctured components, the proof of $\partial_{\text{ECH}}^2 = 0$ in [21] could be copied to conclude that the number of signed gluings is 1 in the case needed to prove (1.20). That is, the proof of [21, Lemma 7.23] would carry over verbatim to show that a broken curve consists of

- a Fredholm/ECH index I curve in the cobordism level (intersecting \bar{z} and $\bar{\eta}$),
- a Fredholm/ECH index 1 curve in the symplectization level, and
- possibly additional levels between them consisting of *connectors*, i.e.

branched covers of \mathbb{R} -invariant cylinders.

The obstruction bundle “gluing analysis” used to prove [21, Theorem 7.20] would then prove (1.20).

Remark 1.3.28. Multiple covers of exceptional spheres do not arise in this process: two curves having nonnegative index cannot glue together to form a negative index curve. This is consistent with the fact that the exceptional spheres are rigid and a sequence of such spheres do not break.

Therefore, it suffices to show that there is no further obstruction bundle gluing needed when our multiply covered curves are introduced. To start, we note that such multiple covers must be unbranched, otherwise the Fredholm index would be too big (see Lemma 1.3.15). Next, we note that an unbranched cover of a plane (respectively, cylinder) is necessarily disjoint copies of a plane (respectively, cylinder), thanks to the Riemann-Hurwitz formula. Let’s first analyze the cylinders, and then the planes:

For each $d > 1$ we can compute the number of ways to glue a d -fold cover of a cylinder using the same reasoning as in the proof of $\partial_{\text{ECH}}^2 = 0$ which computes the number of gluings for two embedded ECH index 1 curves along hyperbolic orbits. In particular, it follows from [21, Lemma 1.7] that there are no connectors, otherwise their Fredholm index would be too big. Thus the multiply covered cylinder would have to glue directly to a curve in the symplectization level, for which the positive partition conditions are satisfied. Thus there are either no ways to glue or there are an even number of ways to glue, given by the number of permutations of the ends of the cylinders that glue (see [21, §1.5]). And as explained in [21,

¹¹There can exist multiple covers of index 0 embedded cylinders at hyperbolic orbits, because the orbit set between the levels of the broken curve need not be an *admissible* orbit set.

Remark 1.5], half of those gluings have one sign and the other half have the other sign, so the total signed count of gluings is zero.

For each $d > 1$ we claim that there is exactly one way to glue a d -fold cover of a special plane. If this multiply covered plane were to glue to an ECH index 1 curve with a connector inbetween, then the gluing of the multiply covered plane and the connector would be another (multiply covered) plane. But the only way this glued curve could have Fredholm index 0 is for the connector to be a disjoint union of \mathbb{R} -invariant cylinders, thanks to Lemma 1.3.15, so there is no nontrivial gluing. Thus the multiply covered plane would have to glue directly to a curve in the symplectization level, for which the positive partition conditions are satisfied. That means the ends of the multiply covered plane satisfy $p_\gamma^+(d) = (1, \dots, 1)$, where γ is the elliptic orbit which the underlying plane is asymptotic to, and therefore there is exactly one possible gluing (the d disjoint copies of the plane are indistinguishable, so there are no permutations for matching the ends of the curves). \square

The cycle $\Phi_{J, \bar{z}}^I(A, \bar{\eta})$ a priori depends on the choice of loops $\bar{\eta}$ and points \bar{z} (and J). We now show that each loop in $\bar{\eta}$ can move around in its homotopy class and each point in \bar{z} can move into the ends of \bar{X} , without affecting the induced homology class in ECH. Our argument follows [24, §2.5] which argues that the U -map does not depend on the choice of base point.

Proposition 1.3.29. *Fix an element $[\bar{\eta}] \in \Lambda^p(H_1(X; \mathbb{Z})/\text{Torsion})$ and an ordering of the zero-circles of ω . For generic J , the homology class $Gr_{X, \omega, J}^I(A, [\bar{\eta}])$ represented by the chain $\Phi_{J, \bar{z}}^I(A, \bar{\eta})$ does not depend on the choice of \bar{z} nor on the choice of representative $\bar{\eta} \subset X_0$ for $[\bar{\eta}]$. Furthermore,*

$$Gr_{X, \omega, J}^I(A, [\bar{\eta}]) = U^{(I-p)/2} \circ Gr_{X, \omega, J}^p(A, [\bar{\eta}]) \in \bigotimes_{k=1}^N ECH_*(S^1 \times S^2, \xi_{\text{Taubes}}, 1)$$

where $U^{(I-p)/2}$ denotes $\frac{1}{2}(I-p)$ compositions of any of the U -maps from components of ∂X_0 .

Proof. Let Θ be an admissible orbit set such that $\mathcal{A}(\Theta) \leq \rho(A)$. Let $\{y_1, \dots, y_{(I-p)/2}\} \subset \partial X_0$ be a collection of base points which do not lie on any Reeb orbit. Pair each $y_k \in \partial X_0$ with $z_k \in \bar{z}$ and choose embedded paths $\gamma_k : [0, 1] \rightarrow X_0$ from z_k to y_k such that the image of γ_k only intersects ∂X_0 in y_k . Define the chains

$$K_{\gamma_k} := \sum_{\Theta} \sum_{\mathcal{C} \in \mathcal{M}_{I-1}(\emptyset, \Theta; A, \bar{z} - z_k, \bar{\eta}, \gamma_k)} q(\mathcal{C}) \cdot \Theta \in \bigotimes_{k=1}^N ECC_*(S^1 \times S^2, \lambda_A, 1)$$

where Θ indexes over the admissible orbit sets. The proofs of Proposition 1.3.14 and Proposition 1.3.17 also work in this setting to show that K_{γ_k} is well-defined: the decrement $I \rightarrow I-1$ and the path constraint of γ_k is compensated by the removal of the point constraint of z_k .

We can check that

$$\partial_{\text{ECH}} \circ K_{\gamma_k} = \Phi_{J, \bar{z}}^I(A, \bar{\eta}) - U_{y_k} \circ \Phi_{J, \bar{z} - z_k}^{I-2}(A, \bar{\eta}) \quad (1.21)$$

by counting ends and boundary points of the moduli space $\mathcal{M}_I(\emptyset, \Theta; A, \bar{z} - z_k, \bar{\eta}, \gamma_k)$, using the same gluing analysis as in the proof of Proposition 1.3.27. Passing to homology and iterating through the recursive equation (1.21) for $1 \leq k \leq \frac{1}{2}(I - p)$, the desired result concerning the U -map follows.

A similar argument yields the independence of the loop η_k in X_0 that represents $[\bar{\eta}_k]$. Note here that $H_1(X; \mathbb{Z}) \cong H_1(X_0; \mathbb{Z})$, which follows from the homological long exact sequence for the pair (X, X_0) in lieu of

$$H_k(X, X_0; \mathbb{Z}) \cong H_k(\text{cl } \mathcal{N}, \partial \mathcal{N}; \mathbb{Z}) \cong H^{4-k}(\text{cl } \mathcal{N}; \mathbb{Z}) \cong H^{4-k}(Z; \mathbb{Z})$$

These isomorphisms are given respectively by excision, Poincaré-Lefschetz duality, and Z being a deformation retraction of $\text{cl } \mathcal{N}$; these homologies are trivial for $k \in \{1, 2\}$. \square

Gradings

The remaining statement to be proved in Theorem 1.1.6 is the fact that the element $\text{Gr}_{X, \omega, J}^I(A, [\bar{\eta}])$, despite being a sum over many orbit sets, is concentrated in a single grading of ECH.

Proposition 1.3.30. $\text{Gr}_{X, \omega, J}^I(A, [\bar{\eta}]) \in \text{ECH}_{g(A, I)}(-\partial X_0, \xi_{\text{Taubes}}, 1)$, where the grading $g(A, I)$ is determined by A and I . In terms of the canonical absolute $\mathbb{Z}/2\mathbb{Z}$ grading on ECH, the parity of $g(A, I)$ is equal to the parity of I .

Proof. We prove this in a slightly more general scenario. Consider a symplectic cobordism (X, ω) from the empty set to a contact 3-manifold (Y, λ) , a homology class $\Gamma \in H_1(Y; \mathbb{Z})$ satisfying

$$c_1(\xi) + 2 \text{PD}(\Gamma) = 0$$

and a relative homology class $A \in H_2(X, -Y; \mathbb{Z})$ satisfying $\partial A = -\Gamma$. Let α and β be admissible orbit sets on (Y, λ) in the class Γ . Let \mathcal{C} and \mathcal{C}' be J -holomorphic currents in the completion \bar{X} which represent A and are asymptotic to α and β , respectively. Since \mathcal{C} and \mathcal{C}' both represent A , the difference

$$[\mathcal{C}'] - [\mathcal{C}] \in H_2(Y, \alpha, \beta)$$

can be used to measure the grading difference

$$|\alpha| - |\beta| = I([\mathcal{C}'] - [\mathcal{C}]) = I(\mathcal{C}') - I(\mathcal{C})$$

where the last equality follows from additivity of the ECH index. Therefore, $I(\mathcal{C}') = I(\mathcal{C})$ if and only if $|\alpha| = |\beta|$.

For an admissible orbit set Θ on (Y, λ) , let $\varepsilon(\Theta)$ be the number of positive hyperbolic orbits in Θ , which determines the parity of $|\Theta|$ as a canonical absolute $\mathbb{Z}/2\mathbb{Z}$ grading. The “index parity” formula [28, Equation 3.7] states that

$$I(\mathcal{C}) \equiv \varepsilon(\Theta) \pmod{2}$$

for $\mathcal{C} \in \mathcal{M}_I(\emptyset, \Theta)$. Therefore, the parities of $g(A, I)$ and I agree. \square

As mentioned in Section 1.1, a given spin-c structure $\mathfrak{s} \in \text{Spin}^c(X)$ determines the relative class $A = A_{\mathfrak{s}} \in H_2(X_0, \partial X_0; \mathbb{Z})$ and also the index $I = d(\mathfrak{s})$. It will be shown in Chapter 2 that $g(A_{\mathfrak{s}}, d(\mathfrak{s}))$ is equal to $N[\xi_*]$, i.e.

$$Gr_{X, \omega, J}^{d(\mathfrak{s})}(A_{\mathfrak{s}}, [\bar{\eta}]) \in \bigotimes_{k=1}^N ECH_{[\xi_*]}(S^1 \times S^2, \xi_{\text{Taubes}}, 1) \cong \mathbb{Z}$$

This fact, $g(A_{\mathfrak{s}}, d(\mathfrak{s})) = N[\xi_*]$, makes sense for a few reasons. First, $[\xi_*]$ has odd parity while N has parity equal to that of $b_+^2(X) - b^1(X) + 1$, which is also the parity of $d(\mathfrak{s})$, so the parity of $N[\xi_*]$ agrees with that of $d(\mathfrak{s})$ hence $g(A_{\mathfrak{s}}, d(\mathfrak{s}))$. Second, it follows from Proposition 1.3.2 that $ECH_j(-\partial X_0, \xi_{\text{Taubes}})$ is a single copy of \mathbb{Z} if and only if $j = N[\xi_*]$.

1.4 Appendix

In this thesis we have assumed that Z consists of only untwisted zero-circles. As we now clarify, a straightforward modification allows us to include twisted zero-circles as long as they are non-contractible in X .

In the presence of a twisted zero-circle, the corresponding boundary component of (X_0, ω) is $(S^1 \times S^2, e^{-1} \lambda_{\text{Taubes}}^\sigma)$. Here, $\lambda_{\text{Taubes}}^\sigma$ is the pushforward of λ_{Taubes} under the double covering map $S^1 \times S^2 \rightarrow S^1 \times S^2$, for which the nontrivial deck transformation is the fixed-point-free involution

$$\sigma(t, \theta, \varphi) = (t + \pi, \pi - \theta, -\varphi)$$

Since λ_{Taubes} is σ -invariant, the orbits in $(S^1 \times S^2, \lambda_{\text{Taubes}})$ descend to the orbits in $(S^1 \times S^2, \lambda_{\text{Taubes}}^\sigma)$, so that $\lambda_{\text{Taubes}}^\sigma$ is a Morse-Bott contact form. The images of the exceptional orbits coincide, giving a single exceptional elliptic orbit with rotation class $\sqrt{\frac{3}{2}} \pmod{1}$. The images of the tori $T(\theta_0)$ for $\theta_0 \neq \frac{\pi}{2}$ are also tori foliated by orbits. However, the image of the torus $T(\frac{\pi}{2})$ is a Klein bottle: it is described as a closed interval family of orbits, whose endpoints lift to orbits in $(S^1 \times S^2, \lambda_{\text{Taubes}})$ that are fixed set-wise by σ . These endpoints are double covers of orbits that have one-half the period of the other orbits in the interval.

The contact structure $\xi_{\text{Taubes}}^\sigma$ associated with $\lambda_{\text{Taubes}}^\sigma$ is also overtwisted. Although ξ_{Taubes} and $\xi_{\text{Taubes}}^\sigma$ are not homotopic over $S^1 \times S^2$ [15, Theorem 10], they have the same Euler class: $e(\xi_{\text{Taubes}}^\sigma)$ can be computed using the pushforward of the section $\sin \theta \partial_\theta \in \Gamma(\xi_{\text{Taubes}})$.

Thus ξ_{Taubes} and $\xi_{\text{Taubes}}^\sigma$ are homotopic over the 2-skeleton of $S^1 \times S^2$ and their corresponding spin-c structures are the same. Subsequently,

$$ECH_j(S^1 \times S^2, \xi_{\text{Taubes}}^\sigma, 1) \cong \widehat{HM}^j(S^1 \times S^2, \mathfrak{s}_\xi + 1) \cong ECH_j(S^1 \times S^2, \xi_{\text{Taubes}}, 1)$$

and

$$J(S^1 \times S^2, \mathfrak{s}_\xi) = \{[\xi_{\text{Taubes}}], [\xi_{\text{Taubes}}^\sigma]\}$$

and the homotopy class $[\xi_*]$ has even parity under the canonical absolute $\mathbb{Z}/2\mathbb{Z}$ grading on $ECH_*(S^1 \times S^2, \xi_{\text{Taubes}}^\sigma, 1)$.

The analog of Lemma 1.3.9 now says that the relevant modification of $\lambda_{\text{Taubes}}^\sigma$ for a given class $A \in \text{Rel}_\omega(X)$ is a nondegenerate contact form λ_A whose orbits of action less than $\rho(A)$ are either $\rho(A)$ -positive elliptic, positive hyperbolic, or negative hyperbolic. Indeed, Bourgeois' perturbation breaks up the Klein bottle of orbits into two doubly covered negative hyperbolic orbits and an embedded elliptic orbit with rotation number slightly positive [4, §2.2, §9.5], and the exceptional orbit of $\lambda_{\text{Taubes}}^\sigma$ can be appropriately modified because the function (1.9) constructed in the proof of Lemma 1.3.7 is σ -invariant.

The existence of these negative hyperbolic orbits may cause problems. The problem with the proof of Proposition 1.3.14 for $\lambda_{\text{Taubes}}^\sigma$ is that the inequality (1.13) is false for an embedded plane C asymptotic to a negative hyperbolic orbit. In particular, the current (C, d) for $d > 1$ has negative ECH index $-\frac{1}{2}d(d-1)$ and may arise in the cobordism level of a broken current. Of course, if the twisted zero-circles are non-contractible in X then these planes cannot exist.

Chapter 2

Seiberg-Witten and Gromov invariants

2.1 Introduction

Continuing with the setup of the previous chapter, (X, g) denotes a closed connected oriented smooth Riemannian 4-manifold with $b_+^2(X) \geq 1$. Furthermore, ω denotes a self-dual near-symplectic form on X whose zero set $Z := \omega^{-1}(0)$ has $N \geq 0$ components, all of which are untwisted zero-circles.¹ Such 2-forms always exist, but the parity of N must be the same as that of $b_+^2(X) - b^1(X) + 1$.

In Chapter 1 we defined the near-symplectic Gromov invariants

$$Gr_{X,\omega} : \text{Spin}^c(X) \rightarrow \Lambda^* H^1(X; \mathbb{Z})$$

in terms of counts of pseudoholomorphic curves in a certain completion of $X - Z$, and it a priori depends on an almost complex structure on $X - Z$. Before we recall how this invariant is defined, we state the main theorem of this paper (from which it follows that $Gr_{X,\omega}$ over $\mathbb{Z}/2\mathbb{Z}$ is indeed a smooth invariant of X).

Theorem 2.1.1. *Given (X, ω) as above and $\mathfrak{s} \in \text{Spin}^c(X)$,*

$$Gr_{X,\omega}(\mathfrak{s}) = SW_X(\mathfrak{s}) \in \Lambda^* H^1(X; \mathbb{Z}) \otimes \mathbb{Z}/2\mathbb{Z}$$

where ω determines the chamber for defining the Seiberg-Witten invariants when $b_+^2(X) = 1$.

Remark 2.1.2. As explained in Section 2.3, a choice of homology orientation of X is the same thing as a choice of ordering of the zero-circles of ω plus a choice of homology orientation of the cobordism obtained from X by removing tubular neighborhoods of the zero-circles. We expect that there is a canonical homology orientation one on the cobordism determined by ω , and that Theorem 2.1.1 can be lifted to \mathbb{Z} coefficients.

¹We may also allow components of Z to be twisted zero-circles which are non-contractible in X (see Section 1.4), but we restrict our attention to untwisted zero-circles for simplicity of notation.

We now explain the setup for defining $Gr_{X,\omega}$, as it sets the stage for the rest of this paper. Let \mathcal{N} denote the union of arbitrarily small tubular neighborhoods of the components of $Z \subset X$, chosen in such a way that the complement

$$(X_0, \omega) := (X - \mathcal{N}, \omega|_{X - \mathcal{N}})$$

is a symplectic manifold with contact-type boundary, where each boundary component is a copy of $(S^1 \times S^2, \xi_{\text{Taubes}})$. We will specify the contact form on $S^1 \times S^2$ momentarily.

It follows from Taubes' work on near-symplectic geometry that ω induces an $H_2(X; \mathbb{Z})$ -equivariant map

$$\tau_\omega : \text{Spin}^c(X) \rightarrow H_2(X_0, \partial X_0; \mathbb{Z}), \quad \mathfrak{s} \mapsto \text{PD}(c_1(E)) \quad (2.1)$$

where $E \rightarrow X_0$ is the complex line bundle that defines the decomposition of the positive spinor bundle associated with the restricted spin-c structure $\mathfrak{s}|_{X_0}$,

$$\mathbb{S}_+(\mathfrak{s}|_{X_0}) = E \oplus K^{-1}E$$

and $K \rightarrow X_0$ is the canonical bundle determined by ω ; see the upcoming Section 2.3 for an elaboration. The following lemma shows that τ_ω gives a canonical identification of $\text{Spin}^c(X)$ with the set

$$\text{Rel}_\omega(X) := \left\{ A \in H_2(X_0, \partial X_0; \mathbb{Z}) \mid \partial A = \mathbb{1} \in H_1(\partial X_0; \mathbb{Z}) \right\}$$

where $\mathbb{1}$ is the oriented generator on each component (the orientation conventions are specified in Section 1.3).

Lemma 2.1.3. *The map τ_ω is injective, and its image consists of the subset of relative homology classes whose boundary is the oriented generator of $H_1(\partial X_0; \mathbb{Z})$, i.e. for each \mathfrak{s} on X*

$$\partial \tau_\omega(\mathfrak{s}) = -(1, \dots, 1) \in -\bigoplus_{k=1}^N H_1(S^1 \times S^2; \mathbb{Z})$$

Proof. The restriction map $\text{Spin}^c(X) \rightarrow \text{Spin}^c(X_0)$ is injective, or in terms of the cohomology actions, the restriction map $H^2(X; \mathbb{Z}) \rightarrow H^2(X_0; \mathbb{Z})$ is injective. This follows from the cohomological long exact sequence applied to the pair (X, X_0) because

$$H^2(X, X_0; \mathbb{Z}) \cong H^2(\text{cl } \mathcal{N}, \partial \mathcal{N}; \mathbb{Z}) \cong H_2(\text{cl } \mathcal{N}; \mathbb{Z}) = 0$$

using excision and Poincaré-Lefschetz duality. Then τ_ω is injective, since

$$\tau_\omega(\mathfrak{s} \otimes E) - \tau_\omega(\mathfrak{s}) = \text{PD } c_1(E|_{X_0})$$

for any complex line bundle $E \rightarrow X$.

Likewise, the determinant line bundle $\det(\mathbb{S}_+(\mathfrak{s}|_{\partial X_0}))$ is trivial. In terms of the cohomology actions, the restriction map $H^2(X; \mathbb{Z}) \rightarrow H^2(\partial X_0; \mathbb{Z})$ is trivial because it factors through $H^2(\text{cl } \mathcal{N}; \mathbb{Z}) = 0$. On each boundary component, this constraint

$$0 = c_1 \left(\det(\mathbb{S}_+(\mathfrak{s}|_{S^1 \times S^2})) \right) = 2c_1(E|_{S^1 \times S^2}) + c_1(K^{-1}|_{S^1 \times S^2})$$

implies

$$c_1(E|_{S^1 \times S^2}) = 1 \in \mathbb{Z} \cong H^2(S^1 \times S^2; \mathbb{Z})$$

because $K^{-1}|_{S^1 \times S^2} = \xi_{\text{Taubes}}$. \square

We now fix a spin-c structure $\mathfrak{s} \in \text{Spin}^c(X)$ and choose \mathcal{N} so that $-\partial X_0 = \partial \mathcal{N}$ is a contact 3-manifold whose contact form is $\lambda_{\mathfrak{s}}$ on each component, as provided in Lemma 1.3.9. Here, $\lambda_{\mathfrak{s}}$ is precisely $\lambda_{\tau_{\omega}(\mathfrak{s})}$ in the notation of Lemma 1.3.9, which is a particular rescaling of Taubes' overtwisted contact form λ_{Taubes} but whose orbits of symplectic action less than $\rho(\tau_{\omega}(\mathfrak{s}))$ are all $\rho(\tau_{\omega}(\mathfrak{s}))$ -flat. The quantity $\rho(\tau_{\omega}(\mathfrak{s})) \in \mathbb{R}$ is spelled out in Section 1.3 and the notion of “flatness” is spelled out in Section 1.2. (The importance of “flatness” becomes evident in Theorem 2.4.5.)

The component of the element $Gr_{X, \omega}(\mathfrak{s})$ in $\Lambda^p H^1(X; \mathbb{Z})$ is defined to be zero if $d(\mathfrak{s}) - p$ is odd or negative, where

$$d(\mathfrak{s}) := \frac{1}{4} (c_1(\mathfrak{s})^2 - 2\chi(X) - 3\sigma(X))$$

and otherwise it is determined by its evaluation on $[\eta_1] \wedge \cdots \wedge [\eta_p]$ for a given ordered set of classes $[\bar{\eta}] := \{[\eta_i]\}_{i=1}^p \subset H_1(X; \mathbb{Z})/\text{Torsion}$. In order to construct this number

$$Gr_{X, \omega}(\mathfrak{s})([\eta_1] \wedge \cdots \wedge [\eta_p]) \in \mathbb{Z}$$

we must first introduce the set $\mathcal{E}_{\omega} \subset H_2(X_0, \partial X_0; \mathbb{Z})$ of classes represented by symplectic exceptional spheres in X_0 (which is empty if X is minimal). Then we fix the following data:

- an ordered set of p disjoint oriented loops $\bar{\eta} \subset X_0$ which represent $[\bar{\eta}]$,
- a set of $\frac{1}{2}(d(\mathfrak{s}) - p + m(\tau_{\omega}(\mathfrak{s})))$ disjoint points $\bar{z} \subset X_0 - \bar{\eta}$,
- a cobordism-admissible almost complex structure J on the completion (\bar{X}, ω) of (X_0, ω) .

Here, $m(\tau_{\omega}(\mathfrak{s}))$ denotes the “algebraic multiplicity” of $\tau_{\omega}(\mathfrak{s})$ as defined in Section 1.3, which equals 0 if $e \cdot \tau_{\omega}(\mathfrak{s}) \geq -1$ for all $e \in \mathcal{E}_{\omega}$ (such as when X is minimal).

Given this data, we form the moduli space $\mathcal{M}_{d(\mathfrak{s})}(\emptyset, \Theta; \tau_{\omega}(\mathfrak{s}), \bar{z}, \bar{\eta})$ of J -holomorphic currents in \bar{X} that satisfy the following properties:

- they are asymptotic to the admissible orbit set Θ ,
- they have ECH index $d(\mathfrak{s})$,
- they represent $\tau_{\omega}(\mathfrak{s})$
- they intersect every point and loop in $\bar{z} \cup \bar{\eta}$.

It turns out that this moduli space is a finite set for generic J (see Proposition 1.3.14) and each Θ satisfies $\mathcal{A}(\Theta) < \rho(\tau_{\omega}(\mathfrak{s}))$ whenever the moduli space is nonempty. Moreover, each

Θ has an absolute grading as a generator of the ECH chain complex $ECC_*(-\partial X_0, \lambda_{\mathfrak{s}}, 1)$ and they are all the same grading (see Section 1.3), denoted $g(\mathfrak{s})$. Thus we can form the *Gromov cycle*²

$$\Phi_{Gr} := \sum_{\Theta \in g(\mathfrak{s})} \mathcal{M}_\Theta \Theta \in ECC_{g(\mathfrak{s})}(-\partial X_0, \lambda_{\mathfrak{s}}, 1) \quad (2.2)$$

where

$$\mathcal{M}_\Theta := \sum_{\mathcal{C} \in \mathcal{M}_{d(\mathfrak{s})}(\emptyset, \Theta; \tau_\omega(\mathfrak{s}), \bar{z}, \bar{\eta})} q(\mathcal{C}) \in \mathbb{Z}$$

As a further reminder, each orbit set Θ comes equipped with a choice of orientation so that the moduli spaces are all coherently oriented, and $q(\mathcal{C})$ is a certain integer weight associated to each current \mathcal{C} (see Section 1.3).

Now, $ECH_*(-\partial X_0, \lambda_{\mathfrak{s}}, 1)$ is the tensor product $\bigotimes_{k=1}^N ECH_*(S^1 \times S^2, \lambda_{\mathfrak{s}}, 1)$. In terms of the absolute grading on $ECH_*(S^1 \times S^2, \lambda_{\mathfrak{s}}, 1)$ by homotopy classes of oriented 2-plane fields on $S^1 \times S^2$, there is a unique class $[\xi_*]$ such that

$$ECH_{[\xi_*]}(S^1 \times S^2, \lambda_{\mathfrak{s}}, 1) \cong \mathbb{Z}$$

while $ECH_{[\xi_*]+n}(S^1 \times S^2, \lambda_{\mathfrak{s}}, 1) = 0$ when $n < 0$ (see Proposition 1.3.2). In the proof of Theorem 2.7.1 it will be shown that $g(\mathfrak{s}) = N[\xi_*]$. That said, $Gr_{X, \omega}(\mathfrak{s})([\eta_1] \wedge \dots \wedge [\eta_p])$ is by definition the coefficient of the class

$$[\Phi_{Gr}] \in ECH_{g(\mathfrak{s})}(-\partial X_0, \xi_{\text{Taubes}}, 1) \cong \bigotimes_{k=1}^N ECH_{[\xi_*]}(S^1 \times S^2, \xi_{\text{Taubes}}, 1) \quad (2.3)$$

as a multiple of the positive generator $1 \in \bigotimes_{k=1}^N ECH_{[\xi_*]}(S^1 \times S^2, \xi_{\text{Taubes}}, 1)$.

What follows is an outline of the remainder of this chapter. In Section 2.2 we briefly describe an application of Theorem 2.1.1 to 3-manifolds. Then we review the relevant gauge theory in Section 2.3, followed by a review of Taubes' isomorphisms between Floer homologies in Section 2.4. In Section 2.5 we reduce Theorem 2.1.1 to the case that there are no multiply covered exceptional spheres, via the “blow-up” formula in Seiberg-Witten theory. The reason for taking this route to prove Theorem 2.1.1 is that we do not know how to directly relate the moduli spaces of pseudoholomorphic curves with the moduli spaces of Seiberg-Witten solutions in the presence of multiply covered exceptional spheres. In Section 2.6 we relate the Gromov cycle to a formally similar “Seiberg-Witten cocycle,” by constructing a correspondence between the moduli spaces of pseudoholomorphic curves and Seiberg-Witten solutions on \bar{X} . This mimics the correspondence from Section 2.4 but the story is much more complicated, the reason being that the existence of multiply covered tori and planes prevents an honest bijection between moduli spaces of pseudoholomorphic curves and Seiberg-Witten solutions (see Section 2.6). In Section 2.7, we relate the “Seiberg-Witten cocycle” (and thus

²In the notation of Chapter 1, $\Phi_{Gr} = \Phi_{J, \bar{z}}^{d(\mathfrak{s})}(\tau_\omega(\mathfrak{s}), \bar{\eta})$.

the Gromov cycle) to the Seiberg-Witten invariant of X , by “stretching the neck” along the zero-circles of the near-symplectic form. Finally, the appendix briefly clarifies some of the differences and complications that occur in Taubes’ constructions of the 3-dimensional isomorphisms (between Floer homologies) versus the 4-dimensional equivalences (between invariants).

2.2 S^1 -valued Morse theory

A basic example of a near-symplectic manifold is $(S^1 \times M, \omega_f)$, where M is a closed oriented Riemannian 3-manifold with $b^1(M) > 0$, the metric on $S^1 \times M$ is the product metric $dt^2 + g_M$, and ω_f is defined momentarily. A result of Honda [14] and Calabi [6] says that for g_M suitably generic, any nonzero class in $H^1(M; \mathbb{Z})$ is represented by a harmonic map $f : M \rightarrow S^1$ (i.e. $d^*df = 0$) with nondegenerate critical points $\text{crit}(f)$ of index 1 or 2, hence a harmonic 1-form df with transversal zeros. Then

$$\omega_f := dt \wedge df + *_3 df$$

is a closed self-dual 2-form which vanishes transversally on $Z_f := S^1 \times \text{crit}(f)$. All zero-circles are untwisted, as can be seen by writing out ω_f in local coordinates and comparing to the standard model on $\mathbb{R} \times \mathbb{R}^3$. There are an even number of zero-circles, i.e. $b_+^2(X) + b^1(M) + 1$ is even, because $b^1(X) = b^1(M) + 1$ and $b^2(X) = 2b^1(M)$ and $b_+^2(X) = b^1(M)$. Here, we note that

$$H^1(M; \mathbb{R}) \rightarrow H_+^2(X; \mathbb{R}), \quad a \mapsto [dt \wedge a]^+ = \frac{1}{2}(dt \wedge a + *_3 a)$$

is an isomorphism.

After equipping $(S^1 \times M) - Z_f = S^1 \times (M - \text{crit}(f))$ with the compatible almost complex structure J determined by ω_f and $dt^2 + g_M$, the S^1 -invariant connected J -holomorphic submanifolds are of the form $C = S^1 \times \gamma$, where γ is a single gradient flowline of ∇f . Since the Morse trajectories in M are either periodic orbits (of some period) or paths between critical points, C is either a torus (with multiplicity) or a cylinder which bounds two zero-circles in $S^1 \times M$ (see Figure 2.1).

In their PhD theses, Hutchings and Lee built a 3-dimensional invariant $I_{M,f}$ of M which suitably counts the gradient flowlines (see [19, 18, 29]), and they further showed that it equals a version of topological (Reidemeister) torsion defined by Turaev [75]. It was subsequently shown by Turaev that this Reidemeister torsion equals the 3-dimensional SW invariant SW_M of M (see [74]). Strictly speaking, there is a required choice of “homology orientation” on M with which to define SW_M , a “chamber” (determined by f) with which to define SW_M when $b^1(M) = 1$, and an ordering of the set $\text{crit}(f)$ with which to define $I_{M,f}$. While SW_M is a function of the set $\text{Spin}^c(M)$ of spin-c structures, $I_{M,f}$ is a function of the set

$$\text{Rel}_f(M) := \{\eta \in H_1(M, \text{crit}(f); \mathbb{Z}) \mid \partial\eta = [\text{crit}(f)]\}$$

and there exists an $H_1(M)$ -equivariant isomorphism τ_f between them [19, Lemma 4.3]. In other words,

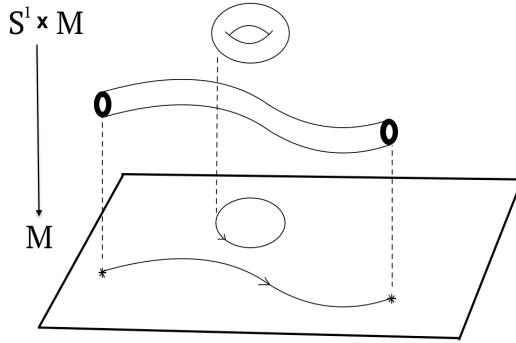


Figure 2.1: Zero set of near-symplectic form in bold

Theorem 2.2.1 (Hutchings-Lee-Turaev). *Let (M, f) be as above, and fix an ordering of the critical points of f . Then*

$$I_{M,f}(\tau_f(\mathfrak{s})) = \pm SW_M(\mathfrak{s}) \in \mathbb{Z}$$

for all $\mathfrak{s} \in \text{Spin}^c(M)$, where f determines the chamber for defining SW_M when $b^1(M) = 1$. The \pm sign is pinned down by a suitable choice of homology orientation of M and ordering of $\text{crit}(f)$.

Now, we can recover Theorem 2.2.1 over $\mathbb{Z}/2\mathbb{Z}$ without having to pass through Reidemeister torsion, via a dimensional reduction of Theorem 2.1.1.³ It was shown in [43, Theorem 3.5] that the 4-dimensional SW invariant recovers the 3-dimensional SW invariant: all solutions to the SW equations on $S^1 \times M$ associated with product spin-c structures are S^1 -invariant. It will be shown elsewhere [11] that the near-symplectic Gromov invariant recovers Hutchings-Lee's flowline invariant.

Corollary 2.2.2 ([11]). *Let (M, f) be as above, and let $\pi : S^1 \times M \rightarrow M$ be the projection map onto the second factor. Fix an ordering of the critical points of f (hence of the zero-circles of ω_f) and a homology orientation of M (hence of $S^1 \times M$). Then*

$$I_{M,f}(\mathfrak{s}) = Gr_{S^1 \times M, \omega_f}(\pi^* \mathfrak{s}) \equiv_{(2)} SW_{S^1 \times M}(\pi^* \mathfrak{s}) = SW_M(\mathfrak{s})$$

for all $\mathfrak{s} \in \text{Spin}^c(M)$. When $b^1(M) = b_+^2(S^1 \times M) = 1$, the chamber is determined by f (hence ω_f).

In fact, when f has no critical points (hence ω_f is symplectic) this was already known [18, Remark 1.10]. The first instance appeared in [48] for the special case of a mapping torus of a symplectomorphism of a Riemann surface, in which Salamon showed that the 3-dimensional SW invariants recover the Lefschetz invariants of the symplectomorphism.

We end this discussion with a brief sketch of the first equality in Corollary 2.2.2. While the gradient flowlines for $I_{M,f}$ sit inside M , the J -holomorphic curves for $Gr_{S^1 \times M, \omega_f}$ do not

³This was in fact an expectation and also the motivation, as explained in [19, §4.2.1].

sit inside $S^1 \times M$ but rather inside the completion of the complement of Z_f . Explicitly, we choose a 3-ball neighborhood $\bigsqcup_k B^3$ of the critical points of f , hence a tubular neighborhood $\mathcal{N} = S^1 \times \bigsqcup_k B^3$ of the zero-circles in Z_f . Then

$$X_0 := (S^1 \times M) - \mathcal{N} = S^1 \times (M - \bigsqcup_k B^3)$$

is a symplectic manifold and $H_2(X_0, \partial X_0; \mathbb{Z})$ is isomorphic to a direct sum of $|\text{crit}(f)|$ copies of

$$H_2(S^1 \times (M - B^3), S^1 \times S^2; \mathbb{Z}) \cong H_2(M - B^3, S^2; \mathbb{Z}) \oplus H_1(M - B^3, S^2; \mathbb{Z})$$

using the relative Künneth formula. The relative 1st homology class $\tau_f(\mathfrak{s}) \in H_1(M, \text{crit}(f); \mathbb{Z})$ has corresponding relative 2nd homology class $[S^1] \times \tau_f(\mathfrak{s}) = \tau_{\omega_f}(\pi^* \mathfrak{s}) \in H_2(X_0, \partial X_0; \mathbb{Z})$.

Since Taubes' contact form is S^1 -invariant, it follows from Chapter 1 that we can choose the 3-balls in such a way that $-\partial X_0$ is a contact boundary of X with ω_f equal to a scalar multiple of $d\lambda_{\text{Taubes}}$ on each component. Although the playground $(X_0, \omega_f, J, \tau_{\omega_f}(\pi^* \mathfrak{s}))$ is S^1 -invariant, the calculation of $Gr_{S^1 \times M, \omega_f}(\tau_{\omega_f}(\pi^* \mathfrak{s}))$ uses a modification Taubes' Morse-Bott contact form λ_{Taubes} (and the boundary ∂X_0) into a non- S^1 -invariant nondegenerate contact form $\lambda_{\mathfrak{s}}$, and so $(-\partial X_0, \lambda_{\mathfrak{s}})$ does not arise from any choice of the 3-balls. This is a complication, because we would like to lift flowlines $\gamma \subset M$ to S^1 -invariant curves $S^1 \times \gamma \subset \overline{X}$ and count them. Nonetheless, a limiting argument will show that we can perturb the S^1 -invariant setup to the non- S^1 -invariant setup and relate the corresponding pseudoholomorphic curves.

2.3 Review of gauge theory

The point of this section is to introduce most of the terminology and notations that appear in the later sections. Further information and more complete details are found in [33, 23].

Closed 3-manifolds

Let (Y, λ) be a closed oriented connected contact 3-manifold, and choose an almost complex structure J on ξ that induces a symplectization-admissible almost complex structure on $\mathbb{R} \times Y$. There is a compatible metric g on Y such that⁴ $|\lambda| = 1$ and $*\lambda = \frac{1}{2}d\lambda$, with $g(v, w) = \frac{1}{2}d\lambda(v, Jw)$ for $v, w \in \xi$.

View a spin-c structure \mathfrak{s} on Y as an isomorphism class of a pair (\mathbb{S}, cl) consisting of a rank 2 Hermitian vector bundle $\mathbb{S} \rightarrow Y$ and Clifford multiplication $\text{cl} : TY \rightarrow \text{End}(\mathbb{S})$. We refer to \mathbb{S} as the *spinor bundle* and its sections as *spinors*. The set $\text{Spin}^c(Y)$ of spin-c

⁴The factors of $\frac{1}{2}$ can be dropped or changed to any other nonzero real number by a particular rescaling of the metric, but they will be left in to be consistent with the papers of Taubes and Hutchings.

structures is an affine space over $H^2(Y; \mathbb{Z})$, defined by

$$(\mathbb{S}, \text{cl}) + x = (\mathbb{S} \otimes E^x, \text{cl} \otimes \mathbb{1})$$

where $E^x \rightarrow Y$ is the complex line bundle satisfying $c_1(E^x) = x \in H^2(Y; \mathbb{Z})$. Denote by $c_1(\mathfrak{s})$ the first Chern class of $\det \mathbb{S}$; it satisfies $c_1(\mathfrak{s} + x) = c_1(\mathfrak{s}) + 2x$.

The contact structure ξ (and more generally, any oriented 2-plane field on Y) picks out a canonical spin-c structure $\mathfrak{s}_\xi = (\mathbb{S}_\xi, \text{cl})$ with $\mathbb{S}_\xi = \underline{\mathbb{C}} \oplus \xi$, where $\underline{\mathbb{C}} \rightarrow Y$ denotes the trivial line bundle, and Clifford multiplication is defined as follows. Given an oriented orthonormal frame $\{e_1, e_2, e_3\}$ for $T_y Y$ such that $\{e_2, e_3\}$ is an oriented orthonormal frame for ξ_y , then in terms of the basis $(1, e_2)$ for \mathbb{S}_ξ ,

$$\text{cl}(e_1) = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}, \quad \text{cl}(e_2) = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad \text{cl}(e_3) = \begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}$$

There is then a canonical isomorphism

$$H^2(Y; \mathbb{Z}) \rightarrow \text{Spin}^c(Y), \quad x \mapsto (E^x \oplus (\xi \otimes E^x), \text{cl})$$

where the 0 class corresponds to \mathfrak{s}_ξ . In other words, there is a canonical decomposition $\mathbb{S} = E \oplus \xi E$ into $\pm i$ eigenbundles of $\text{cl}(\lambda)$. Here and in what follows, the tensor product notation is suppressed.

A *spin-c connection* is a connection \mathbf{A} on \mathbb{S} which is compatible with Clifford multiplication in the sense that

$$\nabla_{\mathbf{A}}(\text{cl}(v)\psi) = \text{cl}(\nabla v)\psi + \text{cl}(v)\nabla_{\mathbf{A}}\psi$$

where ∇v denotes the covariant derivative of $v \in TY$ with respect to the Levi-Civita connection. Such a connection is equivalent to a Hermitian connection (also denoted by \mathbf{A}) on $\det(\mathbb{S})$, and determines a *Dirac operator*

$$D_{\mathbf{A}} : \Gamma(\mathbb{S}) \xrightarrow{\nabla_{\mathbf{A}}} \Gamma(T^*Y \otimes \mathbb{S}) \xrightarrow{\text{cl}} \Gamma(\mathbb{S})$$

With respect to the decomposition $\mathbb{S} = E \oplus \xi E$, the determinant line bundle is $\det \mathbb{S} = \xi E^2$ and any spinor can be written as

$$\psi = (\alpha, \beta)$$

There is a unique connection A_ξ on ξ such that its Dirac operator kills the spinor $(1, 0) \in \Gamma(\mathbb{S}_\xi)$, and there is a canonical decomposition

$$\mathbf{A} = A_\xi + 2A$$

with Hermitian connection A on E . We henceforth refer to a spin-c connection as a Hermitian connection on E , and denote its Dirac operator by D_A .

The gauge group $C^\infty(Y, S^1)$ acts on a given pair (A, ψ) by

$$u \cdot (A, \psi) = (A - u^{-1}du, u\psi)$$

In this paper, a *configuration* \mathbf{c} refers to a gauge-equivalence class of such a pair, and the set of configurations is denoted by

$$\mathcal{B}(Y, \mathfrak{s}) := (\text{Conn}(E) \times \Gamma(\mathbb{S})) / C^\infty(Y, S^1)$$

Fix a suitably generic exact 2-form $\mu \in \Omega^2(Y)$ as described in [23, §2.2], and a positive real number $r \in \mathbb{R}$. A configuration \mathbf{c} solves *Taubes' perturbed SW equations* when

$$D_A \psi = 0, \quad *F_A = r(\tau(\psi) - i\lambda) - \frac{1}{2} *F_{A_\xi} + i * \mu \quad (2.4)$$

where F_{A_ξ} is the curvature of A_ξ and $\tau : \mathbb{S} \rightarrow iT^*Y$ is the quadratic bundle map

$$\tau(\psi)(\cdot) = \langle \text{cl}(\cdot)\psi, \psi \rangle$$

An appropriate change of variables recovers the usual SW equations (with perturbations) that appear in [33].

Remark 2.3.1. There are additional “abstract tame perturbations” to these equations required to obtain transversality of the moduli spaces of its solutions (see [33, §10]), but they have been suppressed because they do not interfere with the analysis presented in this paper. Further clarification on this matter can be found in [23, §2.1] and [55, §3.h Part 5], where the same suppression occurs.

Denote by $\mathfrak{M}(Y, \mathfrak{s})$ the set of solutions to (2.4), called *(SW) monopoles*. A solution is *reducible* if its spinor component vanishes, and is otherwise *irreducible*. After attaching orientations (this being a $\mathbb{Z}/2\mathbb{Z}$ choice for each monopole, see Section 2.3), the monopoles freely generate the SW Floer chain complex $\widehat{CM}^*(Y, \lambda, \mathfrak{s}, J, r)$. The chain complex differential will not be reviewed here. Of importance to this paper are irreducible monopoles with certain bounds on their energy

$$E(\mathbf{c}) := i \int_Y \lambda \wedge F_A$$

Denote by $\widehat{CM}_L^*(Y, \lambda, \mathfrak{s}, J, r)$ the submodule generated by irreducible monopoles \mathbf{c} with energy $E(\mathbf{c}) < 2\pi L$. When r is sufficiently large, $\widehat{CM}_L^*(Y, \lambda, \mathfrak{s}, J, r)$ is a subcomplex of $\widehat{CM}^*(Y, \lambda, \mathfrak{s}, J, r)$ and the homology $\widehat{HM}_L^*(Y, \lambda, \mathfrak{s}, J, r)$ is well-defined and independent of r and μ . Taking the direct limit over $L > 0$, we recover the ordinary $\widehat{HM}^*(Y, \mathfrak{s})$ in [33] which is independent of λ and J . It is sometimes convenient to consider the group $\widehat{HM}^*(Y) := \bigoplus_{\mathfrak{s} \in \text{Spin}^c(Y)} \widehat{HM}^*(Y, \mathfrak{s})$ over all spin-c structures at once.

Symplectic cobordisms

Let (X, ω) be a strong symplectic cobordism between (possibly disconnected or empty) closed oriented contact 3-manifolds (Y_\pm, λ_\pm) . Due to the choice of metric g_\pm on Y_\pm in

Section 2.3 (and following [23, §4.2]), we do not extend ω over \overline{X} using $d(e^s \lambda_{\pm})$ on the ends $(-\infty, 0] \times Y_-$ and $[0, \infty) \times Y_+$. Instead, we extend ω using $d(e^{2s} \lambda_{\pm})$ as follows. Fix a smooth increasing function $\phi_- : (-\infty, \varepsilon] \rightarrow (-\infty, \varepsilon]$ with $\phi_-(s) = 2s$ for $s \leq \frac{\varepsilon}{10}$ and $\phi_-(s) = s$ for $s > \frac{\varepsilon}{2}$, and fix a smooth increasing function $\phi_+ : [-\varepsilon, \infty) \rightarrow [-\varepsilon, \infty)$ with $\phi_+(s) = 2s$ for $s \geq -\frac{\varepsilon}{10}$ and $\phi_+(s) = s$ for $s \leq -\frac{\varepsilon}{2}$, where $\varepsilon > 0$ is such that $\omega = d(e^s \lambda_{\pm})$ on the ε -collars of Y_{\pm} . Then the desired extension is

$$\tilde{\omega} := \begin{cases} e^{\phi_- \lambda_-} & \text{on } (-\infty, \varepsilon] \times Y_- \\ \omega & \text{on } X \setminus \left(([0, \varepsilon] \times Y_-) \cup ([-\varepsilon, 0] \times Y_+) \right) \\ e^{\phi_+ \lambda_+} & \text{on } [-\varepsilon, \infty) \times Y_+ \end{cases} \quad (2.5)$$

Now choose a cobordism-admissible almost complex structure J on $(\overline{X}, \tilde{\omega})$. As in [23, §4.2], we can equip \overline{X} with a metric g so that it agrees with the product metric with g_{\pm} on the ends $(-\infty, 0] \times Y_-$ and $[0, \infty) \times Y_+$ and so that $\tilde{\omega}$ is self-dual. Finally, define $\widehat{\omega} := \sqrt{2} \tilde{\omega} / |\tilde{\omega}|_g$ and note that J is still cobordism-admissible.

The 4-dimensional gauge-theoretic scenario is analogous to the 3-dimensional scenario. View a spin-c structure \mathfrak{s} on X as an isomorphism class of a pair (\mathbb{S}, cl) consisting of a Hermitian vector bundle $\mathbb{S} = \mathbb{S}_+ \oplus \mathbb{S}_-$, where \mathbb{S}_{\pm} have rank 2, and Clifford multiplication $\text{cl} : TX \rightarrow \text{End}(\mathbb{S})$ such that $\text{cl}(v)$ exchanges \mathbb{S}_+ and \mathbb{S}_- for each $v \in TX$. We refer to \mathbb{S}_+ as the *positive spinor bundle* and its sections as *(positive) spinors*. The set $\text{Spin}^c(X)$ of spin-c structures is an affine space over $H^2(X; \mathbb{Z})$, and we denote by $c_1(\mathfrak{s})$ the first Chern class of $\det \mathbb{S}_+ = \det \mathbb{S}_-$. A spin-c connection on \mathbb{S} is equivalent to a Hermitian connection \mathbf{A} on $\det(\mathbb{S}_+)$ and defines a Dirac operator $D_{\mathbf{A}} : \Gamma(\mathbb{S}_{\pm}) \rightarrow \Gamma(\mathbb{S}_{\mp})$.

A spin-c structure \mathfrak{s} on X restricts to a spin-c structure $\mathfrak{s}|_{Y_{\pm}}$ on Y_{\pm} with spinor bundle $\mathbb{S}_{Y_{\pm}} := \mathbb{S}_+|_{Y_{\pm}}$ and Clifford multiplication $\text{cl}_{Y_{\pm}}(\cdot) := \text{cl}(v)^{-1} \text{cl}(\cdot)$, where v denotes the outward-pointing unit normal vector to Y_+ and the inward-pointing unit normal vector to Y_- . There is a canonical way to extend \mathfrak{s} over \overline{X} , and the resulting spin-c structure is also denoted by \mathfrak{s} . There is a canonical decomposition $\mathbb{S}_+ = E \oplus K^{-1}E$ into $\mp 2i$ eigenbundles of $\text{cl}_+(\widehat{\omega})$, where K is the canonical bundle of (\overline{X}, J) and $\text{cl}_+ : \bigwedge^2 T^* \overline{X} \rightarrow \text{End}(\mathbb{S}_+)$ is the projection of Clifford multiplication onto $\text{End}(\mathbb{S}_+)$. This agrees with the decomposition of $\mathbb{S}_{Y_{\pm}}$ on the ends of \overline{X} .

The symplectic form ω picks out the canonical spin-c structure $\mathfrak{s}_{\omega} = (\mathbb{S}_{\omega}, \text{cl})$, namely that for which E is trivial, and the $H^2(X; \mathbb{Z})$ -action on $\text{Spin}^c(X)$ becomes a canonical isomorphism. There is a unique connection $A_{K^{-1}}$ on K^{-1} such that its Dirac operator annihilates the spinor $(1, 0) \in \Gamma((\mathbb{S}_{\omega})_+)$, and we henceforth identify a spin-c connection with a Hermitian connection A on E and denote its Dirac operator D_A .

In this paper, a *configuration* \mathfrak{d} refers to a gauge-equivalence class of a pair (A, Ψ) under the gauge group $C^{\infty}(X, S^1)$ -action. A connection \mathbf{A} on $\det(\mathbb{S}_+)$ is in *temporal gauge* on the ends of \overline{X} if

$$\nabla_{\mathbf{A}} = \frac{\partial}{\partial s} + \nabla_{\mathbf{A}(s)}$$

on $(-\infty, 0] \times Y_-$ and $Y_+ \times [0, \infty)$, where $\mathbf{A}(s)$ is a connection on $\det(\mathbb{S}_{Y_\pm})$ depending on s . Any connection can be placed into temporal gauge by an appropriate gauge transformation. Given monopoles \mathbf{c}_\pm on Y_\pm , the set of configurations which are asymptotic to \mathbf{c}_\pm (in temporal gauge on the ends of \overline{X}) is denoted by

$$\mathcal{B}(\mathbf{c}_-, \mathbf{c}_+; \mathfrak{s}) \subset (\text{Conn}(E) \times \Gamma(\mathbb{S}_+)) / C^\infty(X, S^1)$$

Fix suitably generic exact 2-forms $\mu_\pm \in \Omega^2(Y_\pm)$, a suitably generic exact 2-form $\mu \in \Omega^2(\overline{X})$ that agrees with μ_\pm on the ends of \overline{X} (with μ_* denoting its self-dual part), and a positive real number $r \in \mathbb{R}$. *Taubes' perturbed SW equations* for a configuration \mathbf{d} are

$$D_A \Psi = 0, \quad F_A^+ = \frac{r}{2}(\rho(\Psi) - i\widehat{\omega}) - \frac{1}{2}F_{A_{K^{-1}}}^+ + i\mu_* \quad (2.6)$$

where F_A^+ is the self-dual part of the curvature of A and $\rho : \mathbb{S}_+ \rightarrow \bigwedge^2 T^*X$ is the quadratic bundle map

$$\rho(\Psi)(\cdot, \cdot) = -\frac{1}{2}\langle [\text{cl}(\cdot), \text{cl}(\cdot)]\Psi, \Psi \rangle$$

Similarly to the 3-dimensional equations, there are additional “abstract tame perturbations” which have been suppressed in this paper (see [33, §24.1]). Denote by $\mathfrak{M}(\mathbf{c}_-, \mathbf{c}_+; \mathfrak{s})$ the set of solutions to (2.6) in $\mathcal{B}(\mathbf{c}_-, \mathbf{c}_+; \mathfrak{s})$, called *(SW) instantons*.

Similarly to ECH, an “index” is associated with each SW instanton, namely the local expected dimension of the moduli space of SW instantons. Denote by $\mathfrak{M}_k(\mathbf{c}_-, \mathbf{c}_+; \mathfrak{s})$ the subset of elements in $\mathfrak{M}(\mathbf{c}_-, \mathbf{c}_+; \mathfrak{s})$ that have index k .

Closed 4-manifolds

The case $(Y_\pm, \lambda_\pm) = (\emptyset, 0)$ recovers Seiberg-Witten theory on closed oriented symplectic 4-manifolds. In general, for closed oriented Riemannian 4-manifolds (X, g) , we can recover Seiberg-Witten theory from the above setup by ignoring the appearance of ω and thus ignoring the canonical decomposition of \mathbb{S}_+ . The set of spin-c structures is then only an $H^2(X; \mathbb{Z})$ -torsor. A configuration $\mathbf{d} = [\mathbf{A}, \Psi] \in \mathcal{B}(X, \mathfrak{s})$ solves *the (perturbed) SW equations* when

$$D_{\mathbf{A}} \Psi = 0, \quad F_{\mathbf{A}}^+ = \frac{1}{4}\rho(\Psi) + i\mu \quad (2.7)$$

where $\mu \in \Omega_+^2(X; \mathbb{R})$ is now a self-dual 2-form. Denote the space of solutions to (2.7) by $\mathfrak{M}(\mathfrak{s})$.

When $b_+^2(X) > 0$, a generic choice of μ makes $\mathfrak{M}(\mathfrak{s})$ into a finite-dimensional compact orientable smooth manifold, where the orientation is determined by a *homology orientation* on X (see also Section 2.3), this being an orientation of

$$\det^+(X) := \det(H^1(X; \mathbb{R})) \otimes \det(H_+^2(X; \mathbb{R}))$$

As explained in [67, §1.c], if X is equipped with a symplectic form then there is a canonical homology orientation.

The dimension of $\mathfrak{M}(\mathfrak{s})$ is given by

$$\dim \mathfrak{M}(\mathfrak{s}) = \frac{1}{4}(c_1(\mathfrak{s}) \cdot c_1(\mathfrak{s}) - 2\chi(X) - 3\sigma(X))$$

where $\sigma(X)$ denotes the signature of X , and the parity of $\dim \mathfrak{M}(\mathfrak{s})$ is equal to the parity of $1 - b^1(X) + b_+^2(X)$.

If $\dim \mathfrak{M}(\mathfrak{s}) < 0$ then $\mathfrak{M}(\mathfrak{s})$ is empty and the *SW invariant* $SW_X(\mathfrak{s})$ is defined to be zero. In the remaining cases, the *SW invariant* $SW_X(\mathfrak{s})$ is an element of $\Lambda^* H^1(X; \mathbb{Z})/\text{Torsion}$ and given by suitable counts of elements in $\mathfrak{M}(\mathfrak{s})$ (see Definition 2.3.3). For example, if $\dim \mathfrak{M}(\mathfrak{s}) = 0$ then $SW_X(\mathfrak{s}) \in \mathbb{Z}$ is the signed count of the finite number of oriented points in $\mathfrak{M}(\mathfrak{s})$.

Choice of “chamber”

When $b_+^2(X) > 1$, the value of the SW invariant is a diffeomorphism invariant of X independent of the choice of generic pairs $(g, \mu) \in \text{Met}(X) \times \Omega_+^2(X; \mathbb{R})$, where $\text{Met}(X)$ denotes the Frechet space of smooth Riemannian metrics on X . When $b_+^2(X) = 1$, there is a “wall-crossing phenomenon” as follows. Denote by ω_g the unique (up to scalar multiplication) nontrivial near-symplectic form with respect to g . The set of pairs (g, μ) satisfying the constraint

$$\int_X \omega_g \wedge \mu + 2\pi[\omega_g] \cdot c_1(\mathfrak{s}) = 0 \tag{2.8}$$

defines a “wall” which separates $\text{Met}(X) \times \Omega_+^2(X; \mathbb{R})$ into two open sets, called $c_1(\mathfrak{s})$ -chambers. The SW invariant is constant on any $c_1(\mathfrak{s})$ -chamber, and the difference between chambers is computable.

If X is equipped with a symplectic structure ω (oriented by $\omega^2 > 0$) then there is a canonical $c_1(\mathfrak{s})$ -chamber, namely those pairs (g, μ) for which the left hand side of (2.8) is negative.

Kronheimer-Mrowka’s formalism

The previous sections concerned the setup of SW theory from the point of view of symplectic geometry (using Taubes’ large perturbations). We now briefly review some relevant aspects of SW theory from the point of view of Kronheimer-Mrowka’s monopole Floer homology.

Let $\mathcal{B}(Y, \mathfrak{s})$ denote the space of configurations $[\mathbf{A}, \psi]$. Since we are not taking large perturbations to the SW equations, we have to deal with the reducible locus $\mathcal{B}^{\text{red}}(Y, \mathfrak{s})$ which prevents $\mathcal{B}(Y, \mathfrak{s})$ from being a Banach manifold. This is done by forming the *blow-up* $\mathcal{B}^\sigma(Y, \mathfrak{s})$, the space of configurations $[\mathbf{A}, s, \psi]$ such that $s \in \mathbb{R}^{\geq 0}$ and $\|\psi\|_2 = 1$, equipped with the map

$$\mathcal{B}^\sigma(Y, \mathfrak{s}) \rightarrow \mathcal{B}(Y, \mathfrak{s}), \quad [\mathbf{A}, s, \psi] \mapsto [\mathbf{A}, s\psi]$$

This is a Banach manifold whose boundary $\partial\mathcal{B}^\sigma(Y, \mathfrak{s})$ consists of reducible configurations (where $s = 0$). The same setup applies to the case that X is a closed 4-manifold. The integral cohomology ring $H^*(\mathcal{B}^\sigma(M, \mathfrak{s}); \mathbb{Z})$, for M either Y or X , is isomorphic to the graded algebra

$$\mathbb{A}(M) := (\Lambda^* H_1(M; \mathbb{Z}) / \text{Torsion}) \otimes \mathbb{Z}[U]$$

where U is a 2-dimensional generator (see [33, Proposition 9.7.1]).

We can construct a certain vector field \mathcal{V}^σ on $\mathcal{B}^\sigma(Y, \mathfrak{s})$ using the pull-back of the gradient of the Chern-Simons-Dirac functional $\mathcal{L}_{\text{CSD}} : \mathcal{B}(Y, \mathfrak{s}) \rightarrow \mathbb{R}$ (see [33, §4.1]). Strictly speaking, the Chern-Simons-Dirac functional is not well-defined on $\mathcal{B}(Y, \mathfrak{s})$ unless $c_1(\mathfrak{s})$ is torsion, but such spin-c structures are the only ones relevant to this paper. Likewise, the perturbed gradient $\text{grad } \mathcal{L}_{\text{CSD}} + \mathfrak{q}$ gives rise to a vector field $\mathcal{V}^\sigma + \mathfrak{q}^\sigma$, where \mathfrak{q} is an “abstract tame perturbation” (see [33, §10]). We always assume that \mathfrak{q} is chosen from a Banach space of tame perturbations so that all stationary points of $\mathcal{V}^\sigma + \mathfrak{q}^\sigma$ are nondegenerate.

The critical points (i.e. stationary points) of $\mathcal{V}^\sigma + \mathfrak{q}^\sigma$ are either *irreducibles* of the form $[\mathbf{A}, s, \psi]$ with $s > 0$ and $[\mathbf{A}, s\psi] \in \text{crit}(\text{grad } \mathcal{L}_{\text{CSD}} + \mathfrak{q})$, or *reducibles* of the form $[\mathbf{A}, 0, \psi]$ with ψ an eigenvector of $D_{\mathbf{A}}$. A reducible is *boundary-stable* (respectively, *boundary-unstable*) if the corresponding eigenvalue is positive (respectively, negative). Denote by

$$\mathfrak{C}(Y, \mathfrak{s}) = \mathfrak{C}^o(Y, \mathfrak{s}) \sqcup \mathfrak{C}^u(Y, \mathfrak{s}) \sqcup \mathfrak{C}^s(Y, \mathfrak{s})$$

the decomposition of the set of critical points into the respective sets of irreducibles and boundary-(un)stable reducibles. We can package these critical points together in various ways to form the monopole Floer (co)homologies, such as $\widehat{HM}^*(Y, \mathfrak{s})$ and $\widehat{HM}^*(Y, \mathfrak{s})$ – the former cochain complex is generated by $\mathfrak{C}^o(Y, \mathfrak{s}) \sqcup \mathfrak{C}^s(Y, \mathfrak{s})$ while the latter complex is generated by $\mathfrak{C}^o(Y, \mathfrak{s}) \sqcup \mathfrak{C}^u(Y, \mathfrak{s})$, both equipped with coherent choices of orientations (see Section 2.3). The differentials will not be reviewed here, but we do assume in this paper that all perturbations \mathfrak{q} are chosen so that the differentials are well-defined.

Remark 2.3.2. If \mathfrak{q} is one of Taubes’ sufficiently large perturbations associated with a contact form (given in Section 2.3), then the image of $\mathfrak{C}(Y, \mathfrak{s})$ under the blow-down map is $\mathfrak{M}(Y, \mathfrak{s})$. In fact, we don’t need to use the blow-up model.

Let X either be a closed 4-manifold or have boundary Y . There is a partially-defined restriction map $r : \mathcal{B}^\sigma(X, \mathfrak{s}) \dashrightarrow \mathcal{B}^\sigma(Y, \mathfrak{s})$ whose domain consists of those configurations $[\mathbf{A}, s, \Psi]$ satisfying $\Psi_Y := \Psi|_Y \neq 0$, such that

$$r([\mathbf{A}, s, \Psi]) = \left[\mathbf{A}|_Y, s\|\Psi_Y\|_2, \Psi_Y/\|\Psi_Y\|_2 \right]$$

Similarly, if $X = [0, 1] \times Y$ then there is a family of restriction maps $r_t : \mathcal{B}^\sigma(X, \mathfrak{s}) \dashrightarrow \mathcal{B}^\sigma(Y, \mathfrak{s})$ for $t \in [0, 1]$. If we instead work over $\mathbb{R} \times Y$ or cylindrical ends such as $(-\infty, 0] \times Y$, then we need to use $L^2_{k, \text{loc}}$ -norms (see [33, §13]).

With respect to a cylindrical completion \overline{X} of X , the unperturbed SW equations (2.7) on $\mathcal{B}(\overline{X}, \mathfrak{s})$ now take the form

$$D_{\mathbf{A}}\Psi = 0, \quad F_{\mathbf{A}}^+ = s^2 \frac{1}{4} \rho(\Psi) \quad (2.9)$$

on $\mathcal{B}^\sigma(\overline{X}, \mathfrak{s})$. In the cylindrical case $\overline{X} = \mathbb{R} \times Y$ with spin-c structure induced from \mathfrak{s} on Y and cylindrical perturbation \mathfrak{p} , any solution \mathfrak{d} to the \mathfrak{p}^σ -perturbed version of (2.9) on $\mathbb{R} \times Y$ determines a path $\check{\mathfrak{d}}(t) := r_t(\mathfrak{d}) \in \mathcal{B}^\sigma(Y, \mathfrak{s})$, because there is a unique continuation theorem which ensures that r_t is defined on each slice $\mathfrak{d}|_{\{t\} \times Y}$ (see [33, §10.8]).

In the general case of a cobordism $(X, \mathfrak{s}) : (Y_+, \mathfrak{s}_+) \rightarrow (Y_-, \mathfrak{s}_-)$, we fix abstract perturbations \mathfrak{q}_\pm on Y_\pm and extend them to a suitable abstract perturbation \mathfrak{p} on \overline{X} . To fix notation, if X is a symplectic cobordism with data (ω, λ_\pm) then we denote by \mathfrak{p}_ω and $\mathfrak{q}_{\lambda_\pm}$ the abstract perturbations which are used in Section 2.3 to define Taubes' perturbed SW equations.

Given $\mathfrak{c}_\pm \in \mathfrak{C}(Y_\pm, \mathfrak{s}_\pm)$, we denote by $M(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s})$ the subset of \mathfrak{p}^σ -perturbed SW solutions $\mathfrak{d} \in \mathcal{B}^\sigma(\overline{X}, \mathfrak{s})$ for which $\check{\mathfrak{d}}$ (on the ends of \overline{X}) is asymptotic to \mathfrak{c}_\pm as $t \rightarrow \pm\infty$. Depending on the context, we may alternatively write $M(\mathfrak{c}_-, X, \mathfrak{c}_+; \mathfrak{s})$ to make the manifold explicit. Note that $M(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s}) = \mathfrak{M}(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s})$ when using Taubes' perturbations in Section 2.3.

In the cylindrical case $\overline{X} = \mathbb{R} \times Y$ there is an \mathbb{R} -action by translation on $M(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s})$. The resulting equivalence class of unparametrized nontrivial trajectories, where a trajectory is *nontrivial* if it is not \mathbb{R} -invariant, is denoted by $\check{M}(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s})$. The moduli space of broken trajectories in the sense of [33, Definition 16.1.2] is denoted by $\check{M}^+(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s})$.

We now revisit Section 2.3, where $\mathfrak{M}(\mathfrak{s}) \subset \mathcal{B}(X, \mathfrak{s})$ for a closed 4-manifold X . As explained in [33, §27], for generic perturbations to the SW equations (2.9) on $\mathcal{B}^\sigma(X, \mathfrak{s})$ the resulting moduli space of SW solutions is diffeomorphic to $\mathfrak{M}(\mathfrak{s})$ via the blow-down map. We will therefore define the SW invariants using the blown-up configuration space, and $\mathfrak{M}(\mathfrak{s})$ will also denote the moduli space of SW solutions in $\mathcal{B}^\sigma(X, \mathfrak{s})$. The moduli space gives a well-defined element $[\mathfrak{M}(\mathfrak{s})] \in H_*(\mathcal{B}^\sigma(X, \mathfrak{s}); \mathbb{Z})$.

Definition 2.3.3. For a given choice of homology orientation of X , and a given choice of $c_1(\mathfrak{s})$ -chamber when $b_+^2(X) = 1$, the *Seiberg-Witten invariant* $SW_X(\mathfrak{s}) \in \Lambda^* H^1(X; \mathbb{Z})$ is defined as follows. Its value on $a \in \Lambda^p H_1(X; \mathbb{Z})/\text{Torsion}$, for $p \leq d(\mathfrak{s})$ such that $d(\mathfrak{s}) - p$ is even, is

$$SW_X(\mathfrak{s})(a) := \left\langle U^{\frac{1}{2}(d(\mathfrak{s})-p)} a, [\mathfrak{M}(\mathfrak{s})] \right\rangle \in \mathbb{Z}$$

and it is defined to be zero for all other integers p .

Homology orientations

To coherently orient the moduli spaces $M(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s})$, as explained in [33, §20, §28.4], we must make a $\mathbb{Z}/2\mathbb{Z}$ choice for each generator \mathfrak{c}_\pm and we must choose a *(cobordism) homology orientation* on X . The latter is an orientation of

$$\det^+(X) := \det(H^1(X; \mathbb{R})) \otimes \det(I^+(X; \mathbb{R})) \otimes \det(H^1(Y_+; \mathbb{R}))$$

where $I^+(X; \mathbb{R})$ is defined as follows: The relative cap-product pairing

$$H^2(X, \partial X; \mathbb{R}) \times H^2(X; \mathbb{R}) \rightarrow H^4(X, \partial X; \mathbb{R}) \cong \mathbb{R}$$

induces a nondegenerate quadratic form on the kernel of the restriction map $H^2(X; \mathbb{R}) \rightarrow H^2(\partial X; \mathbb{R})$, and $I^+(X; \mathbb{R}) \subset H^2(X; \mathbb{R})$ is a maximal nonnegative subspace for this quadratic form. The set of homology orientations is denoted by $\Lambda(X)$. In the case that $X = [0, 1] \times Y$ there is a canonical homology orientation $\mathfrak{o}(X) \in \Lambda(X)$, and it is implicitly used when coherently orienting the moduli spaces of trajectories on $\mathbb{R} \times Y$ to define the monopole Floer differentials. In the case that $Y_\pm = \emptyset$, we recover the notion of homology orientation of a closed 4-manifold in Section 2.3.

Likewise, the $\mathbb{Z}/2\mathbb{Z}$ set of orientations for a configuration $\mathfrak{c} \in \mathcal{B}^\sigma(Y)$ is denoted by $\Lambda(\mathfrak{c})$ and defined in [33, §20.3]. These sets are defined so that, when \mathfrak{c}_0 is a reducible critical point of the unperturbed Chern-Simons-Dirac functional, there is a canonical choice $\mathfrak{o}(\mathfrak{c}_0) \in \Lambda(\mathfrak{c})$.

We now explain these choices in a bit more detail, for the case that (X, ω) is a symplectic cobordism and the moduli spaces are defined using Taubes' large perturbations.

Any $\mathfrak{c} \in \mathcal{B}(Y, \mathfrak{s})$ determines a self-adjoint operator $\mathcal{L}_\mathfrak{c}$ which, roughly speaking, is the linearization of Taubes' perturbed SW equations and the gauge group action. A monopole $\mathfrak{c} \in \mathfrak{M}(Y, \mathfrak{s})$ is *nondegenerate* if the kernel of $\mathcal{L}_\mathfrak{c}$ is trivial. Similarly, the linearization of Taubes' perturbed SW equations and the gauge group action at a given configuration $\mathfrak{d} \in \mathcal{B}(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s})$ between monopoles \mathfrak{c}_\pm determines its *deformation operator*

$$\mathfrak{D}_\mathfrak{d} : L_1^2(iT^*\overline{X} \oplus \mathbb{S}_+) \rightarrow L^2(i\bigwedge_+^2 T^*\overline{X} \oplus \mathbb{S}_- \oplus i\mathbb{R})$$

When \mathfrak{c}_\pm are irreducible and nondegenerate, this operator is Fredholm.

Fix spin-c structures \mathfrak{s}_\pm and nondegenerate monopoles \mathfrak{c}_\pm on Y_\pm . Let $\mathcal{B}(\mathfrak{c}_-, \mathfrak{c}_+)$ denote the union of $\mathcal{B}(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s})$ over all spin-c structures on X which restrict to \mathfrak{s}_\pm on Y_\pm , and let $\Lambda(\mathfrak{c}_-, \mathfrak{c}_+)$ denote the orientation sheaf of the determinant line bundle $\det \mathfrak{D} \rightarrow \mathcal{B}(\mathfrak{c}_-, \mathfrak{c}_+)$. The collection $\{\Lambda(\mathfrak{c}_-, \mathfrak{c}_+)\}$ over all nondegenerate monopoles $\mathfrak{c}_\pm \in \mathfrak{M}(Y_\pm, \mathfrak{s}_\pm)$ satisfies the following property: Each nondegenerate monopole \mathfrak{c} has an associated $\mathbb{Z}/2\mathbb{Z}$ -module $\Lambda(\mathfrak{c})$ such that there is a canonical isomorphism

$$\Lambda(\mathfrak{c}_-, \mathfrak{c}_+) \cong \Lambda(\mathfrak{c}_-) \otimes_{\mathbb{Z}/2\mathbb{Z}} \Lambda(X) \otimes_{\mathbb{Z}/2\mathbb{Z}} \Lambda(\mathfrak{c}_+)$$

and the orientations $\{\mathfrak{o}(\mathfrak{c}_-, \mathfrak{c}_+) \in \Lambda(\mathfrak{c}_-, \mathfrak{c}_+)\}_{\mathfrak{c}_\pm \in \mathfrak{M}(Y_\pm, \mathfrak{s}_\pm)}$ are *coherent* if, after fixing a homology orientation $\mathfrak{o}(X) \in \Lambda(X)$, there exists a corresponding set of choices $\{\mathfrak{o}(\mathfrak{c}) \in \Lambda(\mathfrak{c})\}_{\mathfrak{c} \in \mathfrak{M}(Y_\pm, \mathfrak{s}_\pm)}$ such that $\mathfrak{o}(\mathfrak{c}_-, \mathfrak{c}_+) = \mathfrak{o}(\mathfrak{c}_-) \mathfrak{o}(X) \mathfrak{o}(\mathfrak{c}_+)$.

If \mathfrak{d} is nondegenerate, i.e. $\text{Coker}(\mathfrak{D}_\mathfrak{d}) = 0$, then the restriction of $\Lambda(\mathfrak{c}_-, \mathfrak{c}_+)$ to \mathfrak{d} 's component of $\mathfrak{M}(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s})$ is canonically isomorphic to the set of orientations of $\text{Ker}(\mathfrak{D}_\mathfrak{d})$.

Remark 2.3.4. It is currently unknown whether there is a canonical homology orientation of (X, ω) , except in the case of a closed symplectic 4-manifold [67, §1.c]. But the search for a canonical choice can be “pushed to the boundary ∂X ,” as follows. Fix the canonical spin-c

structures $\mathfrak{s}_{\xi_{\pm}}$ on Y_{\pm} and the canonical spin-c structure \mathfrak{s}_{ω} on X . Consider the canonical configurations

$$\mathfrak{c}_{\xi_{\pm}} := [A_{\xi_{\pm}}, (1, 0)] \in \mathcal{B}(Y_{\pm}, \mathfrak{s}_{\xi_{\pm}})$$

and the canonical configuration

$$\mathfrak{d}_{\omega} := [A_{K^{-1}}, (1, 0)] \in \mathcal{B}(\mathfrak{c}_{\xi_{-}}, \mathfrak{c}_{\xi_{+}}; \mathfrak{s}_{\omega})$$

There are perturbations to these configurations, still denoted $\mathfrak{c}_{\xi_{\pm}}$ and \mathfrak{d}_{ω} , which are non-degenerate solutions to Taubes' perturbed SW equations for r sufficiently large, and the deformation operator $\mathfrak{D}_{\mathfrak{d}_{\omega}}$ has trivial kernel and cokernel. Thus there is a canonical orientation of $\det(\mathfrak{D}_{\mathfrak{d}_{\omega}})$, i.e. a canonical choice in $\Lambda(\mathfrak{c}_{\xi_{-}}, \mathfrak{c}_{\xi_{+}})$. If it can be shown that there are canonical choices in $\Lambda(\mathfrak{c}_{\xi_{\pm}})$, then there is a canonical choice in $\Lambda(X)$.

Choice of near-symplectic homology orientation

The case relevant to this paper is a (closed) near-symplectic manifold (X, ω) and the induced symplectic cobordism (X_0, ω) . As explained in [33, §3.4, §26.1], there is a composition law for (cobordism) homology orientations. Namely, we view X as the composition of cobordisms

$$\emptyset \xrightarrow{\mathcal{N}} \bigsqcup_{i=1}^N S^1 \times S^2 \xrightarrow{X_0} \emptyset$$

and then there is a specification $\Lambda(X) = \Lambda(\mathcal{N}) \otimes_{\mathbb{Z}/2\mathbb{Z}} \Lambda(X_0)$ so that a choice of homology orientation for two objects in $\{X, \mathcal{N}, X_0\}$ determines a homology orientation of the third object in that set.

Now, \mathcal{N} is the disjoint union of N copies of $S^1 \times B^3$, the tubular neighborhoods of the zero-circles of ω . Since each $S^1 \times B^3$ has a canonical homology orientation, a homology orientation of \mathcal{N} is equivalent to a choice of ordering of the zero-circles of ω . Therefore, once an ordering of the zero-circles and a homology orientation of X have been fixed, there is an induced homology orientation of X_0 . (Likewise, if it turns out that ω determines a canonical homology orientation of X_0 , then a homology orientation of X is determined by a choice of ordering of the zero-circles.)

Gradings and U -maps

The group $\widehat{HM}^{-*}(Y)$ has an absolute grading by homotopy classes of oriented 2-plane fields on Y (see [33, §28] or [30, §3]), the set of which is denoted by $J(Y)$. This grading of a critical point $\mathfrak{c} \in \mathfrak{C}(Y, \mathfrak{s})$ is denoted by $|\mathfrak{c}| \in J(Y)$.

As described in [33, §28] and [12, §4], there is a well-defined map $J(Y) \rightarrow \text{Spin}^c(Y)$ with the following properties. If $H^2(Y; \mathbb{Z})$ has no 2-torsion then the Euler class of the given 2-plane field uniquely determines the corresponding spin-c structure. There is a transitive \mathbb{Z} -action on $J(Y)$ whose orbits correspond to the spin-c structures: If $[\xi] \in J(Y)$ then $[\xi] + n$

is the homotopy class of a 2-plane field which agrees with ξ outside a small ball $B^3 \subset Y$ and disagrees with ξ on B^3 by a map $(B^3, \partial B^3) \rightarrow (SO(3), \{1\})$ of degree $2n$.⁵ A given orbit $J(Y, \mathfrak{s})$ is freely acted on by \mathbb{Z} if and only if the corresponding Euler class is torsion. In particular, there is an induced relative $\mathbb{Z}/d\mathbb{Z}$ grading on $\widehat{HM}^{-*}(Y, \mathfrak{s})$, where d denotes the divisibility of $c_1(\mathfrak{s})$ in $H^2(Y; \mathbb{Z})/\text{Torsion}$.

It is useful to write out the relative \mathbb{Z} grading on $\widehat{HM}^{-*}(Y, \mathfrak{s})$ when \mathfrak{s} is torsion, as follows. Given $\mathfrak{c}_\pm \in \mathfrak{C}(Y, \mathfrak{s})$, each trajectory $\mathfrak{d} \in M(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s})$ over $\mathbb{R} \times Y$ has a Fredholm operator $Q_{\mathfrak{d}}$ which, roughly speaking, is the linearization of the perturbed version of (2.9) and the gauge group action (see [33, §14.4]). The *relative grading* $\text{gr}(\mathfrak{c}_-, \mathfrak{c}_+)$ between \mathfrak{c}_- and \mathfrak{c}_+ is defined to be the Fredholm index of $Q_{\mathfrak{d}}$ for any $\mathfrak{d} \in M(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s})$,⁶ and

$$|\mathfrak{c}_+| = |\mathfrak{c}_-| + \text{gr}(\mathfrak{c}_-, \mathfrak{c}_+)$$

as expected.

The chain complex $\widehat{CM}_j(Y, \mathfrak{s})$ for $\widehat{HM}_j(Y, \mathfrak{s})$ in grading j is a finitely generated free abelian group (see [33, Lemma 22.3.3]), and the cochain complex for $\widehat{HM}^j(Y, \mathfrak{s})$ in grading j is then defined by $\widehat{CM}^j = \text{Hom}(\widehat{CM}_j, \mathbb{Z})$. If Y is disconnected, then $\widehat{CM}_*(Y, \mathfrak{s})$ and $\widehat{CM}^*(Y, \mathfrak{s})$ are the tensor products of the respective (co)chain complexes of the components of Y . The same applies to the other flavors of monopole Floer (co)homology.

Fix a base point $y \in Y$ and consider SW instantons on $\mathbb{R} \times Y$ for a given spin-c structure \mathfrak{s} on Y . Denote by $\mathfrak{M}_2(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s}, y)$ the subset of SW instantons $[A, (\alpha, \beta)] \in \mathfrak{M}_2(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s})$ for which $\alpha \in \Gamma(E)$ vanishes at $(0, y) \in \mathbb{R} \times Y$. There is a degree -2 chain map⁷

$$U_y : \widehat{HM}^{-*}(Y, \mathfrak{s}) \rightarrow \widehat{HM}^{2-*}(Y, \mathfrak{s})$$

that counts the elements of $\mathfrak{M}_2(\mathfrak{c}_-, \mathfrak{c}_+; \mathfrak{s}, y)$, and on the level of cohomology this U -map does not depend on the choice of base point.

2.4 Review of Taubes' isomorphisms

More details about the relation between pseudoholomorphic curve theory and gauge theory are found in [63, 23, 55, 56, 57, 58, 59].

⁵This convention is opposite to that used in [33].

⁶Indeed, by [33, Proposition 14.4.5, Lemma 14.4.6] the index does not depend on the choice of \mathfrak{d} .

⁷This definition is given in [59, §1.b] and agrees with that in [33, §25.3] as well as that in [34, §4.11]. See [35, §2.5] for details.

Vortices and orbits and curves

The isomorphism between SW Floer cohomology and ECH was inspired by the equivalence of the 4-dimensional invariants, the SW invariants and the Gromov invariants. But these relations were preceded (and depended on) the analogous correspondence in two dimensions, between vortices and points in the complex plane.

A pair (A, α) consists of a Hermitian connection on the trivial complex line bundle $\mathbb{C} \rightarrow \mathbb{C}$ and a section of it, and \mathbf{c} denotes its gauge-equivalence class under the gauge group $C^\infty(\mathbb{C}, S^1)$. Given a nonnegative integer n , the *n-vortex equations* for a configuration \mathbf{c} are

$$*F_A = -i(1 - |\alpha|^2), \quad \bar{\partial}_A \alpha = 0, \quad |\alpha| \leq 1, \quad \int_{\mathbb{C}} (1 - |\alpha|^2) d\text{vol} = 2\pi n \quad (2.10)$$

The solutions are called *n-vortices*, and their moduli space is denoted by \mathfrak{C}_n . For $n = 0$ this space is the single point $(0, 1)$ up to gauge-equivalence, and when $n > 0$ this space has the structure of a complex manifold that is biholomorphic to \mathbb{C}^n . In fact,

Theorem 2.4.1 ([31, 53]). *Given a nonnegative integer n and a collection of (not necessarily distinct) points $z_1, \dots, z_n \in \mathbb{C}$, there exists a unique solution (A, α) of the vortex equation (up to gauge-equivalence) having finite energy and vortex number n and satisfying*

$$\alpha^{-1}(0) = \bigcup_{j=1}^n \{z_j\}$$

Conversely, all finite energy solutions having vortex number $n \geq 0$ are gauge-equivalent to a solution of this form.

The biholomorphisms $\mathfrak{C}_n \approx \text{Sym}^n(\mathbb{C}) \approx \mathbb{C}^n$ are given by

$$\mathbf{c} \mapsto \{z_1, \dots, z_n\} \mapsto (\sigma_1, \dots, \sigma_n), \quad \sigma_k = \sum_{j=1}^n z_j^k = \frac{1}{2\pi} \int_{\mathbb{C}} z^k (1 - |\alpha|^2) d\text{vol}$$

and the origin $0 \in \mathbb{C}^n$ corresponds to the unique symmetric vortex (A, α) satisfying $\alpha^{-1}(0) = 0$.

Given $\mu \in C^\infty(S^1, \mathbb{C})$ and $\nu \in C^\infty(S^1, \mathbb{R})$, the function

$$\mathfrak{h}_{\mu\nu} : \mathfrak{C}_n \rightarrow \mathbb{R}, \quad (A, \alpha) \mapsto \frac{1}{4\pi} \int_{\mathbb{C}} (2\nu|z|^2 + \mu\bar{z}^2 + \bar{\mu}z^2)(1 - |\alpha|^2) d\text{vol} \quad (2.11)$$

induces a time-dependent Hamiltonian vector field for a particular Kähler metric on \mathfrak{C}_n , whose closed integral curves $\mathbf{c}(t) : S^1 \rightarrow \mathfrak{C}_n$ satisfy

$$\frac{i}{2} \mathfrak{c}_*(\partial_t)^{(1,0)} + \nabla^{(1,0)} \mathfrak{h}_{\mu\nu}|_{\mathbf{c}} = 0 \quad (2.12)$$

where $\nabla^{(1,0)}$ denotes the holomorphic part of the gradient. A solution is *nondegenerate* if the linearization of this equation, with respect to a certain covariant derivative, at the solution has trivial kernel (see [56, §2.b Part 3]).

Theorem 2.4.2. ([56, §2.b]) Let (μ, ν) denote the pair associated with an L -flat nondegenerate Reeb orbit γ of a contact 3-manifold, and n a positive integer. If γ is elliptic then there is a single (nondegenerate) solution to (2.12), the constant map to the unique symmetric vortex. The same result holds if γ is hyperbolic and $n = 1$. If $n > 1$ and γ is hyperbolic, there are no solutions to (2.12).

Remark 2.4.3. The solutions granted by this theorem are used to define the isomorphism between ECH and a version of Seiberg-Witten Floer homology. The fact that there are no solutions when $n > 1$ and γ is hyperbolic is not a problem, because such a pair (γ, n) does not arise in an admissible orbit set.

Now we consider J -holomorphic curves in the completion (\bar{X}, ω, J) of a symplectic cobordism. Let $\pi : N_C \rightarrow C$ denote the (holomorphic) normal bundle of an immersed connected J -holomorphic curve C in $\mathcal{M}(\Theta^+, \Theta^-)$, and let $S_{N_C} \subset N_C$ denote the unit circle subbundle. Form the n -vortex bundle

$$\mathfrak{C}_{N_C, n} := S_{N_C} \times_{S^1} \mathfrak{C}_n$$

whose projection onto C will also be denoted by π . With respect to a Hermitian metric and compatible connection on N_C , the $(1, 0)$ -part of its vertical tangent space is

$$T_{1,0}^{\text{vert}} \mathfrak{C}_{N_C, n} = (\text{Ker } d\pi)_{1,0} = S_{N_C} \times_{S^1} T_{1,0} \mathfrak{C}_n$$

Sections $\mathfrak{c} \in \Gamma(\mathfrak{C}_{N_C, n})$ can be viewed as S^1 -invariant maps $S_{N_C} \rightarrow \mathfrak{C}_n$, so their covariant derivative can be taken and restricted to the horizontal subspace $T^{\text{hor}} S_{N_C}$. This defines a “del-bar” operator

$$\mathfrak{c} \mapsto \bar{\partial} \mathfrak{c} \in \Gamma(\mathfrak{c}^* T_{1,0}^{\text{vert}} \mathfrak{C}_{N_C, n} \otimes T^{0,1} C)$$

Given the pair (ν_C, μ_C) associated with the deformation operator D_C of C , define the following section of $\pi^* T^{0,1} C \rightarrow \mathfrak{C}_{N_C, n}$,

$$\mathfrak{h}_{\nu_C \mu_C} = \frac{1}{4\pi} \int_{\mathbb{C}} [2\nu_C |z|^2 + (\mu_C \bar{z}^2 + \bar{\mu}_C z^2)] (1 - |\alpha|^2) \quad (2.13)$$

and denote by $\nabla^{1,0} \mathfrak{h}_{\nu_C \mu_C}$ the corresponding section of $T_{1,0}^{\text{vert}} \mathfrak{C}_{N_C, n} \otimes \pi^* T^{0,1} C$.

Denote by $\Gamma_0(\mathfrak{C}_{N_C, n})$ the space of sections that are asymptotic to zero on the ends of C . Of interest to this paper are those sections $\mathfrak{c} \in \Gamma_0(\mathfrak{C}_{N_C, n})$ which satisfy the equation

$$\bar{\partial} \mathfrak{c} + \mathfrak{c}^* \nabla^{1,0} \mathfrak{h}_{\nu_C \mu_C} = 0 \quad (2.14)$$

because, as will become evident later, they are “halfway” between J -holomorphic curves and SW instantons. Denote the space of such solutions by

$$\mathcal{Z}_0 \subset \Gamma_0(\mathfrak{C}_{N_C, n})$$

Using the identification $\mathfrak{C}_n \approx \mathbb{C}^n$ there is a bundle isomorphism

$$\mathfrak{C}_{N_C, n} \cong \bigoplus_{j=1}^n N_C^j$$

and (2.14) then takes the form

$$\bar{\partial}\eta + \nu_C \aleph(\eta) + \mu_C \mathbb{F}(\eta) = 0 \quad (2.15)$$

for sections $\eta \in \Gamma_0(\bigoplus_{j=1}^n N_C^j)$. Here, $\bar{\partial}$ is the del-bar operator with respect to the Hermitian connection on N_C ,

$$\mathbb{F} : \Gamma\left(\bigoplus_{j=1}^n N_C^j\right) \rightarrow \Gamma\left(\bigoplus_{j=1}^n N_C^{j-2}\right)$$

is some fiber-preserving bundle map that is not \mathbb{R} -linear unless $n = 1$, and

$$\aleph : \Gamma\left(\bigoplus_{j=1}^n N_C^j\right) \rightarrow \Gamma\left(\bigoplus_{j=1}^n N_C^j\right)$$

is the map that multiplies the j^{th} summand by j .

For $n = 1$, \mathbb{F} is the complex conjugation operator and so (2.15) becomes

$$\bar{\partial}\eta + \nu_C \eta + \mu_C \bar{\eta} = 0 \quad (2.16)$$

for sections $\eta \in \Gamma_0(N_C)$. The space of solutions to (2.16) is thus equal to $\text{Ker}(D_C)$. For example, if C is an index 0 curve cut out transversely then the space of solutions to (2.16) is a point, the constant map to the unique symmetric vortex.

Remark 2.4.4. In the setting of [55] where \overline{X} is a symplectization, the only multiply covered curves to be considered were \mathbb{R} -invariant cylinders, for which $\mu_C = 0$. The nonlinear map \mathbb{F} only played a role in the setting of [63] where \overline{X} is a closed manifold, due to the existence of multiply covered tori. In this paper, the space of solutions to (2.15) for $n > 1$ will be of concern whenever C is either a special torus or a special plane.

In [56, §2.f] Taubes spells out the appropriate Morrey spaces to be used for compactly supported sections of the bundles $\mathfrak{c}^* T_{1,0}^{\text{vert}} \mathfrak{C}_{N_C, n}$ and $\mathfrak{c}^* T_{1,0}^{\text{vert}} \mathfrak{C}_{N_C, n} \otimes T^{0,1} C$. They are denoted by \mathcal{K}_c (or \mathcal{K}_{c*} for a different norm) and \mathcal{L}_c , respectively. At any given $c \in \Gamma_0(\mathfrak{C}_{N_C, n})$, the linearization of the operator in (2.15) which cuts out \mathcal{Z}_0 is the \mathbb{R} -linear operator

$$\Delta_c : \mathcal{K}_c \rightarrow \mathcal{L}_c, \quad \zeta \mapsto \bar{\partial}\zeta + \nu_C \aleph(\zeta) + \mu_C d\mathbb{F}_c(\zeta) \quad (2.17)$$

and an element $c \in \mathcal{Z}_0$ is called *regular* if $\text{Coker}(\Delta_c) = 0$.

If all points were regular and \mathcal{Z}_0 was compact, then \mathcal{Z}_0 would be a finite set of points and there would be an associated SW instanton in the image of Ψ_r for each point in \mathcal{Z}_0 . When $d = 1$, the operator (2.17) is precisely the deformation operator D_C which has trivial cokernel (for generic J), and hence \mathcal{Z}_0 consists of a single regular point. But in general, there is no guarantee that any point is regular when $d > 1$ and $\mu_C \neq 0$. Luckily for ECH, if $d > 1$ then C is an \mathbb{R} -invariant cylinder with $\mu_C = 0$, in which case \mathcal{Z}_0 consists of a single regular point.

Maps between ECH and SW Floer cohomology

In [55], Taubes defined a canonical isomorphism of relatively graded $\mathbb{Z}/d\mathbb{Z}$ -modules (with $\mathbb{Z}/2\mathbb{Z}$ coefficients)

$$ECH_*^L(Y, \lambda, \Gamma, J) \cong \widehat{HM}_L^{-*}(Y, \lambda, \mathfrak{s}_\xi + \text{PD}(\Gamma), J, r) \quad (2.18)$$

under the assumption that r is sufficiently large and (λ, J) is a generic L -flat pair, where d denotes the divisibility of $c_1(\mathfrak{s}_\xi + \text{PD}(\Gamma)) = c_1(\xi) + 2\text{PD}(\Gamma)$ in $H^2(Y; \mathbb{Z})/\text{Torsion}$. Taking the direct limit as $L \rightarrow \infty$, Taubes' isomorphism becomes

$$ECH_*(Y, \lambda, \Gamma) \cong \widehat{HM}^{-*}(Y, \mathfrak{s}_\xi + \text{PD}(\Gamma))$$

A detailed explanation can be found in [23], specifically the proof of [23, Theorem 1.3]. As shown in [7], Taubes' isomorphism also preserves the absolute gradings by homotopy classes of oriented 2-plane fields

$$ECH_j(Y, \lambda, J) \cong \widehat{HM}^j(Y)$$

where $ECH_*(Y, \lambda) := \bigoplus_{\Gamma \in H_1(Y; \mathbb{Z})} ECH_*(Y, \lambda, \Gamma)$ and $j \in J(Y)$. As shown in [59, Theorem 1.1], Taubes' isomorphism also intertwines the respective U-maps. We now briefly explain how this isomorphism (2.18) was constructed.

Theorem 2.4.5 (Taubes). *Fix $L > 0$ and a generic L -flat pair (λ, J) on the nondegenerate contact 3-manifold (Y, λ) . Then for all r sufficiently large and $\Gamma \in H_1(Y; \mathbb{Z})$, there is a canonical bijection from the set of generators of $\widehat{CM}^*(Y, \lambda, \mathfrak{s}_\xi + \text{PD}(\Gamma), J, r)$ to the set of generators of $ECC_*^L(Y, \lambda, \Gamma, J)$.*

As shown in [55, Theorem 4.2], the isomorphism is actually established on generators of the chain complexes, so that for any given admissible orbit set

$$\Theta \in ECC_*^L(Y, \lambda, \Gamma, J)$$

there exists a unique irreducible monopole

$$\mathfrak{c}_\Theta = [A_\Theta, (\alpha_\Theta, \beta_\Theta)] \in \widehat{CM}_L^{-*}(Y, \lambda, \mathfrak{s}_\xi + \text{PD}(\Gamma), J, r)$$

To construct \mathfrak{c}_Θ , an arbitrary smooth map $\mathfrak{c}_{\Theta_i} : S^1 \rightarrow \mathfrak{C}_{m_i}$ is first assigned to each pair $(\Theta_i, m_i) \in \Theta$. From such a map, a “potential candidate” $(A_r, (\alpha_r, 0))$ for a monopole is constructed which almost solves Taubes' perturbed SW equations when r is large; this would not be possible if Θ_i was not a Reeb orbit (see [56, Lemma 3.4]). When \mathfrak{c}_{Θ_i} satisfies (2.12), perturbation theory is then used to find an honest solution nearby this candidate. It turns out that there is precisely one such map \mathfrak{c}_{Θ_i} (see Theorem 2.4.2).

As shown in [55, Theorem 4.3], the chain complex differentials also agree: Given two generators

$$\Theta^\pm \in ECC_*^L(Y, \lambda, \Gamma, J)$$

and their corresponding generators

$$\mathfrak{c}_{\Theta^\pm} \in \widehat{CM}_L^{-*}(Y, \lambda, \mathfrak{s}_\xi + \text{PD}(\Gamma), J, r)$$

there is an orientation-preserving diffeomorphism between $\mathcal{M}_1(\Theta^+, \Theta^-)$ and $\mathfrak{M}_1(\mathfrak{c}_{\Theta^-}, \mathfrak{c}_{\Theta^+}; \mathfrak{s})$, which is \mathbb{R} -equivariant with respect to the translation actions. As a reminder, the moduli $\mathfrak{M}_1(\mathfrak{c}_{\Theta^-}, \mathfrak{c}_{\Theta^+}; \mathfrak{s})$ depends on r and we assume r is sufficiently large. In analogy with the construction for chain complex generators, the map

$$\Psi_r : \mathcal{M}_1(\Theta^+, \Theta^-) \rightarrow \mathfrak{M}_1(\mathfrak{c}_{\Theta^-}, \mathfrak{c}_{\Theta^+}; \mathfrak{s}) \quad (2.19)$$

is constructed as follows (see [56, §5] for more details). Given $\mathcal{C} \in \mathcal{M}_1(\Theta^+, \Theta^-)$ there exists a complex line bundle $E \rightarrow \overline{X}$ and a “potential candidate”

$$(A^*, \Psi^*) \in \text{Conn}(E) \oplus \Gamma(E \oplus K^{-1}E)$$

for a SW instanton which almost solves Taubes’ perturbed SW equations when r is large. Here, E has a section whose zero set (with multiplicity) is \mathcal{C} , and away from \mathcal{C} the bundle E is identified with the trivial bundle. Away from \mathcal{C} the pair (A^*, Ψ^*) is close to $(A_0, (1, 0))$, where A_0 is the flat connection on E coming from the product structure. Near a component $(C, d) \in \mathcal{C}$ the pair (A^*, Ψ^*) is determined by a section

$$\mathfrak{c}_{C,d} \in \Gamma_0(\mathfrak{C}_{N_{C,d}})$$

Roughly speaking, when $\mathfrak{c}_{C,d}$ is in \mathcal{Z}_0 for all components $(C, d) \in \mathcal{C}$, a gluing construction perturbs the pair (A^*, Ψ^*) to a SW instanton $\Psi_r(\mathcal{C})$.

The assertion that Ψ_r maps $\mathcal{M}_1(\Theta^+, \Theta^-)$ onto a union of components of $\mathfrak{M}_1(\mathfrak{c}_{\Theta^-}, \mathfrak{c}_{\Theta^+}; \mathfrak{s})$ that contain solely nondegenerate SW instantons is given in [57, §3.a]. The proof that Ψ_r is surjective consists of three main arguments spelled out in [58, §3-7]. First, certain global properties of SW instanton solutions to Taubes’ perturbed SW equations are established in [58, §3] and in [58, Lemma 5.2, Lemma 5.3]. Second, these global results are used to assign an element in $\mathcal{M}_1(\Theta^+, \Theta^-)$ to a given SW instanton in $\mathfrak{M}_1(\mathfrak{c}_{\Theta^-}, \mathfrak{c}_{\Theta^+}; \mathfrak{s})$ and is done so in [58, §4] and in [68, §5-7]. Third, these assignments are given by the map Ψ_r and is done so in [58, §6-7].

Given a basepoint $y \in Y$, the analogous bijection

$$\Psi_r^y : \mathcal{M}_2(\Theta^+, \Theta^-; y) \rightarrow \mathfrak{M}_2(\mathfrak{c}_{\Theta^-}, \mathfrak{c}_{\Theta^+}; \mathfrak{s}, y)$$

is established in [58, Theorem 2.6] and closely follows the construction of Ψ_r .

2.5 The blow-up formula

Recall that an *exceptional sphere* in X is an embedded smooth sphere of self-intersection -1 , and X is *minimal* if there are no exceptional spheres. If X is not minimal then it is a

blow-up of the form

$$X = X_{\min} \# \overbrace{\mathbb{CP}^2 \# \cdots \# \mathbb{CP}^2}^n$$

where X_{\min} is minimal and $n \leq b_2(X)$, though this decomposition is not necessarily unique. Examples of exceptional spheres include each exceptional divisor per copy of $\overline{\mathbb{CP}}^2$ with respect to this decomposition.

Suppose for the moment that $X = X_{\min} \# \overline{\mathbb{CP}}^2$. A homology orientation for X_{\min} determines one also for X , and we suppose that these have been fixed. There is an identification of algebras

$$\mathbb{A}(X) \cong \mathbb{A}(X_{\min}) = (\Lambda^* H_1(X_{\min}; \mathbb{Z}) / \text{Torsion}) \otimes \mathbb{Z}[U]$$

and the respective SW invariants will both be defined over $\mathbb{A}(X_{\min})$. We make the identification $H_2(X; \mathbb{Z}) \cong H_2(X_{\min}; \mathbb{Z}) \oplus \mathbb{Z}[E]$, where $E \subset X$ is the exceptional divisor satisfying $[E] \cdot [E] = -1$. Since S^3 has a unique spin-c structure, we have a bijection

$$\text{Spin}^c(X) \rightarrow \text{Spin}^c(X_{\min}) \oplus \text{Spin}^c(\overline{\mathbb{CP}}^2)$$

induced by restriction. Given $\mathfrak{s} \in \text{Spin}^c(X)$ with $d(\mathfrak{s}) \geq 0$, we write $\mathfrak{s}_{\min} \in \text{Spin}^c(X_{\min})$ for its restriction to X_{\min} . The *blow-up formula* [44, Theorem 2.2] then reads

$$SW_X(\mathfrak{s})(a) = SW_{X_{\min}}(\mathfrak{s}_{\min}) \left(U^{\frac{1}{2}(d(\mathfrak{s}_{\min}) - d(\mathfrak{s}))} a \right) \quad (2.20)$$

for any homogenous element $a \in \mathbb{A}(X_{\min})$ of degree $d(\mathfrak{s})$. For example, if \mathfrak{s} is the spin-c structure that satisfies $c_1(\mathfrak{s}) = c_1(\mathfrak{s}_{\min}) \pm \text{PD}[E]$ then $SW_X(\mathfrak{s})(a) = SW_{X_{\min}}(\mathfrak{s}_{\min})(a)$.

Let $\mathcal{E}_\omega \subset H_2(X_0, \partial X_0; \mathbb{Z})$ denote the set of classes represented by *symplectic* exceptional spheres in (X_0, ω) . The goal of this section is to reduce Theorem 2.1.1 to the case that X is either minimal or that the relative class $\tau_\omega(\mathfrak{s})$ is not represented by a pseudoholomorphic curve with a multiply covered exceptional sphere component. That is,

Theorem 2.5.1. *Given (X, ω, J) ,*

$$Gr_{X, \omega}(\mathfrak{s}) = \pm SW_X(\mathfrak{s}) \in \Lambda^* H^1(X; \mathbb{Z}) \otimes \mathbb{Z}/2\mathbb{Z}$$

for any $\mathfrak{s} \in \text{Spin}^c(X)$ satisfying $e \cdot \tau_\omega(\mathfrak{s}) \geq -1$ for all $e \in \mathcal{E}_\omega$. Here, ω determines the chamber for defining the Seiberg-Witten invariants when $b_+^2(X) = 1$.

Proof of Theorem 2.5.1 \Rightarrow Theorem 2.1.1. Given any $\mathfrak{s} \in \text{Spin}^c(X)$ and its corresponding relative class $A := \tau_\omega(\mathfrak{s}) \in H_2(X_0, \partial X_0; \mathbb{Z})$, consider the relative class

$$A' := A + \sum_{e \in \mathcal{E}_\omega \mid e \cdot A < -1} (e \cdot A) e$$

and its corresponding spin-c structure

$$\mathfrak{s}' := \mathfrak{s} + \sum_{e \in \mathcal{E}_\omega \mid e \cdot A < -1} (e \cdot A) e$$

noting that τ_ω is equivariant with respect to the homology action. The set $\{e \in \mathcal{E}_\omega \mid e \cdot A < -1\}$ is represented by a disjoint collection of symplectic exceptional spheres $\{E_1, \dots, E_n\}$ in X .⁸

We claim now that these spheres can be blown down to obtain a minimal near-symplectic manifold $(X_{\min}, \omega_{\min})$. This is because McDuff's symplectic blow-down operation [39, Lemma 2.1] is constructed locally in the 4-manifold and thus applies to the near-symplectic category (the operation avoids $Z \subset X$).⁹ That is, by the symplectic neighborhood theorem we can excise a neighborhood of a symplectic exceptional sphere and glue in a Darboux 4-ball (whose radius depends on the symplectic area of the sphere), and there is an explicit near-symplectic form ω_{\min} whose pull-back agrees with ω away from the exceptional spheres.

Assume $n = 1$ without loss of generalization. Let $E \subset X$ be the (symplectic) exceptional divisor, let $\pi : (X, \omega) \rightarrow (X_{\min}, \omega_{\min})$ be the blow-down map, and make \mathcal{N} small enough to avoid a neighborhood of E . Defining

$$X_{\min,0} := X_{\min} - \pi(\mathcal{N})$$

and noting that $\partial X_0 = \partial X_{\min,0}$, the relative Mayer-Vietoris sequence gives

$$H_2(X_0, \partial X_0; \mathbb{Z}) \cong H_2(X_{\min,0}, \partial X_{\min,0}; \mathbb{Z}) \oplus \mathbb{Z}[E]$$

We can construct compatible almost complex structures on X_0 and $X_{\min,0}$ in such a way that $\pi|_{X_0}$ is holomorphic, and $c_1(K_{X_0}^{-1}) = \pi^* c_1(K_{X_{\min,0}}^{-1}) - \text{PD}[E]$, where K_{X_0} and $K_{X_{\min,0}}$ are the canonical bundles of (X_0, ω) and $(X_{\min,0}, \omega_{\min})$, respectively.

By hypothesis we know that $[E] \cdot A = -m$ for some positive integer $m \geq 2$. We then compute

$$d(\mathfrak{s}') = d(\mathfrak{s}) + m^2 - m(\text{PD}[E] \cdot c_1(\mathfrak{s}'))$$

Since $E \cap \mathcal{N} = \emptyset$ and $H^2(X) \hookrightarrow H^2(X_0)$, we also compute

$$\begin{aligned} \text{PD}[E] \cdot c_1(\mathfrak{s}') &= \text{PD}[E] \cdot c_1(\mathfrak{s}'|_{X_0}) = \text{PD}[E] \cdot (2 \text{PD}(A') + c_1(K_{X_0}^{-1})) \\ &= 0 + \text{PD}[E] \cdot \pi^* c_1(K_{X_{\min,0}}^{-1}) - [E] \cdot [E] = 1 \end{aligned}$$

Therefore, using the blow-up formula (2.20) we see that

$$SW_X(\mathfrak{s})(a) = SW_X(\mathfrak{s}') \left(U^{\frac{1}{2}m(m-1)} a \right)$$

⁸This is a symmetry argument: If we write $A = m_1 e_1 + m_2 e_2 + B$ such that $e_i \cdot B \geq 0$, then $e_1 \cdot A < 0$ implies $m_2(e_1 \cdot e_2) < m_1$ while $e_2 \cdot A < 0$ implies $m_1(e_1 \cdot e_2) < m_2$, so $(m_1 + m_2)(e_1 \cdot e_2) < m_1 + m_2$ and hence $e_1 \cdot e_2 = 0$.

⁹Likewise, given any near-symplectic 4-manifold (X, ω) we can choose a Darboux ball away from $\omega^{-1}(0)$ and construct a near-symplectic blow-up $(X \# \overline{\mathbb{CP}}^2, \omega_{\text{blow}})$ such that $\omega_{\text{blow}}^{-1}(0)$ is identified with $\omega^{-1}(0)$.

Separately, it follows from Corollary 1.3.20 that

$$Gr_{X,\omega}(A)(a) = Gr_{X,\omega}(A') \left(U^{\frac{1}{2}m(m-1)} a \right)$$

Since $e \cdot A' \geq -1$ for all $e \in \mathcal{E}_\omega$, we are done. \square

We end this section with an example of a blow-up formula for the near-symplectic Gromov invariants, which follows immediately from Theorem 2.5.1. Again, we expect the equivalence to also hold with \mathbb{Z} coefficients.

Corollary 2.5.2. *Suppose (X, ω) is a near-symplectic manifold with a symplectic exceptional sphere E and its near-symplectic blow-down $(X_{\min}, \omega_{\min})$. Given a spin-c structure $\mathfrak{s}_\pm \in \text{Spin}^c(X)$ satisfying $c_1(\mathfrak{s}_\pm) = c_1(\mathfrak{s}_{\min}) \pm \text{PD}[E]$, where \mathfrak{s}_{\min} denotes the restriction of \mathfrak{s}_\pm to X_{\min} ,*

$$Gr_{X,\omega}(\mathfrak{s}_\pm) \equiv_{(2)} Gr_{X_{\min}, \omega_{\min}}(\mathfrak{s}_{\min})$$

after fixing an ordering of the zero-circles of ω (hence of ω_{\min}). In terms of relative homology classes, $\tau_\omega(\mathfrak{s}_+) = \tau_{\omega_{\min}}(\mathfrak{s}_{\min}) + \text{PD}[E]$ and $\tau_\omega(\mathfrak{s}_-) = \tau_{\omega_{\min}}(\mathfrak{s}_{\min})$.

2.6 Equating ECH and HM cobordism counts

In light of the reduction from Theorem 2.1.1 to Theorem 2.5.1, we now fix $(X, \omega, \mathfrak{s})$ such that $e \cdot \tau_\omega(\mathfrak{s}) \geq -1$ for all $e \in \mathcal{E}_\omega$. Subsequently, $(-\partial X_0, \lambda_{\mathfrak{s}})$ and $(\overline{X}, \omega, J)$ are also fixed. We also fix an ordering of the zero-circles of ω , as well as a homology orientation of X_0 (hence of X). Denote by

$$\mathfrak{s}_{\mathfrak{s}} := \mathfrak{s}_\omega + \tau_\omega(\mathfrak{s})$$

the spin-c structure on X_0 (and on \overline{X}) that corresponds to the relative class $\tau_\omega(\mathfrak{s})$. This spin-c structure is identified with $\mathfrak{s}_\xi + 1$ on the ends of \overline{X} . Denote by $E \rightarrow \overline{X}$ the complex line bundle for the spinor decomposition $\mathbb{S}_+ = E \oplus K^{-1}E$ associated with $\mathfrak{s}_{\mathfrak{s}}$. Now, make the following choices:

- an integer $I \geq 0$,
- an integer $p \in \{0, \dots, I\}$ such that $I - p$ is even,
- an ordered set of p disjoint oriented loops $\bar{\eta} := \{\eta_1, \dots, \eta_p\} \subset X_0$,
- a set of $\frac{1}{2}(I - p)$ disjoint points $\bar{z} := \{z_1, \dots, z_{(I-p)/2}\} \subset X_0 - \bar{\eta}$.

Denote by $\mathfrak{M}_I(\mathfrak{c}_\Theta, \emptyset; \mathfrak{s}_{\mathfrak{s}}, \bar{z}, \bar{\eta})$ the subset of SW instantons $\mathfrak{d} = [A, \Psi = (\alpha, \beta)] \in \mathfrak{M}_I(\mathfrak{c}_\Theta, \emptyset; \mathfrak{s}_{\mathfrak{s}})$ for which $\alpha \in \Gamma(E)$ vanishes at each point $z_i \in \bar{z}$ and at some point w_i along each loop $\eta_i \in \bar{\eta}$.

To define the sign $q(\mathfrak{d})$ attached to each \mathfrak{d} in this subset, we build the \mathbb{R}^I -vector space

$$V_{\mathfrak{d}} := \bigoplus_{i=1}^{\frac{1}{2}(I-p)} E_{z_i} \oplus \bigoplus_{i=1}^p \left(E_{w_i} / \nabla_{A_{\mathfrak{d}}} \alpha(T_{w_i} \eta_i) \right) \quad (2.21)$$

as in [67, §2.c]. For generic choices of perturbations $\mu \in \Omega^2(\overline{X})$ that define Taubes' perturbed SW equations (2.6), the covariant derivative of α along η_i at w_i is nonzero and the restriction map

$$\text{Ker}(\mathfrak{D}_{\mathfrak{d}}) \rightarrow V_{\mathfrak{d}}$$

is an isomorphism. Then $q(\mathfrak{d}) = \pm 1$ depending on whether this restriction map is orientation-preserving or orientation-reversing. Here, $\text{Ker}(\mathfrak{D}_{\mathfrak{d}})$ is oriented by the coherent orientations (see Section 2.3), and $V_{\mathfrak{d}}$ is naturally oriented because the complex bundle E is oriented, the loops γ_i are oriented, and the points w_i are ordered.

Notation 2.6.1. Denote $\mathcal{M} := \mathcal{M}_I(\emptyset, \Theta; \tau_{\omega}(\mathfrak{s}), \bar{z}, \bar{\eta})$ and $\mathfrak{M} := \mathfrak{M}_I(\mathfrak{c}_{\Theta}, \emptyset; \mathfrak{s}_{\mathfrak{s}}, \bar{z}, \bar{\eta})$ unless otherwise specified, since these moduli spaces will appear often. We remind the reader that both \mathcal{M} and \mathfrak{M} depend on (λ, J) while \mathfrak{M} also depends on r (and abstract perturbations).

Theorem 2.6.2. *For an admissible orbit set Θ with action less than $\rho(\tau_{\omega}(\mathfrak{s}))$, generic J , and sufficiently large r ,*

$$\sum_{\mathcal{C} \in \mathcal{M}} q(\mathcal{C}) \equiv_{(2)} \sum_{\mathfrak{d} \in \mathfrak{M}} q(\mathfrak{d})$$

The proof of Theorem 2.6.2 is spelled out in the following sections. The key point is that all relevant constructions which occur in a symplectization and a closed symplectic manifold generalize to (completed) cobordisms, because the analysis takes place local to the J -holomorphic curves.

We will construct a multi-valued map (see Definition 2.6.11)

$$\Psi_r : \mathcal{M} \rightarrow \mathfrak{M}$$

almost the same way as in Section 2.4, but the analysis associated with \mathbb{R} -invariant cylinders in [55] disappears and the analysis associated with holomorphic tori in [63] appears. The map Ψ_r is multi-valued because, for a multiply covered plane or torus, the associated space \mathcal{Z}_0 may consist of more than a single (possibly non-regular) point.

All relevant estimates about Taubes' perturbed SW equations have been established in the literature already. For example:

Lemma 2.6.3. *There exists $\kappa \geq 1$ such that if $r \geq \kappa$ and if $(A, \Psi = (\alpha, \beta))$ solves Taubes' perturbed SW equations (2.6) on \overline{X} then*

$$|\alpha| \leq 1 + \kappa r^{-1}$$

$$|\beta|^2 \leq \kappa r^{-1}(1 - |\alpha|^2) + \kappa^2 r^{-2}$$

Proof. This was already stated in [23, Lemma 7.3] for exact symplectic cobordisms, and there are no changes in general. In fact, the proof of this lemma also appears in [65, §6] for our precise setup. The proof follows the arguments in [68, Proposition 2.1, Proposition 2.3] on the compact region X_0 of \overline{X} . This argument extends over the ends $(-\infty, 0] \times \partial X_0$

of \overline{X} , as in the case of symplectizations [58, Lemma 3.1], because the desired bounds are guaranteed on ∂X_0 by [72, Lemma 2.2].

Briefly, from $D_{\mathbf{A}}\Psi = 0$ we know that $D_{\mathbf{A}}^*D_{\mathbf{A}}\Psi = 0$. We rewrite this equation using the Bochner-Weitzenböck formula,

$$\nabla_{\mathbf{A}}^*\nabla_{\mathbf{A}}\Psi + \frac{R_g}{4}\Psi + \frac{1}{2}\text{cl}_+(F_{\mathbf{A}}^+)\Psi = 0 \quad (2.22)$$

where R_g denotes the Ricci scalar curvature of g on \overline{X} . Introducing the components $\Psi = (\alpha, \beta)$,

$$\alpha = \frac{1}{2} \left(1 + \frac{i}{2} \text{cl}_+(\widehat{\omega}) \right) \Psi, \quad \beta = \frac{1}{2} \left(1 - \frac{i}{2} \text{cl}_+(\widehat{\omega}) \right) \Psi$$

we take the inner products of (2.22) with Ψ and $(\alpha, 0)$ and $(0, \beta)$ separately. We then apply the maximum principle to each of the resulting equations, and then we combine all of the resulting inequalities to get the asserted bounds (as well as others). \square

The following proposition is the analog of [23, Proposition 7.1] (which in turn is the analog of [58, Proposition 5.5]), asserting that SW instantons on a symplectic cobordism give rise to pseudoholomorphic curves. Its proof uses Lemma 2.6.3 numerous times along with other estimates established in [68] on closed symplectic manifolds and in [58] on symplectizations. It is explained in [23, §7] how these estimates from [58] on a symplectization carry over, with minor modifications, to *exact* symplectic cobordisms. But we must make an additional modification for everything to carry over to *strong* symplectic cobordisms; (X_0, ω) is such a cobordism. We will elaborate on these modifications immediately after introducing some notation and stating the proposition.

Fix data $(S^1 \times S^2, \lambda_{\mathfrak{s}}, J, \mu, r, \mathfrak{s}_{\xi} + 1)$ as needed to write down Taubes' perturbed SW equations (2.4). Another way to say that a configuration \mathbf{c} is a solution of (2.4) is that it is a critical point of a certain functional, the “Seiberg-Witten action”

$$\mathfrak{a} : \mathcal{B}(S^1 \times S^2, \mathfrak{s}_{\xi} + 1) \rightarrow \mathbb{R}$$

whose definition will not be repeated here (see [23, Equation 97]) save for the remarks that \mathfrak{a} differs from \mathcal{L}_{CSD} by an $O(r)$ constant, and \mathfrak{a} is a gauge-invariant functional because $\mathfrak{s}_{\xi} + 1$ is a torsion spin-c structure. Likewise, fix data $(\overline{X}, \widehat{\omega}, J, \mu, r, \mathfrak{s}_{\mathfrak{s}})$ which extends the data on each $S^1 \times S^2$ boundary component of X_0 as needed to write down Taubes' perturbed SW equations (2.6), hence $\mathfrak{M}(\mathbf{c}, \emptyset; \mathfrak{s}_{\mathfrak{s}}, \bar{z}, \bar{\eta})$. Let

$$s_* : \overline{X} \rightarrow (-\infty, 0]$$

denote the piecewise-smooth function that agrees with the $(-\infty, 0]$ coordinate on each cylindrical end of \overline{X} and equals 0 on X_0 .

Proposition 2.6.4. *Fix $(S^1 \times S^2, \lambda_{\mathfrak{s}}, J, \mu, r, \mathfrak{s}_{\xi} + 1)$ as above on each end of $(\overline{X}, \widehat{\omega}, J, \mu, r, \mathfrak{s}_{\mathfrak{s}})$. Given $\mathcal{K} \geq 1$ and $\delta > 0$, there exist constants $\kappa \geq 1$ and $\kappa_{\delta} \geq 1$ such that the following holds:*

Take $r \geq \kappa_\delta$ and let $\mathfrak{d} = [A, (\alpha, \beta)]$ be an element of $\mathfrak{M}(\mathfrak{c}, \emptyset; \mathfrak{s}_\delta, \bar{z}, \bar{\eta})$ satisfying either $\mathfrak{a}(\mathfrak{c}) \leq \mathcal{K}r$ or $\text{ind}(\mathfrak{D}_\mathfrak{d}) > -\mathcal{K}r$. Then

- $E(\mathfrak{c}) \leq 2\pi\rho(\tau_\omega(\mathfrak{s})) + \delta$.
- Each point in \bar{X} where $|\alpha| \leq 1 - \delta$ has distance less than $\kappa r^{-1/2}$ from $\alpha^{-1}(0)$.
- There exists
 - (a) a positive integer $n \leq \kappa$ and a partition of $(-\infty, 0]$ into intervals $I_1 < \dots < I_n$, each of length at least $2\delta^{-1}$, with $[-1, 0] \subset I_n$, and
 - (b) a broken¹⁰ J -holomorphic curve $\{\mathcal{C}_k\}_{1 \leq k \leq n}$ in \bar{X} (with cobordism level \mathcal{C}_n) asymptotic to the orbit set Θ (in $-\partial X_0$) determined by \mathfrak{c} in Theorem 2.4.5 such that for each $k = 1, \dots, n$ we have

$$\sup_{z \in \mathcal{C}_k \cap s_*^{-1}(I_k)} \text{dist}(z, \alpha^{-1}(0)) + \sup_{z \in \alpha^{-1}(0) \cap s_*^{-1}(I_k)} \text{dist}(\mathcal{C}_k, z) < \delta$$

and $\bar{z} \cup \bar{\eta} \subset \mathcal{C}_n \cap \alpha^{-1}(0) \subset \bar{X}$.

As remarked above, there are some differences between this proposition and [23, Proposition 7.1]. The minor differences are the additional point/loop constraints (as handled in [59, Lemma 4.4]) and the fact that the cobordism X_0 has only negative ends (so the quantity $A_\mathfrak{d}$ in [23, Proposition 7.1] is replaced by $\mathfrak{a}(\mathfrak{c})$). The major difference is the appearance of the homomorphism $\rho : \text{Rel}_\omega(X) \rightarrow \mathbb{R}$. This is ultimately due to the fact that in the proof of [23, Proposition 7.1], Stokes' theorem is used under the assumption that $\widehat{\omega}$ is exact over X_0 , while in our case ρ measures the failure of our non-exact 2-form $\widehat{\omega}$ to satisfy Stokes' theorem with the boundary 1-form $\lambda_\mathfrak{s}$ (this is also explained in [27]).

Remark 2.6.5. If we are to compare Proposition 2.6.4 with [58, Proposition 5.5] then, as in [23], we must replace the manifold $\mathbb{R} \times M$ by \bar{X} , the subset $[s_1, s_2] \times M$ by $s_*^{-1}([s_1, s_2])$, the 2-form $ds \wedge a + \frac{1}{2} * a$ by $\widehat{\omega}$, the 2-form $\frac{\partial}{\partial s} A \pm B_A$ by F_A^\pm , and the spectral flow $f_\mathfrak{d}$ by $\text{ind}(\mathfrak{D}_\mathfrak{d})$.

SW instantons from multiply covered tori and planes

With respect to the analogous bijection $\mathcal{M}_1(\Theta^+, \Theta^-) \rightarrow \mathfrak{M}_1(\mathfrak{c}_{\Theta^-}, \mathfrak{c}_{\Theta^+}; \mathfrak{s})$ over a symplectization (see (2.19)), the analysis in [55] to handle (covers of) \mathbb{R} -invariant cylinders is simple: for an \mathbb{R} -invariant cylinder $C = \mathbb{R} \times \gamma$, the section μ_C vanishes and hence the unique symmetric vortex which solves (2.12) over γ extends to the unique symmetric vortex solution to (2.15) over C . If μ_C did not vanish, the bundle map \mathbb{F} would prevent such an extension (since $\mathbb{F}(0) \neq 0$). This unfortunately turns out to be the case for the special planes and special tori in the scenario at hand. The appropriate analysis to handle the tori is given in [63] and, as will now be shown, can be mimicked to handle the planes. A key point is

¹⁰Strictly speaking, we get a “generalized” broken curve, whose definition only differs from an honest broken curve by requiring that we do *not* mod out by \mathbb{R} -translation of the curves in the symplectization levels of the broken curve.

that there will be no issues with asymptotics, because all perturbations along the curves are required to decay to zero along the curves' ends.

Let (C, d) be a component of $\mathcal{C} \in \mathcal{M}$ where C is a special plane and $d > 1$. As a reminder, $\mathcal{Z}_0 \subset \Gamma_0(\bigoplus_{j=1}^d N_C^j)$ denotes the subspace of elements in the kernel of $\bar{\partial} + \nu_C \aleph + \mu_C \mathbb{F}$ which are asymptotic to zero on the ends of C .

Proposition 2.6.6. \mathcal{Z}_0 is compact.

Proof. The proof copies [67, Proposition 2.7] and its proof in [67, §3] for the analogous case that C is a holomorphic torus. Strictly speaking, [67, Proposition 2.7] argues that \mathcal{Z}_0 is compact if and only if certain Cauchy-Riemann type operators associated to multiple covers of C have trivial kernel and cokernel. As the relevant analysis is performed locally along the curve, the argument extends almost verbatim to the case that C is a holomorphic special plane. What follows is the brief setup and discussion about the necessary changes required to handle covers of holomorphic special planes.

Let $\pi : N_C \rightarrow C$ denote the normal bundle of C with respect to \bar{X} , and let $s : N_C \rightarrow \pi^* N_C$ denote the tautological section. Define the maps

$$R : \Gamma_0\left(\bigoplus_{q=1}^d N_C^q\right) \rightarrow \mathbb{R}, \quad y = (y_1, \dots, y_d) \mapsto \sup_C \sup_q |y_q|^{1/q}$$

$$p : \Gamma_0\left(\bigoplus_{q=1}^d N_C^q\right) \rightarrow \Gamma(\pi^* N_C^n), \quad y = (y_1, \dots, y_d) \mapsto s^d + \pi^* y_1 \cdot s^{d-1} + \dots + \pi^* y_n$$

and suppose there exists a sequence $\{y_j\}_{j \in \mathbb{N}} \subset \mathcal{Z}_0$ with $R_j := R(y_j)$ increasing and unbounded. Then define the subsets

$$\Sigma_j := \{\eta \in N_C \mid p(y_j)(R_j \cdot \eta) = 0\}$$

The proof of [67, Proposition 3.1] shows that a subsequence of $\{\Sigma_j\}_{j \in \mathbb{N}}$ converges pointwise to the image of a somewhere-injective J -holomorphic map $\varphi : C' \rightarrow N_C$ which does not factor through the zero-section of N_C . Here, J on N_C is the unique almost complex structure whose restriction to the fibers of π agrees with the almost complex structure J on \bar{X} . Now, the composition $\pi \circ \varphi : C' \rightarrow C$ is a d -fold holomorphic covering and φ defines a nontrivial element in the kernel of the Cauchy-Riemann type operator $D_{\pi \circ \varphi}$ induced from D_C (see [25, 42, 76]). But this contradicts super-rigidity of C Lemma 1.3.24. Thus \mathcal{Z}_0 is compact. \square

Remark 2.6.7. Since the holomorphic special plane has index zero, it may be possible to modify J in a small neighborhood of its image so that μ becomes 0, following the methodology of Taubes' L -flat approximations (this was brought to the author's attention by Taubes). The a priori issue is that, although all special planes satisfy automatic transversality, there may be other nearby J -holomorphic curves which pop into existence at some point along the smooth path of modifications of J . If this issue can be ruled out, as in the case of the L -flat approximations, then \mathcal{Z}_0 is a single (regular) point and there is no need for Kuranishi structures along special planes.

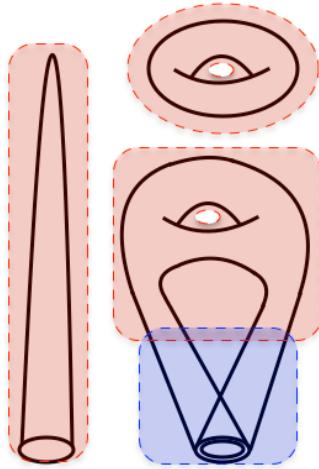


Figure 2.2: C_k in black, U_{C_k} in red, U_{Θ_i} in blue (U_0 not shown)

Kuranishi structures

The point of [56, §5] is to construct families of configurations $(A^*, \Psi^*) \in \text{Conn}(E) \times \Gamma(\mathbb{S}_+)$ for $\mathcal{C} = \{(C_k, d_k)\}$, such that away from $\bigcup_k C_k \subset \overline{X}$ the curvature of A^* is flat and the E -component of Ψ^* is covariantly constant. The construction involves patching together local configurations on open sets in \overline{X} . The difference between compact and noncompact \overline{X} is that, in the noncompact case, there need not be an embedding of a fixed-radius tubular neighborhood of the curves. This complication occurs because the curves may have multiple ends approaching multiple covers of the same Reeb orbit, or an end approaching a multiple cover of a Reeb orbit. But by Proposition 1.3.17 this complication does not occur for the special planes!

The cover of \overline{X} is prescribed by the bullet points of [56, Equation 5-3], and further described as follows (a schematic example is depicted in Figure 2.2):

- There is an open set U_{Θ_i} for each orbit Θ_i that is not approached by a special plane, such that U_{Θ_i} lives far out on the ends of \overline{X} . These sets are pairwise disjoint.
- There is an open set U_{C_k} for each component C_k that is not a special plane, such that U_{C_k} lives in a compact region of \overline{X} (this set does not completely cover the C_k unless this component is closed). These sets are pairwise disjoint.
- There is an open set U_{C_k} for each special plane C_k , given by a fixed-radius tubular neighborhood of C_k . These sets are pairwise disjoint and furthermore disjoint from all preceding open sets.
- There is an open set U_0 to cover the remainder of \overline{X} .

Notation 2.6.8. Use k' to index the components of \mathcal{C} that are special planes or special tori, and use k'' to index the remaining components. Use the index k when it doesn't matter

what the components of \mathcal{C} are.

The constructions in [56, §5] do not require $\mathbf{c} = (\mathbf{c}_k) \in \bigoplus_k \mathcal{Z}_0^{(k)}$ nor for \mathbf{c} to be the zero element. Instead, it can and will be assumed that \mathbf{c} belongs to

$$\mathcal{Y} := \bigoplus_{k'} \mathcal{K}_\Lambda^{(k')} \oplus \bigoplus_{k''} \mathcal{Z}_0^{(k'')}$$

where $\mathcal{K}_\Lambda^{(k')}$ is a suitably small open neighborhood of $\mathcal{Z}_0^{(k')}$ in $\Gamma_0(\bigoplus_{j=1}^{d_{k'}} N_{C_{k'}}^j)$. This neighborhood will be specified in the upcoming Lemma 2.6.10.

As in [56, §5], we begin the search for a solution to the large r version of Taubes' perturbed SW equations by considering families of configurations $(A(\mathbf{c}, \xi, \mathbf{b}, r), \Psi(\mathbf{c}, \xi, \mathbf{b}, r))$ parametrized by

$$\begin{aligned} \mathbf{c} &\in \mathcal{Y} \\ \mathbf{b} &\in \Gamma(iT^*\bar{X} \oplus \mathbb{S}_+) \\ \xi &\in \mathcal{K} := \bigoplus_{k'} \mathcal{K}_{\mathbf{c}_{k'}} \oplus \mathcal{K}'' \end{aligned}$$

Here, \mathcal{K}'' is a Banach space described in [56, §5.b.4] – while $\mathcal{K}_{\mathbf{c}_{k'}}$ is constructed using the bundle $T_{1,0}^{\text{vert}} \mathfrak{C}_{N_{C_{k'}}, d_{k'}}$ along all of $C_{k'}$, the space \mathcal{K}'' is constructed using the pullback of the vortex bundle $T_{1,0} \mathfrak{C}_{m_i}$ along the ends of each component $C_{k''} \in \mathcal{C}$ at each $(\Theta_i, m_i) \in \Theta$ and using the restriction of the vortex bundle $T_{1,0}^{\text{vert}} \mathfrak{C}_{N_{C_{k''}}, 1}$ along the remainder of $C_{k''}$.

Such a configuration $(A(\mathbf{c}, \xi, \mathbf{b}, r), \Psi(\mathbf{c}, \xi, \mathbf{b}, r))$ solves the large r version of Taubes' perturbed SW equations when [56, Equation 5-20] is satisfied, written schematically as

$$\begin{aligned} \mathfrak{D}\mathbf{b} + r^{1/2} \mathbf{b} * \mathbf{b} - \mathbf{v} &= 0 \\ \lim_{s \rightarrow -\infty} \mathbf{b} &= \mathbf{b}_\Theta \end{aligned} \tag{2.23}$$

where \mathbf{b}_Θ is determined by \mathbf{c}_Θ . Here, $\mathbf{b} \mapsto \mathbf{b} * \mathbf{b}$ denotes a (r -independent) quadratic fiber-preserving map from $iT^*\bar{X} \oplus \mathbb{S}_+$ to $i\wedge_+^2 T^*\bar{X} \oplus \mathbb{S}_-$, \mathfrak{D} is the deformation operator associated with $(A(\mathbf{c}, \xi, 0, r), \Psi(\mathbf{c}, \xi, 0, r))$, and \mathbf{v} is the remainder (it is an error term determined by the failure of $(A(\mathbf{c}, \xi, 0, r), \Psi(\mathbf{c}, \xi, 0, r))$ to solve the SW equations). The idea now is to first project (2.23) onto a certain subspace, solve for \mathbf{b} in terms of \mathbf{c} and ξ (and r), and then use the remaining part of (2.23) to solve for ξ in terms of \mathbf{c} .

For each $\mathbf{c} \in \mathcal{Y}$ and $\xi \in \mathcal{K}$, [56, §6 Part 6] introduces a map $\mathfrak{t}_\xi : \mathcal{K}^2 \rightarrow L^2(iT^*\bar{X} \oplus \mathbb{S}_+)$ (see [56, Equation 6-9]), where \mathcal{K}^2 is the version of \mathcal{K} using the L^2 -norm. Let Π_ξ denote the L^2 orthogonal projection onto the image of \mathfrak{t}_ξ . The result of [56, §6] is a solution to

$$(1 - \Pi_\xi) (\mathfrak{D}\mathbf{b} + r^{1/2} \mathbf{b} * \mathbf{b}) = (1 - \Pi_\xi) \mathbf{v} \tag{2.24}$$

Specifically, using the contraction mapping principle, [56, Proposition 6.4] solves for \mathbf{b} as a smooth function of $\{\mathbf{c}, \xi, r\}$, given appropriate bounds on ξ and further assuming bounds on

a certain term $(1 - \Pi_\xi)(\mathbf{v} - \mathbf{v}_h)$ that appears in (2.24). This assumed bound on $(1 - \Pi_\xi)(\mathbf{v} - \mathbf{v}_h)$ is specified in [56, Lemma 6.3] and guaranteed when $\mathbf{c} \in \bigoplus_k \mathcal{Z}_0^{(k)}$, but in general we have to check whether this bound is actually satisfied.

It remains to solve

$$\Pi_\xi(\mathcal{D}\mathbf{b} + r^{1/2}\mathbf{b} * \mathbf{b} - \mathbf{v}) = 0 \quad (2.25)$$

and to find those $\mathbf{c} \in \mathcal{Y}$ for which the term $(1 - \Pi_\xi)(\mathbf{v} - \mathbf{v}_h)$ is suitably bounded. When $\mathbf{c} \notin \bigoplus_k \mathcal{Z}_0^{(k)}$, the constructions in [56, §7] need to be augmented by the constructions in [66]. Indeed, it is shown in [56, §7.d] that the term

$$\theta := r^{1/2}\Pi_\xi(\mathbf{v} - \mathbf{v}_h)$$

is large when \mathbf{c}_k is far away from $\mathcal{Z}_0^{(k)}$, so that (2.25) need not be solvable. What follows are the modifications to [56, §7] that come directly from [66, §5]. In this regard, as the planes and tori $C_{k'}$ are disjoint from the other components of \mathcal{C} and do not approach multiply covered Reeb orbits, the analysis in [66] complements (and does not interfere with) the analysis in [56].

Remark 2.6.9. For the convenience of the reader, here is a dictionary between [66, §5] and our setup (which follows [56, §5-7]). Our variables $\{\mathbf{c}, \xi, \mathbf{b}, \mathbf{v} - \mathbf{v}_h\}$ appear as $\{y^k, x^k, h^k, \text{err}^k\}$ in [66]. Our equations (2.23), (2.24), (2.25) appear as [66, Equation 5.3, Equation 5.12, Equation 5.16]. Also, [56, Proposition 6.4] appears as [66, Lemma 5.5], while the argument of [56, §7.d] corresponds to the argument of [66, §5.f] in which θ appears as [66, Equation 5.17].

While we're giving a dictionary, here is a parallel between the above constructions and the analogous constructions for Reeb orbits (i.e. the construction of \mathbf{c}_Θ from Θ). Our equations (2.23), (2.24), (2.25) are the analogs of [56, Equation 3-6, Equation 3-16, Equation 3-35]. Solving these equations amounts to a proof of Theorem 2.4.5, for the following reason: A key lemma [56, Lemma 3.8] guarantees a unique solution to [56, Equation 3-35] whenever $\mathbf{c} \in \mathcal{Z}_0$ with $\text{Coker } \Delta_\mathbf{c} = 0$, while Theorem 2.4.2 guarantees that $\mathcal{Z}_0 = \{\mathbf{c} = 0\}$ with $\text{Coker } \Delta_0 = 0$, so there is a unique monopole \mathbf{c}_Θ .

As in [56, §7], we view the left hand side of (2.25) as an operator

$$\mathcal{Y} \oplus \mathcal{K} \rightarrow \mathcal{L}, \quad (\mathbf{c}, \xi) \mapsto r^{-1/2}\mathcal{T}(\mathbf{c}, \xi)$$

Here, the Banach space \mathcal{L} is defined in [56, §6.7]; it is the analog of \mathcal{K} in which all bundles are tensored with $T^{0,1}C_k$. We then take the Taylor expansion

$$\mathcal{T}(\mathbf{c}, \xi) = \mathcal{T}_0(\mathbf{c}) + \mathcal{T}_1(\mathbf{c}) \cdot \xi + \mathcal{T}_2(\mathbf{c}, \xi)$$

where $\mathcal{T}_0(\mathbf{c}) := \mathcal{T}(\mathbf{c}, 0)$ and $\mathcal{T}_1(\mathbf{c})$ is linear and $\mathcal{T}_2(\mathbf{c}, \cdot)$ is the remainder.

In the absence of special planes and special tori, the kernel of $\mathcal{T}(\mathbf{c}, \xi)$ is described by [56, Proposition 7.1]. The existence of such elements hinges on appropriate bounds on the \mathcal{T}_j and requires $\mathcal{T}_1(\mathbf{c})$ to satisfy additional properties, granted as follows:

- the bounds on \mathcal{T}_2 are established in [56, §7.e.3] and granted by the small norms on \mathcal{K} ,
- the bounds on $\mathcal{T}_0(\mathbf{c})$ are established in [56, §7.e.1] and granted by the constraint $\mathbf{c} \in \bigoplus_k \mathcal{Z}_0^{(k)}$,
- the properties of $\mathcal{T}_1(\mathbf{c})$ are ultimately granted by the constraint $\bigoplus_k \text{Coker } \Delta_{\mathbf{c}_k} = 0$, which in turn is granted by the fact that $\bigoplus_k \text{Coker } D_{C_k} = 0$ for generic J .

However, for a plane or torus $C_{k'}$ it can be the case that $\text{Coker } \Delta_{\mathbf{c}_{k'}} \neq 0$, so the proof of [56, Proposition 7.1] needs to be modified. The resolution is to relax the constraint $\mathbf{c}_{k'} \in \mathcal{Z}_0^{(k')}$ and appeal to Kuranishi structures as in the proofs of [66, Lemma 5.1, Proposition 5.2]. We do this now.

Using the fact that $\mathcal{Z}_0^{(k')}$ is compact (see Proposition 2.6.6 for special planes and [67, Proposition 2.7] for special tori), we have the following analog of [66, Lemma 5.1] whose proof is contained in [66, §5.a, §5.g.2].¹¹

Lemma 2.6.10. *Let $(C_{k'}, d_{k'})$ be a plane or torus as above. There exists a finite dimensional vector subspace $\Lambda_{k'} \subset \Gamma_0(\bigoplus_{j=1}^{d_{k'}} N_{C_{k'}}^j \otimes T^{0,1}C_{k'})$ such that for all $\mathbf{c}_{k'} \in \mathcal{Z}_0^{(k')}$, the projection of $\Lambda_{k'}$ onto $\text{Coker } \Delta_{\mathbf{c}_{k'}}$ is surjective. Denote by $Q_{\Lambda}^{(k')} : \Gamma_0(\bigoplus_{j=1}^{d_{k'}} N_{C_{k'}}^j \otimes T^{0,1}C_{k'}) \rightarrow \Lambda_{k'}$ the L^2 orthogonal projection. There also exists a smooth $\dim(\Lambda_{k'})$ -dimensional submanifold $\mathcal{K}_{\Lambda}^{(k')} \subset \Gamma_0(\bigoplus_{j=1}^{d_{k'}} N_{C_{k'}}^j)$ with compact closure, satisfying the following properties:*

- 1) *If $\mathbf{c}_{k'} \in \mathcal{K}_{\Lambda}^{(k')}$ then*

$$(1 - Q_{\Lambda}^{(k')}) (\bar{\partial} \mathbf{c}_{k'} + \nu_{C_{k'}} \mathbf{N}(\mathbf{c}_{k'}) + \mu_{C_{k'}} \mathbf{F}(\mathbf{c}_{k'})) = 0$$

- 2) *$\mathcal{Z}_0^{(k')}$ embeds in $\mathcal{K}_{\Lambda}^{(k')}$ as the zero set of the map*

$$\psi_{\Lambda}^{(k')} : \mathcal{K}_{\Lambda}^{(k')} \rightarrow \Lambda_{k'}, \quad \mathbf{c}_{k'} \mapsto Q_{\Lambda}^{(k')} (\bar{\partial} \mathbf{c}_{k'} + \nu_{C_{k'}} \mathbf{N}(\mathbf{c}_{k'}) + \mu_{C_{k'}} \mathbf{F}(\mathbf{c}_{k'}))$$

- 3) *For a suitably chosen r -independent constant $\varepsilon_{k'} > 0$, for each $\mathbf{c}_{k'} \in \mathcal{K}_{\Lambda}^{(k')}$ there is some element in $\mathcal{Z}_0^{(k')}$ that is within $\varepsilon_{k'}$ distance (with respect to the L^2 norm) to $\mathbf{c}_{k'}$, and $(1 - Q_{\Lambda}^{(k')}) \Delta_{\mathbf{c}_{k'}}$ is surjective.*

We now show how this lemma is used to modify the arguments in [56, §7]. Following [56, §7.b.5], we write the components of \mathcal{T} as $((\mathcal{T}_{C_k})_{C_k \in \mathcal{C}}, (\mathcal{T}_{\Theta_j})_{\Theta_j \in \Theta})$. In terms of the Taylor expansion, each component $\mathcal{T}_{C_k}(\mathbf{c}_k, \xi_k)$ equals the sum of the 0th order term

$$\mathcal{T}_{0C_k}(\mathbf{c}_k, \xi_k) = \bar{\partial} \mathbf{c}_k + \nu_{C_k} \mathbf{N}(\mathbf{c}_k) + \mu_{C_k} \mathbf{F}(\mathbf{c}_k)$$

and the linear term

$$\mathcal{T}_{1C_k}(\mathbf{c}_k, \xi_k) = \Delta_{\mathbf{c}_k} \xi_k$$

and the remainder¹² $\mathcal{T}_{2C_k}(\mathbf{c}_k, \xi_k)$. For planes or tori, the linear term $\mathcal{T}_{1C_{k'}}$ may have nonzero

¹¹The submanifold \mathcal{K}_{Λ} in [66, Lemma 5.1] has dimension $d + \dim(\Lambda)$, where $d = \text{ind}_{\mathbb{R}} \Delta$. In our case, $\text{ind}_{\mathbb{R}} \Delta = \text{ind}_{\mathbb{R}} D_{C_{k'}} = 0$. For the reader's benefit, we point out that the integer d is misstated as $d = 2m(n+1-g) + m(m-1)$ in [66, Lemma 5.1], though correctly stated as $d = 2m(1-g) + m(m+1)n = 2m(n+1-g) + m(m-1)n$ in [66, Proposition 3.2].

¹²The remainder is denoted by \mathcal{R}_k in [66].

(co)kernel. To deal with this, we first apply the operator $1 - Q_\Lambda^{(k')}$ to the $\mathcal{T}_{C_{k'}}$ components, and we subsequently abuse notation to denote the resulting operator by

$$\mathcal{T} := \left(\left((1 - Q_\Lambda^{(k')}) \mathcal{T}_{C_{k'}} \right)_{C_{k'} \in \mathcal{C}}, (\mathcal{T}_{C_{k''}})_{C_{k''} \in \mathcal{C}}, (\mathcal{T}_{\Theta_j})_{\Theta_j \in \Theta} \right)$$

Note that the 0th order term of \mathcal{T} vanishes by definition of the space \mathcal{Y} . We then restrict the domain of \mathcal{T} to $\mathcal{Y} \oplus (\bigoplus_{k'} \mathcal{L}_{\mathfrak{c}_{k'}}^\perp) \oplus \mathcal{K}''$, where

$$\mathcal{L}_{\mathfrak{c}_{k'}} := \text{Ker}(1 - Q_\Lambda^{(k')}) \Delta_{\mathfrak{c}_{k'}} \subset \Gamma_0 \left(\bigoplus_{j=1}^{d_{k'}} N_{C_{k'}}^j \right)$$

We now rerun [56, §7] to solve this projected equation $\mathcal{T}(\mathfrak{c}, \xi) = 0$ uniquely for $\xi(\mathfrak{c}) \in \mathcal{B}$ as a function of \mathfrak{c} , where \mathcal{B} is a sufficiently small ball in $(\bigoplus_{k'} \mathcal{L}_{\mathfrak{c}_{k'}}^\perp) \oplus \mathcal{K}'' \subset \mathcal{K}$. This association $\mathfrak{c} \mapsto \xi(\mathfrak{c})$ induces the smooth map

$$\Psi_{\mathcal{C},r} : \mathcal{Y} \rightarrow \mathcal{B}(\mathfrak{c}_\Theta, \emptyset; \mathfrak{s}_\mathfrak{s}), \quad \mathfrak{c} \mapsto \left(A(\mathfrak{c}, \xi(\mathfrak{c}), \mathfrak{b}(\mathfrak{c}, \xi(\mathfrak{c})), r), \Psi(\mathfrak{c}, \xi(\mathfrak{c}), \mathfrak{b}(\mathfrak{c}, \xi(\mathfrak{c})), r) \right) \quad (2.26)$$

and it remains to find those $\mathfrak{c} \in \mathcal{Y}$ that will satisfy $\Psi_{\mathcal{C},r}(\mathfrak{c}) \in \mathfrak{M}$.

The evaluation of the projection $Q_\Lambda^{(k')} \mathcal{T}_{C_{k'}}$ defines the smooth map

$$\psi_{\mathcal{C},r} : \mathcal{Y} \rightarrow \bigoplus_{k'} \Lambda_{k'}, \quad \mathfrak{c} \mapsto \bigoplus_{k'} \left[\psi_\Lambda^{(k')}(\mathfrak{c}_{k'}) + Q_\Lambda^{(k')} (\Delta_{\mathfrak{c}_{k'}} \xi_{k'} + \mathcal{T}_{2C_{k'}}(\mathfrak{c}_{k'}, \xi_{k'})) \right] \quad (2.27)$$

By construction,

$$\Psi_{\mathcal{C},r}(\psi_{\mathcal{C},r}^{-1}(0) \cap \mathcal{Y}_{\bar{z}, \bar{\eta}}) \subset \mathfrak{M}$$

where

$$\mathcal{Y}_{\bar{z}, \bar{\eta}} := \Psi_{\mathcal{C},r}^{-1}(\mathcal{B}(\mathfrak{c}_\Theta, \emptyset; \mathfrak{s}_\mathfrak{s}, \bar{z}, \bar{\eta})) \subset \mathcal{Y}$$

As explained in [66, §6] and [59, §4], $\Psi_{\mathcal{C},r}$ is a smooth embedding for sufficiently large r and it maps $\psi_{\mathcal{C},r}^{-1}(0) \cap \mathcal{Y}_{\bar{z}, \bar{\eta}}$ homeomorphically onto an open subset of \mathfrak{M} .

Finally, the proof of [67, Proposition 2.10] can be copied and combined¹³ with the proof of [58, Theorem 1.2] to show that the inverse image of \mathfrak{M} under $\Psi_{\mathcal{C},r}$ for any given $\mathcal{C} \in \mathcal{M}$ is precisely $\psi_{\mathcal{C},r}^{-1}(0) \cap \mathcal{Y}_{\bar{z}, \bar{\eta}}$, and that each element in \mathfrak{M} lies in the image of $\Psi_{\mathcal{C},r}$ for some $\mathcal{C} \in \mathcal{M}$. In other words, each image $\text{Im } \Psi_{\mathcal{C},r}$ constitutes a “Kuranishi model” for an open subset of \mathfrak{M} .

We can now state the anticipated correspondence between \mathcal{M} and \mathfrak{M} .

¹³The relevant arguments involve “special sections” of powers of N_{C_k} associated with the component (C_k, d_k) , denoted by \mathfrak{o} in [58, Equation 7-9] and by h in [67, Equation 5.26]; they differ by a factor of $\frac{1}{d_k \pi}$. In this regard, Taubes remarks in [58] that [67, Lemma 5.5] is flawed and must be replaced by [58, Lemma 7.1].

Definition 2.6.11. For an admissible orbit set Θ with action less than $\rho(\tau_\omega(\mathfrak{s}))$, generic J , and sufficiently large r , the multi-valued bijection $\Psi_r : \mathcal{M} \rightarrow \mathfrak{M}$ is given by the composition

$$\mathcal{C} \mapsto \psi_{\mathcal{C},r}^{-1}(0) \cap \mathcal{Y}_{\bar{z},\bar{\eta}} \mapsto \Psi_{\mathcal{C},r}(\psi_{\mathcal{C},r}^{-1}(0) \cap \mathcal{Y}_{\bar{z},\bar{\eta}})$$

where the first map is multi-valued and the second map is injective.

Remark 2.6.12. The expected dimension of the moduli of SW instantons agrees with the ECH index of the relevant J -holomorphic currents, i.e. the association of a current $\mathcal{C} \in \mathcal{M}(\emptyset, \Theta; \tau_\omega(\mathfrak{s}))$ to a given $\mathfrak{d} \in \mathfrak{M}_I(\mathfrak{c}_\Theta, \emptyset; \mathfrak{s}_\mathfrak{s})$ satisfies $I(\mathcal{C}) = I$. The analogous statement on symplectizations is proven in [57, §2.b.6], and it can be mimicked to yield the generalized statement on symplectic cobordisms, as done in [7, Theorem 5.1].

Proof of Theorem 2.6.2

First, suppose that \mathcal{M} contains only elements of the form $\{(C_k, d_k = 1)\}$. Then the multi-valued map Ψ_r in (2.6.11) is an honest bijection

$$\mathcal{M} \longleftrightarrow \mathfrak{M}$$

(Such a bijection is the analog of both [55, Theorem 4.3] and [67, Proposition 2.6].) In this case, Theorem 2.6.2 is proved.

Next, suppose that each $\mathcal{Z}_0^{(k)}$ corresponding to each $d_k > 1$ pair (C_k, d_k) from any $\mathcal{C} \in \mathcal{M}$ consists of only regular points. Each $\mathcal{Z}_0^{(k)}$ is subsequently compact, hence a finite set of points. Then we can rephrase the multi-valued map Ψ_r in (2.6.11) as a bijection

$$\bigcup_{\{(C_k, d_k)\} \in \mathcal{M}} \bigoplus_k \mathcal{Z}_0^{(k)} \longleftrightarrow \mathfrak{M}$$

(Such a bijection is the analog of [67, Proposition 2.9].) To prove Theorem 2.6.2 in this case, we need to count the points in $\mathcal{Z}_0^{(k)}$ and show that the resulting number is $r(C_k, d_k)$. We instead handle this in the general scenario, where $\mathcal{Z}_0^{(k)}$ may also contain non-regular points.

Fix a special torus or special plane $C_{k'}$ with multiplicity $d_{k'}$, representing a component of a given current $\mathcal{C} \in \mathcal{M}$. If $\mathcal{Z}_0^{(k')}$ were a finite set of regular points, then we could count the points (with appropriate signs) to define a *weight*

$$r'(C_{k'}, d_{k'}) \in \mathbb{Z}$$

attached to $(C_{k'}, d_{k'}) \in \mathcal{C}$ (see [67, §2.f]). In general, the weight $r'(C_{k'}, d_{k'}) \in \mathbb{Z}$ is defined in [67, §2.g] and roughly speaking, it is the signed count of zeros of a small perturbation of the map $\psi_\Lambda^{(k')} : \mathcal{K}_\Lambda^{(k')} \rightarrow \Lambda_{k'}$ from Lemma 2.6.10. In fact, $\psi_{\mathcal{C},r}$ is such a small perturbation of $\bigoplus_{k'} \psi_\Lambda^{(k')}$, so by construction of (2.6.11) it suffices to show that

$$r'(C_{k'}, d_{k'}) = r(C_{k'}, d_{k'})$$

in order to complete the general proof of Theorem 2.6.2.

If $C_{k'}$ is a torus then [67, Proposition 2.15] asserts that $r'(C_{k'}, d_{k'}) = r(C_{k'}, d_{k'})$. The analogous result for planes is as follows.

Proposition 2.6.13. *For generic J , the weight $r'(C_{k'}, d_{k'})$ is equal to $+1$ when $C_{k'}$ is a J -holomorphic special plane. In particular, $r'(C_{k'}, d_{k'}) = r(C_{k'}, d_{k'})$.*

Proof. [67, Proposition 2.12] implies that $r'(C_{k'}, d_{k'})$ is independent of $\mu_{C_{k'}}$ for $C_{k'}$ a special plane. In particular, $r'(C_{k'}, d_{k'})$ is equal to the signed count of points of the version of $\mathcal{Z}_0^{(k')}$ defined with $\mu_{C_{k'}} = 0$. This version consists of a single regular point with positive sign (namely, the constant map to the unique symmetric vortex) because the operator $\bar{\partial} + \nu_{C_{k'}} \aleph$ is complex linear and has trivial kernel and cokernel for generic $\nu_{C_{k'}}$. Thus $r'(C_{k'}, d_{k'}) = +1$. \square

2.7 Relation of SW counts

In this section we are always using $\mathbb{Z}/2\mathbb{Z}$ coefficients. The goal of this section is to prove Theorem 2.5.1 (and therefore Theorem 2.1.1). It follows from Theorem 2.4.5 and Theorem 2.6.2 using $I = d(\mathfrak{s})$ that the Gromov cycle Φ_{Gr} is chain-isomorphic (over $\mathbb{Z}/2\mathbb{Z}$) to the *Seiberg-Witten cocycle*

$$\Phi_{SW} := \sum_{\mathfrak{c}_\Theta \in g(\mathfrak{s})} \mathfrak{M}_{\mathfrak{c}_\Theta} \mathfrak{c}_\Theta \in \widehat{CM}^{g(\mathfrak{s})}(-\partial X_0, \mathfrak{s}_\xi + 1)$$

where

$$\mathfrak{M}_{\mathfrak{c}_\Theta} := \sum_{\mathfrak{d} \in \mathfrak{M}_{d(\mathfrak{s})}(\mathfrak{c}_\Theta, \emptyset; \mathfrak{s}_\xi, \bar{z}, \bar{\eta})} q(\mathfrak{d}) \in \mathbb{Z}/2\mathbb{Z} \quad (2.28)$$

The number $Gr_{X, \omega}(\mathfrak{s})([\eta_1] \wedge \cdots \wedge [\eta_p])$ is therefore equal to the coefficient of the class

$$[\Phi_{SW}] \in \bigotimes_{k=1}^N \widehat{HM}^{[\xi_*]}(S^1 \times S^2, \mathfrak{s}_\xi + 1)$$

as a multiple of the positive generator $\mathbb{1} \in \bigotimes_{k=1}^N \widehat{HM}^{[\xi_*]}(S^1 \times S^2, \mathfrak{s}_\xi + 1)$. In fact, the following theorem shows that this coefficient is the corresponding Seiberg-Witten invariant.

Theorem 2.7.1. *Fix (X, ω) and assume $\mathfrak{s} \in \text{Spin}^c(X)$ is such that $e \cdot \tau_\omega(\mathfrak{s}) \geq -1$ for all $e \in \mathcal{E}_\omega$. Fix an integer p such that $0 \leq p \leq d(\mathfrak{s})$ and $d(\mathfrak{s}) - p$ is even, and fix an ordered set of homology classes $[\bar{\eta}] := \{[\eta_1], \dots, [\eta_p]\} \subset H_1(X; \mathbb{Z})/\text{Torsion}$. Then*

$$SW_X(\mathfrak{s})([\eta_1] \wedge \cdots \wedge [\eta_p]) \equiv_{(2)} Gr_{X, \omega}(\mathfrak{s})([\eta_1] \wedge \cdots \wedge [\eta_p])$$

In particular, $g(\mathfrak{s}) = N[\xi_]$ as an absolute grading of the N -fold tensor product of $\widehat{HM}^*(S^1 \times S^2)$.*

Before we prove this theorem we make the following remarks. The work of Kronheimer-Mrowka (specifically, [33, Proposition 27.4.1]) recovers the Seiberg-Witten invariant $SW_X(\mathfrak{s})$ using their monopole Floer (co)homologies, explicitly by removing two copies of the 4-ball B^4 from X and counting certain¹⁴ SW instantons on the resulting cobordism $S^3 \rightarrow S^3$. Although not provided in [33], the same result could have been obtained by removing two copies of $S^1 \times B^3$, because both spaces (B^4 and $S^1 \times B^3$) have positive scalar curvature and a unique spin-c structure extending the fixed torsion spin-c structure on their boundary (see [33, Proposition 22.7.1]). There would necessarily be more work to do when using $S^1 \times B^3$ because there exists a circle's worth of reducible monopoles on $S^1 \times S^2$ (to the unperturbed SW equations), compared to a single reducible monopole on S^3 .¹⁵

In the proof of Theorem 2.7.1 we will remove N copies of $S^1 \times B^3$ from X , namely, the tubular neighborhoods of the zero-circles. This is a “neck stretching” argument along the contact hypersurfaces $(S^1 \times S^2, \lambda_{\mathfrak{s}})$ in X , and we analyze the SW equations under this deformation. It is important to note that on X_0 we can use Taubes’ large perturbations to the SW equations, for which there are no reducible solutions, but on each $S^1 \times B^3$ we cannot do this because there is no symplectic form (or said another way, the near-symplectic form ω degenerates somewhere inside $S^1 \times B^3$). We therefore interpolate, on the “neck region” of X , between Taubes’ large perturbations on X_0 and very small perturbations on each $S^1 \times B^3$, where the small perturbations are chosen in such a way that we can understand the SW instantons on each $S^1 \times B^3$ completely.

A final remark is that in this setup, we do not run into the usual difficulties that Kronheimer-Mrowka have when defining monopole Floer cobordism maps for cobordisms with disconnected and empty ends. These difficulties are ultimately due to the (stratified) space of reducible monopoles on the (positive and negative) boundary components of the cobordism, and are avoided in our setup thanks to Taubes’ large perturbations.

Proof of Theorem 2.7.1. We closely follow the arguments in [33, §26, §27.4, §36.1] that recover the Seiberg-Witten invariant and establish the composition law for monopole Floer (co)homology. These arguments involve judicious choices of Riemannian metrics and abstract perturbations on X to “stretch the neck” and compare the resulting moduli spaces of SW instantons.

Let \mathbb{T} denote the circle $H^1(S^1 \times S^2; i\mathbb{R})/H^1(S^1 \times S^2; 2\pi i\mathbb{Z}) \cong S^1$ which parametrizes reducible monopoles to the unperturbed SW equations over $S^1 \times S^2$. In fact, all monopoles are reducible because $S^1 \times S^2$ has a metric of positive scalar curvature (see [33, Proposition 22.7.1]). After fixing a reference connection \mathbf{A}_0 on $\det(\mathbb{S})$ so that any other Hermitian connection can be written as $\mathbf{A} = \mathbf{A}_0 + 2a$ for some $a \in \Omega^1(Y; i\mathbb{R})$, there is a retraction map

¹⁴They build a “mixed” map $\overrightarrow{HM} : \widehat{HM}_*(S^3) \rightarrow \widetilde{HM}^*(S^3)$ and pair the image of a generator $\mathbf{1} \in \widehat{HM}_*(S^3)$ with a generator $\check{\mathbf{1}} \in \widetilde{HM}^*(S^3)$. A homology orientation of the cobordism is identified with a homology orientation of X .

¹⁵We must also choose a homology orientation of $S^1 \times S^2$, i.e. an orientation of the vector space $H^1(S^1 \times S^2; \mathbb{R}) \cong \mathbb{R}$, in order to identify a homology orientation of X with that on the cobordism.

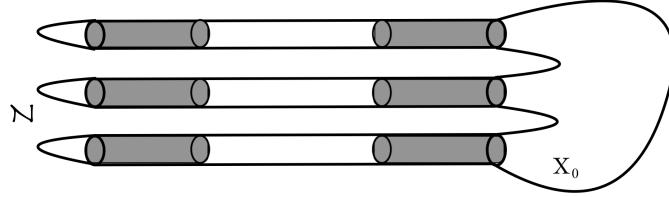


Figure 2.3: “Stretching the necks” of X , depicted in gray

$p : \mathcal{B}(S^1 \times S^2, \mathfrak{s}_\xi + 1) \rightarrow \mathbb{T}$ sending $[\mathbf{A}_0 + 2a, \Psi]$ to the equivalence class of the harmonic part a_{harm} of a (see [33, §11.1]).

Let f be the “height” Morse function on \mathbb{T} with two critical points, and let $f_1 = f \circ p$ be the corresponding function on $\mathcal{B}(S^1 \times S^2, \mathfrak{s}_\xi + 1)$. The gradient of f_1 is an abstract perturbation $\mathfrak{q}_f := \text{grad } f_1$ (assumed small by re-scaling f), and the reducible critical points of $\text{grad } \mathcal{L}_{\text{CSD}} + \mathfrak{q}_f$ are the maximum and minimum critical points $\{\alpha, \beta\}$ of f on \mathbb{T} . The perturbed Dirac operators associated with α and β in the blow-up $\mathcal{B}^\sigma(S^1 \times S^2, \mathfrak{s}_\xi + 1)$ do not have kernel, but to guarantee that their spectrums are simple we add a further small perturbation to \mathfrak{q}_f (still denoted \mathfrak{q}_f) which vanishes on $\mathcal{B}^{\text{red}}(S^1 \times S^2, \mathfrak{s}_\xi + 1)$. Label the corresponding critical points in $\mathcal{B}^\sigma(S^1 \times S^2, \mathfrak{s}_\xi + 1)$ as \mathfrak{a}_i and \mathfrak{b}_i in increasing order of the index, where \mathfrak{a}_0 and \mathfrak{b}_0 correspond to the first positive eigenvalues of the perturbed Dirac operator at α and β (the critical points are boundary-stable for $i \geq 0$ and boundary-unstable for $i < 0$).

Consider a component $\mathcal{N}_k \approx S^1 \times B^3$ of $\mathcal{N} = \bigsqcup_{k=1}^N \mathcal{N}_k$, equipped with a metric having positive scalar curvature and containing a collar region of its boundary in which the metric is cylindrical. Choose a small perturbation $\mathfrak{p}_{\mathcal{N}}$ on $\overline{\mathcal{N}_k}$ equal to \mathfrak{q}_f on the end, so that the corresponding moduli spaces $M(\emptyset, \mathcal{N}_k, \mathfrak{a}_i; \mathfrak{s})$ and $M(\emptyset, \mathcal{N}_k, \mathfrak{b}_i; \mathfrak{s})$ are regular. Here, \mathfrak{s} on \mathcal{N}_k is the unique spin-c structure which extends $\mathfrak{s}_\xi + 1$ on its boundary $S^1 \times S^2$.

Let $X(T)$ be the closed manifold (diffeomorphic to X) obtained by attaching, for each end of X_0 , two copies of the cylinder $[0, T] \times S^1 \times S^2$ and one copy of another cylinder $[0, 1] \times S^1 \times S^2$ and the k th component \mathcal{N}_k of \mathcal{N} (see Figure 2.3):

$$X(T) := \mathcal{N} \cup \bigcup_{k=1}^N \left(([0, T] \times S^1 \times S^2) \cup ([0, 1] \times S^1 \times S^2) \cup ([0, T] \times S^1 \times S^2) \right) \cup X_0$$

The perturbed SW equations on $X(T)$ carry the following perturbations:

- Taubes’ perturbation \mathfrak{p}_ω on X_0 which extends over the adjacent copies of $[0, T] \times S^1 \times S^2$ using Taubes’ perturbation \mathfrak{q}_λ on ∂X_0 ,
- the perturbation $\mathfrak{p}_{\mathcal{N}}$ on each \mathcal{N}_k which extends over the adjacent copies of $[0, T] \times S^1 \times S^2$ using the perturbation \mathfrak{q}_f on $\partial \mathcal{N}_k$,
- an “interpolating” perturbation $\mathfrak{p}_{\text{cyl}}$ on each copy of $[0, 1] \times S^1 \times S^2$ which agrees with \mathfrak{q}_λ near $\{0\} \times S^1 \times S^2$ and with \mathfrak{q}_f near $\{1\} \times S^1 \times S^2$.

To simplify notation we write \mathcal{I}_k for the k th copy of $[0, 1] \times S^1 \times S^2$ and $\mathcal{I} := \bigsqcup_{k=1}^N \mathcal{I}_k$. Consider the moduli space $\mathfrak{M}(X(T), \mathfrak{s})$ for the manifold equipped with this perturbation. As $T \in [0, \infty)$ varies, these form a parametrized moduli space

$$\mathcal{M}(X, \mathfrak{s}) := \bigcup_{T \in [0, \infty)} \{T\} \times \mathfrak{M}(X(T), \mathfrak{s})$$

This has a compactification

$$\mathcal{M}^+(X, \mathfrak{s}) := \bigcup_{T \in [0, \infty]} \{T\} \times \mathfrak{M}(X(T), \mathfrak{s})$$

formed by attaching a fiber at $T = \infty$, where $\mathfrak{M}(X(\infty), \mathfrak{s})$ is defined to be the set of quintuples $(\mathfrak{d}_0, \check{\mathfrak{d}}_1, \mathfrak{d}_2, \check{\mathfrak{d}}_3, \mathfrak{d}_4)$ such that

$$\begin{aligned} \mathfrak{d}_0 &\in M(\emptyset, \mathcal{N}, \mathfrak{c}_{i_1}; \mathfrak{s}) \\ \check{\mathfrak{d}}_1 &\in \check{M}^+(\mathfrak{c}_{i_1}, \mathfrak{c}_{i_2}) \\ \mathfrak{d}_2 &\in M(\mathfrak{c}_{i_2}, \mathcal{I}, \mathfrak{c}_{i_3}; \mathfrak{s}_\xi + 1) \\ \check{\mathfrak{d}}_3 &\in \check{M}^+(\mathfrak{c}_{i_3}, \mathfrak{c}_{i_4}) \\ \mathfrak{d}_4 &\in \mathfrak{M}(\mathfrak{c}_{i_4}, X_0, \emptyset; \mathfrak{s}_\mathfrak{s}) \end{aligned}$$

The space $\mathcal{M}^+(X, \mathfrak{s})$ is stratified by manifolds, its codimesion-1 strata consisting of the fiber $\mathfrak{M}(X, \mathfrak{s})$ over $T = 0$ and those strata over $T = \infty$ with $\mathfrak{c}_{i_1} = \mathfrak{c}_{i_2}$ and $\mathfrak{c}_{i_3} = \mathfrak{c}_{i_4}$ (so that $\check{\mathfrak{d}}_1$ and $\check{\mathfrak{d}}_3$ belong to point moduli spaces). The latter strata are of the form

$$M(\emptyset, \mathcal{N}, \mathfrak{c}_1; \mathfrak{s}) \times M(\mathfrak{c}_1, \mathcal{I}, \mathfrak{c}_2; \mathfrak{s}_\xi + 1) \times \mathfrak{M}(\mathfrak{c}_2, X_0, \emptyset; \mathfrak{s}_\mathfrak{s})$$

where each of $\mathfrak{c}_1 = \{\mathfrak{c}_1^k\}_{1 \leq k \leq N}$ and $\mathfrak{c}_2 = \{\mathfrak{c}_2^k\}_{1 \leq k \leq N}$ is a critical point on \mathcal{I} associated with the perturbations \mathfrak{q}_f and \mathfrak{q}_λ , respectively.

We now incorporate the point and loop constraints that are used to define the closed SW invariant. As \bar{z} and $\bar{\eta}$ sit inside $X_0 \subset X$ and X may be written as the composition of cobordisms $X_0 \circ \mathcal{I} \circ \mathcal{N}$, we may decompose the element

$$u := U^{\frac{1}{2}(d(\mathfrak{s})-p)}[\eta_1] \wedge \cdots \wedge [\eta_p] \in \mathbb{A}(X)$$

as the product

$$u = R_{\mathcal{N}}^*(1) \smile R_{\mathcal{I}}^*(1) \smile R_{X_0}^*(u_0)$$

where $1 \in H^0(\mathcal{B}^\sigma(\mathcal{N}); \mathbb{Z})$, $1 \in H^0(\mathcal{B}^\sigma(\mathcal{I}); \mathbb{Z})$, $u_0 \in H^{d(\mathfrak{s})}(\mathcal{B}^\sigma(X_0); \mathbb{Z})$, and $R_W : \mathcal{B}^\sigma(X) \rightarrow \mathcal{B}^\sigma(W)$ is the restriction map associated with $W \in \{\mathcal{N}, \mathcal{I}, X_0\}$; the product operation is defined in [33, §23.2].

There is a continuous map defined in [33, §26.1],

$$r : \mathcal{M}^+(X, \mathfrak{s}) \rightarrow [0, \infty] \times \mathcal{B}^\sigma(\mathcal{N}, \mathfrak{s}) \times \mathcal{B}^\sigma(\mathcal{I}, \mathfrak{s}_\xi + 1) \times \mathcal{B}^\sigma(X_0, \mathfrak{s}_\mathfrak{s})$$

which, in particular, is given by $(T, \mathfrak{d}) \mapsto (T, \mathfrak{d}|_{\mathcal{N}}, \mathfrak{d}|_{\mathcal{I}}, \mathfrak{d}|_{X_0})$ for $T < \infty$. Likewise, the image $r(\mathcal{M}^+(X, \mathfrak{s}))$ is also stratified by manifolds, and the only relevant strata which pair nontrivially with the cocycle $1 \times 1 \times u_0$ are determined by

$$\dim M(\emptyset, \mathcal{N}_k, \mathfrak{c}_1^k; \mathfrak{s}) = 0 \quad (2.29)$$

$$\dim M(\mathfrak{c}_1^k, \mathcal{I}_k, \mathfrak{c}_2^k; \mathfrak{s}_\xi + 1) = 0 \quad (2.30)$$

$$\dim \mathfrak{M}(\mathfrak{c}_2, X_0, \emptyset; \mathfrak{s}_\mathfrak{s}) = d(\mathfrak{s}) \quad (2.31)$$

for each k . The dimensions of $M(\emptyset, \mathcal{N}_k, \mathfrak{c}_1^k; \mathfrak{s})$ are computed in Lemma 2.7.4 below, from which it follows that (2.29) forces $\mathfrak{c}_1^k = \mathfrak{a}_{-1}$. Then (2.30) forces $|\mathfrak{c}_2^k| = |\mathfrak{c}_1^k|$, which follows immediately from the definition of the grading (see [33, §22.3]), or less directly from the fact that such a product cobordism induces an isomorphism on all monopole Floer (co)homologies. As explained in [33, §36], $\widehat{HM}^{[\xi_*]}(S^1 \times S^2, \mathfrak{s}_\xi + 1)$ is isomorphic to $\widetilde{HM}_{[\xi_*]}(S^1 \times S^2, \mathfrak{s}_\xi + 1)$, with the former generated by \mathfrak{a}_{-1} and the latter generated by \mathfrak{b}_0 .¹⁶ Therefore, each \mathfrak{c}_2^k must be one of the finitely many irreducible generators $\mathfrak{c} \in \widehat{CM}^{[\xi_*]}(S^1 \times S^2, \mathfrak{s}_\xi + 1, \mathfrak{q}_\lambda)$ such that $\mathfrak{M}(\mathfrak{c}, X_0, \emptyset; \mathfrak{s}_\mathfrak{s})$ has dimension $d(\mathfrak{s})$, and hence $g(\mathfrak{s}) = N[\xi_*]$.

Notation 2.7.2. Here and in what follows, we abuse notation by letting \mathfrak{c} and \mathfrak{a}_{-1} denote either the respective monopoles \mathfrak{c}^k and \mathfrak{a}_{-1} on a single component of ∂X_0 or the respective collections $\{\mathfrak{c}^k\}_{1 \leq k \leq N}$ and $\{\mathfrak{a}_{-1}\}_{1 \leq k \leq N}$. Also, we refer to \mathfrak{c} as both a monopole and a cochain in \widehat{CM}^* while $\hat{\mathfrak{c}}$ denotes the corresponding chain in \widehat{CM}_* which pairs nontrivially with the cochain \mathfrak{c} , i.e. $\mathfrak{c}(\hat{\mathfrak{c}}) = 1$.

It follows from Lemma 2.7.4 below that the moduli space $M(\emptyset, \mathcal{N}_k, \mathfrak{a}_{-1}; \mathfrak{s})$ is a point. The version of Stokes' theorem in [33], applied to the $\mathbb{Z}/2\mathbb{Z}$ -pairing of $r(\mathcal{M}^+(X, \mathfrak{s}))$ with $\delta(1 \times 1 \times u_0) = 0$, therefore implies

$$\begin{aligned} SW_X(\mathfrak{s})([\eta_1] \wedge \cdots \wedge [\eta_p]) &= \langle u, [\mathfrak{M}(\mathfrak{s})] \rangle \\ &= \sum_{\mathfrak{c} \in N[\xi_*]} \langle 1, [M(\mathfrak{a}_{-1}, \mathcal{I}, \mathfrak{c}; \mathfrak{s}_\xi + 1)] \rangle \cdot \langle u_0, [\mathfrak{M}(\mathfrak{c}, X_0, \emptyset; \mathfrak{s}_\mathfrak{s})] \rangle \\ &= \sum_{\mathfrak{c} \in N[\xi_*]} \mathbb{M}_\mathfrak{c} \mathfrak{M}_\mathfrak{c} \in \mathbb{Z}/2\mathbb{Z} \end{aligned} \quad (2.32)$$

where $\mathbb{M}_\mathfrak{c}$ denotes the count of points in the 0-dimensional moduli space $M(\mathfrak{a}_{-1}, \mathcal{I}, \mathfrak{c}; \mathfrak{s}_\xi + 1)$, and $\mathfrak{M}_\mathfrak{c}$ is defined by (2.28).

The k th chain complex is $\widehat{CM}_{[\xi_*]}(S^1 \times S^2, \mathfrak{s}_\xi + 1, \mathfrak{q}_f) = \mathbb{Z}/2\mathbb{Z}\langle \hat{\mathfrak{a}}_{-1} \rangle$. The cobordism \mathcal{I}_k induces a chain map

$$\hat{m}_k : \widehat{CM}_*(S^1 \times S^2, \mathfrak{s}_\xi + 1, \mathfrak{q}_f) \rightarrow \widehat{CM}_*(S^1 \times S^2, \mathfrak{s}_\xi + 1, \mathfrak{q}_\lambda)$$

¹⁶Both \mathfrak{a}_{-1} and \mathfrak{b}_0 belong to the same absolute grading in $J(S^1 \times S^2, \mathfrak{s}_\xi + 1)$ because their \mathbb{Z} -grading difference is $\text{gr}[\mathfrak{a}_{-1}, \mathfrak{b}_0] = 0$ (see [33, Equation 16.9, Equation 36.1]).

which is a quasi-isomorphism (and the isomorphism on homology is canonical, see [33, Corollary 23.1.6]). In grading $[\xi_*]$ it is given by

$$z\hat{\mathbf{a}}_{-1} \mapsto z \sum_{\mathbf{c}^k \in [\xi_*]} \mathbb{M}_{\mathbf{c}^k}^k \hat{\mathbf{c}}^k$$

where $z \in \mathbb{Z}/2\mathbb{Z}$ and $\mathbb{M}_{\mathbf{c}^k}^k$ denotes the count of points in the 0-dimensional moduli space $M(\mathbf{a}_{-1}, \mathcal{I}_k, \mathbf{c}^k; \mathfrak{s}_\xi + 1)$. Thus \mathcal{I} induces the chain map

$$\hat{m} = \bigotimes_{k=1}^N \hat{m}_k : \bigotimes_{k=1}^N \widehat{CM}_*(S^1 \times S^2, \mathfrak{s}_\xi + 1, \mathfrak{q}_f) \rightarrow \bigotimes_{k=1}^N \widehat{CM}_*(S^1 \times S^2, \mathfrak{s}_\xi + 1, \mathfrak{q}_\lambda)$$

which is also a quasi-isomorphism, and in grading $N[\xi_*]$ it sends $\hat{\mathbf{a}}_{-1}$ to $\sum_{\mathbf{c} \in N[\xi_*]} \mathbb{M}_{\mathbf{c}} \hat{\mathbf{c}}$ because $\mathbb{M}_{\mathbf{c}} = \prod_{k=1}^N \mathbb{M}_{\mathbf{c}^k}^k$ (the index $\mathbf{c} \in N[\xi_*]$ here means that $\mathbf{c}^k \in [\xi_*]$ for all k). Since $[\hat{m}(\hat{\mathbf{a}}_{-1})]$ is the positive generator of $\bigotimes_{k=1}^N \widehat{HM}_{[\xi_*]}(S^1 \times S^2, \mathfrak{s}_\xi + 1) \cong \mathbb{Z}/2\mathbb{Z}$, the near-symplectic Gromov invariant is the evaluation of the SW cocycle $\Phi_{SW} \in \bigotimes_{k=1}^N \widehat{CM}^{[\xi_*]}(S^1 \times S^2, \mathfrak{s}_\xi + 1, \mathfrak{q}_\lambda)$ at the cycle $\hat{m}(\hat{\mathbf{a}}_{-1})$,

$$Gr_{X, \omega}(\mathfrak{s})([\eta_1] \wedge \cdots \wedge [\eta_p]) = \Phi_{SW} \left(\sum_{\mathbf{c} \in N[\xi_*]} \mathbb{M}_{\mathbf{c}} \hat{\mathbf{c}} \right) = \sum_{\mathbf{c} \in N[\xi_*]} \mathbb{M}_{\mathbf{c}} \mathfrak{M}_{\mathbf{c}} \in \mathbb{Z}/2\mathbb{Z}$$

This number is precisely that in (2.32), so the proof is complete.

Remark 2.7.3. Equivalently, because the chain map \hat{m}_k is a quasi-isomorphism, we know that the cochain map in grading $[\xi_*]$

$$\hat{m}_k^* : \widehat{CM}^{[\xi_*]}(S^1 \times S^2, \mathfrak{s}_\xi + 1, \mathfrak{q}_\lambda) \rightarrow \widehat{CM}^{[\xi_*]}(S^1 \times S^2, \mathfrak{s}_\xi + 1, \mathfrak{q}_f), \quad c \mapsto c \circ \hat{m}_k$$

is also quasi-isomorphism. Then $Gr_{X, \omega}(\mathfrak{s})([\eta_1] \wedge \cdots \wedge [\eta_p])$ is equal to the evaluation of $\hat{m}^*(\Phi_{SW})$ at the cocycle \mathbf{a}_{-1} .

□

Lemma 2.7.4. *For sufficiently small perturbations \mathfrak{p}_N and \mathfrak{q}_f , the moduli spaces $M(\emptyset, S^1 \times B^3, \mathbf{a}_i; \mathfrak{s})$ and $M(\emptyset, S^1 \times B^3, \mathbf{b}_i; \mathfrak{s})$ are empty for $i \geq 0$. The moduli space $M(\emptyset, S^1 \times B^3, \mathbf{a}_{-i}; \mathfrak{s})$ has dimension $2i - 2$ for $i \geq 1$, such that $M(\emptyset, S^1 \times B^3, \mathbf{a}_{-1}; \mathfrak{s})$ is a point, and the moduli space $M(\emptyset, S^1 \times B^3, \mathbf{b}_{-i}; \mathfrak{s})$ has dimension $2i - 1$ for $i \geq 1$.*

Proof. We mimic the analogous proof for a 4-ball B^4 , given by [33, Lemma 27.4.2]. The key point is that both $S^1 \times B^3$ and B^4 have metrics of positive scalar curvature and have trivial 2nd (co)homology. We have already discussed the \mathfrak{q}_f^σ -perturbed SW solutions over $(S^1 \times S^2, \mathfrak{s}_\xi + 1)$ at the beginning of Theorem 2.7.1, and we continue to use that notation.

With respect to the unperturbed SW equations (without blowing up) over $\overline{S^1 \times B^3}$, there are no solutions with nonzero spinor Φ that decay to zero on the cylindrical end $[0, \infty) \times S^1 \times S^2$ because the scalar curvature is positive (see the integration-by-parts trick of [33, Proposition 4.6.1]). For \mathfrak{p}_N sufficiently small the reducible solutions persist and are asymptotic to $\alpha = [\mathbf{A}_\alpha, 0]$ or $\beta = [\mathbf{A}_\beta, 0]$.

Now, the restriction of a reducible solution $\mathfrak{d} \in M(\emptyset, S^1 \times B^3, \mathfrak{a}_i; \mathfrak{s})$ to the cylindrical end is a path

$$\check{\mathfrak{d}}(t) = [\mathbf{A}_\mathfrak{a}, 0, \psi(t)]$$

satisfying $\psi(t) \rightarrow \psi_i$, where $\mathfrak{a}_i = [\mathbf{A}_\alpha, 0, \psi_i]$. Following the argument of [33, Proposition 14.6.1], we would then obtain a nonzero solution to the perturbed Dirac equation on $\overline{S^1 \times B^3}$ with asymptotics $Ce^{-\lambda_i t} \psi_i$ (nonzero constant C) as $t \rightarrow \infty$ on the cylindrical end. If $\lambda_i > 0$ then we just argued that such spinors cannot exist, so $M(\emptyset, S^1 \times B^3, \mathfrak{a}_i; \mathfrak{s})$ is empty. If $\lambda_{-i} < 0$ then, as in [33, Proposition 14.6.1], such spinors with growth bound $Ce^{-\lambda_{-i} t}$ have the form $\sum_{k=-i}^{-1} c_k e^{-\lambda_k t} \psi_k$ on the cylindrical end. Thus $\dim M(\emptyset, S^1 \times B^3, \mathfrak{a}_{-i}; \mathfrak{s}) = 2(i-1)$ and $M(\emptyset, S^1 \times B^3, \mathfrak{a}_{-1}; \mathfrak{s})$ is a single point.

On the other hand, the restriction of a reducible solution $\mathfrak{d} \in M(\emptyset, S^1 \times B^3, \mathfrak{b}_i; \mathfrak{s})$ to the cylindrical end is a path

$$\check{\mathfrak{d}}(t) = [\mathbf{A}(t), 0, \psi(t)]$$

with $\mathbf{A}(t)$ a trajectory lying over a Morse flowline of f on $\mathbb{T} \subset \mathcal{B}(S^1 \times S^2, \mathfrak{s}_\xi + 1)$ and asymptotic to \mathbf{A}_β . We compute $\dim M(\emptyset, S^1 \times B^3, \mathfrak{b}_{-i}; \mathfrak{s})$ for $i \geq 1$ indirectly, thanks to the formal dimension formula

$$\dim M(\emptyset, S^1 \times B^3, \mathfrak{b}_{-i}; \mathfrak{s}) = \dim M(\emptyset, S^1 \times B^3, \mathfrak{a}_{-i}; \mathfrak{s}) + \text{gr}[\mathfrak{a}_{-i}, \mathfrak{b}_{-i}]$$

given by [33, Proposition 24.4.6]. Since $\text{gr}[\mathfrak{a}_{-i}, \mathfrak{b}_{-i}] = \text{ind}_f(\alpha) - \text{ind}_f(\beta) = 1$, the dimension of $M(\emptyset, S^1 \times B^3, \mathfrak{b}_{-i}; \mathfrak{s})$ must be $2(i-1)+1$. Finally, by [33, Proposition 24.4.3] we see that $M(\emptyset, S^1 \times B^3, \mathfrak{b}_i; \mathfrak{s})$ must be empty for $i \geq 0$ because \mathfrak{b}_i is boundary-stable and there are no irreducible SW instantons on $S^1 \times B^3$. \square

Remark 2.7.5. Mrowka mentioned to the author the following heuristic, which can be made precise. The holonomy map

$$\text{hol} : \mathfrak{M}(X, \mathfrak{s}) \rightarrow \prod_{k=1}^N U(1)$$

along the N zero-circles of ω is cobordant to the restriction map

$$\text{res} : M(X_0, \mathfrak{s}_\mathfrak{s}) \rightarrow \prod_{k=1}^N \mathcal{B}^{\text{red}}(S^1 \times S^2, \mathfrak{s}_\xi + 1)$$

where $M(X_0, \mathfrak{s}_\mathfrak{s})$ is the moduli space of SW solutions to the unperturbed equations on X_0 without blowing up. The cobordism is defined by “stretching the neck,” and we identify $\mathcal{B}^{\text{red}}(S^1 \times S^2, \mathfrak{s}_\xi + 1)$ with $U(1)$ by taking the holonomy of a flat connection along the S^1 -factor of $S^1 \times S^2$.

2.8 Appendix – some explanation to Taubes’ analysis

This paper can be viewed as a sort of amalgam of [55] and [63]. There was a complicated feature of [63] that did not arise in [55], and a complicated feature of [55] that did not arise in [63], and both features appeared in this paper. Namely, the multiply covered tori in [63] had to be delicately counted and the map to the Seiberg-Witten moduli space required the use of Kuranishi structures, and \mathbb{R} -invariance in [55] played a complicating role for the analysis associated with non- \mathbb{R} -invariant holomorphic curves. This latter complication came from the existence of multiple ends of a curve hitting the same orbit, or a single end hitting an orbit with multiplicity. An elaboration is given in [55, §5.c.1], and a slightly different elaboration is given below.

In the compact scenario, in order to build a Seiberg-Witten solution from a pseudoholomorphic curve C with n ends approaching a single orbit γ , an n -vortex solution is “grafted” into the normal bundle N_C , and a disk-subbundle of N_C is embedded into the ambient 4-manifold. In the noncompact scenario, that means a 1-vortex solution would be “grafted” into N_C for each end of C , but to ensure embeddedness of a disk-subbundle the radii of the fibers would need to get smaller as γ is approached. Subsequently, the Dirac operator of the corresponding spinor would involve derivatives of the radial coordinate of the fibers, ultimately preventing us from getting the appropriate bounds on the spinor needed to obtain a nearby Seiberg-Witten solution.

To resolve this issue, a 1-vortex solution is not “grafted” into N_C for each of its n ends. Instead, consider the normal bundle $N_{\mathbb{R} \times \gamma}$ of the cylinder $\mathbb{R} \times \gamma$. Then N_C and $N_{\mathbb{R} \times \gamma}$ are “nearby” to each other along the ends of C and $\mathbb{R} \times \gamma$, and objects defined on them can be compared using cutoff-functions and a change of variables. Over a point in the cylinder, the curve hits the disk-fiber in n points, and an n -vortex solution is “grafted” into $N_{\mathbb{R} \times \gamma}$ whose zeros are those n points. This solution is then compared to the would-be solution from the original approach, and seen to be approximately the same except for the worry of varying radial coordinates.

This approach is inspired by the following fact in vortex theory: Given two 1-vortices spaced far apart in \mathbb{R}^2 and one 2-vortex in \mathbb{R}^2 whose zeros are located at the two 1-vortices, the difference between the pair of 1-vortices and the single 2-vortex is exponentially small with respect to the distance between the two 1-vortices.

Chapter 3

“Riemann-Roch” for punctured curves

In [54], Taubes proved the Riemann-Roch theorem for compact Riemann surfaces, as a by-product of taking clever perturbations of the Cauchy-Riemann operator in order to define his Gromov invariant for pseudoholomorphic curves. We will do the same for noncompact surfaces, that is, we will recover the formula for the Fredholm index of a Cauchy-Riemann operator that is asymptotic to nondegenerate asymptotic operators. What follows is a sketch.

Acknowledgements. The main idea stems from a chat with Cliff Taubes on how to extend his transversality results from closed symplectic manifolds to symplectic cobordisms, and builds off of related work with Chris Wendl on closed symplectic manifolds [10]. I appreciate them and their inspiration. Subsequently, a detailed version of this result (with a different style) now appears in Chris Wendl’s book on Symplectic Field Theory [78, §5].

3.1 Setup

Suppose that C is a connected Riemann surface with punctures, and that $E \rightarrow C$ is a holomorphic line¹ bundle with fixed trivialization τ near the punctures. Compactify C so that the neighborhood of the punctures are modeled on $[0, \infty) \times S^1$, and for each puncture denote the corresponding circle by γ . Take a Cauchy-Riemann type operator

$$\mathbf{D} : \Gamma(E) \rightarrow \Gamma(T^{0,1}C \otimes E)$$

with its asymptotic operators

$$\mathbf{A}_\gamma : \Gamma(\gamma^* E) \rightarrow \Gamma(\gamma^* E)$$

for all punctures. With respect to $\tau : \gamma^* E \xrightarrow{\cong} S^1 \times \mathbb{C}$ the asymptotic operator takes the form

$$\mathbf{A}_\gamma = i\partial_t + A_\gamma(t)$$

¹The technique also works for higher rank bundles: take determinants and use the Splitting Lemma.

where $A_\gamma(t)$ is a smooth loop of symmetric matrices. Assume all asymptotic operators are nondegenerate (0 is not an eigenvalue) so that \mathbf{D} is a Fredholm operator.

Example 3.1.1. In practice, C is a pseudoholomorphic curve in a 4-dimensional symplectic cobordism between contact 3-manifolds, with punctures asymptotically approaching non-degenerate Reeb orbits. Here, E is its normal bundle and \mathbf{D} is the (normal) deformation operator.

Theorem 3.1.2 (Schwarz [49]). *The Fredholm index of \mathbf{D} is given by*

$$\text{ind}(\mathbf{D}) = \chi(C) + 2c_1(E, \tau) + \sum_{\gamma} CZ_{\tau}(\mathbf{A}_{\gamma})$$

where $CZ_{\tau}(\mathbf{A}_{\gamma})$ is the Conley-Zehnder index and $c_1(E, \tau)$ is the relative 1st Chern class.

Here is an outline of a novel proof of the index formula, using analytic perturbation theory as in [32]:

- 1) Construct an “L-flat approximation” for all \mathbf{A}_{γ} , i.e. a suitably nice asymptotic operator which \mathbf{A}_{γ} is homotopic to, such that $\text{ind}(\mathbf{D})$ doesn’t vary.
- 2) Prove $\text{ind}(\mathbf{D}) = \text{ind}(\mathbf{D} + B)$ for some $B \in \Gamma(T^{0,1}C \otimes E^2)$ whose winding number along the ends of C with respect to τ satisfies $\text{wind}_{\tau}(B_{\gamma}) = CZ_{\tau}(\mathbf{A}_{\gamma})$.
- 3) Prove $\text{ind}(\mathbf{D} + rB) = \#B^{-1}(0)$ by taking sufficiently large $r \in \mathbb{R}_+$, this being a concentration principle.

We put these steps together, noting that

$$\#B^{-1}(0) = c_1(T^{0,1}C \otimes E^2, \tau) + \text{wind}_{\tau}(B)$$

by definition, to get

$$\text{ind}(\mathbf{D}) = c_1(T^{0,1}C, \tau) + 2c_1(E, \tau) + \sum_{\gamma} \text{wind}_{\tau}(B_{\gamma}) = \chi(C) + 2c_1(E, \tau) + \sum_{\gamma} CZ_{\tau}(\mathbf{A}_{\gamma})$$

The proof of (1) is given in [55, Appendix]. The proof of (3) is a regurgitation of [54, §7], because the argument is local in nature. A new contribution is (2), which did not arise in [54] because there were no punctures (hence no asymptotic constraints).

Remark 3.1.3. The Riemann-Roch formula for the case of closed Riemann surfaces follows by additivity of the Fredholm index with respect to gluing the punctured surfaces at their cylindrical ends. The Riemann-Roch formula can be recasted in the following form: the index equals twice the degree of the bundle plus the Euler characteristic of the surface. When we introduce punctures, the definition of the degree (as a relative 1st Chern class) requires some choice of trivialization of the bundle along the ends of the surface, and a different choice might change the degree. To get an invariant, a ‘boundary correction’ term is needed to compensate – this is the Conley-Zehnder index.

Remark 3.1.4. In using the Conley-Zehnder index, we will suppose the circles γ are non-degenerate (possibly multiply covered) Reeb orbits.

3.2 L -flat approximation

The point here is to reduce everything WLOG to the case that our Fredholm operator \mathbf{D} has asymptotic operators with “nice” explicit descriptions, so that we can do hands-on computations in the subsequent section. As explained in [55, Appendix], there is a homotopy through nondegenerate operators (with fixed Conley-Zehnder index) from a given asymptotic operator \mathbf{A}_γ to an L -flat asymptotic operator. The notion of L -flatness is explained in Section 1.2. The homotopy extends to a homotopy of Fredholm operators of constant index from \mathbf{D} to a Cauchy-Riemann type operator with L -flat asymptotics, and we abuse notation by denoting the resulting operator \mathbf{D} with asymptotic operators $\{\mathbf{A}_\gamma\}$.

We fix L greater than the symplectic action of γ . We also fix the trivialization τ so that the Conley-Zehnder indices are as specified in Section 1.2. For multiply covered orbits we can pull back the asymptotic operator of the underlying embedded orbit, so we assume γ is embedded. The result of how an L -flat $A_\gamma(t)$ acts on $\eta(t) \in \Gamma(\gamma^* E)$ is given as follows (see [55, Lemma 2.3]):

(1) If γ is an embedded elliptic orbit (irrational rotation number θ) then

$$\eta \mapsto \theta \eta$$

(2) If γ is an embedded positive hyperbolic orbit (rotation number 0) then for some sufficiently small positive constant ε .

$$\eta \mapsto \varepsilon i \bar{\eta}$$

(3) If γ is an embedded negative hyperbolic orbit (rotation number 1) then for some sufficiently small positive constant ε

$$\eta \mapsto \frac{1}{2}\eta + \varepsilon i e^{it} \bar{\eta}$$

3.3 Spectral flow for the perturbed asymptotic operator

This section prescribes the asymptotic limits (with respect to the fixed trivialization τ) of the desired section $B \in \Gamma(T^{0,1}\Sigma \otimes E \otimes E)$. All asymptotic operators are assumed to be L -flat from now on. The nonzero complex numbers \mathbb{C}^* will be identified as a subset of $GL_2(\mathbb{R})$ via $a + bi \mapsto \begin{pmatrix} a & -b \\ b & a \end{pmatrix}$.

The perturbed operator

$$\mathbf{D}_r := \mathbf{D} + rB$$

will remain Fredholm of constant index for all $r \in \mathbb{R}$ if and only if each perturbed asymptotic operator

$$\mathbf{A}_{\gamma,r} := \mathbf{A}_\gamma + rB_\gamma(t)$$

remains nondegenerate. In other words, $\mathbf{A}_{\gamma,r}$ must not have any spectral flow as a function of r .

Theorem 3.3.1. $\mathbf{A}_{\gamma,r}$ has no spectral flow if and only if $\text{wind}_\tau(B_\gamma) = CZ_\tau(\mathbf{A}_\gamma)$, with the additional constraint that $B_\gamma(0) \notin i\mathbb{R}$ when γ is a positive hyperbolic orbit or an even cover of a negative hyperbolic orbit.

Proof. As the focus is now on a single asymptotic operator, the subscript γ will be dropped. The Reeb orbit γ is an m -fold cover of some embedded orbit α , and is parametrized by $t \in \mathbb{R}/2\pi m\mathbb{Z}$. Note that the perturbed asymptotic operators act on complex-valued sections $\eta(t) \in L^2(\mathbb{R}/2\pi m\mathbb{Z}, \mathbb{C})$, and the study of A_γ reduces to the study of A_α by pull-back via $\mathbb{R}/2\pi m\mathbb{Z} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$.

Suppose α is **elliptic** with rotation number $0 < \theta < 1$. Then $CZ_\tau(\gamma) = 2\lfloor m\theta \rfloor + 1$ and

$$A(t) \cdot \eta(t) = \begin{pmatrix} \theta & 0 \\ 0 & \theta \end{pmatrix} \eta(t) = \theta \eta(t).$$

We may assume that B satisfies $\text{wind}_\tau(B_\gamma) = n \in \mathbb{Z}$ and takes the form

$$B(t) \cdot \eta(t) = B_0 e^{int/m} \bar{\eta}(t), \quad B_0 \in \mathbb{C}^*.$$

We now look for solutions to $\mathbf{A}_r \eta = 0$. This first-order complex differential equation is

$$i \frac{d\eta}{dt} + \theta \eta + r B_0 e^{int/m} \bar{\eta} = 0.$$

Use the Fourier expansion $\eta(t) = \sum_{k \in \mathbb{Z}} a_k e^{ikt/m}$ and compare modes to obtain the recurrence relation

$$(\theta - \frac{k}{m}) a_k + r B_0 \bar{a}_{n-k} = 0.$$

Swapping $k \rightarrow n - k$ gives another recurrence relation (concerning the same coefficients a_k and a_{n-k}), and these two relations combine to give the constraint

$$r^2 = \frac{1}{|B_0|^2} \left(\theta - \frac{k}{m} \right) \left(\theta - \frac{n-k}{m} \right).$$

If $r^2 > 0$ then there is spectral flow, and if $r^2 < 0$ then there is no spectral flow. Indexing over $l \in \{0, \dots, m-1\}$, suppose $\frac{l}{m} < \theta < \frac{l+1}{m}$, so that $CZ_\tau(\gamma) = 2l + 1$. If $n \geq 2l + 2$ then there is spectral flow (choosing $k = l + 1$) because $(\theta - \frac{k}{m})$ and $(\theta - \frac{n-k}{m})$ are both negative. If $n \leq 2l$ then there is spectral flow (choosing $k = l$) because $(\theta - \frac{k}{m})$ and $(\theta - \frac{n-k}{m})$ are both positive. If $n = 2l + 1$ then there is no spectral flow because $(\theta - \frac{k}{m})$ and $(\theta - \frac{n-k}{m})$ are of opposite sign for any k (if $k \leq l$ then $\frac{n-k}{m} \geq \frac{l+1}{m}$, and if $k \geq l + 1$ then $\frac{n-k}{m} \leq \frac{l}{m}$). We’re done.

Suppose α is **positive hyperbolic** with rotation number 0. Then $CZ_\tau(\gamma) = 0$ and

$$A(t) \cdot \eta(t) = \begin{pmatrix} 0 & \varepsilon \\ \varepsilon & 0 \end{pmatrix} \eta(t) = \varepsilon i \bar{\eta}(t)$$

for $\varepsilon \in \mathbb{R}_+$ small. Again we may assume that B satisfies $\text{wind}_\tau(B_\gamma) = n \in \mathbb{Z}$ and takes the form

$$B(t) \cdot \eta(t) = B_0 e^{int/m} \bar{\eta}(t), \quad B_0 \in \mathbb{C}^*.$$

The first-order complex differential equation to solve is now

$$i \frac{d\eta}{dt} + (rB_0 e^{int/m} + \varepsilon i) \bar{\eta} = 0.$$

Use the Fourier expansion $\eta(t) = \sum_{k \in \mathbb{Z}} a_k e^{ikt/m}$ to obtain the recurrence relation

$$\frac{-k}{m} a_k + rB_0 \bar{a}_{n-k} + \varepsilon i \bar{a}_{-k} = 0.$$

There is spectral flow when $n \neq 0$ (see Appendix 3.5), but when $n = 0$ the recurrence relation contradicts itself (unless B_0 is purely complex). Indeed,

$$\frac{-k}{m} a_k + (rB_0 + \varepsilon i) \bar{a}_{-k} = 0$$

which by swapping $k \rightarrow -k$ gives another recurrence relation (concerning the same coefficients a_k and a_{-k}), and these two relations combine to give the constraint

$$|B_0|^2 r^2 + (\bar{B}_0 - B_0) \varepsilon i r + \varepsilon^2 = -\frac{k^2}{m^2}.$$

This is never satisfied unless $B_0 = b_0 i$ (where $b_0 \in \mathbb{R}^*$) and $k = 0$, in which case $r = -\frac{\varepsilon}{b_0}$ (hence $\eta(t) = 1$ is in the kernel). We’re done.

Suppose α is **negative hyperbolic** with rotation number 1. Then $CZ_\tau(\gamma) = m$ and

$$A(t) \cdot \eta(t) = \frac{1}{2} \eta(t) + \varepsilon i e^{it} \bar{\eta}(t)$$

for $\varepsilon \in \mathbb{R}_+$ small. Again we may assume that B satisfies $\text{wind}_\tau(B_\gamma) = n \in \mathbb{Z}$ and takes the form

$$B(t) \cdot \eta(t) = B_0 e^{int/m} \bar{\eta}(t), \quad B_0 \in \mathbb{C}^*.$$

The first-order complex differential equation to solve is now

$$i \frac{d\eta}{dt} + \frac{1}{2} \eta + (rB_0 e^{int/m} + \varepsilon i e^{it}) \bar{\eta} = 0.$$

Use the Fourier expansion $\eta(t) = \sum_{k \in \mathbb{Z}} a_k e^{ikt/m}$ to obtain the recurrence relation

$$\left(\frac{1}{2} - \frac{k}{m}\right) a_k + rB_0 \bar{a}_{n-k} + \varepsilon i \bar{a}_{m-k} = 0.$$

There is spectral flow when $n \neq m$ (see Appendix 3.5), but when $n = m$ the recurrence relation contradicts itself (unless B_0 is purely complex and m is even). Indeed,

$$\left(\frac{1}{2} - \frac{k}{m}\right)a_k + (rB_0 + \varepsilon i)\bar{a}_{m-k} = 0$$

which by swapping $k \rightarrow m - k$ gives another recurrence relation (concerning the same coefficients a_k and a_{m-k}), and these two relations combine to give the constraint

$$|B_0|^2 r^2 + (\bar{B}_0 - B_0)\varepsilon ir + \varepsilon^2 = -\left(\frac{1}{2} - \frac{k}{m}\right)^2.$$

This is never satisfied unless $B_0 = b_0 i$ (where $b_0 \in \mathbb{R}^*$) and $k = \frac{m}{2}$ with m even, in which case $r = -\frac{\varepsilon}{b_0}$ (hence $\eta(t) = e^{it/2}$ is in the kernel). We’re done. \square

3.4 Localization argument

Using homotopy theory (i.e. obstruction theory), we can extend the choices of B_γ given in Theorem 3.3.1 to a global $B \in \Gamma(T^{0,1}C \otimes E^2)$ having nondegenerate zeros. Note that the zeros are bounded away from the punctures due to the existence of the trivialization τ .

Theorem 3.4.1. *For $r \gg 0$ and B as above, $\text{ind}(\mathbf{D} + rB) = \#B^{-1}(0)$.*

Proof. Using a sequence of cutoff functions to invoke a noncompact version of Stokes’ theorem, there is a Bochner-Weitzenböck formula for $\|\mathbf{D}\eta + rB\bar{\eta}\|^2$ in terms of the quantities $r^2 \int_C |B\bar{\eta}|^2$ and $r\Re \int_C \partial B \cdot \bar{\eta}^2$ (here, \Re means the real part). The same localization argument as in [54, §7] shows that the positive zeros of B contribute to the kernel of $\mathbf{D} + rB$ for sufficiently large $r \in \mathbb{R}_+$, while the negative zeros of B contribute to the cokernel. Roughly speaking, the support of any sequence $\eta_r \in \text{Ker}(\mathbf{D} + rB)$ must concentrate near the (positive) zeros of B . The result follows. \square

Remark 3.4.2. This proof hinges on the asymptotic behavior of the perturbation object B . If B had compact support or exponentially decayed to zero at the punctures of C , we could have a sequence $\{(\eta_r, r)\} \in \Gamma(E) \oplus \mathbb{R}_+$ with $\eta_r \in \text{Ker}(\mathbf{D} + rB)$ and $r \rightarrow \infty$ such that $\sup_C(\eta_r)$ “runs away” towards the punctures.

3.5 Appendix

We supply the missing computation in Section 3.3 but only in the case that γ is an embedded positive hyperbolic orbit and the perturbation is $B_\gamma(t) = e^{it}$ (so $n = m = B_0 = 1$), because the other cases are handled in the same way. Solving for the kernels of the perturbed asymptotic operators $\mathbf{A}_{\gamma,r}$ is equivalent to solving the following 1st order complex differential equations for smooth functions $\eta : \mathbb{R}/2\pi\mathbb{Z} \rightarrow \mathbb{C} - \{0\}$ defined on the circle,

$$i \frac{\partial \eta}{\partial t} + (re^{it} + \varepsilon i)\bar{\eta} = 0$$

where $\varepsilon \in \mathbb{R}_+$ is sufficiently small and fixed and $r \in \mathbb{R}_+$ is a nonzero positive real parameter. We would like to find the set of all such r for which $\text{Ker } \mathbf{A}_{\gamma,r} \neq \emptyset$, as well as $\dim_{\mathbb{R}} \text{Ker } \mathbf{A}_{\gamma,r}$, though it suffices to prove the existence of one such r .

There are at least three ways to do this. The result is a unique solution to $\mathbf{A}_{\gamma,r} \eta = 0$ which occurs around $r \approx \sqrt{\varepsilon}$. Two methods are given in a MathOverflow post [17]. We now give the third method:

When γ is a hyperbolic orbit, the perturbed asymptotic operator $\mathbf{A}_{\gamma,r}$ has trivial kernel when $r = 0$ and $r = r_* \gg 0$ (use the Bochner-Weitzenböck formula to see this). The Conley-Zehnder index is a \mathbb{Z} -valued invariant of homotopy classes of such operators, and the *net* spectral flow of $r \in [0, r_*] \mapsto \mathbf{A}_{\gamma,r}$ is

$$CZ_\tau(\mathbf{A}_{\gamma,r_*}) - CZ_\tau(\mathbf{A}_\gamma) = \text{wind}_\tau(B_\gamma) - CZ_\tau(\gamma)$$

Thus when $\text{wind}_\tau(B_\gamma) \neq CZ_\tau(\gamma)$ there must exist at least one $r > 0$ for which $\text{Ker } \mathbf{A}_{\gamma,r} \neq 0$.

Bibliography

- [1] D. Auroux, S. Donaldson, and L. Katzarkov. “Singular Lefschetz pencils”. In: *Geom. Top.* 9 (2005), pp. 1043–1114.
- [2] F. Bourgeois and K. Mohnke. “Coherent orientations in symplectic field theory”. In: *Math. Z.* 248.1 (2004), pp. 123–146.
- [3] F. Bourgeois et al. “Compactness results in symplectic field theory”. In: *Geom. Topol.* 7 (2003), pp. 799–888.
- [4] Frédéric Bourgeois. “A Morse-Bott approach to contact homology”. Ph.D. Thesis. 2002.
- [5] Frédéric Bourgeois. “A survey of contact homology”. In: *New perspectives and challenges in symplectic field theory, CRM Proc. Lecture Notes* 49 (2009), pp. 45–71.
- [6] Eugenio Calabi. “An intrinsic characterization of harmonic one-forms”. In: *Global Analysis* (1969), pp. 101–117.
- [7] Daniel Cristofaro-Gardiner. “The absolute gradings on embedded contact homology and Seiberg-Witten Floer cohomology”. In: *Alg. Geom. Top.* 13 (2013), pp. 2239–2260.
- [8] S. Donaldson and P. Kronheimer. *The geometry of four-manifolds*. Oxford Mathematical Monographs. Oxford Science Publications. The Clarendon Press Oxford University Press, 1990, pp. x+440. ISBN: 0-19-853553-8.
- [9] D. Gay and R. Kirby. “Constructing symplectic forms on 4-manifolds which vanish on circles”. In: *Geom. Top.* 8 (2004), pp. 743–777.
- [10] C. Gerig and C. Wendl. “Generic transversality for unbranched covers of closed pseudoholomorphic curves”. In: *Comm. Pure Appl. Math.* 70.3 (2017), pp. 409–443.
- [11] Chris Gerig. “Circle-valued Morse theory and Seiberg-Witten invariants of 3-manifolds”. in preparation.
- [12] Robert Gompf. “Handlebody construction of Stein surfaces”. In: *Annals of Mathematics* 148.2 (1998), pp. 619–693.
- [13] Mikhael Gromov. “Pseudoholomorphic curves in symplectic manifolds”. In: *Invent. Math.* 82.2 (1985), pp. 307–347.

- [14] Ko Honda. “A note on Morse theory of harmonic 1-forms”. In: *Topology* 38.1 (1998), pp. 223–233.
- [15] Ko Honda. “Local properties of self-dual harmonic 2-forms on a 4-manifold”. In: *J. Reine Angew. Math.* 577 (2004), pp. 105–116.
- [16] Ko Honda. “Transversality theorems for harmonic forms”. In: *Rocky Mountain J. Math.* 34.2 (2004), pp. 629–664.
- [17] fedja (<https://mathoverflow.net/users/1131/fedja>). *Solution to at least one ODE in a family of ODE's*. URL: <https://mathoverflow.net/q/280953> (version: 2017-09-13). MathOverflow.
- [18] M. Hutchings and Y.J. Lee. “Circle-valued Morse theory and Reidemeister torsion”. In: *Geom. Top.* 3 (1999), pp. 369–396.
- [19] M. Hutchings and Y.J. Lee. “Circle-valued Morse theory, Reidemeister torsion, and Seiberg-Witten invariants of 3-manifolds”. In: *Topology* 38(4) (1999), pp. 861–888.
- [20] M. Hutchings and M. Sullivan. “Rounding corners of polygons and the embedded contact homology of T^3 ”. In: *Geom. Top.* 10 (2006), pp. 169–266.
- [21] M. Hutchings and C. Taubes. “Gluing pseudoholomorphic curves along branched covered cylinders. I”. In: *J. Symplectic Geom.* 5.1 (2007), pp. 43–137.
- [22] M. Hutchings and C. Taubes. “Gluing pseudoholomorphic curves along branched covered cylinders. II”. In: *J. Symplectic Geom.* 7.1 (2009), pp. 29–133.
- [23] M. Hutchings and C. Taubes. “Proof of the Arnold chord conjecture in three dimensions II”. In: *Geom. Top.* 17 (2013), pp. 2601–2688.
- [24] M. Hutchings and C. Taubes. “The Weinstein conjecture for stable Hamiltonian structures”. In: *Geom. Top.* 13 (2009), pp. 901–941.
- [25] Michael Hutchings. *A magic trick for defining obstruction bundles*. Blog post: <https://floerhomology.wordpress.com/2013/09/04/a-magic-trick-for-defining-obstruction-bundles/> (version: 2013-09-04).
- [26] Michael Hutchings. “Beyond ECH capacities”. In: *Geom. Top.* 20 (2016), pp. 1085–1126.
- [27] Michael Hutchings. “Embedded contact homology as a (symplectic) field theory”. in preparation.
- [28] Michael Hutchings. “Lecture notes on embedded contact homology”. In: *Contact and Symplectic Topology, Bolyai Society Mathematical Studies* 26 (2014), pp. 389–484.
- [29] Michael Hutchings. “Reidemeister torsion in generalized Morse theory”. 2002, pp. 209–244.
- [30] Michael Hutchings. “The embedded contact homology index revisited”. In: *New perspectives and challenges in symplectic field theory, CRM Proc. Lecture Notes* 49 (2009), pp. 263–297.

- [31] A. Jaffe and C. Taubes. *Vortices and monopoles. Structure of static gauge theories.* Progress in Physics 2. Birkhäuser, Boston, 1980.
- [32] Tosio Kato. *Perturbation theory for linear operators.* Classics in Mathematics. Reprint of the 1980 edition. Springer-Verlag, Berlin, 1995, pp. xxii+619. ISBN: 3-540-58661-X.
- [33] P. Kronheimer and T. Mrowka. *Monopoles and three-manifolds.* Vol. 10. New Mathematical Monographs. Cambridge University Press, 2007, pp. xii+796. ISBN: 978-0-521-88022-0.
- [34] P. Kronheimer et al. “Monopoles and lens space surgeries”. In: *Annals of Mathematics* 165 (2007), pp. 457–546.
- [35] C. Kutluhan, Y.-J. Lee, and C. Taubes. “HF=HM V: Seiberg-Witten-Floer homology and handle addition”. In: (2017). preprint arXiv:1204.0115.
- [36] Claude LeBrun. “Yamabe constants and the perturbed Seiberg-Witten equations”. In: *Comm. Anal. Geom.* 5.3 (1997), pp. 535–553.
- [37] T. Li and A. Liu. “The equivalence between SW and Gr in the case where $b^+ = 1$ ”. In: *IMRN* 7 (1999), pp. 335–345.
- [38] Jerrel Mast. “Pseudoholomorphic punctured spheres in the symplectization of a quotient”. Ph.D. Thesis. 2005.
- [39] Dusa McDuff. “Blow ups and symplectic embeddings in dimension 4”. In: *Topology* 30.3 (1992), pp. 409–421.
- [40] Dusa McDuff. “Lectures on Gromov invariants for symplectic 4-manifolds”. In: *NATO ASI Series* 488 (1997).
- [41] Dusa McDuff. “The structure of rational and ruled symplectic 4-manifolds”. In: *J. Amer. Math. Soc.* 3.3 (1990), pp. 679–712.
- [42] Jo Nelson. “Automatic transversality in contact homology II: invariance and computations”. In: (2017). preprint arXiv:math/1708.07220.
- [43] C. Okonek and A. Teleman. “3-dimensional Seiberg-Witten invariants and non-Kählerian geometry”. In: *Math. Ann.* 312 (1998), pp. 261–288.
- [44] P. Ozsváth and Z. Szabó. “The symplectic Thom conjecture”. In: *Annals of Mathematics* 151.1 (2000), pp. 93–124.
- [45] Timothy Perutz. “Surface-fibrations, four-manifolds, and symplectic Floer homology”. Ph.D. Thesis. 2005.
- [46] Timothy Perutz. “Zero-sets of near-symplectic forms”. In: *J. Symp. Geom.* 4.3 (2007), pp. 237–257.
- [47] Yongbin Ruan. “Symplectic topology and complex surfaces”. In: *Topology on Complex Manifolds, (T. Mabuchi, J. Noguchi and T. Ochiai, eds.), World Scientific Publ., Singapore* (1994).

- [48] Dietmar Salamon. “Seiberg-Witten invariants of mapping tori, symplectic fixed points, and Lefschetz numbers”. In: *Turkish J. Math.* 23.1 (1999), pp. 117–143.
- [49] Matthias Schwarz. “Cohomology operations from S^1 -cobordisms in Floer homology”. Ph.D. Thesis. 1995.
- [50] Ralph Scott. “Closed self-dual two-forms on four-dimensional handlebodies”. Ph.D. Thesis. 2003.
- [51] Clifford Taubes. “A compendium of pseudoholomorphic beasts in $\mathbb{R} \times (S^1 \times S^2)$ ”. In: *Geom. Topol.* 6 (2002), pp. 657–814.
- [52] Clifford Taubes. “A proof of a theorem of Luttinger and Simpson about the number of vanishing circles of a near-symplectic form on a 4-dimensional manifold”. In: *Math. Res. Lett.* 13.4 (2006), pp. 557–570.
- [53] Clifford Taubes. “Arbitrary N-Vortex Solutions to the First Order Ginzburg-Landau Equations”. In: *Commun. Math. Phys.* 72 (1980).
- [54] Clifford Taubes. “Counting pseudo-holomorphic submanifolds in dimension 4”. In: *J. Diff. Geom.* 44.4 (1996), pp. 818–893.
- [55] Clifford Taubes. “Embedded contact homology and Seiberg-Witten Floer cohomology I”. In: *Geom. Topol.* 14.5 (2010), pp. 2497–2581.
- [56] Clifford Taubes. “Embedded contact homology and Seiberg-Witten Floer cohomology II”. In: *Geom. Topol.* 14.5 (2010), pp. 2583–2720.
- [57] Clifford Taubes. “Embedded contact homology and Seiberg-Witten Floer cohomology III”. In: *Geom. Topol.* 14.5 (2010), pp. 2721–2817.
- [58] Clifford Taubes. “Embedded contact homology and Seiberg-Witten Floer cohomology IV”. In: *Geom. Topol.* 14.5 (2010). Erratum-ibid. *Correction to a lemma in: Embedded contact homology and Seiberg-Witten Floer cohomology IV*, preprint arXiv:1801.07556, pp. 2819–2960.
- [59] Clifford Taubes. “Embedded contact homology and Seiberg-Witten Floer cohomology V”. In: *Geom. Topol.* 14.5 (2010), pp. 2961–3000.
- [60] Clifford Taubes. “Moduli spaces and Fredholm theory for pseudoholomorphic subvarieties associated to self-dual, harmonic 2-forms”. In: *Asian J. Math.* 3.1 (1999), pp. 275–324.
- [61] Clifford Taubes. “Pseudoholomorphic punctured spheres in $\mathbb{R} \times (S^1 \times S^2)$: moduli space parametrizations”. In: *Geom. Top.* 10 (2006), pp. 1855–2054.
- [62] Clifford Taubes. “Pseudoholomorphic punctured spheres in $\mathbb{R} \times (S^1 \times S^2)$: properties and existence”. In: *Geom. Top.* 10 (2006), pp. 785–928.
- [63] *Seiberg-Witten and Gromov invariants for symplectic 4-manifolds*. Vol. 2. First International Press Lecture Series. Edited by Richard Wentworth. International Press, 2000, pp. vi+401. ISBN: 1-57146-061-6.

- [64] Clifford Taubes. “Seiberg-Witten invariants and pseudo-holomorphic subvarieties for self-dual, harmonic 2-forms”. In: *Geom. Top.* 3 (1999), pp. 167–210.
- [65] Clifford Taubes. “Seiberg-Witten invariants, self-dual harmonic 2-forms and the Hofer-Wysocki-Zehnder formalism”. In: *Surv. Diff. Geom.* VII (2000), pp. 625–672.
- [66] Clifford Taubes. “ $\text{Gr} \Rightarrow \text{SW}$: from pseudo-holomorphic curves to Seiberg-Witten solutions”. In: *J. Diff. Geom.* 51 (1999), pp. 203–334.
- [67] Clifford Taubes. “ $\text{Gr} = \text{SW}$: counting curves and connections”. In: First Int. Press Lect. Ser. 2 (2000), pp. 275–401.
- [68] Clifford Taubes. “ $\text{SW} \Rightarrow \text{Gr}$: from the Seiberg-Witten equations to pseudo-holomorphic curves”. In: *J. Amer. Math. Soc.* 9.3 (1996), pp. 845–918.
- [69] Clifford Taubes. “The geometry of the Seiberg-Witten invariants”. In: *Surv. Diff. Geom.* III (1996), pp. 299–339.
- [70] Clifford Taubes. “The geometry of the Seiberg-Witten invariants”. In: *Doc. Math. J. DMV Extra Volume ICM II* (1998), pp. 493–504.
- [71] Clifford Taubes. “The Seiberg-Witten and Gromov invariants”. In: *Math. Res. Lett.* 2.2 (1995), pp. 221–238.
- [72] Clifford Taubes. “The Seiberg-Witten equations and the Weinstein conjecture”. In: *Geom. Topol.* 11 (2007), pp. 2117–2202.
- [73] Clifford Taubes. “The structure of pseudo-holomorphic subvarieties for a degenerate almost complex structure and symplectic form on $S^1 \times B^3$ ”. In: *Geom. Topol.* 2 (1998), pp. 221–332.
- [74] Vladimir Turaev. “A combinatorial formulation for the Seiberg-Witten invariants of 3-manifolds”. In: *Math. Res. Lett.* 5 (1998), pp. 583–598.
- [75] Vladimir Turaev. “Torsion invariants of Spin^c -structures on 3-manifolds”. In: *Math. Res. Lett.* 4 (1997), pp. 679–695.
- [76] Chris Wendl. “Automatic transversality and orbifolds of punctured holomorphic curves in dimension four”. In: *Comment. Math. Helv.* 85.2 (2010), pp. 347–407.
- [77] Chris Wendl. “Holomorphic curves in low dimensions: from symplectic ruled surfaces to planar contact manifolds”. In: *to appear in Springer Lecture Notes in Mathematics, available at <https://www.mathematik.hu-berlin.de/~wendl/publications.html>* (2018).
- [78] Chris Wendl. “Lectures on symplectic field theory”. In: *preprint arXiv:1612.01009* (2016).
- [79] Chris Wendl. “Transversality and super-rigidity for multiply covered holomorphic curves”. In: *preprint arXiv:1609.09867* (2017).
- [80] Edward Witten. “Monopoles and four-manifolds”. In: *Math. Res. Lett.* 1 (1994), pp. 769–796.