

HEEGAARD FLOER HOMOLOGY, DOUBLE POINTS AND NICE DIAGRAMS

ROBERT LIPSHITZ

ABSTRACT. The Heegaard Floer chain complexes are defined by counting embedded curves in $\Sigma \times [0, 1] \times \mathbb{R}$. In [Lip06], it was shown that the chain complex \widehat{CF} can be elaborated by taking into account curves with double points. In this note, we extend results of Manolescu-Ozsváth-Sarkar and Sarkar-Wang on computing CFK and \widehat{CF} to this elaborated complex. The extension has a particularly nice form for grid diagrams: while the original Heegaard Floer differential counts “empty” rectangles in the grid diagram, the elaborated differential also counts rectangles containing some of the x_i .

1. INTRODUCTION.

In [OS04a], [Ras03] and [OS05], Ozsváth-Szabó and Rasmussen constructed certain “categorifications” of the Alexander polynomial called *knot Floer homology*. These invariants, which come in various flavors, take the form of chain complexes, well defined up to homotopy equivalence. Their original definitions were in terms of holomorphic curves. In [MOS06], Manolescu-Ozsváth-Sarkar showed that these chain complexes have simple combinatorial descriptions in terms of toroidal grid diagrams (see Figure 1). For instance, in an $n \times n$ toroidal grid diagram the complex \widetilde{CFK} is generated over \mathbb{F}_2 by *matchings*, i.e., n -tuples of intersection points between the horizontal and vertical circles with no two points on the same circle. (The generators are, thus, in one-to-one correspondence with permutations of $\{1, \dots, n\}$, though the correspondence depends on a choice of where one cuts the diagram.) If \mathbf{x} and \mathbf{y} are matchings, the coefficient of \mathbf{y} in $d(\mathbf{x})$ is zero unless all but two of the points of \mathbf{x} and \mathbf{y} are the same. If \mathbf{x} and \mathbf{y} differ at exactly two points then the coefficient of \mathbf{y} in $d(\mathbf{x})$ is the number of rectangles with corners at \mathbf{x} and \mathbf{y} not containing any x_i, y_i, w_i or z_i . (See Figure 1 for an example of such a rectangle.)

It is natural to try to elaborate the complex by counting more rectangles. As discussed in [MOS06], relaxing the condition that rectangles

Date: August 18, 2008.




FIGURE 1. Left: A toroidal grid diagram for the figure 8 knot. (Opposite edges of the diagram are identified.) Center: two generators \mathbf{x} (empty circles) and \mathbf{y} (solid squares) of \widetilde{CFK} , and a rectangle (shaded) contributing to the coefficient of \mathbf{y} in $d(\mathbf{x})$. Right: different generators \mathbf{x} , \mathbf{y} and a rectangle contributing to the coefficient of \mathbf{y} in $d'(\mathbf{x})$ but not in $d(\mathbf{x})$.

not cover w_i or z_i leads to other variants of knot Floer homology defined in [OS04a], [Ras03] and [OS05]. In this paper, we discuss what happens if one relaxes instead the condition that rectangles not cover x_i or y_i , giving a new differential d' on \widetilde{CFK} . (The homological grading on (\widetilde{CFK}, d) becomes a filtration on (\widetilde{CFK}, d') .) It turns out that, in this case, one is essentially computing a chain complex defined by the author in [Lip06, Section 14.1], originally constructed by counting holomorphic curves with prescribed singularity. This result, one of the goals of the paper, is Proposition 9.

Knot Floer homology is, roughly, a relative version of the three-manifold invariant *Heegaard Floer homology*. As discussed by Sarkar-Wang in [SW], the essential feature of grid diagrams which allows one to compute the Heegaard Floer differential is that all of the regions in the grid diagram have non-negative Euler measure, i.e., are bigons

or rectangles (in particular, rectangles).¹ For one variant of Heegaard-Floer homology, \widehat{CF} , only regions not containing “basepoints” matter, so one can compute the differential on \widehat{CF} for more general diagrams than grid diagrams. In particular, a diagram is called *nice* if all regions in it not containing basepoints are bigons or squares. It was shown in [SW] that in a nice diagram, the differential on \widehat{CF} corresponds to counting embedded bigons and rectangles in the diagram. (They also showed that, remarkably, every 3-manifold admits a nice diagram.) As with grid diagrams, there is an analogous result for holomorphic curves with double points in nice diagrams, so we work in this setting for most of the paper. The main result on computability of the elaborated differential on \widehat{CF} is Proposition 8.

It remains unknown whether the elaborated complexes contain more information than ordinary Heegaard Floer homology. On the one hand, in every example known to the author, (\widehat{CFK}, d') and (\widehat{CFK}, d) are filtered chain homotopy equivalent. On the other hand, the computational complexity of the Manolescu-Ozsváth-Sarkar algorithm and the Sarkar-Wang algorithm are such that finding examples in which the (\widehat{CFK}, d') and (\widehat{CFK}, d) are not filtered chain homotopy would require some insight. (In particular, for knots which are *HF*K-thin, or close to being *HF*K-thin, as all knots with small crossing number are, (\widehat{CFK}, d') and (\widehat{CFK}, d) are necessarily chain homotopy equivalent.)

This paper is organized as follows. In Section 2, we begin by reviewing the cylindrical formulation of the Heegaard Floer homology group $\widehat{CF}(Y)$, as described in [Lip06]. Following this, we re-introduce the elaborated chain complex $\widehat{CF}_{big}(Y)$ of a closed 3-manifold. This is followed by a discussion of the corresponding elaboration of the link Floer homology groups and a discussion of invariance. In Section 3 we turn to the main results of this paper, on holomorphic curves with double points in “nice diagrams,” generalizing [SW, Theorem 3.2]. We conclude by specializing these results to toroidal grid diagrams, justifying the claims made in the introduction.

Throughout the paper there are a number of remarks, which may be of interest to the reader but are not necessary for the main narrative.

Acknowledgments. I thank Anthony Licata and Peter Ozsváth for many interesting, detailed discussions related to this paper. I also thank Ciprian Manolescu, Dylan Thurston and Zoltán Szabó for clarifying and encouraging conversations, and the referee for several helpful

¹Actually, at least in some situations, a weaker and seemingly more natural condition suffices; see [Bel07].

comments. The main result of this paper (Proposition 8) has been discovered independently by Sucharit Sarkar.

The idea to consider holomorphic curves with double points in Heegaard Floer homology, as well as a great many of the other ideas in the background of this paper, was suggested to me by my patient thesis advisor, Yasha Eliashberg.

2. CURVES WITH DOUBLE POINTS...

2.1. The cylindrical setting. We begin by recalling the “cylindrical formulation” of Heegaard-Floer homology from [Lip06]; for omitted details, we refer the reader there. Fix a closed, oriented 3-manifold Y , a Riemannian metric on Y , and a self-indexing Morse-Smale function $f: Y \rightarrow \mathbb{R}$ with k index 0 critical points and k index 3 critical points. Let n be the number of index 1 critical points of f , which is also the number of index 2 critical points of f . Then $\Sigma = f^{-1}(3/2)$ is an orientable surface of genus $g = n - k + 1$. The ascending spheres $\alpha_1, \dots, \alpha_n$ of the index 1 critical points of f are pairwise disjoint embedded circles in Σ , as are the descending spheres β_1, \dots, β_n of the index 2 critical points of f . Choose also k flow lines $\gamma_1, \dots, \gamma_k$ of ∇f from the index 0 to the index 3 critical points, so that one flow line originates at each index 0 critical point and one terminates at each index 3 critical point. The flow lines $\{\gamma_i\}$ intersect Σ in k points z_1, \dots, z_k , none of which lie on α - or β -circles. The data $(\Sigma_g, \boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_n\}, \boldsymbol{\beta} = \{\beta_1, \dots, \beta_n\}, \mathbf{z} = \{z_1, \dots, z_k\})$ is called a *k-pointed Heegaard diagram for Y* . See [OS05] for more details, or Figure 2 (page 12) for an example. Abusing notation slightly, we will also use $\boldsymbol{\alpha}$ to denote $\alpha_1 \cup \dots \cup \alpha_n$, and $\boldsymbol{\beta}$ to denote $\beta_1 \cup \dots \cup \beta_n$.

If one specifies a second, completely different set of k flow lines η_1, \dots, η_k then $\gamma_1 \cup \dots \cup \gamma_k \cup \eta_1 \cup \dots \cup \eta_k$ is a link L in Y . Letting $w_i = \eta_i \cap \Sigma$ and $\mathbf{w} = \{w_1, \dots, w_k\}$, we call $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}, \mathbf{w})$ a *2k-pointed Heegaard diagram for L* . Again, see [OS05] for more details, or Figure 3 (page 15) for an example.

From now on, we fix a k -pointed Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$ for Y ; we will comment on the extension to link invariants in Section 2.3. We will associate to $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$, together with a little additional data, a chain complex $\widehat{CF}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}) = \widehat{CF}(Y, k)$; as the notation suggests, up to chain homotopy, the chain complex depends only on the three-manifold Y and the integer k , and not on the particular Heegaard diagram. We denote $\widehat{CF}(Y, 1)$ by $\widehat{CF}(Y)$. (The chain complex $\widehat{CF}(Y)$ and $\widehat{CF}(Y, k)$ are essentially the same as in [OS04b]. In Section 2.2, we will introduce the elaborations of $\widehat{CF}(Y)$ which will be our main

objects of study.) In fact, $\widehat{CF}(Y, k)$ decomposes as a direct sum of subcomplexes indexed by $spin^c$ -structures on Y ,

$$\widehat{CF}(Y, k) = \bigoplus_{s \in spin^c(Y)} \widehat{CF}(Y, k, s),$$

and the elaborated complexes decompose in the same way. We shall, however, usually suppress this fact.

In order to define the chain complexes, we need to fix one piece of auxiliary data: an almost complex structure J on the manifold $\Sigma \times [0, 1] \times \mathbb{R}$. We require that J satisfy certain properties, as described in [Lip06, Section 1]; in particular,

- (J1) J is invariant under \mathbb{R} -translation,
- (J2) the projection map $\pi_{\mathbb{D}}: \Sigma \times [0, 1] \times \mathbb{R} \rightarrow [0, 1] \times \mathbb{R} \subset \mathbb{C}$ is (J, i) -holomorphic and
- (J3) the fibers of the projection map $\pi_{\Sigma}: \Sigma \times [0, 1] \times \mathbb{R} \rightarrow \Sigma$ are J -holomorphic.

By a *matching* we mean an n -tuples of points $\mathbf{x} = \{x_i\}_{i=1}^n \subset (\boldsymbol{\alpha} \cap \boldsymbol{\beta})$ such that (exactly) one x_i lies on each α_i and one lies on each β_j . In this note, we are interested in the Heegaard-Floer chain complex \widehat{CF} , which is generated over \mathbb{F}_2 by the set of matchings. The differential $d: \widehat{CF} \rightarrow \widehat{CF}$ is defined by counting holomorphic curves in $\Sigma \times [0, 1] \times \mathbb{R}$, with asymptotics specified by matchings. More precisely, we consider J -holomorphic maps

$$(1) \quad u: (S, \partial S) \rightarrow (\Sigma \times [0, 1] \times \mathbb{R}, (\boldsymbol{\alpha} \times \{1\} \times \mathbb{R}) \cup (\boldsymbol{\beta} \times \{0\} \times \mathbb{R}))$$

where S is a (smooth) Riemann surface with boundary and punctures on the boundary. We impose several conditions on these maps:

- (M1) the map u is asymptotic to $\mathbf{x} \times [0, 1] \times \mathbb{R}$ near $-\infty$ and $\mathbf{y} \times [0, 1] \times \mathbb{R}$ near $+\infty$ for some matchings \mathbf{x} and \mathbf{y} ,
- (M2) the map u is proper,
- (M3) For each i , exactly one component of ∂S is mapped to $\alpha_i \times \{1\} \times \mathbb{R}$ and exactly one component of ∂S is mapped to $\beta_i \times \{0\} \times \mathbb{R}$,
- (M4) the source S has no closed components and
- (M5) the map u is an embedding.

We do *not* require that S be connected – in general, S will have many connected components. Note that these conditions are somewhat redundant. For instance, conditions (M5) and (M1) imply condition (M3). Note also that conditions (M4) and (M2) imply that $\pi_{\mathbb{D}} \circ u$ is nonconstant on every component of S .

Every topological map u as in Formula (1) satisfying properties (M1) and (M2) has a well-defined (relative) homology class in $H_2(\Sigma, \boldsymbol{\alpha} \cup \boldsymbol{\beta})$, also referred to (unfortunately) as the *domain* $D(u)$ of the map u . For

fixed asymptotics (matchings) \mathbf{x} and \mathbf{y} , these relative homology classes form a subset of $H_2(\Sigma, \boldsymbol{\alpha} \cup \boldsymbol{\beta})$ which we denote $\pi_2(\mathbf{x}, \mathbf{y})$. It turns out that $\pi_2(\mathbf{x}, \mathbf{y})$ is either empty or an affine copy of $\mathbb{Z} \oplus H_2(Y)$; see [Lip06, Section 2]. (The notation π_2 is used for consistency with the original formulation [OS04b].)

By definition, each homology class $A \in \pi_2(\mathbf{x}, \mathbf{y})$ can be written as a linear combination of connected components of $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ (i.e., a linear combination of *regions* in $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$). We let $\hat{\pi}_2(\mathbf{x}, \mathbf{y}) \subset \pi_2(\mathbf{x}, \mathbf{y})$ denote the set of those homology classes such that all regions containing elements of \mathbf{z} occur with coefficients 0. We call homology classes in $\hat{\pi}_2(\mathbf{x}, \mathbf{y})$ *unpunctured*.

The fact that $\pi_2(\mathbf{x}, \mathbf{y})$ can be empty corresponds to the fact, mentioned earlier, that the complex $\widehat{CF}(Y, k)$ which we will define decomposes as a direct sum of subcomplexes. This also will mean that there is no obvious grading between generators of different subcomplexes. The situation is simpler in the case that Y has the same integral homology as S^3 , in which case $\hat{\pi}_2(\mathbf{x}, \mathbf{y})$ always consists of a single element. The reader unfamiliar with Heegaard Floer homology may wish to focus on this case.

Now, given matchings \mathbf{x} and \mathbf{y} and a homology class $A \in \pi_2(\mathbf{x}, \mathbf{y})$ we can form the moduli space $\mathcal{M}^A(\mathbf{x}, \mathbf{y})$ of holomorphic curves u as in Formula (1) satisfying conditions (M1)–(M5). These moduli spaces depend on the complex structure J we have chosen; for a generic choice of J (subject to the conditions (J1)–(J3)), the moduli spaces $\mathcal{M}^A(\mathbf{x}, \mathbf{y})$ will be smooth manifolds, transversely cut out by the $\bar{\partial}$ -equation. Each moduli space $\mathcal{M}^A(\mathbf{x}, \mathbf{y})$ comes with an action by \mathbb{R} , induced by the translation action on $\Sigma \times [0, 1] \times \mathbb{R}$.

Choosing a generic J , we can define the differential $d: \widehat{CF} \rightarrow \widehat{CF}$ by

$$(2) \quad d\mathbf{x} = \sum_{\mathbf{y}} \sum_{\substack{A \in \hat{\pi}_2(\mathbf{x}, \mathbf{y}) \\ \dim(\mathcal{M}^A(\mathbf{x}, \mathbf{y}))=1}} \#(\mathcal{M}^A(\mathbf{x}, \mathbf{y})/\mathbb{R}) \mathbf{y}.$$

In principle, since $\hat{\pi}_2(\mathbf{x}, \mathbf{y})$ is infinite if $b_1(Y) > 0$, this may be an infinite sum. In Heegaard Floer homology, this is typically resolved by assuming an *admissibility condition* on the Heegaard diagram. In particular, we will always assume that our Heegaard diagrams are weakly admissible, in the sense of [OS04b, Definition 4.10]. This is enough to ensure that the sums in Formula 2 are finite; see [OS04b, Lemma 4.13] or [Lip06, Lemma 5.4].

It is shown in [Lip06] that Formula (2) does, indeed, define a differential. Further, it is shown that this definition is equivalent to the

original one (in terms of disks in $Sym^g(\Sigma)$.) Indeed, if one chooses the complex structures appropriately then the two chain complexes are isomorphic. (This is a relative version of the *tautological correspondence* discussed, say, in [Smi03] or [Ush04].)

Given a map $u: S \rightarrow \Sigma \times [0, 1] \times \mathbb{R}$ satisfying conditions (M1)–(M4), let $\text{ind}(u)$ denote the expected dimension of the moduli space of holomorphic curves near u , i.e., the index of the linearized $\bar{\partial}$ -operator at u . As explained in [Lip06, Section 4.1], a simple doubling argument shows that if u lies in the homology class A then

$$(3) \quad \text{ind}(u) = n - \chi(S) + 2e(A) =: \text{ind}(A, S)$$

Here, $e(A)$ denotes the Euler measure of A . That is: the homology class A is a linear combination of regions of $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$. Each region C is a surface with corners. We set

$$e(C) = \chi(C) - \frac{1}{4} (\text{number of corners of } C)$$

and extend the definition of e linearly to elements of $H_2(\Sigma, \boldsymbol{\alpha} \cup \boldsymbol{\beta})$. (The number e occurs in the Gauss-Bonnet theorem for surfaces with 90° corners – the reason for its name.) Note that formula (3) holds even if u is not an embedding.

Remark 1. *It is straightforward to deduce Formula (3) from Rasmussen’s formula [Ras03, Theorem 9.1] and the Riemann-Hurwitz formula.*

It is shown in [Lip06, Section 4.2] that the Euler characteristic of an embedded holomorphic curve in a homology class $A \in \pi_2(\mathbf{x}, \mathbf{y})$ is determined by A , by the formula

$$(4) \quad \chi(S) = n - n_{\mathbf{x}}(A) - n_{\mathbf{y}}(A) + e(A).$$

Here, $n_{\mathbf{x}}(A)$ is the sum over the $x_i \in \mathbf{x}$ of the average multiplicity of A at x_i . Combining Formulas (3) and (4) implies that, for an *embedded* holomorphic curve u ,

$$\text{ind}(u) = e(A) + n_{\mathbf{x}}(A) + n_{\mathbf{y}}(A).$$

Actually, the exact form that this formula takes will be unimportant for our purposes: all we need is that there is some combinatorial formula for $\text{ind}(u)$ at an embedded curve u .

2.2. The elaborated chain complex \widehat{CF}_{big} . In this section, we will define an elaboration of \widehat{CF} by relaxing condition (M5). Before we do so, a few more words about expected dimensions are in order. Let $\mathcal{M}_p^A(\mathbf{x}, \mathbf{y})$ denote the moduli space of holomorphic curves $u: S \rightarrow \Sigma \times$

$[0, 1] \times \mathbb{R}$ satisfying conditions (M1)–(M4) with exactly p double points (and no other singularities).² As noted above, on the one hand, the expected dimension of \mathcal{M}_p^A near a curve u is still given by Formula (3). On the other hand, general results on holomorphic curves imply that the expected dimension of \mathcal{M}_p^A is $2p$ less than the expected dimension of \mathcal{M}_0^A . (Roughly, near each double point the holomorphic curve is modeled on $\{zw = 0\} \subset \mathbb{C}^2$. This can be deformed to $\{zw = \epsilon\}$ ($\epsilon \in \mathbb{C}$, small) – a two-dimensional family of deformations.) Consequently, if $u \in \mathcal{M}_p^A$ then

$$(5) \quad \text{ind}(u) = e(A) + n_{\mathbf{x}}(A) + n_{\mathbf{y}}(A) - 2p.$$

(Alternately, this can be proved directly by imitating the proof of Formula (4) in [Lip06, Section 4.2].) In particular, if $u: S \rightarrow \Sigma \times [0, 1] \times \mathbb{R}$ is a holomorphic curve in some fixed homology class A , from any one of the three numbers

- $\text{ind}(u)$,
- the number of double points of u and
- $\chi(S)$

it is easy to compute the other two.

Now, we define a chain complex $(\widehat{CF}_{big}, d_{big})$ over $\mathbb{F}_2[[t]]$ as follows. As a module, let $\widehat{CF}_{big} = \widehat{CF} \otimes_{\mathbb{F}_2} \mathbb{F}_2[[t]]$. (That is, \widehat{CF}_{big} is generated over $\mathbb{F}_2[[t]]$ by the set of matchings.) If \mathbf{x} is a matching and p is a non-negative integer, define

$$d_p(\mathbf{x}) = \sum_{\mathbf{y}} \sum_{\substack{A \in \hat{\pi}_2(\mathbf{x}, \mathbf{y}) \\ \dim(\mathcal{M}_p^A) = 1}} \#(\mathcal{M}_p^A(\mathbf{x}, \mathbf{y})) \mathbf{y}$$

$$d_{big}(\mathbf{x}) = \sum_{p=0}^{\infty} t^p d_p(\mathbf{x}) = d_0(\mathbf{x}) + td_1(\mathbf{x}) + t^2 d_2(\mathbf{x}) + \cdots .$$

Again, assuming weak admissibility, the sum defining $d_p(\mathbf{x})$ is finite.

Proposition 1. *The map d_{big} is a differential, i.e., $d_{big}^2 = 0$.*

Proof. This is explained in [Lip06, Section 14.1]. One considers the ends of the one-dimensional ($\text{ind} = 2$) moduli spaces of curves with p double points. Using the compactness result [BEH⁺03, Theorem 10.1], together with an appropriate gluing result (see, e.g., [Lip06, Appendix A]), the ends of the one-dimensional moduli space correspond to two-story holomorphic buildings, the analogue of Morse theory's broken

²For a treatment of singularities of pseudoholomorphic curves, see [McD94] and [MW95].

flow lines. In a two-story building, the double points will be distributed between the two levels, but it is not hard to show that the total number of double points will be p . The proposition then follows from the usual arguments in Morse theory or Floer homology. \square

Note that the fact that $d_{big}^2 = 0$ corresponds to a list of relations among the d_i , the first two being that d_0 is a differential – which we already knew – and that d_1 is a chain map on \widehat{CF} . As we will see in Proposition 2, the chain homotopy type, over $\mathbb{F}_2[[t]]$, of \widehat{CF}_{big} is an invariant of the 3-manifold Y together with the integer k ; this essentially follows from results in [Lip06, Section 14.1].

Remark 2. *Suppose that the manifold Y is a rational homology sphere, or more generally that we are working with a torsion $spin^c$ -structure. In this case, we may replace $\mathbb{F}_2[[t]]$ by $\mathbb{F}_2[t]$. Also, there is a well-defined relative \mathbb{Z} -grading on $(\widehat{CF}(Y), d)$ (in each $spin^c$ -structure). Setting $t = 1$ in \widehat{CF}_{big} , we obtain a differential $d': \widehat{CF}(Y) \rightarrow \widehat{CF}(Y)$. In other words, $d': \widehat{CF}(Y) \rightarrow \widehat{CF}(Y)$ is defined by*

$$d'(\mathbf{x}) = \sum_{\mathbf{y}} \sum_{\substack{A \in \tilde{\pi}_2(\mathbf{x}, \mathbf{y}) \\ \dim(\mathcal{M}_p^A) = 1}} \#(\mathcal{M}_p^A(\mathbf{x}, \mathbf{y})/\mathbb{R}) \mathbf{y}.$$

The relative \mathbb{Z} -grading by the index (the “Maslov grading”) becomes a relative \mathbb{Z} -filtration of (\widehat{CF}, d') . If we define \widehat{CF}_{big} over $\mathbb{F}_2[t]$ rather than $\mathbb{F}_2[[t]]$ then the filtered chain complex (\widehat{CF}, d') contains the same information as $(\widehat{CF}_{big}, d_{big})$.

Remark 3. *The reader may find it interesting to compare the definition of d_{big} with the “Taubes series” as defined by Ionel-Parker in [IP97, Section 2]. This was the original motivation for the definition of d_{big} .*

Remark 4. *As noted above, in the original formulation of Heegaard Floer homology, instead of counting holomorphic curves in $\Sigma \times [0, 1] \times \mathbb{R}$ one counts holomorphic disks in $Sym^g(\Sigma)$. Working with appropriate (but still sufficiently generic) almost complex structures on $\Sigma \times [0, 1] \times \mathbb{R}$ and $Sym^g(\Sigma)$, there is a “tautological” bijective correspondence between holomorphic curves in $\Sigma \times [0, 1] \times \mathbb{R}$ and holomorphic disks in $Sym^g(\Sigma)$; see [Lip06, Section 13]. For these complex structures, the diagonal $\Delta \subset Sym^g(\Sigma)$ is an almost complex submanifold. With respect to the tautological correspondence, a double point of a curve in $\Sigma \times [0, 1] \times \mathbb{R}$ corresponds to a simple tangency to the top-dimensional stratum of the diagonal in $Sym^g(\Sigma)$. This interpretation suggests forming even larger chain complexes by taking into account intersections with or tangencies*

to other strata of the diagonal, and/or higher order tangencies to the top-dimensional stratum of the diagonal. It might be interesting to study these elaborations; the techniques of this paper do not seem to apply to them.

2.3. Links. To a $2k$ -pointed Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}, \mathbf{w})$ for an ℓ -component nullhomologous link L in a manifold Y one can associate various refinements of $\widehat{CF}(Y)$. We will focus on the one denoted \widehat{CFK} in [MOST06]; in the case that $k = \ell$, this is also the invariant denoted \widehat{CFK} in [OS04a].³

For the moment, assume that Y is a homology sphere. The complex $\widehat{CFK}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}, \mathbf{w})$ is obtained from $\widehat{CFK}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$ as follows. Given a domain $A \in \pi_2(\mathbf{x}, \mathbf{y})$ let $n_{w_i}(A)$ denote the coefficient in A of the region containing w_i . Let $n_{\mathbf{w}}(A) = \sum_{i=1}^k n_{w_i}(A)$. Since we assume the link L is null-homologous, it is not hard to show that given two domains $A, A' \in \hat{\pi}_2(\mathbf{x}, \mathbf{y})$, $n_{\mathbf{w}}(A) = n_{\mathbf{w}}(A')$. Further, since the fibers of π_{Σ} are J -holomorphic, if A admits a holomorphic representative then $n_{\mathbf{w}}(A) > 0$. Consequently, \mathbf{w} induces a relative \mathbb{Z} -filtration \mathcal{F} on $\widehat{CF}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$, given by $\mathcal{F}(\mathbf{x}, \mathbf{y}) = n_{\mathbf{w}}(A)$, where A is any element of $\hat{\pi}_2(\mathbf{x}, \mathbf{y})$. (More generally, if Y is not a homology sphere, $\pi_2(\mathbf{x}, \mathbf{y})$ may be empty, so we instead obtain a filtration \mathcal{F} on each (*spin^c*-) summand of $\widehat{CF}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$.) Then, we let $\widehat{CFK}(Y, L, k) = \widehat{CFK}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}, \mathbf{w})$ be the associated graded complex to the filtered complex $(\widehat{CF}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}), \mathcal{F})$. That is, as a module \widehat{CFK} is the same as \widehat{CF} ; the differential on \widehat{CFK} counts only holomorphic curves which do not cover any of the w_i . We set $\widehat{CFK}(Y, L) = \widehat{CFK}(Y, L) = \widehat{CFK}(Y, L, \ell)$.

³Recall that for links there are two different constructions of the invariant \widehat{CFK} . The first construction, used in [OS04a], replaces an ℓ -component link in Y with a knot K in $Y \# (\ell - 1)S^1 \times S^2$, and defines $\widehat{CFK}(Y, L)$ to be $\widehat{CFK}(Y \# (\ell - 1)S^1 \times S^2, K)$. The second construction, used in [OS05], works with multi-pointed Heegaard diagrams for L in Y . It is a nontrivial theorem ([OS05, Theorem 1.1]) that for appropriate choices of diagram and almost complex structure, the resulting complexes are isomorphic. When we talk about the knot Floer complex for a link, we will generally mean the construction from [OS05], using multi-pointed Heegaard diagrams. The only exceptions are towards the end of Section 2.4, where we use the fact that the two constructions are equivalent in order to prove invariance, and in Remark 5, where we discuss generalizations to other flavors on knot Floer homology.

There is, then, an obvious extension of the elaboration d_{big} to \widehat{CFK} : the map \mathcal{F} also gives a filtration on CF_{big} , and we let \widehat{CFK}_{big} be the associated graded complex.

Our main reason for introducing the link invariants is because of the special case $Y = S^3$. In this case, as described in [MOS06], one can take the Heegaard diagram to be a *toroidal grid diagram* – i.e., a diagram for L in which Σ is a torus and each α_i intersects each β_j in a single point. (For example, a toroidal grid diagram of the figure 8 knot is shown in Figure 1.) In this case, as noted in Remark 2, without loss of information one can specialize $t = 1$, giving a new differential d' on the complex \widehat{CFK} . The \mathbb{Z} -grading given by the index (the “Maslov” or “homological” grading) becomes a \mathbb{Z} -filtration instead. We will see in Proposition 9 that in a grid diagram, the differential d' is a particularly natural extension of d .

Remark 5. *There are several fancier versions of \widehat{CFK} or \widehat{CFK}_{big} . One of these is the (relatively) \mathbb{Z} -filtered chain homotopy type of \widehat{CF} , where the filtration is the map \mathcal{F} defined above. The filtration \mathcal{F} obviously induces a filtration on \widehat{CF}_{big} as well. In the case $k = 1$, the arguments in Section 2.4 will apply to prove that the filtered homotopy type of \widehat{CF}_{big} is an invariant of the knot K in Y . (As discussed in [OS04a, Section 2.1], this gives a link invariant, as well, which is \mathbb{Z} -filtered rather than $\mathbb{Z}^{|L|}$ -filtered.) Proving either invariance of the link invariant CFL^- or of the filtered knot invariant for $k > 1$ would require additional work. It seems unlikely, however, that either invariance would fail.*

2.4. Invariance. The complexes we have introduced would not be of much interest if they were not topological invariants. Fortunately, they are, as the next three propositions show:

Proposition 2. *Let $(\Sigma, \alpha, \beta, z)$ and $(\Sigma', \alpha', \beta', z')$ be 1-pointed Heegaard diagrams for Y . Let J be an almost complex structure on $\Sigma \times [0, 1] \times \mathbb{R}$ and J' an almost complex structure on $\Sigma' \times [0, 1] \times \mathbb{R}$ satisfying conditions (J1)–(J3). Then the chain complexes $\widehat{CF}_{big}(\Sigma, \alpha, \beta, z)$ and $\widehat{CF}_{big}(\Sigma', \alpha', \beta', z')$, computed with respect to J and J' respectively, are homotopy equivalent over $\mathbb{F}_2[[t]]$. If Y is a rational homology sphere then the complexes $(\widehat{CF}(\Sigma, \alpha, \beta, z), d')$ and $(\widehat{CF}(\Sigma', \alpha', \beta', z'), d')$ defined in Remark 2 are homotopy equivalent as relative \mathbb{Z} -filtered complexes over \mathbb{F} .*

Proof. As explained in [Lip06, Section 14.1], this proposition follows from the same arguments used in [OS04b] to prove invariance of $\widehat{HF}(Y)$.




FIGURE 2. Left: a 2-pointed, genus 1 Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{z})$ for S^3 . The two shaded ovals are the feet of a 1-handle, itself shown in translucent gray. Right: the corresponding 1-pointed, genus 2 Heegaard diagram $(\Sigma_{\#}, \alpha_{\#}, \beta_{\#}, z)$ for $S^3 \# (S^1 \times S^2)$. There are now two 2-handles, whose feet are two pairs of shaded ovals. Note that the only differences between the two diagrams are in the regions containing z 's, so the chain complexes \widehat{CF}_{big} are the same for the two diagrams.

In particular, one associates a chain homotopy equivalence to each of the three Heegaard moves (isotopies, handleslides and stabilizations), as well as to change of almost complex structure.

It is routine to adapt these maps to \widehat{CF}_{big} . The hardest part to check is that the handleslide map is well-defined and gives a homotopy equivalence. The proof of handleslide invariance ([OS04b, Section 9]; see [Lip06, Section 11] for the proof in the cylindrical setting) leverages the special case of a particular handleslide for $(S^1 \times S^2) \# (S^1 \times S^2)$ to prove invariance in general. The form of the argument remains the same, but we must check that there are no unwanted holomorphic triangles with double points. It turns out ([Lip06, p. 1068]) that there are thirteen domains to check. Seven can not have representatives with double points because no region has multiplicity greater than one. The remaining six can be ruled out by considering the Euler characteristic of a representative with double points. \square

We are now justified in writing $\widehat{CF}_{big}(Y)$ to denote the chain homotopy type of $\widehat{CF}_{big}(\Sigma, \alpha, \beta, z)$ for any 1-pointed Heegaard diagram $(\Sigma, \alpha, \beta, z)$ for Y .

We next argue that the invariance of \widehat{CF}_{big} in the 1-pointed case implies the appropriate kind of invariance in the multi-pointed case.

Proposition 3. *Let $(\Sigma, \alpha, \beta, \mathbf{z})$ be a k -pointed Heegaard diagram for Y . Then $\widehat{CF}_{big}(\Sigma, \alpha, \beta, \mathbf{z})$ is homotopy equivalent (over $\mathbb{F}[[t]]$) to*

$$\widehat{CF}_{big}(Y) \otimes_{\mathbb{F}_2[[t]]} (H_*(S^1; \mathbb{F}_2[[t]])^{\otimes k}),$$

where $H_*(S^1; \mathbb{F}_2[[t]])$ denotes the ordinary homology of S^1 with coefficients in $\mathbb{F}_2[[t]]$. If Y is a rational homology sphere then

$$(\widehat{CF}(\Sigma, \alpha, \beta, \mathbf{z}), d') \simeq (\widehat{CF}(Y), d') \otimes_{\mathbb{F}} (H_*(S^1; \mathbb{F}_2)^{\otimes k}),$$

as relatively \mathbb{Z} -filtered complexes. In particular, the homotopy types of $\widehat{CF}_{big}(\Sigma, \alpha, \beta, \mathbf{z})$ and $(\widehat{CF}(\Sigma, \alpha, \beta, \mathbf{z}), d')$ depend only on Y and k .

Proof. From the k -pointed Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{z})$ one can obtain a 1-pointed Heegaard diagram $(\Sigma_{\#}, \alpha_{\#}, \beta_{\#}, z)$ for $Y \# (k-1)(S^1 \times S^2)$ by gluing a handle to Σ between z_i and z_1 for $i = 2, \dots, k$; see Figure 2 for an example. (Gluing a handle to Σ , not intersecting any of the α - or β -curves corresponds to surgering $\mathbb{D}^3 \times S^0$ from Y and replacing it with $S^2 \times \mathbb{D}^1$.) It is immediate from the definition that $\widehat{CF}(\Sigma, \alpha, \beta, \mathbf{z}) \cong \widehat{CF}(\Sigma_{\#}, \alpha_{\#}, \beta_{\#}, z)$ and $\widehat{CF}_{big}(\Sigma, \alpha, \beta, \mathbf{z}) \cong \widehat{CF}_{big}(\Sigma_{\#}, \alpha_{\#}, \beta_{\#}, z)$. Consequently, it follows from invariance of \widehat{CF}_{big} for 1-pointed diagrams that if $(\Sigma', \alpha', \beta', \mathbf{z}')$ is another Heegaard diagram for Y with $|\mathbf{z}| = |\mathbf{z}'|$ then $\widehat{CF}_{big}(\Sigma, \alpha, \beta, \mathbf{z})$ and $\widehat{CF}_{big}(\Sigma', \alpha', \beta', \mathbf{z}')$ are chain homotopy equivalent over $\mathbb{F}_2[[t]]$.

Now, recall that given 3-manifolds Y_1 and Y_2 ,

$$\widehat{HF}(Y_1 \# Y_2) \cong \widehat{HF}(Y_1) \otimes_{\mathbb{F}_2} \widehat{HF}(Y_2).$$

This is proved by choosing a Heegaard diagram $(\Sigma, \alpha, \beta, z)$ for $Y_1 \# Y_2$ which itself decomposes as a connect sum $\Sigma = \Sigma_1 \# \Sigma_2$, $\alpha = \alpha_1 \amalg \alpha_2$, $\beta = \beta_1 \amalg \beta_2$, and then placing the basepoint z in the region of Σ where the connect sum occurs. For such a diagram,

$$\begin{aligned} \widehat{CF}(\Sigma, \alpha, \beta, z) &\cong \widehat{CF}(\Sigma_1, \alpha_1, \beta_1, z) \otimes_{\mathbb{F}_2} \widehat{CF}(\Sigma_2, \alpha_2, \beta_2, z), \text{ and} \\ \widehat{CF}_{big}(\Sigma, \alpha, \beta, z) &\cong \widehat{CF}_{big}(\Sigma_1, \alpha_1, \beta_1, z) \otimes_{\mathbb{F}_2[[t]]} \widehat{CF}_{big}(\Sigma_2, \alpha_2, \beta_2, z). \end{aligned}$$

Putting this together with the fact that $\widehat{CF}_{big}(\Sigma, \alpha, \beta, \mathbf{z})$ is isomorphic to $\widehat{CF}_{big}(\Sigma_{\#}, \alpha_{\#}, \beta_{\#}, z)$ one obtains that $\widehat{CF}_{big}(\Sigma, \alpha, \beta, \mathbf{z})$ is chain homotopy equivalent to

$$\widehat{CF}_{big}(Y) \otimes_{\mathbb{F}_2[[t]]} \left(\bigotimes_{i=1}^{k-1} \widehat{CF}_{big}(S^1 \times S^2) \right).$$

It is straightforward to check that $\widehat{CF}_{big}(S^1 \times S^2)$ is (chain homotopy equivalent to) $H_*(S^1; \mathbb{F}_2[[t]]) \cong \mathbb{F}_2[[t]] \oplus \mathbb{F}_2[[t]]$. This proves the first part of the result. The second part follows by the same argument. \square

Finally, we turn to link Floer homology.

Proposition 4. *Let $(\Sigma, \alpha, \beta, \mathbf{z}, \mathbf{w})$ be a $2k$ -pointed Heegaard diagram for an l -component link L in Y . Then the chain homotopy type, over $\mathbb{F}_2[[t]]$, of $\widehat{CFK}_{big}(Y, K, k) = \widehat{CFK}_{big}(\Sigma, \alpha, \beta, \mathbf{z}, \mathbf{w})$ depends only on Y , K and k . Further,*

$$\widehat{CFK}_{big}(Y, K, k) \simeq \widehat{CFK}_{big}(Y, K, l) \otimes_{\mathbb{F}[[t]]} (H_*(S^1; \mathbb{F}[[t]])^{\otimes k-l}).$$

Analogous statements hold for the filtered theory $(\widehat{CFK}(Y, K, k), d')$.

Proof. In the case $k = 1$, i.e., the case treated in [OS04a], invariance of $\widehat{CFK} = \widehat{CFK}$ follows from the same argument used to prove invariance of \widehat{CF} . That is, one can move between different 2-pointed Heegaard diagrams of the same knot via isotopies, handleslides, and stabilizations / destabilizations – and one can choose the isotopies and handleslides to be supported in the complement of $\{z, w\}$; see [OS04a, Proposition 3.5]. It is straightforward to check that the maps induced on \widehat{CF} by these moves respect the filtration \mathcal{F} , and in fact induce homotopy equivalences on the filtered complexes. The same holds for the maps induced on \widehat{CF}_{big} .

To pass from $k = 1$ to general k , we adapt the handle attaching construction we used before; see also [OS05, Section 10]. Let $(\Sigma, \alpha, \beta, \mathbf{z}, \mathbf{w})$ be a $2k$ -pointed Heegaard diagram for a link L in Y . Attaching a handle to Σ between z_i and w_{i+1} ($i = 1, \dots, k-1$) gives a Heegaard diagram $(\Sigma_{\#}, \alpha_{\#}, \beta_{\#}, z, w)$, where $z = z_k$ and $w = w_1$, for a link L in $Y \# (k-1)(S^1 \times S^2)$; see Figure 3 for an example. It is not obvious that $\widehat{CFK}(\Sigma, \alpha, \beta, \mathbf{z}, \mathbf{w}) \cong \widehat{CFK}(\Sigma_{\#}, \alpha_{\#}, \beta_{\#}, z, w)$: in the latter chain complex, holomorphic curves can, in principle, cover the regions which used to contain z_i ($i = 1, \dots, k-1$) and w_j ($j = 2, \dots, k$). However, it is proved in [OS05, Section 10, proof of Theorem 1] that for suitable choices of (equivalent) Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{z}, \mathbf{w})$ and suitable complex structures on $\Sigma_{\#} \times [0, 1] \times \mathbb{R}$, the two chain complexes are, in fact, isomorphic. In fact, for these choices, no holomorphic curves in $(\Sigma_{\#} \setminus \{z, w\}) \times [0, 1] \times \mathbb{R}$ cover any of the regions which used to contain z_i ($i = 1, \dots, k-1$) and w_j ($j = 2, \dots, k$). Consequently, the same proof applies to our elaborated chain complexes \widetilde{CFK}_{big} .

If $k = \ell$ then the knot Floer homology of $(\Sigma_{\#}, \alpha_{\#}, \beta_{\#}, z, w)$ is, by definition, the knot Floer homology of (Y, L) , as defined in [OS04a]. In




FIGURE 3. Left: a 4-pointed, genus 1 Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{z}, \mathbf{w})$ for the Hopf link in S^3 . The shaded ovals are the feet of a 1-handle, itself shown in translucent gray. Right: the corresponding 2-pointed, genus 2 Heegaard diagram $(\Sigma_{\#}, \alpha_{\#}, \beta_{\#}, z, w)$ for the corresponding knot in $S^3 \# (S^1 \times S^2)$. Note that, unlike in Figure 2, not all of the difference occurs in regions containing z 's: in principle, curves could cover the region labeled A in the diagram on the right.

general, as was the case for $\widehat{CF}(Y, k)$, up to chain homotopy equivalence one has $\widehat{CFK}(\Sigma_{\#}, \alpha_{\#}, \beta_{\#}, z, w) \cong \widehat{CFK}(Y, K) \otimes_{\mathbb{F}_2} (H_*(S^1))^{\otimes k-\ell}$, and $\widehat{CFK}_{big}(\Sigma_{\#}, \alpha_{\#}, \beta_{\#}, z, w) \cong \widehat{CFK}_{big}(Y, K) \otimes_{\mathbb{F}_2[[t]]} (H_*(S^1; \mathbb{F}_2[[t]]))^{\otimes k-\ell}$. The proposition follows. \square

3. ...IN NICE DIAGRAMS.

Recall from [SW, Section 3] that a pointed Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{z})$ is called *nice* if every component of $\Sigma \setminus (\alpha \cup \beta)$ not containing an element of \mathbf{z} is either a bigon or a square. In other words, in a nice diagram, all unpunctured regions have non-negative Euler measure, as defined at the end of Section 2.1.

The following is [SW, Theorem 3.2], stated in our language:

Proposition 5. *Fix a nice pointed Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{z})$ and a generic almost complex structure J on $\Sigma \times [0, 1] \times \mathbb{R}$ satisfying (J1)–(J3). Let $u: S \rightarrow \Sigma \times [0, 1] \times \mathbb{R}$ be a J -holomorphic curve satisfying Properties (M1)–(M5), in an unpunctured homology class. Suppose that $\text{ind}(u) = 1$. Then:*

- (1) *There is a unique component S_0 of S on which $\pi_\Sigma \circ u$ is non-constant.*
- (2) *The component S_0 is either a bigon or a rectangle.*
- (3) *The map $(\pi_\Sigma \circ u)|_{S_0}: S_0 \rightarrow \Sigma$ is an embedding.*
- (4) *If S_1 is any component of S other than S_0 then $\pi_\Sigma \circ u(S_1)$ is disjoint from $\pi_\Sigma \circ u(S_0)$.*

Consequently, the domain $D(u) \in \pi_2(\mathbf{x}, \mathbf{y})$ of u is an embedded bigon or rectangle in Σ , with no x_i or y_j lying in the interior of $D(u)$. Conversely, any such domain has a unique J -holomorphic representative satisfying Properties (M1)–(M5).

Remark 6. *Point (1) of Proposition 5 is true for any generic, index 1 holomorphic curve in $\Sigma \times [0, 1] \times \mathbb{R}$, whether or not the diagram is nice. Indeed, if $\pi_\Sigma \circ u$ were nonconstant on two components of S then the \mathbb{R} -action on each component would give a 2-dimensional family of holomorphic curves near u – a contradiction.*

The main result of this section generalizes Proposition 5 to curves with double points. Before stating the generalization, we give a technical lemma which will allow us to work with split almost complex structures:

Lemma 6. *Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$ be a nice multi-pointed Heegaard diagram and J an almost complex structure on $\Sigma \times [0, 1] \times \mathbb{R}$ satisfying (J1)–(J3). Fix an unpunctured homology class A and a surface S such that $\text{ind}(A, S) = 0$. Let $u: S \rightarrow \Sigma \times [0, 1] \times \mathbb{R}$ be a J -holomorphic curve in the homology class A . Then $\pi_\Sigma \circ u$ is constant.*

Note. We do *not* assume that J is generic.

Proof. Since the fibers of π_Σ are J -holomorphic, all of the coefficients of components of $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ in the domain $D(u)$ of any holomorphic curve u are non-negative. Consequently, since the coefficients of the regions containing elements of \mathbf{z} are zero, $e(D(u)) \geq 0$. So, by Formula (3), $\chi(S) \geq n$, with equality if and only if $e(D(u)) = 0$. However, by Properties (M3) and (M4), S can have at most n connected components. Hence S consists of n bigons, and also $e(D(u)) = 0$.

Let S' denote the union of those components of S on which $\pi_\Sigma \circ u$ is nonconstant. Since the fibers of π_Σ are J -holomorphic, $\pi_\Sigma \circ u|_{S'}$ is a branched map. By the Riemann-Hurwitz formula, $e(S) = e(D(u)) - r$, where r is the total ramification degree of $\pi_\Sigma \circ u$. In particular, $e(S') \leq e(D(u)) = 0$. Since S' consists entirely of bigons, S' must in fact be empty. \square

Corollary 7. *For any matchings \mathbf{x} and \mathbf{y} , homology class $A \in \hat{\pi}_2(\mathbf{x}, \mathbf{y})$, and integer $p \geq 0$, the mod 2 count of $\mathcal{M}_p^A(\mathbf{x}, \mathbf{y})/\mathbb{R}$ is independent of the choice of generic almost complex structure J .*

Proof. This follows by a standard continuation argument, which we recall. Fix generic almost complex structures J_0 and J_1 , and a generic path of almost complex structures J_s between them. Let $\mathcal{M}_p^A(\mathbf{x}, \mathbf{y}; J_s)$ denote the moduli space computed with respect to J_s . It follows from the compactness theorem (in the cylindrical setting, [BEH⁺03, Theorem 10.1]), together with standard transversality and gluing arguments, that $\overline{\bigcup_{s \in [0,1]} \mathcal{M}_p^A(\mathbf{x}, \mathbf{y}; J_s)/\mathbb{R}}$ has a natural compactification $\overline{\bigcup_{s \in [0,1]} \mathcal{M}_p^A(\mathbf{x}, \mathbf{y}; J_s)/\mathbb{R}}$ is a compact one-manifold with boundary. The spaces $\mathcal{M}^p(\mathbf{x}, \mathbf{y}; J_0)/\mathbb{R}$ and $\mathcal{M}_p^A(\mathbf{x}, \mathbf{y}; J_1)/\mathbb{R}$ form part of the boundary of $\overline{\bigcup_{s \in [0,1]} \mathcal{M}_p^A(\mathbf{x}, \mathbf{y}; J_s)/\mathbb{R}}$. The rest of the boundary consists of two-story holomorphic buildings, i.e., pairs of curves

$$(u_1, u_2) \in \mathcal{M}_{p_1}^{A_1}(\mathbf{x}, \mathbf{w}; J_s)/\mathbb{R} \times \mathcal{M}_{p_2}^{A_2}(\mathbf{w}, \mathbf{y}; J_s)/\mathbb{R}$$

where $A = A_1 + A_2$ and $p = p_1 + p_2$. (This is not quite automatic: we must rule-out bubbling and other boundary degenerations. Doing so reduces to elementary complex analysis; see [Lip06, Proposition 7.1].)

The index is additive, so $\text{ind}(u_1) + \text{ind}(u_2) = 1$. Hence either $\text{ind}(u_1) \leq 0$ or $\text{ind}(u_2) \leq 0$. For generic paths of almost complex structures, all moduli spaces of negative index (expected dimension ≤ -2) will be empty. Thus, Lemma 6 implies that either $\mathcal{M}_{p_1}^{A_1}(\mathbf{x}, \mathbf{w}; J_s)/\mathbb{R}$ or $\mathcal{M}_{p_2}^{A_2}(\mathbf{w}, \mathbf{y}; J_s)/\mathbb{R}$ is empty. Thus, the only ends of $\overline{\bigcup_{s \in [0,1]} \mathcal{M}_p^A(\mathbf{x}, \mathbf{y}; J_s)/\mathbb{R}}$ are $\mathcal{M}^p(\mathbf{x}, \mathbf{y}; J_0)/\mathbb{R}$ and $\mathcal{M}_p^A(\mathbf{x}, \mathbf{y}; J_1)/\mathbb{R}$. The result follows. \square

Proposition 8. *Fix a nice pointed Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$ and a generic almost complex structure J on $\Sigma \times [0, 1] \times \mathbb{R}$ satisfying (J1)–(J3). Let $u: S \rightarrow \Sigma \times [0, 1] \times \mathbb{R}$ be a J -holomorphic curve satisfying Properties (M1)–(M4) with p double points (and no other singularities), in an unpunctured homology class. Suppose that $\text{ind}(u) = 1$. Then:*

- (1) *There is a unique component S_0 of S on which $\pi_\Sigma \circ u$ is non-constant.*
- (2) *The component S_0 is either a bigon or a rectangle.*
- (3) *The map $(\pi_\Sigma \circ u)|_{S_0}: S_0 \rightarrow \Sigma$ is a local diffeomorphism.*

Consequently, the domain $D(u) \in \pi_2(\mathbf{x}, \mathbf{y})$ of u is the image under an orientation-preserving local diffeomorphism of a (not necessarily embedded) bigon or rectangle in Σ , and by Formula 5 satisfies $e(D(u)) + n_{\mathbf{x}}(D(u)) + n_{\mathbf{y}}(D(u)) = 2p + 1$. Conversely, any such domain has an odd number of J -holomorphic representatives satisfying Properties (M1)–(M4) with exactly p double points.

Proof. The proof we will give is somewhat different from the original proof of Proposition 5 in [SW]. It is similar to, but easier than, the proof of [LMW06, Proposition 2.3].

As noted in Remark 6, Point (1) of the proposition is true for generic index 1 curves in any Heegaard diagram.

The proof of the Point (2) of the proposition is similar to the proof of Lemma 6. Since the fibers of π_Σ are J -holomorphic, all of the coefficients of components of $\Sigma \setminus (\alpha \cup \beta)$ in the domain $D(u)$ of any holomorphic curve u are non-negative. Consequently, since the coefficients of the regions containing elements of \mathbf{z} are zero, $e(D(u)) \geq 0$. Thus, by Formula (3), $\chi(S) \geq n - 1$. However, by Properties (M3) and (M4), S can have at most n connected components. Consequently, S consists of either n or $n - 1$ topological disks, perhaps together with some annuli. By Property (M1), S has $2n$ punctures on its boundary. A holomorphic annulus must have at least 4 boundary punctures, and hence is prohibited. Consequently, either S consists of n bigons or S consists of $n - 2$ bigons and one rectangle. This implies Point (2) of the proposition. Note that we have also shown that $e(D(u))$ is either 0 or $1/2$.

We next show that $\pi_\Sigma \circ u|_{S_0}$ is a local diffeomorphism. Recall that, since the fibers of π_Σ are J -holomorphic, $\pi_\Sigma \circ u|_{S_0}$ is a branched map. By the Riemann-Hurwitz formula, $e(S_0) = e(D(u)) - r$, where r is the total ramification degree of $\pi_\Sigma \circ u$. (Here, a branch point on the boundary contributes $1/2$ to r .) Since $e(D(u)) \leq 1/2$, the only case in which $r \neq 0$ is if S_0 is a rectangle and $\pi_\Sigma \circ u|_{S_0}$ has a single boundary branch point. It is not hard to see that, in this case, a split complex structure on $\Sigma \times [0, 1] \times \mathbb{R}$ achieves transversality. But then varying the branch point of $\pi_\Sigma \circ u$ gives a family of holomorphic curves near u – a contradiction. This proves Point (3) of the proposition.

Finally, for the existence statement (converse), one must show that any domain which is the image of an orientation-preserving local diffeomorphism from a bigon or rectangle has a unique holomorphic representative with the specified number of double points. Let us assume, for definiteness, that the domain is the image of a rectangle. We work first with a split almost complex structure $j_\Sigma \times j_\mathbb{D}$ on $\Sigma \times [0, 1] \times \mathbb{R}$. Let S be the disjoint union of one rectangle and $n - 2$ bigons. There is an obvious map $u_\Sigma: S \rightarrow \Sigma$ sending the rectangle by the orientation-preserving local diffeomorphism to the domain and the bigons by constant maps. Requiring that u_Σ be holomorphic induces a complex structure on the rectangle. It then follows by elementary complex analysis that there is a unique holomorphic map $u_\mathbb{D}: S \rightarrow [0, 1] \times \mathbb{R}$ so that the corresponding map $u = (u_\Sigma, u_\mathbb{D}): S \rightarrow \Sigma \times [0, 1] \times \mathbb{R}$ satisfies Properties

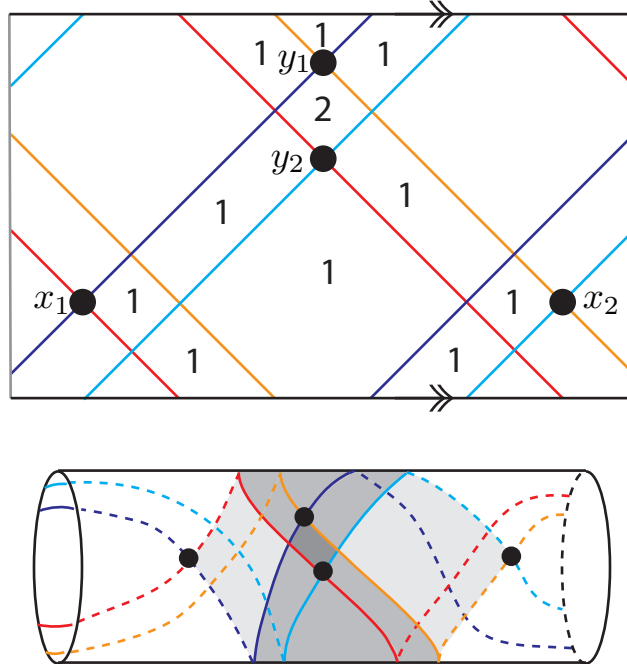


FIGURE 4. Top: a cylinder (top and bottom edges are identified). The curves on the cylinder represent a (hypothetical) part of a Heegaard diagram. The numbers in the regions correspond to a domain in that Heegaard diagram, connecting $\mathbf{x} = \{x_1, x_2, \dots\}$ to $\mathbf{y} = \{y_1, y_2, \dots\}$. (All regions without numbers have coefficient 0.) This domain is represented by a non-embedded, rigid holomorphic rectangle with a single double point. Bottom: the same cylinder, curves and domain, drawn differently to help with visualization of the rectangle.

(M1)–(M4); see, for instance, [Ras03, Section 9.5]. It follows, say, from Oh’s “boundary injectivity” criterion ([Oh96]) as discussed in [Lip06, Proposition 3.9] or in the proof of [OS04b, Proposition 3.9] that the linearized $\bar{\partial}$ -operator is surjective at u . The map u is rigid, so by Formula (5), u has p double points. This proves the result for the split complex structure. For general J , the result then follows from Corollary 7. \square

Remark 7. *Note that if $\pi_\Sigma \circ u|_{S_0}$ is an embedding then exactly p of the $x_i \in \mathbf{x}$ lie in the interior of the domain of u . It is not hard to show that if S_0 is a bigon then, in fact, $\pi_\Sigma \circ u|_{S_0}$ is an embedding. On the other hand, it is not true that $\pi_\Sigma \circ u|_{S_0}$ is necessarily an embedding if S_0 is*

a rectangle; see Figure 3. We will show below that, in a grid diagram, $\pi_\Sigma \circ u|_{S_0}$ is always an embedding.

Remark 8. In [Sar06] and [LMW06], Proposition 5 is extended to the “holomorphic triangles” needed to define cobordism maps. These methods readily extend to prove the analog of Proposition 8, i.e., that rigid holomorphic triangle with double points in a nice Heegaard triple diagram correspond to domains which are the images of n triangles each mapped orientation-preserving diffeomorphically to Σ . However, without substantial effort, the methods from [LMW06] do not compute the number of double points of the corresponding holomorphic map. For this, applying Sarkar’s combinatorial formula for the index of triangles from [Sar06] is much easier.

As discussed in the introduction, in a grid diagram the differential d_{big} or d' , as defined in Section 2.3, is somewhat simpler. It follows from Proposition 5 that, in a grid diagram, the coefficient of \mathbf{y} in $d(\mathbf{x})$ counts rectangles between \mathbf{x} and \mathbf{y} whose interiors do not contain points of \mathbf{x} , \mathbf{y} , \mathbf{z} or \mathbf{w} . The following proposition implies that the differential d' is obtained simply by dropping the condition that rectangles may not contain points of \mathbf{x} or \mathbf{y} .

Proposition 9. *Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}, \mathbf{w})$ be a grid diagram for a link L in S^3 . Suppose that $u: S \rightarrow (\Sigma \setminus \mathbf{z}) \times [0, 1] \times \mathbb{R}$ is a rigid holomorphic curve connecting \mathbf{x} and \mathbf{y} , with p double points, not covering any z_i . Then the domain of u is an embedded rectangle in Σ whose interior contains exactly p elements of \mathbf{x} . Conversely, any such domain has a unique rigid holomorphic representative with p double points.*

Proof. By Proposition 8, it suffices to prove: if u is a holomorphic curve with p double points and S_0 is the (unique) component of u on which $\pi_\Sigma \circ u$ is nonconstant then $\pi_\Sigma \circ u|_{S_0}$ is an embedding. (Then, all double points of u must correspond to intersections of $u|_{S_0}$ and other components of u , which implies p of the x_i must lie in the domain of u .) But since $\pi_\Sigma \circ u|_{S_0}$ is a local diffeomorphism and there is one z_i in every row and column of the grid diagram, $\pi_\Sigma \circ u|_{S_0}$ must be an embedding. \square

REFERENCES

- [BEH⁺03] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder. Compactness results in symplectic field theory. *Geom. Topol.*, 7:799–888 (electronic), 2003.
- [Bel07] Anna Beliakova. Simplification of the combinatorial link Floer homology. 2007, arXiv:0705.0669.

- [IP97] Eleny-Nicoleta Ionel and Thomas H. Parker. The Gromov invariants of Ruan-Tian and Taubes. *Math. Res. Lett.*, 4(4):521–532, 1997.
- [Lip06] Robert Lipshitz. A cylindrical reformulation of Heegaard Floer homology. *Geom. Topol.*, 10:955–1097 (electronic), 2006.
- [LMW06] Robert Lipshitz, Ciprian Manolescu, and Jiajun Wang. Combinatorial cobordism maps in hat Heegaard Floer theory. 2006, math/0611927.
- [McD94] Dusa McDuff. Singularities and positivity of intersections of J -holomorphic curves. In *Holomorphic curves in symplectic geometry*, volume 117 of *Progr. Math.*, pages 191–215. Birkhäuser, Basel, 1994. With an appendix by Gang Liu.
- [MOS06] Ciprian Manolescu, Peter Ozsváth, and Sucharit Sarkar. A combinatorial description of knot Floer homology. 2006, math/0607691.
- [MOST06] Ciprian Manolescu, Peter Ozsváth, Zoltán Szabó, and Dylan Thurston. On combinatorial link Floer homology. 2006, math/0610559.
- [MW95] Mario J. Micallef and Brian White. The structure of branch points in minimal surfaces and in pseudoholomorphic curves. *Ann. of Math. (2)*, 141(1):35–85, 1995.
- [Oh96] Yong-Geun Oh. Fredholm theory of holomorphic discs under the perturbation of boundary conditions. *Math. Z.*, 222(3):505–520, 1996.
- [OS04a] Peter Ozsváth and Zoltán Szabó. Holomorphic disks and knot invariants. *Adv. Math.*, 186(1):58–116, 2004.
- [OS04b] Peter Ozsváth and Zoltán Szabó. Holomorphic disks and topological invariants for closed three-manifolds. *Ann. of Math. (2)*, 159(3):1027–1158, 2004.
- [OS05] Peter Ozsváth and Zoltán Szabó. Holomorphic disks and link invariants. 2005, math/0512286.
- [Ras03] Jacob Rasmussen. *Floer Homology and Knot Complements*. PhD thesis, Harvard University, 2003, math.GT/0306378.
- [Sar06] Sucharit Sarkar. Maslov index of holomorphic triangles. 2006, math/0609673.
- [Smi03] Ivan Smith. Serre-Taubes duality for pseudoholomorphic curves. *Topology*, 42(5):931–979, 2003.
- [SW] Sucharit Sarkar and Jiajun Wang. An algorithm for computing some Heegaard Floer homologies. math/0607777.
- [Ush04] Michael Usher. The Gromov invariant and the Donaldson-Smith standard surface count. *Geom. Topol.*, 8:565–610 (electronic), 2004.