COMBINATORIAL COBORDISM MAPS IN HAT HEEGAARD FLOER THEORY

ROBERT LIPSHITZ, CIPRIAN MANOLESCU, and JIAJUN WANG

Abstract

In a previous article, Sarkar and Wang [15] gave a combinatorial description of the hat version of Heegaard Floer homology for three-manifolds. Given a cobordism between two connected three-manifolds, there is an induced map between their Heegaard Floer homologies. Assume that the first homology group of each boundary component surjects onto the first homology group of the cobordism (modulo torsion). Under this assumption, we present a procedure for finding the rank of the induced Heegaard Floer map combinatorially, in the hat version.

Contents

1. Introduction

In their articles [9]–[11], Ozsváth and Szabó constructed a decorated topological quantum field theory (TQFT) in $(3+1)$-dimensions, called *Heegaard Floer theory*. (Strictly speaking, the axioms of a TQFT need to be altered slightly.) In its simplest version (called *hat*), to a closed, connected, oriented three-manifold Y and a Spin^c-structure $\mathfrak{s}$ on Y, one associates a vector space $\widehat{HF}(Y,\mathfrak{s})$ over the field $\mathbb{F}=\mathbb{Z}/2\mathbb{Z}$. Also, to a connected, oriented four-dimensional cobordism from Y_1 to Y_2 decorated

DUKE MATHEMATICAL JOURNAL
Vol. 145, No. 2, © 2008 DOI 10.1215/00127094-2008-050
Received 15 December 2006. Revision received 30 January 2008.
2000 *Mathematics Subject Classification.* Primary 57R58; Secondary 57R56.
Lipshitz's work partially supported by a National Science Foundation Mathematical Sciences Postdoctoral Research Fellowship.
Manolescu's work partially supported by a Clay Research Fellowship.
Wang's work partially supported by National Science Foundation Holomorphic Curves Focused Research Group grant DMS-0244663.

with a Spin^c-structure $\mathfrak{t}$, one associates a map

$$\hat{F}_{W,\mathfrak{t}} : \widehat{HF}(Y_1, \mathfrak{t}|_{Y_1}) \to \widehat{HF}(Y_2, \mathfrak{t}|_{Y_2}).$$

The maps $\hat{F}_{W,\mathfrak{t}}$ can be used to detect exotic smooth structures on four-manifolds with boundary. For example, this can be seen by considering the nucleus $X = N(2)$ of the elliptic surface $E(2) = K3$, that is, a regular neighborhood of a cusp fiber and a section (cf. [2, page 74]). Let $X' = N(2)_p$ be the result of a log transform with multiplicity p ($p > 1$, odd) on a regular fiber $T^2 \subset X$ (cf. [2, Section 3.3]). Then X and X' are homeomorphic four-manifolds (with $\pi_1 = 1$) having as boundary the Brieskorn sphere $\Sigma(2, 3, 11)$. However, they are not diffeomorphic; this can be shown using the Donaldson or Seiberg-Witten invariants (see [4], [6], [1]) but also by comparing the hat Heegaard Floer invariants $\hat{F}_{W,\mathfrak{t}}$ and $\hat{F}_{W',\mathfrak{t}}$, where W and W' are the cobordisms from S^3 to $\Sigma(2, 3, 11)$ obtained by deleting a four-ball from X and X', respectively. Indeed, the arguments of Fintushel and Stern [1] and Stipsicz and Szabó [16] can be easily adapted to show that W and W' have different hat Heegaard Floer invariants; one needs to use the computation of $\widehat{HF}(\Sigma(2, 3, 11))$ due to Ozsváth and Szabó [7, page 47] and the rational blow-down formula of Roberts [13, Theorem 3]. (It is worth noting that the maps $\hat{F}$ give no nontrivial information for closed four-manifolds; this is a consequence of [9, Proposition 3.1] and [11, Remark 3.2, Lemma 8.2]. Exotic structures on closed four-manifolds can be detected with the mixed Heegaard Floer invariants of [11].)

The original definitions of the vector spaces $\widehat{HF}$ and the maps $\hat{F}$ involved counting pseudoholomorphic disks and triangles in symmetric products of Riemann surfaces; the Riemann surfaces are related to the three-manifolds and cobordisms involved via Heegaard diagrams. In [15, Section 4], Sarkar and Wang showed that every three-manifold admits a Heegaard diagram that is nice in the following sense: the curves split the diagram into elementary domains, all but one of which are bigons or rectangles. Using such a diagram, holomorphic disks in the symmetric product can be counted combinatorially, and the result is a combinatorial description of $\widehat{HF}(Y)$ for any Y as well as of the hat version of Heegaard Floer homology of null homologous knots and links in any three-manifold Y. A similar result was obtained in [5] for all versions of the Heegaard Floer homology of knots and links in the three-sphere.

The goal of this article is to give a combinatorial procedure for calculating the ranks of the maps $\hat{F}_{W,\mathfrak{s}}$ when W is a cobordism between Y_1 and Y_2 with the property that the induced maps $H_1(Y_1; \mathbb{Z})/\text{torsion} \to H_1(W; \mathbb{Z})/\text{torsion}$ and $H_1(Y_2; \mathbb{Z})/\text{torsion} \to H_1(W; \mathbb{Z})/\text{torsion}$ are surjective. Note that this case includes all cobordisms for which $H_1(W; \mathbb{Z})$ is torsion as well as all those consisting of only two-handle additions.

Roughly, the computation of the ranks of $\hat{F}_{W,\mathfrak{s}}$ goes as follows. The cobordism W is decomposed into a sequence of one-handle additions, two-handle additions,

and three-handle additions. Using the homological hypotheses on the cobordism and the (H_1/torsion)-action on the Heegaard Floer groups, we reduce the problem to the case of a cobordism map corresponding to two-handle additions only. Then, given a cobordism made of two-handles, we show that it can be represented by a multi-pointed triple Heegaard diagram of a special form, in which all elementary domains that do not contain base points are bigons, triangles, or rectangles. In such diagrams, all holomorphic triangles of Maslov index zero can be counted algorithmically, thus giving a combinatorial description of the map on $\widehat{HF}$.

We note that in order to turn $\widehat{HF}$ into a fully combinatorial TQFT (at least for cobordisms satisfying our hypothesis), one ingredient is missing: naturality. Given two different nice diagrams for a three-manifold, the invariance theorem of Ozsváth and Szabó [9, Theorem 1.1] shows that the resulting groups $\widehat{HF}$ are isomorphic. However, there is not yet a combinatorial description of this isomorphism. Thus, while the results of this article give an algorithmic procedure for computing the rank of a map $\hat{F}_{W,\mathfrak{s}}$, the map itself is determined combinatorially only up to automorphisms of the image and the target. In fact, if one were to establish naturality, then one could automatically remove the assumption on the maps on H_1/torsion (and compute $\hat{F}_{W,\mathfrak{s}}$ for any W) simply by composing the maps induced by the two-handle additions (computed in this article) with the ones induced by the one- and three-handle additions, which are combinatorial by definition (cf. [11, Section 4.3]).

The article is organized as follows. In Section 2, we define a multipointed triple Heegaard diagram to be nice if all nonpunctured elementary domains are bigons, triangles, or rectangles, and we show that in a nice diagram, holomorphic triangles can be counted combinatorially.* We then turn to the description of the map induced by two-handle additions. In Section 3, for the sake of clarity, we explain in detail the case of adding a single two-handle; we show that its addition can be represented by a nice triple Heegaard diagram with a single base point and, therefore, that the induced map on $\widehat{HF}$ admits a combinatorial description. We then explain how to modify the arguments to work in the case of several two-handle additions. This modification uses triple Heegaard diagrams with several base points. In Section 4, we discuss the additions of one- and three-handles and put the various steps together. Finally, in Section 5, we present the example of $+1$-surgery on the trefoil.

Throughout the article, all homology groups are taken with coefficients in $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$, unless otherwise noted.

2. Holomorphic triangles in nice triple Heegaard diagrams

The goal of this section is to show that under an appropriate condition (what we call *niceness*) on triple Heegaard diagrams, the counts of holomorphic triangles in the symmetric product are combinatorial.

*Sarkar [14] has independently obtained this result (see Proposition 2.5) using slightly different methods.

2.1. Preliminaries

We start by reviewing some facts from Heegaard Floer theory. A triple Heegaard diagram $\mathscr{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ consists of a surface Σ of genus g together with three $(g+k)$-tuples of pairwise disjoint embedded curves $\boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_{g+k}\}$, $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_{g+k}\}$, $\boldsymbol{\gamma} = \{\gamma_1, \dots, \gamma_{g+k}\}$ in Σ such that the span of each $(g+k)$-tuple of curves in $H_1(\Sigma)$ is g-dimensional. If we forget one set of curves (e.g., $\boldsymbol{\gamma}$), the result is an (ordinary) Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$.

By the condition on the spans, $\Sigma \setminus \boldsymbol{\alpha}$, $\Sigma \setminus \boldsymbol{\beta}$, and $\Sigma \setminus \boldsymbol{\gamma}$ each have $k+1$ connected components. By a multipointed triple Heegaard diagram $(\mathscr{H}, \mathbf{z})$, then, we mean a triple Heegaard diagram $\mathscr{H}$ as above together with a set $\mathbf{z} = \{z_1, \dots, z_{k+1}\} \subset \Sigma$ of $k+1$ points in Σ so that exactly one z_i lies in each connected component of $\Sigma \setminus \boldsymbol{\alpha}$, $\Sigma \setminus \boldsymbol{\beta}$, and $\Sigma \setminus \boldsymbol{\gamma}$.

To a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, one can associate a three-manifold $Y_{\alpha,\beta}$. To a triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$, in addition to the three-manifolds $Y_{\alpha,\beta}$, $Y_{\beta,\gamma}$ and $Y_{\alpha,\gamma}$, one can associate a four-manifold $W_{\alpha,\beta,\gamma}$ such that $\partial W_{\alpha,\beta,\gamma} = -Y_{\alpha,\beta} \cup -Y_{\beta,\gamma} \cup Y_{\alpha,\gamma}$ (see [9, Section 8.1]).

Associated to a three-manifold Y is the Heegaard Floer homology group $\widehat{HF}(Y)$. This was defined in [9] using a Heegaard diagram with a single base point. In [12], Ozsváth and Szabó associated to the data $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$—called a *multipointed Heegaard diagram*—a Floer homology group $\widehat{HF}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$ by counting holomorphic disks in $\mathrm{Sym}^{g+k}(\Sigma \setminus \mathbf{z})$ with boundary on the tori $\mathbb{T}_\alpha = \alpha_1 \times \cdots \times \alpha_{g+k}$ and $\mathbb{T}_\beta = \beta_1 \times \cdots \times \beta_{g+k}$. It is not hard to show that

$$\widehat{HF}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}) \cong \widehat{HF}(Y_{\alpha,\beta}) \otimes H_*(T^k). \tag{1}$$

(Here, T^k is the k-torus, and $H_*(T^k)$ means ordinary (singular) homology.) The decomposition (1) is not canonical; it depends on a choice of paths in Σ connecting z_i to z_1 for $i = 2, \dots, k+1$.

The Heegaard Floer homology groups decompose as a direct sum over Spin^c-structures on $Y_{\alpha,\beta}$,

$$\widehat{HF}(Y_{\alpha,\beta}) \cong \bigoplus_{\mathfrak{s} \in \mathrm{Spin}^c(Y_{\alpha,\beta})} \widehat{HF}(Y_{\alpha,\beta}, \mathfrak{s}).$$

More generally, there is a decomposition

$$\begin{aligned} \widehat{HF}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}) &\cong \bigoplus_{\mathfrak{s} \in \mathrm{Spin}^c(Y_{\alpha,\beta})} \widehat{HF}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}, \mathfrak{s}) \\ &\cong \bigoplus_{\mathfrak{s} \in \mathrm{Spin}^c(Y_{\alpha,\beta})} \left(\widehat{HF}(Y_{\alpha,\beta}, \mathfrak{s}) \otimes H_*(T^k)\right). \end{aligned}$$

Associated to the triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ together with a Spin^c-structure $\mathfrak{t}$ on $W_{\alpha,\beta,\gamma}$ is a map

$$\hat{F}_{\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\mathbf{z},\mathfrak{t}} : \widehat{HF}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}, \mathfrak{t}|_{Y_{\alpha,\beta}}) \otimes \widehat{HF}(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z}, \mathfrak{t}|_{Y_{\beta,\gamma}}) \to \widehat{HF}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{z}, \mathfrak{t}|_{Y_{\alpha,\gamma}}). \tag{2}$$

The definition involves counting holomorphic triangles in $\mathrm{Sym}^{g+k}(\Sigma \setminus \mathbf{z})$ with boundary on $\mathbb{T}_\alpha$, $\mathbb{T}_\beta$, and $\mathbb{T}_\gamma$ (cf. [9, Section 8] and [12]).

Two triple Heegaard diagrams are called *strongly equivalent* if they differ by a sequence of isotopies and handleslides. It follows from [9, Proposition 8.14], the associativity theorem [9, Theorem 8.16], and the definition of the handleslide isomorphisms that strongly equivalent triple Heegaard diagrams induce the same map on homology.

Call a $(k+1)$-pointed triple Heegaard diagram *split* if it is obtained from a singly pointed Heegaard triple diagram $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', z')$ by attaching (by connect sum) k spheres with one base point and three isotopic curves (one alpha, one beta, and one gamma) each to the component of $\Sigma' \setminus (\boldsymbol{\alpha}' \cup \boldsymbol{\beta}' \cup \boldsymbol{\gamma}')$ containing z'. We call $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', z')$ the *reduction* of the split diagram.

The following lemma is a variant of [12, Proposition 3.3].

LEMMA 2.1

Every triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ is strongly equivalent to a split one.

Proof

Reorder the alpha circles so that $\alpha_1, \ldots, \alpha_g$ are linearly independent. Let R be the connected component of $\Sigma \setminus \boldsymbol{\alpha}$ containing z_2. Since $\alpha_1, \ldots, \alpha_g$ are linearly independent, one of the curves α_i, $i > g$, must appear in the boundary of R with multiplicity exactly one. By handlesliding this curve over the other boundary components of R, we can arrange that the resulting α_i bounds a disk containing only z_2. Repeat this process with the other z_i, $i = 3, \ldots, k+1$, being sure to use a different alpha curve in the role of α_i at each step. We reorder the curves $\alpha_{g+1}, \ldots, \alpha_{g+k}$ so that α_i encircles z_{i-g+1}. Now, repeat the entire process for the beta and gamma curves. Finally, choose a path ζ_i in Σ from z_i to z_1 for each $i = 2, \ldots, k+1$. Move the configuration $(z_i, \alpha_i, \beta_i, \gamma_i)$ along the path ζ_i by handlesliding (around α_i, β_i, or γ_i) the other alpha, beta, and gamma curves that are encountered along the path ζ_i. The result is a split triple Heegaard diagram. Note that its reduction is obtained from the original diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ by simply forgetting some of the curves and base points. □

The maps from (2) are compatible with the isomorphism (1) in the following sense. Given a triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$, let $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', z')$ be the reduction

of a split diagram strongly equivalent to $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$. Then the following diagram commutes:

$$\begin{array}{ccc} \left(\widehat{HF}(Y_{\alpha,\beta}, \mathfrak{t}|_{Y_{\alpha,\beta}}) \otimes H_*(T^k)\right) \otimes \left(\widehat{HF}(Y_{\beta,\gamma}, \mathfrak{t}|_{Y_{\beta,\gamma}}) \otimes H_*(T^k)\right) & \longrightarrow & \widehat{HF}(Y_{\alpha,\gamma}, \mathfrak{t}|_{Y_{\alpha,\gamma}}) \otimes H_*(T^k) \\ \downarrow \cong & & \downarrow \cong \\ \widehat{HF}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z}, \mathfrak{t}|_{Y_{\alpha,\beta}}) \otimes \widehat{HF}(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z}, \mathfrak{t}|_{Y_{\beta,\gamma}}) & \xrightarrow{\hat{F}_{\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\mathbf{z},\mathfrak{t}}} & \widehat{HF}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{z}, \mathfrak{t}|_{Y_{\alpha,\gamma}}) \end{array} \tag{3}$$

Here, the map in the first row is $\hat{F}_{\Sigma',\boldsymbol{\alpha}',\boldsymbol{\beta}',\boldsymbol{\gamma}',z',\mathfrak{t}}$ on the $\widehat{HF}$-factors and the usual intersection product $H_*(T^k) \otimes H_*(T^k) \to H_*(T^k)$ on the $H_*(T^k)$-factors. The vertical isomorphisms are induced by the strong equivalence. The proof that the diagram commutes follows from the same ideas as in [12, Theorem 4.5]: each of the k spherical pieces in the split diagram contributes an $H_*(S^1)$ to the $H_*(T^k) = (H_*(S^1))^{\otimes k}$ factors above; moreover, a local computation shows that the triangles induce intersection product maps $H_*(S^1) \otimes H_*(S^1) \to H_*(S^1)$ which, tensored together, give the intersection product on $H_*(T^k)$.

Remark 2.2
It is tempting to assert that the map $\hat{F}_{\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},z,\mathfrak{t}}$ induced by a singly pointed triple Heegaard diagram depends only on $W_{\alpha,\beta,\gamma}$ and $\mathfrak{t}$. However, this seems not to be known.

2.2. Index formulas
Fix a triple Heegaard diagram $\mathscr{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ as above. The complement of the $3(g + k)$ curves in Σ has several connected components, which we denote by $D_1, \ldots, D_N$ and call *elementary domains*.

The *Euler measure* of an elementary domain $D \subset \Sigma$ is

$$e(D) = \chi(D) - \frac{\#\text{ vertices of } D}{4}.$$

A *domain* in Σ is a two-chain $D = \sum a_i D_i$ with $a_i \in \mathbb{Z}$. Its Euler measure is simply

$$e(D) = \sum_{i=0}^{N} a_i e(D_i).$$

As mentioned above, the maps $\hat{F}_{\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\mathbf{z},\mathfrak{t}}$ induced by the triple Heegaard diagram $\mathscr{H}$ are defined by counting holomorphic triangles in $\mathrm{Sym}^{g+k}(\Sigma)$ with respect to a suitable almost complex structure. According to the cylindrical formulation from [3, Section 10], this is equivalent to counting certain holomorphic embeddings

$u : S \to \Delta \times \Sigma$, where S is a Riemann surface (henceforth called the *source*) with some marked points on the boundary (which we call *corners*) and Δ is a fixed disk with three marked points on the boundary. The maps u are required to satisfy certain boundary conditions and to be generically $(g + k)$-to-1 when postcomposed with the projection $\pi_\Delta : \Delta \times \Sigma \to \Delta$. More generally, we consider such holomorphic maps $u : S \to \Delta \times \Sigma$ which are generically m-to-1 when postcomposed with π_Δ; these correspond to holomorphic triangles in $\mathrm{Sym}^m(\Sigma)$, where m can be any positive integer. We are interested in the discussion of the index from [3]. Although this discussion was carried out in the case $k = 0$, $m = g$, it applies equally well in the case of arbitrary k and m with only notational changes.

In the cylindrical formulation, one works with an almost complex structure on $\Delta \times \Sigma$ so that the projection π_Δ is holomorphic and the fibers of π_Σ are holomorphic. It follows that for $u : S \to \Delta \times \Sigma$ holomorphic, $\pi_\Delta \circ u$ is a holomorphic branched cover. The map $\pi_\Sigma \circ u$ need not be holomorphic, but since the fibers are holomorphic, $\pi_\Sigma \circ u$ is a branched map. Fix a model for Δ in which the three marked points are 90° corners and a conformal structure on Σ with respect to which the intersections between alpha, beta, and gamma curves are all right angles. Since u is holomorphic, the conformal structure on S is induced via $\pi_\Delta \circ u$ from the conformal structure on Δ. It makes sense, therefore, to talk about branch points of $\pi_\Sigma \circ u$ on the boundary and at the corners as well as in the interior. Generically, while there may be branch points of $\pi_\Sigma \circ u$ on the boundary of S, there are not branch points at the corners.

Suppose that $u : S \to \Delta \times \Sigma$ is as above. Denote by $\pi_\Sigma : \Delta \times \Sigma \to \Sigma$ the projection to Σ. There is an associated domain $D(u)$ in Σ, where the coefficient of D_i in $D(u)$ is the local multiplicity of $\pi_\Sigma \circ u$ at any point in D_i. By [3, page 1018], the *index* of the linearized $\bar{\partial}$ operator at the holomorphic map u is given by

$$\mu(u) = 2e\big(D(u)\big) - \chi(S) + \frac{m}{2}. \tag{3}$$

For simplicity, we call this the *index* of u.

Note that by the Riemann-Hurwitz formula,

$$\chi(S) = e\big(D(u)\big) + \frac{3m}{4} - \mathrm{br}(u), \tag{4}$$

where $\mathrm{br}(u)$ is the ramification index (number of branch points counted with multiplicity) of $\pi_\Sigma \circ u$. (Here, branch points along the boundary count as half an interior branch point.) From (3) and (4), we get an alternate formula for the index:

$$\mu(u) = e\big(D(u)\big) + \mathrm{br}(u) - \frac{m}{4}. \tag{5}$$

Note that it is not obvious how to compute $\mathrm{br}(u)$ from $D(u)$. A combinatorial formula for the index, purely in terms of $D(u)$, was found by Sarkar in [14]. However, we do not use it here.

2.3. Nice triple diagrams

Fix a multipointed triple Heegaard diagram $(\mathscr{H}, \mathbf{z}) = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$. Recall that a domain is a linear combination of connected components of $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma})$. The *support* of a domain is the union of those components with nonzero coefficients. If the support of a domain D contains at least one z_i, then D is called *punctured*; otherwise, it is called *unpunctured*.

Definition 2.3

An elementary domain is called *good* if it is a bigon, a triangle, or a rectangle and *bad* otherwise. The multipointed triple Heegaard diagram $(\mathscr{H}, \mathbf{z})$ is called *nice* if every unpunctured elementary domain is good.

This is parallel to the definition of nice Heegaard diagrams (with just two sets of curves) from [15, Section 3]. A multipointed Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$ is called *nice* if, among the connected components of $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$, all unpunctured ones are either bigons or squares.

Note that a bigon, a triangle, and a rectangle have Euler measure $1/2, 1/4$, and zero, respectively. Since e is additive, every unpunctured positive domain (not necessarily elementary) in a nice diagram must have nonnegative Euler measure. A quick consequence of this is the following.

LEMMA 2.4

If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ is a nice triple Heegaard diagram, then if we forget one set of curves (for example, $\boldsymbol{\gamma}$), the resulting Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$ is also nice.

In order to define the triangle maps, it is necessary to assume that the triple Heegaard diagram is weakly admissible in the sense of [9, Definition 8.8]. In fact, nice diagrams are automatically weakly admissible (cf. Corollary 3.2).

Our goal is to give a combinatorial description of the holomorphic triangle counts for nice triple diagrams.

PROPOSITION 2.5

Let $(\mathscr{H}, \mathbf{z})$ be a nice multipointed triple Heegaard diagram. Fix a generic almost complex structure J on $\Delta \times \Sigma$, as in [3, Section 10.2]. Let $u : S \to \Delta \times \Sigma$ be a J-holomorphic map of the kind occurring in the definition of $\hat{F}_{\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\mathbf{z},\mathfrak{t}}$. In particular, assume that u is an embedding of index zero and such that the image of $\pi_\Sigma \circ u$ is an

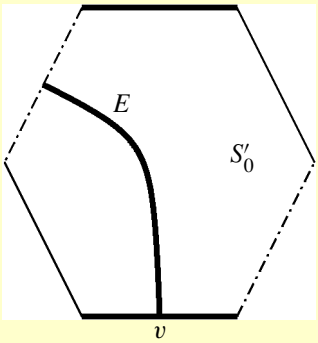

Figure 1. A hexagon component of the source. The preimages of the alpha, beta, and gamma curves (here shown as thick, thin, and interrupted lines) give an embedded graph in the source S_0. A boundary branch point in the image corresponds to a valence-three vertex v in the source.

unpunctured domain. Then S is a disjoint union of m triangles, and the restriction of $\pi_\Sigma \circ u$ to each component of S is an embedding.

Proof
Since the image of $\pi_\Sigma \circ u$ is unpunctured and positive, we have $e(D(u)) \geq 0$. By (4), we get

$$\chi(S) \geq \frac{m}{2} > 0.$$

This means that at least one component of S is topologically a disk. Let S_0 be such a component. It is a polygon with $3l$ vertices. We show that $l = 1$ and that $\pi_\Sigma \circ u|_{S_0}$ is an embedding.

Let us first show that S_0 is a triangle. The index of the $\bar{\partial}$-operator at a disconnected curve is the sum of the indices of its restrictions to each connected component. Therefore, in order for an index-zero holomorphic curve to exist generically, the indices at every connected component, and in particular at S_0, must be zero. Applying (4) to $u|_{S_0}$, we get

$$l = 2 - 4e\big(D(u|_{S_0})\big) \leq 2.$$

If $l = 2$, then by (6), we have $\mathrm{br}(u) \leq 1/2$. Hence, the map $\pi_\Sigma \circ u$ has no interior branch points. If $\mathrm{br}(u) = 0$, then S_0 is mapped locally diffeomorphically by $\pi_\Sigma \circ u$ to Σ. The image must have negative Euler measure, which is a contradiction. So, suppose that $\mathrm{br}(u) = 1/2$. The preimages of the alpha, beta, and gamma curves cut S_0 into several connected components. Without loss of generality, assume that the boundary branch point is mapped to an alpha circle. Then, along the corresponding edge of S_0, there is a valence-three vertex v, as shown in Figure 1. Let E denote the

edge in the interior of S_0 meeting v. Since there is only one boundary branch point, the other intersection point of the edge E with ∂S_0 is along the preimage of a beta or gamma circle. It follows that one of the connected components S_0' of $S_0 \setminus E$ is a hexagon or heptagon. Smoothing the vertex v of S_0', we obtain a pentagon or hexagon that is mapped locally diffeomorphically by $\pi_\Sigma \circ u$ to Σ. The image, then, has negative Euler measure, again a contradiction.

Therefore, $l = 1$, so S_0 is a triangle. Furthermore, by (6), $\mathrm{br}(u) \leq 1/4$, which means that there are no (interior or boundary) branch points at all. Thus, just as in the hypothetical hexagon case in Figure 1, the preimages of the alpha, beta, and gamma curves must cut S_0 into bigons, triangles, and rectangles, all of which have nonnegative Euler measure. Since the Euler measure of S_0 is $1/4$, there can be no bigons; in fact, S_0 must be cut into several rectangles and exactly one triangle. It is easy to see that the only possible tiling of S_0 of this type is as in Figure 2, with several parallel preimages of segments on the alpha curves, several parallel beta segments, and several parallel gamma segments. We call the type of a segment (α, β, or γ) its *color*.

The tiling consists of one triangle and six different types of rectangles, according to the coloring of their edges in clockwise order (namely, $\alpha\beta\alpha\gamma$, $\gamma\alpha\gamma\beta$, $\beta\alpha\beta\gamma$, $\alpha\beta\alpha\beta$, $\beta\gamma\beta\gamma$, and $\gamma\alpha\gamma\alpha$). We claim that the images of the interiors of each of these rectangles by $\pi_\Sigma \circ u$ are disjoint.

Because of the coloring scheme, only rectangles of the same type can have the same image. Suppose that two different $\alpha\beta\alpha\gamma$-rectangles from S_0 have the same image in Σ. (The cases $\gamma\alpha\gamma\beta$, $\beta\alpha\beta\gamma$ are exactly analogous.) Let r_1 and r_2 be the two rectangles; suppose that r_1 is closer to the central triangle than r_2, and suppose that r_2 is closer to the α-boundary of S_0. Because of the way the rectangles are colored, the upper edge of r_1 must have the same image as the upper edge of r_2. Hence, the $\alpha\beta\alpha\gamma$-rectangle right above r_1 has the same image as the one right above r_2. Iterating this argument, at some point we get that the central triangle has the same image as some $\alpha\beta\alpha\gamma$-rectangle, which is impossible.

Now, suppose that two different $\alpha\beta\alpha\beta$-rectangles, r_1 and r_2, have the same image. (The cases $\beta\gamma\beta\gamma$ and $\gamma\alpha\gamma\alpha$ are exactly analogous.) There are two cases, according to whether the upper edge of r_1 has the same image as the upper edge of r_2 or as the lower edge of r_2.

Suppose first that the upper edge of r_1 has the same image as the upper edge of r_2. By the *β-height* of r_i we mean the minimal number of beta arcs which an arc in $S \setminus \boldsymbol{\alpha}$ starting in r_i, going right, and ending at a gamma arc must cross. (The diagram is positioned in the plane, as in Figure 2.) Since $\pi_\Sigma \circ u$ is a local homeomorphism and r_1 and r_2 have the same image, it is clear that the β-height of r_1 and the β-height of r_2 are equal. By the *α-height* of r_i we mean the minimal number of alpha arcs which an arc in $S \setminus \boldsymbol{\beta}$ starting in r_i, going up, and ending at a gamma arc must cross. Again,

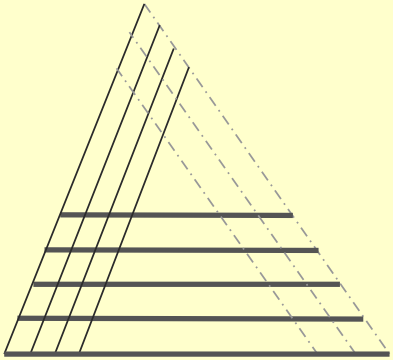

Figure 2. An embedded triangle. The thick lines are α's, the thin ones β's, and the interrupted ones γ's. We first show that this is the picture in the source S_0 and then that the same is true for the image of S_0 in Σ (i.e., the images of all the pieces in the tiling are disjoint).

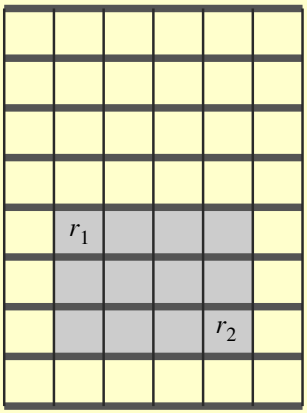

Figure 3. Two $\alpha\beta\alpha\beta$-rectangles in a grid. If the rectangles r_1 and r_2 have the same image in Σ with a 180° turn, then the whole shaded rectangle is mapped to Σ with a branch point.

it is clear that the α-heights of r_1 and r_2 must be equal. But this implies that r_1 and r_2 are equal.

Now, suppose that the upper edge of r_1 has the same image as the lower edge of r_2. There is a unique rectangle R in S with boundary contained in $\boldsymbol{\alpha} \cup \boldsymbol{\beta}$, containing r_1 and r_2 and with one corner the same as a corner of r_1 and the opposite corner the same as a corner of r_2. It is easy to see that $\pi_\Sigma \circ u$ maps antipodal points on the boundary of R to the same point in Σ. It follows that $\pi_\Sigma \circ u|_{\partial R}$ is a twofold covering map. But then, $\pi_\Sigma \circ u$ must have a branch point somewhere inside R—a contradiction (see Figure 3).

Finally, suppose that some arc A on ∂S has the same image as some other arc A' in S. If A' is in the interior of S, then any rectangle (or triangle) adjacent to A has the same image as some rectangle (or triangle) adjacent to A'. We have already ruled this out. If A' is on ∂S, then either any rectangle (or triangle) adjacent to A has the same image as some rectangle (or triangle) adjacent to A' or there is a branch point somewhere on ∂S. We have already ruled out both of these cases.

We have thus established that S_0 is an embedded triangle. By forgetting S_0, we obtain a holomorphic map to $\Delta \times \Sigma$ still of index zero but such that its postcomposition with π_Δ is generically $(m-1)$-to-1 rather than m-to-1. The result then follows by induction on m. □

Observe that in Proposition 2.5, even though each of the m triangles is embedded, some of their domains may overlap. It turns out that they may do so only in a specific way, however.

LEMMA 2.6

Suppose that A is an index-zero homology class represented by a union of embedded holomorphic triangles in a nice triple diagram. Suppose that the union of triangles corresponds to an embedded holomorphic curve in $\Delta \times \Sigma$. Then any two triangles in A are either disjoint in Σ or overlap in Σ head-to-tail, as shown in Figure 4.

Proof

Let T and T' be two of the triangles in the domain A. For a generic representative of A to exist, the pair must also have index zero and be embedded in $\Delta \times \Sigma$.

We already know that T and T' are tiled as in Figure 2. This strongly restricts how T and T' can overlap.

One way for T and T' to overlap is for T to be entirely contained inside T'. In this case, it is not hard to see that the two holomorphic triangles in $\Delta \times \Sigma$ intersect in one interior point. Indeed, the intersection number of two holomorphic curves in a four-manifold is invariant in families. If we deform the Heegaard diagram so that the boundary of T in Σ is a single point (i.e., the alpha, beta, and gamma circles involved in ∂T intersect in an asterisk, with vertex the triangle T), then obviously $(T \cap T') \subset (\Delta \times \Sigma)$ is a single point. It follows that the same is true for the original triangles T and T'.

Another way in which T and T' may overlap is head-to-head, as shown on the left-hand side of Figure 5. It is then possible to decompose $T \cup T'$ into a pair of rectangles R and R' and two new embedded triangles T and T', as shown in Figure 5. An immersed rectangle in Σ has index at least one since it admits a generic holomorphic representative. So, each of R_1 and R_2 has index at least one. Similarly, the pair of triangles $T_1 \cup T_2$ has index at least zero. So, by additivity of the index, the whole domain has index at least two—a contradiction.

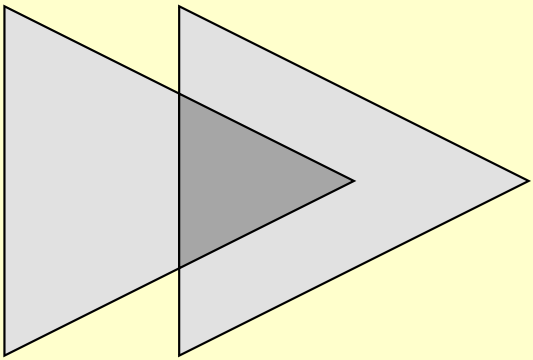

Figure 4. Head-to-tail overlap. A pair of embedded triangles T and T' overlap head-to-tail if their intersection $T \cap T'$ consists of a single connected component, itself a triangle, which contains one vertex of one of the two triangles and no vertices of the other. In the picture, the two triangles are shaded; the intersection is darkly shaded. We have not colored the figure to indicate that any of the three possible coloring schemes is allowed; however, in all three cases, parallel segments in the figure are of the same type (α, β, or γ).

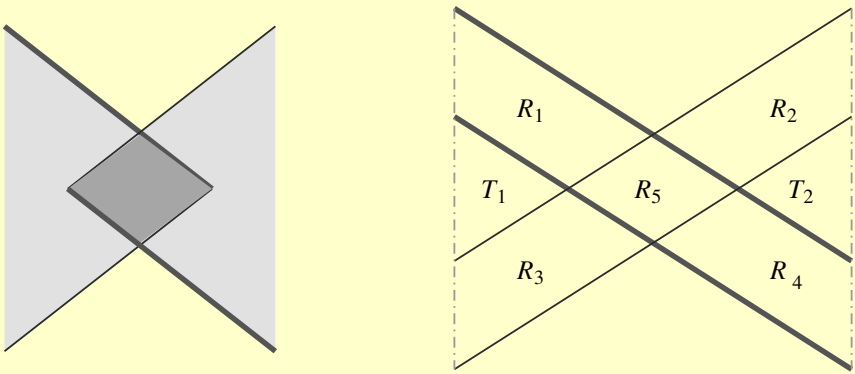

Figure 5. Head-to-head overlap. If T and T' intersect around their $\alpha\beta$-vertices, as on the left-hand side of the figure (where the intersection is shown by a darker shading), their union $T \cup T'$ does not have index zero. This can be seen using an alternate decomposition of the domain $T \cup T'$ as $T_1 \cup T_2 \cup R \cup R'$, where $R = R_1 \cup R_4 \cup R_5$ and $R' = R_2 \cup R_3 \cup R_5$ and where T_i ($i = 1, 2$) and R_j ($j = 1, \ldots, 5$) are the domains shown on the right-hand side. (Note that on the right-hand side, there may be overlaps in other parts of the diagram; e.g., we can have the situation in Figure 6.) The cases of $\alpha\gamma$ or $\beta\gamma$ head-to-head overlaps are similar.

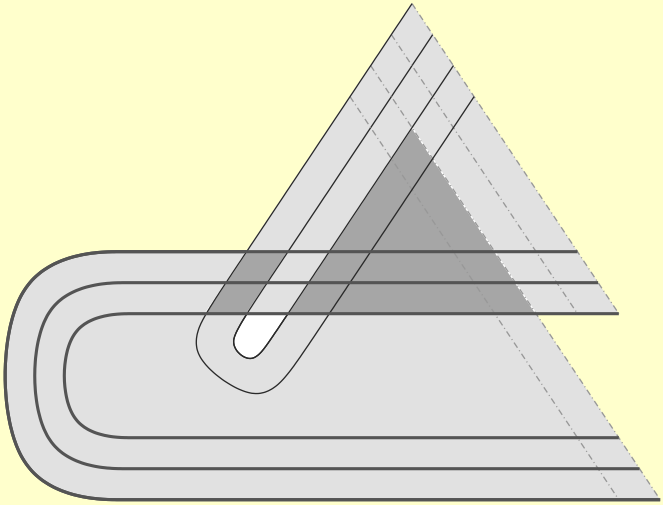

Figure 6. A double overlap. The two triangles (each shown lightly shaded) have a (darkly shaded) overlap with two connected components.

Using these two observations and the rulings of T and T', it is then elementary to check that the only possible overlap in index zero is head-to-tail, as in Figure 4. □

It follows that for a nice Heegaard diagram, we can combinatorially describe the generic holomorphic curves of index zero. If D is the domain of a generic holomorphic curve of index zero, then ∂D has m components, each of which bounds an embedded triangle in Σ. Each pair of triangles must either be disjoint or overlap, as shown in Figure 4. Any such D clearly has a unique holomorphic representative with respect to a split complex structure $j_\Delta \times j_\Sigma$ on $\Delta \times \Sigma$. Further, it is well known that these holomorphic curves are transversally cut out, and so they persist if one takes a small perturbation of $j_\Delta \times j_\Sigma$. In summary, to count index-zero holomorphic curves in $\Delta \times \Sigma$ with respect to a generic perturbation of the split complex structure, it suffices to count domains D that are sums of m embedded triangles in Σ, overlapping as allowed in the statement of Lemma 2.6.

3. Nice diagrams for two-handle additions

In [11, Definition 4.2], Ozsváth and Szabó associate to a four-dimensional cobordism W consisting of two-handle additions certain kinds of triple Heegaard diagrams. The cobordism W from Y_1 to Y_2 corresponds to surgery on some framed link $L \subset Y_1$. Denote by l the number of components of L. Fix a base point in Y_1. Let $B(L)$ be the union of L with a path from each component to the base point. The boundary of a

regular neighborhood of $B(L)$ is a genus l surface, which has a subset identified with l-punctured tori F_i, one for each link component. A singly pointed triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, z)$ is called *subordinate to* $B(L)$ if

- $(\Sigma, \{\alpha_1, \dots, \alpha_g\}, \{\beta_1, \dots, \beta_{g-l}\})$ describes the complement of $B(L)$;
- $\gamma_1, \dots, \gamma_{g-l}$ are small isotopic translates of $\beta_1, \dots, \beta_{g-l}$;
- after surgering out the $\{\beta_1, \dots, \beta_{g-l}\}$, the induced curves β_i and γ_i (for $i = g-l+1, \dots, g$) lie in the punctured torus F_i;
- for $i = g-l+1, \dots, g$, the curves β_i represent meridians for the link components, disjoint from all γ_j for $i \neq j$ and meeting γ_i in a single transverse intersection point;
- for $i = g-l+1, \dots, g$, the homology classes of the γ_i correspond to the framings of the link components.

A related construction is as follows. Given Y_1 and L as above, choose a multipointed Heegaard diagram $(\Sigma_L, \boldsymbol{\alpha}_L, \boldsymbol{\beta}_L, \mathbf{z}_L, \mathbf{w}_L)$ for $L \subset Y_1$, as in [12], of some genus g; here, $\boldsymbol{\alpha}_L = \{\alpha_1, \dots, \alpha_{g+l-1}\}$ and $\boldsymbol{\beta}_L = \{\beta_1, \dots, \beta_{g+l-1}\}$. Precisely, the sets $\mathbf{z} = \{z_1, \dots, z_l\}$ and $\mathbf{w} = \{w_1, \dots, w_l\}$ are collections of distinct points on Σ disjoint from the alpha and the beta curves and with the following two properties. First, each connected component of $\Sigma_L \setminus \boldsymbol{\alpha}_L$ and $\Sigma_L \setminus \boldsymbol{\beta}_L$ contains a single z_i and a corresponding w_i. Second, if c is an l-tuple of embedded arcs in $\Sigma_L \setminus \boldsymbol{\beta}_L$ connecting z_i to w_i $(i = 1, \dots, l)$ and c' is an l-tuple of embedded arcs in $\Sigma_L \setminus \boldsymbol{\alpha}_L$ connecting z_i to w_i $(i = 1, \dots, l)$, then the link L is the union of small pushoffs of c and c' into the two handle bodies (induced by the beta and alpha curves, resp.). Next, we attach handles h_i to Σ_L, connecting z_i to w_i (for $i = 1, \dots, l$), and obtain a new surface Σ. We choose a new β_i $(i = g+l, \dots, g+2l-1)$ to be the belt circle of the handle $h_{i-g-l+1}$ and a new α_i to be the union of the core of $h_{i-g-l+1}$ with c_i' so that α_i intersects β_i in one point. Choose γ_i $(i = 1, \dots, g+l-1)$ to be a small isotopic translate of β_i intersecting β_i at two points. Let γ_i^0 $(i = g+l, \dots, g+2l-1)$ be the union of c_i with a core of the handle $h_{i-g-l+1}$. Obtain γ_i $(i = g+l, \dots, g+2l-1)$ by applying Dehn twists to γ_i^0 around β_i; the framing of the link is determined by the number of Dehn twists.

In this fashion, we obtain a multipointed triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ with $g+2l-1$ curves of each kind. Note that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$ is an l-pointed Heegaard diagram for Y_1, $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{z})$ is an l-pointed Heegaard diagram for Y_2, and $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ is an l-pointed Heegaard diagram for $\#^g(S^1 \times S^2)$.

The new circles $\alpha_i, i = g+l, \dots, g+2l-1$, are part of a maximal homologically linearly independent subset of $\boldsymbol{\alpha}$, and similarly for β_i and γ_i $(i = g+l, \dots, g+2l-1)$. Consequently, by the proof of Lemma 2.1, there is a split diagram strongly equivalent to $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ whose reduction $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', z')$ is obtained from $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ by forgetting $l-1$ of the α_i (resp., β_i, γ_i), $1 \leq i \leq g+l-1$, as well as $z_2, \dots, z_k$. It

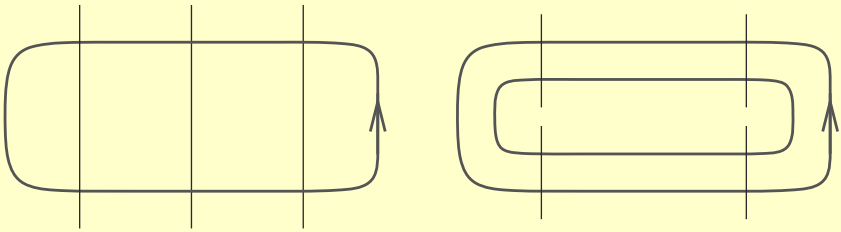

Figure 7. Alpha curves not adjacent to the bad region. The thick curves are alpha curves, and the thin ones are beta curves.

is then not hard to see that $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', z')$ is a triple Heegaard diagram subordinate to a bouquet for L.

In this section, we show that for any two-handle addition, one can construct a *nice* triple Heegaard diagram strongly equivalent to a diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$, as above. This involves finessing the diagram for the link L inside Y and then, after adding the handles and the new curves, modifying the diagram in several steps to make it nice. For the most part, we focus on the case where we add a single two-handle. In Section 3.3, we explain how the arguments generalize to several two-handles.

To keep language concise, we refer to elementary domains as *regions* in this section.

3.1. A property of nice Heegaard diagrams

Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$ be a multipointed Heegaard diagram. Recall that the diagram is nice if all unpunctured regions are either bigons or squares.

LEMMA 3.1

On a nice Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$, for any alpha circle α_i with an arbitrary orientation, there exists a punctured region D that contains an edge e belonging to α_i and such that D is on the left-hand side of α_i. The same conclusion holds for each beta circle.

Proof

Suppose that a half-neighborhood on the left-hand side of the alpha circle α_i is disjoint from all the punctured regions. Then, immediately to the left-hand side of α_i, we only have good regions. There are two possibilities, as indicated in Figure 7.

If there is a bigon region on the left-hand side of α_i, then the other edge is some beta edge β_j. The region on the other side of β_j must be a bigon region or a square since, otherwise, we would have a punctured region on the left-hand side of α_i. If we reach a square, we continue to consider the next region. Eventually, we reach a bigon

region since the number of regions are finite and we do not reach the same region twice. All regions involved form a disk bounded by α_i, as in Figure 7(a). In particular, this means that α_i is null-homologous. This contradicts the fact that the $g + k - 1$ alpha circles represent linearly independent classes in $H_1(\Sigma \setminus \mathbf{z})$.

In the second case, there are no bigon regions. Then on the left-hand side of α_i, we see a chain of squares, as in Figure 7(b). The opposite edges on these squares give another alpha circle, say, α_j. Then α_i and α_j are homologous to each other in $H_1(\Sigma \setminus \mathbf{z})$. This contradicts the same fact as in the previous case. □

Recall that in order to define the triangle maps, it is necessary for the triple Heegaard diagram to be weakly admissible in the sense of [9, Definition 8.8].

COROLLARY 3.2
If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ *is a nice multipointed triple Heegaard diagram, then* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ *is weakly admissible.*

Proof
By definition, the diagram is weakly admissible if there are no nontrivial domains D supported in $\Sigma\setminus\mathbf{z}$ with nonnegative multiplicity in all regions and whose boundary is a linear combination of alpha, beta, and gamma curves. Suppose that such a domain D exists, and consider a curve appearing with a nonzero multiplicity in ∂D. Without loss of generality, we can assume that this is an alpha curve and that all regions immediately to its left have positive multiplicity in D. By Lemma 2.4, the diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$ is nice. Lemma 3.1 now gives a contradiction. □

3.2. A single two-handle addition
Let (Y, K) be a three-manifold together with a knot $K \subset Y$. We choose a singly pointed Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ for Y together with an additional base point $w \neq z \in \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ so that the two base points determine the knot, as in [8, Section 2.2]. After applying the algorithm from [15, Section 4] to the Heegaard diagram, we can assume that the Heegaard diagram is nice with D_z the (usually bad) region containing the base point z. Furthermore, the algorithm in [15, Section 4] also ensures that D_z is a polygon.

We denote by D_w the region containing w; note that either $D_w = D_z$ or D_w is good. Throughout this section, we suppose that D_w and D_z are two different regions and that D_w is a rectangle. The case when $D_w = D_z$ corresponds to surgery on the unknot, which is already well understood. The case when D_w is a bigon can be avoided by modifying the original diagram by a finger move. (Alternately, this case can be treated similarly to the case when D_w is a rectangle.)

Let W be the four-manifold with boundary obtained from $Y \times [0, 1]$ by adding a two-handle along K in $Y \times \{1\}$ with some framing. W gives a cobordism between Y and Y', where Y' is obtained from Y by doing the corresponding surgery along K.

Now, we are ready to describe our algorithm to get a nice triple Heegaard diagram for the cobordism W.

Step 1: Making the knot embedded in the Heegaard diagram. Let c be an embedded arc in Σ connecting z and w in the complement of beta curves, and let c' be an embedded arc connecting z and w in the complement of alpha curves. The union of c and c' is a projection of the knot $K \subset Y$ to the surface Σ, where Σ is viewed as a Heegaard surface in Y. For convenience, we always assume that c and c' do not pass through any bigon regions and that they never leave a rectangle by the same edge through which they entered; this can easily be achieved.

In this step, we modify the doubly pointed Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z, w)$ to make $c \cup c'$ embedded in Σ, while preserving the niceness of the Heegaard diagram.

Typically, c and c' have many intersections. We modify the diagram inductively by stabilization at the first intersection $p \in D_p$ on c' (going from z to w) to remove that intersection, making sure that the new diagram is still nice.

A neighborhood of c and the part on c' from z to p are shown in Figure 8. In the same picture, if we continue the chain of rectangles containing c, we end up with a region D' that is either a bigon or the punctured region D_z.

To get rid of the intersection point p, we stabilize the diagram, as in Figure 9. More precisely, we do a stabilization followed by some handleslides of the beta curves and an isotopy of the new beta curve. After these moves, the number of intersection points decreases by one, and the diagram is still nice.

If we iterate this process, in the end we get a nice Heegaard diagram in which c and c' only intersect at their end points. Furthermore, the bad region D_z is still a polygon.

Step 2: Adding twin gamma curves. Our goal in steps 2 and 3 is to describe a particular triple Heegaard diagram for the cobordism W. Starting with the alpha and the beta curves we already have, for each beta curve β_i we add a gamma curve γ_i (called its *twin*) which is isotopic to β_i and intersects it in exactly two points. (After this, we add some more curves in the next step.)

For any beta curve β_i, by Lemma 3.1 we can choose a region D_i so that D_i is adjacent to the punctured region D_z with the common edge on β_i. If $D_i = D_z$, then we add γ_i close and parallel to β_i, as in Figure 10(a), and make a finger move, as in Figure 10(b).

Suppose now that D_i is different from D_z. Then D_i has to be a good region.

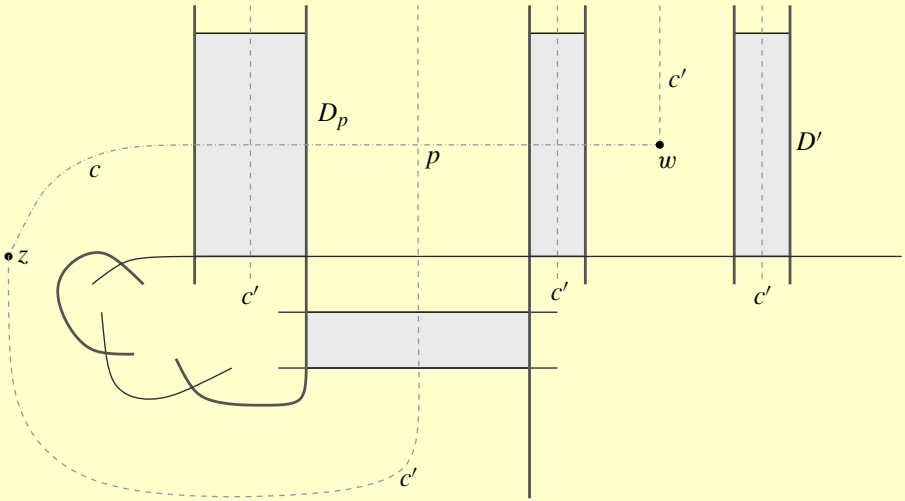

Figure 8. Projection of the knot to the Heegaard surface. As usual, the thick lines are alpha curves, and the thin lines are betas. The point p is the first on the dashed c'-curve (starting from z) where an intersection with c takes place. The light shading in the upper three domains indicates that there may be more parallel copies of alpha curves or parts of the c'-curve. Similarly, the light shading in the lower domain indicates the presence of an arbitrary number of parallel beta segments there. Note that some of the regions in the rightmost shaded domain on the top can coincide with some of the regions in the bottom shaded domain; however, this fact does not create any difficulties.

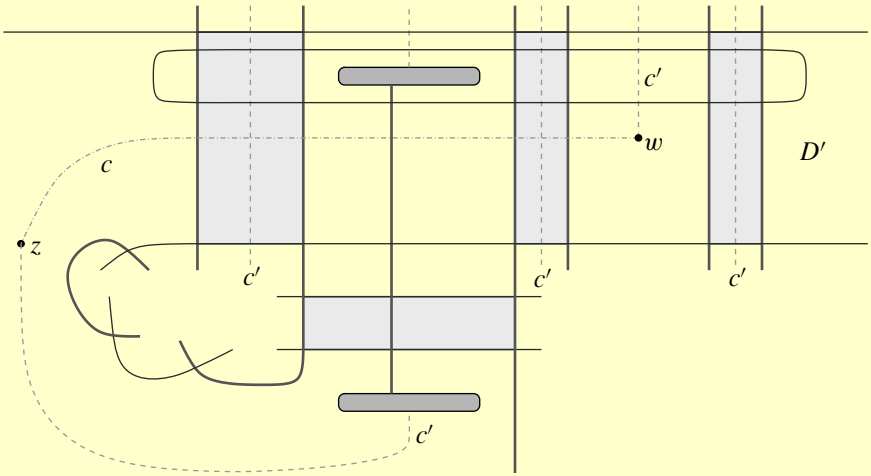

Figure 9. Stabilizing at the intersection. The two darkly shaded ovals are the two feet of the handles which we added. In the stabilized picture, we stretch the new beta curve until it reaches the regions D_z and D' so that the diagram is still nice.

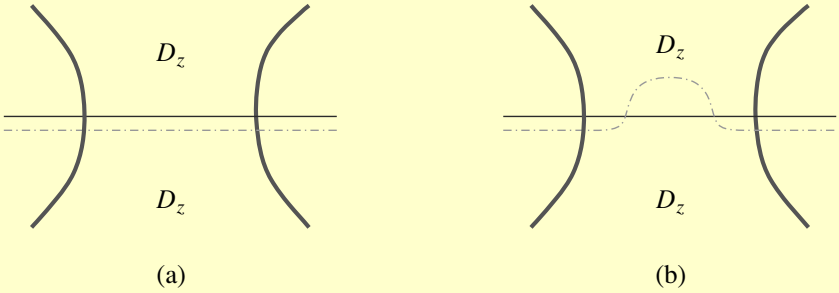

Figure 10. Adding twin gamma curves: Case 1. Here and after, without further specification, we make the convention that the thick arcs are alpha arcs, the thin ones are beta arcs, and the interrupted ones are gamma arcs.

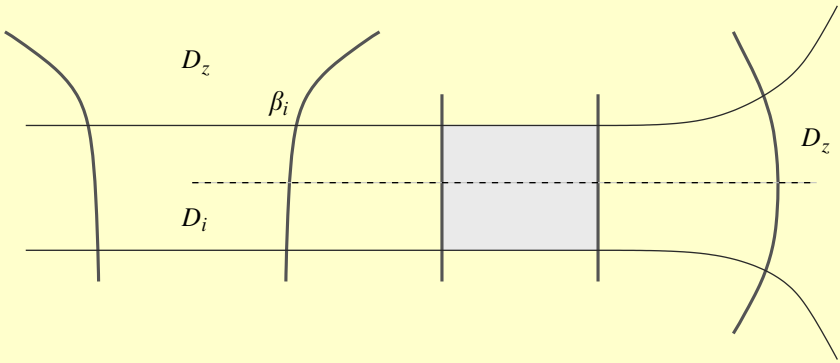

Figure 11. An arc connecting D_i to D_z. The shaded domain may contain several parallel alpha edges, and the dashed arc is the connecting arc.

If D_i is not a bigon since the complement of the beta curves in Σ is connected, we can connect D_i with D_z without intersecting beta curves via an arc traversing a chain of rectangles, as indicated in Figure 11. Then we do a finger move of the curve β_i, as indicated in Figure 12. Note that the knot remains embedded in Σ.

Now, we have a bigon region. We then add the gamma curve γ_i, as shown in Figure 12.

Note that for each pair β_i and γ_i, we either have one subdiagram of the form in Figure 10(b) or one subdiagram of the form in Figure 13. Observe also that during this process, no bad region other than D_z is created.

Step 3: Stabilization and two-handle addition. After step 2, the knot is still embedded in the Heegaard diagram. In other words, we can use arcs to connect z to w by paths in the complement of alpha curves and in the complement of beta curves so that the

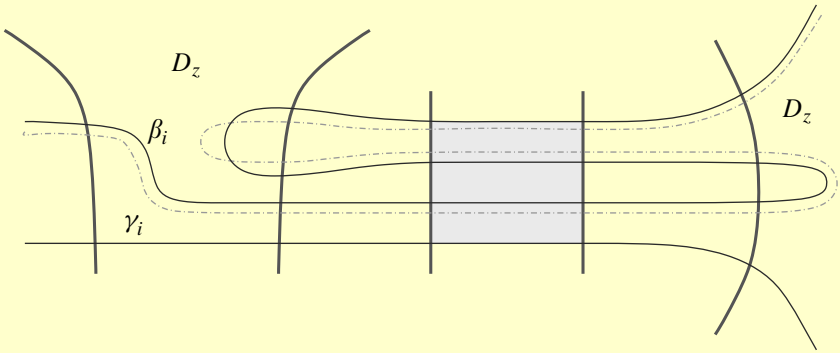

Figure 12. Adding twin gamma curves: Case 2. Before adding γ_i, we do the finger move shown here. Again, the shaded domain may contain several parallel alpha segments.

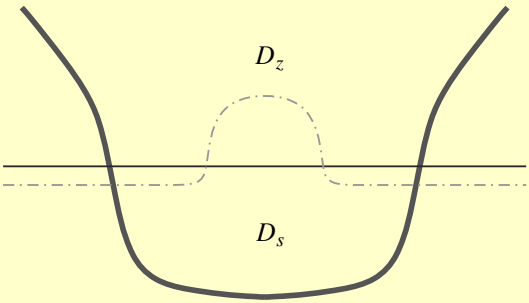

Figure 13. Bigon between beta and gamma curves. The region D_s is dealt with in a special way in later steps.

two arcs do not intersect except for the end points w and z and do not pass through any bigons. We see two chains of squares, as indicated in Figure 14.

We do a stabilization of the Heegaard diagram by adding a handle with one foot in each of D_z and D_w. We add the additional beta circle β_{g+1} to be the meridian of the handle, which we push along c until it reaches D_z. We also push β_{g+1} through the opposite alpha edge of D_w into the adjacent region. Then, we connect the two feet in the complement of alpha curves along c' and get a new alpha circle α_{g+1}. Finally, we add the surgery gamma circle γ_{g+1}, as in Figure 14. The result is a triple Heegaard diagram (with $3(g+1)$ curves) that represents surgery along the knot $K \subset Y$ with a particular framing; the framing is the sum of the number of twists of γ_{g+1} around the handle and a constant depending only on the original Heegaard diagram.

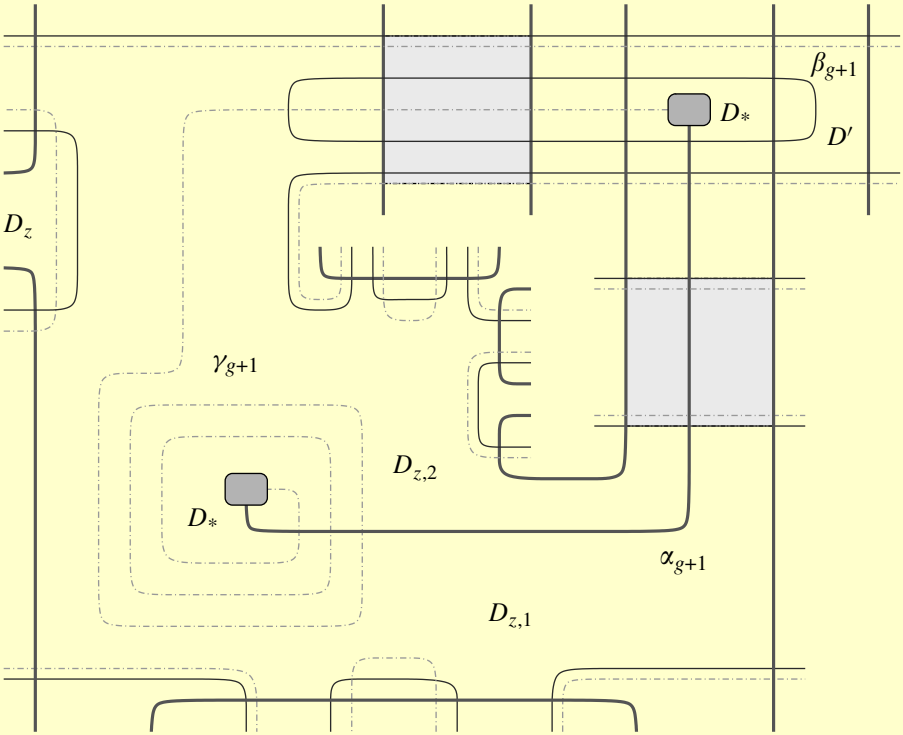

Figure 14. Stabilization: Case 1. The two small darkly shaded ovals are the two feet of the handle. The upper large lightly shaded area may have several parallel copies of alpha arcs, while the lower one may have several copies of beta and gamma arcs. The D_z on the upper left-hand side denotes either $D_{z,1}$ or $D_{z,2}$; such a domain can occur at various places on the boundaries of $D_{z,1}$ or $D_{z,2}$, and it corresponds to case 1 in step 2.

Note that, depending on the framing, the local picture around the two feet of the handle may also look like Figure 15, in which case, instead of the octagon region D_* from Figure 14, we have two hexagon bad regions $D_{*,1}$ and $D_{*,2}$.

After the stabilization, α_{g+1} and γ_{g+1} separate D_z into several regions; among these, $D_{z,1}$ and $D_{z,2}$ are (possibly) bad, but all other regions are good. We end up with a diagram with four (or five) bad regions: $D_{z,1}$, $D_{z,2}$, D_* (or $D_{*,1}$ and $D_{*,2}$), and D'. (In some cases, $D_{z,1}$ or $D_{z,2}$ might be good, or if there is little winding of γ_{g+1}, some of D_*, $D_{z,1}$, and $D_{z,2}$ might coincide. The argument in these cases is a simple adaptation of the one we give below.) We kill the badness of D', $D_{z,2}$, and D_* (or $D_{*,1}$ and $D_{*,2}$), while the region $D_{z,1}$ is the one containing the base point z for our final triple Heegaard diagram.

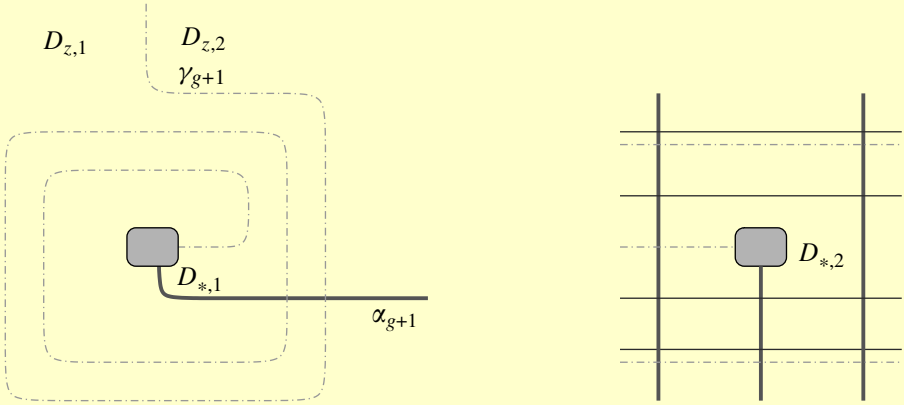

Figure 15. Stabilization: Case 2. If the γ_{g+1} twists around the handle in the opposite direction, we still have a picture similar to Figure 14. The only differences are in the neighborhoods of the two feet of the handle, which are shown here.

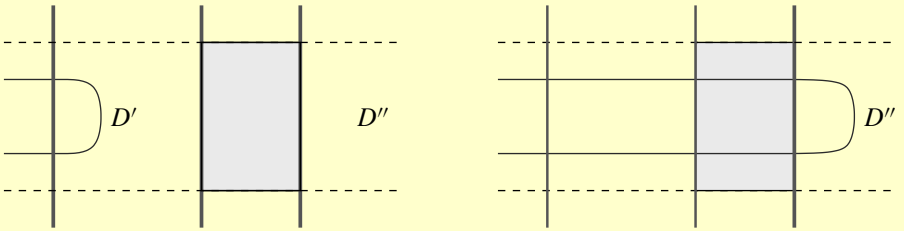

Figure 16. Killing the badness of D': Bigon, $D_{z,1}$-case or $D_{z,2}$-case. The thick arcs are alpha curves, the thin arcs represent β_{g+1}, and the dashed arcs can be either betas or gammas. The shaded part has several parallel alpha arcs. The rightmost region D'' is a bigon, $D_{z,1}$, or $D_{z,2}$.

Step 4: Killing the bad region D'. We push the finger in D' across the opposite alpha edge until we reach a bigon, $D_{z,1}$, $D_{z,2}$, or a region of type D_s, as in Figure 13.

Case 1: A bigon is reached. In this case (Figure 16(a)), our finger move kills the badness of D', as indicated in Figure 16(b), and does not create any new bad regions.

Case 2: $D_{z,1}$ or $D_{z,2}$ is reached. This is completely similar to case 1. The finger move kills the badness of D' and does not create any new bad regions.

Case 3: A region of type D_s is reached. Let us suppose that the topmost region in Figure 13 is $D_{z,1}$. The case when the topmost region is $D_{z,2}$ is completely similar.

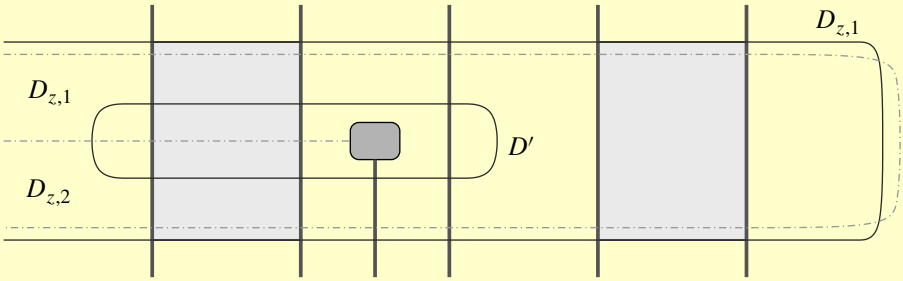

Figure 17. Killing the badness of D': Type D_s region, before. The smaller shaded region is the foot of the handle inside D_w, while the larger two shaded regions may contain several parallel alpha arcs.

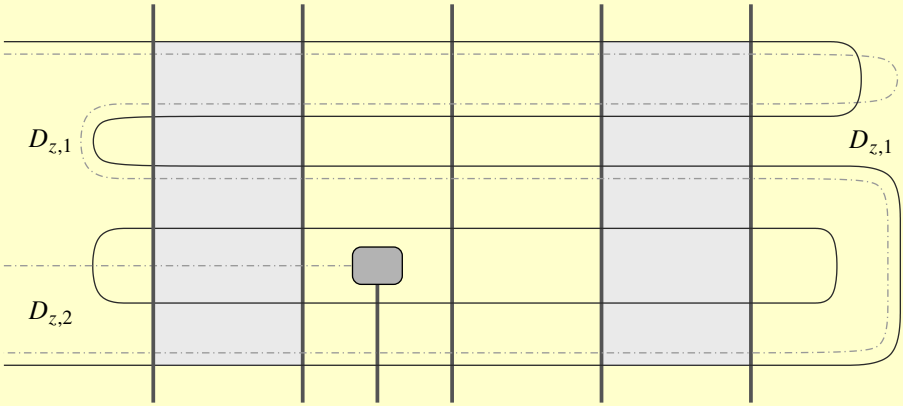

Figure 18. Killing the badness of D': Type D_s region, after. Conventions and shaded regions are the same as in Figure 17. This figure is obtained from Figure 17 via a double finger move.

The regions involved look like those in Figure 17. If on the left-hand side $D_{z,1}$ is on top of $D_{z,2}$, we isotope the diagram to look as in Figure 18. On the left-hand side of Figure 13, the case when $D_{z,2}$ is on top of $D_{z,1}$ is similar, except that we do the double finger move on the other side of β_{g+1}.

We have now killed the badness of D'.

Step 5: Killing the badness of $D_{z,2}$.

If there are any bigons between beta and gamma curves adjacent to $D_{z,2}$ (as in Figures 13 or 10(b) and also in Figure 19(a) (resp., Figure 19(c)), we do a handleslide (more precisely, a handleslide followed by an isotopy) of γ_{g+1} over each γ_i ($i \leq g$) involved, as indicated in Figure 19(b) (resp., Figure 19(d)).

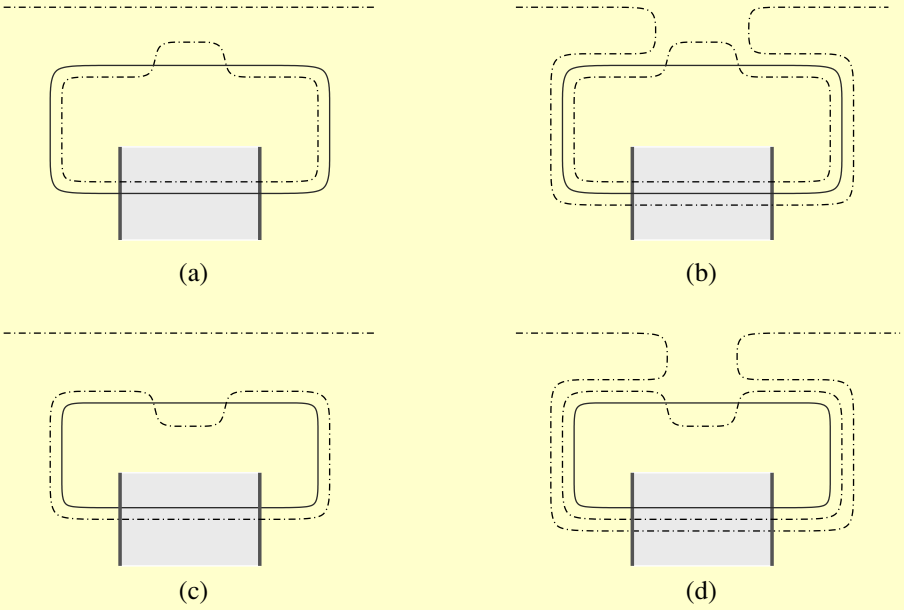

Figure 19. Killing the badness of $D_{z,2}$: Special handleslides. There may be more alpha arcs in the shaded regions.

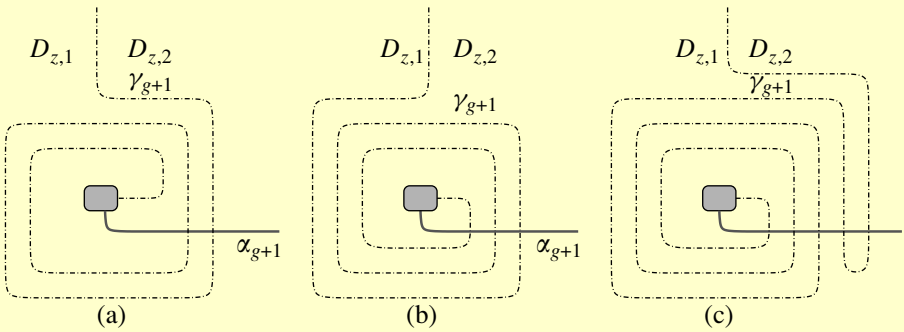

Figure 20. Killing the badness of $D_{z,2}$: Two patterns. The curve γ_{g+1} can rotate around the shaded oval either as in (a) or as in (b). If (b) occurs, we replace it with (c).

The intersection of γ_{g+1} and α_{g+1} has the pattern of Figure 20(a) or (b). In case (a), we do nothing. In case (b), we do the finger move, as in Figure 20(c).

Now, among the possibly bad regions generated from $D_{z,2}$, we have a unique one whose boundary has an intersection of a beta curve with a gamma curve, namely, the one near β_{g+1}, as in Figure 21(a) (see also Figures 14, 18).

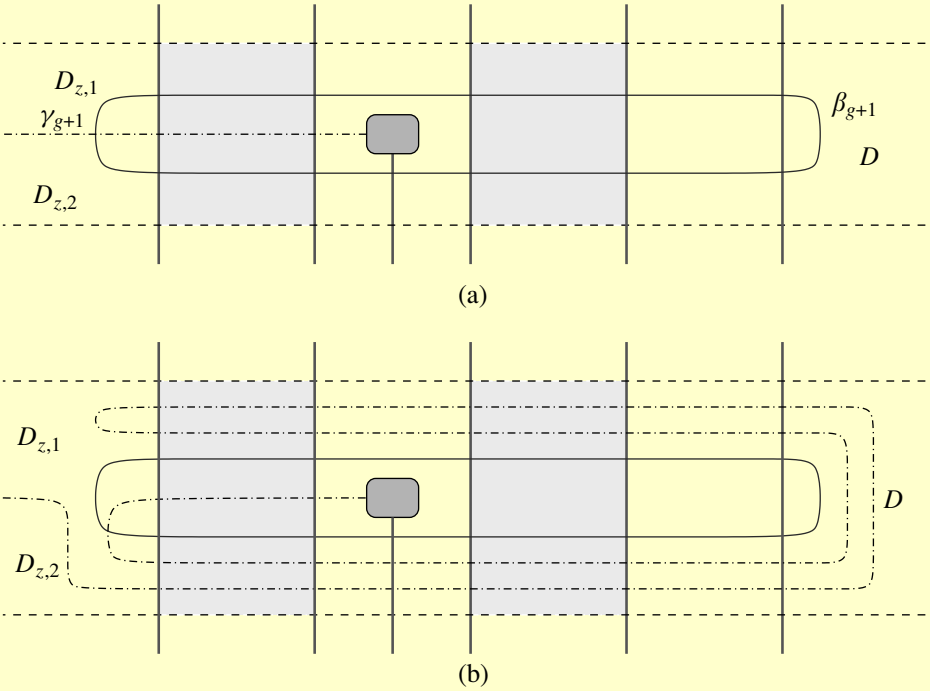

Figure 21. Killing the badness of $D_{z,2}$: Special beta-gamma crossing. The shaded area may contain several parallel alpha arcs. The dashed arcs can be beta arcs or gamma arcs, depending on different cases. The region D is a bigon, $D_{z,1}$, or $D_{z,2}$. We isotope (a) into (b) in order to remove the beta-gamma intersection point on the boundary of $D_{z,2}$.

We then do a finger move, as in Figure 21(b). Note that this finger move does not create any badness other than that of $D_{z,1}$.

After these special handleslides and finger moves, the region $D_{z,2}$ is divided into several possibly bad regions $R_1, \dots, R_m$. These bad regions are all adjacent to $D_{z,1}$ via arcs on γ_{g+1}; furthermore, there are no intersection points of beta and gamma curves on their boundaries. We seek to kill the badness of $R_1, \dots, R_m$ by using the algorithm in [15, Section 4]. The algorithm there consists of inductively decreasing a complexity function defined using the unpunctured bad regions. In our situation, we apply a simple modification of the algorithm to the Heegaard diagram made of the alpha and gamma curves; the modification consists of the fact that we do not deal with the bad region(s) D_* (or $D_{*,1}$ and $D_{*,2}$) but seek only to eliminate the badness of $R_1, \dots, R_m$. Thus, in the complexity function, we do not include terms that involve the badness and distance of D_* (or $D_{*,1}$ and $D_{*,2}$).

Since all the R_i's are adjacent to the preferred (punctured) region $D_{z,1}$ via arcs on γ_{g+1}, the algorithm in [15, Section 4] prescribes doing finger moves of γ_{g+1} through

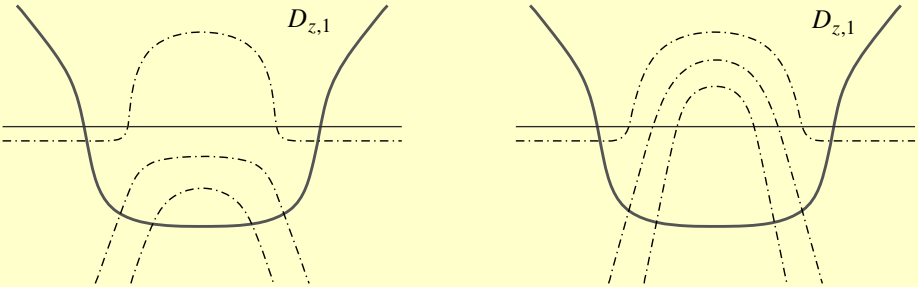

Figure 22. Killing the badness of special bad regions. We isotope the gamma curves to kill the hexagon in (a).

alpha curves and (possibly) handleslides of γ_{g+1} over other gamma curves. We do all these moves in such a way as not to tamper with the arrangements of twin beta-gamma curves, that is, so as not to introduce any new intersection points between γ_{g+1} and β_i for any $i \leq g$. (In other words, we can think of fattening $\gamma_1, \ldots, \gamma_g$ before applying the algorithm, so that they include their respective twin beta curves.) In particular, regions of type D_s are treated as bigons.

The fact that the algorithm in [15, Section 4] can be applied in this fashion is based on the following two observations.

- Our fingers or handleslides do not pass through the regions adjacent to β_{g+1}, except possibly $D_{z,1}$ itself. (This is one benefit of the modification performed in Figure 21.)
- We do not reach any squares between β_i and γ_i ($i \leq g$), nor any of the narrow squares created in Figure 19.

In the end, all the badness of $R_1, \ldots, R_m$ is killed. We arrive at a Heegaard diagram that may still have some bad regions coming from regions of type D_s, as in Figure 22(a). We kill these bad regions using the finger moves indicated in Figure 22(b).

After these moves, the only remaining bad regions are D_* (or $D_{*,1}$ and $D_{*,2}$) and the preferred bad region $D_{z,1}$.

Step 6: Killing the badness of D_* (or $D_{*,i}$). Our remaining task is to kill the badness of D_* or $D_{*,i}$. Recall that depending on the pattern of the intersection of α_{g+1} and γ_{g+1} (cf. Figure 20), there are two cases: either we have an octagon bad region D_*, or we have two hexagon bad regions $D_{*,1}$ and $D_{*,2}$.

In the first case, one possibility is that a neighborhood of $\alpha_{g+1} \cup \beta_{g+1} \cup \gamma_{g+1}$ looks as in Figure 23. We then do the finger moves indicated in Figure 24. It is routine to check that the new diagram is isotopic to the one in Figure 23.

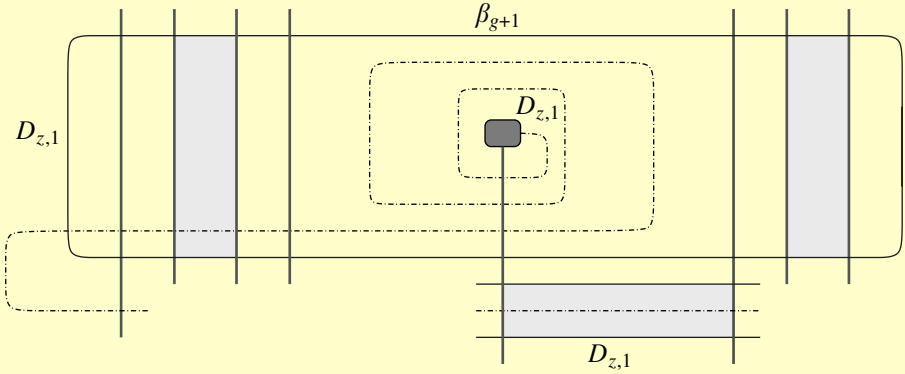

Figure 23. The bad region D_*. Conventions are as before. Lightly shaded regions mean several parallel arcs.

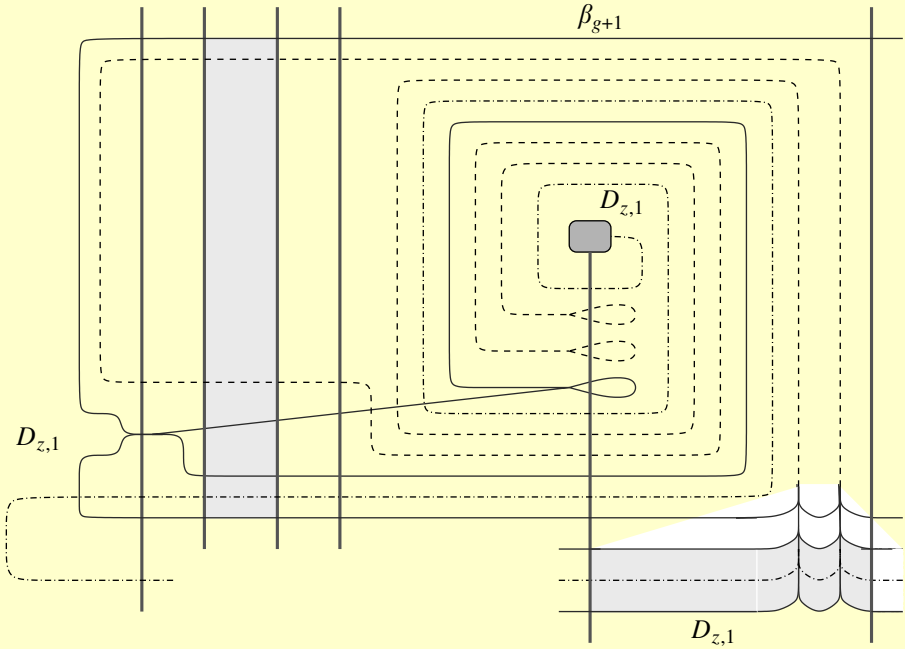

Figure 24. Killing the badness of D_*. We push two multiple fingers (containing several beta and gamma curves) from the bottom right-hand side of the diagram and a single finger from the left-hand side. We are using the train-track convention (cf. Figure 25.)

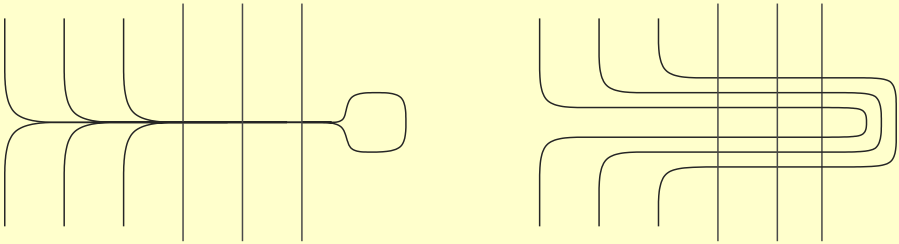

Figure 25. The train-track convention. We use the left-hand side diagram to denote a multiple finger (i.e., the situation pictured on the right-hand side). The curves involved can be of various kinds.

Similarly, in the second case, one possibility is that a neighborhood of $\alpha_{g+1} \cup \beta_{g+1} \cup \gamma_{g+1}$ looks as in Figure 26. In this case, we do the finger moves indicated in Figure 27.

However, the actual picture on the Heegaard diagram may differ from Figures 23 or 26 in several (nonessential) ways.

One possible difference is that at the bottom of Figure 26, the extra gamma curve on top of $D_{z,1}$ may be on the right-hand side rather than on the left-hand side; however, we can still push the two fingers starting from $D_{z,1}$ on each side of α_{g+1}.

Another possible difference is that at the very left-hand side of Figures 23 and 26, the curve γ_{g+1} may have an upward rather than a downward hook (i.e., it may look as in Figure 28(c) rather than Figure 28(a). If so, instead of the beta finger from the left-hand side in Figures 23 and 26 (cf. also Figure 28(b)), we push a beta-gamma finger, as in Figure 28(d).

Finally, instead of the situations shown in Figures 23 and 26, we might have the same pictures reflected in a horizontal axis. If so, we apply similar finger moves and arrive at the reflections of Figures 24 and 27.

In all cases, the finger moves successfully kill the badness of all regions other than $D_{z,1}$, in which we keep the basepoint z. The result is a nice triple Heegaard diagram for the cobordism W.

3.3. Several two-handle additions

We now explain how the arguments in this section can be extended to a cobordism W that consists of the addition of several two-handles. We view W as surgery along a link $L \subset Y$ of l components.

We start with a multipointed Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{z})$, together with another set of base points $\mathbf{w} = \{w_1, \dots, w_l\}$ describing the pair $L \subset Y$, as in [12, Section 3.5]. Each of the two sets of curves ($\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$) has $g + l - 1$ elements.

Applying the algorithm in [15, Section 4], we can make this diagram nice (i.e., such that all regions not containing one of the z's are either bigons or rectangles). For $i = 1, \dots, l$, we denote by D_{z_i} the region containing z_i.

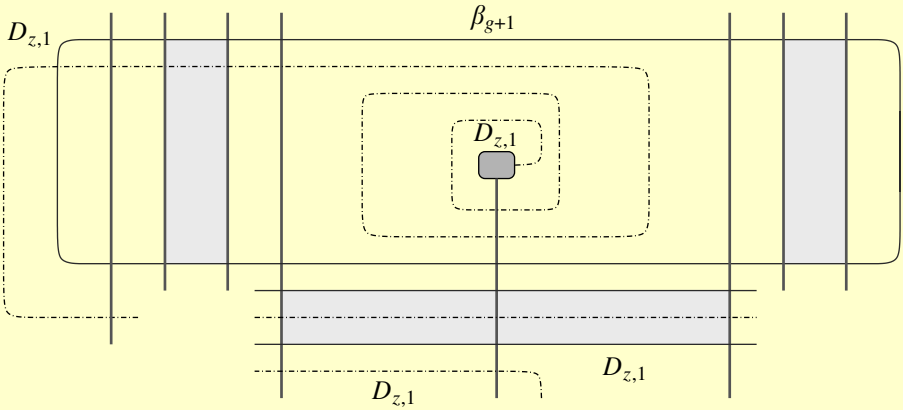

Figure 26. The bad regions $D_{*,1}$ and $D_{*,2}$. Going from each of the two hexagons down through beta and gamma curves, we eventually reach $D_{z,1}$. The path from one of the two hexagons (in the picture, the one on the left-hand side) encounters one additional beta curve before reaching $D_{z,1}$.

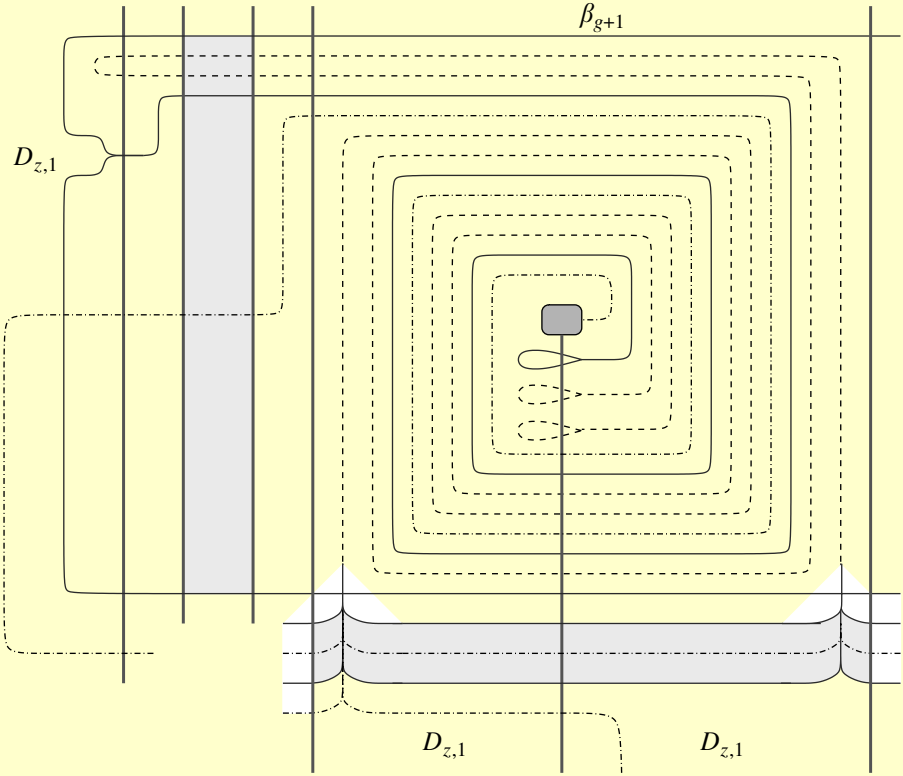

Figure 27. Killing the badness of $D_{*,1}$ and $D_{*,2}$. In this case, we again push two multiple fingers from the bottom but starting from two different sides of α_{g+1}. As before, we also push a finger from the left-hand side.

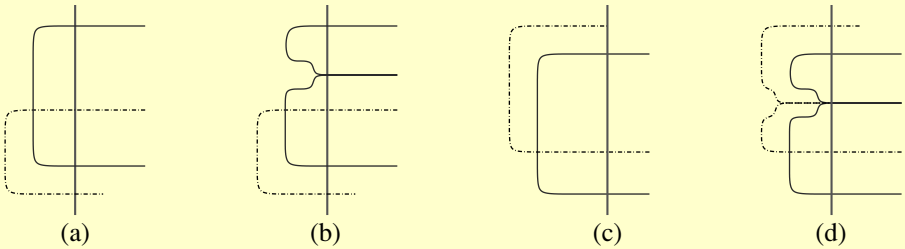

Figure 28. A variation. On the left-hand side of Figures 23 and 26, we may have the picture (c) rather than (a). We then do the finger move in (d) instead of (b). The region on the left-hand side is always $D_{z,1}$.

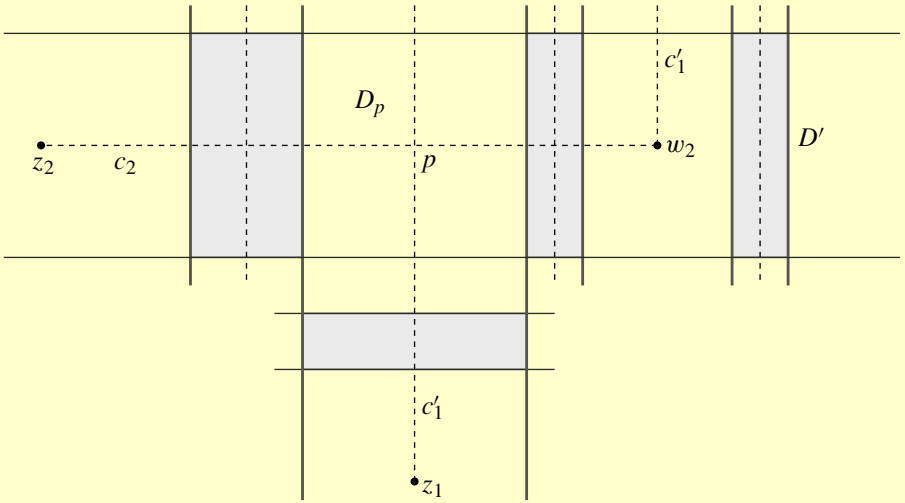

Figure 29. Projection of the link to the Heegaard surface. The point p is the point on the dashed c_1'-curve closest to z_1 where an intersection with any of the c_i-curves (in our case, a c_2-curve) takes place.

As in step 1 of Section 3.2, we inductively remove intersection points between the various components of the projection of L to Σ. This projection consists of arcs c_i and c_i' with end points at z_i and w_i $(i = 1, \dots, l)$, so that each c_i is disjoint from the beta curves, and each c_i' is disjoint from the alpha curves. Instead of Figure 8, we have the situation in Figure 29. Again, we stabilize and perform an isotopy to obtain a good diagram with one fewer intersection point, as in Figure 9. Iterating this process (on all link components), we can assume that the projection of L is embedded in the Heegaard surface.

We then add twin gamma curves, as in step 2 of Section 3.2. For this, we need to do several isotopies of the beta curves, as in Figure 12. In that figure, if the region on the top left-hand side is D_{z_i}, the one on the right-hand side may be D_{z_j} for $j \neq i$; however, the isotopy can be done as before.

Next, we stabilize the Heegaard diagram l times (once for each link component) to obtain a triple diagram for the cobordism, as in step 3 of Section 3.2. We then do the analogue of step 4 by pushing l fingers to kill the badness of the regions of type D'. Since the fingers only pass through rectangles, they do not intersect each other. The only change is that in Figures 17 and 18, the region on the very right-hand side may contain a different puncture z_i than the one on the left-hand side. By contrast, in case 2 of step 4, the region on the right-hand side in Figure 16 contains the same puncture as the region on the left-hand side since they lie in the same connected component of the complement of the beta circles.

At the end of step 4, the beta curves split the Heegaard surface into l connected components $C_1, \dots, C_l$. We then do the analogue of step 5 in Section 3.2. Note that this step (except for the very last bit; see Figure 22) only involves moving gamma curves through alpha curves. (Here, we think of the move in Figure 19 as a single step rather than as a handleslide followed by an isotopy.) Therefore, we can perform the moves in this step once for each connected component of L, independently of each other, because the moves take place in the corresponding component C_i. In the situation considered in Figure 22, the gamma curves cross a beta curve; however, the special region D_s is part of a unique C_i, so we can perform the isotopy of the gamma curves as before, without interference from another C_j.

Finally, for step 6, note that in all the previous steps we have not destroyed the property that the projection of L to the Heegaard surface is embedded. More precisely, in the part of the stabilized Heegaard diagram shown in Figure 14, we take a component of the link projection to be a loop starting in D_*, near the upper foot of the handle, going down along α_{g+1} until it reaches $D_{z,1}$, then going inside $D_{z,1}$ until it reaches the intersection of γ_{g+1} and β_{g+1}, and then going along a subarc of γ_{g+1} to its original departure. These paths remain embedded and disjoint from each other throughout steps 4 and 5. (Indeed, neither β_{g+1} nor this subarc of γ_{g+1} is moved during these steps.) It then suffices to note that the finger moves in step 6 take place in a neighborhood of the projection of the corresponding link component (the path considered above). Therefore, these finger moves can be done without interfering with each other. The result is a nice multipointed triple Heegaard diagram for W.

4. Computing maps induced by cobordisms

Let W be a four-dimensional cobordism from Y_0 to Y_3, and let $\mathfrak{t}$ be a Spin^c-structure on W. Choose a self-indexing Morse function on W. This decomposes W as a

collection of one-handle additions forming a cobordism W_1, followed by some two-handle additions forming a cobordism W_2, and then following by three-handle additions forming a cobordism W_3, in that order. Let Y_1 and Y_2 be the intermediate three-manifolds so that

$$W = W_1 \cup_{Y_1} W_2 \cup_{Y_2} W_3.$$

The map $\hat{F}_{W,\mathfrak{t}} : \widehat{HF}(Y_0, \mathfrak{t}|_{Y_0}) \to \widehat{HF}(Y_3, \mathfrak{t}|_{Y_3})$ from [11, Section 4] is defined as the composition $\hat{F}_{W_3,\mathfrak{t}|_{W_3}} \circ \hat{F}_{W_2,\mathfrak{t}|_{W_2}} \circ \hat{F}_{W_1,\mathfrak{t}|_{W_1}}$ of maps associated to each of the pieces W_1, W_2, and W_3. We review the definitions of these three maps in Sections 4.2 and 4.3. First, we review a few facts about the, $(H_1(Y;\mathbb{Z})/\text{torsion})$-action on the hat Heegaard Floer invariants.

4.1. $(H_1(Y;\mathbb{Z})/torsion)$-actions

In [9, Section 4.2.5], Ozsváth and Szabó constructed an action of the group $H_1(Y;\mathbb{Z})/\text{torsion} \cong \text{Hom}(H^1(Y;\mathbb{Z}),\mathbb{Z})$ on $\widehat{HF}(Y)$. In [11] (see also [7, Section 2]), they also showed that the cobordism maps $\hat{F}_{W,\mathfrak{t}}$ extend to maps $\Lambda^*(H^1(W;\mathbb{Z})/\text{torsion}) \otimes \widehat{HF}(Y_0, \mathfrak{t}|_{Y_0}) \to \widehat{HF}(Y_3, \mathfrak{t}|_{Y_3})$. Moreover, if W is a cobordism from Y_1 to Y_2 (endowed with a Spin^c-structure $\mathfrak{t}$) and we denote by $j_i : H_1(Y_i;\mathbb{Z})/\text{torsion} \to H_1(W;\mathbb{Z})/\text{torsion}\,(i = 1,2)$ the natural inclusions, then for any $\zeta \in H_1(W;\mathbb{Z})/\text{torsion}$ of the form $\zeta = j_1(\zeta_1) - j_2(\zeta_2)$, $\zeta_i \in H_1(Y_i;\mathbb{Z})/\text{torsion}\,(i = 1,2)$, one has

$$(\zeta \otimes \hat{F}_{W,\mathfrak{t}})(x) = \hat{F}_{W,\mathfrak{t}}(\zeta_1 \cdot x) - \zeta_2 \cdot \hat{F}_{W,\mathfrak{t}}(x).$$

This equality has the following immediate corollaries.

LEMMA 4.1
If $\zeta_1 \in \ker(j_1)$, then $\zeta_1 \cdot x \in \ker(\hat{F}_{W,\mathfrak{t}})$ for any $x \in \widehat{HF}(Y_1, \mathfrak{t}|_{Y_1})$.

LEMMA 4.2
If $\zeta_2 \in \ker(j_2)$, then $\zeta_2 \cdot y = 0$ for any $y \in \text{image}(\hat{F}_{W,\mathfrak{t}})$.

Consider now a three-manifold of the form $Y\#(\#^n S^1 \times S^2)$, and let $\mathfrak{s}_0$ denote the torsion Spin^c-structure on $\#^n S^1 \times S^2$. Then, for any Spin^c-structure $\mathfrak{s}$ on Y, there is an isomorphism

$$\widehat{HF}(Y,\mathfrak{s}) \otimes H_*(T^n) \xrightarrow{\cong} \widehat{HF}\big(Y\#(\#^n S^1 \times S^2), \mathfrak{s}\#\mathfrak{s}_0\big) \tag{6}$$

as $\mathbb{F}[(H_1(Y;\mathbb{Z})/\text{torsion}) \oplus H_1(\#^n S^1 \times S^2;\mathbb{Z})]$-modules (cf. [10, Theorem 1.5]). The action of $H_1(\#^n S^1 \times S^2;\mathbb{Z})$ on $\widehat{HF}(Y,\mathfrak{s})$ here is trivial, as is the action of $H_1(Y;\mathbb{Z})/\text{torsion}$ on $\widehat{HF}(\#^n S^1 \times S^2, \mathfrak{s}_0)$. Further, the action of $H_1(\#^n S^1 \times$

$S^2;\mathbb{Z})/\text{torsion} \cong H^1(T^n;\mathbb{Z})$ on $\widehat{HF}(\#^n S^1\times S^2,\mathfrak{s}_0)\cong H_*(T^n;\mathbb{F})$ is exactly that given by the cap product $\bigcap : H^1(T^n;\mathbb{Z})\otimes H_*(T^n;\mathbb{F})\to H_{*-1}(T^n;\mathbb{F})$.

4.2. Maps associated to one- and three-handle additions

Next, we review the definition of the Heegaard Floer maps induced by one- and three-handle additions (cf. [11, Section 4.3]).

Suppose that W_1 is a cobordism from Y_0 to Y_1 built entirely from one-handles. Let $\mathfrak{t}$ be a Spin^c-structure on W_1. The map $\hat{F}_{W_1,\mathfrak{t}} : \widehat{HF}(Y_0,\mathfrak{t}|_{Y_0})\to\widehat{HF}(Y_1,\mathfrak{t}|_{Y_1})$ is constructed as follows. If $h_1,h_2,\dots,h_n$ are the one-handles in the cobordism, for each $i=1,\dots,n$, pick a path ξ_i in Y_0, joining the two feet of the handle h_i. This induces a connected sum decomposition $Y_1\cong Y_0\#(\#^n S^1\times S^2)$, where the first homology of each $(S^1\times S^2)$-summand is generated by the union of ξ_i with the core of the corresponding handle. Further, the restriction of $\mathfrak{t}$ to the $(S^1\times S^2)$-summands in Y_1 is torsion. It follows that $\widehat{HF}(Y_1,\mathfrak{t}|_{Y_1})\cong\widehat{HF}(Y_0,\mathfrak{t}|_{Y_0})\otimes H_*(T^n)$. Let θ be the generator of the top-graded part of $H_*(T^n)$. Then the Heegaard Floer map induced by W_1 is given by

$$\hat{F}_{W_1,\mathfrak{t}}(x)=x\otimes\theta.$$

It is proved in [11, Lemma 4.13] that up to composition with canonical isomorphisms, $\hat{F}_{W_1,\mathfrak{t}}$ does not depend on the choices made in its construction, such as the choice of the paths ξ_i.

Dually, suppose that W_3 is a cobordism from Y_2 to Y_3 built entirely from three-handles. Let $\mathfrak{t}$ be a Spin^c-structure on W_3. The map $\hat{F}_{W_3,\mathfrak{t}} : \widehat{HF}(Y_2,\mathfrak{t}|_{Y_2})\to\widehat{HF}(Y_3,\mathfrak{t}|_{Y_3})$ is constructed as follows. One can reverse W_3 and view it as attaching one-handles on Y_3 to get Y_2. After choosing paths between the feet of these one-handles in Y_3, we obtain a decomposition $Y_2\cong Y_3\#(\#^m S^1\times S^2)$ (where m is the number of three-handles of W_3). Further, the restriction of $\mathfrak{t}$ to the $(S^1\times S^2)$-summands in Y_2 is torsion. It follows that $\widehat{HF}(Y_2,\mathfrak{t}|_{Y_2})\cong\widehat{HF}(Y_3,\mathfrak{t}|_{Y_3})\otimes H_*(T^m)$. Let η be the generator of the lowest-graded part of $H_*(T^m)$. Then the Heegaard Floer map induced by W_3 is given by

$$\hat{F}_{W_3,\mathfrak{t}}(x\otimes\eta)=x$$

and

$$\hat{F}_{W_3,\mathfrak{t}}(x\otimes\omega)=0$$

for any homogeneous generator ω of $H_*(T^m)$ not lying in the minimal degree. Again, the map is independent of the choices made in its construction.

4.3. Maps associated to two-handle additions

Let W_2 be a two-handle cobordism from Y_1 to Y_2, corresponding to surgery on a framed link L in Y_1, and let $\mathfrak{t}$ be a Spin^c-structure on W_2. Let $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', z')$ be a triple Heegaard diagram subordinate to a bouquet $B(L)$ for L, as in the beginning of Section 3. Then, in particular, $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', z')$ is a Heegaard diagram for Y_1, $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\gamma}', z')$ is a Heegaard diagram for Y_2, and $(\Sigma', \boldsymbol{\beta}', \boldsymbol{\gamma}', z')$ is a Heegaard diagram for $\#^{g-l}(S^1 \times S^2)$. (Here, g is the genus of Σ', and l is the number of components of L.) The Spin^c-structure $\mathfrak{t}$ induces a Spin^c-structure (still denoted $\mathfrak{t}$) on the four-manifold $W_{\alpha',\beta',\gamma'}$ specified by $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', z')$. (Note that $W_{\alpha',\beta',\gamma'}$ can be viewed as a subset of W_2.) Consequently, there is an induced map

$$\hat{F}_{\Sigma',\boldsymbol{\alpha}',\boldsymbol{\beta}',\boldsymbol{\gamma}',z',\mathfrak{t}} : \widehat{HF}(Y_1, \mathfrak{t}|_{Y_1}) \otimes \widehat{HF}\left(\#^{g-l}(S^1 \times S^2), \mathfrak{t}|_{\#^{g-l}(S^1\times S^2)}\right) \to \widehat{HF}(Y_2, \mathfrak{t}|_{Y_2}),$$

as discussed in Section 2.1.

The Spin^c-structure $\mathfrak{t}|_{\#^{g-l}(S^1\times S^2)}$ is necessarily torsion, so

$$\widehat{HF}\left(\#^{g-l}(S^1 \times S^2), \mathfrak{t}|_{\#^{g-l}(S^1\times S^2)}\right) \cong H_*(T^{g-l}) = H_*(S^1)^{\otimes(g-l)}.$$

Let θ denote the generator for the top-dimensional part of $\widehat{HF}(\#^{g-l}(S^1 \times S^2), \mathfrak{t}|_{\#^{g-l}(S^1\times S^2)})$. Then we define (cf. [11, Section 4.1]) the map

$$\hat{F}_{W_2,\mathfrak{t}} : \widehat{HF}(Y_1, \mathfrak{t}|_{Y_1}) \to \widehat{HF}(Y_2, \mathfrak{t}|_{Y_2})$$

by $\hat{F}_{W_2,\mathfrak{t}}(x) = \hat{F}_{\Sigma',\boldsymbol{\alpha}',\boldsymbol{\beta}',\boldsymbol{\gamma}',z',\mathfrak{t}}(x \otimes \theta)$.

Now, consider instead the nice, l-pointed triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ constructed in Section 3. As discussed in the beginning of Section 3, $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ is strongly equivalent to a split triple Heegaard diagram whose reduction $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', z')$ is subordinate to a bouquet $B(L)$ as above. Let Θ be the generator for the top-dimensional part of $\widehat{HF}(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z}) \cong H_*(T^{g+l-1})$. Then, by diagram (3), with $k = l - 1$, we have

$$\operatorname{rank}(\hat{F}_{W_2,\mathfrak{t}}) = \frac{1}{2^{l-1}} \operatorname{rank}\left(\hat{F}_{\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\mathbf{z},\mathfrak{t}}(\cdot \otimes \Theta)\right).$$

In light of Corollary 3.2 and Proposition 2.5, the rank of the map $\hat{F}_{\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\mathbf{z},\mathfrak{t}}$ can be computed combinatorially. Further, since the triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ is nice, so are each of the three (ordinary) Heegaard diagrams that it specifies. Consequently, by [15, Theorem 1.1], the element $\Theta \in \widehat{HF}(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ can be explicitly identified (as can a representative for Θ in $\widehat{CF}(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$). Therefore, the rank of $\hat{F}_{W_2,\mathfrak{t}}$ can be computed combinatorially.

4.4. Putting it all together

Recall that W is a four-dimensional cobordism from Y_0 to Y_3, that $\mathfrak{t}$ is a Spin^c-structure on W, and that W is decomposed as a collection of one-handle additions W_1, followed by some two-handle additions W_2 and three-handle additions W_3, with Y_1 and Y_2 the intermediate three-manifolds, so that

$$W = W_1 \cup_{Y_1} W_2 \cup_{Y_2} W_3.$$

As in Section 4.1, we consider the maps

$$\begin{aligned} j_1 &: H_1(Y_0;\mathbb{Z})/\text{torsion} \to H_1(W;\mathbb{Z})/\text{torsion}, \\ j_2 &: H_1(Y_3;\mathbb{Z})/\text{torsion} \to H_1(W;\mathbb{Z})/\text{torsion}. \end{aligned} \tag{7}$$

LEMMA 4.3

If j_1 is surjective, then $\text{image}(\hat{F}_{W_2,\mathfrak{t}|_{W_2}} \circ \hat{F}_{W_1,\mathfrak{t}|_{W_1}}) = \text{image}(\hat{F}_{W_2,\mathfrak{t}|_{W_2}})$.

Proof

The cobordism W_1 consists of the addition of some one-handles $h_1, \dots, h_n$. As in Section 4.2, we choose paths ξ_i in Y_0 joining the two feet of the handle h_i. The union of ξ_i with the core of h_i produces a curve in W_1, which in turn gives an element $e_i \in H_1(W_1 \cup W_2;\mathbb{Z})/\text{torsion}$. Since W is obtained from $W_1 \cup W_2$ by adding three-handles, we have $H_1(W;\mathbb{Z}) \cong H_1(W_1 \cup W_2;\mathbb{Z})$, so the hypothesis implies that the map

$$j_1' : H_1(Y_0;\mathbb{Z})/\text{torsion} \to H_1(W_1 \cup W_2;\mathbb{Z})/\text{torsion}$$

is surjective. Hence, there exist disjoint, embedded curves c_i in Y_0 (disjoint from all the ξ_j) such that $j_1'([c_i]) = -e_i, i = 1, \dots, n$. We can connect sums ξ_i and c_i to get new paths ξ_i' in Y_0 between the two feet of h_i. Using the paths ξ_i', we get a connected sum decomposition $Y_1 \cong Y_0 \# (\#^n S^1 \times S^2)$, as in Section 4.2, with the property that the inclusion of the summand $H_1(\#^n S^1 \times S^2;\mathbb{Z}) \subset H_1(Y_1;\mathbb{Z})$ in $H_1(W_1 \cup W_2;\mathbb{Z})/\text{torsion}$ is trivial. Since we can view $W_1 \cup W_2$ as obtained from W_2 by adding three-handles (which do not affect H_1), it follows that the inclusion of $H_1(\#^n S^1 \times S^2;\mathbb{Z})$ in $H_1(W_2;\mathbb{Z})/\text{torsion}$ is trivial. Lemma 4.1 then says that

$$\hat{F}_{W_2,\mathfrak{t}|_{W_2}}(\zeta \cdot x) = 0$$

for any $\zeta \in H_1(\#^n S^1 \times S^2;\mathbb{Z})$, $x \in \widehat{HF}(Y_1, \mathfrak{t}|_{Y_1}) \cong \widehat{HF}(Y_0, \mathfrak{t}|_{Y_0}) \otimes H_*(T^n)$. Thus, the kernel of $\hat{F}_{W_2,\mathfrak{t}|_{W_2}}$ contains all elements of the form $y \otimes \omega$, where $y \in \widehat{HF}(Y_0, \mathfrak{t}|_{Y_0})$ and $\omega \in H_*(T^n)$ is any homogeneous element not lying in the top grading of $H_*(T^n)$. On the other hand, we know from Section 4.2 that the image of $\hat{F}_{W_1,\mathfrak{t}|_{W_1}}$ consists exactly

of the elements $y \otimes \theta$, where θ is the top-degree generator of $H_*(T^n)$. Therefore,

$$\text{image}(\hat{F}_{W_1,\mathfrak{t}|_{W_1}}) + \ker(\hat{F}_{W_2,\mathfrak{t}|_{W_2}}) = \widehat{HF}(Y_1, \mathfrak{t}|_{Y_1}).$$

This gives the desired result. □

LEMMA 4.4
If j_2 is surjective, then $\text{image}(\hat{F}_{W_2,\mathfrak{t}|_{W_2}}) \cap \ker(\hat{F}_{W_3,\mathfrak{t}|_{W_3}}) = 0$.

Proof
This is similar to the proof of Lemma 4.3. A suitable choice of paths enables us to view Y_2 as $Y_3\#(\#^m S^1 \times S^2)$, so that the inclusion of the summand $H_1(\#^m S^1 \times S^2; \mathbb{Z}) \subset H_1(Y_2; \mathbb{Z})$ in $H_1(W_2; \mathbb{Z})/\text{torsion}$ is trivial. Lemma 4.2 then says that $\zeta \cdot y = 0$ for any $y \in \text{image}(\hat{F}_{W_2,\mathfrak{t}|_{W_2}})$ and $\zeta \in H_1(\#^m S^1 \times S^2)$. In other words, every element in the image of $\hat{F}_{W_2,\mathfrak{t}|_{W_2}}$ must be of the form $y = x \otimes \eta$, where $x \in \widehat{HF}(Y_3, \mathfrak{t}|_{Y_3})$ and η is the lowest-degree generator of $H_*(T^n)$. On the other hand, we know from Section 4.2 that the kernel of the map $\hat{F}_{W_3,\mathfrak{t}|_{W_3}}$ does not contain any nonzero elements of the form $x \otimes \eta$. □

THEOREM 4.5
Let W be a cobordism from Y_0 to Y_3, and let $\mathfrak{t}$ be a Spin^c*-structure on W. Assume that the maps j_1 and j_2 from formula (8) are surjective. Then in each (relative) grading i, the rank of $\hat{F}_{W,\mathfrak{t}} : \widehat{HF}_i(Y_0, \mathfrak{t}|_{Y_0}) \to \widehat{HF}_*(Y_3, \mathfrak{t}|_{Y_3})$ can be computed combinatorially.*

Proof
The map $\hat{F}_{W,\mathfrak{t}}$ is, by definition, the composition $\hat{F}_{W_3,\mathfrak{t}|_{W_3}} \circ \hat{F}_{W_2,\mathfrak{t}|_{W_2}} \circ \hat{F}_{W_1,\mathfrak{t}|_{W_1}}$. Lemmas 4.3 and 4.4 imply that

$$\text{image}(\hat{F}_{W_2,\mathfrak{t}|_{W_2}} \circ \hat{F}_{W_1,\mathfrak{t}|_{W_1}}) \cap \ker(\hat{F}_{W_3,\mathfrak{t}|_{W_3}}) = 0$$

or, equivalently,

$$\text{rank}(\hat{F}_{W_3,\mathfrak{t}|_{W_3}} \circ \hat{F}_{W_2,\mathfrak{t}|_{W_2}} \circ \hat{F}_{W_1,\mathfrak{t}|_{W_1}}) = \text{rank}(\hat{F}_{W_2,\mathfrak{t}|_{W_2}} \circ \hat{F}_{W_1,\mathfrak{t}|_{W_1}}).$$

Using Lemma 4.3 again, the expression on the right-hand side is the same as the rank of $\hat{F}_{W_2,\mathfrak{t}|_{W_2}}$. Thus, the maps $\hat{F}_{W,\mathfrak{t}}$ and $\hat{F}_{W_2,\mathfrak{t}|_{W_2}}$ have the same rank. As explained in Section 4.3, the rank of $\hat{F}_{W_2,\mathfrak{t}|_{W_2}}$ can be computed combinatorially. Note that the relative gradings on the generators of the chain complexes are also combinatorial, using the formula for the Maslov index in [3, Corollary 4.3]. This completes the proof. □

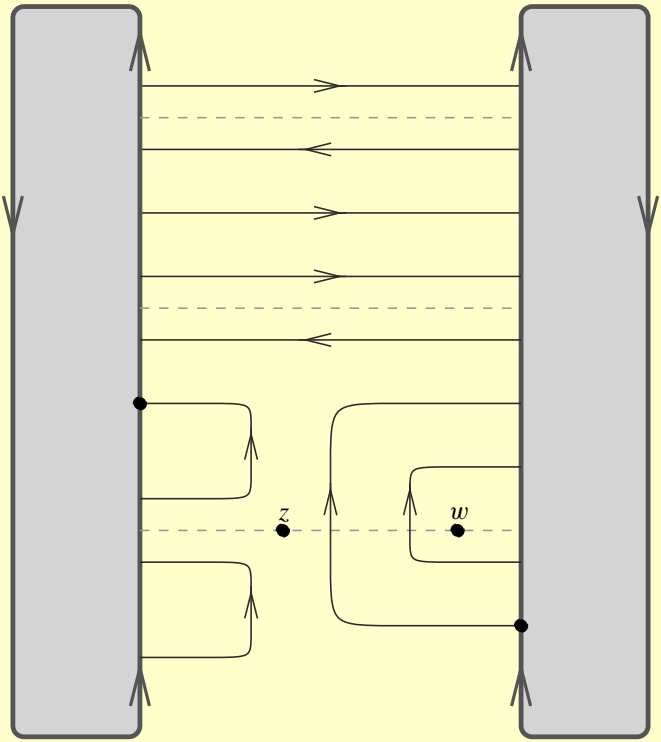

Figure 30. A nice Heegaard diagram for the trefoil knot. The thick curves are alpha curves, and the thin ones are beta curves. The two shaded areas in the diagram are glued together via a reflection and rotation, so that the two black dots on the alpha curves are identified. The knot is given by the dashed curves. Its projection is already embedded in the Heegaard surface.

Remark 4.6

In fact, using Sarkar's remarkable formula for the Maslov index of triangles, [14, Theorem 4.1], the absolute gradings on the Heegaard Floer complexes can be computed combinatorially, and so the rank of $\hat{F}_{W,\mathfrak{t}}$ in each absolute grading can be computed as well.

5. An example

We give a nice triple Heegaard diagram for the cobordism from the three-sphere to the Poincaré homology sphere, viewed as the $+1$-surgery on the right-handed trefoil. The right-handed trefoil knot admits the nice Heegaard diagram shown in Figure 30, which is isotopic to [15, Figure 14]. Applying the algorithm described in Section 3, we obtain the nice triple Heegaard diagram shown in Figure 31. We leave the actual computation of the cobordism map to the interested reader.

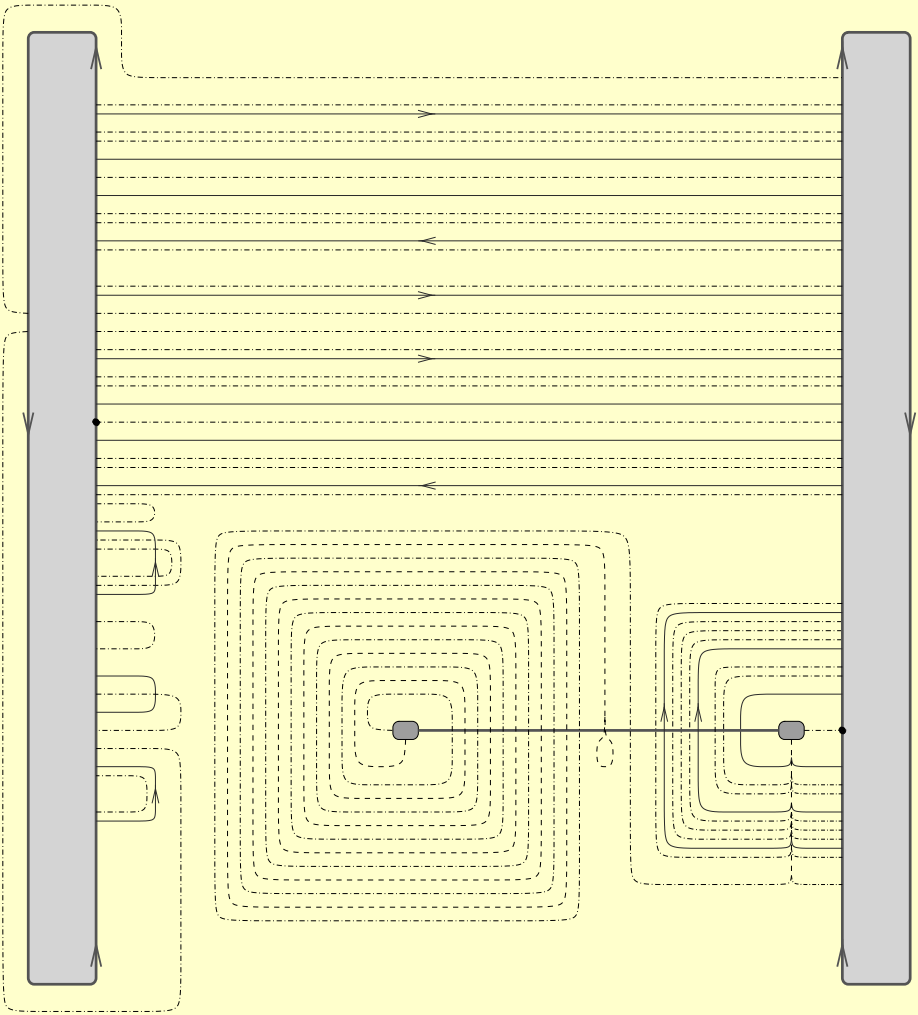

Figure 31. A nice triple Heegaard diagram for $+1$-surgery on the right-handed trefoil. The thick curves are alpha curves, the thin curves are beta curves, and the interrupted curves are gamma curves. The lightly shaded areas are identified as in Figure 30, while the darkly shaded areas represent the new handle that has been attached. Again, we are using the traintrack convention, as in Figure 25. The base point is the point at infinity.

Acknowledgments. We thank Peter Ozsváth and Zoltán Szabó for helpful conversations and encouragement. In particular, several key ideas in the proof were suggested to us by Ozsváth.

This work was done while Jiajun Wang was an exchange graduate student at Columbia University; he is grateful to that university's mathematics department for

its hospitality. He also thanks his advisors, Robion Kirby and Peter Ozsváth, for their continuous guidance and support.

Finally, we thank the referees for many helpful comments, particularly for finding a critical error in Section 4 of a previous version of this article.

References

[1] R. FINTUSHEL and R. STERN, *Rational blowdowns of smooth* 4*-manifolds*, J. Differential Geom. **46** (1997), 181 – 235. MR 1484044 208

[2] R. E. GOMPF and A. I. STIPSICZ, 4*-manifolds and Kirby Calculus*, Grad. Stud. Math. **20**, Amer. Math. Soc., Providence, 1999. MR 1707327 208

[3] R. LIPSHITZ, *A cylindrical reformulation of Heegaard Floer homology*, Geom. Topol. **10** (2006), 955 – 1097. MR 2240908 212, 213, 214, 243

[4] P. LISCA, *On the Donaldson polynomials of elliptic surfaces*, Math. Ann. **299** (1994), 629 – 639. MR 1286889 208

[5] C. MANOLESCU, P. OZSVÁTH, and S. SARKAR, *A combinatorial description of knot Floer homology*, preprint, arXiv:math/0607691v2 [math.GT] 208

[6] J. W. MORGAN and T. S. MROWKA, *On the diffeomorphism classification of regular elliptic surfaces*, Internat. Math. Res. Notices **1993**, no. 6, 183 – 184. MR 1224116 208

[7] P. S. OZSVÁTH and Z. SZABÓ, *Absolutely graded Floer homologies and intersection forms for four-manifolds with boundary*, Adv. Math. **173** (2003), 179 – 261. MR 1957829 208, 239

[8] ———, *Holomorphic disks and knot invariants*, Adv. Math. **186** (2004), 58 – 116. MR 2065507 223

[9] ———, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2) **159** (2004), 1027 – 1158. MR 2113019 207, 208, 209, 210, 211, 214, 223, 239

[10] ———, *Holomorphic disks and three-manifold invariants: Properties and applications*, Ann. of Math. (2) **159** (2004), 1159 – 1245. MR 2113020 239

[11] ———, *Holomorphic triangles and invariants for smooth four-manifolds*, Adv. Math. **202** (2006), 326 – 400. MR 2222356 207, 208, 209, 220, 239, 240, 241

[12] ———, *Holomorphic disks and link invariants*, preprint, arXiv:math/0512286v2 [math.GT] 210, 211, 212, 221, 235

[13] L. ROBERTS, *Rational blow-downs in Heegaard-Floer homology*, preprint, arXiv:math/0607675v1 [math.GT] 208

[14] S. SARKAR, *Maslov index of holomorphic triangles*, preprint, arXiv:math/0609673v2 [math.GT] 209, 214, 244

[15] S. SARKAR and J. WANG, *An algorithm for computing some Heegaard Floer homologies*, preprint, arXiv:math/0607777v3 [math.GT] 207, 208, 214, 223, 232, 233, 235, 241, 244

[16] A. I. STIPSICZ and Z. SZABÓ, *Gluing* 4*-manifolds along* $\Sigma(2, 3, 11)$, Topology Appl. **106** (2000), 293 – 304. MR 1775711 208

Lipshitz
Department of Mathematics, Columbia University, New York, New York 10027, USA;
lipshitz@math.columbia.edu

Manolescu
Department of Mathematics, University of California, Los Angeles, Los Angeles, California 90095, USA; cm@math.ucla.edu

Wang
Department of Mathematics, California Institute of Technology, Pasadena, California 91125, USA; wjiajun@caltech.edu