msp

Geometry & Topology 23 (2019) 1491–1547

A deformation of instanton homology for webs

PETER B KRONHEIMER
TOMASZ S MROWKA

A deformation of the authors' instanton homology for webs is constructed by introducing a local system of coefficients. In the case that the web is planar, the rank of the deformed instanton homology is equal to the number of Tait colorings of the web.

57R58; 05C15

1 Introduction

1.1 Statement of results

In an earlier pair of papers [14; 13], the authors studied an $SO(3)$ instanton homology for "webs" (embedded trivalent graphs) in closed, oriented 3–manifolds. In particular, to a web $K \subset \mathbb{R}^3$, the authors associated a vector space $J^\sharp(K)$ over the field $\mathbb{F} = \mathbb{Z}/2$. One of the reasons for being interested in $J^\sharp$ is that, in conjunction with the other results of [14], the following conjecture implies the four-color theorem; see Appel and Haken [1].

Conjecture 1.1 [14] *If K lies in the plane $\mathbb{R}^2 \subset \mathbb{R}^3$, then the dimension of $J^\sharp(K)$ is equal to the number of Tait colorings of K.*

The homology $J^\sharp(K)$ is constructed from the Morse theory of the Chern–Simons functional on a space of connections $\mathcal{B}$ associated with K. In this paper, we introduce a system of local coefficients Γ on $\mathcal{B}$ and use it to define a variant, $J^\sharp(K;\Gamma)$, which is a module over the ring $R = \mathbb{F}[\mathbb{Z}^3]$ (elements of which we write as finite Laurent series in variables T_1, T_2, T_3). The property that is conjectural for $J^\sharp(K)$ is a theorem for its deformation $J^\sharp(K;\Gamma)$.

Theorem 1.2 *If K lies in the plane, then the rank of $J^\sharp(K;\Gamma)$ as an R–module is equal to the number of Tait colorings of K.*

Published: 28 May 2019 DOI: 10.2140/gt.2019.23.1491

From the construction of $J^\sharp(K;\Gamma)$ as a Morse homology with coefficients in a local system, it is apparent that it is the homology of a complex (C,∂) of free R–modules. Furthermore, the original $J^\sharp(K)$ can be recovered as the homology of the complex $(C\otimes_R \mathbb{F}, \partial\otimes 1)$, where $\mathbb{F}$ is made into an R–module by evaluation at $T_i=1$. From this it follows that there is (for all K) an inequality

$$\dim_{\mathbb{F}} J^\sharp(K) \geq \operatorname{rank}_R J^\sharp(K;\Gamma), \tag{1}$$

because the rank of the differential $\partial\otimes 1$ cannot be larger than the rank of ∂. From the theorem above, we now obtain:

Corollary 1.3 *If K lies in the plane, then the dimension of $J^\sharp(K)$ is greater than or equal to the number of Tait colorings of K.*

Although this corollary is "half" of Conjecture 1.1, it does not have any implication for the four-color theorem. It does, however, put Conjecture 1.1 into perspective. The inequality (1) can be refined by constructing a spectral sequence, which allows us to interpret Conjecture 1.1 as saying that, in the planar case, a certain spectral sequence collapses. For nonplanar webs, the spectral sequence may not collapse, and the inequality (1) can be strict. Indeed, in Section 6.4, we give an example of a web in $\mathbb{R}^3$ for which $J^\sharp(K;\Gamma)$ has rank 0 while $J^\sharp(K)$ is nonzero. Whether or not the web is planar, the differentials d_1, d_2 and d_3 in our spectral sequence are always zero (as a consequence of a nontrivial calculation, Proposition 6.13), so the first interesting question arises with d_4.

1.2 Ingredients of the proof

The content of Theorem 1.2 is that the deformed instanton homology can be effectively computed when K is planar. As an introduction to why this is so, recall that for each edge e of K there is an operator u_e defined on $J^\sharp(K)$ [14] which satisfies $u_e^3=0$. We will see that there is a similar operator on the deformed instanton homology $J^\sharp(K;\Gamma)$ and that in the deformed case there is a relation

$$u_e^3 + Pu_e = 0, \tag{2}$$

where $P\in R$ is the nonzero element

$$P = T_1T_2T_3 + T_1T_2^{-1}T_3^{-1} + T_2T_3^{-1}T_1^{-1} + T_3T_1^{-1}T_2^{-1}. \tag{3}$$

The polynomial satisfied by u_e therefore has two roots 0 and $P^{1/2}$, the second of which is repeated (because we are in characteristic 2). After replacing R by its field of fractions and adjoining $P^{1/2}$, we may therefore decompose $J^\sharp(K;\Gamma)$ into the generalized eigenspaces of the operator u_e. As e runs through all edges, we obtain a collection of commuting operators whose generalized eigenspaces provide a finer decomposition of the instanton homology. This eigenspace decomposition leads to Tait colorings in the planar case.

The proof of the relation (2), and the calculation of P, involves an explicit understanding of some small instanton moduli spaces. The particular form of P is also central to the proof that the differentials d_1, d_2 and d_3 are zero in the spectral sequence (Proposition 6.13), the essential point being that P vanishes to order 4 at the point $(1,1,1)$.

1.3 Remarks

We have restricted our exposition in this introduction to the case that K lies in $\mathbb{R}^3$ or S^3. But the construction of $J^\sharp(K;\Gamma)$, and the spectral sequence which leads to the inequality (1), both extend without modification to the more general case of webs in a closed, oriented 3–manifold Y. We will develop the construction with this greater generality, returning to the case of $\mathbb{R}^3$ for the proof of Theorem 1.2.

Acknowledgement The work of Kronheimer was supported by the National Science Foundation through NSF grants DMS-1405652 and DMS-1707924. The work of Mrowka was supported by NSF grant DMS-1406348, and by a grant from the Simons Foundation, grant number 503559 TSM.

2 An equivalent construction of $J^\sharp(K)$

2.1 Summary of the original construction

The basic objects of study in [14] are closed, oriented, three-dimensional orbifolds $\check{Y}$ whose local isotropy groups are all either $\mathbb{Z}/2$ or $\mathbb{Z}/2\times\mathbb{Z}/2$. We call such an orbifold a *bifold*. The underlying topological space $|\check{Y}|$ is a 3–manifold, and the singular set is an embedded trivalent graph, or web, $K\subset|\check{Y}|$. In the other direction, given an oriented 3–manifold Y and a web $K\subset Y$, there is a corresponding bifold, which we may denote simply by (Y,K).

By a *bifold connection* (E, A) over $\check{Y}$, we mean an orbifold vector bundle E with fiber $\mathbb{R}^3$ and structure group $SO(3)$, with a connection A, such that the action of the isotropy groups on the fibers of E at the singular points is nontrivial. *Marking data* μ on $\check{Y}$ consists of an open set U_μ and an $SO(3)$ bundle $E_\mu \to U_\mu \setminus K$. A bifold bundle E is *marked* by μ if there is given an isomorphism $\sigma\colon E_\mu \to E|_{U_\mu}$. An *isomorphism* τ between μ–marked bundles with connection, (E, A, σ) and (E', A', σ') is an isomorphism of bifold bundles with connection such that the automorphism $\sigma^{-1}\tau\sigma'\colon E_\mu \to E_\mu$ lifts to the determinant-1 gauge group. The marking data μ is *strong* if the automorphism group of every μ–marked bifold connection is trivial.

Two simple examples of strong marking data are highlighted in [14]. The first is on the bifold (S^3, H) where H is a Hopf link. In this example, the marking data μ_H has U_μ a ball containing H and E_μ is a bundle with $w_2(E_\mu)$ nonzero on the boundary of the tubular neighborhood of either component of H. The second example is (S^3, θ) where θ is a standardly embedded theta graph: two vertices joined by three arcs lying in a plane $\mathbb{R}^2 \subset S^3$. (This orbifold is the global quotient of S^3 by an action of the Klein 4–group.) In this example, the marking data μ_θ again has U_μ a ball containing θ, and E_μ is the trivial bundle.

Given strong marking data μ on $\check{Y}$, there is a Banach manifold $\mathcal{B}_l(\check{Y}; \mu)$ parametrizing isomorphism classes of marked bundles with connection, of Sobolev class L^2_l. In this case, one may define an instanton homology group $J(\check{Y}; \mu)$ as the Morse homology of a perturbed Chern–Simons functional on $\mathcal{B}_l(\check{Y}; \mu)$ with coefficients in the prime field $\mathbb{F}$ of characteristic 2. The two examples in the previous paragraph both have one-dimensional instanton homology:

$$J((S^3, H); \mu_H) = \mathbb{F},$$
$$J((S^3, \theta); \mu_\theta) = \mathbb{F}.$$

In [14], the authors defined $J^\sharp(\check{Y})$ for an arbitrary bifold $\check{Y}$ without marking data by forming a connected sum with (S^3, H), with its own marking. That is,

$$J^\sharp(\check{Y}) := J(\check{Y} \# (S^3, H); \mu_H).$$

Here, when writing a connected sum, we mean that the connected sum is formed at nonsingular points of the two bifolds. The definition is valid, because the marking data μ_H is strong on the connected sum. In the case that $\check{Y}$ arises from a web $K \subset \mathbb{R}^3 \subset S^3$, we simply write $J^\sharp(K)$.

In the general case, to make $J^\sharp$ natural, we should take $\check{Y}$ to be a closed bifold with a framed basepoint in the nonsingular part, in order to form the connected sum. Having done this, $J^\sharp$ becomes a functor from a category in which a morphism is a 4–dimensional bifold cobordism with a framed arc joining the basepoints at the two ends of the cobordism. As in dimension three, the underlying topological space $|\check{X}|$ of a 4–dimensional bifold is a 4–manifold, and the singular set is a restricted type of two-dimensional subcomplex called a *foam*. The space of morphisms in the category can be extended by allowing the foams to carry extra data in the form of *dots* on the faces of the foam. See [14].

2.2 Replacing the Hopf link with the theta graph

While this was not pursued in [14], one could form a connected sum with (S^3, θ) instead of (S^3, H) above. Let us temporarily write

$$J^\dagger(\check{Y}) := J(\check{Y} \# (S^3, \theta); \mu_\theta). \tag{4}$$

We shall show that $J^\dagger(\check{Y})$ and $J^\sharp(\check{Y})$ are isomorphic for all $\check{Y}$. We then have:

Proposition 2.1 *On the cobordism category of bifolds with framed basepoint, the two functors $J^\dagger$ and $J^\sharp$ are naturally isomorphic.*

Proof The proof is an application of an excision principle, stated in two variants as Propositions 4.1 and 4.2 in [14]. As an application of the first version, a connected sum theorem is given as Proposition 4.3 in [14], and we it restate here as an isomorphism

$$J(\check{Y}_1 \# \check{Y}_2 \# (S^3, H); \mu) = J(\check{Y}_1 \# (S^3, H); \mu) \otimes J(\check{Y}_2 \# (S^3, H); \mu). \tag{5}$$

In the version described in [14], the marking μ was μ_H throughout. But we could also take additional marking data μ_i on $\check{Y}_i$ for $i = 1, 2$. So, on the left of the above isomorphism we could take

$$\mu = \mu_1 \cup \mu_2 \cup \mu_H$$

and so on. While this connected sum theorem is an application of Proposition 4.1 of [14] and involves the Hopf link H, one can similarly use the other version, Proposition 4.2, and replace H with θ. Thus, we also have

$$J(\check{Y}_1 \# \check{Y}_2 \# (S^3, \theta); \mu) = J(\check{Y}_1 \# (S^3, \theta); \mu) \otimes J(\check{Y}_2 \# (S^3, \theta); \mu). \tag{6}$$

We now apply (5) with $\check{Y}_1 = \check{Y}$ and $\check{Y}_2 = (S^3, \theta)$. For marking data, we take $\mu = \mu_2 \cup \mu_H$, where $\mu_2 = \mu_\theta$. From this we obtain

$$J(\check{Y} \# (S^3, \theta) \# (S^3, H); \mu_\theta \cup \mu_H) = J^\sharp(\check{Y}) \otimes J((S^3, \theta) \# (S^3, H); \mu_\theta \cup \mu_H).$$

In a similar way, from (6), we obtain

$$J(\check{Y} \# (S^3, \theta) \# (S^3, H); \mu_\theta \cup \mu_H) = J^\dagger(\check{Y}) \otimes J((S^3, \theta) \# (S^3, H); \mu_\theta \cup \mu_H).$$

Combining the last two isomorphisms, we have an isomorphism of finite-dimensional $\mathbb{F}$–vector spaces of the form

$$J^\sharp(\check{Y}) \otimes I = J^\dagger(\check{Y}) \otimes I, \tag{7}$$

where $I = J(S^3, H \cup \theta; \mu)$. In the notation of [14], the vector space I is $I^\sharp(\theta)$, which is nonzero by the results of Section 7.1 of [14]. It follows that $J^\sharp(\check{Y})$ and $J^\dagger(\check{Y})$ are isomorphic.

The excision isomorphisms in [14] are natural with respect to bifold cobordisms, and it follows that the isomorphism (7) expresses a functorial isomorphism between $J^\sharp$ and $J^\dagger$. □

From this point on, we will use θ rather than H in the construction of $J^\sharp$. That is, we will drop the notation $J^\dagger$ and take (4) as the *definition* of $J^\sharp$. For brevity, we will sometimes write $\mathcal{B}^\sharp(\check{Y})$ for the relevant space of marked bifold connections,

$$\mathcal{B}^\sharp(\check{Y}) = \mathcal{B}(\check{Y} \# (S^3, \theta); \mu_\theta). \tag{8}$$

3 Instanton homology with local coefficients

3.1 Maps to the circle

Consider again the standard theta-graph $\theta \subset S^3$. The corresponding orbifold (S^3, θ) is a quotient of a standard 3–sphere $\widehat{S}^3$ by the Klein 4–group V_4. Take the marking data μ_θ to have $U_\mu = S^3$ with E_μ trivial. (This is equivalent to the previous description, in which U_μ was a ball in S^3 containing θ.) The space of marked connections $\mathcal{B}_\theta = \mathcal{B}((S^3, \theta); \mu)$ parametrizes bifold $SO(3)$ connections (E, A) equipped with a homotopy class of trivializations of E on the nonsingular part, $S^3 \setminus \theta$, or, equivalently, a lift of the structure group of E to $SU(2)$. If we pull back the principal $SU(2)$ bundle to $\widehat{S}^3$ and take the associated vector bundle with fiber $\mathbb{C}^2$, we obtain a pair $(\widehat{E}, \widehat{A})$ consisting of a vector bundle with a connection with structure group $SU(2)$, and an action of the quaternion group Q_8 on the bundle, preserving the connection. The quotient $V_4 = Q_8/(\pm 1)$ acts on the base, while the kernel $\pm 1 \subset Q_8$ acts on the fibers of the bundle.

Let I_1, I_2 and I_3 be the nontrivial elements of V_4, and let $\hat{I}_m$ be lifts of these in Q_8 satisfying $\hat{I}_1 \hat{I}_2 = \hat{I}_3$. Let s_+ and s_- be the two vertices of the theta graph, and let $\hat{s}_\pm$ be their unique preimages in $\hat{S}^3$. Let γ_m for $m = 1, 2, 3$ be the arcs of θ where the isotropy group is generated by I_m, and let $\hat{\gamma}_m$ be chosen lifts of these, as paths from $\hat{s}_-$ to $\hat{s}_+$. The lifts can be chosen so that their tangent vectors at $\hat{s}_-$ are a right-handed triad.

Let $(\hat{E}, \hat{A})$ be an equivariant bundle on $\hat{S}^3$, as above. Let $\hat{E}_+$ and $\hat{E}_-$ be the fibers of $\hat{E}$ over $\hat{s}_\pm$. Parallel transport along $\hat{\gamma}_m$ gives an isomorphism

$$\iota_m(\hat{A})\colon \hat{E}_- \to \hat{E}_+.$$

Since $\hat{\gamma}_m$ is fixed by $\hat{I}_m$, the isomorphism $\iota_m(\hat{A})$ commutes with the action of $\hat{I}_m$ on $\hat{E}_-$ and $\hat{E}_+$. Let

$$\tau_-\colon \hat{E}_- \to \mathbb{C}^2, \quad \tau_+\colon \hat{E}_+ \to \mathbb{C}^2$$

be Q_8–equivariant isomorphisms of determinant 1. Each of these is unique up to ± 1. The isomorphism

$$\tau_+ \circ \iota_m(\hat{A}) \circ \tau_-^{-1}\colon \mathbb{C}^2 \to \mathbb{C}^2$$

is then an element of $SU(2)$ that commutes with $\hat{I}_m$, and which has an ambiguity of ± 1 resulting from the ambiguity in $\tau_\pm$. The commutant $C(\hat{I}_m)$ in $SU(2)$ is a copy of the circle group, and can be identified with the standard circle $S^1 = \mathbb{R}/\mathbb{Z}$ by a unique isomorphism

$$L\colon C(\hat{I}_m) \to \mathbb{R}/\mathbb{Z}$$

sending $\hat{I}_m$ to $\frac{1}{4}$. To remove the ambiguity of ± 1, we take the square (written as $2L$ in additive notation), and we define

$$h_m(\hat{A}) = 2L(\tau_+ \circ \iota_m(\hat{A}) \circ \tau_-^{-1}),$$

to obtain a well-defined map

$$h_m\colon \mathcal{B}_\theta \to \mathbb{R}/\mathbb{Z}, \quad m = 1, 2, 3. \tag{9}$$

The maps h_1, h_2 and h_3 depend on some universal choices (essentially the labeling of some standard arcs in $\hat{S}^3$), but we regard these choices as having been made, once and for all.

The maps h_m constructed here for the orbifold (S^3, θ) can be defined in the same way for connected sum with (S^3, θ). That is, if we take an arbitrary bifold $\check{Y}$ and consider the space of marked connections $\mathcal{B}^\sharp(\check{Y})$ defined as in (8), then we have maps

$$h_m\colon \mathcal{B}^\sharp(\check{Y}) \to \mathbb{R}/\mathbb{Z}, \quad m = 1, 2, 3, \tag{10}$$

defined in just the same way, using θ.

3.2 The local system

Given a topological space B and a continuous map

$$h\colon B \to \mathbb{R}/\mathbb{Z}.$$

we can construct a local system Γ of rank-1 R–modules Γ on B, where R is the ring $\mathbb{F}[\mathbb{Z}] = \mathbb{F}[T, T^{-1}]$ of finite Laurent series. The local system we want is pulled back from the universal example on $\mathbb{R}/\mathbb{Z}$. We do this concretely by regarding R as contained in the larger ring $\mathbb{F}[\mathbb{R}]$, and for each $b \in B$ we set

$$\Gamma_b = T^{\tilde{h}(b)} R,$$

where $\tilde{h}(b)$ is any lift of $h(b)$ to $\mathbb{R}$. If ζ is a path from a to b, then

$$\Gamma_\zeta\colon \Gamma_a \to \Gamma_b$$

is multiplication by $T^{\tilde{h}(b)-\tilde{h}(a)}$, where $\tilde{h}$ is any continuous lift of h along the path. Given a collection of maps $h_1, \dots, h_n$ from B to $\mathbb{R}/\mathbb{Z}$ we can similarly construct a local system of rank-1 R–modules, where R is now the ring of finite Laurent series in n variables $T_1, \dots, T_n$.

We apply this now to these three circle-valued maps h_1, h_2 and h_3, defined above at (10), to obtain a local system Γ on $\mathcal{B}^\sharp(\check{Y})$ for any bifold $\check{Y}$. This is a local system of free R–modules of rank 1, where $R = \mathbb{F}[\mathbb{Z}^3]$ is the ring of finite Laurent series in variables T_1, T_2, T_3. (Local systems of this sort were constructed in [11; 12], using maps of a similar sort, but see the remarks at the end of this section for a brief discussion.)

We can now construct the instanton homology groups $J^\sharp(\check{Y}; \Gamma)$ with coefficients in the local system Γ. That is, after equipping $\check{Y}$ with an orbifold Riemannian metric and perturbing the Chern–Simons functional on $\mathcal{B}^\sharp(\check{Y})$ to achieve Morse–Smale transversality, we define an R–module

$$C^\sharp(\check{Y}; \Gamma) = \bigoplus_\beta \Gamma_\beta, \tag{11}$$

where the sum is over critical points (a finite set), and define a boundary map

(12) $$\partial_\Gamma = \sum_{(\alpha,\beta,\zeta)} \Gamma_\zeta,$$

where the sum is over pairs of critical points (α, β) and gradient trajectories ζ of index 1.

Lemma 3.1 *The square $\partial_\Gamma \circ \partial_\Gamma$ is zero.*

Proof Recall from [14] that the usual proof that $\partial^2 = 0$ needs extra care for the instanton homologies such as $J^\sharp$. The reason is that the proof involves analyzing moduli spaces of flow-lines of index 2, and there is a codimension-2 bubbling phenomenon.

So consider a moduli space $M_2(\alpha, \beta)$ of flow-lines of index 2, between critical points α and β. Let $M'_2(\alpha, \beta)$ be the 1–dimensional quotient obtained by dividing by translations. Each connected component of $M'_2(\alpha, \beta)$ belongs to a particular homotopy class of paths ζ from α to β, to which corresponds a map in the local system,

$$\Gamma_\zeta \colon \Gamma_\alpha \to \Gamma_\beta.$$

If we pick one representative $[A]$ for each end of $M'_2(\alpha, \beta)$, we therefore have

(13) $$\sum_{\text{ends}} \Gamma_{\zeta[A]} = 0$$

because the ends come in pairs.

The ends of $M'_2(\alpha, \beta)$ are of two sorts. First, there are ends corresponding as usual to broken trajectories,

$$M'_1(\alpha, \gamma) \times M'_1(\gamma, \beta).$$

The contribution of these ends to the above sum is the matrix entry for ∂_Γ^2 from α to β. The remaining ends belong to antiselfdual bifold connections on the cylinder where (as we approach the end) there is bubbling, in which an instanton of charge $\frac{1}{4}$ bubbles off at a point of the cylinder corresponding to a vertex of the web $K \cup \theta$.

As explained in [14, Section 3.3], the bubble-ends come in groups of four. This is the reason that $\partial^2 = 0$ when the coefficients are $\mathbb{F}$. With local coefficients, however, the four ends in each group can each contribute differently to the sum (13), because they belong to different homotopy classes of paths ζ. In more detail, the weak limit of a sequence of connections where a bubble occurs with Yang–Mills action $\frac{1}{4}$ belongs to a zero-dimensional moduli space on the cylinder, and this can therefore only be a moduli space $M_0(\alpha,\alpha)$ of constant trajectories. So these extra ends occur only when $\alpha=\beta$. The contribution of these ends therefore has the form of a sum over vertices v of $K\cup\theta$ of the form

$$\sum_v (\Gamma_{\zeta(v,1)}+\Gamma_{\zeta(v,2)}+\Gamma_{\zeta(v,3)}+\Gamma_{\zeta(v,4)}),$$

where each $\zeta(v,i)$ is a path from α to α. From the construction of Γ it is clear that $\Gamma_{\zeta(v,i)}=1$ unless v is one of the two vertices of θ. Furthermore, when v is one of the vertices of θ, the term $\Gamma_{\zeta(v,i)}$ is a universal quantity, independent of K and α, which depends only the curvature of the instanton at the bubble point. Therefore there is a universal relation of the form

$$\partial_\Gamma\circ\partial_\Gamma=W\mathbf{1},$$

where $W\in R$ is independent of K. Finally, if we consider the special case that K is empty, we see that W must be zero. Indeed, in this special case, ∂_Γ is zero, because the complex has only one generator and is $\mathbb{Z}/2$ graded. □

The lemma tells us that we have a complex (or more properly a differential module). The finitely generated R–module $J^\sharp(\check{Y};\Gamma)$ is defined as its homology. We summarize the definition:

Definition 3.2 For a closed oriented bifold $\check{Y}$, we define $\mathcal{B}^\sharp(\check{Y})$ as the space of marked $SO(3)$ connections on the marked bifold $(\check{Y}\#(S^3,\theta);\mu_\theta)$. Here the marking region of μ_θ is an open ball B^3 containing the graph $\theta\subset S^3$, with the trivial bundle. We write R for the ring $\mathbb{F}[\mathbb{Z}^3]$ of finite Laurent series in variables T_1, T_2, T_3, and we define Γ as the local system of R–modules on $\mathcal{B}^\sharp(\check{Y})$ constructed from the three circle-valued functions (10). We define $J^\sharp(\check{Y};\Gamma)$ as the Floer homology with these local coefficients on $\mathcal{B}^\sharp(\check{Y})$. □

As usual, given a web $K\subset Y$, where Y is a 3–manifold with framed basepoint, we write $J^\sharp(Y,K;\Gamma)$ for the case that $\check{Y}=(Y,K)$. If $Y=S^3$ and the framed basepoint is at infinity, we simply write $J^\sharp(K;\Gamma)$.

We can give a concrete interpretation of the maps Γ_ζ for a path ζ from α to β in $\mathcal{B}^\sharp(\check{Y})$, which will be helpful later. First, our definition means that

$$\Gamma_\zeta=T_1^{\Delta\tilde{h}_1}T_2^{\Delta\tilde{h}_2}T_3^{\Delta\tilde{h}_3}, \tag{14}$$

where $\Delta\tilde{h}_m$ is the change in a continuous lift $\tilde{h}_m$ of h_m along the path ζ. Since h_m is defined in terms of the holonomy of an S^1 connection, the change in h_m along a path can be expressed as the integral of the curvature of an S^1 bundle with connection. Let γ_m for $m = 1, 2, 3$ again be the three arcs of θ and $\hat{\gamma}_m$ their chosen lifts to $\widehat{S}^3$. The path ζ gives rise to an $SO(3)$ connection on $\mathbb{R} \times \widehat{S}^3$, which we may restrict to $\mathbb{R} \times \hat{\gamma}_m$, where its structure group reduces to $C(I_m)_1$ (the identity component of the commutant of I_m in $SO(3)$). The latter is a circle group, so there are two possible isomorphisms

$$C(I_m)_1 \to S^1 \subset \mathbb{C}.$$

The marking gives us a preferred lift $\hat{I}_m$ in $SU(2)$ and picks out a preferred isomorphism of $C(I_m)_1$ with S^1. Via this isomorphism, our bundle on $\mathbb{R} \times \hat{\gamma}_m$ becomes an S^1 bundle with connection K. The change in the holonomy, $\Delta\tilde{h}_m$, is then the Chern–Weil integral,

$$\Delta\tilde{h}_m = \frac{i}{2\pi} \int_{\mathbb{R}\times\hat{\gamma}_m} F_K. \tag{15}$$

Remarks As mentioned above, instanton Floer homology with local coefficients was applied previously to knots and links in the authors' earlier papers [11; 12]. The local systems used there (also denoted by the generic letter Γ) were defined in a similar manner, but the circle-valued functions h that were used were obtained from holonomies along the components of the knot or link K itself. By contrast, the circle-valued functions in the present paper are defined using the holonomy along the edges of the auxiliary graph θ.

3.3 Functoriality and basic properties

Maps from cobordisms The functorial properties of $J^\sharp(Y, K)$ carry over to the case of local coefficients. As in [14, Definition 3.10], we define a category $\mathcal{C}^\sharp$ whose objects are pairs (Y, K) where K is a web in a closed, oriented 3–manifold Y equipped with a framed basepoint $y_0 \in Y \setminus K$. The morphisms are isomorphism classes of triples (X, Σ, γ), where X is a 4–dimensional cobordism, Σ is an embedded foam with "dots", and γ is a framed arc joining the basepoints. Then we have a functor

$$J^\sharp(-;\Gamma)\colon \mathcal{C}^\sharp \to (\text{finitely generated } R\text{–modules}), \tag{16}$$

where $R = \mathbb{F}[\mathbb{Z}^3]$. To the empty web in $\mathbb{R}^3 \subset S^3$, for example, this functor assigns the free rank-1 module R, and to a closed foam in $\mathbb{R}^4$ it assigns a module map $R \to R$, ie an element of R itself.

The changes that are necessary, to adapt the definition from the case of constant coefficients to the case of the local system Γ, are modeled on the definitions in

[11, Section 5.2], and we outline them here. Given a cobordism (X, Σ, γ) from (Y_0, K_0, y_0) to (Y_1, K_1, y_1), we first use the framing along γ to identify a neighborhood of γ with $[0,1]\times B^3$, and we use this to include a copy of the product foam $[0,1]\times\theta$ in X. Writing $\Sigma^\sharp$ for the union of Σ with the product foam, we have a cobordism of pairs $(X, \Sigma^\sharp)$ from $(Y_0, K_0 \cup \theta)$ to $(Y_1, K_1 \cup \theta)$. Let α and β be critical points for the perturbed Chern–Simons functional in $\mathcal{B}^\sharp(Y_0)$ and $\mathcal{B}^\sharp(Y_1)$, respectively (Definition 3.2). We attach cylindrical ends to the cobordism, and consider a corresponding moduli space of marked antiselfdual connections $M^\sharp(X, \Sigma; \alpha, \beta)$. These are the antiselfdual connections on the cylindrical-end orbifold obtained from $(X, \Sigma^\sharp)$, with the marking region $[0,1]\times\mu_\theta$. As usual, we can regard the moduli space as contained in a larger Banach manifold of connections, $\mathcal{B}^\sharp(X, \Sigma; \alpha, \beta)$.

To define the map (16) at the chain level, what we need formally is a homomorphism of R–modules

$$\Delta_z \colon \Gamma_\alpha \to \Gamma_\beta$$

for each pair of critical points (α, β) and each connected component z of $\mathcal{B}^\sharp(X, \Sigma; \alpha, \beta)$. The chain-level map can then be defined as

$$j^\sharp = \bigoplus_\alpha \bigoplus_\beta \Big(\sum_\eta \Delta_{[\eta]} \Big), \tag{17}$$

where η runs through points in zero-dimensional moduli spaces in $\mathcal{B}^\sharp(X, \Sigma; \alpha, \beta)$.

For this to be a chain map (following the usual argument involving broken trajectories as in [11]), what is required is a composition law, as follows. Given a path ζ_0 in $\mathcal{B}^\sharp(Y_0)$ from α to α', and given z as above, we form a composite $z' \in \pi_0(\mathcal{B}^\sharp(X, \Sigma; \alpha', \beta))$ by concatenating $[\zeta_0]$ and z. What we require is then

$$\Delta_{z'} = \Delta_z \circ \Gamma_{\zeta_0}, \tag{18}$$

with a corresponding composition law also for the Y_1 end.

Finally, we can complete the definition of $J^\sharp(-;\Gamma)$ by giving the appropriate formula for Δ_z. In the cylindrical-end orbifold obtained from $(X, \Sigma^\sharp)$, we have a product region containing 2–dimensional facets $\mathbb{R}\times\hat{\gamma}_m$ for $m = 1, 2, 3$, as in (15). Given a connection $\check{A}$ representing the connected component z in $\mathcal{B}^\sharp(X, \Sigma; \alpha, \beta)$, we obtain an S^1 bundle with connection, K, over $\mathbb{R}\times\hat{\gamma}_m$, and just as in (15) we can evaluate a Chern–Weil integral,

$$g_m = \frac{i}{2\pi} \int_{\mathbb{R}\times\hat{\gamma}_m} F_K, \tag{19}$$

and we define

$$\Delta_z = T_1^{g_1} T_2^{g_2} T_3^{g_3}. \tag{20}$$

The composition law (18) holds because the Chern–Weil integral is additive for broken trajectories. Essentially the same composition law, but for composite cobordisms, is what is used also to prove the functorial properties of $J^\sharp(-;\Gamma)$.

There is a natural extension of this construction which is applied (for example) in the proof of the "exact triangles", Proposition 3.4 below. Suppose that the cylindrical end bifold corresponding to $(X, \Sigma^\sharp)$ is equipped with a family of Riemannian metrics parametrized by a compact, smooth manifold G with boundary. We suppose that all the metrics are the same outside a compact region. One then has *parametrized moduli spaces*

$$M_G^\sharp(X, \Sigma; \alpha, \beta) \to G.$$

By taking a sum over all points in zero-dimensional moduli spaces, one obtains a map of R–modules, just as in (16). Thus, we obtain a map at the chain level,

$$m_G\colon C^\sharp(Y_0, K_0; \Gamma) \to C^\sharp(Y_1, K_1; \Gamma),$$

by

$$m_G = \bigoplus_\alpha \bigoplus_\beta \Big(\sum_\eta \Delta_{[\eta]} \Big).$$

(So the map $j^\sharp$ above now arises as the special case that G is a point.) When G has positive dimension, the map m_G is a chain map *if G has no boundary*. Since the proof involves counting boundary points of 1–dimensional moduli spaces, the general case, when the boundary of G is nonempty, has an extra term. The general formula is

$$\partial \circ m_G + m_G \circ \partial = m_{\partial G}. \tag{21}$$

We have followed here the notation and exposition of [10, Section 3.9]. The adaptation of the proof of this formula to the case of local coefficients follows again from the additivity of the integral (19) along composite cobordisms.

Remark It is useful to note that the verification here of the functorial properties of $J^\sharp$ with local coefficients does not run into any complications of same sort that arise in the proof that $\partial^2 = 0$ (Lemma 3.1). In that earlier lemma, the proof involved moduli spaces of index 2, an index which was large enough to allow bubbling to occur at the vertices of θ. In the definition of the chain maps (17), however, the moduli spaces that are used are zero-dimensional, and the verification that the $j^\sharp$ commutes with ∂ involves only 1–dimensional moduli spaces, as does the argument to prove the functorial property for composite cobordisms.

Applications of excision We have an excision property for $J^\sharp(-;\Gamma)$, just as we do for $J^\sharp$ with constant coefficients. The prototype that is most relevant here is [14, Theorem 4.2], which concerns the following situation. Let $\check{Y}$ be a bifold containing two orbifold 2–spheres S_1 and S_2, each of which has three orbifold points. We allow $\check{Y}$ to have more than one component, and we consider the situation that S_1 and S_2 belong to distinct components. By cutting along S_1 and S_2 and regluing, we obtain a new bifold $\check{Y}'$. The theorem from [14] gives an isomorphism (with coefficients $\mathbb{F}$)

$$J(\check{Y},\mu) = J(\check{Y}',\mu') \tag{22}$$

for appropriate marking data μ and μ'. Here, for a bifold with more than one component, J is defined as the homology of the tensor product of the complexes corresponding to the components. (So $J(\check{Y},\mu)$ is also a tensor product, by the Künneth theorem, if the coefficients are $\mathbb{F}=\mathbb{Z}/2$.)

The main application of this excision isomorphism in the constant-coefficient case is the multiplicative property for connected sums of bifolds (or split unions of webs), stated above as (6). When written for the case of a split union of webs $K = K_1 \cup K_2$ (meaning that there is an embedded 2–sphere S which separates K_1 from K_2), the multiplicative property becomes

$$J^\sharp(K) \cong J^\sharp(K_1) \otimes J^\sharp(K_2) \tag{23}$$

when the coefficients are $\mathbb{F}$.

In order to understand the adaptation for the case of a local coefficient system, we first recall how (23) is deduced from the excision isomorphism (22). We let $\check{Y}_1$ be the bifold (S^3,θ), and $\check{Y}_2$ be the bifold $(S^3, K\cup\theta)$. Let S_1 be a 2–sphere in S^3 which is the boundary of a standard neighborhood of one of the two vertices of θ, meeting θ in three points. We apply the excision isomorphism with $\check{Y}=\check{Y}_1\cup\check{Y}_2$, taking $S_1\subset\check{Y}_1$ as above, and taking S_2 to be the internal connect sum of S_1 with a 2–sphere which separates K_1 from K_2. In the latter case, we enlarge the marking region U_μ so that it includes a neighborhood of S_2, without altering the first homology of the marking region. (Such an enlargement of the marking region does not alter $\mathcal{B}(\check{Y}_2;U_\mu)$, so the Floer homology is also unchanged. See [14].) After cutting and gluing, the new bifold $\check{Y}'$ is the disjoint union of $(S^3, K_1 \cup \theta)$ and $(S^3, K_2 \cup \theta)$. Interpreting (22) in this case, we obtain

$$J^\sharp(\varnothing) \otimes J^\sharp(K) \cong J^\sharp(K_1) \otimes J^\sharp(K_2),$$

which is the desired result (23), because $J^\sharp(\varnothing)$ is $\mathbb{F}$.

We now have the following counterpart to (23) for the case of the local system Γ. As in [14, Corollary 4.4], we extend the statement by including a description of the naturality of the isomorphism.

Proposition 3.3 *Suppose $K = K_1 \cup K_2$ is a split web in $\mathbb{R}^3$, meaning that there is an embedded 2–sphere S which separates K_1 from K_2. Suppose that at least one of the $J^\sharp(K_i;\Gamma)$ is a free R–module. Then there is an isomorphism,*

$$J^\sharp(K;\Gamma) = J^\sharp(K_1;\Gamma) \otimes_R J^\sharp(K_2;\Gamma).$$

Moreover, if $\Sigma \subset [0,1] \times S^3$ is a split cobordism, meaning that $\Sigma = \Sigma_1 \cup \Sigma_2$ and Σ is disjoint from $[0,1] \times S$, then

$$J^\sharp(\Sigma;\Gamma) = J^\sharp(\Sigma_1;\Gamma) \otimes_R J^\sharp(\Sigma_2;\Gamma).$$

In general, if neither is free, then $J^\sharp(K;\Gamma)$ is related to $J^\sharp(K_1;\Gamma)$ and $J^\sharp(K_2;\Gamma)$ by a Künneth theorem (a spectral sequence). □

Proof We form the bifold $\check{Y} = \check{Y}_1 \cup \check{Y}_2$, the orbifold 2–spheres $S_i \subset \check{Y}_i$, and the new bifold $\check{Y}'$ as described above in the outline of the result from [14]. Each of $\check{Y}_1$ and $\check{Y}_2$ contains a theta graph in its singular set, giving rise to local systems Γ_i on the configuration spaces of marked connections $\mathcal{B}(\check{Y}_i, \mu_i)$. On the configuration space $\mathcal{B}(\check{Y}, \mu)$, which we define as the product space, there is a product local system

$$\Gamma = \Gamma_1 \otimes_R \Gamma_2.$$

The local system Γ can be seen as arising from the map

$$\mathcal{B}(\check{Y}, \mu) \to T^3$$

given by the *sum* of the maps on the two factors $\mathcal{B}(\check{Y}_i, \mu_i)$. Having a local system Γ on the product configuration space $\mathcal{B}(\check{Y}, \mu)$, we can again form a Morse complex with local coefficients for the perturbed Chern–Simons functional: it is again the tensor product of the two complexes that compute $\mathcal{B}(\check{Y}_1, \mu_1; \Gamma_1)$ and $\mathcal{B}(\check{Y}_2, \mu_2; \Gamma_2)$.

We define $J(\check{Y}, \mu; \Gamma)$ as the homology of this product complex. We make the same constructions for $\check{Y}'$. The isomorphism in the proposition is equivalent to showing

$$J(\check{Y}, \mu; \Gamma) \cong J(\check{Y}', \mu'; \Gamma') \tag{24}$$

(just as (23) follows from (22)).

Recall from [3; 9; 7] that the proof of such an excision isomorphism is achieved by constructing (bifold) cobordisms $\check{X}$ and $\check{X}'$, from $\check{Y}$ to $\check{Y}'$ and back, in order to construct maps both ways, say $J(\check{X})$ and $J(\check{X}')$. One then proves that the composite of the maps $J(\check{X})$ and $J(\check{X}')$, in either order, is the identity, by a surgery argument. The only essential change now is to understand that $\check{X}$ and $\check{X}'$ give rise to maps of the Floer homology groups with local coefficients. As in the discussion of functoriality above, to use the cobordism $\check{X}$ to construct a chain map, we must define a homomorphism of R–modules

$$\Delta_z \colon \Gamma_\alpha \to \Gamma'_\beta$$

for every connected component z of the space $\mathcal{B}(\check{X}, \mu, \alpha, \beta)$ of marked connections on $\check{X}$.

To understand how to define z, we need to recall the geometry of $\check{X}$, which follows the model in [10] for example. Let U be the 2–manifold with corners in Figure 1, viewed as a cobordism between manifolds with boundary, from $I \times \partial I$ to $\partial I \times I$. Let Z be the 4–dimensional bifold $S \times U$, where S is the three-pointed orbifold 2–sphere. The singular set of $S \times U$ consists of three copies of U, say $\{p_1, p_2, p_3\} \times U$. The

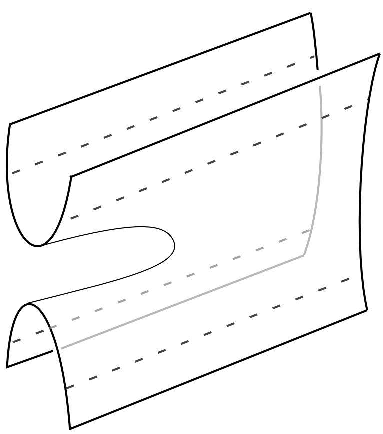

Figure 1: The saddle U, viewed as a surface with corners: a cobordism from $I \times \partial I$ to $\partial I \times I$.

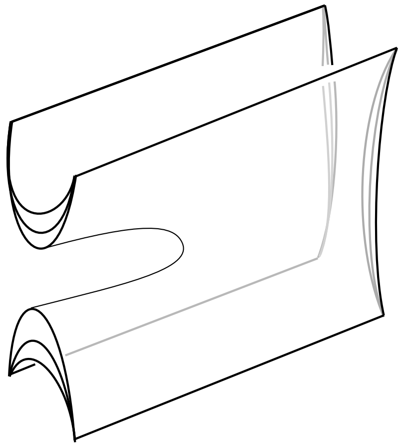

Figure 2: The foam Φ formed from three copies of U, as a cobordism from $\theta_1 \cup \theta_2$ to $\theta_1' \cup \theta_2'$.

nontrivial part of the cobordism $\check{X}$ is exactly a copy of $S \times U$, and $\check{X}$ itself is the union of $S \times U$ with four product cobordisms.

The singular set of $\check{X}$ contains a foam Φ which is a cobordism from two theta graphs, $\theta_1 \cup \theta_2 \subset \check{Y}$, to a different pair of theta graphs, $\theta_1' \cup \theta_2' \subset \check{Y}'$. This Φ is the union of $\{p_1, p_2, p_3\} \times U$ with four copies of $[0,1] \times \in$, where $\in$ is a "Y" graph. See Figure 2, which depicts Φ equivalently as obtained from $\{p_1, p_2, p_3\} \times U$ by identifying $\{p_i\} \times l$ with $\{p_j\} \times l$ for each of four arcs l. We write Φ_1, Φ_2 and Φ_3 for the three copies of U in Φ.

The situation is now very much the same as in the case of the functoriality discussion above. Instead of a trivial cobordism $[0,1] \times \theta$ from θ to itself, we have a nontrivial cobordism Φ from $\theta_1 \cup \theta_2$ to $\theta_1' \cup \theta_2'$. But the construction of Δ_z is essentially unchanged. Corresponding to z, we have an S^1 bundle with connection, K, over Φ, and we define three real numbers

$$g_m = \frac{i}{2\pi} \int_{\Phi_m} F_K, \tag{25}$$

just as in the previous situation (19). Finally, we define Δ_z by the corresponding formula,

$$\Delta_z = T_1^{g_1} T_2^{g_2} T_3^{g_3}.$$

Once again, Stokes' theorem shows that this is a well-defined homomorphism from Γ_α to Γ'_β.

As in [3; 9; 7], the proof is completed by a neck-stretching argument. The composite cobordisms $X_+ = \check{X} \cup_{\check{Y}'} \check{X}'$ and $X_- = \check{X}' \cup_{\check{Y}} \check{X}$ each contain a copy of $R_\pm$ of $S^1 \times S$. Here again S is the bifold 2–sphere with three bifold points. The marking regions contain $R_\pm$. The argument that each composite is the identity is the same. The key feature of the of proof is that $S^1 \times S$, with the marking region being the whole manifold, has a unique marked flat bifold connection ρ, which is in addition irreducible. By the chain-level composition law arising from (21), the composite chain map used to define $J(X') \circ J(X)$ is chain homotopic to the one induced by X_+. Doing the neck stretching along the copy of R_+ breaks X_+ into two disjoint bifolds $\check{X}_1$ and $\check{X}_2$, each with three cylindrical ends. The three ends of $\check{X}_i$ are modeled cylinders whose cross-sections are two oppositely oriented copies of $\check{Y}_i$ and a copy of R_+. Using moduli spaces on these manifolds which are asymptotic to ρ on R_+ and integrating over corresponding surfaces in the $\check{X}_i$, we can define chain maps that give rise to a map

$$J(\check{X}_i)\colon J(\check{Y}_i, \mu_i; \Gamma) \to J(\check{Y}_i, \mu_i; \Gamma).$$

Another application of (21) and additivity of the corresponding curvature integrals used in the defining twisted coefficient map tells us that the chain maps used to define $J(\check{X}') \circ J(\check{X})$ and $J(\check{X}_1 \amalg \check{X}_2)$ are chain homotopic. If we take a standard $Z = D^2 \times S$ with cylindrical end and glue together with either of the $\check{X}_i$, we obtain a manifold diffeomorphic to the product $\mathbb{R} \times \check{Y}_i$. Note that on Z, with the marking set being all of the manifold, there is again a unique marked bifold flat connection $\tilde{\rho}$. This $\tilde{\rho}$ is irreducible and extends ρ. Since $\tilde{\rho}$ is flat, the curvature integrals needed to compute local coefficient maps are zero. We conclude, from a final application of (21) and additivity of the curvature integrals used in the definition of the twisted coefficient map, that (up to chain homotopy) there is no difference in using $\check{X}_i$ as we have or the product $I \times Y_i$ in defining twisted coefficient maps. That is,

$$J(\check{X}') \circ J(\check{X}) = J(\check{X}_+) = J(\check{X}_1 \amalg \check{X}_2) = J(I \times \check{Y}_1 \amalg I \times \check{Y}_2) = J(I \times \check{Y}), \tag{26}$$

which is the identity on $J(\check{Y}, \mu; \Gamma)$. □

Exact triangles The other general properties that we wish to restate in the local-coefficient case are the exact triangles from [13]. The setup here is that we have six webs (in $\mathbb{R}^3$ or in a general 3–manifold) which are the same outside a ball and which differ inside the ball as shown in Figure 3.

There are standard elementary foam-cobordisms from K_{i+1} and L_{i+1} to K_i and L_i.

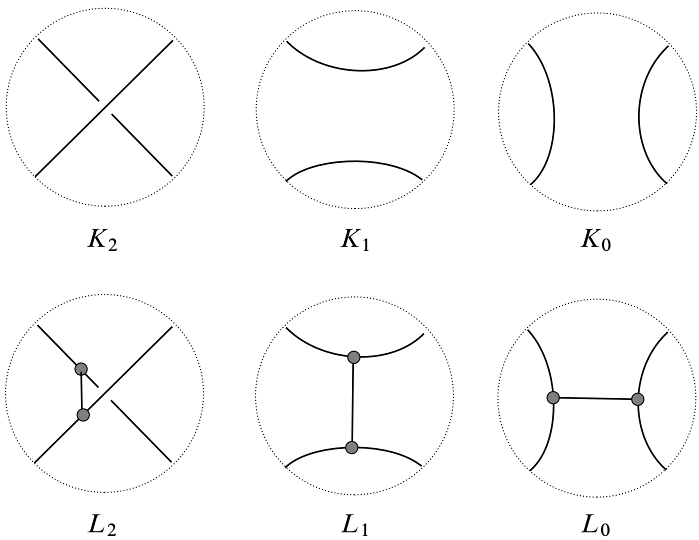

Figure 3: Six webs in Y differing inside a ball

Proposition 3.4 *Given K_i and L_i as in Figure 3, the sequence of R–modules obtained by applying the functor $J^\sharp(-;\Gamma)$ to the 3–periodic sequence*

$$\cdots \to L_2 \to L_1 \to L_0 \to L_2 \to \cdots \tag{27}$$

is exact. So too is the sequence obtained from

$$\cdots \to L_2 \to K_1 \to K_0 \to L_2 \to \cdots, \tag{28}$$

as well as the two other sequences obtained by cyclically permuting the indices (though there is no essential difference).

Proof The proof for the sequence (27) is given in [13] for constant coefficients. Let us write C_i for the chain group corresponding to L_i with local coefficients, so that $J^\sharp(L_i;\Gamma)$ is $H_*(C_i)$. The elementary cobordisms give chain maps of R–modules

$$F_i\colon C_i \to C_{i-1}$$

for all i, by the formula (16). To prove exactness, following the same argument as before, one must first show that $F_i \circ F_{i-1}$ is chain homotopic to zero:

$$F_i \circ F_{i+1} = \partial \circ J_i + J_{i-1} \circ \partial. \tag{29}$$

Having constructed such a $J_i\colon C_i \to C_{i-2}$ (the "first chain homotopy"), one must verify a second chain-homotopy formula: the existence of $K_i\colon C_i \to C_{i-3}$ such that

(omitting the indices)

$$FJ + JF + DK + KD\colon C_i \to C_{i-3} \tag{30}$$

is an isomorphism. (See [13, Section 6].) For the definition of J_i and K_i, one again examines zero-dimensional moduli spaces; but now these moduli spaces lie over families of metrics, Q and P, of dimension 1 and 2, respectively. Thus, given a pair of critical points corresponding to generators α and β for the chain groups C_i and C_{i-2} (respectively C_{i-3}) we have zero-dimensional moduli spaces

$$M_Q(\alpha,\beta) \to Q, \quad M_P(\alpha,\beta) \to P,$$

respectively. (So the fiber dimensions over Q and P are -1 and -2, respectively.) When local coefficients are used, the definition of J_i and K_i is no longer given by just counting the points in these moduli spaces, but instead taking the sum of their contributions in the same way as (16). Thus,

$$J = \bigoplus_\alpha \bigoplus_\beta \Bigl(\sum_\eta \Delta_{[\eta]}\Bigr),$$

with a formula of the same shape for K. The chain-homotopy formulas (29) and (30) arise eventually from application of the principle (21) to compactifications of the families of metrics P and Q. We have already observed that (21) extends without change to the local coefficient case, so this completes the outline of the proof for the exactness of (27).

For the second sequence (28), a shortcut was taken in [13]: the exactness of (28) was deduced from the exactness of (27), using additional relations satisfied by $J^\sharp$. Nevertheless, as indicated in [13], one can ignore the shortcut and give a direct proof for (28), along just the same lines as (27). This direct proof adapts without essential change to the case of local coefficients. □

4 The cubic relation for *u*

4.1 Statement of the result

To each edge e of a web $K \subset Y$, there is an operator

$$u_e\colon J^\sharp(Y,K;\Gamma) \to J^\sharp(Y,K;\Gamma).$$

In the language of foams with dots, this is the operator corresponding to a cylindrical foam $[0,1]\times K$ with a dot on the face $[0,1]\times e$. With constant coefficients $\mathbb{F}$, the operator satisfies $u_e^3=0$, as shown in [14]. But the result with the local system of coefficients Γ is different.

Proposition 4.1 *The operator u_e satisfies*

$$u_e^3+Pu_e=0,$$

where $P\in R$ is the element (3).

This result is of a very similar form as the $N=3$ case of a general result for $SU(N)$ gauge theory, treated in [16]. Our argument follows [16] quite closely, though in the end the argument here is considerably simpler. The proof has some setup required, which we present first.

4.2 Characteristic classes of the basepoint bundles

Let $\check{X}$ be a connected, 4–dimensional bifold cobordism from $\check{Y}'$ to $\check{Y}$. Let ν be marking data on $\check{X}$ (possibly empty) and let $\mathcal{B}^*(\check{X};\nu)$ denote the space of marked bifold connections which are fully irreducible (ie have trivial automorphism group). Depending on the bundle and the marking, this may be all of $\mathcal{B}(\check{X};\nu)$. Write $\check{X}^o$ for the nonsingular part of the orbifold. Then there is a universal $SO(3)$ bundle [4; 2]

$$\mathbb{E}\to\check{X}^o\times\mathcal{B}^*(\check{X};\nu).$$

If we pick a basepoint $x\in\check{X}^o$, then we obtain a bundle $\mathbb{E}_x\to\mathcal{B}^*(\check{X};\nu)$ and characteristic classes

$$\mathsf{w}_{i,x}=w_i(\mathbb{E}_x)\in H^i(\mathcal{B}^*(\check{X};\nu);\mathbb{F}).$$

If s belongs to a 2–dimensional face of the singular set of $\check{X}$, then we can pass to a smooth $\mathbb{Z}/2$ cover of a chart around s and obtain a bundle

$$\widetilde{\mathbb{E}}_s\to\mathcal{B}^*(\check{X};\nu)$$

carrying an involution on the fibers. The $+1$ and -1 eigenspaces of the involution are respectively a line bundle

$$\mathbb{L}_s\to\mathcal{B}^*(\check{X};\nu) \tag{31}$$

and a rank-2 bundle

$$\mathbb{W}_s\to\mathcal{B}^*(\check{X};\nu). \tag{32}$$

Although phrased differently, the real line bundle $\mathbb{L}_s$ is the same one that is used to define the operators u in [14], and we write

$$\mathsf{u}_s = w_1(\mathbb{L}_s) \in H^1(\mathcal{B}^*(\check{X};\nu)).$$

Lemma 4.2 *We have the relation*

$$\mathsf{u}_s^3 + \mathsf{w}_{1,x}\mathsf{u}_s^2 + \mathsf{w}_{2,x}\mathsf{u}_s + \mathsf{w}_{3,x} = 0.$$

Remark Notwithstanding the notation in the lemma, the class $\mathsf{w}_{1,x}$ is zero, because our bundles have structure group $SO(3)$ rather than $O(3)$.

Proof of the lemma Although our exposition asked that x be in the smooth part, we may use the $\mathbb{Z}/2$ cover of a coordinate chart at s in the singular set of $\check{X}$, and we can therefore compute the classes w_i as

$$\mathsf{w}_{i,x} = w_i(\widetilde{\mathbb{E}}_s),$$

where $\widetilde{\mathbb{E}}_s$ is the bundle which has $\mathbb{L}_s$ as a subbundle. The relation in the lemma is a universal relation for a rank-3 bundle containing a line subbundle: it expresses the fact that the complementary rank-2 bundle has $w_3 = 0$. □

4.3 Properties of the point-class operators

Next we parlay the relation from Lemma 4.2 into a relation among operators on the instanton Floer homology groups. We recall how certain cohomology classes in the space of connections give rise to operators, following constructions that go back to [5].

To put ourselves in the framework for $J^\sharp(\check{Y};\Gamma)$, we take $\check{X}$ to be the product cobordism from $\check{Y} \# (S^3,\theta)$ to itself, and we write $\check{X}^+$ for the infinite cylinder. We take the framed arc γ in $\check{X}$ joining the basepoints to be $[0,1]\times y_0$, and the strong marking data ν also to be a product $[0,1]\times\mu_\theta$. We suppose that a holonomy perturbation of the Chern–Simons functional has been chosen so that the critical points are nondegenerate and the moduli spaces $M(\alpha,\beta)$ of marked solutions on $\check{X}^+$ are cut out transversely.

Now let $Z \subset \check{X}$ be a compact 4–dimensional subbifold with boundary. By its intersection with ν, this bifold inherits marking data, and we impose on Z the condition that the restriction map

$$H^1(U_\nu;\mathbb{F}) \to H^1(U_\nu \cap Z;\mathbb{F}) \tag{33}$$

is injective. (This condition is easily fulfilled: we may arrange that Z contains tubular neighborhoods of loops in U_ν that generate $H_1(U_\nu;\mathbb{F})$.) This condition is sufficient to ensure that if (E,A) is a solution in any of the moduli spaces $M(\alpha,\beta)$, then the restriction of (E,A) to Z has trivial automorphism group. We therefore have restriction maps

$$M(\alpha,\beta)\to\mathcal{B}^*(Z;\nu\cap Z). \tag{34}$$

Let $\mathsf{v}\in H^d(\mathcal{B}^*(Z;\nu\cap Z);\mathbb{F})$ be a cohomology class represented as the dual of a codimension-d subvariety V, which we take to a be subset stratified by Banach submanifolds with a smooth open stratum of codimension d and other strata of codimension $d+2$ or more. We suppose these are chosen so that all strata are transverse to the maps (34).

The intersections $M_d(\alpha,\beta)\cap V$ are compact by Uhlenbeck's theorem provided that either

(1) $d\le 7$ and Z does not meet the singular set of the orbifold; or

(2) $d\le 3$ and Z meets the singular set of $\check{X}$ only in the strata with $\mathbb{Z}/2$ stabilizer.

If $d\le 6$ or $d\le 2$, respectively, in the above cases, then we obtain a chain map by summing over the points in these compact intersections. (The proof of the chain property involves moduli spaces of dimension $d+1$, which is the reason for the stronger restriction on d.) The chain map on $C^\sharp(\check{Y};\Gamma)$ is defined by

$$\sum_{\alpha,\beta}\ \sum_{\zeta\in M_d(\alpha,\beta)\cap V}\Gamma_\zeta,$$

and it gives rise to an operator

$$v\colon J^\sharp(\check{Y};\Gamma)\to J^\sharp(\check{Y};\Gamma).$$

A priori, the operator depends on all the choices made.

We apply this construction to the cohomology class u_s described above, for s a basepoint in a face $[0,1]\times e$ of the singular set. To do so, we choose Z containing the point s, not meeting the seams of the foam, and satisfying the constraint that (33) is injective. We can regard u_s as a class in $H^1(\mathcal{B}^*(Z;\nu\cap Z);\mathbb{F})$, where it is the first Stiefel–Whitney class of the line bundle $\mathbb{L}_s$. We define V as the zero-set of a section of $\mathbb{L}_s$ which is transverse to zero and transverse also to the inclusions of (33) in the zero set. The construction of such a transverse section requires smooth partitions of unity on $\mathcal{B}^*(Z;\nu\cap Z)$, which we can always ensure by using L^2 Sobolev spaces (so that our Banach manifolds are Hilbert manifolds). We are in case (2) above, so we obtain an operator

$$u_e\colon J^\sharp(\check{Y};\Gamma)\to J^\sharp(\check{Y};\Gamma).$$

We proceed similarly for the classes $\mathsf{w}_{i,x}$. In this case, since x is not on the foam, we can choose Z to be disjoint from the foam (so that we are in case (1) above). The class $\mathsf{w}_{i,x}$ can be regarded as the Stiefel–Whitney classes of the basepoint bundle $\mathbb{E}_x$ on $\mathcal{B}^*(Z;Z\cap\nu)$. We can define V as the subset where i sections of $\mathbb{E}_x$ fail to be independent, having chosen these sections to achieve transversality to the stratification by rank. In this way, the classes $\mathsf{w}_{i,x}$ give rise to operators

$$w_{i,x}\colon J^\sharp(\check{Y};\Gamma)\to J^\sharp(\check{Y};\Gamma)$$

for $i=1,2,3$.

These operators are independent of choices made (except for the choice of edge e in the case of u_e). In particular, there is no dependence on the choice of Z or the representatives V, by standard chain-homotopy arguments.

Lemma 4.3 *The above operators satisfy the relation*

$$u_e^3+w_{1,x}u_e^2+w_{2,x}u_e+w_{3,x}=0. \tag{35}$$

Proof Let s_1, s_2 and s_3 be distinct basepoints on $[0,1]\times e$, let Z_1, Z_2 and Z_3 be three disjoint subbifolds of $\check{X}$ containing these, and let U_j be a subvariety dual to u_{s_j} in $\mathcal{B}^*(Z_j;\nu\cap Z_j)$. A standard gluing argument [8], used to treat composite cobordisms in general, shows that the composite operator u_e^3 can be computed from the chain map defined by the moduli spaces

$$M_3(\alpha,\beta)\cap U_1\cap U_2\cap U_3.$$

Using further disjoint subsets Z_4, Z_5 and Z_6 and three distinct basepoints in X^o, we construct dual representatives for the classes w_1, w_2 and w_3. Let Z be a larger subset of $\check{X}$ that contains all the Z_i but still meets the singular set of $\check{X}$ only in the codimension-2 faces. The operator on the left-hand side of (35) is then computed from the intersection

$$M_3(\alpha,\beta)\cap V,$$

where V has codimension 3 and is dual to the zero class in $H^3(\mathcal{B}^*(Z;\nu);\mathbb{F})$. We can therefore construct a codimension-2 stratified subvariety W with $\partial W=V$. Let H be the operator defined on chains by the intersections $M_2(\alpha,\beta)\cap W$. At the chain level, we then obtain a chain homotopy formula of the shape

$$u_e^3 + w_{1,x}u_e^2 + w_{2,x}u_e + w_{3,x} = \partial H + H\partial$$

by counting the ends of the 1–dimensional moduli spaces $M_3(\alpha,\beta)\cap W$. Because no moduli spaces of dimension 4 or more are involved, there is no bubbling, and the ends of the 1–dimensional moduli space are all of the form $M_3(\alpha,\beta)\cap V$ (giving the left-hand side) or arise from simple broken trajectories (giving the right-hand side). □

4.4 Calculating the Stiefel–Whitney operators

To complete the proof of Proposition 4.1 from Lemma 4.3, we need to compute the operators $w_{i,x}$.

Proposition 4.4 *On $J^\sharp(\check{Y};\Gamma)$, the operators $w_{1,x}$ and $w_{3,x}$ are both zero, while $w_{2,x}$ is multiplication by the element $P\in R$.*

Proof We first observe that it is sufficient to prove this in the case that $\check{Y}=S^3$. From the definition of $J^\sharp$, we recall that

$$J^\sharp(S^3;\Gamma) = J((S^3,\theta);\mu_\theta;\Gamma)$$

and that this is a copy of R. To understand why the special case is sufficient, consider the disjoint union $\check{Y}=\check{Y}_1\cup\check{Y}_2$, where

$$\check{Y}_1 = \check{Y}\,\#\,(S^3,\theta), \quad \check{Y}_2 = S^3\,\#\,(S^3,\theta),$$

with marking data μ_θ implied in both cases. We have local systems Γ_1 and Γ_2 from the two copies of θ, and we form the local system $\Gamma=\Gamma_1\otimes_R\Gamma_2$ on $\mathcal{B}(\check{Y})$. Because the second is a free R–module, the instanton homology is a tensor product,

$$J(\check{Y};\Gamma) = J(\check{Y}_1;\Gamma_1)\otimes J(\check{Y}_2;\Gamma_2).$$

We can take basepoints x_1 and x_2 in the nonsingular parts of either component. By naturality of the isomorphism, the resulting operators on the tensor product are $w_{i,x_1}\otimes 1$ and $1\otimes w_{i,x_2}$. By an application of excision and its naturality, these two operators are equal. Because $J(\check{Y}_2;\Gamma_2)=R$, the operator w_{i,x_2} is multiplication by an element of R. The same therefore holds for w_{i,x_1}.

So let us consider the operators

$$w_{i,x}\colon J((S^3,\theta);\mu_\theta;\Gamma)\to J((S^3,\theta);\mu_\theta;\Gamma),$$

or in the shorthand notation of webs, the operators

$$w_{i,x}\colon J^\sharp(\varnothing;\Gamma)\to J^\sharp(\varnothing;\Gamma).$$

By Proposition 8.11 of [14], this instanton homology group is $\mathbb{Z}/2$ graded. Being a free rank-1 R–module, it is nonzero in only one of the two gradings. The operators $w_{1,x}$ and $w_{3,x}$ have odd degree and are therefore zero.

It remains to calculate $w_{2,x}$ on $J^\sharp(\varnothing;\Gamma)$. In this Morse homology, there is a unique nondegenerate critical point α (without the need for holonomy perturbation), and to calculate $w_{2,x}$ we need to describe the 2–dimensional moduli space $M_2(\alpha,\alpha)$ of marked antiselfdual orbifold connections on the cylinder. As in Section 3.1, the three-dimensional orbifold can be described as the quotient of a round sphere $\widehat{S}^3$ by $V_4 = Q_8/(\pm 1)$, and marked connections are Q_8–equivariant $SU(2)$ connections on $\widehat{S}^3$. A conformal compactification of the cylinder $\mathbb{R}\times\widehat{S}^3/V_4$ is the orbifold 4–sphere $\widehat{S}^4/V^4$. The fixed-point set of V_4 on $\widehat{S}^4$ is a circle $\widehat{S}^1\subset\widehat{S}^4$. This circle can be identified with the union of the lines $\mathbb{R}\times\hat{s}_+$ and $\mathbb{R}\times\hat{s}_-$ in the cylinder, together with two points at infinity.

The two-dimensional moduli space has Yang–Mills action $\frac{1}{4}$ on the orbifold cylinder, and its pullback to the round $\widehat{S}^4$ therefore has Yang–Mills action 1. In this way, we have identified $M_2(\alpha,\alpha)$ as the space of Q_8–equivariant 1–instantons in an $SU(2)$ bundle on $\widehat{S}^4$. The 1–instanton moduli space on the round 4–sphere is a 5–ball, and the subspace of instantons that are invariant under the action of V_4 is an open 2–ball inside the 5–ball (the one that spans the circle $\widehat{S}^1\subset\widehat{S}^4$). The Uhlenbeck compactification of this V_4–invariant moduli space is the closed 2–disk, obtained by attaching $\widehat{S}^1$.

Each V_4–invariant 1–instanton becomes Q_8–equivariant by lifting the action to the total space of the $SU(2)$ bundle. However, the lift is not unique. Given one lift, we can obtain others by multiplying by any character $Q_8\to\{\pm 1\}$. Since Q_8 has four characters, the moduli space $M_2(\alpha,\alpha)$ consists of four open disks,

$$M_2(\alpha,\alpha)=\bigcup_{i=0}^{3}\mathcal{D}_i^2,$$

each of which has $\widehat{S}^1$ as its Uhlenbeck boundary.

Let $\overline{M}$ be the "small" compactification of the equivariant moduli space on $\widehat{S}^4$, obtained by collapsing $\widehat{S}^1$ to a point. This is a bouquet of four 2–spheres,

$$\overline{M}=\bigvee_{i=0}^{3}\mathcal{S}_i^2.$$

If x is a basepoint in $\widehat{S}^4/V_4$ which does not lie on $\widehat{S}^1$, then the basepoint $SO(3)$ bundle $\mathbb{E}_x$ on the moduli space extends as a bundle on $\overline{M}$.

Lemma 4.5 *The Stiefel–Whitney class of the basepoint bundle is nonzero on each sphere in the bouquet:*

$$w_2(\mathbb{E}_x)[\mathcal{S}_i^2] = 1, \quad i = 0, 1, 2, 3.$$

We postpone the proof of this lemma until the end of this subsection, presenting it after the proof of Lemma 4.6 below.

To continue the proof of Proposition 4.4 we can choose a subset $Z \subset \widehat{S}^4/V_4$ disjoint from the singular set of the orbifold, and we can choose therein a dual representative W_2 for $w_2(\mathbb{E}_x)$ such that the $\overline{M} \cap W_2$ is disjoint from the vertex of the bouquet and meets each of the spheres $\mathcal{S}_i^2$ transversely. The lemma tells us that $\mathcal{S}_i^2 \cap W_2$ is an odd number of points. Returning to the viewpoint of the cylinder $\mathbb{R} \times (\widehat{S}^3/V_4)$, we learn that the transverse intersection $M_2(\alpha, \alpha) \cap W_2$ consists of an odd number of points in each of the disks $\mathcal{D}_i^2$ for $i = 0, 1, 2, 3$. So if ζ_i is a single point of $\mathcal{D}_i$ for each i, then we have that $w_{2,x}$ is multiplication by the element

$$\Gamma_{\zeta_0} + \Gamma_{\zeta_1} + \Gamma_{\zeta_2} + \Gamma_{\zeta_3} \in \mathrm{Hom}(\Gamma_\alpha, \Gamma_\alpha) = R.$$

To complete the proof of Proposition 4.4, we must compute each Γ_{ζ_i} and show that the above sum is P. The following lemma therefore finishes the argument. □

Lemma 4.6 *Let $\zeta_i \in M_2(\alpha, \alpha)$ be a point of $\mathcal{D}_i^2$ for $i = 0, 1, 2, 3$. Then, with suitable conventions and numbering of the components, we have*

$$\begin{aligned}
\Gamma_{\zeta_0} &= T_1 T_2 T_3, \\
\Gamma_{\zeta_1} &= T_1 T_2^{-1} T_3^{-1}, \\
\Gamma_{\zeta_2} &= T_1^{-1} T_2 T_3^{-1}, \\
\Gamma_{\zeta_3} &= T_1^{-1} T_2^{-1} T_3.
\end{aligned}$$

Proof of Lemma 4.6 Consider the case of $\zeta_0 \in \mathcal{D}_0^2$. We take this point to be the center of the disk, which is the standard 1–instanton on $\widehat{S}^4$ with $SO(5)$ symmetry. Each of the disks corresponds to one choice of how to lift the action of V_4 on the standard $SO(3)$ 1–instanton to an action of Q_8 on the $SU(2)$ bundle. We take $\mathcal{D}_0^2$ to be the one obtained by identifying the $SU(2)$ bundle of the centered and scaled 1–instanton with the spin bundle S^- on $\hat{S}^4$ and making $Q_8 \subset \mathrm{Spin}(5)$ act on the spin bundle in the standard way. We use the formulas (14) and (15) to compute Γ_{ζ_0}.

After conformally compactifying $\mathbb{R} \times \hat{S}^3$ to get $\hat{S}^4$, the space $\mathbb{R} \times \hat{\gamma}_m$ becomes one hemisphere in the 2–sphere $\hat{S}^2_m \subset \hat{S}^4$ which is the fixed point set of $I_m \in V_4$. The $SO(3)$ instanton bundle on $\hat{S}^4$ is the rank-3 bundle $\Lambda^- \subset \Lambda^2$. Along $\hat{S}^2_m$ this bundle decomposes under the action of I_m as

$$\Lambda^-|_{\hat{S}^2_m} = \mathbb{R} \oplus K_m,$$

where the circle bundle $K_m \to \hat{S}^2_m$ can be identified with the tangent bundle, as a bundle with connection. Because we are integrating over half of the sphere, we see from (15) that

$$\Delta\tilde{h}_m(\zeta_0) = \tfrac{1}{2}\deg K_m = \tfrac{1}{2}e(S^2_m) = 1.$$

(There is some sign ambiguity remaining in our construction of Γ, but the final sign here fixes our conventions.) From (14), we therefore obtain

$$\Gamma_{\zeta_0} = T_1 T_2 T_3.$$

The other three disks $\mathcal{D}^2_i$ are obtained by changing the lift of the V_4 action to Q_8 by a nontrivial character of Q_8. The each nontrivial character changes two of the three lifts $\hat{I}_1$, $\hat{I}_2$ and $\hat{I}_3$ by -1. Changing $\hat{I}_m$ by -1 changes the identification $C(I_m)_1 \to S^1$ by $z \mapsto -z$, and so changes $\tilde{h}_m$ to $-\tilde{h}_m$. Thus, Γ_{ζ_i} differs from Γ_{ζ_0} by changing the sign of the exponent of T_m for two of the three values of m. This proves the lemma. □

Proof of Lemma 4.5 We return to the postponed proof of Lemma 4.5. The sphere $\mathcal{S}^2_i$ is obtained from a the closed disk $\bar{\mathcal{D}}^2_i$ by collapsing the boundary to a point. We need to describe both the bundle $\mathbb{E}_x \to \bar{\mathcal{D}}^2_i$ and the trivialization that is used when collapsing the boundary.

Let $\hat{x}$ be a lift of x to the round sphere $\hat{S}^4$. There is a unique 2–disk $\hat{D}^2 \subset \hat{S}^4$ which has boundary $\hat{S}^1$ and contains $\hat{x}$. We may take it that $\hat{x}$ is the center of the disk. We identify the centered 1–instanton bundle with Λ^- as in the previous lemma. The disk $\mathcal{D}^2_i$ in the moduli space M is obtained by applying conformal transformations, and we can use these to identify $\mathcal{D}^2_i$ with $\hat{D}^2$ in such a way that the basepoint bundle

$$\mathbb{E}_x \to \mathcal{D}^2_i$$

becomes the bundle

$$\iota^*(\Lambda^-)|_{\hat{D}^2},$$

where ι is the involution on the disk that fixes the center x. The trivialization of Λ^- that we must use on the boundary of $\hat{D}^2$ is the V_4–invariant trivialization of Λ^- on the circle $\hat{S}^1$. This is the same trivialization as is obtained by parallel transport around $\hat{S}^1$.

The question of whether w_2 is zero or not now becomes the question of whether the trivialization of Λ^- on $\hat{S}^1$ obtained by parallel transport can be extended to a trivialization on a disk in $\hat{S}^4$ spanned by $\hat{S}^1$. The answer is that it cannot, and this is essentially the same point as occurs in the proof of the previous lemma: on a suitable disk, we can reduce the structure group of Λ^- to $SO(2)$, and the $SO(2)$ bundle has degree 1 with respect to the trivialization. □

4.5 Three-edge relations

Before moving on, we consider a different relation involving the operators u_e, which can be proved in the same manner. Let $\check{S}$ be an orbifold 2–sphere with three singular points with local isotropy group $\mathbb{Z}/2$. Suppose S is embedded in $\check{Y}$ as a suborbifold. (In terms of the web K in the three-manifold $Y = |\check{Y}|$, this is a sphere meeting K transversely in three points belonging to edges of K.) Corresponding to the three points of intersection, we have three operators

$$u_1, u_2, u_3\colon J^\sharp(\check{Y};\Gamma) \to J^\sharp(\check{Y};\Gamma).$$

Lemma 4.7 *The three operators u_1, u_2 and u_3 satisfy the relations*

$$u_1 + u_2 + u_3 = 0, \quad u_1 u_2 u_3 = 0.$$

Proof Let $\check{X} = [0,1] \times \check{Y}$, and $Z \subset \check{X}$ be a regular neighborhood of the sphere $\{\frac{1}{2}\} \times S$, meeting the singular set only in the codimension-2 strata. On $\mathcal{B}^*(Z)$ we have three cohomology classes u_i for $i = 1, 2, 3$, and the first relation in the lemma will follow, just as in the proof of Lemma 4.3, if we establish a relation in the cohomology of $\mathcal{B}^*(Z)$,

$$\mathsf{u}_1 + \mathsf{u}_2 + \mathsf{u}_3 = 0.$$

This is equivalent to showing that for any bifold $SO(3)$ connection on $S^1 \times \check{S}$, the product of the three real line bundles $\mathbb{L}_i \to S^1$ obtained from the three orbifold points of $\check{S}$ is trivial. The triviality of $\mathbb{L}_i$ is equivalent to the bifold bundle E having $w_2 = 0$ on the torus $S^1 \times \delta_i$, where $\delta_i \subset \check{S}$ is a circle linking the i^{th} singular point. Since the sum of the three tori bounds in the smooth part of $S^1 \times \check{S}$, the above relation follows.

The second relation is an algebraic consequence of the first relation and the relations $u_i^3 + P u_i = 0$. □

5 Proof of the main theorem

In this section, we present the proof of Theorem 1.2. Although the final statement of the theorem involves webs K in $\mathbb{R}^3$ or S^3, we continue to treat the case of a general bifold $\check{Y}$ corresponding to a web $K \subset Y$, for as long as possible. Nevertheless, the 3–manifold Y will be omitted from our notation, and we write $J^\sharp(K;\Gamma)$ for $J^\sharp(\check{Y};\Gamma)$. We will make clear that $Y = S^3$ in the statements that require it.

5.1 The edge-decomposition

Let R' be obtained from R by adjoining an inverse $1/P$ for the element for P of Equation (3). For any bifold $\check{Y}$, let Γ' be the local system $\Gamma \otimes_R R'$. The corresponding instanton homology group $J^\sharp(K;\Gamma')$ is an R'–module. The polynomial

$$u^3 + Pu$$

which annihilates the edge-operators u_e is the product of factors u and u^2+P which are coprime in $R'[u]$. That is, we have

$$au + b(u^2+P) = 1$$

in $R'[u]$, where $a = u/P$ and $b = 1/P$. It follows that the module $J^\sharp(K;\Gamma')$ has a direct-sum decomposition,

$$J^\sharp(K;\Gamma') = \ker(u_e) \oplus \ker(u_e^2+P) = \ker(u_e) \oplus \operatorname{im}(u_e).$$

This is the (generalized) eigenspace decomposition: if we were to adjoin a square root of P, then these summands would become the eigenspace and generalized eigenspace for the two eigenvalues 0 and $P^{1/2}$. There is one such decomposition for each edge e of K, and since the operators belonging to different edges commute, we obtain a decomposition of the R'–module into simultaneous generalized eigenspaces. To introduce notation for this, let us write

$$E(K) = \{\text{edges of } K\},$$

and given a subset $s \subset E(K)$, let us write

$$V(K;s) = \Bigl(\bigcap_{e\in s} \ker(u_e)\Bigr) \cap \Bigl(\bigcap_{e\notin s} \operatorname{im}(u_e)\Bigr).$$

Then we have a decomposition of the R'–module into $2^{\#E(K)}$ direct summands,

$$J^\sharp(K;\Gamma') = \bigoplus_{s\subset E(K)} V(K;s). \tag{36}$$

Definition 5.1 For any web $K \subset Y$, we refer to the above direct-sum decomposition of the R'–module $J^\sharp(K;\Gamma')$ as the *edge-decomposition*.

5.2 The case of the unknot

We calculate the edge decomposition for the unknot K in $\mathbb{R}^3$, which has a single, circular edge e.

Proposition 5.2 *For the unknot $K = U$, the R–module $J^\sharp(U;\Gamma)$ is free of rank 3, and with respect to a standard basis the matrix of the operator u_e is*

$$\begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & P \\ 0 & 1 & 0 \end{pmatrix}.$$

Over the ring R' the module $J^\sharp(U;\Gamma')$ decomposes as a direct sum of $\ker(u_e)$ and $\operatorname{im}(u_e)$, which have ranks 1 and 2, respectively.

As in [14], we obtain information about the unknot by considering the unknotted 2–sphere $S \subset \mathbb{R}^4$ as a foam. We write $S(m)$ for the 2–sphere with m dots, considered as a cobordism from the empty web to itself. As a cobordism, it has an evaluation in $J^\sharp(-;\Gamma)$, which we write as

$$\langle S(m)\rangle \in R.$$

Lemma 5.3 *For $m = 0, 1$ and 2, we have $\langle S(m)\rangle = 0, 0, 1$, respectively. For $m > 2$, we have $\langle S(m)\rangle = P\langle S(m-2)\rangle$.*

Proof This follows [14]. For $m < 2$, the formula for the dimension of the moduli spaces tells us that the zero-dimensional moduli spaces have negative Yang–Mills action, and are therefore empty. For $m = 2$, we are evaluating u_e^2 on a compact moduli space of flat connections on the 4–dimensional cylinder. The moduli space is equal to $\mathbb{RP}^2$, and the evaluation of u_e^2 is 1 because u_e restricts to the generator in cohomology. When using local coefficients, the evaluation $\langle S(2)\rangle$ is therefore equal to $1 \cdot \Delta_z$, where Δ_z is given by the formula (20) as in Section 3.3. Because the connections are flat, the exponents in the formula (20) are zero, and therefore $\Delta_z = 1$. The claim for $m > 2$ is a consequence of the relation (2). □

Proof of Proposition 5.2 Let us write S as the union of two disks D^+ and D^-, viewed as cobordisms from the empty web to U and back. Let $D^\pm(m)$ be the same foams, but with m dots. Let us write

$$\begin{aligned} \boldsymbol{v}_m &= J^\sharp(D^+(m);\Gamma) \in J^\sharp(U;\Gamma), \\ \boldsymbol{a}_m &= J^\sharp(D^-(m);\Gamma) \in \mathrm{Hom}(J^\sharp(U;\Gamma), R). \end{aligned}$$

We have $\boldsymbol{a}_m(\boldsymbol{v}_n) = \langle S(m+n)\rangle$, so from the lemma we can read off these pairings. If we set

$$\begin{aligned} \boldsymbol{b}_0 &= \boldsymbol{a}_2 + P\boldsymbol{a}_0, \\ \boldsymbol{b}_1 &= \boldsymbol{a}_1, \\ \boldsymbol{b}_2 &= \boldsymbol{a}_0, \end{aligned}$$

then we can read off that

$$\boldsymbol{b}_m(\boldsymbol{v}_n) = \delta_{mn}$$

for $m, n \le 2$.

The representation variety of U is $\mathbb{RP}^2$, and there is a holonomy perturbation with three critical points. The complex that computes $J^\sharp(U;\Gamma)$ therefore has three generators, so $J^\sharp(U;\Gamma)$ is either free of rank 3, or has rank strictly less. But the above relation shows that $\boldsymbol{v}_0$, $\boldsymbol{v}_1$, $\boldsymbol{v}_2$ generate a free submodule. So $J^\sharp(U;\Gamma)$ has rank 3, with the $\boldsymbol{v}_i$ as a basis.

To compute the matrix of u_e in the basis $\boldsymbol{v}_0$, $\boldsymbol{v}_1$, $\boldsymbol{v}_2$, we must compute $\boldsymbol{b}_n(u_e\boldsymbol{v}_m)$, which again can be interpreted as evaluations of 2–spheres with dots:

$$\begin{aligned} \boldsymbol{b}_0(u_e\boldsymbol{v}_m) &= \langle S(m+3)\rangle + P\langle S(m+1)\rangle, \\ \boldsymbol{b}_1(u_e\boldsymbol{v}_m) &= \langle S(m+2)\rangle, \\ \boldsymbol{b}_2(u_e\boldsymbol{v}_m) &= \langle S(m+1)\rangle. \end{aligned}$$

From this one may compute the matrix shown in the proposition. The kernel of u_e is the free rank-1 module spanned by the element

$$\begin{pmatrix} P \\ 0 \\ 1 \end{pmatrix}.$$

The image of u_e is the free rank-2 module spanned by the elements

$$\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}.$$

The same applies after tensoring with R' to make P invertible, and in that case the free module of rank 3 is the direct sum of these two submodules. □

In terms of the notation of (36) and Definition 5.1, the conclusion of the proposition is that

$$\begin{aligned} \operatorname{rank} V(U;\{e\}) &= 1, \\ \operatorname{rank} V(U;\varnothing) &= 2. \end{aligned} \tag{37}$$

5.3 The case of the theta graph

Next, we calculate the edge-decomposition for the theta graph θ, with its three edges e_1, e_2, e_3.

Proposition 5.4 *For the theta graph, the deformed instanton homology $J^\sharp(\theta;\Gamma)$ is free of rank 6. In the edge-decomposition of the R'–module $J^\sharp(\theta;\Gamma')$, the nonzero summands are the summands $V(\theta;s)$ where $s \subset \{e_1,e_2,e_3\}$ is a singleton. Each of the three nonzero summands has rank 2.*

Following [14] again, we look at the closed foam $\Theta \subset \mathbb{R}^4$ consisting of three disks with a common circle as boundary. Let $\Theta(m_1,m_2,m_3)$ denote this foam with m_i dots on the i^{th} disk, and let

$$\langle \Theta(m_1,m_2,m_3)\rangle \in R$$

be the evaluation of the closed foam in $J^\sharp(-;\Gamma)$. These evaluations are entirely determined by the following lemma:

Lemma 5.5 *The evaluation $\langle\Theta(m_1,m_2,m_3)\rangle$ is symmetric in the three variables. It is zero if all the m_i are positive, or if the sum of the m_i is either even or less than three. We have*

$$\langle\Theta(0,1,2)\rangle = 1.$$

Finally, if $m_1 \geq 3$, then

$$\langle\Theta(m_1,m_2,m_3)\rangle = P\langle\Theta(m_1-2,m_2,m_3)\rangle.$$

Proof The symmetry is clear. The assertion that the evaluation is zero if all the m_i are positive follows from the relation $u_1u_2u_3 = 0$ in Lemma 4.7. The assertion that the evaluation is zero if the sum of the m_i is even follows from the dimension formula [14] and the fact that bifold bundles for this foam have Yang–Mills action a multiple of $\frac{1}{4}$. If the sum of the m_i is less than three, then the moduli spaces that contribute have negative Yang–Mills action, so the evaluation is zero. The evaluation for $\Theta(0,1,2)$ holds because the moduli space of flat connections is the three-dimensional flag manifold: this calculation proceeds as in [14] of the case of constant coefficients $\mathbb{F}$, with the same adaptation to the case of local coefficients, Γ, as presented in the proof of Lemma 5.3 above. □

Proof of Proposition 5.4 Mimicking the calculation for the case of the unknot, we introduce the half-foams Θ^+ and Θ^- as cobordisms from $\varnothing$ to θ and from θ to $\varnothing$, and we write

$$\begin{aligned} \boldsymbol{v}(l,m,n) &= J^\sharp(\Theta^+(l,m,n);\Gamma) \in J^\sharp(\theta;\Gamma), \\ \boldsymbol{a}(l,m,n) &= J^\sharp(\Theta^-(m);\Gamma) \in \mathrm{Hom}(J^\sharp(\theta;\Gamma), R). \end{aligned}$$

The pairing between these can be computed from the lemma, and we can compute that the pairing between the six elements

$$\boldsymbol{v}(0,0,0), \quad \boldsymbol{v}(0,0,1), \quad \boldsymbol{v}(0,0,2), \quad \boldsymbol{v}(0,1,0), \quad \boldsymbol{v}(0,1,1), \quad \boldsymbol{v}(0,1,2) \tag{38}$$

and the similarly dotted $\boldsymbol{a}$'s is upper triangular with entries 1 on the diagonal. It follows that these six elements $\boldsymbol{v}(l,m,n)$ generate a free submodule of rank 6. On the other hand, the representation variety of θ is the flag manifold, and the Chern–Simons functional has a perturbation with six critical points. So (much as in the case of the unknot), we conclude that $J^\sharp(\theta;\Gamma)$ is free of rank 6 and the elements (38) are a basis.

Four of the six basis elements $\boldsymbol{v}(l,m,n)$ have $n>0$, from which it follows that the rank of $\mathrm{im}(u_3)$ (and hence also the rank of $\ker(u_3^2+P)$) is at least four. The same applies to all the u_i by symmetry. The relation $u_1+u_2+u_3=0$ from Lemma 4.7 implies $u_1^2+u_2^2+u_3^2=0$, so if

$$x \in \ker(u_1^2+P) \cap \ker(u_2^2+P) \cap \ker(u_3^2+P)$$

then $Px=0$. It follows that

$$\mathrm{im}(u_1) \cap \mathrm{im}(u_2) \cap \mathrm{im}(u_3)$$

also consists of P–torsion elements, so this submodule has rank 0.

If we pass now to the free, R'–module $J^\sharp(\theta;\Gamma')$, we see that each submodule $\mathrm{im}(u_i)$ has rank at least 4, and that the intersection of all three is zero. Since the whole module has rank 6, it follows that $\mathrm{im}(u_2)\cap\mathrm{im}(u_3)$ has rank exactly 2, and that

$$\mathrm{im}(u_2)\cap\mathrm{im}(u_3) = \ker(u_1) = V(\theta;\{e_1\}).$$

In the edge-decomposition, the instanton homology $J^\sharp(\theta;\Gamma')$ is now the direct sum of the three rank-2 modules $V(\theta;\{e_i\})$: the other summands are zero because all the rank is accounted for. □

There is an additional relation that can be extracted from this computation:

Lemma 5.6 *For the operators on $J^\sharp(\theta;\Gamma)$, we have the relation*

$$u_2u_3 + u_3u_1 + u_1u_2 = P.$$

Proof Since both are free modules, it is sufficient to check this on the R'–module $J^\sharp(\theta;\Gamma')$. Using the edge-decomposition and symmetry, we may reduce to checking it on the summand $V(\theta;\{e_1\})$. This rank-2 submodule is $\mathrm{im}(u_2)\cap\mathrm{im}(u_3)$ and is spanned by $\boldsymbol{v}(0,1,1)$ and $\boldsymbol{v}(0,1,2)$. Since $u_1 = 0$ on this summand, the operator on the left of the relation in the lemma is just u_2u_3. So we are left to check that

$$u_2u_3\boldsymbol{v}(0,1,n) = P\boldsymbol{v}(0,1,n)$$

for $n=1$ and $n=2$. This is equivalent to

$$\boldsymbol{v}(0,2,n+1) = P\boldsymbol{v}(0,1,n).$$

A basis for the dual of this summand is provided by $\boldsymbol{a}(0,0,0)$ and $\boldsymbol{a}(0,0,1)$. So we have to show

$$\langle\Theta(0,2,m+1)\rangle = P\langle\Theta(0,1,m)\rangle$$

for $m=1,2$ and 3. Both sides are zero unless $m=2$, in which case both sides are P, by Lemma 5.5. □

5.4 Edge-decompositions and 1–sets

The relation of Lemma 5.6, for operators on the theta graph, extends to the case of a general web $K\subset Y$, as follows. At each vertex, there are three incident edges (e_1,e_2,e_3), where we allow that two of the three may be the same. Let u_1, u_2 and u_3 be the corresponding operators on the R–module $J^\sharp(K;\Gamma)$. Then we have the following relations:

Proposition 5.7 *The three operators u_1, u_2, u_3 on $J^\sharp(K;\Gamma)$ corresponding to the edges incident at a vertex satisfy the three relations*

$$u_1+u_2+u_3 = 0, \tag{39}$$

$$u_2u_3+u_3u_1+u_1u_2 = P, \tag{40}$$

$$u_1u_2u_3 = 0. \tag{41}$$

Proof The first and third relations are special cases of the relations of Lemma 4.7. The middle relation (40) follows from the case of the theta graph (Lemma 5.6) by an application of excision, as in [14, Proposition 5.8]. □

Remark The three relations of this proposition reflect the fact that, at a vertex x of the singular set of the orbifold, we have a direct-sum decomposition of the basepoint bundle $\widetilde{\mathbb{E}}_x$ as a sum of three real line bundles,

$$\widetilde{\mathbb{E}}_x = \mathbb{L}_1 \oplus \mathbb{L}_2 \oplus \mathbb{L}_3,$$

and corresponding relations among the characteristic classes,

$$\begin{aligned} \mathsf{u}_1 + \mathsf{u}_3 + \mathsf{u}_3 &= \mathsf{w}_{1,x}, \\ \mathsf{u}_2\mathsf{u}_3 + \mathsf{u}_3\mathsf{u}_1 + \mathsf{u}_1\mathsf{u}_2 &= \mathsf{w}_{2,x}, \\ \mathsf{u}_1\mathsf{u}_2\mathsf{u}_3 &= \mathsf{w}_{3,x}. \end{aligned}$$

One could convert these relations among cohomology classes (in particular, the second relation) into relations between corresponding operators, using the same ideas as the proof of Proposition 4.1. This would provide an alternative proof of the operator relations. However, the argument needs to be carried out in the space of connections on a neighborhood of a vertex, where codimension-2 bubbling can occur, making the proof more difficult.

We say that a subset $s \subset E(K)$ is a *k–set* (for $k = 1$ or 2) if, at each vertex of K, exactly k of the three incident edges belong to s. The complement of a 1–set is a 2–set. This is standard terminology, but since our webs are allowed to have (for example) no vertices, it is possible for a subset s to be both a 1–set and a 2–set. (In particular, both $\varnothing$ and $\{e\}$ are 1–sets in the case of the unknot.) A 2–set is the same as a disjoint union of cycles which includes every vertex. A 1–set is also called a *perfect matching*. Every bridgeless trivalent graph admits a perfect matching [15].

Proposition 5.8 *In the edge-decomposition*

$$J^\sharp(K;\Gamma') = \bigoplus_{s\subset E(K)} V(K;s),$$

the summand $V(K;s)$ is zero if s is not a 1–set.

Proof Let u_1, u_2, u_3 be the operators corresponding to the three edges incident at a vertex of K. Over the ring R', consider the operators

$$\pi_1 = (1/P)u_2u_3, \quad \pi_2 = (1/P)u_3u_1, \quad \pi_3 = (1/P)u_1u_2$$

on $J^\sharp(K;\Gamma')$. From Proposition 5.7, we have

$$\pi_1 + \pi_2 + \pi_3 = 1$$

and

$$\pi_i \pi_j = 0$$

for $i \neq j$. From this it follows that the π_i are projections ($\pi_i^2 = \pi_i$) and that their images give a direct-sum decomposition

$$J^\sharp(K;\Gamma') = \operatorname{im}(\pi_1) \oplus \operatorname{im}(\pi_2) \oplus \operatorname{im}(\pi_3).$$

It is also clear that $\operatorname{im}(\pi_1)$ for example is contained in $\operatorname{im}(u_2) \cap \operatorname{im}(u_3)$, and also $\operatorname{im}(\pi_1) \subset \ker(u_1)$ by another application of (41). So we have

$$\operatorname{im}(\pi_1) \subset \ker(u_1) \cap \operatorname{im}(u_2) \cap \operatorname{im}(u_3). \tag{42}$$

The reverse inclusion holds also, because if $x \in \ker(u_1)$ then $\pi_2 x = \pi_3 x = 0$, from which it follows that $x = \pi_1 x$. We learn that

$$J^\sharp(K;\Gamma') = \ker(u_1) \cap \operatorname{im}(u_2) \cap \operatorname{im}(u_3) \oplus \ker(u_2) \cap \operatorname{im}(u_3) \cap \operatorname{im}(u_1) \oplus \ker(u_3) \cap \operatorname{im}(u_1) \cap \operatorname{im}(u_2). \tag{43}$$

But the first term on the right is the direct sum of all those summands $V(K;s)$ for which s contains e_1 but not e_2 or e_3, and the sum of the three terms subspaces on the right is the sum of all $V(K;s)$ for which s contains exactly one of the three edges. □

Corollary 5.9 *For any web K in Y, we have a direct-sum decomposition*

$$J^\sharp(K;\Gamma') = \bigoplus_{1\text{–sets } s} V(K;s). \qquad \square$$

Corollary 5.10 *For any web K in Y and any subset $t \subset E(K)$, we have*

$$\bigcap_{e \in t} \ker(u_e) = 0$$

in $J^\sharp(K;\Gamma')$ if t is not contained in a 1–set. Similarly,

$$\bigcap_{e \in t} \operatorname{im}(u_e) = 0$$

if t is not contained in a 2–set. □

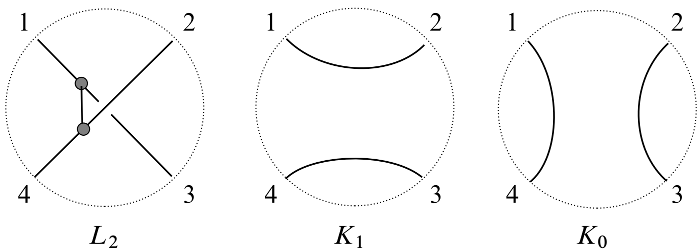

Figure 4: Webs in Y differing inside a ball with edges labeled

5.5 Applications of the exact triangle and excision

The excision isomorphism which gives rise to the product rule, Proposition 3.3, respects the edge-decompositions of the webs that are involved. As a simple application, we have the following:

Lemma 5.11 *Let $K \subset Y$ be a web, and let $s \subset E(K)$ be 1–set. Let K' be the union $K \cup U$, where U is an unknot contained in a ball disjoint from K. Let e be the single edge of U, and $s' \subset E(K')$ be the 1–set $s \cup \{e\}$. Then*

$$V(K; s) \cong V(K'; s').$$

Proof Proposition 5.2 tells us that $V(U; \{e\})$ is a free module of rank 1. So we apply Proposition 3.3 to obtain

$$V(K'; s') \cong V(K; s) \otimes V(U; \{e\}) \cong V(K; s). \qquad \square$$

Consider next the exact triangle of Proposition 3.4 for the case of L_2, K_1, K_0. The proof works for any local system, so the exact sequence holds for the R'–modules $J^\sharp(-; \Gamma')$. For each of L_2, K_1 and K_0, let $p_1, \dots, p_4$ be points lying at the indicated locations on the webs, shown in Figure 4.

Corresponding to these marked points, we have operators $u_1, \dots, u_4$ on $J^\sharp(K; \Gamma')$ for each of the three webs, and the homomorphisms in the exact sequence commute with these operators, because they come from cobordisms in which the points lie on common faces. Of course, some of the operators are equal: say $u_1 = u_2$, for example, as operators on $J^\sharp(K_1; \Gamma')$.

For each of $u_1, \dots, u_4$, we have a decomposition of $J^\sharp(K; \Gamma')$ into $\ker(u_i) \oplus \operatorname{im}(u_i)$, and the long exact sequence of Proposition 3.4 decomposes into a direct sum of 16 exact sequences, although many terms are zero. We give two applications of this idea.

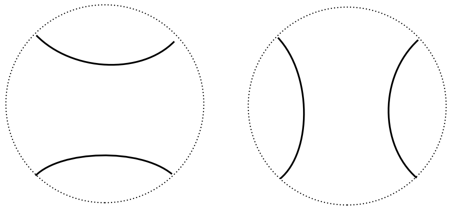

Figure 5: Webs K and K' in Y as in Corollary 5.14. The edges shown belong to 1–sets s and s'.

Lemma 5.12 *For the webs K_1 and K_0, the summands*

$$\bigcap_{i=1}^{4} \ker(u_i)$$

in $J^\sharp(K_1;\Gamma')$ and $J^\sharp(K_0;\Gamma')$ are isomorphic.

Proof The corresponding summand of $J^\sharp(L_2;\Gamma')$ is zero by Corollary 5.10, because the edges of L_2 on which the points $p_1, \ldots, p_4$ lie are not part of a 1–set. So the exact sequence for these summands becomes an isomorphism between the other two terms. □

Lemma 5.13 *For the webs L_2 and K_0, the summands*

$$(\mathrm{im}(u_1) \cap \mathrm{im}(u_4)) \cap (\ker(u_2) \cap \ker(u_3))$$

in $J^\sharp(L_2;\Gamma')$ and $J^\sharp(K_0;\Gamma')$ are isomorphic.

Proof The corresponding summand of $J^\sharp(K_1;\Gamma')$ is zero because $u_1 = u_2$, so $\ker(u_1) \cap \mathrm{im}(u_2) = 0$. So again the exact sequence becomes an isomorphism between the other two terms. □

We can draw the following corollaries of these two lemmas, in the language of the edge-decomposition, Definition 5.1.

Corollary 5.14 *Let s and s' be 1–sets for the webs K and K' in Y. Suppose that K and K' differ only in a ball, as in Figure 5, and that the edges of K and K' which meet the ball belong to s and s', respectively. Then*

$$V(K;s) \cong V(K';s').$$

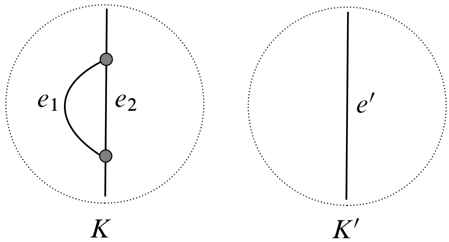

Figure 6: Webs K and K' in Y as in Corollary 5.15. The edge e_1 belongs to the 1–set s. The other edges do not.

Proof This is an immediate consequence of the definition of $V(K;s)$ and Lemma 5.12. □

The second corollary is a little more elaborate, but still straightforward.

Corollary 5.15 *Let s and s' be 1–sets for the webs K and K' in Y. Suppose that K and K' differ only in a ball, as in Figure 6. Suppose that the edge e_1 of K belongs to s* (so that the other edges of K the figure do not). *Suppose that the edge e' of K' does not belong to s'. Then*

$$V(K;s) \cong V(K';s').$$

Proof Set $L_2 = K$ and apply Lemma 5.13 to the ball which forms a neighborhood of the edge e_2. Then the web K_0 in Lemma 5.13 becomes a union of K' and an unknotted circle U. Writing e for the single edge of U, we learn that

$$V(K;s) = V(K' \cup U; s' \cup \{e\}).$$

Now apply Lemma 5.11. □

We can summarize Corollaries 5.14 and 5.15 and Lemma 5.11 together as saying that $V(K,s)$ is unchanged if we alter only the edges of K that belong to s, by the addition of 1–handles, 0–handles or 2–handles (saddle-moves, and births and deaths of circles).

Proposition 5.16 *Let (K,s) and (K',s') be two webs in Y, each equipped with 1–sets $s \subset E(K)$ and $s' \subset E(K')$. Let*

$$C = \bigcup_{e \in E(K)\setminus s} \bar{e}, \quad C' = \bigcup_{e' \in E(K')\setminus s'} \bar{e}'$$

be the closed loops in Y formed by the edges of the complementary 2–sets. Suppose that $C = C'$ as subsets of Y, and that the 1–sets s and s' define the same relative *homology class in* $H_1(Y, C; \mathbb{Z}/2)$*. Then*

$$V(K, s) \cong V(K', s').$$

Proof The homology between s and s' is a composition of isotopies, births, deaths and saddle moves, and the addition and subtraction of bigons as in Corollary 5.15. □

5.6 Calculation for planar webs

We now turn to the special case of webs K in the plane $\mathbb{R}^2 \subset \mathbb{R}^3$. Let s be a 1–set for K, and let C be the union of closed cycles formed by the complementary 2–set. We say that s is *even* if its homology class in $H_1(\mathbb{R}^3, C; \mathbb{Z}/2)$ is zero. This is equivalent to saying that s has an even number of endpoints on each of the connected components of C.

Proposition 5.17 *Let K be a planar web and let s be an even 1–set. Then $V(K, s)$ is a free R'–module of rank 2^n, where n is the number of components in C. If s is not even, then $V(K, s) = 0$.*

Proof If s is even, then we can apply Proposition 5.16 to see that $V(K, s)$ is isomorphic to $V(K', \varnothing)$, where K' is the disjoint union of the circles that constitute C, with no vertices. Since K is planar, K' is an unlink. If it has n components, then by the product formula in Proposition 3.3 and the unknot calculation in Proposition 5.2, we have

$$V(K', \varnothing) = V(U, \varnothing)^{\otimes n} = (R' \oplus R')^{\otimes n}.$$

This establishes the first claim.

For the second claim, suppose s is not even, and let $C_1 \subset C$ be a connected component on which s has an odd number of endpoints. Then, using Proposition 5.16, we can replace (K, s) with (K', s'), where the cycles of the complementary 2–set are unchanged, but s' has exactly one endpoint on C_1, and $V(K, s) = V(K', s')$. Furthermore, we can arrange that C_1 bounds a disk that is disjoint from the rest of K'. In the language of [14], the web K' now has an embedded bridge; that is, there is a 2–sphere in $\mathbb{R}^3$ meeting K' transversely in a single point. It follows that the bifold representation variety of K' is empty and $J^\sharp(K'; \Gamma') = 0$. So $V(K', s') = 0$ a fortiori. □

Corollary 5.18 *For a planar web K, the R'–module $J^\sharp(K,\Gamma')$ is a free module of rank*

$$\sum_{s\in\{\text{even 1–sets}\}} 2^{n(s)},$$

where, for each s, the number $n(s)$ is the number of cycles in the complementary 2–set. □

Corollary 5.19 *For a planar web K, the rank of the R'–module $J^\sharp(K,\Gamma')$ is equal to the number of Tait colorings of K.*

Proof It is an elementary fact that the number of Tait colorings is equal to the sum that appears in the previous corollary. Indeed, given a Tait coloring, the edges colored "red" are a 1–set, and to complete the coloring we must alternate "blue" and "green" along the cycles of the complementary 2–set. This will not be possible if the 1–set is not even, and can be done in $2^{n(s)}$ ways when it is. □

Finally, we return to the ring R and the local system Γ. We do not know whether $J^\sharp(K;\Gamma)$ is a free R–module, but we do know that its rank is the same as that of the R'–module $J^\sharp(K;\Gamma')$. This is because both ranks are equal to the dimension of the F–vector space $J^\sharp(K;\Gamma\otimes_R F)$, where F is the field of fractions of both R and R'. This proves Theorem 1.2. □

6 Comparison of the deformed and undeformed homology

6.1 The spectral sequence of the $\mathfrak{m}$–adic filtration

Recall that R is the ring of finite Laurent series in variables T_1, T_2, T_3, and that Γ is a local system of free rank-1 R–modules on $\mathcal{B}^\sharp(\check{Y})$. Let $m \lhd R$ be the ideal generated by

$$\{T_i - 1 \mid i = 1,2,3\};$$

let R_m be the localization of R at this maximal ideal: the ring of rational functions whose denominator is nonzero at $(1,1,1)$. Let $\mathfrak{m} \lhd R_m$ be the unique maximal ideal in the localization. Let $\Gamma_m = \Gamma\otimes_R R_m$ be the local system of R_m–modules obtained from Γ.

Using the local system Γ_m in place of Γ, we can form the instanton homology group $J^\sharp(\check{Y};\Gamma_m)$. Its rank as an R_m–module is equal to the rank of $J^\sharp(\check{Y};\Gamma_m)$ as an

R–module, because R_m and R have the same field of fractions. In particular, if $\check{Y} = (S^3, K)$ and K is planar, then the rank is equal to the number of Tait colorings of K.

Write

$$C(\Gamma_m) = C^\sharp(\check{Y}; \Gamma_m) = \bigoplus_\beta \Gamma_{m,\beta} \tag{44}$$

for the chain complex whose homology is $J^\sharp(\check{Y}; \Gamma_m)$, as in (11). Although the terminology "complex" is traditional, it should be remembered that $C(\Gamma_m)$ has no grading in general. It is a differential R_m–module. We write ∂_m for the differential, so

$$J^\sharp(\check{Y}; \Gamma_m) = H(C(\Gamma_m), \partial_m).$$

The rank-1 local system Γ_m has the $\mathfrak{m}$–adic filtration

$$\Gamma_m \supset \mathfrak{m}\Gamma_m \supset \mathfrak{m}^2\Gamma_m \supset \cdots,$$

and there is the corresponding filtration of the differential group,

$$C(\Gamma_m) \supset \mathfrak{m}C(\Gamma_m) \supset \mathfrak{m}^2 C(\Gamma_m) \supset \cdots.$$

By its construction, $\mathfrak{m}^p C(\Gamma_m)$ is the same as $C(\mathfrak{m}^p\Gamma_m)$. The $\mathfrak{m}$–adic filtration of the differential group gives rise to an induced filtration of the homology,

$$J^\sharp(\check{Y}; \Gamma_m) = \mathcal{F}^0 \supset \mathcal{F}^1 \supset \mathcal{F}^2 \cdots, \tag{45}$$

where as usual $\mathcal{F}^p$ is the subset of the homology that can be represented by cycles in $\mathfrak{m}^p C(\Gamma_m)$.

For the statement of the next proposition, recall that $J^\sharp(\check{Y})$ denotes the instanton homology with coefficients $\mathbb{F} = \mathbb{Z}/2$.

Proposition 6.1 *There is a convergent spectral sequence of differential R_m–modules whose E_1 page is the filtered module*

$$J^\sharp(\check{Y}) \otimes \operatorname{gr} R_m$$

with the filtration obtained from the $\mathfrak{m}$–adic filtration of R_m, and which abuts to the filtered module $J^\sharp(\check{Y}; \Gamma_m)$ with the filtration $\mathcal{F}^p$ induced by the $\mathfrak{m}$–adic filtration of $C(\Gamma_m)$. Thus,

$$J^\sharp(\check{Y}) \otimes \operatorname{gr} R_m \Rightarrow \operatorname{gr} J^\sharp(\check{Y}; \Gamma_m).$$

Furthermore, the filtration $\mathcal{F}^p$ is $\mathfrak{m}$–stable, in that $\mathfrak{m}\mathcal{F}^p \subset \mathcal{F}^{p+1}$ with equality for large enough p.

Proof Since $C(\Gamma_m)$ has finite rank as an R_m–module, the existence of a convergent spectral sequence from the filtered differential module $C(\Gamma_m)$ is standard. See for example [6, Theorem A3.22]. Note that, although the situation most often considered is a graded differential module (a complex in the traditional sense), we can create a complex from the differential module by placing $C(\Gamma_m)$ in every degree,

$$\cdots \xrightarrow{\partial_m} C(\Gamma_m) \xrightarrow{\partial_m} C(\Gamma_m) \xrightarrow{\partial_m} C(\Gamma_m) \xrightarrow{\partial_m} \cdots .$$

In the present example, the E_1 page is the homology of the associated graded differential module,

$$(\mathfrak{m}^p C(\Gamma_m))/(\mathfrak{m}^{p+1} C(\Gamma_m)) = C((\mathfrak{m}^p \Gamma_m)/(\mathfrak{m}^{p+1}\Gamma_m)).$$

The local system $(\mathfrak{m}^p\Gamma_m)/(\mathfrak{m}^{p+1}\Gamma_m)$ is isomorphic to the constant coefficient system $\mathfrak{m}^p/\mathfrak{m}^{p+1}$, because each T_i acts trivially on the quotient. This identifies the E_1 page as $J^\sharp(\check{Y})\otimes \operatorname{gr} R_m$, as claimed. For the claim of $\mathfrak{m}$–stability, see [6, Exercise A3.42]. □

We have the following corollary (which can also be proved simply and directly, as pointed out in the introduction):

Corollary 6.2 *There is an inequality of ranks,*

$$\dim_{\mathbb{F}} J^\sharp(\check{Y}) \ge \operatorname{rank}_{R_m} J^\sharp(K;\Gamma_m) = \operatorname{rank}_R J^\sharp(K;\Gamma).$$

Proof The spectral sequence tells us that

$$\dim_{\mathbb{F}}(J^\sharp(\check{Y}) \otimes (\mathfrak{m}^p/\mathfrak{m}^{p+1})) \ge \dim_{\mathbb{F}}(\mathcal{F}^p/\mathcal{F}^{p+1}),$$

for these are the dimensions of the associated graded modules on the E^1 page and the limit, respectively. Because the filtration is $\mathfrak{m}$–stable, we have

$$\dim_{\mathbb{F}}(\mathcal{F}^p/\mathcal{F}^{p+1}) \sim \dim_{\mathbb{F}}(\mathfrak{m}^p J^\sharp(\check{Y};\Gamma_m)/\mathfrak{m}^{p+1} J^\sharp(\check{Y};\Gamma_m)),$$

and the right-hand side is asymptotic to $n \dim(\mathfrak{m}^p/\mathfrak{m}^{p+1})$, where n is the rank of $J^\sharp(\check{Y};\Gamma_m)$. □

Corollary 6.3 (Corollary 1.3 of the introduction) *If $K \subset \mathbb{R}^2 \subset \mathbb{R}^3$ is a planar web, then the dimension of $J^\sharp(K)$ is greater than or equal to the number of Tait colorings of K.*

Proof This is now an immediate consequence of Corollary 6.2 and Theorem 1.2. □

The simplest example where this provides a new bound is the case of the 1–skeleton of the dodecahedron. Previous work [14; 13] provided a lower bound of 58 in this case, but the corollary provides a lower bound of 60. The authors do not know whether the inequality of Corollary 1.3 continues to hold if K is not planar.

6.2 Criteria for equality of ranks

We examine when equality can occur for the inequality of ranks in Corollary 6.2. For this purpose, we introduce a variant of the local system Γ. First, let us extend the ground field $\mathbb{F}$ by transcendentals z_1, z_2, z_3, and write

$$\widetilde{\mathbb{F}} = \mathbb{F}(z_1, z_2, z_3).$$

We replace our ring R and its localization R_m by $\widetilde{R} = R \otimes \widetilde{\mathbb{F}}$ and its localization $\widetilde{R}_m$. The latter is the local ring at the point $(1,1,1)$ in $\mathbb{A}^3$, and we now wish to restrict to a generic line through $(1,1,1)$, namely a line

$$(T_1, T_2, T_3) = (1 + z_1 t, 1 + z_2 t, 1 + z_3 t).$$

Thus, we introduce the ring S which is the localization at $0 \in \mathbb{A}^1$ of the polynomial ring $\widetilde{\mathbb{F}}[t]$, and we regard S as a module over $\widetilde{R}$ by

$$q_z\colon T_i \mapsto 1 + z_i t. \tag{46}$$

We write

$$\mathfrak{n} \lhd S$$

for the maximal ideal in this local ring. We have a local system of S–modules

$$\Gamma_S = \Gamma \otimes_R S$$

and instanton homology groups $J^\sharp(K; \Gamma_S)$.

The analysis of $J^\sharp(K; \Gamma_S)$ runs in the same way as $J^\sharp(K; \Gamma)$ and $J^\sharp(K; \Gamma_m)$. The main point is that the image of P under the homomorphism q_z is a nonzero element $P_S \in S$. It has the form

$$P_S = \Bigl(\sum_{i<j} z_i^2 z_j^2\Bigr) t^4 + O(t^5). \tag{47}$$

Our edge operators u_e now satisfy $u_e^3 + P_S u_e = 0$, and the fact that P_S is nonzero is sufficient for us to repeat the previous arguments. The ranks of $J^\sharp(K; \Gamma_S)$ and $J^\sharp(K; \Gamma)$ as an S–module and an R–module, respectively, are the same, because the field of fractions of S is obtained by adjoining t to the field of fractions of R. This allows us to carry over Theorem 1.2 directly; or we could repeat the same proof. Either way, we have:

Proposition 6.4 *If K lies in the plane, then the rank of $J^\sharp(K;\Gamma_S)$ as an S–module is equal to the number of Tait colorings of K.* □

Consider the $\mathfrak{n}$–adic filtration of the local system Γ_S, and the corresponding filtration of the differential S–module $C^\sharp(\Gamma_S)$ which computes $J^\sharp(K;\Gamma_S)$, as in Section 6.1. We have an induced filtration of $J^\sharp(K;\Gamma_S)$,

$$J^\sharp(K;\Gamma_S) = \mathcal{G}^0 \supset \mathcal{G}^1 \supset \mathcal{G}^2 \cdots, \tag{48}$$

as the counterpart to the filtration (45) of $J^\sharp(K;\Gamma_m)$. We therefore have a spectral sequence, just as in Proposition 6.1.

Proposition 6.5 *There is a convergent spectral sequence of ungraded differential S–modules whose E_1 page is the filtered module*

$$J^\sharp(\check{Y}) \otimes \operatorname{gr} S$$

with the filtration obtained from the $\mathfrak{n}$–adic filtration of S, and which abuts to the filtered module $J^\sharp(\check{Y};\Gamma_S)$ with the filtration $\mathcal{G}^p$ induced by the $\mathfrak{n}$–adic filtration of $C(\Gamma_S)$. Thus,

$$J^\sharp(\check{Y}) \otimes \operatorname{gr} S \Rightarrow \operatorname{gr} J^\sharp(\check{Y};\Gamma_S).$$

Furthermore, the filtration $\mathcal{G}^p$ is $\mathfrak{n}$–stable, in that $\mathfrak{n}\mathcal{G}^p \subset \mathcal{G}^{p+1}$ with equality for large enough p. □

The reason for replacing the 3–dimensional regular local ring R_m with the 1–dimensional local ring S is that the induced filtration $\mathcal{G}^p$ on the homology is much easier to understand in the case that S is a principal ideal domain. The above spectral sequence is now an example of a Bockstein spectral sequence, associated in this case to the exact coefficient sequence

$$0 \to S \xrightarrow{t} S \to \widetilde{\mathbb{F}} \to 0.$$

The next simplification in the case of a principal ideal domain is that the induced filtration on the homology of a complex is easier to understand.

Lemma 6.6 *Let S be a domain, and let (C,∂) be a differential S–module, with C a free S–module. Let $H(C)$ be its homology. Let $\mathfrak{n}=(t)\lhd S$ be a principal ideal generated by a prime t, and let $\mathcal{G}^p$ be the p^{th} step of the filtration of $H(C)$ induced by the $\mathfrak{n}$–adic filtration of C. Then $\mathcal{G}^p=\mathfrak{n}^p H(C)$.*

Proof Membership of $\mathcal{G}^p$ means that a homology class α can be written in the form $[a]$, where $\partial a=0$ and $a\in\mathfrak{n}^p C$. The latter condition means that $a=t^p b$ for some b, and the condition $t^p\partial b=0$ implies $\partial b=0$. So there is a homology class $\beta=[b]$ with $\alpha=t^p\beta$. So $\alpha\in\mathfrak{n}^p H(C)$. The reverse inclusion is straightforward, whether or not the ideal is principal. □

So the induced filtration of $J^\sharp(\check{Y};\Gamma_S)$ is the $\mathfrak{n}$–adic filtration of the homology group as an S–module, and the associated graded object to which the spectral sequence of Proposition 6.5 abuts has terms

$$\mathrm{gr}_p\, J^\sharp(\check{Y};\Gamma_S)=\frac{\mathfrak{n}^p J^\sharp(\check{Y};\Gamma_S)}{\mathfrak{n}^{p+1} J^\sharp(\check{Y};\Gamma_S)}.$$

So the spectral sequence in Proposition 6.5 implies an inequality of ranks,

$$\dim_{\mathbb{F}} J^\sharp(\check{Y})\geq \dim_{\widetilde{\mathbb{F}}}\left(\frac{\mathfrak{n}^p J^\sharp(\check{Y};\Gamma_S)}{\mathfrak{n}^{p+1} J^\sharp(\check{Y};\Gamma_S)}\right).$$

Equality holds if and only if all differentials d_r in the spectral sequence vanish. By Nakayama's lemma, we have

$$\dim_{\widetilde{\mathbb{F}}}\left(\frac{J^\sharp(\check{Y};\Gamma_S)}{\mathfrak{n} J^\sharp(\check{Y};\Gamma_S)}\right)\geq \mathrm{rank}_S\, J^\sharp(\check{Y};\Gamma_S),$$

with equality only if $J^\sharp(\check{Y};\Gamma_S)$ is a free S–module. Further, in the case that $J^\sharp(\check{Y};\Gamma_S)$ is free, all the pieces of the associated grade object have dimension equal to the rank of the module. This proves the following proposition:

Proposition 6.7 *We have*

$$\dim_{\mathbb{F}} J^\sharp(\check{Y})\geq \mathrm{rank}_S\, J^\sharp(\check{Y};\Gamma_S),$$

and equality holds if and only if all differentials d_r with $r\geq 1$ in the spectral sequence of Proposition 6.5 are zero, and in that case the module $J^\sharp(\check{Y};\Gamma_S)$ is free. □

Although we have obtained a criterion from the Bockstein spectral sequence, the other tool one can use in the case of principal ideal domain is the universal coefficient theorem for homology. Our differential module $C^\sharp(\Gamma_S)$ is a free S–module, from which $C^\sharp(\widetilde{\mathbb{F}})$ is obtained by reducing mod (t). The universal coefficient theorem therefore tells us that we have split short exact sequence and an isomorphism

$$J^\sharp(\check{Y}) \otimes \widetilde{\mathbb{F}} \cong J^\sharp(\check{Y}; \Gamma_S) \otimes \widetilde{\mathbb{F}} \oplus \operatorname{Tor}(J^\sharp(\check{Y}; \Gamma_S), S/(t)).$$

Concretely, the finitely generated S–module $J^\sharp(\check{Y}; \Gamma_S)$ has a decomposition

$$J^\sharp(\check{Y}; \Gamma_S) \cong S^r \oplus \frac{S}{(t^{a_1})} \oplus \cdots \oplus \frac{S}{(t^{a_l})},$$

and the universal coefficient theorem tells us that

$$\dim_{\mathbb{F}} J^\sharp(\check{Y}) = r + 2l.$$

In particular, equality holds in Proposition 6.7 if and only if $J^\sharp(\check{Y}; \Gamma_S)$ is torsion free. As a corollary, combining this with Proposition 6.4, we obtain the following:

Corollary 6.8 *Let K be a planar web. Then $\dim_{\mathbb{F}} J^\sharp(K)$ is greater than or equal to the number of Tait colorings, and equality holds if and only if one of the following two equivalent conditions holds:*

(1) *the spectral sequence of Proposition 6.5 collapses at the E_1 page; or*

(2) *the instanton homology $J^\sharp(K; \Gamma_S)$ is torsion free.* □

Remark While we chose to introduce three indeterminates z_i in order to have a general line through $(1,1,1)$, it is apparent from the formula (47) that P_S is nonzero if we make either of the substitutions $(z_1, z_2, z_3) = (1,1,1)$, or $(z_1, z_2, z_3) = (1,1,0)$. In place of the ring S, we could have used the smaller ring $\bar{S}$ obtained as the localization of $\mathbb{F}[t]$ at $t = 0$, made into an R–module by either of the two homomorphisms

$$q_{(1,1,1)}\colon R \to \bar{S}, \quad q_{(1,1,0)}\colon R \to \bar{S}$$

given by

$$q_{(1,1,1)}\colon (T_1, T_2, T_3) \mapsto (1+t, 1+t, 1+t),$$
$$q_{(1,1,0)}\colon (T_1, T_2, T_3) \mapsto (1+t, 1+t, 1),$$

respectively.

6.3 Nonvanishing for the homology with local coefficients

Returning to the three-dimensional local ring R_m or to R itself and the original local coefficient system Γ, the universal coefficient theorem provides not a long exact sequence but a Tor spectral sequence, because we are not dealing with a principal ideal domain. Namely, there is a spectral sequence with

$$E_p^2 = \mathrm{Tor}_p^R(J^\sharp(\check{Y};\Gamma), \mathbb{F})$$

abutting to $J^\sharp(\check{Y})$. As an R–module, $\mathbb{F}$ has a resolution by a complex of free R–modules of length 3, so the Tor groups that appear on the E^2 page are zero for $p > 3$. (So only the d_2 and d_3 differentials are potentially nonzero.) A useful consequence is that, if $J^\sharp(\check{Y};\Gamma)$ is zero, then so is $J^\sharp(\Gamma)$. From the nonvanishing theorem in [14] we therefore obtain a nonvanishing theorem for $J^\sharp(\check{Y};\Gamma)$.

Theorem 6.9 *Let $K \subset \mathbb{R}^3$ be a web with no spatial bridge. Then $J^\sharp(K;\Gamma)$ is a nonzero R–module.* □

6.4 Two nonplanar examples

The diagram on the left in Figure 7 is an example of a nonplanar web $K_1 \subset \mathbb{R}^3$ for which the ranks of $J^\sharp(K_1)$ and $J^\sharp(K_1;\Gamma)$ are different.

The web K_1 admits only one 1–set, which is the singleton $\{e_1\}$ consisting of the S–shaped "chain" of the handcuffs. From the edge-decomposition, we therefore learn that $J^\sharp(K_1;\Gamma') = V(K_1;\{e_1\})$. Similarly, for the other web in the picture, $J^\sharp(K_2;\Gamma') = V(K_2;\{e_2\})$. From an application of Proposition 5.16, we have $V(K_1;\{e_1\}) = V(K_2;\{e_2\})$, and it follows that there is an isomorphism

$$J^\sharp(K_1;\Gamma') \cong J^\sharp(K_2;\Gamma').$$

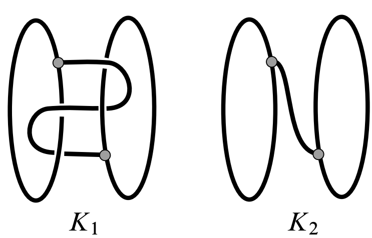

Figure 7: The tangled handcuffs and the standard handcuffs

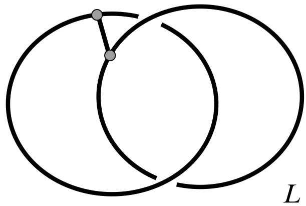

Figure 8: The linked handcuffs L

But the web K_2 has an embedded bridge, so $J^\sharp(K_2;\Gamma')=0$. It follows that for the "tangled handcuffs" K_1, we have $J^\sharp(K_1;\Gamma')=0$. For the coefficient system Γ, we do not have a complete calculation, but we can record at least the following consequence:

Proposition 6.10 *For the tangled handcuffs K_1 in Figure 7, we have*

$$\operatorname{rank}_R J^\sharp(K_1;\Gamma)=0.$$

It is a finitely generated torsion module.

The web K_1 does not have an embedded bridge. (Indeed, its $SO(3)$ representation variety is nonempty and consists of one fully irreducible representation whose image in $SO(3)$ is the octahedral group.) So Theorem 6.9 tells us that $J^\sharp(K_1;\Gamma)$ is nonzero. From the corresponding nonvanishing theorem for $J^\sharp(K_1)$ proved in [14], we also learn that

$$\dim_{\mathbb{F}} J^\sharp(K_1)>0.$$

The tangled handcuffs are therefore an example where there are nontrivial differentials in the spectral sequence of Proposition 6.1.

A related example is the "linked handcuffs" L in Figure 8. There is an exact triangle (Proposition 3.4) in which the role of L_2 is played by the linked handcuffs and the role of both K_1 and K_0 is played by the unknot U. The connecting homomorphism is provided by a cobordism, Σ, from one unknot to the other. This cobordism is obtained by starting with the cylindrical cobordism $[0,1]\times U$ in $[0,1]\times\mathbb{R}^3$ and taking a connected sum with the pair $(S^4,\mathbb{RP}^2)$, where the $\mathbb{RP}^2$ is standardly embedded in S^4 with self-intersection -2. A gluing argument shows that the induced map

$$J^\sharp(U;\Gamma)\to J^\sharp(U;\Gamma)$$

is equal to the map arising from the 2–dimensional cohomology class

$$\mathsf{z}\in H^2(\mathcal{B}^\sharp(U);\mathbb{F})$$

given by

$$\mathsf{z} = w_2(\mathbb{W}_s),$$

where $\mathbb{W}_s$ is the rank-2 bundle (32) corresponding to the point s on the cylinder where the connected sum is made. From the Whitney sum formula, we have

$$\mathsf{z} = \mathsf{u}_s^2 + \mathsf{w}_{2,x}(\mathbb{E}_x),$$

so the corresponding operator on $J^\sharp(U;\Gamma)$ is $u_e^2 + P$. So there is a short exact sequence

$$0 \to \operatorname{coker}(u_e^2 + P) \to J^\sharp(L;\Gamma) \to \ker(u_e^2 + P) \to 0.$$

The calculation in Proposition 5.2 tell us that both the kernel and cokernel of $u_e^2 + P$ are free of rank 2. So the exact sequence of R–modules is

$$0 \to R \oplus R \to J^\sharp(L;\Gamma) \to R \oplus R \to 0,$$

which necessarily splits. We record this result:

Proposition 6.11 *For the linked handcuffs L in Figure 8, we have*

$$J^\sharp(L;\Gamma) \cong R^4$$

as an R–module. □

This example shows that the rank of $J^\sharp(L;\Gamma)$ is not always equal to the number of Tait colorings for nonplanar webs. (The linked handcuffs have no Tait colorings.) With coefficients in the field $\mathbb{F}$, the original $J^\sharp(L)$ has rank 4 also, as can be seen using essentially the same exact sequence, because the kernel and cokernel of u_e^2 have dimension 2 in $J^\sharp(U)$.

6.5 Vanishing of the first three differentials

It is not hard to identify the differential on the E_1 page of the spectral sequence, Proposition 6.1, arising from the filtration of the local ring R_m. Recall that the local system is defined using maps to the circle,

$$h_i \colon \mathcal{B}^\sharp(\check{Y}) \to \mathbb{R}/\mathbb{Z}, \quad i = 1, 2, 3,$$

described at (10). If s is the generator of $H^1(\mathbb{R}/\mathbb{Z};\mathbb{F})$, then by pullback we obtain classes

$$\mathsf{t}_i \in H^1(\mathcal{B}^\sharp(\check{Y});\mathbb{F}), \quad i = 1, 2, 3.$$

As in Section 4.3, these classes give rise to operators

$$\tau_i\colon J^\sharp(\check{Y}) \to J^\sharp(\check{Y}). \tag{49}$$

Explicitly, let $h_i^{-1}(\delta) \subset \mathcal{B}^\sharp(\check{Y})$ be a generic level set of h_i transverse to all 1–dimensional moduli spaces $M_1(\alpha,\beta)$. Then the matrix entries of the corresponding chain-map $\tilde{\tau}_i$ on the chain level are the mod 2 count of the intersection points $M_1(\alpha,\beta) \cap h_i^{-1}(\delta)$.

Lemma 6.12 *The differential on the E_1 page of the spectral sequence of Proposition 6.1 is the operator*

$$d_1\colon J^\sharp(\check{Y}) \otimes (\mathfrak{m}^p/\mathfrak{m}^{p+1}) \to J^\sharp(\check{Y}) \otimes (\mathfrak{m}^{p+1}/\mathfrak{m}^{p+2}),$$

given by

$$d_1 = \sum_{i=1}^{3} \tau_i \otimes (1 - \overline{T}_i),$$

where $\overline{T}_i$ is the homomorphism $\mathfrak{m}^p/\mathfrak{m}^{p+1} \to \mathfrak{m}^{p+1}/\mathfrak{m}^{p+2}$ given by multiplication by T_i. In particular, d_1 is zero if and only if each τ_i is zero.

Proof The formula for d_1 arises by expanding the expression for ∂_m around $T_i = 1$. Let us trivialize the local system Γ_m on the complement of $V_1 \cup V_2 \cup V_3$, so that we identify $\Gamma_{m,\beta} = R_m$ for all critical points β, and for a path ζ which is transverse to the three V_i we have

$$\Gamma_{m,\zeta} = T_1^{n_1} T_2^{n_2} T_3^{n_3},$$

where n_i is the signed intersection number of ζ with V_i. Modulo $\mathfrak{m}^2$, this is equal to

$$1 + \sum_1^3 \bar{n}_i (1 - T_i),$$

where $\bar{n}_i$ is the mod 2 residue of n_i. Thus, at the chain level, we have

$$\partial_m = \partial + \sum_1^3 \tilde{\tau}_i (1 - T_i) + x, \tag{50}$$

where the matrix entries of x belong to $\mathfrak{m}^2$ and ∂ is the ordinary differential on $J^\sharp(\check{Y})$, extended to the trivial local system with fiber R_m. □

It turns out that the operators τ_i on $J^\sharp(\check{Y})$ are identically zero. In fact, we have more:

Proposition 6.13 *For any bifold $\check{Y}$, the differentials d_1, d_2 and d_3 in the spectral sequence of Proposition 6.1 are zero.*

Proof Fix an integer $k > 0$, and consider the R–module

$$\delta_k = m^k / m^{k+4},$$

where m is the maximal ideal at $T_i = 1$. Let Δ_k be the local coefficient system obtained from Γ by tensoring with this module,

$$\Delta_k = \Gamma \otimes_R \delta_k.$$

For each k, we then have instanton homology groups $J^\sharp(\check{Y}; \Delta_k)$. The local system is a system of R/m^4–modules, and this ring has an m–adic filtration, of finite length, as does the module m^k/m^{k+4}. The latter filtration has associated graded

$$\mathrm{gr}_p(\delta_k) = \begin{cases} m^{k+p}/m^{k+p+1} & \text{if } 0 \le p \le 3, \\ 0 & \text{if } p \ge 4. \end{cases}$$

For any $\boldsymbol{Y}$, we have the usual spectral sequence,

$$J^\sharp(\check{Y}) \otimes \mathrm{gr}(\delta_k) \Rightarrow \mathrm{gr}\, J^\sharp(\check{Y}; \Delta_k), \tag{51}$$

and the assertion that the differentials d_r are zero for $r \le 3$ in the original spectral sequence of Proposition 6.1 is equivalent to saying that *all* the differentials in the spectral sequence (51) are zero. This in turn is equivalent to an equality of dimensions of finite-dimensional $\mathbb{F}$–vector spaces,

$$\dim J^\sharp(\check{Y}; \Delta_k) = \dim J^\sharp(\check{Y}) \times \dim \delta_k. \tag{52}$$

To prove the equality (52) and so complete the proof of Proposition 6.13, we will draw on the material of Section 2.2, so we again consider the orbifold (S^3, H) corresponding to the Hopf link, and the marking data μ_H with $w_2(E_\mu)$ nonzero. In order to keep the notation more compact, we write

$$\boldsymbol{H} = ((S^3, H); \mu_H)$$

for this marked orbifold, and we similarly introduce

$$\boldsymbol{T} = ((S^3, \theta); \mu_\theta).$$

We write $\boldsymbol{Y} = (\check{Y}, \mu_Y)$ for an arbitrary auxiliary orbifold, with possibly empty marking. We consider the instanton homology group $J(\boldsymbol{H} \# \boldsymbol{T})$, and the variant with local coefficients, $J(\boldsymbol{H}\,\#\,\boldsymbol{T};\Gamma)$, where the local system as usual is obtained from the marked (S^3,θ) summand by the usual circle-valued functions.

Lemma 6.14 *The Morse complex whose homology is $J(\boldsymbol{H}\,\#\,\boldsymbol{T};\Gamma)$ is quasi-isomorphic to a free module of rank 4 with the differential given by*

$$\partial_\Gamma = \begin{pmatrix} 0 & 0 & P & 0 \\ 0 & 0 & 0 & P \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

In particular, $J(\boldsymbol{H}\,\#\,\boldsymbol{T};\Gamma)$ is isomorphic to a sum of two copies of $R/(P)$.

Proof The proof uses a skein exact sequence, and is very close to the calculation of $J^\sharp(L;\Gamma)$ for the "linked handcuffs" in Proposition 6.11. The skein sequence that holds when the crossing is entirely contained in the marking region is the usual skein sequence for links, as developed in [10]. We deduce that the Morse complex which computes $J(\boldsymbol{H}\,\#\,\boldsymbol{T};\Gamma)$ is quasi-isomorphic to the mapping cylinder of a certain chain map

$$C(\boldsymbol{U}\,\#\,\boldsymbol{T};\Gamma) \to C(\boldsymbol{U}\,\#\,\boldsymbol{T};\Gamma), \tag{53}$$

where $\boldsymbol{U}$ is the unknot. The chain map arises from the same cobordism as in the proof of Proposition 6.11, namely the connect sum of the cylindrical cobordism with $(S^4,\mathbb{RP}^2)$, where the $\mathbb{RP}^2$ has self-intersection -2. Now, however, the marking region encompasses the whole of the cobordism, and the marking data has w_2 nonzero on the complement of the $\mathbb{RP}^2$ in S^4. This map (53) is the operator corresponding to the class

$$\tilde{\mathsf{z}} = w_2(\widetilde{\mathbb{W}}_s),$$

just as in Proposition 6.11, but now $\widetilde{\mathbb{W}}_s$ is the *orientable* rank-2 bundle corresponding to a basepoint on $\boldsymbol{U}$. This is the same as the basepoint operator $w_{2,x}$, which acts by multiplication by P. The instanton homology $J(\boldsymbol{U}\,\#\,\boldsymbol{T};\Gamma)$ is free of rank 2, and arises from a complex with trivial differential, because the representation variety is a 2–sphere. (This is the calculation of $I^\sharp(U)$ from [10].) So the complex $C(\boldsymbol{H}\,\#\,\boldsymbol{T};\Gamma)$ is quasi-isomorphic to the mapping cone of multiplication by P,

$$R \oplus R \xrightarrow{P} R \oplus R,$$

which is what the lemma states. □

Corollary 6.15 *With the local coefficient system Δ_k, the instanton homology*

$$J(\boldsymbol{H} \# \boldsymbol{T}; \Delta_k)$$

is the direct sum of four copies of the module δ_k.

Proof The calculation (47) shows that the element P belongs to the ideal $m^4 \lhd R$. So the action of P on δ_k is zero. The previous lemma therefore tells us that the Morse complex that computes $J(\boldsymbol{H} \# \boldsymbol{T}; \Delta_k)$ is four copies of the fiber δ_k and the differential is zero. □

To continue the proof of Proposition 6.13, we now pass from the special case $\boldsymbol{H} \# \boldsymbol{T}$ to the general case $J^\sharp(\boldsymbol{Y}) = J(\boldsymbol{Y} \# \boldsymbol{T})$, by excision. As in Section 2.2, there are excision cobordisms of marked bifolds, from (a) to (b) and from (b) to (c), inducing isomorphisms on J in each case, where (a)–(c) are

(a) $$\boldsymbol{Y}_a = (\boldsymbol{T} \# \boldsymbol{H}) \cup (\boldsymbol{T} \# \boldsymbol{Y}) \cup (\boldsymbol{H}),$$

(b) $$\boldsymbol{Y}_b = (\boldsymbol{T}) \cup (\boldsymbol{T} \# \boldsymbol{H} \# \boldsymbol{Y}) \cup (\boldsymbol{H}),$$

(c) $$\boldsymbol{Y}_c = (\boldsymbol{T}) \cup (\boldsymbol{T} \# \boldsymbol{H}) \cup (\boldsymbol{H} \# \boldsymbol{Y}).$$

In each case, two of the connected components contain a marked theta graph $\boldsymbol{T}$, so for each of (a)–(c) we may define a map to the torus T^3 by using the sum of the maps coming from the two copies. In this way, we have a local system Γ over the configuration spaces in all three cases. The maps obtained from the excision cobordisms give maps on instanton homology with local coefficients, and the composite of the two gives an isomorphism

$$J(\boldsymbol{Y}_a; \Delta_k) \cong J(\boldsymbol{Y}_c; \Delta_k).$$

The contribution of $\boldsymbol{H}$ to the calculation of the instanton homology of $\boldsymbol{Y}_a$ is trivial, and the Morse complex $C(\boldsymbol{Y}_a; \Delta_k)$ that computes $J(\boldsymbol{Y}_a; \Delta_k)$ can be described as a tensor product,

$$C(\boldsymbol{T} \# \boldsymbol{H}; \Delta_k) \otimes_R C(\boldsymbol{T} \# \boldsymbol{H}; \Gamma).$$

Corollary 6.15 therefore gives an isomorphism

$$C(\boldsymbol{Y}_a; \Delta_k) \cong C(\boldsymbol{T} \# \boldsymbol{Y}; \Delta_k)^{\oplus 4},$$

and hence an isomorphism

$$J(\boldsymbol{Y}_a; \Delta_k) \cong J(\boldsymbol{T} \# \boldsymbol{Y}; \Delta_k)^{\oplus 4}.$$

On the other hand, the contribution of the first T in the calculation for the instanton homology of Y_c is trivial, and the local coefficient system comes only from the $T \# H$ term. So the Morse complex has a description,

$$C(Y_c; \Delta_k) \cong (\delta_k)^{\oplus 4} \otimes_{\mathbb{F}} C(H \# Y).$$

So we have isomorphisms

$$J(Y_c; \Delta_k) \cong (\delta_k)^{\oplus 4} \otimes_{\mathbb{F}} J(H \# Y) \cong (\delta_k)^{\oplus 4} \otimes_{\mathbb{F}} J(T \# Y).$$

Comparing these expressions for Y_a and Y_c, we see that

$$\dim J(T \# Y; \Delta_k) = \dim J(T \# Y) \times \dim(\delta_k),$$

which becomes the desired inequality (52) when we take Y to be the unmarked bifold $\check{Y}$. This completes the proof of Proposition 6.13. □

References

[1] **K Appel**, **W Haken**, *Every planar map is four colorable*, Contemp. Math. 98, Amer. Math. Soc., Providence, RI (1989) MR

[2] **M F Atiyah**, **I M Singer**, *Dirac operators coupled to vector potentials*, Proc. Nat. Acad. Sci. U.S.A. 81 (1984) 2597–2600 MR

[3] **P J Braam**, **S K Donaldson**, *Floer's work on instanton homology, knots and surgery*, from "The Floer memorial volume" (H Hofer, C H Taubes, A Weinstein, E Zehnder, editors), Progr. Math. 133, Birkhäuser, Basel (1995) 195–256 MR

[4] **S K Donaldson**, *Connections, cohomology and the intersection forms of 4–manifolds*, J. Differential Geom. 24 (1986) 275–341 MR

[5] **S K Donaldson**, *Polynomial invariants for smooth four-manifolds*, Topology 29 (1990) 257–315 MR

[6] **D Eisenbud**, *Commutative algebra: with a view toward algebraic geometry*, Graduate Texts in Mathematics 150, Springer (1995) MR

[7] **A Floer**, *Instanton homology, surgery, and knots*, from "Geometry of low-dimensional manifolds, 1" (S K Donaldson, C B Thomas, editors), London Math. Soc. Lecture Note Ser. 150, Cambridge Univ. Press (1990) 97–114 MR

[8] **P Kronheimer**, **T Mrowka**, *Monopoles and three-manifolds*, New Mathematical Monographs 10, Cambridge Univ. Press (2007) MR

[9] **P Kronheimer**, **T Mrowka**, *Knots, sutures, and excision*, J. Differential Geom. 84 (2010) 301–364 MR

[10] **P B Kronheimer**, **T S Mrowka**, *Khovanov homology is an unknot-detector*, Publ. Math. Inst. Hautes Études Sci. 113 (2011) 97–208 MR

[11] **P B Kronheimer**, **T S Mrowka**, *Knot homology groups from instantons*, J. Topol. 4 (2011) 835–918 MR

[12] **P B Kronheimer**, **T S Mrowka**, *Gauge theory and Rasmussen's invariant*, J. Topol. 6 (2013) 659–674 MR

[13] **P B Kronheimer**, **T S Mrowka**, *Exact triangles for $SO(3)$ instanton homology of webs*, J. Topol. 9 (2016) 774–796 MR

[14] **P B Kronheimer**, **T S Mrowka**, *Tait colorings, and an instanton homology for webs and foams*, J. Eur. Math. Soc. 21 (2019) 55–119 MR

[15] **J Petersen**, *Die Theorie der regulären graphs*, Acta Math. 15 (1891) 193–220 MR

[16] **Y Xie**, *On the framed singular instanton Floer Homology from higher rank bundles*, PhD thesis, Harvard University (2016) MR `https://search.proquest.com/docview/1931403147`

Department of Mathematics, Harvard University
Cambridge, MA, United States

Department of Mathematics, Massachusetts Institute of Technology
Cambridge, MA, United States

`kronheim@math.harvard.edu`, `mrowka@math.mit.edu`

Proposed: Simon Donaldson — Received: 8 December 2017
Seconded: András I Stipsicz, Ian Agol — Revised: 20 March 2018