Symmetry in monotone Lagrangian Floer theory

Jack Edward Smith

Trinity College
University of Cambridge

This dissertation is submitted for the degree of Doctor of Philosophy

September 2017

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as specified in the text.

It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution.

Abstract

In this thesis we study the self-Floer theory of a monotone Lagrangian submanifold L of a closed symplectic manifold X in the presence of various kinds of symmetry.

First we consider the group $\mathrm{Symp}(X, L)$ of symplectomorphisms of X preserving L setwise, and extend its action on the Oh spectral sequence to coefficients of arbitrary characteristic, working over an enriched Novikov ring. This imposes constraints on the differentials in the spectral sequence which force them to vanish in certain situations.

We then specialise to the case where L is K-homogeneous for a compact Lie group K, meaning roughly that X is Kähler, K acts on X by holomorphic automorphisms, and L is a Lagrangian orbit. By studying holomorphic discs with boundary on L we compute the image of low codimension K-invariant subvarieties of X under the length zero closed–open string map. This places restrictions on the self-Floer cohomology of L which generalise and refine the Auroux–Kontsevich–Seidel criterion. These often result in the need to work over fields of specific positive characteristics in order to obtain non-zero cohomology. The disc analysis is then developed further, with the introduction of the notion of poles and a reflection mechanism for completing holomorphic discs into spheres.

This theory is applied to two main families of examples. The first is the collection of four Platonic Lagrangians in quasihomogeneous threefolds of $\mathrm{SL}(2, \mathbb{C})$, starting with the Chiang Lagrangian in $\mathbb{CP}^3$. These were previously studied by Evans and Lekili, who computed the self-Floer cohomology of the latter. We simplify their argument, which is based on an explicit construction of the Biran–Cornea pearl complex, and deal with the remaining three cases.

The second is a family of $\mathrm{PSU}(n)$-homogeneous Lagrangians in products of projective spaces. Here the presence of both discrete and continuous symmetries leads to some unusual properties: in particular we obtain non-displaceable monotone Lagrangians which are narrow in a strong sense. We also discuss related examples including applications of Perutz's symplectic Gysin sequence and quilt functors.

The thesis concludes with a discussion of directions for further research and a collection of technical appendices.

Acknowledgements

My first thanks must go to Ivan Smith, whose supervision has provided encouragement, wisdom and direction at the right moments, always with care and wit. He has been extremely generous with his time, insights and feedback, and for that I am truly grateful.

Jonny Evans and Yankı Lekili set this circle of ideas in motion with their paper on the Chiang Lagrangian, and I thank them for their continued interest in my work and for many valuable discussions. Special thanks are also due to Oscar Randal-Williams, through conversations with whom the idea of viewing a spectral sequence as a representation emerged, to Nick Sheridan for his clear and detailed explanations of numerous technical points, to Dmitry Tonkonog for sharing so much of his knowledge and experience with me, to Frol Zapolsky for answering my frequent questions about his work (any remaining mistakes in my understanding are my own), and to my examiners—Paul Biran and Ailsa Keating—for their helpful comments and suggestions.

Much of the material in this thesis is taken from my first two papers [135] and [134], and the former was improved greatly by the contributions of the anonymous referee. This work has also benefited from the comments of Mohammed Abouzaid, Denis Auroux, Benjamin Barrett, François Charette, Kenji Fukaya, Sheel Ganatra, Luis Haug, Michael Hutchings, Momchil Konstantinov, Stuart Martin, Kaoru Ono, Tim Perutz, Tobias Sodoge, Brunella Torricelli and Chris Woodward. Wolfram Mathematica has been an invaluable tool for algebraic manipulation and experimentation.

I am grateful to EPSRC for funding my PhD, and to Trinity College, Cambridge for their hospitality over the last seven years.

Finally, I would like to thank my teachers (in the broadest sense), friends, and family—especially my dad—for their kindness and patience. This thesis is dedicated to my parents and sister.

Contents

Chapter 1

Introduction

1.1 History and context

1.1.1 Algebra and geometry

A common technique in geometry and topology is to take a geometric object which one wants to understand and build from it an algebraic surrogate, which captures some relevant features and through which questions about the original object can be answered. Prototypical examples include singular cohomology in algebraic topology and sheaf cohomology in algebraic and complex geometry.

In order for this approach to be most effective, one needs two key inputs. First, a conceptual understanding—what type of algebraic object one is dealing with, how it is constructed, and what properties it enjoys. And second, computational tools and examples in order to apply the technology in concrete situations.

Of course, these two projects are closely related: the theory will most likely be motivated by specific problems or empirical insights, whilst rigorous computations are often impossible without the use of some abstract machinery, like an exact sequence or vanishing theorem; indeed, a recurring theme in modern mathematics is that the properties of a construction are often more useful than a direct definition. They do, however, have slightly different flavours, and it is possible for developments in one to outpace the other.

Section 1.1 sketches how this approach has been applied in symplectic topology, and gives an informal overview of Floer theory. The contribution of this thesis is outlined in Section 1.2.

1.1.2 Pseudoholomorphic curves

Although symplectic geometry has been around, implicitly at least, since Lagrange's papers on mechanics in the early nineteenth century [101], it wasn't until the visionary work of Floer in the late 1980's that the first algebraic invariants of a truly symplectic nature emerged [46, 47]. Inspired by the Arnol'd conjecture [23, Problem XX] bounding the number of fixed points of Hamiltonian diffeomorphisms, and building on Gromov's groundbreaking introduction of pseudoholomorphic curves in symplectic manifolds [78], Floer defined a cohomology theory for a symplectic manifold X—now known as quantum cohomology—by counting certain such curves connecting Hamiltonian orbits. By showing that the resulting cohomology group $QH^*(X)$ is isomorphic to singular cohomology, Floer was able to prove the Arnol'd conjecture under the assumption that $\pi_2(X)$ vanishes.

This hypothesis is rather restrictive, but subsequently Floer made the important observation that his construction could be extended to symplectic manifolds X which are merely *monotone*, meaning that the area and first Chern class homomorphisms $\pi_2(X) \to \mathbb{R}$ are positively proportional [51]. The purpose of this condition is to establish the compactness of the moduli spaces of pseudoholomorphic curves which one is counting: Gromov's paper introduced his celebrated *compactness theorem* [78, Theorem 1.5.B], which roughly asserts that the moduli spaces of curves of bounded area can be compactified by adjoining boundary strata comprising 'bubbled' configurations, and monotonicity both provides the area bound and ensures that the boundary strata have sufficiently high codimension that they can be ignored (vanishing of $\pi_2(X)$, meanwhile, guarantees that the boundaries are actually empty).

Even before the pioneering work of Gromov and Floer, it had been observed that where symplectic manifolds arise, Lagrangian submanifolds are often objects of natural significance and interest. Perhaps most strikingly, Lagrangian submanifolds of the product

$$(X_1, -\omega_1) \times (X_2, \omega_2)$$

behave in some ways like morphisms from (X_1, ω_1) to (X_2, ω_2), generalising the graphs of symplectomorphisms in the case where $X_1 = X_2$ [80, 143] (this idea has been systematically developed recently in the work of Wehrheim–Woodward [142], Ma'u–Wehrheim–Woodward [102], Lekili–Lipyanskiy [98] and, possibly definitively, Fukaya [56]; see Section 5.2.2 for a brief discussion). In fact, Weinstein even went as far as to formulate his 'symplectic creed' [143]: *everything is a Lagrangian submanifold*. It is not surprising therefore that Lagrangians arose naturally as boundary conditions for Gromov's pseudoholomorphic curves, allowing him to prove for example that certain Lagrangians were non-displaceable from each other by Hamiltonian diffeomorphisms [78, Theorem 2.3.B$'_3$]. In fact, by viewing fixed points of a symplectomorphism as intersections of its graph with the diagonal, Gromov used this result to deduce a weak form of the Arnol'd conjecture (namely that Hamiltonian diffeomorphisms have at least one fixed point, assuming ω vanishes on π_2).

One of Floer's motivational insights was that the quantum cohomology of X can be interpreted as the Morse homology of its space of contractible free loops, for the functional which assigns to a loop γ the area of a capping disc $\widehat{\gamma}$, perturbed by the integral of a (possibly time-dependent) Hamiltonian function along γ [49, 50, 52]. Transferring his theory to the space of paths connecting two Lagrangians L_0 and L_1, in [48] he constructed an invariant—*Lagrangian Floer cohomology* $HF^*(L_0, L_1)$—which is the central object of study in this thesis. In simple terms, it is generated by intersection points of the Lagrangians and the differential counts pseudoholomorphic strips between them. Floer worked under a hypothesis of vanishing relative π_2, but the extension to the monotone setting followed shortly afterwards in the work of Oh [109].

We have glossed over many technical details here, and will continue to do so for the rest of this introduction. However, the following issues are worthy of comment. The moduli spaces occurring in the construction of Lagrangian Floer cohomology do not carry canonical orientations (in contrast to the case of quantum cohomology), so in general one has to work over coefficient rings of characteristic 2. In order to move outside this setting, which is crucial in many situations and is a key part of this thesis, one has to assume that the Lagrangians involved are relatively spin (or something similar), and fix choices of relative spin structures. In addition to this complication, once one weakens the assumption of vanishing relative π_2 (or the essentially equivalent condition of *exactness*) the Floer complex may be *obstructed*, meaning that the differential no longer squares to zero. To each Lagrangian L there is an associated count $\mathfrak{m}_0(L)$ of index 2 pseudoholomorphic discs through a generic point, and if L_0 and L_1 are monotone then the complex $CF^*(L_0, L_1)$ is unobstructed if and only if $\mathfrak{m}_0(L_0) = \mathfrak{m}_0(L_1)$. This is automatically the case when L_0 and L_1 are both Hamiltonian isotopic to a single monotone Lagrangian L, and we call the resulting $HF^*(L_0, L_1)$ the *self-Floer cohomology* of L.

1.1.3 Modern Floer theory

The intervening years have seen enormous advances in Lagrangian Floer theory. First, there have been technical developments in the foundations of pseudoholomorphic curve theory, and in particular in the construction of virtual fundamental cycles on moduli spaces, which have in principle eliminated the monotonicity hypotheses from the definitions, albeit at the expense of considerable hard work and ingenuity. These will not concern us, but it would be remiss not to mention the monumental work in this direction by Fukaya–Oh–Ohta–Ono [63], based on the use of Kuranishi structures, which were introduced by Fukaya–Ono [70] to extend Floer's proof of the Arnol'd conjecture to the general case (which was also achieved independently, at the same time and along similar lines, by Liu–Tian [99]). Recently Pardon has introduced a more abstract, algebraic approach [114], which should enable significant simplification and unification of such constructions.

Second, an increasingly rich algebraic structure has been built around Floer cohomology groups and their underlying cochain complexes. By counting pseudoholomorphic triangles, Donaldson introduced an associative product

$$HF^*(L_1, L_2) \otimes HF^*(L_0, L_1) \to HF^*(L_0, L_2) \tag{1.1}$$

for triples of Lagrangians, described in detail in the thesis of de Silva [39], which makes each group $HF^*(L, L)$ into unital ring. He also observed that this enables the Lagrangian submanifolds of X to be assembled into a category—the *Donaldson category*—with $HF^*(L_0, L_1)$ forming the morphism space $\mathrm{Hom}(L_0, L_1)$. This explains the slightly unusual ordering of the factors on the left-hand side of (1.1).

Fukaya [57] then realised that this construction could be generalised from triangles to other polygons in order to define higher multiplication operations on the Floer cochain complexes, which satisfy the A_∞-relations. The Donaldson category is thus merely a shadow of a truly formidable object: an A_∞-category, defined only up to an appropriate notion of homotopy equivalence (since the constructions are at chain level), now known as the *Fukaya category* of X. The use of categorical methods to answer symplectic-topological questions has been pioneered by Seidel and his collaborators and proved to be very fruitful [121, 122, 123, 54, 126]. Studying a symplectic manifold in this way is somewhat analogous to studying a ring through its category of modules, or, more generally, a variety through its category of sheaves (more of which later).

Of particular interest to us is the *closed–open string map*

$$\mathcal{CO}\colon QH^*(X) \to HH^*(\mathcal{F}(X)),$$

a unital ring homomorphism from the quantum cohomology of X (this carries a 'pair of pants' product, which is the loop space version of the Morse-theoretic cup product counting Y-shaped trajectories) to the Hochschild cohomology of (an appropriate version of) the Fukaya category of X. More specifically, we will be interested in the projection $\mathcal{CO}^0$ to Hochschild cochains of length zero, which concretely amounts to a unital ring homomorphism

$$\mathcal{CO}^0\colon QH^*(X) \to HF^*(L, L)$$

for each L, making the Donaldson category $QH^*(X)$-linear. We study this length-zero closed–open map in detail in Section 3.2.

The other key piece of structure for us is the Oh spectral sequence [111]

$$E_1 \cong H^*(L) \implies HF^*(L, L), \tag{1.2}$$

whose differentials count certain configurations of pseudoholomorphic discs with boundary on L. In the exact setting, where there can be no non-constant discs, this spectral sequence degenerates at the first page and we obtain an additive isomorphism between the singular and self-Floer cohomologies of L (because of the extension problem relating E_∞ to $HF^*(L, L)$, strictly this argument only holds over fields). Taking L to be the diagonal in $X^- \times X$, where $^-$ denotes reversal of the sign of the symplectic form, this recovers Floer's isomorphism between the singular and quantum cohomologies of X. This spectral sequence is discussed at length in Section 2.2.

The third major front of advance in Floer theory has been in applications and connections with other fields. Applications within symplectic topology are so profuse that we can give no more than a brief flavour, but two recurring themes are topological restrictions on Lagrangian submanifolds and the construction of interesting elements in symplectomorphism groups. In the first vein, for monotone Lagrangians there are classic results due to Seidel [120] and Biran–Cornea [14], as well as more recent developments by Damian [38] and others, but much is still unknown. For (closed) exact Lagrangians in cotangent bundles, meanwhile, there has been remarkable progress towards the *nearby Lagrangian conjecture*, which states that any such Lagrangian is Hamiltonian isotopic to the zero section. This is a long story, but key steps were taken by Viterbo [140], Seidel [122], Fukaya–Seidel–Smith [71], Nadler [106] (by an approach *not* based on Floer theory; the fact that this caused such a stir is indicative of how ubiquitous Floer-theoretic techniques have become), Abouzaid [3] and Kragh [96]. In the second vein, see the papers of Seidel on Dehn twists [124] and on invertibles in quantum cohomology [119].

As with developments on the technical side, the connections with other fields are not directly relevant to the present work, but as an illustration we point out that Lagrangian Floer cohomology now plays a central role in low dimensional topology and knot theory, underlying, for instance, the Heegaard Floer homology of Ozsváth–Szabó [113], the closely related knot Floer homology of Ozsváth–Szabó [112] and Rasmussen [118], and the symplectic Khovanov homology of Seidel–Smith [128].

We must also mention the revolutionary *homological mirror symmetry* conjecture of Kontsevich [95], which asserts that the analogy between Fukaya categories and categories of sheaves is not merely superficial: Calabi–Yau symplectic manifolds and complex manifolds should occur in mirror pairs, X and $\check{X}$, such that the split-closed derived Fukaya category $D^\pi\mathcal{F}(X)$ of one is equivalent to the derived category of coherent sheaves $D^b\mathrm{Coh}(\check{X})$ on the other. This conjecture—and its subsequent extensions to other (i.e. non-Calabi–Yau) settings, typically involving many bells, whistles and subtleties—has inspired a great deal of work in symplectic topology, both in rigorously establishing the equivalence in various situations, and in exploring its potential consequences. For a small sample of this, see the papers of Polishchuk–Zaslow [117], Abouzaid–Smith [4], Seidel [125, 127], Sheridan [129] and Fukaya–Oh–Ohta–Ono [68], which consider respectively the 2-torus, 4-torus, genus 2 curve, quartic surface, general Calabi–Yau hypersurfaces in projective space, and compact toric manifolds. For a more leisurely overview, see the surveys of Auroux [9] and Smith [133].

In hindsight, with a view of the enormous breadth and sophistication of the techniques involved, it is not surprising that the algebraic revolution arrived later in symplectic topology than in other areas of geometry. To take a rather crude measure, the definition of singular cohomology can be given in a single page, whilst the full construction of Lagrangian Floer cohomology by Fukaya–Oh–Ohta–Ono stretches to eight hundred. Of course, it is possible that this gulf in apparent complexity may shrink once the underlying machinery is fully integrated into the consciousness of the working mathematician, but the current state of Floer theory nevertheless represents a landmark achievement of modern mathematics.

1.1.4 Existing computations

Despite all of the conceptual advances, and the central role it plays in symplectic topology, the self-Floer cohomology of a Lagrangian submanifold L of a symplectic manifold X remains extremely difficult to compute in most cases. There is no general procedure even to decide whether or not it vanishes, or equivalently whether or not L is (quasi-isomorphic to) the zero object in the Fukaya category of X.

In cases where the calculation of $HF^*(L, L)$ *has* been possible, a common strategy is to consider the Oh spectral sequence (1.2), and combine knowledge of $E_1 \cong H^*(L)$ with enough understanding of pseudoholomorphic discs on L to see how the differentials behave. This is particularly effective in the case of Lagrangian torus fibres in toric varieties (see the early papers of Cho [32] and Cho–Oh [35]), where the spectral sequence degenerates at the second page, since the classical cohomology is generated as a ring by classes of degree at most 1, and the discs can be explicitly written down. Or in examples where for purely topological reasons no discs can contribute, for example when L is exact (this is essentially Floer's original result [48]), or of high minimal Maslov index [111, Theorem II], [14, Proposition 6.1.1]. The product on Floer cohomology, and its compatibility with the spectral sequence (originally proved by Buhovsky [24]), is often a key ingredient.

These arguments are usually a matter of coupling general theory with techniques specific to the situation at hand. One such technique is applicable when L is the fixed locus of an antisymplectic involution τ of X, for example the real locus of an algebraic variety defined by real equations. This has been studied by Givental [75], Oh [108], Frauenfelder [53], Fukaya–Oh–Ohta–Ono [69], Haug [83] and others. Heuristically, the idea is that pseudoholomorphic discs should pair up with their reflections under τ and hence cancel modulo 2, so over a coefficient field R of characteristic 2 all differentials in the Oh spectral sequence vanish, and we obtain an additive isomorphism $HF^*(L, L) \cong H^*(L)$.

Remark 1.1.1. Note that the spectral sequence tells us that $H^*(L)$ is 'the largest' $HF^*(L, L)$ can be, and L is called *wide* if this bound is attained (this is formulated precisely in Definition 2.2.1). At the other extreme, there are Lagrangians with $HF^*(L, L) = 0$; these are called *narrow* over the ground ring. Working over a field it is an open conjecture that every monotone Lagrangian (with irreducible local system) is either wide or narrow [16, Conjecture 1].

This conjecture is known to be false if the conditions are weakened: if one works over $\mathbb{Z}$, rather than a field, the real projective spaces $L = \mathbb{RP}^{2N+1} \subset \mathbb{CP}^{2N+1}$ have $\mathbb{Z}/(2N + 2)$-graded self Floer cohomology which is isomorphic to $\mathbb{Z}/2$ in even degrees and 0 in odd degrees [69, Section 6.4], whilst if one drops the monotonicity condition then a counterexample has been constructed by Fukaya–Oh–Ohta–Ono [62, Theorem 56.32]. //

The literature also contains a handful of direct computations of self-Floer cohomology in the monotone setting, based on the *pearl complex* of Biran–Cornea [14]. This is an alternative model for computing $HF^*(L, L)$ in which, morally, we take the limit of $CF^*(L, \varphi(L))$ as the Hamiltonian diffeomorphism φ tends to the identity. Sufficiently close to this limit we can view $\varphi(L)$ as the graph of an exact 1-form $\mathrm{d}f$ in a Weinstein neighbourhood of L, and the generators of the Floer complex—which are intersection points of L and $\varphi(L)$—are simply critical points of f. Meanwhile, the pseudoholomorphic strips counted by the differential degenerate to Morse flowlines of f interrupted by pseudoholomorphic discs. The power of this model is that one can work with a single Lagrangian L, rather than the pair $(L, \varphi(L))$, as the boundary condition for pseudoholomorphic curves, and in particular one does not need to break any symmetries of L by making an explicit choice of φ. (Its close relationship with the Morse complex also provides the cleanest construction of the Oh spectral sequence.)

Remark 1.1.2. An alternative viewpoint is that whilst the standard Floer complex formally computes the Morse homology of the space of contractible paths from L to itself, for the functional which assigns to a path γ the area of a capping half-disc $\widehat{\gamma}$ perturbed by the integral of a Hamiltonian function along γ, the pearl complex computes the Morse–Bott version where this Hamiltonian is zero. //

Using the pearl complex, and explicit enumeration of holomorphic discs for a special integrable complex structure, Chekanov–Schlenk [30] and Auroux [9] computed the self-Floer cohomology of the (non-toric) *Chekanov tori* in $\mathbb{CP}^1 \times \mathbb{CP}^1$ and $\mathbb{CP}^2$ respectively (see Section 5.2.2 and Section 5.2.3 for further discussion), whilst more recently Evans–Lekili [45] computed the self-Floer cohomology of the *Chiang Lagrangian* [31, Section 2] $L_\triangle$ in $\mathbb{CP}^3$. The disc analysis for the Chekanov tori is based on the use of Lefschetz fibrations, whilst that for the Chiang Lagrangian involves novel techniques introduced by Evans–Lekili to study *homogeneous Lagrangians*. Roughly speaking, these are Lagrangian orbits of Lie group actions on Kähler manifolds, and the work of Evans and Lekili was the first investigation of how much of the toric theory extends to this more general case.

Beyond these computations, there is one major structural result in monotone Floer theory which requires comment. This is the *Auroux–Kontsevich–Seidel criterion* [9, Proposition 6.8], which states that the self-Floer cohomology of an orientable monotone Lagrangian L can be non-zero over a field R only if the count $\mathfrak{m}_0(L)$ of index 2 discs is an eigenvalue of the endomorphism of $QH^*(X;R)$ given by quantum multiplication $c_1(X)*$ by the first Chern class of X. This is proved by showing that $\mathcal{CO}^0$ sends $c_1(X)$ to $\mathfrak{m}_0(L)$ times the unit 1_L in $HF^*(L,L)$ [9, Proposition 6.7], [130, Lemma 2.7].

Remark 1.1.3. The closed–open map has also been computed on 'Seidel elements' in quantum cohomology (which have in turn been calculated in certain situations by McDuff–Tolman [105]) by Charette–Cornea [28] and Tonkonog [137]. //

One surprising feature of the Evans–Lekili calculation is that $HF^*(L_\triangle, L_\triangle)$ is non-zero over a field R if and only if $\operatorname{char} R = 5$, and this is partially explained by these eigenvalue considerations [45, Remark 1.2]: $\mathfrak{m}_0(L_\triangle)$ is 3, and this is an eigenvalue of $c_1(\mathbb{CP}^3)* = 4H*$ acting on $QH^*(\mathbb{CP}^3;R) = R[H]/(H^4-1)$ only if R has characteristic 5 or 7. It is natural to wonder whether there is a simple way in which the 7 can be ruled out.

1.2 The content of this thesis

1.2.1 Aims

Whilst self-Floer cohomology has been determined in a variety of cases, some of which we have just seen, we are still a long way from a general understanding, and are lacking a diverse armoury of systematic computational tools. The aim of this thesis is to provide new techniques and examples to address this problem. A secondary aim is that these techniques should reveal some underlying structure and suggest new directions for theoretical development. We restrict attention throughout to the case where L is monotone. This is to avoid technical baggage and allow us to focus on the geometry.

The unifying theme of our study is *symmetry*. Many naturally occurring Lagrangians possess symmetry of some kind, and our strategy is to exploit it to simplify the analysis of their Floer theory. The use of symmetry to render problems more tractable is of course widespread in mathematics, perhaps especially in applied mathematics and theoretical physics, and it has already made an appearance in Section 1.1.4, in the form of Lie group actions of which L is an orbit (in the toric and general homogeneous cases), and antisymplectic $\mathbb{Z}/2$-actions fixing L. It is therefore a natural starting point for developing new methods of calculation.

Our main results utilise related symmetries of L, both discrete and continuous, in order to study $HF^*(L,L)$. We show how these methods can provide novel perspectives on well-known calculations, and apply them to study new examples which display some surprising properties. The picture that emerges is that discrete symmetries, like an antisymplectic involution, give a method for proving HF^* is non-zero (this underlies Theorem 5 and parts of Theorem 3), whilst continuous symmetries give rise to constraints—analogous to the Auroux–Kontsevich–Seidel criterion—which force HF^* to vanish (as illustrated by Theorem 2 and Theorem 4). In examples where both are present, the tension between them leads to subtle dependence on the coefficient ring and local system. A recurring idea is to make computations in a special integrable complex structure, where they are manageable, and then to argue that the answers agree with those for a *generic* almost complex structure; the first step is essentially geometric in flavour, whilst the second is technical.

1.2.2 Outline of results

The thesis is divided into two main parts: Chapters 2 and 3 cover most of the new theory, whilst Chapters 4 and 5 focus on two different families of examples. In Chapter 6 we briefly discuss some potential future developments suggested by our study, after which there follow four appendices.

Section 2.1 reviews the basics of monotone Floer theory, including the details of the Biran–Cornea pearl complex which is used throughout the thesis. This material is all standard, but we collect it together here in order to fix notation and terminology. In Section 2.2 we recall the construction of the Oh spectral sequence for a monotone Lagrangian $L \subset X$, spelling out the details as it is central in what follows. We then consider the group $\mathrm{Symp}(X, L)$ of symplectomorphisms of the ambient symplectic manifold X which preserve L setwise, and its action on the spectral sequence, which can be expressed purely topologically. This action was constructed in characteristic 2 by Biran–Cornea [14, Section 5.8], and by working over an enlarged Novikov ring $\Lambda^\dagger$ (also considered by Biran–Cornea [17, Section 2.1]) which records the homology classes of discs we extend their construction to arbitrary characteristic in Proposition 2.2.7. The modified coefficient ring is necessary in order to account for changes of relative spin structure on L under the action of $\mathrm{Symp}(X, L)$, but turns out to be extremely useful in its own right since it enables one to exploit the action of $\mathrm{Symp}(X, L)$ on the disc classes. The section concludes with a discussion of H_2^D local systems and B-fields, which allow more flexibility than conventional local systems and are used in our applications.

As an example we show how to reconstruct the results of Cho's computation [32] of the self-Floer cohomology of the Clifford torus in $\mathbb{CP}^n$ without ever having to think about holomorphic discs, just using its symmetries—see Remark 2.2.8. In this case the relevant differential in the spectral sequence has a concrete geometric interpretation as a sum of boundaries of index 2 discs, and one can argue directly about the shape of this class. The power of the spectral sequence method is that one can constrain the differentials even when they have no such obvious geometric meaning. It can be viewed as a coarse (but more general) version of the mod 2 cancellation in the presence of an antisymplectic involution: instead of establishing transversality results which allow one to prove cancellation at the level of actual pseudoholomorphic discs, we aim to achieve cancellation (over an appropriate coefficient ring; in our main examples this will have positive characteristic related to the orders of the symmetries, just as in the involution case) at the level of homology classes.

In Section 3.1 we reintroduce (with minor variations) the notion of a *K-homogeneous Lagrangian* L from [45], and discuss the complexifications of Lie groups and their actions—again this is standard but is summarised for later use. We then recall (again from [45]) the definition of an *axial disc* on such a Lagrangian L, which is a holomorphic disc whose boundary is swept by a one-parameter subgroup of K. These discs can be described and studied rather concretely, and are amenable to explicit enumeration in examples.

Section 3.2.2 contains a crucial new ingredient: we consider moduli spaces of holomorphic discs (for the standard integrable complex structure) with boundary on a K-homogeneous Lagrangian $L \subset X$, and prove that the evaluation map at an interior marked point is transverse to any K-invariant subvariety $Z \subset X$. This enables us to define smooth moduli spaces of holomorphic discs meeting Z and use dimension arguments to prove axiality under restrictions on the Maslov index. In Section 3.2.3 we apply these results to compute part of the closed–open map

$$\mathcal{CO}^0 \colon QH^*(X) \to HF^*(L, L).$$

Letting N_X^+ denote the minimal Chern number of (holomorphic) rational curves in X, and α in $H^{2k}(X) = QH^{2k}(X)$ denote the Poincaré dual class to Z (where k is the complex codimension of Z), the main result is as follows:

Theorem 1 (Corollary 3.2.21, Proposition 3.2.24)**.** *If $k \leq N_X^+ + 1$ then, setting the Novikov variable to 1, $\mathcal{CO}^0(\alpha)$ is a multiple $\lambda \cdot 1_L$ of the unit in $HF^*(L, L)$, where λ is a $\mathbb{Z}$-linear combination of monodromies of the local system on L. If the inequality is strict then this linear combination is simply the sum of the boundary monodromies of certain axial discs. Moreover, if $\dim K = \dim L$ and L is equipped with the 'standard' spin structure (see Definition 3.1.10) then all of the discs contribute to this count with positive sign.*

The proof of the orientation statement is slightly technical so is deferred to Appendix B. For monotone toric fibres Theorem 1 can be combined with standard disc classification results of Cho–Oh [35] to recover the well-known fact that with the trivial local system $\mathcal{CO}^0$ gives 1_L on each component of the toric divisor [64, Equation (6.27)] (this can also be shown by combining results of Charette–Cornea [28], Tonkonog [137] and McDuff–Tolman [105]; see Remark 5.1.8). In Corollary 3.3.49 we prove a generalisation of this fact without using the results of Cho–Oh.

Remark 1.2.1. Using recent work of Fukaya on equivariant Kuranishi structures [55] it is possible that under appropriate hypotheses the codimension condition is not needed to prove that $\mathcal{CO}^0(Z)$ is proportional to the unit [60]. However it really is necessary for the statement about axial discs—in Remark 3.3.51 we give an example with $k = N_X^+ + 1$ where the value of λ predicted by the theorem is incorrect because of certain bubbling. //

In Section 3.3 we develop the detailed analysis of holomorphic discs bounded by a *sharply* K-homogeneous Lagrangian L, for which $\dim K = \dim L$. In this case, the ambient complex manifold X is the union of an open orbit W of the complexification G of K and a compactification divisor Y. We focus on the *poles* of discs, which are the points where they meet Y, and by studying the behaviour at these points we are able to recover our earlier axiality results by local methods, rather than by global considerations about moduli spaces. Through the concept of *obliging* poles and subvarieties we are able to go much further, and significantly simplify our computation of $\mathcal{CO}^0$. For instance, if the K-action on L is free then the count appearing in Theorem 1 involves either no discs or a single disc, so with the standard spin structure and trivial local system we have either $\mathcal{CO}^0(\alpha) = 0$ or $\mathcal{CO}^0(\alpha) = 1_L$.

We also introduce an antiholomorphic involution on the open orbit W, fixing L pointwise, which extends to (complex) codimension 1 orbits in Y—which it need not preserve, and where it need not be involutive or even injective—but does not in general extend to the whole of X. Despite the fact that it is only partially defined, it nevertheless enables holomorphic discs on L to be reflected and thus completed to holomorphic spheres. This allows us to employ tools from algebraic geometry to count such curves, which is a crucial simplification.

Our first applications, in Chapter 4, are to the family of four 'Platonic' Lagrangian SU(2)-orbits L_C in a sequence of Fano threefolds X_C, parametrised by configurations C of points on the sphere (C can be a triangle $\triangle$, tetrahedron T, octahedron O, or icosahedron I, and the respective threefolds are $\mathbb{CP}^3$, the quadric, the threefold known as V_5, and the Mukai–Umemura threefold V_{22}). The first of these, with $C = \triangle$, is the Chiang Lagrangian studied by Evans–Lekili. Section 4.1 reviews the general construction of these Lagrangians and sets out their basic properties. By interpreting the antiholomorphic involution geometrically, we show that it does actually extend to the whole of X_C when C is the octahedron or icosahedron (but not when $C = \triangle$ or T), so L_O and L_I are automatically wide in characteristic 2 (in Section 4.1.6 we rigorously establish the mod 2 cancellation of discs in this setting).

Equip L_C with an arbitrary relative spin structure and the trivial local system to give a brane $L_C^\flat$. Applying Theorem 1, we show:

Theorem 2 (Proposition 4.2.10(ii))**.** *$HF^*(L_C^\flat, L_C^\flat; \mathbb{Z})$ is annihilated by 5, 4, 2 and 8 for C equal to $\triangle$, T, O and I respectively. In particular, if $HF^*(L_C^\flat, L_C^\flat; R)$ is non-zero over a field R then R has characteristic 5 if $C = \triangle$ and characteristic 2 in each of the other cases.*

The significance of characteristic 2 for the octahedron and icosahedron is clear in light of the involution. The characteristics are less obvious for the triangle and tetrahedron, but note that we have ruled out the characteristic 7 possibility for the former. The fact that 2 occurs again for the tetrahedron, making it appear to fall into the same pattern as the octahedron and icosahedron, with the triangle as the lone exceptional case, seems to be a coincidence arising from the fact that the numbers involved are fairly small. It also seems to be a coincidence that there is exactly one possible prime in each case.

We then go on and explicitly compute the self-Floer cohomology groups with the pearl complex, using our disc analysis to make the index 4 counts tractable. We obtain:

Theorem 3 (Proposition 4.1.10, Corollary 4.4.2, Corollary 4.4.5, Proposition 4.4.6)**.** *Working over a field R of characteristic 5, 2, 2 and 2 in the four cases respectively, the Floer cohomology groups are given as $\mathbb{Z}/2$-graded R-vector spaces by*

$$\begin{gathered} HF^0(L_\triangle^\flat, L_\triangle^\flat; R) \cong HF^1(L_\triangle^\flat, L_\triangle^\flat; R) \cong R \\ HF^0(L_T^\flat, L_T^\flat; R) \cong HF^1(L_T^\flat, L_T^\flat; R) \cong R \\ HF^0(L_O^\flat, L_O^\flat; R) \cong HF^1(L_O^\flat, L_O^\flat; R) \cong R^2 \\ HF^0(L_I^\flat, L_I^\flat; R) \cong HF^1(L_I^\flat, L_I^\flat; R) \cong R. \end{gathered}$$

Working over $\mathbb{Z}$, the $\mathbb{Z}/2$-graded Floer cohomology rings are concentrated in degree 0 with

$$HF^0(L_\triangle^\flat, L_\triangle^\flat; \mathbb{Z}) \cong \mathbb{Z}/(5)$$

$$HF^0(L_T^\flat, L_T^\flat; \mathbb{Z}) \cong \mathbb{Z}/(4)$$
$$HF^0(L_O^\flat, L_O^\flat; \mathbb{Z}) \cong \mathbb{Z}[x]/(2, x^2 + x + 1).$$

Using the techniques of this thesis it is also possible to prove that $HF^0(L_I^\flat, L_I^\flat; \mathbb{Z}) \cong \mathbb{Z}/(8)$, but this requires a lengthy analysis of degree 4 curves in X_I which we do not cover (it can be found in Appendix B of the paper [135] on which this chapter is based).

The results for $L_\triangle$ were proved by Evans–Lekili, but the others are new. In each case the Lagrangian is wide over fields of the special characteristic, whilst the Floer cohomology over $\mathbb{Z}$ is as large as is allowed by the restrictions we derive from the closed–open map. Note that $HF^0(L_O^\flat, L_O^\flat; \mathbb{Z})$ is the field $\mathbb{F}_4$ of four elements.

The results of Evans–Lekili imply that $L_\triangle$ is not Hamiltonian-displaceable from itself or from the monotone Clifford torus in $\mathbb{CP}^3$ [45, Corollary B] (recent work by Konstantinov [94, Corollary 1.2] using higher rank local systems shows that it is also non-displaceable from the standard $\mathbb{RP}^3$). Similarly the fact that L_T, L_O and L_I are Floer cohomologically non-trivial, with appropriate coefficients, immediately shows that they are also non-displaceable from themselves. In their subsequent paper [44, Section 7.1] Evans–Lekili showed that the real locus of the quadric X_T, which is a monotone Lagrangian sphere (homogeneous for a different SU(2)-action on X_T), split-generates the Fukaya category over any field R of characteristic 2, so L_T is not displaceable from this sphere either (note that the ring $QH^*(X_T; R)$ is isomorphic to $R[E]/(E^4)$, so already every element is invertible or nilpotent; in particular, the whole Fukaya category forms one of their summands $D^\pi\mathcal{F}(X_T; R)_0$). By [44, Corollary 6.2.6] we actually deduce that L_T also split-generates the Fukaya category of the quadric over R.

The second family of examples forms the bulk of Chapter 5. We construct, for each $N \geq 3$, a PSU$(N-1)$-homogeneous monotone Lagrangian L in $X = (\mathbb{CP}^{N-2})^N$, and use Theorem 1 to compute the closed–open map on the classes Poincaré dual to invariant subvarieties Z_I (indexed by subsets $I \subset \{1, \dots, N\}$) comprising those points $([z_1], \dots, [z_N])$ in X such that $([z_j])_{j \in I}$ is linearly dependent. We obtain equalities in $HF^*(L, L)$ involving symmetric polynomials in certain classes, from which we deduce:

Theorem 4 (Proposition 5.1.9)**.** *Working over a coefficient ring R of characteristic p (prime or 0), L has vanishing self-Floer cohomology with any relative spin structure and local system (possibly of higher rank) unless p is prime and N is a power of p or twice a power of p, or $N = 3$ and $p = 5$.*

In particular, from those N which are not prime powers or twice prime powers we obtain a large family of examples of monotone Lagrangians which are narrow in a rather robust sense: it isn't just that we are working over the wrong coefficient ring or with the wrong local system. To the best of the author's knowledge, the only previously known examples with this property, when the ambient manifold X is closed and simply connected, are the displaceable Lagrangians constructed by Oakley and Usher [107].

We then consider the action of the symmetric group S_N on (X, L) by permuting the N factors of X. Applying some representation theory to the induced action on the Oh spectral sequence, we obtain a criterion for wideness of L which implies:

Theorem 5 (Corollary 5.1.20, Corollary 5.1.22). *(i) If R is a field of prime characteristic p and N is either a power of p or twice a power of p then there exists a relative spin structure on L for which it is wide when equipped with the trivial local system over R.*

(ii) For all values of N there exists a B-field on X for which L is wide over $\mathbb{C}$. In particular, L is non-displaceable.

Thus our narrow examples are of a fundamentally different character from Oakley–Usher's. We also deduce that, even in the monotone case, ordinary Floer cohomology with local systems is not sufficient to detect non-displaceability. Combining our results we can almost completely determine the wide variety [17, Section 3] of L—see Remark 5.1.23.

In Section 5.2 we discuss various related results and constructions. First we focus on the $N = 3$ case of our main family, which is a Lagrangian $\mathbb{RP}^3$ in $(\mathbb{CP}^1)^3$, and an analogous Lagrangian in $\mathbb{CP}^1 \times \mathbb{CP}^2$. These examples are analysed using Perutz's symplectic Gysin sequence [116] and are related to the Chekanov tori in $\mathbb{CP}^1 \times \mathbb{CP}^1$ and $\mathbb{CP}^2$ using quilt theory [142]. In particular we see why these tori must be narrow (with the standard spin structure and trivial local system) except in characteristics 3 and 7 respectively. Exploiting this relationship in the other direction, using explicit knowledge of the discs on the Chekanov torus in $\mathbb{CP}^1 \times \mathbb{CP}^1$ we prove that with an appropriate choice of relative spin structure our $\mathbb{RP}^3$ is wide in the $p = 5$ case left open by Theorem 4 and Theorem 5 (we also obtain this result by considering changes of relative spin structure in the Gysin sequence).

We then consider a family of $\mathrm{PSU}(N)$-homogeneous Lagrangians in $\mathbb{CP}^{N^2-1}$, previously studied by Amarzaya–Ohnita [7], Chiang [31, Section 4], Iriyeh [91] and Evans–Lekili [44, Example 7.2.3]. Evans and Lekili showed that they are wide in prime characteristic p whenever N is a power of p, and narrow if p does not divide N. Earlier Iriyeh had obtained this result for $p = 2$, and shown that when $N = 3$ the Lagrangian is non-narrow over $\mathbb{Z}$. We prove that they are narrow whenever N is not a prime power, over any field and with any choice of relative spin structure and rank 1 local system. One of the arguments we give involves an observation (Proposition 5.2.20) about periodicity in gradings which is particularly useful when the classical cohomology of a Lagrangian L decomposes additively (i.e. ignoring its multiplication) as a tensor product, for example when it's an exterior algebra, so is especially relevant when L is diffeomorphic to a Lie group.

Finally we consider a family of Lagrangian embeddings of complete flag varieties which can be combined with our earlier examples to produce counterexamples to a conjecture of Biran–Cornea about Lagrangian intersections in point invertible manifolds [17, Conjecture 2]. These flag varieties can be obtained from the $\mathrm{PSU}(N)$ family in $\mathbb{CP}^{N^2-1}$ by symplectic reduction, and the chapter ends with some remarks about other Lagrangians that can be constructed in this way.

Chapter 6 indicates some new directions, and in particular outlines the construction of a 'quantum Pontryagin comodule structure'

$$HF^*(L, L) \to H^*(K) \otimes HF^*(L, L)$$

for K-homogeneous monotone Lagrangians L. For toric fibres this exhibits Clifford algebras as comodules over the corresponding exterior algebras—a structure well-known to algebraists. We hope that this comultiplication can be upgraded to the level of Fukaya categories and thus place constraints on the symplectic topology of quasihomogeneous varieties. As a rudimentary application, we combine the new operation with the Oh spectral sequence to complete the computation of the wide varieties for the main family in Chapter 5, and reprove the $N = 3$, $p = 5$ result which we previously obtained using the Chekanov torus.

The thesis concludes with four appendices. The first three cover slightly technical issues: Appendix A contains a detailed discussion of Zapolsky's notion of local system, introduced in [146], which puts our non-displaceability results involving B-fields on a solid footing; Appendix B deals with orientation calculations; Appendix C describes how to achieve transversality for the pearl complex when working with a special integrable complex structure. Appendix D comprises a list of coordinate expressions for specific points in the Platonic Lagrangians L_C, which are needed in Chapter 4.

1.2.3 Orientation schemes

In order to do Floer theory outside characteristic 2 one has to make a choice of orientation scheme. There are several different constructions of such schemes in the literature. In particular, in [17, Appendix A] Biran–Cornea construct explicit coherent orientations on the moduli spaces involved in the pearl complex, whilst in [146] Zapolsky sets up canonical orientation systems for the pearl and Floer complexes and carefully proves their compatibility with the algebraic operations and PSS isomorphisms. Unfortunately, however, the relationships between these approaches, and with other constructions available—particularly that of Seidel [123] (in the Lagrangian intersection picture, rather than the pearl model), which is the best candidate for a 'standard'—are currently unclear.

For most of our results we only use basic properties which would be expected of any reasonable orientation scheme for the pearl complex. Specifically, recalling that the set of relative spin structures on $L \subset X$ is a torsor for $H^2(X, L; \mathbb{Z}/2)$, we assume (in addition to the existence of the differential):

Assumption 1.2.2. (i) (Algebraic structure) Counting Y-shaped pearly trajectories (with appropriate signs) defines a unital product satisfying the graded Leibniz rule. Counting trajectories of the shape shown in Fig. 3.1 defines a map from quantum cochains on X (using a Morse or pseudocycle model, say) which induces a unital ring homomorphism on cohomology. These structures are all compatible with comparison maps between different choices of auxiliary data.

(ii) (PSS isomorphisms) Pearl complex cohomology is canonically identified with self-Floer cohomology via a PSS-type isomorphism, which intertwines the product and the ring homomorphism from quantum cohomology with the Floer product and $\mathcal{CO}^0$ respectively.

(iii) (Morse comparison) The associated graded complex of the filtration by powers of the Novikov variable is canonically identified with the corresponding Morse complex on L, and the pearl product reduces to the classical cup product.

(iv) (Change of relative spin structure) Changing the relative spin structure on L by a class $\varepsilon \in H^2(X, L; \mathbb{Z}/2)$ corresponds to changing the sign attached to a trajectory by $(-1)^{\langle \varepsilon, A \rangle}$, where $A \in H_2(X, L; \mathbb{Z})$ is the total homology class of the discs appearing in the trajectory (and $\langle \cdot, \cdot \rangle$ denotes the pairing between homology and cohomology). //

Biran–Cornea sketch a proof of (i), (iii) and the first part of (ii) for their scheme, whilst Zapolsky verifies (i)–(iii) for his, but without explicit mention of $\mathcal{CO}^0$ (his comparison maps are defined indirectly by composing PSS isomorphisms with Floer continuation maps). The fourth property is easily verified for both schemes (strictly Biran–Cornea only talk about absolute spin structures but the extension to the relative case causes no difficulty), using:

Proposition 1.2.3 (de Silva [39, Theorem Q], Cho [32, Theorem 6.4], Fukaya–Oh–Ohta–Ono [63, Proposition 8.1.16])**.** *Under the change of relative spin structure by $\varepsilon \in H^2(X, L; \mathbb{Z}/2)$, the orientation on the moduli space of pseudoholomorphic discs in class $A \in H_2(X, L; \mathbb{Z})$ changes by $(-1)^{\langle \varepsilon, A \rangle}$.*

This result holds for the raw moduli spaces of parametrised, unmarked discs, but quotienting out by reparametrisations or adding interior or boundary marked points makes no difference under any sensible conventions.

When we come to compute part of $\mathcal{CO}^0$ we have to consider two different sets of auxiliary data: one generic and one special. In order for us to be able to use the latter to say something about the former we need:

Assumption 1.2.4. The signs attached to trajectories in the construction of $\mathcal{CO}^0$ can be defined for any such trajectories which are transversely cut out—even if the auxiliary data are not generic—and are compatible with cobordisms built from paths of auxiliary data in the sense used in the proof of Proposition 3.2.24. //

Any reasonable system is likely to have this property as such cobordisms are the standard approach to proving that Floer-theoretic constructions are independent of choices, but we feel it is worth making explicit. To pin down the actual signs in Proposition 3.2.24 we further make Assumption 3.2.6, which states that certain simple trajectories count with the sign that one might naïvely guess. We prove in Section B.1 that Assumption 3.2.6 is satisfied by both the Biran–Cornea and Zapolsky schemes, and in Section B.2 that it in fact implies Assumption 1.2.4 in the setting we need.

The $\mathcal{CO}^0$ results are applied in Chapter 4 to the Platonic Lagrangians, and in Section 5.1 to the second family of examples. In the latter case we can argue indirectly that the signs we compute using Assumption 3.2.6 are either correct or are wrong in an essentially unimportant way—see Remark 5.1.27. The upshot is that for this application we don't actually need to appeal to any direct sign calculations, only the general principle of Assumption 1.2.4.

Orientation issues in Section 5.2.1–Section 5.2.3 also require comment. There we use various different techniques to study Floer cohomology, including Perutz's Gysin sequence and quilt theory, for which orientation schemes exist but whose compatibility with each other and with the schemes we use in the rest of the thesis we do not verify.

In Section B.4 we compute the signs for the count $\mathfrak{m}_0$ of index 2 discs on a sharply homogeneous Lagrangian. There we follow the conventions of Fukaya–Oh–Ohta–Ono [63, Chapter

8], but we note in Example 3.2.29 that the result is consistent with our $\mathcal{CO}^0$ computation with regard to the observation $\mathcal{CO}^0(c_1) = \mathfrak{m}_0$ of Auroux–Kontsevich–Seidel.

In the long term it would be desirable to understand the relationship between the different schemes in the literature to ensure consistency between the work of different authors.

1.2.4 Notation and conventions

Here we collect together various pieces of notation, and fix certain conventions for the reader to refer to as needed.

We denote by D the closed unit disc in $\mathbb{C}$, with boundary ∂D. The boundary will always be oriented anticlockwise. As is usual we write concatenation of paths from left to right. Compatibly with these two conventions, if X is a space, $L \subset X$ a subspace and $* \in L$ a base point, we think of $\pi_2(X, L, *)$ as homotopy classes of maps of the closed unit square to X which map the bottom side to L and all other sides to $*$, with composition by juxtaposition left to right.

Lie groups will be denoted in uppercase, for example $\mathrm{GL}(n, \mathbb{C})$, G, H, K. The corresponding Lie algebras will be denoted by the same names but in lowercase Fraktur, for example $\mathfrak{gl}(n, \mathbb{C})$, $\mathfrak{g}$, $\mathfrak{h}$, $\mathfrak{k}$ respectively. The exponential map from a Lie algebra to the corresponding Lie group will be denoted by $e^{\cdot}$, whilst if a Lie group G acts on a manifold M the infinitesimal action of a Lie algebra element $\xi \in \mathfrak{g}$ on a point $p \in M$, meaning

$$\frac{\mathrm{d}}{\mathrm{d}t}\bigg|_{t=0} e^{t\xi} p,$$

will usually be denoted by $\xi \cdot p$. We will sometimes also use $g \cdot p$ to denote the action of a group element $g \in G$ on p, although we will often just write gp.

For a totally real submanifold L inside an almost complex manifold (X, J')—for example, a Lagrangian submanifold of a symplectic manifold with a choice of compatible almost complex structure—and a homotopy class β of map from $(D, \partial D)$ to (X, L), we define the moduli space of parametrised J'-holomorphic discs in class β to be

$$\widetilde{\mathcal{M}}_{0,0}(J', \beta) = \{u \colon (D, \partial D) \to (X, L) : [u] = \beta \text{ and } \overline{\partial} u = 0\}$$

(we will try to reserve the notation J for when we are dealing with complex manifolds and have a specific integrable complex structure in mind). This moduli space has virtual dimension $n + \mu(\beta)$, where $\mu(\beta)$ is the Maslov index of β and $n = \dim L$. For an integer j, let $\mathcal{M}_{\mu=j}$ be the union of these moduli spaces over classes β of index j. To emphasise: the discs in $\mathcal{M}_{\mu=j}$ are *parametrised*; we drop the usual tilde from the notation to save clutter. The choice of J' should be clear from the context.

For a moduli space $\mathcal{M}$ of parametrised pseudoholomorphic curves in X, and a point z in the domain, let ev_z denote the evaluation map at the point z. If z is interior then the codomain of ev_z will be X whilst if z is boundary the codomain will be L. Let $\mathcal{M}_{2,\mu=j}$ be the quotient of $\mathcal{M}_{\mu=j}$ by the group of reparametrisations fixing ± 1, and for a relative homology class A of index j, let $\mathcal{M}_2(A)$ be the union of the components of $\mathcal{M}_{2,\mu=j}$ comprising discs in class A. These spaces carry well-defined evaluation maps ev_1 and ev_{-1}.

We identify the group of holomorphic automorphisms of the disc D with the group $\mathrm{PSL}(2, \mathbb{R})$ of Möbius maps preserving the upper half-plane (with infinity adjoined) by identifying this space with D via the map $z \mapsto (iz+1)/(z+i)$ sending 0, 1 and ∞ to $-i$, 1 and i respectively. We view the Lie algebra $\mathfrak{psl}(2, \mathbb{R}) \cong \mathfrak{sl}(2, \mathbb{R})$ as sitting inside the algebra $\mathrm{Mat}_{2\times 2}(\mathbb{C})$ of 2×2 complex matrices. Under these identifications, the rotation $z \mapsto e^{i\theta} z$ of D is generated by the matrix

$$\rho \coloneqq \begin{pmatrix} 0 & \frac{1}{2} \\ -\frac{1}{2} & 0 \end{pmatrix}.$$

For a smooth real-valued function H on a symplectic manifold (X, ω), its Hamiltonian vector field X_H is defined by $\omega(X_H, \cdot) = -\mathrm{d}H$ (rather than $+\mathrm{d}H$). Similarly, a moment map $\mu \colon X \to \mathfrak{k}^*$ for a Hamiltonian K-action on X satisfies

$$\omega(X_\xi, \cdot) = -\langle \mathrm{d}\mu, \xi \rangle$$

for all ξ in $\mathfrak{k}$, where X_ξ denotes the vector field induced by the infinitesimal action of ξ and $\langle \cdot, \cdot \rangle$ is the pairing $\mathfrak{k}^* \times \mathfrak{k} \to \mathbb{R}$.

The symplectic form we take on cotangent bundles is, in the usual notation, $\sum_j \mathrm{d}p_j \wedge \mathrm{d}q_j$, rather than $\sum_j \mathrm{d}q_j \wedge \mathrm{d}p_j$.

The scalar multiplication action of $\mathrm{U}(1)$ on a complex inner product space V is Hamiltonian with moment map $(\|z\|^2 - 1)/2$. Reduction at the zero level set defines a symplectic structure on $\mathbb{P}V$, and this is the one we shall use. For $\mathbb{CP}^n$ we take $V = \mathbb{C}^{n+1}$ with the standard inner product. The resulting symplectic structure corresponds to the Fubini–Study form normalised so that the area of a projective line is π.

There are various choices one has to make in setting up Floer cohomology and the Fukaya category. For us $CF^*(L_0, L_1)$ will be generated by Hamiltonian chords from L_0 to L_1, whilst the differential will count holomorphic strips coordinatised by $(s, t) \in \mathbb{R} \times [0, 1]$ which map $t = 0$ to L_0 and $t = 1$ to L_1, with input at $s = -\infty$ and output at $s = +\infty$. The former is standard, but the latter is strictly a homological (rather than *co*homological) convention; we adopt it for consistency with Zapolsky [146], who works with homology twisted by the orientation bundle, which is canonically identified with cohomology by Poincaré duality. We write the product as a map $CF^*(L_0, L_1) \otimes CF^*(L_1, L_2) \to CF^*(L_0, L_2)$, counting surfaces as shown in Fig. 1.1. The

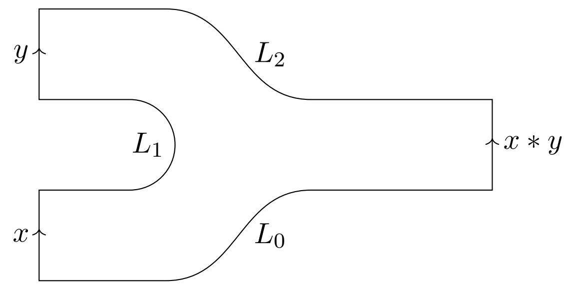

Figure 1.1: A surface defining the Floer product.

order of tensor factors is usually reversed in the Fukaya categories literature, but again we aim for consistency with Zapolsky.

For a (possibly vector-valued) meromorphic function f defined on a neighbourhood of a point p in a Riemann surface Σ, write $o(f)$ for the vanishing order of f at p. Explicitly,

$$o(f) = \sup\{r \in \mathbb{Z} : z^{-r+1} f \to 0 \text{ as } z \to 0\} \in \mathbb{Z} \cup \{\infty\},$$

where z is a local holomorphic coordinate about p, so $o(f)$ is negative if f has a pole at p and ∞ if f is zero on a neighbourhood of p. The point p should be clear from context. This definition obviously also applies to *germs* of functions at p.

Given an isomorphism f between oriented (finite-dimensional real) vector spaces, let $\varepsilon(f)$ be $+1$ if f is orientation-preserving and -1 if it is orientation-reversing. We'll call $\varepsilon(f)$ the *orientation sign* of f.

Chapter 2

Monotone Floer theory and the Oh spectral sequence

This chapter begins with a review of monotone Floer theory, focusing on the use of the pearl complex. We then describe the Oh spectral sequence in detail and discuss the action of the relative symplectomorphism group $\mathrm{Symp}(X, L)$ on it.

2.1 Preliminaries

2.1.1 Monotone Floer theory

Recall that a symplectic manifold (X, ω) is *monotone* if ω is positively proportional to the first Chern class $c_1(X)$ in $H^2(X; \mathbb{R})$, and a Lagrangian submanifold $L \subset X$ is monotone if ω is positively proportional to the Maslov class μ in $H^2(X, L; \mathbb{R})$. Since the image of μ in $H^2(X; \mathbb{R})$ is $2c_1(X)$, a Lagrangian can only be monotone if the ambient symplectic manifold is monotone. We will usually use μ to denote the Maslov index homomorphism $H_2(X, L; \mathbb{Z}) \to \mathbb{Z}$ obtained by pairing with the Maslov class (the reader is warned that it will sometimes also be used to denote a moment map, but there should be no chance of confusion).

Remark 2.1.1. Alternatively one could impose the weaker condition that μ and ω are positively proportional as homomorphisms $\pi_2(X, L) \to \mathbb{R}$, as in Lemma 5.2.29; the distinction between these two definitions will not be important for our purposes. If one is interested in constructing the monotone Fukaya category of X [130] then one also has to require that $H_1(L)$ has trivial image in $H_1(X)$, or $\pi_1(L)$ has trivial image in $\pi_1(X)$ if using the weaker $\pi_2(X, L)$ definition. In all of the examples we consider X is a smooth Fano variety and so is simply connected (it is rationally connected by [93] and hence simply connected by [25, Theorem 3.5]), so this requirement is vacuous. //

We now recollect the essentials of monotone Floer theory. Suppose that (X, ω) is a closed symplectic manifold and $L \subset X$ is a closed, connected, monotone Lagrangian. The *minimal Chern number* N_X of X is defined to be the unique non-negative integer satisfying

$$\langle c_1(X), \pi_2(X) \rangle = N_X \mathbb{Z}.$$

Similarly the *minimal Maslov number* N_L of L is the unique non-negative integer with

$$\mu(\pi_2(X, L)) = N_L\mathbb{Z}.$$

Note that N_L divides $2N_X$ since $\langle 2c_1(X), \pi_2(X)\rangle \subset \mu(\pi_2(X, L))$. If $N_L = 0$ then L bounds no non-constant pseudoholomorphic discs and so its Floer cohomology immediately reduces to its classical cohomology, whilst if $N_L = 1$ then problematic bubbling of index 1 discs can occur (in this case L is not 'monotone enough', and boundary strata in compactifications of moduli spaces do not have sufficiently high codimension). We therefore assume that $N_L \geq 2$. If L is orientable then all Maslov indices are even so the condition $N_L \neq 1$ is trivially satisfied. Let $\nu\colon H_2(X, L; \mathbb{Z}) \to \mathbb{Z}$ denote the 'normalised' Maslov index μ/N_L.

For a ring R let $\Lambda = R[T^{\pm 1}]$ denote the corresponding Novikov ring, with T given grading N_L. Occasionally we will write Λ_R for Λ, to emphasise the ground ring. Working over R, by a *local system* on L we mean a locally constant sheaf $\mathscr{F}$, modelled on an R-module F, with monodromy described by a homomorphism $\overline{m}\colon \pi_1(L, p)^{\mathrm{op}} \to \mathrm{Aut}_R \mathscr{F}_p$ for an arbitrary base point p. The op appears because we write concatenation of paths from left to right, but composition of functions from right to left. In a slight abuse of notation, which seems fairly standard, we'll say that $\mathscr{F}$ *has rank* 1 if $F = R$. Otherwise we'll say F has *higher rank*.

Remark 2.1.2. When higher rank local systems are involved the Floer complex may be obstructed, meaning that d^2 is not equal to zero—see Section A.4. This problem is considered by Konstantinov in [94], where he defines certain subcomplexes which are always unobstructed. Whenever we mention higher rank local systems (in particular in Proposition 5.1.9) we assume that either d^2 really is zero, or that we have passed to one of his subcomplexes. Since these subcomplexes are closed under the Floer product and quantum module action, our arguments still apply to them. //

We'll say that a Lagrangian L and a ring R are *compatible* if L is orientable and relatively spin, meaning that $w_2(L)$ lies in the image of $H^2(X; \mathbb{Z}/2)$ in $H^2(L; \mathbb{Z}/2)$, or R has characteristic 2. The significance of this condition is that when constructing Floer-theoretic invariants associated to L over a ground ring characteristic not equal to 2, an orientation and relative spin structure on L are sufficient to define the necessary orientations on moduli spaces of pseudoholomorphic discs on L. The actual choice of orientation on L is ultimately irrelevant, in that it only affects intermediate steps in the constructions, but the choice of relative spin structure really does matter in general.

Definition 2.1.3. A *monotone Lagrangian brane* over a ring R is a compatible, closed, connected, monotone Lagrangian L inside a closed symplectic manifold X, with $N_L \geq 2$, equipped with a choice of relative spin structure s (if $\mathrm{char}\, R \neq 2$) and a local system $\mathscr{F}$. We abbreviate $(L, s, \mathscr{F})$ to $L^\flat$. Our notation and discussions will implicitly assume the existence of s, but in characteristic 2 all mentions of it can be ignored. We may also assume that an orientation of L has been chosen, although by the above comment this choice doesn't matter. In Section 2.2.3 we broaden the definition, to allow a more general kind of local system. A *monotone Lagrangian pre-brane* $L^\natural$ is a brane without a choice of local system; if we talk about Floer theory of a pre-brane then we are implicitly giving it the trivial local system. //

Remark 2.1.4. The choices of relative spin structure and local system are coupled together in the following sense. Recall that if L is relatively spin then the set of relative spin structures is a torsor for $H^2(X; L; \mathbb{Z}/2)$, and those with fixed background class in $H^2(X; \mathbb{Z}/2)$ correspond to the image of $H^1(L; \mathbb{Z}/2)$ in $H^2(X, L; \mathbb{Z}/2)$. Changing relative spin structure by the image of a class $\varepsilon \in H^1(X, L; \mathbb{Z}/2)$ is equivalent to modifying the local system by twisting $\overline{m}$ by $(-1)^{\varepsilon}$, viewed as a homomorphism $\pi_1(L, p) \to \{\pm 1\}$. One could quotient by this relation in Definition 2.1.3 but in practice it is helpful to have actual choices of s and $\mathscr{F}$ in mind. //

For a monotone Lagrangian brane $L^\flat \subset X$ over R, the quantum cohomology $QH^*(X; \Lambda)$ and self-Floer cohomology $HF^*(L^\flat, L^\flat; \Lambda)$ are both defined (see, for example, [103] and [14]); we will often abbreviate them to QH^* and HF^* respectively when there is no risk of confusion. The former is a $\mathbb{Z}$-graded graded-commutative unital Λ-algebra, whose underlying additive group is $H^*(X; \Lambda)$. The latter is a $\mathbb{Z}$-graded unital Λ-algebra, but is in general neither graded-commutative nor determined additively by the singular cohomology of L. We denote the units in these rings by 1_X and 1_L respectively.

Remark 2.1.5. Multiplication by T gives an automorphism of the Floer complex (and its cohomology) of degree N_L. We can exchange the $\mathbb{Z}$-grading-with-periodicity for a $\mathbb{Z}/N_L$-grading by setting $T = 1$ (i.e. tensoring over Λ with $\Lambda/(T-1) = R$, either at chain level, or after passing to cohomology), and we denote the resulting cohomology ring by $HF^*(L^\flat, L^\flat; R)$. However, it is often more convenient to keep track of the Novikov variable T explicitly.

The quantum cohomology of X can actually be defined over $R[S^{\pm 1}]$, where S has degree $2N_X$. Since N_L divides $2N_X$, working over Λ instead is equivalent to adjoining T as a $(2N_X/N_L)$th root of S. As for T we could set S to 1 and collapse all gradings to $\mathbb{Z}/N_X$, and we denote this quantum cohomology ring by $QH^*(X; R)$. //

The (length zero) *closed–open (string) map*

$$\mathcal{CO}^0 \colon QH^* \to HF^*$$

is a unital homomorphism of graded Λ-algebras, with image contained in the centre of HF^* (i.e. the part which graded-commutes with everything). Of course one can set the Novikov parameters to 1 and view $\mathcal{CO}^0$ as a map of $\mathbb{Z}/N_L$-graded R-algebras (in future we will leave it to the reader to make the obvious translations between constructions over Λ and over R).

Alternatively, the closed–open map may be phrased in terms of the *quantum module structure* (see [14, Theorem 2.1.1]), which is a multiplication

$$QH^j \otimes HF^k \to HF^{j+k},$$

making HF^* into a two-sided algebra over QH^* (meaning that elements of QH^* graded-commute with everything in HF^*). The map $\mathcal{CO}^0$ is simply 'quantum module action on 1_L'. Conversely, the module structure can be reconstructed from $\mathcal{CO}^0$ by combining it with the multiplication in HF^*.

2.1.2 The pearl complex

Throughout the thesis we will make use of the *pearl complex* model for Floer (co)homology, developed by Biran–Cornea [14], so we now give a brief summary.

Fix a Morse function f and metric g on L such that the pair (f, g) is Morse–Smale, and a generic almost complex structure J' on X compatible with the symplectic form. We abbreviate the data (f, g, J') to $\mathscr{D}$. The pearl complex for $L^\flat = (L, s, \mathscr{F})$ defined using $\mathscr{D}$ then has underlying Λ-module

$$C^*(L^\flat; \mathscr{D}; \Lambda) = \bigoplus_{x \in \mathrm{Crit}(f)} (\Lambda \cdot x) \otimes_R \mathrm{End}_R \mathscr{F}_x, \tag{2.1}$$

graded by Morse indices of critical points and the grading on Λ, and the differential counts rigid *pearly trajectories* or *strings of pearls*. These are Morse flowlines between critical points, which may be interrupted by any number (including zero) of non-constant J'-holomorphic discs in X whose boundaries lie on L and carry two marked points—one 'incoming', which we assume is at -1 in ∂D, and one 'outgoing' at 1. Some example trajectories are illustrated in Fig. 2.1.

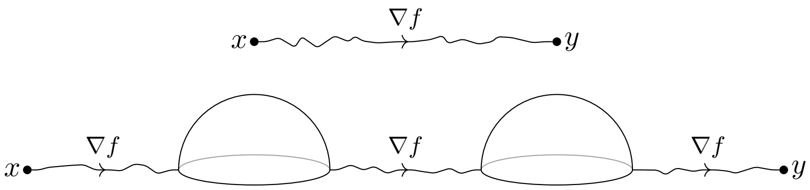

Figure 2.1: Pearly trajectories from x to y.

Explicitly, for critical points x and y the action of the differential from the x-summand of (2.1) to the y-summand is a sum over pearly trajectories γ from x to y which occur in zero-dimensional moduli spaces after quotienting out by reparametrisation of discs and flowlines. For each such γ we can define two homomorphisms, $\mathcal{P}(\gamma_{\mathrm{t}})$ and $\mathcal{P}(\gamma_{\mathrm{b}})$, from $\mathscr{F}_x$ to $\mathscr{F}_y$ by parallel transporting along γ, traversing each disc around the top and bottom half of the boundary respectively, as shown in Fig. 2.2. The contribution of γ to the differential of $x \otimes \theta$, for $\theta \in \mathrm{End}\, \mathscr{F}_x$, is then

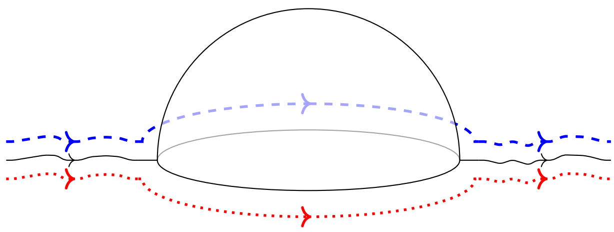

Figure 2.2: The parallel transport paths γ_{t} (dashed) and γ_{b} (dotted) defining the maps $\mathcal{P}(\gamma_{\mathrm{t}})$ and $\mathcal{P}(\gamma_{\mathrm{b}})$.

defined to be $T^{\nu(A)} y \otimes \mathcal{P}(\gamma_{\mathrm{t}}) \circ \theta \circ \mathcal{P}(\gamma_{\mathrm{b}})^{-1}$, where $A \in H_2(X, L; \mathbb{Z})$ is the sum of the homology classes of the discs occurring in γ. From here, the differential is extended Λ-linearly. The virtual dimension of the space of trajectories from x to y in class A, modulo reparametrisation, is $|y| - |x| - 1 + \mu(A)$, where $|\cdot|$ denotes Morse index, so d has degree 1.

Remark 2.1.6. Biran–Cornea work with downward Morse flows and homological notation, whilst we prefer to work with upward flows and *co*homological notation. This is not important (except for defining orientations), and there is Poincaré duality between the two conventions [14, 146]. //

If $\mathscr{F}$ has rank 1 then the space of endomorphisms of each fibre is canonically identified with R, so we can just view the complex as the free Λ-module generated by the critical points. In this case the group of fibre automorphisms, $R^\times$, is abelian so the monodromy representation $\overline{m}\colon \pi_1(L,p)^{\mathrm{op}} \to \mathrm{Aut}_R \mathscr{F}_p$ factors through $H_1(L;\mathbb{Z})$. Then in the map $\theta \mapsto \mathcal{P}(\gamma_\mathrm{t}) \circ \theta \circ \mathcal{P}(\gamma_\mathrm{b})^{-1}$ the contributions of the parallel transports along the Morse flowlines all cancel away, and we are left with $\theta \mapsto \overline{m}(-\partial A)\theta$, where $A \in H_2(X,L;\mathbb{Z})$ is again the total class of the discs in γ and ∂ denotes the boundary map $H_2(X,L;\mathbb{Z}) \to H_1(L;\mathbb{Z})$. Letting m denote the composition of $\overline{m}$ with the inverse map on $\pi_1(L,p)$, this simplifies to $\theta \mapsto m(\partial A)\theta$.

By considering the boundaries of the compactifications of moduli spaces of pearly trajectories of virtual dimension 1, it can be shown that the pearl complex differential satisfies $\mathrm{d}^2 = 0$, at least when $\mathscr{F}$ has rank 1 (otherwise see Remark 2.1.2). The details are spelt out in [14, Section 5.1.1] in characteristic 2 and without local systems, and with a small gap involving bubbling of index 2 discs. These issues are all addressed in Zapolsky's excellent paper [146], which contains a thorough treatment of orientations, and of local systems of a more abstract form which we discuss in Appendix A. Biran–Cornea call the resulting (co)homology the *Lagrangian quantum (co)homology* of L.

The pearl complexes constructed with different choices of data $\mathscr{D}$ can be shown to be quasi-isomorphic [14, Section 5.1.2] by constructing a comparison map using Morse cobordisms, which were introduced by Cornea–Ranicki in [37]. Moreover, the comparison map is unique up to chain homotopy, so the induced identification of cohomology groups is canonical. The quantum cohomology is therefore an intrinsic symplectic invariant of $L^\flat$, and in [14, Section 5.6] (expanded and clarified in [146]) it is shown to be canonically isomorphic to the self-Floer cohomology $HF^*(L^\flat, L^\flat; \Lambda)$ via a so-called PSS map. Biran–Cornea also showed [14, Section 5.2] how to construct a product on $HF^*(L^\flat, L^\flat; \Lambda)$ in the pearl model, where the contribution of a critical point z to $x * y$ is given by counting Y-shaped pearly trajectories from x and y (at the two 'input legs' of the Y) to z. Zapolsky [146, Section 5.2.4] proved that this agrees with the standard Floer product, counting pseudoholomorphic triangles.

Remark 2.1.7. For Morse critical points of index less than $N_L - 1$, the pearl differential coincides with the Morse differential (twisted by the local system $\mathrm{End}\,\mathscr{F}$), so we obtain a map

$$i\colon H^{<N_L-1}(L; \mathrm{End}\,\mathscr{F}) \to HF^*(L^\flat, L^\flat).$$

This is compatible with the comparison maps so is canonical. Moreover, if x and y are Morse cocycles whose degrees sum to less than $N_L - 1$ then the Y-shaped pearly trajectories defining the Floer product are precisely the standard Morse product trajectories, so $i(x) * i(y) = i(x \smile y)$. //

Remark 2.1.8. More generally, one can choose two different local systems $\mathscr{F}^0$ and $\mathscr{F}^1$ on L, and construct the pearl complex for $(L, s, \mathscr{F}^0, \mathscr{F}^1)$, given by

$$\bigoplus_{x \in \mathrm{Crit}(f)} (\Lambda \cdot x) \otimes_R \mathrm{Hom}_R(\mathscr{F}^0_x, \mathscr{F}^1_x).$$

In the twisting of the differential, $\mathcal{P}(\gamma_\mathrm{b})$ is now parallel transport for $\mathscr{F}^0$, whilst $\mathcal{P}(\gamma_\mathrm{t})$ is for

$\mathscr{F}^1$. This complex computes $HF^*((L, s, \mathscr{F}^0), (L, s, \mathscr{F}^1); \Lambda)$. //

Remark 2.1.9. The approach of Biran–Cornea is to fix a Morse–Smale pair and then choose a generic J'. In Appendix C, we show that if X admits a compatible J which is integrable, and has the property that all J-holomorphic discs have all partial indices non-negative, then the pearl complex can be defined using this specific complex structure, possibly after pulling back the pair (f, g) by a diffeomorphism of L which is arbitrarily C^∞-close to the identity. In particular, this is the case when (X, L) is K-homogeneous (see Section 3.1.1). //

2.2 The Oh spectral sequence and $\mathrm{Symp}(X, L)$

In this section we describe the construction of the spectral sequence (1.2), which starts at the classical cohomology of a monotone Lagrangian and converges to its self-Floer cohomology. We then discuss how interesting discrete symmetries of the Lagrangian can be used to gain control over the differentials.

2.2.1 The spectral sequence

Let $L^\flat = (L, s, \mathscr{F}) \subset X$ be a monotone Lagrangian brane over a ring R. Recall from Section 2.1.2 that $\Lambda = R[T^{\pm 1}]$ is the associated Novikov ring (with T given grading N_L), and ν is the normalised Maslov index μ/N_L. We make the following definition (the general idea is standard, but the specifics vary among different authors):

Definition 2.2.1. $L^\flat$ is *wide* (over R) if its Floer cohomology $HF^*(L^\flat, L^\flat; \Lambda)$ is isomorphic to its singular cohomology $H^*(L; \Lambda \otimes_R \mathrm{End}\,\mathscr{F})$ as a graded Λ-module (not necessarily canonically), and *narrow* (over R) if $HF^*(L^\flat, L^\flat; \Lambda) = 0$. Note that this is equivalent to $HF^*(L^\flat, L^\flat; R)$ being isomorphic to $H^*(L; \mathrm{End}\,\mathscr{F})$ (respectively vanishing) as a $\mathbb{Z}/N_L$-graded R-module. //

Fix auxiliary data $\mathscr{D} = (f, g, J')$ and consider the associated pearl complex $C = C^*(L^\flat; \mathscr{D}; \Lambda)$, with underlying Λ-module given by (2.1). For simplicity we'll assume $\mathscr{F}$ has rank 1, so we can drop the $\mathrm{End}\,\mathscr{F}_x$ factors, and for brevity we denote $HF^*(L^\flat, L^\flat; \Lambda)$ simply by HF^*.

The grading on the chain complex C can be refined to a bigrading $C^{*,*}$, by powers of T and by Morse index. Explicitly, for an integer p and a critical point x of index $|x| = q$ we say that the term $T^p x$ (of degree $N_L p + q$) has *polynomial degree* p and *Morse degree* q, and hence lies in $C^{p,q}$. The differential d decomposes as $\partial_0 + \partial_1 + \partial_2 + \dots$, where each ∂_j is of bidegree $(j, 1 - N_L j)$ (and $\partial_j = 0$ for $j > (n+1)/N_L$), so we have a decreasing filtration

$$\cdots \supset F^{-1}C \supset F^0 C \supset F^1 C \supset \cdots$$

of C by subcomplexes $F^p C = C^{\geq p, *}$ comprising elements of polynomial degree at least p: the energy (or index) filtration. Note that this is a filtration of R- or $R[T]$-modules, but not of Λ-modules. The standard construction for a filtered complex then gives a spectral sequence—the *Oh spectral sequence*, named after [111], although it was introduced in the present form by Biran–Cornea in [14]—which will be our key tool. A clear exposition of this construction, on which our treatment is based, is given by Eisenbud in [41, Appendix A, Section A.3.13].

The spectral sequence is a sequence $(E_r, \mathrm{d}_r)_{r=1}^{\infty}$ of differential R-modules (pages) along with an identification of the cohomology $H(E_r, \mathrm{d}_r)$ of the rth page with the $(r+1)$th page E_{r+1}. There is an explicit description of E_r, given by

$$E_r = \bigoplus_{p\in\mathbb{Z}} \frac{\{z \in F^pC : \mathrm{d}z \in F^{p+r}C\} + F^{p+1}C}{(F^pC \cap \mathrm{d}F^{p-r+1}C) + F^{p+1}C}, \tag{2.2}$$

and d_r is the differential mapping the pth summand ('column') to the $(p+r)$th by

$$x + y \in \{z \in F^pC : \mathrm{d}z \in F^{p+r}C\} + F^{p+1}C \mapsto \mathrm{d}x.$$

In particular, E_1 is the homology of the associated graded complex $\operatorname{gr} C = \oplus_p F^pC/F^{p+1}C$ coming from our filtration. Since ∂_0 is the Morse differential for (f, g) on L, this complex is simply the Morse complex over the ring $\operatorname{gr}\Lambda$, which is canonically identified with Λ itself. In other words,

$$E_1 \cong H^*(L;\Lambda) \cong H^*(L;R) \otimes_R \Lambda \cong \bigoplus_{p\in\mathbb{Z}} T^p \cdot H^*(L;R).$$

Since $\partial_j = 0$ for $j \gg 0$ the spectral sequence degenerates after finitely many steps, meaning that d_r vanishes for r sufficiently large and hence that the pages stabilise to a limit $E_\infty = E_{\gg 0}$. There is a filtration on $HF^* = H(C)$, given by the images of the $H(F^pC)$, which we denote by F^pHF^*. The limit page E_∞ is equal to $\operatorname{gr} HF^*$: the associated graded module of this filtration.

The pages of the spectral sequence naturally inherit a Λ-module structure from the filtration, with multiplication by T^j acting from the pth column of (2.2) to the $(p+j)$th. On E_1 this coincides with the Λ-module structure on $H^*(L;\Lambda)$. The differentials d_r are all manifestly Λ-linear. Similarly, $\operatorname{gr} HF^*$ has the structure of a Λ-module, with T^j shifting up j levels of the filtration, and the isomorphism with E_∞ is Λ-linear.

Pick a perturbation $\mathscr{D}'$ of $\mathscr{D}$ and use this to define a pearl complex C'. Using a comparison map $C' \to C$, which is a quasi-isomorphism respecting the filtration, we can then construct a product

$$C \otimes_\Lambda C \xrightarrow{\sim} C \otimes_\Lambda C' \to C,$$

defining the ring structure on HF^*. This product also respects the filtration, so we obtain a multiplicative structure on our spectral sequence. Explicitly, there is a product $*_r$ on each E_r (for $r = 1, 2 \dots, \infty$), satisfying a Leibniz rule with respect to d_r, and such that $*_{r+1}$ is the product induced on $E_{r+1} = H(E_r)$ by $*_r$. Note that $*_1$ is precisely the (Morse-theoretic) classical cup product $\smile$, whilst $*_\infty$ is the product induced on $E_\infty \cong \operatorname{gr} HF^*$ by the Floer product. Since the chain-level multiplication is Λ-bilinear, all of the $*_r$ also have this property.

A typical application of this spectral sequence, which will be a key step in some of our later arguments, is the following well-known result due to Biran–Cornea [14, Proposition 6.1.1].

Proposition 2.2.2. *Suppose that R is a field, $\mathscr{F}$ has rank 1, and that $H^*(L;R)$ is generated as an R-algebra by elements of degree at most m, with $m \leq N_L - 1$. Then $L^\flat$ is either wide or narrow over R, and the latter is only possible if we have equality.*

Proof. Suppose that $m < N_L - 1$ and that $H^*(L;R)$ is generated as an R-algebra by $H^{\leq m}(L;R)$.

Consider the first page $E_1 = H^*(L;R) \otimes \Lambda$ of the spectral sequence, as shown in Fig. 2.3; the solid arrows indicate the directions of the differential d_1. By hypothesis, the subset $H^{\leq m}(L;R)$ of the zeroth column $H^*(L;R)$ generates the whole page as a Λ-algebra. It is clear from the diagram (which is essentially just a graphical way of representing bidegree considerations) that d_1 vanishes on this set, so by the Leibniz rule d_1 is identically zero.

$$
\begin{tikzcd}
 & \vdots & \vdots & \vdots & \vdots & \\
\cdots \arrow[r] & T^{-1}H^{3N_L-3}(L;R) \arrow[r] \arrow[drr, dashed, blue] & H^{2N_L-2}(L;R) \arrow[r] \arrow[drr, dashed, blue] & TH^{N_L-1}(L;R) \arrow[r] & T^2H^0(L;R) \arrow[r] & \cdots \\
\cdots \arrow[r] & T^{-1}H^{3N_L-4}(L;R) \arrow[r] & H^{2N_L-3}(L;R) \arrow[r] & TH^{N_L-2}(L;R) \arrow[r] & 0 \arrow[r] & \cdots \\
 & \vdots & \vdots & \vdots & \vdots & \\
\cdots \arrow[r] & T^{-1}H^{2N_L-1}(L;R) \arrow[r] \arrow[drr, dashed, blue] & H^{N_L}(L;R) \arrow[r] \arrow[drr, dashed, blue] & TH^1(L;R) \arrow[r] & 0 \arrow[r] & \cdots \\
\cdots \arrow[r] & T^{-1}H^{2N_L-2}(L;R) \arrow[r] \arrow[drr, dashed, blue] & H^{N_L-1}(L;R) \arrow[r] \arrow[drr, dashed, blue] & TH^0(L;R) \arrow[r] & 0 \arrow[r] & \cdots \\
\cdots \arrow[r] & T^{-1}H^{2N_L-3}(L;R) \arrow[r] & H^{N_L-2}(L;R) \arrow[r] & 0 \arrow[r] & 0 \arrow[r] & \cdots \\
 & \vdots & \vdots & \vdots & \vdots & \\
\cdots \arrow[r] & T^{-1}H^{N_L}(L;R) \arrow[r] \arrow[drr, dashed, blue] & H^1(L;R) \arrow[r] \arrow[drr, dashed, blue] & 0 \arrow[r] & 0 \arrow[r] & \cdots \\
\cdots \arrow[r] & T^{-1}H^{N_L-1}(L;R) \arrow[r] \arrow[drr, dashed, blue] & H^0(L;R) \arrow[r] \arrow[drr, dashed, blue] & 0 \arrow[r] & 0 \arrow[r] & \cdots \\
\cdots \arrow[r] & T^{-1}H^{N_L-2}(L;R) \arrow[r] & 0 \arrow[r] & 0 \arrow[r] & 0 \arrow[r] & \cdots \\
 & \vdots & \vdots & \vdots & \vdots &
\end{tikzcd}
$$

Figure 2.3: The first page of the Oh spectral sequence.

We deduce that $E_2 = H(E_1, \mathrm{d}_1) = E_1$, and $*_2 = *_1$. The differential d_2 maps right by two columns, and down by one row, as shown by the dashed arrows, so like d_1 it annihilates $H^{\leq m}(L;R)$ in the zeroth column. Again this subset generates the whole page as a Λ-algebra (because it's the same algebra as E_1), so $\mathrm{d}_2 = 0$. Repeating this argument, we see inductively that all differentials d_r vanish, so the pages of the spectral sequence are all equal, and we obtain an isomorphism of Λ-algebras

$$H^*(L;\Lambda) \cong \operatorname{gr} HF^*(L,L;\Lambda). \tag{2.3}$$

At this point we recall that C^* carries an overall grading, which is inherited by HF^* (indexed by the *) and by every page of the spectral sequence, where it is constant along descending diagonals. This is not to be confused with the 'associated graded' construction used earlier, which arises from the filtration by polynomial degree. For each j the space C^j is finite-dimensional over R, so we can count dimensions in degree j in (2.3) to get

$$\dim_R H^j(L;\Lambda) = \dim_R \operatorname{gr} HF^j = \dim_R HF^j,$$

and hence $H^*(L;\Lambda) \cong HF^*$ as graded R-vector spaces. Fixing choices of isomorphisms for $0 \leq * < N_L$, and extending to all $*$ by asking that the isomorphisms commute with multiplication by $T^{\pm 1}$ (which shifts degree by $\pm N_L$), we obtain wideness of $L^\flat$ over R.

The argument for $m = N_L - 1$ is identical, except that now the differential

$$\mathrm{d}_1 \colon H^m(L;R) \to T \cdot H^0(L;R)$$

may be non-zero. However, if this is the case then, since R is a field and $H^0(L;R)$ is one-dimensional, d_1 must hit the unit in $H^*(L;\Lambda) = E_1$. This forces every d_1-cocycle to be exact, and hence $E_2 = 0$ so L is narrow. Otherwise d_1 vanishes and we argue as before to get wideness, completing the proof. □

Remark 2.2.3. The Oh spectral sequence constructed in this way using the pearl complex is canonical, so it makes sense to talk about *the* spectral sequence associated to a particular Lagrangian. Explicitly, suppose we choose different auxiliary data $\mathscr{D}'$, obtain the corresponding pearl complex C', and construct the associated spectral sequence $(E'_r, \mathrm{d}'_r)_{r=1}^\infty$. Then the comparison map $c \colon C \to C'$ respects the filtration, so it induces chain maps $(E_r, \mathrm{d}_r) \to (E'_r, \mathrm{d}'_r)$ for each r, and since c is canonical up to homotopy these pagewise chain maps are canonical. Moreover, since the comparison map agrees with the classical Morse comparison map up to higher order corrections, the induced map between the first pages $E_1 \to E'_1$ is the identity map on $H^*(L;R) \otimes \Lambda$; in fact this alone is enough to see that the comparison maps on later pages are canonical, as they are induced from the map on the first page. //

The key point of the above proof is that the classical cup product on the first page of the spectral sequence is compatible with the differentials, and this constrains them in a useful way. In the following subsection we shall show that the action of the group $\mathrm{Symp}(X, L)$ of symplectomorphisms of X fixing L setwise is also compatible with the differentials, but first let us consider an instructive example. The result is well-known from the work of Cho [32], but we aim to show how it can be derived from very little input.

Example 2.2.4. Take the monotone Clifford torus

$$L = \{[z_0 : \cdots : z_n] : |z_j| = 1 \text{ for all } j\} \subset X = \mathbb{CP}^n,$$

equipped with the standard spin structure s (this is defined in general in Definition 3.1.10, but in the present situation one can just think of it as the spin structure arising from the obvious trivialisation of TL as $T\mathbb{R}^n/\mathbb{Z}^n$). Fix a coefficient ring R, and choose a rank 1 local system $\mathscr{F}$ on L. Recall that by varying $\mathscr{F}$ we can, in effect, also vary the spin structure. There are $n+1$ obvious discs on L, namely

$$u_j \colon z \mapsto [1 : \cdots : 1 : z : 1 : \cdots : 1]$$

for $j = 0, \dots, n$, where the z occurs in the jth position, and their boundary classes $[\partial u_j]$ generate $H_1(L;\mathbb{Z})$, subject only to the relation $\sum_j [\partial u_j] = 0$. Their sum in $H_2(X, L;\mathbb{Z})$ therefore lifts to $H_2(X;\mathbb{Z})$, and clearly meets the class of a coordinate hyperplane once, so represents the class of a line. From the long exact sequence of the pair (X, L) we thus see that the u_j have index 2 (their indices are all equal by symmetry) and their classes A_j form a basis for $H_2(X, L; \mathbb{Z})$.

By the argument of Proposition 2.2.2, $L^\flat = (L, s, \mathscr{F})$ is wide if the differential

$$\mathrm{d}_1 \colon H^1(L; R) \to T \cdot H^0(L; R) \tag{2.4}$$

in the Oh spectral sequence vanishes. This map is given by pairing a class in $H^1(L; R)$ with the homology class $D_1 \in H_1(L; R)$ representing the boundaries of the pseudoholomorphic index 2 discs on L through a generic point [14, Proposition 6.1.4], weighted by the corresponding monodromies of $\mathscr{F}$ (and then multiplying by T). The crucial point is that D_1 is invariant under symplectomorphisms φ of X preserving $L^\flat$. This is because it is independent on the choice of generic almost complex structure used to define it, and the discs contributing to it for a given J' are carried bijectively by φ to those discs contributing to it for $\varphi_* J'$.

Now consider the action of $\mathbb{Z}/(n+1) \subset \mathrm{Symp}(X, L)$ by cycling the homogeneous coordinates, and suppose that $\mathscr{F}$ is invariant, or equivalently that the monodromies of $\mathscr{F}$ around the circles ∂A_j are equal. This preserves the spin structure so leaves D_1 invariant, but on the other hand it also cycles the classes ∂A_j. If $n+1$ and $\operatorname{char} R$ are coprime (in particular if R has characteristic 0), the only possibility for D_1 is zero, so $\mathrm{d}_1 = 0$ and $L^\flat$ is wide.

Of course the vanishing of D_1 over $\mathbb{Z}$ immediately implies its vanishing over any other ring, but with a little more care we can see this directly. To do this, consider the class $D_1' \in H_2(X, L; \mathbb{Z})$ swept by the *whole* of each index 2 disc through a generic point of L, not just its boundary, again weighted by the monodromy of $\mathscr{F}$ around the boundary. This class is still well-defined, independent of the choice of generic almost complex structure, so is invariant under the $\mathbb{Z}/(n+1)$-action. Therefore $D_1' = \lambda \sum A_j$ for some $\lambda \in R$, and hence $D_1 = \partial D_1' = 0$. In other words, over any ring, and with any invariant spin structure and rank 1 local system, $L^\flat$ is wide. This refinement, where we keep track of the $H_2(X, L; \mathbb{Z})$ class of each disc, will be crucial later. //

Remark 2.2.5. As mentioned above, this wideness result was shown by Cho in [32], but his proof is by direct computation of the class D_1. This involves showing that—up to translation by the torus action—the u_j comprise all holomorphic discs on L of index 2 for the usual integrable complex structure, that the discs are regular for this J, and that they all count with sign $+1$ for the standard spin structure. Thus

$$D_1 = \sum_j m(\partial A_j) \partial A_j \in H_1(L; R),$$

and hence $L^\flat$ is wide if and only if the monodromies $m(\partial A_j)$ are all equal. So in fact our method identifies *all* rank 1 local systems for which L is wide. //

We would like to generalise arguments like those of Example 2.2.4 to situations where we may not have an easily defined class like D_1 or D_1' to work with. Roughly speaking (this is made precise in Proposition 2.2.7) we will prove that the differentials in the Oh spectral sequence are equivariant with respect to the standard pullback action of $\mathrm{Symp}(X, L)$, with modifications to account for the relative spin structure and local system, which automatically forces (2.4) to vanish when s and $\mathscr{F}$ are invariant and $n+1$ and $\operatorname{char} R$ are coprime.

In analogy with the argument given for the non-coprime case, we will show that one can work over an enlarged Novikov ring which records the total $H_2(X, L; \mathbb{Z})$ classes of discs (in fact one *has* to do something like this in general, in order to keep track of changes of relative spin structure and local system), so that the map (2.4) is really a composition of $\mathbb{Z}/(n+1)$-equivariant maps

$$H^1(L; R) \xrightarrow{\alpha} \bigoplus_j T^{A_j} \cdot H^0(L; R) \xrightarrow{\beta} T \cdot H^0(L; R).$$

Here the second map, β, sends each T^{A_j} to T. If σ represents a generator of $\mathbb{Z}/(n+1)$ then $\overline{\sigma} := 1 + \sigma + \cdots + \sigma^n$ acts as zero on the left-hand side, since $\sum \partial A_j = 0$, and hence as zero on the image of the first map, α. But this means precisely that $\operatorname{im} \alpha \subset \ker \beta$, so the composite is zero. Explicitly, given an element $a = \sum a_j T^{A_j}$ in the middle group, we have $\overline{\sigma} \cdot a = (\sum a_j)(\sum T^{A_j})$ and $\beta(a) = (\sum a_j)T$. Clearly the vanishing of either of these expressions is equivalent to $\sum a_j = 0$, so elements annihilated by the action of $\overline{\sigma}$ are in the kernel of β.

The topological expression for the action of $\operatorname{Symp}(X, L)$ on the spectral sequence is essentially due to Biran–Cornea [14, Section 5.8], at least in characteristic 2, without local systems and over the standard Novikov ring. What is novel about our approach is the use of its representation theory to prove that certain differentials in the spectral sequence vanish, or can be made to vanish by an appropriate change of coefficient ring. Since the description of the $\operatorname{Symp}(X, L)$-action is topological, it only sees symplectomorphisms up to homotopies through homeomorphisms preserving L, so in particular it factors through the subgroup $\operatorname{MCG}_\omega(X, L)$ of the mapping class group of the pair (X, L) comprising those classes which can be represented by symplectomorphisms.

2.2.2 The action of $\operatorname{MCG}_\omega(X, L)$

Following Biran–Cornea, let H_2^D be the image of $\pi_2(X, L)$ in $H_2(X, L; \mathbb{Z})$. In other words, it is the group of homology classes of topological discs on L. Let $(H_2^D)^+$ be the subset of H_2^D comprising those classes positive Maslov index, and let

$$\Lambda^\dagger = R\big[(H_2^D)^+ \cup \{0\}\big] = R \oplus \bigoplus_{A \in (H_2^D)^+} R \cdot T^A$$

be the monoid ring associated to $(H_2^D)^+ \cup \{0\}$, with T^A carrying grading $\mu(A)$. As remarked in [14, Section 2.2.4], the pearl complex can be defined over this ring (or any algebra over it), with counts of pearly trajectories now weighted by T^A rather than $T^{\nu(A)}$. We shall refer to this as the *enriched* pearl complex. We could equally work with $H_2(X, L; \mathbb{Z})$ in place of H_2^D, but the classes not realised by discs are essentially irrelevant.

Remark 2.2.6. In situations where one can use an integrable complex structure to construct the pearl complex, positivity of intersections can be used to restrict to a smaller ring. Specifically, any holomorphic disc must have non-negative intersection number with every divisor disjoint from L, so we could replace H_2^D by the subset of classes with this property. //

Note that (again as pointed out in [14, Section 2.2.4]) the homology of the enriched pearl complex is not isomorphic to a Floer cohomology group unless we adjoin inverses to each T^A. Even if this is done the resulting Floer cohomology, thought of as $HF^*(L_0, L_1)$ with $L_0 = L$ and L_1 a small Hamiltonian pushoff, is invariant under global symplectomorphisms—i.e. replacing both L_0 and L_1 with $\varphi(L)$ for some φ in $\mathrm{Symp}(X)$—but *not* under Hamiltonian isotopies of L_0 and L_1 separately. This is because one cannot keep track of $H_2(X, L; \mathbb{Z})$ classes for pseudoholomorphic strips between L and a Hamiltonian perturbation when the perturbation becomes large. However, we shall only use this enriched complex as an intermediate step, before projecting to the pearl complex over the standard Novikov ring Λ.

In order for the effect of an element of $\mathrm{Symp}(X, L)$ to be expressible topologically, without knowing the specifics of the pearly trajectories appearing in the differentials, we must restrict to the case of rank 1 local systems, but in practice this covers many situations of interest. With this assumption, recall from Section 2.1.2 that each rigid pearly trajectory γ from x to y contributes $y \otimes m(\partial A)$ to $\mathrm{d}x$, where m encodes the monodromy and $A \in H_2^D$ is the total homology class of the discs in γ.

Taking auxiliary data $\mathscr{D} = (f, g, J')$, our enriched pearl complex is now denoted by

$$C(L^\flat; \mathscr{D}; \Lambda^\dagger) = \bigoplus_{x \in \mathrm{Crit}(f)} \Lambda^\dagger \cdot x,$$

so a general element is an R-linear combination of terms of the form $T^A x$, for $x \in \mathrm{Crit}(f)$ and $A \in H_2^D$. Given $\varphi \in \mathrm{Symp}(X, L)$, we shall define a chain map $\hbar(\varphi)$ on this complex in a way which respects the filtration and has a topological description on the first page E_1 of the spectral sequence, which extends the ordinary pullback φ^* on the $H^*(L; R)$ factor.

The map $\hbar(\varphi)$ is constructed in several steps, and as mentioned above is basically a straightforward modification of Biran–Cornea's construction in [14, Section 5.8]. As usual, when comparing our discussion with that of Biran–Cornea one has to translate between homological and cohomological conventions, so in particular our maps go in the opposite directions to theirs and give rise to a *right* action of $\mathrm{Symp}(X, L)$. For notational brevity we shall blur the line between pushforwards and inverse pullbacks, and vice versa.

The stages of the construction are illustrated in the following diagram of R-linear chain maps:

$$\begin{aligned} C((L, s, \mathscr{F}); \mathscr{D}; \Lambda^\dagger) \xrightarrow{\delta_s} C((L, \varphi_* s, \mathscr{F}); \mathscr{D}; \Lambda^\dagger) \xrightarrow{\delta_{\mathscr{F}}} C((L, \varphi_* s, \varphi_* \mathscr{F}); \mathscr{D}; \Lambda^\dagger) \\ \xrightarrow{h^\varphi} C((L, s, \mathscr{F}); \varphi^* \mathscr{D}; \Lambda^\dagger) \xrightarrow{c} C((L, s, \mathscr{F}); \mathscr{D}; \Lambda^\dagger). \end{aligned}$$

The first map, δ_s, encodes the effect of the change of relative spin structure and we define it by

$$T^A x \mapsto (-1)^{\langle \varphi_* s - s, A \rangle} T^A x,$$

extended R-linearly, where $\varphi_* s - s$ denotes the class in $H^2(X, L; \mathbb{Z}/2)$ describing the difference from s to $\varphi_* s$. Assumption 1.2.2 ensures that this is indeed a chain map.

The second map, $\delta_{\mathscr{F}}$, captures the change of local system. The monodromy of the pushforward $\varphi_* \mathscr{F}$ of $\mathscr{F}$ is simply the pushforward of the monodromy of $\mathscr{F}$, so a pearly trajectory of class $A \in H_2^D$, which is twisted by $m(\partial A)$ under $\mathscr{F}$, is twisted by $\varphi_* m(\partial A) = m(\partial \varphi^* A)$ under $\varphi_* \mathscr{F}$. The map

$$T^A x \mapsto m(\partial \varphi^* A) m(\partial A)^{-1} T^A x,$$

extended R-linearly, therefore intertwines the old and new differentials, so this is how we define $\delta_{\mathscr{F}}$.

Next we define the map h^φ. This acts on critical points simply by sending $x \in \mathrm{Crit}(f)$ to $\varphi^{-1}(x)$ in $\mathrm{Crit}(\varphi^* f)$, and trajectories contributing to the differential for $\mathscr{D}$ are carried bijectively by φ^{-1} to trajectories contributing to the differential for $\varphi^* \mathscr{D}$. To get a chain map we can extend the map on critical points R-linearly, but *not* $\Lambda^\dagger$*-linearly.* The problem is that trajectories of class $A \in H_2^D$, which are weighted by T^A, are carried to trajectories of class $\varphi^* A$, which are weighted by $T^{\varphi^* A}$. The solution is to define h^φ to act non-trivially on $\Lambda^\dagger$, by $T^A \mapsto T^{\varphi^* A}$. In other words, h^φ is defined by

$$T^A x \mapsto T^{\varphi^* A} \varphi^{-1}(x),$$

extended R-linearly. We assume that the orientations chosen on the ascending and descending manifolds of critical points of $\varphi^* f$ are those carried from f by φ^{-1}.

Finally, c is the pearl complex comparison map interpolating between different choices of auxiliary data, constructed using a Morse cobordism. This *is* $\Lambda^\dagger$-linear. Putting everything together, we define $\hbar(\varphi)$ by

$$\hbar(\varphi) \colon T^A x \mapsto (-1)^{\langle \varphi_* s - s, A \rangle} m(\partial \varphi^* A) m(\partial A)^{-1} T^{\varphi^* A} c(\varphi^{-1}(x)),$$

extended R-linearly. By construction this is a chain map, and respects the filtration, so defines an endomorphism of the spectral sequence over $\Lambda^\dagger$ (which is canonical, just as for the spectral sequence over Λ). Moreover, to order zero in T the map $x \mapsto c(\varphi^{-1}(x))$ is simply the classical pullback map on Morse complexes. We have therefore proved the following:

Proposition 2.2.7. *Let $L^\flat \subset X$ be a monotone Lagrangian brane over a ring R, whose local system has rank* 1*. For any $\varphi \in \mathrm{Symp}(X, L)$, the differentials in the Oh spectral sequence for $L^\flat$ over $\Lambda^\dagger$ commute with the R-linear page endomorphisms induced by the map on $E_1 = H^*(L; R) \otimes \Lambda^\dagger$ defined by*

$$a \otimes T^A \mapsto (-1)^{\langle \varphi_* s - s, A \rangle} m(\partial \varphi^* A) m(\partial A)^{-1} \varphi^*(a) \otimes T^{\varphi^* A} \tag{2.5}$$

for all $a \in H^*(L; R)$ and $A \in H_2^D$.

The expression for the action on E_1 involves only pushforwards and pullbacks in classical cohomology, and the pushforward on the relative spin structure, so is purely topological (and depends on φ only up to homotopies through homeomorphisms preserving L) as desired. It is also easily checked to respect composition of symplectomorphisms, so it defines a representation of $\mathrm{MCG}_\omega(X, L)$ on the spectral sequence.

Remark 2.2.8. Proposition 2.2.7 allows one to derive the results of Example 2.2.4 by the methods outlined after Remark 2.2.5. In particular, one can deduce wideness of the Clifford torus $L \subset X = \mathbb{CP}^n$ over any ring, and with any rank 1 local system for which it *is* wide, using just the representation theory of $\mathrm{Symp}(X, L)$ on the Oh spectral sequence. The argument is particularly clean over $\mathbb{C}$, where the domain of (2.4) is the quotient of the regular representation of $\mathbb{Z}/(n+1)$ by the trivial summand, which admits no non-zero equivariant map to the trivial representation on the codomain. //

Remark 2.2.9. There is an alternative way of looking at rank 1 local systems for the enriched pearl complex: rather than building them into the definition of the complex, one can view them as a post hoc modification of the map reducing $\Lambda^\dagger$ to Λ, whereby the monomial T^A is sent to $m(\partial A)T^{\nu(A)}$ rather than just $T^{\nu(A)}$. From this point of view one would use Proposition 2.2.7 with the trivial local system to constrain the images of the differentials over $\Lambda^\dagger$, and then try to choose a local system so that the reduction to Λ kills these images. //

Remark 2.2.10. Working over $\Lambda^\dagger$, the product on E_1 is given by

$$(a_1 \otimes T^{A_1}) *_1 (a_2 \otimes T^{A_2}) = (a_1 \smile a_2) \otimes T^{A_1+A_2}.$$

It is straightforward to check that this product is invariant under (2.5), so the product structure on the spectral sequence is compatible with the representation of $\mathrm{MCG}_\omega(X, L)$. //

Remark 2.2.11. The result of Proposition 2.2.7 can be extended to diffeomorphisms φ which are *antisymplectic*, meaning that they pull back the symplectic form ω to $-\omega$. In this case we have to replace $\varphi^*\mathscr{D} = (\varphi^*f, \varphi^*g, \varphi^*J')$ with $(\varphi^*f, \varphi^*g, -\varphi^*J')$, since φ^*J' is not compatible with ω, and instead of h^φ simply carrying pearly trajectories by φ^{-1} it now has to replace each disc u in the trajectory with $z \mapsto \varphi^{-1} \circ u(\overline{z})$ in order to preserve holomorphicity. This leads to several changes in (2.5).

Firstly, introducing the complex conjugation flips the signs of the homology classes of the discs, so T^{φ^*A} becomes $T^{-\varphi^*A}$ and $m(\partial\varphi^*A)$ becomes $m(\partial\varphi^*A)^{-1}$. And secondly it changes the orientations on the disc moduli spaces, so introduces a factor of $(-1)^{\mu(A)/2}$. This was computed by Fukaya–Oh–Ohta–Ono in [69, Theorem 1.3], or can be seen directly from (B.9) by 'complex conjugating' the orientations on the two spaces oriented by complex structures (this deals with the spaces of *parametrised* discs; we want to mod out by reparametrisations fixing the two marked points, but this action is unaffected by complex conjugation). Recall that if the Maslov indices of discs are not all even then L is non-orientable and we must be working in characteristic 2, so this factor can be ignored. The resulting expression is

$$a \otimes T^A \mapsto (-1)^{\mu(A)/2+\langle \varphi_* s - s, A\rangle} m(\partial\varphi^*A)^{-1} m(\partial A)^{-1} \varphi^*(a) \otimes T^{-\varphi^*A}. \qquad //$$

2.2.3 H_2^D local systems and B-fields

In the preceding discussions we have seen the value in building the pearl complex over an enlarged coefficient ring like $R[H_2^D]$, which records the total homology class of the discs in a trajectory. We would like to use non-vanishing of pearl complex cohomology of a Lagrangian L over such a ring to deduce that L is Hamiltonian non-displaceable (in fact, this has already been used in the literature—see, for example, [30]), but since there is no obvious analogue of these coefficients in the Lagrangian intersection picture of Floer cohomology it is not clear how to justify this deduction.

Zapolsky's paper [146] introduces a general notion of local system and quotient complex which is applicable equally to the pearl and intersection models, and is compatible with the PSS isomorphisms between them. The use of these systems is described in detail in Appendix A. A special case of this construction, which appears in [146, Section 7.3], corresponds precisely (in the pearl case) to using the coefficient ring $R[H_2^D]$. In particular, we see that non-vanishing pearl complex cohomology over this ring does indeed imply non-displaceability.

However, it does not naturally give rise to a non-trivial object in the Fukaya category of the ambient symplectic manifold X, as there is still no (obvious) way to make sense of these coefficients when talking about Floer cohomology between two distinct Lagrangians. The purpose of this subsection is to describe ways in which information contained in the pearl complex over $R[H_2^D]$ can be translated into information which is seen by the monotone Fukaya category of X.

Suppose we take a homomorphism ρ from H_2^D to the group $R^\times$ of invertible elements in our coefficient ring, and consider the induced R-algebra homomorphism $\hat{\rho}\colon R[H_2^D] \to \Lambda_R = R[T^{\pm 1}]$. Explicitly, for each $A \in H_2^D$ the map $\hat{\rho}$ sends the element T^A to $\rho(A)T^{\nu(A)}$, where $\nu = \mu/N_L$ is the normalised Maslov index. Using $\hat{\rho}$ we can view Λ_R as an $R[H_2^D]$-module, by which we can then tensor the pearl complex over $R[H_2^D]$. This is considered in the case $R = \mathbb{C}$ by Biran–Cornea in [17, Section 2.4], and we refer to it as twisting by an H_2^D *local system*. As usual, one can also set $T = 1$ and throw away the Novikov variable.

Remark 2.2.12. H_2^D local systems can be incorporated into the $\mathrm{Symp}(X, L)$-action on the spectral sequence by modifying the map $\delta_{\mathscr{F}}$ of the previous subsection in an obvious way. //

By the universal coefficient theorem, non-vanishing of the cohomology with an H_2^D local system implies the non-vanishing of the cohomology over $R[H_2^D]$, and hence, in turn, non-displaceability. Alternatively, one can directly set up H_2^D local systems in Zapolsky's framework, as outlined in Appendix A. They can be combined with local systems on L (in the usual sense, as described in Section 2.1.2) in a straightforward manner, and we call the result a *composite local system*. From now on we allow such composite local systems to be incorporated into our notion of monotone Lagrangian brane.

Definition 2.2.13. Given an H_2^D local system defined by a homomorphism $\rho\colon H_2^D \to R^\times$, and an ordinary local system $\mathscr{F}$ with monodromy representation $\overline{m}\colon \pi_1(L,p)^{\mathrm{op}} \to \mathrm{Aut}\,\mathscr{F}_p$, let

$$\widetilde{m}\colon \pi_2(X, L, p) \to \mathrm{Aut}\,\mathscr{F}_p$$

be the *composite twisting*, defined by $\widetilde{m}(A) = \rho(A) \cdot \overline{m}(-\partial A)$. //

Remark 2.2.14. The two individual systems are coupled analogously to Remark 2.1.4: the homomorphisms ρ and $\overline{m}$ can be twisted in the same direction by any homomorphism $H_1(L;\mathbb{Z}) \to R^\times$ without altering the net effect of the composite system. //

Remark 2.2.15. A homomorphism ρ from H_2^D to $R^\times$ induces a homomorphism from the image of $\pi_2(X)$ in $H_2(X;\mathbb{Z})$ to $R^\times$, and this can be used to deform the quantum corrections to the cup product appearing in $QH^*(X;R)$. There is a corresponding deformed closed–open map, which gives a unital ring homomorphism between this modified quantum cohomology and the self-Floer cohomology of L equipped with the H_2^D local system. See Section 3.2.1 and the discussion of the module structure in Section A.3. //

The reasons for introducing H_2^D local systems are twofold. First, Assumption 1.2.2 ensures that modifying the relative spin structure by a class $\varepsilon \in H^2(X, L; \mathbb{Z}/2)$ changes the orientation on the moduli space of trajectories in class $A \in H_2^D$ by $(-1)^{\langle \varepsilon, A \rangle}$. Extending Remark 2.1.4 we see that changes of relative spin structure can in effect be realised by H_2^D local systems which factor through the subgroup $\{\pm 1\} \subset R^\times$. Monotone Lagrangian branes equipped with such systems therefore give rise to honest objects in $\mathcal{F}(X)$. Of course if we change the background class in $H^2(X; \mathbb{Z}/2)$ then we really obtain a different category, but we don't explicitly indicate this in our notation.

The second reason is special to the case $R = \mathbb{C}$, so we restrict to this for the rest of the subsection. Note that here the group $R^\times = \mathbb{C}^*$ is divisible and hence injective, so the functor $\mathrm{Hom}(\cdot, \mathbb{C}^*)$ is exact on the category of abelian groups. In particular, it commutes with taking homology, and any homomorphism from a subgroup of an abelian group to $\mathbb{C}^*$ can be extended to the whole group.

Suppose we fix a closed complex-valued 2-form B on X: this is called a *B-field*. Cho [34, Section 2], following Fukaya [58], describes how the Lagrangian intersection Floer complex for a pair of Lagrangians L_0 and L_1 can be deformed by such a B whose restriction defines an integral cohomology class on each L_j (i.e. a class in the image of $H^2(L_j; \mathbb{Z})$ in $H^2(L; \mathbb{C})$). We now briefly recap this construction. See also recent work of Siegel [131].

Since $[B]$ restricts to an integral class on L_j there exists a complex line bundle $\mathcal{L}^j$ on L_j with $-[B]|_{L_j}$ as its first Chern class, and hence we can put a connection ∇^j on $\mathcal{L}^j$ whose curvature is $2\pi i B$. We now equip each L_j with a relative spin structure (with the same background class) and the pair $(\mathcal{L}^j, \nabla^j)$. Just as in the case $B = 0$, where the connections ∇^j are flat and hence define locally constant structures on the $\mathcal{L}^j$, the Floer complex $CF^*(L_0, L_1)$ is generated by the direct sum over intersection points $x \in L_0 \cap L_1$ of the spaces $\mathrm{Hom}(\mathcal{L}^0_x, \mathcal{L}^1_x)$, assuming these intersections are transverse. Again the differential counts pseudoholomorphic strips u which contribute according to the parallel transport maps of $(\mathcal{L}^j, \nabla^j)$ along $u|_{t=j}$, but now there is an extra factor of $e^{-2\pi i \int u^* B}$.

This can be extended to larger collections of Lagrangians, and higher A_∞-operations, by combining parallel transport maps around the boundary of each disc with a factor coming from the integral of B over the interior. We obtain a deformation $\mathcal{F}^B(X)$ of the full subcategory of the monotone Fukaya category of X comprising those Lagrangians on which the class $[B]$ is integral.

Remark 2.2.16. Note that by Stokes's theorem the map which sends a smooth 2-simplex $\sigma \colon \Delta^2 \to X$ to $e^{-2\pi i \int \sigma^* B}$ defines a class in

$$\mathrm{Hom}(H_2(X; \mathbb{Z}), \mathbb{C}^*) = H^2(X; \mathbb{C}^*),$$

which we denote by $[e^{-2\pi i B}]$, and the condition that $[B]|_L$ is integral can be phrased more cleanly as $[e^{-2\pi i B}]|_L = 1$ (the identity in $\mathbb{C}^*$). If B' is another closed 2-form such that $[e^{-2\pi i B}] = [e^{-2\pi i B'}]$ then there exists a complex line bundle $\mathcal{L}_X$ admitting a connection ∇_X with curvature $2\pi i(B' - B)$. Tensoring with $(\mathcal{L}_X, \nabla_X)$ gives a bijection between line bundles $(\mathcal{L}^j, \nabla^j)$ on L_j for the field B and line bundles $((\mathcal{L}')^j, (\nabla')^j)$ for the field B'. Moreover, this bijection induces natural isomorphisms between the corresponding Floer complexes, including A_∞-structures. In this sense, the deformation of the Fukaya category introduced by B depends on B only via the class $[e^{-2\pi i B}]$. //

The analogous construction in the pearl model is to take our single Lagrangian L and equip it with a line bundle $\mathcal{L}$ and connection ∇ of curvature $2\pi i B$. We could take two different bundles as in Remark 2.1.8, but this just corresponds to combining a B-field with local systems in the usual sense, so we restrict to the case of a single bundle. Working over $\Lambda_{\mathbb{C}}$ a trajectory with discs $(u_1 \dots, u_r)$ is then weighted by

$$\prod_{j=1}^{r} e^{-2\pi i \int_D u_j^* B} \operatorname{hol}_{-\partial u_j}(\nabla) T^{\nu(u_j)},$$

where $\operatorname{hol}_{-\partial u_j}(\nabla)$ denotes the monodromy (holonomy) of the connection around the boundary of u_j (with orientation reversed, just as in Section 2.1.2).

Note that the parallel transport maps for the bundle $(\mathcal{L}, \nabla)$ depend in general on the actual paths taken, and not just their homotopy classes relative to end points (the same is true for the bundles $(\mathcal{L}^j, \nabla^j)$ in the Lagrangian intersection picture). The factor of $e^{-2\pi i \int u^* B}$ precisely compensates for this in the following sense:

Lemma 2.2.17. *The choices of B, $\mathcal{L}$ and ∇ determine a homomorphism*

$$\rho_{B,\mathcal{L},\nabla} \colon H_2(X, L; \mathbb{Z}) \to \mathbb{C}^*$$

(and hence a class in $H^2(X, L; \mathbb{C}^*)$) which sends a smooth map $\iota \colon (P, \partial P) \to (X, L)$, for P a compact oriented surface with boundary, to

$$e^{-2\pi i \int_P \iota^* B} \operatorname{hol}_{-\partial\iota}(\nabla) \in \mathbb{C}^*.$$

This class lifts $[e^{-2\pi i B}]$, and the map $(\mathcal{L}, \nabla) \mapsto \rho_{B,\mathcal{L},\nabla}$ defines a bijection between the set of isomorphism classes of choice of $(\mathcal{L}, \nabla)$ and the set of such lifts.

Proof. Take a smooth 2-chain $\sigma = \sum n_j f_j$ in X, representing a class $[\sigma]$ in $H_2(X, L; \mathbb{Z})$, where the n_j are integers and the f_j are smooth maps $\Delta^2 \to X$. Let $v_1, \dots, v_l \in X$ be the distinct images of the vertices of the 2-simplex Δ^2 under the f_j, and fix an identification $\mathcal{L}_{v_j} \cong \mathbb{C}$ for each v_j in L. We define

$$\rho_{B,\mathcal{L},\nabla}([\sigma]) = \prod \lambda_j^{n_j},$$

where $\lambda_j \in \mathbb{C}^*$ is defined as follows.

For each j consider the simplex f_j. Let its edges be $e_1, e_2, e_3 \colon [0, 1] \to X$, so that $\partial f_j = e_1 - e_2 + e_3$. If e_k lies inside L then set $\lambda_{jk} \in \mathbb{C}^*$ to be the monodromy of the parallel transport map for $(\mathcal{L}, \nabla)$ along e_k, using our fixed identifications of the fibres over the end points with $\mathbb{C}$. Otherwise set $\lambda_{jk} = 1$. Now define λ_j to be

$$e^{-2\pi i \int f_j^* B} \lambda_{j1}^{-1} \lambda_{j2} \lambda_{j3}^{-1}.$$

Since $\partial\sigma$ is a *cycle* in L, the choices of identification of the fibres do not affect the final answer. It is easy to show using Stokes's theorem that our construction assigns 1 to 2-simplices contained in L and to boundaries of 3-simplices in X, so it descends to a well-defined map $H_2(X, L; \mathbb{Z}) \to \mathbb{C}^*$. This map is clearly a homomorphism, and assigns the claimed value to a map $\iota\colon (P, \partial P) \to (X, L)$. When σ is actually closed (in $C_*(X)$), so defines a cycle in X, our definition recovers that of $[e^{-2\pi i B}]$, so $\rho_{B,\mathcal{L},\nabla}$ lifts this class.

To see that $(\mathcal{L}, \nabla) \mapsto \rho_{B,\mathcal{L},\nabla}$ gives the claimed bijection, first note that the set of isomorphism classes of choice of $(\mathcal{L}, \nabla)$ forms a torsor for the group $H^1(L; \mathbb{C}^*)$ of isomorphism classes of rank 1 local systems $\mathscr{F}$ on L (over $\mathbb{C}$): thinking of $\mathscr{F}$ as a line bundle $\mathcal{L}'$ with flat connection ∇', it acts by sending the pair $(\mathcal{L}, \nabla)$ to $(\mathcal{L} \otimes \mathcal{L}', \nabla \otimes \nabla')$. This corresponds to twisting $\rho_{B,\mathcal{L},\nabla}$ by $\mathscr{F}$ via the connecting homomorphism $H^1(L; \mathbb{C}^*) \to H^2(X, L; \mathbb{C})$, and under this action the long exact sequence of the pair

$$
\begin{array}{ccccccccc}
\cdots \longrightarrow & H^1(L; \mathbb{C}^*) & \longrightarrow & H^2(X, L; \mathbb{C}^*) & \longrightarrow & H^2(X; \mathbb{C}^*) & \longrightarrow & H^2(L; \mathbb{C}^*) & \longrightarrow \cdots \\
& & & \rho_{B,\mathcal{L},\nabla} & \dashmapsto & [e^{-2\pi i B}] & \longmapsto & 1 &
\end{array}
$$

exhibits the set of lifts of $[e^{-2\pi i B}]$ in $H^2(X, L; \mathbb{C}^*)$ as another torsor for $H^1(L; \mathbb{C}^*)$. In other words, both sets of interest are naturally torsors for $H^1(L; \mathbb{C}^*)$, and our map intertwines the two actions, so it is a bijection. □

From this result we see that introducing a B-field and pair $(\mathcal{L}, \nabla)$ to the pearl complex for L is equivalent to deforming by the H_2^D local system defined by the restriction ρ of $\rho_{B,\mathcal{L},\nabla}$ to H_2^D. Note that any torsion class in $H_2(X; \mathbb{Z})$ pairs to 0 with $[B]$, so $\rho_{B,\mathcal{L},\nabla}$ sends the image To of the torsion subgroup of $H_2(X; \mathbb{Z})$ in $H_2(X, L; \mathbb{Z})$ to 1. Conversely, we have:

Proposition 2.2.18. *Any H_2^D local system represented by a homomorphism $\rho\colon H_2^D \to \mathbb{C}^*$ which is trivial on* $\mathrm{To} \cap H_2^D$ *can be expressed in terms of a B-field.*

Proof. Since ρ maps $\mathrm{To} \cap H_2^D$ to 1 it extends uniquely to a homomorphism from $\mathrm{To} + H_2^D \subset H_2(X, L; \mathbb{Z})$ to $\mathbb{C}^*$ sending all of To to 1. We can then further extend ρ to a homomorphism $\widetilde{\rho}\colon H_2(X, L; \mathbb{Z}) \to \mathbb{C}^*$, which corresponds to a class in $H^2(X, L; \mathbb{C})$. All other extensions are obtained from $\widetilde{\rho}$ by multiplying by a homomorphism $H_2(X, L; \mathbb{Z})/(\mathrm{To} + H_2^D) \to \mathbb{C}^*$.

By construction, the pullback of $\widetilde{\rho}$ to $H_2(X; \mathbb{Z})$ is trivial on the torsion subgroup so lifts to a homomorphism $H_2(X; \mathbb{Z}) \to \mathbb{C}$ under the covering $\mathbb{C} \to \mathbb{C}^*$ given by $z \mapsto e^{-2\pi i z}$, which can then be expressed as pairing with a class in $H^2(X; \mathbb{C})$. Fix a closed complex-valued 2-form B which represents this class: this will be our B-field. Note that by construction $[B]$ pairs to an integer with any class in $H_2(L; \mathbb{Z})$, and hence it lies in the image of $H^2(L; \mathbb{Z})$ in $H^2(L; \mathbb{C})$. We can now apply the last part of Lemma 2.2.17 to see that up to isomorphism there is a unique choice of $\mathcal{L}$ and ∇ such that the resulting class $\rho_{B,\mathcal{L},\nabla}$ coincides with $\widetilde{\rho}$. This shows that the H_2^D local system defined by ρ is expressible within the framework of B-fields, and this expression is essentially unique up to twisting by a homomorphism $H_2(X, L; \mathbb{Z})/(\mathrm{To} + H_2^D) \to \mathbb{C}^*$ and shifting B by a closed 2-form δB representing a class in the image of $H^2(X; \mathbb{Z})$ in $H^2(X; \mathbb{C})$. □

We end this subsection by mentioning the connection with the *bulk deformations* of Fukaya–Oh–Ohta–Ono [63, Section 3.8]. In their construction a class $\mathfrak{b}$ in $QH^{\mathrm{even}}(X)$ is used to deform the Fukaya category by replacing each count of pseudoholomorphic discs appearing in the A_∞-operations by a sum over $l \in \mathbb{Z}_{\geq 0}$ of such counts, weighted by $1/l!$, with l interior marked points introduced and constrained to lie on cycles representing $\mathfrak{b}$. Because of the infinite sums involved, in order to ensure convergence one has to work over the Novikov ring and ensure that any codimension 2 cycles appearing in $\mathfrak{b}$ carry strictly positive powers of the Novikov variable.

Formally, however, deforming by a codimension 2 cycle $Z \subset X \setminus L$—without a Novikov variable—should just correspond to weighting the count of discs in class u by

$$1 + u \cdot Z + (u \cdot Z)^2/2! + \cdots = e^{u \cdot Z}.$$

In other words, it is exactly the deformation introduced by the B-field Poincaré dual to $iZ/2\pi$ when we take B to be supported on the complement of L and then choose $(\mathcal{L}, \nabla)$ to be trivial (see [65, Remark 1.5]). In this sense, B-fields allow us to fill in the missing bulk deformations by hand. This is analogous to the formal correspondence between rank 1 local systems and codimension 1 bounding cochains without a Novikov variable: the former have to be introduced separately to work around convergence issues for the latter. Just as the Fukaya category deformed by a B-field is a module over a deformed quantum cohomology ring, the bulk deformed Fukaya category is a module over the fibre of the big quantum cohomology ring over the class $\mathfrak{b}$.

One final issue worthy of comment is that in order to deform L by a B-field we have to make a choice of the pair $(\mathcal{L}, \nabla)$. These choices correspond to lifts of $[e^{-2\pi i B}]$ from $H^2(X; \mathbb{C}^*)$ to $H^2(X, L; \mathbb{C}^*)$ and form a torsor for the space of $H^1(L; \mathbb{C}^*)$ of rank 1 local systems on L. This is analogous to choosing cycles representing a bulk class $\mathfrak{b}$; changing cycles corresponds to shifting the bounding cochain. The simplest example is the bulk deformation of the equator in $\mathbb{CP}^1$ by the cycle $\lambda \cdot$ (north pole $-$ south pole) representing the zero class, which acts by twisting local systems by e^{λ} times a generator of $H^1(S^1; \mathbb{Z})$.

Chapter 3

Homogeneous Lagrangians and their holomorphic discs

In this chapter we study homogeneous Lagrangians, introduced by Evans and Lekili [45], which will be the setting for most of the rest of the thesis. After computing part of the closed–open map we develop the detailed local analysis of holomorphic discs with boundary on a sharply K-homogeneous Lagrangian.

3.1 Homogeneous Lagrangians

3.1.1 Homogeneity

First the definition is in order. Although it is phrased in terms of totally real submanifolds of arbitrary complex manifolds, our main interest is in the special case of Lagrangian submanifolds of Kähler manifolds.

Definition 3.1.1. If X is a compact complex manifold of complex dimension n, with complex structure J, $L \subset X$ is a compact totally real submanifold, and K is a compact Lie group acting on X by holomorphic automorphisms such that L is an orbit of the K-action, then we say (X, L) is *K-homogeneous*. We will always assume that X and K—and hence also L—are connected. //

When talking about pseudoholomorphic curves in such an X we will always be taking the standard integrable complex structure J (and will simply call the curves holomorphic), except where explicitly stated otherwise.

Definition 3.1.2. If (X, L) is K-homogeneous and $\dim K$ is equal to $\dim L = n$ then we say (X, L) is *sharply K-homogeneous*. If, moreover, L is a free K-orbit then we'll say it's *freely K-homogeneous*. Given a K-homogeneous (X, L), which we'll often just abbreviate to L, we'll say 'the homogeneity is sharp (free)' to mean that it is actually sharply (respectively freely) K-homogeneous. //

Definition 3.1.3. A *(sharply/freely) K-homogeneous monotone Lagrangian (pre-)brane* is a monotone Lagrangian (pre-)brane $L^\flat \subset X$, where (X, ω, J) is a Kähler manifold such that $((X, J), L)$ is (sharply/freely) K-homogeneous. Note that it is J, rather than ω, that is required to be K-invariant. //

When (X, L) is K-homogeneous, Evans–Lekili showed [45, Lemma 2.11, Lemma 3.2] that all holomorphic discs u have all partial indices non-negative (the definition of partial indices is recalled in Section B.3), by using the action of K to generate sections of u^*TX. In particular, the moduli spaces $\mathcal{M}_{\mu=j}$ defined in Section 1.2.4 are all smooth manifolds of the expected dimension $n+j$. Strictly they work with a slightly different definition of K-homogeneity, but their proof carries over immediately to our setting.

3.1.2 Complexifications

Next we summarise the basic theory of universal complexifications of Lie groups, as described in [85, Section 15.1].

Definition 3.1.4. For a Lie group H, a *universal complexification* is a complex Lie group $H_{\mathbb{C}}$ and Lie group morphism $\eta_H \colon H \to H_{\mathbb{C}}$ through which any other morphism to a complex Lie group factors uniquely. //

The prototypical example is the inclusion of the unit circle $\mathrm{U}(1)$ in $\mathbb{C}^*$, or more generally the inclusion of the torus T^n in $(\mathbb{C}^*)^n$ or of the unitary group $\mathrm{U}(n)$ (or special unitary group $\mathrm{SU}(n)$) in $\mathrm{GL}(n, \mathbb{C})$ (respectively $\mathrm{SL}(n, \mathbb{C})$).

We will need the following properties of the universal complexification:

Proposition 3.1.5 (See [85, Theorem 15.1.4])**.** *For any Lie group H:*

(i) H has a universal complexification $(\eta_H, H_{\mathbb{C}})$.

(ii) η_H induces a bijection $\pi_0(H) \to \pi_0(H_{\mathbb{C}})$.

(iii) There exists a unique antiholomorphic group involution $\hat{\tau}$ of $H_{\mathbb{C}}$ satisfying $\hat{\tau} \circ \eta_H = \eta_H$. The fixed-point set $H_{\mathbb{C}}^{\hat{\tau}}$ of $\hat{\tau}$ is a closed subgroup of $H_{\mathbb{C}}$ and hence a Lie subgroup.

(iv) The identity component of $H_{\mathbb{C}}^{\hat{\tau}}$ is the image under η_H of the identity component of H.

Proposition 3.1.6 (See [85, Theorem 15.2.1])**.** *For any compact Lie group K:*

(i) $\eta_K \colon K \to K_{\mathbb{C}}$ is injective so we can view K as a subgroup of $K_{\mathbb{C}}$.

(ii) K is a maximal *compact subgroup of $K_{\mathbb{C}}$, so any other compact subgroup of $K_{\mathbb{C}}$ is conjugate to a subgroup of $K \subset K_{\mathbb{C}}$.*

(iii) If $\mathfrak{k}$ denotes the Lie algebra of K then the Lie algebra of $K_{\mathbb{C}}$ is canonically identified with $\mathfrak{k} \otimes \mathbb{C}$, and the polar maps $\mathfrak{k} \times K \to K_{\mathbb{C}}$ defined by $(\xi, k) \mapsto \exp_{K_{\mathbb{C}}}(i\xi)k$ and $k \exp_{K_{\mathbb{C}}}(i\xi)$ are diffeomorphisms.

(iv) If $\rho \colon K \to \mathrm{GL}(V)$ is a complex representation of K, and $\rho_{\mathbb{C}}$ denotes the extension of ρ to $K_{\mathbb{C}}$, then the kernel of $\rho_{\mathbb{C}}$ is the complexification of the kernel of ρ and the image of $\rho_{\mathbb{C}}$ is a complex Lie subgroup of $\mathrm{GL}(V)$.

Remark 3.1.7. For us, a representation of a Lie group means a *continuous* representation, or equivalently a smooth representation (by [85, Theorem 9.2.16]). //

When talking about a K-homogeneous (X, L), so K is a compact, connected Lie group, we will denote the complexification $K_{\mathbb{C}}$ by G, with Lie algebra $\mathfrak{g}$. In light of Proposition 3.1.6(i) we will treat K as a subgroup of G, and by Proposition 3.1.5(iii) and (iv) G admits an antiholomorphic group involution $\hat{\tau}$ fixing precisely K. In polar form, $\hat{\tau}$ is given by $ke^{i\xi} \mapsto ke^{-i\xi}$, or equivalently $e^{i\xi}k \mapsto e^{-i\xi}k$, where $e^{\cdot}$ is the exponential map $\mathfrak{g} \to G$.

Lemma 3.1.8. *If (X, L) is K-homogeneous then the K-action on X complexifies to a holomorphic G-action. L is contained in a dense G-orbit W whose complement Y is analytically closed.*

Proof. The first part is standard, appearing for example in [79, Section 4], and goes as follows. Since X is compact, the group $\mathrm{Aut}(X)$ of holomorphic automorphisms of X has the structure of a complex Lie group [92, Section III Theorem 1.1]. The action of K on X gives a Lie group morphism $K \to \mathrm{Aut}(X)$, and then by the universal property of the universal complexification this extends to a unique complex Lie group morphism $G \to \mathrm{Aut}(X)$, which is the required action.

If $\xi_1, \dots, \xi_m$ is a basis for $\mathfrak{k}$ then for each subset $I = \{i_1 < \cdots < i_n\}$ of $\{1, \dots, m\}$ there is a holomorphic section σ_I of $\Lambda^n_{\mathbb{C}} TX$ (i.e. the anticanonical bundle of X) defined by

$$\sigma = X_{\xi_1} \wedge \cdots \wedge X_{\xi_n},$$

recalling that X_ξ represents the holomorphic vector field generated by the action of ξ. The common zeros of the σ_I form a proper analytically closed subset Y of X, whose complement W is therefore dense and connected (X itself is connected, and Y has real codimension at least 2 so cannot disconnect it). W is partitioned in G-orbits, each of which is open since the infinitesimal action on W is surjective, and so by connectedness it is a single orbit. It clearly contains L. $\square$

If (X, L) is sharply K-homogeneous then there is only one n-tuple I, and we denote the corresponding σ_I simply by σ.

Lemma 3.1.9. *If (X, L) is sharply K-homogeneous (so $\dim K = \dim L$) then:*

(i) L is parallelisable, so in particular is orientable and spin.

(ii) Y is a divisor (meaning an analytic subvariety of complex codimension 1), and L is special Lagrangian in W (in the sense of [9]).

(iii) The Maslov index of a holomorphic disc $u\colon (D, \partial D) \to (X, L)$ is twice the sum of the vanishing orders of $\sigma \circ u$ at the intersection points of u with Y.

Proof. (i) The infinitesimal action of $\mathfrak{k}$ on L exhibits an isomorphism between TL and the trivial bundle $\mathfrak{k} \times L$.

(ii) Y is defined by the vanishing of σ, so has (complex) codimension 1 in X. Moreover $\sigma|_W$ is nowhere-zero and $\sigma|_L$ is real so L is special Lagrangian in W.

(iii) This is equivalent to Auroux's result [9, Lemma 3.1], which essentially amounts to equating the twisting of $\Lambda^n TL$ around the boundary of u with that of $\mathbb{R}\sigma$, and using the argument principle (i.e. the fact that for a continuous function $f\colon D \to \mathbb{C}$, which is holomorphic over the interior and non-zero on the boundary, the winding number of $f|_{\partial D}\colon \partial D \to \mathbb{C}^*$ is the number of zeros of f in D). $\square$

Definition 3.1.10. If (X, L) is sharply K-homogeneous then by Lemma 3.1.9(i) L is automatically orientable and spin, so is compatible with any ring. In fact TL is canonically isomorphic to $\mathfrak{k} \times L$, so for each choice of orientation there is a particularly natural choice of spin structure, induced by this parallelisation. We call this the *standard* spin structure, following Cho [32] and Cho–Oh [35] in the toric case (although their construction is superficially different). //

3.1.3 Axial discs

Fix a K-homogeneous (X, L). We will be particularly interested in holomorphic discs on L which are of a certain simple form, introduced in [45, Section 3.4] by Evans–Lekili and termed *axial*.

Definition 3.1.11. A holomorphic disc $u\colon (D, \partial D) \to (X, L)$ in a K-homogeneous (X, L) is axial if, possibly after reparametrisation, there exists a Lie algebra element $\xi \in \mathfrak{k}$ such that u is of the form

$$z \mapsto e^{-i\xi \log z} u(1) \tag{3.1}$$

for non-zero z in D. //

Remark 3.1.12. This is not quite the definition given by Evans–Lekili, which is in terms of the existence of a Lie group morphism $R\colon \mathbb{R} \to K$ such that $u(e^{i\theta}z) = R(\theta)u(z)$ for all $z \in D$ and $\theta \in \mathbb{R}$. Clearly any disc satisfying our definition also satisfies theirs, with $R(t) = e^{\xi t}$ for all t. Conversely suppose such a Lie group morphism R exists, and let $\xi = R'(0)$. We claim that u is of the form (3.1) with $\xi = R'(0)$.

Note first that we have $R(\theta) = e^{\theta R'(0)}$ for all $\theta \in \mathbb{R}$, so u is of the claimed form on ∂D. Moreover, we see that $e^{2\pi R'(0)}$ fixes $u(1)$, and hence the right-hand side of (3.1) is well-defined on $\mathbb{C}^*$. Take a point $z \in \partial D$ and pick vectors $\xi_1, \dots, \xi_n$ in the Lie algebra $\mathfrak{k}$ of K whose infinitesimal actions at $u(z)$ form a basis for $T_{u(z)}L$. Then the map

$$(z_1, \dots, z_n) \mapsto e^{\sum z_i \xi_i} u(z)$$

defines a holomorphic parametrisation of a neighbourhood of $u(z)$, under which L corresponds to $\mathbb{R}^n \subset \mathbb{C}^n$ in coordinate space. So in our chart both sides of (3.1) are given on a neighbourhood of z in $D \setminus \{0\}$ by n continuous functions, holomorphic off ∂D, and equal and real on ∂D. The standard Schwarz reflection argument then proves that they agree on the whole neighbourhood of z (or at least the component containing z), and hence on all of $D \setminus \{0\}$. //

It is not obvious that expressions of the form on the right-hand side of (3.1) extend over $z = 0$ for arbitrary choices of ξ, but this is indeed the case for Hamiltonian actions:

Proposition 3.1.13. *Suppose X is Kähler and the K-action is Hamiltonian with moment map $\mu\colon X \to \mathfrak{k}^*$. If x is a point in L and $\xi \in \mathfrak{k}$ satisfies $e^{2\pi\xi}x = x$, so that $u\colon (D \setminus \{0\}, \partial D) \to (X, L)$ given by $z \mapsto e^{-i\xi \log z}x$ is well-defined, then u extends continuously, and hence holomorphically, over* 0.

Proof. By removal of singularities [103, Theorem 4.2.1] it suffices to show u has finite energy. But taking polar coordinates r and θ on $D\setminus\{0\}$, and recalling our sign convention from Section 1.2.4, the energy density is given by

$$\omega(\partial_r u, \partial_\theta u)\ \mathrm{d}r \wedge \mathrm{d}\theta = \omega(\partial_r u, \xi \cdot u)\ \mathrm{d}r \wedge \mathrm{d}\theta = \langle \partial_r u \lrcorner \mathrm{d}\mu, \xi\rangle\ \mathrm{d}r \wedge \mathrm{d}\theta,$$

where $\langle \cdot, \cdot \rangle$ again denotes the pairing between $\mathfrak{k}^*$ and $\mathfrak{k}$ (Cho–Oh did this computation in the toric case in [35, Theorem 8.1]). The right-hand side is simply $u^*\mathrm{d}\langle\mu,\xi\rangle \wedge \mathrm{d}\theta$, so the total energy is at most

$$2\pi(\max_X \langle \mu, \xi\rangle - \min_X \langle \mu, \xi \rangle) < \infty. \qquad \square$$

Remark 3.1.14. There are situations where this 'patching over 0' property fails. For example, if X is the complex torus $\mathbb{C}/(\mathbb{Z} \oplus i\mathbb{Z})$, L is $\mathbb{R}/\mathbb{Z} \subset X$, K is $\mathrm{U}(1)$ acting by translations in the real direction, and ξ is $i \in i\mathbb{R} = \mathfrak{u}(1)$, then the punctured disc $z \mapsto e^{-i\xi \log z} \cdot 0$ wraps around the torus infinitely many times as $z \to 0$. //

Evans–Lekili showed [45, Corollary 3.10] that under certain additional hypotheses on (X, L) all index 2 discs are axial. This result can be simplified and generalised, but first let us introduce some terminology. Given a smooth manifold M, a Lie group H acting on M, a vector subspace $\mathfrak{h}'$ of the Lie algebra $\mathfrak{h}$, and a submanifold $N \subset M$, we say N is $\mathfrak{h}'$-*invariant* if for all $p \in N$ we have $\mathfrak{h}' \cdot p \subset T_pN$, where the left-hand side represents the infinitesimal action of $\mathfrak{h}'$ on M at p. Note that the moduli space of parametrised discs in a given homotopy class carries a left action of K by translation, and a right action of $\mathrm{PSL}(2, \mathbb{R})$ by reparametrisation (recalling our identification $\mathrm{Aut}(D) \cong \mathrm{PSL}(2, \mathbb{R})$ from Section 1.2.4). It therefore makes sense to talk about a moduli space of parametrised discs being invariant under a subspace of $\mathfrak{k}$ or $\mathfrak{psl}(2, \mathbb{R})$ (for instance $\mathfrak{u}(1) \subset \mathfrak{psl}(2, \mathbb{R})$, acting by rotational reparametrisation).

Before stating the generalised axiality criterion, we need two straightforward results about $\mathfrak{psl}(2, \mathbb{R})$:

Lemma 3.1.15. *For $\eta \in \mathfrak{psl}(2, \mathbb{R}) \subset \mathrm{Mat}_{2\times 2}(\mathbb{C})$ the following are equivalent:*

(i) η acts on ∂D without fixed points, i.e. $\eta \cdot z \neq 0$ for all $z \in \partial D$.

(ii) $\det \eta > 0$.

(iii) Some real multiple of η is conjugate by an element of $\mathrm{SL}(2, \mathbb{R})$ to the generator

$$\rho = \begin{pmatrix} 0 & \frac{1}{2} \\ -\frac{1}{2} & 0 \end{pmatrix}$$

of a rotation, from Section 1.2.4.

Proof. (i) $\Longrightarrow$ (ii): Assuming η acts without fixed points, by compactness of ∂D we can pick $\varepsilon > 0$ such that $\|\eta \cdot z\| \geq \varepsilon$ for all $z \in \partial D$ (using the standard metric on ∂D). This means that as t increases from 0 the point $e^{t\eta} \cdot 1$ moves around the unit circle at speed at least ε, so at some time $T \in (0, 2\pi/\varepsilon]$ it returns to its starting point. In other words, there exists $T \in (0, 2\pi/\varepsilon]$ such that $e^{T\eta} \cdot 1 = 1$. An explicit calculation of $e^{T\eta} \cdot 1$ shows that $\det \eta > 0$.

(ii) $\implies$ (iii): Given $\eta \in \mathfrak{psl}(2,\mathbb{R})$ with $\det \eta > 0$, scale η to make its determinant $\frac{1}{4}$. Then η and ρ both have eigenvalues $\pm i/2$ (they are both trace-free), so are conjugate over $\mathbb{C}$. It is well-known that real matrices conjugate over $\mathbb{C}$ are conjugate over $\mathbb{R}$, so we have $\eta = g\rho g^{-1}$ for some $g \in \mathrm{GL}(n,\mathbb{R})$. Replacing g by $g \cdot \left(\begin{smallmatrix}-1 & 0\\ 0 & 1\end{smallmatrix}\right)$ and reversing the sign of η, if necessary, we may assume that $\det g > 0$. Dividing g by the square root of its determinant then ensures $g \in \mathrm{SL}(2,\mathbb{R})$ as required.

(iii) $\implies$ (i): This is immediate from the fact that ρ acts on ∂D without fixed points. $\square$

We'll say that a subspace of $\mathfrak{psl}(2,\mathbb{R})$ has a *global fixed point* (in ∂D) if there exists a $z \in \partial D$ which is fixed by the action of each element of the subspace.

Lemma 3.1.16. *If $\mathfrak{h}$ is a Lie subalgebra of $\mathfrak{psl}(2,\mathbb{R})$ without a global fixed point in ∂D then it contains an element η without a fixed point (i.e. satisfying the three conditions in Lemma 3.1.15).*

Proof. If $\mathfrak{h}$ is 1-dimensional then any non-zero element will do, whilst if it is all of $\mathfrak{psl}(2,\mathbb{R})$ the result is clear, so we are left to deal with the case where $\dim \mathfrak{h} = 2$, and we claim that this is actually impossible. So suppose that $\mathfrak{h}$ is a 2-dimensional Lie subalgebra of $\mathfrak{psl}(2,\mathbb{R})$. Let η_H, η_X, and η_Y be the standard basis elements

$$\begin{pmatrix}1 & 0\\ 0 & -1\end{pmatrix}, \begin{pmatrix}0 & 1\\ 0 & 0\end{pmatrix}, \text{and } \begin{pmatrix}0 & 0\\ 1 & 0\end{pmatrix}$$

of $\mathfrak{psl}(2,\mathbb{R})$, and let $\langle\cdot\rangle$ denote linear span. We cannot have $\mathfrak{h} = \langle \eta_X, \eta_Y\rangle$, since this space is not closed under the Lie bracket, so we can pick a basis for $\mathfrak{h}$ of the form $\eta_H + a\eta_X + b\eta_Y, c\eta_X + d\eta_Y$, with $a, b, c, d \in \mathbb{R}$ and c and d not both zero. Without loss of generality we may assume c is non-zero, and then change basis to be of the form $\eta_H + \kappa\eta_Y, \eta_X + \lambda\eta_Y$, with $\kappa, \lambda \in \mathbb{R}$. Closure under the Lie bracket forces $\kappa^2 = 4\lambda$, but then every η in $\mathfrak{h}$ fixes the point $(\kappa i/2 + 1)/(\kappa/2 + i)$ in ∂D. Thus $\mathfrak{h}$ has a global fixed point, so this case cannot occur, as claimed. $\square$

We can now prove the result we want:

Proposition 3.1.17. *Suppose (X, L) is K-homogeneous, and $\mathcal{M}$ is a smooth moduli space of parametrised non-constant holomorphic discs on L which is $\mathfrak{k}$- and $\mathfrak{psl}(2,\mathbb{R})$- (or $\mathfrak{u}(1)$-) invariant and of dimension at most $n+2$ (respectively n). Then all discs in $\mathcal{M}$ are axial.*

Proof. We deal with the $\mathfrak{psl}(2,\mathbb{R})$ case; the $\mathfrak{u}(1)$ case is analogous but simpler. Fix a disc u in $\mathcal{M}$ and let $\mathfrak{h}$ be the subspace of $\mathfrak{psl}(2,\mathbb{R})$ comprising those η for which there exists ξ in $\mathfrak{k}$ such that $\xi \cdot u = u \cdot \eta$. Note that $K \times \mathrm{PSL}(2,\mathbb{R})$ acts on $\mathcal{M}$ by $(k, \varphi) \cdot v = kv \circ \varphi^{-1}$ and $\mathfrak{h}$ is precisely the projection to $\mathfrak{psl}(2,\mathbb{R})$ of the infinitesimal stabiliser of u, so is actually a Lie subalgebra. We claim that it has no global fixed point.

Assuming this, by Lemma 3.1.16 there exists η in $\mathfrak{h}$ which acts without fixed points, and is thus conjugate to the generator of a rotation by Lemma 3.1.15. By reparametrising u we may assume η actually *does* generate a rotation, and by rescaling it we may assume that it rotates ∂D anticlockwise with period 2π. Moreover, by definition of $\mathfrak{h}$ there exists ξ in $\mathfrak{k}$ with $\xi \cdot u = u \cdot \eta$. This means that $e^{\theta\xi}u(z) = u(e^{i\theta}z)$ for all z in ∂D and all real θ, and hence that u is axial by Remark 3.1.12.

We are therefore left to show that $\mathfrak{h}$ has no global fixed point, so suppose for contradiction that z in ∂D is such a point. The tangent space $T_u\mathcal{M}$ consists of smooth sections of u^*TX, which are holomorphic over the interior of D and which lie in $u|_{\partial D}^*TL$ when restricted to ∂D, and we have the infinitesimal evaluation map $D\,\mathrm{ev}_z \colon T_u\mathcal{M} \to T_{u(z)}L$. By hypothesis this vanishes on the image of the action of $\mathfrak{h}$, which is precisely $(\mathfrak{k} \cdot u) \cap (u \cdot \mathfrak{psl}(2, \mathbb{R}))$, and so we have

$$(\mathfrak{k} \cdot u) \cap (u \cdot \mathfrak{psl}(2, \mathbb{R})) \subset (\mathfrak{k} \cdot u) \cap \ker D\,\mathrm{ev}_z\,. \tag{3.2}$$

Now, the dimension of the left-hand side is at least

$$\dim(\mathfrak{k} \cdot u) + \dim(u \cdot \mathfrak{psl}(2, \mathbb{R})) - \dim \mathcal{M} \geq \dim(\mathfrak{k} \cdot u) + 1 - n.$$

Here we are assuming that u is non-constant, and hence that the action of $\mathrm{PSL}(2, \mathbb{R})$ on u is infinitesimally free. The dimension of the right-hand side of (3.2), however, is

$$\dim(\mathfrak{k} \cdot u) - \dim D\,\mathrm{ev}_z(\mathfrak{k} \cdot u) = \dim(\mathfrak{k} \cdot u) - n.$$

This is because, by homogeneity, the $\mathfrak{k}$-action is surjective at $u(z)$. Therefore the containment (3.2) cannot hold. This gives the desired contradiction, and we deduce that no global fixed point z can exist, completing the proof. □

We recover the original result simply by setting $\mathcal{M} = \mathcal{M}_{\mu=2}$, the space of all index 2 discs. One can also show, again modifying an argument originating with Evans–Lekili:

Lemma 3.1.18 (See [45, Corollary 3.3])**.** *All non-constant holomorphic discs u on L have index at least 2. More generally, any smooth moduli space $\mathcal{M}$ of parametrised holomorphic discs on L which is $\mathfrak{k}$- and $\mathfrak{psl}(2, \mathbb{R})$- (or $\mathfrak{u}(1)$-) invariant and of dimension at most $n + 1$ (respectively $n - 1$) contains only constant discs.*

Proof. The first assertion follows from the second by taking $\mathcal{M}$ to be the moduli space $\mathcal{M}_{\mu=\mu(u)}$ of parametrised holomorphic discs of the same index as u, so suppose for contradiction that $\mathcal{M}$ satisfies the hypotheses of the second assertion and that u is a non-constant disc in $\mathcal{M}$ (again restricting to the $\mathfrak{psl}(2, \mathbb{R})$ case). This time, for all z in ∂D the dimension of the left-hand side of (3.2) exceeds that of the right- by at least 2, so

$$\dim D\,\mathrm{ev}_z(u \cdot \mathfrak{psl}(2, \mathbb{R})) \geq 2.$$

However, the image of $D\,\mathrm{ev}_z$ restricted to $u \cdot \mathfrak{psl}(2, \mathbb{R})$ is contained in the span of $izu'(z)$ (where $'$ denotes the derivative), which has dimension at most 1, giving a contradiction. □

3.1.4 Linear homogeneity

For some of our later analysis in Section 3.3 we will require a stronger notion of homogeneity:

Definition 3.1.19. If (X, L) is K-homogeneous, and there exist a complex inner product space V, with a holomorphic embedding $\iota \colon X \hookrightarrow \mathbb{P}V$, and a Lie group morphism $\rho \colon K \to \mathrm{GL}(V)$ with respect to which ι is K-equivariant, then we say (X, L) is *linearly K-homogeneous.* //

When we describe a pair (X, L) as linearly K-homogeneous we will implicitly fix a choice of V, ι and ρ as in the definition. Note that the representation ρ complexifies to a representation of G, which we also denote by ρ.

Remark 3.1.20. If (X, L) is K-homogeneous and the canonical or anticanonical bundle of X is ample then by the Kodaira embedding theorem (X, L) is automatically linearly K-homogeneous; see [92, Section III Theorem 9.1]. //

Remark 3.1.21. Suppose (X, L) is K-homogeneous and $c_1(X)$ is a positive (or negative) multiple of a Kähler class $[\omega]$. Standard arguments in Hodge theory [77, Proposition, page 148] show that in this case the anticanonical (respectively canonical) bundle is ample, so the preceding remark gives linear homogeneity. One can easily verify that the pullback of the Fubini–Study form ω_{FS} is proportional to ω in cohomology, and if ω is K-invariant then we claim the action of K on X is in fact Hamiltonian.

Indeed, all we need to show is that there exists a K-equivariant moment map $\mu\colon X \to \mathfrak{k}^*$ generating the K-action with respect to ω. Averaging over K we may assume that ω_{FS} is K-invariant, so the K-action has a moment map μ_{FS} with respect to ω_{FS} (because the action of $P\mathrm{U}(V)$ on $\mathbb{P}V$ is Hamiltonian). Since ω and $\omega_{\mathrm{FS}}|_X$ are proportional in $H^2(X, \mathbb{R})$, there exist a non-zero real number λ and a 1-form α on X with

$$\omega = \lambda\omega_{\mathrm{FS}} + \mathrm{d}\alpha.$$

Averaging over K again, we may assume that α is also K-invariant. Now define μ by

$$\langle \mu, \xi\rangle = \lambda\langle \mu_{\mathrm{FS}}, \xi\rangle - X_\xi \lrcorner \alpha$$

for all ξ in $\mathfrak{k}$. This is the required moment map.

In this situation we know from Proposition 3.1.13 that axial expressions of the form (3.1) patch over zero. This is true in general for linearly K-homogeneous (X, L) by considering the decomposition of $u(1)$ into eigenspaces of $-i\xi$: the value of $u(0)$ is the projection of $u(1)$ onto the eigenspace of lowest eigenvalue (the lowest eigenvalue for which this projection is non-zero). //

Before discussing some further consequences of linearity, we need a result about algebraicity of complexifications. This seems to be more or less well-known (although there is considerable variation in precisely what different authors mean by 'complexification'; we shall always use Definition 3.1.4), so we omit a proof. All of the necessary ideas can be found in [22, Chapter III].

Proposition 3.1.22. *The complexification of a compact subgroup of a general linear group is an algebraic subgroup.*

Using this we can prove:

Lemma 3.1.23. *If (X, L) is linearly K-homogeneous, with $\iota\colon X \hookrightarrow \mathbb{P}V$ and $\rho\colon G \to \mathrm{GL}(V)$ the corresponding embedding and Lie group morphism, then:*

(i) The image $\rho(G)$ is an algebraic group.

(ii) The G-orbits on X are locally closed in the Zariski topology induced by ι, and hence are complex submanifolds of X.

(iii) If the homogeneity is also sharp then for x in W the stabiliser $\rho(G)_x$ of x in $\rho(G)$ is finite.

Proof. (i) Apply Proposition 3.1.22 to $\rho(K)$.

(ii) The orbits of G on X coincide under ι with the orbits of the algebraic group $\rho(G)$, which acts linearly and hence algebraically on $\iota(X)$. By a standard application of Chevalley's theorem on constructible sets, these orbits are locally closed (see, for example, [89, Proposition in Section 8.3]). Each orbit is therefore open in its Zariski closure, and consists entirely of smooth points of the closure (since G acts transitively on it), so is an embedded complex submanifold of X.

(iii) For each x in W we have an algebraic orbit map $\varphi\colon \rho(G) \to W$ with

$$\varphi^{-1}(x) = \rho(G)_x.$$

The left-hand side is clearly Zariski closed, whilst the right-hand side is zero-dimensional if the homogeneity is sharp, and hence in this case both are finite. $\square$

Remark 3.1.24. If we replace K and G by their images in $\mathrm{GL}(V)$ then the stabiliser G_x must be finite, so is conjugate to a subgroup of K by Proposition 3.1.6(ii)—say $gG_xg^{-1} \subset K$ where g is an element of G. This means that if we also replace $L = Kx$ by Kgx, and our base point x by gx, then the new stabiliser $G_{gx} = gG_xg^{-1}$ is contained in K. //

Remark 3.1.25. Suppose (X, L) is K-homogeneous (not necessarily linearly) with X Kähler, and the action is Hamiltonian with L contained in a level set of the moment map. For any point x in L the set $\{g \in G : gx \in L\}$ is contained in (and therefore equal to) K, by an argument of Cieliebak–Gaio–Salamon [36, Remark 2.1(iii)]. If the homogeneity is sharp then we deduce that the stabiliser G_x is a zero-dimensional subgroup of K. It is therefore again a finite subgroup of K. //

Remark 3.1.26. If (X, L) is K-homogeneous then by replacing K with its quotient by the kernel of its action on L it is clear that we may assume that this action is faithful. We may make the same assumption for linear homogeneity, although it is less obvious: if X is embedded in $\mathbb{P}V$ then let W be the subspace of V which it spans; it is easy to check that $\mathbb{P}W$ is preserved by the action of K and that the image of K in $\mathrm{PSL}(W)$ acts faithfully on X; the $(\dim W)$-uple Veronese map then gives an embedding of X in $\mathbb{P}S^{\dim W}W$ and an embedding of $\mathrm{PSL}(W)$ in $\mathrm{SL}(S^{\dim W}W)$; replacing K by its image in the latter, we get the claimed faithful linear action. Note also that the Veronese map respects the Kähler forms up to overall scaling. //

3.2 The closed–open string map

3.2.1 The pearl model for $\mathcal{CO}^0$

In this section we formulate and prove a result which provides an effective tool for computing part of the length zero closed–open map. We begin by recapping the definition in terms of the pearl complex model for the quantum module action, given by Biran–Cornea in [14, Section 5.3], so fix a (not necessarily homogeneous) monotone Lagrangian brane $L^\flat \subset X$ over a ring R.

Suppose we wish to compute $\mathcal{CO}^0$ on a class α in $H^j(X;R) \subset QH^j(X;\Lambda)$. Fix a Morse–Smale pair (f,g) on L, which we assume has a unique local minimum p, and choose a pseudocycle $\check{\alpha}\colon M \to X$ representing the Poincaré dual class $\mathrm{PD}(\alpha) \in H_{2n-j}(X;R)$. Of course $\mathrm{PD}(\alpha)$ may really be an R-linear combination of pseudocycles but we suppress this in our notation. Pick a generic ω-compatible almost complex structure J' on X, and consider the pearl complex $C^* = C^*(L^\flat; \mathscr{D} = (f,g,J'); \Lambda)$. The element $p \otimes \mathrm{id}_{\mathscr{F}_p}$ is automatically closed and its class represents the unit 1_L, which may be zero (if and only if $HF^* = 0$). In order to achieve the necessary transversality we should also pick a generic Hamiltonian perturbation H (see [14, Section 5.3.7] and [104, Chapter 8] for a discussion of this notion), which we notationally incorporate into $\mathscr{D}$.

To find the coefficient of a critical point y in $\mathcal{CO}^0(\alpha)$, we consider rigid pearly trajectories γ from p to y in which one of the discs u is (J',H)-holomorphic and has $u(0) \in \check{\alpha}(M)$. This (J',H)-holomorphic disc is allowed to be constant, since it carries three marked points and so is stable, although the presence of the perturbation means that constant discs will probably not be pseudoholomorphic; all other discs are J'-holomorphic (no perturbation) but must be non-constant. Now the virtual dimension is given by $|y| - j + \mu(A)$, where $A \in H_2(X,L;\mathbb{Z})$ is the total class of the discs as usual. The signed count of such trajectories γ, weighted by $T^{\nu(A)}$ and

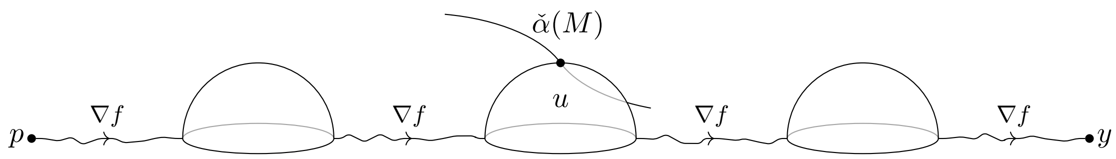

Figure 3.1: A pearly trajectory contributing to $\langle \mathcal{CO}^0(\alpha), y\rangle$.

$\mathcal{P}(\gamma_{\mathrm{t}}) \circ \mathcal{P}(\gamma_{\mathrm{b}})^{-1} \in \mathrm{End}(\mathscr{F}_y)$, is the required coefficient. If there is an H_2^D local system present, defined by a homomorphism $\rho\colon H_2^D \to R^\times$, then one also has to include a factor of $\rho(A)$.

Remark 3.2.1. Biran–Cornea and Zapolsky use a Morse model for QH^*, rather than a pseudocycle model, in which the disc with the interior marked point is required to meet the descending manifold of an input critical point on X. //

As an illustration, we now recall the the Auroux–Kontsevich–Seidel criterion:

Proposition 3.2.2 (See [9, Lemma 6.7, Proposition 6.8], [130, Lemma 2.7])**.** *Let $L^\flat \subset X$ be a monotone Lagrangian brane with rank 1 local system over a field R, and let $\mathfrak{m}_0$ be the signed count of pseudoholomorphic index 2 discs through a generic point of L, weighted by their boundary monodromies with respect to the local system. If $HF^*(L^\flat, L^\flat; R)$ is non-zero then $\mathfrak{m}_0$ is an eigenvalue of quantum multiplication by the first Chern class*

$$c_1(X)*\colon QH^*(X;R) \to QH^*(X;R).$$

Proof. Since L is orientable its Maslov class $\mu \in H^2(X,L;\mathbb{Z})$ is divisible by 2, and by Poincaré duality we can pick a cycle $Y \subset X \setminus L$ representing $\mu/2$. Note that since $\mu/2$ maps to $c_1(X)$ in $H^2(X;\mathbb{Z})$ the cycle Y represents $\mathrm{PD}(c_1(X))$. We now compute $\mathcal{CO}^0(c_1)$ using this representative.

The contributions are either trajectories of index 0 with outputs of index 2 or trajectories of index 2 with outputs of index 0. By construction, each index 2 disc bounded by L has intersection number 1 with Y, so for the index 2 trajectories we can just ignore the incidence condition with Y and count things of the form 'flow up from the minimum of the Morse function, enter an index 2 disc, and exit at the minimum'. Generically this amounts to simply counting index 2 discs through the minimum, twisted by the local system, which gives $\mathfrak{m}_0$. Each index 0 disc has intersection number 0 with Y, so all index 0 trajectories cancel and we obtain $\mathcal{CO}^0(c_1) = \mathfrak{m}_0 \cdot 1_L$.

We therefore have $\mathcal{CO}^0(c_1 - \mathfrak{m}_0 \cdot 1_X) = 0_L$, so if $HF^*(L^\flat, L^\flat; R)$ is non-zero then $c_1 - \mathfrak{m}_0 \cdot 1_X$ cannot be invertible in $QH^*(X; R)$: $\mathcal{CO}^0$ maps invertibles in $QH(X; R)$ to invertibles in $HF^*(L^\flat, L^\flat; R)$, and if the latter is non-zero then 0_L is not invertible. So quantum multiplication $(c_1 - \mathfrak{m}_0 \cdot 1_X)*$ is singular, and thus $\mathfrak{m}_0$ is an eigenvalue of c_1*. □

We will revisit this result later, in Example 3.2.29, and verify that it is consistent with our results for homogeneous Lagrangians. In the sharply homogeneous case we will also observe that the relation $\mathcal{CO}^0(c_1) = \mathfrak{m}_0$ is true *taking into account orientations*, an issue which we glossed over above and which (to the best of the author's knowledge) has not been explicitly addressed in the literature.

Remark 3.2.3. Strictly of course one needs to be careful about the genericity of the auxiliary data: the Morse function, metric, almost complex structure and Hamiltonian perturbation. The machinery of Biran–Cornea (for example [15, Proposition 3.1.2]) allows one to fix the first two, and then choose the latter two (J' and H) generically, so if one is using a particular integrable complex structure J to compute $\mathfrak{m}_0$ it is enough to know that it gives the same answer as a generic (J', H). And for this it is enough to know that the index 2 discs are regular for J (i.e. all of their partial indices are at least -1), and that there are no J-holomorphic index 2 spheres passing through our Morse minimum on L. A standard cobordism argument shows that the values of $\mathfrak{m}_0$ calculated with respect to J and (J', H) agree. //

Remark 3.2.4. As outlined in [14] after Remark 5.3.11, if we restrict to perturbations with curvature bounded by an appropriate constant c, which is sufficient for our purposes, then Gromov compactness remains valid for discs of bounded energy (despite the Hamiltonian perturbation) but now the limit bubbled configuration comprises a single disc component satisfying the perturbed Cauchy–Riemann equation and possibly other disc and sphere components satisfying the unperturbed equation. //

Remark 3.2.5. Since the ascending manifold of the minimum p is a dense open subset of L, if any discs occur before the marked disc in Fig. 3.1 then generically we can delete them and obtain a new trajectory of strictly lower index. We therefore see that for the purpose of counting *rigid* trajectories we may assume that there are no such discs. The condition that there is a flowline from p to the incoming boundary point on the marked disc is also (generically) vacuous, so in fact we can just count trajectories starting from a marked disc and ending at y. //

Suppose the class α has degree $2k$, so our representing pseudocycle has codimension $2k$. For notational simplicity we view this pseudocycle as a genuine submanifold Z of X. Let $\mathcal{M}$ temporarily denote the moduli space of (J', H)-holomorphic discs of index $2k$, carrying evaluation maps ev_0 and ev_1 to X and L respectively. There is a contribution—which in general may be zero—to $\mathcal{CO}^0(\alpha)$ from trajectories which output the minimum p and contain a single disc u, which must lie in $\mathrm{ev}_0^{-1}(Z) \cap \mathrm{ev}_1^{-1}(p) \subset \mathcal{M}$. For our later sign computations we assume the following:

Assumption 3.2.6. Such a trajectory contributes towards $\mathcal{CO}^0(\alpha)$ with sign given by the orientation sign of the isomorphism

$$D_u \operatorname{ev}_0 \oplus D_u \operatorname{ev}_1 \colon T_u\mathcal{M} \to (T_{u(0)}X/T_{u(0)}Z) \oplus T_pL,$$

where $T_{u(0)}X/T_{u(0)}Z$ is oriented by the short exact sequence

$$0 \to T_{u(0)}Z \to T_{u(0)}X \to T_{u(0)}X/T_{u(0)}Z \to 0. \tag{3.3}$$

//

In Section B.1 we verify this assumption for the Biran–Cornea and Zapolsky orientation schemes.

Remark 3.2.7. The reader is warned that there are two natural conventions for orienting a short exact sequence

$$0 \to A \to B \to C \to 0,$$

both of which are in use in the literature. One can either use the induced isomorphism $B \cong A \oplus C$, which is canonical up to homotopy, or the corresponding isomorphism with the order of A and C reversed. We'll call these 'left-to-right' and 'right-to-left' respectively. Biran and Cornea use the latter, but when A or C is even-dimensional, as in (3.3), the distinction is irrelevant. //

3.2.2 Discs meeting invariant subvarieties

Suppose now that (X, L) is K-homogeneous. We have just seen that computing the closed–open map requires understanding moduli spaces of discs on L with a constrained interior marked point, so let us introduce the following notation:

Definition 3.2.8. For a non-negative integer j and a subset $Z \subset X$ let

$$\mathcal{M}^Z_{\mu=j} = \{u \in \mathcal{M}_{\mu=j} : u(0) \in Z\}$$

be the space of parametrised holomorphic discs of index j mapping 0 to Z. Here holomorphicity is with respect to the fixed integrable complex structure J on X. //

In order for such moduli spaces to be smooth manifolds, we would like Z to be a submanifold of X to which the evaluation at 0 map, ev_0, is transverse. In the homogeneous setting we automatically have this property in a natural class of situations:

Lemma 3.2.9. *If Z is a submanifold of X which is $i\mathfrak{k}$-invariant (using the complexified action from Lemma 3.1.8), and $u \colon (D, \partial D) \to (X, L)$ is a holomorphic disc with $u(0) \in Z$, then the evaluation at 0 map $\mathrm{ev}_0 \colon \mathcal{M}_{\mu=\mu(u)} \to X$ is transverse to Z at u.*

Proof. Let p denote $u(0) \in Z$, and pick $\xi_1, \dots, \xi_n \in \mathfrak{k}$ such that the holomorphic sections $v_1, \dots, v_n$ of u^*TX, given by $v_j(z) = \xi_j \cdot u(z) \in T_{u(z)}X$ for all $z \in D$, span $T_{u(1)}X$ at 1 as a complex vector space; we can do this since the infinitesimal K-action on L is surjective. The section $\sigma \coloneqq v_1 \wedge \cdots \wedge v_n$ of $u^*\Lambda^n_{\mathbb{C}}TX$ is then holomorphic and not identically zero, so if it vanishes at 0 then this zero must be isolated. Therefore the v_j are fibrewise linearly independent on a punctured open neighbourhood of 0 in D.

Now take a vector $w \in T_pX$. We wish to find a tangent vector to u in $\mathcal{M}_{\mu=\mu(u)}$, i.e. a holomorphic section v of u^*TX which restricts to a section of $u|_{\partial D}^*TL$ on the boundary, such that $w - v(0)$ lies in T_pZ.

First we extend w arbitrarily to a holomorphic local section $\widetilde{w}$ of u^*TX. Then on a punctured open neighbourhood of 0 there exist holomorphic functions $f_1, \ldots, f_n$ such that

$$\widetilde{w} = \sum_j f_j v_j.$$

For each j we have

$$v_1 \wedge \ldots v_{j-1} \wedge \widetilde{w} \wedge v_{j+1} \wedge \cdots \wedge v_n = f_j \sigma$$

and the left-hand side is holomorphic over 0, so if k denotes the vanishing order of σ at 0 then f_j has at worst a pole of order k.

Let the principal part of the Laurent series of f_j (including the constant term) be

$$\sum_{l=0}^{k} a_{jl} z^{-l},$$

and let a_{j0} have real and imaginary parts b_j and c_j respectively. Now define a meromorphic function F_j for each j by

$$F_j = b_j + \sum_{l=1}^{k} \Big(a_{jl} z^{-l} + \overline{a}_{jl} z^{l} \Big).$$

Note that $F_j|_{\partial D}$ is real, and on a punctured open neighbourhood of 0 we have $f_j - F_j = ic_j + r_j$, where r_j is a holomorphic function which vanishes at 0.

Finally, let v be the meromorphic section of u^*TX given by

$$v = \sum_j F_j v_j.$$

On a punctured neighbourhood of 0 we have

$$\widetilde{w} - v = \sum_j (f_j - F_j) v_j = \sum_j (ic_j + r_j) \xi_j \cdot u,$$

so v is in fact holomorphic, and

$$w - v(0) = \sum_j ic_j \xi_j \cdot p \in i\mathfrak{k} \cdot p \subset T_pZ.$$

Since $F_j|_{\partial D}$ is real for all j, $v|_{\partial D}$ lies in $u|_{\partial D}^*TL$. This v is therefore a tangent vector to $\mathcal{M}_{\mu=\mu(u)}$ at u with the desired properties. $\square$

We immediately deduce:

Corollary 3.2.10. *If Z is an $i\mathfrak{k}$-invariant submanifold of X of codimension m, then $\mathcal{M}^Z_{\mu=j}$ is a smooth manifold of dimension $n+j-m$.*

In order to apply Proposition 3.1.17 and Lemma 3.1.18 to the spaces $\mathcal{M}^Z_{\mu=j}$ we need them to be $\mathfrak{k}$-invariant. Letting G denote the complexification of K as before, we see that a natural family of choices for Z is given by $\mathfrak{g}$-invariant submanifolds of X.

Proposition 3.2.11. *Let Z be a $\mathfrak{g}$-invariant submanifold of X disjoint from L. If m denotes the (real) codimension of Z then for all integers $j < m$ we have $\mathcal{M}^Z_{\mu=j} = \emptyset$, and for $j = m$ all discs in $\mathcal{M}^Z_{\mu=j}$ are axial.*

Proof. Combine Corollary 3.2.10 with the $\mathfrak{u}(1)$ versions of Proposition 3.1.17 and Lemma 3.1.18. (The reason for requiring Z to be disjoint from L in this result is to rule out constant discs in $\mathcal{M}^Z_{\mu=0}$.) □

Remark 3.2.12. Since $\mathcal{M}^Z_{\mu=m}$ has dimension n we see that the action of $\mathfrak{k}$ on any disc u in this space cannot have rank greater than n. In particular, if $u\colon z \mapsto e^{-i\xi \log z}x$ is such a disc and $\mathfrak{k}_x$ is the Lie algebra of the stabiliser of x in K, then $\mathfrak{k}_x$ must stabilise the whole boundary of u, so we must have

$$\mathfrak{k}_x \cdot e^{\theta\xi}x = 0$$

for all $\theta \in \mathbb{R}$. Therefore $e^{-\theta\xi}\mathfrak{k}_x e^{\theta\xi} = \mathfrak{k}_x$ for all θ, so $[\xi, \mathfrak{k}_x] \subset \mathfrak{k}_x$; in other words, ξ is in the normaliser of $\mathfrak{k}_x$ in $\mathfrak{k}$. If this normaliser is trivial, meaning equal to $\mathfrak{k}_x$ itself, then we see that u must be constant, and hence that $\mathcal{M}^Z_{\mu=m}$ is in fact empty if Z is disjoint from L.

More generally, suppose u is an element of $\mathcal{M}^Z_{\mu=j}$ with $u(1) = x$ and that

$$\left.\frac{\mathrm{d}}{\mathrm{d}\theta}\right|_{\theta=0} u(e^{i\theta}) = \xi \cdot x.$$

Then the kernel of the $\mathfrak{k}$-action on u is contained in $\mathfrak{k}_x \cap \mathrm{ad}(\xi)^{-1}(\mathfrak{k}_x)$, so if

$$d := \dim \mathfrak{k}_x - \max_{\xi \in \mathfrak{k} \setminus \mathfrak{k}_x} \dim\bigl(\mathfrak{k}_x \cap \mathrm{ad}(\xi)^{-1}(\mathfrak{k}_x)\bigr) \tag{3.4}$$

then this action has rank at least $n + d$. So Proposition 3.2.11 holds with $j < m + d$ and $j = m + d$ in place of $j < m$ and $j = m$. The condition that $\mathfrak{k}_x$ has trivial normaliser is equivalent to $d \geq 1$. //

Example 3.2.13. To illustrate this remark consider the $\mathrm{SO}(n+1)$-homogeneous Lagrangian $\mathbb{RP}^n$ in $\mathbb{CP}^n$, and without loss of generality let $x = [1 : 0 : \cdots : 0]$. Taking ξ to have top left-hand corner

$$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix},$$

with zeros elsewhere, we obtain $\mathfrak{k}_x \cong \mathfrak{so}(n)$ and $\mathfrak{k}_x \cap \mathrm{ad}(\xi)^{-1}(\mathfrak{k}_x) \cong \mathfrak{so}(n-1)$, so from (3.4) we deduce that $d \leq n - 1$. This shows there are no non-constant holomorphic discs of index less than $n+1$—which is of course well-known in this case for purely topological reasons—and all discs with index equal to $n+1$ are axial. //

We can now prove the main result of this subsection. Note that we switch from talking about *real* codimension to *complex* codimension.

Corollary 3.2.14. *Let (X, L) be K-homogeneous and $Z \subset X$ an analytic subvariety of complex codimension k, which is invariant (setwise) under the action of the complexification G of K. Then:*

(i) No holomorphic discs of index less than $2k$ hit Z.

(ii) All index $2k$ discs hitting Z are axial and meet Z inside its smooth locus.

(iii) The moduli space $\mathcal{M}^Z_{\mu=2k}$ of parametrised index $2k$ discs mapping 0 to Z is a smooth n-manifold, and evaluation at 1, $\mathrm{ev}_1 \colon \mathcal{M}^Z_{\mu=2k} \to L$, is a covering map.

Proof. If $Z = X$ then the result is trivial, so we may as well assume that this is not the case, and thus that Z is disjoint from L. It is well-known (see, for example, the first Proposition on page 21 of [77]) that the singular locus of an analytic variety in a complex manifold is a proper subvariety. Moreover, if the original variety is G-invariant then so is the singular locus. Starting with Z, breaking it into its smooth locus and the connected components of its singular locus, and then applying the same process recursively to these components, we obtain a partition of Z into smooth $\mathfrak{g}$-invariant submanifolds $Z_1, \dots, Z_l$ of X, with Z_1 the smooth locus of Z and Z_r of real codimension greater than $2k$ in X for all $r > 1$.

Now consider the moduli spaces

$$\mathcal{M}^{Z_r}_{\mu=j},$$

for $j \leq 2k$. By Proposition 3.2.11, these are all empty except possibly when $r = 1$ and $j = 2k$, in which case the moduli space contains only axial discs. This proves (i) and (ii). We see that

$$\mathcal{M}^Z_{\mu=2k} = \mathcal{M}^{Z_1}_{\mu=2k},$$

which is a smooth n-manifold by Lemma 3.2.9, and the only thing left to show is that ev_1 is a covering map. This is a straightforward piece of differential topology using equivariance of ev_1 and the fact that the infinitesimal action of K on L is surjective. □

Remark 3.2.15. Note that:

(i) A complex submanifold of X is G-invariant if and only if it is K-invariant, so (applying this to each Z_j) we could have stated the result for subvarieties Z which are K-invariant. This would eliminate any mention of G from the statement.

(ii) There is no requirement that the K-homogeneity of (X, L) be sharp or linear, that the action be Hamiltonian, or that X or L be monotone. We haven't even assumed that X is symplectic and L Lagrangian, only that X is complex and L totally real.

(iii) Compactness of X was only used to establish the existence of the complexified action, which in applications one may know for other reasons. //

Remark 3.2.16. Evans–Lekili proved a special case of this result [45, Corollary 3.11], which says roughly that index 4 discs cleanly intersecting a K-invariant smooth subvariety of real codimension 4 are axial. Note that Lemma 3.2.9 shows that the clean intersection condition is automatically satisfied. ∥

If $\mathcal{M}$ is a single component of the moduli space occurring in Corollary 3.2.14(iii) then the action of K on $\mathcal{M}$ by translation is transitive: the orbits are closed, since K is compact, and open since the infinitesimal action is surjective (the infinitesimal action on L is surjective and factors through $\mathcal{M}$). If we also assume that the homogeneity is sharp, and fix a point $x \in L$ and a disc $u \in \mathcal{M} \cap \mathrm{ev}_1^{-1}(x)$, then we see that the corresponding orbit maps $K \to \mathcal{M}$ and $K \to L$ are covering maps. We thus get a tower of covers $K \to \mathcal{M} \to L$.

Proposition 3.2.17. *In this situation, if K_x and K_u denote the (finite) stabilisers of x and u in K then the (unsigned) degree of $\mathrm{ev}_1|_{\mathcal{M}}$ is given by the index of K_u in K_x.*

Proof. The degrees of the covers $K \to \mathcal{M}$ and $K \to L$ are given by the orders of K_u and K_x respectively. Considering the tower $K \to \mathcal{M} \to L$, we get the result. □

3.2.3 $\mathcal{CO}^0$ for homogeneous Lagrangians

We are now ready to apply this disc analysis to the computation of $\mathcal{CO}^0$, so assume that $L^\flat$ is a K-homogeneous monotone Lagrangian brane over a ring R. Let Z be a G-invariant analytic subvariety of X of complex codimension k; its Poincaré dual $\mathrm{PD}(Z)$ defines a class α in $QH^{2k}(X; \Lambda)$. If Z' denotes the smooth locus of Z then the homology class corresponding to Z is represented by the pseudocycle $\check{\alpha}$ defined to be the inclusion of Z' in X.

First we introduce some notation for the auxiliary data that we need. As in Section 3.2.1, fix a Morse–Smale pair (f, g) on L, with unique local minimum p. If J' is a smooth almost complex structure on X compatible with ω, and H is a smooth Hamiltonian perturbation with curvature bounded above by c (the bound from Remark 3.2.4), then for each j let

$$\mathcal{M}^Z_{\mu=j}(J', H)$$

denote the moduli space of (J', H)-holomorphic discs $u\colon (D, \partial D) \to (X, L)$ with $u(0) \in Z$. Recall that we are reserving the symbol J for the standard integrable complex structure on X, and note that $\mathcal{M}^Z_{\mu=j}(J, 0)$ is the space we called $\mathcal{M}^Z_{\mu=j}$ earlier. Each $\mathcal{M}^Z_{\mu=j}(J', H)$ carries an evaluation map ev_1 to L, sending u to $u(1)$.

Let $\mathcal{A}_c$ denote the space of pairs (J', H) as above. Recall that a pseudoholomorphic curve is called *regular* if the linearisation of the $\overline{\partial}$-operator at that curve is surjective. Let $\mathcal{A}_{\mathrm{trans}}$ be the set of those $(J', H) \in \mathcal{A}_c$ such that the following configurations (of any index) are transversely cut out, meaning that the underlying curves are regular and the evaluation maps and submanifolds which define the incidence conditions are transverse:

(i) (J', H)-holomorphic discs u on L with $u(0) \in Z$ and $u(1) = p$.

(ii) Simple $(J', 0)$-holomorphic spheres u_1 in X and index 0 (J', H)-holomorphic discs u_2 on L with $u_1(0) \in Z$, $u_1(\infty) = u_2(0)$ and $u_2(1) = p$.

When talking about transversality to Z we are implicitly taking a stratification of Z by smooth subvarieties, as in the proof of Corollary 3.2.14, and asking for transversality to each stratum. Really we only need this transversality for curves of index at most $2k$, but we may as well ask for it for curves of arbitrary index.

For each $(J', H) \in \mathcal{A}_{\text{trans}}$, the moduli spaces of configurations of types (i) and (ii), of a given index at most $2k$, are smooth manifolds of the expected dimension. The reason for the index restriction is so that all of these curves avoid the singular locus of Z. Standard arguments along the lines of [104, Theorem 3.1.5] show that $\mathcal{A}_{\text{trans}}$ forms a second category subset of $\mathcal{A}_c$. Since we are allowing Hamiltonian perturbations, we do not need to restrict to somewhere injective discs to achieve this transversality; see [14, Remark 5.3.11].

Lemma 3.2.18. *$\mathcal{A}_{trans}$ contains $(J, 0)$.*

Proof. The results of Section 3.2.2 imply that the condition is satisfied for configurations of type (i), so consider now configurations of type (ii). These are just simple $(J, 0)$-holomorphic spheres u with $u(\infty) = p$ and $u(0) \in Z$. So let $\mathcal{M}$ denote the moduli space of parametrised holomorphic spheres $u\colon \mathbb{CP}^1 \to X$ of fixed index with $u(\infty) = p$. Using the G-action to generate sections of u^*TX, we see that the fibre over ∞ is generated by global sections and hence that u is regular and ev_∞ is transverse to p at u. Thus $\mathcal{M}$ is a smooth manifold of the correct dimension, and we are left to show that $\mathrm{ev}_0\colon \mathcal{M} \to X$ is transverse to each stratum of Z.

To do this, we argue analogously to Lemma 3.2.9. First take global sections $v_j : z \mapsto \xi_j \cdot u(z)$ of u^*TX which generate the fibre over ∞. Extend a given vector $w \in T_{u(0)}X$ (with $u(0) \in Z$) to a local holomorphic section $\widetilde{w}$ of u^*TX, and express $\widetilde{w}$ as $\sum_j f_j v_j$ on a punctured neighbourhood of 0 for some meromorphic f_j. Now let F_j be the principal part of f_j, excluding the constant term, and define a global section v of u^*TX by $\sum_j F_j v_j$. This vanishes at ∞, so lies in $T_u\mathcal{M}$, and at 0 it differs from w by a vector in $\mathfrak{g} \cdot u(0) \in T_{u(0)}Z$. □

Finally, let $\mathcal{A}_{\mathcal{CO}^0} \subset \mathcal{A}_c$ denote those (J', H) which can be used to compute $\mathcal{CO}^0(\mathrm{PD}(Z))$ using the pearl model described above. The pearl complex machinery of Biran–Cornea [14], with the straightforward modifications to work with pseudocycles (as done by Sheridan in [130, Section 2.5], for example), means that $\mathcal{A}_{\mathcal{CO}^0}$ is also of second category in $\mathcal{A}_c$. In particular, $\mathcal{A}_{\text{trans}} \cap \mathcal{A}_{\mathcal{CO}^0}$ is dense in $\mathcal{A}_c$.

Definition 3.2.19. Let N_X^+ denote the minimal Chern number of rational curves in X, i.e. the unique non-negative integer with

$$\langle c_1(X), \text{holomorphic maps } \mathbb{CP}^1 \to X\rangle = N_X^+\mathbb{Z}_{\geq 0}.$$

This is clearly divisible by N_X, so it cannot be smaller unless it's zero, but in general it may be strictly greater. For example, $\mathbb{CP}^1 \times \mathbb{CP}^2$ has $N_X = 1$ but $N_X^+ = 2$. //

Lemma 3.2.20. *If (X, L) and Z are as above, and $N_X^+ = 0$ or $k \leq N_X^+ + 1$, then there exists a C^∞-open neighbourhood U of $(J, 0)$ in $\mathcal{A}_c$ such that for all $(J', H) \in U$ and all $j < 2k$ the space $\mathcal{M}^Z_{\mu=j}(J', H)$ is empty.*

Proof. Suppose for contradiction that no such U exists. Then there exists a $j < 2k$ and a sequence (J_l, H_l) in $\mathcal{A}_c$, converging to $(J, 0)$ in the C^∞-topology, such that for each l the space

$\mathcal{M}^Z_{\mu=j}(J_l, H_l)$ is non-empty; pick a disc u_l in this space. By standard Gromov compactness results, as mentioned above in Remark 3.2.4, there exists a (J-)holomorphic stable map u to which a subsequence of (u_l) Gromov converges. Since each u_l meets Z, which is compact, u must also meet Z.

Because any non-constant holomorphic disc has index at least 2, and any non-constant holomorphic sphere has index at least $2N_X^+ \geq 2(k-1) \geq j-1$, we see that either u consists of a tree of discs, or is a sphere, attached to a constant 'ghost' disc. Suppose first that u consists of a tree of discs, and let $\hat{u}$ be a component of u hitting Z. After reparametrising we may assume that $\hat{u}(0) \in Z$. Then we have

$$\hat{u} \in \mathcal{M}^Z_{\mu=\mu(\hat{u})},$$

but the index $\mu(\hat{u})$ of $\hat{u}$ is strictly less than $2k$, and so this moduli space is empty by Corollary 3.2.14(i). Therefore this case is impossible.

Otherwise u consists of a single holomorphic sphere $u \colon \mathbb{CP}^1 \to X$ meeting Z. The sphere must also meet L (at the ghost disc), at some point $q \in L$. By the argument used to prove Lemma 3.2.18 the moduli space of parametrised holomorphic spheres taking two fixed marked points to Z and q is transversely cut out. Since its virtual dimension is $j - 2k < 0$, we see that it must in fact be empty, completing the proof. □

From this lemma we deduce:

Corollary 3.2.21. *Under the hypotheses of Lemma 3.2.20, $\mathcal{CO}^0(\alpha)$ lies in*

$$[\Lambda \cdot p \otimes \operatorname{End} \mathscr{F}_p] \subset HF^*.$$

In particular, if the local system $\mathscr{F}$ has rank 1 then $\mathcal{CO}^0(\alpha)$ is a (Λ-)scalar multiple of 1_L.

Proof. Let U be the neighbourhood of $(J, 0)$ provided by Lemma 3.2.20, and pick a $(J', H) \in \mathcal{A}_{\mathcal{CO}^0} \cap U$. Computing $\mathcal{CO}^0(\alpha)$ using the pearl model with auxiliary data (J', H) and pseudocycle $\breve{\alpha} \colon Z' \hookrightarrow X$, we see that for any critical point y of positive Morse index there are no trajectories contributing to its coefficient. Hence $\mathcal{CO}^0(\alpha)$ lies in the subspace of HF^* spanned by the fibre of $\operatorname{End} \mathscr{F}$ over p. If $\mathscr{F}$ has rank 1 then this fibre is just R, so $\mathcal{CO}^0(\alpha) \in [\Lambda \cdot p] = \Lambda \cdot 1_L$. □

Remark 3.2.22. The surprising significance of characteristic 5 in relation to the Chiang Lagrangian, as discovered by Evans–Lekili [45], and the need to use coefficients of other positive characteristics for several other classes of examples [44], suggest that working modulo a prime may frequently be important when homogeneous Lagrangians are involved.

This phenomenon was partially explained for the Platonic Lagrangians using the closed–open map, in the form of the Auroux–Kontsevich–Seidel criterion in [45, Remark 1.2], and then extended to certain orbits of complex codimension 2 in [135, Corollary 4.25]. The results of this section generalise these ideas further, and the key point is really Corollary 3.2.21: after setting the Novikov parameter to 1 and collapsing the grading, a significant part of the image of $\mathcal{CO}^0$ is forced to lie in the subring of HF^* generated by the unit, which is just the image of the unique ring homomorphism $\mathbb{Z} \to R$ if $HF^* \neq 0$, and hence restrictions are placed on the arithmetic in this ring. //

Remark 3.2.23. Evans and Lekili have shown [44] that under appropriate hypotheses (in particular, the homogeneity should be free) if HF^* is non-zero then L split-generates a certain summand $D^\pi \mathcal{F}(X)_\lambda$ of the derived Fukaya category of X, and that the Hochschild cohomology of this summand is isomorphic to the corresponding piece $QH^*(X)_\lambda$ of the quantum cohomology of X. It is natural to conjecture that this isomorphism is induced by the full closed–open map

$$\mathcal{CO} : QH^*(X)_\lambda \to HH^*(D^\pi \mathcal{F}(X)_\lambda).$$

In contrast, we have just seen that the composition of this map with the projection to $HF^*(L, L)$ only has rank 1 on the subring of $QH^*(X)_\lambda$ generated by low codimension K-invariant subvarieties. //

Now we prove the main result of this subsection, which expresses $\mathcal{CO}^0(\alpha)$ in terms of a disc count that can be computed in practice, under a slightly stronger codimension condition on Z (or, from a different perspective, a stronger positivity condition on $c_1(X)$). Recall that by Corollary 3.2.14(iii) the moduli space $\mathcal{M}^Z_{\mu=2k}$ is a smooth, R-oriented n-manifold and that its boundary point evaluation map ev_1 to L is a covering map.

Proposition 3.2.24. *Suppose $L^\flat \subset X$ is a K-homogeneous monotone Lagrangian brane and $Z \subset X$ is a K-invariant analytic subvariety complex codimension k, defining a homology class with Poincaré dual $\alpha \in H^{2k}(X; R) \subset QH^{2k}(X; \Lambda)$. If $N_X^+ = 0$ or $k \leq N_X^+$, then the moduli space $\mathcal{M}^Z_{\mu=2k}$ is compact and*

$$\mathcal{CO}^0(\alpha) = T^{2k/N_L} \sum_{u \in \mathrm{ev}_1^{-1}(p)} \pm [p \otimes \widetilde{m}(\partial u)] \in [\Lambda \cdot p \otimes \operatorname{End} \mathscr{F}_p],$$

where $\widetilde{m}$ is the composite twisting homomorphism from Definition 2.2.13.

Under Assumption 3.2.6, if $L^\flat$ is sharply homogeneous and equipped with the standard spin structure then each sign on the right-hand side is $+1$.

Proof. Let U be the neighbourhood from Lemma 3.2.20. Take a (J', H) in

$$\mathcal{A}_{\mathrm{trans}} \cap \mathcal{A}_{\mathcal{CO}^0} \cap U,$$

and pick a homotopy $(J_t, H_t)_{t \in [0,1]}$ from $(J, 0)$ to (J', H) in U which is regular in the following sense. We require that the $\mathcal{A}_{\mathrm{trans}}$ transversality conditions are satisfied at the end points $t = 0$ and $t = 1$ (these hold by Lemma 3.2.18 and by construction respectively), and that the relevant derivatives remain surjective for $t \in (0, 1)$ when one also allows input tangent vectors in the t-direction. Standard arguments as in [104, Theorem 3.1.7] show that such homotopies are generic.

The regularity condition ensures that the space

$$\mathcal{M}_{\mathrm{cob}} \coloneqq \{(u \colon (D, \partial D) \to (X, L), t \in [0, 1]) : u \text{ is } (J_t, H_t)\text{-holomorphic}, \\ \mu(u) = 2k,\ u(0) \in Z, \text{ and } u(1) = p\}$$

is transversely cut out, and thus defines a one-dimensional cobordism between the corresponding moduli spaces for $t = 0$ and $t = 1$, which we denote by $\mathcal{M}_{t=0}$ and $\mathcal{M}_{t=1}$ respectively. Explicitly, if ev_1 and ev'_1 are the 'evaluation at 1' maps on $\mathcal{M}^Z_{\mu=2k}$ and $\mathcal{M}^Z_{\mu=2k}(J', H)$ (we called all such evaluation maps ev_1 previously, but we introduce the notation ev'_1 temporarily to distinguish between the two choices of auxiliary data) then

$$\mathcal{M}_{t=0} = \mathrm{ev}_1^{-1}(p) \subset \mathcal{M}^Z_{\mu=2k}(J, 0) = \mathcal{M}^Z_{\mu=2k}$$

and

$$\mathcal{M}_{t=1} = \mathrm{ev}_1'^{-1}(p) \subset \mathcal{M}^Z_{\mu=2k}(J', H).$$

Arguing as in the proof of Corollary 3.2.21, each trajectory contributing to $\mathcal{CO}^0(\alpha)$ has output p, and comprises a single disc represented by a point of $\mathcal{M}_{t=1}$. Each of these points carries a sign indicating whether the corresponding trajectory counts positively or negatively towards $\mathcal{CO}^0$. By the first part of Assumption 1.2.4 there are similar signs attached to the points of $\mathcal{M}_{t=0}$, and by the second part these signs on $\mathcal{M}_{t=0}$ and $\mathcal{M}_{t=1}$ are equal and opposite respectively (or vice versa) to the signs induced by a certain orientation on $\mathcal{M}_{\mathrm{cob}}$. Assuming that $\mathcal{M}_{\mathrm{cob}}$ is compact (which means $\mathcal{M}_{t=0}$ is also compact; compactness of $\mathcal{M}_{t=1}$ is automatic by choice of (J', H)), and therefore that its boundary consists of zero points (counting with sign), we conclude that the count of trajectories defining $\mathcal{CO}^0$ is equivalent to a signed count of points in $\mathcal{M}_{t=0}$.

If the composite local system is trivial then the result follows, modulo computing the signs. In the general case we split the moduli spaces $\mathcal{M}_{t=0}$, $\mathcal{M}_{t=1}$ and $\mathcal{M}_{\mathrm{cob}}$ according to the homotopy classes of the discs and apply the above argument to each homotopy class separately. We then just have to compute the contribution $\mathcal{P}(\gamma_{\mathrm{t}}) \circ \mathcal{P}(\gamma_{\mathrm{b}})^{-1}$ of each trajectory γ. But since each of these trajectories comprises a single flowline followed by a single disc, we can homotope away the parallel transport back and forth along the flowline (compare with Remark 3.2.5) and are left with the transport around the disc boundary, with reverse orientation, which is exactly what the map $\widetilde{m}$ describes. The fact that all signs are positive for the standard spin structure under Assumption 3.2.6 follows from Corollary B.5.2.

We are therefore left to prove compactness of $\mathcal{M}_{\mathrm{cob}}$, so suppose that we are given a sequence $(t_l) \subset [0, 1]$, and for each l an element u_l in the $t = t_l$ fibre of $\mathcal{M}_{\mathrm{cob}}$. Passing to a subsequence, we may assume that the t_l converge to some t_∞ and that the u_l Gromov converge to some stable map u. Write J_∞ and H_∞ for the values of J_t and H_t at $t = t_\infty$.

Note that this limit curve u must still meet Z, either on a disc component or a sphere component. In the former case, the disc component is (J_∞, H_∞)-holomorphic and since (J_∞, H_∞) lies in U the disc must have index $2k$. There can therefore be no other components, and we deduce that $u \in \mathcal{M}_{\mathrm{cob}}$. In the latter case, u meets Z on a $(J_\infty, 0)$-holomorphic sphere component u_1, and by our assumption on N^+_X there must be exactly one other component, which is an index 0 (J_∞, H_∞)-holomorphic disc u_2. Reparametrising if necessary we have that $u_1(0) \in Z$, $u_1(\infty) = u_2(0)$ and $u_2(1) = p$. We therefore have an index $2k$ configuration of type (ii) from the definition of $\mathcal{A}_{\mathrm{trans}}$ (note u_1 is simple by our hypothesis on N^+_X; even without this hypothesis we could always replace u_1 by a simple curve, without increasing the total index of the configuration). Such configurations occur in virtual dimension -2 (for fixed t), and since our homotopy is of dimension 1 and is regular we conclude that they cannot exist, completing the proof of compactness. □

Remark 3.2.25. The condition on N_X^+ is automatically satisfied if Z is a divisor. //

Definition 3.2.26. Let $\widetilde{m}(Z)$ denote the count appearing on the right-hand side in Proposition 3.2.24. This is an element of $\operatorname{End} \mathscr{F}_p$, perhaps more correctly viewed as a horizontal section of $\operatorname{End} \mathscr{F}$—compare [94, Section 2.2]. //

Remark 3.2.27. It is easy to see that the orientation sign of the isomorphism in Assumption 3.2.6 is the local degree of the evaluation map $\mathrm{ev}_1 \colon \mathcal{M}^Z_{\mu=2k} \to L$, where $\mathcal{M}^Z_{\mu=2k}$ is oriented by the short exact sequence

$$0 \to T_u\mathcal{M}^Z_{\mu=2k} \to T_u\mathcal{M}_{\mu=2k} \to T_{u(0)}X/T_{u(0)}Z \to 0.$$

Under Assumption 3.2.6, $\widetilde{m}(Z)$ is therefore the degree of this evaluation map, twisted by $\widetilde{m}$ in the obvious sense. //

Corollary 3.2.28. *Assume the hypotheses of the first part of Proposition 3.2.24. If R is a field, $\widetilde{m}(Z) = \lambda \cdot \mathrm{id}_{\mathscr{F}_p}$ is a scalar operator, and $HF^*(L^\flat, L^\flat; R)$ is non-zero then λ is an eigenvalue of the R-linear endomorphism $\alpha*$ of $QH^*(X; R)$ given by quantum multiplication by α.*

Proof. Take the result of 3.2.24 and apply the argument from the second part of Proposition 3.2.2. □

Example 3.2.29. Suppose that $L^\flat$ is sharply K-homogeneous and equipped with the standard spin structure, and that the compactification divisor Y has irreducible components $Y_1, \ldots, Y_r$. Let d_j denote the vanishing order of the holomorphic section σ of the anticanonical bundle (from Lemma 3.1.9) on Y_j. Each index 2 disc on L hits Y exactly once, on the smooth locus of a component with $d_j = 1$, and we see that

$$\mathcal{CO}^0(c_1(X)) = \sum_{j:d_j=1} \mathcal{CO}^0(\mathrm{PD}(Y_j)) + \sum_{j:d_j>1} d_j\, \mathcal{CO}^0(\mathrm{PD}(Y_j)) = \sum_{j:d_j=1} \widetilde{m}(Y_j) \cdot 1_L$$

is the count of J-holomorphic index 2 discs on L, twisted by $\widetilde{m}$, all contributing with positive sign. By a cobordism argument analogous to that used earlier (but simpler), the obstruction term $\mathfrak{m}_0$ counting index 2 discs through a generic point of L can also be computed using the standard J, and by Proposition B.4.1 again all discs contribute with positive sign. We conclude that $\mathcal{CO}^0(c_1(X)) = \mathfrak{m}_0 \cdot 1_L$. This recovers the Auroux–Kontsevich–Seidel result, with signs. Changing relative spin structure affects both counts in the same way—it can be viewed simply as changing the H_2^D local system—so the restriction to the standard spin structure is unnecessary. //

3.3 Poles

Throughout Section 3.3 we assume that (X, L) is sharply K-homogeneous and, as usual, let G denote the complexification of K. Recall that G acts holomorphically on X with dense open orbit W, whose complement $X \setminus W$ is a divisor we call Y, cut out by the vanishing of a holomorphic section σ of the anticanonical bundle of X. We also fix a choice of base point x in L, and assume that the stabiliser G_x of x in G is a finite subgroup of K (this is automatic if the homogeneity is linear—modulo Remark 3.1.24—or the action is Hamiltonian as in Remark 3.1.25). Note that since the homogeneity is sharp L is automatically orientable (and spin) so Maslov indices of discs are all even.

For certain stretches we will also assume that the homogeneity is linear. These regions will be clearly demarcated in the text, and we will not repeat this hypothesis in the statements of all results within them. Usually we will include it explicitly in propositions and corollaries, but not lemmas.

The purpose of this section is to study holomorphic discs on L (or other holomorphic curves in X) locally about the points where they hit Y, which we term *poles*. This is useful for several reasons. Firstly, we know that the index of a disc is determined by the vanishing of σ at these points, so there is an immediate link between certain global properties of a disc and the local behaviour at its poles. Secondly, the complement W of Y is a single G-orbit, and the orbit map $G \to W$ (with respect to our base point x) is a covering map. This means that away from Y a disc can locally be lifted to G, and any two discs differ locally by the action of a G-valued holomorphic function, so we may expect that in fact *most* of the interesting information is concentrated near the poles. On a more practical level, focusing on poles lets us study curves without worrying about boundary conditions, or whether they close up. Finally, meromorphic functions on compact Riemann surfaces are well understood—any such function is determined up to the addition of constants by the positions and principal parts of its poles—and we aim to develop an analogy with this and understand its limitations.

Note that if X is $\mathbb{CP}^1 = \mathbb{C} \cup \{\infty\}$, with $K = \mathbb{R}$ and $G = \mathbb{C}$ acting by addition, then points at which holomorphic curves meet the compactification divisor $\{\infty\}$ are simply poles in the usual sense. However, this is a slightly misleading picture since $\mathbb{R}$ is not compact. A better toy model is with $X = \mathbb{CP}^1$ still, but now $K = \mathrm{U}(1)$ and $G = \mathbb{C}^*$ acting by multiplication. Then a pole in our new sense corresponds to a pole or zero in the traditional sense.

3.3.1 First notions

We begin with the formal definitions:

Definition 3.3.1. Given a point p in a Riemann surface Σ, a *pole germ at p* is the germ (at p) of a holomorphic map u from a neighbourhood of p to X, such that $u^{-1}(Y)$ contains p as an isolated point. Pole germs u_1 and u_2 at p are *equivalent*, denoted by $\sim$, if there exists a holomorphic map g from a neighbourhood of p to G such that $u_1 = gu_2$ near p. The equivalence class of a pole germ is its *principal part*.

If u is a pole germ at p and there exist a positive integer k, a Lie algebra element ξ in $\mathfrak{k}$ such that $\{t \in \mathbb{R} : e^{2\pi t\xi}x = x\} = \mathbb{Z}$, and a local holomorphic coordinate z about p (with p corresponding to $z = 0$), such that u is equivalent to the pole germ

$$z \mapsto e^{-ik\xi \log z}x \tag{3.5}$$

then u is *quasi-axial* of *type* ξ and *order* k. //

Remark 3.3.2. (i) If we don't specify 'at p' then we are implicitly working at 0 in $\mathbb{C}$.

(ii) If $u\colon \Sigma \to X$ is an actual holomorphic curve in X, with no component contained in Y, it defines pole germs at each point of intersection with Y. In a slight abuse of terminology we shall use the expression 'poles of u' to refer both to the intersection points $u^{-1}(Y)$ in Σ and to the corresponding pole germs and their properties. If all of the poles of u are quasi-axial then we'll say that u itself is quasi-axial.

(iii) If u is a non-constant axial disc on L then its unique pole is quasi-axial. If the pole is of type ξ and order k then we say u itself has type ξ and order k.

(iv) The type of a quasi-axial pole depends on the choice of base point x.

(v) For the sake of brevity we will often suppress explicit mention of the fact that we are working with *germs* rather than actual maps, and instead just choose representatives of these germs. We will use informal phrases like 'near p' to signify that an assertion holds at the level of germs (or, equivalently, on sufficiently small neighbourhoods of p), or 'away from p' to mean everywhere except p, or on small punctured neighbourhoods of p. Hopefully these will be clear from the context.

(vi) Suppose we are given a pole germ u at p in Σ. Since the orbit map $G \to W$ is a covering, locally away from p we can lift u from W to G. There may be non-trivial monodromy in G_x if we try to lift along a loop encircling p, but we may eliminate it by passing to a cover of u—given in terms of a local coordinate z about p by $z \mapsto z^d$—since by assumption G_x is finite. Then the pole germ lifts to G on a whole punctured neighbourhood of p. We write ψ_d as a shorthand for such a local branched covering map, and assume that all ψ_d are defined in terms of the same local coordinate so that $\psi_{d_1} \circ \psi_{d_2} = \psi_{d_1 d_2}$ for all d_1 and d_2. //

Much of the work in this section centres on understanding when poles are quasi-axial, and what its consequences are. Along the way we recover some of our earlier results by independent methods (but under slightly stronger hypotheses).

We now prove some basic properties of pole germs.

Lemma 3.3.3. *If u is a quasi-axial pole germ of type ξ and order k at a point p in a Riemann surface Σ then u is equivalent to the pole germ at p defined by* (3.5) *for* any *choice of local coordinate z about p. For any non-negative integer d, $u \circ \psi_d$ is quasi-axial of type ξ and order dk.*

Proof. By definition u is equivalent to (3.5) for *some* choice of local coordinate z. If w is any other choice then we have

$$u \sim e^{-ik\xi \log z} x = e^{-ik\xi \log(z/w)} e^{-ik\xi \log w} x \sim e^{-ik\xi \log w} x.$$

For the second assertion, simply work in a local coordinate z in which ψ_d is $z \mapsto z^d$. □

Remark 3.3.4. There exist (necessarily non-quasi-axial) pole germs which are not equivalent to their reparametrisations. For example, take X to be the projectivisation of the space of 2×2

complex matrices, isomorphic to $\mathbb{CP}^3$, carrying the obvious action of SU(2) by multiplication on the left. Fix the identity matrix as our base point, and consider the pole germ u_1 at 0 in $\mathbb{C}$ defined by

$$u_1 : z \mapsto \left[\begin{pmatrix} 1 & \frac{1}{z} \\ 0 & 1 \end{pmatrix}\right].$$

Setting u_2 to be the pole germ $u_1 \circ (z \mapsto 2z)$, we see that each u_j lifts to the map g_j from a punctured neighbourhood of 0 to $\mathrm{SL}(2, \mathbb{C})$ given by

$$g_j : z \mapsto \begin{pmatrix} 1 & \frac{1}{jz} \\ 0 & 1 \end{pmatrix}.$$

Then one can easily compute that $g_1 g_2^{-1}$ does not extend holomorphically over 0, so u_1 and u_2 are inequivalent.

One might complain that u_1 does not extend to a disc with boundary on L, and wonder whether this is the source of the problem. However, u_1 is equivalent to the pole germ

$$z \mapsto \left[\begin{pmatrix} 1 & \frac{1}{z} \\ -z & 1 \end{pmatrix}\right],$$

which does extend to a disc on L, so this is not the case. //

Lemma 3.3.5. *Suppose u_1 and u_2 are pole germs at p in Σ, and that d is a positive integer.*

(i) If $u_1 \circ \psi_d \sim u_2 \circ \psi_d$ then $u_1 \sim u_2$.

(ii) If $u_1 \circ \psi_d$ is quasi-axial of type ξ and order k then d divides k and u_1 is quasi-axial of type ξ and order k/d.

(iii) The order of a quasi-axial pole is unique.

Proof. (i) Suppose $u_1 \circ \psi_d \sim u_2 \circ \psi_d$. Replacing d by a multiple if necessary, we may assume that the $u_j \circ \psi_d$ lift to holomorphic maps h_j to G away from p (i.e. on small punctured neighbourhoods of p). Fixing a local coordinate z about p, such that ψ_d is $z \mapsto z^d$, we have $u_2(z^d) = g(z)u_1(z^d)$ near p for some holomorphic map g from a neighbourhood of p to G. Modifying our lift h_2 by multiplying it on the right by an element of G_x, if necessary, we then have that $h_2(z) = g(z)h_1(z)$ near p.

Since h_1 and h_2 are lifts of d-fold covers of poles, if ζ denotes a primitive dth root of unity then there exist m_1 and m_2 in G_x such that for each j we have $h_j(\zeta z) = h_j(z)m_j$ near p. Writing g as $h_2 h_1^{-1}$, we deduce that

$$g(\zeta z)g(z)^{-1} = h_2(z)m_2 m_1^{-1} h_2(z)^{-1} \tag{3.6}$$

near p.

As G is the complexification of a compact group, it admits a faithful representation ρ (take a faithful representation of K, extend to G and use Proposition 3.1.6(iv)). Applying ρ to (3.6) and taking characteristic polynomials χ we see that

$$\chi(\rho(g(\zeta z)g(z)^{-1})) = \chi(\rho(m_2 m_1^{-1}))$$

near p. As $z \to 0$ the left-hand side tends to $(T-1)^N$, where N is the dimension of the representation, and hence the right-hand side (which is constant) also takes this value. On the other hand, $\rho(m_2 m_1^{-1})$ has finite order since $m_2 m_1^{-1}$ lies in the finite stabiliser G_x, so is diagonalisable. This forces $\rho(m_2 m_1^{-1})$ to be the identity, and as ρ is faithful this in turn forces m_1 and m_2 to be equal.

Plugging this back into (3.6), we see that $g(\zeta z) = g(z)$ near p, and hence g can be written as $\widetilde{g} \circ \psi_d$ for some holomorphic map $\widetilde{g}$ to G on a neighbourhood of p. We then have $u_2 = \widetilde{g} u_1$, and so $u_2 \sim u_1$ as claimed.

(ii) Let z be as in (i). Then there exists a holomorphic map g from a neighbourhood of p to G such that

$$g(z)u_1(z^d) = e^{-ik\xi \log z} x.$$

Replacing d and k by appropriate multiples we may assume that $u_1 \circ \psi_d$ and $e^{-ik\xi \log z} x$ lift to holomorphic maps h_1 and $h_2 = e^{-ik\xi \log z}$ to G away from p, chosen compatibly so that $h_2 = gh_1$.

For ζ a primitive dth root of unity, away from p we have $h_1(\zeta z) = h_1(z)m$ for some m in G_x, and $h_2(\zeta z) = h_2(z)e^{2\pi(k/d)\xi}$. Applying the argument used in (i) we see that $e^{2\pi(k/d)\xi}$ and m must be equal, so by definition of the scaling of ξ we have that k/d is an integer. If u_2 denotes a quasi-axial pole of type ξ and order k/d then we obtain

$$u_1 \circ \psi_d \sim u_2 \circ \psi_d,$$

and the result follows from (i).

(iii) Suppose u is quasi-axial of type ξ and order k, and again let z be as in (i). If $u \sim u_1 \circ \psi_d$ for some u_1 and d then d divides k by (ii). On the other hand, if u_1 is the pole germ defined by $z \mapsto e^{-i\xi \log z} x$ then by Lemma 3.3.3 we have $u \sim u_1 \circ \psi_k$. We conclude that k is the maximal d such that u is equivalent to a pole germ of the form $u_1 \circ \psi_d$. $\square$

Lemma 3.3.6. *Suppose u is a quasi-axial pole germ of type ξ and order k at p in Σ. Then u is also of type η and order l if and only if $k = l$ and η is conjugate to ξ by an element of G_x.*

Proof. The 'if' direction is clear, so suppose conversely that u also has type η and order l. The fact that $k = l$ follows immediately from Lemma 3.3.5(iii), so with respect to an arbitrary choice of local coordinate z we have

$$e^{-ik\xi \log z} x \sim e^{-ik\eta \log z} x.$$

Replacing k by a multiple if necessary, we may assume that $e^{-ik\xi \log z}$ and $e^{-ik\eta \log z}$ define holomorphic maps to G on a punctured neighbourhood of p—call these maps h_1 and h_2 respectively. Since they are lifts of equivalent poles, there exists a holomorphic map g from a neighbourhood of p to G and an element m of G_x with $h_2 = gh_1 m^{-1}$ away from p. Replacing ξ by $m\xi m^{-1}$ we may assume that m is the identity.

As in Lemma 3.3.5 take a faithful representation of $\rho\colon G \to \mathrm{GL}(V)$ of G, now equipped with a K-invariant inner product so that $\rho(K) \subset \mathrm{U}(V)$. We'll also denote the corresponding Lie algebra representation by ρ. Then $-ik\rho(\xi)$ and $-ik\rho(\eta)$ are diagonalisable, with integer eigenvalues, and distinct eigenspaces are orthogonal. If $E_j(\xi)$ and $E_j(\eta)$ denote the respective eigenspaces with eigenvalue j, then for each integer l we have

$$\bigoplus_{j \geq l} E_j(\xi) = \{v \in V : o(\rho(h_1(z))v) \geq l\} \tag{3.7}$$

and similarly

$$\bigoplus_{j \geq l} E_j(\eta) = \{v \in V : o(\rho(h_2(z))v) \geq l\}. \tag{3.8}$$

Here o denotes vanishing order at 0 as in Section 1.2.4. Since $h_2 = gh_1$, the right-hand sides of (3.7) and (3.8) coincide, and hence so do the left-hand sides. Using orthogonality of the eigenspaces, we deduce that $E_j(\xi) = E_j(\eta)$ for all j. Therefore $\rho(\xi) = \rho(\eta)$, and thus $\xi = \eta$, proving the claim. □

Proposition 3.3.7. *If u_1 and u_2 are holomorphic discs with boundary on L which each have a unique pole, at the point p in D, and the two poles are equivalent, then the discs are translates of each other under the action of K. In particular, if the poles are quasi-axial then the discs are axial.*

Proof. By reparametrising we may assume the unique pole is at 0 with respect to the standard coordinate z on D, and write D^* for $D \setminus \{0\}$. We can also choose d such that $u_1 \circ \psi_d$ and $u_2 \circ \psi_d$ lift to holomorphic maps h_1 and h_2 from $(D^*, \partial D)$ to (G, K). Since their poles are equivalent, there exists a holomorphic map g from a neighbourhood of 0 to G such that $u_2 = gu_1$ near 0. Modifying our lift h_2 by an element of G_x if necessary, we may assume that $h_2 = gh_1$ on a small punctured neighbourhood of 0. Since G_x is contained in K, the modified h_2 still takes ∂D to K.

Now consider the holomorphic map $h_2 h_1^{-1} \colon D^* \to G$. It maps ∂D to K, and by the previous paragraph it extends over 0 (by $g(0)$), so it can be completed to a holomorphic map $\widetilde{g} \colon \mathbb{CP}^1 \to G$ by Schwarz reflection, using the antiholomorphic involution $\hat{\tau}$ from Proposition 3.1.5. Composing $\widetilde{g}$ with a faithful representation $\rho \colon G \to \mathrm{GL}(V)$ of G, to give a holomorphic map from $\mathbb{CP}^1$ to the vector space $\mathrm{End}(V)$, we deduce that $\widetilde{g}$ is constant; call this constant g (this really is the same as the g from earlier). Since $g = h_2(1)h_1(1)^{-1}$ lies in K we see that $u_2 = gu_1$ is a K-translate of u_1.

Now suppose the pole of u_1 is quasi-axial, of type ξ and order k. Taking u_2 to be $z \mapsto e^{-ik\xi \log z}x$, we see from the first part that $u_1(z) = ge^{-ik\xi \log z}x$ for some g in K, and hence that

$$u_1(z) = e^{-ik(g\xi g^{-1}) \log z} u_1(1),$$

which is axial. □

Next we introduce another definition:

Definition 3.3.8. If u is a pole germ at p in Σ, the *index* of u, denoted $\mu_p(u)$, is twice the vanishing order of $\sigma \circ u$ at p. //

Remark 3.3.9. The index of a holomorphic disc on L is the sum of the indices of its poles. //

Lemma 3.3.10. *The index of a pole germ u at p in Σ depends only on its principal part.*

Proof. Let $\mathcal{L}$ denote the anticanonical bundle of X, and for g in G let L_g denote the automorphism of X corresponding to action by g. We get pushforward maps L_{g*}, either from space of sections of $\mathcal{L}$ to itself or from the fibre $\mathcal{L}_p$ over p to the fibre $\mathcal{L}_{gp}$ over gp. We also have the adjoint action Ad of G on $\mathfrak{g}$, and it is easy to check that for all g in G we have

$$L_{g*}\sigma = \det(\operatorname{Ad} g)\sigma,$$

where σ is the usual section of $\mathcal{L}$.

Introduce a metric on X, and consider the induced metric on $\mathcal{L}$. If $S(\mathcal{L})$ denotes the unit sphere bundle then the map $S(\mathcal{L}) \times G \to \mathbb{R}$ given by

$$(s \in S(\mathcal{L})_p, g) \mapsto \|L_{g*}s \in \mathcal{L}_{gp}\|$$

is continuous and strictly positive. When restricted to compact subsets in G it is therefore bounded between ε and $1/\varepsilon$ for some small $\varepsilon > 0$.

Let g be a holomorphic map from a neighbourhood of p to G. Shrinking the neighbourhood if necessary, we may assume that the image of g is contained in a compact set so that this boundedness property holds. If o denotes vanishing order at p, as in Section 1.2.4, we thus have

$$\mu_p(gu) = o\big(\|\sigma \circ (gu)\|^2\big) = o\big(\|(L_{g^{-1}}{}_*\sigma) \circ u\|^2\big) = o\big(\det(\operatorname{Ad} g^{-1})^2\|\sigma \circ u\|^2\big) = \mu_p(u),$$

where the second equality follows from the boundedness property, and the final equality holds because $\det(\operatorname{Ad} g^{-1})$ is also bounded near p. Strictly we only defined the vanishing order for functions which are holomorphic, but it clearly makes sense for the squared norms of such functions too. □

Definition 3.3.11. If u is a pole germ at p in Σ and Z is a G-invariant subset of X then we'll say u *evaluates to* Z if $u(p)$ lies in Z. This is preserved by equivalence of poles. //

Note that in general a G-orbit $\mathcal{O}$ in X need not be a submanifold. By the *dimension* of such an $\mathcal{O}$ we simply mean the rank of the infinitesimal G-action at an arbitrary point of $\mathcal{O}$. Orbits of complex codimension 1 are open subsets of the compactification divisor Y, and thus are in fact submanifolds. Pole germs evaluating to these orbits are especially well-behaved:

Lemma 3.3.12. *Suppose that $\mathcal{O}$ is a G-orbit in X of complex codimension 1, and that the vanishing order of σ on $\mathcal{O}$ is d. For any point p in a Riemann surface Σ:*

(i) There exists an index $2d$ pole germ $\hat{u}$ at p which evaluates to $\mathcal{O}$.

(ii) If u is an index $2kd$ pole germ at p which evaluates to $\mathcal{O}$, then there exists a local cover ψ_k such that u is equivalent to $\hat{u} \circ \psi_k$. In particular, if $\hat{u}$ is quasi-axial of type ξ (and necessarily of order 1) then u is quasi-axial of type ξ and order k.

Proof. (i) Fix a coordinate z on Σ about p, and coordinates $z_1, \ldots, z_n$ on X about a point q in $\mathcal{O}$, with $\mathcal{O}$ given locally by $z_n = 0$. Define a pole germ $\hat{u}$ at p by setting $z_j \circ \hat{u}(z) = 0$ for $j < n$ and $z_n \circ \hat{u}(z) = z$. Then $\hat{u}$ meets $\mathcal{O}$ transversely, so $\sigma \circ \hat{u}$ vanishes to order d at p and $\hat{u}$ has index $2d$.

(ii) Acting by an element of G on u (which doesn't affect its principal part), we may assume that $u(p) = \hat{u}(p)$. Consider the infinitesimal action of G at this point. Pick a complement $\mathfrak{g}'$ in $\mathfrak{g}$ to the one-dimensional kernel of this action, and define a holomorphic map $\Phi \colon \mathfrak{g}' \times U \to X$ by $(\xi, z) \mapsto e^{\xi}\hat{u}(z)$, where U is a small open neighbourhood of p.

By construction, Φ is a holomorphic local diffeomorphism at $(0, p)$, so it has a local holomorphic inverse. Near p we can therefore write

$$u(z) = \Phi(\xi_u(z), \varphi(z)) = e^{\xi_u(z)}\hat{u} \circ \varphi(z)$$

for some holomorphic maps ξ_u and φ. Replacing u by $e^{-\xi_u}u$, which still has the same principal part, we may assume that $\xi_u = 0$ and thus that $u(z) = \hat{u} \circ \varphi(z)$. Since u has index $2kd$, φ vanishes to order k at p so is of the form ψ_k, defined with respect to some local coordinate about p which in general will be different from z.

For the final part suppose $\hat{u}$ is quasi-axial of type ξ. Then it is of order 1, otherwise it would be a multiple cover of a pole of index less than $2d$ which evaluates to $\mathcal{O}$, and no such poles exist. The result then follows from Lemma 3.3.3. □

3.3.2 A closer look at pole indices

Suppose u is a pole germ; we'll assume it's at 0 in $\mathbb{C}$, although this is not important. Let R denote the ring of germs of holomorphic functions at 0, and let Γ denote the R-module comprising germs of holomorphic section of u^*TX. There is a map $\mathfrak{g} \to \Gamma$, which sends a Lie algebra element ξ to the section X_ξ generated by its infinitesimal action, and we denote the R-linear span of its image by M. If $\xi_1, \ldots, \xi_n$ is a basis for $\mathfrak{g}$ then $u^*X_{\xi_1}, \ldots, u^*X_{\xi_n}$ is an R-basis for M (they are independent since in u^*TX they are fibrewise linearly independent on a punctured neighbourhood of 0), which is therefore a free R-module of rank n.

Now consider the quotient $V = M/zM$, which is a free $R/(z) \cong \mathbb{C}$-module of rank n, canonically identified with $\mathfrak{g}$ by the map sending ξ to the image of u^*X_ξ. M carries a filtration by vanishing order at 0,

$$M = M_0 \supset M_1 \supset M_2 \supset \cdots,$$

where $M_j = \{s \in M : o(s) \geq j\}$, and the images of the M_j define a filtration

$$V = V_0 \supset V_1 \supset V_2 \supset \cdots$$

of V. The dimensions of the V_j determine the index of u as follows:

Lemma 3.3.13. *The index $\mu_0(u)$ is given by $2\sum_{j\geq 1} \dim V_j$.*

Proof. Pick a basis for V compatible with the filtration, and lift to sections $s_1, \ldots, s_n$ in M. Because (z) is the unique maximal ideal in R, Nakayama's lemma ensures that the s_j span M, and since R is a Euclidean domain (with Euclidean function given by vanishing order o)

any spanning set of size n must form a free basis. It is easy to see that the vanishing order of $s_1 \wedge \cdots \wedge s_n$ is equal to the vanishing order of the corresponding wedge product for any other R-basis for M, and so in particular—considering the basis $X_{\xi_1}, \ldots, X_{\xi_n}$ from above—it is equal to half the index of u. It therefore suffices to show that

$$o(s_1 \wedge \cdots \wedge s_n) = \sum_{j \geq 1} \dim V_j. \tag{3.9}$$

For each j let o_j be the vanishing order of s_j, so we can write s_j as $z^{o_j} t_j$ for some holomorphic section t_j of u^*TX which is non-vanishing at 0. Note that the right-hand side of (3.9) is precisely $o_1 + \cdots + o_n$, so we need to show that

$$o(t_1 \wedge \cdots \wedge t_n) = 0.$$

This is equivalent to proving that $t_1(0), \ldots, t_n(0)$ are linearly independent, so suppose for contradiction that this is not the case.

We thus have a non-trivial linear dependence $\sum a_j t_j(0) = 0$ in $T_{u(0)}X$, where the a_j are complex numbers, not all zero. This means that $o(\sum a_j t_j) \geq 1$. Let N denote the maximal value of o_j over those j for which $a_j \neq 0$, and reorder the s_j so that those with $o_j = N$ are precisely $s_1, \ldots, s_r$. We then have

$$\sum_{j \leq r} a_j s_j = z^N \sum a_j t_j - \sum_{j : o_j < N} a_j z^{N - o_j} s_j \in M_{N+1} + zM.$$

In particular, when we project to V the image of the left-hand side is in V_{N+1}, so there is a non-trivial linear dependence between the projections of $s_1, \ldots, s_r$ to V_N / V_{N+1}. But this is impossible, since by construction these projections form a basis for the latter space. We therefore have the desired contradiction, and conclude that the $t_j(0)$ are linearly independent and hence that the index of u is as claimed. □

Corollary 3.3.14. *The index of a pole germ u at a point p in a Riemann surface Σ is at least twice the complex corank of the infinitesimal G-action at $u(p)$.*

Proof. We may assume p is 0 in $\mathbb{C}$ and apply Lemma 3.3.13, noting that the corank of the infinitesimal action is precisely the dimension of V_1. □

Remark 3.3.15. This is a local version of Proposition 3.2.11 which says discs of index less than $2k$ cannot meet invariant subvarieties of complex codimension k. //

Suppose u_1 and u_2 are pole germs with $u_1(0) = u_2(0)$, and fix a trivialisation of TX on a neighbourhood of this point. In this way we can view (germs of) sections of u_j^*TX as (germs of) $\mathbb{C}^n$-valued holomorphic functions of z. If u_1 and u_2 have the same r-jet, then for each ξ the sections $u_1^* X_\xi$ and $u_2^* X_\xi$ agree up to and including terms of order z^r. This means that their corresponding R-modules M/M_{r+1} are canonically identified, and hence their spaces

$$V_j / V_{j+1} = M_j / (M_{j+1} + zM \cap M_j)$$

are also identified for $j \leq r$. We deduce:

Lemma 3.3.16. *For a pole germ u the numbers $\dim V_0, \dots \dim V_{r+1}$ depend only on the r-jet of u.*

3.3.3 Obliging poles and subvarieties

Throughout this subsection Z will denote an irreducible K-invariant analytic subvariety of X, of complex codimension k, with singular locus S. For simplicity all pole germs we consider will be at 0 in $\mathbb{C}$, although again it doesn't really matter.

An important role is played by those u for which we have equality in Corollary 3.3.14, so we introduce some terminology for this:

Definition 3.3.17. A pole germ u is *obliging* if its index is twice the complex corank of the infinitesimal G-action at $u(0)$. An invariant analytic subvariety Z, as above, is obliging if there exists a pole germ of index $2k$ which evaluates to Z. Such a pole germ will be called *Z-obliging.* //

Remark 3.3.18. A Z-obliging pole germ is necessarily obliging. Conversely, an obliging pole germ u is Z-obliging if and only if Z is the analytic closure of the G-orbit through $u(0)$, and is of the same dimension. //

Remark 3.3.19. Each component of the singular locus S of Z is a K-invariant subvariety of X of codimension greater than k, so any Z-obliging pole germ must evaluate to the smooth locus $Z \setminus S$. //

Remark 3.3.20. If Z is a divisor, i.e. k is 1, then it is obliging if and only if σ vanishes to order 1 along it. The 'only if' direction is clear, whilst for the converse note that any pole germ meeting the smooth locus of Z transversely has index given by twice the vanishing order of σ along Z. //

Remark 3.3.21. By Lemma 3.3.13 u is obliging if and only if the space V_2 from Section 3.3.2 is zero, and by Lemma 3.3.16 this depends only on the 1-jet of u (which is a point in TX). In particular, if $u'(0)$ lies in the image of the infinitesimal G-action at $u(0)$ then u has the same 1-jet as the germ of a curve which is contained entirely in the G-orbit of $u(0)$ and thus in the compactification divisor Y. Clearly $\dim V_2 = n$ for such a curve, and hence u is not obliging. This means that any obliging pole germ u must meet the G-orbit of $u(0)$ cleanly (meaning transversely to the subspace of $T_{u(0)}X$ spanned by the infinitesimal G-action). //

We shall see later, in Corollary 3.3.41, that if the homogeneity is linear then all obliging pole germs are quasi-axial, so by Proposition 3.3.7 any index $2k$ holomorphic disc on L which maps an interior marked point to Z is axial, giving an alternative proof of Proposition 3.2.11. Counting such discs through a generic point of L is the main ingredient in the computation of the length zero closed–open string map in Section 3.2.3. The key result of this subsection shows that essentially there is a unique obliging principal part evaluating to each obliging Z, which reduces the required count to a stabiliser calculation, which in turn is trivial if the homogeneity is free. The idea is to convert Z into a divisor by blowing up (inspired by [45, Corollary 3.11]), then further simplify to a codimension 1 G-orbit and apply Lemma 3.3.12.

Proposition 3.3.22. *All Z-obliging pole germs are equivalent up to reparametrisation. If one such germ is quasi-axial then all of them are equivalent, without needing reparametrisation.*

Proof. Consider the blowup X' of $X \setminus S$ along the smooth submanifold $Z \setminus S$. Note that the action of G on X restricts to $X \setminus S$ and lifts to X'. If u_1 and u_2 are Z-obliging pole germs in X then by Remark 3.3.19 they evaluate to $Z \setminus S$ and so we can consider their proper transforms $\hat{u}_j$, which are pole germs in X' evaluating to the exceptional divisor E. Moreover by Remark 3.3.21 each u_j evaluates *cleanly* to $Z \setminus S$, so the index of $\hat{u}_j$ is given by

$$\mu_0(\hat{u}_j) = \mu_0(u_j) - 2(k-1) = 2.$$

The $\hat{u}_j$ are E-obliging so evaluate to its smooth locus, which is connected by [77, second Proposition, page 21], and the infinitesimal G-action at each $\hat{u}_j(0)$ must have full rank inside this smooth locus. The argument used for the second part of Lemma 3.1.8 says that the set of points where this rank is full comprises a single G-orbit, so the $\hat{u}_j$ evaluate to the same G-orbit $\mathcal{O}$ of codimension 1. Lemma 3.3.12(ii) then implies that $\hat{u}_1$ and $\hat{u}_2$ are equivalent up to reparametrisation (and reparametrisation is unnecessary if one of them is quasi-axial), and we deduce that the same is therefore true of u_1 and u_2. □

Corollary 3.3.23. *The moduli space $\mathcal{M}^Z_{\mu=2k}$ is either empty or comprises a single K-orbit.*

Proof. Suppose u_1 and u_2 are discs in $\mathcal{M}^Z_{\mu=2k}$. By Proposition 3.2.11 they are axial and by Proposition 3.3.22 their poles are equivalent. Proposition 3.3.7 now completes the proof. □

Remark 3.3.24. It is not obvious a priori that $\mathcal{M}^Z_{\mu=2k}$ is even compact. //

Remark 3.3.25. If $\mathcal{M}^Z_{\mu=2k}$ is non-empty then Z must be obliging, but it is not obvious that the converse holds. The problem is that the existence of a pole germ of a given index evaluating to a given invariant subvariety does not immediately imply the existence of an entire disc with these properties. However, quasi-axial pole germs *do* always complete to discs (up to equivalence; in other words, every quasi-axial pole germ is equivalent to the pole of an axial disc), so once we have shown that obliging pole germs are quasi-axial then we will know that $\mathcal{M}^Z_{\mu=2k}$ is non-empty if and only if Z is obliging. //

Remark 3.3.26. One can show using Lemma 3.3.13 that if Z is obliging then the set of Z-obliging 1-jets is open and dense in $TX|_Z$, so in this sense a generic pole evaluating to an obliging subvariety is obliging. //

3.3.4 Reflection of discs

We now introduce a powerful tool for studying holomorphic discs with boundary on L. Recall that W is the dense open G-orbit in X.

Lemma 3.3.27. *W admits an antiholomorphic involution τ whose fixed locus is precisely L.*

Proof. Given a point q in W, we can write it in the form gx for some g in G, where x is our base point in L. The element g is uniquely determined up to multiplication on the right by an element of the stabiliser G_x, which lies in K by assumption. Let $\tau(q)$ be the point $\hat{\tau}(g)x$, where $\hat{\tau}$ is the antiholomorphic group involution on G from Proposition 3.1.5. This is independent of the choice of g since $\hat{\tau}$ fixes K and therefore G_x, so τ is well-defined. Given this, τ is clearly antiholomorphic and involutive, and fixes L.

Now suppose that $q = gx$ is fixed by τ, and write $g = e^{i\xi}k$ for ξ in $\mathfrak{k}$ and k in K. Then $\hat{\tau}(g) = gk'$ for some k' in $G_x \subset K$, and hence $e^{-i\xi}k = e^{i\xi}kk'$. Since the polar map is a diffeomorphism, we must have that $\xi = 0$ (and k' is equal to the identity) and hence g lies in K. This means q is in L, so the fixed locus is exactly L. □

In order to take this further, we will assume for the remainder of this subsection that (X, L) is sharply linearly K-homogeneous, so X is embedded in a projective space $\mathbb{P}V$ and the K-action factors through $\mathrm{GL}(V)$. By averaging, we assume that K preserves the inner product on V, so replacing K and G by their images in $\mathrm{GL}(V)$ we may assume that K is a subgroup of $\mathrm{U}(V)$ and, by Lemma 3.1.23, that G is an algebraic subgroup of $\mathrm{GL}(V)$. This replacement does not affect the main result of this subsection, Proposition 3.3.29, that holomorphic discs complete to spheres. Since the homogeneity is sharp we have $\dim X = \dim G$, and since G is connected in the classical topology (because K is) it is irreducible in the Zariski topology (distinct irreducible components of G cannot intersect as G is a group and thus is homogeneous).

First, we have a technical result limiting how badly behaved pole germs can be:

Lemma 3.3.28. *Suppose u is a pole germ at p in Σ without monodromy, so it lifts to a holomorphic map g from a punctured neighbourhood U of p to G. Then g extends meromorphically over p as a map to* $\mathrm{End}(V)$*, i.e. its components have, at worst, poles at p.*

Proof. Fix a local coordinate z about p. We shall show that each component of $g(z)$ satisfies a non-zero polynomial whose coefficients are holomorphic functions of z, which then implies that $g(z)$ is meromorphic over 0 in the claimed sense. In order to do this one has to make a slightly fiddly reduction to exploit the fact that the action of G 'sees all of the components of g'. As an example of the kind of bad behaviour that would ruin our argument, suppose G is a product $G_1 \times G_2$ with G_1 acting trivially (of course this is ruled out by our hypotheses): then one could say nothing about the G_1-components of g.

Let $\varphi\colon G \to W$ denote the orbit map determined by our base point x. Note that W is actually *Zariski* open in X (since by Lemma 3.1.23(ii) the union of the other orbits is Zariski closed) and that φ is a morphism of varieties which is a finite covering map in the classical topology. Let W' be the Zariski closure of $u(U)$ in W and let G' be the component of $\varphi^{-1}(W')$ containing the image of g; we may as well assume U is connected so that both W' and G' are irreducible (in the Zariski topology). Since they have the same dimension, their function fields, $k_{G'}$ and $k_{W'}$, have the same transcendence degree over $\mathbb{C}$. Therefore the pullback map $\varphi^*\colon k_{W'} \to k_{G'}$ exhibits $k_{G'}$ as an algebraic extension of $k_{W'}$.

Let $X_0, \dots, X_m$ be affine coordinates on V such that in the corresponding homogeneous coordinates the point $u(p)$ is $[1 : 0 : \cdots : 0]$. If $x_1, \dots, x_m$ denote the rational functions $X_1/X_0, \dots, X_m/X_0$ on W' then $k_{W'}$ is generated as an extension of $\mathbb{C}$ by the x_j (W' is locally closed in $\mathbb{P}V$ so its function field is equal to that of its closure, which is a closed subvariety of $\mathbb{P}V$ not contained in $\{X_0 = 0\}$). Let $\{Y_{jk}\}_{j,k=0}^{m}$ be the coordinates on $\mathrm{End}(V)$ corresponding to the X_j on V, and let the y_{jk} denote their images in $k_{G'}$. We want to show that each $y_{jk} \circ g(z)$ satisfies a non-zero polynomial with coefficients holomorphic in z.

Fix indices j and k. Since $k_{G'}/k_{W'}$ is algebraic, y_{jk} satisfies a non-zero polynomial with coefficients in $k_{W'}$. These coefficients are rational functions in the x_l (strictly the $\varphi^* x_l$), and by clearing denominators we may assume that they are in fact polynomials. We thus have

$$\sum_{r=0}^{s} p_r(x_1, \dots, x_m) y_{jk}^r = 0$$

for some polynomials $p_r(x_1, \dots, x_m)$ not all representing the zero function in $k_{W'}$. Composing with g, we get a polynomial equation for $y_{jk} \circ g(z)$ whose coefficients are $p_r \circ u(z)$. Since $u(z)$ is holomorphic over 0 in the local affine coordinates $X_1/X_0, \dots, X_m/X_0$, these coefficients are holomorphic. We are just left to check that the coefficients are not all identically zero, but this follows from the fact that the p_r are not all zero on W', which is the Zariski closure of $u(U)$. $\square$

Before going on to show that holomorphic discs complete to spheres, we need a simple observation about the involution $\hat{\tau}$. Let $\ddagger$ denote the adjoint-inverse map on $\mathrm{GL}(V)$, so that

$$M^{\ddagger} = (M^{\dagger})^{-1}$$

for each M in $\mathrm{GL}(V)$. This is easily seen to be an antiholomorphic group involution fixing $\mathrm{U}(V)$, and hence K. Any given g in G can be written in the form $e^{i\xi}k$ for some ξ in $\mathfrak{k}$ and k in K, and we then have

$$g^{\ddagger} = \left(e^{i\xi}\right)^{\ddagger} k^{\ddagger} = e^{-(i\xi)^{\dagger}} k = e^{-i\xi} k,$$

which itself lies in G. Therefore $\ddagger$ restricts to an antiholomorphic group involution of G fixing K, and is thus equal to $\hat{\tau}$. We are now ready to prove:

Proposition 3.3.29. *If (X, L) is sharply linearly K-homogeneous then every holomorphic disc $u\colon (D, \partial D) \to (X, L)$ doubles to a holomorphic sphere $\widetilde{u}$.*

Proof. Suppose u is a holomorphic disc on L, let $P = u^{-1}(Y)$ be the set of poles of u, and let $c\colon \mathbb{CP}^1 \to \mathbb{CP}^1$ denote the map $z \mapsto 1/\overline{z}$, corresponding to reflection in the equator ∂D. By Lemma 3.3.27 there exists an antiholomorphic involution τ of W fixing L, so by Schwarz reflection we can extend u to a holomorphic map $\widetilde{u}\colon \mathbb{CP}^1 \setminus c(P) \to X$ by

$$\widetilde{u}(z) = \begin{cases} u(z) & \text{if } z \in D \\ \tau \circ u \circ c(z) & \text{if } z \in c(D \setminus P). \end{cases}$$

It is now left to show that $\widetilde{u}$ extends continuously and hence holomorphically over $c(P)$.

So take a point in P. By reparametrising, we may assume it is 0. To show that $\widetilde{u}$ patches over ∞, it suffices to prove that the cover $\widetilde{u} \circ \psi_d$ patches over for some d, so pick d such that $u \circ \psi_d$ lifts on a punctured neighbourhood of 0 to a map g to G. By Lemma 3.3.28 g is meromorphic over 0. We then have for z small that

$$\widetilde{u}(1/z) = \tau \circ u(\overline{z}) = \hat{\tau}(g(\overline{z}))x = g(\overline{z})^{\ddagger} x,$$

and $g(\overline{z})^{\ddagger}$ is meromorphic over 0. Cancelling any common factors of $z^{\pm 1}$ from the homogeneous coordinates of $g(\overline{z})^{\ddagger} x$, we see that it does indeed tend to a limit in $\mathbb{P}V$ (and hence in X) as $z \to 0$, and so $\widetilde{u}$ extends continuously over $c(P)$, as desired. $\square$

Remark 3.3.30. Clearly a similar argument applies to higher genus curves with boundary on L. //

Remark 3.3.31. We can also apply this reflection process to pole germs, not just to whole curves. If the pole germ is based at p in a Riemann surface (Σ, j), where j is the complex structure, then the reflected pole is based at p in $(\Sigma, -j)$. If Σ admits an antiholomorphic involution c then alternatively one can think of the reflected pole as being based at $c(p)$ in (Σ, j). //

Proposition 3.3.29 allows us to reduce the study of discs on L to the study of rational curves in X, which are automatically algebraic. This is not especially useful if one cannot control the degrees of the resulting curves, but when u is quasi-axial we have the following result in this direction:

Lemma 3.3.32. *If u is a quasi-axial pole germ at p in Σ of type ξ and order k then the reflected pole germ is quasi-axial of type $-\xi$ and order k.*

Proof. Take a local coordinate z about p. There exists a holomorphic map g from a neighbourhood of 0 to G such that $u(z)$ is given by $g(z)e^{-ik\xi \log z}x$. The reflected pole germ is then given by

$$z \mapsto \hat{\tau}(g(\overline{z}))e^{(ik\xi \log \overline{z})^\dagger}x = \hat{\tau}(g(\overline{z}))e^{-ik\xi^\dagger \log z}x,$$

so is quasi-axial of type $\xi^\dagger = -\xi$ and order k (ξ is skew-adjoint since it lies in $\mathfrak{k}$, which is contained in $\mathfrak{u}(V)$). □

We can also extend the involution τ over large G-orbits:

Corollary 3.3.33. *If (X, L) is sharply linearly K-homogeneous and Z is a G-orbit of complex codimension 1 then τ extends antiholomorphically over Z as a map to X.*

Proof. Take an arbitrary point q in Z. As in the proof of Lemma 3.3.12 we can pick a transverse pole germ $\hat{u}$, with $\hat{u}(0) = q$, and define a holomorphic map $\Phi \colon \mathfrak{g}' \times U' \to X$ by $(\xi, z) \mapsto e^\xi \hat{u}(z)$, where $\mathfrak{g}'$ is a complement in $\mathfrak{g}$ to the kernel of the infinitesimal action at q, and U' is a small open neighbourhood of 0. This has a local holomorphic inverse Φ^{-1} defined on an open neighbourhood U of q in X, and in a slight abuse of notation we denote its components by ξ and z (i.e. if y is in U then $y = \Phi(\xi(y), z(y))$).

If $\hat{u}_{\text{ref}}$ denotes the pole germ at 0 obtained by reflecting $\hat{u}$, then for points y in $U \setminus Z$ we have $z(y) \neq 0$ and thus

$$\tau(y) = e^{-\xi(y)^\dagger}\hat{u}_{\text{ref}}\Big(\overline{z(y)}\Big). \tag{3.10}$$

The right-hand side of (3.10) is perfectly well-behaved over Z, so we use it to define a local antiholomorphic extension of τ. These local definitions all patch together since any local continuous extension of τ over Z is unique. □

3.3.5 Group derivatives

In this subsection we define a meromorphic Lie algebra-valued notion of the derivative of a holomorphic curve in X, which is closely related to the logarithmic (or Darboux) derivative of a smooth map to a Lie group, and thus to the pullback of the Maurer–Cartan form on G [85, page 311]. In [87] and subsequent papers, Hitchin constructed holomorphic curves in quasihomogeneous threefolds of $\mathrm{SL}(2,\mathbb{C})$ and used the Maurer–Cartan pullback to produce meromorphic connections on the Riemann sphere, in order to build solutions to isomonodromic deformation problems and the Painlevé equations. Our approach here is in the opposite direction—we use properties of our derivative to constrain holomorphic curves—although we hope that some of our ideas may be applicable to the study of related isomonodromic deformations, as studied by Boalch and others [19].

Definition 3.3.34. If $u\colon \Sigma \to X$ is a holomorphic curve with no component contained in Y, then away from its poles (which form a discrete subset of Σ) we can lift u locally to a holomorphic map g to G. We can then define a holomorphic $\mathfrak{g}$-valued 1-form locally by $(\mathrm{d}g)g^{-1}$, and this is independent of the choice of lift g. These local definitions therefore all patch together to define a holomorphic $\mathfrak{g}$-valued 1-form on $\Sigma \setminus u^{-1}(Y)$, which we call the *group derivative* of u, denoted $\mathcal{D}u$. //

Remark 3.3.35. (i) The Darboux derivative of a smooth map g into a Lie group is defined to be $g^{-1}\mathrm{d}g$. In this case we don't have to worry about deleting poles or choosing local lifts, and the order of the derivative and inverse terms is the reverse of that used in the group derivative.

(ii) We will often choose a local coordinate z on Σ (usually about a pole) and replace $\mathrm{d}A$ by A', where $'$ denotes $\partial/\partial z$, to obtain a function rather than a 1-form.

(iii) The derivative $\mathrm{d}u$ of u satisfies $\mathrm{d}u = \mathcal{D}u \cdot u$, and this uniquely determines $\mathcal{D}u$.

(iv) We can similarly define the group derivative of a pole germ, which is the *germ* of a $\mathfrak{g}$-valued 1-form. //

Although the group derivative is not defined at the poles of a curve, it is not too badly behaved there:

Lemma 3.3.36. *If u is a pole germ at p in Σ then $\mathcal{D}u$ extends meromorphically over p. If u is obliging then $\mathcal{D}u$ has at worst a simple pole.*

Proof. Fix a local coordinate z about p and work with $'$ in place of d, as in Remark 3.3.35(ii). Following the proof of Lemma 3.3.13 and its notation, pick a basis for V compatible with the filtration V_j and lift to sections $s_1, \dots, s_n$ of u^*TX with vanishing orders $o_1, \dots, o_n$. The sections $t_j = z^{-o_j}s_j$ form a local holomorphic frame for u^*TX near 0, so we can write the germ $u'(z)$ uniquely in the form $\sum f_j t_j$ for some holomorphic functions f_j.

By construction, each s_j is given by the infinitesimal action of some holomorphic $\mathfrak{g}$-valued function $\xi_j(z)$. We then have

$$\mathcal{D}u(z) \cdot u(z) = u'(z) = \left(\sum \frac{f_j(z)\xi_j(z)}{z^{o_j}} \right) \cdot u(z)$$

and thus

$$\mathcal{D}u(z) = \sum \frac{f_j(z)\xi_j(z)}{z^{o_j}}$$

extends meromorphically over p. In fact, $\mathcal{D}u$ has at worst a pole of order $\max\{j : V_j \neq 0\}$, so if u is obliging then it has at worst a simple pole. □

Remark 3.3.37. (i) The first part of this result follows from Lemma 3.3.28 if the homogeneity is linear.

(ii) If u also has the property that no component is constant, then we can projectivise the $\mathfrak{g}$-valued 1-form $\mathcal{D}u$ to get a holomorphic curve $[\mathcal{D}u] : \Sigma \setminus u^{-1}(Y) \to \mathbb{P}\mathfrak{g}$, and by Lemma 3.3.36 it extends over the poles. We call this the *projectivised group derivative*, but will not pursue its study further here. //

We now show that $\mathcal{D}u$ can detect poles of u, and is as nice as possible in the quasi-axial case:

Lemma 3.3.38. *Suppose u is a pole germ at p in Σ.*

(i) $\mathcal{D}u$ has a pole at p.

(ii) If u is quasi-axial of type ξ and order k then $\mathcal{D}u$ has a simple pole at p with residue conjugate to $-ik\xi$ in $\mathfrak{g}$.

Proof. (i) Suppose for contradiction that $\mathcal{D}u$ is in fact holomorphic over p. Then $-\mathcal{D}u$ is the connection 1-form of a germ of a holomorphic (and hence flat) connection on the trivial principal G-bundle $\Sigma \times G$ over Σ, and if we quotient this bundle on the right by the stabiliser G_x of our base point x in L then we can identify the fibres with the open orbit W in X. Under this identification, the graph of u defines a germ of a horizontal section.

By assumption the connection is holomorphic over p, and hence horizontal sections extend over this point. In particular we see that $u(p)$ lies in W, contradicting the fact that u has a pole at p. This proves the claim.

(ii) If u has a quasi-axial pole of type ξ and order k then, fixing a coordinate z about p, there exists a germ g at p of a holomorphic map to G such that u is locally of the form

$$u(z) = g(z)e^{-ik\xi \log z}x$$

on a punctured neighbourhood of p. It is then straightforward to compute that

$$\mathcal{D}u(z) = \mathrm{d}g(z)g(z)^{-1} - \frac{1}{z}ikg(z)\xi g(z)^{-1}\mathrm{d}z,$$

and hence that $\mathcal{D}u$ has a simple pole at p of residue $-ikg(0)\xi g(0)^{-1}$, which is manifestly conjugate to $-ik\xi$. □

The main goal of this subsection is to prove the converse to Lemma 3.3.38(ii), namely that if $\mathcal{D}u$ has a simple pole (and the homogeneity is linear) then u is quasi-axial. The first step is the next result, which says roughly that to prove quasi-axiality one needn't check that the generator ξ lies in $\mathfrak{k}$:

Lemma 3.3.39. *Suppose V is a complex inner product space, K is a compact subgroup of $\mathrm{U}(V)$ with complexification $G \subset \mathrm{GL}(V)$, U is an open neighbourhood of 0 in $\mathbb{C}$, and $g : U^* \to G$ is a* *holomorphic map, where U^* denotes $U \setminus \{0\}$. If there exist ξ in $\mathfrak{u}(V)$ and a holomorphic map $h\colon U \to \mathrm{GL}(V)$, such that $e^{2\pi\xi}$ is the identity I and*

$$g(z) = h(z)e^{-i\xi \log z}$$

for all z in U^*, then ξ lies in $\mathfrak{k}$.

Proof. First note that there exist positive real numbers a and b such that on U (shrinking if necessary) we have $\|h(z)\|^2 \leq 1/a$ and $\|h(z)^{-1}\|^2 \leq b$; here we are using the operator norm on $\mathrm{End}(V)$. By polar decomposition in G we can write g uniquely in the form $g(z) = k(z)e^{H(z)}$, for smooth functions k and H on U^* taking values in K and $i\mathfrak{k}$ respectively ('H' is for 'Hermitian'). Then for all v in V and all z in U^* we have

$$a\|e^{H(z)}v\|^2 \leq \|e^{-i\xi \log z}v\|^2 \leq b\|e^{H(z)}v\|^2. \tag{3.11}$$

We claim that

$$\frac{H(z)}{\log|z|} \to -i\xi$$

as $z \to 0$, and since the left-hand side lies in $\mathfrak{k}$ for all z, the right-hand side must also lie in $\mathfrak{k}$.

The operators $-i\xi$ and $H(z)$ are both self-adjoint, so are diagonalisable with real eigenvalues and orthogonal eigenspaces (in fact the eigenvalues of $-i\xi$ are integers, since $e^{2\pi\xi}$ is the identity). For $\lambda \in \mathbb{R}$ let E_λ and E'_λ denote their respective λ-eigenspaces, and π_λ and π'_λ the orthogonal projections onto them. We can then rewrite (3.11) as

$$a\sum_\nu e^{2\nu}\|\pi'_\nu v\|^2 \leq \sum_\mu |z|^{2\mu}\|\pi_\mu v\|^2 \leq b\sum_\nu e^{2\nu}\|\pi'_\nu v\|^2. \tag{3.12}$$

Suppose the eigenvalues of $-i\xi$ and $H(z)$ are $\mu_1 \leq \cdots \leq \mu_m$ and $\nu_1(z) \geq \cdots \geq \nu_m(z)$ respectively (note the opposite orderings), listing them *with multiplicity*. By counting dimensions, for all j and z there exist unit vectors

$$v_j(z) \in \Big(\bigoplus_{\mu \geq \mu_j} E_\mu\Big) \cap \Big(\bigoplus_{\nu \geq \nu_j(z)} E'_\nu\Big)$$

and

$$w_j(z) \in \Big(\bigoplus_{\mu \leq \mu_j} E_\mu\Big) \cap \Big(\bigoplus_{\nu \leq \nu_j(z)} E'_\nu\Big).$$

These need not be continuous in z.

Plug $v_j(z)$ into the inequality (3.12). The left-hand term is at least $ae^{2\nu_j(z)}$, whilst when $|z|$ is at most 1 the central term is at most $|z|^{2\mu_j}$, so we deduce that $|z|^{2\mu_j}e^{-2\nu_j(z)}$ is at least a. Similarly, plugging $w_j(z)$ into the central and right-hand terms we see that $|z|^{2\mu_j}e^{-2\nu_j(z)}$ is at most b. Therefore

$$a \leq |z|^{2\mu_j}e^{-2\nu_j(z)} \leq b \tag{3.13}$$

for all z in U^* with $|z| \leq 1$. In particular, we must have $\nu_j(z)/\log|z| \to \mu_j$ as $z \to 0$.

We have thus shown that the eigenvalues of $H(z)/\log|z|$ converge (with multiplicity) to those of $-i\xi$, so we are left to show that the corresponding eigenspaces converge. To this end, for each j let $F_j(z)$ be the space

$$\sum_{l:\mu_l=\mu_j} E'_{\nu_l}.$$

For z sufficiently small this is simply the direct sum of the eigenspaces of $H(z)/\log|z|$ with eigenvalues close to μ_j, and has the same dimension as E_{μ_j}. We are done if we can show that $F_j(z)$ converges to E_{μ_j} as $z \to 0$.

So fix a j and consider (3.12) restricted $F_j(z)$. By (3.13) the right-hand term is at most $b|z|^{2\mu_j}\|v\|^2/a$ for z in U^* with $|z| \leq 1$ hence for such z we have

$$\sum_{\mu} |z|^{2(\mu-\mu_j)}\|\pi_\mu v\|^2 \leq \frac{b}{a}\|v\|^2.$$

For each $\mu < \mu_j$ the quantity $|z|^{2(\mu-\mu_j)}$ blows up as $z \to 0$, so $\|\pi_\mu v\|$ must tend to zero. In other words, $F_j(z)$ and E_μ 'tend to orthogonal' as $z \to 0$.

Take $j = m$. For z in U^* small, the spaces

$$\bigoplus_{\mu<\mu_m} E_\mu \text{ and } F_m(z)$$

are of complementary dimensions, and by the above argument they tend to orthogonal. This forces $F_m(z)$ to converge to E_{μ_m}. Now apply this argument to

$$\bigoplus_{\mu<\mu_{m-1}} E_\mu \text{ and } F_{m-1}(z) + F_m(z),$$

to see that $F_{m-1}(z)$ converges to $E_{\mu_{m-1}}$. Continuing in this way we conclude that each $F_j(z)$ converges to E_{μ_j}, which is exactly what we want. □

Using this lemma we can prove the main result of this subsection:

Proposition 3.3.40. *Suppose (X, L) is sharply linearly K-homogeneous and u is a pole germ at a point p in a Riemann surface Σ. If $\mathcal{D}u$ has a simple pole at p, then u is quasi-axial.*

Proof. Suppose X is embedded in $\mathbb{P}V$ as usual, and fix a local coordinate z about p. Note that we may replace K and G by their images in $\mathrm{GL}(V)$, and we may also replace u by its composition with an appropriate multiple cover ψ_d about p so that it lifts to a meromorphic map g to G; here we are using Lemma 3.3.28 to ensure that g has at worst poles at p. Then g satisfies

$$\mathrm{d}g - \mathcal{D}u \cdot g = 0,$$

so is horizontal with respect to the germ of the meromorphic connection on $\Sigma \times G$ defined by $-\mathcal{D}u$ as in Lemma 3.3.38(i). Equip the associated vector bundle $\Sigma \times V$ with this connection. Note that the germs of meromorphic horizontal sections of this vector bundle about p are exactly expressions of the form gv for v in V.

Now define a filtration

$$V \supset \cdots \supset V_{-1} \supset V_0 \supset V_1 \supset \cdots$$

of V by vanishing order at p:

$$V_j = \{v \in V : o(gv) \geq j\}.$$

Let W_j be the orthogonal complement to V_{j+1} in V_j, and define ξ in $\mathfrak{u}(V)$ to act on each W_j as the scalar ij. We claim that $g(z)e^{i\xi \log z}$ extends holomorphically over 0, from which the result then follows by Lemma 3.3.39.

Well, for each w in W_j we have $g(z)e^{i\xi \log z}w = z^{-j}g(z)w$, and this extends holomorphically over 0 by definition of V_j (which contains W_j). Therefore $g(z)e^{i\xi \log z}$ extends over 0 as a holomorphic map to $\mathrm{End}(V)$ and we are left to show that it actually maps to $\mathrm{GL}(V)$, i.e. that the limit endomorphism l over 0 is invertible. For each j let L_j denote the linear map $W_j \to V$ which sends a vector w to the leading term $\lim_{z\to 0} z^{-j}g(z)w$ of the section gw. Note that L_j has kernel $W_j \cap V_{j+1} = 0$, so it is injective, and its image is precisely the image of the limit endomorphism l restricted to W_j. Thus it suffices to show that the images of the L_j span V.

Let R in $\mathfrak{g}$ denote the residue of $\mathcal{D}u$ at p, and let U be a small simply connected open neighbourhood of p on which u is defined and in which p itself is the only pole. If q is a second point in U then parallel transport by our connection around a loop at q encircling p (once, anticlockwise) defines an automorphism of the fibre over q corresponding to multiplication on the left by some monodromy element M_q in G. A standard result, proved by Deligne in [40, Théorème 1.17] by passing to polar coordinates, states that M_q tends to $e^{2\pi i R}$ as $q \to p$ (note that R is the residue of *minus* the connection 1-form, which explains the apparent sign discrepancy with [40]). But since we already have a horizontal section of the original principal bundle which has no monodromy—namely g—we deduce that M_q is the identity, I, for all q, and hence that $e^{2\pi i R} = I$. We can thus define a representation of $\mathrm{U}(1)$ on V by $e^{i\theta} \mapsto e^{i\theta R}$, and obtain a decomposition of V into eigenspaces of R with integer eigenvalues: let E_j denote the j-eigenspace.

Fix a j and a vector w in W_j, and consider the parallel transport equation $\mathrm{d}s = \mathcal{D}u \cdot s$ satisfied by the section $s = gw$. By looking at the Laurent expansions of the two sides, we see that the leading coefficient $L_j w$ of s must lie in E_j. Therefore each L_j maps W_j to E_j. We have already seen that each L_j is injective, so by counting dimensions ($\sum \dim W_j$ and $\sum \dim E_j$ are both equal to $\dim V$) we conclude that each L_j is in fact *onto* E_j. Thus the images of the L_j do indeed span V, completing the proof. □

This characterisation of quasi-axiality in terms of the group derivative has many non-trivial consequences, including:

Corollary 3.3.41. *If (X, L) is sharply linearly K-homogeneous then all obliging pole germs in X are quasi-axial. In particular, all pole germs of index 2 are quasi-axial.*

Proof. Combine Lemma 3.3.36 with Proposition 3.3.40. □

Remark 3.3.42. This is like a local version of Proposition 3.2.11, and gives an alternative proof of it when the homogeneity is linear. //

Corollary 3.3.43. *If (X, L) is sharply linearly K-homogeneous and $Z \subset X$ is an irreducible K-invariant analytic subvariety of complex codimension k then $\mathcal{M}^Z_{\mu=2k}$ is a single K-orbit of axial discs if Z is obliging and is empty otherwise.*

Proof. Suppose Z is obliging—the result is trivial otherwise—so there exists a Z-obliging pole germ. By Corollary 3.3.41 this germ is quasi-axial so can be realised as the pole of an axial disc. Now apply Corollary 3.3.23. $\square$

Corollary 3.3.44. *If (X, L) is sharply linearly K-homogeneous then the obliging components of the compactification divisor Y (which are precisely those on which σ vanishes to order 1) are linearly independent in $H_{2n-2}(X \setminus L; \mathbb{Z})$.*

Proof. Let the obliging components be $Y_1, \dots, Y_m$. By Corollary 3.3.43, for each j there exists an index 2 disc u_j meeting Y_j. We then have classes $[u_1], \dots, [u_m]$ in $H_2(X, L; \mathbb{Z})$ and $[Y_1], \dots, [Y_m]$ in

$$H_{2n-2}(X \setminus L; \mathbb{Z}) \cong H^2_c(X \setminus L; \mathbb{Z}) \cong H^2(X, L; \mathbb{Z}),$$

where H^*_c denotes cohomology with compact support, which satisfy $[u_j] \cdot [Y_k] = \delta_{jk}$. $\square$

Remark 3.3.45. Let Y have components $Y_1, \dots, Y_r$, and let the vanishing order of σ along Y_j be d_j. Then

$$\sum_j d_j Y_j$$

is the divisor of zeros of σ, so represents the anticanonical class $c_1(X)$. If Y itself is anticanonical, meaning that $\sum_j Y_j$ also represents $c_1(X)$, then we conclude that $d_j = 1$ for all j, and hence that every component of Y is obliging.

In this case any pole germ of index at most 4 is quasi-axial. To see this note that any pole germ of index 4 is either obliging, in which case it is quasi-axial, or it evaluates to a codimension 1 G-orbit. In the latter case, by Lemma 3.3.12(ii) it is a double cover of an index 2 (and thus quasi-axial) pole germ, and so is itself quasi-axial. //

3.3.6 Some simple applications

We conclude this section with some simple applications to freely T^n- and SO(3)-homogeneous Lagrangians. First we have the following extension of Proposition 3.3.40:

Proposition 3.3.46. *Suppose (X, L) is sharply linearly K-homogeneous and u is a pole germ at p in Σ. If the principal part of $\mathcal{D}u$ at p is diagonalisable (this requires a choice of local coordinate about p, but is independent of this choice) then $\mathcal{D}u$ has a simple pole and hence u is quasi-axial.*

Proof. Fix a local coordinate z about p. Replacing u by a cover $u \circ \psi_d$ does not affect quasi-axiality, and it replaces $\mathcal{D}u$ by $dz^{d-1}\mathcal{D}u \circ \psi_d$ so doesn't affect diagonalisability of the principal part. Choosing d appropriately, we may therefore assume that u lifts to a germ of a meromorphic map g to G.

Pick a basis of V in which the principal part of $\mathcal{D}u$ is actually diagonal, say

$$\mathcal{D}u(z) = \left(\frac{D_k}{z^k} + \frac{D_{k-1}}{z^{k-1}} + \cdots + \frac{D_1}{z} + \text{holomorphic part}\right) \mathrm{d}z,$$

for some positive integer k and some diagonal matrices D_j with $D_k \neq 0$. Without loss of generality we may assume that the first entry of D_k is non-zero, and let d_j denote the first entry of each D_j. In the same basis our map g can be expressed as a meromorphic $\mathrm{GL}(n, \mathbb{C})$-valued function, which we denote by A. Since $A(z)$ is non-singular for all z, there exists an entry in the first row which is not identically zero. If f is this entry, which is a meromorphic function of z, then from the equation $\mathrm{d}g = \mathcal{D}u \cdot g$ we obtain

$$f' = \left(\frac{d_k}{z^k} + \frac{d_{k-1}}{z^{k-1}} + \dots\right) f,$$

with $d_k \neq 0$. Considering the Laurent expansion of each side we see that k must be 1. In other words, $\mathcal{D}u$ has a simple pole at p and so u is quasi-axial. $\square$

From this we obtain:

Corollary 3.3.47. *Suppose (X, L) is freely linearly T^n-homogeneous. Then every pole germ is quasi-axial and every irreducible T^n-invariant subvariety $Z \subset X$ is obliging.*

Proof. A linear representation of a torus is diagonalisable so the first part follows immediately from Proposition 3.3.46. Now take an arbitrary irreducible invariant subvariety Z, of complex codimension k, let S denote its singular locus, and let $2l$ be the minimal index of a pole germ evaluating to $Z \setminus S$. We need to show that $l = k$.

Consider the moduli space $\mathcal{M}^{Z\setminus S}_{\mu=2l}$ of index $2l$ holomorphic discs on L which map 0 to $Z \setminus S$. By Lemma 3.2.9 this is cut out transversely so is a smooth manifold of dimension $n+2(l-k)$. It carries a T^n-equivariant evaluation map ev_1 to L, and we define M to be the preimage $\mathrm{ev}_1^{-1}(x)$ of our base point x. This is a smooth manifold of dimension $2(l - k)$.

Any disc in M has a pole at 0 of index at least $2l$, so can have no other poles. Its unique pole is quasi-axial by the first part, so the disc is axial by Proposition 3.3.7, of the form

$$z \mapsto e^{-i\xi \log z} x \tag{3.14}$$

for some ξ in $\mathfrak{t}^n$ with $e^{2\pi\xi}$ equal to the identity. The set of such ξ is countable, so M must also be countable. On the other hand, M is non-empty since we know there exists a pole germ of index $2l$ evaluating to $Z \setminus S$, and it is quasi-axial, of type ξ say (it is necessarily of order 1 as it has minimal index); then (3.14) defines a point of M.

We conclude that M is a non-empty smooth manifold of dimension $2(l-k)$, but is countable. This is only possible if $l = k$. $\square$

Remark 3.3.48. A freely linearly T^n-homogeneous pair (X, L) is essentially a compact toric manifold with a choice of free torus orbit. The holomorphic discs bounded by such orbits were computed by Cho–Oh [35] using a different approach. //

Corollary 3.3.49. *Suppose $L^\flat \subset X$ is a freely linearly T^n-homogeneous monotone Lagrangian brane equipped with the standard spin structure and trivial local system. Under Assumption 3.2.6, for any irreducible T^n-invariant subvariety Z of complex codimension at most N^+_X (as in Definition 3.2.19) we have*

$$\mathcal{CO}^0(\mathrm{PD}(Z)) = 1_L.$$

Proof. Combine Proposition 3.2.24 with Corollary 3.3.43 and Corollary 3.3.47. $\square$

There are obvious modifications of this result in the presence of local systems, but they are less clean to state.

Remark 3.3.50. When Z is a divisor this gives an alternative proof that $\mathcal{CO}^0$ maps toric divisor classes to 1_L, which is well-known. For higher codimension subvarieties the result can be extracted from the work of Cho–Oh and is known to experts, but the author is not aware of its explicit appearance in the literature. //

Remark 3.3.51. Take X to be the toric monotone blowup of $\mathbb{CP}^2$ at three points, with hexagonal moment polytope, and L to be the monotone fibre. The superpotential (recapped in Section 6.2.4) is

$$W = x + xy + y + \frac{1}{x} + \frac{1}{xy} + \frac{1}{y},$$

and the six terms correspond, in order, to the weights attached by the local system to the basic index 2 classes in $H_2(X, L; \mathbb{Z})$ which are dual to the divisors over the edges of the hexagon (taken in cyclic order). We have $N_X^+ = 1$, so our $\mathcal{CO}^0$ result does not apply to the point class, and we claim that if it did then one would obtain a contradiction.

Indeed, by the previous results we know that each vertex of the moment polytope is hit by a unique (up to translation) index 4 disc, and if the $\mathcal{CO}^0$ result applied then $\mathcal{CO}^0(\mathrm{PD}(\text{point}))$ would simply be $\lambda \cdot 1_L$, where λ is the weight attached by the local system to any of these discs. The homology class of the disc entering the vertex between two edges is the sum of the classes of the discs which are dual to these edges, so the six index 4 discs are weighted by

$$x^2y,\ xy^2,\ \frac{y}{x},\ \frac{1}{x^2y},\ \frac{1}{xy^2},\ \text{and}\ \frac{x}{y}$$

(products of consecutive terms in W). These six numbers must all therefore agree, for any local system for which L is wide, i.e. for any critical point of W. It is easy to check that $(x, y) = (-1, -1)$ is a critical point for which these numbers do not coincide, proving that the $\mathcal{CO}^0$ result must fail. The bad bubbled configurations comprise an index 2 disc attached to the dual divisor, which is a rational curve of Chern number 1.

For $\mathbb{CP}^2$ and $\mathbb{CP}^1 \times \mathbb{CP}^1$, which have minimal Chern numbers 3 and 2 respectively, one can easily verify that the products of consecutive terms in the superpotential (when written in the natural cyclic order) are all equal at critical points. Analogous considerations in higher dimensions may give constraints on critical points of superpotentials which are easier to solve than the equation $\mathrm{d}W = 0$ itself. //

For freely $\mathrm{SO}(3)$-homogeneous Lagrangians meanwhile, we can say the following:

Proposition 3.3.52. *Suppose $L \subset X$ is a freely $\mathrm{SO}(3)$-homogeneous monotone Lagrangian, and let m denote the number of components of the compactification divisor Y on which the section σ of the anticanonical bundle vanishes to order* 1*. If $m \neq 1$ then L is non-displaceable.*

Proof. Equip L with an H_2^D local system and consider its self-Floer cohomology over the field $R = \mathbb{Z}[x]/(2, x^2 + x + 1)$ of four elements (there is no need to choose a relative spin structure as we're working in characteristic 2). Since the classical cohomology of L over R is generated in degree 1, just as for tori we see that it is wide if and only if the sum of the boundary classes of the index 2 discs vanishes.

Now any index 2 disc u on L is axial, of the form $z \mapsto e^{-ik\xi \log z} u(1)$ where

$$\{t \in \mathbb{R} : e^{2\pi t\xi} x = x\} = \mathbb{Z},$$

and is not a multiple cover, so $k = 1$. This means that its boundary sweeps a circle subgroup of $\mathrm{SO}(3)$ exactly once, so lies in the non-trivial class in $H_1(L; \mathbb{Z})$. We therefore see that the relevant boundary sum vanishes if and only if the sum of the weights assigned to the index 2 discs by the H_2^D local system vanishes. By Corollary 3.3.44 these weights can be chosen arbitrarily (i.e. independently of each other), as long as they are non-zero, and we know that the number of index 2 discs is exactly m. Therefore L can be made wide over R by an appropriate choice of H_2^D local system—and is thus non-displaceable—as long as there exist m non-zero weights $t_1, \dots, t_m$ in R whose sum is zero. This is the case for all m not equal to 1: if m is even then take $t_j = 1$ for all j, whilst if m is odd and at least 3 then take $t_1 = x^2$, $t_2 = x$ and $t_j = 1$ for all other j. $\square$

Question 3.3.53. Do displaceable freely $\mathrm{SO}(3)$-homogeneous monotone Lagrangians exist? //

Chapter 4

The Platonic Lagrangians

In this chapter we study a family of SU(2)-homogeneous Lagrangians and apply some of the methods of the previous chapter. We make Assumption 3.2.6 throughout, which enables us to compute $\mathcal{CO}^0$ with signs. Recall that this holds for both the Biran–Cornea and Zapolsky orientation schemes, as proved in Section B.1.

4.1 The spaces and their properties

4.1.1 The Chiang Lagrangian

Given a finite-dimensional complex inner product space W, the symplectic form ω induced by the metric and complex structure has primitive 1-form

$$\lambda = \operatorname{Im}(z^\dagger \mathrm{d}z)/2$$

(meaning $\omega = \mathrm{d}\lambda$), where z is a vector of coordinates with respect to an orthonormal basis and $\dagger$ denotes conjugate transpose. The unitary group $\mathrm{U}(W)$ clearly preserves this 1-form, and hence its action on W is Hamiltonian, with moment map $\widetilde{\mu} \colon W \to \mathfrak{u}(W)^*$ given by

$$\langle \widetilde{\mu}(z), \xi \rangle = X_\xi \lrcorner \lambda_z = \frac{1}{2} \operatorname{Im} z^\dagger \xi z = -\frac{i}{2} z^\dagger \xi z$$

for all z in W and all ξ in $\mathfrak{u}(W)$. Here $\langle \cdot, \cdot \rangle$ denotes the pairing between $\mathfrak{u}(W)^*$ and $\mathfrak{u}(W)$, X_ξ is the vector field generated by the infinitesimal action of ξ, and $\lrcorner$ denotes contraction.

Now consider the fundamental representation V of SU(2). This is (tautologically) unitary with respect to the standard inner product g, so all of its tensor powers $V^{\otimes d}$ are also unitary with respect to the corresponding tensor powers $g^{\otimes d}$. Inside $V^{\otimes d}$ we have the subrepresentation comprising totally symmetric tensors, which is isomorphic to the dth symmetric power S^dV of V, and we deduce that S^dV is unitary with respect to the restriction of $g^{\otimes d}$. Fix a basis of S^dV which is orthonormal with respect to this inner product, let $\varphi \colon \mathfrak{su}(2) \to \mathrm{Mat}_{(d+1)\times(d+1)}(\mathbb{C})$ describe the infinitesimal action in this basis, and let z denote a corresponding coordinate vector.

Taking $W = S^dV$ above, we see that the SU(2)-action on S^dV is Hamiltonian, with moment map $\widetilde{\mu}\colon S^dV \to \mathfrak{su}(2)^*$ defined by

$$\langle \widetilde{\mu}(z), \xi \rangle = -\frac{i}{2} z^\dagger \varphi(\xi) z$$

for all z in S^dV and all ξ in $\mathfrak{su}(2)$. This action commutes with the diagonal U(1)-action on S^dV, and the moment map is U(1)-invariant, so it descends to a Hamiltonian action on the projective space $\mathbb{P}S^dV$ with moment map μ given by

$$\langle \mu([z]), \xi \rangle = -\frac{i}{2} \frac{z^\dagger \varphi(\xi) z}{z^\dagger z} \tag{4.1}$$

for all z in S^dV, representing $[z]$ in $\mathbb{P}S^dV$, and all ξ in $\mathfrak{su}(2)$.

It is well-known (see, for example, [31, Proposition 1.5]) that an orbit of a Hamiltonian action of a compact Lie group is isotropic if it is contained in the moment map preimage of a fixed point of the coadjoint representation of the group. In particular, orbits contained in the zero set of the moment map are isotropic. In [31] Chiang considered the case of the above SU(2)-action on S^dV with $d = 3$. In her example the set $\mu^{-1}(0)$ is a single three-dimensional orbit inside $\mathbb{CP}^3$, and hence is Lagrangian: this is the so-called Chiang Lagrangian.

4.1.2 Coordinates on projective space

Let x and y be the standard basis vectors for the fundamental representation V of SU(2), which we now think of as being extended to a representation of $\mathrm{SL}(2, \mathbb{C})$, with respect to which a group element $\left(\begin{smallmatrix} t & u \\ v & w \end{smallmatrix}\right) \in \mathrm{SL}(2, \mathbb{C})$ acts as the matrix itself. We then have an induced basis for S^dV given by $\{x^iy^j : 0 \leq i, j \text{ and } i + j = d\}$. We'll refer to this as *the standard basis* for S^dV and the corresponding coordinates (and their projective counterparts) as *standard coordinates* on S^dV (respectively $\mathbb{P}S^dV$). This is the identification we will always use between $\mathbb{P}V$ and $\mathbb{CP}^1$. We will also use the identifications

$$\mathbb{CP}^1 \cong \mathbb{C} \cup \{\infty\} \cong \{x \in \mathbb{R}^3 : \|x\| = 1\}$$

given by viewing a point λ in $\mathbb{C} \cup \{\infty\}$ as both the point $[\lambda x + y] = [\lambda : 1]$ in $\mathbb{CP}^1$ and as the point on the unit sphere given by stereographic projection through the north pole from the complex (equatorial) plane. For example $[i : 1]$ in $\mathbb{CP}^1$ corresponds to i in $\mathbb{C}$ and $(0, 1, 0)$ in $\mathbb{R}^3$.

The vectors x and y are orthonormal with respect to the standard inner product g, and under our embedding of S^dV in $V^{\otimes d}$ as totally symmetric tensors (normalised, so, for example, xy embeds as $(x \otimes y + y \otimes x)/2$) we see that with respect to $g^{\otimes d}$ the x^iy^j are pairwise orthogonal and satisfy

$$\|x^iy^j\| = \sqrt{\frac{i!\,j!}{d!}}.$$

We thus have a *unitary basis*

$$\left\{ \sqrt{\frac{d!}{i!\,j!}}\, x^iy^j : 0 \leq i, j \text{ and } i + j = d \right\}$$

and corresponding *unitary coordinates*.

A point in $\mathbb{P}S^dV$ is a non-zero homogeneous polynomial of degree d in x and y, modulo $\mathbb{C}^*$-scalings. Since $\mathbb{C}$ is algebraically closed, we can express such a polynomial as a product of d linear combinations of x and y, which are uniquely determined up to scaling and reordering. Moreover, the action of $\mathrm{SL}(2,\mathbb{C})$ on $\mathbb{P}S^dV$ induced by the representation S^dV corresponds precisely to expressing elements in this factorised form and acting via the fundamental representation on each factor. In other words, we have an $\mathrm{SL}(2,\mathbb{C})$-equivariant identification between $\mathbb{P}S^dV$ and $\mathrm{Sym}^d\,\mathbb{P}V \cong \mathrm{Sym}^d\,\mathbb{CP}^1$.

In this way, an unordered d-tuple of points on $\mathbb{CP}^1$ can be viewed as a point of $\mathbb{P}S^dV$ and thus expressed in terms of either standard or unitary coordinates. As an example, consider the equilateral triangle on the real axis in $\mathbb{CP}^1$, with one vertex at ∞. Its two other vertices are at $\pm 1/\sqrt{3}$, so it is represented by the point $[x(x+\sqrt{3}y)(x-\sqrt{3}y)] = [x^3 - 3xy^2]$ in $\mathbb{P}S^3V$. It is therefore given by $[1 : 0 : -3 : 0]$ in standard coordinates and $[1 : 0 : -\sqrt{3} : 0]$ in unitary coordinates. Note that the expression (4.1) for the moment map is valid only in unitary coordinates.

4.1.3 From the triangle to the Platonic solids

With these notions fixed, there is another, more geometric, way to describe Chiang's construction. If we fix a value of $d \geq 3$, and a configuration C of d distinct points in $\mathbb{CP}^1$, then the $\mathrm{SL}(2,\mathbb{C})$-orbit of C in $\mathrm{Sym}^d\,\mathbb{CP}^1 \cong \mathbb{P}S^dV$ is a three-dimensional complex submanifold, of which the $\mathrm{SU}(2)$-orbit is a three-dimensional totally real submanifold. In [6] Aluffi and Faber identified those C for which the $\mathrm{SL}(2,\mathbb{C})$-orbit has smooth closure X_C in $\mathbb{P}S^dV$. There are four cases, namely the orbits of the configurations C given by (using the notation of Evans–Lekili): $\triangle$, the vertices of an equilateral triangle on a great circle in $\mathbb{CP}^1$; T, O and I, respectively the vertices of a regular tetrahedron, octahedron and icosahedron in $\mathbb{CP}^1$. These are quasihomogeneous threefolds of $\mathrm{SL}(2,\mathbb{C})$, in the sense that they carry an $\mathrm{SL}(2,\mathbb{C})$-action with dense Zariski open orbit. The stabiliser of C in $\mathrm{SL}(2,\mathbb{C})$ is a finite subgroup of $\mathrm{SU}(2)$ which we denote by Γ_C. The group $\Gamma_\triangle$ is the binary dihedral group of order 12, whilst Γ_T, Γ_O and Γ_I are the binary tetrahedral, octahedral and icosahedral groups respectively, of orders 24, 48 and 120.

In each case the restriction of the $\mathrm{SU}(2)$-action to X_C, with the symplectic form induced from the ambient projective space, is Hamiltonian with moment map of the form (4.1). The representative configurations $\triangle$, T, O and I all lie in the zero sets of the respective moment maps, and hence their $\mathrm{SU}(2)$-orbits are Lagrangian; we denote these 'Platonic' Lagrangians by L_C. The Chiang Lagrangian itself can then be described as $L_\triangle$ in $X_\triangle = \mathbb{P}S^3V \cong \mathbb{CP}^3$. By construction, each L_C is sharply linearly $\mathrm{SU}(2)$-homogeneous.

4.1.4 Basic properties of the spaces X_C

In this subsection we collect together some of the properties of the quasihomogeneous threefolds X_C. These results are not original and most are contained in [45, Section 4]. We follow the notation of Evans–Lekili.

For each C let W_C denote the Zariski open $\mathrm{SL}(2,\mathbb{C})$-orbit in X_C, and Y_C its complement, the compactification divisor. Y_C consists of those d-point configurations in X_C where at least $d-1$

of the points coincide. Inside Y_C we have the subvariety N_C consisting of those configurations where all d points coincide.

The cohomology ring of X_C is

$$H^*(X_C;\mathbb{Z}) = \mathbb{Z}[H,E]/(H^2 = k_C E, E^2 = 0),$$

where k_C is 1, 2, 5, 22 for C equal to $\triangle$, T, O, I respectively, and H is the class of a hyperplane section. The first Chern class of X_C is $c_1(X_C) = l_C H$, where l_C is 4, 3, 2, 1 for the four choices of C, so X_C is Fano with minimal Chern number l_C. Since L_C has finite fundamental group it is automatically monotone.

Remark 4.1.1. The numbers 1, 2, 5, 22 come about as follows. The value of k_C is the intersection product of three transverse hyperplane sections of X_C. We can take these hyperplane sections to be of the form $X_C \cap \Pi_z$ for z equal to 0, 1 and ∞, where Π_z consists of those d-point configurations containing the point $z \in \mathbb{CP}^1$, and then each $p \in X_C \cap \Pi_0 \cap \Pi_1 \cap \Pi_\infty$ can be described by choosing three ordered vertices of C to send to 0, 1 and ∞. This can be done in $|\Gamma_C|/2$ different ways for each p, corresponding to rotating C before choosing the points (we divide by 2 as we are interested in the image of Γ_C in SO(3)). Any such triple gives rise to some p, so we conclude that the triple intersection consists of $2d(d-1)(d-2)/|\Gamma_C|$ points, which works out to be 1, 2, 5, 22 in the four cases respectively. This argument appears in [6, Section 0]. //

In quantum cohomology the product is deformed to give a $\mathbb{Z}/2l_C$-graded ring

$$QH^*(X_C;\mathbb{Z}) = \mathbb{Z}[H,E]/(H^2 = k_C E + R_C, E^2 = Q_C),$$

where R_C and Q_C are as given in Table 4.1; see [11, Section 2]. We collapse the grading to $\mathbb{Z}/2$, and the ring is then concentrated in degree 0.

C	$\triangle$	T	O	I
R_C	0	0	3	$2H+24$
Q_C	1	H	$E+1$	$2E+H+4$

Table 4.1: Quantum corrections to the cup product.

The group $H_1(L_C;\mathbb{Z})$ is the abelianisation of $\pi_1(L_C, C) \cong \Gamma_C$, and is isomorphic to $\mathbb{Z}/l_C$. It follows from Proposition 4.2.5 that each L_C has minimal Maslov index 2, and since the minimal Chern number of X_C is l_C we conclude that $H_2(X_C, L_C;\mathbb{Z})$—which is an extension of $H_1(L_C;\mathbb{Z})$ by $H_2(X_C;\mathbb{Z}) \cong \mathbb{Z}$—is isomorphic to $\mathbb{Z}$ via $\mu/2$, as shown in the diagram below.

$$\begin{tikzcd}
H_2(X_C;\mathbb{Z}) \cong \mathbb{Z} \arrow[r] \arrow[d, "\langle c_1(X_C), \cdot\rangle"] & H_2(X_C, L_C;\mathbb{Z}) \arrow[r] \arrow[d, "\mu/2"] & H_1(L_C;\mathbb{Z}) \cong \mathbb{Z}/l_C \\
\mathbb{Z} \arrow[r] & \mathbb{Z} &
\end{tikzcd}$$

Remark 4.1.2. This means that each L_C admits exactly one non-standard relative spin structure, and compared with the standard spin structure the orientations of moduli spaces of discs or pearly trajectories differ by a factor of $(-1)^{\mu/2}$. Therefore discs and trajectories of index 2 count with opposite signs with respect to the different relative spin structures, whilst those of index 4 count with the same signs. //

4.1.5 Seeing the antiholomorphic involution

We know from Section 3.3.4 that the orbit W_C carries an antiholomorphic involution τ fixing L_C. Moreover this involution extends over $W_C \cup Y_C = X_C \setminus N_C$ and can be used to reflect poles and complete holomorphic discs on L_C to spheres. It arises from the antiholomorphic group involution $\hat{\tau}$ on $\mathrm{SL}(2,\mathbb{C})$ which fixes $\mathrm{SU}(2)$, given by the conjugate-transpose-inverse map, which we denote by $A \mapsto A^{\ddagger}$.

The involution τ has a useful geometric interpretation. First note that if we define J_0 to be $\left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ then for any A in $\mathrm{SL}(2,\mathbb{C})$ we have

$$A^{\ddagger} = J_0 \overline{A} J_0^{-1}.$$

For z in $\mathbb{CP}^1$, the map $z \mapsto J_0^{\pm 1} \cdot \overline{z}$ is precisely the antipodal map $\alpha \colon z \mapsto -1/\overline{z}$, so if p in W_C is described by $A \cdot C$ for some A in $\mathrm{SL}(2,\mathbb{C})$ then $A^{\ddagger} \cdot C$ is obtained by taking C, applying the antipodal map α (to each factor of $\mathrm{Sym}^d \mathbb{CP}^1 \cong \mathbb{P}S^dV$), acting by A, and then applying α again.

The configurations O and I are invariant under α, so in these cases τ acts on W_C simply as the antipodal map itself. Therefore τ extends to all of X_C, and we obtain a *global* antiholomorphic involution. The Floer theory of the fixed loci of such involutions has been studied previously, as mentioned in Section 1.1.4, and we shall see some of these ideas shortly.

For the triangle and tetrahedron, which are not preserved by α, τ is rather more subtle. It collapses down $Y_C \setminus N_C$ to N_C—explicitly, the point in $Y_C \setminus N_C$ corresponding to a d-point configuration on $\mathbb{CP}^1$ with a $(d-1)$-fold point at p and single point at q is sent to the configuration with a d-fold point at $\alpha(p)$—so it cannot possibly extend to a global involution since it is not injective. Evans–Lekili remark that $L_{\triangle}$ can't be the fixed-point set of *any* antiholomorphic involution, since by Lemma 4.2.8 the count of index 2 discs is odd (this count was also computed by Evans–Lekili [45, Lemma 6.2]).

For C equal to O or I, the involution on X_C is the restriction of the antipodal involution on $\mathbb{P}S^dV$ and it is easy to see that it is antisymplectic: the point (i.e. homogeneous polynomial of degree d in x and y, modulo scaling)

$$[(a_1x + b_1y) \dots (a_dx + b_dy)]$$

maps to

$$[(\overline{b}_1x - \overline{a}_1y) \dots (\overline{b}_dx - \overline{a}_dy)]$$

so in standard coordinates we have

$$[z_0 : z_1 : z_2 : \cdots : z_d] \mapsto [\overline{z}_d : -\overline{z}_{d-1} : \overline{z}_{d-2} : \cdots : (-1)^d \overline{z}_0],$$

which flips the sign of the Fubini–Study form. For C equal to $\triangle$ or T, however, the involution is *not* antisymplectic. In fact we shall see later that for a holomorphic disc u on L_C, the reflection of u by τ often has different Maslov index from u itself. By monotonicity of L_C, this means the reflected disc has different area.

4.1.6 Exploiting the antiholomorphic involution

In this subsection we discuss some of the consequences of the existence of τ. In particular, we shall see that because τ can be defined globally on X_O and X_I, wideness of L_O and L_I in characteristic 2 is automatic, without computing any holomorphic discs. Since we will not use many of the specific features of these examples, we temporarily move to a slightly more general setting.

Let X be a complex manifold, $Y \subset X$ an analytic subvariety, $L \subset X \setminus Y$ a totally real submanifold which is closed as a subset of X, and τ an antiholomorphic involution of $X \setminus Y$ which fixes L pointwise. Suppose moreover that τ enables us to reflect holomorphic discs with boundary on L, in the sense that for any holomorphic disc $u\colon (D, \partial D) \to (X, L)$ there exists a holomorphic disc v on L such that $v(z) = \tau \circ u(\overline{z})$ for all z in D with $\overline{z} \notin u^{-1}(Y)$. Using this we can double any disc u on L to a sphere $\widetilde{u}$.

We now introduce a new definition:

Definition 4.1.3. A holomorphic disc $u\colon (D, \partial D) \to (X, L)$ is *strongly simple* if its double $\widetilde{u}$ is not multiply-covered. //

Holomorphic discs can always be replaced by strongly simple discs, in the following sense:

Lemma 4.1.4. *Given a non-constant holomorphic disc u with boundary on L, there exists a strongly simple disc v on L such that:*

(i) $u(\partial D) \subset v(\partial D)$.

(ii) $\widetilde{u}(\mathbb{CP}^1) = \widetilde{v}(\mathbb{CP}^1)$.

(iii) If u is not itself strongly simple and every non-constant holomorphic disc on L has Maslov index at least 2 then $\mu(v) \leq \mu(u) - 2$.

Proof. If u is already strongly simple then we can just take $v = u$, so suppose this is not the case. Then $\widetilde{u}$ is a non-simple sphere, and hence is given by $w \circ \psi$ for some branched cover $\psi\colon \mathbb{CP}^1 \to \mathbb{CP}^1$ of degree $d > 1$ and some simple holomorphic sphere $w\colon \mathbb{CP}^1 \to X$. Pick three points a_1, a_2 and a_3 in ∂D whose images under ψ are distinct injective points of w. Reparametrising w if necessary, and correspondingly modifying ψ, we may assume that $\psi(a_i) \in \partial D$ for each i.

As in Proposition 3.3.29, let $c\colon \mathbb{CP}^1 \to \mathbb{CP}^1$ denote the map $z \mapsto 1/\overline{z}$, with fixed-point set ∂D, and let $\overline{w}$ denote the reflection of w, i.e. the holomorphic sphere given by $\tau \circ w \circ c$ whenever this is defined. Note that $\overline{w}$ is simple (if not then it would be a multiple cover and hence $w = \tau \circ \overline{w} \circ c$ would be a multiple cover) and that

$$\overline{w}(\mathbb{CP}^1) = w(\mathbb{CP}^1) = \widetilde{u}(\mathbb{CP}^1).$$

In particular, w and $\overline{w}$ are simple holomorphic spheres with the same image, and therefore differ by reparametrisation. To see this, let U and $\overline{U}$ be the cofinite subsets of their common image comprising the images of injective points of w and $\overline{w}$ respectively. Then $w^{-1} \circ \overline{w}$ defines a biholomorphism between the cofinite sets $\overline{w}^{-1}(U \cap \overline{U})$ and $w^{-1}(U \cap \overline{U})$, and considering the effect of this biholomorphism on the ends of these sets (it must pair them up) we deduce that it extends to an automorphism φ of $\mathbb{CP}^1$ satisfying $\overline{w} = w \circ \varphi$.

Now note that we have

$$w \circ \psi = \widetilde{u} = \tau \circ \widetilde{u} \circ c = \overline{w} \circ c \circ \psi \circ c$$

on the cofinite subset $\widetilde{u}^{-1}(Y)$ of $\mathbb{CP}^1$, and hence

$$w \circ \psi = w \circ \varphi \circ c \circ \psi \circ c \tag{4.2}$$

on all of $\mathbb{CP}^1$. Applying this at our points a_i we deduce that φ fixes the three points $\psi(a_i)$ and thus is the identity. Then (4.2) tells us that ψ coincides with $c \circ \psi \circ c$ at injective points of w and therefore everywhere. In other words, we have shown that $\overline{w} = w$ and that ψ commutes with c (so $\psi(\partial D) \subset \partial D$).

Using this, we see that $w(\partial D)$ contains $u(\partial D) = w(\psi(\partial D))$ and lies in the fixed locus of τ. This fixed locus contains L as an isolated component (locally about a fixed point of an antiholomorphic involution one can choose holomorphic coordinates in which the involution is given by complex conjugation), so $w^{-1}(L)$ is open in ∂D. Since L is closed in X, $w^{-1}(L)$ is also closed in ∂D, and hence $w(\partial D) \subset L$. This means that $v_1 := w|_D$ and $v_2 := w \circ (z \mapsto 1/z)|_D$ are holomorphic discs on L whose boundaries contain $u(\partial D)$. Their doubles are w and $\overline{w}$ respectively, so they are strongly simple and satisfy $\widetilde{u}(\mathbb{CP}^1) = \widetilde{v}_i(\mathbb{CP}^1)$. If we can show that $\mu(v_i) \leq \mu(u) - 2$ for some i then we can take v to be this v_i and we're done.

Well, in $H_2(\mathbb{CP}^1, \partial D)$ we have $\psi_*([D]) = d_1[D] + d_2[c(D)]$ for some non-negative integers d_1 and d_2, which sum to d since ψ commutes with c. Then in $H_2(X, L; \mathbb{Z})$ we have $[u] = d_1[v_1] + d_2[v_2]$, and hence

$$\mu(u) = d_1\mu(v_1) + d_2\mu(v_2).$$

Since each $\mu(v_i)$ is at least 2 (by our assumption on Maslov indices of discs on L), and the sum of the d_i is $d > 1$, we must have $\mu(v_i) \leq \mu(u) - 2$ for some i, proving the lemma. $\square$

In order to apply this to Floer theory, suppose now that in fact X is a compact Kähler manifold, $L^\flat \subset X$ a monotone Lagrangian brane, and that every holomorphic disc u bounded by L has all partial indices non-negative. We are still assuming the existence of Y (disjoint from L) and τ as above. In particular, these conditions are satisfied if L is sharply linearly K-homogeneous.

If f is a Morse function on L, and g is a metric such that (f, g) is Morse–Smale, Proposition C.3.1 ensures that, possibly after replacing f and g by their pullbacks under a diffeomorphism of L, we may use the data (f, g, J) to compute the self-Floer cohomology of L using the pearl complex. Using Lemma 4.1.4 we can restrict our attention to pearly trajectories in which all discs are strongly simple:

Lemma 4.1.5. *In every pearly trajectory contributing to the differential on the pearl complex given by Proposition C.3.1, all holomorphic discs are strongly simple.*

Proof. Suppose for contradiction that there is a trajectory in which some disc u is not strongly simple. Let its two marked points be at ± 1—note that $u(-1) \neq u(1)$, otherwise we could delete the disc u from the trajectory and obtain a trajectory in negative virtual dimension, which is impossible (in the proof of Proposition C.3.1 it is ensured that the relevant moduli space is transversely cut out). By Lemma 4.1.4 there exists a holomorphic disc v, of index strictly less than u, with $u(\partial D) \subset v(\partial D)$. In particular, we may reparametrise v such that $v(\pm 1) = u(\pm 1)$, and then replace u by v to obtain again a trajectory in negative virtual dimension, giving the desired contradiction. □

This is particularly useful when τ extends to a global involution, so we assume from now on that Y is empty and that our coefficient ring R has characteristic 2. We also assume that the local system on $L^\flat$ is trivial. Since the relative spin structure is irrelevant, we can just write L for $L^\flat$. In this situation we have:

Proposition 4.1.6. *L is wide over R. In other words, after collapsing the grading of $H^*(L; R)$ to $\mathbb{Z}/2$, we have an isomorphism of $\mathbb{Z}/2$-graded R-modules*

$$HF^*(L, L; R) \cong H^*(L; R).$$

Proof. We argue analogously to Haug [83], and show that all positive index contributions to the pearl complex differentials (which we also refer to as 'quantum corrections') occur in pairs and hence cancel over R. This makes the self-Floer cohomology of L, as computed by the pearl complex, agree with the Morse cohomology, which is in turn isomorphic to the singular cohomology.

The way we pair up the positive index contributions is by constructing a fixed-point-free involution τ_* on the space of such trajectories. For each sequence $(u_1, \dots, u_l)$ of non-constant holomorphic discs comprising a pearly trajectory (with $l \geq 1$), we define a new trajectory by

$$\tau_*(u_1, \dots, u_l) = (\overline{u}_1, \dots, \overline{u}_l),$$

where $\overline{u}_i$ is the disc given by $\overline{u}_i(z) = \tau \circ u_i(\overline{z})$ for all z. If we can show that this 'reflect the discs' map has no fixed points then we're done.

Well, if it *did* have a fixed point $(u_1, \dots, u_l)$ then the disc u_1 would be equal to its reflection, up to reparametrisation. In particular, its double $\widetilde{u}_1$ would hit every point in its image at least twice. This forces $\widetilde{u}_1$ to be a multiple cover (see [103, Section 2.3]), contradicting the fact that u_1 is strongly simple. Therefore τ_* has no fixed points, and thus we have the required cancellation. □

In Appendix C we also establish (Proposition C.4.2) that any triple $(f_j, g_j)_{j=1}^3$ of Morse–Smale pairs on L can be perturbed in order to be used to define the product on self-Floer cohomology, by counting Y-shaped pearly trajectories with one Morse–Smale pair used on each leg, i.e. each branch of the Y. And by an argument analogous to Lemma 4.1.5 the discs appearing in the legs, rather than at the centre of the Y, are all strongly simple. Using this we get:

Proposition 4.1.7. *The product on $HF^*(L, L; R)$ is commutative. The only positive index trajectories whose contributions do not cancel have a single disc, at the centre of the Y, and no others.*

Proof. By reflecting the leg discs, we see that trajectories with such discs cancel out. Reflecting the disc at the centre of the Y reverses the order of the three boundary marked points, and we get a bijection between trajectories contributing to the coefficient of z in $x * y$ and those contributing to the coefficient of z in $y * x$. □

Remark 4.1.8. Fukaya–Oh–Ohta–Ono proved a very similar result [69, Corollary 1.6], that the Floer cohomology ring is graded-commutative with rational Novikov coefficients, under hypotheses that ensure the Maslov index is trivial modulo 4, so that one can control the signs of reflected discs. //

Remark 4.1.9. In general the product on $HF^*(L, L; R)$ is different from that on $H^*(L; R)$. A simple example is provided by the equator in $\mathbb{CP}^1$, whose self-Floer cohomology ring over $\mathbb{Z}/(2)$ is isomorphic to $\mathbb{Z}[x]/(2, x^2 - 1)$, whereas $H^*(S^1; \mathbb{Z}/(2)) \cong \mathbb{Z}[x]/(2, x^2)$ (and $|x| = 1$ in both cases); these are not isomorphic as $\mathbb{Z}/2$-graded rings. //

We shall not use Proposition 4.1.7 in what follows, but from Proposition 4.1.6 we obtain:

Proposition 4.1.10. *For a field R of characteristic 2 we have isomorphisms of R-vector spaces*

$$HF^0(L_O, L_O; R) \cong HF^1(L_O, L_O; R) \cong R^2$$

and

$$HF^0(L_I, L_I; R) \cong HF^1(L_I, L_I; R) \cong R.$$

Proof. Applying Proposition 4.1.6 to these Lagrangians, we reduce the problem to computing their singular cohomology. We saw in Section 4.1.4 that $H_1(L_O; \mathbb{Z}) \cong \mathbb{Z}/2$ and $H_1(L_I; \mathbb{Z}) \cong 0$, so by Poincaré duality and the universal coefficient theorem we deduce that L_O has integral cohomology groups $\mathbb{Z}$, 0, $\mathbb{Z}/2$ and $\mathbb{Z}$ in degrees 0 to 3 respectively, whilst L_I is an integral homology 3-sphere (in fact it is the well-known Poincaré 3-sphere). The result now follows by passing to R-coefficients using the universal coefficient theorem again, and collapsing the gradings to $\mathbb{Z}/2$. □

Remark 4.1.11. Even outside characteristic 2, τ is useful because it induces an involution on moduli spaces of discs. The effect of this involution on orientations was computed in [69, Theorem 1.3]: on the space of unparametrised index $2i$ discs with j boundary marked points, the involution changes the orientation by a factor of $(-1)^{i+j}$. For example, on the space of unmarked index 2 discs, it is orientation-reversing.

In the context of the pearl complex we are interested in discs with 2 marked points. In this case, the involution reverses orientations in index 2 and preserves orientations in index 4. Hence, since we have shown that there are no contributing discs fixed by the involution, all trajectories in the pearl complex which contain an index 2 disc cancel out, whilst the count of index 4 discs through two points is always even. //

Combining Remark 4.1.2 and Remark 4.1.11, we deduce:

Corollary 4.1.12. *Equipping L_O or L_I with the trivial local system over an arbitrary ring R, its self-Floer cohomology is independent (additively, at least) of the choice of relative spin structure.*

Remark 4.1.13. If one tries to apply the argument of Proposition 4.1.6 to $L_\triangle$ and L_T, the problem is that the reflection of a disc generally has different index from the original disc, so the map τ_* does not act on individual moduli spaces of trajectories. Instead it mixes up moduli spaces of different virtual dimensions, and the argument falls apart. //

4.2 Holomorphic discs

4.2.1 Special pole types

For the purpose of computing the self-Floer cohomology of L_C we need to understand discs of index 2 and 4. In this subsection we classify the poles of these indices. From the results of Section 3.3.5 we know that index 2 poles exist if and only if Y_C is obliging, in which case they are quasi-axial and there is a unique such pole type. Any index 4 pole either evaluates to $Y_C \setminus N_C$, in which case it is equivalent to a double cover of an index 2 pole (by Lemma 3.3.12), or it is an obliging pole evaluating to N_C, in which case it is again quasi-axial and of a unique type.

Let ξ_v, ξ_e and ξ_f in $\mathfrak{su}(2)$ be generators of rotations about a vertex of C, the midpoint of an edge, and the centre of a face respectively, scaled so that $\{t \in \mathbb{R} : e^{2\pi t} \cdot C = C\} = \mathbb{Z}$, and directed so that the vertex, midpoint and centre are at the 'top' of the axis of rotation. Here the top of the axis is taken with right-handed convention, so, for example, the rotation $(\theta, z) \mapsto e^{i\theta} z$ has ∞ at the top of its axis, whilst $e^{-i\theta} z$ has 0 at the top (this is right-handed in the sense that when the fingers of the right hand are curled around the axis in the direction of rotation, the outstretched thumb points towards the top). One choice of such rotations for $C = \triangle = [x^3 + y^3]$ is shown in Fig. 4.1; we think of the triangle as having two faces—one for each side.

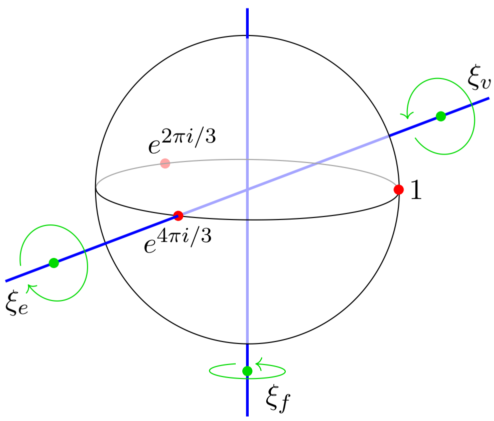

Figure 4.1: Examples of choices for ξ_v, ξ_e and ξ_f for $C = \triangle$.

Let r_C be 2, 3, 4, 5 for C equal to $\triangle$, T, O, I, denoting the number of faces meeting at a vertex. Then $e^{2\pi\xi_v}$ represents a rotation through angle $2\pi/r_C$—the smallest angle through which one can rotate C about a vertex to return it to its original position. Similarly $e^{2\pi\xi_e}$ and $e^{2\pi\xi_f}$ represent rotations through angle π and $2\pi/3$ respectively.

Definition 4.2.1. For $\bullet$ equal to v, e or f let $u_\bullet$ denote the quasi-axial pole germ

$$z \mapsto e^{-i\xi_\bullet \log z} C.$$

Taking C as our base point in L_C, this is of type $\xi_\bullet$ and order 1. Recall that the type is defined up to conjugation by Γ_C, which corresponds to changing the particular vertex, edge or face about which $\xi_\bullet$ rotates. //

Note that as z winds around 0 the configuration $u_\bullet(z)$ traces out the rotation generated by $\xi_\bullet$. And as z moves towards 0 the configuration stretches towards the point $w \in \mathbb{CP}^1$ at the bottom of the axis, meaning that all of the points of the configuration, except for the top of the axis if this is one of them, move towards w. The model for this stretching when $w = 0$ is multiplication by a positive real number t on $\mathbb{CP}^1$: as $t \to 0$ all points except ∞ converge to 0. In particular, u_v evaluates to Y_C, whilst u_e and u_f evaluate to N_C and hence have index at least 4.

Example 4.2.2. For the choices shown in Fig. 4.1, with vertices at 1, $\zeta \coloneqq e^{2\pi i/3}$ and ζ^2, we see that for all $z \in \mathbb{C}^*$ the configuration representing $u_v(z)$ has a vertex at ζ^2. As $z \to 0$, the other two vertices of $u_v(z)$ tend to $-\zeta^2$. As z moves around ∂D, these two vertices rotate around the axis of ξ_v (which fixes ζ^2). //

Remark 4.2.3. Recall from Lemma 3.3.32 that a pole of type ξ and order k reflects to a pole of type $-\xi$ and order k. In SU(2), reversing the sign of a Lie algebra element corresponds to reversing the direction of the rotation it generates, or equivalently to exchanging the top and bottom of the axis of rotation. For $C = \triangle$ we see that a quasi-axial pole of type ξ_v reflects to one of type ξ_e—since vertices are antipodal to midpoints of edges—and one of type ξ_f reflects to another of type ξ_f—since the centres of the two faces are antipodal. Similarly for $C = T$ we see that type ξ_v reflects to type ξ_f whilst type ξ_e reflects to type ξ_e. For O and I, the types of quasi-axial poles are preserved under reflection. //

Lemma 4.2.4. *The pole germs u_v, u_e and u_f have indices* 2, 6 *and* 4 *respectively.*

Proof. First note that when $C = C_\bullet$ from Appendix D, where $\bullet$ is v, e or f, then we can take $\xi_\bullet$ to be the generator of the right-handed rotation about the axis from 0 to ∞. Then $u_\bullet$ is given by

$$u_\bullet(z) = \begin{pmatrix} z^\kappa & 0 \\ 0 & z^{-\kappa} \end{pmatrix} C_\bullet,$$

where κ is the smallest positive rational number such that the right-hand side is single-valued in z. For example, for $C = O_e$ we have have $\kappa = 1/4$ and

$$u_e(z) = [z^3 : 0 : -5z^2 : 0 : -5z : 0 : 1].$$

Now focus on $C = \triangle$. The expressions we obtain for u_v and u_f extend to holomorphic spheres of degree 1 and hence index 8 (since $c_1(X_\triangle) = 4H$). They have poles at 0 of type ξ_v and ξ_f (and order 1) respectively, and reflected poles at ∞ of types ξ_e and ξ_f by Remark 4.2.3. From the second sphere we see that u_f has index 4, and since u_e evaluates to N_C but is not equivalent to u_f we conclude that it must have index at least 6. Then the first sphere shows that u_v has index 2 and u_e has index 6, as claimed.

For $C = T$, the completion of u_v has index 6 (degree 1) and shows that the poles u_v and u_f have indices 2 and 4 respectively, whilst the completion of u_e has index 12 (degree 2) and shows that the pole u_e has index 6. For O and I pole type is preserved under reflection, so the index of $u_\bullet$ is just half the index of its completion, which can be read off from the degree. □

Putting everything together, we obtain:

Proposition 4.2.5. *There is a unique family of index 2 discs, which are axial of type ξ_v and order 1. There is a unique family of index 4 discs meeting N_C, which are axial of type ξ_f and order 1. An index 4 disc not meeting N_C either has a single pole of type ξ_v and order 2 (so is a double cover of an index 2 disc) or has two poles of type ξ_v and order 1.*

4.2.2 Counting the discs

We have just seen that index 2 discs are axial of type ξ_v, whilst index 4 discs meeting N_C are axial of type ξ_f. It is straightforward to check that the stabilisers of such discs through C are precisely the subgroups of Γ_C which fix a vertex and a face respectively. By the orbit-stabiliser theorem, Proposition 3.2.17 then yields:

Lemma 4.2.6. *The degrees of the evaluation maps*

$$\mathrm{ev}_1 \colon \mathcal{M}_{\mu=2} \to L_C$$

and

$$\mathrm{ev}_1 \colon \mathcal{M}_{\mu=4}^{N_C} \to L_C$$

are, respectively, the numbers v_C and f_C of vertices and faces in the configuration C.

Remark 4.2.7. Alternatively, one can think of the degree of the evaluation map as the number of distinct conjugates of ξ_v (respectively ξ_f) by Γ_C. //

By Proposition B.4.1 we deduce:

Lemma 4.2.8. *With the standard spin structure and trivial local system we have*

$$\mathfrak{m}_0(L_C) = v_C.$$

More importantly we can use the degrees of these evaluation maps to compute $\mathcal{CO}^0$ on $\mathrm{PD}(Y_C)$ and $\mathrm{PD}(N_C)$, subject to Chern class constraints, so it is important to know what these two classes are. Since Y_C is obliging we automatically know that it is anticanonical and hence represents the class $l_C H$. Calculating $\mathrm{PD}(N_C)$ is not hard either:

Lemma 4.2.9. *For each C we have $\mathrm{PD}(N_C) = v_C E$, where E is the generator of $H^4(X_C)$ as in Section 4.1.4.*

Proof. In $H^*(X_C)$ we know that $E \smile H$ is Poincaré dual to a point, so $\mathrm{PD}(N_C) = n_C E$ where n_C is the intersection number of N_C with a hyperplane section of X_C. Taking our hyperplane section to be the set of d-point configurations in $\mathbb{P}V$ containing the point $[y]$, we see that there is a single intersection at $[y^d]$ with multiplicity $d = v_C$. $\square$

Assembling the information above gives:

Proposition 4.2.10. *For the brane $L_C^\flat$ obtained by equipping L_C with either relative spin structure and the trivial local system over $\mathbb{Z}$ we have:*

(i) The closed–open map satisfies:

C	$\mathcal{CO}^0(\mathrm{PD}(Y_C))$	$\mathcal{CO}^0(\mathrm{PD}(N_C))$
$\triangle$	$\mathcal{CO}^0(4H) = 3\varepsilon \cdot 1_{L_\triangle}$	$\mathcal{CO}^0(3E) = 2 \cdot 1_{L_\triangle}$
T	$\mathcal{CO}^0(3H) = 4\varepsilon \cdot 1_{L_T}$	$\mathcal{CO}^0(4E) = 4 \cdot 1_{L_T}$
O	$\mathcal{CO}^0(2H) = 6\varepsilon \cdot 1_{L_O}$	$\mathcal{CO}^0(6E) = 8 \cdot 1_{L_O}$
I	$\mathcal{CO}^0(H) = 12\varepsilon \cdot 1_{L_I}$	

where ε is 1 if the relative spin structure is standard and -1 if it is non-standard.

(ii) $HF^*(L_C^\flat, L_C^\flat; \mathbb{Z})$ is annihilated by 5, 4, 2 and 8 for C equal to $\triangle$, T, O and I respectively. In particular, if $HF^*(L_C^\flat, L_C^\flat; R)$ is non-zero over a field R then R has characteristic 5 if $C = \triangle$ and characteristic 2 in each of the other cases.

Proof. (i) With the standard (relative) spin structure this follows from substituting the values of c_1, v_C and f_C into the preceding results and applying Proposition 3.2.24. The reason we cannot compute $\mathcal{CO}^0(\mathrm{PD}(N_C))$ for $C = I$ is that the minimal Chern number of holomorphic spheres in X_I is 1: less than the complex codimension of N_C. With the non-standard relative spin structure recall from Remark 4.1.2 that the signs of index 2 discs are reversed whilst those of index 4 discs are preserved.

(ii) For $L_\triangle$ we have $\mathcal{CO}^0(3E) = 2 \cdot 1_{L_\triangle}$. Squaring, and using the fact that $E^2 = 1$ in $QH^*(X_\triangle)$, we obtain $5 \cdot 1_{L_\triangle} = 0$.

For L_T, square each side of $\mathcal{CO}^0(3H) = 4\varepsilon \cdot 1_{L_T}$ and apply $H^2 = 2E$ to get $\mathcal{CO}^0(18E) = 16 \cdot 1_{L_T}$. Doubling this and combining with $\mathcal{CO}^0(4E) = 4 \cdot 1_{L_T}$ yields $4 \cdot 1_{L_T} = 0$.

For L_O we have $\mathcal{CO}^0(2H^2) = \mathcal{CO}^0(6\varepsilon H) = 18 \cdot 1_{L_O}$, so from the relation $H^2 = 5E + 3$ we see that $\mathcal{CO}^0(10E) = 12 \cdot 1_{L_O}$. Now using $\mathcal{CO}^0(6E) = 8 \cdot 1_{L_O}$ we get

$$\mathcal{CO}^0(2E) = 2\,\mathcal{CO}^0(10E) - 3\,\mathcal{CO}^0(6E) = 24 \cdot 1_{L_O} - 24 \cdot 1_{L_O} = 0.$$

Applying this to twice $E^2 = E + 1$ we get $2 \cdot 1_{L_O} = 0$.

For L_I, recall from Corollary 4.1.12 that the self-Floer cohomology groups for the two choices of relative spin structure are canonically isomorphic (additively), so any additive constraint we derive on 1_{L_I} for one choice must actually apply to both choices. With respect to the basis $1, H, E, EH$ of $QH^*(X_I; \mathbb{Z})$, the matrix for quantum multiplication by $c_1(X_I) = H$ is given (using Table 4.1) by

$$\begin{pmatrix} 0 & 24 & 0 & 88 \\ 1 & 2 & 0 & 22 \\ 0 & 22 & 0 & 68 \\ 0 & 0 & 1 & 2 \end{pmatrix}.$$

The characteristic polynomial of this matrix is

$$\chi(T) = (T + 4)(T^3 - 8T^2 - 56T - 76),$$

and applying $\mathcal{CO}^0$ to the Cayley–Hamilton relation $\chi(H) = 0$, we obtain $\chi(12\varepsilon) \cdot 1_{L_I} = 0$. The greatest common divisor of $\chi(12)$ and $\chi(-12)$ is 32 so we must have $32 \cdot 1_{L_I} = 0$ for either choice of relative spin structure. Considering the Oh spectral sequence for L_I (which, recall, is an integral homology 3-sphere), we see that its self-Floer cohomology ring is concentrated in degree 0 and is isomorphic to $\mathbb{Z}/(D)$, where D is a factor of 32.

Suppose for contradiction that D is greater than 8, and assume without loss of generality that the relative spin structure is the non-standard one. From $H^2 = 22E + 2H + 24$ we obtain $\mathcal{CO}^0(22E) = 144 \cdot 1_{L_I}$, which means that $\mathcal{CO}^0(E)$ must be divisible by 8 in $\mathbb{Z}/(D)$. Now apply $\mathcal{CO}^0$ to the relation $E^2 = 2E + H + 4$. The left-hand side vanishes (since D divides 32, which divides $\mathcal{CO}^0(E)^2$), but the right-hand side is $8 \mod 16$, which is impossible. We conclude that D is actually a factor of 8, and hence 8 annihilates HF^*. ☐

Remark 4.2.11. For the icosahedron the characteristic constraint just amounts to the Auroux–Kontsevich–Seidel criterion (plus the trick of comparing the different relative spin structures). For the other three cases, however, we have obtained strictly more information: the computation of $\mathcal{CO}^0(c_1)$ alone still leaves open the possibilities of $p = 7$ for $C = \triangle$ (as mentioned by Evans–Lekili), $p = 11$ and 43 for $C = T$ (with the standard and non-standard relative spin structures respectively), and $p = 19$ for $C = O$. //

4.2.3 Bubbled configurations

In order to compute the self-Floer cohomology of L_C we need to study the moduli space $\mathcal{M}_{2,\mu=4}$ of index 4 discs with 2 boundary marked points, and the evaluation map

$$\mathrm{ev}_\pm \coloneqq (\mathrm{ev}_{-1}, \mathrm{ev}_1) \colon \mathcal{M}_{2,\mu=4} \to L_C^2.$$

In general this moduli space has boundary components comprising bubbled configurations, and when we compute the local degree of $\mathrm{ev}_\pm$ at a point $(q, p) \in L_C^2$ we need to ensure that this point does not lie in the image of the boundary. This is so that the local degree is locally constant on a neighbourhood of (q, p), and hence that we can perturb p and q if necessary to ensure transversality in the pearl complex.

So take distinct points $p, q \in L_C$ and suppose that there exists a bubbled configuration evaluating to (q, p). Either the two marked points are in a single index 2 disc component of the bubble tree, or they are in adjacent index 2 disc components. In either case, there exists a point $r \in L_C$ and index 2 discs u_1 and u_2 such that the boundary of u_1 passes through p and r whilst that of u_2 passes through q and r (in the first case we just take $u_1 = u_2$).

By the classification of index 2 discs, this means that there exist vertices v_1 and v_2 of the d-point configuration representing r such that p and q are obtained from r by rotating around v_1 and v_2 respectively. So v_1 lies in the configuration representing p whilst v_2 lies in that representing q. In other words, there exist a vertex of p and a vertex of q whose angle (or, equivalently, distance) of separation coincides with the angle between two vertices of C, which need not be distinct. And conversely if there are two such vertices then a bubbled configuration does exist.

Now note that if w_1 and w_2 are non-zero vectors in the fundamental representation V of SU(2) then the angle θ between the points $[w_1]$ and $[w_2]$ on the sphere $\mathbb{P}V$ satisfies

$$\cos\frac{\theta}{2} = \frac{|\langle w_1, w_2\rangle|}{\|w_1\|\|w_2\|}. \tag{4.3}$$

This can be verified easily when $w_1 = (0, 1)$, and then the general result follows from the SU(2)-invariance of both sides. So there exists an index 4 bubbled configuration through p and q, if and only if the sets

$$\left\{\frac{|\langle w_1, w_2\rangle|^2}{\|w_1\|^2\|w_2\|^2} : [w_1] \text{ a vertex of } p \text{ and } [w_2] \text{ a vertex of } q\right\} \tag{4.4}$$

and

$$\left\{\frac{|\langle w_1, w_2\rangle|^2}{\|w_1\|^2\|w_2\|^2} : [w_1] \text{ and } [w_2] \text{ vertices of } C\right\} \tag{4.5}$$

intersect.

For example, if $p, q \in L_\triangle$ are represented by the triangles with vertices $\blacktriangle$ and $\bullet$ respectively on the left-hand diagram in Fig. 4.2 then there *is* a bubbled configuration through p and q because the vertices in the southern hemisphere are distance $2\pi/3$ apart. The third vertex of the configuration r mentioned above would be at ∞. In contrast, there is *no* bubbled configuration through the points p and q shown in the right-hand diagram since in this case the distances between vertices of p and vertices of q are $\pi/3$ and π, neither of which appears as a distance between vertices in a single equilateral triangle. In terms of the two sets above, (4.5) is clearly

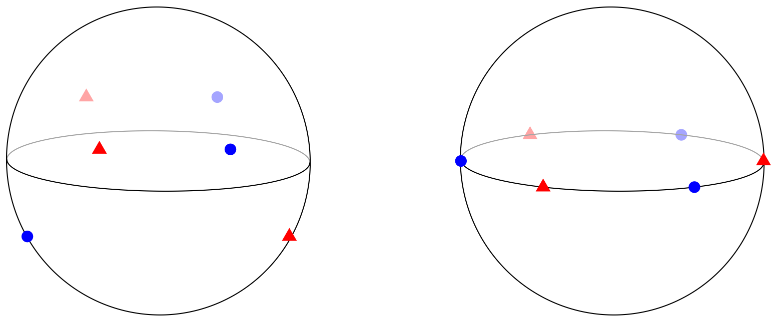

Figure 4.2: Choices of $p, q \in L_\triangle$ demonstrating existence and non-existence of bubbled configurations.

$\{1/4, 1\}$ in both cases, whilst (4.4) is easily seen to contain $1/4$ for the left-hand diagram but is given by $\{0, 3/4\}$ for the right-hand diagram.

4.3 The Morse and pearl complexes

In this section we construct the pearl complexes for the L_C, up to some unknown disc counts which we deal with later. We equip each L_C with an arbitrary relative spin structure and the trivial local system (over $\mathbb{Z}$) to give a brane $L_C^\flat$.

4.3.1 Stereographic projection

To understand the topology of the Lagrangian $L_C \cong \mathrm{SU}(2)/\Gamma_C$ we seek the fundamental domain for the action of SU(2) on C, in other words the set of points of SU(2) which are no further from the identity, I, than from any other element of Γ_C. We think of SU(2) as the unit sphere in the quaternions $\mathbb{H}$, via

$$\mathrm{SU}(2) = \left\{ \begin{pmatrix} u & -\overline{v} \\ v & \overline{u} \end{pmatrix} : u, v \in \mathbb{C} \text{ and } |u|^2 + |v|^2 = 1 \right\},$$

and identify $\mathbb{H}$ with $\mathbb{R}^4$ by $(u, v) \leftrightarrow (\operatorname{Re} u, \operatorname{Im} v, \operatorname{Re} v, \operatorname{Im} u)$. The left- and right-multiplication actions of SU(2) on $\mathbb{H}$ clearly preserve the standard inner product, and hence the induced round metric on SU(2) is bi-invariant. In particular, one-parameter subgroups of SU(2) are geodesics in the round metric. We view the fundamental domain as a subset of $\mathbb{R}^3 = \{0\} \times \mathbb{R}^3 \subset \mathbb{R}^4$ by stereographic projection from $-I$:

$$\begin{pmatrix} u & -\overline{v} \\ v & \overline{u} \end{pmatrix} \in \mathrm{SU}(2) \mapsto \frac{1}{1 + \operatorname{Re} u} (\operatorname{Im} v, \operatorname{Re} v, \operatorname{Im} u). \tag{4.6}$$

Recall from Section 4.1.2 the identification of $\mathbb{CP}^1$ with the unit sphere in $\mathbb{R}^3$, which gives the action of SU(2) as rotations of the sphere that we have been using throughout. With these conventions, the rotation through angle $\theta \in [0, \pi]$ about a unit vector $\mathbf{n} \in \mathbb{R}^3$ lifts to $\exp(\mathbf{n} \cdot i\theta\overline{\sigma}/2)$ in SU(2), where σ is the vector of Pauli matrices

$$\sigma = \left(\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \right).$$

It also lifts to $-I$ times this, but we only need consider the representative closest to I. Geodesics on SU(2) through I correspond to intersections of 2-planes through $\pm I$ in $\mathbb{R}^4$ with SU(2), and hence their stereographic projections are straight lines in $\mathbb{R}^3$. For each fixed $\mathbf{n}$ the one-parameter subgroup $\theta \mapsto \exp(\mathbf{n} \cdot i\theta\overline{\sigma}/2)$ is therefore sent by (4.6) to a straight line in $\mathbb{R}^3$. It is easy to check that the line is in the direction $\mathbf{n}$, and by restricting the stereographic projection to the 2-plane through $\pm I$ containing this direction it is easy to compute that in fact $\exp(\mathbf{n} \cdot i\theta\overline{\sigma}/2)$ projects to $\tan(\theta/4)\mathbf{n}$—see Fig. 4.3.

The set of points in SU(2) which are equidistant from $\exp(\mathbf{n} \cdot i\theta\overline{\sigma}/2)$ and the identity is an equatorial 2-sphere on $\mathrm{SU}(2) \cong S^3$ which cuts the geodesic generated by $\mathbf{n} \cdot i\overline{\sigma}/2$ orthogonally at the points

$$\exp(\mathbf{n} \cdot i(\theta/2)\overline{\sigma}/2) \text{ and } \exp(\mathbf{n} \cdot i(\theta/2 + 2\pi)\overline{\sigma}/2).$$

Under stereographic projection this maps to a sphere in $\mathbb{R}^3$ which cuts the line in direction $\mathbf{n}$ orthogonally at $\tan(\theta/8)\mathbf{n}$ and $\tan(\theta/8 + \pi/2)\mathbf{n} = -\cot(\theta/8)\mathbf{n}$. Its centre is thus at

$$\frac{1}{2}(\tan(\theta/8) - \cot(\theta/8))\mathbf{n} = -\cot(\theta/4)\mathbf{n}$$

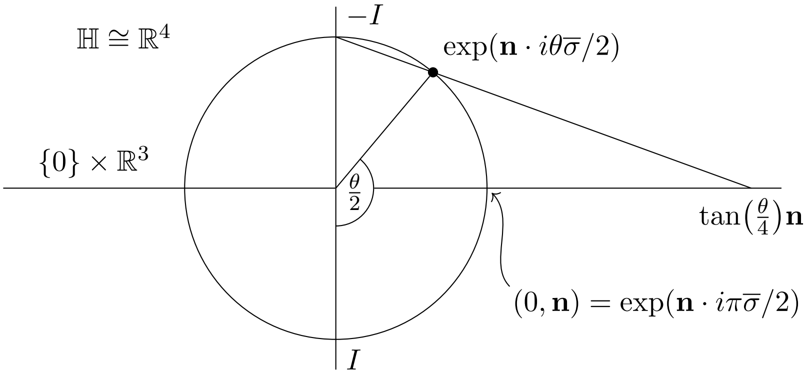

Figure 4.3: Stereographic projection on SU(2).

and its radius is $(\tan(\theta/8) + \cot(\theta/8))/2 = \operatorname{cosec}(\theta/4)$.

In fact, the fundamental domain is contained in the unit ball in $\mathbb{R}^3$, since this corresponds to the set of points in SU(2) closer to I than to $-I$, and it is convenient to compose (4.6) with the diffeomorphism $\mathbf{x} \mapsto \mathbf{2x}/(1 - \|\mathbf{x}\|^2)$ from the unit ball to the whole of $\mathbb{R}^3$; this corresponds to replacing the denominator in (4.6) by $\operatorname{Re} u$, and results in $\exp(\mathbf{n} \cdot i\theta\overline{\sigma}/2)$ projecting to $\tan(\theta/2)\mathbf{n}$ rather than $\tan(\theta/4)\mathbf{n}$. If $\mathbf{x}$ denotes the coordinate on the unit ball, and $\mathbf{y}$ the coordinate on the codomain $\mathbb{R}^3$, the ball

$$\left\{\|\mathbf{x} + \cot(\theta/4)\mathbf{n}\|^2 \leq \operatorname{cosec}^2(\theta/4)\right\} = \left\{2\cot(\theta/4)\mathbf{x} \cdot \mathbf{n} \leq 1 - \|\mathbf{x}\|^2\right\}$$

(or strictly its intersection with the unit ball) is sent to the half-space

$$\left\{\mathbf{y} \cdot \mathbf{n} \leq \tan(\theta/4)\right\}. \tag{4.7}$$

For any α, the right action of $\exp(\mathbf{n} \cdot i\theta\overline{\sigma}/2)$ on SU(2) sends $\{\mathbf{y} \cdot \mathbf{n} = \tan\alpha\}$ to $\{\mathbf{y} \cdot \mathbf{n} = \tan(\alpha + \theta/2)\}$, by translation in the direction $\mathbf{n}$ combined with a left-handed rotation about $\mathbf{n}$ through angle $\theta/2$. This is easy to check when $\mathbf{n} = (0, 0, 1)$, and the general case then follows from the fact that the projection intertwines the action of SU(2) on itself by conjugation with its action on $\mathbb{R}^3$ by rotation (this in turn is easy to check infinitesimally). We refer to the map

$$\begin{pmatrix} u & -\overline{v} \\ v & \overline{u} \end{pmatrix} \in \mathrm{SU}(2) \mapsto \frac{1}{\operatorname{Re} u}(\operatorname{Im} v, \operatorname{Re} v, \operatorname{Im} u)$$

as *modified stereographic projection*, and will usually use coordinates (x, y, z) on $\mathbb{R}^3$ (instead of the $\mathbf{y}$ used above).

This projection also has the advantage that boundaries of axial discs are sent to straight line segments in the fundamental domain. To see this, note that the boundary of an axial disc is (the image in L_C of) a right translate of one-parameter subgroup of SU(2). We have seen that such a path is a geodesic on SU(2), i.e. the intersection of $\mathrm{SU}(2) \cong S^3 \subset \mathbb{R}^4$ with a 2-plane through the origin. We can therefore write it as the intersection of two equatorial 2-spheres, which we know project to planes.

4.3.2 Triangle

The constructions of the fundamental domain, the Heegaard splitting and the Morse function are based on [45, Section 5], though are not identical. In particular, our projection leads to left-handed face identifications rather than right-handed. For convenience of comparison, we employ matching notation. We fix our choice of representative configuration $\triangle$ as $[x^3+y^3]$, i.e. $\triangle_f$ from Appendix D. Then the stabiliser of $\triangle$ under the SU(2)-action consists of (the lifts to SU(2) of) rotations about $(0,0,1)$ through angle $\pm 2\pi/3$, and rotations about $(1,0,0)$, $(-1/2,\sqrt{3}/2,0)$ and $(-1/2,-\sqrt{3}/2,0)$ through angle $\pm\pi$.

Plugging these into (4.7), we see that under modified stereographic projection the image of the fundamental domain is $H\times[-1/\sqrt{3},1/\sqrt{3}]$, where H is the regular hexagon with vertices at $2/\sqrt{3}$ times the sixth roots of -1. Each square face of this hexagonal prism is identified with the opposite face via a left-handed rotation through angle $\pi/2$, whilst the hexagonal faces are identified by a left-handed rotation through angle $\pi/3$.

We take a genus 3 Heegaard splitting of $L_\triangle$ whose handlebodies are a thickening of the edges and hexagonal faces of the prism and a thickening of the three lines joining the centres of opposite square faces. Figure 4.4 shows the prism with these two sets marked in the left and right diagrams respectively. It also shows the critical points of the Morse function built by Evans–Lekili from this splitting: the maximum is at m, the index 2 critical points at x_1, x_2 and x_3, the index 1 critical points at x_1', x_2' and x_3' and the minimum at m'.

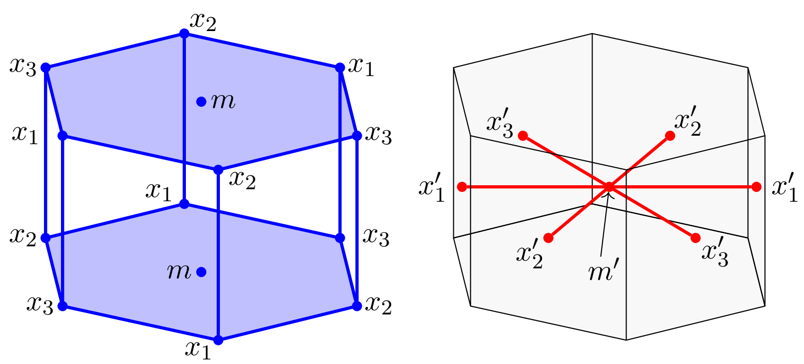

Figure 4.4: The fundamental domain, Heegaard splitting and critical points for $C=\triangle$.

Figure 4.5 shows the point x_3', the front face of the prism centred on it, and the shape of trajectories in its ascending manifold close to this face. The solid trajectories belong to the descending manifolds of x_1, x_2 (twice) and x_3, the dotted trajectories flow into the fundamental domain towards m, whilst the dashed trajectories flow out of the domain, and so back into the opposite face with a twist of $\pi/2$, again towards m. Of course we really need to choose a metric on $L_\triangle$ in order to talk about trajectories, but we have a natural choice: that induced from the standard round metric on $\mathrm{SU}(2)\cong S^3$.

The point m' at the centre of the fundamental domain represents the identity in SU(2). We saw earlier that our modified stereographic projection sends the one-parameter subgroup of SU(2) comprising (the lifts of) rotations about an axis $l\subset\mathbb{R}^3$ to l itself. The boundaries of the three index 2 discs through m' correspond to the one-parameter subgroups of (lifts of) rotations of $\triangle$ about vertices, and therefore project to the axes through the vertices. These are the lines

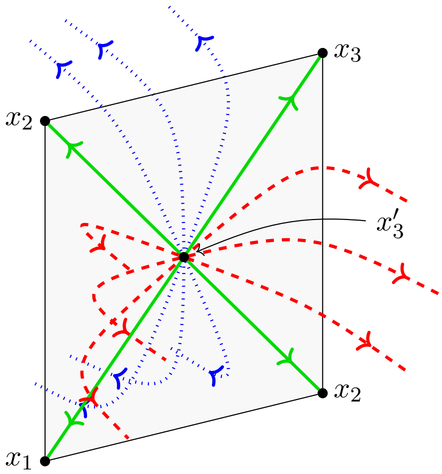

Figure 4.5: Trajectories in the ascending manifold of x_3'.

joining the centres of the opposite square faces of the fundamental domain, i.e. the core circles of the second handlebody. The point m represents the rotation of $\triangle$ through angle $\pi/3$ about a vertical axis, so the lifts to SU(2) of the boundaries of the index 2 discs through m are obtained from those through m' by multiplying on the right by (a lift of) this rotation. We have seen that this right-multiplication action corresponds to translating m' to m and rotating by angle $\pi/6$ about a vertical axis, so the boundaries of the index 2 discs through m are the diagonals of the hexagonal faces of the fundamental domain.

The rotational symmetry group of the triangle $\triangle$ in SO(3) acts on the fundamental domain respecting the Heegaard splitting. This corresponds to the action of Γ_C on SU(2) by conjugation, so preserves the round metric, and we may assume that it also preserves the Morse function. It therefore permutes the descending manifolds, and we may choose orientations on them which are invariant under this action. To see this note that it trivially preserves the orientations of the descending manifolds of m' and m, which are a point and a dense open subset of L_C respectively, and by inspection we can choose invariant orientations on the descending manifolds of the x_i', which we can take to be the boundaries of the index 2 discs through m' (minus the point m' itself). Similarly we can choose invariant orientations on the *ascending* manifolds of the x_i, which we take to be the boundaries of the index 2 discs through m (minus the point m). This gives invariant coorientations on their descending manifolds, and since the orientation of L_C is itself invariant, these invariant coorientations can be turned into invariant orientations.

In order to ensure transversality in the pearl complex we perturb the auxiliary data. Proposition C.3.1 shows that we can pull back the Morse function and metric by a diffeomorphism φ arbitrarily C^∞-close to the identity in order to achieve the necessary genericity, so from now on we assume this has been done. We take φ sufficiently close to $\mathrm{id}_{L_\triangle}$ that the later arguments involving intersections of discs with various ascending and descending manifolds, and the index 4 count in Section 4.4.1, are valid. Although the Morse function and metric may themselves no longer be invariant under the action of Γ_C, they are small perturbations of symmetric data, and the perturbations do not affect the symmetry of the orientations.

Choosing the invariant orientations on the descending manifolds appropriately, the Morse differentials d_M are

$$\begin{aligned}\mathrm{d}_M m' &= 0\\ \mathrm{d}_M x_i' &= x_i + x_{i+1} + 2x_{i+2}\\ \mathrm{d}_M x_i &= 0,\end{aligned}$$

with subscripts understood modulo 3. The coefficient 2 in $\mathrm{d}_M x_i'$ corresponds to the two solid flowlines towards x_2 in Fig. 4.5. They both count with the same sign as they differ by rotation about the axis through the centre of the face shown in the figure, whilst the overall differential is invariant under cycling the i because this corresponds to rotations through angle $2\pi/3$ about a vertical axis. There is no quantum correction to $\mathrm{d}_M m'$ in the Floer (pearl) differential $\mathrm{d}m'$ for degree reasons, whilst that to $\mathrm{d}_M x_i'$ is $X_i m'$, where X_i counts flows upward from x_i' and then along an index 2 disc to m' (in other words, intersections of index 2 discs through m' with the ascending manifold of x_i'). There is one such trajectory for each i, so each X_i is ± 1, and by the symmetry we can replace all of the X_i by a single $X \in \{\pm 1\}$.

The correction for m can be written (again using approximate cyclic symmetry) in the form $Y(x_1 + x_2 + x_3) + \hat{Z}m'$, for $Y, \hat{Z} \in \mathbb{Z}$. These count respectively the index 2 trajectories $m \rightsquigarrow x_i$ and the index 4 trajectories $m \rightsquigarrow m'$. The former comprise intersections of the index 2 discs through m with the descending manifolds of the x_i; there is one of these for each i so $Y \in \{\pm 1\}$. The latter comprise index 4 discs through m and m', which we count in Section 4.4.1. There are no index 4 contributions of the form 'index 2 disc through m, flow, index 2 disc through m'', since the boundaries of these index 2 discs stay within their respective handlebodies and the Morse flow goes from the the m' handlebody to the m handlebody and not vice versa. Note that Evans–Lekili write $2Z$ for the count we are calling $\hat{Z}$.

Putting everything together, the $\mathbb{Z}/2$-graded Floer (pearl) cochain complex is

$$CF^0(L^\flat_\triangle, L^\flat_\triangle; \mathbb{Z}) = \langle m', x_1, x_2, x_3 \rangle \text{ and } CF^1(L^\flat_\triangle, L^\flat_\triangle; \mathbb{Z}) = \langle x_1', x_2', x_3', m \rangle$$

(where $\langle \cdot \rangle$ now indicates the free $\mathbb{Z}$-module generated by $\cdot$), and the Floer differentials

$$\mathrm{d}^0 \colon CF^0 \to CF^1 \text{ and } \mathrm{d}^1 \colon CF^1 \to CF^0$$

are given in these bases by

$$\mathrm{d}^0 = \begin{pmatrix} 0 & & & \\ 0 & & A & \\ 0 & & & \\ 0 & 0 & 0 & 0 \end{pmatrix} \text{ and } \mathrm{d}^1 = \begin{pmatrix} X & X & X & \hat{Z} \\ 1 & 2 & 1 & Y \\ 1 & 1 & 2 & Y \\ 2 & 1 & 1 & Y \end{pmatrix}.$$

The matrix A describes the as yet unknown index 2 corrections to $\mathrm{d}_M x_i$, but from the fact that $\mathrm{d}^1 \circ \mathrm{d}^0 = 0$ it is easy to see that actually $A = 0$. The Smith normal form of d^1 is thus diagonal with entries 1, 1, 1 and $\det \mathrm{d}^1 = 3XY - 4\hat{Z}$. This is essentially the argument given by Evans–Lekili.

4.3.3 Tetrahedron

We follow a similar strategy for the tetrahedron, taking $T = [x^4 + 2\sqrt{3}x^2y^2 - y^4]$, with vertices at $\pm\sqrt{2}/(\sqrt{3}-1)$ and $\pm(\sqrt{3}-1)i/\sqrt{2}$—this is exactly T_e from Appendix D. Under modified stereographic projection there are 14 possible planes contributing to the boundary of the fundamental domain: 8 from rotations about vertices (or equivalently about the centres of faces) through angle $\pm 2\pi/3$, and 6 from rotations about the mid-points of edges through angle $\pm\pi$. In fact, only the former are needed. To see this note that the constraint coming from rotation about the bottom edge is $z \geq -1$ (applying (4.7)), whilst rotations about the two lower vertices give $z \pm \sqrt{2}x \geq -1$. Adding the latter two inequalities gives the edge rotation inequality for free.

The fundamental domain is therefore a regular octahedron, with vertices at

$$(1/\sqrt{2}, \pm 1/\sqrt{2}, 0),\ (-1/\sqrt{2}, \pm 1/\sqrt{2}, 0) \text{ and } (0,0,\pm 1).$$

Opposite faces are identified by a left-handed rotation through angle $\pi/3$, and in particular all of the vertices are identified. We take a genus 4 Heegaard splitting with handlebodies given by thickening the edges of the octahedron and the line segments joining opposite pairs of faces, shown in Fig. 4.6 in the left and right diagrams respectively. From this splitting we construct a Morse function with maximum at m, index 2 critical points at x_1, x_2, x_3 and x_4, index 1 critical points at x_1', x_2', x_3' and x_4' and minimum at m'.

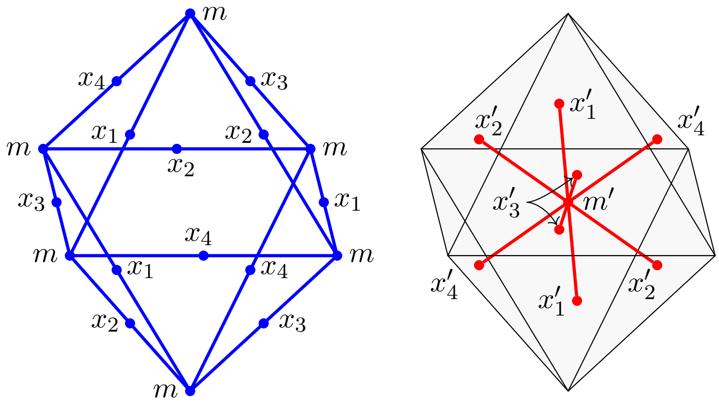

Figure 4.6: The fundamental domain, Heegaard splitting and critical points for $C = T$.

The ascending manifolds of the x_i' are the faces of the octahedron, whilst the ascending manifolds of the x_i are the edges. The descending manifolds of the x_i are, locally, small discs orthogonal to the edges. Looking down from above onto the top half of the octahedron, we orient the faces and discs as indicated in Fig. 4.7 by the dotted and dashed arrows respectively. As for $\triangle$ we perturb the Morse function and metric to ensure transversality, by pulling them back along a diffeomorphism near the identity.

Up to an overall sign depending on the chosen orientation of L_C, which we can eliminate if necessary by reversing the orientations of the faces defined above, the Morse differentials d_M

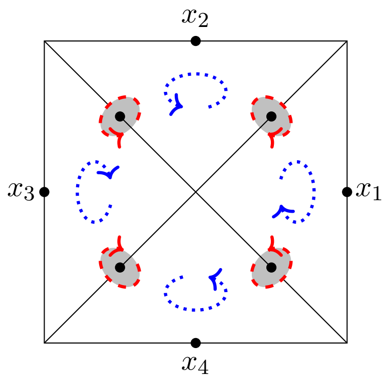

Figure 4.7: Orientations of the ascending manifolds of the x_i' (dotted) and of the descending manifolds of the x_i (dashed).

are then

$$\begin{aligned} \mathrm{d}_M m' &= 0 \\ \mathrm{d}_M x_i' &= x_{i+1} + x_{i+2} + x_{i+3} \\ \mathrm{d}_M x_i &= 0, \end{aligned}$$

with subscripts modulo 4.

The quantum correction to $\mathrm{d}_M x_i$ vanishes by the same $\mathrm{d}^1 \circ \mathrm{d}^0 = 0$ argument we used earlier, and there is no correction to $\mathrm{d}_M m'$ for degree reasons. The index 2 correction to $\mathrm{d}_M x_i'$ counts upward flows from x_i' into index 2 discs through m', i.e. intersections between index 2 discs through m' and the ascending manifold of x_i': there is one such intersection for each i, given by the thick lines hitting the faces of the octahedron in the right-hand diagram Fig. 4.6. So

$$\mathrm{d}x_i' = x_{i+1} + x_{i+2} + x_{i+3} + Xm'$$

for some $X \in \{\pm 1\}$. The reason why the intersections all carry the same sign is that the relative orientations of the disc boundaries match up with the relative orientations of the faces.

Finally, the correction to $\mathrm{d}_M m$ counts index 4 discs through m and m'—of which there are Z, say—and intersections between index 2 discs through m with the descending manifolds of the x_i (Z is analogous to the count we called $\hat{Z}$ for the triangle, but we drop the $\hat{}$ to reduce clutter). There are no index 4 trajectories of the form 'index 2, flow, index 2' since index 2 discs through m and m' remain in their respective handlebodies as before. The count of index 2 discs through m hitting the descending manifold of x_i is ± 1 for each i, so we have

$$\mathrm{d}m = Y(x_1 + x_2 + x_3 + x_4) + Zm'$$

for some $Y \in \{\pm 1\}$. Again the index 2 contributions all carry the same sign, because the relative orientations of disc boundaries and descending manifolds match up.

Thus the Floer cochain complex is

$$CF^0(L_T^\flat, L_T^\flat; \mathbb{Z}) = \langle m', x_i \rangle \text{ and } CF^1(L_T^\flat, L_T^\flat; \mathbb{Z}) = \langle x_i', m \rangle,$$

and with respect to these bases the Floer differentials d^0 and d^1 are given by $\mathrm{d}^0 = 0$ and

$$\mathrm{d}^1 = \begin{pmatrix} X & X & X & X & Z \\ 0 & 1 & 1 & 1 & Y \\ 1 & 0 & 1 & 1 & Y \\ 1 & 1 & 0 & 1 & Y \\ 1 & 1 & 1 & 0 & Y \end{pmatrix}.$$

The Smith normal form of d^1 is diagonal with entries 1, 1, 1, 1 and $4XY - 3Z$.

4.3.4 Octahedron

For $C = O$ one could similarly construct a Heegaard splitting and Morse function under stereographic projection, but we don't need to do this explicitly. We just need to take a Morse function and metric on L_O, with a unique local maximum (at m, say) and a unique local minimum (at m'), and replace them with appropriate pullbacks as given by Proposition C.3.1.

Let the index 2 critical points be at $x_1, \ldots, x_k$, and the index 1 critical points be at $x'_1, \ldots, x'_k$; there are equal numbers of each by considering the Euler characteristic of the Morse complex. Then the Floer cochain complex is

$$CF^0(L_O^\flat, L_O^\flat; \mathbb{Z}) = \langle m', x_i \rangle \text{ and } CF^1(L_O^\flat, L_O^\flat; \mathbb{Z}) = \langle x'_i, m \rangle,$$

and the Floer differentials d^0 and d^1 have the form

$$\mathrm{d}^0 = \begin{pmatrix} 0 & & & \\ \vdots & & A & \\ 0 & & & \\ 0 & 0 & \cdots & 0 \end{pmatrix} \text{ and } \mathrm{d}^1 = \begin{pmatrix} b_1 & \cdots & b_k & D \\ & & & c_1 \\ & M & & \vdots \\ & & & c_k \end{pmatrix},$$

where M is the Morse differential $\langle x'_i \rangle \to \langle x_i \rangle$, the matrix A and the vectors $B = (b_i)$ and $C = (c_i)$ represent index 2 corrections, and the number D (not the unit disc!) is the index 4 correction to $\mathrm{d}_M m$. By the comments at the end of Section 4.1.6, A, B and C all vanish, whilst the count D involves only index 4 discs and is even. The cokernel of M is exactly the Morse cohomology group $H^2(L_O; \mathbb{Z}) \cong \mathbb{Z}/2$, so its Smith normal form has diagonal entries $1, \ldots, 1, 2$. The Smith normal form of d^1 therefore has diagonal entries $1, \ldots, 1, 2, D$.

Remark 4.3.1. From Proposition 4.2.10(ii) we know that D divides 2, and since it is even we deduce it must in fact be equal to ± 2. //

Remark 4.3.2. The fact that the signed count of index 4 discs through two generic points p and q of L_O is ± 2 can be understood as follows. Any such disc doubles to a rational curve of index 8 and hence degree 2 (since $2c_1(X_O) = 4H$), which must therefore be contained in some 2-plane in the ambient projective space $\mathbb{P}S^6V$ in which X_O lives. This curve passes through p and q tangent to X_O (since the whole curve is contained in X_O), and generically these two tangent 3-planes $T_pX_O, T_qX_O \subset \mathbb{P}S^6V$ meet in a single point, r. The plane of the curve is then spanned by p, q and r. The fact that the count is ± 2 tells us that the intersection of the plane $\langle p, q, r \rangle$

with X_O is indeed a degree 2 curve with equator on L_O. We already remarked in Section 4.1.6 that the two hemispheres should count with the same sign. //

4.4 Index 4 discs and computation of Floer cohomology

We now compute the unknown counts in the complexes. For Sections 4.4.1 to 4.4.3 we give each Lagrangian the trivial local system and standard spin structure. Other relative spin structures are considered in Section 4.4.4.

4.4.1 Triangle

Let $p = [x^3 + y^3]$ and $q = [x^3 - y^3]$, representing equilateral triangles on the equator of $\mathbb{P}V$ which differ by a rotation through angle π about a vertical axis. These are the (unperturbed) points m' and m respectively from Section 4.3.2.

Proposition 4.4.1. *(q, p) is a regular value of the two-point index 4 evaluation map*

$$\mathrm{ev}_{\pm 1} \colon \mathcal{M}_{2,\mu=4} \to L_{\triangle}^2,$$

with exactly two preimages.

Proof. There are two axial discs of type ξ_f and order 1 passing through p and q. They are reflections of each other and their boundaries sweep the rotation from p to q to p about a vertical axis in either direction, through a total angle of $2\pi/3$. Since the infinitesimal action of $\mathfrak{su}(2)$ and $\mathfrak{sl}(2, \mathbb{C})$ on $N_{\triangle}$ both have real rank 2, one can readily verify from the splitting constructed in Section B.3 that these discs have partial indices 1, 1 and 2, so the evaluation map $\mathrm{ev}_{\pm 1}$ is a submersion at the corresponding points of $\mathcal{M}_{2,\mu=4}$.

We now check that there are no other index 4 discs through p and q. It is easy to see that there can be no axial discs of type ξ_v and order 2 (since the configurations p and q do not have a common vertex), so by Proposition 4.2.5 we are left to rule out discs with two poles of type ξ_v and order 1.

Suppose for contradiction then that $u \colon (D, \partial D) \to (X_{\triangle}, L_{\triangle})$ is such a disc passing through p and q. Recall from Remark 4.2.3 that the poles of type ξ_v reflect to poles of type ξ_e, and hence of index 6, so the double $\widetilde{u}$ is a rational curve in $X_{\triangle}$ of index 16 and thus degree 2 (since $2c_1(X_{\triangle}) = 8H$). It is therefore either a smooth conic or a double cover of a line. If the latter, the image of $\widetilde{u}$ would be the line through p and q, but this line does not intersect $Y_{\triangle} \setminus N_{\triangle}$. We know, however, that $\widetilde{u}$ *does* meet this set, at the poles of u, so this case is impossible. Therefore $\widetilde{u}$ must be a smooth conic, and hence an embedding into a 2-plane Π in $X_{\triangle} = \mathbb{CP}^3$.

The poles of $\widetilde{u}$ of type ξ_e evaluate to N_C, say to points

$$P \coloneqq [(ax + y)^3] \text{ and } Q \coloneqq [(bx + y)^3] \in N_{\triangle}$$

with $a, b \in \mathbb{CP}^1$. Note that a and b must be distinct as $\widetilde{u}$ is injective, and the points p, q, P and Q must be coplanar (they're all in Π). These conditions are inconsistent with a and b both being finite and non-zero. We may therefore assume without loss of generality that $a = \infty$ and that b is finite—the case where one of a and b is zero is analogous. Then the poles of type ξ_v evaluate to

$$R \coloneqq [(x+cy)y^2] \text{ and } S \coloneqq [((bx+y)+d(x-\bar{b}y))(x-\bar{b}y)^2] \in Y_C \setminus N_C,$$

for some $c, d \in \mathbb{C}$, using that R and S lie in $Y_\triangle \setminus N_C$ and satisfy $\tau(R) = P$ and $\tau(S) = Q$.

From coplanarity of p, q, P, Q and R (again, all lie in Π) we deduce that $b = 0$. But then P, Q, R and S have standard coordinates $[1:0:0:0]$, $[0:0:0:1]$, $[0:0:1:c]$ and $[d:1:0:0]$, so cannot be coplanar, giving a contradiction. Hence no such two-pole index 4 disc u can exist, and we're done. $\square$

Note also that there is no index 4 bubbled configuration through p and q. In fact this is precisely the right-hand example given in Fig. 4.2 in Section 4.2.3. Hence there is an open neighbourhood U of (q,p) in $L_\triangle^2$ such that each point in U is a regular value of $\mathrm{ev}_{\pm 1}$, and the local degree (i.e. signed count of preimages) remains constant on U.

We can now compute:

Corollary 4.4.2. *The index* 4 *count* $\widehat{Z}$ *appearing in Section 4.3.2 is* ± 2*, and the determinant* $3XY - 4\widehat{Z}$ *of the pearl complex differential* d^1 *is* ± 5*. The self-Floer cohomology ring of* $L_\triangle^\flat$ *over* $\mathbb{Z}$ *satisfies*

$$HF^0(L_\triangle^\flat, L_\triangle^\flat; \mathbb{Z}) \cong \mathbb{Z}/(5) \text{ and } HF^1(L_\triangle^\flat, L_\triangle^\flat; \mathbb{Z}) = 0.$$

If R *is a field of characteristic* 5 *then we have additive isomorphisms*

$$HF^0(L_\triangle^\flat, L_\triangle^\flat; R) \cong HF^1(L_\triangle^\flat, L_\triangle^\flat; R) \cong R.$$

Proof. First note that the determinant $3XY - 4\widehat{Z}$ must be ± 1 or ± 5, since 5 annihilates HF^* by Proposition 4.2.10(ii). By Proposition 4.4.1, and the absence of bubbled configurations, we have $\widehat{Z} = \pm 2$ (if the two discs count with the same sign) or 0 (if they count with opposite signs). Recalling that X and Y are equal to ± 1, the only possibility is that $\widehat{Z}$ is ± 2 and $3XY - 4Z$ is ± 5. Plugging the latter into the Smith normal form of the Floer differential calculated in Section 4.3.2 gives the claimed cohomology. $\square$

Remark 4.4.3. Note that the Riemann–Hilbert pairs $(u^*TX_\triangle, u|_{\partial D}^* TL_\triangle)$ associated to the two axial discs of type ξ_f are reflections of each other, so by Remark 4.1.11 we should expect these discs to contribute with the same sign. //

4.4.2 Tetrahedron

Now let $p = [x^4 + 2\sqrt{3}x^2y^2 - y^4]$ and $q = [x^4 - 2\sqrt{3}x^2y^2 - y^4]$, representing regular tetrahedra with an opposite pair of horizontal edges, differing by rotation through angle $\pi/2$ about a vertical axis. These are the points m' and m respectively from Section 4.3.3.

Proposition 4.4.4. (q,p) *is not in the image of of the two-point index* 4 *evaluation map*

$$\mathrm{ev}_{\pm 1}\, \mathcal{M}_{2,\mu=4} \to L_T^2$$

(and so is vacuously a regular value).

Proof. Since the configurations p and q have no vertex in common there are no axial discs of type ξ_v passing through them both. Similarly since there is no face of p which differs from a face of q by rotation about its centre there are no axial discs of type ξ_f passing through p and q. So now suppose for contradiction that $u\colon (D, \partial D) \to (X_T, L_T)$ is a two-pole index 4 disc through p and q with double $\widetilde{u}$. This time the two poles of type ξ_v from u reflect to poles of type ξ_f. Again $\deg \widetilde{u} = 2$ but now we can rule out the double cover of a line for more trivial reasons: the line in $\mathbb{P}S^4V$ through p and q does not lie in X_T (it contains the point $[x^2y^2]$ for example). Hence $\widetilde{u}$ is an embedding.

Considering the points

$$P \coloneqq [(ax+y)^4] \text{ and } Q \coloneqq [(bx+y)^4] \in N_T,$$

to which the poles of type ξ_f evaluate, and the fact that they must be coplanar with p and q, we get either $a = 0$ and $b = \infty$ (or vice versa) or that a is equal to $-b$ and is a fourth root of -1. Each of these cases leads to a contradiction by looking at the possible reflections of P and Q, as with $C = \triangle$. $\square$

Again there is no bubbled configuration through p and q. To see this recall from Section 4.3.3 that the vertices of p are at $\pm\sqrt{2}/(\sqrt{3}-1)$ and $\pm(\sqrt{3}-1)i/\sqrt{2}$. Those of q differ by multiplication by i (rotation by $\pi/2$ about a vertical axis), so we can explicitly compute the sets (4.4) and (4.5) from Section 4.2.3. The former is $\{0, 2/3\}$, whilst the latter is $\{1/3, 1\}$, and these are clearly disjoint. We can therefore perturb p and q slightly without introducing any preimages. We get:

Corollary 4.4.5. *The index* 4 *count* Z *in Section 4.3.3 is* 0*, and the determinant* $4XY - 3Z$ *of* d^1 *is* ± 4*. The self-Floer cohomology ring of* $L_T^\flat$ *over* $\mathbb{Z}$ *satisfies*

$$HF^0(L_T^\flat, L_T^\flat; \mathbb{Z}) \cong \mathbb{Z}/(4) \text{ and } HF^1(L_T^\flat, L_T^\flat; \mathbb{Z}) = 0.$$

If R *is a field of characteristic* 2 *then we have additive isomorphisms*

$$HF^0(L_T^\flat, L_T^\flat; R) \cong HF^1(L_T^\flat, L_T^\flat; R) \cong R.$$

Proof. The value of Z follows immediately from Proposition 4.4.4 (plus the absence of bubbled configurations), and then, recalling that X and Y are ± 1, the determinant must be ± 4. Substituting into the Floer differential in Section 4.3.3 gives the cohomology. $\square$

4.4.3 Octahedron

For the octahedron the self-Floer cohomology over $\mathbb{Z}$ is completely determined by the shape of the pearl complex computed in Section 4.3.4 and the constraints imposed by the closed–open map—one doesn't need to calculate any index 4 discs beyond the axial ones appearing in $\mathcal{CO}^0$. The result is:

Proposition 4.4.6. *We have an isomorphism of unital rings*

$$HF^0(L^\flat_O, L^\flat_O; \mathbb{Z}) \cong \mathbb{Z}[x]/(2, x^2 + x + 1) \text{ and } HF^1(L^\flat_O, L^\flat_O; \mathbb{Z}) = 0.$$

Proof. From Section 4.3.4 we know that this result is true additively, and that the 'reduction mod 2' map $HF^0(L^\flat_O, L^\flat_O; \mathbb{Z}) \to HF^0(L^\flat_O, L^\flat_O; \mathbb{Z}/(2))$ is an isomorphism (of rings). Now consider the restriction of the mod 2 closed–open map to the subalgebra of $QH^*(X_O; \mathbb{Z}/(2))$ generated by E. This ring is isomorphic to $\mathbb{Z}[E]/(2, E^2 + E + 1)$, which is the field of 4 elements, so $\mathcal{CO}^0$ must map it *injectively* to $HF^0(L^\flat_O, L^\flat_O; \mathbb{Z}/(2))$. By counting elements, this map must be surjective as well, so we deduce that $HF^0(L^\flat_O, L^\flat_O; \mathbb{Z}/(2))$, and hence $HF^0(L^\flat_O, L^\flat_O; \mathbb{Z})$, is also isomorphic to the field of 4 elements. □

4.4.4 Other relative spin structures

By Remark 4.1.2, when we switch to the non-standard relative spin structure we simply reverse the signs of the index 2 contributions to the pearl complex differential. For the triangle and tetrahedron this simply amounts to flipping the signs of X and Y, which doesn't affect the crucial quantities $3XY - 4\widehat{Z}$ and $4XY - 3Z$ respectively. For the octahedron and icosahedron the index 2 contributions all cancel out anyway, so again changing the relative spin structure has no effect on the cohomology.

This only applies to the additive structure on HF^*, but for $\triangle$ and T this completely determines the multiplication when working over $\mathbb{Z}$ (which is the only setting where we computed it). For O the ring structure arises from the relation $E^2 = E + 1$ in quantum cohomology, and this argument is unaltered by the change of relative spin structure.

Chapter 5

A family of $\mathrm{PSU}(n)$-homogeneous Lagrangians

In this chapter we introduce a new family of monotone Lagrangians in products of projective spaces, and use them to demonstrate the spectral sequence method of Chapter 2. Combining this with the disc analysis and $\mathcal{CO}^0$ computation from Chapter 3 we deduce some surprising properties. We then study some related examples, illustrating the use of some powerful techniques from the literature for which few explicit computations are available.

5.1 The main family

5.1.1 Constructing the family

Fix an integer $N \geq 3$ and consider the standard action of $\mathrm{SU}(N-1)$ on $\mathbb{C}^{N-1}$. This descends to a $\mathrm{PSU}(N-1)$-action on $\mathbb{CP}^{N-2}$ which, as discussed in Section 4.1.1, is Hamiltonian with moment map $\mu_{\mathbb{CP}^{N-2}} \colon \mathbb{CP}^{N-2} \to \mathfrak{psu}(N-1)^* = \mathfrak{su}(N-1)^*$ given by

$$\langle \mu_{\mathbb{CP}^{N-2}}([z]), A \rangle = -\frac{i}{2} \frac{z^\dagger A z}{z^\dagger z} \tag{5.1}$$

for all $z \in \mathbb{C}^{N-1}$ representing a point $[z] \in \mathbb{CP}^{N-2}$ and all $A \in \mathfrak{su}(N-1)$. Now take X to be the product $(\mathbb{CP}^{N-2})^N$, carrying the diagonal action of $\mathrm{PSU}(N-1)$, for which the moment map μ satisfies

$$\langle \mu([z_1], \dots, [z_N]), A \rangle = -\frac{i}{2} \sum_{j=1}^{N} \frac{z_j^\dagger A z_j}{z_j^\dagger z_j} \tag{5.2}$$

for all $([z_1], \dots, [z_N]) \in X$ and all A as above. Using the map $\mathfrak{u}(N-1) \to \mathfrak{su}(N-1)^*$ induced by the inner product $\langle A, B \rangle = \operatorname{Tr} A^\dagger B$ on $\mathfrak{u}(N-1)$, we can express μ as the map

$$Z \in (\mathbb{CP}^{N-2})^N \mapsto \frac{i}{2} Z Z^\dagger \in \mathfrak{u}(N-1), \tag{5.3}$$

where Z is an $(N-1) \times N$ matrix whose columns are the homogeneous coordinates of the components of the corresponding point in $(\mathbb{CP}^{N-2})^N$, scaled to have norm 1 (so we can ignore the denominators in (5.2)). The phases of the columns of Z are undetermined but do not affect the quantity $ZZ^\dagger$.

Consider the vectors $v_1, \ldots, v_N$ in $\mathbb{C}^{N-1}$ defined by

$$v_j = (\zeta^j, \zeta^{2j}, \ldots, \zeta^{(N-1)j}),$$

where $\zeta = e^{2\pi i/N}$ is a primitive Nth root of unity, and let x be the point $([v_1], \ldots, [v_N])$ in X. Note that one choice of matrix Z representing this point has components $Z_{jk} = \zeta^{jk}/\sqrt{N-1}$, so

$$(ZZ^\dagger)_{jk} = \frac{1}{N-1}\sum_{l=1}^{N} \zeta^{(j-k)l} = \frac{N}{N-1}\delta_{jk}.$$

In particular, $iZZ^\dagger/2$ is proportional to i times the identity matrix, which spans the orthogonal to $\mathfrak{su}(N-1)$ inside $\mathfrak{u}(N)$. This means that x lies in the zero level set of μ, and hence its orbit is an isotropic submanifold of X [31, Proposition 1.5].

If a matrix M in $\mathrm{SU}(N-1)$ stabilises x then each of the vectors $(\zeta^j, \zeta^{2j}, \ldots, \zeta^{(N-1)j}) \in \mathbb{C}^{N-1}$, for $j = 1, \ldots, N$, is an eigenvector of M. Since no pair of these vectors is orthogonal their eigenvalues must all coincide, and since they span $\mathbb{C}^{N-1}$ we deduce that M must be a scalar. In other words, the $\mathrm{PSU}(N-1)$-orbit of x is free. In particular, this orbit has dimension $(N-1)^2 - 1 = \dim_{\mathbb{C}} X$, so it is a Lagrangian submanifold L. Setting $K = \mathrm{PSU}(N-1)$, we have thus constructed a freely K-homogeneous (X, L) with $X = (\mathbb{CP}^{N-2})^N$. Moreover, by the Segre embedding of X in $\mathbb{CP}^{N^2-N-1}$ the homogeneity is linear (this gives an action of $\mathrm{PSU}(N-1)$ via $\mathrm{PSU}(N^2-N)$; in order to upgrade it to a non-projective representation, apply a suitable Veronese embedding as in Remark 3.1.26).

Let $H_j \in H^2(X; \mathbb{Z})$ denote the pullback of the hyperplane class on the jth $\mathbb{CP}^{N-2}$ factor of X. We then have $c_1(X) = (N-1)(H_1 + \cdots + H_N)$, so X is monotone with minimal Chern number $N_X = N-1$. Moreover $\pi_1(L, x) \cong \mathbb{Z}/(N-1)$ is finite so L is also monotone, with minimal Maslov number N_L at least 2 (using either orientability of L or the fact that $N_L \geq N_X/|\pi_1(L, x)|$).

5.1.2 Computing $\mathcal{CO}^0$

From the remainder of this section we make Assumption 3.2.6. We show in Remark 5.1.27 that this hypothesis is in fact unnecessary for our main results.

Let G denote the complexification $\mathrm{PSL}(N-1, \mathbb{C})$ of K, and for a proper subset $I \subset \{1, \ldots, N\}$ of size at least 2 let

$$Z_I = \Big\{([z_1], \ldots, [z_N]) \in X : (z_j)_{j \in I} \text{ is linearly dependent}\Big\}. \tag{5.4}$$

For $k = 1, \ldots N-2$ we shall use Z_k as shorthand for Z_I where I is $\{1, 2, \ldots, N-k\}$. Note that each Z_I is a K-invariant subvariety of X of complex codimension $N - |I|$, which is less than or equal to N_X^+. Our goal is to compute $\mathcal{CO}^0$ on the Poincaré dual classes, using Proposition 3.2.24, which requires us to count axial discs meeting the Z_I. Using the theory of obliging subvarieties from Section 3.3.3 and Section 3.3.5 we can reduce most of the problem to the following:

Lemma 5.1.1. *Each Z_I is obliging.*

Proof. Suppose I is a proper subset of $\{1, \ldots, N\}$ of size $N-k \geq 2$, and pick another subset J of size $k+1$ (which is also at least 2, since I is a *proper* subset) which intersects I in a singleton $\{m\}$. Consider the line l in X passing through

$$P = \left([v_1], \ldots, [v_{m-1}], \left[\sum_{j \in I \setminus \{m\}} v_j\right], [v_{m+1}], \ldots, [v_N]\right)$$

and

$$Q = \left([v_1], \ldots, [v_{m-1}], \left[\sum_{j \in J \setminus \{m\}} v_j\right], [v_{m+1}], \ldots, [v_N]\right).$$

Note that P and Q lie in Z_I and Z_J respectively, but every other point on l lies in the dense open orbit W. Therefore l has poles precisely at P and Q, and these are of indices at least $2k$ and $2(N-k-1)$ respectively, with equality if and only if Z_I and Z_J are obliging. On the other hand, l has Chern number $N-1$ and hence total index $2(N-1)$, so we do indeed have equality. $\square$

Corollary 5.1.2. *Each Z_I meets a unique family of index $2(N-|I|)$ discs and the evaluation map*

$$\mathrm{ev}_1 \colon \mathcal{M}^{Z_I}_{\mu=2(N-|I|)} \to L$$

has degree 1 (with the standard spin structure).

Proof. Since (X, L) is sharply linearly K-homogeneous, we can apply Corollary 3.3.43 in conjunction with Lemma 5.1.1 to deduce the existence and uniqueness of the family of discs, or equivalently to see that the moduli space $\mathcal{M}^{Z_I}_{\mu=2(N-|I|)}$ is a single K-orbit. Because ev_1 is K-equivariant and the homogeneity is free, the evaluation map must have degree ± 1, and the sign is pinned down by the calculations of Section B.5. $\square$

This tells us that with the standard spin structure and trivial local system we have

$$\mathcal{CO}^0(\mathrm{PD}(Z_I)) = 1_L$$

for each proper subset I of $\{1, \ldots, N\}$ of size at least 2. In order to make use of this computation, we need to understand the class Poincaré dual to Z_k:

Lemma 5.1.3. *For each $k = 1, \ldots, N-2$ we have*

$$\mathrm{PD}(Z_k) = h_k(H_1, \ldots, H_{N-k}),$$

where h_k denotes the complete homogeneous symmetric polynomial of degree k, i.e.

$$\sum_{\substack{r_1, \ldots, r_{N-k} \geq 0 \\ r_1 + \cdots + r_{N-k} = k}} H_1^{r_1} \cdots H_{N-k}^{r_{N-k}}.$$

(Recall that H_j denotes the pullback of the hyperplane class on the jth factor of X.)

Proof. Since $\mathrm{PD}(Z_k)$ lies in $H^{2k}(X; \mathbb{Z})$, it is a linear combination of monomials $H_1^{r_1} \cdots H_N^{r_N}$ with total degree $r_1 + \cdots + r_N$ equal to k. The coefficient of such a monomial is obtained by counting intersection points of Z_k with

$$\prod_{j=1}^{N} \Pi_j \subset \prod_{j=1}^{N} \mathbb{CP}^2 = X,$$

where Π_j is a generic r_j-plane in $\mathbb{CP}^{N-2}$ (here 'generic' means that the intersection takes place in the smooth locus of Z_k and is transverse). All such intersections count positively since all cycles involved carry complex orientations.

We therefore wish to count N-tuples of points in $\mathbb{CP}^{N-2}$ such that the jth point is constrained to an r_j-plane and the first $N-k$ points are linearly dependent, where $\sum r_j = k$. It is easy to see that if $r_1 + \cdots + r_{N-k} \leq k-1$ then the r_j-planes can be chosen so that

$$\Pi_j \cap \Big(\Pi_1 + \cdots + \widehat{\Pi}_j + \cdots + \Pi_{N-k}\Big) = \emptyset \tag{5.5}$$

for all j between 1 and $N-k$, where $\hat{}$ denotes omission, and so in this case there are no such N-tuples. When $r_1 + \cdots + r_{N-k} = k$, on the other hand, one can choose the planes so that the left-hand side of (5.5) is a single transversely cut out point for $j = 1, \ldots, N-k$, and so the coefficient of $H_1^{r_1} \ldots H_N^{r_N}$ is 1. This gives the claimed expression for $\mathrm{PD}(Z_k)$. $\square$

To explore the consequences of the resulting $\mathcal{CO}^0$ constraints in a fairly general setting we equip L with the standard spin structure s and an H_2^D local system $\mathscr{F}$ as discussed in Section 2.2.3—let $L^\flat = (L, s, \mathscr{F})$ be the obtained K-homogeneous monotone Lagrangian brane. Recall that such a local system corresponds to a homomorphism $\rho\colon H_2^D \subset H_2(X, L; \mathbb{Z}) \to R^\times$ to the multiplicative group of our coefficient ring R, so our next task is to identify a convenient description of the group $H_2(X, L; \mathbb{Z})$.

First note that the compactification divisor $Y \subset X$ is the union of the invariant subvarieties

$$Z_{\{1,\ldots,\widehat{j},\ldots,N\}} \tag{5.6}$$

from (5.4). By Lemma 5.1.1 each of these subvarieties is obliging, and hence by Corollary 3.3.43 each meets a unique family of index 2 holomorphic discs on L. Let $A_j \in H_2(X, L; \mathbb{Z})$ be the homology class of such a disc meeting the jth.

Lemma 5.1.4. *The A_j freely generate $H_2(X, L; \mathbb{Z})$.*

Proof. All (co)homology groups in this proof are over $\mathbb{Z}$. By pairing with the classes (5.6) in $H_2(X \setminus L)$ we obtain a map $H_2(X, L) \to \mathbb{Z}^N$ which exhibits the A_j as a free basis for a direct summand of $H_2(X, L)$. We also have the (surjective) normalised Maslov index homomorphism $\nu = \mu/2\colon H_2(X, L) \to \mathbb{Z}$, which tells us that $H_2(X, L) \cong \mathbb{Z} \oplus \ker \mu$. The reduction $\overline{\nu}$ of this homomorphism modulo $N-1$ annihilates the image of $H_2(X)$ and so factors as the boundary map $\partial\colon H_2(X, L) \to H_1(L)$ followed by a surjection (and thus isomorphism) $H_1(L) \to \mathbb{Z}/(N-1)$. We deduce that $\ker \overline{\nu} = \ker \partial$, and hence that $\ker \mu$ is contained in the image of $H_2(X)$, and thus is a subquotient of

$$\ker(\langle c_1(X), \cdot \rangle \colon H_2(X) \to \mathbb{Z}) \cong \mathbb{Z}^{N-1}.$$

Putting everything together, we see that $H_2(X, L)$ contains a subgroup freely generated by the A_j as a direct summand, but also that it is a direct sum of $\mathbb{Z}$ and a subquotient of $\mathbb{Z}^{N-1}$.

Considering ranks, we see that $H_2(X, L)$ must be isomorphic to $\mathbb{Z}^N$, and that the A_j form a basis. $\square$

With this in hand, for each j let $t_j \in R^\times$ denote the image of A_j under the homomorphism ρ corresponding to $\mathscr{F}$, and let

$$\eta_j = \mathcal{CO}^0(H_j) \in HF^*(L^\flat, L^\flat; R).$$

Let $\boldsymbol{\eta}$ denote the N-tuple of all η_j, and for a subset $I \subset \{1, \dots, N\}$ let $\boldsymbol{\eta}_I$ denote the tuple $(\eta_j)_{j \in I}$, so that

$$h_k(\boldsymbol{\eta}_I) = \sum_{\substack{j_1, \dots, j_k \in I \\ j_1 \leq \dots \leq j_k}} \eta_{j_1} \dots \eta_{j_k}$$

for each k. Using this notation we can now state:

Lemma 5.1.5. *For any non-empty subset $I \subset \{1, \dots, N\}$ we have*

$$h_{N-|I|}(\boldsymbol{\eta}_I) = \left(\prod_{j \notin I} t_j \right) \cdot 1_L$$

in $HF^*(L^\flat, L^\flat; R)$. The classes $\eta_j + t_j \cdot 1_L$ are all equal to $\eta_1 + \dots + \eta_N$. We denote this common value by $\overline{\eta}$.

Proof. By definition, the basis (A_j) for $H_2(X, L; \mathbb{Z})$ is dual to the basis for $H^2(X, L; \mathbb{Z}) \cong H_{2n-2}(X \setminus L; \mathbb{Z})$ given by the classes of the $Z_{\{1, \dots, \widehat{j}, \dots, N\}}$. Since the map $I \mapsto Z_I$ is inclusion-preserving, for any $k \in \{1, \dots, N-2\}$ the homology class of any holomorphic disc meeting Z_k has positive $A_1, \dots, A_k$-components. All other components are non-negative (by positivity of intersections) so if the disc has index $2k$ then its class must be exactly $A_1 + \dots + A_k$. In particular, this is the class of each member of the unique family of such discs whose existence is guaranteed by the fact that each Z_I is obliging.

Computing $\mathcal{CO}^0(\mathrm{PD}(Z_k))$ using Proposition 3.2.24, then applying Lemma 5.1.3, we therefore see that

$$h_k(\eta_1, \dots, \eta_{N-k}) = (t_{N-k+1} \dots t_N) \cdot 1_L.$$

Note that although the statement of Lemma 5.1.3 is in terms of the classical ring structure on the cohomology of X, the expressions are all of degree strictly less than the minimal Chern number of X so the classical and quantum products coincide. This proves the claimed equality for $I = \{1, \dots, N-k\}$ and clearly the same argument applies to any other tuple I of size $2, \dots, N-1$. Taking $|I| = N-1$ we see that the classes $\eta_j + t_j \cdot 1_L$ are all equal to the claimed quantity.

The case $|I| = N$ is vacuous, so it remains to deal with $|I| = 1$, for which we can use the quantum cohomology relations. The curve contributing to the quantum product H_j^{N-1} lies in the class of a line on the jth $\mathbb{CP}^{N-2}$ factor, which is expressed as

$$A_1 + \dots + \widehat{A_j} + \dots + A_N$$

in terms of our basis for $H_2(X, L; \mathbb{Z})$ (by intersecting with the subvarieties (5.6)). It is therefore weighted by $t_1 \cdots \widehat{t_j} \cdots t_n$ by $\mathscr{F}$, so in the deformed quantum cohomology ring we have

$$H_j^{N-1} = \left(\prod_{k \neq j} t_k \right) \cdot 1_X.$$

Applying $\mathcal{CO}^0$ gives exactly what we want. $\square$

Remark 5.1.6. If $\mathscr{F}$ also incorporates twisting by a higher rank local system $\mathscr{F}'$, in the ordinary sense of Section 2.1.2, then each t_j should carry with it a factor of $[1_L \otimes M]$, where M is the automorphism of the fibre $\mathscr{F}_x = \mathscr{F}'_x$ describing the monodromy of $\mathscr{F}'$ around the boundary of an index 2 disc. //

5.1.3 Analysing the constraints

It will be helpful for us to transform the equalities computed in the preceding subsection into a more digestible form, via the yoga of symmetric polynomials:

Lemma 5.1.7. *Treating the η_j, $\overline{\eta}$ and t_j as formal variables over $\mathbb{Z}$, the following are equivalent:*

(i) The equalities from Lemma 5.1.5.

(ii) $\eta_j + t_j = \overline{\eta}$ for all j and $h_k(\boldsymbol{\eta}) = \overline{\eta}^k$ for $k = 1, \dots, N-1$.

(iii) $\eta_j + t_j = \overline{\eta}$ for all j, $e_1(\boldsymbol{\eta}) = \overline{\eta}$, and $e_k(\boldsymbol{\eta}) = 0$ for $k = 2, \dots, N-1$, where the e_j are the elementary symmetric polynomials.

Proof. We assume throughout that $\eta_j + t_j = \overline{\eta}$ for all j as this is entailed by all three sets of equalities. Note that for any k in $\{1, \dots, N-1\}$ and any subset $I \subset \{1, \dots, N\}$ of size $N-k$ we have

$$h_k(\boldsymbol{\eta}) = h_k(\boldsymbol{\eta}_I) + \sum_{j_1 \notin I} \eta_{j_1} h_{k-1}(\boldsymbol{\eta}_{I \cup \{j_1\}}) + \sum_{\substack{j_1, j_2 \notin I \\ j_1 < j_2}} \eta_{j_1} \eta_{j_2} h_{k-2}(\boldsymbol{\eta}_{I \cup \{j_1, j_2\}}) + \dots.$$

Here the first term on the right-hand side contains precisely the monomials from the left-hand side which only involve variables t_j for $j \in I$, the second term (the sum over j_1) contains those monomials involving exactly one variable outside this set, the third term those monomials involving exactly two variables outside the set, and so on. An easy induction on k using this decomposition shows that (i) and (ii) are equivalent.

Now use the fact that

$$\sum_{r \geq 0} h_r(\boldsymbol{\eta}) z^r = \prod_j \frac{1}{1 - \eta_j z} = \frac{1}{\sum_{r \geq 0} e_r(\boldsymbol{\eta}) z^r}$$

as power series in z. Both (ii) and (iii) are equivalent to this quantity being $1/(1 - \overline{\eta} z) + O(z^N)$. $\square$

Remark 5.1.8. The closed–open map for Lagrangians invariant under a loop γ of Hamiltonian diffeomorphisms was studied by Charette–Cornea in [28] and more recently by Tonkonog in

[137]. If $S(\gamma)$ denotes the Seidel element in $QH^*(X)$ defined by γ, they showed that without local systems $\mathcal{CO}^0(S(\gamma))$ is equal to $\pm 1_L$, where the sign depends on the choice of spin structure.

Taking γ to be the action of the one-parameter subgroup (g_θ) of $\mathrm{PSU}(N-1)$ comprising diagonal matrices with diagonal entries $e^{-i\theta/(N-1)}(e^{i\theta}, 1, \ldots, 1)$, and applying a result of McDuff–Tolman [105, Theorem 1.10] to each $\mathbb{CP}^{N-2}$ factor, we obtain $S(\gamma) = H_1 \ldots H_N$. Their result then yields $e_N(\boldsymbol{\eta}) = \pm 1_L$. On the other hand, from Lemma 5.1.7 we see that with the standard spin structure

$$\left(\prod_{j=1}^{N} t_j\right) \cdot 1_L = \prod_{j=1}^{N}(\overline{\eta} - \eta_j) = \sum_{j=0}^{N}(-1)^j e_j(\boldsymbol{\eta})\overline{\eta}^{N-j} = (-1)^N e_N(\boldsymbol{\eta}).$$

Setting the t_j to 1, which corresponds to choosing the trivial local system (with no H_2^D local system), we obtain $e_N(\boldsymbol{\eta}) = (-1)^N \cdot 1_L$.

The sign $(-1)^N$ can be understood by comparing spin structures, as follows. Let $\overline{\gamma}$ denote an orbit of γ on L. A choice of spin structure on L gives rise to a homotopy class of trivialisation of $\overline{\gamma}^*TL$, and Tonkonog shows that the sign of $\mathcal{CO}^0(S(\gamma))$ is $+1$ if and only if this homotopy class agrees with that defined by transport by γ. In our case, this corresponds to the identification

$$T_{\overline{\gamma}(\theta)}L = (g_\theta)_* T_{\overline{\gamma}(0)}L = g_\theta \mathfrak{su}(N-1) \cdot \overline{\gamma}(0)$$

for all θ. The standard spin structure, meanwhile, uses the identification

$$T_{\overline{\gamma}(\theta)}L = \mathfrak{su}(N-1) \cdot \overline{\gamma}(\theta) = \mathfrak{su}(N-1) g_\theta \cdot \overline{\gamma}(0).$$

The two trivialisations therefore differ by the loop $\widetilde{\gamma}$ in $\mathrm{SL}(\mathfrak{su}(N-1))$ defined by the adjoint action of (g_θ).

Under this action $\mathfrak{su}(N-1)$ decomposes as $N-2$ two-dimensional subspaces which g_θ rotates by angle θ, and a complement which is fixed. The loop $\widetilde{\gamma}$ therefore represents $N-2$ times the generator of $\pi_1(\mathrm{SL}(\mathfrak{su}(N-1))) \cong \mathbb{Z}/2$. This means that the two homotopy classes of trivialisation coincide if and only if N is even: hence the factor of $(-1)^N$.

Applying the same ideas of Charette–Cornea, Tonkonog and McDuff–Tolman to circle subgroups of the torus acting on a Fano toric variety, one sees that $\mathcal{CO}^0$ maps each component of the toric divisor to $\pm 1_L$ in the self-Floer cohomology of the monotone fibre. To compute the sign in the standard spin structure one again has to consider the adjoint action of the circle subgroup, but since the torus is abelian this action is trivial. This means the sign is always $+1$, in agreement with Proposition 3.2.24. $/\!\!/$

We now reach the punchline:

Proposition 5.1.9. *Let R be a ring of characteristic p (prime or 0), s' a relative spin structure on L, and $\mathscr{F}'$ a local system (in the usual sense, possibly of higher rank—see Remark 2.1.2). The monotone Lagrangian brane $(L, s', \mathscr{F}')$ has vanishing self-Floer cohomology over R unless:*

(i) p is prime, N is a power of p, and s' has signature $(N, 0)$ or $(0, N)$ if $p \neq 2$.

(ii) p is prime, N is twice a power of p, and s' has signature $(N/2, N/2)$ if $p \neq 2$.

(iii) $N = 3$*,* $p = 5$ *and* s' *has signature* $(2, 1)$ *or* $(1, 2)$*.*

Here the signature of s' *is defined by writing the difference between* s' *and the standard spin structure* s *(see Definition 3.1.10) as an element* $\varepsilon \in H^2(X, L; \mathbb{Z}/2)$ *and counting the number of basic disc classes* A_j *on which* ε *takes the value* 0*; the signature is then*

$$(\textit{number of zeros}, \textit{number of ones}).$$

Note that the relative spin structure, and hence signature, is irrelevant in characteristic 2*.*

Proof. Assume the self-Floer cohomology is non-vanishing, and let $\varepsilon \in H^2(X, L; \mathbb{Z}/2)$ be as above. By replacing $\mathscr{F}'$ by a composite local system $\mathscr{F}$ which incorporates both the H_2^D local system $A_j \mapsto (-1)^{\langle \varepsilon, A_j \rangle}$ and the (possibly higher rank) twisting of $\mathscr{F}'$, we may reduce to the case of the brane $L^\flat = (L, s, \mathscr{F})$ considered in the preceding analysis, with $t_j = (-1)^{\langle \varepsilon, A_j \rangle}$ for each j. In other words, we may assume that we are working with the standard spin structure, at the cost of introducing an H_2^D local system. By reordering the $\mathbb{CP}^{N-2}$ factors, we may assume without loss of generality that $t_1 = \cdots = t_a = 1$ and $t_{a+1} = \cdots = t_N = -1$. Let $b = N - a$ (so the signature is (a, b)), let M denote the monodromy of $\mathscr{F}'$ as in Remark 5.1.6, and let M' denote the class $[1_L \otimes M]$ in $HF^*(L^\flat, L^\flat; R)$.

Suppose first that $N \geq 4$, $p \neq 2$ and $a, b > 0$. From Lemma 5.1.5, via Lemma 5.1.7(iii), we obtain (from the e_1 condition)

$$a(\overline{\eta} - M') + b(\overline{\eta} + M') = \overline{\eta} \tag{5.7}$$

and, for k in $\{2, \dots, N-1\}$,

$$\sum_{j=0}^{k} \binom{a}{j} \binom{b}{k-j} (\overline{\eta} - M')^j (\overline{\eta} + M')^{k-j} = 0. \tag{5.8}$$

These are all equalities in $HF^*(L^\flat, L^\flat; R)$. Note that the quantities $\overline{\eta} - M'$ and $\overline{\eta} + M'$ commute since they are the images under $\mathcal{CO}^0$ of the commuting variables H_1 and H_{a+1} in $QH^*(X)$. Since $p \neq 2$ this forces $\overline{\eta}$ and M' to commute.

The equations $a + b = N$ and (5.7) give

$$aM' = \frac{NM' + (N-1)\overline{\eta}}{2} \quad \text{and} \quad bM' = \frac{NM' - (N-1)\overline{\eta}}{2}. \tag{5.9}$$

Substituting these into (5.8) with $k = 2$ and $k = 3$ (using the fact that $N \geq 4$), and simplifying, we get

$$\overline{\eta}^2(N-1) = NM'^2 \tag{5.10}$$

and

$$\overline{\eta}\big(2M'^2 - \overline{\eta}^2(N-1) + NM'^2\big) = 0. \tag{5.11}$$

Now substituting (5.10) into (5.11), and using the fact that M' is invertible (its inverse is $[1_L \otimes M^{-1}]$), we conclude that $\overline{\eta} = 0$. Plugging this back into (5.10) and (5.9) we see that $N \cdot 1_L$, $a \cdot 1_L$ and $b \cdot 1_L$ vanish in the subring of $HF^*(L^\flat, L^\flat; R)$ generated by the unit, which is isomorphic to $\mathbb{Z}/(p)$. This means $N \equiv a \equiv b \equiv 0$, where $\equiv$ denotes equality modulo p. In particular, p must be prime, not 0.

We can now express the equations (5.8) as

$$\sum_{j=0}^{k} \binom{a}{j} \binom{b}{k-j} (-1)^j \equiv 0 \tag{5.12}$$

for k in $\{2, \ldots, N-1\}$. Since we are assuming $a, b > 0$ we can write $a = p^\alpha A$ and $b = p^\beta B$ for unique positive integers α, β, A and B, with A and B not divisible by p. Suppose for contradiction that a and b are not both powers of p, so A and B are not both 1. This means that $p^\alpha + p^\beta$ lies in $\{2, \ldots, N-1\}$, so we can take $k = p^\alpha + p^\beta$ in (5.12). All terms on the left-hand side are divisible by p except for that with $j = p^\alpha$, which is not, so we get the desired contradiction and conclude that a and b are indeed both powers of p.

Now suppose for contradiction that $a \neq b$ and take $k = \min(a, b)$ in (5.12). All terms on the left-hand side are divisible by p except the first (if $a > b$) or the last (if $b > a$), so again we obtain a contradiction and now conclude that $a = b$. Putting everything together, we have shown that p is prime and that $N = a + b = 2a$ is twice a power of p.

Next we deal with the case $N \geq 4$, $p \neq 2$, with $b = 0$; the $a = 0$ case is analogous. The argument proceeds exactly as above until (5.12), which now reads

$$\binom{a}{k} \equiv 0 \tag{5.13}$$

for all k in $\{2, \ldots, N-1\}$. Again we write $a = p^\alpha A$, and this time $k = p^\alpha$ gives a contradiction unless $A = 1$ and hence $N = a$ is a power of p.

Now consider the case $p = 2$ with $N \geq 3$. Since here there is no distinction between 1 and -1 we may assume that $b = 0$. From (5.7) we see that $(N-1)\overline{\eta} = NM'$, which forces N to be even and hence $\overline{\eta}$ to vanish (otherwise the left-hand side is zero but the right-hand side is invertible). This turns (5.8) into (5.13), and the argument is completed as above.

Finally we turn to the exceptional case $N = 3$ and $p \neq 2$. The equations (5.7) and (5.10) are still valid, and they give $2\overline{\eta} = (a-b)M'$ and $2\overline{\eta}^2 = 3M'^2$ respectively. Combining these we get $(a-b)^2 \equiv 6$, and the left-hand side can only take the values 9 and 1, so a solution is only possible if p is 3 or 5. $\square$

We shall see later, in Corollary 5.1.20 and then in Remark 5.2.3 and Proposition 5.2.10, that this result is sharp in terms of the values of p, N and s' for which the self-Floer cohomology must vanish.

Question 5.1.10. Can the anomalous appearance of characteristic 5 for $N = 3$ be related to the significance of the same characteristic for the Chiang Lagrangian (see Section 5.2.1 for its relation with our family)? //

Corollary 5.1.11. *Over any coefficient ring R, the self-Floer cohomology of L with any relative spin structure and local system vanishes unless N is a prime power or twice a prime power.*

Proof. If the self-Floer cohomology is non-zero then the subring generated by the unit is of the form $\mathbb{Z}/(m)$ for some non-negative integer m. Take a prime p dividing m, replace R by $R \otimes_{\mathbb{Z}} \mathbb{Z}/(p)$ and apply Proposition 5.1.9. □

Question 5.1.12. Must the self-Floer cohomology vanish if we allow L to be equipped with a general bounding cochain? //

5.1.4 Non-displaceability and wideness

Having seen that the $\mathrm{PSU}(N-1)$-homogeneous Lagrangian $L \subset X = (\mathbb{CP}^{N-2})^N$ is narrow in many circumstances, we now use the methods of Section 2.2 to show that for the trivial local system L is in fact wide in all but one of the situations not ruled out by Proposition 5.1.9. We shall also prove that L is non-displaceable for all values of N, and, moreover, that it is wide over $\mathbb{C}$ after turning on an appropriate B-field.

The first observation is that $\mathrm{Symp}(X, L)$ contains an obvious subgroup isomorphic to S_N, the symmetric group on N objects, which acts by permuting the factors of X (this action preserves L as it preserves the moment map (5.2), of which L is the zero set). In a slight abuse of notation we will refer to this subgroup simply as S_N.

Lemma 5.1.13. *S_N acts on $H_2(X, L; \mathbb{Z})$ by permuting the A_j. Its action on L is trivial up to isotopy, and hence is trivial on $H^*(L; \mathbb{Z})$ and preserves the standard spin structure.*

Proof. The action on $H_2(X, L; \mathbb{Z})$ can be easily computed by considering the action on the dual subvarieties (5.6). For the second statement it suffices to show that the transposition of two factors is isotopic to the identity on L, and without loss of generality we may assume they are the first two.

Recall the vectors $v_j = (\zeta^j, \zeta^{2j}, \dots, \zeta^{(N-1)j})$ from Section 5.1.1, where ζ is a primitive Nth root of unity. These are the components of our base point $x = ([v_1], \dots, [v_N])$. It is easy to compute that the inner product $\langle v_j, v_k \rangle$ is $(N-1)$ if $j = k$ and -1 otherwise, so v_1 and v_2 have equal projections u to the span U of $v_3, \dots, v_N$, and equal length. We can therefore write $v_1 = u + w$ and $v_2 = u + e^{i\alpha} w$ for some w in $U^\perp$ and some phase $e^{i\alpha}$. In fact, since the v_j sum to zero we must have

$$u = -\frac{1}{2}(v_3 + \dots + v_N)$$

and $e^{i\alpha} = -1$. Let R_θ be the one-parameter subgroup of $\mathrm{U}(N-1)$ which fixes U and rotates w by $e^{i\theta}$. Note that R_π acts on the v_j precisely by transposing v_1 and v_2.

Now consider the family of diffeomorphisms f_θ of L, parametrised by $\theta \in [0, \pi]$, defined by

$$f_\theta(Ax) = AR_\theta x$$

for all A in $\mathrm{PSU}(N-1)$. Clearly $f_0 = \mathrm{id}_L$ and f_π acts by transposing the first two components, so this is the desired isotopy. (With some more thought, one can see that the orbits of $(f_\theta)_{\theta=0}^{2\pi}$ are actually the boundaries of the index $2(N-1)$ discs meeting Z_{N-2}.) □

Since the A_j freely generate $H_2(X, L; \mathbb{Z})$ the group ring $R[H_2^D]$ is simply $R[T_1^{\pm 1}, \dots, T_N^{\pm 1}]$, where we identify T^{A_j} with T_j, and the extended Novikov ring $\Lambda_R^\dagger$ from Section 2.2.2 is the subring generated by 1 and the monomials of strictly positive total degree (so, for example,

$T_1T_2^{-1}$ is not allowed but $T_1^2T_2^{-1}$ is). For our purposes there is no benefit in using $\Lambda_R^\dagger$ over $R[H_2^D]$ so for simplicity we'll stick to the latter. S_N acts by permuting the T_j.

Let P_R denote the polynomial ring $R[T_1, \dots, T_N]$ and let $e_1, \dots, e_N$ be the elementary symmetric polynomials in the T_j. Let Λ'_R denote $R[H_2^D]/(e_1, \dots, e_{N-1})$, and let F^0 be the image of P_R in Λ'_R. Setting $F^r = e_N^r F^0$ for $r \in \mathbb{Z}$ we obtain a descending filtration of Λ'_R for which each quotient F^r/F^{r+1} is isomorphic (via multiplication by e_N^{-r}) as a graded $R[S_N]$-module to $F^0/F^1[-2Nr]$, where $[-2Nr]$ denotes the grading shift (all ideals and submodules involved are homogeneous, so everything inherits a natural grading—see Remark 5.1.25 for what goes wrong when this isn't the case).

Lemma 5.1.14. *The kernel of the surjection $P_R \to F^0/F^1$ is the ideal generated by $e_1, \dots, e_N$.*

Proof. The kernel is precisely the intersection in $R[H_2^D]$ of the subring P_R with $I + e_N P_R$, where I is the ideal (of $R[H_2^D]$) generated by $e_1, \dots, e_{N-1}$. It therefore suffices to show that $I \cap P_R$ is the ideal of P_R generated by $e_1, \dots, e_{N-1}$.

It is well-known that $e_1, \dots, e_N$ are algebraically independent over R and that P_R is a free $R[e_1, \dots, e_N]$-module of rank $N!$—let B be a free basis. Note that $R[H_2^D]$ can be obtained from P_R by adjoining an inverse to e_N (since we can express T_1^{-1} as $T_2 \dots T_N e_N^{-1}$, with similar expressions for the other T_j^{-1}), so B also forms a free basis for $R[H_2^D]$ as an $R[e_1, \dots, e_{N-1}, e_N^{\pm 1}]$-module.

We thus have that $R[H_2^D]$ is a free R-module with basis $\{e_1^{j_1} \dots e_N^{j_N} b\}$, where $j_1, \dots, j_{N-1}$ are non-negative integers, j_N is any integer, and b is an element of B. The ideal I is the R-linear span of those elements with j_l strictly positive for some l in $\{1, \dots, N-1\}$, whilst P_R is the span of those elements with $j_N \geq 0$. Their intersection is thus the span of those elements for which some j_l is positive *and* $j_N \geq 0$, which is precisely the ideal of P_R generated by $e_1, \dots, e_{N-1}$. $\square$

We therefore have an isomorphism of graded $R[S_N]$-modules $F^0/F^1 \cong P_R/(e_1, \dots, e_N)$. If R is a field of characteristic coprime to $N!$ then the group ring $R[S_N]$ is semisimple, and from our filtration we conclude that there is an isomorphism of graded $R[S_N]$-modules

$$\Lambda'_R \cong \bigoplus_{r \in \mathbb{Z}} P_R/(e_1, \dots, e_N)[-2Nr]. \tag{5.14}$$

We would like to understand the S_N-action on $\Lambda'_{\mathbb{Q}}$, so in light of this decomposition it suffices to consider the quotient $P_{\mathbb{Q}}/(e_1, \dots, e_N) = \mathbb{Q}[T_1, \dots, T_N]/(e_1, \dots, e_N)$. The identification of this representation is standard (see, e.g. [100, Corollary 2.5.8] for the statement over $\mathbb{C}$), and can be seen as a piece of straight algebra, but there is also a beautiful geometric argument, explained to the author by Oscar Randal-Williams:

Lemma 5.1.15. *Over any field R of characteristic coprime to $N!$ the representation*

$$R[T_1, \dots, T_N]/(e_1, \dots, e_N)$$

of S_N is isomorphic to the regular representation $R[S_N]$.

Proof. The ring $\mathbb{Z}[T_1, \dots, T_N]/(e_1, \dots, e_N)$ is well-known to be isomorphic to the cohomology of the variety F of complete flags in $\mathbb{C}^N$, which is diffeomorphic to the unitary group $\mathrm{U}(N)$ modulo the group T^N of diagonal matrices. This can be seen as follows. For $r = 0, \dots, N$ let F_r denote the space of 'first r columns of elements of $\mathrm{U}(N)/T^N$'. We obtain an obvious chain of fibrations

$$F = F_N \twoheadrightarrow F_{N-1} \twoheadrightarrow \cdots \twoheadrightarrow F_1 \twoheadrightarrow F_0 = *$$

with fibres $\mathbb{CP}^1, \dots, \mathbb{CP}^{N-1}$ respectively, and applying the Serre spectral sequence to each step we see that the cohomology group $H^*(F;\mathbb{Z})$ is free of rank $N!$. Moreover, if $\mathcal{L}_r$ denotes the tautological complex line bundle over F whose fibre over a matrix $Q \in \mathrm{U}(N)/T^N$ is the span of the rth column of Q inside $\mathbb{C}^N$, then the spectral sequence argument shows that in fact the first Chern classes $T_1, \dots, T_N$ of the $\mathcal{L}_r$ generate the cohomology as a ring. The sum of these bundles is trivial, so $\prod_j (1 + T_j) = 1$ and hence $H^*(F;\mathbb{Z})$ is a quotient of $\mathbb{Z}[T_1, \dots, T_N]/(e_1, \dots, e_N)$. Since both are free of rank $N!$ we deduce that in fact they are equal.

From this perspective we also see that the S_N-action on $\mathbb{Z}[T_1, \dots, T_N]/(e_1, \dots, e_N)$ corresponds to the action on $H^*(F;\mathbb{Z})$ induced by permutation of the columns of matrices in $\mathrm{U}(N)/T^N$. This action on F is clearly free, so there is a smooth quotient manifold F/S_N. Lifting a cell structure from this quotient, over R we obtain a cellular cochain complex of free $R[S_N]$-modules which computes $H^*(F;R)$ as a representation of S_N.

Since R has characteristic coprime to $|S_N|$, the ring $R[S_N]$ is semisimple, so it makes sense to talk about the Euler characteristic of a complex of $R[S_N]$-modules: this is a tuple of multiplicities indexed by isomorphism classes of simple modules. For our complex this Euler characteristic is manifestly an integer multiple of the multiplicities occurring in the regular representation. On the other hand, the cohomology is concentrated in even degrees, so the Euler characteristic is just the multiplicities occurring in

$$H^*(F;R) \cong R[T_1, \dots, T_N]/(e_1, \dots, e_N).$$

Therefore the latter is isomorphic to a direct sum of copies of $R[S_N]$, and counting dimensions we see that the number of copies is exactly 1. □

Combining this with (5.14) we deduce:

Corollary 5.1.16. *The S_N-invariant subring of $\Lambda'_{\mathbb{Q}}$ is spanned by elements in degrees $2N\mathbb{Z}$.*

We also need to understand the cohomology ring of $L \cong \mathrm{PSU}(N-1)$. The standard reference is Borel's paper [21], but we shall follow Baum and Browder [10]. The crucial idea is to replace the quotient $\mathrm{PSU}(N-1) = \mathrm{U}(N-1)/S^1$ by the homotopy quotient $\mathrm{U}(N-1) \times_{S^1} ES^1$, which is a principal $U(N-1)$-bundle over $BS^1 = \mathbb{CP}^\infty$, so that we can employ the Serre spectral sequence. The second page

$$E_2 \cong H^*(\mathbb{CP}^\infty; H^*(\mathrm{U}(N-1);\mathbb{Z}))$$

is isomorphic as a ring to $\mathbb{Z}[b] \otimes \Lambda_{\mathbb{Z}}(a_1, a_3, \dots, a_{2N-3})$, where b has degree 2 and a_{2j-1} has degree $2j-1$, and the generators a_{2j-1} of $H^*(\mathrm{U}(N-1);\mathbb{Z})$ can be chosen so that each transgresses to the Chern class c_j of the associated $\mathbb{C}^{N-1}$-bundle over $\mathbb{CP}^\infty$ (this can be seen by viewing the bundle as a pullback of the universal bundle over $B\mathrm{U}(N-1)$ and using naturality of the Serre spectral sequence). These Chern classes were computed in [10, Section 4] to be

$$c_j = \binom{N-1}{j} b^j.$$

There are two important points to take from this discussion. The first is that each page is generated as a ring by elements of degree at most $2N-3$, so the same is true of $H^*(\mathrm{PSU}(N-1);\mathbb{Z})$. The second is that if p is a prime not dividing $N-1$ then in characteristic p the class b is hit by the E_2 differential, from which we see that

$$\dim_{\mathbb{Z}/(p)} H^*(\mathrm{PSU}(N-1);\mathbb{Z}/(p)) = 2^{N-2} = \dim_{\mathbb{Q}} H^*(\mathrm{PSU}(N-1);\mathbb{Q}),$$

and hence that $H^*(\mathrm{PSU}(N-1);\mathbb{Z})$ has no p-torsion.

We can now prove the key result:

Proposition 5.1.17. *Equipping L with the standard spin structure s and trivial local system, all differentials in the Oh spectral sequence over $\Lambda'_{\mathbb{Q}}$ (starting from E_1) vanish. The same holds with $\mathbb{Q}$ replaced by any field R of characteristic p (prime or 0) not dividing $N-1$.*

Proof. Consider the first page $H^*(L;\mathbb{Q}) \otimes \Lambda'_{\mathbb{Q}}$ of the spectral sequence over $\Lambda'_{\mathbb{Q}}$. Since S_N acts trivially on $H^*(L;\mathbb{Q})$ and on the spin structure (and local system, which is trivial), its action on this page is purely on $\Lambda'_{\mathbb{Q}}$. It therefore acts trivially on the subset $H^*(L;\mathbb{Q}) \otimes 1$ of the zeroth column, but by Corollary 5.1.16 it has no non-zero fixed points in columns 1 to $N-1$. This forces the differentials $\mathrm{d}_1, \dots, \mathrm{d}_{N-1}$ to vanish on this subset, and hence on the whole of the first $N-1$ pages respectively. All later differentials vanish by an argument analogous to Proposition 2.2.2 since $H^*(L;\mathbb{Q})$ is generated in degrees less than $2N-1$.

Letting $L^\flat$ denote our Lagrangian brane, we deduce that $HF^*(L^\flat, L^\flat; \Lambda'_{\mathbb{Q}}) \cong H^*(L;\Lambda'_{\mathbb{Q}})$ as graded $\mathbb{Q}$-vector spaces. Since $\Lambda'_{\mathbb{Q}}$ is finite-dimensional in each degree (it has rank $N!$ over $\mathbb{Q}[e_N^{\pm 1}]$), we can count dimensions and use the universal coefficient theorem (passing via $\Lambda'_{\mathbb{Z}}$) to see that

$$\dim_R HF^j(L^\flat, L^\flat; \Lambda'_R) \geq \dim_{\mathbb{Q}} HF^j(L^\flat, L^\flat; \Lambda'_{\mathbb{Q}}) = \dim_{\mathbb{Q}} H^j(L;\Lambda'_{\mathbb{Q}}) = \dim_R H^j(L;\Lambda'_R)$$

for all j. The final equality uses the fact that $H^*(L;\mathbb{Z})$ has no p-torsion. Considering the spectral sequence for $L^\flat$ over Λ'_R we see that all differentials must vanish (and for each j the above inequality is in fact an equality). $\square$

Using this we can deduce wideness in many cases:

Proposition 5.1.18. *Suppose R is a field of characteristic p (prime or 0) not dividing $N-1$, s is the standard spin structure on L, and $\mathscr{F}$ is an H_2^D local system with corresponding homomorphism $\rho\colon H_2^D \to R^\times$ given by $\rho\colon A_j \mapsto t_j$ for $j = 1, \dots, N$. If the elementary symmetric polynomials $e_k(t_1, \dots, t_N) \in R$ vanish for $k = 1, \dots, N-1$ then $L^\flat = (L, s, \mathscr{F})$ is wide over R.*

Proof. Recall from Section 2.2.3 that ρ induces a ring homomorphism $\hat{\rho}\colon R[H_2^D] \to \Lambda_R$ given by $T^A \mapsto \rho(A) T^{\nu(A)}$, which amounts to $T_j \mapsto t_j T$ in our case, and that the pearl complex twisted by the H_2^D local system $\mathscr{F}$ is obtained from the untwisted complex over $R[H_2^D]$ by tensoring with

Λ_R viewed as an $R[H_2^D]$-module via $\hat{\rho}$. The assumption $e_k(t_1, \dots, t_N) = 0$ for $k = 1, \dots, N-1$ means that $\hat{\rho}$ factors through Λ'_R, so in fact we can construct the complex twisted by $\mathscr{F}$ from the untwisted complex over Λ'_R—instead of $R[H_2^D]$—by tensoring with Λ_R. The differentials in the Oh spectral sequence E'_* over Λ'_R all vanish by Proposition 5.1.17, whilst the first page of the spectral sequence $E_*^{\mathscr{F}}$ twisted by $\mathscr{F}$ is given by tensoring E'_1 with Λ_R (via $\hat{\rho}$ still), and the differentials are the induced maps. The differentials in $E_*^{\mathscr{F}}$ must therefore also all vanish, proving the claimed wideness. $\square$

Remark 5.1.19. (i) As in the proof of Proposition 5.1.9 we can transform any pair $(s', \mathscr{F}')$ comprising a relative spin structure and H_2^D local system into one of the form $(s, \mathscr{F})$, where s is the standard spin structure.

(ii) The condition on the $e_j(t_1, \dots, t_N)$ can be expressed as the equality

$$\prod_j (z + t_j) = z^N + \prod_j t_j$$

in the polynomial ring $R[z]$. //

Corollary 5.1.20. *Suppose N, p, and the relative spin structure s' on L are of the form described in Proposition 5.1.9(i) or (ii). Then (L, s') equipped with the trivial local system is wide over any field R of characteristic p.*

Proof. We interpret s' as the standard spin structure s modified by the H_2^D local system sending A_j to t_j with $t_1 = \dots = t_a = 1$ and $t_{a+1} = \dots = t_N = -1$. Let $b = N - a$ so the signature of s' is (a, b). By Proposition 5.1.18 it suffices to prove that $e_k(t_1, \dots, t_N) = 0$ for $k = 1, \dots, N-1$, which by Remark 5.1.19(ii) is equivalent to

$$(z+1)^a (z-1)^b = z^N + (-1)^b$$

in $R[z]$. This is easily seen to hold in the cases of interest, using the fact that $(x+y)^r = x^r + y^r$ in characteristic p when r is a power of p. $\square$

The case of Proposition 5.1.9(iii) will be dealt with in Remark 5.2.3 and Proposition 5.2.10.

Question 5.1.21. In the prime power case of Corollary 5.1.20 each of the wide relative spin structures is invariant under the action of S_N, so there is an induced action on the ring $HF^*(L^\flat, L^\flat; R)$. When R is a field of characteristic p (the prime of which N is a power), we obtain a 2^{N-2}-dimensional representation of S_N over R; what is this representation? From the spectral sequence we see that it carries a filtration such that the associated graded representation is trivial, but since the characteristic of our field divides the order of the group the representation need not be semisimple.

When N is twice a prime power each wide relative spin structure is no longer invariant under the whole of S_N. Instead its stabiliser is a subgroup $S_{N/2} \times S_{N/2}$, and again it is natural to wonder whether the corresponding representation on Floer cohomology is trivial. //

Corollary 5.1.22. *For any $N \geq 3$, any complex number λ, and any permutation $\sigma \in S_N$, L is wide over $\mathbb{C}$ when equipped with the standard spin structure and the H_2^D local system which* *maps A_j to $e^{2\pi i(\lambda+\sigma(j)/N)}$. In particular, L is non-displaceable and supports a non-zero object in the monotone Fukaya category of X deformed by any B-field representing a class of the form*

$$\sum_{j=1}^{N}\left(\lambda+\frac{\sigma(j)}{N}+\mathbb{Z}\right)H_j \in H^2(X;\mathbb{C})/H^2(X;\mathbb{Z}).$$

Proof. By Proposition 5.1.18 and Remark 5.1.19(ii), the first claim can be reduced to showing that

$$\prod_{j=1}^{N}\left(z+e^{2\pi i(\lambda+\sigma(j)/N)}\right)=z^N+\prod_{j=1}^{N}e^{2\pi i(\lambda+\sigma(j)/N)}$$

in $\mathbb{C}[z]$, and both sides are readily seen to be equal to $z^N-(-e^{2\pi i\lambda})^N$. Non-displaceability then follows from the discussion in Section 2.2.3, and it is just left to compute the class of a background B-field for our H_2^D local system (since $H_2(X;\mathbb{Z})$ is torsion-free, by Proposition 2.2.18 any H_2^D-local system over $\mathbb{C}$ can be realised using a B-field). To do this we recall from Lemma 5.1.5 that the class of a line on the jth $\mathbb{CP}^{N-2}$ factor decomposes as

$$A_1+\cdots+\widehat{A_j}+\cdots+A_N,$$

so is given weight $e^{2\pi i\theta_j}$ by our H_2^D local system, where

$$\theta_j=(N-1)\lambda+\frac{\sigma(1)+\cdots+\widehat{\sigma(j)}+\ldots\sigma(N)}{N}=(N-1)\lambda+\frac{N+1}{2}-\frac{\sigma(j)}{N}.$$

Letting $\lambda'=-(N-1)\lambda-(N+1)/2\in\mathbb{C}$, we have $\theta_j=-(\lambda'+\sigma(j)/N)$ for all j, and hence our B-field is in class $\sum_j(\lambda'+\sigma(j)/N+\mathbb{Z})H_j$. Clearly any value of λ' can be achieved by appropriate choice of λ, so for any B-field representing a class of the claimed form L supports a non-zero object in the corresponding deformed Fukaya category. $\square$

Remark 5.1.23. In [17, Section 3], Biran and Cornea defined the *wide variety* $\mathcal{W}_2$ of a monotone Lagrangian pre-brane $L^\natural\subset X$ to be the set of homomorphisms $\rho\colon H_2^D\to\mathbb{C}^*$ such that $L^\natural$ is wide over $\mathbb{C}$ when equipped with the corresponding H_2^D local system (strictly they only allow homomorphisms whose kernels contain the torsion subgroup, but this detail is irrelevant in our case). They proved that this is an algebraic subvariety of $(\mathbb{C}^*)^{\operatorname{rank}H_2^D}$ which is cut out (in some coordinate system) by equations over $\mathbb{Z}$. We now have almost enough information to identify $\mathcal{W}_2$ for the $\mathrm{PSU}(N-1)$-homogeneous Lagrangian L we have been considering, equipped with the standard spin structure.

Let us coordinatise $\operatorname{Hom}(H_2^D,\mathbb{C}^*)$ by the t_j used throughout the above discussions. Proposition 5.1.18 tells us that $\mathcal{W}_2$ contains the variety defined by the equations

$$e_1(t_1,\ldots,t_N)=\cdots=e_{N-1}(t_1,\ldots,t_N)=0. \tag{5.15}$$

On the other hand, Lemma 5.1.5 and Lemma 5.1.7(iii) tell us that for any point $(t_1,\ldots,t_N)\in\mathcal{W}_2$ we have $e_1(\boldsymbol{\eta})=\overline{\eta}$ and $e_k(\boldsymbol{\eta})=0$ for $k=2,\ldots,N-1$. The former implies that

$$\overline{\eta}=\frac{e_1(t_1,\ldots,t_N)}{N-1}\cdot 1_L,$$

and substituting into the remaining equations we deduce that

$$e_k\left(t_1 - \frac{e_1(t_1,\dots,t_N)}{N-1},\dots,t_N - \frac{e_1(t_1,\dots,t_N)}{N-1}\right) = 0 \text{ for } k = 2,\dots,N-1. \tag{5.16}$$

Therefore $\mathcal{W}_2$ lies between the varieties defined by equations (5.15) and (5.16). //

Question 5.1.24. What *is* $\mathcal{W}_2$ in this case? Can the class $\overline{\eta}$ ever be non-zero (over any ring, with any H_2^D local system)? It follows from Remark 5.2.3 and Proposition 5.2.10 that the answer to the latter is 'yes' when $N = 3$. This question is answered fully by the new methods introduced in Chapter 6—see Proposition 6.1.9 and Corollary 6.1.11. //

Remark 5.1.25. If $e_1(t_1,\dots,t_N)$ is equal to some $\lambda \in \mathbb{C}^*$ then we cannot just apply our earlier wideness argument with the ideal $I_\lambda = (e_1 - \lambda, e_2, \dots, e_{N-1})$ in place of $I = (e_1,\dots,e_{N-1})$. The reason is that I_λ is not homogeneous, and the filtration induced on $\Lambda_\lambda \coloneqq \mathbb{Q}[H_2^D]/I_\lambda$ by the grading filtration on $\mathbb{Q}[H_2^D]$ is trivial. Explicitly, for each $p \in \mathbb{Z}$ we take $F^p\mathbb{Q}[H_2^D]$ to be the Laurent polynomials in $T_1,\dots,T_N$ which only involve monomials of total degree at least p, and $F^p\Lambda_\lambda$ to be the image of this subset in the quotient. But for any f in $\mathbb{Q}[H_2^D]$ and any positive integer r we have $f = fe_1^r/\lambda^r$ in Λ_λ, and hence the image of f is contained in the image of $F^p\mathbb{Q}[H_2^D]$ for all p. In other words, for all p we have $F^p\Lambda_\lambda = \Lambda_\lambda$, and hence the spectral sequence arising from the induced filtration on the pearl complex over Λ_λ is zero, and contains no information. //

Remark 5.1.26. When $N = 3$, the argument of Proposition 3.3.52 shows that the wide variety over fields of characteristic 2 is defined by $t_1 + t_2 + t_3 = 0$. //

Remark 5.1.27. Suppose that we do not make Assumption 3.2.6, so that the equalities

$$h_{N-|I|}(\boldsymbol{\eta}_I) = \left(\prod_{j\notin I} t_j\right) \cdot 1_L$$

from Lemma 5.1.5 (where I is a non-empty subset of $\{1,\dots,N\}$) have to be replaced by

$$h_{N-|I|}(\boldsymbol{\eta}_I) = \left(\varepsilon_{N-|I|}\prod_{j\notin I} t_j\right) \cdot 1_L$$

for some unknown signs $\varepsilon_0,\dots,\varepsilon_{N-1} \in \{\pm 1\}$ (the sign depends only on $|I|$, rather than the whole subset I, by considering its behaviour under permuting the factors). Moreover, we know that $\varepsilon_{N-1} = 1$ since the expressions for h_{N-1} come from quantum cohomology relations rather than the closed–open map.

Now work over $\mathbb{C}$ and set $t_j = \zeta^j$ for each j, where $\zeta = e^{2\pi i/N}$. By Corollary 5.1.22 (which doesn't rely on any explicit sign calculations) L is wide in this case, and from the $|I| = 1$ equalities we obtain $\eta_j = -\varepsilon_1\zeta^j \cdot 1_L$ for all j. Our relations then become

$$\varepsilon_1^{N-|I|}h_{N-|I|}(-\boldsymbol{\zeta}_I) = \varepsilon_{N-|I|}\prod_{j\notin I}\zeta^j, \tag{5.17}$$

where $\boldsymbol{\zeta}_I$ denotes the tuple $(\zeta^j)_{j\in I}$. In particular, all ε_j are determined by ε_1.

It is easy to see from Remark 5.1.19(ii) that $e_j(\boldsymbol{\zeta}) = 0$ for $j = 1,\dots,N-1$, so using the equivalence of (i) and (iii) in Lemma 5.1.7 (with $\eta_j = -\zeta^j$) we conclude that

$$h_{N-|I|}(\boldsymbol{\zeta}_I) = \prod_{j \notin I} \zeta^j$$

for all I. Plugging this into (5.17) we obtain $\varepsilon_j = \varepsilon_1^j$ for all j. If $\varepsilon_1 = 1$ we see that all discs do indeed count with sign $+1$ in our $\mathcal{CO}^0$ computations. Otherwise, $\varepsilon_j = (-1)^j$ for all j (this is only possible if N is odd, since we know $\varepsilon_{N-1} = 1$), so the contributions are still all $+1$ if we twist the standard spin structure by the class in $H^2(X, L; \mathbb{Z}/2)$ which assigns $1 \mod 2$ to each of the basic disc classes A_j. This just has the effect of transforming the signatures of relative spin structures from (a, b) to (b, a), so in particular the conclusions of our $\mathcal{CO}^0$ constraints, namely Proposition 5.1.9 and Corollary 5.1.11, remain valid. //

5.2 Further discussion and related examples

The purpose of this section is to describe an assortment of explicit calculations related to the ideas of the rest of the thesis. In Section 5.2.1 and Section 5.2.2 we showcase some general machinery (Perutz's symplectic Gysin sequence and Wehrheim–Woodward's quilt theory) for which there are few concrete applications in the literature, using it in the context of our Lagrangian PSU(2) in $(\mathbb{CP}^1)^3$ to recover the results we saw earlier and deal with the missing characteristic 5 case from Proposition 5.1.9. We spell out the geometry in detail, and chase through the crucial orientations and relative spin structures involved. An important ingredient is a close connection between our PSU(2) and the Chekanov torus T_{Ch} in $\mathbb{CP}^1 \times \mathbb{CP}^1$, and known computations of the index 2 discs which the latter bounds. One can also view the relationship between the two Lagrangians as a rationalisation of the significance of characteristic 3 for T_{Ch}. Section 5.2.3 extends this analysis to a related Lagrangian in $\mathbb{CP}^2 \times \mathbb{CP}^1$ and the corresponding Chekanov torus in $\mathbb{CP}^2$.

In Section 5.2.4 we consider a family of homogeneous Lagrangians in projective spaces which display similar properties to our main family, and fill in some computations missing from the literature by a variety of methods, including the application of periodicity in gradings. Finally, Section 5.2.5 discusses Lagrangians which can be obtained from symplectic reduction. We construct some new examples, compute some Floer cohomology rings, and combine these with our main family to build counterexamples to a conjecture of Biran and Cornea.

It is known that orientation schemes exist for the technology we use in Section 5.2.1–Section 5.2.3, but it is not obvious how they relate to each other. The reader is warned that the Floer cohomology groups appearing in these subsections may not be the same as those elsewhere in the thesis. Reassuringly, the results are all consistent.

5.2.1 The symplectic Gysin sequence

The $N = 3$ case of the construction considered in the previous section gives a Lagrangian $L \subset (\mathbb{CP}^1)^3$ diffeomorphic to $\mathrm{PSU}(2) \cong \mathrm{SO}(3) \cong \mathbb{RP}^3$. It is easy to see that it consists of all ordered triples of points on the sphere which form the vertices of an equilateral triangle on a great circle, so is in fact the lift of the Chiang Lagrangian under the branched cover $(\mathbb{CP}^1)^3 \to \operatorname{Sym}^3 \mathbb{CP}^1 \cong \mathbb{CP}^3$.

It was observed by Ivan Smith that L can be realised as a circle bundle over the antidiagonal sphere $\overline{\Delta}$ in $\mathbb{CP}^1 \times \mathbb{CP}^1$ and thus probed using the Gysin sequence. There are two distinct approaches to this theory in the literature, using different methods but leading to similar results: the Lagrangian circle bundle construction and Floer–Gysin sequence of Biran [12], Biran–Cieliebak [13] and Biran–Khanevsky [18], and Perutz's symplectic Gysin sequence associated to a spherically fibred coisotropic submanifold [116]. We shall follow the latter for convenience and because Perutz explicitly deals with coefficient rings of characteristic other than 2. Strictly Perutz works with a Novikov variable that can have arbitrary real exponents, see [116, Notation 1.5], but the monotonicity hypotheses mean that this is not strictly necessary and we can restrict to integer exponents as we have been using.

In order to describe this circle bundle perspective, we begin with an explicit Weinstein neighbourhood of the antidiagonal. Recall that our normalisation conventions are such that $\mathbb{CP}^1$ has area π, so we shall view it as a sphere of radius $1/2$ in $\mathbb{R}^3$, equipped with the standard Riemannian metric, which defines an isomorphism between the cotangent bundle and the tangent bundle

$$T\mathbb{CP}^1 = \left\{(q, v) \in \mathbb{R}^3 \times \mathbb{R}^3 : \|q\| = \frac{1}{2},\ \langle q, v\rangle = 0\right\}.$$

Consider the map

$$f \colon (q, v) \in T\mathbb{CP}^1 \mapsto \left(q \cos g + \frac{q \times v}{\|v\|} \sin g, -q \cos g + \frac{q \times v}{\|v\|} \sin g\right) \in \mathbb{CP}^1 \times \mathbb{CP}^1,$$

illustrated in Fig. 5.1 (with f_1 and f_2 denoting the two components), where g is a function of $\|v\|$. We shall see that for a particular choice of g the map f gives a symplectomorphism between

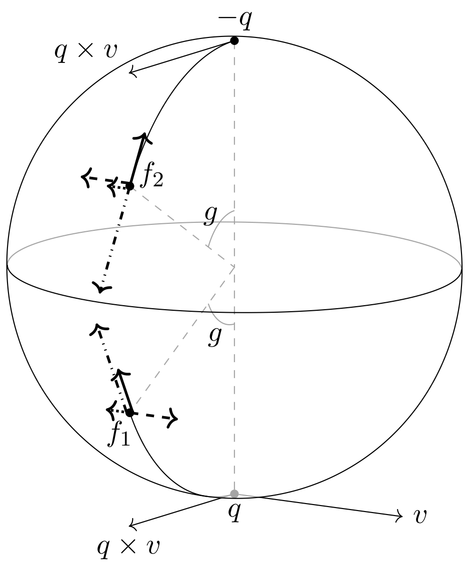

Figure 5.1: The map f.

the unit disc bundle $D\mathbb{CP}^1$ in the tangent (or, equivalently, cotangent) bundle of the sphere, and the complement of the diagonal Δ in $\mathbb{CP}^1 \times \mathbb{CP}^1$. Moreover it takes the zero section to the antidiagonal.

Remark 5.2.1. Our convention is that spheres in $\mathbb{R}^3$ are oriented such that the vertical projection from the tangent space at the south pole to the (x, y)-plane is orientation-preserving, where the latter carries symplectic form $\mathrm{d}x \wedge \mathrm{d}y$, and that (as stated in Section 1.2.4) the symplectic form on a cotangent bundle is $\sum_j \mathrm{d}p_j \wedge \mathrm{d}q_j$, rather than $\sum_j \mathrm{d}q_j \wedge \mathrm{d}p_j$. //

We consider four infinitesimal perturbations of (q, v) spanning the tangent space to $T\mathbb{CP}^1$:

$$\begin{aligned} &q \mapsto q && v \mapsto v + \varepsilon \frac{v}{\|v\|} \\ &q \mapsto q && v \mapsto v + \varepsilon \frac{q \times v}{\|v\|} \\ &q \mapsto q + \varepsilon \frac{v}{\|v\|} && v \mapsto v - \varepsilon \|v\| q \\ &q \mapsto q + \varepsilon \frac{q \times v}{\|v\|} && v \mapsto v, \end{aligned}$$

where ε is an infinitesimal parameter. Their images under the differential of f are indicated by the bold arrows in Fig. 5.1, with dash-dotted, dotted, dashed and solid tails respectively, and have lengths $g'/2$, $(\sin g)/4\|v\|$, $\cos g$ and $1/2$ (these lengths refer to each of the two vectors of a given type, with one lying in $T_{f_1}\mathbb{CP}^1$ and the other in $T_{f_2}\mathbb{CP}^1$). In order for f to be symplectic, the two parallelograms spanned by the dash-dotted and dashed vectors should have total area 1, whilst those spanned by the dotted and solid vectors should have total area 1/4. All other pairs should give total area 0. The latter is automatic, so we are left to solve

$$g' \cos g = 1 \text{ and } \sin g = \|v\|.$$

The solution is clearly $g = \arcsin\|v\|$ (which satisfies $g(0) = 0$, so the zero section does indeed map to the antidiagonal), giving

$$f \colon (q, v) \in D\mathbb{CP}^1 \mapsto \left(\sqrt{1 - \|v\|^2}\, q + q \times v, -\sqrt{1 - \|v\|^2}\, q + q \times v\right)$$

as our embedding of a Weinstein neighbourhood. This is almost the same as the formula obtained by Oakley–Usher [107, Proof of Lemma 2.3]. The difference is that their spheres are oppositely oriented to ours (beware that they also use different normalisations for the symplectic forms).

Now consider the circle bundle V_r of radius $r \in (0, 1)$ in $T\mathbb{CP}^1$. It is coisotropic, and the characteristic foliation corresponds to the geodesic flow on $\mathbb{CP}^1$. This flow is periodic and each orbit has the same period, so the foliation is actually a fibration with fibre S^1. The leaf space B is also a sphere, and for positive real numbers λ an explicit embedding in $\mathbb{R}^3$ is given by $(q, v) \mapsto \sqrt{\lambda}(q \times v)/r$. B inherits a symplectic structure and we claim that for a suitable choice of λ this embedding is a symplectomorphism.

A global section of the fibration over $B \setminus \{\text{north and south poles}\}$ is given by

$$q = \frac{1}{2}(\cos\varphi, \sin\varphi, 0) \text{ and } v = r(-\sin\theta \sin\varphi, \sin\theta\cos\varphi, \cos\theta),$$

where θ parametrises $(0, \pi)$ and φ parametrises $[0, 2\pi)$. A straightforward computation shows that the area of the parallelogram in $T\mathbb{CP}^1$ spanned by ∂_θ and ∂_φ is $(r\cos\theta)/2$, whilst the area of the image in B under our embedding in $\mathbb{R}^3$ is $(\lambda\cos\theta)/4$. We deduce that the embedding $B \hookrightarrow \mathbb{R}^3$ is a symplectomorphism when $\lambda = 2r$. The leaf space B is therefore $\mathbb{CP}^1$ with symplectic form scaled by $2r$, which we denote by $2r\mathbb{CP}^1$ (to emphasise, it is the *areas* that are scaled by $2r$, not the lengths). The graph of the projection of V_r to B^- (the $^-$ denotes that the sign of the symplectic form is reversed) exhibits $V_r \cong \mathbb{RP}^3$ as a Lagrangian submanifold of $D\mathbb{CP}^1 \times (2r\mathbb{CP}^1)^-$ [116, Proposition 1.1]. We can realise the reversal of the symplectic form on B by applying the antipodal map, and embed $D\mathbb{CP}^1$ in $\mathbb{CP}^1 \times \mathbb{CP}^1$ as above to obtain a Lagrangian in $\mathbb{CP}^1 \times \mathbb{CP}^1 \times 2r\mathbb{CP}^1$:

$$(q, v) \in V_r \mapsto \left(\sqrt{1-r^2}\; q + q \times v, -\sqrt{1-r^2}\; q + q \times v, -\sqrt{\frac{2}{r}}\; q \times v\right).$$

Note that V_r is isomorphic (as a circle bundle) to the unit tangent bundle of the sphere, which is diffeomorphic to $\mathbb{RP}^3$. When $r = 1/2$ we get exactly our earlier monotone Lagrangian L in $(\mathbb{CP}^1)^3$.

To summarise: the boundary of an appropriately-sized Weinstein neighbourhood of the antidiagonal in $\mathbb{CP}^1 \times \mathbb{CP}^1$ is a coisotropic submanifold fibred over $\mathbb{CP}^1$ by circles, and the graph of the projection to the base is our $L \subset (\mathbb{CP}^1)^3$. Such graphs of spherically fibred coisotropic submanifolds are precisely the framework of Perutz's symplectic Gysin sequence. Letting $\pi_j \colon (\mathbb{CP}^1)^3 \to \mathbb{CP}^1$ denote projection to the jth factor, the fibration relevant to us—which Perutz calls ρ—is $\pi_3|_L \colon L \to \mathbb{CP}^1$.

Remark 5.2.2. For other values of r we get non-monotone Lagrangians parametrising triangles on great circles which are isosceles, rather than equilateral. More generally, one could consider the zero level set of the moment map for the SU(2)- or SO(3)-action on the product $\alpha\mathbb{CP}^1 \times \beta\mathbb{CP}^1 \times \gamma\mathbb{CP}^1$, which parametrises scalene triangles on great circles with specific side lengths. Assuming that the areas are all positive (we can always do this by applying the antipodal map to any factor of negative area), this level set is a non-empty Lagrangian if and only if α, β and γ satisfy the triangle inequality, in which case the ratio of the side lengths of the parametrised triangles is

$$\sqrt{\frac{(\beta+\gamma)^2-\alpha^2}{\beta\gamma}} : \sqrt{\frac{(\gamma+\alpha)^2-\beta^2}{\gamma\alpha}} : \sqrt{\frac{(\alpha+\beta)^2-\gamma^2}{\alpha\beta}}.$$

//

Returning to the monotone case, Perutz tells us [116, Section 6.3] how to define a relative spin structure on L, in the form of a spin structure on $\xi|_L \oplus TL$, where ξ is a specified vector bundle on X. Explicitly, let $\xi = \pi_1^* T\mathbb{CP}^1 \oplus \pi_2^* T\mathbb{CP}^1$, and consider the short exact sequence

$$0 \to TL \oplus TL \xrightarrow{((\pi_1)_* \oplus (\pi_2)_*, \mathrm{id})} \xi|_L \oplus TL \to N \to 0$$

of vector bundles over L. The outer terms carry canonical spin structures (N is simply an orientable real line bundle, so admits a canonical homotopy class of trivialisation up to reversing the orientation), and the induced spin structure on $\xi|_L \oplus TL$ is the one we should take.

Now, the bundle ξ is stably trivial (explicitly we view each $\mathbb{CP}^1$ as a sphere in $\mathbb{R}^3$ with trivial normal bundle, and obtain an isomorphism $\xi \oplus \mathbb{R}^2 \cong \mathbb{R}^6$), so our relative spin structure on L is stably conjugate—in the sense of [63]—to an *absolute* spin structure. One could check whether this spin structure is the standard or unique non-standard one, but this is essentially irrelevant for our purposes since changing between the two simply reverses the signs of discs of index congruent to 2 mod 4, which is equivalent to transforming the Novikov variable $T \mapsto -T$. The upshot is that Perutz's results are valid in any characteristic if we use the standard spin structure on L, although there is some ambiguity in the sign of the Novikov variable.

We therefore see from [116, Section 1.2], after equipping L with the standard spin structure and trivial local system to give a monotone Lagrangian brane $L^\flat$, that for any ground ring R the Floer cohomology $HF^*(L^\flat, L^\flat; \Lambda_R)$ over $\Lambda_R = R[T^{\pm 1}]$ (with T in degree 2) fits into a long exact sequence of $QH^*(\mathbb{CP}^1; \Lambda_R)$-modules

$$\begin{tikzcd} QH^{*-2}(\mathbb{CP}^1; \Lambda_R) \arrow[rr, "*\widehat{e}"] & & QH^*(\mathbb{CP}^1; \Lambda_R) \arrow[dl] \\ & HF^*(L^\flat, L^\flat; \Lambda_R). \arrow[ul, "{[1]}"] & \end{tikzcd}$$

The horizontal arrow is quantum product with the class $\widehat{e} = e + \nu T$, where e is the Euler class of the circle bundle and ν is the signed count of index 2 discs through a point x of L which send a second boundary marked point to a *global angular chain*. The latter is a chain on L which intersects a generic circle fibre of $\pi_3|_L$ in a single point and whose boundary is the union of the fibres over a chain in the base representing the Poincaré dual of the Euler class.

In our case, the quantum cohomology groups $QH^*(\mathbb{CP}^1; \Lambda_R)$ are isomorphic to $\Lambda_R[H]/(H^2 - T^2)$, where H is the degree 2 hyperplane class, the Euler class e is $2H$, and the global angular chain therefore has boundary given by two fibres. We can take these two fibres to be those over the north and south poles of the base $\mathbb{CP}^1$, and then define the global angular chain to meet each other fibre in the south-pointing direction, as shown in Fig. 5.2. Setting our base point x

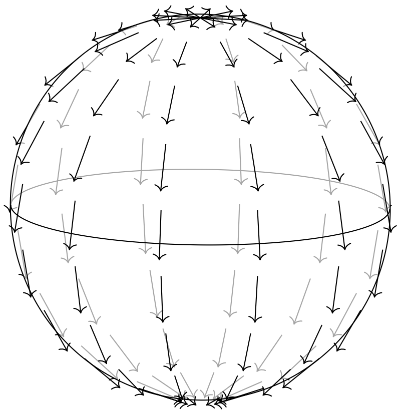

Figure 5.2: The global angular chain inside the unit tangent bundle of the sphere.

to be an equilateral triangle on the equator, two of the three index 2 discs through x (namely those whose boundaries are swept by the rotations about the first two vertices, i.e. the vertices in the first two $\mathbb{CP}^1$ factors) count with the opposite sign to the third. Therefore $\nu = \pm 1$, and hence $\widehat{e} = 2H \pm T$. The sign ambiguity in the count ν subsumes our earlier uncertainty about the sign of T.

We then have

$$(2H \mp T) * \widehat{e} = 3T^2$$

so, supposing R is a field, we see that quantum product with $\widehat{e}$ has rank 1 if $\operatorname{char} R = 3$ and is invertible otherwise. We deduce that $HF^*(L^\flat, L^\flat; \Lambda_R)$ is non-zero exactly when $\operatorname{char} R = 3$, and in this case it is one-dimensional in each degree. In characteristic 3 we in fact have

$$HF^{2j} \cong \langle T^{j-1}H, T^j \rangle / \langle 2T^{j-1}H \pm T^j \rangle$$

and

$$HF^{2j+1} \cong \langle 2T^{j-1}H \mp T^j \rangle,$$

where $\langle \cdot \rangle$ denotes R-linear span. Clearly $2H \pm T$ acts as zero between all pairs

$$HF^* \to HF^{*+2},$$

so H acts on HF^* as $\pm T$. This is consistent with our computation of $\mathcal{CO}^0$ in Section 5.1.2, where we saw that with the standard spin structure the hyperplane class H on a factor acts as $-T$ (and with the other spin structure it is easy to see that H acts as T).

Remark 5.2.3. In fact, for any classes $b_{12} \in H^2(\mathbb{CP}^1 \times \mathbb{CP}^1; \mathbb{Z}/2)$ and $b_3 \in H^2(\mathbb{CP}^1; \mathbb{Z}/2)$ we can equip L with a relative spin structure s with background class $(\pi_1 \times \pi_2)^* b_{12} + \pi_3^* b_3$ and Perutz's result remains valid as long as we incorporate the background class b_3 into $QH^*(\mathbb{CP}^1)$. Note, however, that the choice of s will in general also affect the count ν.

For example, if we switch to the relative spin structure which flips the sign of the index 2 disc whose boundary is swept by rotation about the third vertex, which has background class $H_1 + H_2$, then the count ν becomes ± 3 and $QH^*(\mathbb{CP}^1)$ is unaffected. Therefore, after setting the Novikov variable to 1, multiplication by $\hat{e}$ has matrix

$$\begin{pmatrix} \pm 3 & 2 \\ 2 & \pm 3 \end{pmatrix}$$

with respect to the obvious basis $1, H$. This has determinant 5, so L is wide precisely in characteristic 5. This deals with the missing $N = 3$, $p = 5$ case from Proposition 5.1.9 and Corollary 5.1.20.

Up to permuting the $\mathbb{CP}^1$ factors, an equivalent relative spin structure can be obtained by instead flipping the sign of one of the other index 2 discs. In this case ν becomes ∓ 1 (i.e. its sign is reversed) but the relation in $QH^*(\mathbb{CP}^1)$ becomes $H^2 = -1$ (rather than $+1$) so $\hat{e}*$ has matrix

$$\begin{pmatrix} \mp 1 & -2 \\ 2 & \mp 1 \end{pmatrix}.$$

This again has determinant 5, as expected. //

5.2.2 Quilt theory and the Chekanov torus in $\mathbb{CP}^1 \times \mathbb{CP}^1$

We now turn to the quilt theory of Wehrheim and Woodward, set out in [142] and subsequent papers by the same authors and by Ma'u-Wehrheim–Woodward (this theory is actually also the basis of Perutz's Gysin sequence).

Recall that given symplectic manifolds (X_j, ω_j) a *Lagrangian correspondence* from X_{j-1} to X_j is a Lagrangian submanifold $L_{(j-1)j}$ of $X_{j-1}^- \times X_j$, where X_{j-1}^- is shorthand for $(X_{j-1}, -\omega_{j-1})$. These generalise both ordinary Lagrangians in X_j, when X_{j-1} is a point, and symplectomorphisms from X_{j-1} to $X_j = X_{j-1}$, when $L_{(j-1)j}$ is the graph. The *composition* of correspondences $L_{(j-1)j}$ and $L_{j(j+1)}$, written $L_{(j-1)j} \circ L_{j(j+1)}$, is the subset

$$\pi_{(j-1)(j+1)}\Big(\big(L_{(j-1)j} \times L_{j(j+1)}\big) \cap \Big(X_{j-1}^- \times \Delta_{X_j} \times X_{j+1}\Big)\Big) \subset X_{j-1}^- \times X_{j+1}, \tag{5.18}$$

where $\pi_{(j-1)(j+1)}$ is the projection

$$X_{j-1}^- \times X_j \times X_j^- \times X_{j+1} \to X_{j-1}^- \times X_{j+1}$$

and Δ_{X_j} is the diagonal in $X_j^- \times X_j$. The correspondence is said to be *embedded* if the intersection in (5.18) is transverse and the restriction of $\pi_{(j-1)(j+1)}$ to this intersection is an embedding, in which case it is a Lagrangian correspondence from X_{j-1} to X_{j+1}.

The idea of Wehrheim–Woodward is to define (under appropriate hypotheses) a 'quilted' Floer cohomology for cycles of Lagrangian correspondences X_0 to X_1 to $\dots$ to $X_{r+1} = X_0$, and prove that it is invariant under replacing consecutive correspondences by their composition when it is embedded. Moreover, when $r = 1$ and X_0 is a point, so the cycle of correspondences is just a pair of Lagrangians in X_1, their theory reproduces the ordinary Lagrangian intersection Floer cohomology of the two Lagrangians. The aspirational guiding principle, which goes back to ideas of Fukaya, is that composing with correspondences should yield a functor

$$\mathcal{F}(X_0^- \times X_1) \to \mathrm{Fun}(\mathcal{F}(X_0), \mathcal{F}(X_1))$$

from the Fukaya category of the product to the category of functors between the Fukaya categories of the factors; see [102] and the enormous [56] for progress towards this goal. For us the important result is:

Theorem 5.2.4 ([142, Theorem 6.3.1])**.** *Suppose we have a Lagrangian correspondence L_{01} from X_0 to X_1 and a Lagrangian L_1 in X_1 such that the composition $L_0 \coloneqq L_{01} \circ L_1$ is embedded. Assume moreover that all of these manifolds are closed, oriented and monotone, with the same monotonicity constant, that $\pi_1(X_0 \times X_1)$ is torsion, and that $HF^*(L_0, L_0) \neq 0$. Then $HF^*(L_0 \times L_1, L_{01})$ is well-defined and isomorphic to $HF^*(L_0, L_0)$ as a $\mathbb{Z}/2$-graded module over the ground ring. In particular, it is a non-zero module over $HF^*(L_{01}, L_{01})$, so the latter is also non-zero.*

Remark 5.2.5. We have been deliberately vague about the coefficients here: as stated the result only applies in characteristic 2, and to move outside this setting we need the orientations constructed in [141]. First we fix relative spin structures on L_{01} and L_1 whose background classes on the $X_1^{(-)}$ factors differ by $w_2(X_1)$. We then take the canonical relative spin structure on the diagonal in $X_1 \times X_1^-$ whose background class is the pullback of $w_2(X_1)$ from X_1^-, and find the induced relative spin structure on the composition L_0 defined by [141, Remark 4.6.2(a)]. It is with respect to these relative spin structures on L_0, L_{01} and L_1 that Theorem 5.2.4 holds over arbitrary coefficient rings (with trivial local systems).

In the simplest situation, we choose spin structures on TL_{01} and on $TL_1 \oplus TX_1^-|_{L_1}$ (representing relative spin structures on L_{01} and L_1 with background classes 0 and $w_2(X_1)$ respectively) and consider the short exact sequence

$$0 \to T\Gamma \to TL_{01} \oplus TL_1 \to \Delta^\perp \to 0$$

of vector bundles over $\Gamma = L_{01} \times_{X_1} L_1 = (X_0^- \times \Delta_{X_1}) \cap (L_{01} \times L_1)$, where

$$\Delta^\perp = \{(v, -v) \in T(X_1 \times X_1^-)|_{\Delta_{X_1}} : v \in TX_1\}.$$

Letting E denote the pullback of TX_1^- to Γ (which we identify with L_0), and adding it to the middle and right-hand terms, we obtain

$$0 \to TL_0 \to TL_{01} \oplus (TL_1 \oplus E) \to \Delta^\perp \oplus E \to 0. \tag{5.19}$$

Spin structures on TL_{01} and $TL_1 \oplus E$ are precisely what we have chosen, whilst $\Delta^\perp \oplus E$ carries a canonical spin structure from the obvious identification $\Delta^\perp \cong E$ over Δ_{X_1} (this identification is obvious up to a choice of sign, which makes no difference). The induced spin structure on TL_0 is the one we need. //

We shall view $L \subset (\mathbb{CP}^1)^3$ as a Lagrangian correspondence L_{01} from $X_0 = (\mathbb{CP}^1 \times \mathbb{CP}^1)^-$ to $X_1 = \mathbb{CP}^1$, and consider its composition L_0 with the Clifford torus (i.e. the equatorial circle) L_1 in X_1. This is equivalent to performing symplectic reduction at the equatorial level set for the S^1-action on the third $\mathbb{CP}^1$ factor by rotation about the vertical axis. This composition is embedded, and L_0 is precisely the monotone Chekanov torus T_{Ch} in X_0 as presented by Entov–Polterovich [42, Example 1.22]. It consists of ordered triples of points on the sphere which form the vertices of an equilateral triangle on a great circle, such that the third point is constrained to the equator. We'll refer to this as the *apex*, and the other two points as the *base vertices*. The embedding of the Lagrangian in $(\mathbb{CP}^1 \times \mathbb{CP}^1)^-$ is given by forgetting the apex.

Remark 5.2.6. This torus was first discovered by Chekanov in $\mathbb{R}^4$ [29], and appears in both $\mathbb{CP}^2$ and $\mathbb{CP}^1 \times \mathbb{CP}^1$ in many Hamiltonian isotopic guises. Comparisons between various different constructions are given by Gadbled [72] and Oakley–Usher [107]. One alternative (though obviously equivalent) viewpoint is Gadbled's description of the torus as the Biran circle bundle over the equatorial circle in the diagonal $\Delta \subset \mathbb{CP}^1 \times \mathbb{CP}^1$, which is in a sense complementary to our construction as a restriction of the circle bundle over the antidiagonal. //

Remark 5.2.7. Taking this composition, or equivalently performing the symplectic reduction, for the Lagrangians obtained as circle bundles V_r of arbitrary radius $r \in (0, 1)$, we obtain a continuous family of Lagrangian tori in $\mathbb{CP}^1 \times \mathbb{CP}^1$, which are non-monotone unless $r = 1/2$. Fukaya–Oh–Ohta–Ono [67] showed that these tori are non-displaceable for $r \in (0, 1/2]$, using Floer cohomology with bulk deformations. //

Question 5.2.8. What are the displaceability properties of the non-monotone Lagrangians in Remark 5.2.2? Perhaps a non-monotone version of Proposition 3.3.52 can be used here. Can these be related, in the isosceles ($\alpha = \beta$) case, to the non-displaceability of the reduced tori for $r \in (0, 1/2]$? //

The hypotheses of Theorem 5.2.4 are easily verified for this example and so non-vanishing of $HF^*(T_{\mathrm{Ch}}, T_{\mathrm{Ch}})$ implies non-vanishing of $HF^*(L, L)$. The former is equivalent, by the argument in Example 2.2.4, to the vanishing of the homology class swept by the boundaries of the index 2 discs through a generic point of T_{Ch}, and these discs were explicitly computed (for a specific regular complex structure) by Chekanov–Schlenk [30, Lemma 5.2]. There are exactly five such discs, in classes D_1, $S_1 - D_1 - D_2$, $S_1 - D_1$, $S_2 - D_1$ and $S_2 - D_1 + D_2$ in $H_2(\mathbb{CP}^1 \times \mathbb{CP}^1, T_{\mathrm{Ch}})$, where S_1 and S_2 are the classes of the spheres in each factor and D_1 and D_2 are discs whose boundaries form a basis for $H_1(T_{\mathrm{Ch}}; \mathbb{Z})$. The sum of their boundaries is therefore $-3\partial D_1$, and we deduce that T_{Ch}, and hence L, has non-vanishing self-Floer cohomology in characteristic 3.

Of course we need to be careful about the orientations, and in particular ensure that when we take the standard spin structure on L we can choose a relative spin structure on the equator L_1 with background class $w_2(\mathbb{CP}^1) = 0$ (which we may as well think of as an absolute spin structure) so that all five discs on L_0 count with the same sign under the induced relative spin structure from Remark 5.2.5. The easiest way to do this is to use the fact that the counts $\mathfrak{m}_0$ of index 2 discs on the Lagrangians satisfy

$$\mathfrak{m}_0(L_0) + \mathfrak{m}_0(L_{01}) + \mathfrak{m}_0(L_1) = 0. \tag{5.20}$$

This is because the proof of [142, Theorem 6.3.1] establishes well-definedness of the quilted Floer cohomology of the cycle (L_0, L_{01}, L_1), and the left-hand side of (5.20) is exactly the square of the differential on this complex. We have already seen that our Lagrangian $L = L_{01}$ bounds three index 2 discs through a generic point, whilst L_1 bounds two, all counting positively with the standard spin structures by Proposition B.4.1. Therefore the five discs above all count with sign -1 and we do indeed get vanishing of the boundary class exactly in characteristic 3.

Remark 5.2.9. Alternatively one can directly compute the induced spin structure on T_{Ch}. The first step is to view the standard spin structure on L_1 as a spin structure on $TL_1 \oplus T(\mathbb{CP}^1)^-|_{L_1}$, which we can do by the 'shift of background class' operation described by Wehrheim–Woodward in [141, Remark 5.1.8]. According to this procedure, we start with a spin structure on TL_1 and then take the direct sum with $T(\mathbb{CP}^1)^-|_{L_1}$, carrying the canonical spin structure arising from the splitting

$$T(\mathbb{CP}^1)^-|_{L_1} \cong TL_1 \oplus T^*L_1 \cong TL_1 \oplus TL_1. \tag{5.21}$$

Viewing relative spin structures as a torsor for $H^2(\mathbb{CP}^1, L_1; \mathbb{Z}/2)$, Wehrheim–Woodward compute that this shift corresponds to half the Maslov class, reduced modulo 2, so in order for the shifted structure to be (stably conjugate to) the standard spin structure we should start with the *non-standard* one.

The next step is to consider the short exact sequence

$$0 \to TT_{\mathrm{Ch}} \to TL \oplus (TL_1 \oplus T(\mathbb{CP}^1)^-) \to \Delta^\perp \oplus T(\mathbb{CP}^1)^- \to 0.$$

of vector bundles over $\Gamma = (L \times L_1) \cap (\mathbb{CP}^1 \times \mathbb{CP}^1 \times \Delta_{\mathbb{CP}^1}) \subset (\mathbb{CP}^1)^4$ from (5.19). If we equip $\Delta^\perp|_{\Delta_{L_1}}$ with the spin structure analogous to that given by (5.21) then we can cancel off the explicit $T(\mathbb{CP}^1)^-$ summands and obtain

$$TT_{\mathrm{Ch}} \oplus \Delta^\perp \cong TL \oplus TL_1 \tag{5.22}$$

over Γ, where TL carries the standard spin structure and TL_1 carries the non-standard one.

There are two obvious S^1-actions on T_{Ch}, given by rotating about a vertical axis and about the apex respectively, and after fixing a base point and orientations for these rotations (all of which are irrelevant for us) we obtain an identification $T_{\mathrm{Ch}} \cong S^1 \times S^1$. This defines a standard spin structure on T_{Ch} and an identification $\pi_1(T_{\mathrm{Ch}}) \cong \mathbb{Z}^2$. If L were equipped with its non-standard spin structure, which is equivalent to trivialising TL by transporting a basis of one tangent space around the whole of L by the PSU(2)-action, and L_1 with its standard one, then (5.22) would induce the standard spin structure on T_{Ch}. Reversing the spin structure on L changes the induced spin structure on T_{Ch} by $(1, 1) \in (\mathbb{Z}/2)^2 = H^1(T_{\mathrm{Ch}}; \mathbb{Z}/2)$, since each S^1 factor is given by a rotation which is non-trivial in $\pi_1(\mathrm{PSU}(2))$, whilst reversing the spin structure on L_1 changes the induced spin structure by $(1, 0)$.

We conclude that the required spin structure on T_{Ch} differs from the standard one by $(0, 1)$, which is the unique non-zero element of $H^1(T_{\mathrm{Ch}}; \mathbb{Z}/2)$ which vanishes on a loop swept by rotation about a vertical axis. This rotation lifts to a global circle action on $\mathbb{CP}^1 \times \mathbb{CP}^1$, using which one can read off from Chekanov–Schlenk's construction that the corresponding loop is (plus or minus) the boundary of D_2 (which they call D_τ). Therefore to pass from the standard to the induced spin structure we simply modify the sign of a disc u according to the parity of the coefficient of D_1 in $[u]$. Since all index 2 discs have an odd coefficient of D_1, we see that switching spin structures just reverses all of their signs.

For the standard spin structure–which can be defined as the spin structure induced by *any* diffeomorphism $T_{\mathrm{Ch}} \cong S^1 \times S^1$—the superpotential of T_{Ch} has been calculated by Auroux [9, Section 5.4] and Fukaya–Oh–Ohta–Ono [67, Proposition 3.1, Theorem 3.1]. More conceptually, one can read it off from the superpotential of the Clifford torus in $\mathbb{CP}^1 \times \mathbb{CP}^1$, with the standard spin structure, using the wall-crossing formula recently proved by Pascaleff–Tonkonog [115]. The key point is that all discs count positively in this case, so once we pass to the induced spin structure all discs count negatively, in agreement with the earlier $\mathfrak{m}_0$ computation. //

More generally, suppose we equip L and L_1 with relative spin structures with background classes $a_1 H_1 + a_2 H_2 + a_3 H_3$ and $a_3 H_3$ respectively, where H_j is the hyperplane class on the jth $\mathbb{CP}^1$ factor of $X = (\mathbb{CP}^1)^3$ and $a_1, a_2, a_3 \in \mathbb{Z}/2$. The induced relative spin structure on T_{Ch} then has background class $a_1 H_1 + a_2 H_2$. Compared with the case of the standard spin structures on L and L_1, where all five discs on T_{Ch} count negatively, we must modify the signs of the discs in classes $S_1 - D_1 - D_2$ and $S_1 - D_1$ by $(-1)^{a_1}$ and those in classes $S_2 - D_1$ and $S_2 - D_1 + D_2$ by $(-1)^{a_2}$ to take into account the change of background class. We must then further modify the signs of all five discs by pairing their boundaries with a class ε in $H^1(T_{\mathrm{Ch}}; \mathbb{Z}/2)$, in order to move to the correct relative spin structure within this background class. Often the required class ε is completely determined by $\mathfrak{m}_0(T_{\mathrm{Ch}})$, which we know is equal to $-(\mathfrak{m}_0(L) + \mathfrak{m}_0(L_1))$.

As an application we can reprove the $N = 3$, $p = 5$ result from Remark 5.2.3:

Proposition 5.2.10. *L is wide over a field of characteristic* 5 *when equipped with a relative spin structure of signature* $(2,1)$ *and the trivial local system.*

Proof. Take the non-standard spin structure on the equator L_1, and the relative spin structure on L which differs from the non-standard spin structure by the class in $H^2(X, L; \mathbb{Z}/2)$ Poincaré dual to the diagonal on the first two $\mathbb{CP}^1$ factors. These have background classes 0 and $H_1 + H_2$ respectively, and give $\mathfrak{m}_0(L_0) = -2$ and $\mathfrak{m}_0(L) = -1$ (the latter comes from the fact that the non-standard spin structure on L has $\mathfrak{m}_0 = -3$, and the diagonal on the first two factors intersects exactly one of the three index 2 discs). After performing the change of background class on T_{Ch} we have one disc in class D_1 counting negatively and discs in classes $S_1 - D_1 - D_2$, $S_1 - D_1$, $S_2 - D_1$ and $S_2 - D_1 + D_2$ counting positively. This has $\mathfrak{m}_0 = 3$, as we need.

An easy calculation shows that if the class ε from the above discussion is non-zero then we get a different $\mathfrak{m}_0$-value (for example, if ε is zero on D_1 but non-zero on D_2 then the signs of two of the four positive discs become negative, giving $\mathfrak{m}_0 = -1$). Therefore $\varepsilon = 0$ and the discs count with the signs just described. In particular, the sum of their boundaries is $-5\partial D_1$, which vanishes in characteristic 5, giving wideness of L in this case. □

Remark 5.2.11. One can view T_{Ch} itself as a Lagrangian correspondence from $X_0 = (\mathbb{CP}^1)^-$ to $X_1 = \mathbb{CP}^1$ and compose it with the equator on the latter. The intersection in (5.18) is transverse but the projection is no longer injective and is instead a 2-to-1 cover. In this case the composition theorem does not apply immediately, but it should be valid if certain bubbling can be ruled out [142, Remark 5.4.2(e)] (see also [98], [56] for an alternative approach). A computation similar to that in Remark 5.2.9 shows that the spin structure induced on L_0 from the standard spin structures on L_{01} and L_1 is the standard one, so $\mathfrak{m}_0(L_0) = 2$. Combining this with $\mathfrak{m}_0(L_{01}) = 5$ and $\mathfrak{m}_0(L_1) = 2$ we see that (5.20) holds over a field R only if $\operatorname{char} R = 3$, so some bad bubbling must take place. Indeed, the composition theorem cannot possibly hold with arbitrary coefficients, otherwise we would deduce wideness of T_{Ch} in characteristics in which we already know it is narrow. //

Question 5.2.12. Can the bubbling be shown to cancel mod 3, proving that T_{Ch} is wide in characteristic 3 without computing all of the discs? Can a similar argument be used to show wideness of T_{Ch} in characteristic 5 with the relative spin structure from the proof of Proposition 5.2.10? //

Remark 5.2.13. Over a field R of characteristic 3, the quantum cohomology of $(\mathbb{CP}^1)^3$ breaks up as the direct sum of 8 copies of R, and we obtain a corresponding decomposition of the monotone Fukaya category according to whether $\mathcal{CO}^0$ of each H_j is 1 or 2. With the standard spin structure and trivial local system, our Lagrangian L lies in the summand on which $\mathcal{CO}^0(H_j) = 1$ for each j, and by Abouzaid's criterion [1] (proved in the monotone case by Sheridan [130]) it must split-generate this summand. The Clifford torus in $(\mathbb{CP}^1)^3$, with its standard spin structure, lies in the same summand and we conjecture that it is quasi-isomorphic to a direct sum of two copies of L. This would mean that the quilt functor it induces from $\mathbb{CP}^1$ to $\mathbb{CP}^1 \times \mathbb{CP}^1$ is quasi-isomorphic to two copies of the quilt functor induced by L.

We saw above that the latter sends the equator to the Chekanov torus, and it is easy to see that—heuristically at least—the former sends the equator to two copies of the Clifford torus in $\mathbb{CP}^1 \times \mathbb{CP}^1$ (this composition is not embedded, so some work would be needed to make this rigorous). Assuming this result, and the above conjecture, we deduce that two copies of the Chekanov torus are quasi-isomorphic to two copies of the Clifford torus in an appropriate summand of the mod 3 Fukaya category of $\mathbb{CP}^1 \times \mathbb{CP}^1$. In particular, T_{Ch} is Floer-theoretically non-trivial in characteristic 3 without computing any discs. //

5.2.3 The Chekanov torus in $\mathbb{CP}^2$

Just as there are Lagrangian orbits for the standard SU(2)-actions on $(\mathbb{CP}^1)^3$ and $\mathrm{Sym}^3 \mathbb{CP}^1 = \mathbb{CP}^3$, there is also such an orbit on $(\mathrm{Sym}^2 \mathbb{CP}^1) \times \mathbb{CP}^1 = \mathbb{CP}^2 \times \mathbb{CP}^1$, which we now analyse briefly. The branched cover $\mathbb{CP}^1 \times \mathbb{CP}^1 \to \mathbb{CP}^2$ is given by

$$([a:b],[c:d]) \to \left[ac : \frac{ad+bc}{\sqrt{2}} : bd\right],$$

using homogeneous coordinates on $\mathbb{CP}^n$ induced by symplectic reduction from $\mathbb{C}^{n+1}$, and the image of the diagonal (respectively antidiagonal) is parametrised by $[z^2 : \sqrt{2}z : 1]$ (respectively $[z : (1-|z|^2)/\sqrt{2} : -\overline{z}]$) for $z \in \mathbb{CP}^1$. We shall refer to these images themselves as the diagonal and antidiagonal in $\mathbb{CP}^2$. Note that although the former is a sphere, the latter is a real projective plane, and is in fact the 'Chiang-type Lagrangian in $\mathbb{CP}^2$' described by Cannas da Silva [26]. Under the identification of $\mathbb{CP}^2$ with $\mathrm{Sym}^2 \mathbb{CP}^1$, a point in $\mathbb{CP}^2$ represents an unordered pair of points on the sphere.

We view $\mathbb{RP}^2$ as the sphere of radius $1/2$ in $\mathbb{R}^3$ modulo the antipodal involution, with the induced metric, and thus view $T\mathbb{RP}^2$ as

$$\left\{(q,v) \in \mathbb{R}^3 \times \mathbb{R}^3 : \|q\| = \frac{1}{2},\ \langle q,v\rangle = 0\right\} \Big/ (q,v) \sim (-q,-v).$$

Let g denote the function

$$\frac{1-\sqrt{1-\|v\|^2}}{2\|v\|^2}$$

on $T\mathbb{RP}^2$, with $g(q,0) = 1/4$ for all q. Adapting the embedding of the Weinstein neighbourhood of $\mathbb{RP}^2 \subset \mathbb{CP}^2$ given by Oakley–Usher [107, Lemma 3.1] to our normalisation, and then modifying by a linear automorphism of $\mathbb{CP}^2$, we see that the map

$$(q,v) \mapsto \left[\begin{pmatrix} 1 & i & 0 \\ 0 & 0 & -\sqrt{2} \\ -1 & i & 0 \end{pmatrix} \left(\sqrt{g}v + \frac{i}{\sqrt{g}}q\right)\right]$$

embeds the unit disc bundle $D\mathbb{RP}^2$ as a Weinstein neighbourhood of the antidiagonal, sending $(\pm q, 0)$ to the unordered pair $\pm q$ on $\mathbb{CP}^1$. Oakley and Usher show that the complement of the image of their embedding is the quadric $z_0^2 + z_1^2 + z_2^2 = 0$, so the complement of the image of our embedding is given by applying the same automorphism of $\mathbb{CP}^2$ to this quadric. We get precisely the diagonal $z_1^2 = 2z_0z_2$.

The pair of points corresponding to the image of $(q = (1/2,0,0), v = (0,r,0))$ under our embedding can be computed directly, and is obtained by sliding the points $\pm q$ along the geodesics with initial direction $q \times v$, just like in Fig. 5.1. It is straightforward to check that the embedding is SU(2)-invariant, and hence this sliding-along-geodesics property holds for all values of q and v. In other words, the picture so far is rather similar to that we obtained before for $\mathbb{CP}^1 \times \mathbb{CP}^1$: there is an SU(2)-equivariant embedding of a Weinstein neighbourhood of the antidiagonal as the complement of the diagonal, given by moving the antipodal pair of points $\pm q$ in the direction $q \times v$.

The angle θ between the pair of points on the sphere corresponding to a point $[z_0 : z_1 : z_2] \in \mathbb{CP}^2$ satisfies

$$\cos\theta = \frac{\|z\|^2 - 3(\!|z|\!)^2}{\|z\|^2 + (\!|z|\!)^2},$$

where $(\!|z|\!)^2$ denotes the absolute value of the SU(2)-invariant quadratic form on $\mathbb{C}^3$ given by $2z_1z_3 - z_2^2$ (since both sides of this angle expression are SU(2)-invariant, it only need be checked when the pair of points is $0 \in \mathbb{CP}^1$ and $x \in \mathbb{R} \subset \mathbb{CP}^1$). Using this relation it is not hard to verify that our embedding sends (q, v) to a pair of points separated by angle θ satisfying

$$\cos\frac{\theta}{2} = \frac{1 - \sqrt{1 - \|v\|^2}}{\|v\|}.$$

This is in contrast to the case of $\mathbb{CP}^1 \times \mathbb{CP}^1$ where we had $\cos(\theta/2) = \|v\|$. Putting everything together, we obtain the following commutative diagram

$$\begin{tikzcd}
(q, v) \arrow[d, mapsto] & D\mathbb{CP}^1 \arrow[r, hook] \arrow[d] & \mathbb{CP}^1 \times \mathbb{CP}^1 \arrow[d] \\
\left(q, \frac{2v}{1+\|v\|^2}\right) & D\mathbb{RP}^2 \arrow[r, hook] & \mathbb{CP}^2,
\end{tikzcd} \tag{5.23}$$

where the horizontal arrows are our Weinstein neighbourhood embeddings and the right-hand vertical arrow is the symmetrisation map.

Now consider the action of SU(2) on $\mathbb{CP}^2 \times \lambda\mathbb{CP}^1$, recalling that the λ represents the scaling of the symplectic form, and take the basis of $\mathfrak{su}(2)^*$ dual to

$$\begin{pmatrix} 0 & \frac{i}{2} \\ \frac{i}{2} & 0 \end{pmatrix}, \begin{pmatrix} 0 & -\frac{1}{2} \\ \frac{1}{2} & 0 \end{pmatrix} \text{ and } \begin{pmatrix} \frac{i}{2} & 0 \\ 0 & -\frac{i}{2} \end{pmatrix},$$

which represent infinitesimal rotations about the x-, y- and z-axes respectively. Restricting to the image of our embedding of $D\mathbb{RP}^2$, we can view the domain of the moment map μ as

$$\left\{\big((q, v), q'\big) \in (\mathbb{R}^3 \times \mathbb{R}^3) \times \mathbb{R}^3 : \|q\| = \|q'\| = \frac{1}{2}, \langle q, v\rangle = 0 \text{ and } \|v\| < 1\right\} \Big/ (q, v) \sim (-q, -v),$$

and a calculation using (5.1) shows that in the chosen basis μ is given by

$$\mu\big((q, v), q'\big) \mapsto \frac{\lambda}{2}q' - v \times q. \tag{5.24}$$

In particular, for $\lambda = 2r \in (0, 1)$ the zero level set is given by $\|v\| = r$ and $q' = 2(v \times q)/r$, and this set is three-dimensional and hence Lagrangian. It is diffeomorphic to $\mathrm{SU}(2)/(\mathbb{Z}/4)$ and thus has finite fundamental group, so is monotone precisely when the ambient manifold is. In our standard normalisation a line has area π in both $\mathbb{CP}^2$ and $\mathbb{CP}^1$, but its index is 6 and 4 respectively, so monotonicity occurs when $r = 1/3$.

Using (5.23) we see that the monotone Lagrangian, which we'll denote by L', therefore parametrises triangles on a great circle which are *not* equilateral; rather they are isosceles with apex angle $\arccos(3 - 2\sqrt{2})$, which is greater than $\pi/3$. Note that the two base vertices are unordered. This Lagrangian L' is not then the image of our previous Lagrangian $L \subset (\mathbb{CP}^1)^3$. We can still use some of the same methods to probe it though.

First we need to understand the index 2 discs on L', which we equip with the standard spin structure and trivial local system. We know that all such discs u are axial, of the form

$$z \mapsto e^{-i\xi \log z} u(1)$$

for some ξ in $\mathfrak{su}(2)$, and from the proof of Proposition 3.1.13 the area of u is given by

$$2\pi \langle \mu(u(0)), \xi \rangle.$$

If ξ generates the rotation through angle 2π about the unit vector n, and this axis of rotation does not contain any of the vertices of the triangle represented by $u(1)$, then $u(0)$ corresponds to $q \times v = q' = -n/2$ and so using (5.24) we see that the area of u is $2\pi/3$. Since a line on $\mathbb{CP}^2$ has area π and index 6, we deduce that the disc u has index 4. In other words, the only possible index 2 discs through a point $x \in L'$ are generated by rotations about axes containing one of the vertices of the triangle represented by x. By similar calculations it is easy to check that each of the three axes give a single index 2 disc. Thus $\mathfrak{m}_0(L') = 3$. Alternatively the indices can be computed using the vanishing of a holomorphic 3-form, as in Lemma 3.1.9(iii).

The compactification divisor has two components—one on which the two base points come together, in class $2H_1$, and the other on which the apex and one of the base points come together, in class $H_1 + 2H_2$, where H_1 is the hyperplane class on $\mathbb{CP}^2$ and H_2 is the hyperplane class on $\mathbb{CP}^1$. One of the index 2 discs meets the former, whilst two meet the latter, so we deduce that

$$\mathcal{CO}^0(2H_1) = 1_{L'} \text{ and } \mathcal{CO}^0(H_1 + 2H_2) = 2 \cdot 1_{L'}; \tag{5.25}$$

we have set the Novikov variable to 1 for simplicity. Cubing the first of these equations and using the quantum cohomology relation $H_1^3 = 1$, we see that L' can only be non-narrow over a coefficient ring of characteristic 7.

Again we can apply Perutz's symplectic Gysin sequence, this time to the circle bundle given by projecting L' to the $\mathbb{CP}^1$ factor, which has Euler number 4. There is a global angular chain analogous to Fig. 5.2, comprising those isosceles triangles which lie in a vertical plane (the plane is not fixed), whose boundary consists of two copies of the fibre over each of the north and south poles. With the standard spin structure the count ν is then ± 3—the 'rotate about the apex' disc hits the angular chain once, whereas each of the other two discs hits it twice with the opposite sign—so the class $\widehat{e}$ is $4H \pm 3$. Multiplication by $\hat{e}$ on $QH^*(\mathbb{CP}^1)$ thus has determinant 7, and so L' is indeed wide in characteristic 7. We also see that H (which is really H_2) acts as ± 1—this is consistent with (5.25) in the -1 case.

Remark 5.2.14. Changing the relative spin structure by flipping the sign of exactly one of the two families of index 2 discs has the effect of changing ν to ± 5 but leaving $QH^*(\mathbb{CP}^1)$ unaltered, and so we get wideness in characteristics dividing

$$\det \begin{pmatrix} \pm 5 & 4 \\ 4 & \pm 5 \end{pmatrix} = 9.$$

This is consistent with our $\mathcal{CO}^0$ constraints as follows.

Suppose we reverse the sign of the 'rotate about the apex' disc. We obtain $\mathcal{CO}^0(2H_1) = -1_{L'}$, with background class $2H_1 = 0 \in H^2(\mathbb{CP}^2 \times \mathbb{CP}^1; \mathbb{Z}/2)$, so cubing gives $9 \cdot 1_{L'} = 0$. If instead we reverse the sign of the other family of discs then we revert to the original expression $\mathcal{CO}^0(2H_1) = 1_{L'}$ but now with background class $H_1 + 2H_2 = H_1 \in H^2(\mathbb{CP}^2 \times \mathbb{CP}^1; \mathbb{Z}/2)$. This changes the relation in $QH^*(\mathbb{CP}^2)$ to $H_1^3 = -1$, and so again when we cube we get $9 \cdot 1_{L'} = 0$. //

Alternatively we can view L' as a Lagrangian correspondence from $(\mathbb{CP}^2)^-$ to $\mathbb{CP}^1$ and once more form the composition with the equator on $\mathbb{CP}^1$. Now we obtain the Chekanov torus in $\mathbb{CP}^2$ (again see the papers of Gadbled [72] and Oakley–Usher [107] for equivalences of various definitions), and, using the methods of Chekanov–Schlenk [30], Auroux [9, Proposition 5.8] computed the superpotential of this torus. This was also computed by Wu [145] using different methods, and can again be calculated by wall-crossing from the Clifford torus. Just as for the torus in $\mathbb{CP}^1 \times \mathbb{CP}^1$ there are five index 2 discs through a generic point, this time in classes D_1, $H - 2D_1 - D_2$, $H - 2D_1$, $H - 2D_1$, $H - 2D_1 + D_2$, where H is the hyperplane class and D_1 and D_2 are classes of discs whose boundaries form a basis of $H_1(L')$, and all count negatively by the $\mathfrak{m}_0$ equation (5.20). We deduce that the torus, and hence L', is non-narrow if the boundary sum $7\partial D_1$ vanishes, i.e. if the coefficient ring has characteristic 7.

Remark 5.2.15. The non-monotone tori obtained by taking other values of r (half the area scaling applied to the $\mathbb{CP}^1$ factor) have been studied by Tonkonog–Vianna using their low-area Floer theory [138]. They conjecture that these tori are non-displaceable for $r \in (0, 1/3]$ and prove that they are non-displaceable from the Clifford torus for $r \in (0, 1/9]$. //

5.2.4 The Amarzaya–Ohnita–Chiang family and periodicity

In [7], Amarzaya and Ohnita introduced, for each $N \geq 2$, a monotone homogeneous Lagrangian in $\mathbb{CP}^{N^2-1}$, which can be described as the projectivisation of the orbit of the identity under the left-multiplication action of $\mathrm{SU}(N)$ on the space of $N \times N$ complex matrices. This family was later rediscovered by Chiang [31, Section 4]. This orbit is clearly diffeomorphic to $\mathrm{PSU}(N)$, and we denote it by L_{AOC}. Iriyeh [91] proved that over a field of characteristic 2 L_{AOC} is wide if N is a power of 2 and narrow if N is odd, and when $N = 3$ L_{AOC} is non-narrow over $\mathbb{Z}$. More generally, Evans and Lekili [44, Example 7.2.3] showed that L_{AOC} is wide over a field of characteristic p when N is a power of p, and narrow if N is not divisible by p. Their arguments work with any relative spin structure and rank 1 local system. Recently this family has also been studied by Torricelli [139].

In the present subsection we deal with the case when N is not a prime power. In this situation we shall show that L_{AOC} is narrow over any field, again with any relative spin structure and rank 1 (composite) local system. This dichotomy between wideness when N is a prime power and narrowness otherwise closely resembles the behaviour of our main family of examples, and in both cases the result can be regarded as a consequence of divisibility properties of binomial or multinomial coefficients. However, it is interesting to note that the indices N differ by 1 for the two families: for our main family it is the $\mathrm{PSU}(p^r - 1)$'s which are wide whereas in the Amarzaya–Ohnita–Chiang family it is the $\mathrm{PSU}(p^r)$'s. The two families are connected by a third family which we discuss in Section 6.2.2.

We present three approaches to proving narrowness, of decreasing explicitness but increasing generality. The final method readily extends to other situations, and gives another instance of symmetry in Floer theory, namely periodicity properties of the Betti numbers of wide Lagrangians.

Remark 5.2.16. Periodicity in gradings has famously been used to constrain Lagrangian embeddings before, as in Seidel's work on graded Lagrangians [120, Section 3]. Related ideas, including the use of invertible elements in quantum cohomology, appear in [14] (for example Corollary 6.2.1), [16] and other papers by Biran. //

Fundamental to all of the approaches is the following result [10, Corollary 4.2]: if R is a field of characteristic p and $N = p^r q$, with q an integer coprime to p, then

$$H^*(\mathrm{PSU}(N); R) \cong \Lambda(x_1, \ldots, \widehat{x}_{2p^r-1}, \ldots, x_{2N-1}) \otimes R[y]/(y^{p^r}), \tag{5.26}$$

where the Λ denotes the exterior algebra over R generated by elements x_{2j-1} of degree $2j - 1$; if $p = 2$ and $r = 1$ then instead of $x_1^2 = 0$ we have $x_1^2 = y$. For $p \neq 2$ this can be easily deduced (using graded-commutativity) from our earlier discussion of the cohomology of $\mathrm{PSU}(N)$ via the Serre spectral sequence for a principal $\mathrm{U}(N)$-bundle over $\mathbb{CP}^\infty$, whilst for $p = 2$ one can use a Hopf algebra structure theorem of Borel [21] (see [27, Theorem C, page 19]) along with the interpretation of $x_1^2 = \mathrm{Sq}^1 x_1$ in terms of a Bockstein homomorphism. Since the Lagrangian L_{AOC} has minimal Maslov index equal to $2N$ (twice the minimal Chern number of $\mathbb{CP}^{N^2-1}$ divided by the order of $\pi_1(L)$), we immediately see from Proposition 2.2.2 that L_{AOC} is either wide or narrow, and is wide if $N = p^r$. For the remainder of the subsection we fix a choice of relative spin structure and rank 1 composite local system on our Lagrangian L_{AOC}, making it into a monotone Lagrangian brane $L^\flat_{\mathrm{AOC}}$.

The first approach is a fairly explicit computation. Consider the Oh spectral sequence for L_{AOC} over a field R of prime characteristic p. Following the argument of Proposition 2.2.2, narrowness is equivalent to non-vanishing of the differential

$$\mathrm{d}_1 \colon H^{2N-1}(\mathrm{PSU}(N); R) \to T \cdot H^0(\mathrm{PSU}(N); R),$$

and analogously to Example 2.2.4 this differential is given by pairing with the homology class swept by the boundaries of the index $2N$ discs through a generic point of L_{AOC}, twisted by the local system (and multiplied by $T \cdot 1$); since $2N$ is the minimal Maslov index there can be no bubbling. In other words, to prove narrowness it is sufficient (and, indeed, necessary) to demonstrate the non-vanishing of this homology class. Since all of the index $2N$ discs are homologous, they are all twisted equally by the local system, and hence this twisting makes no difference to whether or not the required class vanishes.

Now, the moduli space of index $2N$ discs through a point of L_{AOC} is diffeomorphic to $\mathbb{CP}^{N-1}$, by identifying a disc u with the kernel of the matrix $u(0) \in \mathbb{CP}^{N^2-1}$. The homology class of interest is then the pushforward of the fundamental class of $S^1 \times \mathbb{CP}^{N-1}$ under the evaluation map at the boundary marked point parametrised by the S^1. This evaluation map to $L_{\mathrm{AOC}} \cong \mathrm{PSU}(N)$ factors through $\mathrm{U}(N)$, and if $\pi\colon \mathrm{U}(N) \to S^{2N-1}$ denotes projection of a matrix onto its first column we obtain the following diagram

$$\begin{tikzcd} S^1 \times \mathbb{CP}^{N-1} \arrow[r] & \mathrm{U}(N) \arrow[r] \arrow[d, "\pi"] & \mathrm{PSU}(N) \\ & S^{2N-1} & \end{tikzcd}$$

We wish to show that the image of the fundamental class along the horizontal composition is non-zero, or equivalently that the pullback map $H^{2N-1}(\mathrm{PSU}(N)) \to H^{2N-1}(S^1 \times \mathbb{CP}^{N-1})$ is surjective (all cohomology groups taken over R). One can show by counting preimages that the composition $S^1 \times \mathbb{CP}^{N-1} \to S^{2N-1}$ has degree ± 1 and hence $H^{2N-1}(\mathrm{U}(N)) \to H^{2N-1}(S^1 \times \mathbb{CP}^{N-1})$ is surjective. This map factors through the quotient of $H^{2N-1}(\mathrm{U}(N))$ by the subspace generated by products of classes of degrees $2, 3, \dots, 2N-1$, and hence for *any* identification $H^*(\mathrm{U}(N)) \cong \Lambda(a_1, \dots, a_{2N-1})$ we see that a_{2N-1} maps to a generator of $H^{2N-1}(S^1 \times \mathbb{CP}^{N-1})$; in particular this holds for the identification under which a_{2j-1} transgresses to the Chern class c_{2j} in the Serre spectral sequence for $\mathrm{U}(N) \times_{S^1} ES^1 \simeq \mathrm{PSU}(N)$. And when N is not a power of p we see from the spectral sequence that under this identification the projection $\mathrm{U}(N) \to \mathrm{PSU}(N)$ (which corresponds to the inclusion of the fibre in the homotopy quotient) sends $x_{2N-1} \in H^{2N-1}(\mathrm{PSU}(N))$ to $a_{2N-1} \in H^{2N-1}(\mathrm{U}(N))$. Putting everything together, x_{2N-1} maps to a generator of $H^{2N-1}(S^1 \times \mathbb{CP}^{N-1})$ so $H^{2N-1}(\mathrm{PSU}(N)) \to H^{2N-1}(S^1 \times \mathbb{CP}^{N-1})$ is surjective, and hence $L^\flat_{\mathrm{AOC}}$ is narrow.

The second approach is to consider the ring homomorphism

$$\mathcal{CO}^0\colon QH^*(\mathbb{CP}^{N^2-1}; \Lambda_R) \to HF^*(L^\flat_{\mathrm{AOC}}, L^\flat_{\mathrm{AOC}}; \Lambda_R),$$

where Λ_R is the Novikov ring $R[T^{\pm 1}]$ over R in the variable T of degree $2N$. Since the minimal Maslov index is $2N$, any Morse cocycle on L_{AOC} of degree (meaning the Morse degree, ignoring any powers of the Novikov variable) less than $2N-1$ is also a cocycle in the pearl complex, and Floer products of total degree less than $2N-1$ coincide with the corresponding Morse products. In particular there is a class $y_F \in HF^2$ corresponding to the class y in (5.26) and if N is not a power of p (so y^{p^r} has degree less than $2N-1$) then it satisfies $y_F^{p^r} = 0$. For degree reasons, we must have

$$\mathcal{CO}^0(H) = \lambda y_F$$

for some $\lambda \in R$, where H is the hyperplane class on $\mathbb{CP}^{N^2-1}$, and then

$$tT^N \cdot 1_{L_{\mathrm{AOC}}} = t\,\mathcal{CO}^0(T^N \cdot 1_X) = \mathcal{CO}^0(H^{N^2}) = y_F^{N^2} = 0,$$

where $t \in R^\times$ depends on the H_2^D local system and the background class of the relative spin structure. Therefore the unit in HF^* vanishes and hence L_{AOC} is narrow. The same argument would apply to any Lagrangian embedding of $\mathrm{PSU}(N)$ of minimal Maslov index $2N$ in a closed monotone symplectic manifold whose classical cohomology contains a degree 2 class whose p^rth power is invertible under the quantum product.

Remark 5.2.17. In fact, on classes of degree less than $2N$ the map $\mathcal{CO}^0$ agrees with the classical pullback $H^*(X) \to H^*(L_{\mathrm{AOC}})$, so the coefficient λ above is actually equal to 1. To see this note that by construction y is (minus) the first Chern class of the pullback of the tautological line bundle on $BS^1 = \mathbb{CP}^\infty$ to $\mathrm{U}(N) \times_{S^1} ES^1$. This can be obtained from the trivial line bundle on $\mathrm{U}(N) \times ES^1$ by identifying the fibres over a point x and its translate $e^{i\theta}x$ by multiplication by $e^{-i\theta}$, and hence can also be viewed as the pullback from $\mathrm{PSU}(N)$ of the line bundle whose fibre over A is the span of A in $\mathbb{C}^{N^2}$. In other words y is precisely the pullback of (minus) first Chern class of the tautological line bundle on $\mathbb{CP}^{N^2-1}$, i.e. H. //

We now describe the third approach. Let X be closed monotone symplectic manifold of minimal Chern number N_X. Its quantum cohomology is $\mathbb{Z}/2N_X$-graded. Let $A \subset \mathbb{Z}/2N_X$ be the subgroup comprising the degrees of homogeneous invertible elements, and let $\widehat{\psi} \in \{1, 2, \ldots, 2N_X\}$ be the minimal generator of A.

Example 5.2.18. Complex projective space and quadrics of even complex dimension [132, Lemma 4.3] have $\widehat{\psi} = 2$. //

Remark 5.2.19. A rich source of invertible elements in quantum cohomology is provided by Seidel's construction [119] from loops of Hamiltonian diffeomorphisms. //

Now suppose $L^\flat \subset X$ is a monotone Lagrangian brane over a compatible ring R, with local system of rank 1 and minimal Maslov index N_L. The local system may be composite, in which case the quantum cohomology should be the correspondingly twisted version, but we suppress explicit mention of this. Consider the quantum cohomology of X with grading reduced modulo N_L, and let $\psi \in \{1, 2, \ldots, N_L\}$ denote the minimal generator of the subgroup of $\mathbb{Z}/N_L$ comprising the degrees of homogeneous invertible elements in this new ring. Note that this number must divide both $\widehat{\psi}$ and N_L. The Floer cohomology $HF^*(L^\flat, L^\flat; \Lambda_R)$ over the Novikov ring $\Lambda_R = R[T^{\pm 1}]$, where T has degree N_L, is $\mathbb{Z}/N_L$-graded but the quantum module action of invertible elements in $QH^*(X; \Lambda_R)$ means that it is ψ-periodic in this grading. In particular, if R is a field and $L^\flat$ is wide over R then the quantity

$$B_k \coloneqq \sum_{j \in k + N_L\mathbb{Z}} b_j(L; R) = \dim_R HF^k(L^\flat, L^\flat; \Lambda_R),$$

where b_j denotes the jth Betti number of L, is ψ-periodic in k. In the ring $\mathbb{Z}[S]/(S^{N_L} - 1)$ we therefore have

$$\sum_{j=0}^{\dim L} S^j b_j(L; R) = \sum_{k=1}^{N_L} S^k B_k = \sum_{k=1}^{\psi} \sum_{l=0}^{N_L/\psi - 1} S^{k + l\psi} B_{k + l\psi} = \left(\sum_{k=1}^{\psi} S^k B_k \right) \left(\sum_{l=0}^{N_L/\psi - 1} S^{l\psi} \right) \tag{5.27}$$

Taking $S = 1$ we see that N_L/ψ divides $\dim H^*(L; R)$, whilst taking S to be a primitive N_Lth root of unity the sum over l vanishes unless $\psi = N_L$.

However, the left-hand side of (5.27)—the Poincaré polynomial of L (reduced modulo $S^{N_L} - 1$)—has a property which makes it easy to compute, or at least see that it is non-zero: it is multiplicative under tensor product decompositions of the cohomology of L. We have therefore proved:

Proposition 5.2.20. *Suppose X is a closed monotone symplectic manifold and $L^\flat \subset X$ is a wide monotone Lagrangian brane over a compatible field R, with rank 1 local system and minimal Maslov index N_L. Let $\psi \in \{1, 2, \ldots, N_L\}$ be minimal such that there exists a homogeneous invertible element in $QH^*(X; \Lambda_R)$—with grading reduced modulo N_L—of degree ψ. If the cohomology $H^*(L; R)$ is additively isomorphic to a tensor product of graded R-vector spaces $V_1^*, \ldots, V_m^*$ then*

$$N_L/\psi \;\Big|\; \prod_k \dim V_l^* \tag{5.28}$$

and if $\psi \neq N_L$ then for some k we have

$$\sum_{j \in \mathbb{Z}} \zeta^j \dim V_k^j = 0, \tag{5.29}$$

where ζ is a primitive N_Lth root of unity.

From this we can deduce the following, which immediately gives the claimed narrowness result:

Corollary 5.2.21. *If there exists a Lagrangian $L \cong \mathrm{PSU}(N)$ of minimal Maslov index $2N$ (with $N \geq 2$) in a closed monotone symplectic manifold X, whose quantum cohomology over a field R contains an invertible element of degree 2, then either:*

(i) L is narrow over R for all choices of spin structure and rank 1 local system.

(ii) R has prime characteristic p and N is a power of p.

Proof. First note that L automatically inherits monotonicity from X since its fundamental group is finite. Supposing (i) doesn't hold, by Proposition 2.2.2 there exists a choice of spin structure and rank 1 local system over R for which L is wide. Therefore Proposition 5.2.20 applies with ψ equal to 1 or 2, since ψ divides $\widehat{\psi} = 2$.

If R has positive characteristic p and $N = p^r q$ then from (5.26) we have a tensor decomposition of $H^*(\mathrm{PSU}(n); R)$, with factors $R[y]/(y^{p^r})$ and $\Lambda(x_{2j-1})$ for $j = 1, \ldots, \widehat{p^r}, \ldots, N$ (this decomposition is only additive if $p = 2$ and $r = 1$). By (5.28) we deduce that

$$p^r q \mid 2^{N-1} p^r,$$

so q divides a power of 2 and hence either (ii) holds (if $q = 1$) or N is even (if $q > 1$). In the latter case we now apply (5.29), and since the sum is non-zero for each of the exterior algebra factors (as N is even) it must vanish for $R[y]/(y^{p^r})$. But this only happens if $N = p^r$, and again we conclude that (ii) holds.

If R instead has characteristic zero then we have a similar decomposition but now just with factors $\Lambda(x_{2j-1})$ for $j = 2, 3, \ldots, N$. This time the divisibility condition gives us that N is a power of 2, so none of the sums (5.29) vanish and we obtain a contradiction. This case is therefore impossible. □

Remark 5.2.22. This result extends to relative spin structures if the quantum cohomology is deformed by the background class. //

Remark 5.2.23. In contrast to the second approach, where we required the existence of a certain class of degree 2 in the $\mathbb{Z}$-graded classical cohomology, the above argument only refers to a class having degree 2 in the $\mathbb{Z}/2N_X$-graded quantum cohomology (and really could be phrased in terms of the $\mathbb{Z}/N_L = \mathbb{Z}/2N$-graded quantum cohomology). //

We can also give a simple proof (which is probably well-known to experts) of the following:

Corollary 5.2.24. *If a closed symplectic manifold X contains a monotone Lagrangian torus L of minimal Maslov index $N_L > 2$ then in the quantum cohomology of X with grading collapsed modulo N_L all homogeneous invertible elements lie in degree 0. In particular if the $\mathbb{Z}/2N_X$-graded quantum cohomology contains an invertible element of degree 2 then $N_L = 2$.*

Proof. Suppose we have such an $L \cong T^n \subset X$. By Proposition 2.2.2 L is wide over any field R and with any spin structure and rank 1 local system. The cohomology of L over R factorises as a tensor product of n copies of an exterior algebra on a degree 1 generator, and applying (5.29) of Proposition 5.2.20 to this decomposition we see that ψ must be equal to N_L, which is exactly what we want. □

Remark 5.2.25. Using his theory of Floer (co)homology on the universal cover, Damian [38, Theorem 1.6] proved that $N_L = 2$ for all monotone orientable aspherical Lagrangians L in the product of $\mathbb{CP}^n$ (for $n \geq 1$) with an arbitrary symplectic manifold W. Fukaya [59, Theorem 6] obtained a similar result, without monotonicity, for aspherical spin Lagrangians in $\mathbb{CP}^n$ and other uniruled symplectic manifolds. //

Remark 5.2.26. Corollary 5.2.21 and Corollary 5.2.24 hold without assuming the existence of the invertible element in quantum cohomology if one can ensure 2-periodicity of Floer cohomology for some other reason. This holds, for example, if X is a symplectic hyperplane section whose complement is subcritical, at least in characteristic 2; see [18, Corollary 1.3], where the relevant definitions are also given. //

5.2.5 Flag varieties and a conjecture of Biran–Cornea

A closed, connected, monotone symplectic manifold X is called *point invertible* over a field R if the class Poincaré dual to a point is invertible in $QH^*(X; R)$. Biran and Cornea conjectured [16, Conjecture 2] in characteristic 2 that any pair L_0, L_1 of non-narrow monotone Lagrangians in a point invertible manifold X intersect. They proved this conjecture when X is complex projective space [16, Corollary 1.2.8] using spectral invariants.

They also remarked [16, Section 6.4.1] that $\mathbb{CP}^2 = \mathrm{Sym}^2\,\mathbb{CP}^1$ contains the monotone Lagrangian antidiagonal $\mathbb{RP}^2$ and the monotone Lagrangian circle bundle over the equator on the diagonal quadric, which are disjoint. In characteristic 2 the former is wide, so they deduce that the latter—which is the Chekanov torus (discussed in Section 5.2.1–Section 5.2.2)—must be narrow. However, we saw that this torus is actually wide in characteristic 7, where $\mathbb{RP}^2$ becomes narrow. This demonstrates that their spectral invariants must depend on the choice of coefficient ring.

In fact, in $\mathbb{CP}^1 \times \mathbb{CP}^1$ the antidiagonal sphere and the Chekanov torus provide a counterexample to their conjecture in characteristic 3. In the language of Fukaya categories, the antidiagonal split-generates the 0-eigenvalue summand (by [1], since $\mathcal{CO}^0$ for the antidiagonal is injective on this summand) whilst the Chekanov torus has $\mathfrak{m}_0 = 5$ so modulo 3 it lies in the 2-eigenvalue summand. We can lift this example to $(\mathbb{CP}^1)^3$ by taking L_0 to be our PSU(2)-homogeneous Lagrangian (whose symplectic reduction gives the Chekanov torus) and L_1 to be the product of the antidiagonal with the equator in $\mathbb{CP}^1$. These are disjoint since the first two components of a point in L_1 are orthogonal (as lines in $\mathbb{C}^2$, which corresponds to being antipodal on $\mathbb{CP}^1$) whilst those of a point in L_0 are not. We now describe how to generalise this construction to obtain counterexamples in all positive characteristics.

Remark 5.2.27. The cubic surface, which is the blowup of $\mathbb{CP}^2$ at six points, provides another source of counterexamples, as pointed out to the author by Dmitry Tonkonog. It contains disjoint Lagrangian spheres, which are trivially wide in any characteristic, and is point invertible in every characteristic except 2 (see the description of the quantum cohomology in [130, Appendix B.2]). Combining this with our counterexamples we see that the conjecture is violated in all characteristics. Other counterexamples may be well-known to experts in the field. //

Take N to be a power of a prime p and let $L_0^\flat$ be our $\mathrm{PSU}(N-1)$-homogeneous Lagrangian in $(\mathbb{CP}^{N-2})^N$, equipped with the standard spin structure and trivial local system. By Corollary 5.1.20 this is wide in characteristic p. Now take $L_1 \subset (\mathbb{CP}^{N-2})^N$ to be the product of L' in $(\mathbb{CP}^{N-2})^{N-1}$ with the Clifford torus in the final $\mathbb{CP}^{N-2}$ factor, where L' is defined by

$$L' = \{([z_1], \dots, [z_{N-1}]) : \text{the } z_j \text{ are pairwise orthogonal as vectors in } \mathbb{C}^{N-1}\}.$$

This is Lagrangian (it is contained in the zero level set of the moment map of the obvious $\mathrm{PSU}(N-1)$-action) and is easily seen to be diffeomorphic to the variety F of complete flags in $\mathbb{C}^{N-1}$.

From the long exact sequence for the fibration $T^{N-1} \hookrightarrow \mathrm{U}(N-1) \twoheadrightarrow F$ we see that L' is simply connected, and hence monotone and orientable. If π denotes the projection $\mathrm{U}(N-1) \to F$ and $\mathfrak{h} \subset \mathfrak{u}(n-1)$ is the complementary subspace to $\mathfrak{t}^{N-1}$ given by skew-Hermitian matrices with vanishing diagonal entries, then the infinitesimal action of $\mathfrak{h}$ on $\mathrm{U}(N-1)$ on the right gives a trivialisation of π^*TF, which is the quotient of $T\mathrm{U}(N-1)$ by the subbundle generated by the infinitesimal T^{N-1}-action on the right. The T^{N-1}-action on $\mathrm{U}(N-1)$ induces an action on this quotient $\pi^*TF = \mathrm{U}(N-1) \times \mathfrak{h}$ which corresponds to the conjugation action on $\mathfrak{h}$. We claim that there exists a spin structure on π^*TF such that this action lifts to the corresponding $\mathrm{Spin}(\mathfrak{h})$-bundle, and hence that the spin structure descends to F, so F is spin.

Well, with respect to our chosen trivialisation of π^*TF we can express any spin structure on it by an element $\varepsilon \in H^1(\mathrm{U}(N-1); \mathbb{Z}/2)$. Similarly we have an element ε' in

$$H^1(T^{N-1}; \mathbb{Z}/2) = \mathrm{Hom}\big(\pi_1(T^{N-1}), \pi_1(\mathrm{SO}(\mathfrak{h}))/\pi_1(\mathrm{Spin}(\mathfrak{h}))\big),$$

defined by taking a loop in T^{N-1}, acting by it on $\mathfrak{h}$ to get a loop in $\mathrm{SO}(\mathfrak{h})$, and then seeing whether the lift to $\mathrm{Spin}(\mathfrak{h})$ closes up. The spin structure defined by ε descends to the quotient if and only if the pullback of ε under the inclusion $T^{N-1} \hookrightarrow \mathrm{U}(N-1)$ coincides with ε'. It is straightforward to compute that under the obvious identification $H^1(T^{N-1}; \mathbb{Z}/2) \cong (\mathbb{Z}/2)^{N-1}$ the class ε' is given by $(0, 0, \dots, 0)$ if N is even and $(1, 1, \dots, 1)$ if N is odd, and that the pullback of ε is—in a slight abuse of notation—just $(\varepsilon, \varepsilon, \dots, \varepsilon)$. So for any N there is a (unique) choice of ε such that the corresponding spin structure descends to F.

Putting all of this together, we conclude that L' is a closed, connected, monotone, orientable and spin Lagrangian in $(\mathbb{CP}^{N-2})^{N-1}$, and hence that its product L_1 with the Clifford torus is a Lagrangian in $(\mathbb{CP}^{N-2})^N$ that also has these properties. Moreover, since L' is simply connected, a choice of spin structure on L_1 really amounts to a choice of spin structure on the Clifford torus. Choosing the standard spin structure, we now equip L_1 with the trivial local system to give a brane $L_1^\flat$ whose Floer cohomology is defined over any ring. We saw in Lemma 5.1.15 that the classical cohomology of F is concentrated in even degree so from the Oh spectral sequence L' is trivially wide over any field. Combining this with the known computation for the Clifford torus, we see that $L_1^\flat$ is also wide over any field.

We therefore have two monotone Lagrangian branes $L_0^\flat$ and $L_1^\flat$ in $(\mathbb{CP}^{N-2})^N$ which are both wide in characteristic p (recalling that N is assumed to be a power of p). However, they are disjoint since the first $N-1$ components of L_1 are pairwise orthogonal whereas those of L_0 are not. Thus we have the claimed counterexample. In the case $N = 3$, L' is simply the antidiagonal sphere in $\mathbb{CP}^1 \times \mathbb{CP}^1$ so we recover our earlier example.

Remark 5.2.28. We can actually read off the product structure on $HF^*(L', L'; \mathbb{Z})$ from the closed–open map. To simplify notation, for the remainder of the discussion of this example we relabel $N-1$ as N, so that $L' \cong \mathrm{PSU}(N)$. We know from the proof of Lemma 5.1.15 that the classical cohomology ring is isomorphic to $\mathbb{Z}[T_1, \dots, T_N]/(e_1, \dots, e_N)$, where T_j is the first Chern class of the jth tautological line bundle. We can also express this class as the pullback of $-H_j$ under our embedding of L' in $(\mathbb{CP}^{N-1})^N$, where H_j is the hyperplane class on the jth factor. Moreover, we know that additively HF^* is isomorphic to this ring, and since the minimal Maslov index of L' is $2N$ there is a canonical injection $H^{<2N-1} \hookrightarrow HF^*$ and algebraic operations in HF^* agree with their classical counterparts up to corrections of Morse degree at least $2N$ lower. In particular, we see that HF^* is generated as a ring by classes $T_j = \mathcal{CO}^0(-H_j)$ which satisfy $e_1 = \dots = e_{N-1} = 0$.

As formal polynomials in z over HF^* we then have

$$\prod_{j=1}^{N} (z - T_j) = z^N + (-1)^N e_N.$$

Setting $z = T_1$, and using the quantum cohomology relation $H_1^N = 1$, we deduce that $e_N = -1$. Therefore as a ring HF^* is a quotient of $\mathbb{Z}[T_1, \dots, T_N]/(e_1, \dots, e_{N-1}, e_N + 1)$. But both are free $\mathbb{Z}$-modules of rank $N!$ and so they must in fact be equal. In other words

$$HF^*(L', L'; \mathbb{Z}) \cong \mathbb{Z}[T_1, \dots, T_N]/(e_1, \dots, e_{N-1}, e_N + 1)$$

as $\mathbb{Z}/2N$-graded rings, with each T_j in degree 2. //

Up to rescaling the symplectic form, L' is the symplectic reduction of $L_{\mathrm{AOC}} \subset \mathbb{CP}^{N^2-1}$ under the Hamiltonian T^{N-1}-action by $N \times N$ diagonal matrices of determinant 1 on the right. One can similarly consider the reduction by other subgroups of $\mathrm{SU}(N)$. In particular, taking the group

$$\mathrm{S}(\mathrm{U}(k_1) \times \cdots \times \mathrm{U}(k_r))$$

of block diagonal matrices ($k_1 + \cdots + k_r = N$) we obtain a Lagrangian embedding of a partial flag variety in a product of Grassmannians. One can easily check that the action of this group on the zero set of its moment map is free. Monotonicity follows from the following:

Lemma 5.2.29. *Suppose X is a symplectic manifold carrying a Hamiltonian action of a Lie group K, and $L \subset X$ is a monotone Lagrangian submanifold which is K-invariant (setwise). Assume moreover that L is contained in the zero level set of the moment map μ, on which K acts freely. Then L/K is monotone for discs as a Lagrangian in the symplectic reduction $X /\!\!/ K$, meaning that the Maslov index and area homomorphisms are positively proportional on $\pi_2(X /\!\!/ K, L/K)$ (for brevity we suppress all mention of base points). In particular, $X /\!\!/ K$ is spherically monotone.*

Remark 5.2.30.
(i) This notion of monotonicity is weaker than that defined in Section 2.1.1 but is sufficient for defining self-Floer cohomology.

(ii) If L is connected then the condition that it lies in the zero level set of the moment map is automatic, possibly after shifting μ by a fixed point of the coadjoint action—see [36, Lemma 4.1]. //

Proof. Let K_0 be a K-orbit inside L. Considering the long exact sequences in homotopy groups for the pair (L, K_0) and the fibration $K_0 \hookrightarrow L \twoheadrightarrow L/K$, and applying the five lemma, we see that the obvious projection induces an isomorphism $\pi_j(L, K_0) \to \pi_j(L/K)$ for $j \geq 1$. Similarly we have isomorphisms $\pi_j(\mu^{-1}(0), K_0) \to \pi_j(\mu^{-1}(0)/K = X /\!\!/ K)$ for $j \geq 1$. Now applying the five lemma to the long exact sequences for the pair $(X /\!\!/ K, L/K)$ and the triple $(\mu^{-1}(0), L, K_0)$ we deduce that $p_* \colon \pi_2(\mu^{-1}(0), L) \to \pi_2(X /\!\!/ K, L/K)$ is an isomorphism, where p is the quotient map $(\mu^{-1}(0), L) \to (X /\!\!/ K, L/K)$.

So take an arbitrary class β in $\pi_2(X /\!\!/ K, L/K)$ and choose a disc $u \colon (D, \partial D) \to (\mu^{-1}(0), L)$ with $[p \circ u] = \beta$. Since $p^*\omega_{X/\!\!/K} = \omega_X|_{\mu^{-1}(0)}$ we immediately have that u and β have equal areas. We now show that they have equal Maslov indices, so that L/K inherits monotonicity for discs from L.

Picking a compatible almost complex structure J' on X, we obtain a short exact sequence of complex vector bundles on $\mu^{-1}(0)$

$$0 \to (\mathfrak{k} \cdot \mu^{-1}(0)) \oplus J'(\mathfrak{k} \cdot \mu^{-1}(0)) \xrightarrow{\iota} TX|_{\mu^{-1}(0)} \to \operatorname{coker} \iota \to 0.$$

Combining this with the corresponding short exact sequence of totally real subbundles over L

$$0 \to \mathfrak{k} \cdot L \to TL \to TL/(\mathfrak{k} \cdot L) \to 0$$

we can express the Maslov index of the pair $u^*(TX, TL)$ (and hence of the disc u) as the sum of the indices of the pairs $u^*(\operatorname{coker}\iota, TL/(\mathfrak{k}\cdot L))$ and

$$u^*\big((\mathfrak{k}\cdot\mu^{-1}(0))\oplus J'(\mathfrak{k}\cdot\mu^{-1}(0)), \mathfrak{k}\cdot L\big).$$

The latter pair is trivial and hence has index 0, so it is enough to show that the indices of $(\operatorname{coker}\iota, TL/(\mathfrak{k}\cdot L))$ and β coincide.

To see that this is indeed the case, note that the index of β is that of the pair $(p\circ u)^*(T(X /\!\!/ K), T(L/K))$, which is canonically identified with

$$u^*\big(T\mu^{-1}(0)/(\mathfrak{k}\cdot\mu^{-1}(0)), TL/(\mathfrak{k}\cdot L)\big).$$

There is a natural isomorphism of real vector bundles $\operatorname{coker}\iota \cong T\mu^{-1}(0)/(\mathfrak{k}\cdot\mu^{-1}(0))$, so we will be done if we can show that the complex structure on the former is compatible with the symplectic structure on the latter. But the subbundle $T\mu^{-1}(0)\cap J'(T\mu^{-1}(0))$ of $TX|_{\mu^{-1}(0)}$—which is both complex and symplectic—projects isomorphically to $\operatorname{coker}\iota$ and $T\mu^{-1}(0)/(\mathfrak{k}\cdot\mu^{-1}(0))$, and, moreover, the complex and symplectic structures on these two quotients are those induced by these projections, and hence are compatible. This shows that the index of β is that of the pair $(\operatorname{coker}\iota, TL/(\mathfrak{k}\cdot L))$, completing the proof. □

These partial flag varieties are simply connected (hence orientable) and have cohomology concentrated in even degree so all differentials in the Oh spectral sequence necessarily vanish, but note that not all of them are spin, so some care is needed outside characteristic 2. In fact, it is easy to generalise the argument given above for the complete flag variety to see that the partial flag variety is spin if and only if all k_j have the same parity. For example, in the case $(k_1, k_2) = (N-1, 1)$ we obtain the standard embedding of $\mathbb{CP}^{N-1}$ (which is spin if and only if N is even) as a Lagrangian in $\mathrm{Gr}(N-1, N)\times\mathrm{Gr}(1, N)\cong(\mathbb{CP}^{N-1})^-\times\mathbb{CP}^{N-1}$.

Remark 5.2.31. Generalising Remark 5.2.28, the product on the self-Floer cohomology of each of these Lagrangians—which we shall denote by $L_{\mathbf{k}} = L_{(k_1,\dots,k_r)}$—is determined by the quantum cohomology of the ambient Grassmannians. First note that $L_{\mathbf{k}}$ is diffeomorphic to G/H where $G = \mathrm{U}(N)$ and $H = \mathrm{U}(k_1)\times\cdots\times\mathrm{U}(k_r)$. We thus have a fibration $BG \twoheadrightarrow BH$ with fibre homotopy equivalent to $L_{\mathbf{k}}$. Applying the Serre spectral sequence (essentially in the form of the Leray–Hirsch theorem) we see that $H^*(BH;\mathbb{Z})$ is a free $H^*(BG;\mathbb{Z})$-module, and that $H^*(L_{\mathbf{k}};\mathbb{Z})$ is isomorphic as a ring to the quotient of $H^*(BH;\mathbb{Z})$ by the ideal generated by the images of elements of $H^*(BG;\mathbb{Z})$ of positive degree.

Algebraically, this means that $H^*(L_{\mathbf{k}};\mathbb{Z})$ is simply the ring of polynomials in variables $T_1,\dots,T_N$ (of degree 2) which are invariant under the permutation group $S_{k_1}\times\cdots\times S_{k_r}$, modulo the ideal generated by the fully symmetric polynomials of positive degree [20, Section 31]. Writing $e_1,\dots,e_N$ for the elementary symmetric polynomials in $T_1,\dots,T_N$, and $e^j_1,\dots,e^j_{k_j}$ for the elementary symmetric polynomials in $T_{k_1+\dots+k_{j-1}+1},\dots,T_{k_1+\dots+k_j}$, we can express this as

$$\mathbb{Z}[e^1_1,\dots,e^1_{k_1},\dots,e^r_1,\dots,e^r_{k_r}]/(e_1,\dots,e_N).$$

Geometrically, if the tautological bundles on $L_{\mathbf{k}}$ are denoted by $E_1,\dots,E_r$, of rank $k_1,\dots,k_r$

respectively, then e^j_l represents the lth Chern class of E_j, whilst $e_1, \dots, e_N$ are the Chern classes of the trivial bundle $E_1 + \dots + E_r$.

Now pass to the quantum world. The Grassmannian $\mathrm{Gr}(k_j, N)$ has minimal Chern number N and its quantum cohomology ring was computed by Witten in [144, Section 3.2] to be freely generated by the Chern classes of E_j and

$$F_j \coloneqq E_1 + \dots + \widehat{E}_j + \dots + E_r,$$

modulo $c(E_j)c(F_j) = 1 + (-1)^{k_j} T$, where T is the Novikov variable of degree $2N$. Note that Witten's relation is in terms of the duals of these bundles so comes with a different sign.

Arguing as in Remark 5.2.28 and applying $\mathcal{CO}^0$ to this relation, we see that as a ring $HF^*(L_{\mathbf{k}}, L_{\mathbf{k}}; \mathbb{Z})$ is a quotient of

$$\mathbb{Z}[e^1_1, \dots, e^1_{k_1}, \dots, e^r_1, \dots, e^r_{k_r}]/(e_1, \dots, e_N - (-1)^{k_j}).$$

But this is a free $\mathbb{Z}$-module of the correct rank, so in fact no further quotienting is necessary. Of course this is only consistent if the k_j all have the same parity, but this is exactly the condition for $L_{\mathbf{k}}$ to be spin. If the k_j have different parities then we must work in characteristic 2, and the result then becomes consistent again. //

Similarly one can consider the symplectic reduction of our $\mathrm{PSU}(N-1)$-homogeneous Lagrangian in $(\mathbb{CP}^{N-2})^N$ by block diagonal subgroups of $\mathrm{SU}(N-1)$. Using (5.3), we see that the action of such a subgroup on the zero set of its moment is again free, so these reductions make sense. Now we obtain Lagrangian embeddings of flag varieties in symplectic manifolds which are much less easy to describe. Even in the simplest case, when $N = 3$ and we reduce by the subgroup $\mathrm{S}(\mathrm{U}(1) \times \mathrm{U}(1))$ representing rotations of $\mathbb{CP}^1$ about a vertical axis, the resulting symplectic manifold is the blowup of $\mathbb{CP}^2$ at three points (or equivalently the blowup of $\mathbb{CP}^1 \times \mathbb{CP}^1$ at two points), as illustrated in [5, Figure 2], and the reduced Lagrangian is a 'ternary' sphere [43, Definition 2.4] in the homology class $H - (E_1 + E_2 + E_3)$.

Up to scaling, the moment polytope is the hexagon defined by $-1 \leq x, y, x + y \leq 1$ in $\mathbb{R}^2$, where a triple $([z_1], [z_2], [z_3]) \in \mu^{-1}(0) \subset (\mathbb{CP}^1)^3$ maps to the heights x and y of $[z_1]$ and $[z_2]$, thought of as points on the unit sphere in $\mathbb{R}^3$. The Lagrangian upstairs comprises those triples forming the vertices of an equilateral triangle on a great circle, and for a fixed value of x the height y is extremised when the triangle lies in a vertical plane. In this case we have $x = \cos\theta$ and $y = \cos(\theta \pm 2\pi/3)$ for some angle θ, so $(y + x/2)^2 = 3(1 - x^2)/4$. From this we see that the reduced Lagrangian projects to the solid ellipse $x^2 + xy + y^2 \leq 3/4$ which is tangent to each edge of the hexagon.

For general N, the $\mathrm{S}(\mathrm{U}(1) \times \dots \times \mathrm{U}(1))$ reduction of our $L \cong \mathrm{PSU}(N-1) \subset (\mathbb{CP}^{N-2})^N$ yields a Lagrangian complete flag variety inside the Fano toric variety with moment polytope given by those $(N-1) \times N$ matrices (in $\mathbb{R}^{(N-1)N}$) whose entries are each at least -1, whose rows and columns each sum to at most 1, and all of whose entries sum to at least -1.

Chapter 6

Future directions

In this chapter we outline some ideas for future research related to the ideas of this thesis. It should all be seen as work in progress, with many details left to iron out, and the results stated here should not be considered fully proved.

6.1 A quantum Pontryagin module structure

6.1.1 The construction

Recall that if K is a topological group then the multiplication $m\colon K \times K \to K$ induces a map $H_*(K) \otimes H_*(K) \to H_*(K)$, making the homology of K into an associative (but not, in general, graded-commutative) ring: the Pontryagin ring of K. The homology of any space L on which K acts then inherits the structure of a module over this ring, from the map $H_*(K) \otimes H_*(L) \to H_*(L)$ induced by the action $K \times L \to L$. We denote these ring and module products by $\odot$.

If L is actually a Lagrangian submanifold of a symplectic manifold X, and K acts in a way which respects a compatible almost complex structure J', then heuristically we should be able to make the same construction on a singular model for the Floer complex of L (as used in [63] for example). Roughly, the Floer differential $\mathrm{d}\sigma$ of a singular chain σ is its usual boundary $\partial\sigma$ plus the sum of the chains swept out by the boundaries of J'-holomorphic discs meeting σ. If we now take a chain κ on K, and consider the Floer differential of the chain $\kappa \odot \sigma$ (which is the union of the κ-translates of σ), we obtain $\partial(\kappa \odot \sigma)$ plus the unions of the κ-translates of the disc boundaries contributing to $\mathrm{d}\sigma$—this is because we are assuming J' is K-invariant. In other words, up to signs we obtain

$$\mathrm{d}(\kappa \odot \sigma) = (\mathrm{d}\kappa) \odot \sigma + (-1)^{|\kappa|} \kappa \odot (\mathrm{d}\sigma).$$

This means that the Pontryagin module structure descends to the self-Floer homology of L.

Remark 6.1.1. This construction relied on choosing a specific K-invariant almost complex structure, so a priori it is not a purely symplectic invariant. //

To make this precise, let us focus on the rather more restricted setting in which K is a compact, connected Lie group and $L^\flat \subset X$ is a K-homogeneous monotone Lagrangian brane over a ring R. We shall construct a chain-level quantum Pontryagin product

$$C_*^M(K; f_1) \otimes C_*(L^\flat; f_2) \to C_*(L^\flat; f_3),$$

of degree 0, where C and C^M denote the pearl and Morse complexes respectively, and the f_j are generic Morse functions. We suppress explicit mention of the metrics we choose, which may vary with the choice of Morse function, the almost complex structure, which will always be the homogeneous integrable J, and the coefficient ring. Note that we have switched from cohomological to homological conventions, so in particular will use *downwards* Morse flows.

This product satisfies a Leibniz rule and an associativity property with the Pontryagin product on K, so makes $HF_*(L^\flat, L^\flat; \Lambda)$ into a graded module over $H_*(K; R)$. Moreover, it is compatible with the energy filtration on the pearl complex, and at lowest order coincides with the classical Pontryagin module structure, so descends to the (homological) Oh spectral sequence and is given by the classical action on the first page $E_1 = H_*(L; R) \otimes \Lambda$. It is also independent of the choices of generic Morse functions and metrics.

Remark 6.1.2. For simplicity we shall assume that the local system on L is trivial, but our discussion extends to the case where the local system $\mathscr{F}$ merely admits a lift of the K-action (which is equivalent to the pullback of $\mathscr{F}$ under an orbit map $K \to L$ being trivial; in particular, it is automatic if K is simply connected), and even to the case of two distinct local systems with this property on the two copies of L. One could allow non-trivial H_2^D local systems as well, since the connectedness of K ensures that relative homology classes of discs are preserved under translation by the action. //

We shall only sketch the basic idea of the construction. In particular, we will not address transversality issues, except to say that they can be dealt with in a similar manner to the Floer product in Section C.4, or signs. One can define orientations on the relevant moduli spaces, analogous to those used by Biran–Cornea in the construction of the pearl complex [17, Appendix A], and verify that these are compatible with taking boundaries, and with the orientations for the classical Pontryagin product. However, this is rather involved, and a simpler approach should be provided by an alternative construction of the quantum Pontryagin structure described in Section 6.2.3 (although we have not yet checked the signs for this alternative method).

Remark 6.1.3. The proof of transversality for generic Morse data is where we use the fact that L is K-homogeneous—i.e. that the K-action on L is transitive—rather than just K-invariant. We can drop this transitivity hypothesis if we instead assume that the partial indices of all holomorphic discs on L are non-negative. In Section 6.3 we introduce a rather different construction, which avoids this issue at the expense of requiring the action to be Hamiltonian, and being far less transparent. //

First we review the Morse-theoretic picture of the classical construction. Our general reference for Morse theory is [97, Section I.2]. The Pontryagin product on $H_*(K)$ is defined as the composition

$$H_*(K) \otimes H_*(K) \xrightarrow{\text{cross product}} H_*(K \times K) \xrightarrow{m_*} H_*(K).$$

At chain level this can be described by counting triples of partial trajectories in K—two outgoing from the input critical points and one incoming to the output critical point—whose ends lie on the graph of m. Explicitly, if i is the inclusion of this graph

$$\begin{aligned} K \times K &\hookrightarrow K \times K \times K \\ (k_1, k_2) &\mapsto (k_1, k_2, k_1 k_2), \end{aligned}$$

and f_1, f_2 and f_3 are generic Morse functions on K, then for x in $\mathrm{Crit}(f_1)$ and y in $\mathrm{Crit}(f_2)$ we have

$$x \odot y = \sum_{\substack{z \in \mathrm{Crit}(f_3) \\ |z| = |x| + |y|}} \#i^{-1}\big(W^d_x(f_1) \times W^d_y(f_2) \times W^a_z(f_3)\big)\ z.$$

Here $W^a_p(f)$ and $W^d_p(f)$ denote the ascending and descending manifolds respectively of a critical point p of a Morse function f, and $\#$ is the (signed) count of points.

We shall represent such trajectories pictorially as shown in Fig. 6.1. The blue blob and arrow

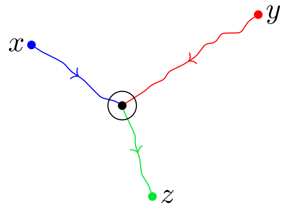

Figure 6.1: Trajectories used to define the Pontryagin product on K.

denote a critical point of f_1 and flowline of $-\nabla f_1$ respectively, whilst red and green are used for f_2 and f_3. The $\odot$ junction represents the incoming paths from x and y arriving at points k_1 and k_2 in K, and the outgoing path to z departing from $k_1 k_2$. This is not to be confused with the junctions of *intersecting* trajectories which define the cup product in cohomology.

By considering the boundaries of compactifications of one-dimensional moduli spaces of trajectories of this form, we see that the product satisfies a Leibniz rule

$$\mathrm{d}(a \odot b) = (\mathrm{d}a) \odot b + (-1)^{|a|} a \odot (\mathrm{d}b).$$

Independence of the auxiliary data is proved using Morse cobordisms, discussed briefly in Section 2.1.2 and described in more detail in Section C.3. Associativity can then be shown by considering trajectories of the form shown in Fig. 6.2, where light blue, orange and brown are used to denote three additional Morse functions. The boundary points of the compactifications

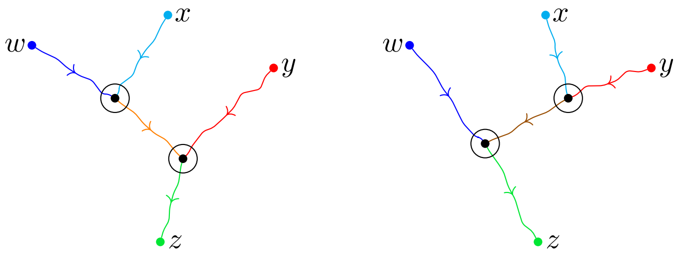

Figure 6.2: Trajectories giving associativity of the Pontryagin product.

of their one-dimensional moduli spaces come in three types: trajectories in which an *external* flowline breaks, giving terms which disappear after passing to homology; those where the *internal* flowline breaks, giving the associativity relation between $(w \odot x) \odot y$ and $w \odot (x \odot y)$; those where the internal flowline shrinks to zero, giving terms which cancel out between the two shapes of diagram.

The module structure on $H_*(L)$ is constructed similarly, but now with the graph of the action $K \times L \to L$. Again we depict the trajectories by Y-shaped trees, where the first incoming leg represents a flowline on K whilst the second incoming leg and the outgoing leg represent flows on L. The central vertex is again marked by a $\odot$, this time denoting a junction of the form $(k, l, kl) \in K \times L \times L$. Independence of auxiliary data and associativity are proved analogously.

To define the quantum Pontryagin module structure using the pearl complex, we count the same configurations but now allow the flowlines on L to be interrupted by holomorphic discs, as shown in Fig. 6.3. The diagram is purely illustrative, and in particular the number of discs can

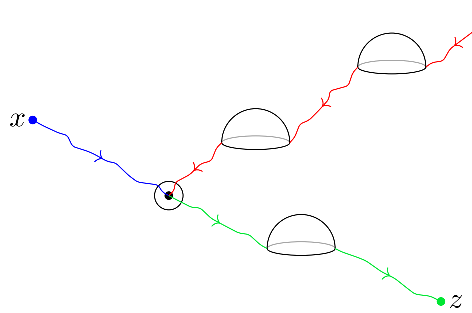

Figure 6.3: Configurations defining the quantum Pontryagin module structure.

be any non-negative integer (including zero, which gives the terms defining the ordinary Morse Pontryagin structure).

The only complication introduced to the classical arguments by this modification is the appearance of new boundary points in the one-dimensional moduli spaces, resulting from discs bubbling into two, or from flowlines between discs—or between a disc and a junction point—shrinking to zero. By standard pearl complex arguments these should all cancel, except perhaps for the shrinking of the flowline between a disc and a junction point. However, if we use the K-invariant integrable complex structure on the ambient Kähler manifold X then such degenerate trajectories again occur in cancelling pairs, since we can use the action of K to transport discs 'across the junction'.

Explicitly, for each configuration in which the flowline between a junction point (k, l, kl) and the final disc u on an incoming leg shrinks to zero, we can form another degenerate configuration by deleting the disc u and adding the disc ku to the start of the outgoing leg (and moving the junction point from (k, l, kl) to (k, l', kl'), where l and l' are the positions of the outgoing and incoming marked points of u respectively), and vice versa, as shown in Fig. 6.4.

6.1.2 An application

We now describe a simple consequence of this structure. As we remarked earlier, a priori the quantum Pontryagin product itself depends on the choice of complex structure so is not necessarily symplectically meaningful. However, we will only use its action on the spectral

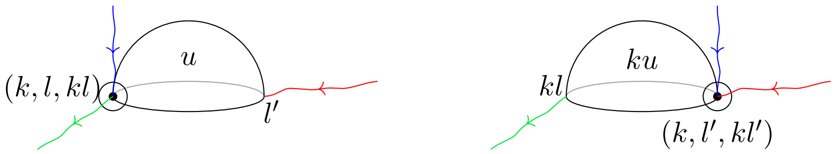

Figure 6.4: Pairing up the novel boundary configurations.

sequence, to constrain the differentials, in a similar way to our use of enlarged coefficient rings in Chapter 2.

Definition 6.1.4. For a field R, we'll say that $H_*(L; R)$ has *maximal unit divisibility* if for any non-zero class $b \in H_*(L; R)$ there exists a class $a \in H_*(K; R)$ such that under the classical Pontryagin module structure we have $a \odot b = [L]$. //

Example 6.1.5. If the coefficient field R has characteristic zero then it is well-known that the Pontryagin ring $H_*(K; R)$ is an exterior algebra [88, Theorem 6.88]. It is not hard to see that in this case any non-zero element divides the top degree element, which means that if L is a free K-orbit then $H_*(L; R)$ has maximal unit divisibility. More generally, if L is sharply K-homogeneous then the orbit map $K \to L$ is a finite cover and hence induces a surjection on homology over R (the map which sends a simplex on L to the average of the lifts of the simplex to K induces a right inverse), so $H_*(L; R)$ still has maximal unit divisibility. //

The significance of this property is that it enables us to establish a wide–narrow dichotomy:

Proposition 6.1.6. *Suppose $L^\flat \subset X$ is K-homogeneous monotone Lagrangian brane over a field R, and assume that the local system has rank 1 and that it carries a lift of the K-action. If $H_*(L; R)$ has maximal unit divisibility then $L^\flat$ is either wide or narrow over R. (Note that by Floer-theoretic Poincaré duality the notions of wide and narrow defined in terms of HF^* are equivalent to those in terms of HF_*.)*

Proof. Consider the homological Oh spectral sequence $H_*(L; R) \otimes \Lambda \implies HF_*(L^\flat, L^\flat; \Lambda)$. As in the proof of Proposition 2.2.2 it suffices to show that on each page of the spectral sequence the differential either vanishes or hits the unit for the cup product, which is the fundamental class $[L]$.

Suppose therefore that the differential vanishes on pages $E_1, \dots, E_{k-1}$—so E_k is isomorphic as an $H_*(K; R)$-module to E_1—but that there exists a non-zero differential $\mathrm{d}_k c = b$ on E_k. We need to show d_k hits the unit. By maximal unit divisibility there exists a in $H_*(K; R) \otimes \Lambda$ with $a \odot b = (-1)^{|a|}[L]$, and we then have $\mathrm{d}_k(a \odot c) = [L]$, so we're done. ☐

To take this further we introduce some new notation. Suppose $L^\natural$ is a monotone Lagrangian pre-brane over a field R and consider the Oh spectral sequence over $R[H_2^D]$. The subalgebra $H^0(L; R) \otimes R[H_2^D]$ of the first page is trivially annihilated by all of the differentials, so it admits a natural map to the limit page. The kernel is of the form $H^0(L; R) \otimes I$ for some ideal I of $R[H_2^D]$, and we denote this ideal by I_L. For each positive integer k let n_k be the dimension of the quotient of $H^{kN_L-1}(L; R)$ by the subspace spanned by cup products of classes of lower degree.

Lemma 6.1.7. *If $L^\natural$ is a K-homogeneous monotone Lagrangian pre-brane over a field R such that $H_*(L;R)$ has maximal unit divisibility, then:*

(i) The differentials all vanish in the Oh spectral sequence over $R[H_2^D]/I_L$.

(ii) I_L is generated by $\sum_k n_k$ elements.

Proof. (i) By construction of I_L, all differentials mapping to the $H^0(L;R) \otimes R[H_2^D]/I_L$ part of each page vanish. Using the Pontryagin module structure as in Proposition 6.1.6 (plus Poincaré duality), *any* non-zero differential gives rise to one which lands in this part.

(ii) On the first page of the spectral sequence over $R[H_2^D]$, the part of the differential which maps to $H^0(L;R) \otimes R[H_2^D]$ comes from $H^{N_L-1}(L;R) \otimes R[H_2^D]$. It vanishes on the subspace generated by cup products of lower degree classes, so its image I_1 is generated as an ideal of $H^0(L;R) \otimes R[H_2^D]$ by n_1 elements.

Now consider the spectral sequence over $R[H_2^D]/I_1$. By the argument used in (i), the differential vanishes on the first page so the second page is simply $H^*(L;R) \otimes R[H_2^D]/I_1$. This time the image of the part of the differential which maps to $H^0(L;R) \otimes R[H_2^D]/I_1$ is generated by n_2 elements. Iterating this argument, we get the result by induction. □

To give a concrete application of this result, first recall the wide variety $\mathcal{W}_2$ from Remark 5.1.23. We have:

Proposition 6.1.8. *If $L^\natural \subset X$ is a sharply K-homogeneous monotone Lagrangian pre-brane over $\mathbb{C}$ then each component of its wide variety $\mathcal{W}_2$ has dimension at least*

$$\operatorname{rank} H_2^D - \sum_k n_k.$$

Proof. Such an L automatically satisfies the hypotheses of the previous result by Example 6.1.5. The (closed) points of the wide variety correspond to maximal ideals $\mathfrak{m}$ in $\mathbb{C}[H_2^D]$ which contain I_L, so its components are cut out inside the $(\operatorname{rank} H_2^D)$-dimensional variety $\operatorname{Spec}\mathbb{C}[H_2^D]$ by minimal prime ideals containing I_L. By [41, Theorem 10.2], the codimension of any such ideal is bounded above by the number of generators of I_L, and thus by $\sum_k n_k$. □

Now consider the main family of examples from Section 5.1. For each $N \geq 3$ we have a freely $\mathrm{PSU}(N-1)$-homogeneous $L^\natural$ in $X = (\mathbb{CP}^{N-2})^N$ (equipped with the standard spin structure), and we saw earlier that H_2^D has rank N. The ring $H^*(L;\mathbb{C})$ is an exterior algebra on $N-2$ odd degree generators, and $N_L = 2$, so the sum of the n_j is exactly $N-2$. By Proposition 6.1.8 we deduce that each component of $\mathcal{W}_2$ has dimension at least 2.

Recall that we identify $\mathbb{Z}[H_2^D]$ with $\mathbb{Z}[T_1, \dots, T_N]$ and coordinatise $\operatorname{Spec}\mathbb{C}[H_2^D]$ by $t_1, \dots, t_N$. In Remark 5.1.23 we computed that $\mathcal{W}_2$ is contained in the subvariety W of $\operatorname{Spec}\mathbb{C}[H_2^D]$ cut out by the ideal I generated by

$$e_k\left(T_1 - \frac{e_1(T_1, \dots, T_N)}{N-1}, \dots, T_N - \frac{e_1(T_1, \dots, T_N)}{N-1}\right) \text{ for } k = 2, \dots, N-1,$$

and contains the subvariety V cut out by the ideal generated by

$$e_k(T_1, \dots, T_N) \text{ for } k = 1, \dots, N-1.$$

We can now show:

Proposition 6.1.9. *$\mathcal{W}_2$ is equal to W.*

Proof. By our dimension bound on components of $\mathcal{W}_2$, it suffices to show that each component of W has dimension 2 and contains a point of V in its interior (meaning the complement of the union of the other components). To see that this is the case note that the map

$$(t_1, \dots, t_N) \mapsto (e_1(t_1, \dots, t_N), e_N(t_1, \dots, t_N))$$

exhibits W as an $N!$-to-1 branched cover of $\mathbb{C} \times \mathbb{C}^*$, and each sheet meets V. $\square$

With some more algebra, we can squeeze further information out of this. First, one can use a Jacobian computation in conjunction with [41, Theorem 18.15] to prove that the ideal I is radical, and hence it is precisely the ideal of functions vanishing on $\mathcal{W}_2$. Consequently, the ideal I_L from earlier—which is contained in every maximal ideal corresponding to a point of the wide variety—is contained in I. Therefore the differentials all vanish in the Oh spectral sequence over $\mathbb{C}[H_2^D]/I$. We will shortly need to consider various changes of the base ring from $\mathbb{C}$, and in each case will let I denote the ideal generated by

$$e_k((N-1)T_1 - e_1(T_1, \dots, T_N), \dots, (N-1)T_N - e_1(T_1, \dots, T_N)) \text{ for } k = 2, \dots, N-1. \quad (6.1)$$

Our goal is the following:

Proposition 6.1.10. *Let R be an algebraic field extension of $\mathbb{Z}/(p)$, for p a prime not dividing $N-1$, and let $t_1, \dots, t_N$ be non-zero elements of R which are zeros of the polynomials* (6.1)*. Then $L^\natural$ is wide over R when equipped with the H_2^D local system defined by the t_j.*

Proof. Let $\alpha \in R$ be $e_1(t_1, \dots, t_N)^N / e_N(t_1, \dots, t_N)$. This is algebraic over $\mathbb{Z}/(p)$, so it has a (monic) minimal polynomial $f(x)$, of degree d say. Let F denote a monic lift of f to $\mathbb{Z}[x]$. We claim first that F can be chosen so that it is irreducible over $\mathbb{Z}$, and hence over $\mathbb{Q}$.

To see this, pick a prime q not equal to p, and note that we can add a $\mathbb{Z}$-linear combination of $px^{d-1}, \dots, px, p$ to F so that the constant term is congruent to $q \mod q^2$ and all other coefficients (except the leading coefficient, which is 1) are divisible by q. This new F is still a lift of f, but it is now irreducible over $\mathbb{Z}$ by Eisenstein's criterion.

Given this polynomial F, consider the ring $S \coloneqq \mathbb{Z}[H_2^D, x]/(I, F(x), e_1^N - xe_N)$, where e_k is shorthand for $e_k(T_1, \dots, T_N)$ and x has degree 0 (note that the ideal we are quotienting by is homogeneous). Tensoring S with $\mathbb{Z}/(p)$ gives $S_p \coloneqq R'[H_2^D]/(I, e_1^N - \alpha e_N)$, where R' is the field obtained by adjoining α to $\mathbb{Z}/(p)$. Similarly, tensoring with $\mathbb{Q}$ gives $S_\mathbb{Q} \coloneqq Q'[H_2^D]/(I, e_1^N - \beta e_N)$, where Q' is the field obtained by adjoining a root β of F to $\mathbb{Q}$. Note that S_p and $S_\mathbb{Q}$ are finite-dimensional in each degree over $\mathbb{Z}/(p)$ and $\mathbb{Q}$ respectively, and their quotients by $e_N - 1$ have the same dimension, namely $dN \times N!$ (the quotient of the Laurent polynomial ring by

$(e_1, \ldots, e_{N-1}, e_N - 1)$ has rank $N!$, and our construction is analogous but with a relation on e_1^N—rather than e_1 itself—and a degree d extension of the base field). Since S_p and $S_{\mathbb{Q}}$ are concentrated in even degree and contain a degree 2 automorphism (for example, multiplication by T_1), we conclude that they in fact have the same dimension over $\mathbb{Z}/(p)$ and $\mathbb{Q}$ respectively in every degree.

We can now apply the dimension counting argument from Proposition 5.1.17: the spectral sequence differentials over $S_{\mathbb{Q}}$ all vanish since they vanish over $\mathbb{C}[H_2^D]/I$; this gives a lower bound on the $\mathbb{Q}$-dimension of the Floer cohomology over $S_{\mathbb{Q}}$ in each degree; by the universal coefficient theorem we obtain a lower bound on the $\mathbb{Z}$-rank of the Floer cohomology over S in each degree and then on the $\mathbb{Z}/(p)$-dimension of the Floer cohomology over S_p in each degree; this shows that the spectral sequence differentials over S_p all vanish; tensoring down to the Novikov ring $R[T^{\pm 1}]$ over R, setting each T_j equal to $t_j T$, we see that with the given H_2^D local system the spectral sequence differentials over R all vanish. □

This gives many points in the characteristic p equivalent of the wide variety $\mathcal{W}_2$. In particular, we can prove wideness in the $N = 3$, $p = 5$ case not covered by Proposition 5.1.9 and Corollary 5.1.20. Recall that we gave two proofs of this earlier, in Remark 5.2.3 and Proposition 5.2.10, using the Gysin sequence and knowledge of the discs on the Chekanov torus in $\mathbb{CP}^1 \times \mathbb{CP}^1$ respectively.

Corollary 6.1.11. *When $N = 3$, and L is equipped with the trivial local system and a relative spin structure of signature $(2, 1)$, it is wide over any field R of characteristic 5.*

Proof. It suffices to consider the case $R = \mathbb{Z}/(5)$, for which we can apply Proposition 6.1.10. The choice of relative spin structure corresponds to $t_1 = t_2 = 1$ and $t_3 = -1$ (or some permutation of these), for which we have

$$e_2(2t_1 - e_1(t_1, t_2, t_3), 2t_2 - e_1(t_1, t_2, t_3), 2t_3 - e_1(t_1, t_2, t_3)) = e_2(1, 1, -3) = -5.$$

Thus the required equation is satisfied in R. □

6.2 The virtual pearl complex

6.2.1 The complex and a quantum pullback map

Suppose $L^\natural \subset X$ is a K-homogeneous monotone Lagrangian pre-brane. We prove in Appendix C that an arbitrary Morse–Smale pair on L can be perturbed by a diffeomorphism arbitrarily C^∞-close to id_L so that it can be used to define the pearl complex on L with respect to the K-invariant integrable complex structure. The crucial properties of the moduli spaces of holomorphic discs on L which are used are that: they, and their bubbled versions, are smooth manifolds of the correct dimension; they satisfy a gluing result, compatible with orientations, so that the bubbled configurations at the boundaries of moduli spaces of trajectories cancel out.

Now suppose that K' is a closed, connected subgroup of K which acts freely on L, and let L' denote L/K', of dimension $n' = n - \dim K'$, with $\pi \colon L \to L'$ the quotient projection. The moduli spaces $\mathcal{M}_{\mu=j}$ of holomorphic discs on L carry free K'-actions, and the quotients $\mathcal{M}'_{\mu=j}$ come with well-defined evaluation maps to L'. These quotient moduli spaces are smooth manifolds of the 'correct' dimension $n' + j$, as if they were moduli spaces of pseudoholomorphic discs on L' of the same index j, and so we can try to use them to define a pearl-type complex on L'.

Explicitly, we take a Morse–Smale pair (f', g') on L' and (after perturbing, if necessary) count its Morse trajectories, interrupted by boundaries of holomorphic discs upstairs on L, as illustrated in Fig. 6.5. The discs on L—of which there may be any number (including

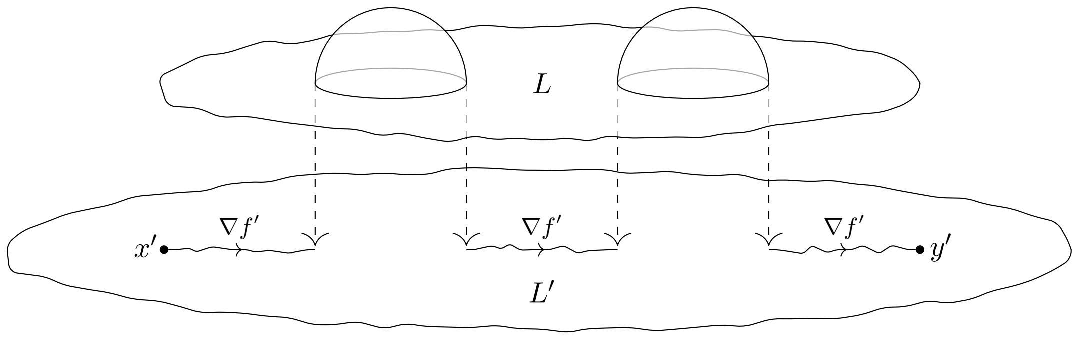

Figure 6.5: A virtual pearly trajectory between critical points x' and y'.

zero), but which must all be non-constant—are only remembered modulo translation by K' and reparametrisations fixing ± 1. More precisely, we count elements of fibre products, analogous to (C.1) built from the Morse flow on L' and the evaluation maps from the quotient moduli spaces $\mathcal{M}'_{2,\mu=j}$. It should be possible to equip the quotient spaces $\mathcal{M}'_{\mu=j}$ with appropriate orientations so that they inherit the necessary gluing property from upstairs.

We thus obtain a complex on L' which we call the *virtual pearl complex*. Using the arguments of Appendix C we should be able to define a product on the complex, and show that the resulting cohomology ring $HF^*_{\text{virt}}(L', L')$ is independent of the initial choice of Morse data.

One reason to be interested in this complex is that it can be related to the pearl complex on L as follows. The proof of genericity of Morse–Smale pairs in [90] can be adapted to show that there exists such a pair (f, g) on L such that the projection $\pi_* \nabla f$ coincides with $\nabla f'$. In other words, the Morse flow in the total space projects to the Morse flow in the base. Each critical point x of f on L projects to a critical point $x' = \pi(x)$ of f' on L', and we define the base and fibre indices of x to be $|x'|$ and $|x| - |x'|$ respectively. Since K' is connected we may arrange it so that there is a unique minimum in each fibre over a critical point (this is really only for convenience; we could just work everywhere with the Morse cocycle representing the fibre unit, instead of with the unique fibre minimum).

Now, this Morse data upstairs may not satisfy the necessary transversality properties for defining the pearl complex. However, we do have the following partial result:

Lemma 6.2.1. *The rigid (i.e. virtual dimension* 0*) pearly trajectories on* L *starting from fibre minima are transversely cut out and end at fibre minima, or are Morse trajectories lying in the fibres.*

Proof. Suppose we have a rigid pearly trajectory on L from a fibre minimum x to another critical point y of f. Let μ denote the total index of the discs in the trajectory. We then have

$$|y| = |x| + 1 - \mu. \tag{6.2}$$

Now project the trajectory to L' and consider the resulting virtual pearly trajectory between critical points $x' = \pi(x)$ and $y' = \pi(y)$ of f'.

Since (f', g') were chosen to ensure transversality for this virtual pearl complex, we must have

$$|y'| \geq |x'| - \mu, \tag{6.3}$$

with equality if and only if the projected trajectory is constant, i.e. $x' = y'$ and all Morse flow segments in the original trajectory lie in the fibre over this point. In the latter case we have $|y| \geq |x|$ since x is the minimum in its fibre, so γ must just be a Morse trajectory (without any discs), and we are left to deal with the case where the inequality in (6.3) is strict.

Comparing the strict form of (6.3) with (6.2) we conclude that the fibre index of y is at most the fibre index of x, and hence that y is also a fibre minimum. To see that the trajectory is transversely cut out, first note that its projection to L' is transversely cut out so it suffices to check 'vertical transversality'. For this one can just use the translation of discs in the fibres along with the fact that the ascending manifold of x is open in its fibre. □

Note that the Morse trajectories which start from a fibre minimum x and remain in the fibre must cancel out since x defines a cycle in its fibre. Note also that if γ is a rigid trajectory on L' from x' to y' then there is a unique lift of the final end point to the fibre minimum y over y' and, tracing backwards along the trajectory, there is a unique lift of the entire trajectory to L, up to and including the initial end point. Generically this lift of the initial end point will land in the ascending manifold of the fibre minimum x over x' and we may set things up so that this is indeed the case. We thus have a bijection between trajectories in the virtual pearl complex and trajectories from the fibre minima upstairs.

We are now ready to perturb the Morse data on L so that it really can be used to compute the pearl complex, and we can do this so that the counts of trajectories from fibre minima are unaffected. We thus obtain an inclusion of the virtual complex on L' as a subcomplex of the pearl complex on $L^\flat$ (which is $L^\natural$ equipped with the trivial local system), and hence a homomorphism

$$\pi^* \colon HF^*_{\mathrm{virt}}(L', L') \to HF^*(L^\flat, L^\flat).$$

This is a quantum analogue of the Morse-theoretic pullback map, and respects the energy filtration, so induces a map of spectral sequences from the 'virtual Oh spectral sequence' $H^*(L') \implies HF^*_{\mathrm{virt}}(L', L')$ of L' to the ordinary Oh spectral sequence of L which agrees with the classical π^* on the first page.

Remark 6.2.2. Similar ideas appear in the work of Biran–Khanevsky [18] on circle bundles. //

Remark 6.2.3. For any monotone Lagrangian brane $L^\flat = (L, s, \mathscr{F})$ there is a canonical map $H^{<N_L}(L) \to HF^*(L^\flat, L^\flat)$, where the cohomology of L is twisted by $\operatorname{End}\mathscr{F}$; this comes directly from the pearl complex. One consequence of the construction of the quantum pullback map as

a deformation of the classical pullback map is that the diagram

$$\begin{tikzcd} H^{<N_L}(L') \arrow[r, "\pi^*"] \arrow[d] & H^{<N_L}(L) \arrow[d] \\ HF^*_{\mathrm{virt}}(L', L') \arrow[r, "\pi^*"] & HF^*(L^\flat, L^\flat) \end{tikzcd}$$

commutes. //

It should be possible to prove that the quantum pullback map is canonical, in the sense of not depending on any of the Morse data, and also that it intertwines the product structures on the virtual and ordinary pearl complexes. The pullback should be functorial in the sense that if one considers a pair of subgroups $K'' \subset K'$ in K then the map

$$HF^*_{\mathrm{virt}}(L/K', L/K') \to HF^*(L, L)$$

should be the composite of the analogous pullback maps

$$HF^*_{\mathrm{virt}}(L/K', L/K') \to HF^*_{\mathrm{virt}}(L/K'', L/K'') \to HF^*(L, L).$$

Everything works equally well with the enriched Novikov ring or H_2^D local systems, and we can also equip L' with an ordinary local system $\mathscr{F}$ if we equip L with the pullback $\pi^*\mathscr{F}$.

6.2.2 A worked example

Recall again our main family from Chapter 5, of sharply $\mathrm{PSU}(N-1)$-homogeneous Lagrangians in $(\mathbb{CP}^{N-2})^N$. They are constructed by letting $\mathrm{SU}(N-1)$ act on the space $\mathrm{Mat}_{(N-1)\times N}$ of $(N-1) \times N$ matrices by multiplication on the left, and then projectivising the columns.

Suppose instead that we projectivise the whole space of matrices, rather than the individual columns. We obtain a left action of $\mathrm{PSU}(N-1)$ on $\mathbb{CP}^{N(N-1)-1}$, but now the orbits are too small to be Lagrangian. However, if we let T^N (in the form of diagonal $N \times N$ matrices) act by multiplication on the right, then we *do* obtain a Lagrangian orbit. The subgroup of T^N comprising the scalar matrices acts trivially, so the resulting Lagrangian is actually (freely) $\mathrm{PSU}(N-1) \times T^{N-1}$-homogeneous.

We could apply the argument we are about to give to this Lagrangian, however it will be helpful for later to consider instead a Lagrangian one dimension greater. To build this we simply replace $\mathrm{Mat}_{(N-1)\times N}$ in the above construction with $\mathrm{Mat}_{(N-1)\times N} \oplus \mathbb{C}$ before projectivising, and replace $\mathrm{SU}(N-1)$ with $\mathrm{U}(N-1)$. The result is a freely $K = \mathrm{U}(N-1) \times T^{N-1}$-homogeneous monotone Lagrangian L in $X \cong \mathbb{CP}^{N(N-1)}$.

The compactification divisor Y has $N+1$ components: one, denoted Y_0, corresponding to hyperplane

$$\mathbb{P}(\mathrm{Mat}_{(N-1)\times N} \oplus 0) \subset \mathbb{P}(\mathrm{Mat}_{(N-1)\times N} \oplus \mathbb{C}),$$

and one, Y_j, for each j in $\{1, \ldots, N\}$ corresponding to the vanishing of the jth $(N-1) \times (N-1)$ minor on the $\mathrm{Mat}_{(N-1)\times N}$ part. Y_0 represents the hyperplane class H, whilst each Y_j represents $(N-1)H$ (since each minor has degree $N-1$ in the homogeneous coordinates). In total, the (reduced) compactification divisor thus represents the class $(N(N-1)+1)H$, which is anticanonical, so by Remark 3.3.45 each component is obliging. There are therefore $N+1$ families of index 2 discs on L, and the boundary point evaluation map on each of these families has degree 1 (since the homogeneity is free).

Let the class of the index 2 disc meeting Y_j be A_j (for $j=0,\dots,N$). Pairing with $Y_0,\dots,Y_N$ exhibits $A_0,\dots,A_N$ as a free basis for a summand of $H_2(X,L;\mathbb{Z})$ of rank $N+1$, and from the long exact sequence we see that this summand must in fact be the whole group. The maps $f\colon H_2(X;\mathbb{Z})\to H_2(X,L;\mathbb{Z})$ and $g\colon H^2(X\setminus L;\mathbb{Z})\to H^2(X;\mathbb{Z})$ are dual, and from our computation of the classes of $Y_0,\dots,Y_N$ we know that the kernel of g is spanned by $Y_j-(N-1)Y_0$ for $j=1,\dots,N$. These therefore span the annihilator of the image of f, so pairing with these classes gives a map to $\mathbb{Z}^N\cong H_1(L;\mathbb{Z})$ which is equivalent to the boundary map.

H_2^D local systems over a ring R are parametrised by $t_0,\dots,t_N\in R^\times$, encoding the weights attached to $A_0,\dots,A_N$ respectively. With such a system, we see that under the above identification of $H_1(L;\mathbb{Z})$ with $\mathbb{Z}^N$ the sum of the boundaries of the index 2 discs is

$$\begin{pmatrix} t_1-(N-1)t_0 \\ \vdots \\ t_N-(N-1)t_0 \end{pmatrix}. \tag{6.4}$$

Now take K' to be the subgroup $\mathrm{U}(1)\times T^{N-1}\subset \mathrm{U}(N-1)\times T^{N-1}$, where $\mathrm{U}(1)$ denotes the subgroup of scalar matrices in $\mathrm{U}(N-1)$. The quotient space L' is precisely $\mathrm{PSU}(N-1)$ and we can apply the same representation theory as for our main family to show—over a field R of characteristic p not dividing $N-1$—that the differentials in the virtual Oh spectral sequence all vanish when L' is equipped with an H_2^D local system for which $e_1(t_1,\dots,t_N),\dots,e_{N-1}(t_1,\dots,t_N)$ all vanish (this condition can be weakened using Section 6.1.2 but we don't need this). In fact the virtual pearl complex for L' is essentially equivalent to the ordinary pearl complex for the $\mathrm{PSU}(N-1)$ Lagrangian in our main family—see Section 6.2.5.

Since $H^*(L;R)$ is generated as a ring by $\pi^*H^*(L';R)$ and $H^1(L;R)$ (still assuming that R is a field of characteristic p not dividing $N-1$), using the virtual pullback map on spectral sequences we see that L is wide if we can ensure that the index 2 boundary sum (6.4) and the differentials in the virtual Oh spectral sequence all vanish. It therefore suffices to ask that $t_j=(N-1)t_0$ for $j=1,\dots,N$ and $e_k(1,\dots,1)=0$ for $k=1,\dots,N-1$. In particular, L is wide if N is a power of p and we set $t_0=-1$ and $t_1,\dots,t_N=1$.

The reason for choosing this particular example is that it connects with both the main family from Section 5.1 and the Amarzaya–Ohnita–Chiang family from Section 5.2.4 via symplectic reduction. If Y is a symplectic manifold carrying a Hamiltonian action of a compact Lie group H, and the action of H on the zero level set of the moment map μ is free, then we can form the reduction $Y/\!\!/H=\mu^{-1}(0)/H$ and there is a Lagrangian correspondence Γ from Y to $Y/\!\!/H$ (or vice versa) defined by

$$\{(y,[y]) : y\in\mu^{-1}(0)\}\subset Y^-\times Y/\!\!/H,$$

where $[\cdot]$ denotes the H-orbit (i.e. equivalence class under the H-action). Under appropriate technical hypotheses (see [142], [102], [98], [56]) this defines a quilt functor from the Fukaya category of $Y /\!\!/ H$ to the (possibly enlarged or bulk deformed) Fukaya category of Y, and similarly from Y to $Y /\!\!/ H$.

A Lagrangian $M \subset Y$ may not interact at all well with the H-action but there are two types of behaviour that are particularly nice. First, suppose M is contained in $\mu^{-1}(0)$ and is H-invariant. There is then a reduced Lagrangian M/H in $Y /\!\!/ H$, and the result of composing Γ with M/H is precisely M. Moreover, this composition is embedded. Modulo technical details, we deduce that the quilt functor from $Y /\!\!/ H$ to Y sends M/H to M. In particular, if M is Floer-theoretically non-trivial then so should M be. This is an algebraic version of the fact that displaceability of M/H implies displaceability of M, as employed by Abreu–Macarini in [5].

In our situation, taking $Y = \mathbb{CP}^{N(N-1)}$ and H and M to be respectively the torus K' and Lagrangian L from above, it is not hard to see that the reduced space $Y /\!\!/ H$ is $(\mathbb{CP}^{N-2})^N$, whilst the reduced Lagrangian M/H is precisely the Lagrangian from Section 5.1 (if the moment map is normalised so that the zero set contains M). The fact that M is wide in characteristic p when N is a power of p thus implies the same of M/H (subject to checking the technicalities). In other words, we recover our result from Chapter 5. Of course, we used this result—or the ideas from its proof—in the computation of the Floer cohomology of M, but it is interesting to note that the virtual pearl complex gives a way to pass information in the opposite direction to that which one would expect from quilt theory.

The other case which is especially nice is when M is transverse to $\mu^{-1}(0)$ and it meets each H-orbit in this level set at most once. Then the composition of Γ with M is embedded, and outputs M/H, so we get a quilt functor going the other way than before. Now if M/H has non-vanishing self-Floer cohomology we can deduce the same about M.

To apply this, take Y to be the projectivisation of the space of $N \times N$ matrices, and let $H = T^{N-1}$ act by rotating the phases of the final entries of the first $N-1$ columns. One can check that the zero level set of the (appropriately normalised) moment map is transverse to the Amarzaya–Ohnita–Chiang Lagrangian $M = L_{\mathrm{AOC}}$ (if we had acted by T^N on all N columns then this would not be the case), and that M meets each orbit in the zero set at most once. In fact, the intersection is the Lagrangian L from above. The quilt functor thus recovers the Evans–Lekili result that L_{AOC} is wide in characteristic p when N is a power of p (although their proof is much easier).

The above computation for the Lagrangian L in $\mathbb{CP}^{N(N-1)}$ therefore 'contains' the results for the two other families, and gives a rationale for the appearance of the same prime power condition in both cases.

Remark 6.2.4. Really one should carefully keep track of the relative spin structures throughout these quilt functor arguments, and we have not checked that the signs all work out. $/\!\!/$

6.2.3 A quantum Pontryagin comodule structure

The virtual pearl complex also gives a new perspective on the Pontryagin structure, which we now describe. First we return to the case of a topological group K acting on a space L. If we take the pullbacks in cohomology induced by the multiplication and action maps, rather than the pushforwards in homology, then we obtain a comultiplication $H^*(K) \to H^*(K) \otimes H^*(K)$ and comodule structure $H^*(L) \to H^*(K) \otimes H^*(L)$ (as long as the Künneth theorem is valid, for example if we work over a field). The benefit of this, when compared with the module structure on homology, is that it respects the cup products, so that $H^*(K)$ becomes a Hopf algebra (see [82, Section 3.C]) and $H^*(L)$ a comodule compatible with the multiplication on both itself and $H^*(K)$.

For a K-homogeneous monotone Lagrangian brane $L^\flat \subset X$, one can interpret this construction Morse theoretically (essentially one takes the diagrams in Section 6.1.1, but views inputs as outputs and vice versa) and define pearly analogues to get a quantum Pontryagin comodule structure

$$HF^*(L^\flat, L^\flat) \to H^*(K) \otimes HF^*(L^\flat, L^\flat)$$

which respects the cup product on $H^*(K)$ and the Floer product on $HF^*(L^\flat, L^\flat)$. This is rather ad hoc and requires the verification of many orientation identities, which we have not done (we have also not checked the orientations for the virtual pearl complex, on which the rest of our discussion is based, but that should be significantly simpler).

An alternative, more systematic, approach is as follows. Consider the Lagrangian $Z \times L$ inside $T^*K \times X$, where Z is the zero section. Note that T^*K is non-compact (assuming K, L and X are not just points!), which we have not allowed before now, but since $Z \subset T^*K$ is exact it bounds no non-constant (pseudo)holomorphic discs and so its pearl complex can be straightforwardly defined, and is just its Morse complex.

The Lagrangian $Z \times L$ is $K \times K$-homogeneous (strictly our definition of homogeneity requires us to think of the ambient spaces as complex, rather than symplectic, manifolds, so we view T^*K as the complexification G of K, although as we just remarked there are no holomorphic discs anyway), and we can consider the virtual pearl complex for the antidiagonal subgroup $K' = \{(k^{-1}, k)\} \subset K \times K$. This clearly acts freely, and the quotient is canonically identified with L itself. Moreover, if we identify Z with K then the quotient projection $\pi \colon Z \times L \to L$ is just the map defining the action. Holomorphic discs on $Z \times L$ are constant in the T^*K factor, so the virtual pearl complex on L is just the ordinary pearl complex (when the K-action on X is Hamiltonian this gives the simplest example of Fukaya's equivalence between equivariant and reduced Floer theory—see Section 6.2.5).

The quantum pullback map thus gives a homomorphism

$$\Delta \colon HF^*(L^\flat, L^\flat) \cong HF^*_{\mathrm{virt}}(L, L) \to HF^*(Z \times L^\flat, Z \times L^\flat) \cong H^*(K) \otimes HF^*(L^\flat, L^\flat).$$

The final isomorphism uses the Künneth theorem (so requires some hypotheses) and the fact that the self-Floer and classical cohomology rings of the zero section $Z = K$ are canonically identified. This is the Floer-theoretic analogue of the map which defines the classical Pontryagin comodule structure, and by general properties of the quantum pullback it should be compatible with the ring structures. The fact that this defines a (coassociative) comodule structure follows from functoriality for the towers of fibrations

$$(K \times K) \times L \to K \times L \to L$$

and

$$K \times (K \times L) \to K \times L \to L.$$

We thus obtain a quantum Pontryagin comodule structure from general constructions with the virtual pearl complex.

The direct pearl complex construction, based on Section 6.1.1 with the inputs and outputs swapped, counts pearly trajectories on L which split, at a junction we denoted by $\odot$, into a Morse trajectory on K and another pearly trajectory on L. We can think of the part after the junction as a pearly trajectory on $Z \times L$, using Morse data which is split according to the product structure. Using the virtual pearl complex, meanwhile, we can think of the part before the junction as a pearly trajectory on $Z \times L$ using (a small perturbation of) Morse data pulled back from L under the action map π, and beginning at a fibre minimum. The junction then corresponds to an intersection of these two trajectories on $Z \times L$.

Counting trajectories for which the Morse data changes part way along like this gives an alternative (to the Morse cobordism) definition of the pearl complex comparison map—see [15, Section 3.4]. We deduce that the direct pearl complex construction of the Pontryagin structure is simply the quantum pullback construction followed by a comparison map for a change of auxiliary data. The two constructions therefore coincide after passing to cohomology. The involvement of this comparison map is what makes the direct construction more cumbersome to work with.

Remark 6.2.5. In Remark 6.1.2 we noted that if L is equipped with a local system $\mathscr{F}$ then, in order to define the Pontryagin structure, the pullback of $\mathscr{F}$ under the orbit map $\varphi\colon K \to L$ induced by an arbitrary base point $x \in L$ should be trivial. Since the inclusion of $K \vee L$ in $K \times L$ as $K \times \{x\} \cup \{e\} \times L$ is surjective on π_1, a local system on $K \times L$ is determined up to isomorphism by its restriction to this subspace. Therefore $\varphi^*\mathscr{F}$ is trivial if and only if $\pi^*\mathscr{F}$ and $\mathrm{pr}_2^*\mathscr{F}$ are isomorphic as local systems on $K \times L$, where $\pi\colon K \times L$ is the action map as above and $\mathrm{pr}_2\colon K \times L \to L$ is projection to the second factor.

This is precisely the property we need in order to define the Pontryagin structure using the virtual complex, as follows. When constructing the quantum pullback map in Section 6.2.1 we remarked that we can equip the base with any local system as long as we equip the total space with the pullback. In our case this means that if L is given the local system $\mathscr{F}$ then $K \times L$ (i.e. $Z \times L$) should be given the local system $\pi^*\mathscr{F}$. When we come to apply the Künneth theorem we then need that

$$\pi^*\mathscr{F} \cong \mathrm{pr}_1^*\mathscr{E} \otimes \mathrm{pr}_2^*\mathscr{F},$$

where $\mathscr{E}$ is the local system we are using on $Z = K$, which we have assumed is trivial throughout our discussions. Note that $\mathscr{E}$ is the local system for $HF^*(Z, Z)$, which corresponds to the local system $\operatorname{End}\mathscr{E}$ for $H^*(K)$. //

Remark 6.2.6. The cotangent bundle T^*K seems somewhat redundant in this construction—we may as well just work with the Morse complex on K and forget that it is the pearl complex on the zero section—although by viewing $H^*(K)$ as $HF^*(Z, Z)$ we put it on the same footing as $HF^*(L^\flat, L^\flat)$. We will see a much more serious benefit of this perspective in Section 6.3. //

6.2.4 Monotone toric fibres

Compatibility with the Floer and cup products makes the Pontryagin comodule structure much more readily computable than the earlier module structure. As an example, in this subsection we calculate it over $\mathbb{C}$ for monotone toric fibres.

Recall that if $L \subset X$ is such a toric fibre, of dimension n, then it has an associated *superpotential*, W, which can be viewed as an element of $\mathbb{Z}[H_1(L;\mathbb{Z})]$. It is the sum of the monomials $T^{[\partial u]}$, over holomorphic index 2 discs u through a generic point of L, and is usually written as a Laurent polynomial in $x_1 = T^{\gamma_1}, \dots, x_n = T^{\gamma_n}$ where $\gamma_1, \dots, \gamma_n$ is a basis for $H_1(L;\mathbb{Z})$. Rank 1 local systems on L over $\mathbb{C}$ are equivalent to homomorphisms $\rho\colon H_1(L;\mathbb{Z}) \to \mathbb{C}^*$, or equivalently elements of $H^1(L;\mathbb{C}^*) \subset H^1(L;\mathbb{C})$, and there is a pairing

$$\mathbb{Z}[H_1(L;\mathbb{Z})] \times H^1(L;\mathbb{C}) \to \mathbb{C}$$

given by

$$\left(\sum n_j T^{\gamma_j}, \rho\right) \mapsto \sum n_j \rho(\gamma_j).$$

Restricting this map to $\{W\} \times H^1(L;\mathbb{C}^*)$, the superpotential defines a complex-valued function f_W on the space of local systems.

It is well-known from the work of Cho–Oh [35] that the (rank 1) local systems ρ for which L is wide, when equipped with the standard spin structure, are precisely those which are critical points of f_W. If $L^\flat$ is the brane corresponding to a critical point ρ, Cho [33] subsequently showed that the second derivatives of f_W at ρ determine the product structure on $HF^*(L^\flat, L^\flat;\mathbb{C})$: there is an isomorphism of algebras

$$HF^*(L^\flat, L^\flat;\mathbb{C}) \cong \mathrm{Cl}(H^1(L;\mathbb{C}), \mathrm{Hess}(f_W)|_\rho).$$

Here Hess is the Hessian, and for a vector space V, equipped with a quadratic form Q, $\mathrm{Cl}(V,Q)$ denotes the corresponding Clifford algebra. Explicitly, this is the quotient of the tensor algebra by the two-sided ideal generated by expressions of the form $v \otimes v - Q(v,v)/2$. For the remainder of this subsection, all cohomology groups will be over $\mathbb{C}$.

Remark 6.2.7. One could instead consider the upgraded superpotential $\widetilde{W} = \sum_u T^{[u]}$, lying in $\mathbb{Z}[H_2(X,L;\mathbb{Z})]$, which records the entire relative homology classes of the index 2 discs. This defines a complex-valued function on the space of H_2^D local systems, but those for which L is wide no longer correspond to critical points. In fact, the wide systems are precisely those on which the function

$$\sum_u \alpha(\partial u) T^{[u]}$$

vanishes for all α in $H^1(X;\mathbb{C})$. //

Since the algebra $HF^*(L^\flat, L^\flat)$ is generated by the obvious image of $H^{\leq 1}(L)$ in the Clifford algebra, the quantum Pontryagin comodule map Δ is completely determined by its restriction to this subspace. By the classical compatibility from Remark 6.2.3 (noting that the map $H^{<N_L}(L) \to HF^*(L^\flat, L^\flat)$ used there coincides with the inclusion of $H^{\leq 1}(L)$ which we are considering here) the following diagram commutes, where the top arrow is the classical comodule

structure map:

$$\begin{array}{ccc} H^{\leq 1}(L) & \xrightarrow{\Delta} & H^{\leq 1}(K) \otimes H^{\leq 1}(L) \\ \downarrow & & \downarrow \\ HF^*(L^\flat, L^\flat) & \xrightarrow{\Delta} & H^*(K) \otimes HF^*(L^\flat, L^\flat) \end{array}$$

Here $K \cong T^n$ is the abstract torus acting on L. After choosing a base point on L we get an identification of K with L, and the induced isomorphism on cohomology is independent of this choice, so there is a canonical isomorphism of algebras

$$H^*(K) \cong H^*(L) \cong \Lambda H^1(L),$$

where Λ denotes the exterior algebra.

This classical comodule structure is well-known, and can be computed as the dual of the Pontryagin product on $H_*(K)$, which is more easy to visualise—it satisfies $\Delta(x) = 1 \otimes x + x \otimes 1$ for all x in $H^1(L)$. The quantum Pontryagin comodule structure is therefore the unique algebra homomorphism

$$\Delta \colon \mathrm{Cl}(H^1(L), \mathrm{Hess}(f_W)|_\rho) \to \Lambda H^1(L) \otimes \mathrm{Cl}(H^1(L), \mathrm{Hess}(f_W)|_\rho)$$

satisfying this relation. This exhibits the Clifford algebra as a comodule over the corresponding exterior algebra, and the homomorphism is well-known to algebraists—see [84, Theorem 4.4.1]—as the map of Clifford algebras induced by the diagonal inclusion of spaces with quadratic forms

$$(H^1(L), \mathrm{Hess}(f_W)|_\rho) \to (H^1(L), 0) \oplus (H^1(L), \mathrm{Hess}(f_W)|_\rho).$$

6.2.5 Equivariant Floer cohomology and a generalised superpotential

Fukaya has recently been developing an equivariant Floer theory in the setting of symplectic manifolds X carrying a symplectic action of a compact Lie group K [61]. His construction is not yet published, so it is not clear what the details of his results are, but the basic idea is to take a K-invariant Lagrangian, choose an invariant compatible almost structure, construct equivariant Kuranishi structures on the corresponding moduli spaces of pseudoholomorphic discs on L [55], and use these to define A_∞-operations on a suitable cochain complex on L/K. Although the original motivation for the virtual pearl complex was unrelated to this equivariant Floer cohomology—it was intended as a way to constrain the pearl complex upstairs, where we have seen its utility—one can view it as a model for computing equivariant Floer cohomology in the homogeneous, monotone setting: we should have

$$HF^*_{\mathrm{virt}}(L/K, L/K) \cong HF^*_K(L, L).$$

Fukaya's construction will presumably work in much greater generality than we have been considering, which is obviously a considerable advantage. In particular, it does not depend on a specific choice of invariant integrable complex structure, so really is a symplectic object, and includes a full A_∞-structure. However, the virtual pearl complex has the benefit of being rather concrete and, in some situations, computable, as well as being defined over $\mathbb{Z}$ (which may or may not be the case in Fukaya's work). In Section 6.3.2 we propose an alternative, abstract construction of equivariant Floer cohomology.

There are two key properties of his theory which the author is aware of. First, there is a morphism from equivariant to non-equivariant theory, on Floer cohomology of an individual Lagrangian and at the level of Fukaya categories [60]. This is analogous to our quantum pullback map, and is again likely to have considerable conceptual advantages at the expense of concreteness and explicit computability. More remarkably, when the K-action on X is Hamiltonian and the action on the zero level set of the moment map is free, there is an equivalence between K-equivariant Floer theory on X and ordinary Floer theory on the symplectic reduction $X /\!\!/ K$, under which an invariant Lagrangian in the zero set on X corresponds to its quotient in the reduction [61]. In general one needs to introduce a bulk deformation to obtain this equivalence, but this should be unnecessary in the monotone setting.

We have already seen an example of this equivalence, when we compared the virtual pearl complexes of the Lagrangians L in Section 6.2.2 with the ordinary pearl complexes of the Lagrangians from Section 5.1. In the remainder of this subsection we sketch a generalisation of the notion of the superpotential of a toric fibre to more general homogeneous Lagrangians, which, as well as being of independent interest, provides an alternative perspective on this phenomenon in certain situations.

Suppose then that (X, L) is K-homogeneous. We saw in Section 3.2.2 that for each index j, the evaluation map $\mathrm{ev}_0 \colon \mathcal{M}_{\mu=j} \to X$ transverse to all K-invariant complex submanifolds. Therefore for each j, and any invariant complex submanifold $Z \subset X$ of complex codimension at least 2, the set of discs in $\mathcal{M}_{2,\mu=j}$ which meet Z form a submanifold of real codimension at least 2. Using the arguments of Appendix C, we can ensure that no such discs appear in the trajectories defining the pearl complex (or the product, comparison maps, or Pontryagin structure), so they are in this sense invisible to Floer theory. The same remains true if we replace Z by any finite collection of invariant analytic subvarieties of complex codimension at least 2.

Now assume that the homogeneity is sharp and linear, and for simplicity that each component $Y_1, \dots, Y_r$ of the compactification divisor Y is obliging—one can do something similar more generally but it is more complicated. Fix a base point x in L, and let $\Gamma \subset K$ be its stabiliser. For each j, the results of Section 3.3 ensure that any pole germ evaluating to the dense open orbit $\mathcal{O}_j$ in Y_j is quasi-axial, and is of a unique type. This type is described by an element ξ_j of $\mathfrak{k}$, which is defined up to conjugation by Γ.

Letting Z denote the union over j of the subvarieties $Y_j \setminus \mathcal{O}_j$, we see that holomorphic discs in $X \setminus Z$ with boundary on L are equivalent to holomorphic maps $h \colon (\Sigma, \partial\Sigma) \to (G/\Gamma, K/\Gamma)$, where Σ is a disc with a finite number (possibly zero) of interior punctures, such that near each puncture the map blows up like one of the above quasi-axial poles. Explicitly, this means that if z is a local coordinate about the puncture then there exist a positive integer k, an element j of $\{1, \dots, r\}$, and a holomorphic map g from a neighbourhood of 0 in $\mathbb{C}$ to G such that

$$h(z) = g(z) e^{-ik\xi_j \log z}$$

on a punctured neighbourhood of 0. Since Z is a K-invariant subvariety of complex codimension at least 2, we may assume that all discs involved in the pearl complex lie in $X \setminus Z$, and hence are of this form.

We conclude that the self-Floer cohomology of L (including product and Pontryagin structures) is entirely determined by a finite collection of data, namely the groups K and Γ, and the Γ-conjugacy classes of the ξ_j, with everything taken modulo overall conjugation by K (corresponding to changing the base point). One can think of this data as a generalised superpotential for L. We can package it up as (K, Γ, W), where W is the element of $\mathbb{Z}[\mathfrak{k}]$ given by the sum of the monomials T^{ξ}, where ξ ranges over the ξ_j and their distinct Γ conjugates.

Remark 6.2.8. Of course, one has to worry about orientations, but we have a standard spin structure and with respect to this the sign attached to a disc can be determined from the corresponding map h. //

The relevance to the virtual pearl complex is as follows. Suppose (in the same setting) that K' is closed, connected subgroup of K which is contained in the centre and acts freely on L. Let H be the quotient K/K', with projection $p \colon K \to H$. The virtual pearl complex of $L' = L/K'$ behaves like the actual pearl complex for the generalised superpotential $(H, p(\Gamma), p_* W)$, where p_* is the homomorphism $\mathbb{Z}[\mathfrak{k}] \to \mathbb{Z}[\mathfrak{h}]$ induced by $p_* \colon \mathfrak{k} \to \mathfrak{h}$. Note that $p_* W$ may contain repeated monomials (which should be interpreted as the requirement to count discs with certain multiplicities), though this did not happen for 'real' generalised superpotentials. If we can find a sharply linearly H-homogeneous Lagrangian L'' with generalised superpotential $(H, p(\Gamma), p_* W)$, then its ordinary pearl complex should compute the virtual pearl complex cohomology for L'.

Remark 6.2.9. Here there is a more serious issue regarding orientations: it is not clear that the orientations upstairs agree with those downstairs. We will ignore this. //

Now consider the example from Section 6.2.2, where $K = \mathrm{U}(N-1) \times T^{N-1}$, Γ is trivial, and there are $N+1$ Lie algebra elements $\xi_0, \dots, \xi_N$ describing the types of the poles evaluating to the divisors $Y_0, \dots, Y_N$ respectively. The group K' is the (central) torus $\mathrm{U}(1) \times T^{N-1}$, and hence the quotient generalised superpotential is

$$\left(\mathrm{PSU}(N-1), \{e\}, \sum_{j=0}^{N} T^{p_* \xi_j}\right).$$

One can check that this is equivalent to the generalised superpotential for our main family from Chapter 5, except there is an extra term $T^{p_* \xi_0} = T^0$.

However, if u is a holomorphic disc on the upstairs Lagrangian L which has a pole of type ξ_0 (which we assume is at 0), and the corresponding punctured curve in G is h, then

$$h(z) e^{i \xi_0 \log z}$$

corresponds to a disc v of index $\mu(v) = \mu(u) - 2$, whose boundary has the same projection to L' as u. This means that the evaluation maps to L', from moduli spaces of discs which have such a pole, factor through moduli spaces 2 dimensions lower. Such poles can therefore be ignored from the perspective of the virtual pearl complex (clearly this applies more generally, and means that terms T^0 in the quotient superpotential can be dropped). We conclude that the virtual pearl complex on L' is equivalent to the ordinary pearl complex for our main family. To tie this back to Fukaya's result relating equivariant Floer theory to symplectic reduction, recall that our $\mathrm{PSU}(N-1)$ Lagrangian in $(\mathbb{CP}^{N-2})^N$ is the quotient of L in the reduction of $\mathbb{CP}^{N(N-1)}$ by K'.

6.3 An A_∞-Pontryagin structure

6.3.1 The moment correspondence

Given a symplectic manifold X, and a Hamiltonian action of a Lie group K on X with moment map μ, there is a natural Lagrangian

$$C_X = \{(\theta, k, x, y) \in \mathfrak{k}^* \times K \times X \times X : \theta = \mu(y) \text{ and } y = kx\} \subset T^*K \times X \times X^-$$

called the *moment correspondence*. Here we are identifying T^*K with $\mathfrak{k}^* \times K$ by identifying each fibre with the fibre $\mathfrak{k}^*$ over the identity by *right*-translation. In other words, for $\xi \in \mathfrak{k}$, $\theta \in \mathfrak{k}^*$ and $k \in K$, the pairing of the point (θ, k) in $T^*_k K$ with the vector $\xi \cdot k$ in $T_k K$ is given by $\langle \theta, \xi \rangle$. Recall also our conventions regarding the sign of the moment map and the symplectic form on the cotangent bundle from Section 1.2.4. This Lagrangian appears in the papers of Weinstein [143] and Guillemin–Sternberg [81], and is one of the key ingredients in the split-generation work of Evans–Lekili [44].

Evans and Lekili view C_X as a correspondence from T^*K to $X^- \times X$ and study the induced quilt functor. In contrast, we shall regard it as a correspondence from X to $T^*K \times X$, and—ignoring technical issues—consider the quilt functor

$$\Phi_{C_X} \colon \mathcal{F}(X) \to \mathcal{F}^\sharp(T^*K \times X).$$

Here $\mathcal{F}^\sharp$ denotes the *extended* Fukaya category in the sense of Ma'u–Wehrheim–Woodward [102], and we should restrict to appropriately nice Lagrangians, for example compact, orientable, relatively spin and exact/monotone. It may be possible to remove some of these constraints, as well as the $\sharp$, using Fukaya's recent work [56]. Under appropriate conditions, we expect there to be a Künneth decomposition of the right-hand side as a tensor product (this is also considered by Fukaya, in [56, Section 16]), so conjecturally we have a functor

$$\Psi_X \colon \mathcal{F}(X) \to \mathcal{F}(T^*K) \otimes \mathcal{F}(X).$$

The left-translation action of K on T^*K is always Hamiltonian, with moment map ν given by $\nu_{(\theta,k)} = \theta$ for all θ in $\mathfrak{k}^*$ and k in K. Taking $X = T^*K$ above, the moment correspondence is

$$C_{T^*K} = \{(\theta_1, k_1, \theta_2, k_2, \theta_3, k_3) : \theta_1 = (\mathrm{ad}^*)_{k_1}\theta_2 = \theta_3 \text{ and } k_3 = k_1 k_2\}$$

and Ψ_{T^*K} gives a functor $\mathcal{F}(T^*K) \to \mathcal{F}(T^*K) \otimes \mathcal{F}(T^*K)$. One can readily check that the compositions of correspondences

$$X \xrightarrow{C_X} T^*K \times X \xrightarrow{\Delta_{T^*K} \times C_X} T^*K \times T^*K \times X$$

and

$$X \xrightarrow{C_X} T^*K \times X \xrightarrow{C_{T^*K} \times \Delta_X} T^*K \times T^*K \times X$$

are embedded and equal, so the corresponding compositions of Ψ. functors should be homotopic (here Δ_Y denotes the identity correspondence from Y to itself). In the case $X = T^*K$, we obtain the A_∞-categorical equivalent of a Hopf algebra structure on $\mathcal{F}(T^*K)$. For general X, $\mathcal{F}(X)$ becomes a comodule over this Hopf algebra, compatibly with the A_∞-relations (up to homotopy).

Remark 6.3.1. We may be able to define a coproduct on the *wrapped* Fukaya category $\mathcal{W}(T^*K)$ in place of the compact category $\mathcal{F}(T^*K)$ if C_{T^*K} is sufficiently well-behaved. Quilt theory for wrapped Floer cohomology has previously been studied by Ganatra [73] and Gao [74]. //

Question 6.3.2. Abouzaid [2] has shown that the triangulated closure of $\mathcal{W}(T^*K)$ is quasi-isomorphic to the category of twisted complexes over the algebra of chains on the based loop space of K (this does not use the group structure on K, only the fact that K is spin; otherwise one needs a background class). Does the latter category admit a natural coproduct? Does this correspond to the conjectural coproduct on $\mathcal{W}(T^*K)$? //

Remark 6.3.3. A similar structure on the Fukaya category of a symplectic manifold carrying a Hamiltonian group action was considered by Teleman in [136] (see in particular Conjecture 2.9). The author was independently led to this idea by the direct pearl complex constructions of a quantum Pontryagin (co)module structure described earlier. //

Slightly less speculatively, restrict now to the case where X is closed and monotone, and suppose L is a K-invariant monotone Lagrangian contained in $\mu^{-1}(0)$. From now on we will ignore all matters related to orientations, relative spin structures and local systems, so will not use the brane notation, but we will assume that the minimal Maslov index of L is at least 2.

It is easy to see that the composition of L with C_X is embedded, and is exactly $Z \times L$, where Z denotes the zero section as before. Quilt theory therefore gives an A_∞-morphism

$$CF^*(L, L) \to CF^*(Z, Z) \otimes CF^*(L, L) \cong C^*(K) \otimes CF^*(L, L).$$

The associativity of the compositions between C_X and C_{T^*K} shows that this defines an comodule structure $CF^*(L, L)$ over $C^*(K)$, compatible with the A_∞-structures up to homotopy. Taking cohomology, and assuming the Künneth theorem holds, we obtain a comodule structure on $HF^*(L, L)$ over $H^*(K)$.

It is not clear how this relates to the Pontryagin comodule structure we defined earlier, but we conjecture that they are equivalent when both make sense. It is not even immediately obvious that the coproduct we obtain on $H^*(K)$ is the classical one. Note that although the new construction requires the K-action to be Hamiltonian, with L contained in the zero set of the moment map, it does not require the K-action to be transitive, or to preserve an integrable complex structure. In fact, it makes no reference to particular complex structures, so is manifestly a symplectic construction.

6.3.2 Equivariant Floer cohomology revisited

This (very brief) subsection has particularly benefited from suggestions of Yankı Lekili, who pointed out the paper [76], and Oscar Randal-Williams.

First recall that a comodule over a coalgebra C can be made into a module over the dual algebra $C^\vee$ by combining the comodule structure map with the pairing between C and $C^\vee$. In particular, for an arbitrary space L carrying an action of a topological group K, the Pontryagin coproduct makes $C^*(L)$ into a module over $C_*(K)$. In an apparently different direction (and under mild conditions on K), there is the Borel construction $EK \times_K L \to BK$, which exhibits $C^*_K(L) \coloneqq C^*(EK \times_K L)$ as a $C^*(BK)$-module. Taking cohomology we obtain precisely the classical equivariant cohomology $H^*_K(L)$ as a $H^*(BK)$-module.

When K is a compact Lie group and the ground ring is $\mathbb{R}$, Goresky–Kottwitz–MacPherson [76, Theorem 1.5.1] showed that the $C_*(K)$-module $C^*(L)$ and the $C^*(BK)$-module $C^*_K(L)$ determine each other up to quasi-isomorphism by Koszul duality. In particular, knowledge of the cochain-level Pontryagin structure is sufficient to recover the equivariant cochain complex, via a functor t. Strictly they work with $H_*(K)$- and $H^*(BK)$-modules, but it seems reasonable to expect an analogous result to hold for the $C_*(K)$- and $C^*(BK)$-module structures (indeed this may be obvious to an author more well-versed in this area).

In the previous subsection we outlined the construction of a cochain-level Floer-theoretic A_∞-Pontryagin structure. It is natural then to propose that equivariant Floer theory may be defined (over $\mathbb{R}$ at least) by applying the functor t to this construction. This avoids the need to deal with equivariant transversality issues, although it is of course rather abstract and indirect. An obvious question is how this relates to Fukaya's construction. Less ambitiously, one might ask whether this definition has the two properties we mentioned in Section 6.2.5: does it admit a map to non-equivariant Floer theory, and is it equivalent to ordinary Floer theory in the symplectic reduction?

Appendix A

Zapolsky's local systems

A.1 The complexes

The original construction of the pearl complex and its (co)homology, as set out in the foundational paper [14] of Biran and Cornea, was in characteristic 2 and without local systems. Subsequently these restrictions have been partially addressed, for instance in [17, Appendix A], but a detailed, unified treatment was lacking until Zapolsky's recent paper [146], which sets out an abstract approach to Floer and quantum (i.e. pearl complex) (co)homology, the PSS isomorphism between them, and its compatibility with various algebraic structures such as products and quantum module structures, focusing on the use of canonical orientations.

Zapolsky constructs chain complexes which are much larger than those normally considered—larger even than the enriched pearl complexes considered in Section 2.2—and (under fairly mild hypotheses) uses the extra information contained in this setup to orient the moduli spaces canonically. These complexes also allow one to incorporate a rather general notion of local system which we shall call a *Zapolsky system*. He then proves that the various moduli space compactifications defined by Biran–Cornea, involving breaking of Morse flowlines and bubbling of holomorphic discs, are compatible with this scheme. His original complexes are not in general modules over a Novikov-type ring, but by taking appropriate quotients one can recover the usual complex over Λ, or over $R[H_2^D]$. It is during the quotient procedure that a choice of relative spin structure becomes important.

The purpose of this appendix is to describe how these Zapolsky systems and their quotients twist the algebraic operations, and how local systems in the usual sense, as outlined in Section 2.1.2, can be translated into his language, as well as B-fields and the more general H_2^D local systems discussed in Section 2.2.3. This justifies our use of the pearl complex to prove things about Floer cohomology, the closed–open map (in the form of the quantum module structure), and non-displaceability. One can view it as a more sophisticated version of the splitting of moduli spaces of pearly trajectories according to the homotopy classes of the paths in Fig. 2.2 or the homology classes of the discs. We will frequently need to talk about various parallel transport maps, so for a local system $\mathscr{E}$ on a space Y, and a path Γ in Y, let $\mathcal{P}_{\mathscr{E}}(\Gamma)$ denote the parallel transport along Γ from $\mathscr{E}_{\Gamma(0)}$ to $\mathscr{E}_{\Gamma(1)}$.

Fix a monotone Lagrangian brane $L^\flat = (L, s, \mathscr{F}) \subset X$ over a ring R. Following [146], let

$$\Omega_L = \{\gamma\colon ([0,1], \{0,1\}) \to (X, L) : [\gamma] = 0 \text{ in } \pi_1(X, L)\}$$

be the space of null-homotopic paths from L to itself, and consider the cover $\widetilde{\Omega}_L$ of Ω_L whose fibre over the point γ consists of the homotopy classes of cappings of γ. We view a capping $\widehat{\gamma}$ as a map from the closed left half-disc, with the diameter parametrising γ from bottom to top and the semicircular boundary mapping to L.

Remark A.1.1. Fix a base point $*$ in L. Given a path $\gamma\colon ([0,1], \{0,1\}) \to (X, L)$, we define a *tailed capping* of γ to be a map from the closed left half-disc with the interval $[-2,-1] \times \{0\}$ ('tail') attached, such that the diameter parametrises γ from bottom to top, the rest of the boundary (including the tail) maps to L, and the end of the tail maps to $*$—see Fig. A.1. We

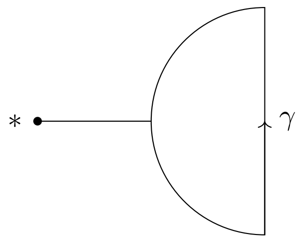

Figure A.1: A tailed capping of a path γ.

consider such tailed cappings modulo homotopies which preserve these three conditions, and use $\overline{*}$ to denote the constant path at $*$ with constant capping (with or without tail).

The space of paths equipped with a homotopy class of tailed capping is easily seen to be a model for the universal cover of Ω_L, in that it is equivalent to the space of homotopy classes of paths in Ω_L with fixed starting point, and we'll denote it by $\overline{\Omega}_L$. There is a natural free action of $\pi_1(L, *)$ on this space by gluing loops onto the tail (it's free by considering the homotopy class of path from $*$ to $\gamma(0)$ given by traversing the tail and lower boundary of the tailed capping), and a projection to $\widetilde{\Omega}_L$ which forgets the tail. It is not hard to see that this projection is precisely the quotient map, so $\pi_1(\widetilde{\Omega}_L)$ is isomorphic to $\pi_1(L)$. Explicit inverse isomorphisms

$$\pi_1(\widetilde{\Omega}_L, \overline{*}) \leftrightarrow \pi_1(L, *)$$

are induced by the projection $[\gamma, \widehat{\gamma}] \mapsto \gamma(0)$ and the inclusion of L as constant loops with constant cappings. //

We denote the homotopy class of a capping $\widehat{\gamma}$ of γ by $\widetilde{\gamma} = [\gamma, \widehat{\gamma}]$. Given a generic time-dependent Hamiltonian $H\colon X \times [0,1] \to \mathbb{R}$, there is an action functional

$$\mathcal{A}_{H:L}\colon \widetilde{\Omega}_L \to \mathbb{R},$$

given by

$$\mathcal{A}_{H:L}(\widetilde{\gamma}) = \int_0^1 H(\gamma(t), t)\,\mathrm{d}t - \int \widehat{\gamma}^*\omega,$$

whose critical set $\mathrm{Crit}(\mathcal{A}_{H:L})$ comprises those $\widetilde{\gamma}$ for which γ is an integral curve of the Hamiltonian flow of H. Zapolsky's Lagrangian Floer complex is then defined by

$$CF^*_{\mathrm{Zap}}(L;H;R) = \bigoplus_{\widetilde{\gamma}\in\mathrm{Crit}(\mathcal{A}_{H:L})} C(\widetilde{\gamma}),$$

where $C(\widetilde{\gamma})$ is a free rank 1 R-module canonically associated to $\widetilde{\gamma}$. The grading is by Conley–Zehnder index.

After choosing a generic compatible almost complex structure and appropriate perturbation data, the differential counts pseudoholomorphic strips asymptotic to integral curves (recall our slightly unusual convention from Section 1.2.4 that when the asymptotic curves go from bottom to top, the strips defining the differential go from left to right), but keeps track of their homotopy classes. Explicitly, for each critical point $\widetilde{\gamma}_0 = [\gamma_0, \widehat{\gamma}_0]$, and each rigid pseudoholomorphic strip u from γ_0 to another integral curve γ_1, we consider the capping class $\widetilde{\gamma}_1 = [\gamma_1, \widehat{\gamma}_1]$ of γ_1 given by the concatenation $\widehat{\gamma}_0\#u$ of u with our capping of γ_0. Zapolsky constructs an isomorphism $C(u)\colon C(\widetilde{\gamma}_0) \to C(\widetilde{\gamma}_1)$, and this gives the contribution of u to the differential on $C(\widetilde{\gamma}_0)$.

Now suppose that we are given auxiliary data $\mathscr{D} = (f, g, J')$ as in Section 2.1.2, and consider Zapolsky's Lagrangian pearl complex

$$C^*_{\mathrm{Zap}}(L;\mathscr{D};R) = \bigoplus_{x\in\mathrm{Crit}(f)} \bigoplus_{\beta\in\pi_2(X,L,x)} C(x,\beta),$$

where $C(x,\beta)$ is a certain canonical free rank 1 R-module (we say *Lagrangian* pearl complex to distinguish it from the pearl-type complex used to compute quantum cohomology of X). Note that we can view the pair (x,β) as the constant path at x equipped with the homotopy class β of capping, and to unify the notations we will allow $\widetilde{\gamma} = [\gamma, \widehat{\gamma}]$ also to denote a pair (x,β). Similarly we can view each pearly trajectory as a strip asymptotic to the constant paths at its end points, by thinking of each Morse flowline as a crushed strip, and each disc as a strip asymptotic to its two marked points, and then concatenating these strips along the trajectory. The differential can then be constructed analogously to the Floer differential: for each rigid pearly trajectory u from x to y, and each β in $\pi_2(X,L,x)$, we have an induced class $\beta\#u$ in $\pi_2(X,L,y)$, and Zapolsky defines an isomorphism $C(u)\colon C(x,\beta) \to C(y,\beta\#u)$ which gives the contribution of u to the differential on $C(x,\beta)$.

Comparison (or continuation) maps between different choices of auxiliary data can be defined in a similar way, as can product operations, with the latter counting surfaces with three strip-like ends in the Floer picture (as in Fig. 1.1) and Y-shaped trajectories in the pearl picture. Just as we view a pearly trajectory as a strip, by collapsing regions of the strip onto the Morse flowlines, we can likewise view a Y-shaped pearly trajectory as a surface with three strip-like ends, so again the pictures are formally the same. Zapolsky also describes how to set up quantum cohomology of X and construct the quantum module action on Floer cohomology, in both pictures. In this case the space Ω of contractible loops $y\colon [0,1] \to X$ in X, and its cover $\widetilde{\Omega}$ in which loops are equipped with homotopy classes of capping $\widehat{y}$, appear in place of Ω_L and $\widetilde{\Omega}_L$. He then constructs the PSS isomorphisms between the two pictures and proves that they intertwine the product and module structures.

A.2 Zapolsky systems

Now we can discuss the use of local systems. First note that in both the Lagrangian Floer and Lagrangian pearl frameworks the complex is a direct sum of modules over points of $\widetilde{\Omega}_L$, and the differential counts certain strips

$$u\colon (\mathbb{R}\times[0,1], \mathbb{R}\times\{0,1\}) \to (X, L),$$

with specified asymptotics. Because of these asymptotics we can always work with compactified domains $[0,1]\times[0,1]$, on which we take coordinates (s', t). Such a strip can be foliated by arcs of fixed s'-value, and hence defines a path $p(u)$ in Ω_L, given by $s' \mapsto u(s', \cdot)$, which connects the two asymptotic arcs $u(0,\cdot)$ and $u(1,\cdot)$.

We can lift $p(u)$ to a path $\widetilde{p}(u)$ in $\widetilde{\Omega}_L$, and the capping defined by the final end point is the concatenation of the strip u with the capping defined by the initial end point. Moreover, the homotopy class of $p(u)$ relative to its end points depends only on the homotopy class of u relative to its end asymptotics, and in particular does not depend on our choice of foliation of $[0,1]\times[0,1]$.

A Zapolsky system is a local system $\mathscr{E}$ (in the usual sense of a locally constant sheaf) of R-modules on $\widetilde{\Omega}_L$. To incorporate such a system into the Lagrangian Floer or pearl complex, we simply tensor each $C(\widetilde{\gamma})$ with the fibre $\mathscr{E}_{\widetilde{\gamma}}$ of $\mathscr{E}$ over the corresponding point. The contribution of a strip u to the differential on $C(\widetilde{\gamma})\otimes_R \mathscr{E}_{\widetilde{\gamma}}$ is then twisted by the parallel transport $\mathcal{P}_{\mathscr{E}}(\widetilde{p}(u))$ of $\mathscr{E}$ along the path $\widetilde{p}(u)$ defined by u starting at $\widetilde{\gamma}$. This is compatible with differentials, comparison maps and PSS isomorphisms (all similarly twisted by $\mathscr{E}$) since breaking of pseudo-holomorphic strips and pearly trajectories preserves the homotopy classes of the corresponding parallel transport paths.

Remark A.2.1. This is the correct thing to do if one views Morse theory on $\widetilde{\Omega}_L$ as the blueprint for Floer theory. //

Zapolsky does not explicitly describe the effect of local systems on the product and module structures, so we briefly do this now. These modifications are straightforward and only affect the $\mathscr{E}_{\widetilde{\gamma}}$ factor which is tensored with each $C(\widetilde{\gamma})$. The basic idea is to cap off all but two ends of each surface being counted in order to make it into a (homotopy class of) strip, which is then viewed as a (homotopy class of) path in Ω_L and used to define a parallel transport map between the required fibres, just as for the differentials. The fact that this respects all of the structures again follows from the fact that breaking of strips and trajectories preserves homotopy classes.

The quantum module structure counts punctured strips u with a cylindrical end. The inputs are a path γ_0 in X with ends on L and a loop γ_X in X, each equipped with a homotopy class of capping, $\widehat{\gamma}_0$ and $\widehat{\gamma}_X$, whilst the output is another path γ_1 on X with ends on L, equipped with capping $\widehat{\gamma}_1 = \widehat{\gamma}_0\#u\#\widehat{\gamma}_X$. This configuration is illustrated in Fig. A.2. We make this into a strip from γ_0 to γ_1 simply by capping the cylindrical end with $\widehat{\gamma}_X$.

For the product, we count surfaces u with three strip-like ends from $\widetilde{\gamma}_0$ and $\widetilde{\gamma}_1$ to

$$[\gamma_2, \widehat{\gamma}_0\#u\#\widehat{\gamma}_1].$$

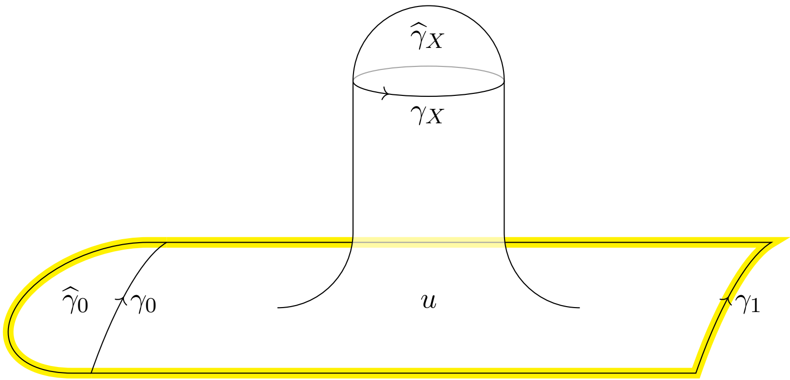

Figure A.2: A pseudoholomorphic curve u defining the quantum module structure, with cappings.

To define the parallel transport from $\mathscr{E}_{\widetilde{\gamma}_0}$ to $\mathscr{E}_{\widetilde{\gamma}_2}$ we cap off the γ_1 end of u with $\widehat{\gamma}_1$ and use the resulting strip $\widehat{\gamma}_1\#u$ to give a path $p(\widehat{\gamma}_1\#u)$ from γ_0 to γ_2, as shown schematically in Fig. A.3, where the blue arcs indicate the value of the path at different values of s' (as mentioned above, the particular choice of foliation by blue arcs is irrelevant, as is the representative $\widehat{\gamma}_1$ of the capping homotopy class). We similarly cap the γ_0 end with $\widehat{\gamma}_0$ to define a path $p(\widehat{\gamma}_0\#u)$ from

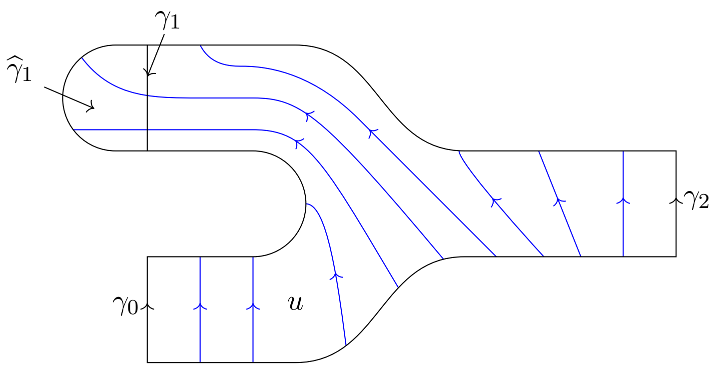

Figure A.3: The path from γ_0 to γ_2 defined by $\widehat{\gamma}_1\#u$.

γ_1 to γ_2. The product is then then given by tensoring Zapolsky's product map

$$C(u)\colon C(\widetilde{\gamma}_0)\otimes C(\widetilde{\gamma}_1)\to C(\widetilde{\gamma}_2)$$

with the parallel transport maps

$$\mathcal{P}_{\mathscr{E}}(\widetilde{p}(\widehat{\gamma}_1\#u))\otimes\mathcal{P}_{\mathscr{E}}(\widetilde{p}(\widehat{\gamma}_0\#u))\colon \mathscr{E}_{\widetilde{\gamma}_0}\otimes\mathscr{E}_{\widetilde{\gamma}_1}\to\mathscr{E}_{\widetilde{\gamma}_2}^{\otimes 2}.$$

This gives a product from the complexes twisted by $\mathscr{E}$ to the complex twisted by $\mathscr{E}^{\otimes 2}$, and there is an obvious generalisation to the case where the two inputs are twisted by different Zapolsky systems but we do not pursue this. If $\mathscr{E}$ is actually a sheaf of R-*algebras* on $\widetilde{\Omega}_L$ then we can use the multiplication $\mathscr{E}^{\otimes 2}\to\mathscr{E}$ to get an output in the same complex as the inputs.

Remark A.2.2. Higher A_∞-operations can be defined similarly by counting pseudoholomorphic curves with more incoming strip-like ends. They are twisted by a Zapolsky system by capping all but one of the incoming ends to define the path from the remaining incoming end to the outgoing end. //

We can similarly introduce a local system $\mathscr{E}_X$ on the space $\widetilde{\Omega}$ appearing in the construction of quantum cohomology, and use it to deform the differential and operations in an analogous way. We now describe briefly how the quantum module structure is twisted. Since our main interest in quantum cohomology is in its product structure, we assume that $\mathscr{E}_X$ is a sheaf of R-algebras.

Let $\widetilde{\Omega}'$ denote the set of pairs $[y, \widehat{y}]$ in $\widetilde{\Omega}$ with the property that $y|_{[1/2,1]}$ maps to L. There is a map

$$f \colon \widetilde{\Omega}' \to \widetilde{\Omega}_L,$$

given by sending $[y, \widehat{y}]$ to the point $\widetilde{\gamma}$ with $\gamma(t) = y(t/2)$ for $t \in [0,1]$, and $\widehat{\gamma} = \widehat{y}$. In other words, we take a disc $\widehat{y} \colon D \to X$, half of whose boundary y maps to L, and view the other half of the boundary as a path from L to itself, capped by the disc. In order to incorporate $\mathscr{E}_X$ into the quantum module structure we assume that its restriction from $\widetilde{\Omega}$ to $\widetilde{\Omega}'$ is identified with the pullback by f of a local system $\mathscr{E}_{X,L}$ on $\widetilde{\Omega}_L$. We also need to assume that our Zapolsky system $\mathscr{E}$ for L is a sheaf of $\mathscr{E}_{X,L}$-modules.

Remark A.2.3. Let E_X denote the fibre of $\mathscr{E}_X$ over $\overline{*}$, $m_X \colon \pi_1(\widetilde{\Omega}, \overline{*}) \to \operatorname{End} E_X$ the homomorphism defining $\mathscr{E}_X$, and $\iota \colon \widetilde{\Omega}' \hookrightarrow \widetilde{\Omega}$ the inclusion map. Concretely, m_X is the composition of the monodromy representation with the inversion map on $\pi_1(\widetilde{\Omega}, \overline{*})$. Equivalently it is the action of $\pi_1(\widetilde{\Omega}, \overline{*})$ on the fibres of the trivial E_X-bundle over the universal cover of $\widetilde{\Omega}$ such that the quotient is $\mathscr{E}_X$. Identifications of $\mathscr{E}_X|_{\widetilde{\Omega}'}$ with the pullback under f of a local system on $\widetilde{\Omega}_L$ correspond to homomorphisms $m_L \colon \pi_1(\widetilde{\Omega}_L, \overline{*}) \to \operatorname{End} E_X$ satisfying $f^* m_L = \iota^* m_X$ as homomorphisms $\pi_1(\widetilde{\Omega}', \overline{*}) \to \operatorname{End} E_X$.

We saw in Remark A.1.1 that the inclusion $L \to \widetilde{\Omega}_L$ of constant loops with constant cappings and the projection $\widetilde{\Omega}_L \to L$ given by $\widetilde{\gamma} \mapsto \gamma(0)$ induce mutually inverse isomorphisms on π_1. A similar argument with tailed cappings shows that the corresponding inclusion $X \to \widetilde{\Omega}$ and projection $\widetilde{\Omega} \to X$ by $[y, \widehat{y}] \mapsto y(0)$ also induce inverse isomorphisms on π_1. We therefore see that for any m_X there exists a unique choice for m_L, namely the restriction of m_X from $\pi_1(\widetilde{\Omega}, \overline{*}) = \pi_1(X, *)$ to $\pi_1(\widetilde{\Omega}_L, \overline{*}) = \pi_1(L, *)$.

We conclude that for any $\mathscr{E}_X$ there exists a unique (up to isomorphism) local system $\mathscr{E}_{X,L}$ on $\widetilde{\Omega}_L$ with an identification $f^*\mathscr{E}_{X,L} = \mathscr{E}_X|_{\widetilde{\Omega}'}$. Explicitly, given two points $[y_0, \widehat{y}_0]$ and $[y_1, \widehat{y}_1]$ in $\widetilde{\Omega}'$ with the same image $\widetilde{\gamma}$ in $\widetilde{\Omega}_L$, we identify the corresponding fibres of $\mathscr{E}_X$ by parallel transport along any path $[y_t, \widehat{y}_t]$ in $\widetilde{\Omega}$ for which $y_t(0)$ is constant. The fibre of $\mathscr{E}_{X,L}$ over $\widetilde{\gamma}$ is defined to be the direct sum of the fibres of $\mathscr{E}_X$ over preimages of $\widetilde{\gamma}$, modulo these identifications. //

Given a curve u contributing to the module structure, as in Fig. A.2, its contribution is defined by tensoring Zapolsky's map

$$C(u) \colon C(\widetilde{\gamma}_X) \otimes C(\widetilde{\gamma}_0) \to C(\widetilde{\gamma}_1)$$

with the composition

$$(\mathscr{E}_X)_{\widetilde{\gamma}_X} \otimes \mathscr{E}_{\widetilde{\gamma}_0} \xrightarrow{a \otimes b} (\mathscr{E}_{X,L})_{\widetilde{\gamma}_1} \otimes \mathscr{E}_{\widetilde{\gamma}_1} \to \mathscr{E}_{\widetilde{\gamma}_1}. \tag{A.1}$$

Here the second arrow is the $\mathscr{E}_{X,L}$-module structure map, and $b \colon \mathscr{E}_{\widetilde{\gamma}_0} \to \mathscr{E}_{\widetilde{\gamma}_1}$ is the parallel transport arising from $\widehat{\gamma}_X \# u$ as discussed earlier. The remaining map,

$$a\colon (\mathscr{E}_X)_{\widetilde{\gamma}_X} \to (\mathscr{E}_{X,L})_{\widetilde{\gamma}_1},$$

is constructed as follows. We view $\widehat{\gamma}_0 \# u$ as a path in $\widetilde{\Omega}$ from $\widetilde{\gamma}_X$ to the pair $\widetilde{y} = [y, \widehat{y}]$, where y is the loop defined by $y(t) = \gamma_1(2t)$ for t in $[0, 1/2]$, closed up around the path highlighted in yellow in the figure, and $\widehat{y}$ is the capping $\widehat{\gamma}_0 \# u \# \widehat{\gamma}_X$. This gives a parallel transport map

$$(\mathscr{E}_X)_{\widetilde{\gamma}_X} \to (\mathscr{E}_X)_{\widetilde{y}},$$

and composing with the identification of $(\mathscr{E}_X)_{\widetilde{y}}$ with $(\mathscr{E}_{X,L})_{f(\widetilde{y})=\widetilde{\gamma}_1}$ defines a.

A.3 Taking quotients

We have now seen how to construct the large complexes and deform them with Zapolsky's general notion of local system. In light of Remark A.1.1 such a system $\mathscr{E}$ is specified by a homomorphism $m\colon \pi_1(L, *) \to \operatorname{End} \mathscr{E}_*$, so is apparently no more general than an ordinary local system on L. However, we now describe Zapolsky's quotient procedure, focusing on the Lagrangian Floer and Lagrangian pearl complexes (the quantum cohomology complexes are dealt with later), which provide an extra level of flexibility. The story is essentially the same for both the Floer and pearl complexes, with the obvious translation $[\gamma, \widehat{\gamma}] \leftrightarrow (x, \beta)$, and for the sake of concreteness we base our discussion on the former.

For each $\gamma \in \Omega_L$ the group $\pi_2(X, L, \gamma(0))$ acts freely and transitively on the set of homotopy classes of cappings of γ by attaching discs at $\gamma(0)$. In order to get a left action, a class β should act by attaching $-\beta$, as shown in Fig. A.4. Here $-\beta$ is viewed as a map from the unit square as in Section 1.2.4, squashed down into a half-disc so that the top and vertical edges (which map to $\gamma(0)$) become the curved boundary. The space above the semicircle in the diagram is foliated by copies of γ, indicated by the lines with arrows. Zapolsky proves that a choice of

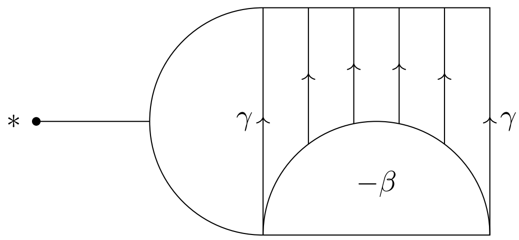

Figure A.4: The action of β on $\widetilde{\gamma}$ by attaching $-\beta$.

relative spin structure on L (or the restriction to char $R = 2$) is sufficient to define identifications $C([\gamma, \widehat{\gamma}]) \to C([\gamma, \widehat{\gamma}'])$ for different cappings, and that these identifications induce the structure of a $\pi_2(X, L, \gamma(0))$-module on

$$C_{\oplus}(\gamma) \coloneqq \bigoplus_{[\widehat{\gamma}]} C([\gamma, \widehat{\gamma}]).$$

Moreover, given a rigid pseudoholomorphic strip u from γ_0 to γ_1 there is a parallel transport map

$$\mathcal{P}_{\pi_2(X,L)}(u_{t=0})\colon \pi_2(X,L,\gamma_0(0)) \to \pi_2(X,L,\gamma_1(0))$$

along the '$t=0$' edge of u (u is a path of paths, and we want to take the path of initial points), and the actions of $\pi_2(X,L,\gamma_j(0))$ on $C_\oplus(\gamma_j)$ are intertwined by $\mathcal{P}_{\pi_2(X,L)}(u_{t=0})$ and

$$C(u)\colon C_\oplus(\gamma_0) \to C_\oplus(\gamma_1).$$

In this sense we say that the Zapolsky Floer complex is a module over the natural local system $\pi_2(X,L)$ of groups on L with fibre $\pi_2(X,L,p)$ over a point p. The identifications are also compatible with the comparison maps, product, module structure and PSS isomorphisms (and presumably also higher operations, although he doesn't claim this and we don't need it) in an analogous way.

Suppose we choose a system of subgroups $G \subset \pi_2(X,L)$. Using the above $\pi_2(X,L)$-action we can form the quotient complex by the action of this subsystem, and by the intertwining property the differential descends, as well as the product, module structure and PSS isomorphisms. If the system G is normal in $\pi_2(X,L)$ then the quotient complex is a module over $\pi_2(X,L)/G$. Note that if $\pi_2(X,L)/G$ is the trivial local system modelled on a group H then a module over this system is the same thing as a module over the group ring $R[H]$. In particular: $G = \ker(\pi_2(X,L) \to H_2(X,L;\mathbb{Z}))$ recovers the pearl complex over $R[H_2^D]$, of which the enriched pearl complex is a subcomplex (obtained by restricting to $(H_2^D)^+ \cup \{0\}$); $G = \ker(\mu\colon \pi_2(X,L) \to \mathbb{Z})$ recovers the pearl complex over the usual Novikov ring Λ (this isn't quite true if $N_L = 0$, but recall we are assuming $N_L \geq 2$); $G = \pi_2(X,L)$ gives the complex in which the Novikov variable is set to 1 and the grading collapsed.

In order for a Zapolsky system $\mathscr{E}$ to descend to the quotient, the action of $G \subset \pi_2(X,L)$ on $\widetilde{\Omega}_L$ must lift to $\mathscr{E}$. In other words, we need to define isomorphisms (of R-modules or -algebras) between each fibre $\mathscr{E}_{\widetilde{\gamma}}$ and its translates $\mathscr{E}_{g\cdot\widetilde{\gamma}}$, which respect the sheaf structure (i.e. are locally constant as we vary the base points $\widetilde{\gamma}$ for $\mathscr{E}$ and $\gamma(0)$ for G), and assemble into a $G_{\gamma(0)}$-module structure on

$$\mathscr{E}_\gamma \coloneqq \bigoplus_{[\widetilde{\gamma}]} \mathscr{E}_{\widetilde{\gamma}}$$

for each γ. Given a path Γ in $\widetilde{\Omega}_L$, these G-module structures are necessarily intertwined by the parallel transport maps $\mathcal{P}_{\mathscr{E}}(\Gamma)$ and $\mathcal{P}_G(\Gamma_{t=0})$, so the quotient system is compatible with the differentials and other structures on the quotient complexes.

Example A.3.1. Take the trivial local system on the universal cover $\widetilde{L}$ of L and push it forward to L via the projection map. Now let $\mathscr{E}$ be the pullback of this system to $\widetilde{\Omega}_L$ under the map $\widetilde{\gamma} \mapsto \gamma(0)$. For any $\widetilde{\gamma}$ in $\widetilde{\Omega}_L$ and any class β in $\pi_2(X,L,\gamma(0))$ the fibres over $\widetilde{\gamma}$ and $\beta \cdot \widetilde{\gamma}$ are canonically identified, which defines a lift of the $\pi_2(X,L)$-action on $\widetilde{\Omega}_L$ to $\mathscr{E}$. The resulting quotient complex computes Damian's Floer theory on the universal cover [38]. //

Remark A.3.2. If the system G is contained in $\ker(\mu\colon \pi_2(X,L) \to \mathbb{Z})$ then the quotient complex retains a $\mathbb{Z}$-grading and a filtration by Maslov index. We can thus construct the Oh spectral sequence as in Section 2.2.1, described by Zapolsky (in the case where $\mathscr{E}$ is trivial and G is just the identity subsystem in $\pi_2(X,L)$) in [146, Section 4.6]. In general, the first page of the spectral sequence is the homology of L with coefficients in the local system obtained as follows: form the quotient system $\mathscr{E}/G$ on $\widetilde{\Omega}_L/G$, push forward to Ω_L, and pull back under the inclusion $L \hookrightarrow \Omega_L$ of constant paths. //

Recall from Remark A.1.1 that after fixing a base point $*$ in L we have a model $\overline{\Omega}_L$ for the universal cover of Ω_L comprising paths $\gamma \in \Omega_L$ equipped with homotopy classes of tailed cappings. There is a natural action of G_* (meaning the fibre of the system G over the point $*$) on this space, under which a disc g acts on a tailed capping of a path γ by attaching the boundary ∂g to the tail and the disc $-g'$ to the capping at $\gamma(0)$, where g' is obtained from g by parallel transport from $*$ along the tail and the lower half of the boundary of the capping. This gives an embedding i of G_* in $\pi_1(\Omega_L, \overline{*})$.

We also have the embedding j of $\pi_1(\widetilde{\Omega}_L, \overline{*}) = \pi_1(L, *)$ in $\pi_1(\Omega_L, \overline{*})$ given by attaching loops at the tail. It is easy to see that for a disc $g \in G_*$ and a loop $\gamma \in \pi_1(L, *)$ we have $j(\gamma)i(g)j(\gamma)^{-1} = i(g')$, where g' is the disc $\gamma \cdot g$ obtained by parallel transporting g around $-\gamma$, and so our two embeddings i and j induce a homomorphism

$$G_* \rtimes \pi_1(L, *) \to \pi_1(\Omega_L, \overline{*}),$$

where the action of $\pi_1(L, *)$ on G_* is as just described. Since G acts freely on the fibres of $\widetilde{\Omega}_L$ over Ω_L this homomorphism is injective, and its image is clearly $\pi_1(\widetilde{\Omega}_L/G, \overline{*})$. In other words, we have a natural identification

$$\pi_1(\widetilde{\Omega}_L/G, \overline{*}) = G_* \rtimes \pi_1(L, *).$$

Remark A.3.3. This decomposition of the fundamental group of $\widetilde{\Omega}_L/G$ corresponds to viewing it as a quotient of $\overline{\Omega}_L$ first by G_* and then by $\pi_1(L, *)$. This is more convenient than quotienting by $\pi_1(L, *)$ first and then by G—which is what seems natural from the notation $\widetilde{\Omega}_L/G$—because the possible non-constancy of the local system G (which corresponds to the non-triviality of the semidirect product) means that the covering $\widetilde{\Omega}_L \to \widetilde{\Omega}_L/G$ need not be regular. //

Armed with this identification of $\pi_1(\widetilde{\Omega}_L/G, \overline{*})$, we see that given a Zapolsky system $\mathscr{E}$ defined by a homomorphism

$$m\colon \pi_1(L, *) = \pi_1(\widetilde{\Omega}_L, \overline{*}) \to \operatorname{End} E,$$

where E is the fibre of $\mathscr{E}$ over $\overline{*}$, lifts of the G-action to $\mathscr{E}$ are described by homomorphisms $q\colon G_* \to \operatorname{End} E$ satisfying

$$m(\gamma)q(g) = q(\gamma \cdot g)m(\gamma)$$

for all $\gamma \in \pi_1(L, *)$ and all $g \in G_*$. Explicitly, given a lift of the G-action to $\mathscr{E}$ and an element g of G_* we construct $q(g)$ by using lift of g to map $E = \mathscr{E}_{\overline{*}} \to \mathscr{E}_{g \cdot \overline{*}}$ and then parallel transporting $\mathscr{E}_{g \cdot \overline{*}}$ back to $\mathscr{E}_{\overline{*}}$ along the path which homotopes away the tailed capping of the constant path at $*$ shown in Fig. A.5; the yellow regions all map to $*$.

We now discuss the two particular types of Zapolsky system we are interested in. First consider the standard twisting by local systems $\mathscr{F}^0$ and $\mathscr{F}^1$ of R-modules on L, as in Remark 2.1.8. We have projections $p_0, p_1\colon \widetilde{\Omega}_L \to L$, given by $p_t(\widetilde{\gamma}) = \gamma(t)$, and we define the Zapolsky system

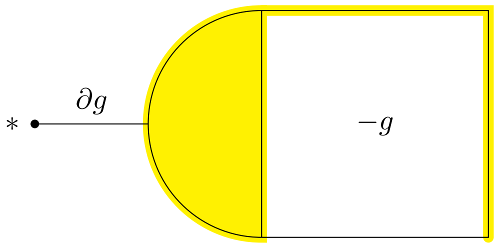

Figure A.5: The tailed capping used to define q.

$\mathscr{E}$ to be

$$\mathscr{H}om(p_0^*\mathscr{F}^0, p_1^*\mathscr{F}^1),$$

where $\mathscr{H}om$ denotes the sheaf of homomorphisms over the constant sheaf $\underline{R}$. Concretely, the fibre over a capped path $\widetilde{\gamma}$ is $\mathrm{Hom}_R(\mathscr{F}^0_{\gamma(0)}, \mathscr{F}^1_{\gamma(1)})$, and this gives a natural identification of fibres over different G-translates (for any subsystem $G \subset \pi_2(X, L)$) which we use to lift the G-action to $\mathscr{E}$. It is immediate from the construction that this coincides with the twisting described in Remark 2.1.8 when G is the kernel of the Maslov index homomorphism.

If $m^j : \pi_1(L, *) \to \mathrm{End}\, F^j$ is the homomorphism defining $\mathscr{F}^j$, where F^j denotes the fibre of $\mathscr{F}^j$ over $*$, then $m\colon \pi_1(L, *) \to \mathrm{End}(\mathrm{Hom}(F^0, F^1))$ is given by

$$m(\gamma) \cdot \theta = m^1(\gamma) \circ \theta \circ (m^0(\gamma))^{-1}$$

for all $\gamma \in \pi_1(L, *)$ and all $\theta \in \mathrm{Hom}(F^0, F^1)$. Recall that $m^0(\gamma)$ and $m^1(\gamma)$ represent the monodromies of $\mathscr{F}^0$ and $\mathscr{F}^1$ around *minus* γ. The map q is defined by dragging $\gamma(1)$ along the top edge of the capping in Fig. A.5 to the end of the tail—which corresponds to the monodromy of $\mathscr{F}^1$ along $-\partial g$, i.e. $m^1(\partial g)$—and similarly dragging $\gamma(0)$ along the bottom edge, which does nothing since this path is nullhomotopic. In other words, $q(g)$ maps $\theta \in \mathrm{Hom}(F^0, F^1)$ to $m^1(g) \circ \theta$.

The other type of Zapolsky system $\mathscr{E}$ we need to consider is that corresponding to an H_2^D local system defined by $\rho \in \mathrm{Hom}(H_2^D, R^\times)$, as discussed in Section 2.2.3. For this we take $\mathscr{E}$ to be the constant sheaf $\underline{R}$ (so the map m above is trivial), and define the map q to be the composition

$$G_* \hookrightarrow \pi_2(X, L, *) \twoheadrightarrow H_2^D \xrightarrow{\rho^{-1}} R^\times = \mathrm{End}\, R.$$

This works for any subsystem $G \subset \pi_2(X, L)$. It is straightforward to verify that when G is the whole of $\pi_2(X, L)$ this coincides with the H_2^D local system defined in Section 2.2.3.

Remark A.3.4. Viewing this quotient as a two step process, going via the quotient by

$$\ker(\pi_2(X, L) \to H_2(X, L; \mathbb{Z})),$$

we also recover the perspective of Remark 2.2.9, where the deformation is viewed as a modification of the reduction map $\Lambda^\dagger \to \Lambda$. //

Remark A.3.5. Zapolsky systems, and G-actions on them, can be tensored together. In this way we can form composite local systems as in Section 2.2.3. //

Finally we discuss quotienting the complexes for quantum cohomology. The ideas are completely analogous to the above, with a system $H \subset \pi_2(X)$ on X replacing $G \subset \pi_2(X, L)$ on L, so for example quotienting by $\ker(c_1 \colon \pi_2(X) \to \mathbb{Z})$ gives the usual quantum cohomology over Λ. For consistency with the Lagrangian case, a class β in $\pi_2(X, y(0))$ acts on $[y, \widehat{y}]$ by gluing $-\beta$ to $\widehat{y}$, even though the minus sign is no longer necessary in order to get a left action (as $\pi_2(X, \gamma(0))$ is abelian). There are, however, two points to note with respect to the quantum module structure.

First, because of the capping by $\widetilde{\gamma}_X$ involved in the module structure (see Fig. A.2), if we quotient the quantum cohomology complexes by a system $H \subset \pi_2(X)$ then we must quotient the Lagrangian complexes by a system $G \subset \pi_2(X, L)$ containing the image of H.

Second, suppose the quantum cohomology and Lagrangian complexes are twisted by Zapolsky systems $\mathscr{E}_X$ and $\mathscr{E}$ respectively. Recall we required that the restriction of $\mathscr{E}_X$ to the set $\widetilde{\Omega}'$ of 'capped loops whose second half lies in L' be pulled back from a local system $\mathscr{E}_{X,L}$ on $\widetilde{\Omega}_L$. For compatibility with the quotient procedure we need the action of the image H' of H in $\pi_2(X, L)$ on $\widetilde{\Omega}_L$ to lift to $\mathscr{E}_{X,L}$ such that the identification with $\mathscr{E}_X$ over $\widetilde{\Omega}'$ is H-equivariant. We also need the module action (A.1) to respect the actions of G and H.

Using the explicit description of $\mathscr{E}_{X,L}$ in Remark A.2.3, an equivariant lift of the H'-action exists (and is necessarily unique) if and only if the elements of $H \cap \pi_2(L) = \ker(H \to H')$ act trivially in the following sense. For every $h \in H_* \cap \pi_2(L, *)$ we can use the lift of h to obtain a homomorphism

$$(\mathscr{E}_X)_{\overline{*}} \to (\mathscr{E}_X)_{h \cdot \overline{*}}$$

and then can map back the other way by parallel transport along any path $[y_t, \widehat{y}_t]$ with $y_t(0) = *$ for all t. We need this composition to be the identity.

If $m_X \colon \pi_1(X, *) = \pi_1(\widetilde{\Omega}, \overline{*}) \to \operatorname{End} E_X$ is the homomorphism defining $\mathscr{E}_X$, where E_X is the fibre over $\overline{*}$, then lifts of the H-action to $\mathscr{E}_X$ correspond to homomorphisms $q_X \colon H_* \to \operatorname{End} E_X$ satisfying

$$m_X(\gamma) q_X(h) = q_X(\gamma \cdot h) m_X(\gamma)$$

for all $\gamma \in \pi_1(X, *)$ and all $h \in H_*$. This is just as in the Lagrangian case considered above, except that ∂h is trivial in $\pi_1(X, *)$ so when we homotope away the tailed capping to define q_X we don't have to touch the tail. The condition for the equivariant H'-lift to exist can therefore be phrased as $H_* \cap \pi_2(L, *) \subset \ker q_X$.

If m and q are the corresponding homomorphisms for the local system $\mathscr{E}$ on $\widetilde{\Omega}_L$, with fibre E over $\overline{*}$, then the module action of $\mathscr{E}_{X,L}$ on $\mathscr{E}$ is described by a map $\varphi \colon E_X \otimes E \to E$ (note that the fibre of $\mathscr{E}_{X,L}$ over $\overline{*}$ is canonically identified with E_X). Respect for the actions amounts to

$$\varphi \circ (\mathrm{id}_{E_X} \otimes q(g)) = q(g) \circ \varphi$$

for all g in G_* and

$$\varphi \circ (q_X(h) \otimes \mathrm{id}_E) = q(h') \circ \varphi$$

for all h in H_*, where h' denotes the image of h in G_*. The fact that $\mathrm{id}_{E_X} \otimes q(g)$ commutes with $q_X(h) \otimes \mathrm{id}_E$ for all g and h is consistent with these constraints, since the image of $\pi_2(X, *)$ in $\pi_2(X, L, *)$ is contained in the centre of the latter.

As a brief example, suppose $\mathscr{E}$ and $\mathscr{E}_X$ are the trivial local systems with fibre $\mathbb{C}$, and m is the standard module action of $\mathbb{C}$ on itself. Given a class $b \in H^2(X;\mathbb{C})$ we can define a homomorphism $q_X \colon \pi_2(X,*) \to \mathbb{C}^*$ by

$$\beta \mapsto e^{2\pi i \langle b, \beta \rangle}.$$

If the restriction of b to L lies in the image of $H^2(L;\mathbb{Z})$ then $\pi_2(L,*)$ lies in the kernel of q_X, and hence there exists a unique equivariant lift of the $\pi_2(X)$-action on $\widetilde{\Omega}'$ to $\mathscr{E}_{X,L}$. In order to define a lift of the $\pi_2(X,L)$-action on $\widetilde{\Omega}_L$ to $\mathscr{E}$ compatible with the module structure, we need to give a homomorphism $q \colon \pi_2(X,L,*) \to \mathbb{C}^*$ which coincides with q_X on the image of $\pi_2(X,*)$. One way to do this is to pick a complex line bundle $\mathcal{L}$ on L with $c_1(\mathcal{L}) = -b|_L$, and a connection ∇ on $\mathcal{L}$ with curvature F, and then define q on a disc β by

$$q(\beta) = e^{\int_\beta F} \operatorname{hol}_{\partial\beta}(\nabla).$$

This is exactly the construction of B-field twisting described in Section 2.2.3.

A.4 Obstruction

Non-triviality of the local system $\mathscr{E}$ can cause the bubbling of index 2 discs not to cancel, obstructing $\mathrm{d}^2 = 0$. Explicitly, suppose $\widetilde{\gamma}$ is a critical point of the action functional (or a capped critical point of the Morse function in the pearl case), corresponding to a generator of the Floer (or pearl) complex. Let E be the fibre of $\mathscr{E}$ over $\widetilde{\gamma}$ and let our base point $*$ in L be $\gamma(0)$.

Fix a second critical point $\widetilde{\gamma}'$ and consider the contribution to $\mathrm{d}^2\widetilde{\gamma}$ coming from $\widetilde{\gamma}'$. This can only be non-zero when $\widetilde{\gamma}'$ is of the form $-\beta \cdot \widetilde{\gamma}$, where $\beta \in \pi_2(X,L,*)$ is a class of index 2—the minus sign occurs because the action of a disc is by gluing on *minus* that disc. In this case, the contribution is given (up to an overall sign) by

$$\mathfrak{m}_{0,\beta} \cdot \big(m(\partial\beta)^{-1} - \mathrm{id}_E\big) \in \operatorname{End} E,$$

where $\mathfrak{m}_{0,\beta}$ is the signed count of holomorphic discs based at $*$ in class β with respect to a generic almost complex structure. If these quantities do not all vanish then d^2 is non-zero and the cohomology is not defined.

If, however, one is only interested in the complex obtained after quotienting by some subsystem $G \subset \pi_2(X,L)$, using a homomorphism $q \colon G_* \to \operatorname{End} E$, then these obstructions may cancel. The relevant condition is now

$$\sum_{\substack{g \in G_* \\ \mu(g)=0}} \mathfrak{m}_{0,\beta g} \cdot q(g) \circ \big(m(\partial\beta\partial g)^{-1} - \mathrm{id}_E\big) = 0 \tag{A.2}$$

for all index 2 classes β in $\pi_2(X,L,*)$. This is still not guaranteed to hold of course.

One way to avoid this problem is to restrict attention to Lagrangians of minimal Maslov index at least 3, as is done by Damian in his Floer theory on the universal cover. Another possibility is to replace $\mathscr{E}$ by the subsystem on which the endomorphisms (A.2) *do* vanish. This corresponds to the central subcomplex introduced by Konstantinov in [94] in the case where $\mathscr{E}$ comes from a pair of local systems on L. His paper also includes a thorough analysis of obstruction issues for higher rank local systems (in the usual sense) in general monotone Lagrangian Floer theory.

Appendix B

Orientation computations

B.1 Verifying Assumption 3.2.6

This appendix is concerned with various orientation issues. In this first subsection we show that the Biran–Cornea and Zapolsky orientation schemes satisfy Assumption 3.2.6, whose notation we follow.

We begin with the Biran–Cornea scheme, for which the necessary definitions are given in [17, Appendix A, Sections A.1.8 and A.2.3], recalling that Biran–Cornea use the quantum module structure in place of the closed–open map and work with downward Morse flows and homology rather than upward flows and *co*homology. From this perspective, the coefficient of 1_L in $\mathcal{CO}^0(\alpha)$ becomes the coefficient of $[L]$ in $[Z] * [L]$.

We see from [17, Equation (107)] that the disc u counts with positive sign if and only the map

$$TZ \oplus TX \oplus TL \oplus TL \oplus T\mathcal{M}_{\mu=2k} \oplus TL \to TX \oplus TX \oplus TL \oplus TL \oplus TL \oplus TL$$

given by

$$(v_1, v_2, v_3, v_4, v_5, v_6) \mapsto (v_1 - v_2, v_2 - D_u \operatorname{ev}_0(v_5), v_3 - v_4, v_4 - D_u \operatorname{ev}_{-1}(v_5), D_u \operatorname{ev}_1(v_5) - v_6, v_6)$$

is orientation-preserving (for brevity we drop the subscripts denoting the base points of the tangent spaces; these can be recovered by considering the domains and targets of the maps $D_u \operatorname{ev}_z$). Here TX carries its symplectic orientation and TZ carries its orientation as a pseudocycle, whilst TL carries the orientation we chose on L. Note that changing the orientation of L makes no difference, as it reverses the orientations on each TL and on $T\mathcal{M}$ (the orientation on disc moduli spaces is described in more detail in Section B.3). Using row and column operations, this can be reduced to the orientation sign of the infinitesimal evaluation map

$$T_u\mathcal{M}_{\mu=2k} \to T_{u(0)}X/T_{u(0)}Z \oplus T_{u(1)}L,$$

which is exactly what we want.

We now deal with Zapolsky's orientation scheme, which requires much more work. Let

$$u\colon (\Sigma, \partial\Sigma) \to (X, L)$$

be a pseudoholomorphic curve, and let D_u denote the Cauchy–Riemann operator obtained from linearising the pseudoholomorphic curve equation at u (throughout this appendix we ignore all analytic details regarding Sobolev spaces and specific smoothness hypotheses, as well as the fact that one needs to make some auxiliary choices in order to define a Cauchy–Riemann operator at curves which are not themselves pseudoholomorphic). Let $p_1, \dots, p_k$ be distinct interior marked points in Σ, and $U_1, \dots, U_k$ be subspaces of $T_{u(p_1)}X, \dots, T_{u(p_k)}X$ respectively. Similarly let $q_1, \dots, q_l$ be boundary marked points, with $V_1, \dots, V_l$ subspaces of the corresponding tangent spaces on L. We denote by $D_u[U_1, \dots, U_k; V_1, \dots, V_l]$ the restriction of D_u to the subspace of its domain comprising those infinitesimal deformations of u whose evaluation at each p_j (respectively q_j) lies in U_j (respectively V_j).

We will need to be able to glue constant discs or spheres on to a given curve, so let us consider this in the above setup. We shall focus on the case of gluing a disc at the boundary, but gluing a sphere in the interior is analogous. For brevity, let $U; V$ denote $U_1, \dots, U_k; V_1, \dots, V_{l-1}$—note the final subscript is $l-1$, not l. Assume that the operator $D_u[U; V]$ is surjective and that the evaluation map from its kernel to $T_{u(q_l)}L$ is transverse to V_l. Pick a complement $V_l^\perp$ to V_l which is contained in the image of this evaluation map, and let v be the constant disc at $u(q_l)$. Choose an arbitrary marked point q' on the boundary of v and view $V_l^\perp$ as a boundary condition for v at q', with corresponding operator $D_v[; V_l^\perp]$.

Now consider the direct sum $D_v[; V_l^\perp] \oplus D_u[U; V, V_l]$. This is naturally isomorphic to the operator obtained by dropping the $V_l^\perp$ and V_l boundary conditions, and replacing the direct sum by the fibre product over the subspace $V_l^\perp \oplus V_l$ of $T_{u(q_l)}L \oplus T_{u(q_l)}L$ under the evaluation maps at q' and q_l. We can then 'deform the incidence condition' (see [146]) from $V_l^\perp \oplus V_l$ to the diagonal to obtain the operator $D_v \sharp D_u[U; V]$, and glue to get $D_{v\sharp u}[U; V]$ (of course, one now has to view the marked points as living in the glued domain). The constant disc v can then be homotoped away. In this way we obtain an isomorphism of determinant lines which is well-defined up to homotopy (which we denote by $\simeq$):

$$\mathrm{d}(D_v[; V_l^\perp]) \otimes \mathrm{d}(D_u[U; V, V_l]) \simeq \mathrm{d}(D_u[U; V]).$$

In other words, we get

$$\Lambda^{\mathrm{top}} V_l^\perp \otimes \Lambda^{\mathrm{top}} \ker D_u[U; V, V_l] \simeq \Lambda^{\mathrm{top}} \ker D_u[U; V]. \tag{B.1}$$

Lemma B.1.1. *This coincides, up to homotopy, with the isomorphism obtained by identifying $V_l^\perp$ with $T_{u(q_l)}L/V_l$ and using the short exact sequence*

$$0 \to \ker D_u[U; V, V_l] \to \ker D_u[U; V] \to T_{u(q_l)}L/V_l \to 0,$$

with the right-to-left convention from Remark 3.2.7.

Proof. The boundary conditions at $p_1, \dots, p_k; q_1, \dots, q_{l-1}$ play no role in the argument so we assume $k = l - 1 = 0$, and drop the subscripts from q_l and V_l. Pick a basis $v_1, \dots, v_k$ for $V^\perp$ (we shall also view this as a basis for $\ker D_v[; V^\perp])$ and $v_{k+1}, \dots, v_n$ for V. Let θ denote the evaluation map from $\ker D_u$ at q, and let $\xi_1, \dots, \xi_k$ be lifts to $\ker D_u$ of $v_1, \dots, v_k$ under θ. Extend to a basis $\xi_1, \dots, \xi_m$ for $\ker D_u$.

To deform the incidence condition we use the family of subspaces $V_t \subset T_{u(q)}L \oplus T_{u(q)}L$ given by

$$V_t = \{(v' + tv, tv' + v) : v \in V, v' \in V^\perp\}$$

for $t \in [0, 1]$, which interpolates between $V^\perp \oplus V$ at $t = 0$ and $\Delta_{T_{u(q)}L}$ at $t = 1$. A family of bases for these spaces is given by $(v_1, tv_1), \dots, (v_k, tv_k), (tv_{k+1}, v_{k+1}), \dots, (tv_n, v_n)$, and a corresponding family of bases for the fibre product of D_v and D_u is

$$(v_1, t\xi_1), \dots, (v_k, t\xi_k), (t\theta\xi_{k+1}, \xi_{k+1}), \dots, (t\theta\xi_m, \xi_m).$$

The identification we obtain between $\mathrm{d}(D_v[; V^\perp]) \otimes \mathrm{d}(D_u[; V])$ and $\mathrm{d}(D_v \sharp D_u)$ is therefore

$$(v_1 \wedge \cdots \wedge v_k) \otimes (\xi_{k+1} \wedge \cdots \wedge \xi_m) \mapsto (v_1 = \theta\xi_1, \xi_1) \wedge \cdots \wedge (\theta\xi_m, \xi_m). \tag{B.2}$$

The final step comes from gluing and homotoping away the constant disc. We start by attaching a disc to Σ at the boundary point q, and then smoothing near this point—really we are identifying collar neighbourhoods of the attaching points in the two domains, with large gluing length. Σ carries the operator D_u, whilst the disc carries the Cauchy–Riemann operator D_v for the constant disc v, and these are glued to give the operator $D_{v\sharp u}$. Assuming the gluing length is sufficiently large, elements of $\ker D_{v\sharp u}$ can be approximated by elements of $\ker D_v \sharp D_u$, which is to say elements of $\ker D_u$ over Σ, extended constantly over the disc (interpolated over the gluing region). This identification is canonical up to homotopy and defines the gluing isomorphism $\ker D_v \sharp D_u \to \ker D_{v\sharp u}$.

We now need to understand what happens when the constant disc is homotoped away, which we can do by shrinking the disc in the domain. If we simultaneously increase the gluing length sufficiently fast, at each stage we preserve the identification of the kernel of the glued operator with constant extensions of elements of $\ker D_u$. In the limit we return to the domain Σ carrying the operator D_u, and we conclude that the required isomorphism $\mathrm{d}(D_{v\sharp u}) \simeq \mathrm{d}(D_u)$ is given by composing the inverse gluing isomorphism $\mathrm{d}(D_{v\sharp u}) \to \mathrm{d}(D_v \sharp D_u)$ with the projection $\mathrm{d}(D_v \sharp D_u) \to \mathrm{d}(D_u)$. In other words, the composite

$$\mathrm{d}(D_v \sharp D_u) \xrightarrow{\text{glue}} \mathrm{d}(D_{v\sharp u}) \xrightarrow{\text{homotope}} \mathrm{d}(D_u)$$

is given by

$$(\theta\xi_1, \xi_1) \wedge \cdots \wedge (\theta\xi_m, \xi_m) \mapsto \xi_1 \wedge \cdots \wedge \xi_m.$$

Combining this with (B.2), the identification (B.1) is

$$(v_1 \wedge \cdots \wedge v_k) \otimes (\xi_{k+1} \wedge \cdots \wedge \xi_m) \mapsto \xi_1 \wedge \cdots \wedge \xi_m.$$

This is easily seen to agree with the identification defined by the short exact sequence. $\square$

In order to compute the sign of the contribution of u to $\mathcal{CO}^0$ in the situation of Assumption 3.2.6, we take our disc u and consider marked points $p_1 = 0$, $q_1 = 1$, $q_2 = -1$ with corresponding boundary conditions $U_1 = T_{u(0)}Z$, $V_1 = 0$, $V_2 = T_{u(-1)}L$. The operator $D_u[TZ; 0, TL]$ is an isomorphism so its determinant line is canonically oriented. Gluing a constant sphere to u at p_1 and a constant disc at q_2 as above, we obtain an identification

$$\Lambda^{\mathrm{top}} TZ^{\perp} \otimes \Lambda^{\mathrm{top}} 0 \otimes \Lambda^{\mathrm{top}} \ker D_u[TZ; 0, TL] \simeq \Lambda^{\mathrm{top}} \ker D_u[; 0],$$

and since the vector space 0 is even-dimensional the order in which we do the two gluings is irrelevant. Using the orientation on $TZ^{\perp} \cong TX/TZ$ given by (3.3) (as Z is even-dimensional it doesn't matter whether we use the left-to-right or right-to left conventions) we obtain an induced orientation on $\mathrm{d}(D_u[; 0])$. Explicitly, by Lemma B.1.1 this orientation comes from

$$0 \to \ker D_u[TZ; 0, TL] \to \ker D_u[; 0] \to TX/TZ \to 0. \tag{B.3}$$

Remark B.1.2. Strictly the gluings are not supposed to take place at $u(-1)$ and $u(0)$: we should deform u along Morse flowlines before gluing, and can then deform back. However, the composite operation 'deform, glue, deform back' can be homotoped just to 'glue' in an obvious way. //

An orientation of $\mathrm{d}(D_u[; 0])$ is by definition a generator for the free rank 1 abelian group which Zapolsky calls $C(p, [u])$, where $p = u(1)$ is the Morse minimum and $[u] \in \pi_2(X, L, p)$ is the homotopy class of u. The sign we require is obtained by identifying this space with $\mathbb{Z}$, so that $1 \in \mathbb{Z}$ corresponds to the orientation of $\mathrm{d}(D_u[; 0])$ induced by the relative spin structure. This orientation is defined by Zapolsky, following Seidel, in [146, Section 7.1], and coincides with the definition by Fukaya–Oh–Ohta–Ono described in Section B.3 (both amount to concentrating the twisting of $(u^*TX, u|_{\partial D}^* TL)$ in a sphere bubble in the interior of the disc, trivialising the untwisted bundle, and homotoping the boundary bundle to be constant). This orientation on $\mathrm{d}(D_u[; 0])$ is related to the orientation on $\mathrm{d}(D_u)$ by the exact sequence

$$0 \to \ker D_u[; 0] \to \ker D_u \to T_{u(1)}L \to 0. \tag{B.4}$$

Combining (B.3) and (B.4) we see precisely that the sign is $+1$ if and only if the isomorphism

$$T_u\mathcal{M} = \ker D_u \to (T_{u(0)}X/T_{u(0)}Z) \oplus T_pL,$$

given by the infinitesimal evaluation maps, is orientation-preserving. This verifies Assumption 3.2.6 for Zapolsky's orientation scheme.

B.2 Verifying Assumption 1.2.4

In this subsection we make Assumption 3.2.6 and show that it implies Assumption 1.2.4 in the setting we need it, in the proof of Proposition 3.2.24. First note that the orientation sign of the isomorphism in Assumption 3.2.6 is well-defined for any regular pseudoholomorphic disc u satisfying $u(0) \in Z$ and $u(1) = p$ for which the derivative of the evaluation map at 0 and 1 is transverse $T_{u(0)}Z \oplus 0$ in $T_{u(0)}X \oplus T_{u(1)}L$. This is precisely what it means for the corresponding trajectory contributing to $\mathcal{CO}^0$ to be transversely cut out, so the first part of Assumption 1.2.4 is proved.

To see the second part of Assumption 1.2.4 consider the path (J_t, H_t) used in the proof of Proposition 3.2.24, and the cobordism $\mathcal{M}_{\text{cob}}$. Consider the trivial fibration E over $[0,1]$ with fibre given by the space of maps of suitable Sobolev regularity from $(D, \partial D)$ to (X, L). At each point (u,t) in E, where u is a disc and $t \in [0,1]$, we have a Cauchy–Riemann operator D_u^E which maps from the tangent space $T_{(u,t)}E$ to an appropriate Sobolev space W_u of u^*TX-valued $(0,1)$-forms on D. We write $T^{\text{vert}}E$ for the vertical tangent space inside TE and D_u for the restriction of D_u^E to this subspace.

The space $\mathcal{M}_{\text{cob}}$ is cut out inside E by the (J_t, H_t)-holomorphic curve equation and evaluation maps at 0 and 1 (which must hit Z and p respectively). The regularity hypothesis on the cobordism says that it is cut out transversely. Concretely this means that for each (u,t) in $\mathcal{M}_{\text{cob}}$ the operator $D_u^E \colon T_{(u,t)}E \to W_u$ is surjective and the infinitesimal evaluation map

$$D\,\mathrm{ev}_0 \oplus D\,\mathrm{ev}_1 \colon \ker D_u^E \to (T_{u(0)}X/T_{u(0)}Z) \oplus T_{u(1)}L \tag{B.5}$$

is surjective. Note that the restriction D_u to the vertical subspace is surjective precisely if u is regular as a (J_t, H_t)-holomorphic disc for fixed t.

Letting V_u denote the right-hand side of (B.5) and ev denote $D\,\mathrm{ev}_0 \oplus D\,\mathrm{ev}_1$ the regularity condition amounts to surjectivity of the Fredholm operator $(D_u^E, \mathrm{ev}) \colon T_{(u,t)}E \to W_u \oplus V_u$. The tangent space $T_{(u,t)}\mathcal{M}_{\text{cob}}$ is exactly the kernel of this operator, so to orient $\mathcal{M}_{\text{cob}}$ it suffices to orient its determinant line. In order to do this, consider the short exact sequences of Fredholm operators:

$$\begin{tikzcd}
0 \arrow[r] & T^{\text{vert}}_{(u,t)}E \arrow[r] \arrow[d, "{(D_u,\mathrm{ev})}"] & T_{(u,t)}E \arrow[r] \arrow[d, "{(D_u^E,\mathrm{ev})}"] & \mathbb{R} \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & W_u \oplus V_u \arrow[r] & W_u \oplus V_u \arrow[r] & 0 \arrow[r] & 0
\end{tikzcd}$$

and

$$\begin{tikzcd}
0 \arrow[r] & T^{\text{vert}}_{(u,t)}E \arrow[r, "i_1"] \arrow[d, "{(D_u,\mathrm{ev})}"] & T^{\text{vert}}_{(u,t)}E \oplus (W_u \oplus V_u) \arrow[r, "j_1"] \arrow[d, "D_u \oplus \mathrm{id}"] & W_u \oplus V_u \arrow[r] \arrow[d, "\mathrm{pr}_1"] & 0 \\
0 \arrow[r] & W_u \oplus V_u \arrow[r, "i_2"] & W_u \oplus (W_u \oplus V_u) \arrow[r, "j_2"] & W_u \arrow[r] & 0
\end{tikzcd}$$

where $i_1 = (\mathrm{id}, (D_u, \mathrm{ev}))$, $i_2 = (\mathrm{pr}_1, \mathrm{id})$, $j_1(x,y) = (D_u, \mathrm{ev})(x) - y$ and $j_2(x,y) = x - \mathrm{pr}_1(y)$ (this second diagram is based on [147, (2.19)]). These give rise to continuous isomorphisms of determinant line bundles

$$\mathrm{d}((D_u^E, \mathrm{ev})) \cong \mathrm{d}((D_u, \mathrm{ev})) \otimes \mathbb{R} \tag{B.6}$$

and

$$\mathrm{d}(D_u) = \mathrm{d}(D_u \oplus \mathrm{id}) \cong \mathrm{d}((D_u, \mathrm{ev})) \otimes \Lambda^{\text{top}} V_u \tag{B.7}$$

over E; for further details see [147]. Moreover, by [147, Normalization II], when (D_u, ev) is surjective (and thus an isomorphism) the orientation on $\mathrm{d}((D_u, \mathrm{ev}))$ induced by (B.7) is exactly the orientation sign of the map $\mathrm{ev} \colon \ker D_u \to W_u$.

The spaces V_u and $\mathbb{R}$ carry natural orientations—the former from (3.3)—whilst $\mathrm{d}(D_u)$ has a (continuous) orientation from the choice of orientation and relative spin structure on L. Combining these with (B.6) and (B.7) we obtain a continuous orientation of $\mathrm{d}((D_u^E, \mathrm{ev}))$ and hence an orientation of $\mathcal{M}_{\mathrm{cob}}$. We are left to show that the induced orientation on the boundary components $\mathcal{M}_{t=0}$ and $\mathcal{M}_{t=1}$ are equal and opposite respectively (or vice versa) to the signs coming from the restriction of (B.5) to the vertical subspaces at $t = 0$ and $t = 1$. This follows from (B.6) and the relationship between $\mathrm{d}((D_u, \mathrm{ev}))$ and the sign of $\mathrm{ev}\colon \ker D_u \to W_u$, which is precisely the isomorphism from Assumption 3.2.6.

B.3 Riemann–Hilbert pairs and splittings

We assume throughout the rest of this appendix that (X, L) is sharply K-homogeneous and monotone. Equip L with an orientation and the corresponding standard spin structure from Definition 3.1.10.

First recall that a Riemann–Hilbert pair (E, F) is a (finite rank) smooth complex vector bundle E over the disc D, holomorphic over the interior, along with a totally real subbundle F of $E|_{\partial D}$. By Oh's splitting theorem [110, Theorem I], there exist integers κ_j, unique up to reordering, such that

$$(E, F) \cong \bigoplus_j (\mathbb{C}, z^{\kappa_j/2}\mathbb{R}).$$

By definition, these are the *partial indices* of the pair (E, F).

Now suppose u is a disc in $\mathcal{M}_{\mu=2k}$. There is an induced Riemann–Hilbert pair (E, F) given by $(u^*TX, u|_{\partial D}^*TL)$, and $T_u\mathcal{M}_{\mu=2k}$ is the kernel of the Cauchy–Riemann operator

$$D_u = \overline{\partial}_{(E,F)}\colon \Gamma\big((D, \partial D), (E, F)\big) \to \Gamma(D, \Omega^{0,1}(E)),$$

where Γ denotes the space of smooth sections. A choice of orientation and relative spin structure on L induces a homotopy class of trivialisation of F, and using such a homotopy class Fukaya–Oh–Ohta–Ono [63, Chapter 8], building on work of de Silva [39], constructed an orientation on $\ker \overline{\partial}_{(E,F)}$ which is used to orient $\mathcal{M}_{\mu=2k}$.

Roughly speaking, their construction proceeds by degenerating the disc D into a nodal curve $D \cup \mathbb{CP}^1$ (joined at the point 0 in each component), itself carrying a bundle pair (E', F') and Cauchy–Riemann operator $\overline{\partial}_{(E',F')}$, such that (E, F) and (E', F') are identified over a collar neighbourhood of ∂D. Gluing results give a bijection between holomorphic sections of E' over each component of the nodal curve separately, which agree at the joining point, and holomorphic sections of the original bundle E. From this we obtain an isomorphism

$$\ker \overline{\partial}_{(E,F)} \cong \ker\Big(\ker \overline{\partial}_{(E'|_D,F')} \oplus \ker \overline{\partial}_{E'|_{\mathbb{CP}^1}} \to E'_0\Big), \tag{B.8}$$

where the map on the right-hand side sends $(s_D, s_{\mathbb{CP}^1})$ to $s_D(0) - s_{\mathbb{CP}^1}(0)$. Since u has all partial indices non-negative [45, Lemma 2.11, Lemma 3.2], this map is in fact surjective, so we have a short exact sequence

$$0 \to \ker \overline{\partial}_{(E,F)} \to \ker \overline{\partial}_{(E'|_D,F')} \oplus \ker \overline{\partial}_{E'|_{\mathbb{CP}^1}} \to E'_0 \to 0. \tag{B.9}$$

The degeneration is done in such a way that $(E'|_D, F')$ is trivialised by our choice of trivialisation of F, so we have an identification of $\ker \overline{\partial}_{(E'|_D,F')}$ with $\mathbb{R}^n$, by evaluating solutions at a boundary point (and reversing the orientation on L switches the orientation of this identification). The other two spaces appearing on the right-hand side of (B.8) carry complex structures and hence canonical orientations. Putting these three orientations together with the short exact sequence (B.9), we see that there is an induced orientation on $\ker \overline{\partial}_{(E,F)}$, which is the space we are really interested in. The auxiliary choices made do not affect this orientation.

The details of this procedure are technical and unimportant for our purposes, but we note two key properties of the construction. First, if the trivialisation of F extends to a holomorphic trivialisation of E then the bundle pair (E', F') is trivial and $\ker \overline{\partial}_{(E,F)}$ is oriented directly by its identification with a fibre of F by evaluation at a boundary point. And second, if (E, F) splits as a direct sum $(E^1, F^1) \oplus (E^2, F^2)$, compatible with the trivialisation of F (i.e. so that fibrewise, over each point z, F^1_z corresponds to the first k components of $F_z \cong \mathbb{R}^n$ and F^2_z to the last $n-k$, and we obtain canonical trivialisations of F^1 and F^2), then we can make the construction respect this splitting and check straightforwardly that the identification

$$\ker \overline{\partial}_{(E,F)} = \ker \overline{\partial}_{(E^1,F^1)} \oplus \ker \overline{\partial}_{(E^2,F^2)}$$

is orientation-preserving.

We also need to make one simple explicit calculation:

Lemma B.3.1. *Suppose (E, F) is a rank 1 Riemann–Hilbert pair of index 2—which can be identified with the tangent bundle of the pair $(D, \partial D)$—with F oriented. Evaluation at 0 and 1 defines an isomorphism*

$$f\colon \ker \overline{\partial}_{(E,F)} \xrightarrow{\sim} E_0 \oplus F_1,$$

and the codomain is oriented by the complex structure on E_0 and the choice of orientation on F. This isomorphism is orientation-preserving.

Proof. Let $(E^{(t)}, F^{(t)}) \to (D^{(t)}, \partial D^{(t)})$ be the family of bundle pairs, parametrised by $t \in [0, 1]$, realising the above degeneration of the domain from $D^{(0)} = D$ to $D^{(1)} = D \cup \mathbb{CP}^1$. Note that for all $t < 1$ the domain is biholomorphic to the disc D, and we may choose a continuous family $p^{(t)}$ of interior marked points in the domain which converge to the point ∞ in $\mathbb{CP}^1 \subset D^{(1)}$ as $t \to 1$, as shown in Fig. B.1 (recall that 0 in $\mathbb{CP}^1$ is the point which is glued to 0 in D to construct $D^{(1)}$). Similarly we can choose a continuous family $q^{(t)}$ of boundary marked points.

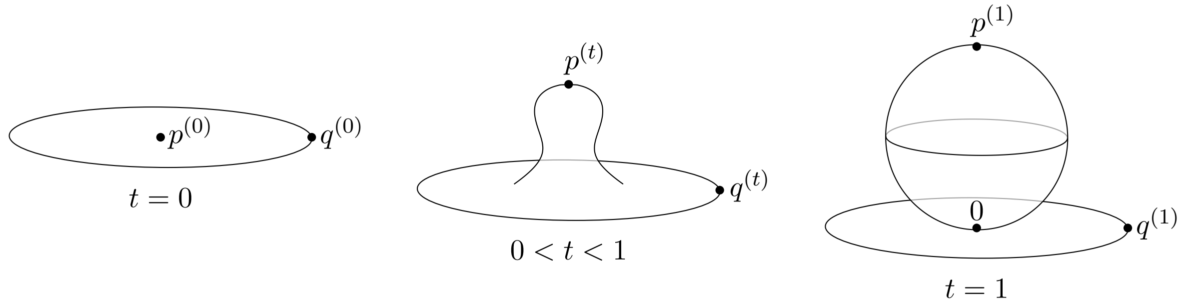

Figure B.1: The degeneration $D^{(t)}$.

We obtain a continuous family of kernels of the corresponding Cauchy–Riemann operators, which we denote by $\ker \overline{\partial}_t$. At $t = 1$ we impose the condition that the sections over D and $\mathbb{CP}^1$ agree at 0, and continuity is in the sense of the above gluing result. This comes equipped with a family of evaluation maps, giving a continuous family of isomorphisms

$$f^{(t)} \colon \ker \overline{\partial}_t \xrightarrow{\sim} E^{(t)}_{p^{(t)}} \oplus F^{(t)}_{q^{(t)}}.$$

The codomains are naturally oriented, and the definition of the orientation on $\ker \overline{\partial}_0$ is such that the orientation sign $\varepsilon(f^{(t)})$ is constant. Taking $p^{(0)} = 0$ and $q^{(0)} = 1$ the statement of the lemma can be expressed as $\varepsilon(f^{(0)}) = +1$, so it suffices to show that $\varepsilon(f^{(1)}) = +1$. To simplify notation we shall assume $q^{(1)} = 1$ and replace $^{(1)}$ by $'$.

Well, the restriction of the bundle E' to the $\mathbb{CP}^1$ component of $D^{(1)}$ has first Chern class given by half the index of our original Riemann–Hilbert pair, so is isomorphic to $\mathcal{O}(1)$. We may therefore think of holomorphic sections as affine linear functions on $\mathbb{C}$, say $z \mapsto a + bz$ for complex numbers a and b. Solutions to the Riemann–Hilbert problem on the D component, meanwhile, are all of the form cs, where c is a real number and s is an arbitrary fixed solution with $s(1)$ pointing in the positive direction in F'_1. The matching condition at 0 forces $a = \lambda c$ for some fixed $\lambda \in \mathbb{C}^*$, so if $b = b_1 + ib_2$ then from (B.9) we see that b_1, b_2, c form a positively oriented set of coordinates on $\ker \overline{\partial}'$. Choosing an appropriate basis vector for E'_∞, the map

$$f' \colon \ker \overline{\partial}' \to E'_\infty \oplus F'_1$$

can be viewed as sending the point (b_1, b_2, c) to $(b_1 + ib_2, c)$. It is therefore orientation-preserving, which is exactly what we needed to show, completing the proof. $\square$

Now assume that Z is a K-invariant subvariety of X of complex codimension k, containing $u(0)$ so that $u \in \mathcal{M}^Z_{\mu=2k}$; by Proposition 3.2.11 we deduce that u is axial. The key ingredient for our later computations is a particularly simple trivialisation of F, in the homotopy class induced by the standard spin structure (this will be essentially tautological, since the spin structure is defined by a parallelisation of L), and the observation that it is compatible with a splitting of (E, F) into rank 1 pairs of index 0 or 2 and rank 2 pairs with partial indices $(1, 1)$. Note that in order for the boundary real subbundle to have a trivialisation it must be orientable, which is equivalent to having even total index, so there is no hope of decomposing the problem further into rank 1 pairs of index 1.

In order to construct the trivialisation of F, consider the infinitesimal action of $\mathfrak{g} = \mathfrak{k} \otimes \mathbb{C}$ at $u(0)$ in Z, and let its kernel be V. Since Z is G-invariant and of complex codimension k we have $\dim_\mathbb{C} V \geq k$. Pick an $\mathbb{R}$-basis $\xi_1, \dots, \xi_a$ for $V \cap \mathfrak{k}$, and extend to a $\mathbb{C}$-basis $\xi_1, \dots, \xi_a, \eta_1, \dots, \eta_b$ for V. Each η_j can be written as $\alpha_j + i\beta_j$ for unique $\alpha_j, \beta_j \in \mathfrak{k}$, and one can verify that $\xi_1, \dots, \xi_a, \alpha_1, \dots, \alpha_b, \beta_1, \dots, \beta_b$ are linearly independent in $\mathfrak{k}$. We can therefore pick $\theta_1, \dots, \theta_c$ which extend this set to a basis of $\mathfrak{k}$. We obtain a trivialising frame for F

$$\xi_1 \cdot u, \dots, \xi_a \cdot u, \alpha_1 \cdot u, \beta_1 \cdot u, \dots, \alpha_b \cdot u, \beta_b \cdot u, \theta_1 \cdot u, \dots, \theta_c \cdot u,$$

where, for example, $\xi_1 \cdot u$ denotes the section $z \mapsto \xi_1 \cdot u(z) \in T_{u(z)}L$, and since the actual choice of orientation on L is irrelevant we may assume this frame is positively oriented. The frame is compatible with the global trivialisation of TL defining the standard spin structure, so its homotopy class is precisely that induced by this spin structure.

There is a related splitting of E given by

$$\left\langle \frac{\xi_1}{z} \cdot u \right\rangle \oplus \cdots \oplus \left\langle \frac{\xi_a}{z} \cdot u \right\rangle \oplus \left\langle \frac{(1+z)\alpha_1 + i(1-z)\beta_1}{z} \cdot u, \frac{(1+z)\beta_1 - i(1-z)\alpha_1}{z} \cdot u \right\rangle \oplus \cdots \oplus \\ \left\langle \frac{(1+z)\alpha_b + i(1-z)\beta_b}{z} \cdot u, \frac{(1+z)\beta_b - i(1-z)\alpha_b}{z} \cdot u \right\rangle \oplus \langle \theta_1 \cdot u \rangle \oplus \cdots \oplus \langle \theta_c \cdot u \rangle, \quad \text{(B.10)}$$

where $\langle \cdot \rangle$ here denotes $\mathbb{C}$-linear span and, for example, $(\xi_1/z) \cdot u$ denotes the holomorphic section of E given by $z \mapsto (\xi_1/z) \cdot u(z)$. Note that the $a + 2b + c = n$ sections of E listed in (B.10) are holomorphic (since the $\xi_j \cdot u$ and $(\alpha_j + i\beta_j) \cdot u$ vanish at 0), and are fibrewise $\mathbb{C}$-linearly independent. To see the latter, consider the wedge product (over $\mathbb{C}$) of the sections: it is

$$\frac{1}{z^{a+b}}(\xi_1 \cdot u) \wedge \cdots \wedge (\xi_a \cdot u) \wedge (\alpha_1 \cdot u) \wedge (\beta_1 \cdot u) \wedge \cdots \wedge (\alpha_b \cdot u) \wedge (\beta_b \cdot u) \wedge (\theta_1 \cdot u) \wedge \cdots \wedge (\theta_c \cdot u),$$

which is clearly non-zero on $D \setminus \{0\}$. Since u has index $2k$, this expression vanishes to order $k - (a + b)$ at 0, but we know that $a + b = \dim_{\mathbb{C}} V \geq k$. Therefore we have equality $a + b = k$ (so $T_{u(0)}Z = \mathfrak{g} \cdot u(0)$) and the sections remain independent at 0.

It is easy to see that F meets each summand E^j of (B.10) (where $j = 1, 2, \ldots, a + b + c$) in a subbundle F^j of the appropriate rank, namely half the real rank of E^j, and moreover one can write down explicit frames for each F^j in terms of the given trivialisations of the E^j:

$$F^j = z\left\langle \frac{\xi_j}{z} \cdot u \right\rangle_{\mathbb{R}} \text{ for } 1 \leq j \leq a$$
$$F^{j+a} = z^{1/2}\left\langle \frac{(1+z)\alpha_j + i(1-z)\beta_j}{z} \cdot u, \frac{(1+z)\beta_j - i(1-z)\alpha_j}{z} \cdot u \right\rangle_{\mathbb{R}} \text{ for } 1 \leq j \leq b$$
$$F^{j+a+b} = \langle \theta_j \cdot u \rangle_{\mathbb{R}} \text{ for } 1 \leq j \leq c,$$

where $\langle \cdot \rangle_{\mathbb{R}}$ denotes $\mathbb{R}$-linear span. In particular, the splitting of E is in fact a splitting of the pair (E, F) compatible with our trivialisation of F, and the first a summands are rank 1, index 2; the next b are rank 2, partial indices $(1, 1)$; the final c are rank 1, index 0. We thus have the desired decomposition of (E, F).

B.4 Signs for $\mathfrak{m}_0$

The quantity $\mathfrak{m}_0(L)$ counts index 2 discs whose boundaries pass through a generic point of L. More precisely, if L is equipped with a brane structure—i.e. a relative spin structure and local system, which may be composite (as defined in Section 2.2.3) or of higher rank—it is the degree of the evaluation map to L from the space of unparametrised index 2 discs with a single boundary marked point [66, Equation (10)], weighted by the local system. Here we follow the conventions of Fukaya–Oh–Ohta–Ono, although note that by Example 3.2.29 these are compatible with Assumption 3.2.6 in the sense that the resulting signs are consistent with the Auroux–Kontsevich–Seidel criterion.

In our sharply homogeneous setting, the signs are as simple as one could hope:

Proposition B.4.1. *If $L^\flat$ is a sharply K-homogeneous monotone Lagrangian brane equipped with the standard spin structure then all index 2 discs on L count towards $\mathfrak{m}_0(L^\flat)$ with positive sign (before twisting by the local system).*

Remark B.4.2. In the toric case this result is a well-known piece of folklore, but to our knowledge a proof has not appeared in the literature. //

Proof. Suppose u is a holomorphic disc on L of index 2. Since it is axial, after reparametrisation we may assume that it is given by $z \mapsto e^{-i\xi \log z}u(1)$ for some ξ in $\mathfrak{k}$. Taking Z to be the whole compactification divisor, we find ourselves in the setting of the previous subsection, except we know that ξ acts trivially at $u(0)$ so the quantity a is equal to 1, b is 0, and we can take $\xi_1 = \xi$. After extending to a basis $\xi, \theta_1, \dots, \theta_{n-1}$ for $\mathfrak{k}$, we may assume that the orientation on L is chosen so that $\xi \cdot u(1), \theta_1 \cdot u(1), \dots, \theta_{n-1} \cdot u(1)$ is positively oriented. By Lemma B.3.1, the following basis of $T_u\mathcal{M}_{\mu=2}$ is then positively oriented:

$$\frac{1+z^2}{z}\xi \cdot u, \frac{i(1-z^2)}{z}\xi \cdot u, \xi \cdot u, \theta_1 \cdot u, \dots, \theta_{n-1} \cdot u. \tag{B.11}$$

We now need to orient the quotient $(\mathcal{M}_{\mu=2} \times \partial D)/\mathrm{PSL}(2,\mathbb{R})$ according to [63, Section 8.3], where $\mathrm{PSL}(2,\mathbb{R})$ acts on the right by reparametrisation, i.e. $\varphi \cdot (u, z) = (u \circ \varphi, \varphi^{-1}(z))$. One can check that a positively oriented basis for infinitesimal reparametrisations of the disc is given by

$$-i(1+z^2), (1-z^2), -iz.$$

Here the orientation is defined by the action on the right on three points z_1, z_2, z_3 in anticlockwise order around ∂D (in other words, a vector field X on the disc maps to $(-X(z_1), -X(z_2), -X(z_3))$ in $T_{z_1}\partial D \oplus T_{z_2}\partial D \oplus T_{z_3}\partial D$). The corresponding tangent vectors to $\mathcal{M}_{\mu=2}$ at u are precisely *minus* the first three vectors of (B.11).

We can identify the tangent space to the quotient at $[(u,1)]$ with the subspace of $T_u\mathcal{M}_{\mu=2} \oplus T_1\partial D$ spanned by

$$(\theta_1 \cdot u, 0), \dots, (\theta_{n-1} \cdot u, 0), (0, i). \tag{B.12}$$

By the ordering convention

$$T(\mathcal{M}_{\mu=2} \times \partial D) = T((\mathcal{M}_{\mu=2} \times \partial D)/\mathrm{PSL}(2,\mathbb{R})) \oplus \mathfrak{psl}(2,\mathbb{R})$$

from [63, Equation (8.2.1.2)] we see that the basis in (B.12) carries orientation sign $(-1)^{n-1}$. Moving the $(i, 0)$ to the start of (B.12) eliminates this sign, so $(i,0), (\theta_1 \cdot u, 0), \dots, (\theta_{n-1} \cdot u, 0)$ represents a positively oriented basis for the quotient. This basis evaluates to $\xi \cdot u(1), \theta_1 \cdot u(1), \dots, \theta_{n-1} \cdot u(1)$, which is by definition a positively oriented basis for $T_{u(1)}L$, so we conclude that the relevant evaluation map is indeed orientation-preserving. ☐

B.5 Signs for $\mathcal{CO}^0$

Finally we return to the setting of the end of Section B.3, with an analytic subvariety $Z \subset X$ of complex codimension k and a holomorphic disc u in $\mathcal{M}^Z_{\mu=2k}$. Our aim is to show that the isomorphism in Assumption 3.2.6 is orientation-preserving.

Recall the splitting (B.10) of the Riemann–Hilbert pair corresponding to u, and that

$$T_{u(0)}Z = \mathfrak{g} \cdot u(0) = \langle (\alpha_1 + i\beta_1) \cdot u(0), \dots, (\alpha_b + i\beta_b) \cdot u(0), \theta_1 \cdot u(0), \dots, \theta_c \cdot u(0) \rangle.$$

We thus have

$$T_{u(0)}Z = \bigoplus_j \Big(E_0^j \cap T_{u(0)}Z \Big),$$

where $E_0^j \cap T_{u(0)}Z$ is a complex vector space of dimension 0, 1 or 1 according to whether the partial indices μ^j of (E^j, F^j) are 2, $(1,1)$ or 0 respectively. Hence the required orientation sign is

$$\varepsilon \left(\bigoplus_j \ker \overline{\partial}_{(E^j,F^j)} \to \left(\bigoplus_j E_0^j \Big/ \Big(E_0^j \cap T_{u(0)}Z \Big) \right) \oplus \left(\bigoplus_j F_1^j \right) \right)$$
$$= \prod_j \varepsilon \Big(\ker \overline{\partial}_{(E^j,F^j)} \to \Big(E_0^j \Big/ \Big(E_0^j \cap T_{u(0)}Z \Big) \Big) \oplus F_1^j \Big).$$

It is left to understand each factor on the right-hand side—let $\varepsilon^{(j)}$ denote the jth, and let $\overline{\partial}_j$ be the corresponding Cauchy–Riemann operator.

Lemma B.5.1. *For all j we have $\varepsilon^{(j)} = +1$*

Proof. If $\mu^j = 2$ then the intersection $E_0^j \cap T_{u(0)}Z$ is zero and the result follows from Lemma B.3.1. On the other hand, if $\mu^j = 0$ then $(E^j, F^j) = (\mathbb{C}\theta \cdot u, \mathbb{R}\theta \cdot u)$, where $\theta = \theta_{j-(a+b)} \in \mathfrak{k}$, and the space $E_0^j / (E_0^j \cap T_{u(0)}Z)$ is zero. A positively oriented basis for $\ker \overline{\partial}_j$ is given by $\theta \cdot u$, whilst a positively oriented basis for F_1^j is given by $\theta \cdot u(1)$, so the required map is indeed orientation-preserving.

Finally suppose $\mu^j = (1,1)$, so E^j and F^j are spanned by appropriate expressions built from $\alpha = \alpha_{j-a}$ and $\beta = \beta_{j-a}$, with $(\alpha + i\beta) \cdot u(0) = 0$. To find a positively oriented basis for $\ker \overline{\partial}_j$ we degenerate our $(1,1)$ pair to a $(2,0)$ as follows. Suppose we replace the condition $(\alpha + i\beta) \cdot u(0) = 0$ by $(\alpha + ti\beta) \cdot u(0) = 0$, and deform the real parameter t from 1 down to 0. We obtain a family of Riemann–Hilbert pairs over $[0,1]$, and bases for the kernels of the corresponding Cauchy–Riemann operators are given by

$$\frac{(1+z^2)\alpha + ti(1-z^2)\beta)}{z} \cdot u, \frac{i(1-z^2)\alpha - t(1+z^2)\beta)}{z} \cdot u, \alpha \cdot u, \beta \cdot u. \tag{B.13}$$

By Lemma B.3.1, and the fact that $\alpha \cdot u(1)$, $\beta \cdot u(1)$ is positively oriented as a basis of F_1^j, this basis is positively oriented in the limit $t \to 0$. Hence by continuity of the orientation construction it is also positively oriented at $t = 1$.

Now (restricting again to the $t = 1$ situation, which is what we actually care about) consider the evaluation map to

$$\left(E_0^j \Big/ \left(E_0^j \cap T_{u(0)} Z\right)\right) \oplus F_1^j = \left(\left\langle \lim_{z \to 0} \frac{\alpha + i\beta}{z} \cdot u(z), \beta \cdot u(0) \right\rangle \Big/ \langle \beta \cdot u(0) \rangle\right) \oplus \langle \alpha \cdot u(1), \beta \cdot u(1) \rangle_{\mathbb{R}}.$$

It is easy to check that the basis (B.13) is sent to a positively oriented basis for this space, and so $\varepsilon^{(j)} = +1$ as claimed. □

This proves:

Corollary B.5.2. *Under Assumption 3.2.6 all discs contributing to Proposition 3.2.24 count with positive sign for the standard spin structure.*

Appendix C

Transversality for the pearl complex

C.1 Preliminaries

In this appendix we discuss the transversality results required to set up the pearl complex. This is mainly to show how to work with a *fixed* complex structure, but also includes a brief review of Biran–Cornea's foundational work in [14] based on *generic* almost complex structures, in order to compare and contrast the two approaches and show that they give the same (co)homology.

We begin by recalling some basic notions in differential topology. A standard reference is [86, Chapter 2], on which we base our terminology. For topological spaces P and Q let $C^0(P, Q)$ denote the space of continuous maps from P to Q, equipped with the compact-open topology. If P and Q are actually smooth manifolds then for each $r \in \{1, 2, \dots, \infty\}$ we can form the r-jet space $J^r(P, Q)$, and the space $C^r(P, Q)$ of r-times continuously differentiable maps from P to Q (or smooth maps in the case $r = \infty$) embeds in $C^0(P, J^r(P, Q))$ via the prolongation map j^r. This gives a natural topology on $C^r(P, Q)$, which we call the (weak) C^r-topology, which is induced by a (non-canonical) complete metric d_r. When P is compact this topology coincides with the strong C^r-topology, sometimes called the Whitney topology (although the reader is warned that terminology varies between different authors), and the set of (smooth) diffeomorphisms $\mathrm{Diff}(P)$ is open in $C^\infty(P, P)$.

Given manifolds P and Q, a submanifold $R \subset Q$, subsets $A \subset P$, $B \subset R$, and a smooth map $f \colon P \to Q$, let $f\ {}_A{\pitchfork}_B\ R$ denote that f is transverse to R along $A \cap f^{-1}(B)$. In other words, for all points p in A such that $f(p)$ is in B we have

$$\mathrm{Im}\, D_p f + T_{f(p)} R = T_{f(p)} Q.$$

Note this differs from the notation used in [86], where $f \pitchfork_K R$ denotes what we are calling $f\ {}_K{\pitchfork}_R\ R$.

Fix now a closed n-manifold L. Later this will be our Lagrangian, but for our present purposes this is irrelevant. For a positive integer s let $\Delta_{L,s}$ denote the big diagonal

$$\{(p_1, \dots, p_s) \in L^s : p_j = p_k \text{ for some } j \neq k\} \subset L^s.$$

The key result in differential topology we shall use is the following:

Lemma C.1.1. *For any countable collections of manifolds (M_j), positive integers (s_j), submanifolds $(N_j \subset L^{s_j} \setminus \Delta_{L,s_j})$, and smooth maps $(f_j \colon M_j \to L^{s_j})$, there exists a diffeomorphism φ of L, arbitrarily C^∞-close to id_L, such that for all j the map f_j is transverse to $\varphi^{\times s_j}(N_j)$.*

This will follow from:

Lemma C.1.2. *In the setup of Lemma C.1.1, for all j and all $p \in M_j$ and $q \in N_j$ there exist neighbourhoods $U_{j,p,q}$ of p in M_j and $V_{j,p,q}$ of q in N_j, such that the set*

$$W_{j,p,q} \coloneqq \{\varphi \in \mathrm{Diff}(L) : (\varphi^{-1})^{\times s_j} \circ f \;{}_{U_{j,p,q}}\!\pitchfork_{V_{j,p,q}} N_j\}$$

is open and dense in $\mathrm{Diff}(L)$ in the C^∞-topology.

To deduce Lemma C.1.1 from Lemma C.1.2, for each j we simply take a countable subcover $\{U_{j,p_k,q_k} \times V_{j,p_k,q_k}\}_k$ of the cover $\{U_{j,p,q} \times V_{j,p,q}\}_{(p,q)}$ of $M_j \times N_j$ and then consider the intersection

$$\bigcap_{j,k} W_{j,p_k,q_k} \subset \mathrm{Diff}(L).$$

Since $\mathrm{Diff}(L)$ is an open subset of the complete metric space $C^\infty(L, L)$, this intersection is dense in $\mathrm{Diff}(L)$ by the Baire category theorem, so in particular it contains elements arbitrarily C^∞-close to id_L. Such elements provide the φ of Lemma C.1.1.

Proof of Lemma C.1.2. Fix arbitrary j, p and q, and metrics on L and M_j. From now on we shall drop all j's from the notation, and just refer to M_j, s_j, N_j and f_j as M, s, N and f respectively. Let $\pi_1, \dots, \pi_s \colon L^s \to L$ denote the projections onto the factors, and choose vectors $v_1, \dots, v_a \in T_qL^s$ which form a basis for a complement to T_qN. For each i use cutoff functions to construct a smooth vector field V_i on L whose value at $\pi_k(q)$ coincides with the projection $D_q\pi_k(v_i)$ for all k.

Now consider the map

$$\begin{aligned} \psi \colon \mathbb{R}^a &\to C^\infty(L, L) \\ \mathbf{t} &\mapsto \exp\left(\sum_i t_i V_i\right), \end{aligned}$$

which sends a vector $\mathbf{t}$ to the time 1 flow of $\sum t_i V_i$. It is easy to check that given three topological spaces P, Q and R, and a continuous map $h \colon P \times Q \to R$, the map $h_{\mathrm{ev}} \colon P \to C^0(Q, R)$ given by $x \mapsto h(x, \cdot)$ is continuous. Since the map

$$\begin{aligned} \mathbb{R}^a \times L &\to L \\ (\mathbf{t}, x) &\mapsto \psi(\mathbf{t})(x) \end{aligned}$$

is smooth, and hence defines a continuous map $j^\infty_L\psi \colon \mathbb{R}^a \times L \to J^\infty(L, L)$ by prolongation along the L factor, we deduce that $\psi = (j^\infty_L\psi)_{\mathrm{ev}}$ is continuous.

By construction of the v_i and V_i, the map

$$\Psi \colon \mathbb{R}^a \times N \to L^s$$

$$(\mathbf{t}, x) \mapsto \psi(\mathbf{t})^{\times s}(x)$$

is a submersion at the point $(0, q)$, and along $\{0\} \times N$ it is simply the inclusion of N. There therefore exist an open ball B in N about q and a positive ε such that Ψ gives a diffeomorphism from $B_0^a(\varepsilon) \times B$ onto an open tubular neighbourhood T of B in L^s, where $B_0^a(\varepsilon)$ is the open ball of radius ε about 0 in $\mathbb{R}^a$. Let $\pi\colon T \to B_0^a(\varepsilon)$ be the composition of the inverse diffeomorphism Ψ^{-1} with projection onto the first factor. Now pick a smooth cutoff function ρ on L^s which has compact support contained in T and takes the value 1 on a compact neighbourhood T' of q in L^s. Let V be a compact neighbourhood of q in N such that T' contains a neighbourhood of V in L^s. Figure C.1 shows this setup.

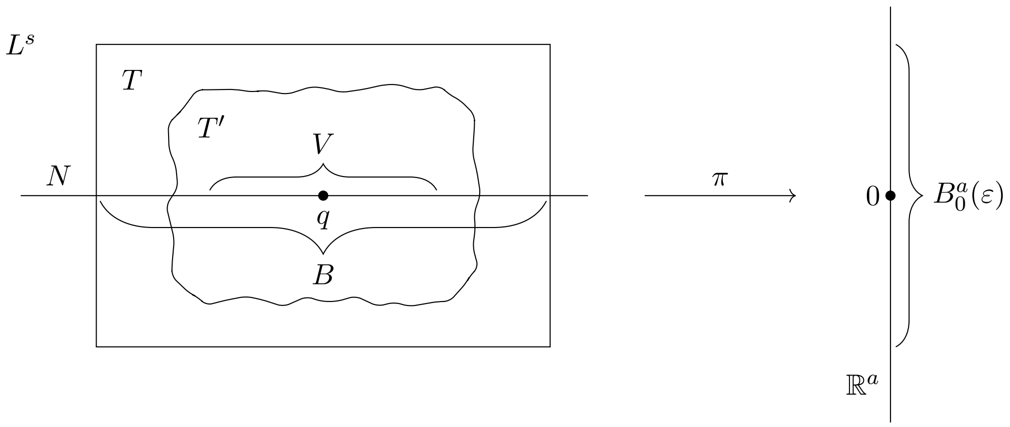

Figure C.1: The tubular neighbourhood T and projection π.

Let U be an arbitrary compact neighbourhood of p in M. We claim that $U_{j,p,q} = U$ and $V_{j,p,q} = V$ have the desired properties, so let

$$W = \{\varphi \in \mathrm{Diff}(L) : (\varphi^{-1})^{\times s} \circ f \;{}_U{\pitchfork}_V\; N\}.$$

We need to show that this set is open and dense in $\mathrm{Diff}(L)$.

First we prove it is open, so take a diffeomorphism $\varphi \in W$ and let $F = (\varphi^{-1})^{\times s} \circ f$. We wish to show that if $\widehat{\varphi}$ is sufficiently C^∞-close to φ then $\widehat{\varphi}$ is also in W. In fact we will prove the stronger statement that if a smooth map $G\colon M \to L^s$ is sufficiently C^1-close to F then $G\,{}_U{\pitchfork}_V\, N$. Let $U' = U \cap F^{-1}(T')$ be the preimage of T' in U. Note that U' is a closed subset of the compact set U, so is itself compact. Since $L^s \setminus T'$ is bounded away from V, if G is sufficiently C^0-close to F then $G(U \setminus U')$—which is contained in $G(F^{-1}(L^s \setminus T'))$—is disjoint from V. To show that $G\,{}_U{\pitchfork}_V\, N$ in this case it therefore suffices to check that $G\,{}_{U'}{\pitchfork}_V\, N$.

Given such a map G (C^0-close to F) and a point x in U', consider the derivative $D_x(\pi \circ G)\colon T_xM \to \mathbb{R}^a$ of $\pi \circ G$ at x. This sends the unit ball in T_xM (with respect to our metric on M) to a subset S_x of $\mathbb{R}^a$ containing 0. Let

$$r_G(x) = \sup\{r \in \mathbb{R}_{\geq 0} : B_0^a(r) \subset S_x\}$$

be the supremum of the radii of the balls about 0 in $\mathbb{R}^a$ which are contained in this subset. Note that the map $r_G\colon U' \to \mathbb{R}_{\geq 0}$ is continuous, and satisfies $r_G(x) > 0$ if and only if $D_x(\pi \circ G)$ is surjective. Now consider the map

$$\begin{aligned} R_G\colon U' &\to \mathbb{R}^2_{\geq 0} \\ x &\mapsto (r_G(x), d(G(x), V)), \end{aligned}$$

where $d(G(x), V)$ denotes the distance (with respect to the metric on L^s coming from our metric on L) between the point $G(x)$ and the set V. This map is continuous and vanishes precisely at points of U' where $G\ {}_{U'}\pitchfork_V N$ fails. Crucially R_G is also continuous in G in the C^1-topology. Since $F\ {}_{U'}\pitchfork_V N$ by hypothesis, R_F is nowhere zero and thus by compactness of U' its image is bounded away from zero. Therefore the same is also true of R_G for G sufficiently C^1-close to F. In other words, such G satisfy $G\ {}_{U'}\pitchfork_V N$, proving our openness claim.

We now show that W is dense in $\mathrm{Diff}(L)$, so take any diffeomorphism φ of L and again let $F = (\varphi^{-1})^{\times s} \circ f$. We need to construct a diffeomorphism $\widehat{\varphi}$, arbitrarily C^∞-close to φ, which is contained in W. Equivalently, we need a $\widehat{\varphi}$ arbitrarily C^∞-close to id_L such that $F\ {}_U\pitchfork_{\widehat{\varphi}^{\times s}(V)} \widehat{\varphi}^{\times s}(N)$. Since the map $\psi\colon \mathbb{R}^a \to C^\infty(L, L)$ is continuous, it is enough to show that $F\ {}_U\pitchfork_{V_{\mathbf{t}}} N_{\mathbf{t}}$ for arbitrarily small choices of $\mathbf{t}$, where $N_{\mathbf{t}} = \psi(\mathbf{t})^{\times s}(N)$ and $V_{\mathbf{t}} = \psi(\mathbf{t})^{\times s}(V)$.

Define a smooth map $\widetilde{\pi}\colon L^s \to B^a_0(\varepsilon)$ by

$$\widetilde{\pi}(x) = \begin{cases} \rho(x)\pi(x) & \text{if } x \in T \\ 0 & \text{otherwise,} \end{cases}$$

where ρ is our cutoff function on L^s. This coincides with π on a neighbourhood of V, so for $\mathbf{t}$ sufficiently small and $x \in V_{\mathbf{t}}$ we have that $\widetilde{\pi}(x) = \mathbf{t}$ and $T_x N_{\mathbf{t}} = \ker D_x\widetilde{\pi}$ (by definition of Ψ, the map $\psi(\mathbf{t})^{\times s}$ on $B \subset N$ corresponds under Ψ^{-1} to translation by $\mathbf{t}$). In particular, for $\mathbf{t}$ small F is transverse to $N_{\mathbf{t}}$ along $V_{\mathbf{t}}$ if and only if $\mathbf{t}$ is a regular value of $\widetilde{\pi} \circ F$. By Sard's theorem, such regular values exist arbitrarily close to 0, completing the proof of density and thus of Lemma C.1.2. ☐

C.2 Constructing the complex

Suppose (X, ω) is a closed symplectic manifold and $L^\natural \subset X$ is a monotone Lagrangian pre-brane over a ring R. Let (f, g) be a Morse–Smale pair on L, and J' an ω-compatible almost complex structure on X. For a tuple $\mathbf{A} = (A_1, \dots, A_r) \in H_2(X, L; \mathbb{Z})^r$, with $r > 0$, let $z(\mathbf{A})$ be the number of j for which $A_j = 0$. Let W^a_x and W^d_y be respectively the ascending and descending manifolds of critical points x and y of f, let Φ_t denote the time t flow of ∇f, and let Q be the submanifold of $L \times L$ given by

$$Q = \{(p, q) \in L \times L : p \notin \mathrm{Crit}(f),\ q = \Phi_t(p) \text{ for some } t > 0\} \cong (L \setminus \mathrm{Crit}(f)) \times \mathbb{R}_{>0}.$$

For a class $A \in H_2(X, L; \mathbb{Z})$, recall that $\mathcal{M}_2(A)$ denotes the moduli space of J'-holomorphic discs $u\colon (D, \partial D) \to (X, L)$ representing A, modulo reparametrisations fixing ± 1. Throughout this appendix we will refer to the evaluation maps ev_1 and ev_{-1} on $\mathcal{M}_2(A)$ as ev_+ and ev_- respectively.

Now define, for x, y and $\mathbf{A}$ as above, the pearly trajectory (or string of pearls) moduli space

$$\mathcal{P}(x, y, \mathbf{A}) = \Big(\big(\mathrm{ev}_-(A_1) \times \mathrm{ev}_+(A_1) \times \mathrm{ev}_-(A_2) \times \dots \\ \times \mathrm{ev}_+(A_r)\big)^{-1}\Big(W^a_x \times Q^{r-1} \times W^d_y\Big)\Big)\Big/\mathbb{R}^{z(\mathbf{A})}, \quad \text{(C.1)}$$

where the $\mathbb{R}^{z(\mathbf{A})}$ acts by translation of constant discs (corresponding to classes $A_j = 0$) along flowlines. None of the flowlines is actually doubly infinite in time so strictly each constant disc can only be translated by a subinterval of $\mathbb{R}$, but we overlook this slight notational imprecision. Such configurations are illustrated in Fig. C.2; the arrows depict the flowlines of ∇f. Really

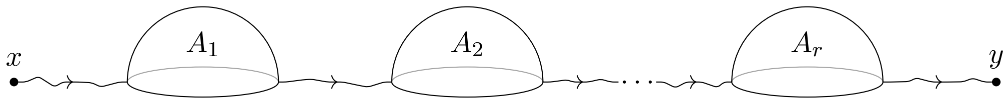

Figure C.2: A pearly trajectory, or string of pearls.

we may restrict our attention to *reduced* strings, for which $A_j \neq 0$ for all j unless $r = 1$, and transversality for these spaces automatically gives transversality for the same spaces with extra constant discs inserted (the case $r = 1$, $A_1 = 0$ gives rise to standard Morse trajectories). However, it is notationally convenient to allow any number of constant discs.

We shall also need moduli spaces of strings of pearls with *loose ends*, of the form

$$W^a_x(\mathbf{A}) \coloneqq \Big(\big(\mathrm{ev}_-(A_1) \times \mathrm{ev}_+(A_1) \times \mathrm{ev}_-(A_2) \times \cdots \times \mathrm{ev}_-(A_r)\big)^{-1}\big(W^a_x \times Q^{r-1}\big)\Big)\Big/\mathbb{R}^{z^a(\mathbf{A})} \quad \text{(C.2)}$$

$$W^d_y(\mathbf{A}) \coloneqq \Big(\big(\mathrm{ev}_+(A_1) \times \mathrm{ev}_-(A_2) \times \mathrm{ev}_+(A_2) \times \cdots \times \mathrm{ev}_+(A_r)\big)^{-1}\Big(Q^{r-1} \times W^d_y\Big)\Big)\Big/\mathbb{R}^{z^d(\mathbf{A})}, \quad \text{(C.3)}$$

where $z^a(\mathbf{A})$ is the number of $j \leq r - 1$ with $A_j = 0$ and $z^d(\mathbf{A})$ is the number of $j \geq 2$ with $A_j = 0$—we quotient out by translation of constant discs except those which are at the ends of trajectories. These configurations are illustrated in Fig. C.3. Note that these spaces carry

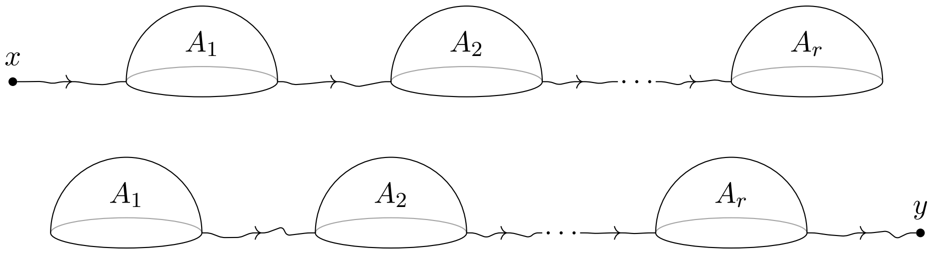

Figure C.3: Strings of pearls with loose ends.

evaluation maps $\mathrm{ev}(x, \mathbf{A})$ and $\mathrm{ev}(\mathbf{A}, y)$ at the loose marked point of the end disc. Again we could restrict to reduced trajectories, where all discs but the end one are non-constant, but we shall not do so for now.

In order to define the pearl complex for $L^{\natural}$ using the data (f, g, J'), we need the following:

(i) J' is regular, meaning that the moduli spaces $\mathcal{M}_2(A)$ are cut out transversely. This ensures that the $\mathcal{M}_2(A)$ are all smooth manifolds of the correct dimension.

(ii) The moduli spaces (C.1) of virtual dimension 0 are cut out transversely, so that the spaces used to define the differential d (as described below) are smooth manifolds of the correct dimension.

(iii) The same requirement as (ii) but in virtual dimension 1, so that we can construct the moduli spaces used to prove $d^2 = 0$ (again, see below). We'll unimaginatively call these '$d^2 = 0$ moduli spaces'.

(iv) The moduli spaces (C.1), with some of the Q factors (possibly none or all of them) replaced by copies of the diagonal $\Delta_L \in L \times L$, are cut out transversely whenever their virtual dimension is at most 0. This means that the moduli spaces in (ii) are compact, and that those in (iii) can be compactified by introducing strings of pearls which are degenerate in exactly one of the following ways: a single Morse flowline has broken, a single disc has bubbled into two (with one marked point in each component), or a single flowline has shrunk to zero.

(v) Given a broken string of pearls γ in virtual dimension 0 (in which a single flowline is broken, but which is otherwise non-degenerate), as illustrated in Fig. C.4, we need the loose end spaces

$$W_x^a(\mathbf{A}' = (A_1, \dots, A_k)) \text{ and } W_y^d(\mathbf{A}'' = (A_{k+1}, \dots, A_r))$$

to be transversely cut out, so that they are smooth manifolds of the correct dimension (note we may always assume that k is not 0 or r, by introducing extra classes equal to zero at the start and end of $\mathbf{A}$). Given this, we automatically have by (ii) and by (iii) that

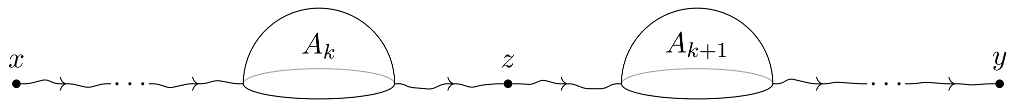

Figure C.4: A broken string of pearls.

$\mathrm{ev}(x, \mathbf{A}') \times \mathrm{ev}(\mathbf{A}'', y)$ is transverse to $W_z^d \times W_z^a$ and to Q respectively. Standard Morse-theoretic gluing arguments, as given in [8, Proposition 3.2.8] for example, then show that every such broken string γ occurs as a unique boundary point in the compactification of the $d^2 = 0$ moduli spaces. In fact, we only need the loose end spaces to be cut out transversely in neighbourhoods of the points appearing in γ, viewed as an element of

$$\big(\mathrm{ev}(x, \mathbf{A}') \times \mathrm{ev}(\mathbf{A}'', y)\big)^{-1}\Big(W_z^d \times W_z^a\Big) \subset W_x^a(\mathbf{A}') \times W_y^d(\mathbf{A}'').$$

(vi) Given a bubbled string of pearls γ in virtual dimension 0, with a single disc—the kth—bubbled into two (of classes A'_k and $A''_k = A_k - A'_k$) but otherwise non-degenerate, as illustrated in Fig. C.5, we need the loose end spaces

$$W_x^a(\mathbf{A}' = (A_1, \dots, A_{k-1}, 0)) \text{ and } W_y^d(\mathbf{A}'' = (0, A_{k+1}, \dots, A_r))$$

to be cut out transversely. By (iv), the maps

$$\mathrm{ev}(x, \mathbf{A}') \times (\text{inclusion of } \Delta_L \subset L \times L) \times \mathrm{ev}(\mathbf{A}'', y)$$

and

$$\mathrm{ev}_-(A'_k) \times \mathrm{ev}_+(A'_k) \times \mathrm{ev}_-(A''_k) \times \mathrm{ev}_+(A''_k)$$

are transverse, so by the gluing theorem for pseudoholomorphic discs [14, Theorem 4.1.2] each such bubbled string occurs as a unique boundary point in the compactification of the $\mathrm{d}^2 = 0$ moduli spaces. Again, we only need this transversality in neighbourhoods of the points of $W^a_x(\mathbf{A}')$ and $W^d_y(\mathbf{A}'')$ appearing in γ.

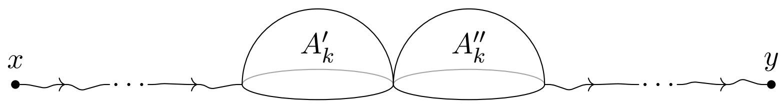

Figure C.5: A once-bubbled string of pearls.

If all of these conditions are satisfied then the differential can be defined by counting rigid reduced strings of pearls, which form compact zero-dimensional manifolds (compactness for fixed $\mathbf{A}$ follows from (iv), whilst the fact that only finitely many reduced choices of $\mathbf{A}$ give non-empty moduli spaces of virtual dimension 0 follows from Gromov compactness). The relation $\mathrm{d}^2 = 0$ is proved by considering one-dimensional moduli spaces of pearly trajectories, which can be compactified to compact one-manifolds by adding in boundary points as described in (iv). The first two types of boundary point appear exactly once each, by (v) and (vi), whilst the third type appear once by the transversality already provided by (iv). Explicitly, collapsing a flowline to zero corresponds to replacing a copy of Q in (C.1) by Δ_L, for which we have achieved transversality, and the end of the $\mathrm{d}^2 = 0$ moduli space which exhibits this collapsing can be seen by instead replacing Q by the manifold-with-boundary Q_0, defined by

$$Q_0 = \{(p, q) \in L \times L : p \notin \mathrm{Crit}(f),\ q = \Phi_t(p) \text{ for some } t \geq 0\} = (Q \cup \Delta_L) \setminus \Delta_{\mathrm{Crit}(f)}.$$

Before discussing how to achieve the necessary transversality, we first introduce some further terminology and notation. Recall that a pseudoholomorphic disc u is *simple* if its set of injective points

$$\{z \in D : u^{-1}(u(z)) = \{z\} \text{ and } u'(z) \neq 0\}$$

contains a dense open subset of D, and that a sequence $(u_1, \dots, u_r)$ of discs is *absolutely distinct* if for all j we have

$$u_j(D) \not\subset \bigcup_{k \neq j} u_k(D).$$

For an r-tuple $\mathbf{B}$ of homology classes we define $\mathcal{M}_2(\mathbf{B})$ to be the moduli space

$$\big(\mathrm{ev}_+(B_1) \times \mathrm{ev}_-(B_2) \times \mathrm{ev}_+(B_2) \times \cdots \times \mathrm{ev}_-(B_r)\big)^{-1}\big(\Delta_L^{r-1}\big) \tag{C.4}$$

of bubbled chains of discs, illustrated in Fig. C.6. These spaces carry evaluation maps $\mathrm{ev}_\pm(\mathbf{B})$

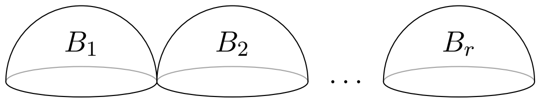

Figure C.6: A bubbled chain of discs.

at the two end marked points. We shall be interested in moduli spaces defined by (C.1) but with each A_j now an r_j-tuple $\mathbf{B}^j = (B^j_1, \dots, B^j_{r_j})$. We denote this modification by (C.1)', and refer to these configurations as *generalised strings of pearls* or *generalised pearly trajectories*. A generalised string is reduced if each disc is non-constant, or there is only one disc. Note that transversality for (C.1) with copies of Q replaced by Δ_L can be expressed in terms of transversality for the $\mathcal{M}_2(\mathbf{B})$ (i.e. for (C.4)) and for (C.1)'.

We'll say a generalised string of pearls γ is *diagonal-avoiding* if the evaluation maps in (C.1)' at γ miss the big diagonal $\Delta_{L,2r}$. A trajectory γ which is *not* diagonal-avoiding is shown

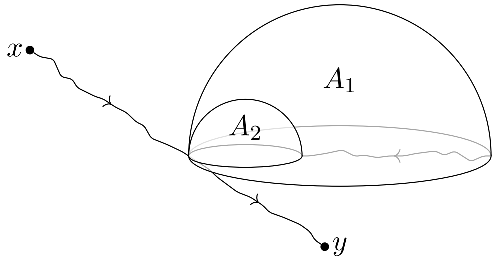

Figure C.7: A non-diagonal-avoiding pearly trajectory.

in Fig. C.7: the entry point of the first disc is equal to the exit point of the second disc, or in other words $\mathrm{ev}_-(A_1)$ and $\mathrm{ev}_+(A_2)$ coincide at γ. The notion of diagonal avoidance clearly also applies to the loose end moduli spaces (C.2) and (C.3) in an obvious way. To be clear, in this case the diagonal avoidance condition applies only to the $2r - 1$ evaluation maps appearing in each of (C.2) and (C.3), not to the evaluation map at the loose end marked point.

We are now ready to describe the attainment of transversality in various settings, so suppose that X is in fact Kähler with integrable complex structure J, and that all J-holomorphic discs in X with boundary on L have all partial indices non-negative (recall that Evans–Lekili showed that these hypotheses are satisfied when (X, L) is K-homogeneous, in the proof of [45, Lemma 3.2]). In particular (i) is automatically satisfied and for all classes $A \in H_2(X, L; \mathbb{Z})$ the evaluation maps $\mathrm{ev}_\pm(A)\colon \mathcal{M}_2(A) \to L$ are submersions. This in turn means that for all tuples of classes $\mathbf{B} = (B_1, \dots, B_r)$ the evaluation map

$$(\mathrm{ev}_+(B_1), \mathrm{ev}_-(B_2), \dots, \mathrm{ev}_+(B_{r-1}), \mathrm{ev}_-(B_r))$$

is transverse to Δ_L^{r-1} (i.e. (C.4) is transverse), so the moduli spaces $\mathcal{M}_2(\mathbf{B})$ of bubbled chains are transversely cut out and thus form smooth moduli spaces of the correct dimension. Given any Morse–Smale pair (f, g), we shall show that it can be pulled back by a diffeomorphism φ of L, which is C^∞-close to id_L, so that the moduli spaces of diagonal-avoiding generalised strings of pearls and of diagonal-avoiding strings with loose ends are transversely cut out.

It is then enough to show that all reduced generalised strings of pearls in virtual dimension at most 1 are diagonal-avoiding. This immediately gives (ii)–(iv), even for non-reduced strings, although the latter aren't actually needed. For the loose end spaces in (v) and (vi) recall that we only need transversality near to the points which actually appear in the degenerate trajectories. In particular, once we have shown that the only degenerate trajectories which occur are diagonal-avoiding, the only trajectories with loose ends which we need consider are those which are also diagonal-avoiding. (Our assumption that all partial indices of holomorphic discs are non-negative actually implies that the loose end moduli spaces here are automatically transversely cut out, even if they are not diagonal-avoiding. We do not make use of this fact in our argument though, as it does not extend to all of the Y-shaped loose end trajectories used later for the Floer product.)

Strictly there *are* non-diagonal-avoiding reduced generalised strings of pearls in virtual dimension at most 1, namely those trajectories with a single disc, which is constant, but we shall show that these are the only exceptions. This issue does not affect the argument for (v) and (vi), and the only potential issue it causes for (ii)–(iv) is with transversality for standard Morse trajectories. However, we assumed that the pair (f, g) we started with was already Morse–Smale, and pulling back by a diffeomorphism does not affect this property, so this potential issue does not actually arise.

We have therefore reduced the problem of constructing a pearl complex using the integrable J to the following two results:

Lemma C.2.1. *For any Morse–Smale pair (f, g) there exists a diffeomorphism φ of L, arbitrarily C^∞-close to id_L, such that the moduli spaces of diagonal-avoiding generalised strings of pearls and of diagonal-avoiding strings with loose ends for auxiliary data $(\varphi^* f, \varphi^* g, J)$ are transversely cut out.*

Proof. Apply Lemma C.1.1, taking the M_j to be products of moduli spaces $\mathcal{M}_2(\mathbf{B})$ of bubbled chains of discs, the f_j to be products of the corresponding evaluation maps, and the N_j to be products of copies of Q with ascending and descending manifolds of critical points of f. □

Lemma C.2.2. *If (f, g) is a Morse–Smale pair for which all moduli spaces of diagonal-avoiding generalised strings of pearls are transversely cut out, then all reduced generalised strings of pearls in virtual dimension at most 1 are diagonal-avoiding unless they have a single disc, which is constant.*

Proof. Suppose that we are given a generalised string of pearls

$$\gamma = \big((u_1^1, \dots, u_{s_1}^1), \dots, (u_1^r, \dots, u_{s_r}^r)\big) \in \mathcal{P}(x, y, \mathbf{A} = (\mathbf{B}^1, \dots, \mathbf{B}^r)) \subset \mathcal{M}_2(\mathbf{B}^1) \times \dots \times \mathcal{M}_2(\mathbf{B}^r)$$

in virtual dimension $d \leq 1$ which has no constant discs. Under $\mathrm{ev}_-(\mathbf{B}^1) \times \dots \times \mathrm{ev}_+(\mathbf{B}^r)$ this evaluates to some point $(p_1, \dots, p_{2r}) \in L^{2r}$; let $N(\gamma)$ be the number of pairs (j, k) with $1 \leq j < k \leq 2r$ and $p_j = p_k$.

If the count $N(\gamma)$ is zero then by definition γ is diagonal-avoiding, so we're done. Otherwise, pick a pair (j, k) contributing to this count. Our aim is to show that we can delete some of the discs to form a new trajectory satisfying the same hypotheses as γ, but now in negative virtual dimension (since the deleted discs are non-constant, and hence of index at least 2) and with a strictly smaller N-value. Repeating this until N reaches 0 we obtain a diagonal-avoiding trajectory—which is therefore cut out transversely—in negative virtual dimension, which is impossible. We thus conclude that γ was diagonal-avoiding to begin with.

There are several cases to consider. First suppose that j and k are both odd—say $j = 2a-1$ and $k = 2b-1$, with $a < b$. In this case we delete the discs u_m^l for $a \leq l < b$, and obtain a generalised string of pearls in virtual dimension

$$d - \sum_{a \leq l < b} \mu(\mathbf{B}^l) \leq d - N_L < 0.$$

Similarly if $j = 2a$ and $k = 2b$ (with $a < b$) or $j = 2a-1$ and $k = 2b$ (with $a \leq b$) then we delete u_m^l for $a < l \leq b$ or $a \leq l \leq b$ respectively. Finally, if $j = 2a$ and $k = 2b-1$ then we must have $b > a+1$ (otherwise (p_j, p_k) lies in Q, which does not meet Δ_L), and we delete u_m^l for $a < l < b$. $\square$

As an aside, we remark that it is possible that one could also achieve transversality by an argument similar to Haug's in [83, Section 7.2], which is itself adapted from the proof of genericity of Morse–Smale metrics for a given Morse function using the Sard–Smale theorem. Rather than restricting to diagonal-avoiding trajectories—which are trajectories whose flowlines don't intersect where they meet discs—one would instead work with trajectories whose flowlines don't intersect at all, and would then have to check that Haug's analysis can all be made to work in this setting. The approach to transversality given above seems preferable however, as it is more concrete and gives better control on the resulting ascending and descending manifolds.

Now return to the case of general (X, L), where there is no longer a preferred choice of almost complex structure. The approach of Biran–Cornea is to *fix* (f, g) and instead choose J' to achieve transversality. Standard arguments with universal moduli spaces show that (i)–(vi) are satisfied for a generic (second category) choice of J', as long as we restrict to trajectories in which the discs are simple and absolutely distinct. This is to ensure that for any trajectory we can perturb J' independently on a neighbourhood of (the image of) an injective point of each disc, and is analogous to the restriction to diagonal-avoiding trajectories in our argument.

One then has to show that, generically, the only trajectories which occur in virtual dimension at most 1 automatically satisfy the simple and absolutely distinct conditions. This is to ensure that imposing these conditions does not destroy compactness: a priori a limit of simple discs need not be simple, for example. Here the argument, described in [14, Section 3.2], splits into two cases ($\dim L \geq 3$ and $\dim L \leq 2$), but both amount to showing that for a second category set of almost complex structures certain additional evaluation map transversality conditions are satisfied (namely [14, Equation (9)] and the two bullet points in the proof of [14, Proposition 3.4.1] in the two cases respectively).

C.3 Invariance of the cohomology

Having constructed the pearl complex, we would like to show that the resulting cohomology is independent of the choice of auxiliary data, up to canonical isomorphism. There are at least three different variants of this independence which one may be interested in. Firstly, if we have a choice of integrable J for which all holomorphic discs have all partial indices non-negative then one may want to prove that the cohomology for this fixed J is independent of the choice of perturbed (f, g) constructed above. Secondly, in the general setting considered by Biran–Cornea one may want to prove that the whole triple (f, g, J') can be varied. And thirdly, one may want to show that the complex we constructed for special J gives the same cohomology as that constructed by Biran–Cornea with generic J'. Of course, the latter really supersedes the first variant, but there are situations where one may want to work at all times with the special J—for instance in the construction of the quantum Pontryagin structure and virtual pearl complex in Chapter 6—and then the first method genuinely is needed.

The first two of these arguments use the method of Morse cobordisms, introduced in [37] and used by Biran–Cornea to prove the second variant of independence in [14, Section 5.1.2]. The idea is as follows. Given two choices (f_0, g_0, J_0) and (f_1, g_1, J_1) of auxiliary data, we choose a Morse cobordism (F, G) comprising a Morse–Smale pair on $L \times [0, 1]$ which coincides with (f_0, g_0) on $L \times \{0\}$ and $(f_1 + C, g_1)$ on $L \times \{1\}$, where $C \gg 0$ is a positive constant, and such that the flow of ∇F is tangent to the boundary $L \times \{0, 1\}$ and points in the direction of strictly increasing t (which denotes the coordinate on the $[0, 1]$ factor) in the interior. We also choose an appropriate homotopy J_t of almost complex structures from J_0 to J_1.

We then consider moduli spaces of discs in $X \times [0, 1]$ with boundary on $L \times [0, 1]$, which are constant—T, say—on the $[0, 1]$ factor, and J_T-holomorphic in the X factor, and carry evaluation maps at two boundary marked points. Using these moduli spaces and the data (F, G) we can build moduli spaces of strings of pearls on $L \times [0, 1]$. Counting rigid strings from $x \in \mathrm{Crit}(f_0) \times \{0\} \subset \mathrm{Crit}(F)$ to $y \in \mathrm{Crit}(f_1) \times \{1\} \subset \mathrm{Crit}(F)$ we obtain a map between the corresponding pearl complexes. The fact that this is a chain map follows from considering one-dimensional moduli spaces of such trajectories, and the boundaries of their compactifications. To prove that this chain map is a quasi-isomorphism, and is independent of the choice of (F, G, J_t) up to chain homotopy, we use Morse cobordisms on $L \times [0, 1]^2$ defined in a similar way.

The key property of all of these constructions, where we work over a parameter space P which is a manifold with corners (equal to $[0, 1]$ or $[0, 1]^2$), is that trajectories (excluding their end points) live entirely in one stratum of P. For example, if γ is a pearly trajectory on $L \times [0, 1]^2$ which contains a disc lying over a point $p \in \{0\} \times (0, 1) \subset [0, 1]^2$ then the whole trajectory, apart from its end points, lies over $\{0\} \times (0, 1)$. The reason for this is that the flow of the Morse cobordism is tangent to each boundary stratum.

Now, to prove the first variant of independence, where J_0 and J_1 are both equal to our special integrable J for which all partial indices are non-negative, we choose an arbitrary Morse cobordism (F, G) and take $J_t = J$ for all t. By construction of (f_0, g_0) and (f_1, g_1) we already have transversality for strings of pearls and loose end spaces which lie over 0 or 1, so we only need worry about transversality for the moduli spaces comprising discs over the interior of $[0, 1]$. And for these we can use the same arguments as for Lemma C.2.1 and Lemma C.2.2 to perturb (F, G) on $L \times (0, 1)$ to achieve transversality. The same approach works for the cobordisms over $[0, 1]^2$, where we start out with transversality for trajectories contained in the boundary strata and perturb the cobordism over the interior.

The second variant is proved by taking an arbitrary cobordism (F, G) and choosing a generic path J_t of almost complex structures from J_0 to J_1 in order to achieve transversality for trajectories of discs which are simple and absolutely distinct (with the images of discs now viewed as subsets of $X \times [0, 1]$ rather than just X), and the additional conditions needed to ensure all trajectories in virtual dimension at most 1 are of this form. Again we can do this since we only need to consider trajectories living over the interior of $[0, 1]$, where we have the freedom to perturb J_t. The $[0, 1]^2$ cobordisms are dealt with similarly.

For the third variant we can actually take a slightly simpler approach. We need not vary the Morse data, and instead can just consider one-parameter families of moduli spaces of strings of pearls in which the almost complex structure varies along a generic path starting at our special J and ending at some generic J_1. The one-dimensional such moduli spaces can be compactified and all boundary points occurring in the interior of the path J_t cancel out. We are left with boundary points occurring at the end of the path—which are pearly trajectories for J_1—and those at the beginning—-which are trajectories for J but counted with a minus sign. The J- and J_1-complexes are thus isomorphic.

Combining the results of the appendix so far with Biran–Cornea's proof that pearl complex (co)homology is self-Floer (co)homology [14, Section 5.6], we have proved the following:

Proposition C.3.1. *Suppose X is a compact Kähler manifold with complex structure J, and $L^\natural \subset X$ is a monotone Lagrangian pre-brane over a ring R, equipped with a Morse–Smale pair (f, g). If every J-holomorphic disc in X with boundary on L has all partial indices non-negative then there exists a diffeomorphism φ of L, arbitrarily C^∞-close to the identity, such that the pearl complex for $L^\natural$ over R can be defined using the auxiliary data $(\varphi^* f, \varphi^* g, J)$, and computes the self-Floer cohomology of $L^\natural$.*

C.4 The Floer product using special J

The pearl complex also carries extra algebraic structures which one may want to compute using a special integrable J (for which all partial indices are non-negative, as above), and we now consider the case of the Floer product. This is defined by taking three Morse–Smale pairs (f_1, g_1), (f_2, g_2) and (f_3, g_3), which are generic in a sense to be made precise later, and defining a map

$$*\colon C_1^p \otimes C_2^q \to C_3^{p+q}$$

(where C_j is the pearl complex constructed using (f_j, g_j, J)) which satisfies the Leibniz rule and hence induces a product on cohomology. This product then has to be shown to be associative and independent of the various choices made.

The arguments involved are fundamentally the same as those used in the preceding subsections, so we focus on the features which require modification. The moduli spaces are more numerous than before and it would be rather cumbersome and unenlightening to express them all individually as fibre products analogous to (C.1), (C.2) and (C.3), so we instead describe them in words and illustrate them with diagrams of examples, which are hopefully easier to digest. It is easy to translate back and forth between these diagrams and fibre product expressions as needed.

The product itself is defined by counting Y-shaped configurations, as shown in Fig. C.8. In

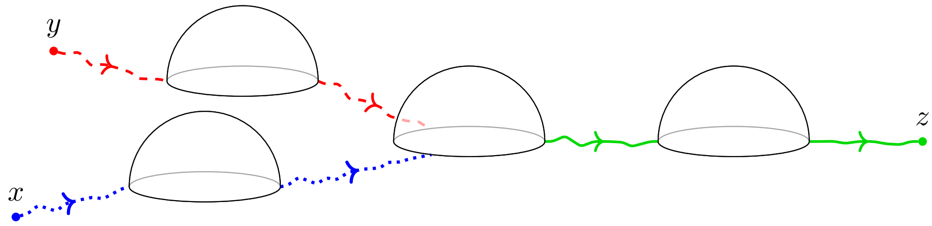

Figure C.8: A Y-shaped pearly trajectory.

the diagram the dotted lines denote flowlines of ∇f_1, whilst dashed is used for f_2 and solid for f_3. Blobs on flowlines denote critical points of the corresponding Morse function. The number of discs shown is purely illustrative: each branch of the Y may have any number of discs, which we may assume to be non-constant, including zero. The central disc *is* allowed to be constant, and terms defining the standard cup product on the Morse cohomology of L come from trajectories in which this is the case and there are no other discs. These moduli spaces have an obvious description analogous to (C.1), in which the central disc carries three marked points, which must be in the order indicated in the diagram (going round clockwise we must have the incoming dotted flowline, then the incoming dashed flowline, and finally the outgoing solid flowline). This restriction on the order of the marked points leads to the failure of the product to be graded-commutative in general.

Just as for the ordinary strings of pearls, it is helpful to consider more general moduli spaces in which we allow discs to be replaced by bubbled chains of discs, or by bubbled Y-shaped configurations of discs at the centre. Some of these bubbled Y-shaped configurations are illustrated in Fig. C.9: the top left diagram shows a non-constant disc with a single bubble at each marked point; the top right shows a constant central disc, bubbled at each marked point; the bottom diagram shows a non-constant central disc which carries a single bubble at one marked point and a chain of two bubbles at another. Note that since the discs are assumed to have all partial indices non-negative these bubbled Y-shaped configurations are also cut out transversely so form smooth moduli spaces of the correct dimension. We shall call trajectories in which discs may be bubbled *generalised Y-shaped strings of pearls* or *generalised Y-shaped pearly trajectories*.

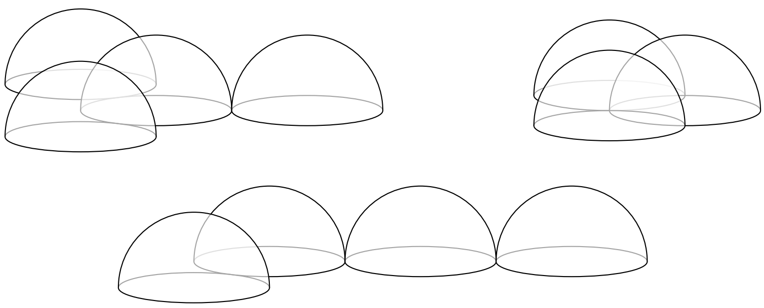

Figure C.9: Bubbled Y-shaped configurations of discs.

We shall also need the corresponding loose end moduli spaces, where there may now be one or two loose ends, as illustrated in Fig. C.10. As before, we insert a constant disc at any loose end point which is bare (i.e. without a disc) to keep track of its position. We say a Y-shaped

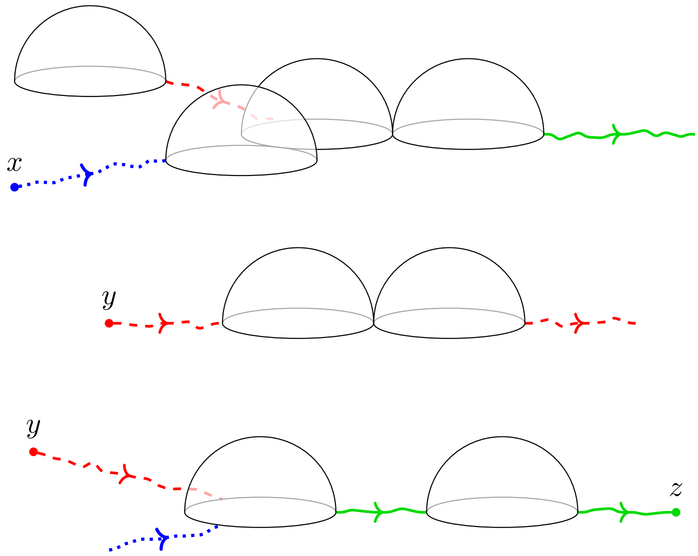

Figure C.10: Y-shaped strings of pearls with loose ends.

pearly trajectory, possibly with loose ends, is *reduced* if the only constant discs are at the centre of a Y or at loose ends.

The notion of diagonal-avoidance has to be slightly modified for these spaces of trajectories involving multiple sets of Morse data. Each evaluation map into L from a moduli space of discs, or more generally of bubbled configurations of discs, comes labelled with a 1, 2 or 3 depending on which function's gradient flow joins up with that evaluation map. For example, in the bottom trajectory in Fig. C.10 the three evaluation maps on the central disc are labelled 1, 2 and 3 clockwise from bottom left (and this will always be the case), the other non-constant disc has both marked points labelled 3 (they join solid flowlines, indicating ∇f_3), whilst the constant disc at the loose end has its evaluation map labelled 1 (as it joins a dotted flowline, meaning ∇f_1). The modified diagonal-avoidance condition is then that the evaluation maps carrying the same label avoid the big diagonal in their corresponding L factors.

With these definitions in place, the transversality we require is that: moduli spaces of diagonal-avoiding generalised strings of pearls and the associated diagonal-avoiding loose end spaces are transversely cut out for each (f_j, g_j); moduli spaces of diagonal-avoiding generalised Y-shaped strings of pearls and the loose end versions are transversely cut out. We'll call these conditions 'product transversality'. Using an obvious modification of Lemma C.1.1 these can be achieved by pulling back the (f_i, g_i) by diffeomorphisms φ_j of L which are C^∞-close to the identity. The first condition ensures that each (f_j, g_j) defines a valid pearl complex, and the second condition lets us define the product by counting rigid reduced Y-shaped strings of pearls. By considering compactifications of moduli spaces of reduced Y-shaped pearly trajectories of virtual dimension 1 we obtain the Leibniz property

$$\mathrm{d}(x * y) = (\mathrm{d}x) * y + (-1)^{|x|} x * (\mathrm{d}y),$$

which means that the product descends to cohomology. The reduced moduli spaces in virtual dimension at most 1 are all automatically diagonal-avoiding, by applying the argument of Lemma C.2.2 to each leg of the Y. Note that the exceptional non-diagonal-avoiding case that occurs for the basic (i.e. non-Y-shaped) trajectories, namely that of standard Morse trajectories, does not occur in the Y-shaped case, since the Morse product trajectories actually *are* diagonal-avoiding in our modified sense (the three flowlines which meet correspond to distinct Morse–Smale pairs which can be perturbed independently).

The only new phenomenon that occurs is bubbling of the thrice-marked central disc at the boundary of one-dimensional moduli spaces, which is taken care of by gluing results analogous to those for twice-marked discs, as in [14, Section 5.2]. Note that convergence of two of the marked points (which can be viewed as bubbling off of a constant 'ghost' disc), is cancelled out by the shrinking of a Morse flowline from a constant central disc to a non-constant disc, as shown from above in Fig. C.11. In particular, if we allow the marked points to appear in both orders around

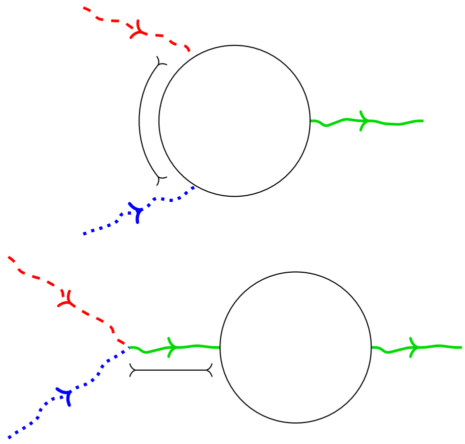

Figure C.11: Convergence of marked points cancels shrinking of a flowline.

the boundary circle then the degenerate configuration occurs at *three* ends of the corresponding one-dimensional moduli space (twice from the convergence of marked points—once in either order—and once from the shrinking of a flowline) and hence does not cancel out. This is why the order of the marked points has to be fixed.

Suppose we replace (f_3, g_3), say, with another Morse–Smale pair (f_3', g_3'). From now on we'll drop explicit mention of the metrics. The key idea for proving invariance of the product is:

Lemma C.4.1. *We can perturb f_3' by a diffeomorphism C^∞-close to the identity in order to achieve product transversality for f_1, f_2 and f_3'.*

Proof. Fix a moduli space $\mathcal{M}$ of trajectories for which we need to achieve transversality. For example, $\mathcal{M}$ could be the space of diagonal-avoiding trajectories of the shape shown in the third diagram in Fig. C.10 (with f_3' in place of f_3), with specified homology classes for the discs. Deleting all flowlines of $\nabla f_3'$, the trajectory breaks into pieces which are either moduli spaces of discs (possibly bubbled chains or bubbled Y-shaped configurations), or loose end trajectories involving only f_1 and f_2. Since we have already attained product transversality for f_1, f_2 and f_3, these loose end trajectories for f_1 and f_2 are transversely cut out.

We can therefore describe $\mathcal{M}$ as a fibre product analogous to (C.1). Instead of taking a product of moduli spaces of discs over the flow spaces (meaning the ascending/descending manifolds, or the space Q) for ∇f, we take a product of moduli spaces of discs (possibly bubbled) or of transversely cut out loose end trajectories for f_1 and f_2, over the flow spaces for $\nabla f_3'$. This is illustrated for our example in Fig. C.12, where the downward arrows represent the evaluation maps at the marked points. Once all of the moduli spaces $\mathcal{M}$ are described in this way, we can

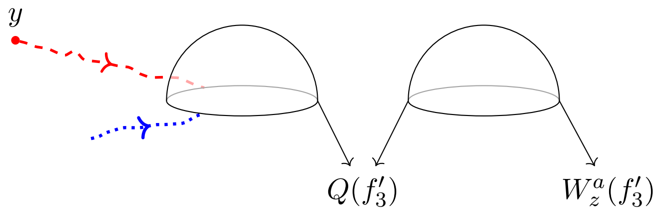

Figure C.12: Expressing $\mathcal{M}$ as a fibre product.

use Lemma C.2.1 as before to see that transversality can be achieved by perturbing f_3'. $\square$

Similarly, if we are given transverse triples of data f_1, f_2, f_3 and f_1, f_2, f_3' then we can perturb any Morse cobordism from f_3 to f_3' as in Section C.3 to ensure the transversality required to get a comparison map from the f_3 complex to the f_3' complex and for this comparison map to respect the product on cohomology. Clearly the same is true if we change either of the other Morse–Smale pairs (f_1, g_1) or (f_2, g_2) instead.

Of course, to prove that the product is independent of the choices of Morse data in general we need a way to compare the products induced by two arbitrary triples f_1, f_2, f_3 and f_1', f_2', f_3'. To do this we introduce an auxiliary triple f_1'', f_2'', f_3'', and Morse cobordisms from each f_j to f_j'' and from f_j'' to f_j', and perturb them all so that we get transverse triples and comparison maps as follows

$$\begin{aligned}(f_1, f_2, f_3) \rightsquigarrow (f_1, f_2, f_3'') \rightsquigarrow (f_1, f_2'', f_3'') \rightsquigarrow (f_1'', f_2'', f_3'')& \\ \rightsquigarrow (f_1'', f_2'', f_3') \rightsquigarrow (f_1'', f_2', f_3') \rightsquigarrow (f_1', f_2', f_3').&\end{aligned}$$

In order to show the product we have just defined using the special integrable J coincides with the product defined using a generic almost complex structure (which obviously gives an indirect proof of the invariance of the former) we proceed as in the third variant of Section C.3: we introduce one-parameter moduli spaces of Y-shaped strings of pearls in which the almost complex structure is allowed to vary along a generic path, and consider the boundaries of the moduli spaces of virtual dimension 1. In Biran–Cornea's work, the same metric is used on each leg of the Y-shaped trajectories, but this is not necessary.

The upshot of this discussion is:

Proposition C.4.2. *In the setting of Proposition C.3.1, but now given three Morse–Smale pairs $(f_j, g_j)_{j=1}^3$ on L, there exist diffeomorphisms $(\varphi_j)_{j=1}^3$ of L, arbitrarily C^∞-close to the identity, such that the Floer product can be computed using the pearl model with auxiliary data $(\varphi_j^* f_j, \varphi_j^* g_j, J)$.*

Appendix D

Coordinates of the Platonic solids

Here we collect together explicit expressions in standard coordinates for the triangle, tetrahedron, octahedron and icosahedron in each of three positions, depending on what feature is pointing vertically upwards: a vertex, the mid-point of an edge, or the centre of a face. We denote these configurations by C_v, C_e and C_f respectively. To remove any ambiguity regarding rotations about a vertical axis, for the edge (respectively face) case we take one end of the top edge (respectively one vertex of the top face) to lie on the positive real axis. In the case of an upward-pointing vertex, we take one of next-northernmost vertices to lie on the positive real axis. With these conventions, we have:

$$\begin{aligned}
\triangle_v &= [1 : 0 : -3 : 0] \\
\triangle_e &= [0 : -3 : 0 : 1] \\
\triangle_f &= [1 : 0 : 0 : 1]
\end{aligned}$$

$$\begin{aligned}
T_v &= [1 : 0 : 0 : 2\sqrt{2} : 0] \\
T_e &= [1 : 0 : 2\sqrt{3} : 0 : -1] \\
T_f &= [0 : 2\sqrt{2} : 0 : 0 : 1]
\end{aligned}$$

$$\begin{aligned}
O_v &= [0 : 1 : 0 : 0 : 0 : -1 : 0] \\
O_e &= [1 : 0 : -5 : 0 : -5 : 0 : 1] \\
O_f &= [1 : 0 : 0 : -5\sqrt{2} : 0 : 0 : -1]
\end{aligned}$$

$$\begin{aligned}
I_v &= [0 : 1 : 0 : 0 : 0 : 0 : -11 : 0 : 0 : 0 : 0 : -1 : 0] \\
I_e &= [\sqrt{5} : 0 : -22 : 0 : -33\sqrt{5} : 0 : 44 : 0 : -33\sqrt{5} : 0 : -22 : 0 : \sqrt{5}] \\
I_f &= [1 : 0 : 0 : -11\sqrt{5} : 0 : 0 : -33 : 0 : 0 : 11\sqrt{5} : 0 : 0 : 1].
\end{aligned}$$

All of these can be computed straightforwardly, with the help of a computer to simplify the algebra, using (4.3). For instance, we know that I_v has vertices at 0, ∞, $\zeta^j a$ and $-\zeta^j/a$ for some real number $a > 1$, where ζ is a primitive fifth root of unity and j ranges from 1 to 5. Since all of the edges of the icosahedron are of equal length, we know that the angle between the points ∞ and a is the same as the angle between a and ζa. We thus have

$$\frac{a}{\sqrt{1+a^2}} = \frac{|1+\zeta a^2|}{1+a^2},$$

and after squaring and simplifying, using the fact that $\cos(2\pi/5) = (\sqrt{5}-1)/4$, we get $a = (\sqrt{5}+1)/2$. Therefore

$$I_v = \left[xy(a^5x^5+y^5)(x^5-a^5y^5)\right] = \left[xy(x^{10}-11x^5y^5-y^{10})\right].$$

The other calculations are similar.

Bibliography

[1] M. Abouzaid, "A geometric criterion for generating the Fukaya category," *Publ. Math. Inst. Hautes Études Sci.* no. 112, (2010) 191–240. `http://dx.doi.org/10.1007/s10240-010-0028-5`.

[2] M. Abouzaid, "A cotangent fibre generates the Fukaya category," *Adv. Math.* **228** no. 2, (2011) 894–939. `http://dx.doi.org/10.1016/j.aim.2011.06.007`.

[3] M. Abouzaid, "Nearby Lagrangians with vanishing Maslov class are homotopy equivalent," *Invent. Math.* **189** no. 2, (2012) 251–313. `http://dx.doi.org/10.1007/s00222-011-0365-0`.

[4] M. Abouzaid and I. Smith, "Homological mirror symmetry for the 4-torus," *Duke Math. J.* **152** no. 3, (2010) 373–440. `http://dx.doi.org/10.1215/00127094-2010-015`.

[5] M. Abreu and L. Macarini, "Remarks on Lagrangian intersections in toric manifolds," *Trans. Amer. Math. Soc.* **365** no. 7, (2013) 3851–3875. `http://dx.doi.org/10.1090/S0002-9947-2012-05791-6`.

[6] P. Aluffi and C. Faber, "Linear orbits of d-tuples of points in $\mathbf{P}^1$," *J. Reine Angew. Math.* **445** (1993) 205–220.

[7] A. Amarzaya and Y. Ohnita, "Hamiltonian stability of certain minimal Lagrangian submanifolds in complex projective spaces," *Tohoku Math. J. (2)* **55** no. 4, (2003) 583–610. `http://projecteuclid.org/euclid.tmj/1113247132`.

[8] M. Audin and M. Damian, *Morse theory and Floer homology.* Universitext. Springer, London; EDP Sciences, Les Ulis, 2014. `http://dx.doi.org/10.1007/978-1-4471-5496-9`. Translated from the 2010 French original by Reinie Erné.

[9] D. Auroux, "Mirror symmetry and T-duality in the complement of an anticanonical divisor," *J. Gökova Geom. Topol. GGT* **1** (2007) 51–91.

[10] P. F. Baum and W. Browder, "The cohomology of quotients of classical groups," *Topology* **3** (1965) 305–336.

[11] A. Bayer and Y. I. Manin, "(Semi)simple exercises in quantum cohomology," in *The Fano Conference*, pp. 143–173. Univ. Torino, Turin, 2004.

[12] P. Biran, "Lagrangian non-intersections," *Geom. Funct. Anal.* **16** no. 2, (2006) 279–326. `http://dx.doi.org/10.1007/s00039-006-0560-0`.

[13] P. Biran and K. Cieliebak, "Symplectic topology on subcritical manifolds," *Comment. Math. Helv.* **76** no. 4, (2001) 712–753. `http://dx.doi.org/10.1007/s00014-001-8326-7`.

[14] P. Biran and O. Cornea, "Quantum structures for Lagrangian submanifolds," `arXiv:0708.4221v1 [math.SG]`.

[15] P. Biran and O. Cornea, "A Lagrangian quantum homology," in *New perspectives and challenges in symplectic field theory*, vol. 49 of *CRM Proc. Lecture Notes*, pp. 1–44. Amer. Math. Soc., Providence, RI, 2009.

[16] P. Biran and O. Cornea, "Rigidity and uniruling for Lagrangian submanifolds," *Geom. Topol.* **13** no. 5, (2009) 2881–2989. `http://dx.doi.org/10.2140/gt.2009.13.2881`.

[17] P. Biran and O. Cornea, "Lagrangian topology and enumerative geometry," *Geom. Topol.* **16** no. 2, (2012) 963–1052. `http://dx.doi.org/10.2140/gt.2012.16.963`.

[18] P. Biran and M. Khanevsky, "A Floer-Gysin exact sequence for Lagrangian submanifolds," *Comment. Math. Helv.* **88** no. 4, (2013) 899–952. `http://dx.doi.org/10.4171/CMH/307`.

[19] P. Boalch, "Symplectic manifolds and isomonodromic deformations," *Adv. Math.* **163** no. 2, (2001) 137–205. `http://dx.doi.org/10.1006/aima.2001.1998`.

[20] A. Borel, "Sur la cohomologie des espaces fibrés principaux et des espaces homogènes de groupes de Lie compacts," *Ann. of Math. (2)* **57** (1953) 115–207. `http://dx.doi.org/10.2307/1969728`.

[21] A. Borel, "Sur l'homologie et la cohomologie des groupes de Lie compacts connexes," *Amer. J. Math.* **76** (1954) 273–342.

[22] T. Bröcker and T. tom Dieck, *Representations of compact Lie groups*, vol. 98 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995. Translated from the German manuscript, Corrected reprint of the 1985 translation.

[23] F. E. Browder, ed., *Mathematical developments arising from Hilbert problems.* Proceedings of Symposia in Pure Mathematics, Vol. XXVIII. American Mathematical Society, Providence, R. I., 1976.

[24] L. Buhovsky, "The Maslov class of Lagrangian tori and quantum products in Floer cohomology," *J. Topol. Anal.* **2** no. 1, (2010) 57–75. `http://dx.doi.org/10.1142/S1793525310000240`.

[25] F. Campana, "On twistor spaces of the class $\mathscr{C}$," *J. Differential Geom.* **33** no. 2, (1991) 541–549. `http://projecteuclid.org/euclid.jdg/1214446329`.

[26] A. Cannas da Silva, "A Chiang-type Lagrangian in $\mathbb{CP}^2$," arXiv:1511.02041v1 [math.SG].

[27] P. Cartier, "A primer of Hopf algebras," in *Frontiers in number theory, physics, and geometry. II*, pp. 537–615. Springer, Berlin, 2007. http://dx.doi.org/10.1007/978-3-540-30308-4_12.

[28] F. Charette and O. Cornea, "Categorification of Seidel's representation," *Israel J. Math.* **211** no. 1, (2016) 67–104. http://dx.doi.org/10.1007/s11856-015-1261-x.

[29] Y. V. Chekanov, "Lagrangian tori in a symplectic vector space and global symplectomorphisms," *Math. Z.* **223** no. 4, (1996) 547–559. http://dx.doi.org/10.1007/PL00004278.

[30] Y. Chekanov and F. Schlenk, "Notes on monotone Lagrangian twist tori," *Electron. Res. Announc. Math. Sci.* **17** (2010) 104–121. http://dx.doi.org/10.3934/era.2010.17.104.

[31] R. Chiang, "New Lagrangian submanifolds of $\mathbb{CP}^n$," *Int. Math. Res. Not.* no. 45, (2004) 2437–2441. http://dx.doi.org/10.1155/S1073792804133102.

[32] C.-H. Cho, "Holomorphic discs, spin structures, and Floer cohomology of the Clifford torus," *Int. Math. Res. Not.* no. 35, (2004) 1803–1843. http://dx.doi.org/10.1155/S1073792804132716.

[33] C.-H. Cho, "Products of Floer cohomology of torus fibers in toric Fano manifolds," *Comm. Math. Phys.* **260** no. 3, (2005) 613–640. http://dx.doi.org/10.1007/s00220-005-1421-7.

[34] C.-H. Cho, "Non-displaceable Lagrangian submanifolds and Floer cohomology with non-unitary line bundle," *J. Geom. Phys.* **58** no. 11, (2008) 1465–1476. http://dx.doi.org/10.1016/j.geomphys.2008.06.003.

[35] C.-H. Cho and Y.-G. Oh, "Floer cohomology and disc instantons of Lagrangian torus fibers in Fano toric manifolds," *Asian J. Math.* **10** no. 4, (2006) 773–814. http://dx.doi.org/10.4310/AJM.2006.v10.n4.a10.

[36] K. Cieliebak, A. R. Gaio, and D. A. Salamon, "J-holomorphic curves, moment maps, and invariants of Hamiltonian group actions," *Internat. Math. Res. Notices* no. 16, (2000) 831–882. http://dx.doi.org/10.1155/S1073792800000453.

[37] O. Cornea and A. Ranicki, "Rigidity and gluing for Morse and Novikov complexes," *J. Eur. Math. Soc. (JEMS)* **5** no. 4, (2003) 343–394. http://dx.doi.org/10.1007/s10097-003-0052-6.

[38] M. Damian, "Floer homology on the universal cover, Audin's conjecture and other constraints on Lagrangian submanifolds," *Comment. Math. Helv.* **87** no. 2, (2012) 433–462. http://dx.doi.org/10.4171/CMH/259.

[39] V. de Silva, *Products in Symplectic Floer Homology of Lagrangian Intersection.* PhD thesis, University of Oxford, 1999.

[40] P. Deligne, *Équations différentielles à points singuliers réguliers.* Lecture Notes in Mathematics, Vol. 163. Springer-Verlag, Berlin-New York, 1970.

[41] D. Eisenbud, *Commutative algebra*, vol. 150 of *Graduate Texts in Mathematics.* Springer-Verlag, New York, 1995. http://dx.doi.org/10.1007/978-1-4612-5350-1. With a view toward algebraic geometry.

[42] M. Entov and L. Polterovich, "Rigid subsets of symplectic manifolds," *Compos. Math.* **145** no. 3, (2009) 773–826. http://dx.doi.org/10.1112/S0010437X0900400X.

[43] J. D. Evans, "Lagrangian spheres in del Pezzo surfaces," *J. Topol.* **3** no. 1, (2010) 181–227. http://dx.doi.org/10.1112/jtopol/jtq004.

[44] J. D. Evans and Y. Lekili, "Generating the Fukaya categories of Hamiltonian G-manifolds," arXiv:1507.05842v2 [math.SG].

[45] J. D. Evans and Y. Lekili, "Floer cohomology of the Chiang Lagrangian," *Selecta Math. (N.S.)* **21** no. 4, (2015) 1361–1404. http://dx.doi.org/10.1007/s00029-014-0171-9.

[46] A. Floer, "Holomorphic curves and a Morse theory for fixed points of exact symplectomorphisms," in *Aspects dynamiques et topologiques des groupes infinis de transformation de la mécanique (Lyon, 1986)*, vol. 25 of *Travaux en Cours*, pp. 49–60. Hermann, Paris, 1987.

[47] A. Floer, "Morse theory for fixed points of symplectic diffeomorphisms," *Bull. Amer. Math. Soc. (N.S.)* **16** no. 2, (1987) 279–281. http://dx.doi.org/10.1090/S0273-0979-1987-15517-0.

[48] A. Floer, "Morse theory for Lagrangian intersections," *J. Differential Geom.* **28** no. 3, (1988) 513–547. http://projecteuclid.org/euclid.jdg/1214442477.

[49] A. Floer, "A relative Morse index for the symplectic action," *Comm. Pure Appl. Math.* **41** no. 4, (1988) 393–407. http://dx.doi.org/10.1002/cpa.3160410402.

[50] A. Floer, "The unregularized gradient flow of the symplectic action," *Comm. Pure Appl. Math.* **41** no. 6, (1988) 775–813. http://dx.doi.org/10.1002/cpa.3160410603.

[51] A. Floer, "Symplectic fixed points and holomorphic spheres," *Comm. Math. Phys.* **120** no. 4, (1989) 575–611. http://projecteuclid.org/euclid.cmp/1104177909.

[52] A. Floer, "Witten's complex and infinite-dimensional Morse theory," *J. Differential Geom.* **30** no. 1, (1989) 207–221. http://projecteuclid.org/euclid.jdg/1214443291.

[53] U. Frauenfelder, "The Arnold-Givental conjecture and moment Floer homology," *Int. Math. Res. Not.* no. 42, (2004) 2179–2269. http://dx.doi.org/10.1155/S1073792804133941.

[54] K. Fukaya, P. Seidel, and I. Smith, "The symplectic geometry of cotangent bundles from a categorical viewpoint," in *Homological mirror symmetry*, vol. 757 of *Lecture Notes in Phys.*, pp. 1–26. Springer, Berlin, 2009.
http://dx.doi.org/10.1007/978-3-540-68030-7_1.

[55] K. Fukaya, "Lie groupoid, deformation of unstable curve, and construction of equivariant Kuranishi charts," arXiv:1701.02840v1 [math.SG].

[56] K. Fukaya, "Unobstructed immersed Lagrangian correspondence and filtered A_∞ functor," arXiv:1706.02131v2 [math.SG].

[57] K. Fukaya, "Morse homotopy, A^∞-category, and Floer homologies," in *Proceedings of GARC Workshop on Geometry and Topology '93 (Seoul, 1993)*, vol. 18 of *Lecture Notes Ser.*, pp. 1–102. Seoul Nat. Univ., Seoul, 1993.

[58] K. Fukaya, "Floer homology for families—a progress report," in *Integrable systems, topology, and physics (Tokyo, 2000)*, vol. 309 of *Contemp. Math.*, pp. 33–68. Amer. Math. Soc., Providence, RI, 2002. http://dx.doi.org/10.1090/conm/309/05341.

[59] K. Fukaya, "Application of Floer homology of Langrangian submanifolds to symplectic topology," in *Morse theoretic methods in nonlinear analysis and in symplectic topology*, vol. 217 of *NATO Sci. Ser. II Math. Phys. Chem.*, pp. 231–276. Springer, Dordrecht, 2006. http://dx.doi.org/10.1007/1-4020-4266-3_06.

[60] K. Fukaya. Personal communication, 2016.

[61] K. Fukaya, "Equivariant Kuranishi structure and its (potential) applications," 2016. https://video.ias.edu/file/42734. IAS video lecture, *Workshop on Homological Mirror Symmetry: Methods and Structures.*

[62] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, "Chapter 10. Lagrangian surgery and holomorphic discs," 2007.
http://www.math.kyoto-u.ac.jp/~fukaya/Chapter10071117.pdf.

[63] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, *Lagrangian intersection Floer theory: anomaly and obstruction. Two-volume set*, vol. 46 of *AMS/IP Studies in Advanced Mathematics*. American Mathematical Society, Providence, RI; International Press, Somerville, MA, 2009.

[64] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, "Lagrangian Floer theory on compact toric manifolds. I," *Duke Math. J.* **151** no. 1, (2010) 23–174.
http://dx.doi.org/10.1215/00127094-2009-062.

[65] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, "Lagrangian Floer theory on compact toric manifolds II: Bulk deformations," *Selecta Math. (N.S.)* **17** no. 3, (2011) 609–711.
http://dx.doi.org/10.1007/s00029-011-0057-z.

[66] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, "Lagrangian Floer theory on compact toric manifolds: Survey," in *Surveys in differential geometry. Vol. XVII*, vol. 17 of *Surv. Differ. Geom.*, pp. 229–298. Int. Press, Boston, MA, 2012. http://dx.doi.org/10.4310/SDG.2012.v17.n1.a6.

[67] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, "Toric degeneration and nondisplaceable Lagrangian tori in $S^2 \times S^2$," *Int. Math. Res. Not. IMRN* no. 13, (2012) 2942–2993.

[68] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, "Lagrangian Floer theory and mirror symmetry on compact toric manifolds," *Astérisque* no. 376, (2016) vi+340.

[69] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, "Antisymplectic involution and Floer cohomology," *Geom. Topol.* **21** no. 1, (2017) 1–106. http://dx.doi.org/10.2140/gt.2017.21.1.

[70] K. Fukaya and K. Ono, "Arnold conjecture and Gromov-Witten invariant," *Topology* **38** no. 5, (1999) 933–1048. http://dx.doi.org/10.1016/S0040-9383(98)00042-1.

[71] K. Fukaya, P. Seidel, and I. Smith, "Exact Lagrangian submanifolds in simply-connected cotangent bundles," *Invent. Math.* **172** no. 1, (2008) 1–27. http://dx.doi.org/10.1007/s00222-007-0092-8.

[72] A. Gadbled, "On exotic monotone Lagrangian tori in $\mathbb{CP}^2$ and $\mathbb{S}^2 \times \mathbb{S}^2$," *J. Symplectic Geom.* **11** no. 3, (2013) 343–361. http://projecteuclid.org/euclid.jsg/1384282840.

[73] S. Ganatra, *Symplectic Cohomology and Duality for the Wrapped Fukaya Category.* ProQuest LLC, Ann Arbor, MI, 2012. http://gateway.proquest.com/openurl?url_ver=Z39.88-2004&rft_val_fmt=info:ofi/fmt:kev:mtx:dissertation&res_dat=xri:pqm&rft_dat=xri:pqdiss:0828785. Thesis (Ph.D.)–Massachusetts Institute of Technology.

[74] Y. Gao, "Wrapped Floer cohomology and Lagrangian correspondences," arXiv:1703.04032v1 [math.SG].

[75] A. B. Givental, "Periodic mappings in symplectic topology," *Funktsional. Anal. i Prilozhen.* **23** no. 4, (1989) 37–52, 96. http://dx.doi.org/10.1007/BF01078943.

[76] M. Goresky, R. Kottwitz, and R. MacPherson, "Equivariant cohomology, Koszul duality, and the localization theorem," *Invent. Math.* **131** no. 1, (1998) 25–83. http://dx.doi.org/10.1007/s002220050197.

[77] P. Griffiths and J. Harris, *Principles of algebraic geometry.* Wiley Classics Library. John Wiley & Sons, Inc., New York, 1994. http://dx.doi.org/10.1002/9781118032527. Reprint of the 1978 original.

[78] M. Gromov, "Pseudoholomorphic curves in symplectic manifolds," *Invent. Math.* **82** no. 2, (1985) 307–347. http://dx.doi.org/10.1007/BF01388806.

[79] V. Guillemin and S. Sternberg, "Geometric quantization and multiplicities of group representations," *Invent. Math.* **67** no. 3, (1982) 515–538. http://dx.doi.org/10.1007/BF01398934.

[80] V. Guillemin and S. Sternberg, "Some problems in integral geometry and some related problems in microlocal analysis," *Amer. J. Math.* **101** no. 4, (1979) 915–955. http://dx.doi.org/10.2307/2373923.

[81] V. Guillemin and S. Sternberg, "The moment map revisited," *J. Differential Geom.* **69** no. 1, (2005) 137–162. http://projecteuclid.org/euclid.jdg/1121540342.

[82] A. Hatcher, *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[83] L. Haug, "On the quantum homology of real Lagrangians in Fano toric manifolds," *Int. Math. Res. Not. IMRN* no. 14, (2013) 3171–3220.

[84] J. Helmstetter and A. Micali, *Quadratic mappings and Clifford algebras.* Birkhäuser Verlag, Basel, 2008.

[85] J. Hilgert and K.-H. Neeb, *Structure and geometry of Lie groups.* Springer Monographs in Mathematics. Springer, New York, 2012. http://dx.doi.org/10.1007/978-0-387-84794-8.

[86] M. W. Hirsch, *Differential topology*, vol. 33 of *Graduate Texts in Mathematics.* Springer-Verlag, New York, 1994. Corrected reprint of the 1976 original.

[87] N. J. Hitchin, "Poncelet polygons and the Painlevé equations," in *Geometry and analysis (Bombay, 1992)*, pp. 151–185. Tata Inst. Fund. Res., Bombay, 1995.

[88] K. H. Hofmann and S. A. Morris, *The structure of compact groups*, vol. 25 of *De Gruyter Studies in Mathematics.* Walter de Gruyter & Co., Berlin, augmented ed., 2006. http://dx.doi.org/10.1515/9783110199772. A primer for the student—a handbook for the expert.

[89] J. E. Humphreys, *Linear algebraic groups.* Springer-Verlag, New York-Heidelberg, 1975. Graduate Texts in Mathematics, No. 21.

[90] M. Hutchings, "Lecture notes on Morse homology (with an eye towards Floer theory and pseudoholomorphic curves),". https://math.berkeley.edu/~hutching/teach/276-2010/mfp.ps.

[91] H. Iriyeh, "Symplectic topology of Lagrangian submanifolds of $\mathbb{C}P^n$ with intermediate minimal Maslov numbers," *Adv. Geom.* **17** no. 2, (2017) 247–264. http://dx.doi.org/10.1515/advgeom-2017-0005.

[92] S. Kobayashi, *Transformation groups in differential geometry.* Classics in Mathematics. Springer-Verlag, Berlin, 1995. Reprint of the 1972 edition.

[93] J. Kollár, Y. Miyaoka, and S. Mori, "Rational connectedness and boundedness of Fano manifolds," *J. Differential Geom.* **36** no. 3, (1992) 765–779. `http://projecteuclid.org/euclid.jdg/1214453188`.

[94] M. Konstantinov, "Higher rank local systems in Lagrangian Floer theory," `arXiv:1701.03624v2 [math.SG]`.

[95] M. Kontsevich, "Homological algebra of mirror symmetry," in *Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994)*, pp. 120–139. Birkhäuser, Basel, 1995.

[96] T. Kragh, "Parametrized ring-spectra and the nearby Lagrangian conjecture," *Geom. Topol.* **17** no. 2, (2013) 639–731. `http://dx.doi.org/10.2140/gt.2013.17.639`. With an appendix by Mohammed Abouzaid.

[97] P. Kronheimer and T. Mrowka, *Monopoles and three-manifolds*, vol. 10 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2007. `http://dx.doi.org/10.1017/CBO9780511543111`.

[98] Y. Lekili and M. Lipyanskiy, "Geometric composition in quilted Floer theory," *Adv. Math.* **236** (2013) 1–23. `http://dx.doi.org/10.1016/j.aim.2012.12.012`.

[99] G. Liu and G. Tian, "Floer homology and Arnold conjecture," *J. Differential Geom.* **49** no. 1, (1998) 1–74. `http://projecteuclid.org/euclid.jdg/1214460936`.

[100] L. Manivel, *Symmetric functions, Schubert polynomials and degeneracy loci*, vol. 6 of *SMF/AMS Texts and Monographs*. American Mathematical Society, Providence, RI; Société Mathématique de France, Paris, 2001. Translated from the 1998 French original by John R. Swallow, Cours Spécialisés [Specialized Courses], 3.

[101] C.-M. Marle, "The inception of symplectic geometry: The works of Lagrange and Poisson during the years 1808–1810," *Lett. Math. Phys.* **90** no. 1-3, (2009) 3–21. `http://dx.doi.org/10.1007/s11005-009-0347-y`.

[102] S. Ma'u, K. Wehrheim, and C. T. Woodward, "A_∞ functors for Lagrangian correspondences," `arXiv:1601.04919v2 [math.SG]`.

[103] D. McDuff and D. Salamon, *J-holomorphic curves and quantum cohomology*, vol. 6 of *University Lecture Series*. American Mathematical Society, Providence, RI, 1994. `http://dx.doi.org/10.1090/ulect/006`.

[104] D. McDuff and D. Salamon, *J-holomorphic curves and symplectic topology*, vol. 52 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2004.

[105] D. McDuff and S. Tolman, "Topological properties of Hamiltonian circle actions," *IMRP Int. Math. Res. Pap.* (2006) 72826, 1–77.

[106] D. Nadler, "Microlocal branes are constructible sheaves," *Selecta Math. (N.S.)* **15** no. 4, (2009) 563–619. `http://dx.doi.org/10.1007/s00029-009-0008-0`.

[107] J. Oakley and M. Usher, "On certain Lagrangian submanifolds of $S^2 \times S^2$ and $\mathbb{C}\mathrm{P}^n$," *Algebr. Geom. Topol.* **16** no. 1, (2016) 149–209. `http://dx.doi.org/10.2140/agt.2016.16.149`.

[108] Y.-G. Oh, "Floer cohomology of Lagrangian intersections and pseudo-holomorphic disks. III. Arnol'd-Givental conjecture," in *The Floer memorial volume*, vol. 133 of *Progr. Math.*, pp. 555–573. Birkhäuser, Basel, 1995.

[109] Y.-G. Oh, "Floer cohomology of Lagrangian intersections and pseudo-holomorphic disks. I," *Comm. Pure Appl. Math.* **46** no. 7, (1993) 949–993. `http://dx.doi.org/10.1002/cpa.3160460702`.

[110] Y.-G. Oh, "Riemann-Hilbert problem and application to the perturbation theory of analytic discs," *Kyungpook Math. J.* **35** no. 1, (1995) 39–75.

[111] Y.-G. Oh, "Floer cohomology, spectral sequences, and the Maslov class of Lagrangian embeddings," *Internat. Math. Res. Notices* no. 7, (1996) 305–346. `http://dx.doi.org/10.1155/S1073792896000219`.

[112] P. Ozsváth and Z. Szabó, "Holomorphic disks and knot invariants," *Adv. Math.* **186** no. 1, (2004) 58–116. `http://dx.doi.org/10.1016/j.aim.2003.05.001`.

[113] P. Ozsváth and Z. Szabó, "Holomorphic disks and topological invariants for closed three-manifolds," *Ann. of Math. (2)* **159** no. 3, (2004) 1027–1158. `http://dx.doi.org/10.4007/annals.2004.159.1027`.

[114] J. Pardon, "An algebraic approach to virtual fundamental cycles on moduli spaces of pseudo-holomorphic curves," *Geom. Topol.* **20** no. 2, (2016) 779–1034. `http://dx.doi.org/10.2140/gt.2016.20.779`.

[115] J. Pascaleff and D. Tonkonog, "The wall-crossing formula for mutations of Lagrangian tori," in preparation.

[116] T. Perutz, "A symplectic Gysin sequence," `arXiv:0807.1863v1 [math.SG]`.

[117] A. Polishchuk and E. Zaslow, "Categorical mirror symmetry in the elliptic curve [MR1633036 (99j:14034)]," in *Winter School on Mirror Symmetry, Vector Bundles and Lagrangian Submanifolds (Cambridge, MA, 1999)*, vol. 23 of *AMS/IP Stud. Adv. Math.*, pp. 275–295. Amer. Math. Soc., Providence, RI, 2001.

[118] J. A. Rasmussen, *Floer homology and knot complements.* ProQuest LLC, Ann Arbor, MI, 2003. `http://gateway.proquest.com/openurl?url_ver=Z39.88-2004&rft_val_fmt=info:ofi/fmt:kev:mtx:dissertation&res_dat=xri:pqdiss&rft_dat=xri:pqdiss:3091665`. Thesis (Ph.D.)–Harvard University.

[119] P. Seidel, "π_1 of symplectic automorphism groups and invertibles in quantum homology rings," *Geom. Funct. Anal.* **7** no. 6, (1997) 1046–1095. http://dx.doi.org/10.1007/s000390050037.

[120] P. Seidel, "Graded Lagrangian submanifolds," *Bull. Soc. Math. France* **128** no. 1, (2000) 103–149. http://www.numdam.org/item?id=BSMF_2000__128_1_103_0.

[121] P. Seidel, "A long exact sequence for symplectic Floer cohomology," *Topology* **42** no. 5, (2003) 1003–1063. http://dx.doi.org/10.1016/S0040-9383(02)00028-9.

[122] P. Seidel, "Exact Lagrangian submanifolds in T^*S^n and the graded Kronecker quiver," in *Different faces of geometry*, vol. 3 of *Int. Math. Ser. (N. Y.)*, pp. 349–364. Kluwer/Plenum, New York, 2004. http://dx.doi.org/10.1007/0-306-48658-X_8.

[123] P. Seidel, *Fukaya categories and Picard-Lefschetz theory.* Zurich Lectures in Advanced Mathematics. European Mathematical Society (EMS), Zürich, 2008. http://dx.doi.org/10.4171/063.

[124] P. Seidel, "Lectures on four-dimensional Dehn twists," in *Symplectic 4-manifolds and algebraic surfaces*, vol. 1938 of *Lecture Notes in Math.*, pp. 231–267. Springer, Berlin, 2008. http://dx.doi.org/10.1007/978-3-540-78279-7_4.

[125] P. Seidel, "Homological mirror symmetry for the genus two curve," *J. Algebraic Geom.* **20** no. 4, (2011) 727–769. http://dx.doi.org/10.1090/S1056-3911-10-00550-3.

[126] P. Seidel, "Abstract analogues of flux as symplectic invariants," *Mém. Soc. Math. Fr. (N.S.)* no. 137, (2014) 135.

[127] P. Seidel, "Homological mirror symmetry for the quartic surface," *Mem. Amer. Math. Soc.* **236** no. 1116, (2015) vi+129. http://dx.doi.org/10.1090/memo/1116.

[128] P. Seidel and I. Smith, "A link invariant from the symplectic geometry of nilpotent slices," *Duke Math. J.* **134** no. 3, (2006) 453–514. http://dx.doi.org/10.1215/S0012-7094-06-13432-4.

[129] N. Sheridan, "Homological mirror symmetry for Calabi-Yau hypersurfaces in projective space," *Invent. Math.* **199** no. 1, (2015) 1–186. http://dx.doi.org/10.1007/s00222-014-0507-2.

[130] N. Sheridan, "On the Fukaya category of a Fano hypersurface in projective space," *Publ. Math. Inst. Hautes Études Sci.* **124** (2016) 165–317. http://dx.doi.org/10.1007/s10240-016-0082-8.

[131] K. Siegel, "Squared Dehn twists and deformed symplectic invariants," arXiv:1609.08545v1 [math.SG].

[132] I. Smith, "Floer cohomology and pencils of quadrics," *Invent. Math.* **189** no. 1, (2012) 149–250. http://dx.doi.org/10.1007/s00222-011-0364-1.

[133] I. Smith, "A symplectic prolegomenon," *Bull. Amer. Math. Soc. (N.S.)* **52** no. 3, (2015) 415–464. `http://dx.doi.org/10.1090/S0273-0979-2015-01477-1`.

[134] J. Smith, "Discrete and continuous symmetries in monotone Floer theory," `arXiv:1703.05343v2 [math.SG]`.

[135] J. Smith, "Floer cohomology of Platonic Lagrangians," `arXiv:1510.08031v3 [math.SG]`.

[136] C. Teleman, "Gauge theory and mirror symmetry," `arXiv:1404.6305v1 [math-ph]`.

[137] D. Tonkonog, "The closed–open string map for S^1-invariant Lagrangians," `arXiv:1504.01621v2 [math.SG]`.

[138] D. Tonkonog and R. Vianna, "Low-area Floer theory and non-displaceability," `arXiv:1511.00891v2 [math.SG]`.

[139] B. Torricelli, "A survey on the Lagrangian $\mathrm{SU}(n)/(\mathbb{Z}/n)$ in $\mathbb{CP}^{n^2-1}$," Master's thesis, University College London, 2016, unpublished.

[140] C. Viterbo, "Exact Lagrange submanifolds, periodic orbits and the cohomology of free loop spaces," *J. Differential Geom.* **47** no. 3, (1997) 420–468. `http://projecteuclid.org/euclid.jdg/1214460546`.

[141] K. Wehrheim and C. T. Woodward, "Orientations for pseudoholomorphic quilts," `arXiv:1503.07803v6 [math.SG]`.

[142] K. Wehrheim and C. T. Woodward, "Quilted Floer cohomology," *Geom. Topol.* **14** no. 2, (2010) 833–902. `http://dx.doi.org/10.2140/gt.2010.14.833`.

[143] A. Weinstein, "Symplectic geometry," *Bull. Amer. Math. Soc. (N.S.)* **5** no. 1, (1981) 1–13. `http://dx.doi.org/10.1090/S0273-0979-1981-14911-9`.

[144] E. Witten, "The Verlinde algebra and the cohomology of the Grassmannian," in *Geometry, topology, & physics*, Conf. Proc. Lecture Notes Geom. Topology, IV, pp. 357–422. Int. Press, Cambridge, MA, 1995.

[145] W. Wu, "On an exotic Lagrangian torus in $\mathbb{C}P^2$," *Compos. Math.* **151** no. 7, (2015) 1372–1394. `http://dx.doi.org/10.1112/S0010437X14007945`.

[146] F. Zapolsky, "The Lagrangian Floer-quantum-PSS package and canonical orientations in Floer theory," `arXiv:1507.02253v2 [math.SG]`.

[147] A. Zinger, "The determinant line bundle for Fredholm operators: Construction, properties, and classification," *Math. Scand.* **118** no. 2, (2016) 203–268. `http://dx.doi.org/10.7146/math.scand.a-23687`.